Top Tipps aus dem Support von IBM zur Informix Datenbank

Dr. Elisabeth Bach
IBM

Contents

- Tips from Support
- A few new onstats
- IDS and Windows
- Roadmap

Tips from Support

Performance PMRs: Stories from support

- It was on one Monday morning…..
- Everything is slow
- We need to fix it IMMEDIAETLY

What has happened?

What was changed?

What is bad now?

So what to do?

- Find out what has happened
- Check cpu/memory/disk
- Check what threads are doing
  - Onstat –g stk
  - Get stacks from running threads
- Check the online.log

Rules of thumb

You cannot extrapolate from one measurement. Always repeat a few times

Where there is a persistent queue, there is likely a bottleneck.

Look at the online.log

Stick to your processes, especially during times of stress

DONTDRAINPOOLS

- Environment variable

  Set DONTDRAINPOOL=1

  Read during initialization

- Writes a message the online.log

  ```
  Server is disabling pools draining
  ```

- Sessions might use more memory as before

Snooping at your system

- System information
  - cpu usage
  - Memory usage
  - disk usage
- Onstats
  - onstat –g glo
  - onstat –g rea
  - onstat –g act
- what else is dependent on your system

BASELINE INFO IS ESSENTIAL FOR PERFORMANCE TUNING

What is my application doing?

- Application development and DBA are different jobs
- Applications might be developed using tools that don't tell you the queries they are using
- What to do when the users are complaining?

SQL Tracing before 11.10

- Methods for examination
  - Set explain
  - Looking at database objects accessed by the SQL for inefficiency
- Lack of simplicity in the process
- Hard to build a repeatable process

SQLtrace

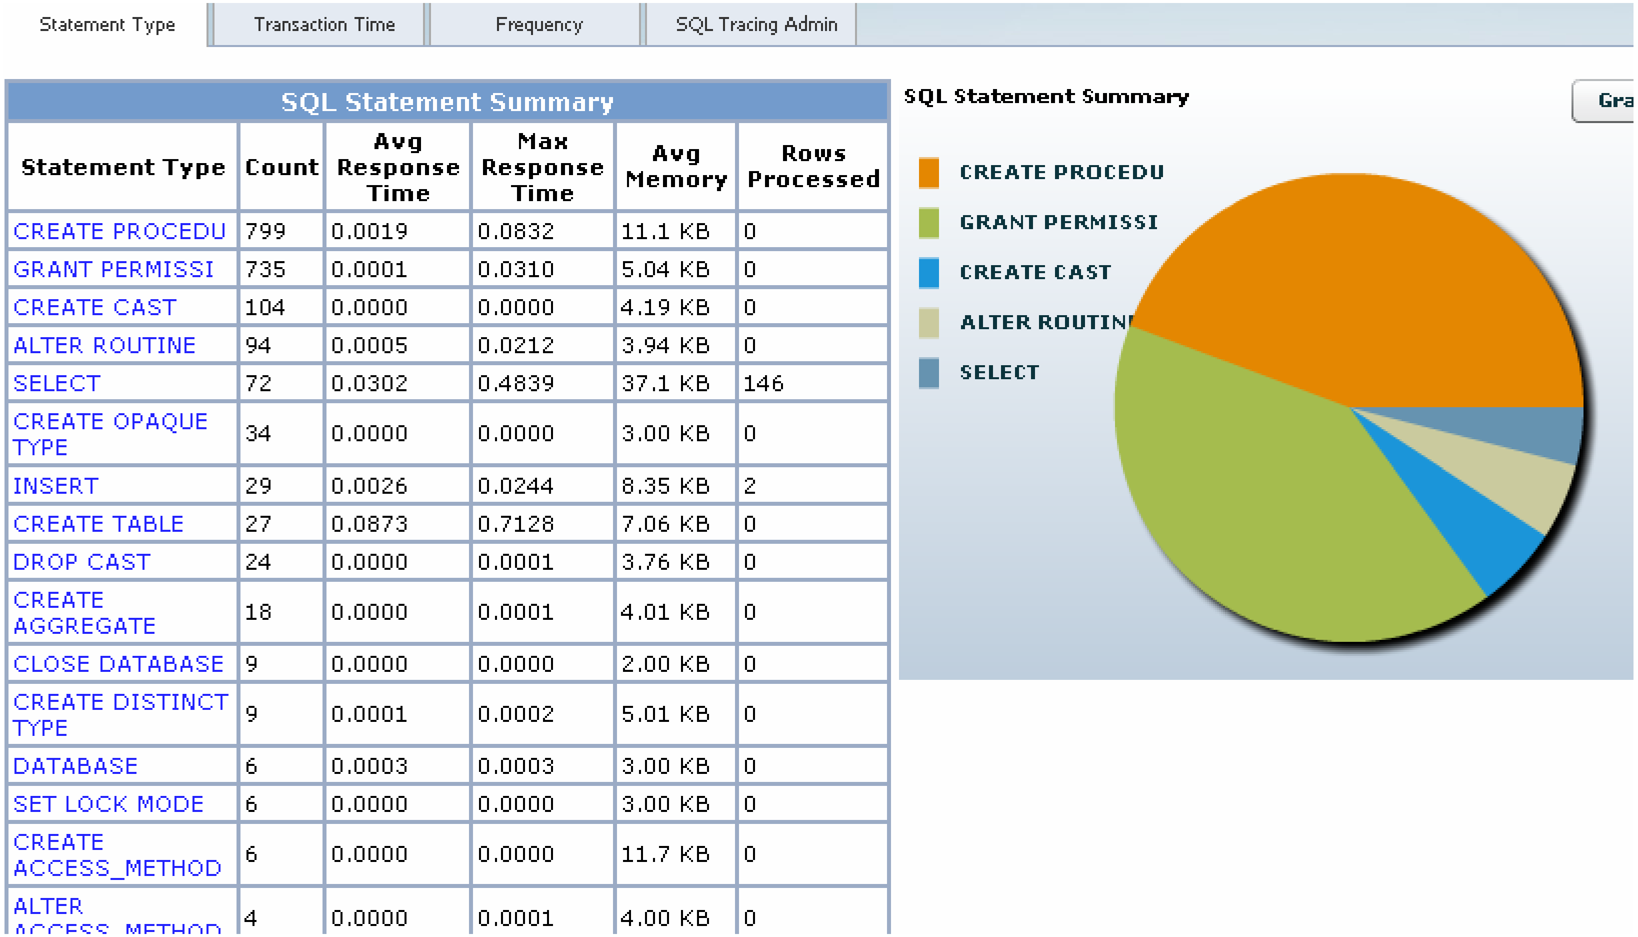

SQL Statement Summary

Statement Type	Count	Avg Response Time	Max Response Time	Avg Memory	Rows Processed
CREATE PROCEDU	799	0.0019	0.0832	11.1 KB	0
GRANT PERMISSI	735	0.0001	0.0310	5.04 KB	0
CREATE CAST	104	0.0000	0.0000	4.19 KB	0
ALTER ROUTINE	94	0.0005	0.0212	3.94 KB	0
SELECT	72	0.0302	0.4839	37.1 KB	146
CREATE OPAQUE TYPE	34	0.0000	0.0000	3.00 KB	0
INSERT	29	0.0026	0.0244	8.35 KB	2
CREATE TABLE	27	0.0873	0.7128	7.06 KB	0
DROP CAST	24	0.0000	0.0001	3.76 KB	0
CREATE AGGREGATE	18	0.0000	0.0001	4.01 KB	0
CLOSE DATABASE	9	0.0000	0.0000	2.00 KB	0
CREATE DISTINCT TYPE	9	0.0001	0.0002	5.01 KB	0
DATABASE	6	0.0003	0.0003	3.00 KB	0
SET LOCK MODE	6	0.0000	0.0000	3.00 KB	0
CREATE ACCESS_METHOD	6	0.0000	0.0000	11.7 KB	0
ALTER ACCESS_METHOD	4	0.0000	0.0001	4.00 KB	0

SQL Query Drill Down Feature

- Provide consolidated detail information about SQL statements through all layers
- •Information available through
  - Onstat
  - sysmaster database
  - OpenAdmin Tool for IDS (aka OAT)
- Dynamically configurable
- Global and User Tracing modes
- Find the slowest query with only three clicks of the mouse

Controlling Query Drill Down

- SQLTRACE level=low, ntraces=2000,size=1,mode=global
  - ONCONFIG variable SQLTRACELevel =[off,low,med,high]
  - Ntraces=[number of traces]
  - Size=[size of each trace buffer in KB]
  - Mode=[global|user]

- What to trace? (new in 11.50.xC3)
  - SET SQL TRACING USER ADD/REMOVE/LIST/CLEAR user_name
  - SET SQL TRACING DATABASE- ADD *database_name*
  - SET SQL SET SQL TRACING SESSION *current_session_id“ ON*

Controlling Query Drill Down

- Dynamically enable or modify SQL Tracing
  - Trace 2000 SQL statements
  - Trace 1024 bytes of data for each SQL statement

```
execute function task("SET SQL TRACING ON",2000,1);
```

- Turn off SQL Tracing

```
execute function task("SET SQL TRACING OFF");
```

- Suspend/ Resume SQL Tracing
  - Keeps definitions what to Trace

```
execute function task("SET SQL TRACING SUSPEND");
execute function task("SET SQL TRACING RESUME);
```

What can you do?

- Change the application
- If you can‘t, check the query plan
- Check update statistics currency
- Add additional indexes
- Use optimizer hints
- If you can‘t
  - Consider to use external directives

Index builds

- How is an index build
- How to improve performance?
  - Smaller indexes
    - DS_NONPDQ_MEMORY
  - Bigger indexes
    - Tweak sort parameters

How is an index build?

- Read through the data
  - Either single or fragmented
- Sort the data
  - Got either one set or a number of sets
  - Merge the data, if necessary

Do as much in memory as possible

Index build: DS_NONPDQ_MEMORY

- Tune DS_NONPDQ_QUERY_MEM
  - memory given to sorts without PDQ
  - Minimum Value=Default 128KB
  - Maximum Value 25% of DS_TOTAL MEMORY
  - Change while instance is online

```
onmode -wf DS_NONPDQ_QUERY_MEM=5000
```

How is an index build (using PDQ)?

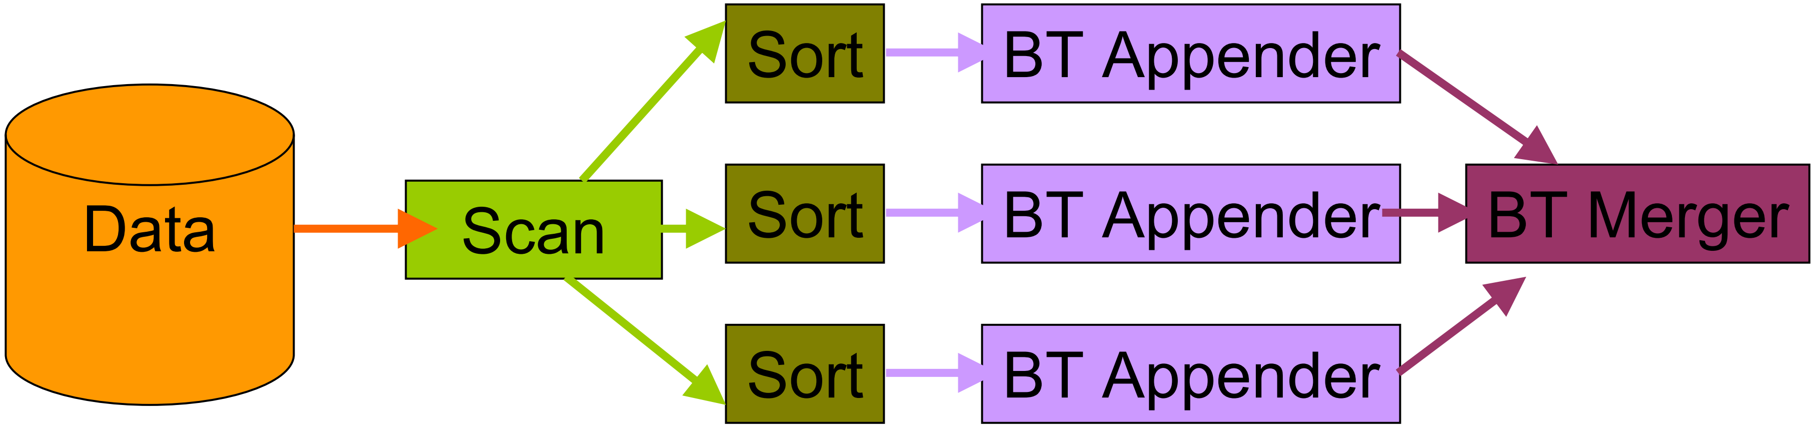

Large Index builds

- Ideal: sort in memory
- Enable PDQ
  - PDQPRIORITY
  - DS_TOTAL_MEMORY
  - DS_MAX_QUERIES
- PSORT_NPROCS

PDQ and Update statistics

- create table t1 (c1 int, c2 int, c3 char(20), c4 char(20))
- add two millions rows
- Without extra sort memory:
- table will be scanned 4 times
- Can be seen in set explain output:

```
PASS #1 c3
PASS #2 c4
PASS #3 c1
PASS #4 c2
```

PDQ with Update Statistics

- PDQ can change the number of table scans needed.
- Use set explain

```
Table:          ebach.t1
Mode:            HIGH
Number of Bins:      267        Bin size    10813
Sort data          166.1 MB
PDQ memory granted         174.4 MB
Estimated number of table scans 1
PASS #1 c1,c2,c3,c4
Light scans enabled
Scan 4 Sort 8 Build 3 Insert 0 Close 0 Total 15
Completed pass 1 in 0 minutes 15 seconds
```

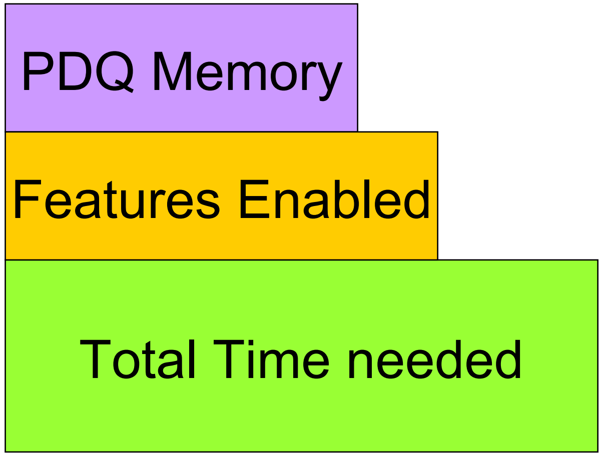

Top Tips

1. Create a simple but tested backup and restore strategy
2. Locking problems: Try “Last Committed” Read
3. DS_NONPDQ_QUERY_MEM
4. DIRECT_IO or raw devices
5. Make Informix directory a symbolic link
6. Use separate dbspaces root, log. logs, phys. log, temp, data/index, sysadmin)
7. Stay up to date - New versions are faster - New versions enable new possibilities
8. Consider OAT

A few new onstats

onstats

- New or improved in 11.x
  - Onstat –g cpu
  - Onstat –g ath
  - Onstat –g ckp

You cannot extrapolate from one measurement.
Always repeat a few times

Onstat –g cpu

- New with version 11.10
- tell you when which thread was scheduled
- Identify what is currently hanging
- Use along with onstat –g stk to see what threads are doing/waiting for

```
IBM Informix Dynamic Server Version 11.10.FC2W1   -- On-Line -- Up 3 days 05:51:27 -- 8649728 Kbytes

Thread CPU Info:

 tid     name              vp          Last Run          CPU Time    #scheds     status
 4      aio vp 0          36aio*      03/26 17:41:55     491.8493      74091     IO Idle
 5      msc vp 0          37msc*      03/26 17:41:59    8544.2635   10100529     IO Idle
 6      aio vp 1          38aio*      03/26 17:12:08       5.8684        426     IO Idle
 7      main_loop()       13cpu       03/26 17:41:59     244.5516     576872     sleeping secs: 1
 8      tlitcppoll        39tli*      03/26 17:41:59  280247.3047  138388591     running
15      flush_sub(0)      13cpu       03/26 17:41:59     121.2926    1401344     sleeping secs:
16      flush_sub(11)     7cpu        03/26 17:41:58      66.5788    1044132     IO Wai
26      kaio              1cpu*       03/26 17:41:59  141010.0865   44337017     running
52      kaio              3cpu*       03/26 17:41:59  135401.9964   40585737     IO Idle
80      kaio              7cpu*       03/26 17:41:59  115871.8548   28019443     running
138     aio vp 4          47aio*      03/23 11:51:36       0.3423         44     IO Idle
139     aio vp 5          48aio*      03/23 11:51:48       0.4359         74     IO Idle
140     aio vp 6          49aio*      03/23 11:52:11       0.7263         80     IO Idle
```

Onstat –g ath

- Now comes with more detailed information
- No longer simply „sleeping“

```
IBM Informix Dynamic Server Version 11.10.FC2W1   -- On-Line (CKPT REQ) -- Up 15 days 00:36:28 -- 197632 Kbytes
Blocked:CKPT

Threads:
 tid      tcb              rstcb            prty status              vp-class       name
 2        1113a1ad8        0                1    running               3lio*        lio vp 0
 3        1113c1c80        0                1    IO Idle               4pio*        pio vp 0
 4        1113e0c80        0                1    running               5aio*        aio vp 0
 5        1113ffc80        0                1    IO Idle               6msc*        msc vp 0
 6        11142ec80        0                1    IO Idle               7aio*        aio vp 1
 7        1113a1d28        1111c9028        1    sleeping secs: 1     11cpu         main_loop()
 40       1116ca348        1111ce930        1    sleeping forever      1cpu*        dbWorker2
 48       111f83b58        1111cc8d0        1    cond wait  bp_cond    9cpu         bf_priosweep()
 242      11224c028        0                1    IO Idle              18aio*        aio vp 2
 336      1126e18e8        1111d19c0        1    cond wait  cp         1cpu         sqlexec
 338      11505fbf8        1111e6db0        1    IO Wait               9cpu         sqlexec
 339      1126c8028        1111e4d50        1    IO Wait               1cpu         sqlexec
```

Onstat –g ckp

- Information about the last 15 checkpoints
- Check for what is triggering the checkpoint, duration, blocking time and flush rates to get an idea what the system is doing

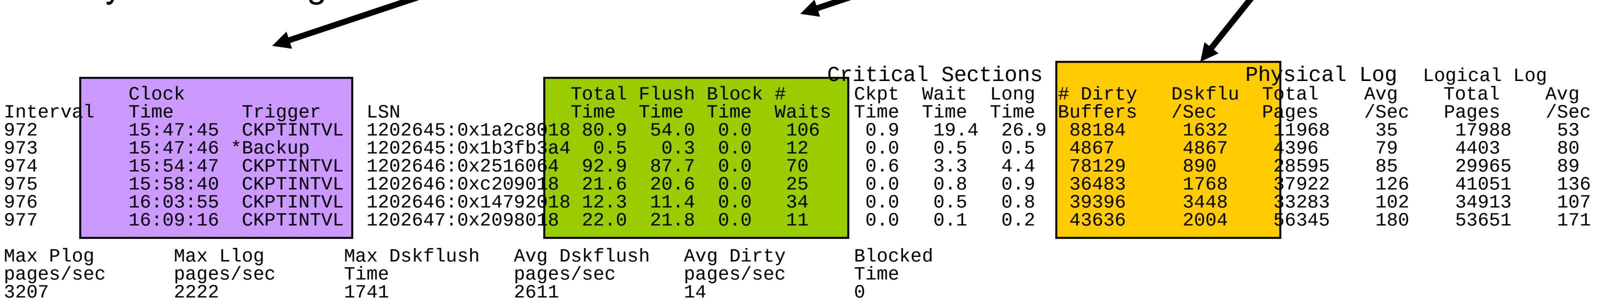

```
                                                                   Critical Sections                      Physical Log   Logical Log
                Clock                            Total Flush Block #     Ckpt  Wait  Long  # Dirty   Dskflu  Total    Avg     Total     Avg
Interval        Time      Trigger    LSN         Time  Time  Time  Waits Time  Time  Time  Buffers   /Sec    Pages    /Sec    Pages     /Sec
972             15:47:45  CKPTINTVL  1202645:0x1a2c8018 80.9  54.0  0.0   106   0.9   19.4  26.9  88184     1632    11968    35      17988     53
973             15:47:46 *Backup     1202645:0x1b3fb3a4  0.5   0.3  0.0   12    0.0   0.5   0.5   4867      4867    4396     79      4403      80
974             15:54:47  CKPTINTVL  1202646:0x2516064  92.9  87.7  0.0   70    0.6   3.3   4.4   78129     890     28595    85      29965     89
975             15:58:40  CKPTINTVL  1202646:0xc209018  21.6  20.6  0.0   25    0.0   0.8   0.9   36483     1768    37922    126     41051     136
976             16:03:55  CKPTINTVL  1202646:0x14792018 12.3  11.4  0.0   34    0.0   0.5   0.8   39396     3448    33283    102     34913     107
977             16:09:16  CKPTINTVL  1202647:0x2098018  22.0  21.8  0.0   11    0.0   0.1   0.2   43636     2004    56345    180     53651     171

Max Plog       Max Llog       Max Dskflush   Avg Dskflush   Avg Dirty      Blocked
pages/sec      pages/sec      Time           pages/sec      pages/sec      Time
3207           2222           1741           2611           14             0
```

Based on the current workload, the physical log might be too small to accommodate the time it takes to flush the buffer pool during checkpoint processing. The server might block transactions during checkpoints.
If the server blocks transactions, increase the physical log size to at least 11391264 KB.

Onstat –g osi

- Prints auf Machine Information

```
IBM Informix Dynamic Server Version 11.50.TC2

Machine Configuration....
OS Name                                 Windows
OS Release                              5.1 Service Pack 2
OS Node Name                            RUACH
OS Version                              Unknown
OS Machine                              x86
Number of processors                    1
Number of online processors             1
System memory page size                 4096 bytes
System memory                           2046 MB
System free memory                      733 MB
Number of open files per process        0
shmmax                                  <<  UnSupported  >>
shmmin                                  <<  UnSupported  >>
shmids                                  <<  UnSupported  >>
shmNumSegs                              <<  UnSupported  >>
```

Onstat –g ppf

- onstat –g ppf or sysmaster:sysptprof now tracks how many index deletes and inserts have occurred
- Example

```
SELECT TRIM(dbsname)||":"||TRIM(tabname) as index_name,
isreads, iswrites, isrewrites, isdeletes
FROM sysmaster:sysptprof
WHERE partnum in
(SELECT partn FROM dbs:sysfragments WHERE fragtype = 'I');
```

- Output

```
index_name i:ix1
isreads 932
iswrites 990
isrewrites 0
isdeletes 1
```

IDS on Windows

IDS on Windows: What's New

- New Platform Support
  - Vista / Server 2008
  - 64-bit Support
  - XP Embedded (planned)
- New Features
  - Multiple Installation Support
  - Local System user support
  - SHMNOACCESS
  - ISM GUI is back
  - Debug libraries included
- Application Development – Visual Studio Add-ins

SHMNOACCESS - Fighting DLL Hell

Problem

- IDS needs contiguous free address space for shared memory segments
- These need to fit around any DLL’s loaded by the process
- Utilities which attach to IDS shared memory like onstat, onmode, onbar may load DLL’s at different addresses
- In large shared memory configurations on 32-bit Windows this can cause these utilities to fail to attach to shared memory

Solution

- Use the **SHMNOACCESS** Onconfig parameter

SHMNOACCESS Onconfig parameter

- Used to specify addresses where shared memory segments should not be allocated
- In IDS 10.00.xC7 and later it worked for 1st Virtual segment
- In 11.50.TC2 and later it works for all segments
- Takes effect when the database server is shutdown and re-started
- 1 – 10 address ranges can be specified

- Example:
  - SHMNOACCESS 0x70000000-0x75000000, 0x7A000000-0x80000000

SHMNOACCESS Example

Consider these Onconfig settings…

```
SHMBASE 0xc000000L
SHMVIRTSIZE 204800
SHMNOACCESS 0x70000000-0x7FFFFFFF
BUFFERPOOL size=4K,buffers=575000,lrus=8,
         lru_min_dirty=50.000000,lru_max_dirty=60.000000
```

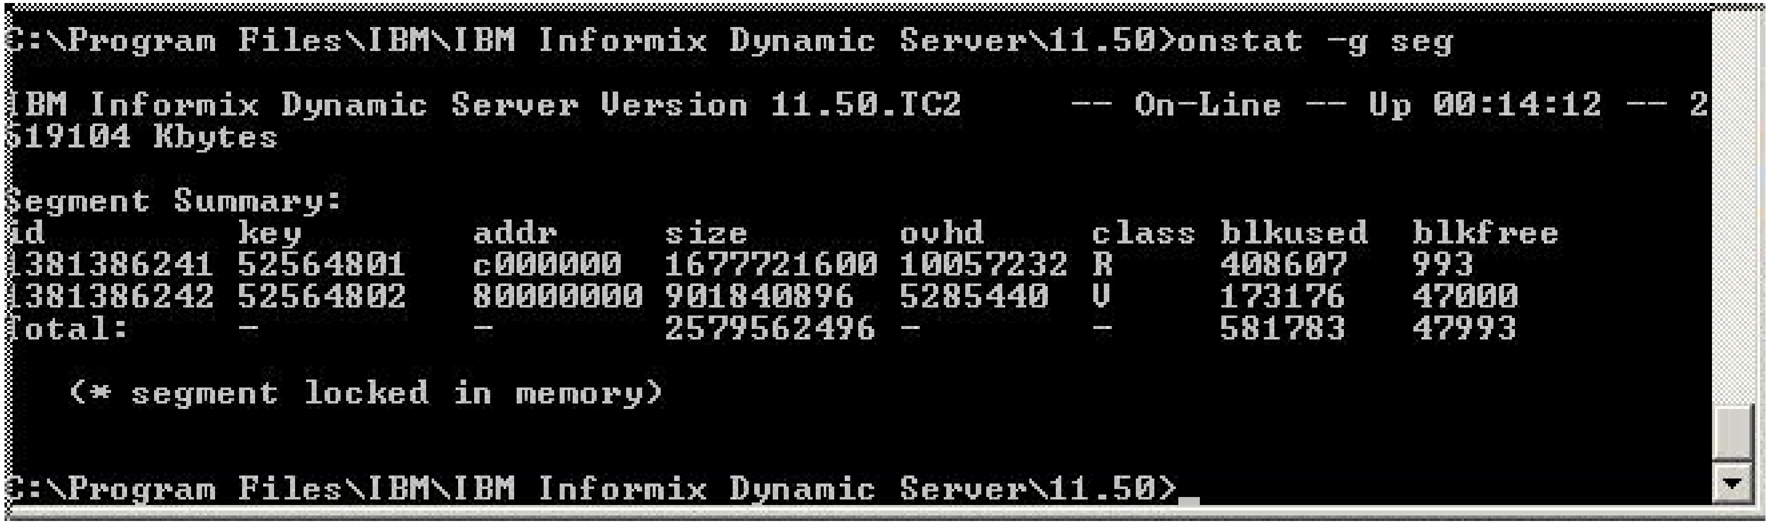

```
C:\Program Files\IBM\IBM Informix Dynamic Server\11.50>onstat -g seg

IBM Informix Dynamic Server Version 11.50.TC2     -- On-Line -- Up 00:14:12 -- 2
519104 Kbytes

Segment Summary:
id          key         addr      size        ovhd       class blkused  blkfree
1381386241  52564801    c000000   1677721600  10057232   R     408607   993
1381386242  52564802    80000000  901840896   5285440    V     173176   47000
Total:      -           -         2579562496  -          -     581783   47993

   (* segment locked in memory)


C:\Program Files\IBM\IBM Informix Dynamic Server\11.50>_
```

The Virtual segments starts safely at 2GB rather than in system DLL range

Windows Performance Tool Kit (WPT)

- Provided by Microsoft to analyze product performance issues
- Requires Windows Vista/Server 2008 Kernel
- Download: http://www.microsoft.com/whdc/system/sysperf/perftools.mspx
- Documentation: http://msdn.microsoft.com/en-us/library/cc305187.aspx
- General flow:
  - Use **xperf.exe** to enable tracing with the Event Tracing for Windows. interface.
  - Perform your set of operations.
  - Use xperf.exe to disable ETW interface. Trace data is saved in an ETL file.
  - Use the performance analyzer tool **xperfview.exe** to process and view ETLTrace files
- The ETL trace files can be copied to a different machine for analysis. All of the information required for performance analysis is saved in the ETL trace file.
- Interesting features in WPT: Symbol Decoding, Sample Profiling and Capturing Call-stacks on Kernel Events

WPT – Example - oninit.exe/oninit.pdb

- If IDS 11.50 is installed and running in C:\IDS1150\
- Make sure you have the correct **oninit.pdb** for oninit.exe in %INFORMIXDIR%\bin
  - Win64 – oninit.pdb is installed with the product
  - Win32 – oninit.pdb is supplied with the install media
- Set **_NT_SYMCACHE_PATH** and **_NT_SYMBOL_PATH**
  - SET _NT_SYMCACHE_PATH=
    %SystemDrive%\SymCache
  - SET _NT_SYMBOL_PATH=
    srv*%SystemDrive%\Symbols*http://msdl.microsoft.com/download/symbols;C:\IDS1150\bin

WPT – Example (Continued)

- Step #1: Start trace
  - **xperf -on latency -stackwalk Profile**
- Step #2: Run your test for performance analysis
- Step #3: Collect & stop trace
  - **xperf -d trace.etl**
  - By default, collects the trace information for all processes in the system
- Step #4: View trace (Windows Performance Analyzer)
  - **xperfview trace.etl**
  - GUI interface is easy to understand and use
  - Select the CPU time window and filter it for your process

WPT – Sample output

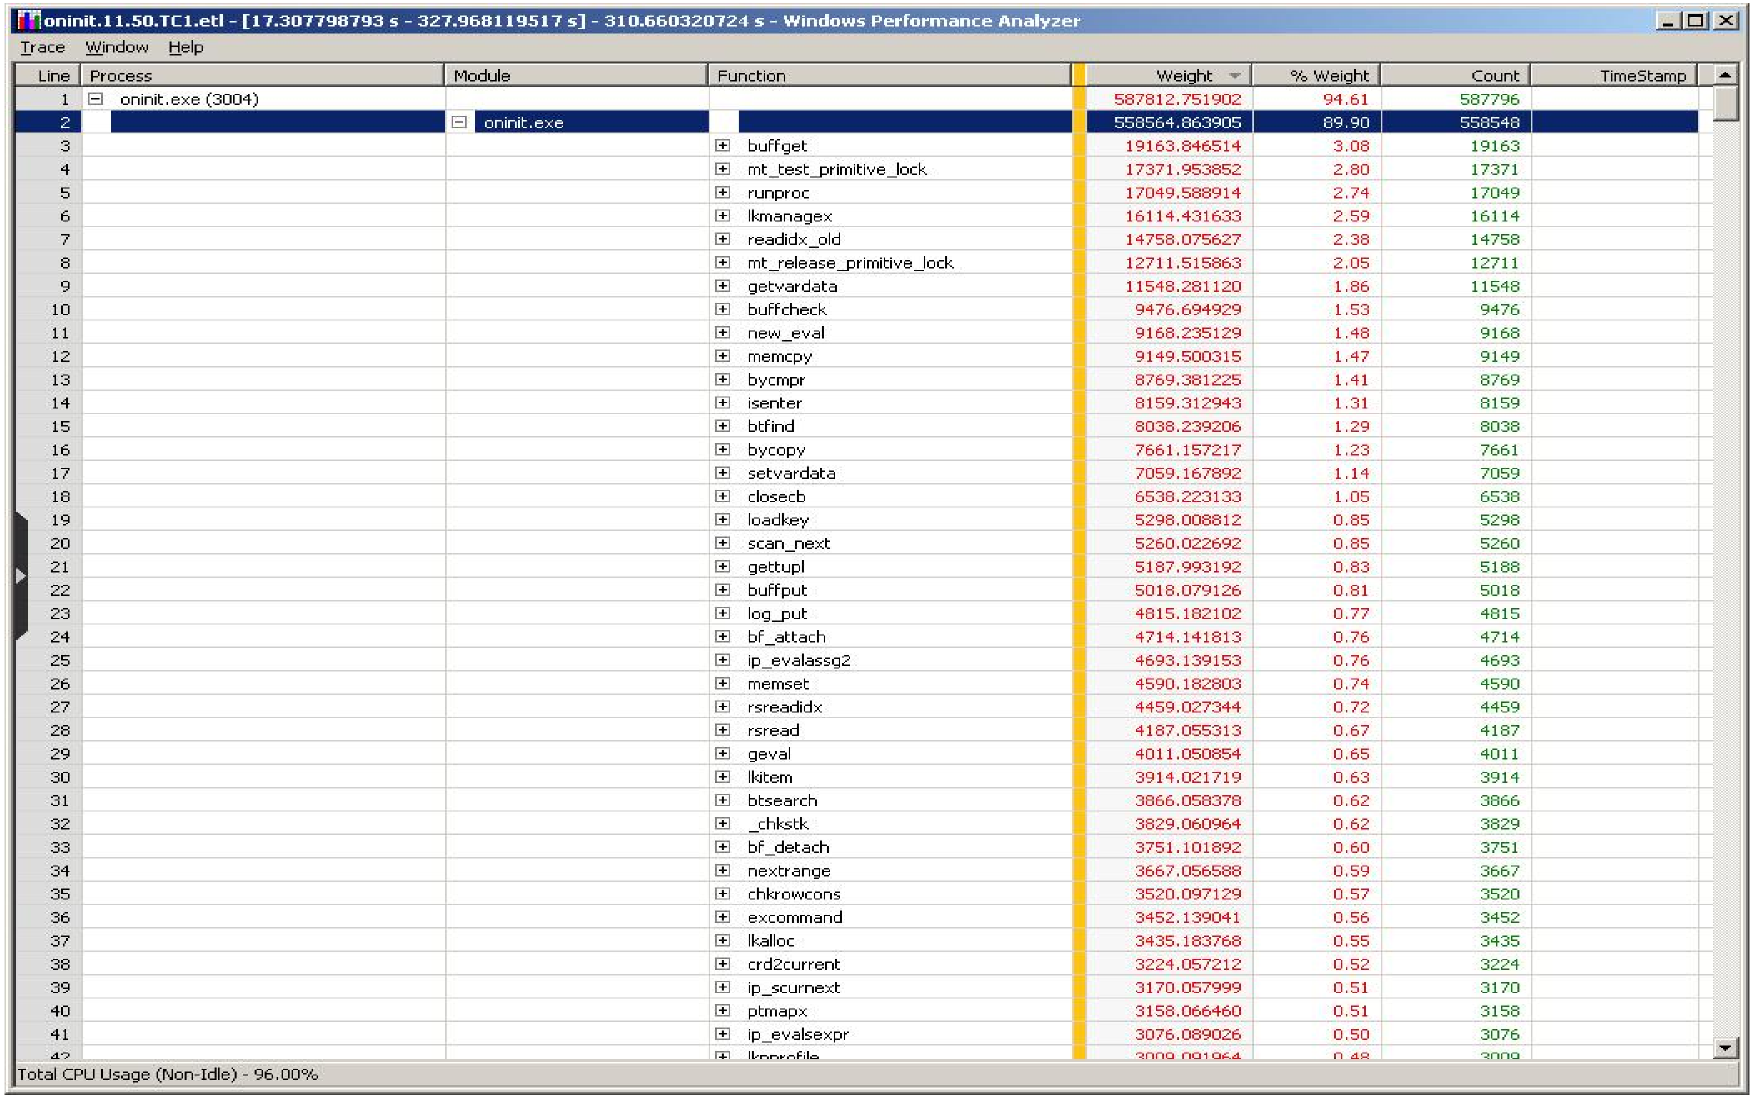

oninit.11.50.TC1.etl - [17.307798793 s - 327.968119517 s] - 310.660320724 s - Windows Performance Analyzer

Line	Process	Module	Function	Weight	% Weight	Count	TimeStamp
1	⊟ oninit.exe (3004)			587812.751902	94.61	587796	
2		⊟ oninit.exe		558564.863905	89.90	558548	
3			⊞ buffget	19163.846514	3.08	19163	
4			⊞ mt_test_primitive_lock	17371.953852	2.80	17371	
5			⊞ runproc	17049.588914	2.74	17049	
6			⊞ lkmanagex	16114.431633	2.59	16114	
7			⊞ readidx_old	14758.075627	2.38	14758	
8			⊞ mt_release_primitive_lock	12711.515863	2.05	12711	
9			⊞ getvardata	11548.281120	1.86	11548	
10			⊞ buffcheck	9476.694929	1.53	9476	
11			⊞ new_eval	9168.235129	1.48	9168	
12			⊞ memcpy	9149.500315	1.47	9149	
13			⊞ bycmpr	8769.381225	1.41	8769	
14			⊞ isenter	8159.312943	1.31	8159	
15			⊞ btfind	8038.239206	1.29	8038	
16			⊞ bycopy	7661.157217	1.23	7661	
17			⊞ setvardata	7059.167892	1.14	7059	
18			⊞ closecb	6538.223133	1.05	6538	
19			⊞ loadkey	5298.008812	0.85	5298	
20			⊞ scan_next	5260.022692	0.85	5260	
21			⊞ gettupl	5187.993192	0.83	5188	
22			⊞ buffput	5018.079126	0.81	5018	
23			⊞ log_put	4815.182102	0.77	4815	
24			⊞ bf_attach	4714.141813	0.76	4714	
25			⊞ ip_evalassg2	4693.139153	0.76	4693	
26			⊞ memset	4590.182803	0.74	4590	
27			⊞ rsreadidx	4459.027344	0.72	4459	
28			⊞ rsread	4187.055313	0.67	4187	
29			⊞ geval	4011.050854	0.65	4011	
30			⊞ lkitem	3914.021719	0.63	3914	
31			⊞ btsearch	3866.058378	0.62	3866	
32			⊞ _chkstk	3829.060964	0.62	3829	
33			⊞ bf_detach	3751.101892	0.60	3751	
34			⊞ nextrange	3667.056588	0.59	3667	
35			⊞ chkrowcons	3520.097129	0.57	3520	
36			⊞ excommand	3452.139041	0.56	3452	
37			⊞ lkalloc	3435.183768	0.55	3435	
38			⊞ crd2current	3224.057212	0.52	3224	
39			⊞ ip_scurnext	3170.057999	0.51	3170	
40			⊞ ptmapx	3158.066460	0.51	3158	
41			⊞ ip_evalsexpr	3076.089026	0.50	3076	
42			⊞ lkprofile	3009.091964	0.48	3009	

Total CPU Usage (Non-Idle) - 96.00%

Roadmap Info

Upgrading IDS and the Product Life Cycle

- ►In order to guarantee continued feature enhancements and functionality improvements, Development needs to focus on the latest key code streams
- ►Clients on Subscription & Support can upgrade for free* from IDS 9.x or 10.x to IDS 11.50
- ►IDS 7 clients must purchase a “trade-up” license top upgrade

IDS Version	EOM	EOS
7.31	Sept 2008	Sept 2009
9.40	Apr 2008	Apr 2009
10.00	Sept 2009	Sept 2010
11.10	Sept 2009	Sept 2012

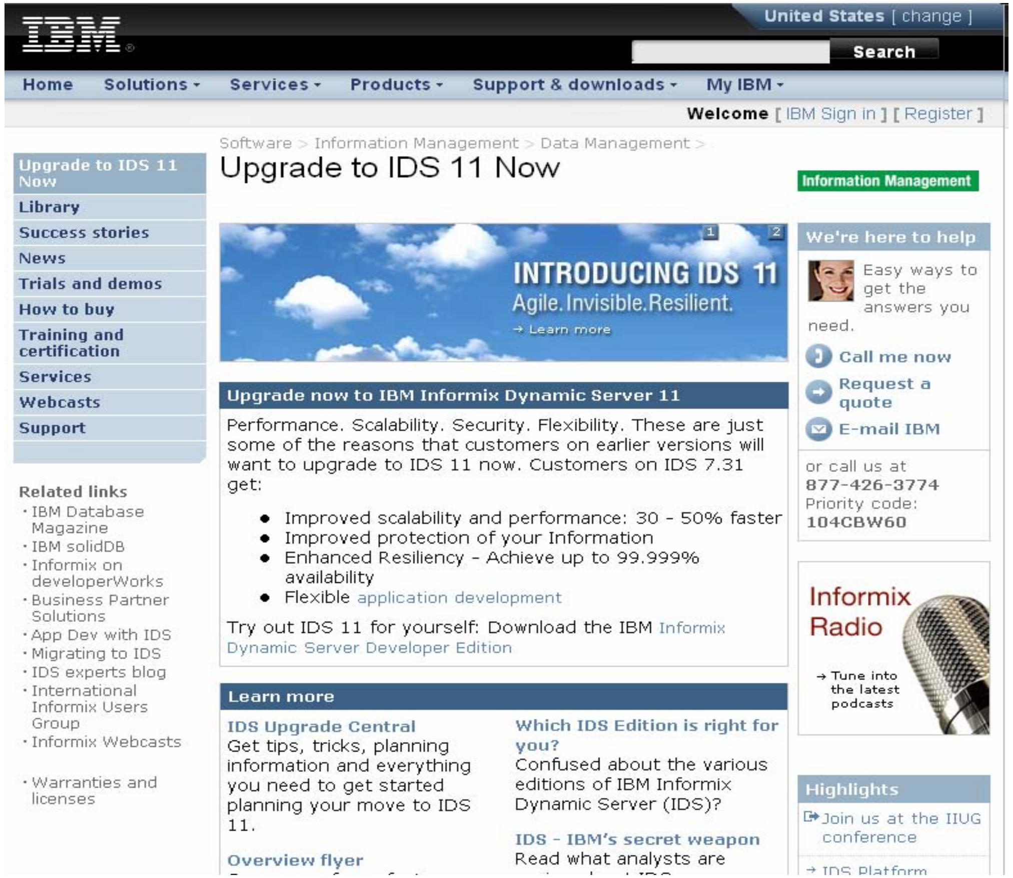

United States [change]

Search

Home | Solutions | Services | Products | Support & downloads | My IBM

Welcome [IBM Sign in] [Register]

- Upgrade to IDS 11 Now
- Library
- Success stories
- News
- Trials and demos
- How to buy
- Training and certification
- Services
- Webcasts
- Support

Related links

- IBM Database Magazine
- IBM solidDB
- Informix on developerWorks
- Business Partner Solutions
- App Dev with IDS
- Migrating to IDS
- IDS experts blog
- International Informix Users Group
- Informix Webcasts

- Warranties and licenses

Software > Information Management > Data Management >

Upgrade to IDS 11 Now

Information Management

Upgrade now to IBM Informix Dynamic Server 11

Performance. Scalability. Security. Flexibility. These are just some of the reasons that customers on earlier versions will want to upgrade to IDS 11 now. Customers on IDS 7.31 get:

- Improved scalability and performance: 30 - 50% faster
- Improved protection of your Information
- Enhanced Resiliency - Achieve up to 99.999% availability
- Flexible application development

Try out IDS 11 for yourself: Download the IBM Informix Dynamic Server Developer Edition

Learn more

IDS Upgrade Central
Get tips, tricks, planning information and everything you need to get started planning your move to IDS 11.

Overview flyer

Which IDS Edition is right for you?
Confused about the various editions of IBM Informix Dynamic Server (IDS)?

IDS - IBM's secret weapon
Read what analysts are

We're here to help

Easy ways to get the answers you need.

- Call me now
- Request a quote
- E-mail IBM

or call us at
877-426-3774
Priority code:
104CBW60

Highlights

- Join us at the IIUG conference
- IDS Platform

Product Roadmap

Product	Q1 2009	Q2 2009	Q3 2009	
10 fixpack	10.00.xC10		10.00.xC11	
11.10 fixpacks				11.10.xC4
11.50 fixpacks	11.50.xC4	11.50.xC5	11.50.xC6	11.50.xC7
10 PIDs	10.00.xC10W1	10.00.xC10W2	10.00.xC11W1	10.00.xC11W2
11.10 PIDs	11.10.xC3W1	11.10.xC3W2	11.10.xC3W3	11.10.xC3W4
11.50 PIDs	11.50.xC3W2 11.50.xC3W3	11.50.xC4W1 11.50.xC4W2	11.50.xC5W1 11.50.xC5W2	11.50.xC6W1 11.50.xC6W2
OAT	OAT		OAT	OAT
vNext	vNext Release Work – Target GA Mid 2010			

Best Thing Since Sliced Bread
… as Easy as a Toaster

Information Management
Informix Dynamic Server 11.5 Discovery Kit
Informix software
International Informix Users Group - www.iiug.org
Informix
Join the Revolution at www.ibm.com/informix

IDS 11 virtualization to simplify a complex world

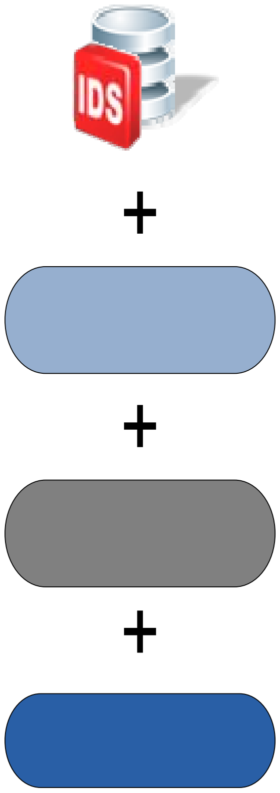

Optimize your IT investments, infrastructure and resources

Investments

Reduce power, cooling and real-estate expenses

Reduce hardware expenses through consolidation and improved utilization

Infrastructure

Simplify server configuration and deployment

Instant High Availability and Scalability

Enhance manageability and responsiveness

Improve server utilization and control

Improve desktop management and administration

Resources

Improve IT professional productivity

Reduce software and maintenance complexity

Available now:

- Developer Edition virtual appliance image for exploration and development
- Offer through ISVs and OEM partners
- Repackaging Guide
- Partner and ISV Assistance

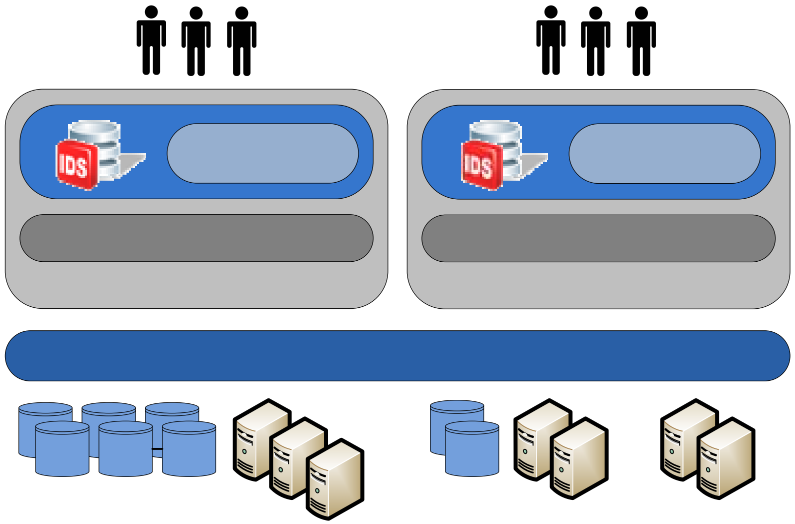

Informix Technical Roadmap

Unlocking the Business Value of Information for Competitive Advantage

IDS 9.3

- Updatable table statistics
- Improved logging algorithm
- Customer control over query plans used

IDS 9.4

Scale with no limits

- Max DB size increased From 4TB to 128 Petabytes
- Increased availability with ER and HDR
- Increased replication performance

IDS 10

Optimized OLTP Engine

- Query performance improvement
- Online index build
- Multiple page size support for better space utilization
- Enhanced buffer management
- Increased security with column encryption
- Disaster recovery with table level restore

IDS 11

Scale-out at lower costs

- Multi-node active cluster for high-availability (MACH) 11 with multiple remote servers and shared disk cluster
- Open source tool for administration, SQL Admin API to embed admin tasks
- Improved checkpoint performance
- Secured data encryption, LBAC, Common Criteria certification
- Enhanced application development for SOA and XML
- Web Feature Service support for geospatial applications

Vnext*

Application Integration

- Preconfigured installs
- Automatic storage provisioning
- Fine-grained Auditing
- Heterogeneous Replication
- In-memory Cache Connect
- Google KML
- External table support
- Performance/Index advisor
- Conversion of ASCII DBs to UTF-8
- Storage Optimization
- Analytics

and more…

2001
Informix acquisition

2003

2005

2007-2008

Informix
Invisible, Agile, Resilient

*Features subject to change

Questions

THANK YOU

Dr. Elisabeth Bach

IBM
elisabeth.bach@de.ibm.com