



White Star Software
WSS.COM

Maximizing Network Throughput

Paul Koufalis – White Star Software

Why listen to me?

pk@wss.com

Progress DBA and UNIX admin since 1994

Expert consulting related to technical aspects of Progress and OpenEdge


Wide range of experience

Small 10 person offices to 2500+ concurrent users

AIX, HPUX, Linux, Windows...if Progress runs on it, I've worked on it

Single malts and American bourbons





White Star Software
wss.com

- The oldest and most respected independent DBA consulting firm in the world
- Four of the world's top OpenEdge DBAs
- Author of ProTop, the #1 FREE OpenEdge Database Monitoring Tool
 - <http://dashboard.dbappraise.com>

Agenda

- a little theory
- tweakables
- stupid programming tricks
- numbers

a little bit of theory

- Two basic database connection types
 - Shared memory or self-service
 - Client server
- Shared memory is always going to be faster
 - Self-service clients reach into shared memory and get the data themselves
 - C/S clients have to ask a proxy process to do this for them
 - AND transport the data across a slow connection
- Most legacy Progress code was written back in the ChUI days
 - All shared memory
 - Doesn't necessarily port well to C/S

the (simplified) process

- Progress 4GL client connects to the database broker through a service port (`_mprosrv`)
- The broker assigns the client to a new or existing server (`_mprosrv -m1`)
- Client submits query to server
- Server executes query
 - A server is a fancy shared memory client
- Server returns data to client
- Client does post-processing on data
 - Ex.: Filtering and sorting
- Data becomes available to program/user

the 4gl server process

- OK – a little *more* than just a fancy shared memory client
- Single-threaded process
- Uses a polling mechanism to determine which connected clients need servicing
- Uses a round-robin allocation algorithm to determine which client to service
- Note: `_sqlsrv2` is multi-threaded



-prefetch*



-Mm
-Mm
-Mm
-Mm
-Mm
-Mm



tweakable parameters

MTU



new –prefetch* parameters

- No-lock queries
- Forward only or scrolling
- 10.2B06 + and 11.1+

- -prefetchPriority: server defers poll for other requests while filling message
 - Current suggested value 100 records added before next poll
- -prefetchDelay: Fills first message. By default first message contains one record
 - In theory this is better. In practice the difference is not significant
- -prefetchNumRecs : How many records are stuffed in a message
 - 100 records is a good start (default is 16)
- -prefetchFactor: How full (%-wise) to fill a message
 - 90-100%

message buffer size (-Mm)

- Default 1024 is too small
- Max 32600 doesn't seem to be warranted
- 8192 is a nice sweet spot

- Pre-11.6: clients and databases all had to have the same -Mm
- 11.6+: Client uses -Mm from database server: i.e. no more client-side -Mm required

maximum transmission unit (MTU)

- A network parameter set at the NIC level
- Enable on the routing infrastructure
- Default is 1500
- “Jumbo Frames” is 9000 bytes

- The advantage lies in the relative size of network header data
 - 1500 byte MTU: 1460 byte payload / 1538 byte total = 95% efficient
 - 9000 byte MTU: 8960 byte payload / 9038 byte total = 99% efficient

- Use ping to test
 - Windows: `ping -l 8900 -f <ip address>`
 - Linux: `ping -s 8900 -M do <ip address>`

Mm – MTU interaction

- Progress generates a message of size $-Mm$
- The message is submitted to the TCP/IP stack
- Stack breaks the message up into MTU-sized bytes to ship across the network

- Even if the originating NIC has an MTU of 9000, every switch and router between the client and server must support jumbo frames
 - Otherwise the jumbo frame will be broken down into 1500 byte (or whatever MTU-size) packets

other parameters

- -Mi / -Ma: min/max users per server
- The more users per server, the less time the server can dedicate to any one user
- -Mi 1 -Ma 5 is a good start
- -Mi 1 -Ma 1 if you have heavy duty users
 - AppServer agents

- Note: _sqlsrv2 is multi-threaded

stupid programming tricks

- Field lists: fields that are sent to the client
- Why send the whole record if you don't need it

`for each order fields (order-num terms) no-lock:`

`for each order except (instructions) no-lock:`

- No-lock is the most important part of this
- Without “no-lock” the server returns one record per message

stupid programming tricks

- Thank <*insert favourite deity*> Progress doesn't send the LOB across the wire unless you ask for it
- The LOB field is really a separate entity to the record
- The *real* record only contains a pointer to the LOB
- The LOB may be in another storage area (and should be)
- When you access the LOB, the client requests it from the server



Monitoring Server Messages

- `_ActServer VST`
 - Key: `_server-ID`
 - Interesting fields: `_Server-ByteSent`, `_ServerMsgSent`, `_ServerRecSent`
 - Calculate send size (Bytes sent / messages sent) and compare to `-Mm`
- Bug: When a record is larger than `-Mm`, only the first msg is counted
 - I.e. if you send a 4K record and `-Mm` is 1024, only 1 msg and 1024 bytes sent recorded
 - Blobs sent in 32,000 byte chunks – each chunk increments `msgSent` by 1 and `byteSent` by `-Mm`
- ProTop Free (<http://dashboard.dbappraise.com>)

Server Activity																
Srv	Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent	v	MB Rcvd	RcvdSz	SendSz
9999	Total	0	1	3	0	0	42	0.00	0	42	0.00	0.04		0.00	116	1020
1	Auto	3000	1	1	0	0	42	0.00	0	42	0.00	0.04		0.00	116	1020
4	Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00		0.00	0	0
3	Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00		0.00	0	0
2	Auto	3001	0	1	0	0	0	0.00	0	0	0.00	0.00		0.00	0	0

Critical Resources

Name	HB Age	Bkp Age	Alert Status	Cnx	Blocked	Trx	Cur Lk	BI Cls	Hit%	Log Rd	OS Rd	LatchTMO	io Resp
stl.stl_db37	11:25:27	181d 23:44	0	3	0	0	0	0	100.00%	0	0	0	0.08
stl.stl_db6	01:30:22	08:17	0	9	0	0	0	0	99.95%	1,871	1	0	9.27
demo49.db_2601	00:00:28	1d 11:42	1	502	0	41	210	0	99.96%	42,903	18	0	4.25
demo49.db_2602	00:00:25	1d 12:06	1	30	0	3	27	0	99.96%	2,760	1	0	6.24
travelco.db_2706	00:00:04	2d 08:28	2	137	0	62	7	0	100.00%	137,940	3	0	0.19

Resources

Name	HB Age	Bkp Age	Alert Status	Cnx	Blocked	Trx	Cur Lk	BI Cls	Hit%	Log Rd	OS Rd	LatchTMO	io Resp
+	abc	00:03:11	12										
+	abcorp	16d 08:25:38	0										
+	bang	00:02:14	0										
+	bib	00:01:11	1										
+	circle	00:01:13	0										
+	contractor	00:00:15	27										
+	dbappraise	00:00:14	0										
+	demo124		0										
+	demo18	00:03:13	0										
+	demo26		0										
+	demo27		0										
+	demo30		0										
+	demo49	00:01:11	2										
+	demo61	00:00:13	1										
+	epic	00:00:12	0										
+	fr1	00:03:12	0										
+	ink	00:00:58	3										
+	ironwork		0										
+	medco	00:01:14	4										
+	mtl	00:00:04	0										
+	plumbco	00:01:13	0										

Critical Alerts

2016/10/11		
04:11:08	demo49.db_2602	(435) SYSTEM ERROR: Ikrels record 2081853 not locked.
02:52:34	travelco.db_2706	(26) ** Array subscript 0 is out of range.
02:37:52	demo49.db_2601	(435) SYSTEM ERROR: Ikrels record 121876435 not locked.
01:43:53	travelco.db_2706	(26) ** Array subscript 84 is out of range.

Alerts Feed

2016/10/11		
04:11:08	demo49.db_2602	(435) SYSTEM ERROR: Ikrels record 2081853 not locked.
04:10:52	tqcan.db5	Usernum 360 terminated abnormally. (794)
04:10:51	tqcan.db6	Usernum 387 terminated abnormally. (794)
04:10:50	tqcan.db4	Usernum 390 terminated abnormally. (794)
04:08:48	contractor.db_3385	(43) ** Impossible de trouver ou d'ouvrir le fichier \...
04:08:47	contractor.db_3384	(43) ** Impossible de trouver ou d'ouvrir le fichier \...
04:08:46	contractor.db_3383	(43) ** Impossible de trouver ou d'ouvrir le fichier \...
04:08:45	contractor.db_3382	(43) ** Impossible de trouver ou d'ouvrir le fichier \...
04:08:44	contractor.db_3381	(43) ** Impossible de trouver ou d'ouvrir le fichier \...
04:08:41	contractor.db_3380	(43) ** Impossible de trouver ou d'ouvrir le fichier \...
04:08:40	contractor.db_3379	(43) ** Impossible de trouver ou d'ouvrir le fichier \...
04:08:38	contractor.db_3378	(43) ** Impossible de trouver ou d'ouvrir le fichier \...
04:08:36	contractor.db_3377	(43) ** Impossible de trouver ou d'ouvrir le fichier \...
04:07:36	stl.stl_db12	(794) Usernum 861 terminated abnormally.
04:07:10	stl.stl_db11	(794) Usernum 885 terminated abnormally.
04:07:09	stl.stl_db3	(794) Usernum 940 terminated abnormally.
04:06:29	abc.abc_db2	pctAlloc Space used 91 > 85
04:06:28	abc.abc_db2	pctAlloc Space used 93 > 85
04:06:27	abc.abc_db2	pctAlloc Space used 95 > 85
04:06:26	abc.abc_db2	pctAlloc Space used 99 > 85
04:06:22	abc.abc_db2	pctAlloc Space used 101 > 85
04:06:20	abc.abc_db2	pctAlloc Space used 111 > 85
04:06:19	abc.abc_db2	pctAlloc Space used 118 > 85
04:06:17	abc.abc_db2	pctAlloc Space used 134 > 85
04:06:16	abc.abc_db2	excessIdx 0 < 10
04:06:15	abc.abc_db2	cnxPerServ 200 > 5
04:06:14	abc.abc_db2	con_WDOG No WDOG running
04:05:14	medco.db_2500	OSRd OS reads/sec 16915 > 15000
04:02:06	tqcan.db3	ai_Locked Locked AI Extents 2 > 1
03:58:03	xyzcorp.naprod-repl	LatchTMO Latch timeouts/sec 28831 > 15000
03:39:17	bang.bang_db1	(794) Usernum 98 terminated abnormally.
03:39:17	bib.bib_db7	Full backup successfully completed. (1364)
03:39:16	bang.bang_db2	(794) Usernum 103 terminated abnormally.
03:39:08	bib.bib_db7	Full backup started. (1362)
03:39:04	bib.bib_db6	Full backup successfully completed. (1364)
03:37:57	bib.bib_db6	Full backup started. (1362)
03:37:54	bib.bib_db5	Full backup successfully completed. (1364)
03:37:53	bib.bib_db5	Full backup started. (1362)



...or Focus on One DB


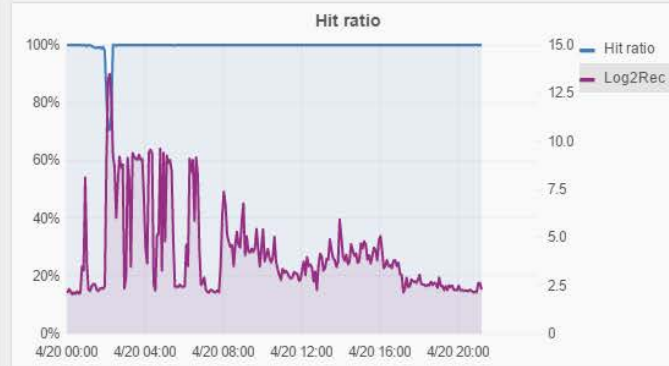
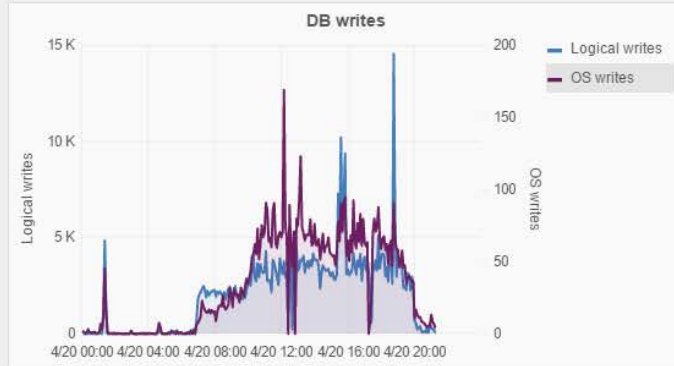

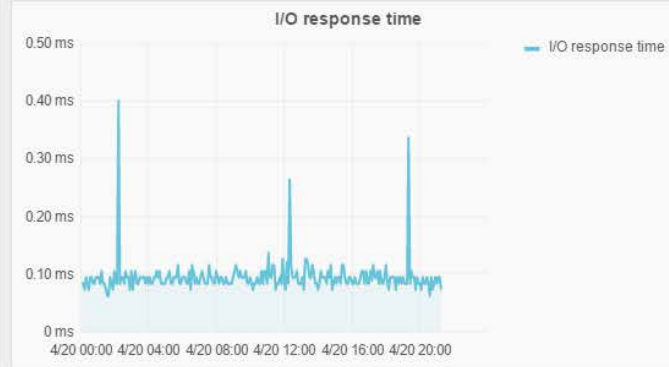
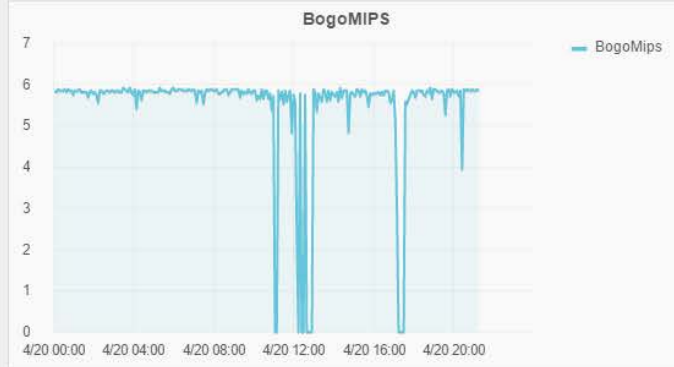
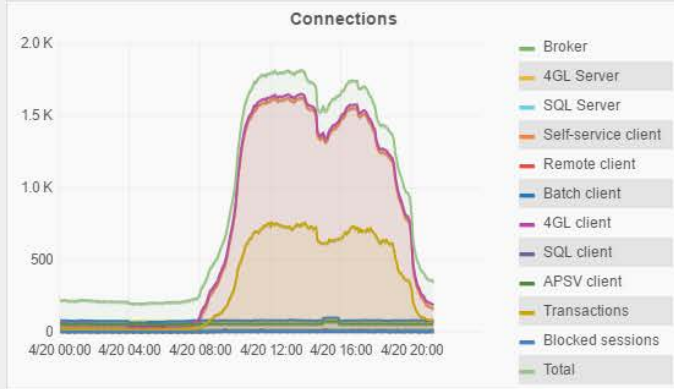
Critical Alerts			🔄 ↗ ^
2016/04/19			
08:36:52	xyzcorp.naprod	(14675) SYSTEM ERROR: Attempt to define too many indexes ...	

Alerts Feed			🔄 ↗ ^
2016/04/20			
20:10:25	xyzcorp.naprod	otrx Old Transaction 7426 > 7200	
18:55:21	xyzcorp.naprod	con_Block Blocked sessions 3 > 2	
18:43:05	xyzcorp.naprod	(439) ** Save file named core for analysis by Progress ...	
18:42:55	xyzcorp.naprod	(49) SYSTEM ERROR: Memory violation.	
17:25:21	xyzcorp.naprod	lkTableSize 2000005 > 1000000	
17:25:21	xyzcorp.naprod	biSize 18496.22 > 8192	
17:20:17	xyzcorp.naprod	lkTableSize 2000005 > 1000000	
17:20:17	xyzcorp.naprod	biSize 18496.22 > 8192	
17:17:37	xyzcorp.naprod	lkTableSize 2000005 > 1000000	
17:17:34	xyzcorp.naprod	biSize 18496.22 > 8192	
17:14:43	xyzcorp.naprod	lkTableSize 2000005 > 1000000	
17:14:40	xyzcorp.naprod	biSize 18496.22 > 8192	

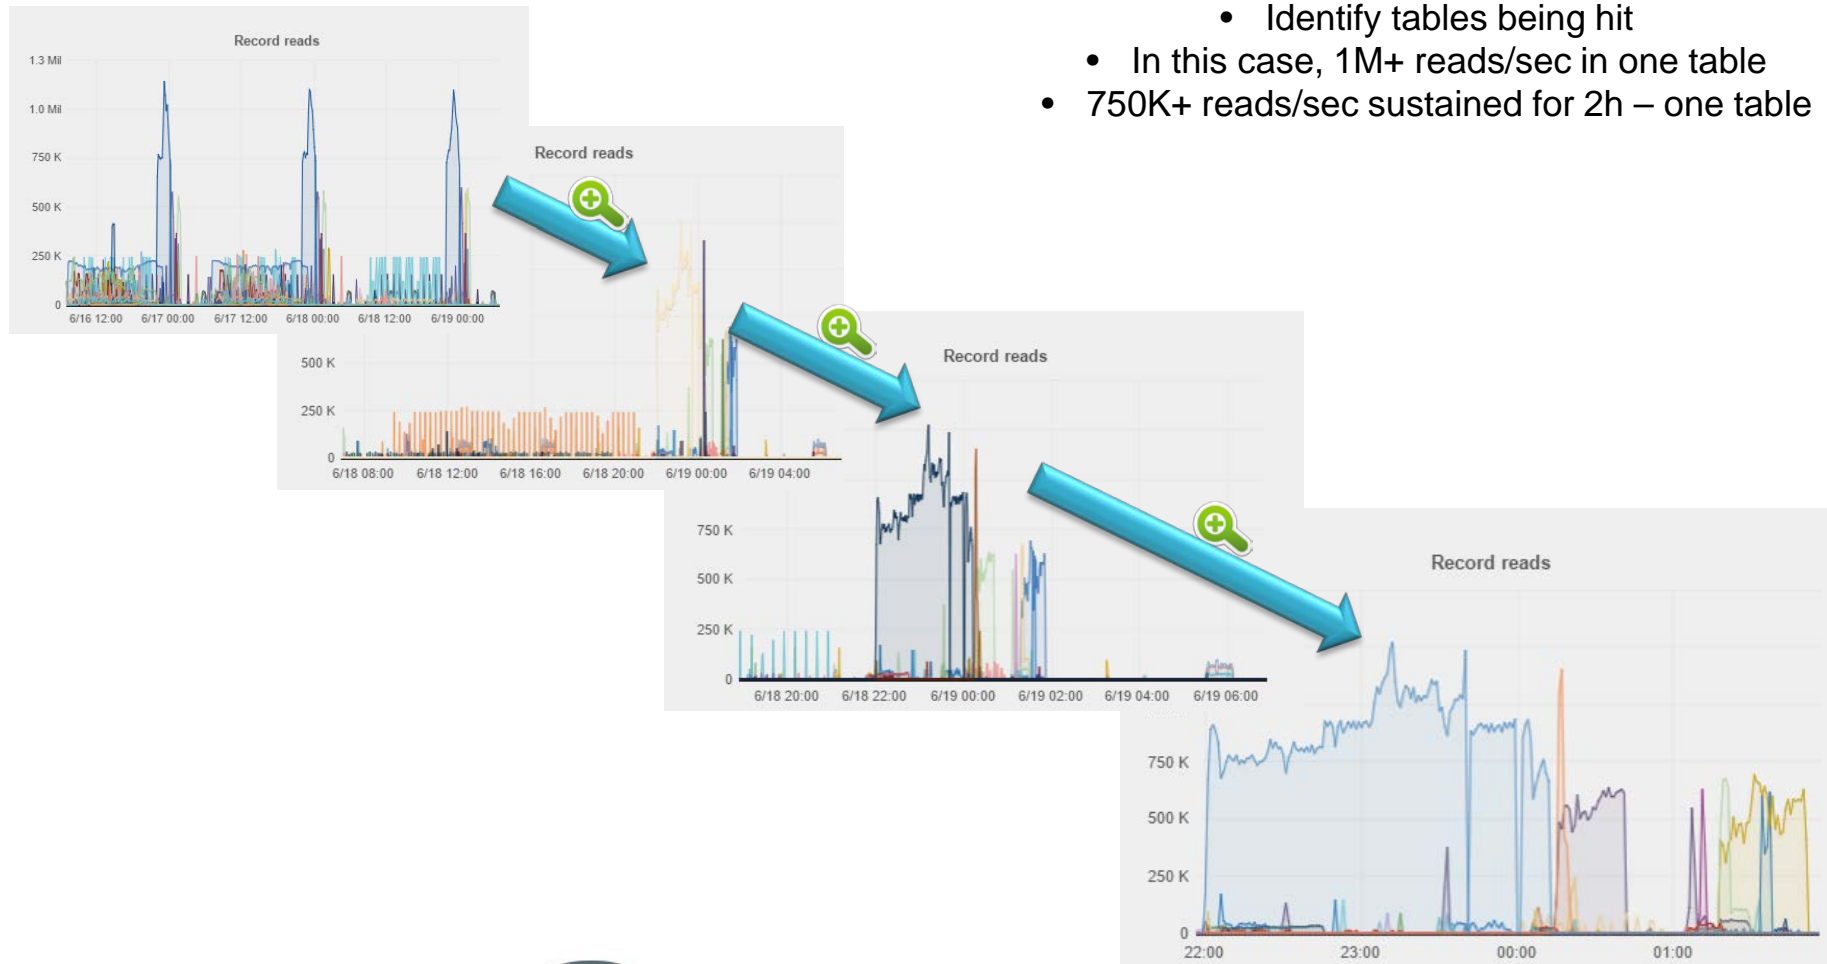




Last full backup	AI Status	Current AI	Full AI extents	Empty AI	Curr. AI seq	Lock HWM	Lock table size	Avail. TRS	Avail. IRS	DB Size	Version
20 hour	Enabled	5_of_8	0	7	18093	1.87 Mil	N/A	291	498	838 GiB	11.3.3_64bit



Narrowing Down the Problem



Local UI for Real-Time Analysis



xyzzz Auto Interval Rate JSON 86214 0 1.014
xyzzz 0 0

ProTop Version 3.3kx
/db/xyzzz

2016/04/21 10:19:36
demo.dbappraise.com

Hit%	100.00	Commits:	112	New RM:	89	Oldest TRX:	00:49:49	Connections:	1,364
Log Reads:	1,344,284	Undos:	32	From RM:	89	Curr BIClstr:	33,902	Brokers:	10
OS Reads:	53	Lock Tbl HWM:	1,869,296	From Free:	0	Oldest BIClstr:	33,883	4gl Servers:	72
Rec Reads:	362,429	Curr # Locks:	912	Examined:	90	Num BIClstrs:	19	SQL Servers:	14
LogRd/RecRd:	3.71	Modified BuFs:	9,283	Front2Bk:	0	BI MB Used:	1,216	4gl Clients:	1,200
Log Writes:	872	IO Response:	0.08	Remove Lk:	32	Curr AI Extent:	2 of 8	SQL Clients:	6
OS Writes:	21	BogoMIPS:	5.15			Curr AI Seq#:	18,146	App Server:	53
Rec Creates:	87	BogoMIP%:	96.91			Empty AI Exts:	7	Web Speed:	0
Rec Updates:	167					Full AI Exts:	0	BIW:	1
Rec Deletes:	10	Notes:	1,675	1,675		Locked AI Exts:	0	AIW:	1
Rec Locks:	1,078	BIW/AIW Write%:	38	97		APW Writes:	21	APWs:	4
Rec Waits:	0	Writes to Log:	15	10		APW Write%:	100	WDOG:	1
Idx Blk Spl:	0	BIW/AIW Writes:	6	10		BuFs Scanned:	17,064	Local:	1,175
Resrc Waits:	3	Partial Buf Wr:	5	0		APW Scan Wrts:	15	Remote:	6
Latch Waits:	80	Busy Buf Waits:	0	0		APW Q Wrts:	0	Batch:	78
		Empty Buf Wts:	0	0		Chkpt Q Wrts:	6	TRX:	472
						Flushed BuFs:	0	Blocked:	1

Replication Agents

Agent	Meth	Remote Host	Port	Stat	Comm	Blks Sent	Blks ACK	Last Block	TRX Behind	Remote DB Name
agent2	Async	10.1.11.16	4387	3049	1	10	0	04/21/16 10:19:36	3	/db/trax/xus61r2

Table Activity

Tbl #	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS Read
930	20	wm-pick	60	441347	0.53	0	235765	1	0	0
645	152	so-pack-log	469	34324829	0.00	1	33647	1	0	0
214	174	customer	63	159037	0.20	0	32490	0	0	0
670	112	so-trans	93793	18779417	0.00	0	12161	7	0	12
450	174	loc-group	35	25	453.62	0	11340	0	0	0

User IO Activity

Usr#	Tenant Name	PID	Flags	Blk Ac v	OS Rd	OS Wr	Hit%	Rec Lck	Rc Wts	Line#	Program Name
794	0 mlwt3	38006	S4B*	498037	21	0	100.00%	383	0	30995	wm/printpa.p
1483	0 xpacrmn	30205	S4	244157	0	0	100.00%	0	0	412	so/bormonit1.p
1803	0 xgtjavar	31928	S4	32664	1	0	100.00%	1	0	9508	ar/csaperf.p

enough already – you promised some numbers!

A little info about the test DB:

- Order table: 61K records (142K fragments) w/ mean size 5500 bytes
 - Most of it in the instructions field
- Customer table: 83 records w/ mean size 174 bytes
 - Field custPDF is a 4.7 MB blob (a pdf)

- `_mprosrv sports -S 5000 -B 250000 -Mi 1 -Ma 15 -Mn 2 -Mpb 1`
 - All users on the same server
 - -B big enough so that all data is in memory – i.e. disk does not come into play

- 10 concurrent `_progres` connections running “for each customer” 10K times
- Our goal is to minimize network messages

- Using ProTop Free to monitor server activity (<http://dashboard.dbappraise.com>)

Server Activity																
Srv	Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent	v	MB Rcvd	RcvdSz	SendSz
9999	Total	0	1	3	0	0	42	0.00	0	42	0.00	0.04		0.00	116	1020
1	Auto	3000	1	1	0	0	42	0.00	0	42	0.00	0.04		0.00	116	1020
4	Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00		0.00	0	0
3	Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00		0.00	0	0
2	Auto	3001	0	1	0	0	0	0.00	0	0	0.00	0.00		0.00	0	0

baseline

- for each customer: 3m30s

Table Activity												
Tbl#	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS Read		
√	2	8 Customer	8	83	471.40	0	39126	0	0	0		
	1	7 Invoice		147	0.00	0	0	0	0	0		
	3	7 Item		55	0.00	0	0	0	0	0		
	4	8 Order	5051	61297	0.00	0	0	0	0	0		
	5	8 Order-Line	7	415147	0.00	0	0	0	0	0		
	6	7 Salesrep		9	0.00	0	0	0	0	0		
	7	7 State		51	0.00	0	0	0	0	0		
	8	7 Local-Default		10	0.00	0	0	0	0	0		

Server Activity															
Srv	Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent v	MB Rcvd	RcvdSz	SendSz
√9999	Total	0	11	30	39595	0	79190	0.00	39124	39595	0.99	10.39	5.81	77	275
1	Auto	3000	11	15	39595	0	79190	0.00	39124	39595	0.99	10.39	5.81	77	275
2	Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0
0	Login	5000	0	15	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0

- for each customer no-lock: 2 min

Table Activity												
Tbl#	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS Read		
√	2	8 Customer	8	83	840.09	0	69727	0	0	0		
	1	7 Invoice		147	0.00	0	0	0	0	0		
	3	7 Item		55	0.00	0	0	0	0	0		
	4	8 Order	5051	61297	0.00	0	0	0	0	0		
	5	8 Order-Line	7	415147	0.00	0	0	0	0	0		
	6	7 Salesrep		9	0.00	0	0	0	0	0		
	7	7 State		51	0.00	0	0	0	0	0		
	8	7 Local-Default		10	0.00	0	0	0	0	0		

Server Activity															
Srv	Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent v	MB Rcvd	RcvdSz	SendSz
√9999	Total	0	11	30	16802	0	17642	0.00	69728	16802	4.15	13.66	1.89	113	853
1	Auto	3000	11	15	16802	0	17642	0.00	69728	16802	4.15	13.66	1.89	113	853
2	Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0
0	Login	5000	0	15	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0

Using field lists

- for each customer fields (cust-num name) no-lock: 1m15s

Table Activity															
Tbl#	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS	Read				
√	2	8 Customer	8	83	1254.49	0	104123	0	0	0	0				
	1	7 Invoice		147	0.00	0	0	0	0	0	0				
	3	7 Item		55	0.00	0	0	0	0	0	0				
	4	8 Order	5051	61297	0.00	0	0	0	0	0	0				
	5	8 Order-Line	7	415147	0.00	0	0	0	0	0	0				
	6	7 Salesrep		9	0.00	0	0	0	0	0	0				
	7	7 State		51	0.00	0	0	0	0	0	0				
	8	7 Local-Default		10	0.00	0	0	0	0	0	0				
Server Activity															
Srv	Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent v	MB Rcvd	RcvdSz	SendSz
√9999	Total	0	11	30	8781	0	10036	0.00	104123	8781	11.86	5.60	1.16	121	668
1	Auto	3000	11	15	8781	0	10036	0.00	104123	8781	11.86	5.60	1.16	121	668
2	Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0
0	Login	5000	0	15	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0

- Mm 8192: 1m15s

Table Activity															
Tbl#	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS	Read				
√	2	8 Customer	8	83	1250.27	0	103773	0	0	0	0				
	1	7 Invoice		147	0.00	0	0	0	0	0	0				
	3	7 Item		55	0.00	0	0	0	0	0	0				
	4	8 Order	5051	61297	0.00	0	0	0	0	0	0				
	5	8 Order-Line	7	415147	0.00	0	0	0	0	0	0				
	6	7 Salesrep		9	0.00	0	0	0	0	0	0				
	7	7 State		51	0.00	0	0	0	0	0	0				
	8	7 Local-Default		10	0.00	0	0	0	0	0	0				
Server Activity															
Srv	Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent v	MB Rcvd	RcvdSz	SendSz
√9999	Total	0	10	30	8752	0	10002	0.00	103772	8752	11.86	5.58	1.15	121	668
1	Auto	3000	10	15	8752	0	10002	0.00	103772	8752	11.86	5.58	1.15	121	668
2	Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0
0	Login	5000	0	15	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0

- prefetchDelay: 1m15s

Table Activity															
Tbl#	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS Read					
√	2	8 Customer	8	83	1304.85	0	108302	0	0	0					
	1	7 Invoice		147	0.00	0	0	0	0	0					
	3	7 Item		55	0.00	0	0	0	0	0					
	4	8 Order	5051	61297	0.00	0	0	0	0	0					
	5	8 Order-Line	7	415147	0.00	0	0	0	0	0					
Server Activity															
Srv Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent v	MB Rcvd	RcvdSz	SendSz	
√9999 Total	0	10	30	7829	0	9134	0.00	108306	7829	13.83	5.70	1.07	123	764	
1 Auto	3000	10	15	7829	0	9134	0.00	108306	7829	13.83	5.70	1.07	123	764	
2 Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0	
0 Login	5000	0	15	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0	

- + -prefetchNumRecs 100: 55s

Table Activity															
Tbl#	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS Read					
√	2	8 Customer	8	83	1808.57	0	150111	0	0	0					
	1	7 Invoice		147	0.00	0	0	0	0	0					
	3	7 Item		55	0.00	0	0	0	0	0					
	4	8 Order	5051	61297	0.00	0	0	0	0	0					
	5	8 Order-Line	7	415147	0.00	0	0	0	0	0					
Server Activity															
Srv Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent v	MB Rcvd	RcvdSz	SendSz	
√9999 Total	0	10	30	1809	0	3617	0.00	150110	1809	83.00	7.07	0.55	160	4098	
1 Auto	3000	10	15	1809	0	3617	0.00	150110	1809	83.00	7.07	0.55	160	4098	
2 Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0	
0 Login	5000	0	15	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0	

- prefetchNumRecs 300: 55s

Table Activity										
Tbl#	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS Read
√	2	8 Customer	8	83	1763.92	0	146405	0	0	0
	1	7 Invoice		147	0.00	0	0	0	0	0
	3	7 Item		55	0.00	0	0	0	0	0
	4	8 Order	5051	61297	0.00	0	0	0	0	0
	5	8 Order-Line	7	415147	0.00	0	0	0	0	0

Server Activity															
Srv	Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent v	MB Rcvd	RcvdSz	SendSz
√9999	Total	0	0	30	1764	0	3528	0.00	146407	1764	83.00	6.89	0.54	160	4098
1	Auto	3000	0	15	1764	0	3528	0.00	146407	1764	83.00	6.89	0.54	160	4098
2	Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0
0	Login	5000	0	15	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0

- prefetchFactor 100 –prefetchNumrecs 16: 1m15s

Table Activity										
Tbl#	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS Read
√	2	8 Customer	8	83	1288.11	0	106913	0	0	0
	1	7 Invoice		147	0.00	0	0	0	0	0
	3	7 Item		55	0.00	0	0	0	0	0
	4	8 Order	5051	61297	0.00	0	0	0	0	0
	5	8 Order-Line	7	415147	0.00	0	0	0	0	0

Server Activity															
Srv	Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent v	MB Rcvd	RcvdSz	SendSz
√9999	Total	0	0	30	7729	0	9017	0.00	106912	7729	13.83	5.63	1.06	123	764
1	Auto	3000	0	15	7729	0	9017	0.00	106912	7729	13.83	5.63	1.06	123	764
2	Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0
0	Login	5000	0	15	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0

increasing -Mm

- Mm 16384 (no prefetch params): 1m15s

Table Activity															
Tbl#	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS Read					
√	2	8 Customer	8	83	1234.26	0	102444	0	0	0					
	1	7 Invoice		147	0.00	0	0	0	0	0					
	3	7 Item		55	0.00	0	0	0	0	0					
	4	8 Order	5051	61297	0.00	0	0	0	0	0					
	5	8 Order-Line	7	415147	0.00	0	0	0	0	0					
Server Activity															
Srv	Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent v	MB Rcvd	RcvdSz	SendSz
√9999	Total	0	10	30	8639	0	9873	0.00	102443	8639	11.86	5.51	1.14	121	668
	1 Auto	3000	10	15	8639	0	9873	0.00	102443	8639	11.86	5.51	1.14	121	668
	2 Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0
	0 Login	5000	0	15	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0

- prefetchNumRecs 100: 59s

Table Activity															
Tbl#	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS Read					
√	2	8 Customer	8	83	1636.12	0	135798	0	0	0					
	1	7 Invoice		147	0.00	0	0	0	0	0					
	3	7 Item		55	0.00	0	0	0	0	0					
	4	8 Order	5051	61297	0.00	0	0	0	0	0					
	5	8 Order-Line	7	415147	0.00	0	0	0	0	0					
Server Activity															
Srv	Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent v	MB Rcvd	RcvdSz	SendSz
√9999	Total	0	10	30	3272	0	4908	0.00	135799	3272	41.50	6.54	0.67	143	2096
	1 Auto	3000	10	15	3272	0	4908	0.00	135799	3272	41.50	6.54	0.67	143	2096
	2 Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0
	0 Login	5000	0	15	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0

adding prefetchDelay

- + -prefetchDelay: 57s

Table Activity															
Tbl#	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS Read					
√ 2	8	Customer	8	83	1805.47	0	149854	0	0	0					
1	7	Invoice		147	0.00	0	0	0	0	0					
3	7	Item		55	0.00	0	0	0	0	0					
4	8	Order	5051	61297	0.00	0	0	0	0	0					
5	8	Order-Line	7	415147	0.00	0	0	0	0	0					
Server Activity															
Srv	Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent v	MB Rcvd	RcvdSz	SendSz
√9999	Total	0	10	30	1805	0	3611	0.00	149854	1806	83.00	7.06	0.55	160	4098
1	Auto	3000	10	15	1805	0	3611	0.00	149854	1806	83.00	7.06	0.55	160	4098
2	Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0
0	Login	5000	0	15	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0

- Notice the reduced message count: 3272 msg/sec vs. 1806 msgs/sec
 - Why?
- So why didn't average message size change with -Mm ?

send the whole 174 byte record

- Same 57 seconds:

Table Activity																
Tbl#	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS Read						
√	2	8 Customer	8	83	1806.82	0	149966	0	0	0						
	1	7 Invoice		147	0.00	0	0	0	0	0						
	3	7 Item		55	0.00	0	0	0	0	0						
	4	8 Order	5051	61297	0.00	0	0	0	0	0						
	5	8 Order-Line	7	415147	0.00	0	0	0	0	0						
Server Activity																
Srv	Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent v	MB Rcvd	RcvdSz	SendSz	
√9999	Total		0	10	30	1807	0	3613	0.00	149962	1807	83.00	26.21	0.54	156	15213
	1 Auto	3000	0	10	15	1807	0	3613	0.00	149962	1807	83.00	26.21	0.54	156	15213
	2 Inact		0	0	0	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0
	0 Login	5000	0	0	15	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0
User IO Activity																

- But now sending 15K per message
- 83 records X 174 bytes = 14.5K

using order (61K records)...

- No prefetch
- One single user running “for each order no-lock ...” 10 times
- 29 sec

Table Activity												
Tbl#	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS	Read	
√ 4	8	Order	5051	61297	6.41	0	392631	0	0	0	0	
1	7	Invoice		147	0.00	0	0	0	0	0	0	
2	8	Customer	8	83	0.00	0	0	0	0	0	0	
3	7	Item		55	0.00	0	0	0	0	0	0	
5	8	Order-Line	7	415147	0.00	0	0	0	0	0	0	

Server Activity															
Srv	Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent v	MB Rcvd	RcvdSz	SendSz
√9999	Total	0	1	30	194855	0	194867	0.00	392618	194860	2.01	2104.69	20.07	108	11326
1	Auto	3000	1	15	194855	0	194867	0.00	392618	194860	2.01	2104.69	20.07	108	11326
2	Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0
0	Login	5000	0	15	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0

- With prefetch: 29 sec – WAIT ? WHAT ?

Table Activity												
Tbl#	Area#	Table Name	RM Chain	#Records	Turns	Create	Read v	Update	Delete	OS	Read	
√ 4	8	Order	5051	61297	7.15	0	438130	0	0	0	0	
1	7	Invoice		147	0.00	0	0	0	0	0	0	
2	8	Customer	8	83	0.00	0	0	0	0	0	0	
3	7	Item		55	0.00	0	0	0	0	0	0	
5	8	Order-Line	7	415147	0.00	0	0	0	0	0	0	


Server Activity															
Srv	Type	Port	Cnx	Max	QryRcvd	RecRcvd	MsgRcvd	rr/msg	RecSent	MsgSent	rs/msg	MB Sent v	MB Rcvd	RcvdSz	SendSz
√9999	Total	0	1	30	217257	0	217270	0.00	438114	217262	2.02	2347.78	22.38	108	11331
1	Auto	3000	1	15	217257	0	217270	0.00	438114	217262	2.02	2347.78	22.38	108	11331
2	Inact	0	0	0	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0
0	Login	5000	0	15	0	0	0	0.00	0	0	0.00	0.00	0.00	0	0

- -Mm makes a difference for lobs
- Send 83 customer lobs @ 4.7 MB each
- -prefetchDelay –prefetchNumRecs 100

- -Mm 1024: 11 sec
- -Mm 16384: 4 sec
- -Mm 32600: 3.5 sec

- With lobs, the prefetch params didn't make much difference
 - Server sends 32000 bytes at a time

Q&A



pk@wss.com