

A graphic at the top of the page featuring two golden, complex mechanical structures resembling gears or stylized globes. These structures are interconnected by a large, curved golden arrow pointing from left to right. The background is a blue and white geometric pattern.

NAV 2009 Scalability

Locking Management Solution
for Dynamics NAV
SQL Server Option

Stress Test Results

White Paper

July 2009

Table of Contents

EXECUTIVE SUMMARY 3

INTRODUCTION..... 3

HARDWARE..... 5

RESULTS 6

CONCLUSION..... 8

PARTNERSHIP..... 9

EXECUTIVE SUMMARY

Locking Management for Dynamics NAV 2009 is a set of modifications and programming methodology that allows minimizing locking in the Dynamics NAV system. “Stress Test” shows a throughput of over 120000 shipped and invoiced orders with over 600000 lines per hour in 1000 user environment with 0 Deadlocks and Lock Timeouts and an ability for other users to work without any noticeable delays.

Locking Management methodology allows Dynamics NAV systems to scale to thousands of users making it competitive with SAP and Oracle applications on the upper midsize and enterprise ERP markets.

INTRODUCTION

Locking Management methodology was developed based on the standard practices for parallel programming. Dynamics NAV locking logic was modified to allow the system to process every user as an independent process that competes with other users for shared resources. We have introduced in the system “abstractions” that define shared system resources and efficient “semaphore” functions that allow locking shared resources when necessary.

From a Parallel Programming point of view, Dynamics NAV system can be represented as a multi-process system with shared resources defined at Table level. Dynamics NAV uses LOCKTABLE command to ensure that two users will not process conflicting read/write transactions. Essentially LOCKTABLE command is used as “semaphore” function from parallel programming point of view. Multiple users are competing for the same resource/table in the standard Dynamics NAV system. Most of the problems arise from tables like: “Reservation Entry”, “Item Leger Entry”, “No. Series Line”.

Locking Management methodology allows defining resources on lower level so a smaller number of users are competing for the same resource. In addition, Locking Management “semaphore” function allows locked resources to be released before a transaction finishes if this does not affect the integrity of the system. This allows multiple users to access the same resource during long transactions.

Standard Dynamics NAV 2009 system was modified using Locking Management methodology. Modifications affected 11 tables and 160 codeunits. Modifications were focused on the Dynamics NAV locking process and practically did not have any effect on the system business logic. All modifications are made in the system without use of any external components.

THE TEST

The “Stress Test” was designed to find any possible software bottlenecks in the Locking Management methodology. We have tested three different server configurations with 1050 users and performed multiple tests continually increasing load on the system until a limit was reached on each server. We measured and compared system throughput for 1050 users processing Create and Post Sales Order (Ship&Invoice) role with different delays between transactions. Each order had 5 lines. USER Profile was configured for 1 ms delay between field entries to emulate auto created orders. In addition manual orders have been created by users to verify acceptable response time and overall user experience (see user experience video).

Each test lasted one hour and average measures were calculated based on the middle 30-40 minutes of the test. Number of Order Lines is the actual number of lines that have been processed in one hour.

Each system was configured with 10000 customers and 50000 items with default CRONUS setup for Taxes and Discounts. Beginning inventory had been posted for each item to insure that each transaction processes Item Application Entries. Before starting the test we processed over 3 million sales orders which created over 18 million item ledger entries. Database size was approximately 145GB before first test and 200GB after all tests were done.

We have identified four possible bottlenecks: Processor, Hard Drive System, RAM and Software. A resource is considered a bottleneck if its load is close to 100% and all other resources are not loaded close to 100% and the system cannot process more transactions with increasing load. We consider Software to be a bottleneck if all hardware resources are not loaded close to 100% and system cannot process more transactions with increasing load.

Tests were performed using Application Benchmark Tool. One modification was made to insure more “randomized” user load. The original tool selected a fixed delay between transactions. In our test 50 second delay means that the system was randomly selecting a delay between 0 and 100 seconds (therefore the Average delay in this case was 50 seconds).

The system was working few hours before each test to make sure that necessary data was properly cached. Effect of SQL “checkpoints” was eliminated to make results comparable and independent from random variations.

Following system metrics were recorded and compared in the tests:

Number of orders Lines per Hour – a common measure for ERP system throughput. In the test each user was creating and processing orders with 5 lines.

Average time per Order – a measure of user experience. Most of the tests consider 5000 ms acceptable for Create and Post Sales Order transaction.

Average Lock Time, Locks Per Second, Total Lock Time per Sec – a measure of user experience. It gives a very good understanding of what users will experience after each field validation or during posting.

Processor Time (%) – a measure of processor load. Processor Time % should be under 85-90%.

Disk Time (%) – a measure of load for a storage system. In our test E: Drive was used for the Log file (so writing only). N: Drive was used for the database files and had both Reads and Writes.

Cache Hit Ratio (%) – a measure of RAM and Storage System efficiency. Higher hit ratio means less data has to be read from hard drives and less load to the database hard drives.

HARDWARE

All hardware provided by Aquarius Production Company, Russia (<http://www.aq.ru>).

Database Server LOW: AquaServer T50 D50

2 Quad Core Intel Xeon E5420 2.50 GHz
12 GB RAM
64bit Windows Server R2
SQL Server 2008 64 Bit
Drive C: System 2x300GB SAS RAID 1

Database Server PRO: AquaServer T50 D50

2 Quad Core Intel Xeon E5420 2.50 GHz
32 GB RAM
64bit Windows Server R2
SQL Server 2008 64 Bit
Drive C: System 2x300GB SAS RAID 1

Database Server HIGH: AquaServer N70 Q42

4 Six Core Intel Xeon E7450 2.40 GHz
32 GB RAM
64bit Windows Server R2
SQL Server 2008 64 Bit
Drive C: System 2x300GB SAS RAID 1

Storage System: Aquarius ARRAY FS5412

External SAN with 12 SAS 15K RPM HDD
Drive E: Log 4x300GB SAS RAID 10
Drive N: Database 8x300GB SAS RAID 10

Terminal Servers: AquaServer N90 X60 (6 blades) and AquaServer T50 D50

2 Quad Core Intel Xeon E5420 2.50 GHz
16 GB RAM
64bit Windows Server R2

Database server and Terminal Servers were connected by 1Gb network with each terminal server running 150 users.

RESULTS

Following table summarizes test results for T50 LOW Server (average statistics were taken from middle 30-40 mins of the test, Number of order Lines is the actual number of lines processed):

Measure/ Test	Number of Order Lines (Per Hour)	Avg. Time per Order (ms)	Process or (%)	Disk E: (%)	Disk N: (%)	Cache Hit Ratio	Locks Per Second	Avg. Lock Time (ms)	Total Lock Time per Sec (ms)	Numb er of Errors
175 Sec Delay	110825	1072	16.33	18.42	112.7	99.75	7.86	78.82	650.63	0
150 Sec Delay	129280	1201	19.38	22.77	152.4	99.77	11.35	97.59	1201.97	0
125Sec Delay	153060	1388	23.93	34.15	209.5	99.854	17.25	148.45	2888.16	0
115 Sec Delay	165640	1719	26.08	34.29	227.8	99.839	21.12	208.07	5148.29	0
100 Sec Delay	186990	2770	29.82	41.70	266.5	99.841	29.62	567.58	18306.21	0

System Throughput is constantly growing with decrease of delay between transactions but Average Lock Time is growing to 567 ms when delay between transactions reaches 100 Seconds. User will notice this delay during data entry time.

Even with Average lock Time of 567 ms average time for Create &Ship Order transaction was just 2.77 seconds what is normally considered acceptable for any ERP system. But further increase of load will lead to exponential growth of average time per order.

Tests show that the bottleneck in current configuration is throughput for Disk N. The Disk was mostly used for data reading. There are two ways to increase throughput for Disk N: increasing number of drives or increasing RAM which will be used for data caching.

Following table summarizes test results for T50 PRO Server (average statistics were taken from middle 30 mins of the test, Number of order Lines is the actual number of lines processed):

Measure/ Test	Number of Order Lines (Per Hour)	Avg. Time per Order (ms)	Process or (%)	Disk E: (%)	Disk N: (%)	Cache Hit Ratio	Locks Per Second	Avg. Lock Time (ms)	Total Lock Time per Sec (ms)	Numb er of Errors
50 Sec Delay	372365	1194.02	61.62	55.21	27.62	99.971	54.81	39.02	2114.34	0
45 Sec Delay	410955	1287.97	68.82	61.84	29.83	99.975	77.49	35.64	2726.65	0
40 Sec Delay	459405	1442.27	77.74	70.29	22.42	99.907	128.5	34.03	4640.08	0

35 Sec Delay	509280	2231.29	87.00	81.53	28.83	99.947	292.61	56.32	18720.89	0
30 Sec Delay	495790	8666.96	87.15	79.75	57.16	99.943	436.54	401.00	173738.4	126 (*)

(*) all errors happened when benchmark tool started recording the results and this caused the system to lock "Client Log Entry" table.

From the results we can see that system throughput increased up to 509208 lines orders (101856) per hour until delay between transactions reached 35 seconds and decreased when delay was 30 seconds.

Server T50 PRO has exactly the same configuration as T50 LOW except RAM was increased from 12GB to 32GB. This change in server configuration leads to more than double system throughput and eliminated Drive N bottleneck.

Average Lock Time stays constant up to 40 second delay between transactions. Then the system reached a hardware bottleneck (at about 35 sec between transactions) thus increasing Average Lock Time.

Processor was the bottleneck for the T50 PRO server configuration and Disk E: (Log File) came close to becoming a bottleneck as well.

Cache Hit ratio was over 99.9%. This shows that 32GB RAM is enough for current database size and test configuration. Disk N: (Database) usage was just below 30%. This means that 0.1% of data reads and data writes have loaded Disk N: almost 30%. In real life environment when users print multiple reports Cache Hit Ratio can decrease and Disk N: (Database) will become a bottleneck again. We recommend using "Data Warehouse" with BI reporting tools for databases with high transaction volumes.

Following table summarizes test results for N70 Server (average statistics were taken from middle 30-40s min of the test, Number of order Lines is the actual number of lines processed):

Measure/Test	Number of Order Lines (Per Hour)	Avg. Time per Order (ms)	Process or (%)	Disk E: (%)	Disk N: (%)	Cache Hit Ratio	Locks Per Second	Avg. Lock Time (ms)	Total Lock Time per Sec (ms)	Number of Errors
40 Sec Delay	461690	1242	26.21	77.10	50.32	99.97	79.30	36.55	2861.56	3
35 Sec Delay	522500	1368	30.69	90.30	52.17	99.98	131.25	34.02	4445.17	1
30 Sec Delay	595380	1865	36.61	111.8	46.52	99.98	291.02	46.74	14716.08	0
25 Sec Delay	612715	4357	37.75	120.8	74.23	99.98	436.84	209.62	98445.25	0

Server N70 has much more processing power in comparison with T50 PRO – 24 cores vs 8 cores. This eliminates the Processor bottleneck from the T50 PRO server. As we can see

from the tests with 40, 35 and 30 seconds delay the system was able to process more transactions and processing time decreased.

Somewhere around 30 seconds new bottleneck was reached: Disk E: (Log File) throughput. Even so the number of processed Sales lines grew up to 612715 before the system started experiencing significant delays and relatively long waits times.

CONCLUSION

Locking Management Solution from Thrifty Software Building Team Ltd can provide highly scalable solution based on the Dynamics NAV system. Tests show that simple improvements in hardware configuration will increase system throughput.

Server performance most of all depends on the size of RAM. In our test, increasing RAM to 15%-20% of database increased system throughput 2.7 times. The most important result is that we have not reached a "Software Bottleneck". Eliminating hardware bottlenecks is relatively cheap and does not require any software modifications. Dynamics NAV modified based on Locking Management Methodology is scalable to thousands of users and millions of transactions per day by simply adding more hardware.

Dynamics NAV modified based on Locking Management Methodology had shown the same or better performance than SAP in two tier configuration with comparable hardware - SAP test results [2009025](#), [2009024](#).

A properly balanced server will process a sales order in less than 2 seconds. Each transaction represents 20 field validations (5 for header and 3 per line) and takes less than 100 ms. This means that an user working in the system will not experience any delays. (See user experience video on our website www.thriftysoftwarebuildingteam.com).

PARTNERSHIP

We invite partners who are interested in implementing large scale solutions to contact us at sales@ThriftySoftwareBuildingTeam.com



Lean4systems Group unites companies in USA, Western and Eastern Europe. Since 2002 we do corporate information systems development, support and upgrades, as well as custom software development, on global Microsoft Dynamics and on local platforms.

The name of the Group unites words 'Lean' and 'Information Systems'. This is the way to reflect our effort to become 'Toyota' in the world of information technologies for business. We mean this as the ability to provide our clients with values of modern information systems in a faster and more reliable way than others do.

Our implementation, support and development processes are ISO 9001:2000 certified.

Thrifty Software Building Team Ltd is Cyprus-based IT company with main specialization in offshore programming for USA and Europe based customers, with focus on Microsoft Dynamics NAV platform. **Thrifty Software Building Team Ltd** is the member of **Lean4systems Group** since 2005.



Aquarius Group, established in 1989, became the first non-state company in Russia to open a large-scale manufacturing plant producing personal computers under its own brand name AQUARIUS.

Aquarius offers an extensive portfolio of products and services, including high-performance desktops and workstations, servers, notebooks, monitors, thin clients, storage systems and clustering solutions. Aquarius manufacturing systems meets world industry standards. Furthermore, its R&D engineers and scientists provide a wide range of hardware/software solutions based on Aquarius products. In addition to product manufacturing Aquarius provides their partners with turnkey IT solutions for various vertical markets taking into account the priority national projects, economic and social development of Russia's regions.

All products under the brand name AQUARIUS are certified according to Russian standards and regulations. In addition, the Aquarius Group has attained various international certificates and licenses, including ISO 9001:2001 international quality management system certification.

Aquarius Group has built a strong partnership with its more than 1300 dealers in 160 cities throughout Russia.

Aquarius Group operates the largest network of more than 365 service centers in 135 cities across the country, providing the highest level of competence and expertise in the Information Technology sector.

The latest surveys conducted by market researchers IDC, IT Research, Gartner revealed that Aquarius Group is ranked among the top five leaders in total PC and server sales on the Russian IT market.

The information contained in this document represents the current view of Thrifty Software Building Team Ltd on the issues discussed as of the date of publication. This White Paper is for informational purposes only. Thrifty Software Building Team Ltd MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT. Tests were performed in the lab environment for one user role. Results cannot be applied to the "Real Life" environment without necessary adjustments.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Thrifty Software Building Team Ltd.

Thrifty Software Building Team Ltd may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Thrifty Software Building Team Ltd, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.