WareLite Limited

WareLite Business Operating Support System

WL BOSS 3.0 Performance Benchmarks

Benchmark	**ROADMAP-A**
Date	1 March 2008
Status	**Final**
Release	1
Site	Milano
Availability	**General**

Sponsors	
ESIET S.r.l.	ESIET SERVIZI PER LA COMUNICAZIONE
Medan S.r.l.	MEDAN

4 Ramridge House, Andover SP11 9BY, UK
Email: info@warelite.com
www.warelite.net

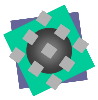

Table of Contents

1. A very hot room

As you will realize reading through this document, the figures we present here have been obtained on a stack of gloriously ancient consumer-oriented personal computers. Even though they have trespassed Obsolescence Land, there are several reasons why we have decided to keep them on benchmarks duty. The most important of all is aimed at keeping the reader focused on architecture and software. From a commercial standpoint, there might be a few unfair advantages in showing how the figures presented here can be achieved on the cheap. However, we believe that modern is better and it would be a bad idea to forget about the marvels of the already-mainstream n-core CPUs and any other technology that improves the efficiency of your processing resources.

Dario Amidani taming the "Grid"

WareLite BOSS 3.0 is the result of some important architectural changes that improve the inner parallelism of its components – the Node Manager and the Global Lock Manager. Thanks to these improvements you will be able to make the most out of more expensive computers, avoiding the drawbacks of an old–but-sturdy infrastructure – and you won't need to wait for winter to run benchmarks like these.

2. Summary

WareLite BOSS (Business Operating Support System) is an Event Driven Application Platform – a grid-based platform for the execution of event-driven, real time business processes requiring extreme transaction processing (XTP). The benchmarks published in this paper measure the performance of WL BOSS version 3.0 and its scalability at the persistence layer. Previously published documents present a detailed discussion on WL BOSS scalability at the execution layer.

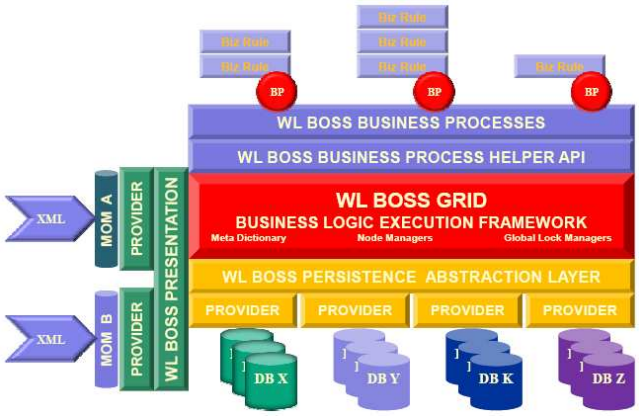

Figure 1 – WL BOSS Architecture

WareLite BOSS architecture has four main components:

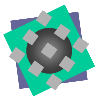

- Node Managers – Execute processes in a transactional framework
- Global Lock Managers – Ensure determinism to parallel processes generating contention
- Persistence Providers – Provide an abstraction layer between process execution and data storage/retrieval, implementing data partitioning and distribution across multiple third party RDBMSs (e.g. IBM DB2, Microsoft SQL Server)
- Presentation Layer – typically based on third party message oriented middleware (e.g. IBM MQSeries, Microsoft MSMQ), it presents events (XML frames) to WL BOSS

The WL BOSS installation used to obtain these benchmarks was made of:

- Up to 5 computers hosting the Node Managers and persistence services devoted to static Tariff repositories. In production implementations this number can be increased as needed to achieve the desired capacity and Tariff repositories can be hosted on a set of dedicated computers; any of the supported third party RDBMS can be deployed to host Tariff repositories.
- 1 computer hosting one Global Lock Manager. In production implementations the minimum recommended number of GLM is 2, to support high availability. This number can be increased as needed to achieve the desired capacity;
- Up to 8 computers hosting the persistence services/RDBMS (Microsoft SQL Server) devoted to maintaining customer profile and account balance records. In production implementations this number can be increased as needed to achieve the desired capacity and any supported RDBMS can be deployed.
- 1 computer hosting the Presentation Layer. In production implementations this number can be increased as needed to achieve the desired capacity and different well-established configurations (Clustering or Network Load Balancing) can be implemented.

The performance benchmarks shown in this document were obtained varying the number of Customer Objects within the WareLite BOSS Grid and monitoring the number of statements per second at the RDBMS level and the number of messages de-queued per second, leveraging third party performance counters.

The scalability benchmarks shown in this document were obtained varying the number of Persistence Providers within the WareLite BOSS Grid and monitoring the number of statements per second at the RDBMS level and the number of messages de-queued per second, leveraging third party performance counters.

The benchmark process used to produce the figures presented in this document is based on a simple real time rating engine requiring contention resolution, provided by the Global Lock Manager. Each process (a transactional unit) generates 4 statements ('read' and 'write') against one or more databases; thus the number of business processes per second is equal to the number of messages de-queued per second and can also be calculated dividing by 4 the number of statements per second measured

with the RDBMS performance counters. This benchmark process has been selected for its closeness to commercial applications. Please note that each business process fires some other statements (such as Begin Transaction and Commit or Rollback) that are not captured by the available performance counters exposed by the RDBMS used for these benchmarks.

The benchmarks show that:

With a hardware infrastructure that in 2006 was worth €17.020 (including networking and KVM equipment) and hosting Node Managers on 5 computers worth about €500 each, with 4,000,000 Customer Objects (represented as 8,000,000 records partitioned over 8 Persistence Providers), we have achieved a sustainable rate of 10,800 TPS (Transactions Per Second) against databases, i.e. 2,700 complete, deterministic Real Time Rating Processes Per Second, at a hardware cost per TPS of €1.58[1]. The Transaction Per Minute cost is thus €0.026. At today's hardware prices – for more powerful computers – these figures would be respectively €1.27 and €0.021.

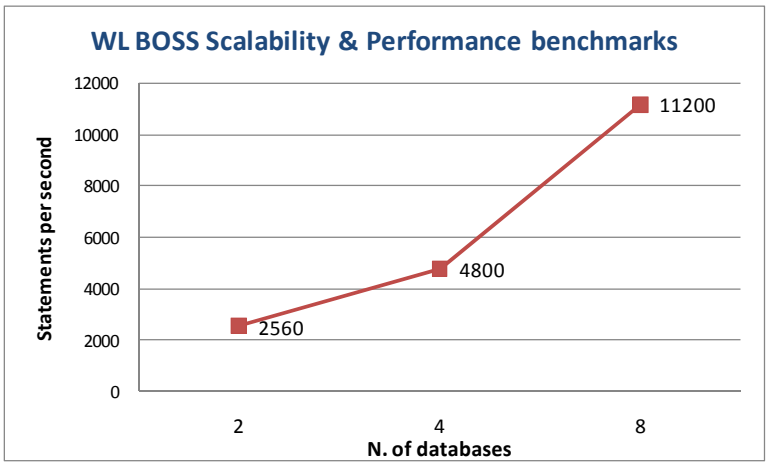

Figure 2 – WL BOSS Scalability Benchmarks

Performance increases linearly with the addition of computers hosting Persistence Providers to the grid, as shown in Figure 2.

[1] This figure is calculated taking into consideration the total cost of hardware, including network equipment, utilised to measure the benchmarks

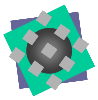

3. ROADMAP-A Description

The benchmark ROADMAP-A is based on the ROADMAP solution set. The ROADMAP solution set is distributed with the WL BOSS run-time and WL BOSS SDK licenses. The ROADMAP solution set comprises the following business processes:

- rdmp_create_product
- rdmp_create_customer
- rdmp_purchase_reseller
- rdmp_purchase
- drop

The ROADMAP-A benchmark utilises only the following business processes from the whole ROADMAP solution set:

- rdmp_create_product
- rdmp_create_customer
- rdmp_purchase

The benchmark ROADMAP-A is divided in two parts: customer creation performance and rating performance. The customer creation performance benchmark is based upon the business process rdmp_create_product whilst the rating benchmark is based upon the business process rdmp_purchase.
Please note that the ROADMAP-A solution set has been designed for benchmarking purposes and as such it is non-optimized. For instance, the rdmp_purchase business process can be implemented in a different way in order to reduce the number of operations (or statements) against the RDBMS from 4 to 3. Also, with a slightly different implementation of the WL.examples.roadmap.customer class, it would be possible to dramatically reduce the number of connections towards the underlying RDBMS. However, we have decided to keep using the ROADMAP-A solution set to allow comparison among the various benchmark documents we have published in the past.

3.1. Customer Creation

During this phase of the benchmark, one or more client agents present events (WL EPL XML documents, each containing one instance of the class WL.examples.roadmap.customer) to a WL BOSS domain presentation component. One or more Node Managers drain the events and propagate the instances to one or more persistence providers.

3.2. Rating

During this phase of the benchmark, a rating workload is presented to the WL BOSS domain. The workload is generated by one or more client agents presenting instances of the class WL.examples.roadmap.purchase to the domain presentation layer. In more detail, for each run a purchase event is presented for each customer object created during the customer creation phase, i.e. if N different customer objects have been created, then during the rating test phase N different purchase events will be presented to WL BOSS, simulating one purchase event for each single customer.

The business process (rdmp_purchase) triggered by each event will:

- load the customer object
- load the tariff of the product being purchased
- calculate the impact
- update the customer object (new balance)

This business process generates 3 read operations against one or more RDBMS (the WL.examples.roadmap.customer class has been declared with inner partitioning) and 1 write operation against one or more RDBMS (in order to update the customer object with the newly calculated balance), for a total of 4 operations against one or more RDBMS (transactional mix: 75% Read, 25% Write).

The rating phase of the benchmarks, differently from the customer creation phase, produces workload against the Global Lock manager component, ensuring determinism whenever the object 'customer' is updated.

Figure 3 shows the rdmp_purchase business process. The arrows with orientation toward a RDBMS represent write operations. The arrows with orientation toward the business rules represent read operations.

Figure 3 - WL BOSS Rating Benchmarks Process

Please note that each business process also fires N "Begin Transaction" and N Rollback or N Commit statements against the involved resource managers, where N is the number of RDBMS instances involved in one single transaction.

4. WL BOSS Domain Configuration (Hardware)

Figure 4 shows the hardware configuration used for the benchmarks described in this document. From top to bottom:

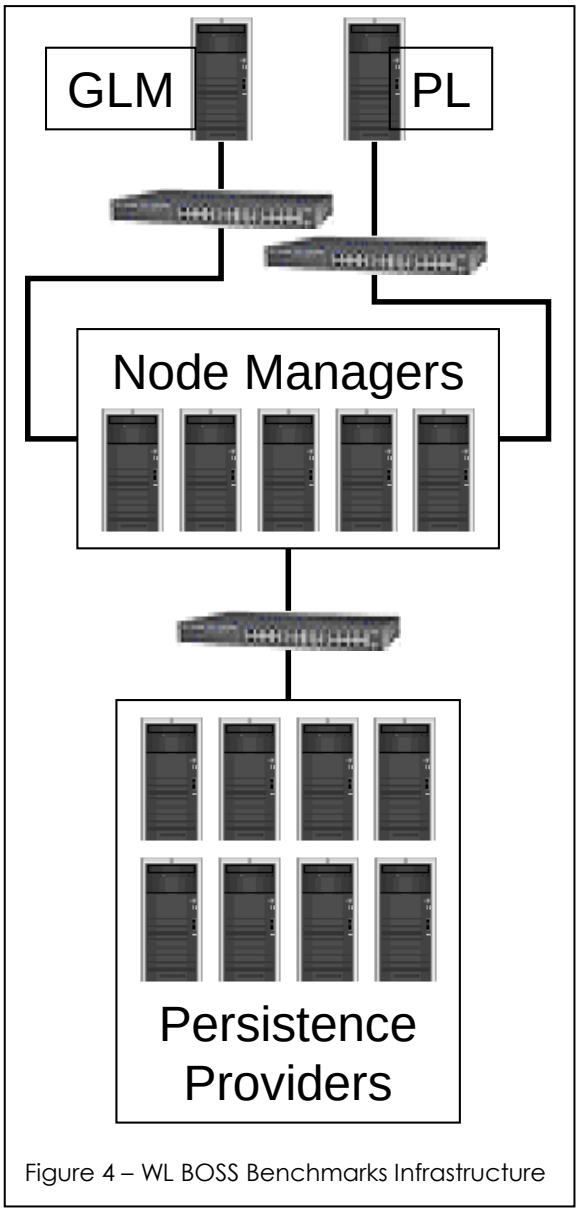

Figure 4 – WL BOSS Benchmarks Infrastructure

- **GLM** is the single Global Lock Manager used for the benchmark. The computer used for the benchmark is a HP ML110 G3 with 512 megabytes RAM.
- **PL** is the single Presentation Layer component used for the benchmark. The computer used for the benchmark is a HP ML110 G3 with two gigabyte RAM.
- The **GLM** is connected via a **Gigabit switch** to the Node Managers (**subnet A**). The Gigabit switch is a Netgear GS724T.
- The **PL** is connected to the **Node Managers** via a **Gigabit switch** (**subnet B**). The Gigabit Switch used for the benchmark is a Netgear GS724T.
- The **five** computers hosting the **Node Manager** processes used for this benchmark. The computers used for the benchmark are HP ML110 G3 with 512 megabytes RAM. Please note that the same computers have been used to host the Tariff repository, too.
- The **five Node Managers** are connected to the **persistence providers** via a **Gigabit switch**. The Gigabit switch used for the benchmark is a Netgear GS724T.
- The **eight persistence providers**. The computers used for the benchmark are HP ML110 G3 with two gigabyte RAM.

The following tables provide details about network devices and computers (brand, configuration, etc.) used for the benchmark. The clients presenting events to the WL BOSS domain shown in the picture connect to the subnet B. We have used several different computers, typically laptops or desktops with a Pentium 4 2 Ghz or equivalent. The performance obtainable by the clients is not covered by these benchmarks.

Computers				
	Node Manager	**Presentation Layer**	**Persistence Provider**	**Global Lock Manager**
Quantity	**5**	**1**	**8**	**1**
Model	HP ML 110g3	HP ML 110g3	HP ML 110g3	HP ML 110g3
CPU	Intel Pentium 4 3.0 GHz, 800MHz FSB	Intel Pentium 4 3.0 GHz, 800MHz FSB	Intel Pentium 4 3.0 GHz, 800MHz FSB	Intel Pentium 4 3.0 GHz, 800MHz FSB
Memory RAM	512MB	2048MB	2048MB	512MB
Disk Drives	80 GB Ultra SATA Non Hot Plug	80 GB Ultra SATA Non Hot Plug	2 x 80 GB Ultra SATA Non Hot Plug + 1 IDE 80 Gb	80 GB Ultra SATA Non Hot Plug

Networking Devices & Other		
Switches		
Quantity	**Model**	**Specifications**
3	Netgear GS724T	24x10/100/1000Base-T autosensing
16	Netgear GA311	10/100/1000Base-T Ethernet card
KVM		
Quantity	**Model**	**Specifications**
1	Avocent AV 2000	2 users – 16 server control
2	Avocent AMX5120	User Station
8	Avocent AMIQ-PS2	Balun PS2
2	HP PX850AT	L1906 TFT 19“

The following tables show the prices for the computers and networking devices utilised for these benchmarks.

COMPUTERS				
Model	**Description**	**Part No.**	**Qty**	**Unit Price in €**
HP ML 110	ML 110 t g3 p4 30 2mb512mbRAM 80gb SATA	HP470063-787	15	468,00
	512MB ADV ECC PC2 4200 DDRII SDRAM	HP390825-B21	25	84,00
	80 gb SATA		8	80,00
	80 gb IDE		8	80,00
			Subtotal A - €	**10.400,00**

NETWORKING				
Brand	**Description**	**Part No.**	**Qty**	**Unit Price in €**
Netgear	Netgear GS724T 24x10/100/100Base-T autosensing	GS724T	3	474,00
	Netgear GA311 10/100/1000Base-T	GA311	20	24,00
HP	VIDEO LCD L1906 TFT 19“	HP PX850AT	2	409,00
AVOCENT	2 local users, 16 system switch, with GUI	AV2000-203	1	900,00
	Avocent AMX5120	AMX5120	2	900,00
	Balun PS2	AMIQ-PS2	8	150,00
			Subtotal B - €	**6.620,00**

Total (A+B) - €	**17.020,00**

Note: *the new ML110 G5, with an E2160 processor, costs €255. This means that, at today's price, our benchmarks infrastructure would be worth €13.725.*

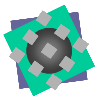

5. WL BOSS Domain Configuration (Software)

SOFTWARE CONFIGURATION	
Node Manager	**Microsoft Windows 2003 Server Standard with SP2** **Microsoft SQL Server 2000 with SP4** **IBM MQSeries client** **WL BOSS**
WLNM01,WLNM02,WLNM03 WLNM04,WLNM05	
Presentation Layer	**Microsoft Windows 2003 Server Standard with SP2** **IBM MQSeries 6.00**
WLPL01	
Global Lock Manager	**Microsoft Windows 2003 Server Standard with SP2** **WL BOSS**
WLGLM01	
Persistence Providers	**Microsoft Windows 2003 Server Standard with SP2** **Microsoft SQL Server 2000 with SP4**
WLPP01,WLPP02,WLPP03, WLPP04, WLPP05, WLPP06, WLPP07, WLPP08	

Notes:
All the benchmarks provided here utilise one single queue for the presentation layer.

Detailed analysis has highlighted that the bottleneck of the setup used for the benchmarks presented in this document is related to the databases hosted by the persistence providers.

Each node manager hosts a copy of an instance of the Tariff repository (identified as Persistence Provider WL03).

The WLPP01-08 computers host the persistence providers WL01 and WL02. Both persistence providers are distributed over all the WLPP01-08. The partitioning of WL01 and WL02 has been provided in order to maximize workload distribution, too.

On each WLPPx computer WL01 and WL02 are two logical databases owned by the same physical instance of Microsoft SQL Server. With this setup and considering a total of 4,000,000 customer objects, each logical instance of WL01 and WL02 contains about 250,000 records. With a total of 100,000 customer objects each WL01 and WL02 contains about 6250 records.

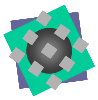

6. Measurement methodology

The performance figures presented in the benchmarks results are always obtained by monitoring the number of statements per second at the RDBMS level and the number of messages de-queued per second using the performance counters exposed by IBM MQSeries. Please note that several third party performance counters use the word "transaction" as a synonym of the word "statement".

The total number of statements (fired against a RDBMS) generated by a single business process varies depending on the business process' scope. For example, the 'rdmp_create_customer' business process generates **two** insert statements against one or more RDBMS. The two insert operations belong to the same transactional unit, as any WL BOSS business process is executed within a transactional executive framework. Thus, a total of N insert statements per second measured on the target RDBMS means N/2 business 'rdmp_create_customer' processes per second executed by WL BOSS. Please note that there are other "hidden" statements fired against the RDBMS such as Begin Transaction, Commit and Rollback. If we counted these statements, then we should add at least 2 statements for each business process to the overall amount of statements per seconds.

To obtain the RDBMS performance figure, a third party monitor facility has been used. In these benchmarks, the number of transactions per second has been obtained using the performance counters exposed by Microsoft SQL Server. The results have also been cross-checked with other methods.

The word "transaction" is used with its technical meaning throughout this document – each transaction comprising one or more "begin transaction" instruction, a set of operations providing the content of the transaction and each transaction ending with one or more "commit" or "rollback" instructions. In this respect, the RDBMS deployed for the benchmarks are always accessed with AUTOCOMMIT OFF. Also, Microsoft SQL Server and IBM MQSeries have been installed and used out of the box without applying any configuration optimization.

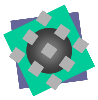

7. How to read the benchmarks detailed results

The following chapter will provide all the results collected through several "runs". Please note that all runs described in this document use the same hardware configuration, as detailed in previous chapters.

For each run, this document provides three different tables:

1. **RUN Details**
2. **RUN results (Customer Creation)**
3. **RUN results (Rating)**

The first table (**RUN Details**) contains details about the run. In this table you will find all relevant information about the kind of partitioning that has been configured, which persistence providers have been used for the run, how many customer objects are being created, how many product objects are being created and how many Global Lock Managers are being used.

The second and third table (RUN Results, Customer Creation and Rating) contain 6 columns. The first 3 columns are representative of the configuration of the deployed domain of WareLite Node Managers, the other columns show the results obtained on such domain.

The first column (**NMHOST**) is the total number of computers used to obtain the results shown on the same row. The second column (**NMPROC**) represents the number of Node Manager processes running on **each** computer (**NMHOST**) in the domain. The third column (**NMTHREAD**) shows the number of concurrent business processes running on **each** Node Manager process.
The following table provides an example of this representation.

RUN Results					
CUSTOMER CREATION					
NMHOST	**NMPROC**	**NMTHREAD**	**P/S AVG**	**P/S MAX**	**DB T/S AVG**
1	1	2	100	120	200
2	6	4	1037	1266	2074

The table above shows two sets of results obtained on two slightly different domain configurations.

The first row describes a domain that includes 1 computer running 1 single instance of a Node Manager process which is configured to run 2 concurrent business process threads. In this configuration, the total number of computers is 1, the total number of Node manager processes is 1 and the total number of concurrent business process threads is 2.

The second row describes a domain that includes 2 computers. Each computer runs 6 instances of the Node Manager process. Each instance of the Node Manager process is configured to run 4 concurrent business process threads. Thus in this domain there are 12 Node Manager processes running and a total of 48 concurrent business process threads.

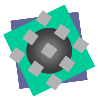

The fourth, fifth and sixth columns respectively show the sustainable number of business processes executed per second, the max number of business processes executed per second and the sustainable number of transactions committed per second against one or more RDBMS (the number of physical RDBMS – Persistence Providers - is shown in the RUN details table).

8. Benchmarks detailed results

8.1. Benchmark RUN 1

RUN details	
Persistence Providers	8, WLPP01-PP08
Partitioning	Horizontal Split
Number of Customer Objects	4,000,000
Number of Product Objects	1,000
Number of Global Lock Managers	1

RUN Results					
CUSTOMER CREATION					
NMHOST	NMPROC	NMTHREAD	P/S AVG	P/S MAX	DB T/S AVG
5	1	100	5500	11000	11000

RUN Results					
RATING					
NMHOST	NMPROC	NMTHREAD	P/S AVG	P/S MAX	DB T/S AVG
5	1	100	2700	2950	10800

8.2. Benchmark RUN 2.A

RUN details	
Persistence Providers	8, WLPP01-PP08
Partitioning	Horizontal Split
Number of Customer Objects	100,000
Number of Product Objects	1,000
Number of Global Lock Managers	1

RUN Results					
CUSTOMER CREATION					
NMHOST	NMPROC	NMTHREAD	P/S AVG	P/S MAX	DB T/S AVG
5	1	100	5500	11000	11000

RUN Results					
RATING					
NMHOST	NMPROC	NMTHREAD	P/S AVG	P/S MAX	DB T/S AVG
5	1	100	2900	3200	11600

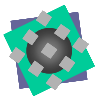

8.3. Benchmark RUN 3 – Persistence Layer Scalability

RUN details	
Persistence Providers	2,4,8, WLPP01-PP08
Partitioning	Horizontal Split
Number of Customer Objects	400,000
Number of Product Objects	1,000
Number of Global Lock Managers	1

RUN Results					
RATING					
NMHOST	NMPROC	NMTHREAD	NMPP	P/S AVG	DB T/S AVG
5	1	100	2	640	2560
5	1	100	4	1200	4800
5	1	100	8	2780	11120

9. Glossary

Glossary	
NMHOST	Number of computers running Node Manager processes
NMPROC	Number of Node Manager processes (Node Manager instances) for each computer (NMHOST)
NMTHREAD	Number of concurrent business processes for each Node Manager instance (NMPROC)
NMPP	Number of Persistence Providers
P/S AVG	Average number of business processes per second (sustainable)
P/S MAX	Maximum number of business processes per second (peak)
DB T/S AVG	Total Number of Transactions per Second committed on one or more RDBMS