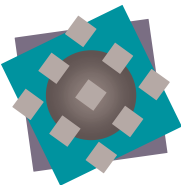


WareLite Limited

WareLite Business Operating Support System

Performance and Scalability Benchmarks

	Benchmark	ROADMAP-A
	Date	19 Apr 2006
	Status	Final
	Release	1
	Site	Milano
	Availability	Public



Sponsors	
ESIET S.r.l.	
Medan S.r.l.	



Table of Contents

1.	Summary.....	3
2.	ROADMAP-A Description.....	5
3.	WL BOSS Domain Configuration (Hardware)	8
4.	WL BOSS Domain Configuration (Software).....	11
5.	How to read the benchmark detailed results	12
6.	Benchmark detailed results	14
7.	Glossary	20



1. Summary

WareLite BOSS (Business Operating Support System) is an Event Driven Application Server – a grid-based platform for the execution of event-driven, real time business processes requiring extreme OLTP capabilities. The benchmarks reported in this paper measure WL BOSS performance and scalability.

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 (e.g. real time rating)
- Persistence Providers – Provide an abstraction layer between process execution and data storage/retrieval, implementing data partitioning and distribution across multiple RDBMSs
- Presentation Layer – Based on standard messaging (MSMQ), presents events to WL BOSS

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

- Up to 4 computers hosting the Node Managers. In production implementations this number can be increased as needed to achieve the desired capacity;
- 1 computer hosting one Global Lock Manager. In production implementations the minimum number of GLM recommended is 2, to support high availability, and it can be increased as needed to achieve the desired capacity;
- Up to 2 computers hosting the Persistence Providers/RDBMS (SQL Server). In production implementations this number can be increased as needed to achieve the desired capacity;
- 1 computer hosting the Presentation Layer (MSMQ). In production implementations this number can be increased as needed to achieve the desired capacity;

The benchmarks shown in this document were obtained varying the number of Node Managers and Persistence Providers (RDBMS) within the WareLite BOSS Grid and monitoring the number of transactions per second at the RDBMS level.

One of the benchmarks we executed is based on a simple real time rating process requiring contention resolution, provided by the Global Lock Manager. Each process generates 4 transactions ('read' and 'write' operations) against one or more databases, thus the number of business processes per second is calculated dividing by 4 the number of transactions per second measured by the RDBMS monitor. This benchmark has been selected for its closeness to commercial applications.

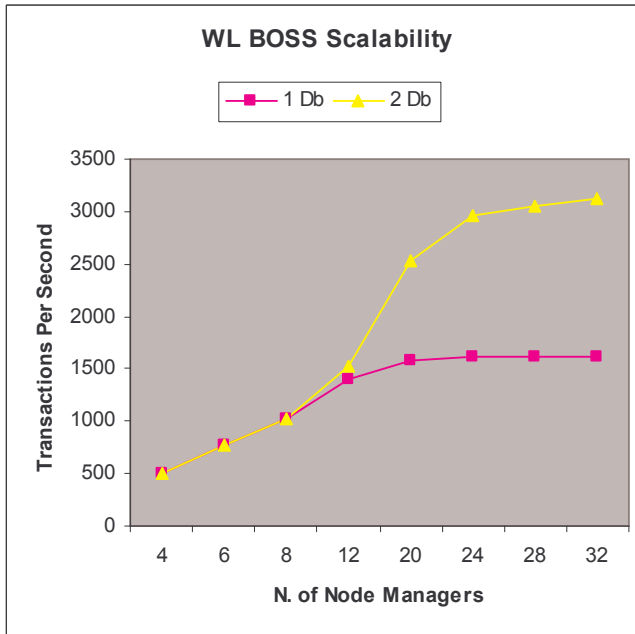
The benchmark shows that:

- **The capacity of WL BOSS is a linear function of the number of computers hosting the Node Managers (Horizontal Scalability).** In addition, multiple Node Managers can be hosted by the same computer, thus increasing the host capacity up to a limit dependent



upon the computer's characteristics (Vertical Scalability). In these benchmarks the capacity limit of a single computer is reached with 16 Node Managers.

- **The capacity of WL BOSS is a linear function of the number of Persistence Providers (i.e. Databases)**



The pink curve in the diagram shows that when using 1 single Db to support persistence, WL BOSS performance hits a plateau at around 1600 TPS, i.e. the addition of further nodes does not increase performance. This is because 1600 TPS is the capacity limit of a single database within our benchmark infrastructure. However, as shown in the yellow curve, horizontal scalability is restored by simply increasing the number of databases. The diagram shows that our

WL BOSS benchmark infrastructure saturates the capacity of 2 databases.

With a hardware infrastructure worth €19.302,00 (including networking and KVM equipment) and hosting Node Managers on 4 computers worth about €500 each, with 2,000,000 records partitioned over 2 Persistence Providers (Databases), we have achieved a sustainable rate of 3160 TPS (Transactions Per Second) against databases, i.e. 790 complete, deterministic Real Time Rating Processes Per Second, at a hardware cost per TPS of € 6.19¹. The cost of the Transaction Per Minute is thus € 0.10.

With the same hardware infrastructure we have achieved 3,800 TPS with 2,000,000 records for an intrinsically deterministic (i.e. with no Global Lock Manager) process. As this process comprises only one transaction, the number of Processes per Second (i.e. of events processed per second) equals the number of Transactions per Seconds. An intrinsically deterministic process might be a process providing transformation, implementing look-ups or appending records to a database. For instance most of the business processes needed to implement an RFID infrastructure are intrinsically deterministic processes.

Our benchmarks prove that WL BOSS provides extreme, scalable OLTP to event-driven business processes, thanks to the complete horizontal scalability of its four main components, Node Managers, Global Lock Managers, Persistence Providers and Presentation Layer:

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



- Performance can be increased by simply adding Node Managers to the infrastructure;
- Global Lock Managers ensure determinism but do not represent a performance bottleneck, as workload can be partitioned across any number of Global Lock Managers;
- The bottleneck typically exposed by databases does not limit overall WL BOSS scalability and performance: thanks to WL BOSS data partitioning and distribution services, performance can be increased by the addition of more physical databases to the infrastructure;
- Performance is not limited by the presentation layer, as multiple queues can be implemented to present events to the grid.

2. 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.

2.1. 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. The benchmark is run based on different configurations, as shown in the detailed results.

The performance figures presented in the benchmark results are always obtained by monitoring the number of transactions per second at the RDBMS level.

The total number of transactions (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 operations 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 single transactional framework. Thus, a total of N insert operations per second measured on the target RDBMS means N/2 business processes per second executed by WL BOSS.

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" instruction. In this respect, the RDBMS deployed for the benchmarks are always accessed with AUTOCOMMIT OFF and with the isolation level (for each session accessing the databases) set to READ COMMITTED, unless otherwise specified.

2.2. 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:

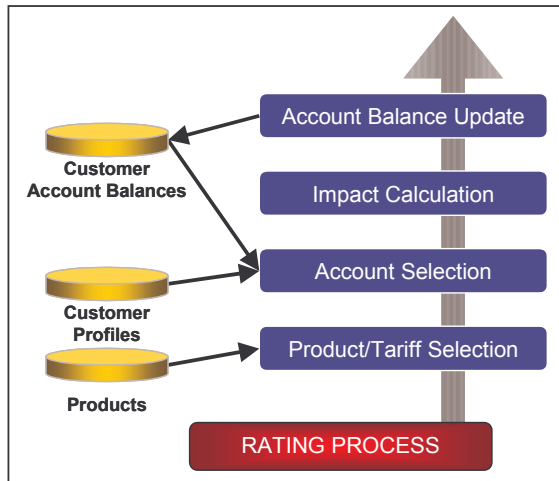
- load the customer object
- load the tariff of the product being purchased
- calculate the impact
- update the customer object (new balance)

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




The following picture 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.





3. WL BOSS Domain Configuration (Hardware)

The picture below shows the hardware configuration used for the benchmarks described in this document. From top to bottom:



- **GLM** is the single Global Lock Manager used for the benchmark. The computer used for the benchmark is a HP ML110 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 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 **four** computers hosting the **Node Manager** processes used for this benchmark. The computers used for the benchmark are HP ML110 with 512 megabytes RAM.
- The **four Node Managers** are connected to the **persistence providers** via a **Gigabit switch**. The Gigabit switch used for the benchmark is a Netgear GS724T.
- The **two persistence providers**. The computers used for the benchmark are HP ML350.



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 this benchmark (the actual RDBMS being the limit for higher throughput).

Computers				
	Node Manager	Presentation Layer	Persistence Provider	Global Lock Manager
Quantity	4	1	2	1
Model	HP ML 110g3	HP ML 110g3	HP ML 350g4	HP ML 110g3
CPU	Intel Pentium 4 3.0 GHz, 800MHz FSB	Intel Pentium 4 3.0 GHz, 800MHz FSB	2 x Intel Xeon Processor 3.2 GHz/800-2 MB L2	Intel Pentium 4 3.0 GHz, 800MHz FSB
Memory RAM	512MB	2048MB	2048MB	512MB
Disk Controllers	80 GB Ultra SATA Non Hot Plug	80 GB Ultra SATA Non Hot Plug	5 x 72.8 GB - ultra320 15k rpm SCSI	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 † g3 p4 30 2mb512mbRAM 80gb SATA	HP470063-787	6	468,00
	512MB ADV ECC PC2 4200 DDRII SDRAM	HP390825-B21	3	84,00
HP ML 350	pl ml350t g4p xeon 3 2 2mb cache 1gb ram raid sa641	HP380196-421	2	1704,00
	Intel 3.2 Ghz 2M-350 G4P Proc. Upgrade	HP382182-B21	2	470,0
	NC7771 10/100/1000T 64/133MHZ PCI-X	HP290563-B21	4	156,00
	HP 512 REG DDR2-3200 SGLDMM Memory	HP384163-B21	4	125,00
	72.8GB U320 15KRPM SCSI	HP286778-B22	10	415,00
Subtotal A - €				12.682,00

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) - €			19.302,00	
------------------------	--	--	------------------	--



4. WL BOSS Domain Configuration (Software)

SOFTWARE CONFIGURATION	
Node Manager	Windows 2000 Server
WLDC01	WL BOSS
Node Manager	Windows 2000 Professional with SP4
WLN02,03,04	WL BOSS
Presentation Layer	Windows 2000 Professional with SP4
WLPL01	
Global Lock Manager	Windows 2000 Server
WLGLM01	WL BOSS
Persistence Providers	Windows 2000 Server
WLPP01,02	MicroSoft SQL Server 2000 with SP3

Notes:

WLDC01 is configured as a MicroSoft Windows domain controller.

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



5. 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 for 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

This table 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).



6. Benchmark detailed results

6.1. Benchmark RUN 1

RUN details	
Persistence Providers	1, WLPP01
Partitioning	None
Number of Customer Objects	5,000
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
1	1	1	55	68	110
1	1	2	100	120	200
1	1	4	144	212	288
1	1	6	170	283	340
1	2	1	92	123	184
1	2	2	144	188	288
1	2	4	250	308	500
1	4	4	440	500	880
1	6	4	550	677	1100
2	6	4	1037	1266	2074
4	6	4	1950	2337	3900

RUN Results					
RATING					
NMHOST	NMPROC	NMTHREAD	P/S AVG	P/S MAX	DB T/S AVG
1	1	1	31	32	124
1	1	2	31	32	124
1	1	4	32	32	128
1	1	6	32	32	128
1	2	1	63	64	252
1	2	2	63	64	252
1	2	4	64	64	256
1	4	4	127	129	508
1	6	4	191	194	764
2	6	4	377	389	1508
4	5	4	400	427	1600
4	6	4	417	460 (1)	1668

(1)Notes: The performance limit of the RDBMS has been hit (the throughput upper limit is about 1600 transactions per second (based on the specific transactional mix of the rating business process)).



6.2. Benchmark RUN 2.A

RUN details	
Persistence Providers	1, WLPP01
Partitioning	None
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
1	1	4	156	196	312
1	2	4	244	283	488
1	3	4	282	362	564
1	4	4	335	487	670
1	5	4	454	595	908
1	6	4	531	665	1062
2	6	4	964	1268	1928
4	6	4	1839	2183	3678
4	7	4	2072	2318	4144
4	8	4	2055	2417	4110

RUN Results					
RATING					
NMHOST	NMPROC	NMTHREAD	P/S AVG	P/S MAX	DB T/S AVG
1	1	4	31	32	124
1	2	4	64	64	256
1	3	4	96	97	384
1	4	4	127	128	508
1	5	4	160	161	640
1	6	4	191	192	764
1	7	4	223	225	892
1	8	4	254	257	1016
2	4	4	254	257	1016
2	6	4	378	384	1512
4	4	4	392	407	1568
4	5	4	395	413	1580
4	6	4	390	416 (1)	1560

(1)Notes: The performance limit of the RDBMS has been hit (the throughput upper limit is about 1600 transactions per second (based on the specific transactional mix of the rating business process)).



6.3. Benchmark RUN 2.B

RUN details	
Persistence Providers	1, WLPP01, WLPP02
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
1	1	4	166	214	332
1	2	4	288	342	576
1	3	4	392	464	784
1	4	4	467	543	934
1	5	4	552	639	1104
1	6	4	626	729	1252
1	7	4	709	817	1418
1	8	4	761	914	1522
2	1	4	299	363	598
2	2	4	514	611	1028
2	3	4	695	791	1390
2	4	4	872	1024	1744
2	5	4	1050	1196	2100
2	6	4	1136	1334	2272
2	7	4	1296	1583	2592
2	8	4	1369	1557	2738
3	1	4	405	513	810
3	2	4	723	893	1446
3	3	4	975	1136	1950
3	4	4	1238	1405	2476
3	5	4	1537	1831	3074
3	6	4	1692	1978	3384
3	7	4	1712	2013	3424
3	8	4	1822	2113	3644
4	1	4	553	713	1106
4	2	4	974	1105	1948
4	3	4	1231	1405	2462
4	4	4	1659	1926	3318
4	5	4	1862	2138	3724
4	6	4	2079	2429	4158
4	7	4	2110	2771	4220
4	8	4	2076	2943	4152



RUN Results					
RATING					
NMHOST	NMPROC	NMTHREAD	P/S AVG	P/S MAX	DB T/S AVG
1	1	4	31	32	124
1	2	4	63	63	252
1	3	4	95	96	380
1	4	4	128	128	512
1	5	4	159	160	636
1	6	4	191	192	764
1	7	4	223	224	892
1	8	4	255	257	1020
1	9	4	287	289	1148
1	10	4	319	321	1276
1	11	4	351	352	1404
1	12	4	383	384	1532
1	13	4	415	419	1660
1	14	4	445	450	1780
1	15	4	470	480	1880
1	16	4	495	504	1980
2	10	4	633	643	2532
3	7	4	660	671	2640
3	8	4	710	749	2840
3	9	4	745	787	2980
3	10	4	755	794	3020

This particular benchmark run shows the upper limit of the capacity offered by a single Node Manager host (the fingerprint is HP ML110 specific). This run is well suited to understand vertical scalability (increase of capacity on a single computer) obtainable with the Node Manager process. It is important to note that the capacity shown in the following runs (based on four Node Manager hosts) can be equalled by a domain with 2 node manager computers of the same type.



6.4. Benchmark RUN 3.A

RUN details	
Persistence Providers	1, WLPP01
Partitioning	None
Number of Customer Objects	200,000
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
1	4	4	415	516	830
1	6	4	574	714	1148
2	6	4	1043	1256	2086
4	6	4	1729	1960	3458

RUN Results					
RATING					
NMHOST	NMPROC	NMTHREAD	P/S AVG	P/S MAX	DB T/S AVG
1	1	4	32	32	128
1	2	4	63	64	252
1	3	4	95	96	380
1	4	4	127	128	508
1	6	4	191	192	764
2	4	4	254	257	1016
2	6	4	350	354	1400
4	5	4	397	412	1588
4	6	4	404	422	1616
4	7	4	405	419	1620
4	8	1	388	401 (1)	1552
4	8	2	399	417	1596
4	8	4	405	427	1620

(1)Notes: The performance limit of the RDBMS has been hit (the throughput upper limit is about 1600 transactions per second (based on the specific transactional mix of the rating business process)).



6.5. Benchmark RUN 3.B

RUN details	
Persistence Providers	2, WLPP01 & WLPP02
Partitioning	Horizontal Split
Number of Customer Objects	200,000
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
1	4	4	470	590	940
1	6	4	630	760	1260
2	6	4	1149	1360	2298
4	6	4	1940	2200	3880

RUN Results					
RATING					
NMHOST	NMPROC	NMTHREAD	P/S AVG	P/S MAX	DB T/S AVG
1	4	4	127	130	508
1	6	4	191	193	764
2	4	4	254	259	1016
2	6	4	382	392	1528
4	5	4	634	647	2536
4	6	4	740	766	2960
4	7	4	765	820	3060
4	8	1	585	615	2340
4	8	2	735	780	2940
4	8	4	780	831 (1)	3120

(1)Notes: The performance limit of the RDBMS has been hit (the throughput upper limit is about 1600 transactions per second (based on the specific transactional mix of the rating business process) or 3200 transactions per second partitioned over two physical RDBMS.



6.6. Benchmark RUN 4

RUN details	
Persistence Providers	2, WLPP01 & WLPP02
Partitioning	Horizontal Split
Number of Customer Objects	2,000,000
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
2	6	4	585	678	1170
4	6	4	1010	1248	2020
4	8	4	1110	1445	2220

RUN Results					
RATING					
NMHOST	NMPROC	NMTHREAD	P/S AVG	P/S MAX	DB T/S AVG
2	4	4	255	257	1020
2	6	4	383	386	1532
4	5	4	634	645	2536
4	6	4	745	764	2980
4	7	4	770	805	3080
4	8	4	790	824 (1)	3160

(1)Notes: The performance limit of the RDBMS has been hit (the throughput upper limit is about 1600 transactions per second (based on the specific transactional mix of the rating business process) or 3200 transactions per second partitioned over two physical RDBMS.

7. 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)
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