

WareLite Limited

WareLite Business Operating Support System

Scaling IT for event intensive business processes



Table of Contents

| | | |
|---|---|---|
| 1 | The scalability challenge facing IT | 3 |
| 2 | Introducing WareLite BOSS | 3 |
| 3 | Persistence Management | 3 |
| 4 | Determinism | 4 |
| 5 | Data Partitioning and Distribution | 5 |
| 6 | WareLite BOSS Scalability: benefits..... | 7 |



1 The scalability challenge facing IT

In an increasingly connected world, where not only computers but business processes are networked and more tightly integrated together, having the resilience to cope with shifts in the load on the IT infrastructure is a real challenge. Whether it is caused by an unpredictable surge in demand created by users connected over the Internet, growth in new services such as cellular voice and data traffic, or new technologies like RFID tagging. They all demand increased levels of flexibility and scalability from an IT infrastructure.

The integration of business processes leads to an increasing need to respond quickly, in order to realise the benefits available. It becomes a very difficult challenge to define just how much impact this will have on the IT infrastructure, and so harder to size a sufficiently capable system, without overstretching budgets. Businesses need a solution that provides the flexibility to increase the scale of their investment in proportion with the scale of demand.

2 Introducing WareLite BOSS

WareLite Business Operating Support System (WL BOSS) is a platform for the execution of real time, event driven processes. Each single event - an order, a phone call, an RFID tagged package moving past an antenna - triggers one instance of a specific business process. Technically, an event is an instance of a class of objects and a business process is a sequence of business rules.

Being able to process each single event requires an infrastructure capable of running any process in N parallel threads, where N must be allowed to grow indefinitely to satisfy growing capacity needs. This means that event driven processing requires scalability.

We wanted to obtain what is known as 'horizontal' scalability, as opposed to 'vertical' scalability, i.e. the possibility to obtain higher processing capacity by simply adding computational units to a 'grid' of computers, rather than by porting the business logic (i.e. all the business processes) to a more powerful machine.

The capacity to increase a system's overall performance by adding computational units provides a number of benefits, specifically:
growing capacity, and thus investments, in small incremental steps
delivering high availability – if a unit fails, others will keep the service active
delivering load balancing – as each unit is independent, the first 'idle' unit will drain an event and process it

To make horizontal scalability a viable option, it was necessary to be able to add computational units to a grid without having to modify the business logic.

3 Persistence Management

As the number of computational units grows, the number and organisation of the data objects that are created, accessed and modified by the business



logic will change. To make it possible to execute the same business logic on any number of computational units, with any number and organisation of data objects, it was necessary to create a strong separation between business logic and 'persistence' management. Technically, we define 'persistence' as the capability of propagating data frames to a third party database engine.

WareLite BOSS provides persistence as a service to the whole enterprise business logic. Within WareLite Business Operating Support System (WL BOSS), the built-in persistence service implements abstraction between the business logic layer and a set of third party database engines. This means that all business rules will be able to retrieve, modify and create persistent data without actually having the responsibility to drive a specific database engine.

Once you have declared a class of objects, the WL BOSS persistence service will take care of propagating any new data frame representing an instance of such class to one or more underlying database engines. In addition, WL BOSS makes any third party database engine transparent to the solution layer: a given set of business processes will work without any change with any chosen database engine. This means that business rules, by consuming the interfaces exposed by the WL BOSS run-time, can retrieve, modify and store any object without being aware of how such object is retrieved from or propagated to a database engine.

4 Determinism

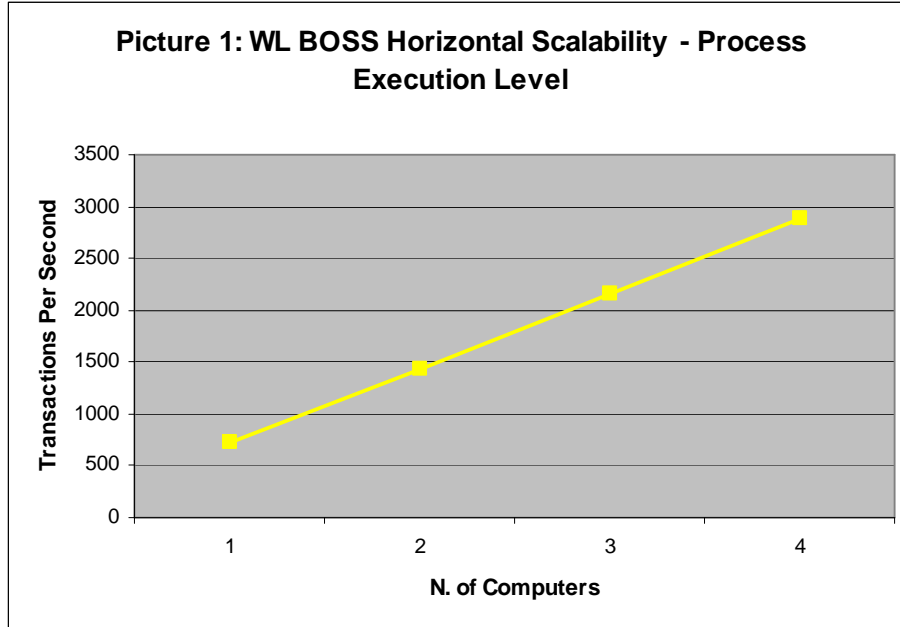
Having thus addressed persistence management, another issue arises: identical and completely parallel business processes, running on different computers, will inevitably try to gain access to the same global data objects at the same time, generating contention.

WareLite BOSS addresses this issue, offering contention resolution ('determinism') as a service to the whole enterprise business logic. So, by adding more computational units to the 'grid', you will not lose determinism: any contention generated by two or more parallel processes will be resolved by the WareLite BOSS Global Lock Management service.

WareLite BOSS therefore provides scalability as a service to the whole enterprise business logic. Pictures 1 and 2 describe the scalability tests executed in our lab.



Picture 1 shows a process - Real Time Rating - requiring the contention resolution services - 'Lock Management' - offered by WareLite BOSS, where



we have achieved a sustainable rate of 720 TPS with Node Managers installed on 1 computer, 1440 TPS on 2 computers, 2160 TPS on 3 computers, 2880 TPS on 4 computers - the growth in throughput is linear. These figures are obtained measuring the transactions committed to a RDBMS. As the rating process comprises 4 transactions, the measured Transactions Per Second equal to 180 Processes/Second on 1 Node, 360 Processes/Second on 2 Nodes, up to 720 Processes/Second on 4 Nodes.

The Node Managers run on computers worth around €500 each – the overall cost of the infrastructure used in our tests, including networking equipment, is around €19,000.

The capacity growth rate remains constant adding up to 52 computers to the grid; after the addition of a 53rd computer, the capacity growth rate for 53 + 1, 2....N computers will decrease, but it will still remain constant.

5 Data Partitioning and Distribution

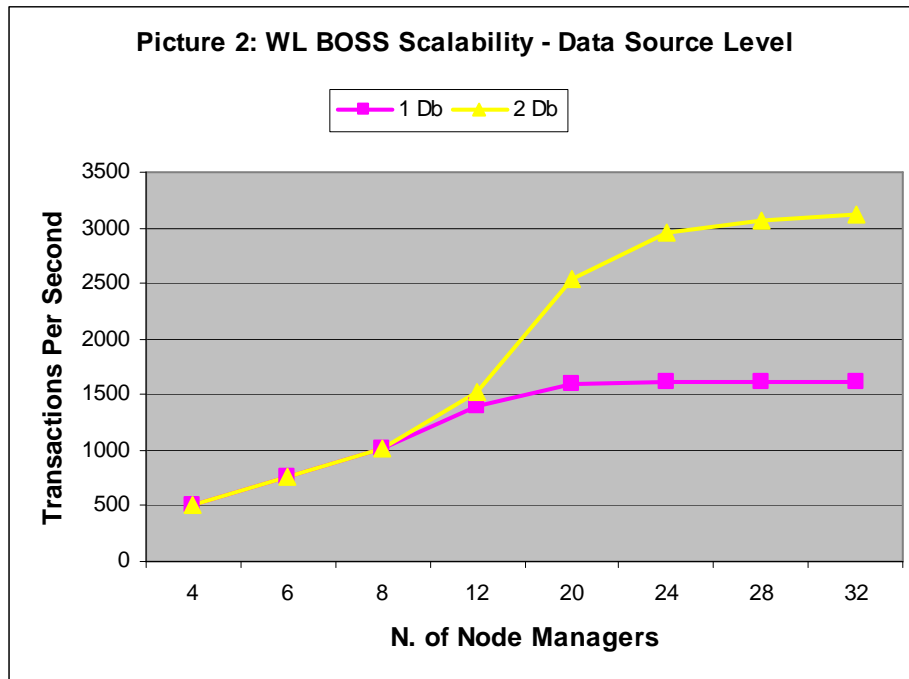
We have discussed Persistence and Determinism as fundamental requirements to achieve horizontal scalability, i.e. to be able to add computational resources to a 'grid', achieving higher capacity without modifying the business logic.

However, just adding node managers - the transactional executive frameworks that take care of the selection of a business process and of its execution - to the 'grid' would not be enough to provide high scalability: without the way to provide high data partitioning, very soon database engines would become a bottleneck not easy to mitigate.



To overcome this issue, the WL BOSS persistence service takes care of data partitioning. In other words, yet again without any impact on the solution layer, data objects can be propagated onto several physical instances of heterogeneous database engines.

Picture 2 shows WareLite BOSS scalability at the Data Source (Persistence Providers) level.



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.

Theoretically, for each declared class of objects, it is possible to deploy up to $N * 16^{32}$ physical database engines, where N is the number of object classes.

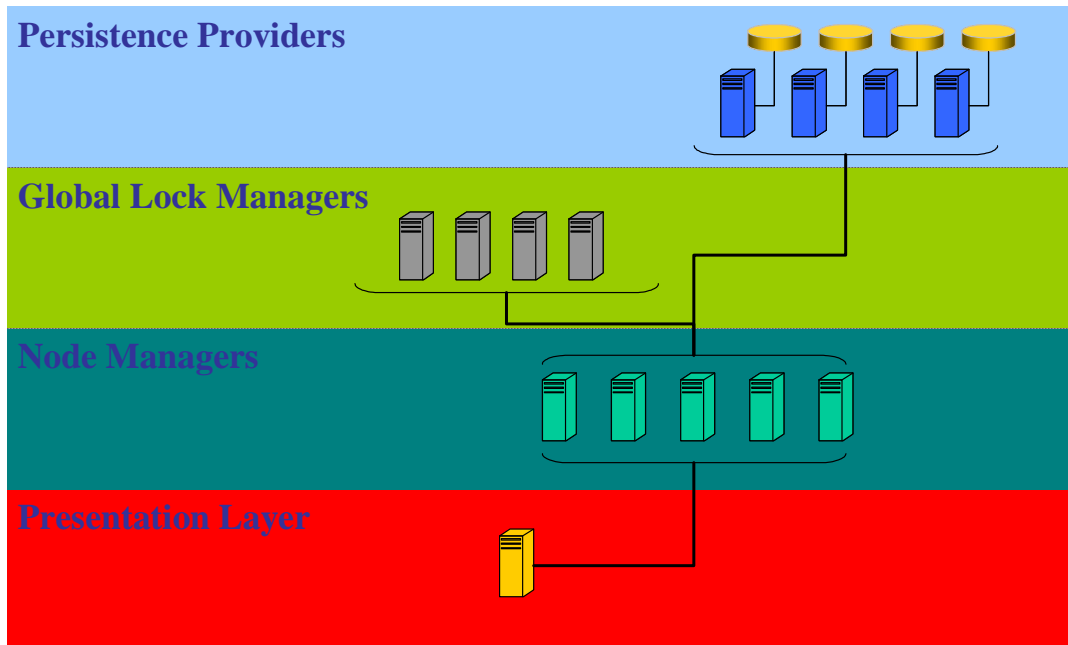
Even in this respect, thanks to the abstraction layer exposed by the WL BOSS run-time, the same business solution, without any modification, can be deployed in a very demanding, high volume environment requiring several physical database engines, or in a less demanding environment, requiring maybe just a single physical database engine.

Our benchmarks prove that WL BOSS provides extreme, scalable OLTP capabilities 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, as shown in Picture 3:

Picture 3: WareLite BOSS Architecture



- 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 databases to the infrastructure;
- Performance is not limited by the presentation layer, as multiple queues can be implemented to present events to the grid.

6 WareLite BOSS Scalability: benefits

WareLite BOSS provides high capacity, scalability and data partitioning at low cost. This is of special relevance in environments generating very high volumes of data and needing both to process them in real time and to aggregate them for reporting purposes, for example Rating & Billing, Service Performance Real Time Monitoring & Reporting, RFID Application Service Provision.

Application service provision requires a flexible approach to both scaling for unpredictable levels of demand, and data separation for speed and integrity of access. For RFID this issue is compounded by the large number of data items being managed and the high frequency of access.



The Data Partitioning and Distribution capabilities make WareLite BOSS able to propagate and access data objects on up to a huge 16^{32} physical database engines. This means that we can provide the level of scalability necessary for an application service provider to globally manage item level data for massive RFID applications. This makes the delivery of multiple large scale RFID projects more straightforward, and therefore more profitable.

Real time processes, such as rating, access database engines in On Line Transaction Processing (OLTP) mode, while data aggregation activities, for example billing, adopt On Line Analytical Processing (OLAP). When data is maintained in a single database, as it often happens, the two different activities compete for access, reducing efficiency. By employing WL BOSS, data can be partitioned and distributed on any number of database engines, eliminating processing bottlenecks.

Currently, the collection of high volumes of data generated as single entities, such as the passage of a data packet on a portion of network, requires the implementation of an *ad hoc* infrastructure to aggregate discrete data in batches and to transfer them to database engines. WareLite BOSS eliminates the requirement for such infrastructure, simply propagating each single data object to the appropriate database, or databases, as it is generated.

The combination of event driven processing with intelligent data storage delivered by the WareLite BOSS platform means that the business logic is insulated from any dependency on how the data is stored. Services can 'push' information to where it can be most effectively used, rather than relying on 'pull' functionality.

Thus, WareLite BOSS represents a unified, single platform for both real time processing and data aggregation in high data volume environments and, as needed by Application Service Providers, for both the provision of revenue generating services in solutions like RFID management and the support of revenue management services for rating and billing.