ORACLE UTILITIES

An Oracle White Paper
November 2012

Oracle Real-Time Scheduler Benchmark

Demonstrates Superior Scalability for Large Service Organizations

ORACLE

Introduction

Large service organizations with greater than 5,000 workers require their scheduling and mobile workforce management applications to perform on two fronts: the building of an optimized work schedule for their workers and the dispatching of activities to the workers. The scheduling of work activities challenges these applications:

- By the sheer volume of activities coming into the application to be scheduled, especially if they are dropped in large batches at a time, or are continuously coming in real-time, including emergencies that must take precedence over less critical work
- By the sheer volume of resources, vehicles, equipment, and depots that can be selected for work activities (see Sidebar – The Scheduling Problem)
- By the number of service territories
- By the need to respond quickly to appointment bookings by customer service representatives (CSRs) talking to customers, or customers directly booking appointments through self service
- By providing CSRs, customers, dispatchers, and other applications timely and accurate work status and alerts from the field

The dispatching of activities to the workers challenges these applications:

- When workers start their shifts at the same time, all needing their activities for the day dispatched to them as fast as possible
- When workers go in and out of network coverage, and needing to send work statuses and receive updates as soon as network connectivity is available

New technologies are dramatically increasing the capability of the service organization to optimize the planning, scheduling, and dispatching of work. As a result, service organization business managers and executives must examine their existing software and hardware infrastructure to ensure that they can process this work and produce results needed to maximize performance. Many of them are concluding that combining best-of-breed applications from different vendors and integrating them on various hardware platforms creates unnecessary cost and risk. They increasingly find that they can achieve better results at lower cost by using software and hardware combinations engineered to work together.

Oracle Real-Time Scheduler software combined with Oracle Exadata (see Sidebar on Exadata) and Oracle Exalogic (see Sidebar on Exalogic) software and hardware is an example of such an engineered system – engineered to work together. Optimized schedules, travel routes, resource cost, and dispatching speed have bottom-line benefits. But because of the volumes of data that must be analyzed concurrently, larger service organizations must also recognize the important benefits of engineered systems that:

- Maximize the availability of the scheduling application with built in redundant database, storage, and connectivity.

- Increase throughput across the middleware and database layers by leveraging the InfiniBand network with Exalogic and Exadata.
- Increase database performance with the use of flash cache that intelligently caches objects in cache and provide an overall gain in read and write operations.
- Reduce IT complexity, lowering the cost to staff and maintain systems.
- Manage the software lifecycle efficiently with the Oracle Enterprise Manager Application Management Pack built for Oracle Real-Time Scheduler.

Because it is critical that service organizations have the assurance that the system they adopt will perform as expected, Oracle has benchmarked realistic volumes to test scalability, reliability, and response rates for this combination of Oracle Real-Time Scheduler, Exadata, and Exalogic.

The Scheduling Problem

Oracle Real-Time Scheduler builds an optimal work schedule based on the requirements and location of the work versus the availability, location, and skills of the workers. The number of scheduling permutations can increase rapidly.

The general calculation is:
1 worker, N jobs = N! possible solutions

Examples:
1 worker, 10 jobs = 3.628×10^{6}
1 worker, 100 jobs = 9.332×10^{157}

With 9,600 workers and 64,000 activities per day, the number of permutations is astronomical.

Exadata

The Oracle Exadata Database Machine is engineered to be the highest performance and most available platform for running the Oracle Database. Built using industry-standard hardware from Sun, and intelligent database and storage software from Oracle, the Exadata Database Machine delivers extreme performance for all types of database workloads including Online Transaction Processing (OLTP), Data Warehousing (DW) and consolidation of mixed workloads. Simple and fast to implement, the Exadata
Database Machine is ready to tackle your largest and most important database applications — and often run them 10x faster, or more.

Exalogic

Oracle Exalogic is an Engineered System, consisting of software, firmware and hardware. It is designed to meet the highest standards of reliability, serviceability and performance under widely varied, performance-sensitive, mission-critical workloads. Because the Exalogic system is fully pre-integrated by Oracle it is also easier to provision, manage and maintain, further reducing ongoing costs and shortening time to value for new projects.

Summary of Results

The tests that Oracle performed demonstrate that Oracle Real-Time Scheduler can optimize the scheduling and dispatching of 1.92 million activities for 9,600 workers, using a quarter rack Exadata X2-2 and a quarter rack Exalogic X2-2 with superior performance standards[1].

Oracle Real-Time Scheduler Performance on Exadata and Exalogic

Average time to	
Schedule an activity	18 seconds
Dispatch activities to a crew	7 - 21 seconds
Respond to an appointment booking	1.08 seconds

Additionally, the test demonstrated that this hardware/software combination provides near-linear scalability. This means that service organizations requiring faster performance, or throughput of increased data volumes, or expanding their services or territories can accomplish that objective through a simple addition of hardware nodes.

[1] Actual results may vary, based on a broad range of implementation-specific factors, such as transaction mix, hardware platform, network parameters, and database size. Oracle does not warrant or guarantee that customers will obtain the same or similar results, even if they use the same or similar equipment and/or software applications. Oracle does not warrant, endorse, or guarantee any performance of any products, any results desired or achieved, or any statements made within this document.

Test Details

Software and Hardware Used in the Tests

The tests used Oracle Real-Time Scheduler (ORS) v2.1.0 installed on a quarter rack Exadata X2-2 server (2 database servers) and a quarter rack Exalogic X2-2 server (2 application servers, 12 mobile data terminal [MDT] simulator servers, and 6 scheduling servers). The test hardware was configured as such based on Oracle's experiences implementing and tuning ORS at customer locations. The operating system was Oracle Enterprise Linux 5 64-bit, the database was Oracle Database Server 11gR2, and WebLogic 10.3.5 64-bit was used.

Data Volume

The tests were designed to be representative of the daily processing needs for a typical large service organization with 9,600 workers, using 9,600 vehicles, with 64,000 activities per day across 144 service territories. The tests simulated the loading, scheduling, and completion of 1.92 million activities over a 30 day period.

In addition, the test used:

- 11.52 million historic activities (to test database performance with 6 months of completed activities)
- Each worker was scheduled for one shift per day, totaling 288,000 shifts over 30 days; each had one break per day, and one period of unavailability per week
- Appointment booking requests were pumped in at a rate of 2,880 per minute
- Emergency / last minute activities were pumped in at a rate of 9.6 activities per minute; emergency and last minute activities require the scheduling engine to re-optimize the current day's schedule to fit in these activities.

The Scheduling and Dispatch Process

The scheduling and dispatch process used the following steps.

1. Created the 11.52 million historical activities.
2. Configured 144 scheduling engines with non-overlapping service territories and scheduling horizon of 30 days*. With the scheduling engines in offline mode, created 9,600 shifts and 1.92 million activities, equally distributed across 144 service territories.
3. Started the scheduling engines.
4. When the scheduling process reached a steady state, 9,600 crew shifts for a day at a time were created and their corresponding mobile devices were started on virtual machines. Activities were automatically dispatched to the virtual mobile devices and in return, they produced the same calls to the server components as would be expected by actual mobile devices in the field.
5. Emergency activities were pumped in at the rate of 9.6 activities per minute.
6. Appointment booking requests were pumped in at a rate of 2,880 per minute.

7. Each dispatched activity was completed on the mobile device, and each shift was closed at the end of the day.
8. Service level agreements and hardware resource utilization were measured.

*Alternatively, overlapping service territories can be configured, but that will increase CPU demand.

The Measurements

The test took several measurements throughout the activity lifecycle, from creation to being dispatched to the mobile device, to the worker going en route to the activity.

1. The time between when the activity was created and when the activity was queued for dispatch.
2. The time between when the activity was queued for dispatch and when the activity reached the mobile device.
3. The time between when the activity was created and when the activity reached the mobile device.
4. The time between when the field resource indicated on the mobile device that they were en route and when that status reached the application server.

Oracle Real-Time Scheduler Test: Activity Lifecycle

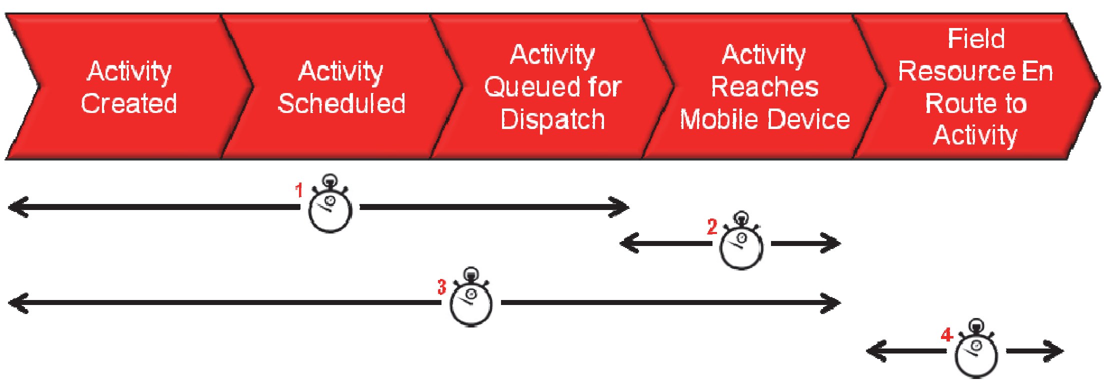

In addition, the following items were measured:

- The response time of an appointment booking request.
- The time to recall activities from the mobile device (in case the host system cancels the activities).
- The time to refresh the dispatcher's Common Dispatching Interface (CDI).
- The time to generate and display an alert in the CDI.
- The time for the application to terminate.

Results

The large service organization tests demonstrated excellent performance for the scheduling and dispatching of activities using a quarter rack Exadata X2-2 server (2 database servers) and a quarter rack Exalogic X2-2 server.

During the design of the tests, target performance goals for each measurement were established based on the experiences of our industry specialists and consulting personnel working with existing customers. The actual results for various measurements taken, including the measurements listed above, were as follows:

MEASUREMENT (ELAPSED TIME)	TARGET PERFORMANCE	ACTUAL RESULT
(1) Activity queued for dispatch	<=30 sec	9.6 seconds
(2) Activity dispatched	<=1 Min	7-21 seconds
(3) Activities reached the mobile device	<=1 Min	26.82 seconds
Mobile device synchronized to application server	<=1 Min	1 seconds
(4) Activity status update from field received in CDI	<= 2 CDI refresh cycles*	1 CDI refresh cycle*
Appointment booking response received	< 3 Sec	1.083 seconds
Activities recalled from mobile device	<=1 Min	6-9 seconds
CDI refresh	2 sec	1.6 seconds
Alert generated and displayed in CDI	<= 1 CDI refresh cycle*	1 CDI refresh cycle*
Application is terminated	<=30 sec	1 second

* The CDI refresh cycle was configured at 60 seconds for the test.

Because each service organization brings a unique combination of the number of resources, sizes and numbers of service territories, amount of work to be scheduled, scheduling horizon, and performance goals, Oracle works with each implementation to recommend the appropriate hardware and software configuration and to fine tune it to achieve the organization's optimal scheduling and dispatching results.

Conclusion

Through these benchmark tests, Oracle has demonstrated superior results which far exceeded the standards set for "acceptable performance" for receiving, scheduling, dispatching, and completing activities as per the daily load of a typical large service organization. In addition, the tests demonstrated near-linear scalability to address any size organization. The testing methodology included scenarios to simulate extreme utilization of the system; including processes that are traditionally bottleneck areas such as when workers all log on at the beginning of the day or when activities are dropped into the scheduling engine in large batches.

Oracle is committed to developing practices and products that help protect the environment

Oracle Utilities Mobile Workforce Management Benchmark
November 2012

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com

Hardware and Software, Engineered to Work Together