Performance brief for Oracle Enterprise Financial Management 8.9 (Order-to-Cash Counter Sales) on HP Integrity BL870c server blades

Executive summary

In January 2008, HP Alliances engineers completed a performance benchmark test to characterize the scalability of Oracle's® PeopleSoft Enterprise Financial Management 8.9 solutions. The specific business process simulated by the benchmark was Order-to-Cash, the aggregation of Order Management, Inventory, and Billing services. The batch processes in the test represent a typical daily workload for a large retail customer.

The objective was to measure the performance of HP's new Integrity BL870c blade servers for this class of application. The complete test scenario included an initial set of parallel jobs (illustrating the multi-core scaling of the server under test) followed by several single-threaded jobs (illustrating consistent single-core processing capacity). Seeding the Financial Management instance with the orders processed by the batch jobs was done by HP LoadRunner test scripts; results for the LoadRunner activity, while not strictly part of the benchmark definition from Oracle, are reported here as well.

The workload was initially designed for a much larger server than the Integrity BL870c blade server. For that reason, tests were done on both a single-server configuration (where both the PeopleSoft Enterprise batch server processes and the Oracle RDBMS database instance are deployed on the same server) and a split-server deployment (where the batch and RDBMS components reside on separate servers). The tables below provide a summary of the observed results. Detailed results are shown in the section Test results on page 7.

A cursory summary of results can be found in the following tables:

Table 1. Summary of batch results (single-server configuration)

Scenario	Processing time (parallel jobs)	Processing time (single threaded)	Throughput
250,000 order lines (5 parallel streams)	84.20 min	13.65 min	153,296 lines/hr
500,000 order lines (10 parallel streams)	108.00 min	23.30 min	228,484 lines/hr
750,000 order lines (15 parallel streams)	160.93 min	38.15 min	226,040 lines/hr

Table 2. Summary of batch results (split-server configuration)

Scenario	Processing time (parallel jobs)	Processing time (single threaded)	Throughput
250,000 order lines (5 parallel streams)	86.10 min	12.90 min	151,515 lines/hr
500,000 order lines (10 parallel streams)	107.85 min	26.30 min	223,630 lines/hr
750,000 order lines (15 parallel streams)	148.28 min	38.90 min	240,410 lines/hr

Table 3. Overview of on-line performance

Scenario	Virtual Users	Transaction Rate (cpu utilization)	Average Response Time (enter/save 5-line order)
Single-server deployment (web/app/db on a single 8-core server)	250	75 txn/min (95%)	11.61 sec
Dual-server deployment (8-core web/app server with 8-core db server)	250	80 txn/min (90%)	10.49 sec
Multi-server deployment (2 web/app servers with a single db server)	500	160 txn/min (90%)	13.14 sec

Target audience: The intended audience for this performance report consists of IT professionals interested in the application characteristics and performance requirements of an Oracle Financial Management 8.9 implementation on HP Integrity Blade servers. The Oracle Enterprise IT community may find the tested configuration useful as a reference architecture for Financial Management 8.9.

Test methodology

The Order-to-Cash workload models the processing of several hundred thousand order lines seeded into the test database. The seeded transactions represent a typical single-day counter sales load for a large retail customer. For batch processing of counter sales orders, traditional inventory operations (reserve, pick, and confirm) are unnecessary, as they are implemented as part of the original counter sale transaction. Figure 1 below details the process flow within the Oracle Enterprise deployment. Batch processes are launched via the web-based PeopleSoft Internet Architecture (PIA) interface and require no operator intervention after they begin. The entire job flow (consisting of the parallel streams followed by the single-stream billing interface activities) was encapsulated as a single PeopleSoft JobSet.

Figure 1. Process Flow

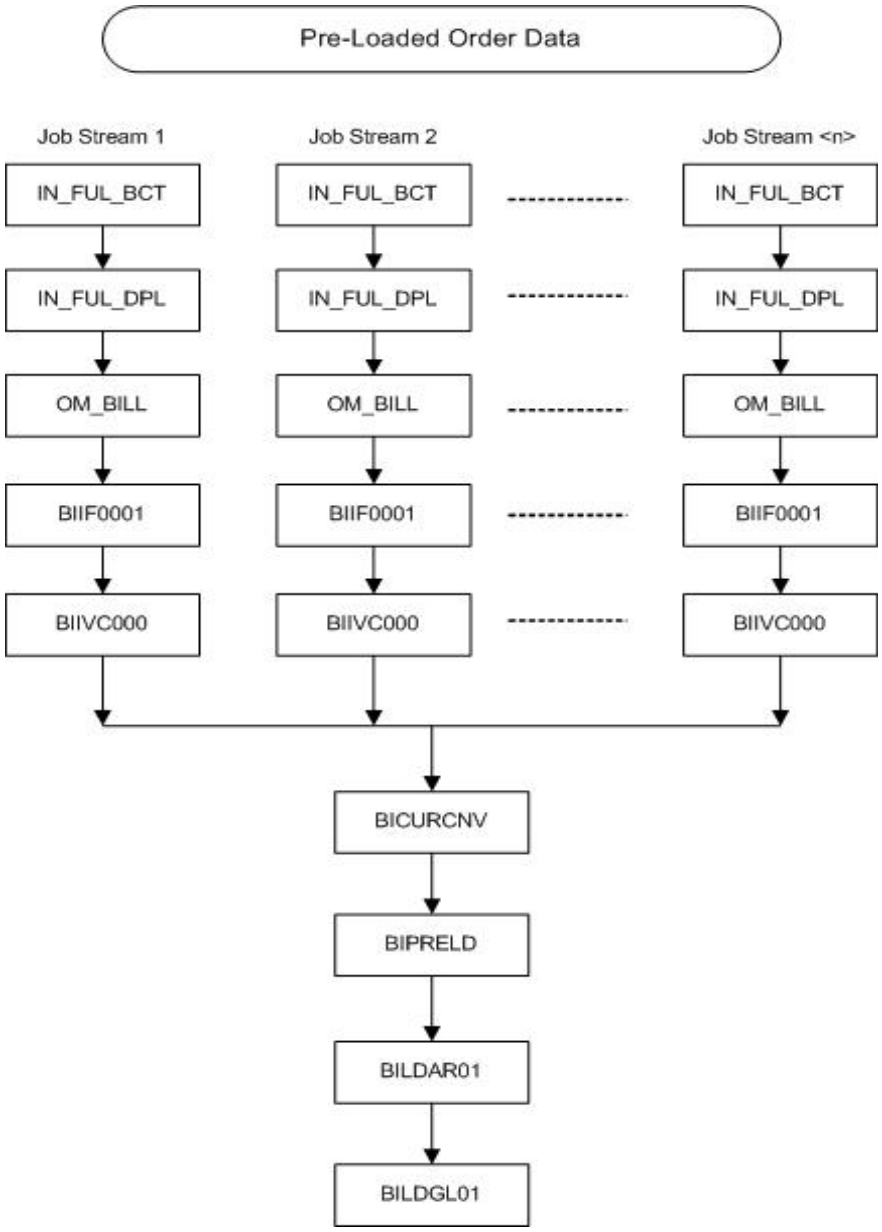

The initial order management and inventory processes can be performed in parallel, modeling individual business units in a single company. The last four billing interface processes, however, must be run singly against the corporate-wide data set to avoid contention. The partitioning of the order data and the specifics of the billing interface processes may be different for real-world customers.

The business processes implemented by the batch jobs can be divided into two categories:

Inventory

Fulfillment Requests (IN_FUL_BCT): Completes the confirmation and shipping processes for individual order lines entered as part of the Counter Sales transactions.

Depletions (IN_FUL_DPL): Depletes the shipped quantities from inventory.

Billing

Populate Billing (OM_BILL): Passes the shipping information on to the billing tables (in preparation for invoicing the orders).

Billing Interface (BIIF0001): Creates the invoice headers and lines, along with the account distributions.

Invoice Finalization (BIIVC000): Creates the final invoices based on all the shipped orders.

Currency Conversion (BICURCNV): Performs any required currency conversions on the generated invoices.

Preload and Load into AR (BIPRELD and BILDAR01): Transfer the invoices into the accounts receivable system.

Load GL Interface (BILDGL01): Loads the General Ledger system with the accounts receivable data from the previous steps.

Test hardware

The Integrity BL870c blade servers under test were deployed in the configuration shown in Figure 2. For the single-server tests, both the Enterprise batch instance and the Oracle database instance were on a single Integrity BL870c blade server. The split-server tests moved the batch instance to a second server, leaving the Oracle database to consume all the resources of a single server. The three HP ProLiant BL480c servers were used only during the LoadRunner activity to seed the initial order data.

Figure 2. Hardware layout

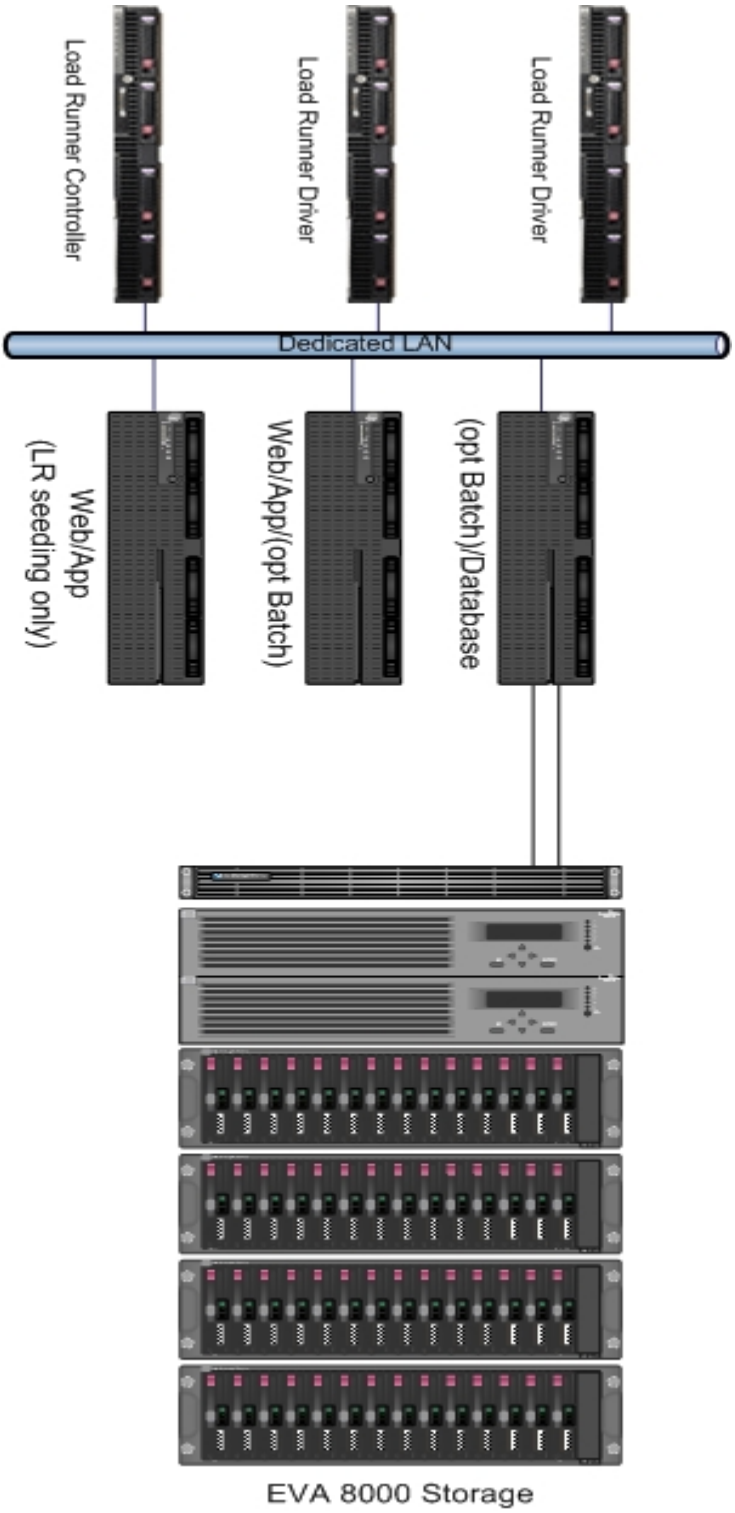

The logical topology of the software deployment is reflected in the hardware layout. For the initial seeding activity, the 2 LoadRunner driver systems interactively added orders via the two Integrity BL870c Web/App server systems. No batch processing was performed during the seeding activity. A third Integrity BL870c server was configured as the database server with Oracle RDBMS 10gR2.

Storage for the database server was provided by an HP StorageWorks 8000 Enterprise Virtual Array (EVA 8000) storage array. For the benchmark runs, the batch server components were deployed either on the database server (the "single-server" configuration) or on one of the web/app servers (the "split-server" configuration). A complete description of the Oracle Enterprise software stack can be found in *Appendix A: Software stack and hardware configurations*.

Important
Care has been taken to perform a test that includes a representative set of business units and inventory locations. As with any laboratory testing, however, the performance metrics quoted in this paper are idealized. In a production environment, these metrics may be impacted by a variety of factors.

HP recommends proof-of-concept testing in a non-production environment using the actual target application as a matter of best practice for all application deployments. Testing the actual target application in a test/staging environment identical to, but isolated from, the production environment is the most effective way to estimate system behavior.

Test results

The tables below summarize the results of the different batch process loads. All processing times are given in minutes. For the parallel jobs, the range of times given shows the minimum and maximum time consumed by the parallel streams. The total time given at the end represents the time from the earliest start of the first process to the completion of last process (and thus may be slightly different than the sum of the different timings).

Table 4. Batch results for varied workloads (single-server configuration)

Process	5 streams (250,000 lines)	10 streams (500,000 lines)	15 streams (750,000 lines)
IN_FUL_BCT	35.97 - 36.73	45.28 – 48.95	68.68 – 71.93
IN_FUL_DPL	20.48 – 21.30	25.20 – 26.78	36.78 – 40.38
OM_BILL	13.17 – 14.48	17.25 – 19.12	22.77 – 30.95
BIIF0001	11.63 – 11.65	13.70 – 15.85	21.60 – 26.90
BIIVC000	1.27 – 1.28	1.53 – 1.63	2.30 – 2.85
BICURCNV	3.07	3.78	5.82
BIPRELD	3.28	6.33	10.10
BILDAR01	1.83	3.77	6.37
BILDGL01	5.40	10.35	15.88

Process	5 streams (250,000 lines)	10 streams (500,000 lines)	15 streams (750,000 lines)
Total	97.85	131.30	199.08
Hourly throughput	153,296	228,484	226,040

Table 5. Batch results for varied workloads (split-server configuration)

Process	5 streams (250,000 lines)	10 streams (500,000 lines)	15 streams (750,000 lines)
IN_FUL_BCT	37.17 - 37.70	43.38 – 45.22	62.05 – 65.55
IN_FUL_DPL	20.73 – 21.30	25.50 – 26.83	36.83 – 38.42
OM_BILL	14.18 – 14.43	18.13 – 20.52	21.15 – 24.97
BIIF0001	12.13 – 12.13	14.98 – 15.78	18.67 – 21.12
BIIVC000	1.27 – 1.27	1.78 – 1.82	1.52 – 2.33
BICURCNV	1.77	3.78	5.55
BIPRELD	3.03	6.83	9.87
BILDAR01	2.48	4.78	7.30
BILDGL01	5.53	10.85	16.13
Total	99.00	134.15	187.18
Hourly throughput	151,515	223,630	240,410

System utilization data was collected using sar(1M) and HP Performance Manager software during all the test runs. That data is included in *Appendix B: CPU Utilization*.

Not surprisingly, the single-stream billing interface processes show linear scaling of duration as a function of processed order lines. The one exception appears to be the BICURCNV process, which behaved sporadically in the single-server tests. Overall throughput scales well from 5-streams to 10-streams, and then flattens out with 15-streams as the 8-core servers are saturated. The saturation effects are less noticeable for the split-server configuration, where the Enterprise batch processes have the additional headroom of the second 8-core server.

Key findings

This benchmark illustrated several key advantages of the HP Integrity BL870c blade servers:

- The Integrity BL870c blade server demonstrates excellent scaling of Oracle Enterprise and RDBMS workloads (both within a single server and between multiple servers).
- The HP C-class blade servers are easier to upgrade, patch, and maintain than traditional server deployments. The familiar HP-UX tools of Ignite-UX and HP Performance Manager have been enhanced with the blade infrastructure management features.
- The HP BladeSystem infrastructure allows for more efficient configuration of network and fibre-channel based storage resources. More importantly, the inherent redundancy in these shared components makes deployment of highly available environments (e.g., Oracle RAC or BEA WebLogic Clusters) less expensive.
- Under CPU saturation conditions at the database tier, disk I/O behavior becomes more critical. Testing of additional storage configurations (not fully documented here for space considerations) showed some interesting differences between file-system-based and raw logical volume based layouts for the database instance. Standard VxFS 4.1 file systems showed marginally better performance for the 5-stream run and only a 5% performance degradation for the 10-stream run. However, using file system storage for the database instance in the 15-stream run resulted in 20% less throughput than when raw logical volume storage was used. Users who wish the flexibility of file system storage with the performance of raw logical volumes should consider the Veritas Oracle Disk Manager product from HP.

Recommendations

HP engineers have extensive experience with Oracle Enterpise applications and Oracle RDBMS. Applying that experience to the HP Integrity BL870c blade platform used in this performance benchmark, it is possible to develop a few basic recommendations for customer deployments:

- Sizing the batch and database server tiers for an Oracle Enterprise deployment can be as simple as identifying the maximum number of parallel batch jobs and the time constraints for job completion.
- The application server tier achieved linear scalability during the LoadRunner seeding activity. The relatively high transaction rates used to minimize benchmark setup time still demonstrated a 2-to-1 ratio of web/app tier CPU's to database tier CPU's. A more typical customer deployment could likely increase that ratio to 4-to-1. The key to tuning the Enterprise application tier is to configure the BEA WebLogic instance with sufficient Java Heap and Execute Threads.
- The web server component of the Oracle Enterprise deployment (the BEA WebLogic instance) consumes relatively few system resources. That component can share the platform with the Enterprise Application server or be deployed on smaller servers to allow for network isolation and security.

Appendix A: Software stack and hardware configurations

Software stack

- HP-UX 11i v2 (11.23), December 2007 Quality Pack
- Oracle's PeopleSoft Enterprise Financial Management 8.9 MP2
- Oracle's PeopleSoft Enterprise PeopleTools 8.47.08
  - BEA Tuxedo 8.1 RP 135 with JOLT 8.1
  - BEA WebLogic 8.1 SP5
- Oracle RDBMS10*g*R2 Enterprise Edition (10.2.0.3 64-bit) {database tier}
- Oracle RDBMS 10*g*R2 Client (10.2.0.3) {app tier}
- HP LoadRunner 8.1

Hardware configuration

10gR2 database server

HP Integrity BL870c blade server, configured as follows:

- Four Intel® Itanium2® Dual-Core Processors Model 9140N, each with 64KB L1 cache, 2MB L2 Instruction Cache, 512KB L2 data cache, and 18MB L3 cache; Hyper-Threading disabled
- 32GB RAM
- Two internal 72GB SAS (Serial Attached SCSI) disk drives
- 2 dual-port 1Gb PCI-X Ethernet host bus adapters
- 1 dual-port 4Gb PCI-e Fibre-channel host-bus adapter

HP StorageWorks EVA 8000 storage array

- Two HSV210 array controllers with 2GB control cache and 2GB data cache each
- 140 72GB 15K RPM fibre-channel hard disk drives (~9TB raw storage available)
- Database utilization: 512 GB Raid 1+0 storage (~1TB of system's raw capacity)
  - database instance used raw logical volumes for tablespaces and archive logs

Financial 8.9 batch server

HP Integrity BL870c blade server, configured as follows:

- Four Intel Itanium2 Dual-Core Processors Model 9140N, each with 64KB L1 cache, 2MB L2 Instruction Cache, 512KB L2 data cache, and 18MB L3 cache; Hyper-Threading disabled
- 32GB RAM
- Two internal 72GB SAS (Serial Attached SCSI) disk drives
- 2 dual-port 1Gb PCI-X Ethernet host bus adapters

Financial 8.9 auxiliary web/application server

HP Integrity BL870c blade server, configured as follows:

- Four Intel Itanium2 Dual-Core Processors Model 9140N, each with 64 KB L1 cache, 2 MB L2 Instruction Cache, 512 KB L2 data cache, and 18 MB L3 cache; Hyper-Threading disabled
- 32GB RAM
- Two internal 72GB SAS (Serial Attached SCSI) disk drives

Appendix B: CPU utilization and disk I/O

The following tables summarize the sar(1M) data collected during the batch test runs. In conjunction with the timing data presented above, they offer insight into the scalability of Oracle Enterprise batch processes and Oracle RDBMS on HP Integrity blade servers. Key details to note are as follows:

- The high percentage of usr time indicates that the HP-UX operating system is imposing minimal resource constraints on the application.
- Each job stream consumes about one core's worth of system processing capacity. This clearly illustrates the nature of Oracle Enterprise batch jobs: a single process on the Batch Server system bound to an Oracle shadow client on the database server. The inter-process workflow is such that the batch job makes a request to the shadow client, waits for the requested data, and then leaves the shadow client idle while it processes the data.
- Under saturation conditions (15 jobs streams), the split-server configuration performs slightly better than the single-server configuration during the parallel portion of the run. This slow-down is counter-balanced somewhat by slightly higher efficiency during the single-threaded portion.

Table B1. Single-server CPU utilization for 5-stream run (250,000 order lines processed)

Process	**Unified Batch/DB**		
	usr%	**sys%**	**wio%**
IN_FUL_BCT	60	1	2
IN_FUL_DPL	58	1	4
OM_BILL	53	5	3
BIIF0001	52	3	2
BIIVC000	49	6	5
BICURCNV	11	0	1
BIPRELD	12	0	0
BILDAR01	11	2	0
BILDGL01	12	0	0

Table B2. Single-server CPU utilization for 10-stream run (500,000 order lines processed)

Process	Unified Batch/DB		
	usr%	sys%	wio%
IN_FUL_BCT	96	2	1
IN_FUL_DPL	95	1	3
OM_BILL	81	8	1
BIIF0001	84	6	1
BIIVC000	74	7	2
BICURCNV	11	0	2
BIPRELD	12	0	0
BILDAR01	11	2	0
BILDGL01	12	0	0

Table B3. Single-server CPU utilization for 15-stream run (750,000 order lines processed)

Process	Unified Batch/DB		
	usr%	sys%	wio%
IN_FUL_BCT	97	2	1
IN_FUL_DPL	97	1	2
OM_BILL	88	8	1
BIIF0001	92	5	1
BIIVC000	81	7	1
BICURCNV	11	0	2
BIPRELD	13	0	1
BILDAR01	11	2	0
BILDGL01	13	0	0

Table B4. Split-server CPU utilization for 5-stream run (250,000 order lines processed)

Process	Batch			Database		
	usr%	sys%	wio%	usr%	sys%	wio%
IN_FUL_BCT	5	1	0	54	1	2
IN_FUL_DPL	1	0	0	57	1	4
OM_BILL	19	4	0	31	2	3
BIIF0001	8	2	0	42	2	2
BIIVC000	11	4	0	37	3	1
BICURCNV	0	0	0	11	0	0
BIPRELD	1	0	0	11	0	0
BILDAR01	4	1	0	4	1	0
BILDGL01	9	0	0	3	0	0

Table B5. Split-server CPU utilization for 10-stream run (500,000 order lines processed)

Process	Batch			Database		
	usr%	sys%	wio%	usr%	sys%	wio%
IN_FUL_BCT	9	1	0	92	1	3
IN_FUL_DPL	1	0	0	95	1	3
OM_BILL	30	6	0	52	3	9
BIIF0001	14	4	0	74	3	4
BIIVC000	21	7	0	62	5	1
BICURCNV	0	0	0	11	0	2
BIPRELD	1	0	0	12	0	0
BILDAR01	5	1	0	5	1	0
BILDGL01	9	0	0	3	0	0

Table B6. Split-server CPU utilization for 15-stream run (750,000 order lines processed)

Process	Batch			Database		
	usr%	sys%	wio%	usr%	sys%	wio%
IN_FUL_BCT	9	1	0	96	1	1
IN_FUL_DPL	1	0	0	97	1	2
OM_BILL	38	7	0	60	3	9
BIIF0001	24	7	0	80	3	4
BIIVC000	17	6	0	66	5	12
BICURCNV	0	0	0	10	0	3
BIPRELD	1	0	0	12	0	0
BILDAR01	5	1	0	5	1	0
BILDGL01	9	0	0	2	0	0

This last table summarizes the I/O statistics for database server under low (5-stream) and high (15-stream) load. The overall rates show the outstanding performance of the StorageWorks EVA 8000 storage array when coupled with the fibre-channel interconnect of the C-class blade enclosure chassis.

Table B7. Disk I/O rates for database server during 5-stream and 15-stream runs

	5-stream (250,000 order lines)		**15-stream run (750,000 order lines)**	
	DB Tables/Indexes	**DB Logs**	**DB Tables/Indexes**	**DB Logs**
I/O per sec (average/**peak**)	464 / **5,666**	96 / **837**	1,006 / **9,702**	156 / **899**
Read/Write ratio of I/O requests	0.97	N/A	0.73	N/A
Read rate (kb/sec) (average/**peak**)	3,760 / **89,243**	N/A	4,982 / **80,319**	N/A
Write rate (kb/sec) (average/**peak**)	6,684 / **114,195**	16,624 / **134,873**	12,751 / **146,318**	18,684 / **171,850**

For more information

More information on HP/Oracle benchmark efforts is available at
http://www.hporacleapps.com/applications/index.asp?app=Enterprise

Information on HP BladeSystem servers and solutions can be found at www.hp.com/go/bladesystem.

For details on the complete range of HP Integrity servers, please visit www.hp.com/go/integrity.

To help us improve our documents, please provide feedback at www.hp.com/solutions/feedback.

4AA1-7863ENW, Revision 1, February 2008