Mercury IT Governance Center™

Mercury Demand Management™ Benchmark for Windows 2000 Server

Executive Summary

This document provides a benchmark that characterizes the performance of Mercury Demand Management™ in a Windows 2000 Server environment. Mercury LoadRunner® version 8.0 was used to simulate a high volume of various request types submitted by non-IT users, IT users, and IT managers.

Test results confirm the scalability of Mercury Demand Management in support of large numbers of concurrent users and transaction volumes.

Other Mercury IT Governance Center™ products such as Mercury Program Management™ leverage the same forms, workflows, and Mercury IT Governance Dashboard™ technologies as Mercury Demand Management. However, these other products usually operate with much lower transaction volumes and fewer concurrent users, so their performance typically exceeds that of Mercury Demand Management.

MERCURY™

Introduction

Mercury has conducted tests that establish a benchmark for performance characterization of Mercury Demand Management and confirm its scalability. Mercury LoadRunner® version 8.0 testing software was used to simulate many concurrent users and high transaction volumes.

Types of Users Simulated in Benchmark

The LoadRunner scripts simulated a business environment in which three types of users—non-IT users, IT users, and IT managers—use Mercury Demand Management to manage business and IT requests.

- **Non-IT users.** Non-IT users typically create and submit requests that IT users will be assigned to work on. Non-IT users log on only when necessary and do not remain logged on. Typically, non-IT users respond to notifications (emails) to process requests, but they also log on and then search requests they created or view the My Requests portlet.

  The Mercury IT Governance Dashboard of every non-IT user is configured to show the My Requests portlet, as well as a Request List (Expanded) portlet with a filter for the Created By field.

- **IT users.** IT users work in IT and typically respond to incoming requests and process them to closure. IT users use the Dashboard throughout the day to monitor their work. They typically log on once and remain logged on until the end of the day.

  Every IT user's Dashboard is configured to show Request List portlets for Benchmark-Simple requests and Benchmark-Complex requests. These request types are explained later. Each portlet contains a filter for the Status and Creation Date fields.

- **IT managers.** Managers in the IT organization are not as directly involved as IT users in processing individual requests. Instead, they monitor the overall status of requests on a regular basis throughout the day.

  The Dashboard for IT managers contains summary portlets for Request Activity, Open Requests by Priority, and Request Summary Bar Chart.

Business Processes Simulated in Benchmark

LoadRunner scripts were run for 3.5 hours, ramping up to a peak load of 302 simulated users concurrently using Mercury Demand Management to create, search, monitor, and process requests. These activities are used in a variety of business processes as described in the following sections. LoadRunner simultaneously ran multiple instances of each business process (script), where each instance represented one active user. (The number of concurrent users is usually less than 10 percent of the number of named users in real-world environments.)

Based on interviews with customers, Mercury categorizes request types and workflows, and thus benchmarks, as simple or complex.

In the LoadRunner simulation, a simple benchmark includes:

- A request type containing 50 fields. Rules are not used for field interdependencies.
- A workflow of 18 steps with no subworkflows. It takes approximately five steps (four approvals and one execution) to close a request.

In the LoadRunner simulation, a complex benchmark includes:

- A request type containing more than 200 fields. Rules are used to drive field interdependencies.
- A workflow of 75 steps, some of which are in subworkflows. It takes approximately 25 steps (18 approvals and 7 executions) to close a request.

Create Simple Request

This business process simulates a non-IT user submitting a new request. The user logs on, creates a request using the Benchmark-Simple request type, enters data for some fields of the request, lets auto-complete fill out the rest, attaches a 1 KB document to the request, and submits the request. Then the user opens the request. Finally, the user logs off.

Create Complex Request

This business process simulates a non-IT user submitting a new request. The user logs on and creates a request using the Benchmark-Complex request type. No document is attached.

The complex request has a field set consisting of six fields that repeat eight times. For 40 percent of the requests, only one field set (out of eight) is assigned data when the request is saved. For another 40 percent of the requests, two field sets are assigned data. For the remaining 20 percent of the requests, all eight field sets are assigned data. After the user submits the request, the user logs off.

Search Requests by IT User

This business process simulates an IT user searching for requests. It does not include logon and logoff, because an IT user typically remains logged on during the day.

For 40 percent of the requests, the user searches by the Assigned To field. For another 40 percent of the requests, the user searches by Status. For the remaining 20 percent of the requests, the user searches by Request Type. The user opens a request at random from the search results.

Search Requests by Non-IT User

This business process simulates a non-IT user searching for requests. It begins by logging on the user. In the simulation, the non-IT user always searches by the Created By field to check on the status of the requests that user created. Finally, the user opens a request at random from the search results, and logs off.

Search Requests (Common)

This business process simulates IT users as well as non-IT users searching for requests. For two-thirds of the requests, the user searches for a known request by request number, opens that request, and remains logged on. For the other one-third of the requests, the user searches by date, opens a request at random from the search results, and logs off.

Process Request

This business process simulates the processing of a request by the IT user to whom it is assigned. The IT user remains logged on, processing multiple requests during the benchmark run. The IT user opens a request at random from the My Requests portlet, then selects at random from the list of possible workflow actions associated with that request, and then updates the request. Eventually, in following the workflow, the request is closed.

Process Request by Smart URL

This business process simulates a non-IT user receiving a notification from Mercury IT Governance Center and clicking on a smart URL (a URL that opens a particular request in the browser). The business process includes logging on and off, since a non-IT user is not expected to remain logged on.

The actions are similar to the Process Request business process. However, before the benchmark run is started, a data file is created with a list of requests that are currently open and awaiting action by a non-IT user. These requests are processed as part of the benchmark by the simulated users.

View Summary Portlets

This business process simulates IT managers logging on and viewing their Dashboards containing summary and graphical portlets in order to check the overall status of requests. Dashboards of IT managers have more complex portlets and therefore take longer to open than Dashboards of other users.

Mix of Scripts Used in Benchmark

For each business process script, *Table 1* shows the number of simulated users running that script and the LoadRunner pacing—the average time between starting up one script of that type and starting up the next script of that type, for each user.

For example (referring to the first table row), on average, every 5 minutes LoadRunner starts a Create Simple Request script for each of the 50 simulated users associated with this script. For the last table row, on average, every 15 minutes LoadRunner starts a View Summary Portals script for each of the 30 simulated users associated with this script.

Table 1. Mix of scripts used in benchmark

Business Process Script	Number of Simulated Users (Total of 302)	Pacing (Time Between Scripts for Each User)
Create Simple Request	50	5 minutes
Create Complex Request	30	5 minutes
Search Requests by IT User	25	5 minutes
Search Requests by Non-IT User	10	5 minutes
Search Requests (Common)	25	5 minutes
Process Request	120	3 minutes
Process Request by Smart URL	12	3 minutes
View Summary Portals	30	15 minutes

Mercury IT Governance Center Benchmark Environment

This section provides details about the environment created for the benchmarking effort. *Figure 1* illustrates the environment.

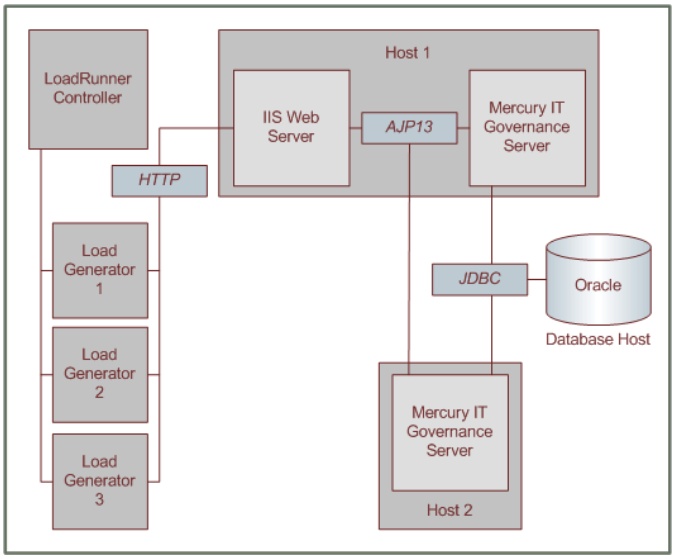

Figure 1. Mercury IT Governance Center environment for benchmarking

- **Hosts with Mercury IT Governance Servers.** Two "clustered" host machines designated Host 1 and Host 2, each with:
  - One 3.06 GHz Intel Pentium processor with Hyper-Threading Technology
  - 2 GB of RAM
  - Microsoft Windows 2003 Server operating system
  - One instance of the Mercury IT Governance Server (application server), with the default `server.conf` file, including 550 MB of heap for the JVM
- **Web Server.** The Host 1 machine also runs the Microsoft IIS 6.0 Web server, which uses the Apache Jakarta JK2 plug-in to balance the load between the two Mercury IT Governance Servers. The Web server communicates with the Mercury IT Governance Centers using the AJP13 protocol.

- **Database.** The database runs on a dedicated host machine and communicates with both Mercury IT Governance Servers using the JDBC protocol. The host machine and database include:
  - Dual 2.8 GHz Intel Xeon processors with Hyper-Threading Technology
  - 4 GB of RAM
  - Microsoft Windows 2003 Server operating system
  - Oracle version 9.2.0.5
  - Configuration parameters as shown in *Appendix A: Database Configuration Parameters* on page 11
  - To better simulate a business environment, before the benchmark scripts were run, the database was initialized with data having the following characteristics:
    - A total of 72,000 open and closed requests
    - 80 percent of all requests of type Benchmark-Simple
    - 20 percent of all requests of type Benchmark-Complex
    - 90 percent of all requests closed
    - 10 percent of all requests open
    - A total of 35,000 named users
- **Client Browser.** Consistent with the default settings of Microsoft Internet Explorer 6.0 SP1, the simulation assumed that the browser:
  - Has caching turned on
  - Supports GZIP encoding, where the application server compresses HTTP replies and thereby reduces network bandwidth utilization

Benchmark Results

Results are shown here for a total period of 3.5 hours. Discussions focus on the peak load period that starts one hour after ramp-up began and ends one hour later, in other words, from 1:00 to 2:00 on the time line (horizontal axis) of each graph.

Figure 2 shows that, as the number of users ramps up to 302, the database CPU utilization grows linearly, then peaks at 62 percent.

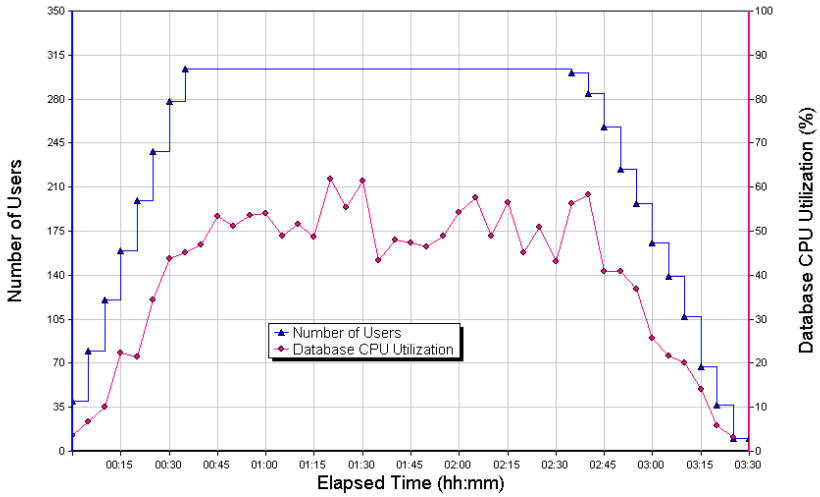

Figure 2. Number of users and database CPU utilization

Figure 3 shows CPU utilization for the application servers and the database.

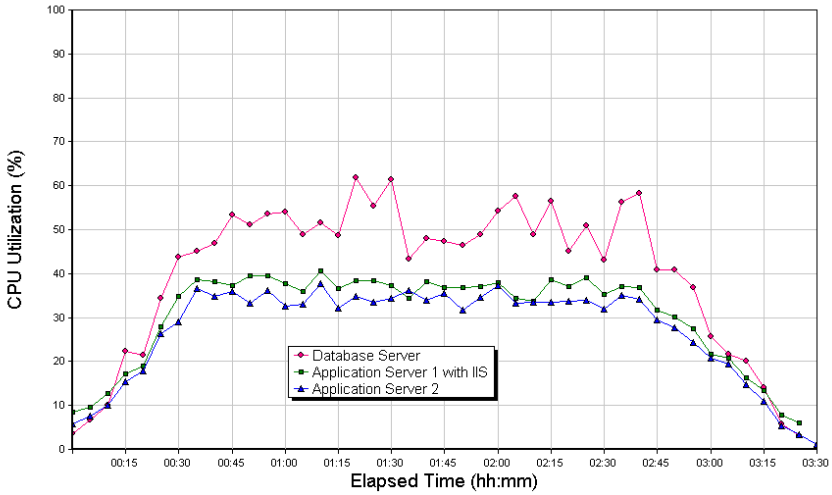

Figure 3. CPU utilization

Table 2 summarizes the CPU utilization during peak load (from 1:00 to 2:00 on the time line).

Table 2. Peak load CPU utilization

Server	Average CPU Utilization During Peak Load
Database server	51%
Host 1 (Application Server 1 with IIS)	37%
Host 2 (Application Server 2)	34%

Figure 4 shows the network utilization for HTTP requests per second and transactions per second. During the hour of peak load, 39,123 user transactions were processed, each one consisting of one or more HTTP requests. The total number of HTTP requests during this hour was 131,111. So, on average, a user transaction included approximately 3.4 HTTP requests.

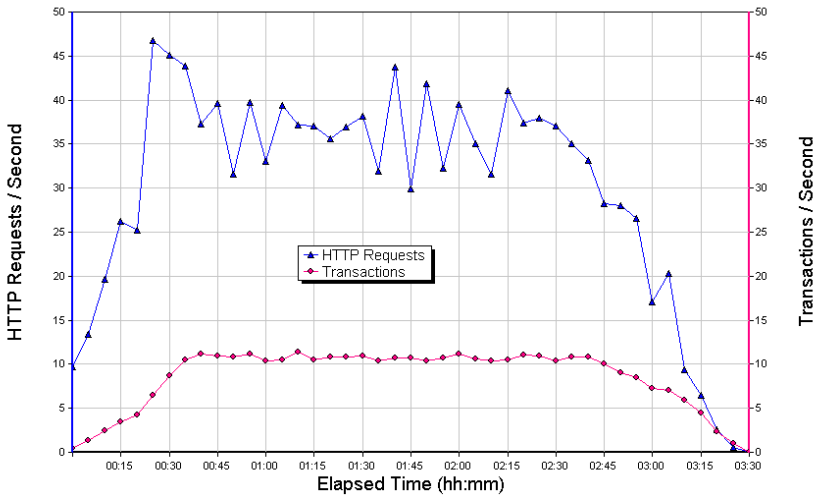

Figure 4. Network utilization

There were 36.4 HTTP requests per second from 1:00 to 2:00 on the time line. The network traffic was 175 kB per second.

With data extracted from the simulation's peak load hour, *Table 3* shows the most important transactions, their average response times, and the number of transactions during that hour.

Table 3. Number of transactions and average response times during peak load

Activity	Average Response Time (seconds)	Number of Transactions
Submit a simple request	0.46	596
Submit a complex request	0.77	363
Search for a request	3.55	513
Open an existing request	0.58	2,858
Update a request and process through workflow	0.96	2,484
View Mercury IT Governance Dashboard	1.20	2,389
Log on with portlets	1.95	813
Log on with summary portlets	7.90	120
Add a 1 KB attachment	0.04	596

Conclusion

Mercury IT Governance Center has been designed with attention to scalability, availability, and performance. The benchmark data presented here illustrate that, even with relatively modest hardware configurations, Mercury Demand Management can scale to support large numbers of concurrent users and transaction volumes.

Because your environment may differ in ways that affect performance, contact Mercury Support for further guidance.

Appendix A: Database Configuration Parameters

The benchmark database parameters were configured as shown below.

Database Parameter	Value
_like_with_bind_as_equality	TRUE
compatible	9.2.0.5
cursor_sharing	EXACT
db_block_checking	FALSE
db_block_checksum	TRUE
db_block_size	8192
db_cache_advice	OFF
db_cache_size	738197504
db_file_multiblock_read_count	16
db_writer_processes	1
dml_locks	748
enqueue_resources	968
job_queue_processes	2
max_commit_propagation_delay	700
open_cursors	1500
optimizer_dynamic_sampling	1
optimizer_features_enable	9.2.0
optimizer_index_caching	50
optimizer_index_cost_adj	50
optimizer_mode	CHOOSE
pga_aggregate_target	1677721600
processes	150
query_rewrite_enabled	TRUE
session_cached_cursors	225
sessions	170
shared_pool_size	629145600
undo_management	AUTO
undo_retention	10800
workarea_size_policy	AUTO

Mercury
379 North Whisman Road
Mountain View, CA 94043
Tel: (650) 603-5200
Toll Free: (800) TEST-911
Customer Support: (877) TEST-HLP
Fax: (650) 603-5300

If you have any comments or suggestions regarding this document, please send email to documentation@mercury.com.