

HP-Oracle Reference Configuration Assessment Service Brief

Data warehouse infrastructure assessment



Overview

If you're an IT manager, you're familiar with the challenges of building or upgrading a data warehouse or data mart platform that will accommodate your organization's changing needs. Sometimes the product selection and sizing process can seem as complicated as the BI systems themselves. Designing an adequate configuration can be time consuming and risky—a delicate balance of tradeoffs between price, performance, scalability, manageability, and reliability.

That's why HP and Oracle have teamed up to create a portfolio of reference configurations designed specifically for data warehousing. HP Reference Configurations for Oracle 10g Data Warehousing define optimized, off-the-shelf configurations that combine software, servers, and storage. This modular approach lets you support a variety of BI workloads and solutions based on your workload profile, data size, architecture choice, and price/performance goals. You can find the configuration that most closely fits your needs, giving you a proven advantage in speeding up design and implementation.

This service brief describes the assessment service that HP provides to reduce the time, effort, and risks associated with deploying a data warehouse to ensure that you can invest wisely and optimize the business results for your organization by accelerating time-to-intelligence.

The HP-Oracle Reference Configuration Assessment Service

The HP-Oracle Reference Configuration Assessment Service is a structured approach that enables a rapid review of your organization's current business intelligence infrastructure landscape and assesses it in comparison to both your short and long term business and technology goals. Upon completion, HP will recommend the HP Reference Configuration (servers and storage) that best meets your current needs, allowing you to keep pace with your projected growth and performance requirements.

The HP-Oracle Reference Configuration Assessment Service is intended to take less than one week to complete. We've taken the expertise gained from implementing successful data warehousing initiatives for thousands of satisfied customers and created a rapid, precise process based on best practices from both HP and Oracle to provide you the insight needed to make smart, strategic decisions.

Challenges

The workload of BI solutions is typically quite different from OLTP (transactional) systems. The necessity of architecting servers and storage for these specific needs reflects the growing amount of data being captured and the increased usage of BI applications. BI data warehouses and data marts typically have long-running queries with large result sets. Unless the server, storage, and database infrastructure is configured properly (i.e., balanced), poor performance will result.

Performance imbalance can cause resource contention and poor utilization in several areas: the numbers of CPUs, the amount of memory, the number of disks, the number of fiber channels, etc. These areas must match the workload to optimize performance. Commonly, the workload of BI databases is "I/O bound." HP Reference Configurations for Oracle 10g Data Warehousing are designed and tested to eliminate this bottleneck and optimize the data flow to match your business priorities.

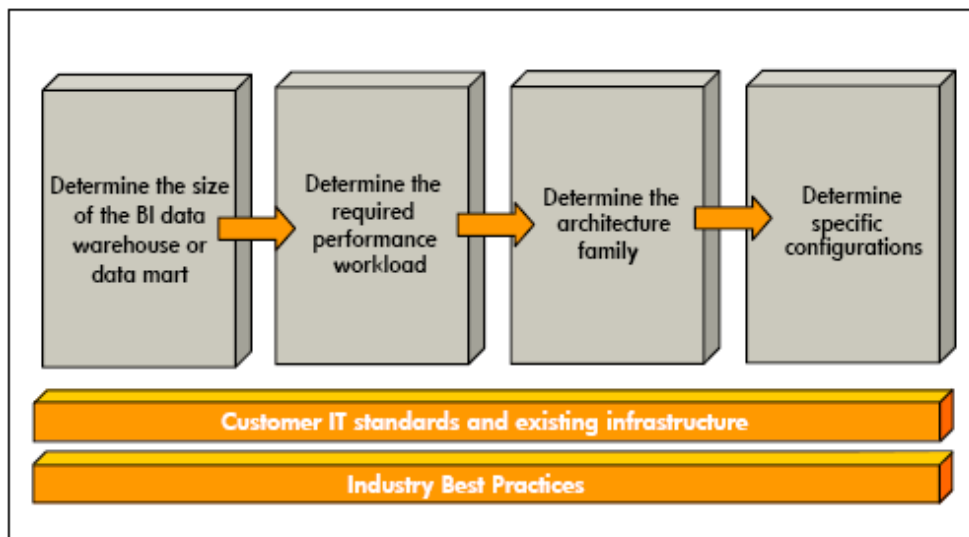
The Service

The HP-Oracle Reference Configuration Assessment can help you leverage HP and Oracle's experience and best practices to meet your demanding BI database requirements. This service provides a recommendation on a complete BI infrastructure configuration regardless of whether it is an existing data warehouse environment or it is a new infrastructure to be built and implemented. You will receive a detailed, recommended configuration that is based on our expertise with other customer BI databases, TPC benchmarks, results obtained at the HP Advance Technology Center (ATC) labs, and best practices for Oracle 10g.

Methodology

HP has adopted a proven, four-step methodology to deliver this assessment service. In keeping with our rapid approach to providing answers, this assessment will provide the key information necessary for a starter configuration. More detailed assessments and improvements may be needed to more closely customize the infrastructure for your BI environment and specific workload profile.

Figure 1. HP four-step methodology for delivering assessments



Step 1: Determine the size of the existing data warehouses/data marts

The size of a data warehouse is best measured in terms of raw data size—the size of the actual data, in the business sense, before it is inserted into the Oracle 10g DBMS, adjusted for data warehouse requirements like indexes, materialized views, and other such structures used to calculate the estimated user data size, and then used together with calculations for growth and redundancy (RAID) allocations to determine the total storage that needs to be provisioned.

If you already have a data warehouse in place, the assessment starts with the size of the existing database, then extrapolates based on your projected growth rate. If, however, this is a new infrastructure, we will work with you to provide an estimate for the raw data size. This can be derived from the size of the record and the number of records to be stored.

Step 2: Determine the required performance workload

A key step in understanding requirements is to determine the workload pattern accurately. This is often the most “mysterious” part of any BI infrastructure. Rules of thumb about the number of concurrent users can be used as a starting point, but are not the only method to determine the workload. Query types and the specific BI tools used have a significant impact on the workload pattern. A best practice is to benchmark using the I/O throughput, CPU utilization, and memory utilization. If you already have a data warehouse, we can obtain the storage system I/O throughput rates, server CPU utilization, and memory utilization through direct measurement. For new installations the performance workload will be estimated based on similar reference configuration installations.

Step 3: Determine the architecture family

HP Reference Configurations for Oracle 10g Data Warehousing includes the choice of:

- Scale-up with HP Integrity, HP-UX, and Oracle 10g
- Scale-out with HP Integrity, HP-UX, and Oracle 10g RAC
- Scale-out with HP ProLiant, Linux, and Oracle 10g RAC

Each architectural configuration is paired with an attached storage subsystem based on price and performance. Choosing the right architecture is a critical decision that will have lasting effects upon the performance of your data warehouse. Based on real-world results, lab testing, and industry benchmarks, we can assist you in making the right choice of architecture to meet your current needs and provide scaling capability to meet future requirements.

Step 4: Determine specific configurations

A specific HP Reference Configuration is selected in the final phase. You receive a detailed recommendation based on the choice of architecture, the raw data size, and the performance workload characteristics discovered and documented during the assessment process. Often, you will receive a choice of options enabling you to compare the cost and performance benefits of each of the relevant configurations and decide which one is ideal for your BI database.

Timeline and Deliverables

A typical HP-Oracle Reference Configuration Assessment takes five days and requires a commitment from your IT and database administration staff to provide key data points and input on your operations to be successful.

We keep you apprised of our progress and recommendations as the assessment takes place. At the end of the assessment you will receive a complete and detailed summary of our findings including:

- Overview summary of your current BI infrastructure
- Summary reports Steps One to Four

- A complete and detailed recommendation for the choice and deployment of an HP Reference Configuration for Oracle 10g Data Warehousing

Why HP?

The assessment service enables you to choose the right data warehouse infrastructure and invest in a configuration that will meet your needs today and scale for future growth. HP Reference Configurations for Oracle 10g Data Warehousing are based on solid architecture principles, industry best practices, customer installation observations, and lab testing. An HP-Oracle Reference Configuration Assessment allows you to make an investment in your BI infrastructure with confidence.

In addition to the assessment service, HP Information Management practice focuses on solving massive and complex data problems for Global 2000 companies. HP is a trusted advisor, strategic business partner, and expert implementer in business intelligence, data warehousing, data integration, and information quality.

With proven methodologies, deep expertise, and industry specialization, HP solves its clients' most pressing data challenges. Our practical solutions deliver actionable and measurable business results that assist critical decision making, optimize IT efficiency, and improve business performance.

For more information

For more information and availability of this offering in your country please contact your local HP representative or contact us via our web page: <http://www.hp.com/hps/intelligence>

© 2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Linux is a U.S. registered trademark of Linus Torvalds. Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. UNIX is a registered trademark of The Open Group.

4AA1-0563ENW, March, 2007

