

Informix I-Spy

Version 1.0

For UNIX

Informix® I-Spy™ is a smart data warehouse monitoring and optimization tool designed for Informix databases. I-Spy helps administrators and architects increase the business value of enterprise data warehouses and data marts through enhanced utilization, efficiency, design improvements, and a lower cost of maintenance. With new and more effective ways to quantify, manage, and control data warehouse usage, Informix I-Spy sits transparently between the database and the client to help the administrator monitor and adjust database resources and client query usage. Regardless of the size—from the smallest data mart to the largest data warehouse, or the function—from simple ad-hoc queries to the most complex online analytical processing, Informix I-Spy ensures the data warehouse is operating at its maximum efficiency and effectiveness—to deliver the highest return on investment.

About Informix I-Spy

In today's competitive business environment, many organizations are investing large amounts of money and resources to develop and deploy data warehouses. However, the effectiveness of a data warehouse investment is difficult to measure since there are few tools available that enable administrators to see the type of query executed, the length of query execution time, and the amount of data returned. As a result, organizations have no way to determine whether their data warehouse resources are being utilized efficiently.

Extending Informix's leadership in the data warehouse arena, Informix has introduced I-Spy. I-Spy offers a wide range of functionality to assist administrators in tracking the overall usage of the data warehouse. I-Spy provides the ability to measure and control utilization of the warehouse as well as to identify and limit inefficient use of resources, enabling customers to obtain the information needed to refine and tune their data warehouse data and models over time. Additionally, I-Spy provides basic metrics required for calculating operational return

on investment, allowing customers to determine the effectiveness of their data warehouse.

Informix I-Spy Advantages

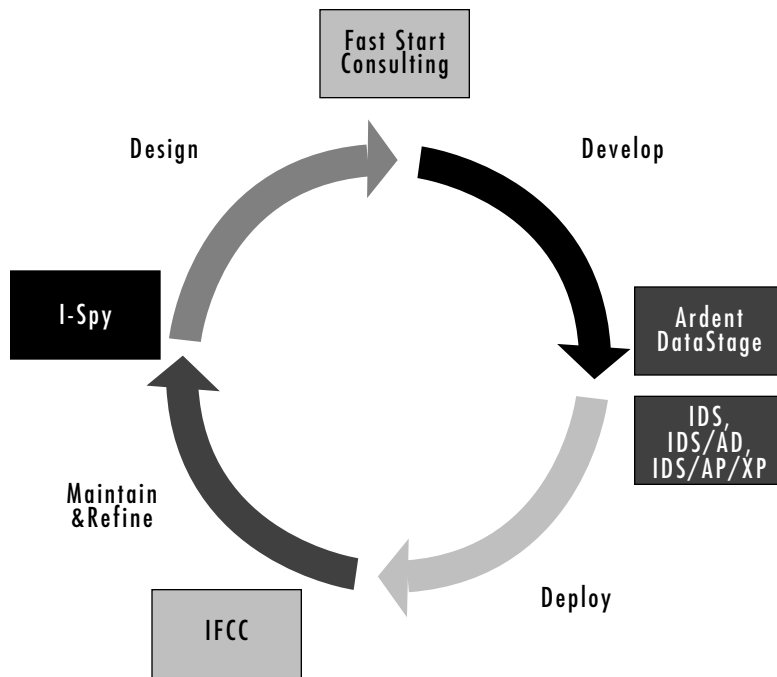
Key advantages of Informix I-Spy include the ability to:

- maximize the data warehouse investment by providing useful information on how the warehouse is being used;
- improve the efficiency of data warehouse design by allowing the data warehouse architects and developers to prototype the data and queries required, which can reduce design cycles and data model tuning; and
- enable customers to take a more proactive approach to capacity planning so that they can better prepare for data warehouse expansion over time.

Maximize Data Warehouse Investment

Developing and deploying a data warehouse is a costly venture. Therefore, it is important to ensure that the data warehouse is operating at maximum efficiency. Unlike online transaction processing

(OLTP) systems, however, the cost/benefit justification of a warehouse is much more difficult to calculate, and the success of a warehouse is often measured by its perceived value to the end user. To ensure that resources are being used efficiently and effectively, customers need to be able to track how the data warehouse is being used. Informix I-Spy lets data warehouse administrators, developers, and architects track all the activities within the warehouse including the type of SQL executed, the type and quantity of data being accessed, and SQL execution time. I-Spy provides continuous logging to keep track of client level activities, allowing administrators to easily identify inefficient or ineffective queries and pinpoint users who might need help or additional training. I-Spy also provides "rule-based" commands that allow administrators to transparently place constraints on client activity. Queries can be canceled based on configuration parameters such as the maximum number of rows returned, total query time, or the actual text of the SQL statement. These constraints can help prevent end-users from consuming excessive data warehouse resources.



Improve Efficiency of Data Warehouse Design

The development of a successful data warehouse is an iterative process. A complete solution not only requires a robust database server and tools for development and deployment, but also tools that monitor and identify how the resources of the warehouse are managed. A resource tracking/monitoring tool enables the warehouse model to be refined over time for maximum performance.

Informix I-Spy provides architects with concise information to help them tune the data warehouse model. By identifying which tables, columns, and indexes are to be used, I-Spy can speed up the development process. Once the data warehouse is operational, I-Spy can provide statistics on what entities in the warehouse the users are actually using versus what they expected to use. This allows developers to eliminate data in the warehouse that is not being used. Additionally, it allows them to make intelligent decisions on how to refine indexing as well as tune database schema, as illustrated on the following page.

The efficiency of the data warehouse design can be further improved through analyzing query plans. I-Spy provides the ability to load query plans into the activity database, therefore

allowing administrators to pinpoint problem queries and educate users issuing these queries on developing more efficient queries. Developers can also use query plan analysis to make further improvements in indexing and fragmentation strategies.

Plan for Growth

Once the data warehouse is deployed, its utilization will grow over time as users become familiar with the data available and the benefits it can provide. This growth can result in performance and capacity problems if demand for information grows beyond the capabilities of the warehouse infrastructure. It is, therefore, critical to monitor the usage and performance of the warehouse over time to plan for growth.

With Informix I-Spy, customers can take a more proactive approach to warehouse capacity planning. I-Spy provides a set of utilization reports that can be used to track usage of the data warehouse by time or by users. When comparing performance trends over time, these reports can help customers plan for growing demand and justify additional resource requirements.

How Informix I-Spy Works

I-Spy is essentially a daemon process that sits between Informix clients and database servers. When a new connection arrives, I-Spy creates a new process to handle the connection while the original I-Spy process goes back to listening for connections. Each client connection, therefore, has a unique I-Spy process handling the server message traffic.

I-Spy is completely transparent and passes all client and server message traffic unaltered. The only intervention taken by I-Spy is in the case of a match with a configuration rule, which may result in an SQL statement being canceled. A single configuration file is used to control I-Spy and is read and processed when the I-Spy daemon is started.

Informix I-Spy Features and Functions

Informix I-Spy provides a complete set of features and functions to track how the data warehouse is being utilized.

Activity Logging

Informix I-Spy maintains statistics on all activities at the client level by trapping SQL statements submitted to the database server. Stored as a flat file or in a separate activity database, I-Spy logs information such as session start and end times, SQL statements executed and their execution time, number of rows returned, and server cost information which can be used for later analysis.

Session Constraints

I-Spy provides “rule-based” commands that allow administrators to transparently place constraints on client-level activities. These rules, which can be user-based, host-based, or global in scope, are maintained within the I-Spy configuration file. They can help prevent excessive consumption of database resources by canceling or prohibiting inefficient queries. Queries can be canceled or blocked based on configuration parameters such as the estimated cost, number of rows returned, query time, or the actual text of the SQL

statement. I-Spy's rules can also be used to set SQL session controls, such as isolation level or "pdqpriority". This is significant because many of the OLAP and ROLAP tools on the market today don't have the capability to change these session controls.

Administration Interface

I-Spy includes a Web-based interface for administering the I-Spy engine. It lets administrators view the status of the I-Spy engine, start and stop the I-Spy daemon, edit the configuration file, view currently connected clients and their activities, run and view reports, and view the I-Spy log file.

Conforming to the standard http/CGI architecture, the I-Spy administration interface provides an open environment that does not require the customer to purchase or use a particular browser, Web server, or middleware product.

Standard Reports

I-Spy comes with a number of standard reports that allow administrators to monitor data warehouse usage and client activities. For analysis of general queries, I-Spy provides reports identifying the longest running queries, queries with the most rows returned, and queries with most total data returned. For user-based statistics, it provides reports tracking users

with the most query executions, users who consumed the most engine time, users with the most SQL errors, etc. For server-based statistics, it provides reports summarizing total queries per day, total sessions per day, and total errors per day. All reports can be executed using the Web-based administration interface and include a form interface for entering report parameters, such as start and end date, user ID, and number of rows to return.

By default, the report output is generated in standard HTML tables that users can view using their browser. Users also have the option to receive the output in the form of an unload format, which allows them to create unload files that can easily be imported into other tools like Microsoft Excel.

Query Plans

One of the most powerful tools used in optimizing queries is the output from the SET EXPLAIN statement. This output contains detailed information on the execution plan that the optimizer has selected to run a given query. It can be used to identify inefficient queries, problems within indexing strategies, and other issues, such as the need to run update statistics.

Recognizing that the data produced from SET EXPLAIN can assist in improving the efficiency of the data warehouse, I-Spy includes an option to load the SET EXPLAIN output into its activity database through a post-execution process. This allows users to view output directly from the I-Spy administration interface and monitor how the optimizer plans for change over time. Coupled with client statistics within the activity database, the SET EXPLAIN output allows data warehouse developers and architects to facilitate performance tuning efforts.

System Requirements

Supported Hardware Platforms

- Sun SPARC running Solaris 2.5.1 or later
- HP running HP-UX 10.20 or later
- IBM running AIX 4.2.1 or later

Supported Database Servers

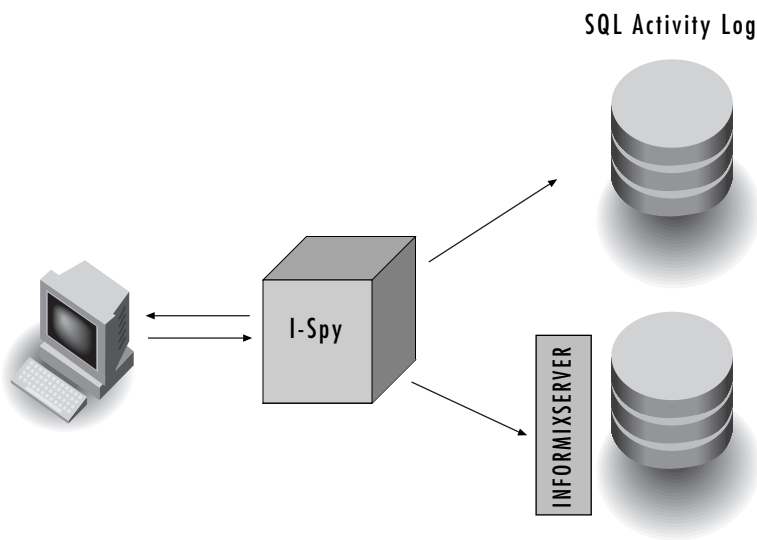
- Informix Dynamic Server™, version 7.2x, 7.3x
- Informix Dynamic Server with Advanced Decision Support Option™, version 8.x
- Informix Dynamic Server with Advanced Decision Support Option and Extended Parallel Option™, version 8.x

Web Server (optional)

The Informix I-Spy product includes a Web interface for administering the I-Spy server. This interface was developed with a standard HTTP/CGI architecture. Therefore, a Web server that supports the Common Gateway Interface (CGI) protocol must be installed on the same server as I-Spy. This is required to run the user interface, but not for the I-Spy product itself (which can be run from the command line).

Memory and Disk Requirements

- 6 MB RAM minimum
- 10 MB of disk space



About Informix

Informix Corporation, the world's database technology leader based in Menlo Park, California, provides innovative database products that assist the world's major corporations attain competitive advantage. Informix is widely recognized as the technology leader for corporate computing environments ranging from small workgroups to very large parallel processing applications. Informix's database server, application development tools, superior customer service, and strong partnerships enable the company to be at the forefront of major information technology solution areas including data warehousing, high performance OLTP, and Web/content management.

For more information, contact the sales office nearest you or visit our Web site at www.informix.com.

INFORMIX REGIONAL SALES OFFICES	
• Asia/Pacific	65 298 1716
• Canada (Toronto)	1 416 730 9009
• Europe/Middle East/Africa	44 181 818 1000
• Federal	1 703 847 2900
• Japan	81 3 5562 4500
• Latin America	1 305 265 7545
• North America	1 800 331 1763

	1 650 926 6300

Informix®

4100 Bohannon Drive
Menlo Park, California 94025
1 650 926 6300

World Wide Web: www.informix.com

© 1999 Informix Corporation. All rights reserved. The following are trademarks of Informix Corporation or its affiliates, one or more of which may be registered in the U.S. or other jurisdictions: Informix®, Informix Dynamic Server™, Advanced Decision Support Option™, Extended Parallel Option™, and Informix I-Spy™.