IBM® Information Management Software

Performance Management

Using IBM InfoSphere Optim Performance Manager and Query Workload Tuner

- Prevent problems before they impact the business
- Align monitoring and business objectives
- Improve query workload performance

Chuck Ballard
Ute Baumbach
Holly Hayes
Marcia Miskimen
Lakshmi Palaniappan
Marichu Scanlon
Yong Hua Zeng

Redbooks

ibm.com/redbooks

International Technical Support Organization

Performance Management Using IBM InfoSphere Optim Performance Manager and Query Workload Tuner

October 2013

SG24-8111-00

Note: Before using this information and the product it supports, read the information in “Notices” on page ix.

First Edition (October 2013)

This edition applies to Version 5.2 of InfoSphere Optim Performance Manager, Version 3.2 of InfoSphere Optim Query Workload Tuner.

Contents

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®
Cognos®
DataStage®
DB2®
developerWorks®
IBM®
InfoSphere®
Netcool®
OMEGAMON®
Optim™
pureQuery®
pureScale®
PureSystems™
Rational®
Redbooks®
Redbooks (logo) ®
System x®
Tivoli®
WebSphere®
z/OS®

The following terms are trademarks of other companies:

Adobe, the Adobe logo, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Preface

IBM® InfoSphere® Optim™ Performance Manager offerings are designed for managing the overall performance of IBM DB2® for Linux, UNIX, and Windows data servers, typically by a database administrator (DBA) or an application owner. It provides 24x7 monitoring of the application database and alerts users to potential problems so that problems can be resolved before affecting the business. IBM InfoSphere Optim Performance Manager provides drill-down capability from alerts to diagnostic detail for determining whether the source of the problem is CPU, I/O, memory, workload, SQL, or another area. Integration with InfoSphere Optim Query Workload Tuner enables seamless analysis of query workloads and recommends resolutions to improve performance. InfoSphere Optim Performance Manager also enables users to collect, prune, and retain historical information facilitating problem determination, trend analysis, and capacity planning. Historical information lets you go back in time to see how a problem arose, mitigate problems first, and do causal analysis later, or compare current values to prior time periods. It also helps staff be proactive about query tuning to optimize resource utilization and to plan for growth.

IBM InfoSphere Optim Performance Manager helps businesses improve performance and reduce costs by giving DBAs and other IT staff the information needed to manage performance proactively to help in these areas:

- Prevent problems before they affect the business.
- Save hours of staff time and stress.
- Align monitoring objectives with business objectives.

Its Extended Insight capability is unique in database performance monitoring, giving database professionals the visibility across the application stack to where database workloads are spending time. It provides problem isolation to the appropriate layer of the stack, extensive introspection into database bottlenecks, and the ability to monitor response time objectives for the highest priority workloads.

InfoSphere Optim reduces risk and costs, speeds solution delivery, boosts performance, and addresses compliance requirements for databases, warehouses and big data environments.

Data lifecycle management is the process of managing business information throughout its lifecycle, from requirements through retirement. Data lifecycle management spans different application systems, databases and storage media and is part of an overall Information Integration and Governance Strategy. By managing data properly over its lifetime, organizations are better equipped support business goals with less risk.

IBM InfoSphere Optim Query Workload Tuner facilitates query and query workload analysis and provides expert recommendations for improving query and query workload performance. Use InfoSphere Optim Performance Manager to identify slow running queries, top CPU consumers, or query workloads needing performance improvements and seamlessly transfer them to InfoSphere Optim Query Workload Tuner for analysis and recommendations. Query Workload Tuner facilitates query analysis using query formatting annotated with relevant statistics, access plan graphical or hierarchical views, and access plan analysis. It further provides recommendations for improving query structure, statistics collection, and indexes including generated command syntax and rationale for the recommendations.

This IBM Redbooks® publication presents an overview of the IBM InfoSphere Performance Management solution for IBM DB2. It includes detailed product overviews for the product components that comprise the solution and also has use cases to understand how to apply the products to common scenarios to quickly get value from solution deployments.

This book is intended for DB2 database administrators or other IT staff directly involved in diagnosing and resolving performance issues that arise with database applications.

Authors

This book was produced by a team of specialists from around the world working with the International Technical Support Organization in San Jose, California.

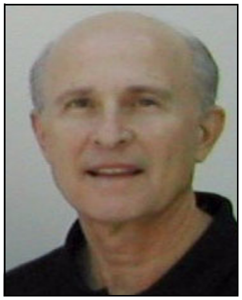

Chuck Ballard is a Project Manager at the International Technical Support organization in San Jose, California. He has over 35 years of experience, holding positions in the areas of product engineering, sales, marketing, technical support, and management. His expertise is in the areas of database, data management, data warehousing, business intelligence, and process re-engineering. He writes extensively about these subjects, teaches classes, and presents at conferences and seminars worldwide. Chuck has both a Bachelor's degree and a Master's degree in Industrial Engineering from Purdue University.

Ute Baumbach has been a Software Developer at the IBM laboratory in Germany for more than 20 years, where she works in various software development projects and roles. Most of her projects were based on DB2. For ten years, she has worked as a member of the development team for DB2 performance tools, and now the Optim Performance Manager development team. In addition, Ute supports customers in setting up Optim Performance Manager and utilizing it to its best advantage in DB2 data server environments. She is an IBM Certified Databases Administrator and a Certified Application Developer for DB2 for Linux, UNIX and Windows.

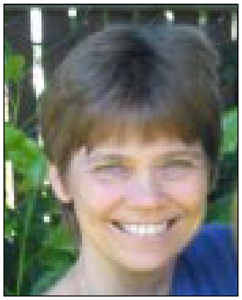

Holly Hayes is a Product Manager with IBM for InfoSphere Optim Solutions. With 30-plus years at IBM, she has held development, strategy, product management, product marketing, and management positions primarily working in and around database management systems, data warehouse infrastructure, information integration, database administration, performance management, and database development technology. She has been a prominent speaker at industry events and customer briefings and a frequent contributor to industry articles, analyst research, and other publications. She holds a U.S. patent in replication technology.

Marcia Miskimen is a member of the Optim Integration QA team, where she focuses on usage-based testing of the Optim Data Studio portfolio of products, such as Optim Performance Manager, Optim Query Workload Tuner, Optim Data Studio, Optim Configuration Manager, and Optim Query Capture Replay. She has worked with Optim Performance Manager since version 1, and has co-authored six IBM Redbooks publications. Marcia has been with IBM for 23 years and has had various development and test roles for both IBM Software Group and Global Services.

Lakshmi Palaniappan is a Senior Software Engineer at the IBM Silicon Valley Lab in San Jose, California. She is part of the development team for IBM InfoSphere Optim Performance Manager, and is the development focal point for the monitoring tooling in IBM PureSystems™. Lakshmi leads the Worldwide Level 3 IBM InfoSphere Optim Performance Manager Support team, which includes enablement support to customers in setting up their environments. She was previously team leader for IBM InfoSphere Replication Server configuration and monitoring tooling. Lakshmi has over15 years of industry experience.

Marichu Scanlon is a member of the Optim Tools Enablement Team at the IBM Silicon Valley Laboratory in San Jose, CA. She has actively participated in presenting, demonstrating, and assisting customers with deploying several Optim Tools products, such as InfoSphere Optim Performance Manager, InfoSphere Optim Query Workload Tuner, and InfoSphere Optim Configuration Manager. Before joining the enablement team, she was the team leader for the Data Studio Continuing Engineering team, and a member of Data Studio SQL and Routine tooling development.

Yong Hua Zeng (Henry) is a Solution Architect of Database Management portfolio, and is located in Beijing, China. He led the development of several IBM Optim database management products and infrastructure. Henry has significant experience in DB2 system administration, monitoring and tuning, and is a certified DB2 Advanced Database Administrator. Henry is now leading the Optim database tools solution enablement, training and advocacy in the Asian Pacific area.

Other contributors

We want to thank others who contributed to this book, in the form of written content, subject expertise, and support.

From IBM Locations Worldwide

- Kevin Beck: Senior Software Engineer, InfoSphere Optim Performance Manager Development, IBM Software Group, Information Management, in Hillsboro, OR, USA
- Sandra Bialek: Cross-platform Product Manager, IBM Software Group, Information Management, Silicon Valley Lab, San Jose CA, USA
- Paul Bird: Senior Technical Staff Member for DB2 LUW Development, IBM Software Group, Information Management, Markham, ON Canada
- Dilip Biswal: OPM/PE Software Developer, IBM Software Group, Information Management, Silicon Valley Lab, San Jose CA, USA
- Marianne Carter: Manager - Optim Admin Tools Enablement, IBM Software Group, Silicon Valley Lab, San Jose CA, USA
- Anson Kokkat: Product Manager - Optim Database Tools, IBM Software Group, Toronto Lab, Markham, ON Canada

- Cliff Leung: Software Architect, Distinguished Engineer, Optim Data Studio Chief Architect, IBM Software Group, Information Management, Silicon Valley Lab, San Jose CA, USA
- Sven Miller: IT Developer, Data Studio Performance Manager, IBM Software Group, Information Management, Boeblingen, Germany
- Steve Rees: STSM, IBM DB2 pureScale® and Optim Data Studio Performance, IBM Software Group, Information Management, Markham, ON Canada
- Sriram Srinivasan: Architect/STSM, Optim Data Studio Web Infrastructure, IBM Software Group, Information Management, Silicon Valley Lab, San Jose CA, USA
- Neela M Suram: Software Engineer for InfoSphere Optim Performance Manager, IBM Software Group, Information Management, Bangalore, India
- Zin Wu: Software Engineer, IBM Software Group, Information Management, Silicon Valley Lab, San Jose, CA, USA
- Shu Ji Wan: Software Engineer, IBM Software Group, Application and Integration Middleware Software, Beijing, China
- Yun Ting Wang: Advisory Software Engineer, Application and Integration Middleware Software, Beijing, China
- Daniel Zilio: Senior Software Developer, InfoSphere Optim Query Tuner, IBM Software Group, Information Management, Markham, ON Canada

From the International Technical Support Organization

- Mary Comianos: Publications Management, San Jose, CA
- Linda Robinson: Graphics Support, Raleigh, NC
- Ann Lund: Residency Administration, Poughkeepsie, NY
- Diane Sherman: Editor, Poughkeepsie, NY

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Part 1

Strategy for performance management

This part describes the InfoSphere Optim performance management products and provides guidance for using and deploying those products to help improve the performance in your particular environment.

This part includes the following chapters:

1. Chapter 1: The intended audience for this topic includes DB2 database administrators or other IT staff directly involved in diagnosing and resolving performance issues that arise with database applications. It includes detailed product overviews for two of the component products of the solution and also use cases to understand how to apply the products to common scenarios to quickly get value from solution deployments.
2. Chapter 2: This chapter contains important information for planning your deployment of InfoSphere Optim Performance Manager. It indicates the requirements of each of the InfoSphere Optim Performance Manager components and answers questions such as these:
   - What products do I want or need to install?
   - How can I best prepare for installation?
   - What kind of system resources are required?
   - How can assure compliance with my company's security requirements?
3. Chapter 3: This chapter describes how to set up and run Optim performance management tools, and in particular InfoSphere Optim Performance Manager and InfoSphere Optim Query Workload Tuner. The chapter includes information about installation, configuration, and several practices that can help the deployment be efficient and smooth.

1

Guided approach to performance management

Poor application and database performance can have a significant effect on your company's bottom line. For e-commerce sites, inadequate performance is a leading cause of lost sales and customer churn. For mission-critical systems, poor performance can result in substantially reduced productivity and revenue. Missed service-level agreements (SLAs) can lead to loss of customers, a damaged brand, and even fines and other penalties.

Addressing application and database performance is a tough challenge: Organizations are under constant pressure to develop and deploy new business applications while ensuring that existing applications maintain optimal performance levels to meet customer expectations. But as usage increases, applications often become less responsive. And with data volumes doubling in size every five years, database performance can become unacceptably sluggish.

Attempting to solve these problems by purchasing more hardware and adding staff causes infrastructure and resource costs to skyrocket and is not a long-term solution. Database administrators (DBAs) can spend significant amounts of time responding to a performance problem-time that can be spent on more strategic work to help the business reduce operational costs and become more efficient.[1]

[1] http://public.dhe.ibm.com/common/ssi/ecm/en/imb14129usen/IMB14129USEN.PDF

This book presents an overview of the IBM InfoSphere Performance Management solution for IBM DB2, herein referred to simply as the performance management solution. The solution can help you implement a proven methodology to identify, diagnose, solve and prevent performance problems while reducing the time and money spent on performance-related tasks. By combining solutions for integrated performance management, DBAs can move out of reaction mode and resolve performance issues before they have a major effect on the business. In addition, it can support organization to move into a proactive approach to performance management improving organizational effectiveness.

1.1 Overall IBM InfoSphere approach

The performance management solution can enable a logical and intuitive approach to managing performance delivering a closed loop process for performance management.

Figure 1-1 depicts the process for rapidly resolving and ultimately preventing performance issues.

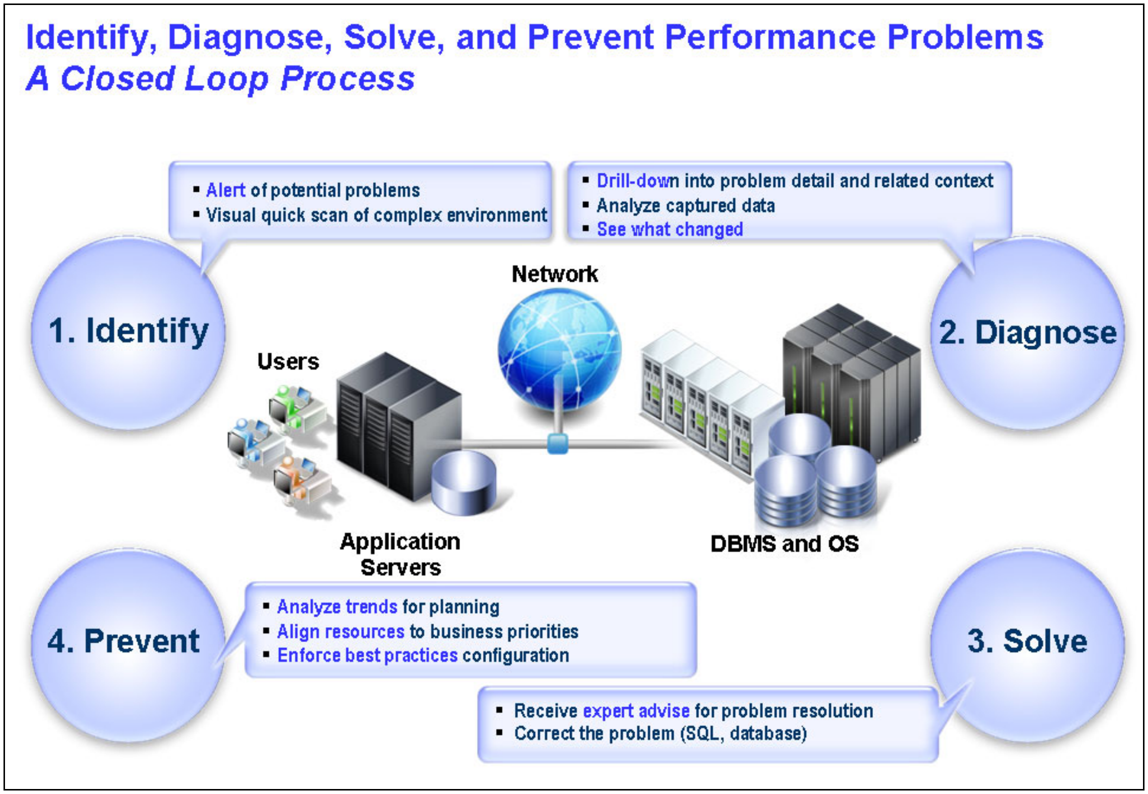

Figure 1-1 A closed loop performance management process

The first step is to identify that a potential problem or emerging bottleneck. Organizations must identify problems before problems impact the business. The performance management solution monitors key health and performance indicators to alert users to resource bottlenecks or degrading performance.

Next, drill down from the alert to more detailed information for problem diagnosis. Here, InfoSphere Optim provides rich diagnostic dashboards help determine root cause. You can view current detail, compare to prior time frames, view trends, and see configuration changes that might be related to performance degradation.

After the problem is isolated, you must solve it. The performance management offerings provide actionable, expert advice to improve the performance of individual queries and query workloads. Analysis of query plans, indexes, statistics, and options that can take days can be reduced to hours providing rapid response.

But ultimately, you want to be ahead and prevent problems from occurring in the first place. This is where IBM can help. First, make good use of the resources you have. The performance management solutions help you take advantage of DB2 Workload Manager so that resources are allocated according to business priorities and ad hoc queries are not allowed to monopolize system resources. Next, the performance management solutions help you to enforce compliance with configuration best practices on both servers and clients so you follow your own best advice. The performance repository aggregates performance history giving you the trending information needed for capacity planning to stay ahead of business growth. Plus, the performance management offerings interface with development tools to help developers and DBAs collaborate to deliver high-performing from the outset.

1.1.1 Products in the InfoSphere Optim performance management solution

The performance management solution includes several complementary and integrated offerings:

- InfoSphere Optim Performance Manager: This comprehensive, proactive solution provides real-time monitoring, alerting, and performance warehousing for DB2 and IBM InfoSphere Warehouse for Linux, UNIX, and Windows.
- InfoSphere Optim Configuration Manager: This innovative solution offers centralized, effective, and efficient management of database and client configurations.
- InfoSphere Optim Query Workload Tuner: This intelligent solution delivers expert recommendations to maximize application performance, helping reduce specialized skill requirements and the total cost of ownership.

- IBM InfoSphere Optim pureQuery® Runtime: This high-performance data access platform helps enhance the performance of existing in-house database client applications and speeds the development and deployment of new applications.
- Data Studio: Included with DB2, Data Studio provides database administration and development features.

As shown in Figure 1-2, InfoSphere Optim Performance Manager is the heart of performance monitoring, providing broad and deep key performance indicator (KPI) monitoring and alert notification. Alerts lead users to resource or workload-specific dashboards that provide the detail needed to isolate the problem. With Integration with InfoSphere Optim Configuration Manager, you can quickly see what changes have happened recently that might be related to a potential performance issue.

All of these products are included with DB2 Advanced Enterprise Server Edition V10.

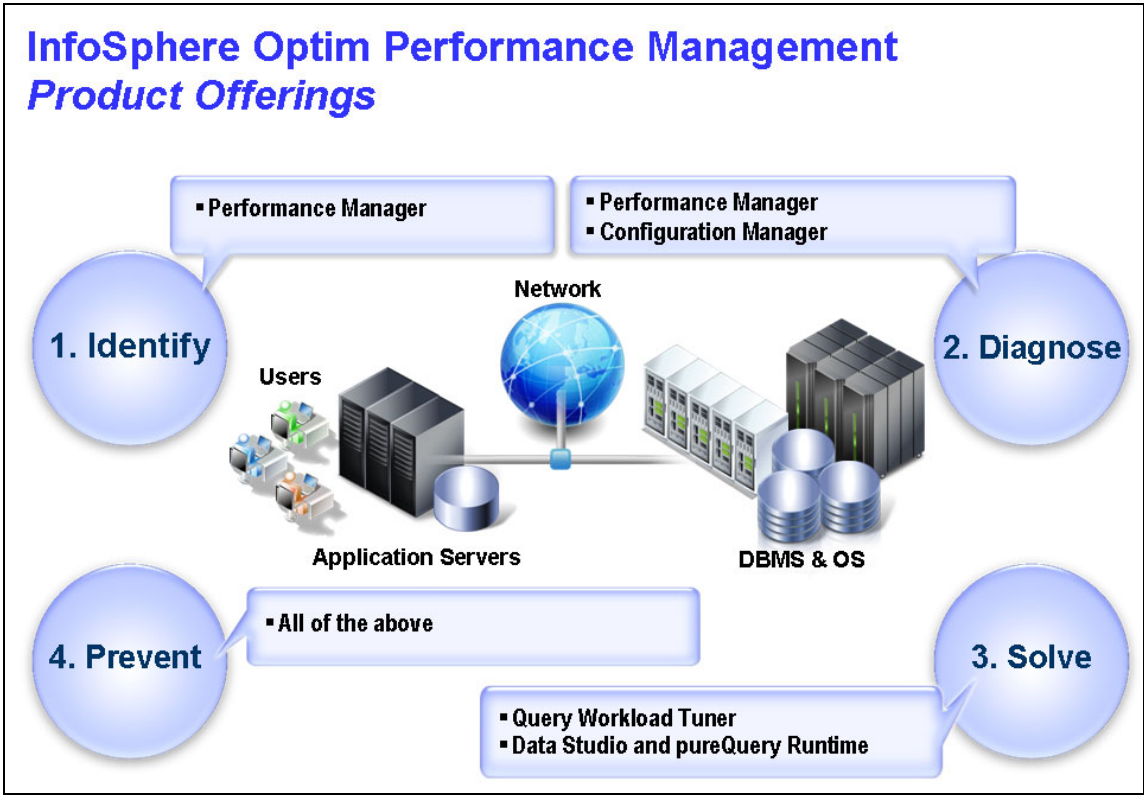

Figure 1-2 Optim PM product offerings

InfoSphere Optim Query Workload Tuner provides expert and actionable recommendations to improve the performance of queries or query workloads. Data Studio, included with DB2, provides basic administration for object management, change management, and maintenance. It also provides the development environment for queries and routines with integrated access to InfoSphere Optim Performance Manager and InfoSphere Optim Query Workload Tuner. Data Studio also provides the development environment for InfoSphere Optim pureQuery Runtime for developing high performance applications or improving performance of existing applications without having to change the code.

All of the tools contribute to the solution's ability to prevent problems before they impact the business.

1.2 What is new in the products

The performance management solution represents a relatively young set of products with rapid growth and frequent feature delivery. Product deliveries tend to be timed with DB2 releases and reflect the key priority for the portfolio: supporting user success with DB2.

This book is based on examples and features associated with specific product releases as specified in this section. *IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows*, SG24-7925 has information about InfoSphere Optim Performance Manager Version 4 architecture. You might be familiar with that offering or with the performance management solution in general. The following sections list recent product additions.

1.2.1 InfoSphere Optim Performance Manager for DB2 for Linux, UNIX, and Windows V5.1

This was a major release in which the server was re-architected to take advantage of the DB2 9.7 state-of-the-art in-memory metrics monitoring capabilities. IBM significantly reduced monitoring overhead and increased scalability for monitoring partitioned systems, reduced disk footprint for performance data, and improved retrieval times.

That release and subsequent modification (mod) releases and fix packs introduced new functions, several of which are described in the following list:

- All dashboards that use in-memory metrics data collection can now display data collected at any point in time (for example, now, yesterday, last week, or last year. You can view trends instantly by selecting the time frame.
- Performance history needed to detect trends, plan capacity and for service reporting now has automated rolling aggregation at four levels with user-defined retention periods so you can keep data longer at lower cost.
- With the At-a-Glance View dashboard, you can view current activity and determine whether it is typical or atypical using baseline support.
- You can list, filter, sort, analyze, and report on database connections and resource consumption across server, I/O, row, transaction, locking, communication, and utilities. Basically about anything you need to see about connections.
- Many improvements are available in SQL dashboards so you can view top executions of SQL statements or aggregated metrics and integrate with InfoSphere Optim Query Workload Tuner.
- Many new reports are available. One example is the SQL comparison report to help you more easily find SQL performance improvements or regressions before and after changes. In addition, custom reporting against the performance repository is enabled.
- Dashboards now have enhanced analysis for database partitions and DB2 pureScale members. View overall performance for a database or at a member of partition level. You can get introspect performance characteristics of the cluster caching facility (CF) or get alerts on potential CF bottlenecks.
- With a new best practices template for DB2 Workload Manager configuration, you can have DB2 Workload Manager for warehouse environments up and running quickly. Additional enhancements support both new and experienced Workload Manager users.
- Alerting extensions to support definition of user-defined was added.

1.2.2 InfoSphere Optim Performance Manager for DB2 for Linux, UNIX, and Windows V5.2

This version of the product delivers many new features:

- Stored procedures may be analyzed and ranked according to aggregate cost. With drill-down capabilities, you can analyze call paths and correlate SQL statements with the procedures that issue them without the overhead of SQL statement tracing.
- Workload-specific monitoring templates are updated to reflect aggregated customer experience and to provide better guidelines on use.
- You can easily get regular reports through the new report scheduler that includes email notification and report storage and retention management in the repository.
- Blackouts for monitoring activities, job execution, and alerts can now be more flexibly scheduled. This includes multi-day periods, repeating windows, and custom scheduled windows.
- You can automate responses to alerts through user-defined jobs that may be triggered when either a predefined or user-defined alert is encountered.
- Database connection information can be kept in sync with the enterprise by using the new synchronization between InfoSphere Optim Configuration Manager centralized database repository and InfoSphere Optim Performance Manager database list.
- The At-a-Glance View in the performance overview dashboard now provides the operating system performance details for HP-UX and Solaris.

1.2.3 InfoSphere Optim Query Workload Tuner for DB2 Linux, UNIX, and Windows

This product supports query analysis and workload tuning recommendations that can significantly reduce time spent analyzing and tuning workloads. It fully subsumes and significantly extends capabilities available in DB2 Design Advisor. It supports query and access plan advice and also workload-based index, statistics, materialized query table, and multidimensional clustering advice.

The most recent release upon which this book is based, V3.2, delivers the following enhancements:

- Improvements in "what-if" analysis allow for determining the impact on a workload of a change to one or more indexes. You can determine the possible benefit or negative impact on a workload by the virtual creation or dropping of an index without actually implementing the change.
- A workload access plan comparator is available to compare access paths using the "what-if" scenario or to compare access paths between two workloads, even on different systems such as test and production. The workload access plan comparator can also be used to assist in locking down access plans for a very large workload.
- Improved index advisor provides recommendations for the elimination of existing non-unique indexes because of the presence of existing near-duplicate, unique indexes. It also now displays the last-used time for existing indexes to further help with the analysis of the use of an index.
- Workloads may now be captured from EXPLAIN tables or ACTIVITY event monitors, along with their access plan.
- You can use the workload summary report to improve collaboration across your organization. This report summarizes the various recommendations provided by InfoSphere Optim Query Workload Tuner.
- InfoSphere Optim Query Workload Tuner now offers native 64-bit support.
- Identify and support singleton select statements used by applications to improve performance by using less expensive methods to retrieve database data.

1.2.4 IBM InfoSphere Optim Configuration Manager for DB2 for Linux, UNIX, and Windows

With V2.2, you can do the following tasks:

- Show enterprise data assets using information gleaned from database application clients and scheduled database server jobs. InfoSphere Optim Configuration Manager V2.2 adds support for CLI and .NET clients, in addition to Java clients.
- Track changes to server and client configurations to help you quickly determine the root cause of performance degradation.
- Automatically monitor database client and server configurations and database objects, alerting you when configurations are out-of-sync.

- Show storage savings opportunities across multiple databases. Savings are achieved by running DB2 V10, compression opportunities, both static and dynamic, trapped storage, and objects that are not frequently accessed.
- Define migration schedules to move data from one storage group on a faster disk to another on a slower disk, or from slower to faster, using DB2 V10 multi-temperature storage capabilities.
- Redirect application clients from one DB2 for Linux, UNIX, and Windows system to another mirrored system for migrations, disaster recovery, or high availability. InfoSphere Optim Configuration Manager V2.2 adds support for CLI and.NET clients in addition to Java clients.
- Enforce application client and driver properties to achieve desired workload balancing in DB2 pureScale systems. InfoSphere Optim Configuration Manager V2.2 adds support for CLI and .NET clients in addition to Java clients.
- Centralize client connection definitions and server data models to improve manageability. InfoSphere Optim Configuration Manager V2.2 adds support for synchronization of connection information with other Optim tools with InfoSphere Optim Configuration Manager serving as a Tools Registry.

1.2.5 IBM InfoSphere Optim pureQuery Runtime for Linux, UNIX, and Windows V3.2

This product is a data access platform that offers improved performance, security, and manageability of database client applications. It provides a runtime environment and an application programming interface that enhances the performance of existing in-house applications without having to modify them. It also helps in rapid development of new applications.

InfoSphere Optim pureQuery Runtime for Linux, UNIX and Windows helps in the following aspects:

- Improves performance of existing applications, including C, Java, and .NET, without the need to modify code.
- Accelerate development of database applications in conjunction with IBM Data Studio.
- Improve manageability with trackback to source, problem isolation, and version control.
- Enhance security by reducing SQL injection risks.

1.2.6 IBM Data Studio V3.2

This version of IBM Data Studio can help improve the productivity of database administrators and developers, with capabilities such as the following examples:

- Streamline database development with advanced query validation, object management, and system deployment features:
  - Use customizable templates and editors to create, test, debug, and deploy routines, such as stored procedures and user-defined functions.
  - Simplify query building, editing, tuning, and formatting (including support for XQuery); view access plan graphs, and get statistics advice.
  - Debug efficiently with additional support for triggers, PL/SQL anonymous blocks and declared routines, persistent breakpoints in both the routine editor and debugger and routine recompilation, without having to redeploy.
  - Manage objects and privileges, analyze dependencies, view data distributions, and update object statistics.
- Maximize collaboration through an integrated environment and shared platforms, as in the following examples:
  - Collaborate roles using an Eclipse based integrated development environment (IDE) to develop new applications, administer the database and monitor database health.
  - Simplify database schema changes, develop and test routines, and generate data-centric web services.
  - Benefit from integration with IBM InfoSphere Data Architect; IBM InfoSphere Optim pureQuery Runtime for Linux, UNIX and Windows; and IBM Rational® application development solutions.
- Save time and reduce errors using advanced data management, and configuration and administration tools and features.
  - Create, schedule, and track command scripts and utilities, and manage jobs by configuring email notifications in command scripts to report upon job completion.
  - Efficiently manage database instances for standard, database partitioning features and IBM DB2 pureScale topologies; and reverse-engineer databases into physical models, and synchronize changes between models and databases.
  - Manage database objects change using model-based forward engineering, develop Java applications that use IBM pureQuery annotated methods, and copy objects from one database to another.

- Run commands on multiple objects, and manage cluster members using the IBM pureScale environment.
- Benefit from enhanced IBM DB2 capabilities such as adaptive compression, time travel query, row and column access control, and multi-temperature data management because IBM Data Studio release is closely aligned with DB2.

1.3 InfoSphere Optim Performance Manager introduction

InfoSphere Optim Performance Manager helps DBAs, development, and IT staff monitor DB2 for Linux, UNIX, and Windows workloads.

1.3.1 Features

InfoSphere Optim Performance Manager offers the following features:

- Monitor any DB2 workloads from small to large, and online transaction processing (OLTP) to business intelligence (BI). Monitor single instance symmetric multi-processors (SMPs), multi-partition warehouses, or pureScale deployments. Get a "bird's-eye" view across all your database systems from a single dashboard. Drill down into deep partition, member, or component analysis within the database. Reduce monitoring overhead on the monitored databases with built-in use of state-of-the-art DB2 9.7 in-memory metrics for performance monitoring.
- Get proactive notification of emergent problems before they affect your business. Quickly configure monitoring using built-in templates for common workloads, and receive early warning of emerging resource bottlenecks. Templates include default values for which metrics to collect, collection frequency, and key performance indicator thresholds and alerts. Visualize problems through the dashboard, get email notification, integrate SNMP alerts with system management software, or provide custom actions or integrations with ticketing or other systems.
- Prioritize those issues that have most impact to the business. Monitoring database server metrics such as buffer pool hit ratio and locking contention are important to fine tuning database performance, but they might not give the DBA early warning of degrading user experience. They also do not help the DBA distinguish performance characteristics of high priority applications, transactions, or users. With the *Extended Insight* feature, DBAs can configure and track adherence to response time objectives for key workloads to align their work with business priorities. Default workloads for IBM Cognos®, SAP,

IBM WebSphere®, IBM InfoSphere DataStage®, and InfoSphere Warehouse are predefined. However, any workloads can be defined by user, by transaction, by application server, by application, or other connection attributes.

- Isolate problems to the correct component immediately. Many times problem isolation is a complex, multi-day, multi-person effort where each administrator (such as for application server, network, and database), brings logs and traces, and the team tries to correlate the metrics to identify the source of the problem. Such time represents lost productivity by the staff, and can represent millions of dollars in lost revenue for customer-facing applications. The Extended Insight feature provides unique database monitoring capability that tracks transaction execution across the layers of the application stack to make it instantly apparent where transactions are spending time (that is within the application, the application server, the database driver, the network, or the database server). DBAs or other IT staff can visualize how much time is being spent in the application, application server, client driver, network, and database server, and then route the problem to the appropriate team for resolution.
- Use seamless analysis of performance data from real time to any time. All dashboards using in-memory metrics data collection can now display data collected at any point in time (for example, now, yesterday, last week, or last year). Compare current values to prior values to determine whether they are out of the norm. Mitigate symptoms first and defer root cause analysis with historical data. Instantly view metric trends on dashboards to see spikes, emergent problems, or significant changes to steady state. Automatic rolling aggregation with user-defined retention provides data necessary for comparisons, trend detection, capacity planning, and service reporting. This facilitates long term performance data retention and analysis with minimal cost.
- Compare current activity to baselines. The At-a-Glance View dashboard provides an activity overview for the system. Here a baseline can be set and deviations from baseline highlighted to quickly see if system activity, database workload, I/O, or other activity is significantly different from the baseline.
- Drill into contextual, resource-specific data for performance analysis. List, filter, sort, analyze, and report on database resources including buffer pool and I/O, system resources, memory, locking, and logging. Drill through contained and container relationships. View connection, SQL, and utility activity across the database or drill into the detail for a specific one. Identify the most costly SQL either from single executions or in aggregate. View stored procedure calling patterns and associated SQL.
- Create interactive reports for proactive tuning, capacity planning, and service reporting. Each interactive report is like multiple reports in one and can be over any time frame. You can browse, print, export, present, and share your

analysis and results. Ready-for-use reporting includes reports on performance by database, disk consumption, connections, configuration, tables, packages, and SQL. The SQL comparison reports make it easy to see SQL performance improvements or regressions across configuration changes, workload tuning enhancements, or application or database upgrade. Custom reporting from the performance repository is also available.

- Improve cross-team collaboration and responsiveness. Transfer a single query or entire query workload in a single click to InfoSphere Optim Query Workload Tuner to get expert advice for improving queries, access plans, statistics, and indexes. Developers using Data Studio can use integration with InfoSphere Optim Performance Manager to quickly identify hot spots for tuning and immediately view and compare performance improvement or degradation. Contextual launching of InfoSphere Optim Configuration Manager allows DBAs to determine whether configuration changes might be the source of performance issues. Integration with IBM Tivoli® offerings provides complete infrastructure monitoring for DB2 applications. with browser-based infrastructure, access to performance information is available anywhere at anytime, and enables deployment to a larger user base, for example developers and testers, to be practical because it does not introduce additional monitoring overhead on the DB2 server.
- Align resource allocation according to business priorities. DB2 Workload Manager allows DBAs to control resource allocation against defined workloads to assure that the most important work gets the highest priority, increase systems utilization, and prevent rogue queries from taking over the system. DBAs can use InfoSphere Optim Performance Manager to get started with DB2 Workload Manager for warehouse systems in a single click. Configuration enhancements result from recent work on best practices for deploying DB2 Workload Manager for warehouse environments. Results have shown that the majority of performance improvements result from straightforward deployments of DB2 Workload Manager. Deploy the default template, tune service class boundaries based on histograms of activity mappings, and enable governance to prevent rogue queries from taking over systems. For advanced users, the Workload Manager configuration interface accepts any existing Workload Manager configuration without modification, and exposes all features and options for modification.

1.3.2 Packaging

InfoSphere Optim Performance Manager is available in the following five offerings:

- InfoSphere Optim Performance Manager Enterprise Edition provides data server monitoring with integrated problem identification and notification, problem diagnosis, performance reporting, and DB2 Workload Manager configuration.
- InfoSphere Optim Performance Manager Workgroup Edition provides the same capabilities as Enterprise Edition, but is priced for departments or small businesses and is restricted to monitoring DB2 Workgroup Sever Edition databases only.

  > **Note:** DB2 Workload Manager functions are subject to their separate licensing on the DB2 server.

- InfoSphere Optim Performance Manager Content Manager Edition provides the same capabilities as Enterprise Edition, but is priced for monitoring the DB2 library associated with IBM Content Manager or IBM Content Manager On Demand.

  > **Note:** DB2 Workload Manager functions are subject to their separate licensing on the DB2 server.

- InfoSphere Optim Performance Manager Extended Edition provides all the capabilities of Enterprise Edition plus *Extended Insight* monitoring, which includes the following items:
  - Database transaction response time monitoring: The Extended Insight feature captures the transaction when it begins at the database client and tracks its response time characteristics across the transaction. This provides problem isolation across the application stack so DBAs can determine if it is their issue to resolve or whether to route to other areas. DBAs can see where transactions are spending their time: in the application, application server, client driver, network, or database. When monitoring DB2 9.7, or later databases, you can even see the time-spent metric in the database server components, for example, sort time and lock wait time. In addition, you can configure workloads to track, for example, high priority users, transactions, or applications, and monitor response time objectives for those workloads.

Thus, you can proactively, quickly, and intuitively identify the following information:

- Who has response time performance issues, or causes them to others, by identifying the specific set of transactions that make up the problem workload. For example, which user ID or application issues those transactions.
- When the response time performance occurred, by identifying the problem periods. For example, monitor minutes, hours, days, weeks, or even years.
- What specific activities were involved in the response time performance problem, by identifying the complete list of involved problem SQL statements. For example, check the top N SQL statements by end-to-end response time or data-server time.
- Why the response time problem occurred, by identifying the exact problem layer that slows the response time. For example, detect whether the transactions are slowed down in the database, network, driver, application, or application server.

> **Note:** The Extended Insight feature collects and displays additional monitoring information for applications running in WebSphere Application Server. Application servers, such as WebLogic or Tomcat, are treated as generic Java applications.

– Automated DB2 Workload Manager configuration adjustments based on response time objectives: The Extended Insight feature adds new configuration options for DB2 Workload Manager. It enables users to set response time objectives for workloads and let InfoSphere Optim Performance Manager adjust resource allocation to achieve those objectives.

– Integration with Tivoli offerings: The Extended Insight feature also provides the correlation needed to integrate InfoSphere Optim Performance Manager with Tivoli offerings, specifically IBM Tivoli Composite Application Manager for Transactions, IBM Tivoli Composite Application Manager for Application Diagnostics, and IBM Tivoli Manager for Servers. Integration enables InfoSphere Optim Performance Manager contextual launching within the Tivoli Enterprise Portal.

- InfoSphere Optim Performance Manager Extended Insight provides the Extended Insight capabilities as an add-on to InfoSphere Optim Performance Manager Enterprise Edition when purchased at a later time. InfoSphere Optim Performance Manager Extended Insight should be purchased only for use with InfoSphere Optim Performance Manager Enterprise Edition.

Data Studio provides a subset of InfoSphere Optim Performance Manager features for health monitoring. Table 1-1 shows a comparison of features.

Table 1-1 Features for health monitoring

Feature detail	Data Studio	IOPM EE, WE, CME[a]	IOPM XE[b]
Multi-system monitoring	Yes	Yes	Yes
Health summary dashboard	Yes	Yes	Yes
Health monitoring customization	Yes	Yes	Yes
Health monitoring	Yes	Yes	Yes
Health alerts drill down	Yes	Yes	Yes
User-defined alerts	Yes	Yes	Yes
Email, SMTP, and SNMP notification	Yes	Yes	Yes
User-defined actions from alerts	Yes	Yes	Yes
View current or historical data	Yes	Yes	Yes
Performance alerts	No	Yes	Yes
Database server KPI templates and customization	No	Yes	Yes
Database server KPI monitoring	No	Yes	Yes
Database server KPI alerts	No	Yes	Yes
Instant dashboard trending	No	Yes	Yes
Database server At-a-glance and KPI dashboard	No	Yes	Yes
Baseline comparison for At-a-glance deviation detection	No	Yes	Yes
SQL analysis dashboard	No	Yes	Yes
Buffer Pool and I/O analysis dashboard	No	Yes	Yes
Memory analysis dashboard	No	Yes	Yes
Locking analysis dashboard	No	Yes	Yes
Logging analysis dashboard	No	Yes	Yes
Workload analysis dashboard	No	Yes	Yes
System analysis dashboard	No	Yes	Yes
Utilities analysis dashboard	No	Yes	Yes

Feature detail	Data Studio	IOPM EE, WE, CME[a]	IOPM XE[b]
Performance history aggregation and retention	No	Yes	Yes
Ready-for-use interactive reporting	No	Yes	Yes
Report scheduling, email, and retention	No	Yes	Yes
Data Studio integration to visual explain	No	Yes	Yes
Optim Query Workload Tuner integration to query analysis and tuning	No	Yes	Yes
DB2 Workload Manager configuration and reporting	No	Yes	Yes
What changed visibility through integration with Optim Configuration Manager	No	Yes	Yes
Application accelerators (SAP, ISW, Cognos, Info Server)	No	Yes	Yes
Response time monitoring for complete database transaction (spans application stack)	No	No	Yes
Response time component analysis	No	No	Yes
DB2 Workload Manager automation based on response time objectives	No	No	Yes
WebSphere connection pool detail	No	No	Yes
See and compare performance metrics in Data Studio	No	No	Yes
Tivoli offering integration (contextual drill-down from IBM Tivoli Composite Application Manager for transactions, contextual drill-through to detailed operating system metrics with IBM Tivoli Monitoring for Servers)	No	No	Yes

a. InfoSphere Optim Performance Manager Enterprise Edition (IOPM EE); InfoSphere Optim Performance Manager Workgroup Edition (IOPM WE); InfoSphere Optim Performance Manager Content Manager Edition (IOPM CME)

b. InfoSphere Optim Performance Manager Extended Edition (IOPM XE)

1.3.3 InfoSphere Optim Performance Manager architecture

Figure 1-3 depicts the architecture for InfoSphere Optim Performance Manager V5.

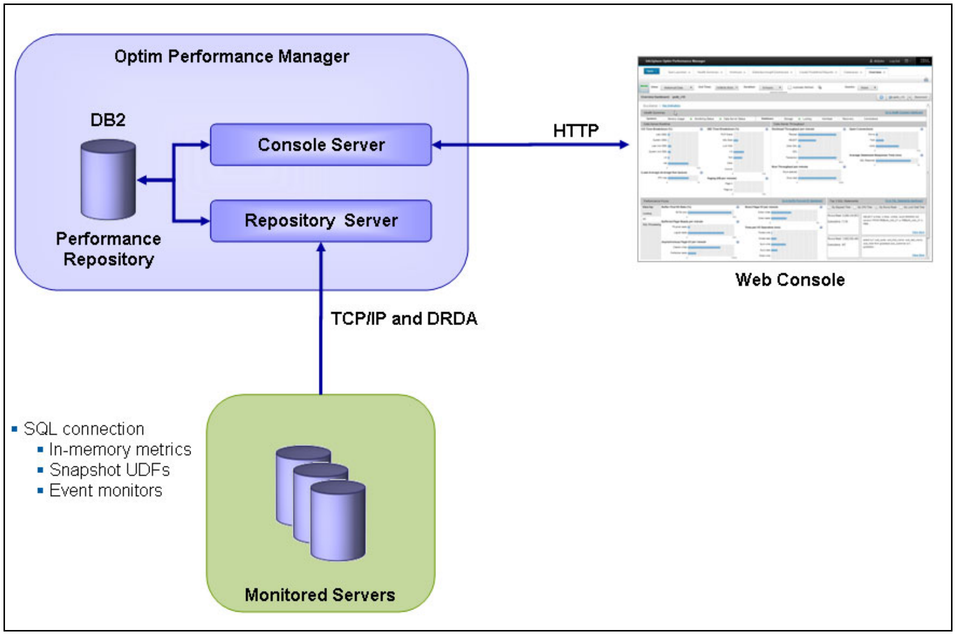

Figure 1-3 InfoSphere Optim Performance Manager architecture

InfoSphere Optim Performance Manager consists of a console server and a repository server component. The repository server collects the data from the monitored database and stores it in a DB2 performance repository database. It also manages aggregation and pruning of the performance data in the repository. A restricted use DB2 V10 license is included with the product to use for the performance repository.

The console server typically reads the data from the repository database and displays it in the web UI. However, the repository access is bypassed in the case of real-time monitoring and data is requested directly from the monitored database through the repository server.

The servers must be co-located with the DB2 performance repository.

Note: The performance repository might not be Distributed Partitioning Feature (DPF) or pureScale instances and it cannot be running in Oracle compatibility mode.

For monitoring DB2 V9.7 or later databases, as it relates to this book, the repository server connects to the monitored database and issues SQL statements to collect the performance metrics using a light-weight in-memory metrics infrastructure. Depending on the configuration, some event monitors may be created on the monitored database, for example, to collect locking information. A limited number of metrics are still collected through snapshot user-defined functions (UDFs). The user specifies how often to collect the data, which data to collect, and how long to retain the data in the monitoring profiles.

The main user interface is the web console. It provides the primary user interface for the product for administering the product, configuring monitoring profiles, viewing both real-time and historical performance data online, and generating reports.

Occasionally other tools may connect directly to the performance repository. InfoSphere Optim Query Workload Tuner and Data Studio connect to the repository server to retrieve performance metrics in support of high cost query identification and metrics retrieval or to populate the performance repository with application metadata captured by InfoSphere Optim pureQuery Runtime. Reporting tools can also access the performance repository through specially designs functions to support custom reporting.

InfoSphere Optim Performance Manager Extended Insight architecture

The Extended Insight architecture is depicted in Figure 1-4 on page 23.

When deploying the Extended Insight feature, users can optionally deploy the Extended Insight client (also known as the Data Tools Runtime Client) on the monitored client environments. These environments are installed on the systems where your database applications are running and support Java, CLI, .Net, and embedded applications clients. The Extended Insight clients intercept database traffic for the monitored database and collect response time data about transactions and SQL statements. This data is then periodically forwarded back to the repository server. The client, however, has to know where to send the information. For DB2 9.7 FP6 (or later) clients, the port information is taken directly from the database itself. For earlier clients, the port must be specified during the client configuration process and is saved in the `pdq.properties` file.

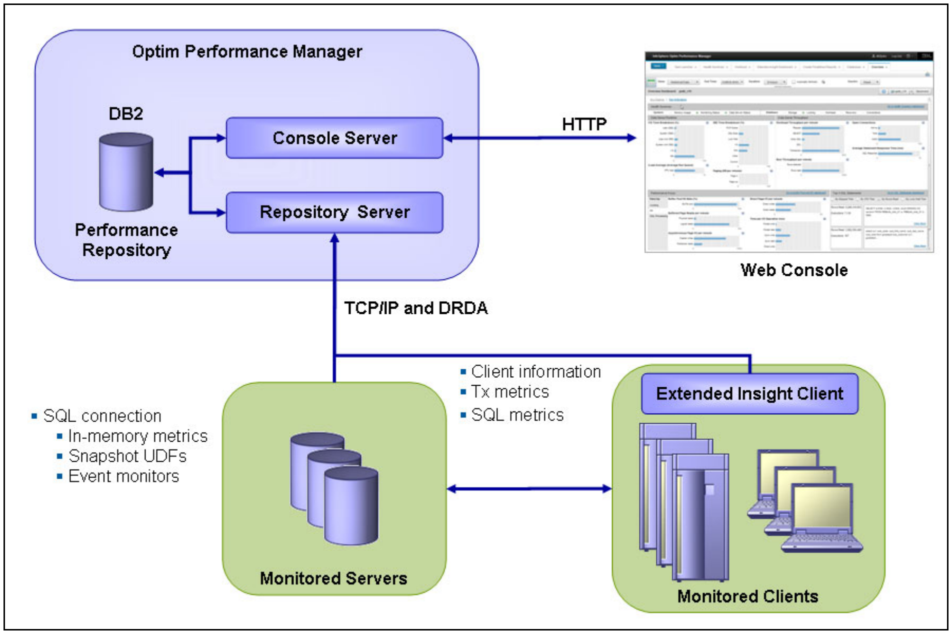

Figure 1-4 InfoSphere Optim Performance Manager Extended Insight architecture

In parallel, the repository server collects data server execution data for SQL and transactions from the monitored database and correlates them to the data sent by the clients. This way you get a complete picture about the transaction and statement response times of your application. To collect the Extended Insight server data, the repository server may start package cache and unit of work event monitors on the monitored database to get the SQL and transaction details.

Note: Extended Insight client and server data collection is independent; you can activate one without the other and you can further collect SQL data without also collecting the transactional detail.

Monitored database

No components are installed on the monitored DB2 database. InfoSphere Optim Performance Manager uses native DB2 monitoring features accessed through SQL. Overhead on the monitored server depends on the monitoring configuration specified and can be negligible or can have impact, depending on the monitoring configuration specified and the utilization on the monitored server. Following a proven methodology for deployment allows users to minimize impact to the monitored database while balancing data collection. The repository server can start event monitors on the monitored server for locking events, package cache collection, and other metric collections. If using monitoring configurations that start event monitors, an advisable approach is to activate Automatic Task Scheduler to watch for repository server failures to mitigate risk of event monitor tables growing.

InfoSphere Optim Performance Manager integration with Tivoli Monitoring

The InfoSphere Optim Performance Manager plug-in for Tivoli Enterprise Portal facilitates integration between InfoSphere Optim Performance Manager Extended Insight, IBM Tivoli Composite Application Manager for Application Diagnostics, and IBM Tivoli Composite Application Manager for Transactions in a Tivoli Enterprise Portal Console for end-to-end transaction monitoring. It allows InfoSphere Optim Performance Manager to send database transaction information directly to IBM Tivoli Composite Application Manager. This data is then surfaced in the Tivoli Enterprise Portal Console, which allows operators of the Tivoli Enterprise Portal Console to be notified when database transactions are not performing well. It also allows operators to contextually launch InfoSphere Optim Performance Manager dashboards for further diagnosis and analysis.

The InfoSphere Optim Performance Manager plug-in for Tivoli Enterprise Portal also delivers extended operating system performance data by launch-in-context capabilities. For instance, when you launch an InfoSphere Optim Performance Manager dashboard from the Tivoli Enterprise Portal Console, you can open the operating system monitoring details for a selected application client by launch-in-context into the Tivoli Enterprise Portal workspace.

1.4 InfoSphere Optim Query Workload Tuner

IBM InfoSphere Optim Query Workload Tuner helps database administrators and SQL developers optimize the performance of SQL statements in applications that query DB2 for Linux, UNIX, and Windows databases. InfoSphere Optim Query Workload Tuner can provide recommendations for single SQL statements or for entire SQL workloads.

InfoSphere Optim Query Workload Tuner includes four major sets of functions:

- Query capture and management features: These features help users capture queries from various sources and store them for analysis purposes.
- Query analysis features: These features help DBAs and developers understand queries, the choices that the DB2 optimizer is making in terms of the access plan, and why it is making those choices.

- Single query advisory features: These features help DBAs and developers improve the performance of a query by making specific recommendations, typically about the access plan or the way in which the query is structured.
- Query workload advisory features: These features provide advice about improving the performance of an overall workload. They focus on recommendations about the database design, statistics, and database features that can improve the aggregate workload performance.

The following topics provide an overview of the functional areas.

Query capture and management features

InfoSphere Optim Query Workload Tuner is designed to work in the context of DBA and developer work environments for DB2, thus there are a variety of means for capturing queries and query workloads. Most important, in the context of performance management, it can capture queries or query workloads from InfoSphere Optim Performance Manager. InfoSphere Optim Performance Manager identifies high cost queries or workloads that are not meeting their service level objectives and transfers them to InfoSphere Optim Query Workload Tuner for analysis and advice. However, many more data sources are available including the following items:

- The SQL editor or routine editor in Data Studio

  Features are available to directly initiate tuning features from within the editors.
- A package, SQL stored procedure, trigger that uses compiled SQL statements, and user-defined function that uses compiled SQL statements in the DB2 catalog

  Features are also available to initiate tuning from a view of these objects in the Data Source Explorer.
- A text-based file
- The DB2 package cache
- An InfoSphere Optim Performance Manager repository
- An XML file that uses the InfoSphere Query Workload Tuner schema
- A cut and paste or directly typing the SQL statement or statements that you want to work with

After you have captured a SQL statement or set of SQL statements, you might want to prune them or arrange into custom workloads and save them.

Query analysis features

Query analysis has the following features:

- Format and annotate SQL statements.

  In the formatted query, each table reference, each column reference under the SELECT clause, and each predicate is shown on its own line. You can expand and collapse sections of complex queries, such as query blocks and subqueries, to see an overview of the query and drill into parts of the query in more detail. When you click any line in the formatted query, other lines of the query that contain column or table references from the same table are also highlighted. You can also customize the formatting by ordering the predicates according to various criteria such as local predicates or join predicates, table references, and highest filter factor.

- Generate and compare visual representations of access plans.

  An access plan graph displays the access plan for any statement for which EXPLAIN data is available.

  In addition to the access path, InfoSphere Optim Query Workload Tuner also annotates the operations that provide insight into the statistics that were used to determine the access path. For a detailed explanation of access paths, see the IBM developerWorks® article "Tuning SQL with Optim Query Tuner, Part 1: Understanding access paths" at the following address:

  https://ltsbwass001.sby.ibm.com/cms/developerworks/data/library/techarticle/dm-1006optimquerytuner1/index.html

  You can also compare access plans to see whether there have been changes to access paths across database migrations or to see whether your tuning efforts have had an impact.

- Browse access plans with the Access Plan Explorer.

  The Access Plan Explorer provides a tabular view of the access plan making it easy to sort by cost and see relationships between components.

Single query advisory features

The most important single query advisory features are the query advisor and access path advisor. Although index advice and statistics advice are available in single query mode, those functions are more commonly used in a workload context, so impact to the entire workload can be considered in the recommendations.

The Query Advisor looks for opportunities to improve a query. It is not focused on semantically identical query rewrite because the DB2 optimizer does that. Rather, Query Advisor looks for opportunities to improve the query where there might be a subtle change to the query semantics. In addition to the actual advice, the Query Advisor generates explanations and examples for the recommendations thus also serving to provide education to DBAs and developers.

The Access Path Advisor examines the access plan that is chosen by the optimizer and identifies certain common access path issues. The warnings that this advisor provides can help you to understand where to look for trouble in an access plan graph or in the Access Plan Explorer.

Query workload advisory features

These features accept an entire SQL workload and consider how suggested changes impact the overall workload. For example, adding an index might increase the speed of queries but slow the insert, update, or delete operations. The workload advisors make recommendations in the following areas:

- Collecting statistics or statistical views that are relevant to a query workload

  The Workload Statistics Advisor recognizes missing, conflicting, or out-of-date statistics, generates RUNSTATS statements to generate statistics that are needed to improve query performance. The Workload Statistics Advisor can also recommend new statistical views and changes to existing views. Statistical views help the DB2 optimizer to estimate cardinality better when data has a non-uniform distribution or when SQL statements have complex predicates, relationships among complex predicates, and relationships across tables. It can also generate statistics profiles in support of DB2 automatic maintenance features.

- Creating and modifying indexes

  The Workload Index Advisor recommends new indexes and changes to existing indexes on tables that are referenced by SQL statements in query workloads. It generates the necessary CREATE INDEX DDL, rationale for making the recommendation, and also provides estimates of performance improvement and required disk for the new or changes index. It also allows you to perform what-if analysis by adding constraints to the recommendations and creating virtual indexes that allow you to analyze resultant access plans.

- Creating and modifying materialized query tables, multi-dimensional clustering, and partitioning data across database partitions

  InfoSphere Optim Query Workload Tuner subsumes all features of the DB2 Design Advisor. Run the Workload Design Advisor to learn how materialized query tables, multidimensional clustering, and data redistribution can improve the performance of query workloads.

1.4.1 Packaging

InfoSphere Optim Query Workload Tuner for Linux, UNIX, and Windows is now available exclusively in the workload tuning offering, which provides both single query and query workload functions. A subset of features is available in Data Studio. It includes query formatting, access plan graph visualization, and single query statistics advice.

Architecture

IBM InfoSphere Optim Query Workload Tuner, Version 3.1.1 consists of two components:

- IBM Data Studio full client Version 3.1.1
- License activation kit

After you install a Data Studio 3.1.1 client, you install the InfoSphere Optim Query Workload Tuner license activation kit to enable the full feature set. In the first connection to the target DB2 database, the InfoSphere Optim Query Workload Tuner license will be pushed to the database. As a result, any other IBM Data Studio full client may use the full feature set without installing the license activation kit.

Before using query tuner features, the target DB2 database must be configured properly. During the configuration, EXPLAIN tables are checked for correct format. InfoSphere Optim Query Workload Tuner creates a set of tables on the database used to store the query workload, explain information, and recommendation results. This set of tables will allow multiple users to share and access all the captured query workloads and recommendation results. In addition, a few stored procedures are deployed to the database as part of the runtime components.

When performing analysis, InfoSphere Optim Query Workload Tuner requires a connection to the target DB2 database on which the query runs on. If the access plan does not exist, an EXPLAIN function will be done. The access plan is kept in the EXPLAIN tables. The query tuning functions then retrieve the EXPLAIN records, analyzes the access plan, and provides recommendations.

1.5 Looking ahead

The remainder of this book describes the following topics:

- Chapter 2: Planning your InfoSphere Optim Performance Manager deployment covers prerequisites, installation roadmaps, capacity planning, storage options, and privilege management for InfoSphere Optim Performance Manager deployments.
- Chapter 3: Installing and configuring Optim performance management tools steps through installing and configuring InfoSphere Optim Performance Manager and InfoSphere Optim Query Workload Tuner.
- Chapter 4: Getting to know InfoSphere Optim Performance Manager provides a dashboard by dashboard and report by report overview of InfoSphere Optim Performance Manager features.
- Chapter 5: Getting to know InfoSphere Optim Query Workload Tuner details analyzing and tuning workflow and discusses each of the key InfoSphere Optim Query Workload Tuner advisors.
- Chapter 6: Finding and fixing database level bottlenecks discusses how to use InfoSphere Optim Performance Manager to recognize disk bottlenecks, which are probably the most common types of bottlenecks. It also shows how users can identify performance problems related to database design issues, and how to remedy them.
- Chapter 7: Performance management for distributed DB2 environments features use cases for how to use InfoSphere Optim Performance Manager to detect bottlenecks that are unique to Distributed Partitioning Facility or IBM pureScale deployments.
- Chapter 8: Implementing Workload Management demonstrates how to configure DB2 Workload Manager and how to use advanced and reporting features for managing workload performance. It has detailed guidance for how to help ensure the stability and predictability of your database system.
- Chapter 9: Monitoring packaged database application systems covers special integrations available, offerings.of ready-for-use configurations for SAP, IBM WebSphere, IBM Cognos, IBM InfoSphere DataStage, and InfoSphere SQL Warehouse applications.
- Appendixes: Contains more information about differences in using the tools with DB2 9.5 versus DB2 9.7 or later, information about troubleshooting, and other frequently asked questions. You can also learn how to access and download additional web material associated specifically with this book.

2

Planning your InfoSphere Optim Performance Manager deployment

This chapter provides important information for planning your deployment of InfoSphere Optim Performance Manager. It establishes the requirements of each of the InfoSphere Optim Performance Manager components and helps you answer several questions:

- What products do I want or need to install?
- How can I best prepare for installation?
- What kind of system resources are required for the InfoSphere Optim Performance Manager server?
- How can InfoSphere Optim Performance Manager comply with my company's security requirements?

2.1 Installation roadmaps

Installation roadmaps provide an overview of the required installation and configuration tasks you must follow based on your goals and the products that you purchase. By first choosing a roadmap, you can decide what components of InfoSphere Optim Performance Manager you need, obtain the correct installation images, and verify that your computers meet installation requirements.

The installation process of InfoSphere Optim Performance Manager installs the server component and sets up the DB2 repository database. After installation, you add databases and configure them for monitoring using the InfoSphere Optim Performance Manager web console. InfoSphere Optim Performance Manager Extended Edition consists of the following components:

- Server component: InfoSphere Optim Performance Manager
- Client components:
  - Data Tools Runtime Client
  - InfoSphere Optim Performance Manager Extended Insight plug-in for Tivoli Enterprise Portal
  - Optional: Performance Expert Client

All components must be installed separately because they are typically installed on separate systems.

The client component installations do not require separate license activation.

The InfoSphere Optim Performance Manager requires a license. You must activate the license whether you install InfoSphere Optim Performance Manager Enterprise Edition or the Extended Edition. The latter installation package contains the activation kit for the Extended Insight feature.

2.1.1 New or direct installation

New installations require different resources depending on the roadmap you use. Your roadmap choices are as follows:

- Installing InfoSphere Optim Performance Manager only

  Use this roadmap as an overview of the tasks for installing only InfoSphere Optim Performance Manager server software. You might use this roadmap if you purchased InfoSphere Optim Performance Manager Enterprise Edition, if you want to do a proof-of-concept installation, or if you want to see results quickly with the option to install additional components later.

- Installing InfoSphere Optim Performance Manager and InfoSphere Optim Performance Manager Extended Insight

  Use this roadmap as an overview of the tasks for installing InfoSphere Optim Performance Manager Extended Edition and Data Tools Runtime Client, which supports the Extended Insight features.

- Installing InfoSphere Optim Performance Manager Extended Edition and all optional components

  Use this roadmap for an overview of the tasks for installing InfoSphere Optim Performance Manager and all optional components, such as Data Tools Runtime Client, DB2 Performance Expert Client, and CIM server.

- Integrating InfoSphere Optim Performance Manager with ITCAM for Application Diagnostics and IBM Tivoli Composite Application Manager (ITCAM) for Transactions in a Tivoli Enterprise Portal Console

  Use this roadmap as an overview of the tasks for integrating InfoSphere Optim Performance Manager with ITCAM for Application Diagnostics and ITCAM for Transactions in a Tivoli Enterprise Portal Console for end-to-end transaction monitoring.

- Installing and configuring InfoSphere Optim Performance Manager to work with IBM Tivoli OMEGAMON® XE for DB2 Performance Expert on z/OS®

  Use this roadmap as an overview of the tasks for installing and configuring InfoSphere Optim Performance Manager Extended Edition and Data Tools Runtime Client to work with Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS.

Installation details for each of these roadmaps are in the InfoSphere Optim Performance Manager information center:

`http://pic.dhe.ibm.com/infocenter/perfmgmt/v5r1/index.jsp?topic=%2Fcom.ibm.datatools.perfmgmt.installconfig.doc%2Froadmaps.html`

2.1.2 Update installation

If you have Optim Performance Manager V5.1.1.1 (and earlier) installed, you can update to Optim Performance Manager V5.2 by installing the Optim Performance Manager Extended Edition V5.2 installation package and continue using the existing repository database.

You obtain the same set of functional features whether you perform a new installation or update. If you are new to Optim Performance Manager, the best approach is to use the direct or new installation option instead of using the update option. The remainder of this chapter describes prerequisites and requirements for a new installation.

Updating from InfoSphere Optim Performance Manager V5.1.x or V4.1.1 and later

If you have an existing installation of Optim Performance Manager V5.x, or V4.1.x, you can upgrade with the V5.2 installation program. After the installation program runs, you must do additional setup tasks:

- Review user access and authentication.

  InfoSphere Optim Performance Manager V5.2 introduced two privileges, *Is Database Owner* and *Can Enable Automatic Collection* privileges. In addition, the OPERATOR role has been redefined. In prior versions, the VIEWER and OPERATOR roles were the same.

  After installation, determine if you need to grant the additional privileges or reassign new roles to your current user ID. See 2.6, “User security” on page 55 for details about the new privileges and the OPERATOR role.

- Review your data collection method.

  You can choose to use the in-memory metrics collection method for any DB2 9.7 (or later) monitored databases that are using the snapshot collection method.

  For monitored databases on DB2 9.7 or later, collecting in-memory metrics is the default collection method. Collecting in-memory metrics has less overhead on the monitored database and on the disk storage in the repository database than the snapshot collection method, which was the only collection method available in Optim Performance Manager V4.1.

  For monitored databases on DB2 9.5 or previous, only the snapshot collection method is available.

Updating from InfoSphere Optim Performance Manager 4.1.0.x

Starting in InfoSphere Optim Performance Manager Version 4.1.1, running the InfoSphere Optim Performance Manager web console application on WebSphere Application Server is no longer supported. The upgraded InfoSphere Optim Performance Manager web console application runs as a separate process from WebSphere Application Server.

Aside from the user access and collection method, you must consider how the change from using WebSphere Application Server to using a separate process might affect your environment. The information center topic ("Replacement of WebSphere Application Server as the web console application in InfoSphere Optim Performance Manager Version 4.1.1") details the changes that can occur when you upgrade to InfoSphere Optim Performance Manager v5.x from Optim Performance Manager 4.1.0. It is in the following location:

http://pic.dhe.ibm.com/infocenter/perfmgmt/v5r1/topic/com.ibm.datatools.perfmgmt.installconfig.doc/app_server_support.html

2.1.3 Migrating from Performance Expert

You can migrate to all InfoSphere Optim Performance Manager versions from Performance Expert. Migration has the following meaning:

- The performance database of Performance Expert is used for Optim Performance Manager and updated to the enhanced database schema of Optim Performance Manager.
- On Linux and UNIX, the working directory of Performance Expert Server is used by Optim Performance Manager. On Windows, the default location is used as the working directory; see "Parameter summary" on page 48. The important property files are copied to the new location.

The migration is possible only if the same DB2 instance used for Performance Expert Server is used for Optim Performance Manager.

> **Important:** You must use the Advanced Installation mode to migrate. See migration option information in *IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows*, SG24-7925, for the results and further considerations of migrating to InfoSphere Optim Performance Manager.

To use DB2 Performance Expert Client, you must ensure that the monitored databases are configured to use the snapshot collection method. You cannot use DB2 Performance Expert Client with any databases that are configured to use the in-memory collection method.

The default collection method for monitored databases on DB2 9.7 or later is in-memory metrics. You can continue to use Performance Expert Client even if the monitored database is DB2 9.7, if during migration, the data collection was not yet changed to in-memory metrics. If during migration, the collection method was changed to in-memory metrics, you can still continue to use Performance Expert Client in a limited view. You can view the following information:

- Available short-term history data based on snapshot data until it is deleted according to retention time
- Available long-term history data in the Performance Warehouse until manually deleted from the Performance Warehouse

2.2 Prerequisites

There are general prerequisites for installing and running InfoSphere Optim Performance Manager.

Today, product changes can happen before documentation is updated, which means information becomes inaccurate after only a few months. System and software requirements for products change as other software and hardware enhancements are made. To provide you with the latest information, this section indicates general requirements and where to find the latest information. Check the IBM support portal for detailed information about downloadable parts and versions for InfoSphere Optim Performance Manager at the following web addresses:

- InfoSphere Optim Performance Manager (Extended Edition):

  http://www.ibm.com/support/docview.wss?uid=swg24033769
- InfoSphere Optim Performance Manager (Extended Insight):

  http://www.ibm.com/support/docview.wss?uid=swg24033773

2.2.1 InfoSphere Optim Performance Manager

The IBM Software Product Compatibility Report (SPCR) tool can help you locate complete lists of supported operating systems, system requirements, prerequisites, and optional supported software for InfoSphere Optim Performance Manager versions. These reports are kept updated with the latest information about the product. The following sections show you how to access these reports.

Hardware and operating system

The best approach is to install Optim Performance Manager in a separate physical or virtual server from the monitored DB2 database. This approach prevents Optim Performance Manager from sharing CPU, memory, and disk resources with a monitored database, thus allowing it to collect unbiased database performance data.

You can install Optim Performance Manager in the UNIX (IBM AIX®, HP-UX, and Solaris), Linux (Red Hat, SUSE), and Windows environments. The size of the server on which it is installed generally depends on the following factors:

- The number of monitored DB2 databases
- The number of partitions on your monitored DB2 databases
- The number of DB2 objects on your monitored DB2 databases
- The number of monitoring functions that are activated in Optim Performance Manager
- The volume of workload against your monitored database (for example, number of SQL statements per minute)

To find the supported platforms and versions for InfoSphere Optim Performance Manager, use the SPCR tool:

http://pic.dhe.ibm.com/infocenter/prodguid/v1r0/clarity-reports/report/html/osForProduct?deliverableId=1336056338621

Then, select the OS family first to find the relevant software for each platform. As an example, to find the supported Linux versions, click **Linux Family**, select the OS, for example, Red Hat Enterprise Linux (RHEL), and then click the icon for additional details (such as fix packs and maintenance levels), as demonstrated in Figure 2-1 on page 38.

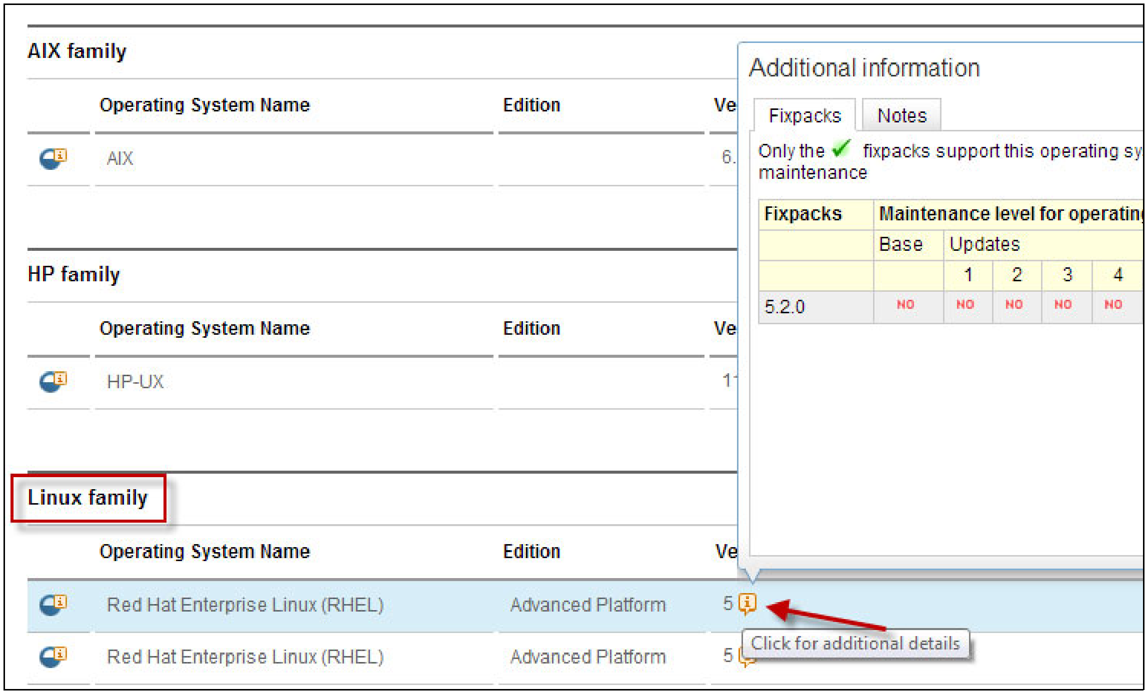

Figure 2-1 Supported platforms

DB2 data server for Optim Performance Manager

Optim Performance Manager uses a DB2 database as a repository for storing collected performance data and also its own configuration information. If the server where Optim Performance Manager will be installed already contains a copy of the DB2 server product, you can use it. Otherwise, the product includes a restricted use license of DB2 Enterprise Server Edition Version 10.1.

Starting with InfoSphere Optim Performance Manager V5.2, the DB2 9.1 and 9.5 for Linux, UNIX, and Windows versions are no longer supported to be used as the repository database. You can specify a DB2 V9.7 or V10.1 as the repository database, with specific fix packs applied. The table of supported system requirements is shown in Figure 2-2.

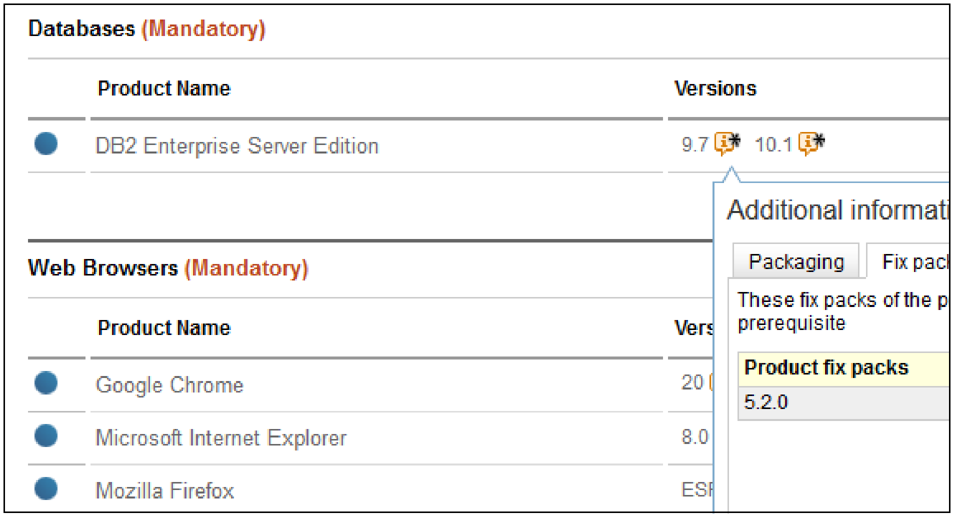

Figure 2-2 SPCR, supported repository database versions

For supported DB2 versions, click **Related Software** at the following address:

http://www.ibm.com/support/docview.wss?uid=swg27035831

The Software Product Compatibility Report (SPCR) contains a list of the supported database connectors, as shown in Figure 2-3. Click the **Additional Information** icon to view related fix pack levels, OS restrictions, packaging and other notes about this DB2 version.

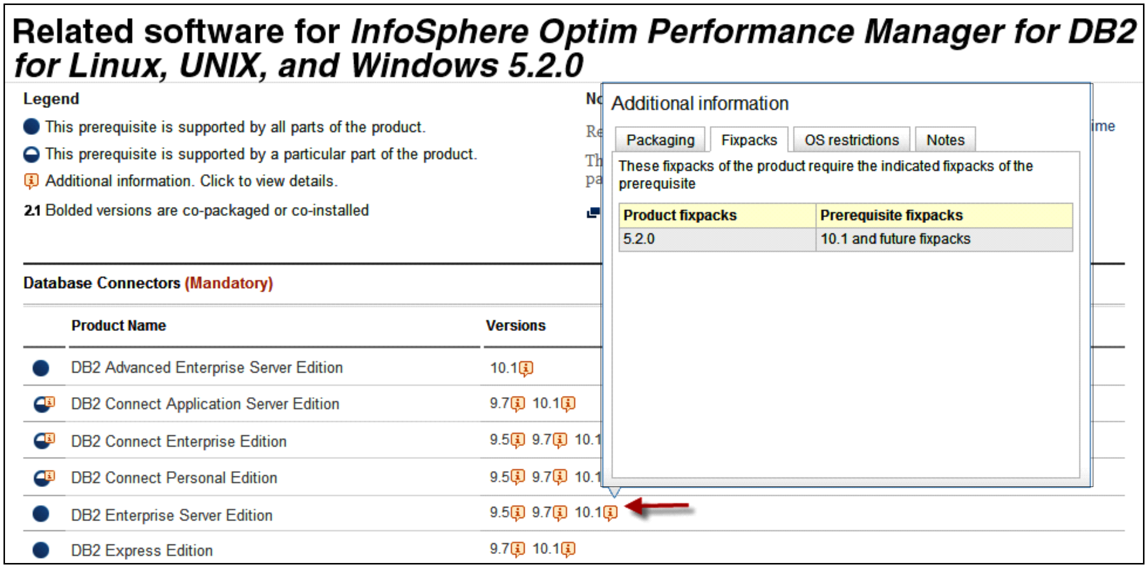

Figure 2-3 Software Product Compatibility Report (SPCR)

> **Compatibility:**
>
> - The DB2 instance where the repository server runs might not run in Oracle compatibility mode. For more information about Oracle compatibility mode, see the DB2 information center:
>
>   `http://pic.dhe.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.apdv.porting.doc/doc/c0052882.html`
>
> - Optim Performance Manager requires the DB2 Enterprise Server Edition product, because its repository database uses DB2 Enterprise Server Edition features such as table partitioning.

Monitored DB2 database

The following DB2 data servers are supported for the DB2 instances to be monitored by InfoSphere Optim Performance Manager. 64-bit DB2 instances are supported on each of these data servers. The 32-bit DB2 instances are supported only on Linux on IBM System x® and Windows.

- DB2 Advanced Enterprise Server Edition 10.1, 9.7 and future mod levels and fix packs for each version
- DB2 Enterprise Server Edition 10.1, 9.7, 9.5 and future mod levels and fix packs for each version
- DB2 Express Edition Version 10.1, 9.7 and future mode levels and fix packs for each version
- DB2 Workgroup Server Edition Version 10.1, 9.7, 9.5, and future mod levels and fix packs for each version
- DB2 for z/OS Version 10.1 and 9.1
- DB2 pureScale Feature 9.8, 10.1 and future fix packs

For the most recent list of supported databases for monitoring, scroll to the Managed Resources section of the System Requirements SPCR (as shown in Figure 2-4):

http://pic.dhe.ibm.com/infocenter/prodguid/v1r0/clarity-reports/report/html/prereqsForProduct?deliverableId=1336056338621

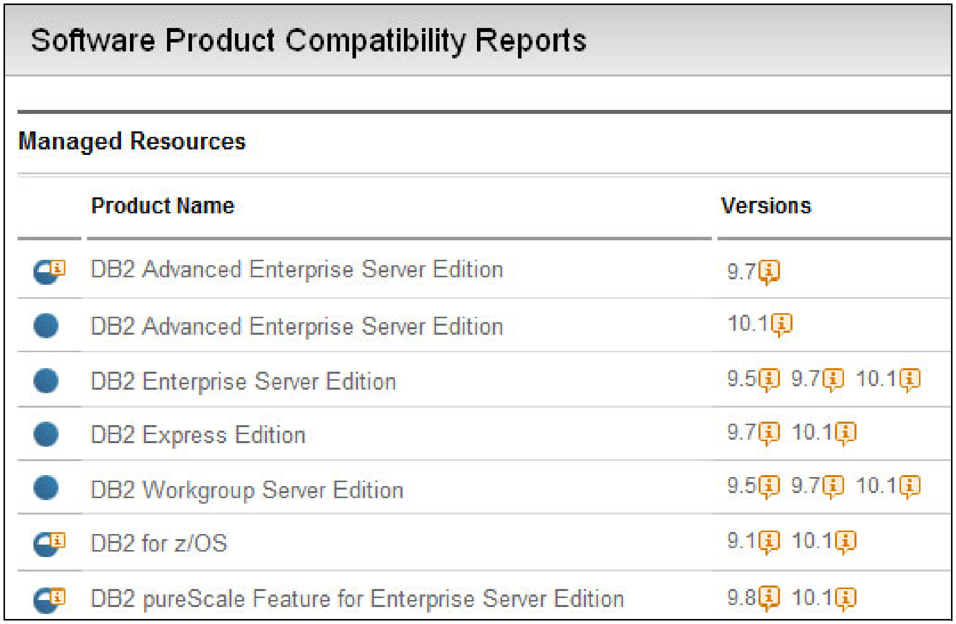

Software Product Compatibility Reports

Managed Resources

Product Name	Versions
DB2 Advanced Enterprise Server Edition	9.7
DB2 Advanced Enterprise Server Edition	10.1
DB2 Enterprise Server Edition	9.5 9.7 10.1
DB2 Express Edition	9.7 10.1
DB2 Workgroup Server Edition	9.5 9.7 10.1
DB2 for z/OS	9.1 10.1
DB2 pureScale Feature for Enterprise Server Edition	9.8 10.1

Figure 2-4 Supported Monitored Databases

Web browsers

The web interface of Optim Performance Manager is supported in the following web browsers:

- Google Chrome Version 20
- Mozilla Firefox Version 3.6 or later with Adobe Flash Player 10 or later
- Microsoft Internet Explorer Version 9.0 with Adobe Flash Player 10.0 or later

For the latest versions of the supported web browsers, scroll to the Web Browsers section of the following System Requirements SPCR address:

http://pic.dhe.ibm.com/infocenter/prodguid/v1r0/clarity-reports/report/html/prereqsForProduct?deliverableId=1336056338621

2.2.2 Optim Performance Manager Extended Insight

Optim Performance Manager Extended Insight enables end-to-end monitoring of DB2 database applications from the following generic database client environments:

- JCC applications V9.5.5 or V9.7.0 or later
- CLI/ODBC applications that use Data Server Client packages for Version 9.7 Fix Pack 2 or later
- NET applications that use the DB2 .NET provider from the IBM Data Server Client packages Version 9.7 Fix Pack 3 or later
- Embedded SQL applications that use Data Server Client Packages for Version 9.7 Fix Pack 6 or later
- WebSphere Application Server Versions 6, 7, and 8 for Linux, UNIX, and Windows, and z/OS

  Some fix packs or patches are required to enable Extended Insight in some versions. The most recent list of supported WebSphere Application Server versions is available in the Application Servers section of the SPCR, as shown in Figure 2-5 on page 43. See the System Requirements SPCR for InfoSphere Optim Performance Manager Extended Insight V5.2 address:

  http://pic.dhe.ibm.com/infocenter/prodguid/v1r0/clarity-reports/report/html/prereqsForProduct?deliverableId=1336057136339 h

Software Product Compatibility Reports

Application Servers

	Product Name	Versions
	WebSphere Application Server	6.1 7.0 8.0 8.5
	WebSphere Application Server - Express	6.1 7.0 8.0 8.5
	WebSphere Application Server Network Deployment	6.1 7.0 8.0
	WebSphere Application Server for Developers	8.5
	WebSphere Application Server for z/OS	6.1 7.0 8.0 8.5

Figure 2-5 Supported Application Servers

> **Note:** Application servers such as WebLogic or Tomcat are treated by InfoSphere Optim Performance Manager Extended Insight as generic Java applications.

- Data Connectors

  Specific versions of the JDBC drivers are required so that InfoSphere Optim Performance Manager can send the information from the client to the Repository Server. At a minimum, the following driver versions are supported:

  – JCC applications V9.5.5 or V9.7.0 or later

  – CLI/ODBC applications that use Data Server Client packages for Version 9.7 Fix Pack 2 or later

  – NET applications that use the DB2 .NET provider from the IBM Data Server Client packages Version 9.7 Fix Pack 3 or later

  – Embedded SQL applications that use Data Server Client Packages for Version 9.7 Fix Pack 6 or later

The current list of DB2 data connectors, depicted in Figure 2-6 on page 44, is available in the Database Connectors section of the SPCR at the following location:

http://pic.dhe.ibm.com/infocenter/prodguid/v1r0/clarity-reports/report/html/prereqsForProduct?deliverableId=1336057136339

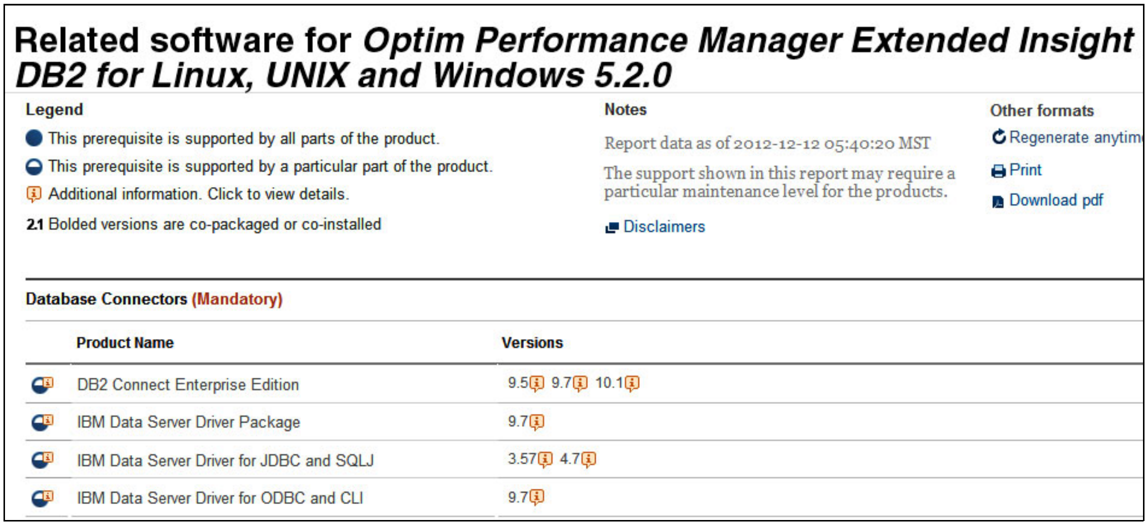

Figure 2-6 Supported Data Connectors

Extended Insight with other systems

InfoSphere Optim Performance Manager supports end-to-end monitoring of several applications, such as SAP applications, Cognos reporting applications, InfoSphere Warehouse, or Information Server with the following versions:

- SAP kernel Version 7.0 SR3 or later
- Cognos Version 8.4 Fix Pack 2 or later
- InfoSphere Warehouse Version 9.7 Fix Pack 1
- Information Server Version 8.5

The System Requirements SPCR for InfoSphere Optim Performance Manager Extended Insight V5.2, Reporting and Analysis section has the latest information for the software. The SPCR link is at the following address:

`http://pic.dhe.ibm.com/infocenter/prodguid/v1r0/clarity-reports/report/html/prereqsForProduct?deliverableId=1336057136339`

2.2.3 Integration with InfoSphere Optim Query Workload Tuner

Installing InfoSphere Optim Query Workload Tuner Version 3.2 installs two components on your workstation:

- IBM Data Studio full client, Version 3.2
- License activation kit for InfoSphere Optim Query Workload Tuner Version 3.2

Although you can launch InfoSphere Optim Query Workload Tuner from InfoSphere Optim Performance Manager by clicking **Tune** or **Tune All** in the InfoSphere Optim Performance Manager web console, the following items must be met:

- IBM Data Studio client must be installed on the system where the browser that you are using InfoSphere Optim Performance Manager is installed.
- The DB2 database or subsystem that your SQL statements run against must be configured for query tuning.
- The IBM Data Studio client must be running, and the embedded HTTP server must be active. To check whether the embedded HTTP server is active, move your cursor over the embedded HTTP server button in the toolbar at the top of the IBM Data Studio client. Read the tooltip that is displayed. If the server is not running, click the button.
- The database or subsystem that the SQL statements run against must be enabled for monitoring.
- The web console user is authorized to access the performance metrics stored in the repository database. To be authorized, your user ID must be a member of the user group that was authorized for accessing the repository database. This user group authorization takes place at the time InfoSphere Optim Performance Manager is installed.

If you have a previous version of IBM Data Studio, or its predecessor, IBM Optim Development Studio, you cannot use the update feature in Installation Manager. You will need to install the Data Studio 3.2 that includes the InfoSphere Optim Query Workload Tuner package.

You have several installation options for InfoSphere Optim Query Workload Tuner. Those options are covered in 3.4, “Installing and configuring IBM InfoSphere Optim Query Workload Tuner” on page 127.

2.2.4 Integration with InfoSphere Optim Configuration Manager

To be able to launch InfoSphere Optim Configuration Manager from InfoSphere Optim Performance Manager, identify the InfoSphere Optim Configuration Manager server in the InfoSphere Optim Performance Manager Web Console services setup.

In the InfoSphere Optim Performance Manager web console, click **Open** → **Services**, and select the row for InfoSphere Optim Configuration Manager. The Configure button will be enabled. Click **Configure** and supply the URL for the InfoSphere Optim Configuration Manager web console, as shown Figure 2-7 on page 46.

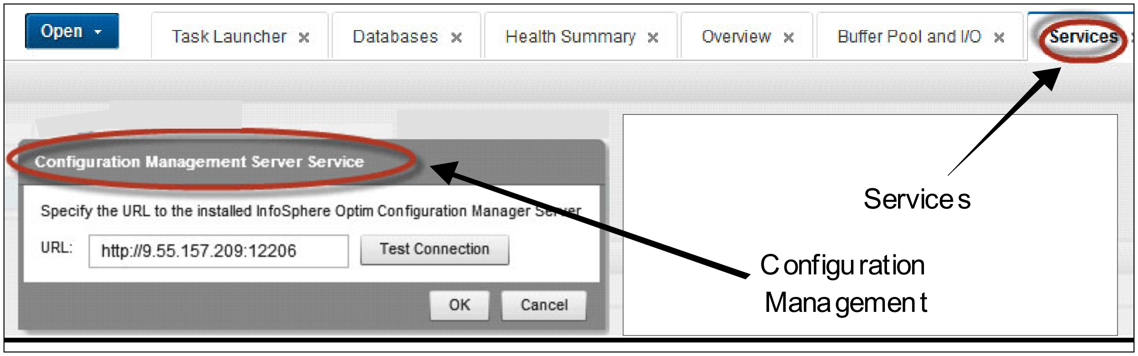

Figure 2-7 Configure the Optim Configuration Manager Service

When installing the Data Tools Runtime Client, described in 2.1.1, "New or direct installation" on page 33, you are prompted whether to configure the client for InfoSphere Optim Configuration Manager, as shown in Figure 2-8.

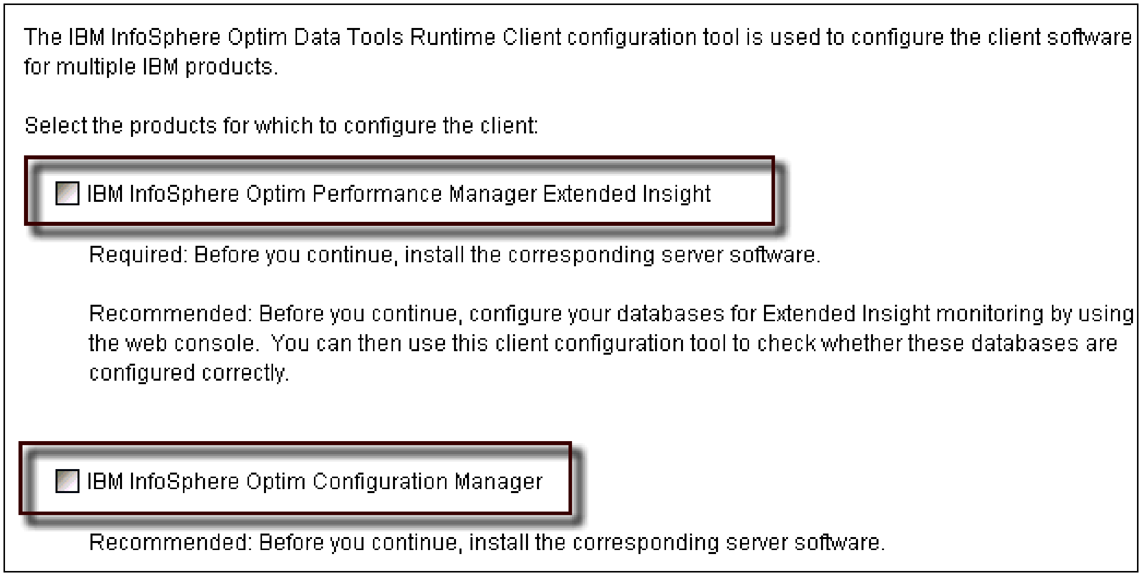

Figure 2-8 Configure the client for InfoSphere Optim Configuration Manager

During this configuration, the host address and port for the InfoSphere Optim Configuration Manager is specified. At the end of the configuration, a URL is generated. This is the URL you specify during the service configuration of the InfoSphere Optim Configuration Manager.

2.3 Installation parameters

During InfoSphere Optim Performance Manager installation or update, the installer requires certain inputs or parameters to complete the installation. You have the choice to accept default values for these parameters or to customize them during install.

Several of the parameters that you set during Optim Performance Manager installation are described in this section.

DB2 Instance selection

Optim Performance Manager requires a DB2 instance to host the repository database. During installation, you can specify which DB2 instance you want to use. If the DB2 instance does not yet exist, the Optim Performance Manager installation creates it.

DB2 user specification

The Optim Performance Manager installer uses the user that you specify to create the repository database in the selected DB2 instance. Later at Optim Performance Manager run time, this user is used by Optim Performance Manager to connect to the repository database to access the collected data. This user must have SYSADM authority on the DB2 instance. Learn more about user privileges in Optim Performance Manager in 2.6, “User security” on page 55.

Advanced only installation

This section describes advanced only installation of the following items:

- Repository database specification (advanced only)
- Table space type selection (advanced only)
- Working directory specification (advanced only)
- Performance Expert Client group specification (advanced only)

Repository database specification (advanced only)

The repository database is the database of Optim Performance Manager to store the collected performance data. The Optim Performance Manager installer creates this database. You can use the advanced installation mode to specify the following settings for the repository database:

- Database name
- Database location
- Table space location for small SMS table spaces storing control and metadata

Table space type selection (advanced only)

For each database that Optim Performance Manager monitors Optim Performance Manager creates a set of table spaces in the repository database. The table spaces are created when you configure a database for monitoring after Optim Performance Manager is installed. During installation you can specify which type of table spaces (SMS, DMS, or Automatic Storage) Optim Performance Manager must create. These table spaces can grow to multiple gigabytes (GBs) in size. See 2.4, “Capacity planning” on page 49 to learn how large these table spaces can get. See 2.5, “Storage options” on page 52 for more details about these table spaces.

Working directory specification (advanced only)

The Optim Performance Manager repository server uses the directory that you specify to write log and trace files during run time. In addition, this directory contains the property files that the Optim Performance Manager repository server uses.

Performance Expert Client group specification (advanced only)

If you want to use Performance Expert Client, specify an existing group account which will have the permission to log on from Performance Expert Client to the InfoSphere Optim Performance Manager repository server. The value is used only if you install and use the DB2 Performance Expert Client.

Parameter summary

Table 2-1 summarizes the installation parameters and describes the defaults for the typical installation mode.

Table 2-1 Installation parameter summary

Parameter	Specify always	Specify in advanced mode	Default value for typical mode
DB2 instance	Yes	-	-
DB2 user	Yes	-	-
Repository database name	-	Yes	PERFDB or PERFDB[x] If PERFDB exists then x is replaced with a positive number
Repository database location	-	Yes	Default database path (DFTDBPATH) from database manager configuration
Repository database table spaces location	-	Yes	By default, this directory is the same directory as the repository database working directory.

Parameter	Specify always	Specify in advanced mode	Default value for typical mode
Table space type	-	Yes	Data-managed space (DMS)
Working directory			► Windows: `<OPM install dir>\RepositoryServer\instances` ► Linux or UNIX: `<Home directory of DB2 instance owner>/opm/v5`

Review the installation planner topic in the information center, which contains a worksheet to facilitate the installation planning process; it has a complete table of installation parameters. See it at the following address:

http://pic.dhe.ibm.com/infocenter/perfmgmt/v5r2/topic/com.ibm.datatools.perfmgmt.installconfig.doc/install_planner.html

2.4 Capacity planning

Each monitored database, database partition (in DPF instances), or member (in pureScale instances) requires operating system resources (for example, memory, CPU, disk space, and network bandwidth) on the InfoSphere Optim Performance Manager server to collect, process, and store the monitoring data. Be sure your InfoSphere Optim Performance Manager server has sufficient capacity to monitor the intended environment.

Certain factors affect resource requirements for InfoSphere Optim Performance Manager, and the InfoSphere Optim Performance Manager Concierge program. Various strategies and best practices are necessary when deploying InfoSphere Optim Performance Manager. This information is described next. This section uses monitoring configuration terminology, which is described in 3.2, “Configuring InfoSphere Optim Performance Manager” on page 93.

2.4.1 Parameters that affect InfoSphere Optim Performance Manager resources

The system resources that Optim Performance Manager needs depend mostly on the following parameters:

- Type of monitoring information that is being collected
- Granularity of the monitoring information that is collected
- Interval at which the monitoring information is being collected

- Retention time that the collected monitoring information is stored
- Number of users concurrently using Optim Performance Manager web console
- Characteristics of the database that is being monitored (for example, the number of database objects), and
- Characteristics of the application or workload that is being monitored (for example, the rate of executing transactions and statements)

2.4.2 Estimating resource requirements

The following sections of *IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows*, SG24-7925, describe how to manually estimate resource requirements for the InfoSphere Optim Performance Manager server:

- 2.4.1 Factors influencing capacity planning of Optim Performance Manager servers
- 2.4.2 Hard disk requirement estimation
- 2.4.3 CPU requirement estimation
- 2.4.4 Memory requirement estimation

Since the publication of that book, InfoSphere Optim Performance Manager has changed its data model; the capacity planning calculation is mostly automated.

2.4.3 Concierge Program for InfoSphere Optim Performance Manager

The best way to estimate the resources required for your InfoSphere Optim Performance Manager deployment is by working with your IBM representative.

IBM provides a no-charge service, the Concierge Program, to assist customers with their initial or first InfoSphere Optim Performance Manager deployment.

You, with the help of your IBM representative, complete a capacity planning questionnaire that contains instructions that describe how to obtain information about your workloads and usage. Download the current Capacity Planning Questionnaire from the following location:

https://www.ibm.com/developerworks/mydeveloperworks/files/app?lang=en#/person/27000403P9/file/f7d3089b-aa5e-4496-b1b8-44760c2ef226

When completed, your IBM representative will then perform capacity planning calculations and return to you the estimated CPU, memory, and disk capacity requirements for an InfoSphere Optim Performance Manager server in your environment. The result of this exercise is a deployment roadmap that is customized to your environment.

2.4.4 Enabling compression

InfoSphere Optim Performance Manager deployments for large scale environments (for example DPF environments, high volume transactional systems, or large numbers of monitored servers) can result in very large storage requirements for collecting and retaining performance data. The DB2 Storage Optimization Feature is highly effective in reducing the storage requirements for performance data.

In InfoSphere Optim Performance Manager v5.1.1 and beyond, compression is automatically enabled for the InfoSphere Optim Performance Manager tables when the DB2 for Linux, UNIX, and Windows is licensed for this feature.

DB2 for Linux, UNIX, and Windows v9.7 and v10.1, Advanced Enterprise Edition includes the compression feature. Both the Workgroup and Enterprise editions do not include this feature. To determine if compression is enabled for the InfoSphere Optim Performance Manager repository database, issue the **`db2licm -l show detail`** command.

2.4.5 Network bandwidth

When monitoring a large number of systems, or few very large and active systems, particular attention has to be paid to the bandwidth of the network adapter card on the InfoSphere Optim Performance Manager server. If the InfoSphere Optim Performance Manager server's network card is running near full capacity, then InfoSphere Optim Performance Manager will be unable to receive the data from the monitored data servers in a timely manner and consequently the data will not be processed or stored. In these cases, you might need to do the following tasks:

- Reduce the amount of monitoring data that is being collected by adjusting monitoring profiles and sampling intervals.
- Increase the capacity of the network adapter card.
- Split the monitoring of your database environment across multiple InfoSphere Optim Performance Manager servers.

2.5 Storage options

For each database that Optim Performance Manager monitors, Optim Performance Manager creates two table spaces in the repository database, one for holding the short-term history data and one for holding the long-term history data. The table spaces are created when you configure a database for monitoring after InfoSphere Optim Performance Manager is installed. During installation, in advanced installation mode, you can specify the type of table spaces (SMS, DMS, or automatic storage) InfoSphere Optim Performance Manager must create. If you use the typical installation, then the table space type defaults to DMS.

2.5.1 Table space type selection

Each table space type has its characteristics. The DB2 topic at the following location provides a comparison:

http://pic.dhe.ibm.com/infocenter/db2luw/v10r1/index.jsp?topic=%2Fcom.ibm.db2.luw.qb.server.doc%2Fdoc%2Fr0006867.html

From a performance point of view, DMS and automatic storage table spaces are faster than SMS table spaces, especially for large tables. Because various tables in the Optim Performance Manager repository database can grow large, the best approach is to use DMS or automatic storage table spaces, considering performance.

Optim Performance Manager creates table spaces as a *large* table space. If you consider the maximum table space size, the DMS and automatic storage table spaces allow a higher maximum size than SMS table spaces when the table spaces are created as large table space. For SMS table spaces, the large option is not available. See the table space size comparison in DB2 information center:

http://pic.dhe.ibm.com/infocenter/db2luw/v10r1/topic/com.ibm.db2.luw.admin.dbobj.doc/doc/c0052381.html

If the way to reclaim storage is important for you, then SMS table spaces are the easiest to handle because the table space size shrinks when data is deleted from the table space. If capacity planning results in disk space shortage, consider the storage reclaiming behavior for your table space type decision. You might find that during InfoSphere Optim Performance Manager run time, you run out of storage although InfoSphere Optim Performance Manager deletes collected data from the repository database regularly and automatically. In this case, you must manually delete data from the repository database to free disk space.

Learn how DMS and Automatic Storage table spaces reclaim storage by going to the following address:

http://pic.dhe.ibm.com/infocenter/db2luw/v10r1/index.jsp?topic=%2Fcom.ibm.db2.luw.qb.server.doc%2Fdoc%2Fr0006867.html

One of the Optim Tooling products, InfoSphere Optim Configuration Manager, can also help you with discovering storage optimization opportunities such as reclaiming storage through REORG, compression, and identifying underused objects. This product integrates with InfoSphere Optim Performance Manager V5.1.1 and later.

2.5.2 Table space naming and usage

For each monitored instance, InfoSphere Optim Performance Manager creates the following table spaces to save the collected performance data:

- `SHORTTERM_<instance_id>`
- `LONGTERM_<instance_id>`

The `<instance_id>` variable is a unique positive number that InfoSphere Optim Performance Manager assigns to each monitored database. The first database that you add for monitoring most likely receives the instance ID 1. For this database, the `SHORTTERM_1` and `LONGTERM_1` table spaces are created.

All the collected performance data that can be displayed on the InfoSphere Optim Performance Manager Web console are saved in the `SHORTTERM_<instance_id>` table space. This table space can grow large. the Calculation of the required space for this table space is described in 2.4, “Capacity planning” on page 49. Retention times are set for the collected data and InfoSphere Optim Performance Manager deletes the data automatically when the retention time is reached.

The `LONGTERM_<instance_id>` table space stores the collected and aggregated data for long-term trend analysis and reporting. Because the data is stored in an aggregated format, the size of this table space is typically much smaller than the `SHORTTERM_<instance_id>` table space. Long-term data, which stores aggregation levels 2, 3 and 4 data, is not deleted automatically, Data pruning is done by InfoSphere Optim Performance Manager when the specified retention time for each aggregation level is reached.

> **Note:** This table space is used only by InfoSphere Optim Performance Manager for monitoring databases in snapshot mode. In version 9.7 databases within-memory metrics mode, this table space is no longer used.

After you configure a database for monitoring, you can check the unique instance ID assigned to the database by InfoSphere Optim Performance Manager to learn which table spaces were created for the database. To check the ID, use the **peconfig -list** command from the RepositoryServer/bin directory of your InfoSphere Optim Performance Manager installation.

Within peconfig, use the **list** command that returns information including the instance ID, as in the following example:

```
Instance ID = 1
----------------
   Enabled = Yes
   Status = Inactive
   CIM Object Manager enabled = No
   Monitored Instance Alias = LOCALHOST_50000_PEDEMO
   Node, Host, Port/Service name = NODE7507, LOCALHOST, 50000
   Database, remote alias, local alias, EVM = PEDEMO, PEDEMO, PMDB3902, OFF
```

In this example, InfoSphere Optim Performance Manager uses instance ID 1 for the configured database. This means that the table spaces for this database are named SHORTTERM_1 and LONGTERM_1.

2.5.3 Table space location

If the DMS or SMS table space type is selected during installation, you specify the table space path during the configuration of a monitored database. InfoSphere Optim Performance Manager creates both table spaces in the specified path. You can either specify an explicit path or choose the default location. The default storage path is the main path of the repository database location.

> **Tip:** If you monitor multiple databases, the best approach is to specify an explicit path for each monitored database and ensure that the specified paths reside on multiple disks to avoid I/O bottlenecks. Also, be sure to keep the table spaces apart from the database log files.

For Automatic Storage table spaces, you already specified the paths during installation. When configuring a monitored database, InfoSphere Optim Performance Manager creates both table spaces using the automatic storage option. If you must add storage paths later, you can use the **ALTER DATABASE** command.

2.5.4 Table space DDL

The following examples are DDL statements for each table space type that InfoSphere Optim Performance Manager uses to create the table space. These examples can help you understand which parameters are used to create the table spaces:

- SMS:

```
CREATE REGULAR TABLESPACE SHORTTERM_1 IN NODEGROUP PENG PAGESIZE 8K MANAGED
BY SYSTEM USING ('SHORTTERM_1 $N')  BUFFERPOOL DATA
```

- DMS:

```
CREATE LARGE TABLESPACE SHORTTERM_1 IN NODEGROUP PENG PAGESIZE 8K MANAGED
BY DATABASE USING ( FILE 'SHORTTERM_1' 5000)  BUFFERPOOL DATA AUTORESIZE
YES
```

- Automatic Storage:

```
CREATE LARGE TABLESPACE SHORTTERM_1 IN NODEGROUP PENG PAGESIZE 8K MANAGED
BY AUTOMATIC STORAGE INITIALSIZE 40M BUFFERPOOL DATA AUTORESIZE YES
```

2.6 User security

InfoSphere Optim Performance Manager uses a combination of operating system or external security to *authenticate* its users, and DB2 system and database privileges to control or *authorize* access to the repository and monitored database objects.

InfoSphere Optim Performance Manager uses a tiered approach to control access to the web console first and then to the performance data that is presented in the web console. This section describes the authentication that is done for each user type. It describes how to authorize or grant the proper privileges to users who need to access the objects in both InfoSphere Optim Performance Manager and the monitored database.

InfoSphere Optim Performance Manager categorizes users:

- Web console users
- Monitored database users

These users are further defined into two categories:

- User IDs used during installation and configuration
- User IDs for monitoring

2.6.1 Web console users

The first tier of user access to InfoSphere Optim Performance Manager is the login ID to the web console. To set up web access, you must have a repository database set up with an authentication method such as local operating system, Lightweight Directory Access Protocol (LDAP), Kerberos, or NIS+.

Default administrative user

During installation, the InfoSphere Optim Performance Manager installer sets the authentication method. The default authentication method is through the repository database. You also specify an existing administrative user for the repository database.

This user has the following requirements:

- It must be an existing local, LDAP, Kerberos or NIS+ user.
- The user, or the group that the user belongs to, must already have SYSADM authority on the DB2 for Linux, UNIX, and Windows instance for the repository database.
- This user will be granted DBADM rights on the repository database.

This user is referred to as the *default administrative user.* This user is often the instance owner of the InfoSphere Optim Performance Manager repository database. When the installation completes, you log in to the web console for the first time with this user ID. During this initial login to the web console, the default administrative user can grant web console access to other users.

Viewer, Operator, Administrator

Immediately after the installation, only the default administrative user can log on to the InfoSphere Optim Performance Manager web console. This user then grants web console access to other users. These authorized users are collectively called *web console users.*

Web console users can be identified by their individual authorization IDs that have CONNECT privileges to the repository database, or by any groups or roles that they are part of in the repository database.

Additionally, web console users can have one of three basic privilege levels (viewer, operator, and administrator) in InfoSphere Optim Performance Manager. These privileges are configured in the Console Security dashboard, and they control the tasks that web console users can perform.

Viewer

The Viewer role is the minimum global privilege for every InfoSphere Optim Performance Manager web console user. A user who is assigned the Viewer role cannot change any global settings (for example, SMTP server settings). Viewers cannot see the historical monitoring information of any monitored databases that are disconnected. In addition, the Viewer role cannot create database connections.

Operator

The Operator user has the same basic privileges as the Viewer web console user with the added privilege to add database connections from the Databases panel. When an Operator adds a database connection, the user is granted *Is Database Owner* privilege for that database. Operators with the Is Database Owner privilege for a database can grant or revoke privileges for other web console users on that database.

Administrator

The Administrator role is a global privilege that allows the user to do any task in the InfoSphere Optim Performance Manager web console. A user who is assigned the Administrator role can do the following tasks:

- View any page.
- Configure SMTP information.
- Can change authentication method.
- View historical monitoring information of all disconnected databases.
- Add a database connection.
- Configure logs.

When an Administrator adds a database for monitoring, the user is granted the Is Database Owner privileges for that database. Administrators with the Is Database Owner privilege for a database can grant or revoke privileges for other web console users on that database.

The default administrative user that you specify during installation automatically receives the Administrator privilege. Therefore, this user can log on to the web console immediately after installation. This user can grant the console privileges to more users at any time after installation. To grant or revoke these privileges to web console users, use the Console Security dashboard in InfoSphere Optim Performance Manager, as shown in Figure 2-9 on page 58.

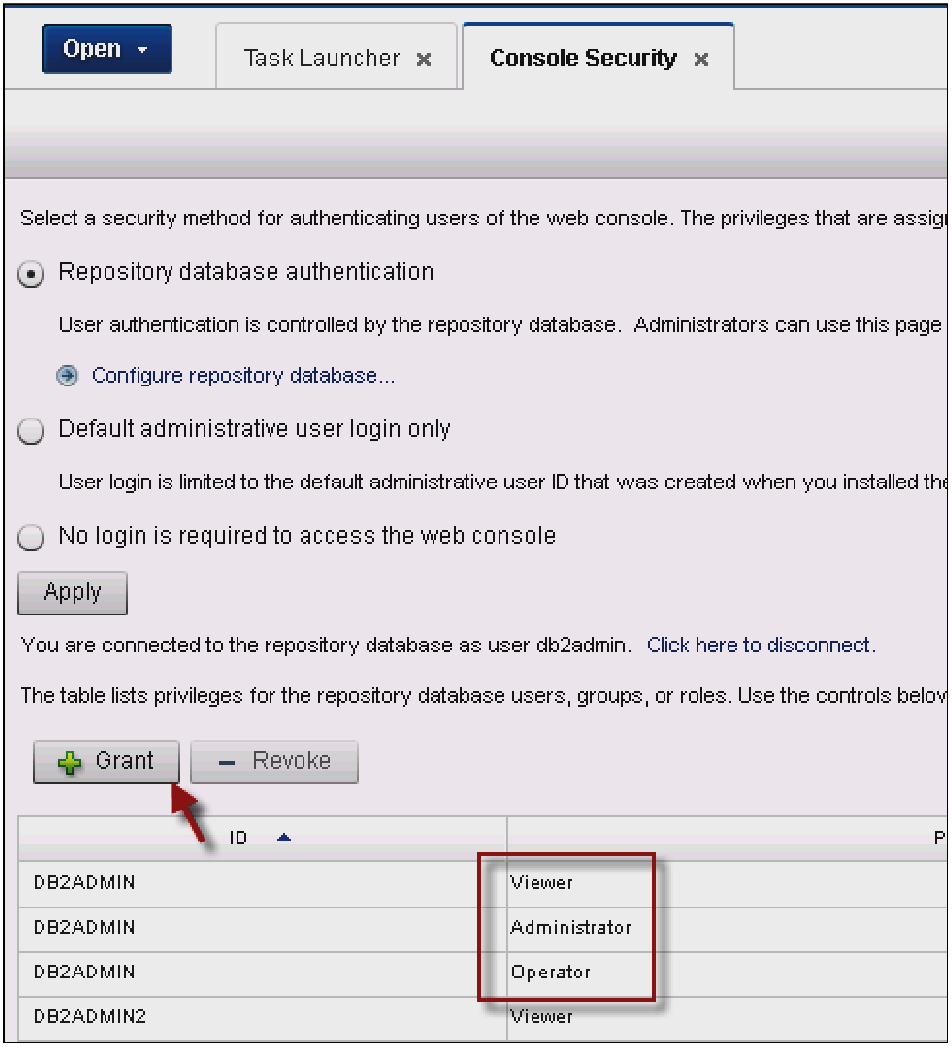

Figure 2-9 Console Security dashboard

By granting the Administrator, Operator, or Viewer privilege to a user, group or role, InfoSphere Optim Performance Manager grants execute rights to one of these user defined functions in the repository database:

- DSWEBSECURITY.CANVIEW
- DSWEBSECURITY.CANOPERATE
- DSWEBSECURITY.CANADMINISTER

Alternatively, you can grant the console privileges directly using the DB2 command GRANT EXECUTE, as in the following example:

```
GRANT EXECUTE ON FUNCTION DSWEBSECURITY.CANVIEW TO USER <user id>
```

For more information about granting and revoking Viewer, Operator, or Administrator privileges, see the “Configuring user access to the InfoSphere Optim Performance Manager web console” topic at the following address:

http://ibm.co/18vbLyX

Is Database Owner

This additional privilege is granted at the level of each monitored database, and is required before the following actions can be performed:

- Edit or delete a database from the Databases page
- Create or change a monitoring configuration for a database
- Run jobs against a database
- Schedule blackout events
- Run scripts for user-defined alert types

This privilege also allows users to enable automatic data collection for a database. The Is Database Owner privilege is automatically granted to Administrator and Operator web console users when they add a database connection.

Special consideration for LDAP and Kerberos authentication

During InfoSphere Optim Performance Manager installation you specify an existing administrative user for the repository database. This must be an existing local user or an Lightweight Directory Access Protocol (LDAP) user.

To enable LDAP authentication, you must configure the DB2 instance where the repository database will be located to use LDAP through the LDAP security plug-in or through transparent LDAP. You must configure this DB2 instance to use LDAP before you install InfoSphere Optim Performance Manager. The latest setup procedure for enabling LDAP authentication is found in the “Setting up LDAP authentication” topic in the information center:

http://pic.dhe.ibm.com/infocenter/perfmgmt/v5r2/topic/com.ibm.datatools.perfmgmt.installconfig.doc/ldap_setup.html

InfoSphere Optim Performance Manager also supports the Kerberos authentication protocol for the repository database and for monitored databases on DB2 for Linux, UNIX, and Windows. You can use the Kerberos authentication protocol if the repository database or the monitored database, or both, and the DB2 server are configured for Kerberos authentication.

The setup procedure to enable LDAP authentication is in the information center:

http://pic.dhe.ibm.com/infocenter/perfmgmt/v5r2/topic/com.ibm.datatools.perfmgmt.installconfig.doc/kerberos_authentication_setup.html

2.6.2 Monitored database users

The second tier of user access defines access and monitoring rights on the monitored databases. The performance data collection user and monitoring users are defined on the monitored database. Web console users select a monitored database and log on with a monitoring user ID to view performance data. InfoSphere Optim Performance Manager requires at least one, but usually two or more users defined on each monitored database.

Performance Data Collection user

When you add a database, you need to specify a *collection user* credential for that database which is used to collect performance, health and other information from the monitored database.

On the monitored database, the collection user must have the following privileges and authorities:

- For DB2 9.7 and later:
  - DBADM (with DATAACCESS and ACCESSCTL)

    This privilege is required to create and manage event monitors, access all monitoring functions, and grant or revoke monitoring privileges.

    Alternatively, the following minimum privileges are needed:

    - SQLADM for managing event monitors
    - WLMADM for Workload Manager Configuration tasks
    - DATAACCESS for accessing monitoring functions
    - ACCESSCTL to grant monitoring privileges such as canMonitor
    - CREATE_EXTERNAL_ROUTINE with IMPLICIT_SCHEMA to create FUNCTIONS and PROCEDURES such as canMonitor
- One of the following authorities, which are required to take snapshots and update DB2 configuration parameters.

  > **Note:** Even if you collect in-memory metrics, some snapshots are still taken.

  - SYSADM

    This privilege is required for taking snapshots and updating database and database manager configuration parameters. InfoSphere Optim

Performance Manager updates the monitor switches in the database manager configuration based on the monitoring settings that you specify. Updating the monitor switches requires SYSADM rights.

– SYSCTRL or SYSMAINT

  This privilege is required for taking snapshots and updating the database configuration parameters. If the specified user does not have SYSADM rights, you must ensure that InfoSphere Optim Performance Manager does not update the monitor switches as described in Setting monitor switches.

– SYSMON

  This privilege is required for taking snapshots.

If the specified user does not have SYSCTRL or SYSMAINT rights, you must ensure that InfoSphere Optim Performance Manager does not update database configuration parameters.

Depending on your monitoring configuration, InfoSphere Optim Performance Manager updates the `MON_*` database configuration parameters of your monitored database. When you configure a database for monitoring, the Configure Monitoring wizard lists the necessary changes for these parameters. To ensure that InfoSphere Optim Performance Manager does not update these parameters, update them manually before you configure the database for monitoring.

Additionally, if you use the Workload Manager monitoring profile, InfoSphere Optim Performance Manager sets the database configuration parameter WLM_COLLECT_INT to 0 if it is set to a different value.

> **Note:** The suggestion is *not* to use the instance owner ID as the collection user ID that is used by InfoSphere Optim Performance Manager for monitoring. If possible, use a unique authorization ID to ensure that InfoSphere Optim Performance Manager will not interfere with your ability to manage the database

To configure a DB2 v9.5 and previous databases for monitoring, consult the InfoSphere Optim Performance Manager information center website for the required authorities:

http://ibm.co/1eH4FR7

To change the monitoring configuration, you must specify the credentials of the InfoSphere Optim Performance Manager collection user, or have the Is Database Owner privilege on that database. Otherwise, any changes to the monitoring configuration cannot be saved.

Monitoring users

Monitoring users are web console users that have specific privileges to view and modify monitoring settings for a given database. A web console user must have the Is Database Owner privilege to assign privileges to other monitoring users. From any one of the dashboards, when a database is selected, you prompted to enter the credentials for a monitoring user. If a web console user already has the Is Database Owner privilege, then grant the individual “Can do” privileges to these users is unnecessary.

Varying levels of monitoring privileges can be granted to a monitoring user. The privilege level determines the access to the monitored data and tasks that the monitoring user can do within the web console. These privileges apply only after the user logs in to the InfoSphere Optim Performance Manager web console. These privileges are as follows:

- Can Monitor

  This privilege allows the user to view all windows and data.

- Can Manage Alerts

  This privilege allows the user to control alerts and their thresholds, which includes pruning and deleting alert instances. This privilege also includes the Can Monitor privilege.

- Can Monitor in Real Time

  This privilege allows a real-time refresh of performance data from the monitored database. This privilege also includes the Can Monitor privilege.

- Can Enable Automatic Data Collection

  This privilege allows the user to enable automatic data collection, which allows for historical monitoring of a monitored database.

> **Note:** The InfoSphere Optim Performance Manager collection user automatically receives the canMonitor and canManageAlerts privilege. Also, the is Database Owner privilege implicitly gives the user the Can Manage Alerts, Can Monitor in Real Time, and Can Enable Automatic Data Collection privileges.

Go to **Product Setup** → **Manage Privileges** dashboard in InfoSphere Optim Performance Manager. This dashboard is used to grant or revoke these privileges. On the Enforce Privileges tab, you can select whether privilege checking should be turned on or off for each database. Privilege checking can be disabled for unimportant databases or for test databases, but enable it otherwise. For example, Figure 2-10 on page 63 shows that only users explicitly granted the privileges listed for the GSDB database can perform those privileges.

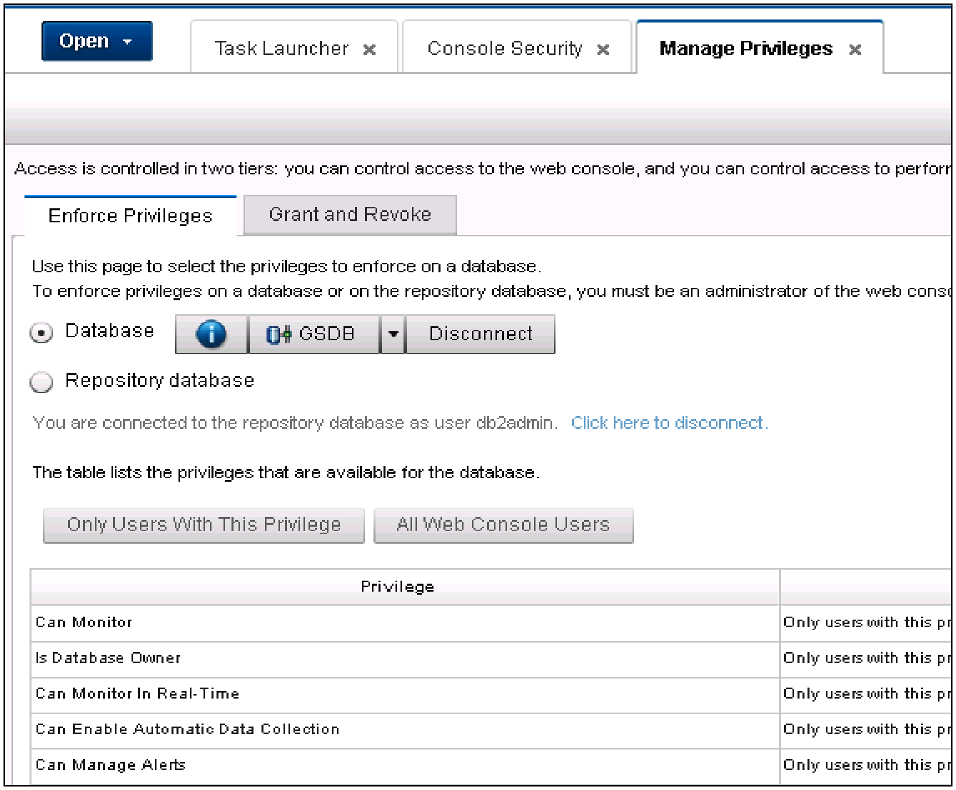

Figure 2-10 Manage Privileges dashboard: enforcing privileges

In the Grant and Revoke tab of the Manage Privileges dashboard, you can explicitly grant privileges to a user, group, or role. Figure 2-11illustrates granting the Monitor privilege to a user.

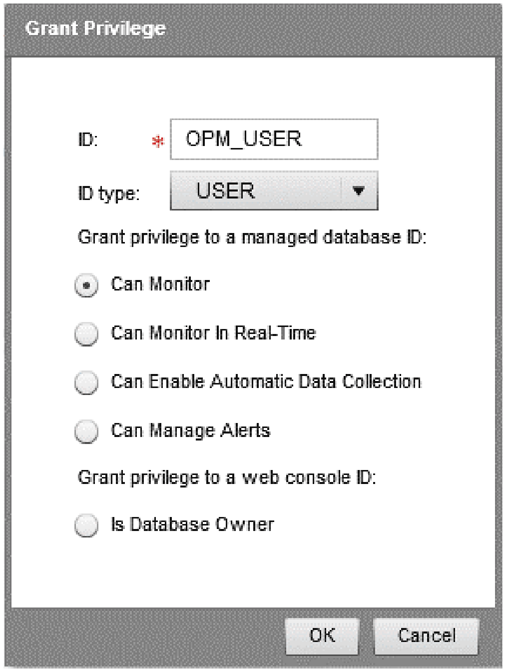

Figure 2-11 Manage Privileges dashboard, Grant Privilege dialog

InfoSphere Optim Performance Manager grants EXECUTE rights to one of these user defined functions that InfoSphere Optim Performance Manager creates in the monitored database:

- OPM.CAN_MONITOR
- OPM.CAN_MONITOR_IN_REALTIME
- OPM.CAN_MANAGE_ALERTS
- OPM.CAN_ENABLE_AUTOMATIC_COLLECTION

Alternatively, you can grant the privileges on the monitored database by using the DB2 command GRANT EXECUTE, as in the following example:

```
GRANT EXECUTE ON FUNCTION OPM.CAN_MONITOR TO USER <user id>
```

The user defined functions do not do any operation; they are in the monitored database to check the EXECUTE rights. So, for example, when you open a monitoring dashboard in the web console, select a database, and specify user credentials for this database, the canMonitor privilege is checked. If the user that you specify has EXECUTE rights on the CAN_MONITOR user defined function, the canMonitor privilege is confirmed and you are allowed to look at the monitoring data of this database. The canManageAlerts privilege is checked on the alert notification and configuration dashboards when you select a database and specify user credentials for this database.

2.6.3 Objects in the monitored database

When monitoring is enabled for a database, InfoSphere Optim Performance Manager creates database objects in the monitored database including event monitors, event monitor tables, user-defined functions (UDFs), and stored procedures.

Depending on the monitoring data that you want to collect, InfoSphere Optim Performance Manager creates the following event monitors after the database is configured and enabled for monitoring:

- For InfoSphere Optim Performance Manager:
  - Lock event monitor (only DB2 V9.7 or later)
  - Deadlock event monitor
  - Statistic event monitor (only DB2 V9.5 or later)
- For InfoSphere Optim Performance Manager Extended Insight:
  - Package cache event monitor (only DB2 V9.7 or later)
  - Transaction event monitor (only DB2 V9.7 or later)

These event monitors write the collected data into tables in the monitored database. InfoSphere Optim Performance Manager maintains these tables by deleting the data after InfoSphere Optim Performance Manager has read and saved it in the repository database. The interval that InfoSphere Optim Performance Manager saves the monitor data is short, for example every minute. By default, the tables are created in the default table space that DB2 chooses. Often it is USERSPACE1.

When configuring the database for monitoring using the InfoSphere Optim Performance Manager Web console, you can specify the table space that InfoSphere Optim Performance Manager must use to create the event monitor tables. The best approach is to specify a dedicated table space for the event monitors that InfoSphere Optim Performance Manager creates.

Note: For a partitioned system, this table space must exist on all partitions that InfoSphere Optim Performance Manager monitors.

The following event monitors are created only if you use Performance Expert Client and start SQL or DB2 Workload Manager (WLM) activity traces:

- Statement event monitor
- Activity event monitor

These event monitors generate lots of data and, therefore, write the collected data into files on the monitored system instead of into tables. You specify the path for these files during configuration of a monitored database if you enable the performance warehouse monitoring profile. If the monitored database is partitioned, this path must be available on each partition. The instance owner and fenced user must have read and write privileges to this path. InfoSphere Optim Performance Manager deletes the files after it reads and stores the contents to the repository database.

InfoSphere Optim Performance Manager creates user-defined functions (UDFs) and stored procedures in the monitored database for the following purposes and does not generate overhead on the monitored database:

- In the InfoSphere Optim Performance Manager schema, UDFs to control *Can Monitor*, *Can Monitor In Realtime, Can Manage Alerts, Can Manage Jobs* and *Can Enable Automatic Data Collection* privileges.
- In DSWEB schema, UDF to control tasks allowed with the `isDatabaseOwner` privilege
- In DSJOBMGR schema, UDF to control *Can Manage Jobs* privilege
- In DSCSTMALERT schema, UDF to control *Can Manage Custom Alerts* privilege
- Control and read the event monitor files resulting from SQL or WLM activity traces.
- Watchdog stored procedures for event monitors to ensure that event monitors are dropped on the monitored database in case InfoSphere Optim Performance Manager loses connection to the monitored database. The Automatic Task Scheduler (ATS) must be activated in the monitored database.

2.6.4 Special considerations for direct repository access

Special considerations exist for direct repository access.

Customized performance reports

To enable users, groups, or roles to execute the table functions for creating a custom report, the users, groups, or roles must be granted membership in the OPM_REPORT role.

A prerequisite is to grant the membership in the OPM_REPORT role, SECADM authority.

To grant the membership to a user, group or role, use one of the following GRANT ROLE statements, where user, group, and role is the name of the user, group, or role to which you want to grant the EXECUTE rights for all table functions:

```
GRANT ROLE OPM_REPORT TO USER user
GRANT ROLE OPM_REPORT TO GROUP group
GRANT ROLE OPM_REPORT TO ROLE role
```

To determine which user, group, or role is already a member, run the AUTH_LIST_ROLES_FOR_AUTHID function as follows:

```
SELECT GRANTOR, GRANTORTYPE, GRANTEE, GRANTEETYPE, ROLENAME,
    CREATE_TIME, ADMIN
  FROM TABLE (SYSPROC.AUTH_LIST_ROLES_FOR_AUTHID ('OMP_REPORT', 'R') )
AS T
```

InfoSphere Optim Query Workload Tuner access to repository

For the InfoSphere Optim Query Workload Tuner integration with InfoSphere Optim Performance Manager, ensure that your user ID is authorized to access the performance metrics stored in the repository database.

The prerequisites for integrating InfoSphere Optim Query Workload Tuner with InfoSphere Optim Performance Manager are described in 2.2.3, "Integration with InfoSphere Optim Query Workload Tuner" on page 44.

Data Studio access to repository

You can use Data Studio to connect to an instance of InfoSphere Optim Performance Manager. InfoSphere Optim Performance Manager collects performance data about your SQL scripts and stores the data in the repository database. You can display this performance data in the Performance Metrics view and use it to optimize your scripts.

The user ID that connects to the monitored database from Data Studio must hold Viewer privileges in InfoSphere Optim Performance Manager to be able to connect to the repository database, and collect performance metrics.

For additional details about collecting performance metrics from Data Studio, go to the Data Studio information center:

http://pic.dhe.ibm.com/infocenter/dstudio/v3r2/topic/com.ibm.datatools.javatool.ui.doc/topics/cpdqtunsqlperf.html

3

Installing and configuring Optim performance management tools

This chapter describes how to get Optim performance management tools, in particular InfoSphere Optim Performance Manager and InfoSphere Optim Query Workload Tuner, up and running, including installation and configuration. It also describes several proven practices to help make the deployment efficient and smooth.

3.1 Installing and running InfoSphere Optim Performance Manager

InfoSphere Optim Performance Manager provides diversified installation options for users. These options are described in 2.6, “User security” on page 55 and can help you decide which option best suits your environment.

Unlike an installation guide, this section provides an example of how you can set up InfoSphere Optim Performance Manager, but does not cover every detail. It briefly describes the major steps of a fresh installation of InfoSphere Optim Performance Manager V5.2 with the default options used in most of the steps. For more details about installation, see the information center:

http://pic.dhe.ibm.com/infocenter/perfmgmt/v5r2/index.jsp?topic=%2Fcom.ibm.datatools.perfmgmt.installconfig.doc%2Fpm_installconfigure.html

Before installing InfoSphere Optim Performance Manager, read the installation requirement described in the guide. Also see 2.2, “Prerequisites” on page 36.

This section covers installation using the following options:

- GUI installer on Windows
- Console installation on AIX
- Activating the license with the activation kit
- Validating the installation and basic start and stop operations

3.1.1 Installing InfoSphere Optim Performance Manager

The GUI, console, and silent installations are supported. This section provides an example of a GUI-based installation on 64-bit Windows Server 2008 and a console-based installation on AIX. During GUI-based or console-based installations, you can create a response file using the installer. That response file can subsequently be used for installations on other instances, which is called a *silent installation* and is a much easier process. We briefly describe that process at the end of this section.

Tip: Make sure you go through a capacity planning calculation of InfoSphere Optim Performance Manager before the installation to ensure that it meets your requirements for database monitoring. For more information, see 2.4, “Capacity planning” on page 49.

GUI installation on a 64-bit Windows server

This section describes the primary steps for installing InfoSphere Optim Performance Manager on Windows using a GUI installation. Before installing, you must have DB2 available on the machine because a repository database is created in DB2 during the installation. You can either create a DB2 instance during the installation or use an existing DB2 instance. The supported DB2 version for a repository is described in "DB2 data server for Optim Performance Manager" on page 39. The installation demonstrated in this section uses an existing DB2 instance.

The InfoSphere Optim Performance Manager installer contains the following files:

- `IOPM.server.v5.2.install-on-win64.exe`
- `opmserver.zip`

The installation steps in GUI mode are as follows:

1. Use the installation starter panel to choose a language.

   Launch the `opm.server.v5.2.install-on-win64.exe` file with administration privilege. The installer panel opens. Choose a language and click **OK** to proceed.

   > **Tip:** The language selected here is only for the installation panel during the installation. The Web UI has been translated into several languages. If you want choose the Web UI language after the installation, you basically need to choose the browser language.

2. See the installation wizard and welcome panel.

   The installation wizard and welcome panel open. On the left side are the steps of the installation, and your current step. In the left content area, you see a description of the installer, and the location where you can see the installation architecture online to help you understand the deployment architecture. That architecture is also described in 2.1, "Installation roadmaps" on page 32.

3. Review the prerequisite panel.

   This prerequisite panel lists the hardware and software perquisites, which are also described in detail in 2.2, "Prerequisites" on page 36. Click **Next** to go to the License Type selection panel.

4. License Type selection panel:

   You can choose to install as either the Try and Buy (provides a 60-day license) or the full licensed edition. If you already have a license and the activate kit, you can extract the license file, choose licensed edition option, and browse for that license file. This section demonstrates the use of a license file for the licensed edition option.

> **Tip:** You can also choose the Try and Buy option at installation time and includes a 60-day trial license. When that license expires, or whenever you choose, you can use the activation kit to activate a full license, which is described in 3.1.2, “Activating the InfoSphere Optim Performance Manager license” on page 83.

5. Read and accept the license agreement.
6. Select an installation method.

   The installation program can create a response file which can be used for subsequent silent installation. The response file contains your input of each step during the installation process. If you choose to create a response file, the file will be created, when the installation completes, in the directory specified in this step. We choose the option to install the production on this machine without creating response file. However, if you need to deploy a similar configuration on several servers in your environment, it is suggested that you generate a response file, and then perform a silent installations for the other servers.
7. Select an installation directory.

   When choosing an installation directory, be sure there is enough space available for installation. In the examples, we accept the default Windows installation location, `c:\Program Files\IBM\OPM`. If you have a previous version of InfoSphere Optim Performance Manager installed on the server, you can choose to update an existing version for upgrade installation.
8. Select an installation type.

   The options are Typical Installation and Advanced Installation. If you do not want to use the default repository database specifications including name, table space type, and location, choose **Advanced Installation** to change them. The table space type to use is one of the planning tasks described in 2.5.1, “Table space type selection” on page 52. You must select the Advanced Installation to specify the table space type. The default type is DMS. We select **Advanced Installation** in this example.
9. Select DB2 instance (Figure 3-1 on page 73).

> **Tip:** If you want to simplify the installation process and do not need to understand details of the configuration, then you can choose the Typical Installation, which uses default settings for your server.

If you have a DB2 instance available, use the **Select an existing DB2 instance** option. To have the installer create a DB2 instance for you, choose **Create a new DB2 instance** and specify the instance user. In this example, we create a new DB2 instance (db2_opm) to hold the performance repository, and thus had to specify the instance owner credential.

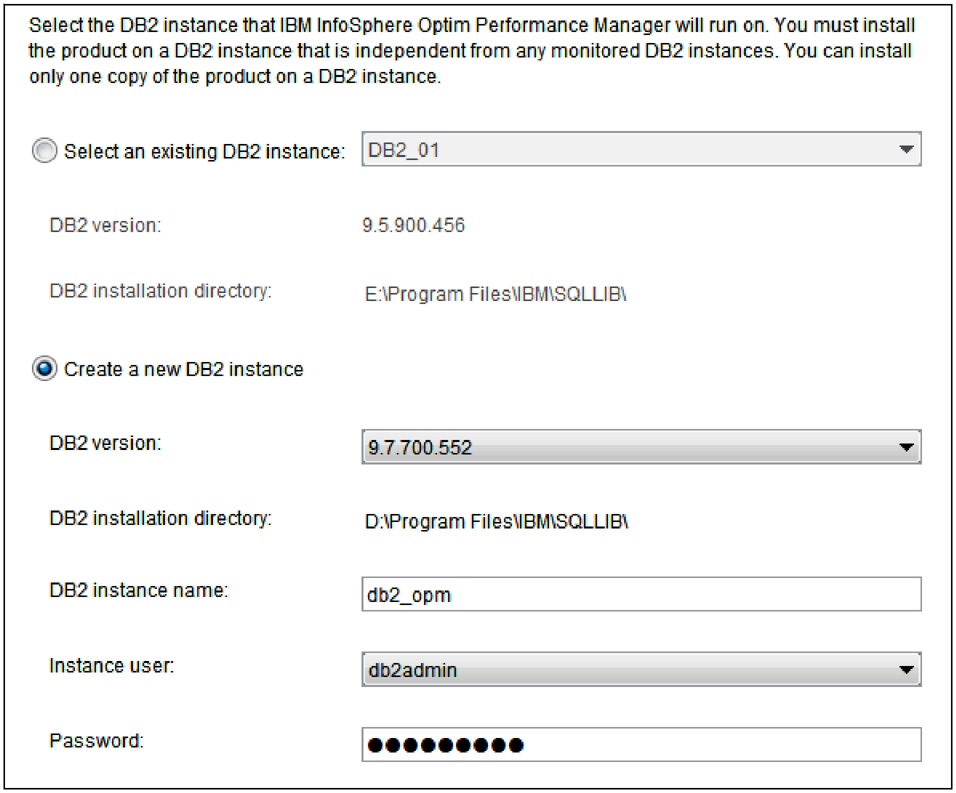

Figure 3-1 DB2 instance selection for InfoSphere Optim Performance Manager repository

10. Select an authentication method to the repository database.

    You can choose the authentication method of the repository database. By default, it is local OS authentication (SERVER), but it also supports Kerberos authentication if that is your enterprise security policy.

11. Create a database.

    In the Database Creation panel, you can specify the repository database name and path of the database, and also the path of table spaces. By default the repository database name is PERFDB, which is also used in this example. The table space here is used to hold the meta tables (which do not take too much space). The database will be created under the instance you selected in the previous panel (in Figure 3-1 on page 73).

12. Specify a working directory for the repository server.

    Here you specify the path for the working files, including log, trace, and configuration files of the repository server for each monitored database. This example uses the default path, under the installation directory of the repository server.

    > **Tip:** Choose the default path for the working directory, because it is easy for you to view the log files, and make some changes to the repository server later, when necessary. In the working directory, you can set the global logs and configuration for the repository server; meanwhile for each monitored database, there is a corresponding folder with odd digits as the folder name to hold the log files for each monitored database. The odd numbers correspond to a monitored database and you can use the **`peconfig -list`** command in the `InfoSphere Optim Performance Manager install directory/Repository Server/bin` folder to get the assigned odd number for each monitored database.

13. Select repository database table space options.

    There are three options for the repository database table spaces. The table space is used to hold all the performance data collected for the monitored database, which consumes most of the space of the repository database. You can select DMS, SMS and Automatic Storage (the default option is DMS), and let the table space be managed by the database manager.

    > **Tip:** If you are not certain about the resources suggested when you did your capacity planning, because you might have more databases to monitor in the future, and hence the current storage might not be enough, then a good approach is to choose the automatic storage option. Then later, if the repository database disk space is nearly full, you can, on-demand, easily add more storage paths for repository database. Otherwise, you might need to add containers manually to the big table space if you choose DMS.

14. Provide connection information.

    Input a DB2 instance user name and a password, which will be used to create and access the repository database. This user must have the SYSADM authority on the DB2 instance. Input a group name nd the appropriate privileges will be granted to this group for two purposes:

    – So the users can log on from the Performance Expert Client to the repository server.

    – If you are using InfoSphere Data Studio or InfoSphere Optim Query Workload Tuner to capture performance data from the performance database.

15. Specify web console port options.

    You can specify the port of web console that can be accessed through a browser to InfoSphere Optim Performance Manager server. You can enable both HTTP and HTTPS; URLs for both HTTP and HTTPS access will be provided. You can specify the port number, and make sure the ports are not blocked by the firewall.

16. Define product startup.

    Next on the Product Startup selection panel you can select whether you want to start IBM InfoSphere Optim Performance Manager when the computer starts. To learn how to start Optim Performance Manger by using the command line, see 3.1.4, "Starting and stopping InfoSphere Optim Performance Manager" on page 90.

17. Review the preinstallation summary.

    This panel shows a summary of your installation selections. When you proceed, the installer will start the installation on your server. The time required for the installation varies based on the available resource of InfoSphere Optim Performance Manager server.

18. Review the installation summary, start InfoSphere Optim Performance Manager and open web console.

    After the installation finishes, the Installation Summary panel indicates that the installation was successful. You can then choose whether to start InfoSphere Optim Performance Manager at that time.

    – If you choose to start InfoSphere Optim Performance Manager at that time and it starts successfully, it also starts the web console server. A web browser session is launched and the InfoSphere Optim Performance Manager web console is displayed when you click **Done**.

    – If you choose *not* to start InfoSphere Optim Performance Manager at that time, be sure to note the two URLs presented in the panel: these are the InfoSphere Optim Performance Manger address and the web console address. You can later start InfoSphere Optim Performance Manager and the Web Console manually as described in "Start and stop the server on Windows" on page 91.

19. Finally you can click **Done** to finish and exit the installation.

> **Note:** The web console does not display if you click **Done** and you did not choose to start InfoSphere Optim Performance Manager. At that point, you must start both InfoSphere Optim Performance Manger and the web console manually. This is documented in 3.1.4, "Starting and stopping InfoSphere Optim Performance Manager" on page 90.

At this point, a complete installation is finished successfully on Windows with GUI with advanced configuration.

Next, you see how to install on an AIX server with console installation. There is a live demo of this installation, which can be a good reference even it is based on previous version. It is at the following address:

http://www.ibm.com/developerworks/offers/lp/demos/summary/im-installoptimpm51-1.html

Console installation on AIX Server

This section describes using a command terminal (SSH, Telnet, and so forth) to install the InfoSphere Optim Performance Manager console on an AIX server. It describes specific key steps, but skips others because you can follow the instructions in the installer.

The key steps are as follows:

1. Log in to the AIX server as root user, and decompress the installation files.

   The AIX installer image is in `tar.gz` format, but you can decompress it with **`unzip`** and **`tar`** commands. You then see the following installation files:

   - `IOPM.server.v5.2.install-on-aix.bin`
   - `IOPM.server.v5.2.install-on-aix.sh`
   - `opmserver.zip`

2. Open the decompressed folder and change the `.sh` and `.bin` file attributes to executable. Run the `./IOPM.server.v5.2.install-on-aix.sh` file to start the installation with root user.

Tip: Before you start the installation, make sure your /tmp folder has enough space (approximately 1.5 GB) to be used by the installer during installation. If you do not have enough space, either add more space for /tmp or create a new tmp folder for the installer as follows:

1. Create a new directory, such as /data/tmp, with sufficient space.
2. Change the /data/tmp directory to have the same access privilege as /tmp directory.
3. Define IATEMPDIR=/data/tmp and EXPORT IATEMPDIR.
4. Start the installation.

5. Follow the instructions on the terminal to choose a language to use during the installation, as shown in Example 3-1.

Example 3-1 Choose console installation language

```
====================================================================
Choose Locale...
----------------
1- Deutsch
  ->2- English
    3- Español
    4- Francais
    5- Italiano
    6- Português  (Brasil)
```

6. In the prerequisite page (Example 3-2), you see the required levels for an AIX system; be sure your AIX system meets the requirements. See 2.2, “Prerequisites” on page 36 for details about the prerequisites.

Example 3-2 Prerequisites

```
Prerequisites------------

Ensure that your AIX operating system is at one of the following
minimum levels before you install the product:

- AIX 5.3
For TL11: APAR IZ74749 or 5300-11-05 or later
For TL12: APAR IZ76228 or 5300-12-02 or later

- AIX 6.1
For TL04: APAR IZ74508 or 6100-04-06 or later
For TL05: APAR IZ76227 or 6100-05-02 or later
```

```
For TL06: APAR IZ74932 or 6100-06-00 or later

- AIX 7.1
....
```

7. In the license type selection (Example 3-3), you may select **Install the Try and Buy option** and then later, use the activation kit to activate the license. See “Activating the InfoSphere Optim Performance Manager license” on page 83 for more details.

Example 3-3 Selecting the license type

```
License Type
------------

If you have a license for the product, you can install the licensed
edition. If you do not have a license for the product, you can
install the Try and Buy option. The Try and Buy option expires in 60
days.
```

> **Tip:** You can change the edition license later by using the InfoSphere Optim Performance Manager License Activation Kit.

```
  → 1- Install the Try and Buy option
    2- Install a licensed edition
:1
```

8. In the License Agreement step, you can review and then enter `1` to accept the license agreement. Then, you can continue forward and choose the **Installation Method**.
9. Choose to install a new product or update an existing product. The installer finds and list them. Example 3-4 shows that `Install a new product` is selected and the directory of the installation is the default value on AIX, which is (`/opt/IBM/OPM`).

Example 3-4 Installation Directory

```
Installation Directory ------------------------------

Accept the default installation directory, or specify a different
installation directory. If you are updating an existing copy of IBM
InfoSphere Optim Performance Manager, select the installation
directory of that existing copy.
```

```
 → 1- Install a new product.

Choose an option by entering a number, or press Enter to accept the
default.
   :
Installation directory: (DEFAULT: /opt/IBM/OPM):
```

10. For installation Type, similar to GUI installation, you are provided with both typical installation and advanced installation. (See “GUI installation on a 64-bit Windows server” on page 71 for the options in advanced installation.)
11. Specify repository database and server information, including the repository database instance, database name, table space options, and the repository server working directory. In the demonstration, we use an existing db2 instance, use the default database name PERFDB, and use DMS for the PERFDB table spaces. The repository server working directory is the instance owner home directory. This is depicted in Example 3-5.

Example 3-5 DB2 Instance

```
DB2 Instance
------------

Select the DB2 instance that IBM InfoSphere Optim Performance Manager will
run
on. You must install the product on a DB2 instance that is independent from
any
monitored DB2 instances. You can install only one copy of the product on a
DB2
instance.
 → 1- Select an existing DB2 instance:
    2- Create a new DB2 instance

Choose an option by entering a number, or press Enter to accept the
default.
   :
Select the existing instance
 → 1- db2inst5
    2- db2inst6
    3- db2inst7
    4- db2inst3
    5- db2inst9
    6- db2inst1

Choose an option by entering a number, or press Enter to accept the
default.
   : 6
====================================================================
```

```
Repository Database
-------------------

IBM InfoSphere Optim Performance Manager stores the data that is collected
for each monitored database in database tables of the repository database.
Specify the information for the repository database and control tables.
Choose an option by entering a number, or press Enter to accept the
default.
   : 1
Specify a name for the database.
Database name: (DEFAULT: PERFDB): PERFDB
===================================================================
Working Directory
-----------------

Specify a directory in which to store all product log and trace
filesSpecify a directory in which to store all product log and trace files.
Specify a working directory for InfoSphere Optim Performance Manager.

Working directory: (DEFAULT: /home/db2inst1/opm/v5):
===================================================================
Table Space Options
-------------------

IBM InfoSphere Optim Performance Manager stores the data that is collected
for each monitored database in separate table spaces. Specify the type of
table spaces to use:
 ->1- Database-managed space (allocated and managed by the database
manager)
    2- System-managed space (allocated and managed by the file system
manager of the operating system)
    3- Automatic storage (allocated and managed by DB2 based on the storage
paths that you specify)

Choose an option by entering a number, or press Enter to accept the
default.
   : 1
```

Tip: The repository database instance owner is the InfoSphere Optim Performance Manager administrator, which is the owner for this installation instance, and is granted to start and stop InfoSphere Optim Performance Manager server, and use other shell commands. After the installation, you cannot use root to start or stop InfoSphere Optim Performance Manager, and will need to use this instance user to do that.

The repository server working directory is under the home directory of repository instance owner. It does not require much space, but depends on the number of databases to be monitored. A good approach is to have at least 200 MB space for the repository server working directory. If you do not have enough space under `/home`, you can specify another directory for which the instance owner will have full access.

You can specify the database name with the organization name convention. The default name is PERFDB. Use Automatic Storage; later, you can add more storage to the table space if is required.

12. Specify the InfoSphere Optim Performance Manager administrator credential, and the web console ports. A useful approach is for the administrator to use the repository database instance owner, or any other user with SYSADM privilege, to specify the ports number used by InfoSphere Optim Performance Manager.

    You can add more console users through the web UI. See 2.6, “User security” on page 55 for more information.

    Make sure the HTTP and HTTPS port numbers you specify are not blocked by the firewall. Consult with the system or network administer for the correct port numbers to avoid conflict and blocking. This is depicted in Example 3-6.

Example 3-6 Connection information

```
Connection Information
----------------------
Repository database and web console

Specify an existing user account that will be given administrator
access to the web console and the repository database to work with
the collected data. This user account must have SYSADM authority on
the DB2 instance for the repository database. DBADM rights will be
granted to this user account during installation.
Tip: You can add more users later from the web console.

User name: (DEFAULT: db2inst1):
```

```
Enter the password.
Password:*********

Specify an existing group whose users will be given permission to
log on from DB2 Performance Expert Client or Data Studio to the
InfoSphere Optim Performance Manager repository server. The group
account must exist on the InfoSphere Optim Performance Manager
server. If you do not use DB2 Performance Expert Client or Data
Studio, then this value is not used.
Group name: (DEFAULT: db2iadm1):

===================================================================
Web Console Ports
------------------------------
Specify the port numbers for the IBM InfoSphere Optim Performance
Manager web console. These ports are used to form the URLs where the
IBM InfoSphere Optim Performance Manager web console can be
accessed. Ensure that no firewall is blocking these ports.

Enable the HTTP port: (Y/N): Enable the HTTP port: (Y/N): Enable the
HTTP port: (Y/N):
Enable the HTTP port: (Y/N): y
HTTP port: (DEFAULT: 55000):
Enable the HTTPS port: (Y/N): y
HTTPS port: (DEFAULT: 55001):

Specify the control port for the web console. This port is used only
locally and does not require firewall configuration.
Control port: (DEFAULT: 55002):
```

13. Choose whether to start InfoSphere Optim Performance Manager automatically, review the preinstallation summary, then press Enter to install.
14. When the installation succeeds, you can start InfoSphere Optim Performance Manager, and keep a copy of the URLs for future access using a browser.

Tip: The access URL template is as follows:

```
[HTTP/HTTPS]://HOSTNAME:PORT/optimdatatools/console
```

HTTP and HTTPS are the protocol, and each protocol has a corresponding PORT number. If you forget to copy the URLs at the end of the installation, you can follow the template to get the correct URL. If you have multiple network interfaces, the generated URLs by the installation will choose the first hostname/external IP address from the system. You can change the URL with the proper IP or host name when necessary.

Summary

This section demonstrated how to install InfoSphere Optim Performance Manager with GUI and console mode on Windows and AIX separately. Next description is how to activate InfoSphere Optim Performance Manager with an activation kit, update it to latest version or fix packs, and how to validate the installation with start and stop.

3.1.2 Activating the InfoSphere Optim Performance Manager license

The next stage is to activate the types of licenses, which are briefly described in this section.

Extended Edition (XE) license

Although this is a licensed product, IBM provides a 60-day *Try and Buy* option at no charge so you can experience the product features and functions. The activation toolkit is for applying the license to the Try and Buy option and enable you to continue using the product after the 60 days.

The InfoSphere Optim Performance Manager activation toolkit has three types:

- Enterprise Edition (EE)

  If you purchase Enterprise Edition, you can use the EE activation kit to activate your server.

- Extended Edition (XE)

  If you purchase Extended Edition, you can use the XE activation kit to activate your server.

- Extended Insight (EI)

  If you have EE, but want to use the features in XE (such as end-to-end monitoring, integration with Tivoli and WLM performance objectives), you can purchase EI and use the EI activation to activate your server.

As for the package and edition of InfoSphere Optim Performance Manager, see section 1.3.2, “Packaging” on page 17 for more information.

The InfoSphere Optim Performance Manager EE or XE activation toolkit contains the license. By having the license, you can apply it during installation or activate the license for the Try and Buy option.

This section describes how to activate the license using the EE activation toolkit on Windows 64-bit server.

The activation toolkit for Windows contains the following files:

- `IOPM.server.v5.2.0.0.activate-on-win64.exe`
- `Enterprise.opm_lic`

Use the following steps to apply the license:

1. Stop InfoSphere Optim Performance Manager server (see 3.1.4, “Starting and stopping InfoSphere Optim Performance Manager” on page 90)
2. Run `IOPM.server.v5.2.0.0.activate-on-win64.exe` file.
3. Choose the installation language.
4. Review and accept the license agreement.
5. Specify the installation directory.

   You might have various copies of InfoSphere Optim Performance Manager installed on the same machine. Choose the one to which you want to apply the license.
6. Review the preinstallation summary.
7. Click **Done** when the installation is finished.
8. Start the InfoSphere Optim Performance Manager server.

As for Enterprise Edition activation kit, you do not need an additional configuration, just apply the license file to activate the InfoSphere Optim Performance Manager server. For the XE (Extended Edition) and EI (Extended Insight) activation kit, you need to do additional configuration for Extended Insight port number and IP addresses.

Activating Extended Edition (XE) and Extended Insight (EI) licenses

The EI option is a separate licensed option of InfoSphere Optim Performance Manager that is installed with it but not activated. You can activate these capabilities with the XE or EI activation kit. If you have an XE license, then you can use the XE activation kit directly soon after you install InfoSphere Optim Performance Manager, or else you can purchase EI with EE to get the same capability and then use the EI activation kit to activate the InfoSphere Optim Performance Manager server.

Extended Insight (EI) contains a server component and a client component. The server component is contained in the installation package and is installed in the installation directory.

For example, on AIX, if you install InfoSphere Optim Performance Manager in the `/opt/IBM/OPM` location, you can see the `/opt/IBM/OPM/pureQuery` subdirectory, which contains the major part of EI server. The property file of the EI Server (`pdq.properties`) is in `<working directory>/<db2 instance>` location.

- The `pdq.properties` file is used in configuring EI monitoring.
- The `<working directory>` name is the working directory of the repository server
- The `<db2 instance>` name is the name of the DB2 instance on which the repository server runs.

Tips:

- On AIX, the default working directory of the InfoSphere Optim Performance Manager repository server is as follows, where `<user home dir>` is the home directory of the DB2 instance owner (the repository user):

  `<user home dir>/opm/v5`

- On Windows, the default working directory is as follows:

  `<OPM install dir>\RepositoryServer`

To use EI, you must activate the server by applying a license to the EI server using the EI activation toolkit.

The EI or XE activation toolkit for AIX contains the following files:

- `IOPM.server.v5.2.0.0.activate-on-aix.bin`
- `ExtendedEdition.opm_lic`

Use the following steps to activate the EI license on AIX:

1. As the instance owner, stop InfoSphere Optim Performance Manager.
2. As root, run `IOPM.server.v5.2.0.0.activate-on-aix.bin` file.
3. Choose the installation language.
4. Read and accept license agreement.
5. Specify the InfoSphere Optim Performance Manager installation for which the EI is to be activated (see Figure 3-16 on page 110).

   The installer detects and lists the valid copies installed on the system. Choose the copy for which you want to activate the EI feature. If the copy is not in the list, click **Browse** to specify the installation directory.
6. Set up EI communication.

   The EI client will communicate with the server using the specified IP and port number of the server. If there is a Network Address Translation (NAT) between the server and the EI client, you can input the external IP for the server that is accessible by the EI client. This step is optional only when your server uses NAT.

   Then, you can specify the port number that the server will open to accept request from EI client (Figure 3-17 on page 111). The port number must be unused, and not blocked by the firewall. The default port numbers are 65000 and 65001. The value of `pdq.cmx.controllerURL` file in `<working directory>/<db2 instance>/pdq.properties` will be updated with the host name and port number that you specify in this step as follows:

   ```
   pdq.cmx.controllerURL=<host name>:<port number>
   ```

   In this line, `<host name>` is the host name identified by the installer and `<port number>` is the port number you specified.

   You can check and update this property value later if necessary.
7. Review the preinstallation summary, and start the activation installation.
8. Finish the installation and start InfoSphere Optim Performance Manager server.

The installation of the client component is described in 3.4, “Installing and configuring IBM InfoSphere Optim Query Workload Tuner” on page 127.

Tip: How can you check the edition and license type for the InfoSphere Optim Performance Manager server?

You can use the following command to verify the current edition and license type for your server. In the command `<OPM_DIR>` is the InfoSphere Optim Performance Manager installation directory:

```
<OPM_DIR>/bin/opmlicm
```

The following example on a Windows server shows the version, edition (XE), and license type (permanent), which means the server has been successfully activated with XE:

```
C:\Program Files\IBM\OPM\bin>opmlicm.bat
Product name:                  "InfoSphere Optim Performance Manager"
Version Information:           "5.2.0.0.6141"
Edition:                       "Extended Edition for LUW"
License type:                  "permanent"
```

The following example on AIX shows a Try and Buy option with a trial license:

```
bash-3.2$ ./opmlicm.sh
Product name:                  "InfoSphere Optim Performance Manager"
Version Information:           "5.2.0.0.6141"
Edition:                       "Enterprise Edition for LUW"
License type:                  "trial"
Days until expiration:         "57"
```

3.1.3 Updating InfoSphere Optim Performance Manager

This section describes how to update InfoSphere Optim Performance Manager after new fix packs and releases are available. This usually happens when you have the server installed and activated for database monitoring, and want to apply the new version or fix packs to use the new features or enhancement.

Tip: The update requires additional space on the server side to install the new product feature on top of existing one, and also to migrate the repository database. Be sure you have the same size of spare space of the server installation directory and repository database.

The repository database takes some time to do the schema and data migration during the update installation. The time depends on the data volume of the repository database and the server capability. So plan a better time window for the update and be patient for the installer to finish the migration.

The updating includes the following steps:

1. Plan the update.

   Plan for the storage and time for windows updating.

2. Stop the InfoSphere Optim Performance Manager server.

   See 3.1.4, "Starting and stopping InfoSphere Optim Performance Manager" on page 90 for the instructions.

3. Optional: Back up the server repository database.

   This is not mandatory, but strongly suggested in case problems occur during the update. Then, you can restore to the existing version.

4. Install the new fix pack or version of the server.

5. Optional: Install the new fix pack or version of EI client on data client side.

   This step is optional, and is necessary only when you have new version of EI client. See 3.4, "Installing and configuring IBM InfoSphere Optim Query Workload Tuner" on page 127 regarding how to install the EI client.

6. Start the InfoSphere Optim Performance Manager server.

   Start the server and be ready to validate the update.

7. Optional: Restart data client.

   Do this step i only when you do the EI client updating on data client side.

8. Validate the update installation.

   Validate the new server to check the version, build number, fixed issues, and new features in the announcement letter.

The next example shows how to do the update installation from a previous version to version 5.2 on Windows.

> **Note:** The assumption is that you have already planned for a good time window and storage on the server for updating, and there has been a backup of the repository database.
>
> Most of the steps are similar to a new installation, so the following description highlights some of the differences.

Complete the following steps:

1. Launch the `IOPM.server.v5.2.0.0.install-on-win64.exe` server installer.
2. Follow the GUI installation guide step by step, and then go to **Installation Directory**.

The installer lists the copy of the installed previous version. Choose **Update an existing product**, and select the correct copy that you want to update. In this example, only one existing copy is installed.

3. The installer reads and reuses all of the configuration of the previous installation, and displays it in the preinstallation panel.

   Before the installation, a message informs you that the repository database will be migrated, and can take a long time depending on the data volume of the repository database.

   > **Attention:** *Do not stop or cancel the installation* during that time, or else the repository database might be corrupted. So when the update installation starts, it is suggested that you do not stop or cancel the installation, but wait till it completes.

4. After the update installation is complete, start the server to validate the updating installation.

5. Validate the installation by checking for the new version as follows:

   Use the **pelevel** command on the server side show the current version of the server and build that is installed. See Example 3-7. The command is in the following location, where <OPM_DIR> is the server installation directory (such as for AIX, the default is /opt/IBM/OPM):

   <OPM_DIR>/RepositoryServer/bin

Example 3-7 pelevel command output

```
bash-3.2$ pwd
/opt/IBM/OPM/RepositoryServer/bin
bash-3.2$ ./pelevel
=========================================================================
IBM InfoSphere Optim Performance Manager
=========================================================================
InfoSphere Optim Performance Manager Server for DB2 for Linux, UNIX and
Windows
Version 5.2.0.0.6141, code levels:
opmc-5.2.0.0.6123,common-5.2.0.0.6141,pdq-3.200.71
```

On the InfoSphere Optim Performance Manager web UI, at the right top part, click **Help** → **About** to see the version and build number.

Tip: Be sure to clear your browser cache that accesses the new sever UI to load the new version of the client applications to your browser. Then, access the server web UI to access the new features in latest server fix pack and version.

3.1.4 Starting and stopping InfoSphere Optim Performance Manager

This section describes the procedures to start and stop the InfoSphere Optim Performance Manager on AIX and Windows including the repository server and the console.

Important: Before starting the server, the DB2 instance on which it will run must be started.

Start or stop InfoSphere Optim Performance Manager server on AIX

To start (or stop) the server (including the web console and the repository server) on AIX, complete the following steps:

1. Log on to AIX server with server instance owner (the administrator) you specified during installation (for example, db2inst1 in this scenario).
2. Go to the server installation directory (for example, `/opt/IBM/OPM`) in the previous installation demonstration.
3. Use `./OPMstart.sh` to start the server, and `./OPMstop.sh` to stop it.

You can also start and stop the console server and the repository server separately for advanced usage.

- Start or stop the *console* server

  Go to the installation directory with the instance owner ID (for example, `db2inst1`), change to the `bin` folder (for example, `/opt/IBM/OPM/bin`), and then issue one of the following commands:

  – Start the console server:

  ```
  ./OPMstart.sh WebClient
  ```

  – Stop the console server:

  ```
  ./OPMstop.sh WebClient
  ```

- Start or stop the *repository* server

  Go to the installation directory with the instance owner ID (for example, `db2inst1`), change to `RepositoryServer/bin` folder (for example,

/opt/IBM/OPM/RepositoryServer/bin), and then issue one of the following commands:

- Start the repository server:

  ```
  ./OPMstart.sh RepositoryServer
  ```

- Stop the repository server:

  ```
  ./OPMstop.sh RepositoryServer
  ```

Tip: Check the status of the server on AIX. You can use the **OPMstatus.sh** command to see whether the repository server and console server are running. The following sample shows that the console server is ACTIVE, and the repository server is running on db2inst1 instance:

```
bash-3.2$ ./OPMstatus.sh
SERVER STATUS: ACTIVE
======================================================================
IBM InfoSphere Optim Performance Manager
======================================================================
Found: Optim Performance Manager
running on the DB2 instance: db2inst1
USER          PID %CPU %MEM   SZ  RSS    TTY STAT    STIME  TIME COMMAND
db2inst1  8716348  0.0  1.0 103460 103532      - A    20:50:53  1:26

.....<some process info omitted here>
```

Start and stop the server on Windows

There are three ways to start and stop the server on Windows:

- Use Windows menu **Start** → **All Program**, and find the server entry. Then choose the start or stop menu to start or stop the server (including both console and the repository server).
- Use the services management

  During the installation of the server on Windows, the console and the repository servers will be registered as a service, and you can start or stop the server in the Windows service management console (Figure 3-2 on page 92).

  Important: In this way, you need to start (or stop) the console and the repository server separately with two services. You can set the service to automatic if would like to get it auto started when the operation system starts. When the registered user for the service has updated the password, you need to upgrade it here in the service properties.

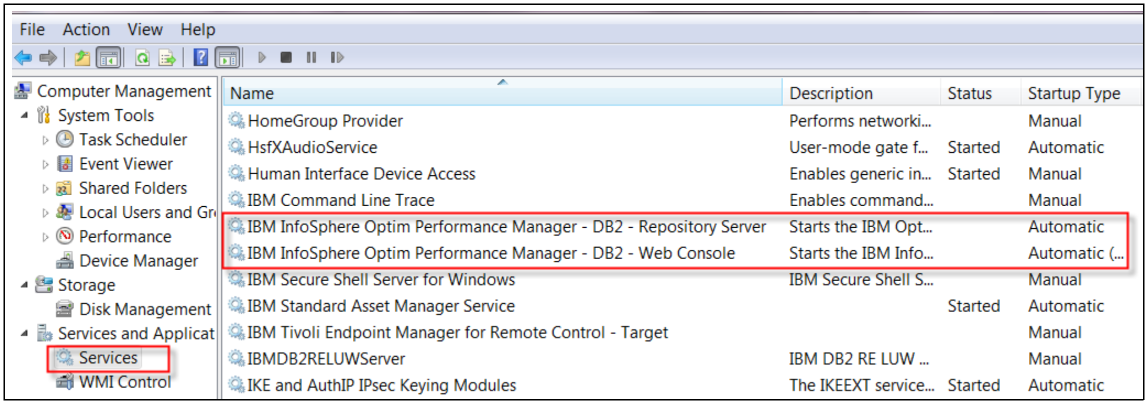

Figure 3-2 Manage Windows services to start (stop) the server

- Use the command line.

  You can also use the command window to start (and stop) the console and the repository servers separately.

  – Start the console server:

    Go to the server installation directory, change to `bin` folder and enter the following command:

    ```
    C:\Program Files\IBM\OPM\bin>start.bat
    ```

    A new command window opens to start the console server.

  – Start the repository server:

    Go to the `RepositoryServer\bin` under installation directory to start the repository server:

    ```
    C:\Program Files\IBM\OPM\RepositoryServer\bin>pestart.bat
    ```

  – Stop the console server

    ```
    C:\Program Files\IBM\OPM\bin>stop.bat
    ```

  – Stop the repository server:

    ```
    C:\Program Files\IBM\OPM\RepositoryServer\bin>pestop.bat
    ```

After installation and activation, configure InfoSphere Optim Performance Manager, which consists of the following steps:

1. Optional: Configure user access.

   After installation, the DB2 user that you specified during installation can log on to the server web console. If other users must have access, then you can give them the privileges. This is an optional step that can also be done after the next two steps.

2. Add or import database connections.

   This step defines the databases that you want to monitor with InfoSphere Optim Performance Manager.

3. Configure the database connections for monitoring.

   In this step, you configure monitoring using monitoring profiles, define monitoring authorizations, and configure partition sets. Partition sets can be configured only for a partitioned database, and member sets can be configured only for DB2 pureScale database.

   > **Note:** InfoSphere Optim Performance Manager still supports the configuration by **`peconfig`** in addition to, or alternatively to, the configuration by web console.

You can log in to the server console using the URL shown at the last step of installation, as in the following example:

```
http://9.30.203.156:55000/optimdatatools/console
```

Enter authentication details to access InfoSphere Optim Performance Manager. Immediately after the installation, only the DB2 user (the server administrator) that was provided during the installation can be used if you use the repository database authentication. For more details, see 2.6, “User security” on page 55. That section introduces the security and authentication concept of InfoSphere Optim Performance Manager.

3.2 Configuring InfoSphere Optim Performance Manager

The next step is to configure InfoSphere Optim Performance Manager to make it ready for daily database performance management. Unlike the installation section, in which most of the actions are done on the server side, for this section, most of the actions are conducted through the server web UI using a browser.

Important: Use the qualified web browser to access the server web UI. See "Web browsers" on page 42 for the prerequisites.

3.2.1 Configuring console security for user access

One aspect of server configuration is to manage which users can authenticate to the server web console and what privileges they have when they are logged in.

The first time you log on to the server web console, you must use the server administrator (the user specified during the installation time).

The privileges that are assigned to a user control the actions that a user can perform in the web console. Assigning privileges is done by opening the Console Security panel under Task Manager on the web console. You must have administrator privileges to use that page.

After a fresh installation of InfoSphere Optim Performance Manager, the authentication method is set to *Repository database authentication* and you can manage user access from the same page. The user, group, or role to which you grant these privileges must already be defined in the repository database and have CONNECT privileges on the repository database.

To each user that requires access to the InfoSphere Optim Performance Manager web console, you can grant one of the following user roles:

- Viewer: The Viewer role is the default global privilege for every server web console user. A user who is assigned the Viewer role cannot change any global settings. Viewers cannot see the historical monitoring information of any monitored databases that are disconnected. Viewers cannot add connections, manage privileges, and so forth.
- Operator: The Operator role in InfoSphere Optim Performance Manager is similar with Viewer role, which has the additional privilege to *Add Database Connections.*
- Administrator: The Administrator role is a global privilege that allows the user to perform any task in the InfoSphere Optim Performance Manager web console. Administrators can also view historical monitoring information of all disconnected databases.

For more details, see 2.6, "User security" on page 55, which introduces the security and authentication concept of InfoSphere Optim Performance Manager. Other user security configuration is described in the 3.2.5, "Other InfoSphere Optim Performance Manager configurations" on page 126.

You can configure the user access in the Console Security page to grant privileges to users, roles, and groups of the web console by clicking **Grant**. Be sure the users, roles, and groups added here can be authenticated by the server repository database. Figure 3-3 and Figure 3-4 show grant privilege to a console user, and list of the console user privileges.

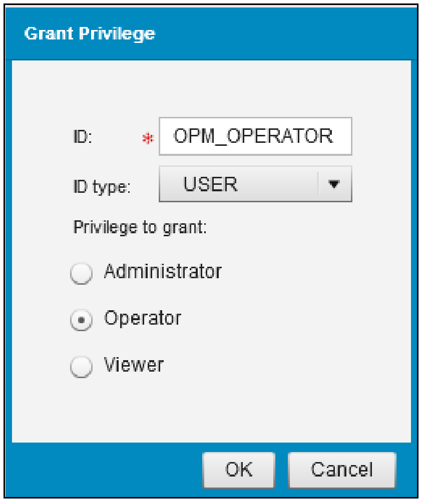

Figure 3-3 Grant privilege to console user

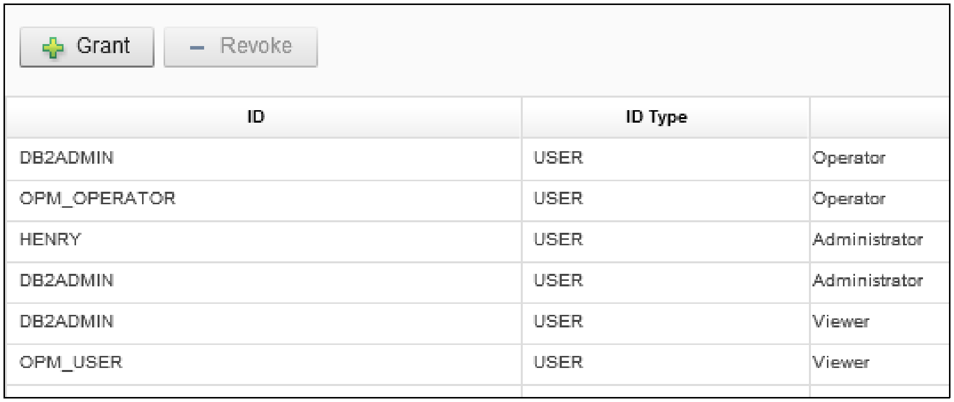

ID	ID Type	
DB2ADMIN	USER	Operator
OPM_OPERATOR	USER	Operator
HENRY	USER	Administrator
DB2ADMIN	USER	Administrator
DB2ADMIN	USER	Viewer
OPM_USER	USER	Viewer

Figure 3-4 List of user privileges of the console user

3.2.2 Adding and configuring database monitoring

After you configure the console users privilege, you can use the administrator role to log in to the server to add and configure databases for monitoring.

The following steps show how to add and configure databases for monitoring with the InfoSphere Optim Performance Manager web UI.

> **Important:** Only the console user with Administrator or Operator privilege can add a database, and only the user *isDatabaseOwner* can configure monitoring. The user adding the database is automatically granted the isDatabaseOwner privilege.

1. You can go to the Database dashboard either from the task launcher dashboard or from the task navigation at the bottom (Figure 3-5).

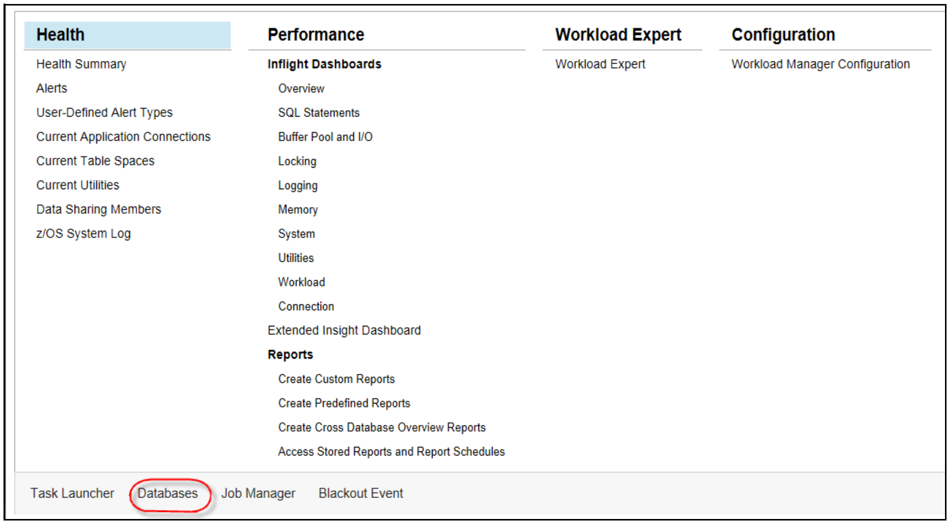

Figure 3-5 Open database dashboard from task navigation

   In the database dashboard, you see a list of database connections that have been added, with and without monitoring configured.

2. Add one connection by clicking the plus sign (**+**), as shown in Figure 3-6 on page 97.

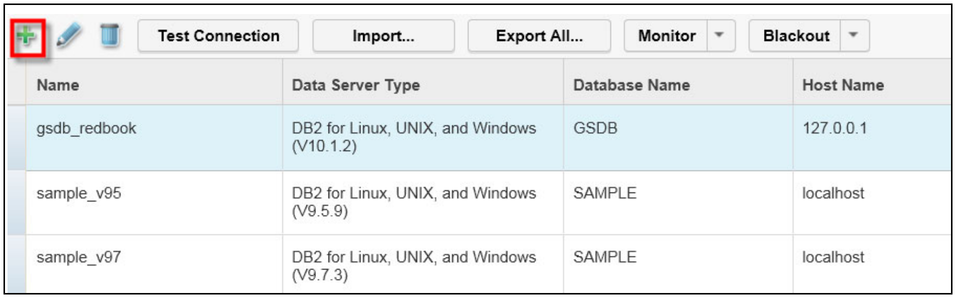

Figure 3-6 Connection Dashboard

3. In the pop-up dialog, input the information for the database you want to monitor; the fields with the asterisk (*) are required information. The connection name is the unique identifier you used within the server, but you can specify a more meaningful name for future use. For the performance data collection, the user must have the privilege for SYSADM and DBADM for the database as described in 2.6, “User security” on page 55.

   Click **Test Connection**. After the connection information is tested successfully, click **OK** to add the connection and continue for monitoring configuration, as shown in Figure 3-7 on page 98.

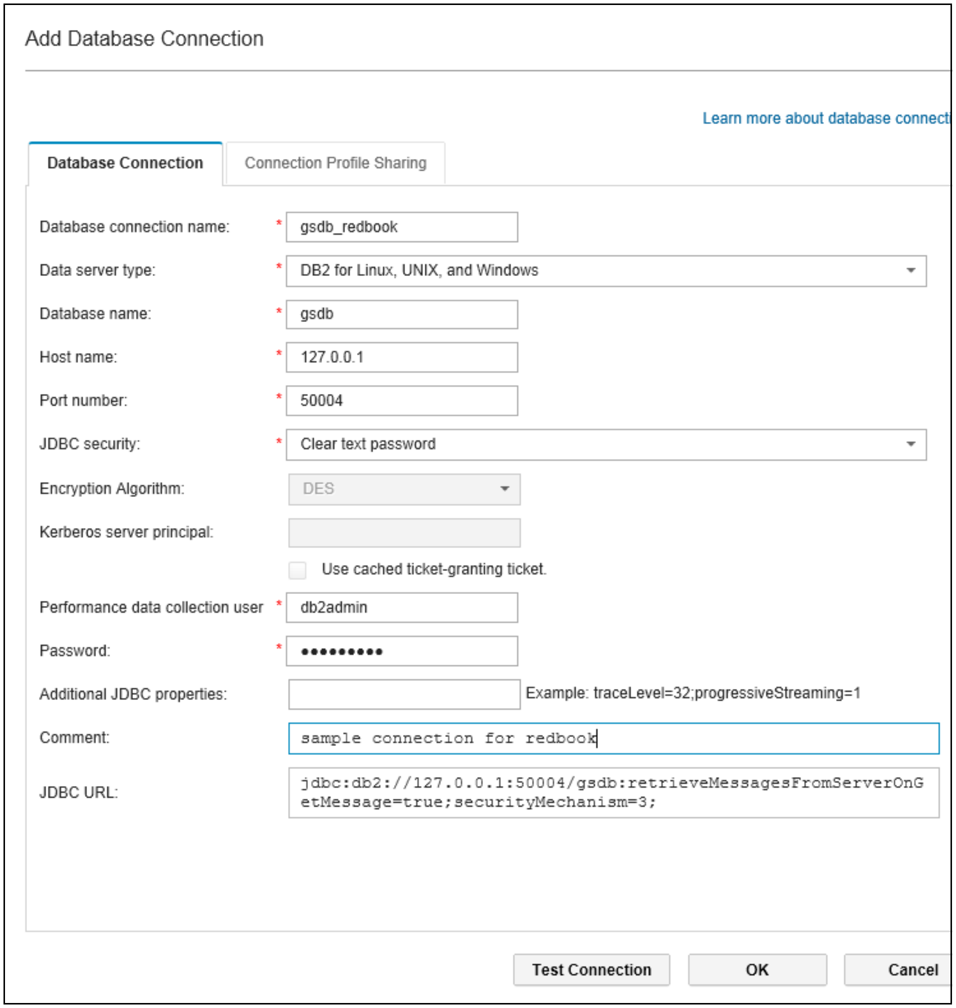

Figure 3-7 Add database connection panel

4. You can choose one of the following options to configure:
   - Only real-time monitoring only
   - Both real-time and automatic data collection monitoring

   The real-time monitoring configuration does not make any changes to the monitored database, only collects performance data as requested by the web UI in real-time mode, and does not store the data as performance data. With automatic data collection enabled, InfoSphere Optim Performance Manager will make corresponding changes to the monitored databases to enable the automatic data collection and will store the collected performance data in the repository database for future analysis and reporting. See 4.3, “Real-time monitoring and history monitoring” on page 176 for more information about real-time and historical monitoring.

> **Note:** If you have database created, but have not yet configured for monitoring, you can go to the Database dashboard, choose the database connection, choose the **Monitor** option on top of the table, and then select **Configure Monitoring**.

5. In the Configure Monitoring panel (step 1 of 4, as shown in Figure 3-8 on page 100), enter the following information:

   – Physical database name

     This is automatically populated with the selected database name.

   – Storage path for collected monitor data

     By default, this is selected to use the default table space path, but you can specify another path. A new path can be entered only if you selected SMS or DMS table space type during installation. For more details about storage options, see 2.5, “Storage options” on page 52, which introduces the storage concept of InfoSphere Optim Performance Manager.

   – InfoSphere Optim Performance Manager collection user

     This is automatically populated with the same user that you specified when you added a database connection. You can edit the user ID only when you edit the connection profile information.

   – Password

     The password of the collection user is automatically populated when you create the connection.

   – Time zone

     This parameter refers to the time zones of the server on which the DB2 instance runs. Make sure that this is correct. InfoSphere Optim Performance Manager uses the time zone information to display collected performance data in the time zone of the monitored database. If the time zone is incorrect, then wrong timestamps are displayed for collected performance data.

   – Starting point

     Now select a starting point for your monitoring configuration profile that can be either a new, user-predefined template, or a clone from another already configured database. Predefined templates are available for various monitored database systems (such as OLTP, pureScale, BI, or SAP systems). For more details about the various predefined templates with description, see “Predefined templates” on page 121. In this example **Use predefined template** is selected, and from the drop-down menu, **Development system** is selected. In a real production system, this profile is not suggested.

– Activate monitoring

  Select the **Enable automatic data collection** check box if you want to collect data for historical monitoring and reporting. Later you can disable (or enable) automatic data collection from the Database dashboard.

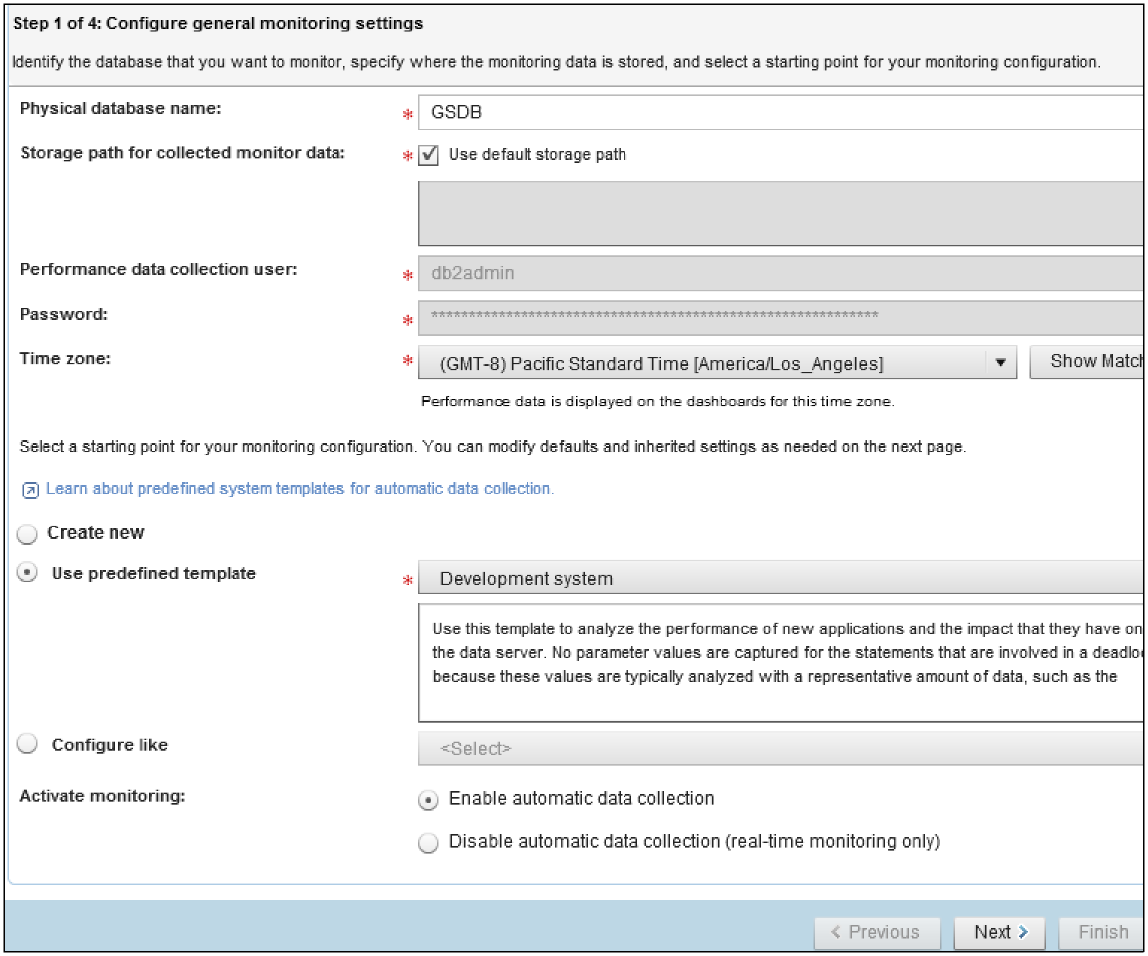

Figure 3-8 Monitoring configuration panel basis

6. Click **Next**.

Because a predefined template is selected in the previous step, InfoSphere Optim Performance Manager enables associated monitoring profiles. Depending on what you select, the system collects various types of data, such as inflight performance, reporting, workload manager, or EI data, as shown in Figure 3-9.

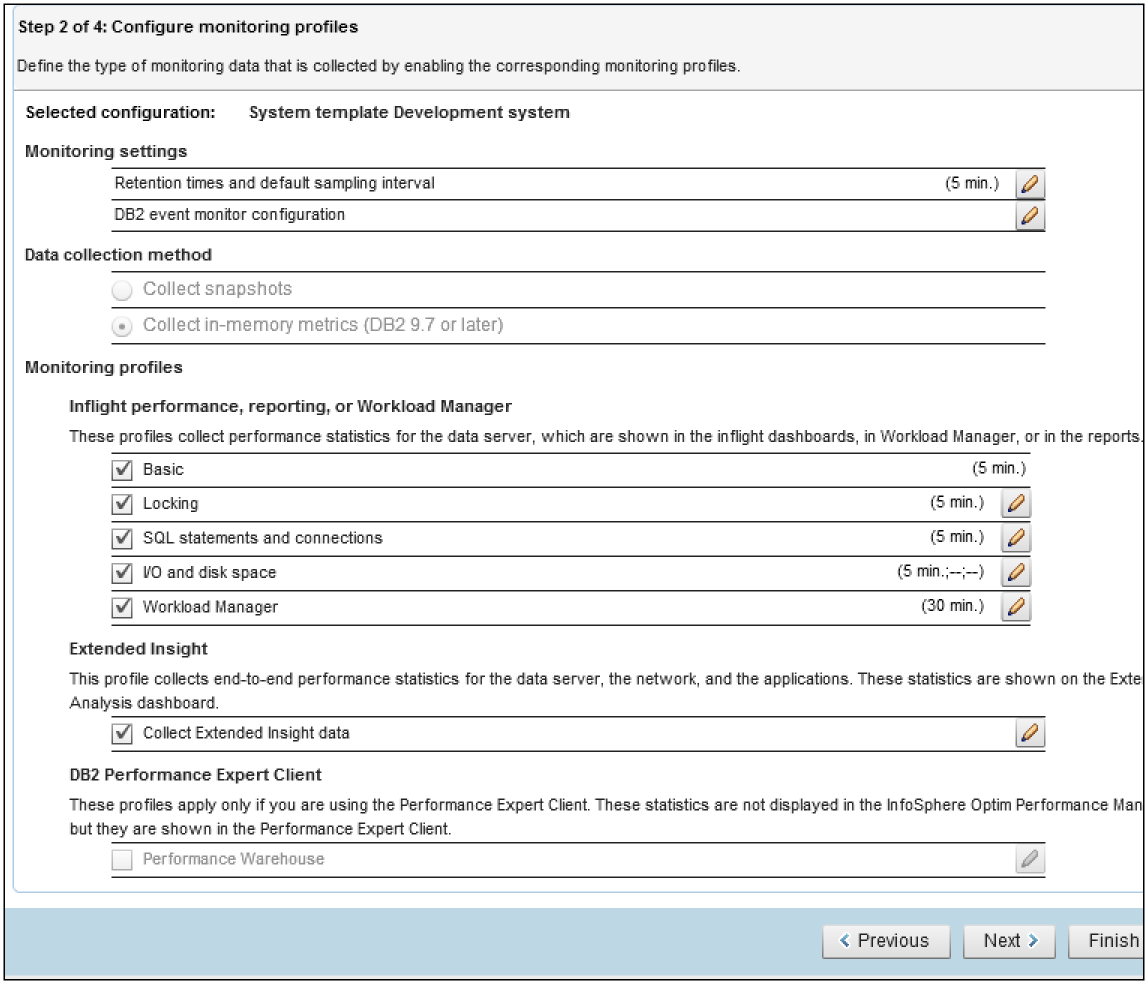

Figure 3-9 Configure monitoring profile

7. Click the pencil icon that is next to each profile to see the default settings for each preconfigured monitoring profile, and edit it if required. Additionally, InfoSphere Optim Performance Manager sets alert thresholds depending on the selected predefined template. For example, for a monitored OLTP system, a buffer pool hit ratio over 90% is more important than for a BI system. You cannot edit the alert thresholds using the pencil icons, but you can edit them

after configuration by using the Performance Alert Configuration dashboard available from Task Manager (see 4.2.3, “Alert configuration” on page 163).

Tip: For each monitoring profile, InfoSphere Optim Performance Manager provides the filters to control the data collected. You can exclude the data that is of no interest or that is collected and is specific only to one application for problem determination (PD), or to reduce the overhead to the monitored database.

Predefined templates are a good way to start with InfoSphere Optim Performance Manager. If you are not familiar with them, follow the monitoring best practices (3.2.3, “Monitoring templates and proven practices” on page 121) and implement deployment stages. Start small using a starter template and grow safely later depending on your monitoring needs.

a. Monitor settings:

- Set the sampling intervals, aggregation level, and report retention time. See Figure 3-10 on page 103.

  Specify how often to collect performance data of the monitored database, how often to aggregate the collected data in the repository database for different aggregation levels, and how long to keep performance data. The selected predefined template will be the basis for these values. You can edit the values at a later time on demand.

  The *sampling rate in minutes* specifies the overall interval in which InfoSphere Optim Performance Manager collects performance data for the target database. In various monitoring profiles, you can increase the basic interval to collect data less often. You can also edit the specific collection interval for each group of metrics such as I/O, statements, and workload manager in the configuration for each set of metrics.

  The *Inflight and Extended Insight data and report retention settings* specify how long to store the collected performance data, and also the report generated through scheduled reports saved on the server side. InfoSphere Optim Performance Manager stores the collected performance data in four aggregation levels: Sampling interval, 15 minutes, 1 hour, and 1 day. Each aggregation level has a storage period, and you can specify the retention time. InfoSphere Optim Performance Manager deletes the data from an aggregation level automatically after the storage period is over. The storage period for aggregation level 1 is the shortest, and for aggregation level 4 the longest. For every sampling interval, InfoSphere Optim Performance Manager receives performance data from the monitored databases and applications, and stores it in aggregation level 1. Aggregation means that average and maximum values of the collected metrics within the aggregation interval are calculated. After 15 minutes of collection, InfoSphere Optim Performance Manager reads the data from aggregation level 1 and aggregates it into aggregation level 2. After one hour of collection, InfoSphere Optim Performance Manager reads the data of aggregation level 2 and aggregates it into aggregation level 3. The aggregation to level 4 happens first after one day of collection. Figure 3-11 on page 104 shows the aggregation concept.

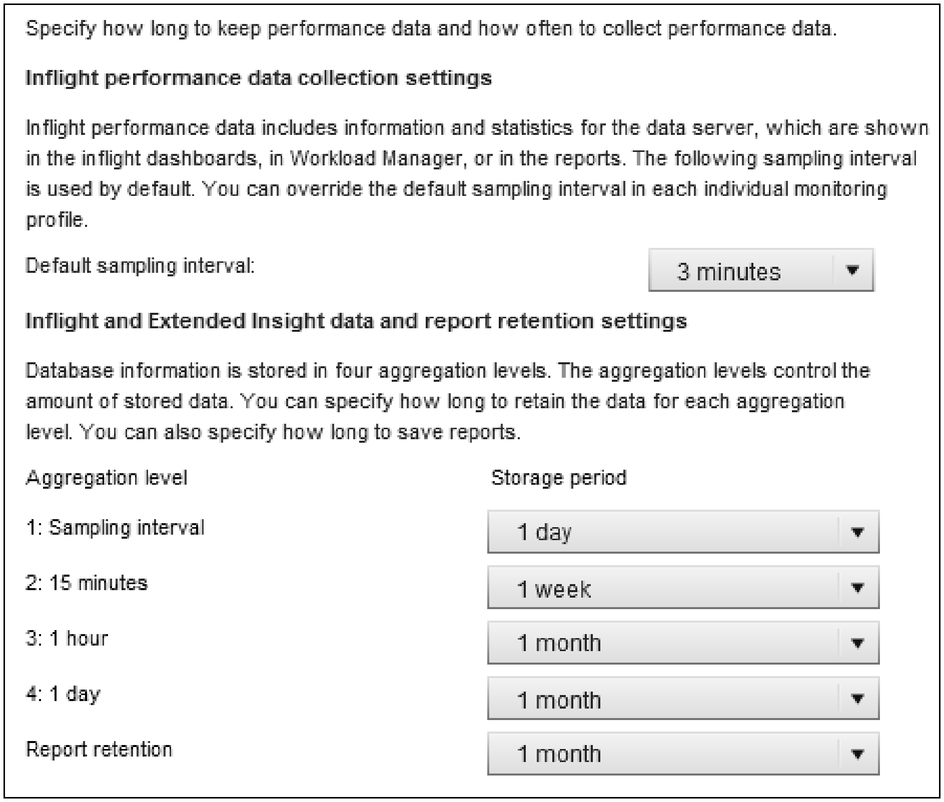

Figure 3-10 Time interval, aggregation and retention selection

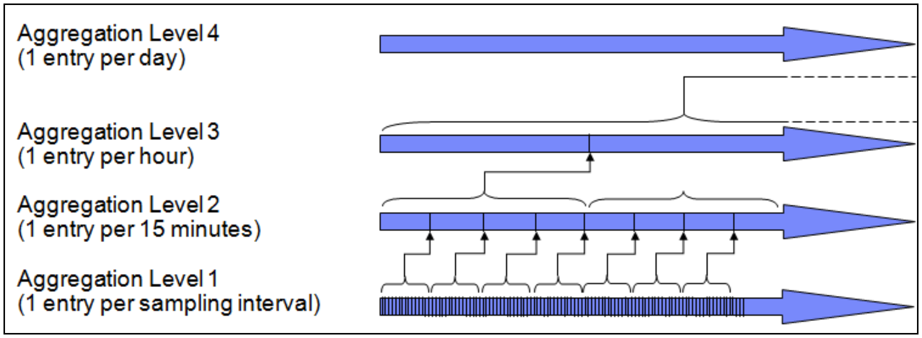

Figure 3-11 Performance Data Aggregation

- DB2 event monitor configuration

  Specify the table space on the monitored database that InfoSphere Optim Performance Manager has to use to create event monitor tables for event monitors. If you do not specify one, DB2 chooses the default table space. The table space that you specify here is used for all event monitors that InfoSphere Optim Performance Manager creates unless you specify another table space for dedicated event monitors in the following monitoring profiles: Locking, Workload Manager, and Extended Insight. *The Maximum tablespace fill size percent* is the percent of table space utilization at which the event monitor will stop after it is reached to avoid using all the table space. The default value is 90%, which means if the table space is 90% used, then all event monitors that use this table space are stopped. If the table space specified here is enabled with *autoresize*, you can specify a higher value (up to 100%) because the table space will increase its size when it becomes full. See Figure 3-12 on page 105.

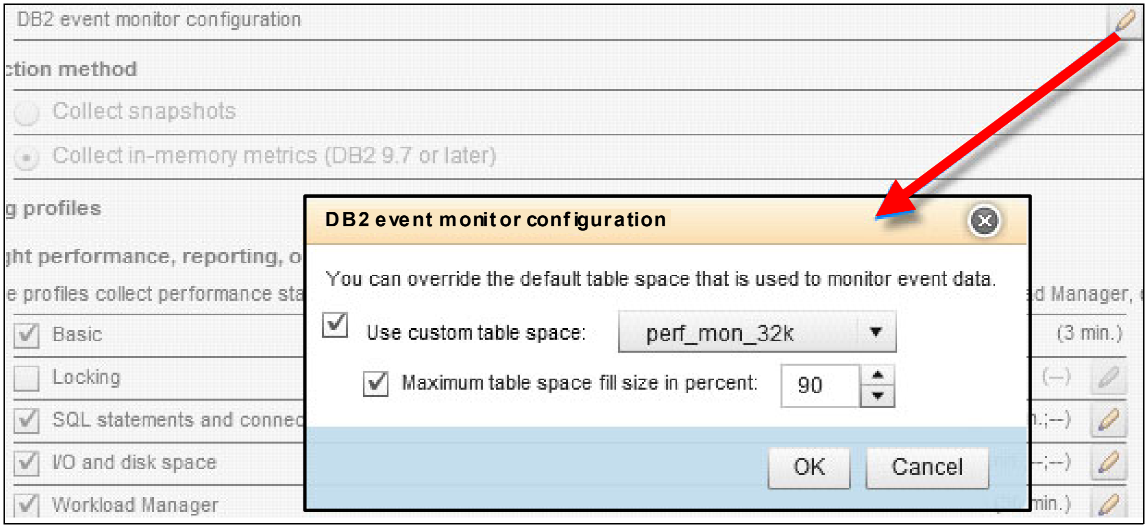

Figure 3-12 Event monitor configuration

> **Tip:** Create a dedicated 32K table space for the event monitors. If your monitored database is a partitioned database, the table space must be created across all partitions that you monitor. You can create new dedicated table space ahead of the monitoring configuration with the data collection user.

b. Monitoring profiles for inflight performance, or Workload Manager.

These profiles collect performance statistics from the data server and present the data in the inflight dashboards, in Workload Manager, or in the reports. You can edit any of these values from the default thresholds at any point.

- Basic: The Basic profile collects data from the database manager, database, buffer pool, configuration, HADR information, utilities through in-memory metric table functions, or snapshot administration views.
- Locking: The Locking profile (Figure 3-13 on page 107) uses the Locking event monitor and lock wait information settings that you can enable separately. The deadlock and lock time-out events are from the monitored database side using the Locking event monitor; the lock wait warning is set and calculated by InfoSphere Optim Performance Manager, with the collected lock-wait information using in-memory metrics. Set the **Use legacy deadlock event monitor** option to use the legacy deadlock event monitor, which was deprecated in DB2 v9.7.

In the configuration, you can also specify the separate interval for lock-waiting information collection, and define (in the filter) the kinds of applications you want to monitor (or not monitor).

Tip: The **Capture event details** provide you the options when the lock event (deadlock or time-out) happens, what information the event monitor will record and, finally, is captured by InfoSphere Optim Performance Manager and shows on the UI. The drop-down list has three options: Without statement history, Statement history, and Values. The more information you get, as a sequence, the more overhead you will have against the monitored database. For daily monitoring, use the **Without statement history** option, if locks are captured. You can change to the **With statement history** option for awhile to capture the detailed statements that cause the lock events when the lock event happens again. If you change to the Statement history option or Statement history and Values, increase the monitored database register variable DB2_MAX_INACT_STMTS and application heap size (applheapsz) database configuration. More details are at the information center.

http://pic.dhe.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.admin.mon.doc/doc/c0024837.html

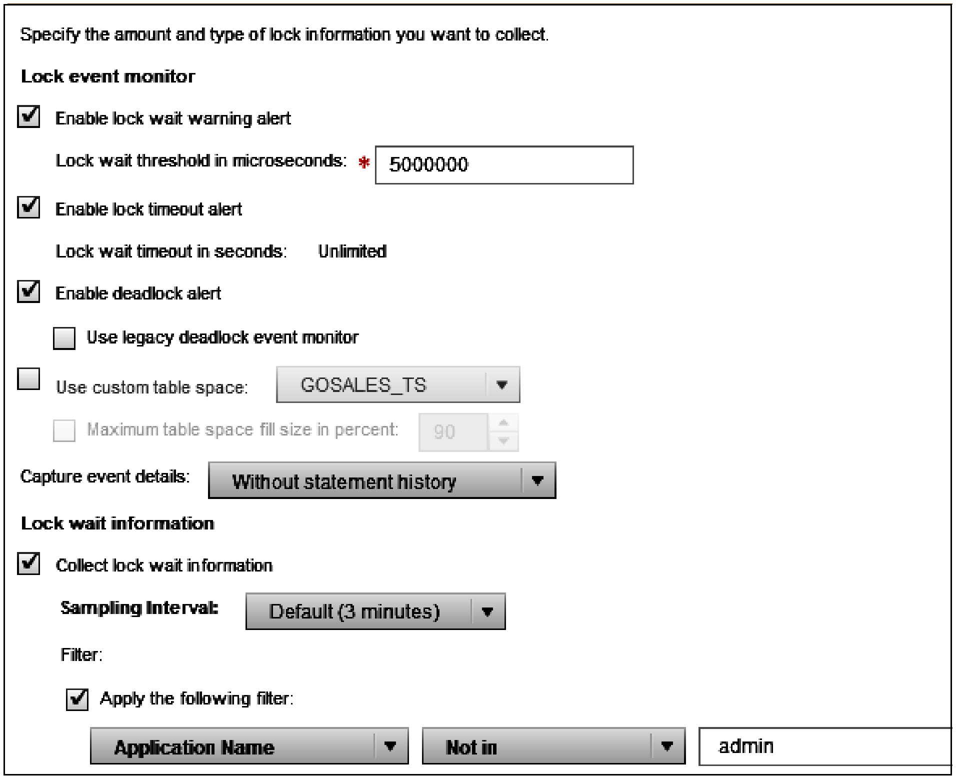

Figure 3-13 Locking related monitor configuration panel

c. SQL Statements and Connections:

The SQL Statement and Connections profile collects data from the application using in-memory metrics with related monitoring table functions. Each table function can allow specific collection interval that will overwrite the global interval. The stored procedure monitoring requires DB2 v10.1 FP2, and later. For a database prior to this version you do not see this option. You can also define filters for SQLs and connections to collect the data you want or filter out the data you do not want. Figure 3-14 on page 108 shows a filter on the SQLs summary that collects only performance data for DML SQLs, and defines another filter on the Top Executions and Connections that collects only top executions and connection information for the applications. where the Client User Id contains `henry`.

Important: The SQL summary uses the MON_GET_PKGCACHE_STMT table function to access the package cache to incrementally retrieve all the runtime information produced by the SQL statements. If the SQLs have been flushed out from the package between two consecutive collection intervals, InfoSphere Optim Performance Manager might not capture it, unless you enable the Extended Insight, which uses package event monitor to capture every SQL that has been flushed from the package cache through event monitor. Top SQLs and connections are accessing the current connections and related activities details at each time interval, so if a connection or related activities start and stop within one collection interval, then InfoSphere Optim Performance Manager might not collect the data. If you want to see more granular data, you can either change the collection interval to a small value or use real-time monitoring, both of which cause more impact to your system.

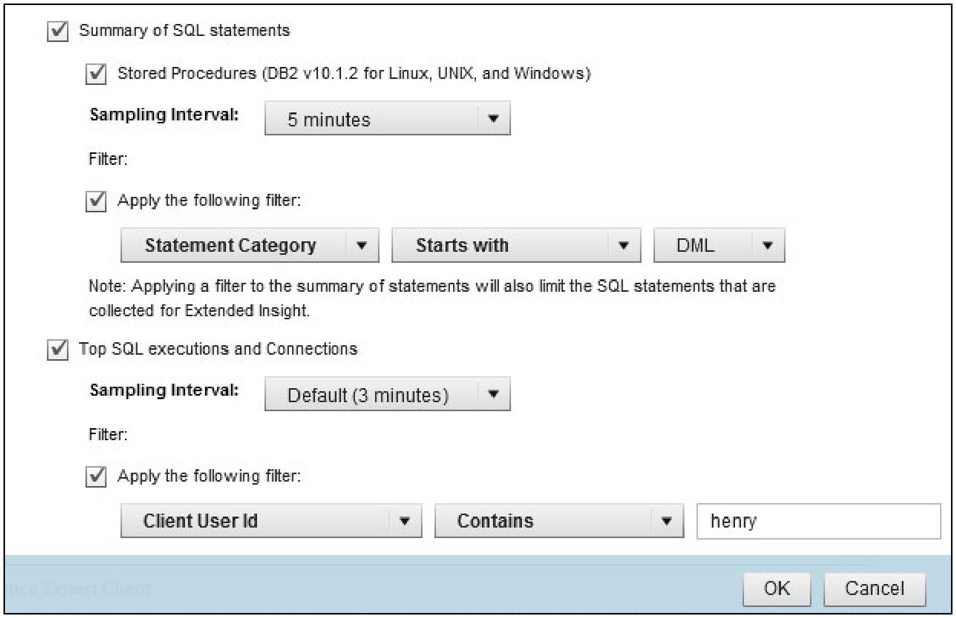

Figure 3-14 SQL statement and connections profile

- I/O and Disk Space:

  The I/O and Disk Space profile, by default, collects information for buffer pools. You can also specify to collect I/O information for tables and table spaces. You can specify the collection interval for buffer pool, table, and table space separately. You can define filters on tables, and table space for performance data collection, as shown in Figure 3-15.

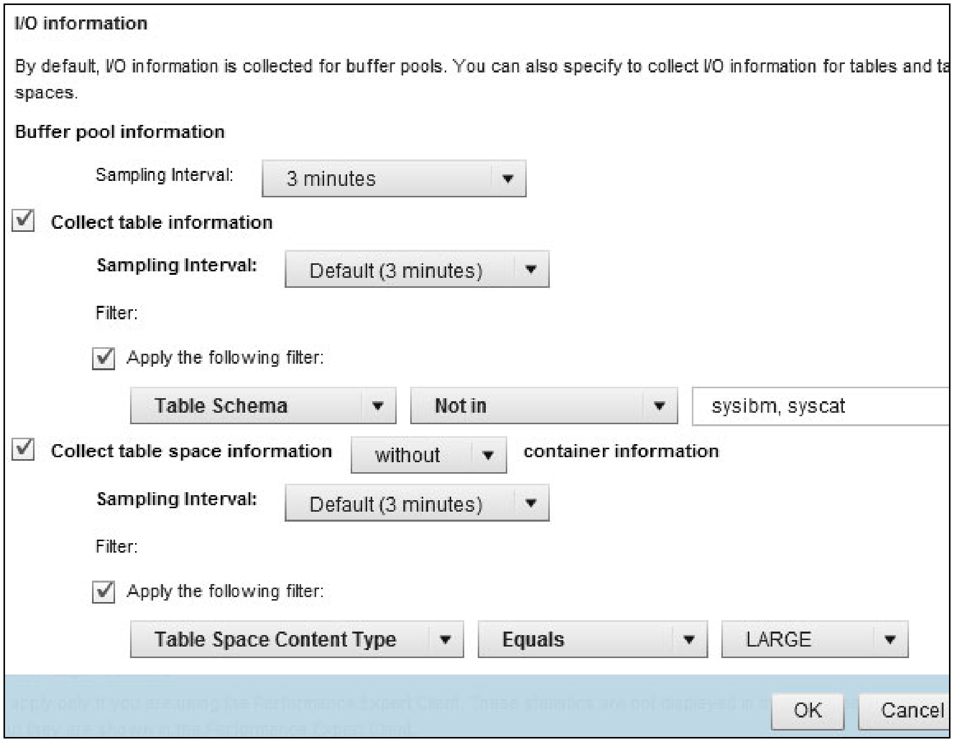

Figure 3-15 IO and Disk Space monitoring profile

- Workload Manager

  The Workload Manager profile, shown in Figure 3-16 on page 110, collects data from the statistic event monitor for workload management statistics and the table functions for WLM configuration data. A good approach is to collect the activity data more often than the configuration data if your configuration of workload, service class becomes stable.

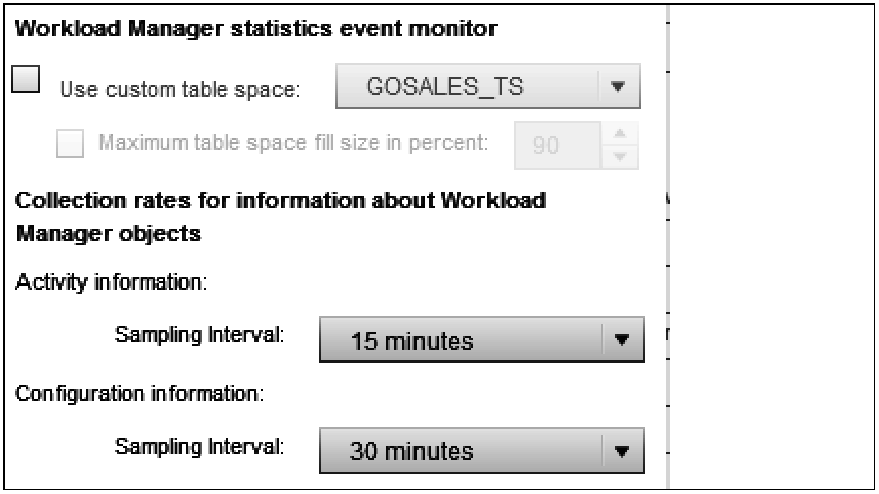

Figure 3-16 Workload Manager monitoring profile

d. Monitoring profile for Extended Insight (Figure 3-17 on page 111):

The Extended Insight profile collects statement and transaction metrics from the Extended Insight Clients and from the data server.

This option is available only if you activated the Extended Insight feature on the InfoSphere Optim Performance Manager server. If the option is listed but is not available to use, check that your Extended Insight activation is completed and is successful (see the Tip at the end of section 3.1.2, "Activating the InfoSphere Optim Performance Manager license" on page 83 for how to check the InfoSphere Optim Performance Manager edition and license).

In the configuration panel, you can choose to collect statement and transaction metrics on the client side and server side. The statement metrics on the server side will create the package cache event monitor, and the transaction metrics on server side will create the Unit of Work (UOW) event monitor. You can mask or exclude the client field information from aggregation using the **Usage of client field information** tab. If you have IBM Tivoli Composite Application Monitor (ITCAM) and plan to integrate the database transaction monitoring data with it, you can set the ITCAM for transaction information in the **Integration with Tivoli Monitoring** tab.

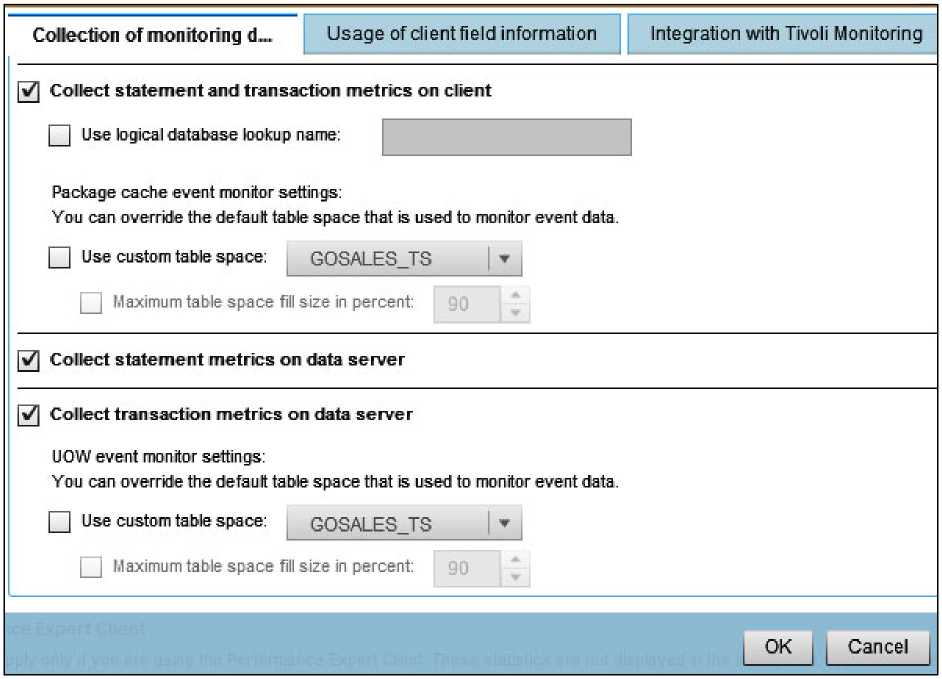

Figure 3-17 Configure the Extended Insight panel

For more details about the various monitoring profiles, see the information center:

http://pic.dhe.ibm.com/infocenter/perfmgmt/v5r2/topic/com.ibm.datatools.perfmgmt.monitor.doc/sys_templates_monitor_profiles.html

8. When you are done with configuring monitoring profiles, click **Next**.
9. View resulting DB2 settings.

   The switches, accessed monitoring functions and views, and configuration settings that are displayed on the EI panel are set for this database based on the monitoring profiles that you specified.

   InfoSphere Optim Performance Manager primarily uses two types of monitors for DB2 v9.7 (and later) monitoring: In-memory table functions and event monitors. InfoSphere Optim Performance Manager requires only one monitor switch (DFT_MON_TIMESTAMP) to be on because it still collects basic snapshot data.

   InfoSphere Optim Performance Manager creates event monitors in a monitored database depending on the configuration of the monitoring profile. Figure 3-18 on page 113 shows the list of functions and views that will be accessed by InfoSphere Optim Performance Manager corresponding to the monitoring profile configuration. Assuming that you have not set a dedicated table space in the monitoring profiles and have not set a PCTDEACTIVATE value in the monitoring profiles, InfoSphere Optim Performance Manager uses the following statements to create event monitors for the monitored database:

```
CREATE EVENT MONITOR <name> FOR LOCKING WRITE TO UNFORMATTED EVENT TABLE
(PCTDEACTIVATE 100) MANUALSTART
```

```
CREATE EVENT MONITOR <name> FOR PACKAGE CACHE WHERE
UPDATED_SINCE_BOUNDARY_TIME WRITE TO UNFORMATTED EVENT TABLE MANUALSTART
```

```
CREATE EVENT MONITOR <name> FOR STATISTICS WRITE TO TABLE
```

If the unit of work event monitor is also set to ON, the result is the following statement:

```
CREATE EVENT MONITOR <name> FOR UNIT OF WORK WRITE TO UNFORMATTED EVENT
TABLE (PCTDEACTIVATE 100) MANUALSTART
```

The configuration might also require some database configuration parameter changes listed in the *Others* section, such as the `mon_*` parameters to get the configured performance data through table functions.

Warning icons indicate any configuration settings that might cause increased overhead. Review these settings and click **Next**.

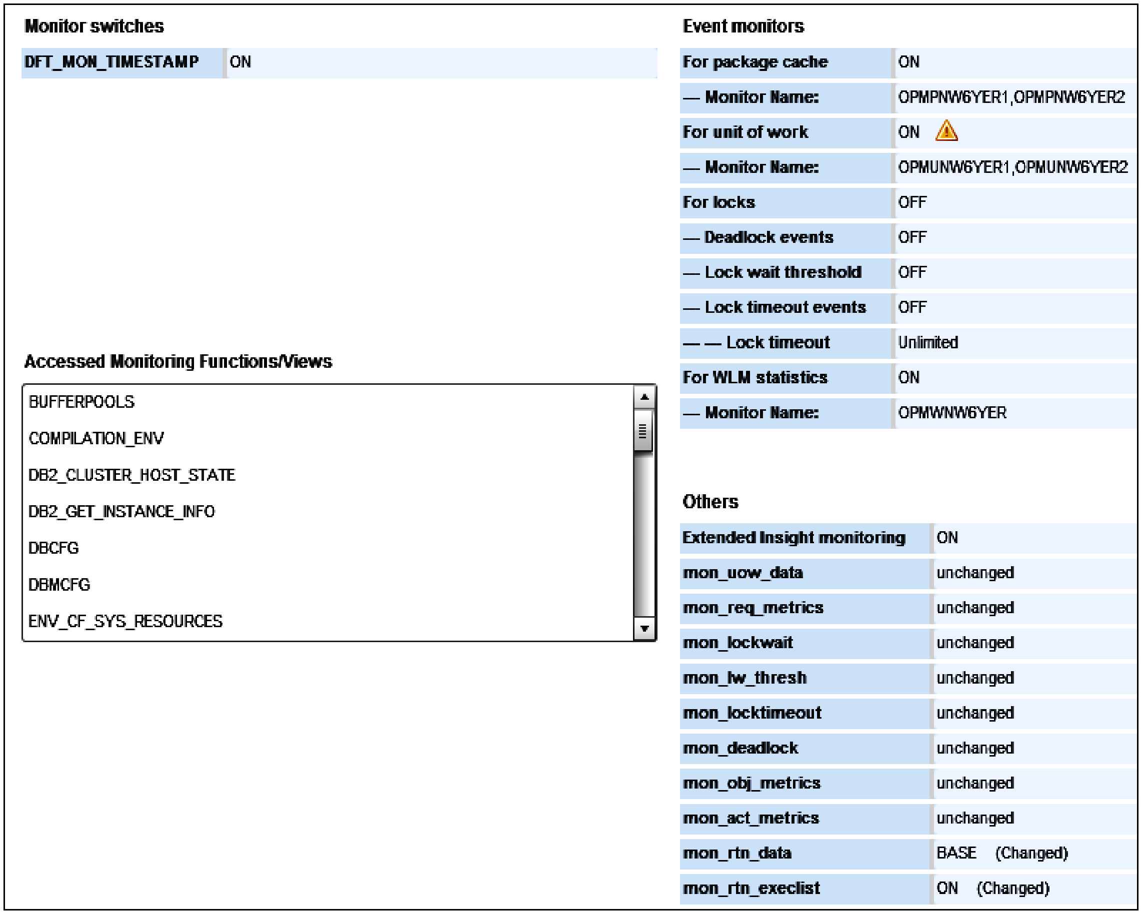

Figure 3-18 View DB2 settings for monitoring configuration

10. Configure partition (member) sets (Figure 3-19 on page 114).

    This configuration is applicable only for partitioned databases (DPF) and DB2 pureScale databases. When you configure a partitioned database (pureScale database) for monitoring, the InfoSphere Optim Performance Manager server discovers that it is a multiple partition (member) database. You can edit the configuration and add the partitions (members) to monitor. You can also assign each partition (member) a role if a partition has various monitoring requirements. InfoSphere Optim Performance Manager will collect and represent the performance data, based on the configuration here, whether to collect and show data for all partitions (members) or only certain partitions (members).

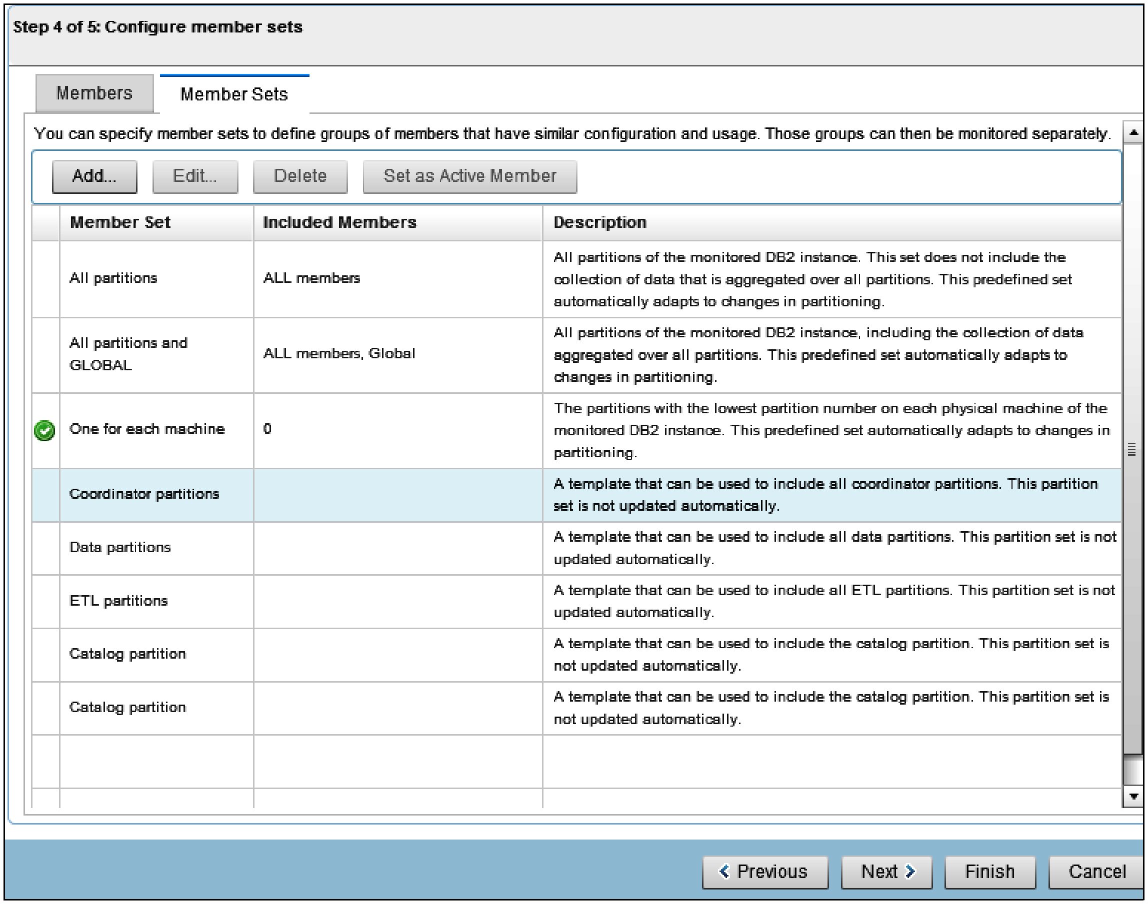

Figure 3-19 Configure partition or member sets

This panel shows a summary of the monitoring configuration that you specified for this database. The monitoring settings take effect when this monitoring configuration is saved. Complete the configuration by clicking **Finish**. It takes a few minutes to complete the configuration, and InfoSphere Optim Performance Manager will first create a set of tables and views in the repository database to hold the collected performance data, and then connect to the monitor database.

When the configuration is successful, you can open various dashboards from the successful configuration message box, or configure monitoring privileges for other users.

If you choose **Define monitoring privileges for other users,** and click **OK**, the privilege configuration panel opens.

11. Configure monitoring privileges (Figure 3-20 on page 116, Figure 3-21 on page 117).

    When you configure a database for monitoring, you can assign various monitoring privileges for users, groups, or roles that require access to the monitored database. These privileges define the monitoring operations a user is allowed to have. Consider the following examples:

    - The *Can Monitor* privilege allows the user to look at collected monitoring data in this database. It is verified when you open a dashboard. Select a database and specify user credentials for that database.
    - The *Can Manage Alerts* privilege allows a user to change alert thresholds, and delete alerts and notifications for this database. It is verified on the alert notification and configuration dashboards when you select a database and specify user credentials for that database.
    - The *Can Monitor in Real-Time* privilege allows a user to do real-time monitoring of the database. It is verified on each dashboard when you change from historical mode to real-time mode and specify user credentials for this database.
    - The *Can Enable Automatic Data Collection* privilege allows a user to enable and disable the atomic data collection for monitored database. It is verified in the Database dashboard when you disable or enable automatic data collection and specify user credentials for this database.
    - The *Is Database Owner* privilege allows a user to edit the connection properties, define blackout, change monitoring configuration, and run jobs against the database. It is verified when you do these actions of the database connection.

    The *Can* privileges are verified on the monitored database side; the *isDatabaseOwner* privilege is verified on the InfoSphere Optim Performance Manager server side. For more details about authorization, see 2.6, “User security” on page 55 that introduces you the security and authentication concept of InfoSphere Optim Performance Manager.

Figure 3-20 shows the options to enforce the privilege or to open to all web console users.

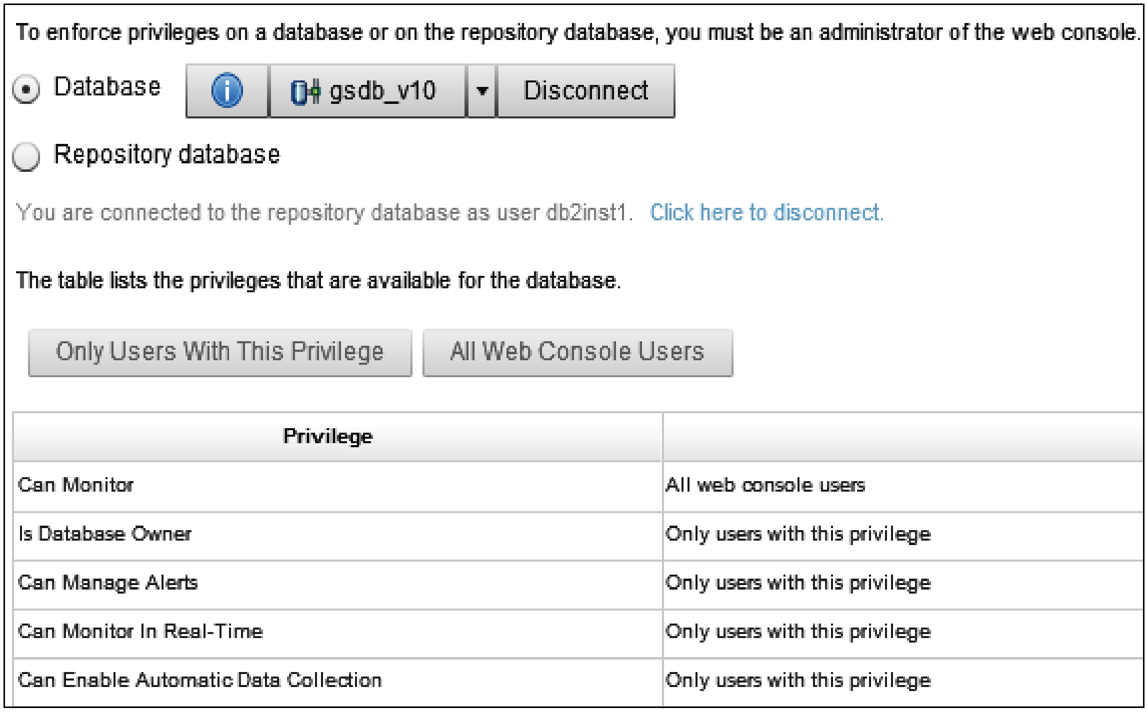

Figure 3-20 Configure database monitoring privilege--enforce privilege

If you enforce the privilege to **Only users with this privilege**, you must then grant and revoke privileges to users on the monitored database (Figure 3-21), where you can grant and revoke those five privileges to monitored database users.

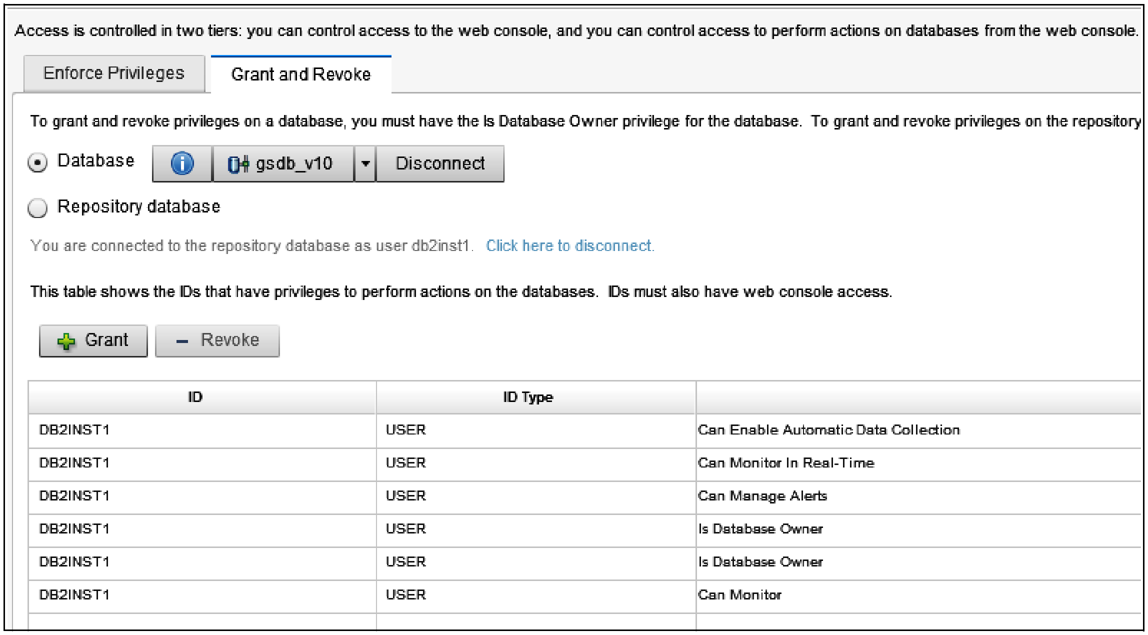

Figure 3-21 Configure database monitoring privilege -- grant and revoke

Control automatic data collection

This section describes types of capabilities for automatic data collection:

- Planned outage defines a blackout for monitoring.
- Ad hoc start and stop data collection in UI; start and stop data collection with the peconfig program.
- Edit configuration monitoring.

In some cases, you might need to control the automatic data collection such as when there is a planned outage for offline back up, or there is a time window dedicated for a task exclusive to any other connections. Meanwhile, you might also need to edit the monitoring configuration to collect more performance data for analysis or reduce the overhead to the monitored database. This section describes how to control the data monitoring, including automatic data collection and the monitoring profiles.

Use blackout for planned outage

If you have a planned outage for the monitored database, or there is a repeated time that you do not want InfoSphere Optim Performance Manager to connect to the monitored database for data collection, you can define a blackout schedule for that database. In the blackout time window, InfoSphere Optim Performance Manager will disconnect from that database, and stop the data collection against the monitored database. In the Database dashboard, you see the Blackout Active column is set to *yes*.

In the Database dashboard, you also see the list of the database connections and a column named Blackout Active to show the blackout status for each connection. You can choose a connection, and click **Blackout** → **Schedule** menu (Figure 3-22) to define a blackout schedule for that monitored database.

You can see the blackout status in the Database dashboard, and also can release the blackout event by choosing **Blackout** → **Release Blackout Now** from the menu on top of the connection list table.

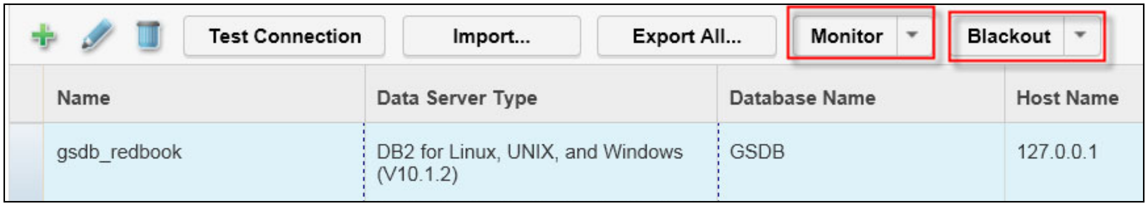

Figure 3-22 Connection menu for monitor and blackout

Figure 3-23 on page 119 shows the UI to define a blackout schedule, which is similar to other scheduled jobs. Define the start, end time, repeat mode, and so forth, and then you can choose several databases to apply the schedule to if they share the same blackout policy.

When the monitored database has been blacked out, InfoSphere Optim Performance Manager will not connect to the database and will block any request from InfoSphere Optim Performance Manager server to that database during the blackout. So, if you test the connection, it fails with the message that the database is in blackout.

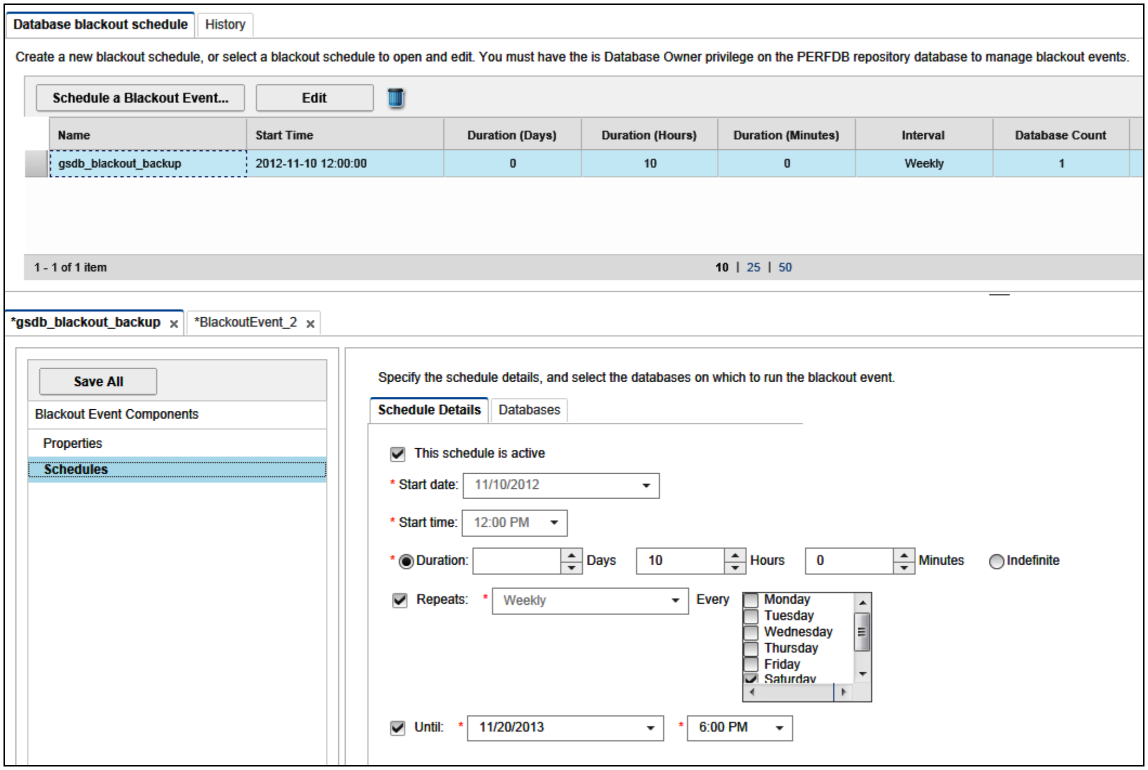

Figure 3-23   Define database monitoring blackout

> **Note:** If the database is in blackout status, InfoSphere Optim Performance Manager will keep the monitoring configuration to the monitored database, and also keep all the historical performance data in repository data. It will however, drop the event monitors created by InfoSphere Optim Performance Manager against the monitored database.

Meanwhile, in the health summary, you see an operation alert for that database saying the database monitoring is in blackout (Figure 3-24). And you also cannot view the performance data for that database if it is blacked out because you cannot connect to that database.

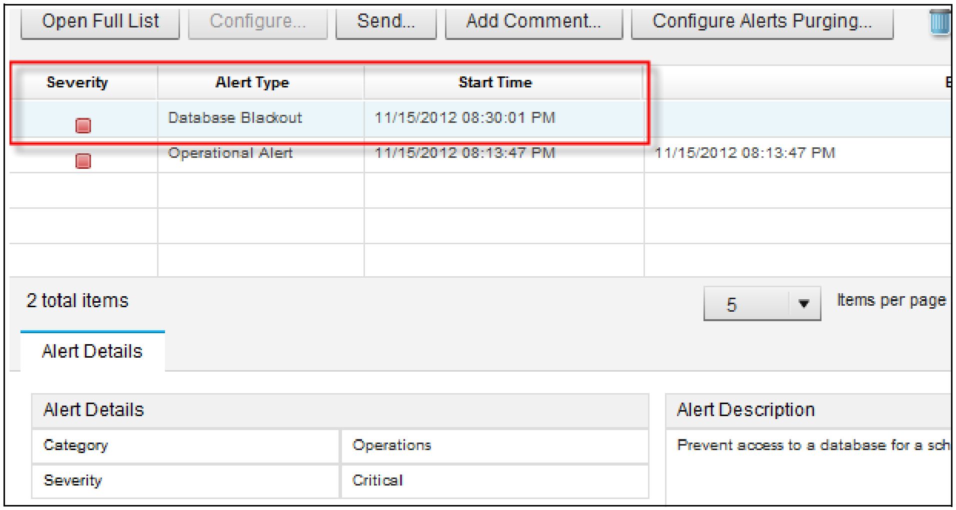

Figure 3-24 Operation alert for blackout database

Stop or start the automatic data collection on demand

In some cases, you may need to stop or start the automatic data collection on demand for problem determination or unplanned actions to the monitored database such that you do not want any connections to that database for a while. You can go to the Database dashboard, choose the connection you plan to stop or start, use the **Monitor** → **Disable** (or **Enable**) **automatic data collection** menu at the t top of the database connection list table.

You can also disable (or enable) automatic data collection in the InfoSphere Optim Performance Manager server side with the **`peconfig`** command. The command is in the installation directory of the `RepositoryServer/bin` folder (Example 3-8).

Example 3-8 Use peconfig to disable (enable) automatic data collection

```
peconfig -disable -id 1 (disable the data collection, you can use
peconfig -list to get the database id.)

peconfig -enable -id 1 (enable the data collection)
```

Tip: What is the difference between blackout and disable automatic data collection? Both actions stop the monitoring of the monitored database, and keep the monitoring configuration.

- With blackout, there will not be any access to the monitored database from InfoSphere Optim Performance Manager, and the user cannot view the database performance data.
- With disable automatic data collection, the user can still do real-time monitoring, and also can view the historical performance data of that monitored database.

Editing the monitoring configuration

You might need to update the monitoring configuration to select or remove some monitoring profile on demand. For monitoring purposes, go to the Database dashboard, choose the database connection, and select the **Monitor** → **Edit monitoring configuration** menu to update the configuration. Also you can unconfigure the monitoring configuration by using the **`Unconfigure Monitoring`** command.

Note: If you unconfigure monitoring of the monitoring database, all changes made by InfoSphere Optim Performance Manager to the database are restored, the database configuration is changed, the monitor switch will be turned off, and the created event monitors, schema, and UDFs will be dropped. Meanwhile, the historical performance collected in the repository database and corresponding tables will be dropped. Therefore, if you want to stop monitoring only for awhile, and still keep the performance data, you can either use the blackout feature or disable automatic data collection action for awhile, through the UI.

3.2.3 Monitoring templates and proven practices

This section describes several monitoring templates and proven practices.

Predefined templates

There are 18 predefined monitoring templates in InfoSphere Optim Performance Manager that cover most of the DB2 system workload to help you get started with the monitoring configuration. Each template defines the monitoring profile to be collected, collection interval, data aggregation time, and data retention time. Because the alerts are determined by what metrics will be collected, the template indirectly configures the alert generation.

There are five major workload types: OLTP, OLAP, BI, SAP BI, and SAP ERP system. For each workload type, InfoSphere Optim Performance Manager offers three templates:

- Starter template: Collects basic information with short retention times suggested for initial deployments. Upgrades to more comprehensive monitoring configurations should be undertaken after validating available resources for the DB2 performance repository using the Monitoring server template.
- Proven practices template: Collects the information needed for steady state monitoring, reporting, and occasional problem determination according to proven practices.
- Detailed diagnostics template: Collects the information needed for detailed problem analysis and reporting. These templates might be turned on during certain time periods for detailed problem determination or reporting scenarios.

Another monitoring option is real-time monitoring for which InfoSphere Optim Performance Manager does not automatically collect the performance data against the monitored database, but collects only upon request through the real-time dashboard.

Figure 3-25 on page 123 illustrates a suggested workflow for the monitoring deployment. A good approach is to impose noticeable overhead on the monitored server starting from real-time monitoring, which has zero configuration against the monitor database, then upgrade to starter history template, which collects the basic and necessary performance data for the monitored database with limited overhead to both monitored database and InfoSphere Optim Performance Manager server. Gradually, when you are in the starter template for awhile, and plan to collect more data for deep analysis, then you can change to the proven practices template and stay with this monitoring template for monitoring. In case, for a period of time, you need do performance problem determination (PD) to resolve some detailed performance issues, you can set the detailed diagnostics to collect detailed performance data. When you are done with the PD, you can change back to proven practices template.

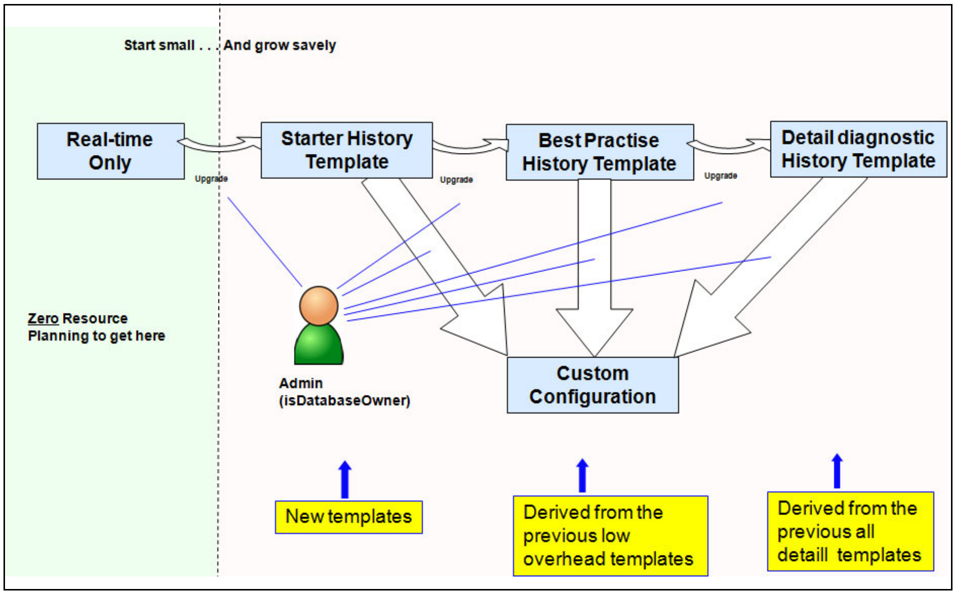

Figure 3-25 Monitoring deployment stages

There is a Test/QA System and Development System template, but use it only in your test and development environment to determine the performance bottlenecks before they arrive in production.

For the detailed description of each template, see the following link:

http://pic.dhe.ibm.com/infocenter/perfmgmt/v5r2/topic/com.ibm.datatools.perfmgmt.monitor.doc/profile_settings.html

Customizing templates

Apart from the predefined templates, you can also update the monitoring configuration based on one template for your own environment, and then take this as the monitoring template for other database monitoring configurations. Then, when you can configure a new database for monitoring, you can select **Configure like** (Figure 3-8 on page 100) to copy an existed database configuration template.

Configuration monitoring proven practices

For each of the monitoring profiles during the configuration (see 3.2.2, “Adding and configuring database monitoring” on page 96), this section provides tips for how to make the proper configuration to meet the monitoring requirement, minimize the overhead to the monitored database, and also make sure the InfoSphere Optim Performance Manager server runs efficiently.

See the following web address for more proven practices when you deploy InfoSphere Optim Performance Manager in large scale environments for database monitoring on DB2 V9.5. The proven practices are based on InfoSphere Optim Performance Manager 4.1, but some of the practices are still valid for InfoSphere Optim Performance Manager 5.2 when you monitor DB2 database version v9.5.

http://ibm.co/17DTK4m

3.2.4 Consider impact on the monitored database and applications

To collect the performance information from the monitored DB2 databases, InfoSphere Optim Performance Manager must query the in-memory table functions, administration view and run event monitors. Depending on the monitoring profile, InfoSphere Optim Performance Manager might introduce overhead on the monitored database and monitored applications.

Configuration impact to the monitored database and applications

Table 3-1 lists that might made against the monitored database for a monitoring configuration, depending on what monitoring profile has been configured.

Table 3-1 List of changes to the monitored database

Category	Created objects and changes
Monitor switches	Requires that DFT_MON_TIMESTAMP be turned on.
Database configuration	The following database monitoring configuration might be changed because of a different monitoring configuration you can get from the configuration summary: MON_REQ_METRICS MON_ACT_METRICS MON_OBJ_METRICS MON_RTN_DATA MON_RTN_EXECLIST MON_UOW_DATA MON_UOW_PKGLIST MON_UOW_EXECLIST MON_LOCKTIMEOUT MON_DEADLOCK MON_LOCKWAIT MON_LW_THRESH MON_PKGLIST_SZ MON_LCK_MSG_LVL
Schema	Create a schema with name `OPM`

Category	Created objects and changes
Table functions	CAN_MONITOR_IN_REALTIME CAN_MONITOR CAN_MANAGE_ALERTS (under the schema of OPM)
Event monitors	► Locking ► Package Cache ► UOW ► Statistics ► Tables for each corresponding event monitor. As for package cache and UOW event monitor, InfoSphere Optim Performance Manager will create a pair of each type, keeping only one active at any time so that data is not lost.
Stored procedures	► DB2MON_LOC, the schema is the collection user ID ► One watchdog stored procedure for each type of event monitor, with the InfoSphere Optim Performance Manager schema. This requires that ATS be enabled for the monitored database. InfoSphere Optim Performance Manager will create the watchdog procedure to drop event monitor and tables when in some cases InfoSphere Optim Performance Manager fails to connect to the monitor database. For details of ATS, see the DB2 information center: http://pic.dhe.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.admin.gui.doc/doc/t0054396.html

For the changes to the monitored applications of the monitoring configuration, see 3.3, “Installing and configuring Extended Insight Client” on page 127.

Daily monitoring impact to the monitored database and applications

During the data collection for daily monitoring of the database and applications, there is also some overhead to the monitored database and applications. For DB2 v9.7 and later, InfoSphere Optim Performance Manager uses in-memory table functions to collect performance data, which has a minimal impact compared with the snapshot methods. Meanwhile, InfoSphere Optim Performance Manager uses the unformatted table for event monitors, which can also minimize the impact to the database. The in-memory metrics overhead is much smaller than the snapshot; you can configure the specific collection intervals for each monitoring profile related table functions, and also define the filter for the collection. Therefore, the major overhead should be on the event monitors. For the details about event monitor overhead, see *IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows*, SG24-7925 (section B.1.2).

As for overhead against the monitored applications, the impact is only when you enable for the EI feature, which collects the SQLs and transactions from the data client side. For more information, see I*BM Optim Performance Manager for DB2 for Linux, UNIX, and Windows*, SG24-7925 (section B.1.3).

3.2.5 Other InfoSphere Optim Performance Manager configurations

Besides the user access configuration, database monitoring configuration, and database monitoring privilege configuration, there are other configurations for InfoSphere Optim Performance Manager through the web UI. The following list indicates examples:

- Alert configuration for the alert threshold, and notification (4.2.3, “Alert configuration” on page 163).
- Repository configuration to configure the repository database information after there are any changes to the repository database. On the task navigator click **Product Setup** → **Configuration Repository** to see the current used repository database. You can then edit it using **Select Repository Database**.

> **Tip:** If the repository database password, port number or IP have been changed, you can use this configuration to update the repository database information to make InfoSphere Optim Performance Manager server aware of such changes or else the server may fail to connect to repository database.

You can also update the Extended Insight communication using this configuration. This update is the same as the Extended Insight communication setup made during the Extended Edition activation step (Figure 3-17 on page 111). On the task navigator click **Product Setup** → **Monitoring Server.**

> **Tip:** If you update the Extended Insight communication port on the InfoSphere Optim Performance Manager server side, ensure that you also make the corresponding update on the Extended Insight client side, otherwise the EI client might not be able to communicate with InfoSphere Optim Performance Manager for EI monitoring data.

3.3 Installing and configuring Extended Insight Client

The Extended Insight Client was renamed to the *IBM Data Tools Runtime Client (DTRC)*. You must install the DTRC on the monitored data client where you plan to have the Extended Insight monitoring feature.

See the demo of how to install and configure the Extended Insight client software for EI monitoring:

http://www.ibm.com/developerworks/offers/lp/demos/summary/im-installoptimpm51-2.html

Also see *IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows*, SG24-7925 for more details about Extended Insight Client installation and configuration.

3.4 Installing and configuring IBM InfoSphere Optim Query Workload Tuner

The IBM InfoSphere Optim Query Workload Tuner (InfoSphere Optim Query Workload Tuner) is software that runs at your desktop. It is built with Eclipse technology, and differs from InfoSphere Optim Performance Manager, which is server-based software. Figure 3-26 on page 128 shows the deployment architecture of InfoSphere Optim Query Workload Tuner, which is now integrated into Data Studio as the query tuning feature. So, as this section describes InfoSphere Optim Query Workload Tuner, it refers to the query tuning component within Data Studio. InfoSphere Optim Query Workload Tuner will connect to the target database for query analysis and tuning using a JDBC connection. This requires the target database to be accessed by the desktop where InfoSphere Optim Query Workload Tuner is installed (typically the user's desktop). In some cases, the production database is not accessible directly through the desktop, so there is an alternative solution for deploying InfoSphere Optim Query Workload Tuner (see 3.4.7, "Deploying InfoSphere Optim Query Workload Tuner in a DMZ environment" on page 138).

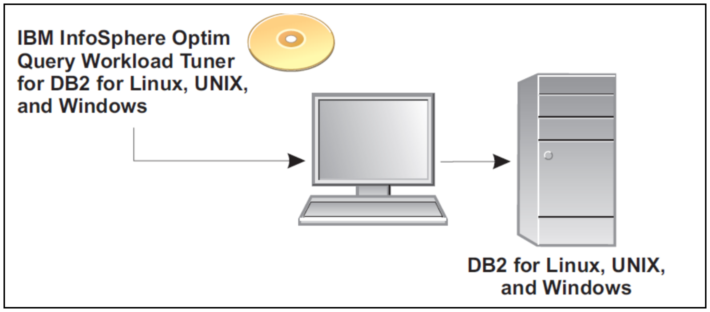

Figure 3-26 Deployment architecture

The installation of InfoSphere Optim Query Workload Tuner client is the same as installation of Data Studio. You can also install the License Activation Kit together if you already have the InfoSphere Optim Query Workload Tuner license. Or you can first use the no-charge query tuning features, and later you can install the License Activation Kit.

Before starting query tuning against the target database, you must configure the database so it is ready for query analysis and tuning (see 3.4.6, "Configure database for query tuning" on page 133). As a query analysis and tuning product, InfoSphere Optim Query Workload Tuner can also be integrated with development tools (such as the Data Studio development component) and performance monitoring product (InfoSphere Optim Performance manager). Some set up is required to make it work (see 3.4.8, "Configuration to capture from InfoSphere Optim Performance Manager repository" on page 139).

3.4.1 InfoSphere Optim Query Workload Tuner information road map

The InfoSphere Optim Query Workload Tuner information roadmap (Information Roadmap) is a good way to get started with the product. The roadmap has the current information for InfoSphere Optim Query Workload Tuner including an overview, and installation, configuration, tuning, and support. Open the Information Roadmap for InfoSphere Optim Query Workload Tuner v3.2:

http://www.ibm.com/developerworks/data/roadmaps/roadmap_ioqt_ioqwt_32.html

Here, you can also find the quick-start guide for InfoSphere Optim Query Workload Tuner to understand the basic and necessary steps to get InfoSphere Optim Query Workload Tuner running in your environment.

3.4.2 Installing InfoSphere Optim Query Workload Tuner client

Before installing the client (as part of the Data Studio), review the system prerequisites of the OS, and hardware of the desktop where the client will be installed.

http://www.ibm.com/support/docview.wss?uid=swg27036261

InfoSphere Optim Query Workload Tuner client support for Windows and Linux, are both GUI-based installations or silent installations. This section shows an example of how to install the client on 64-bit Windows with administrator role.

A good approach is to install Data Studio with administrator privilege. If you do not have the OS administrator privilege, you can use **setup_win7_nonadmin64** executable to start the installation. When you use that executable, first the welcome page opens.

First, the welcome page opens. If you do not already have the Installation Manager installed, the installer will install it.

> **Note:** IBM Installation Manager is a program for installing, updating, and modifying packages. You can get detailed information at the following location:
>
> http://publib.boulder.ibm.com/infocenter/dstudio/v3r2/topic/com.ibm.datatools.base.install.doc/topics/c_plan_imover.html

As you begin the installation, go to **Install Package**, and choose the **Administrative Installation**.

> **Tip:** If you use non-administrative installation, there might be a problem to create shortcuts of the installer. If so, you can use the steps in the following link to resolve it:
>
> http://www-01.ibm.com/support/docview.wss?uid=swg21594734

In the Install Package panel, you can choose the component to be installed. If you have an InfoSphere Optim Query Workload Tuner license, you see the *IBM InfoSphere Optim Query Workload Tuner for Db2 for Linux, UNIX, and Windows License Activation Kit*. Choose all of the package, and click **Next**. You can also install the License Activation Kit later by using Installation Manager.

Accept the license, specify the directory of the installation and the architecture (32-bit or 64-bit), and then select the languages to be installed.

Tip: InfoSphere Optim Query Workload Tuner client is translated into multiple languages. After the installation, when you launch InfoSphere Optim Query Workload Tuner client, by default it uses your OS-configured language and locale. However, you can specify a UI language if you have other languages installed, as follows:

1. Go to the Data Studio installation directory, such as `C:\Program Files\IBM\DS3.2.0` (default path on Windows)
2. Open `eclispe.ini` file with a text editor.
3. Append one line at the end of the file, such as `-Duser.language=zh` (*zh* is for Simplified Chinese, or other languages such as *fr*, as examples).
4. Save the file, and restart InfoSphere Optim Query Workload Tuner client.

Finish the installation. For other installation methods and information, see the information center:

http://publib.boulder.ibm.com/infocenter/dstudio/v3r2/topic/com.ibm.datatools.qrytune.installconfig.doc/topics/ioqwt32_top.html

3.4.3 Upgrade InfoSphere Optim Query Workload Tuner and apply fix packs

You can upgrade to a new version or apply fix packs using Installation Manager; see how at the following location:

http://publib.boulder.ibm.com/infocenter/dstudio/v3r2/topic/com.ibm.datatools.base.install.doc/topics/c_plan_consider_upgrade_product.html

Tip: The upgrade to the new version using Installation Manager is the suggested approach. This approach can keep the existing package and install the new one on it. If you have a concern about the desktop disk space, an alternative solution is as follows:

1. Uninstall the current version while keeping the workspace files.
2. Install the new version.
3. For the new version, start it and choose the previous workspace.
4. Validate the updated client. If it is successful and you can get all the configurations, projects, and connection information, it is completed.

3.4.4 Activate InfoSphere Optim Query Workload Tuner license

Now you have the latest InfoSphere Optim Query Workload Tuner client installed either from a fresh installation or an upgrade. The next step is to activate the license. InfoSphere Optim Query Workload Tuner actually checks the license on each target database to determine whether the database has the license for query analysis and tuning. So the license is finally applied to each target database. You might need to have the proper privilege to activate the license to the target database. See 3.4.5, “User security for InfoSphere Optim Query Workload Tuner” on page 132 for details.

There are two ways to activate the license on target database. For both you must ensure that the License Activation Kit has been installed on the client side.

- Activate the license from the client side

  If a license for IBM InfoSphere Optim Query Workload Tuner is not active on a DB2 database, when you connect to that database to start tuning for the first time, InfoSphere Optim Query Workload Tuner client will install and activate the license on the target database.

- Activate the license to the database directly with scripts

  You can use scripts to activate the license. Go to the InfoSphere Optim Query Workload Tuner installation directory (`C:\Program Files\IBM\DS3.2.0`), then to the `\QueryTunerServerConfig\all_features\LUW\License` folder. There are two executable script files:

  - `License.bat` for a database on Windows
  - `License.sh` for a database on Linux and AIX.

  You can copy the script files and the *`qwt_license.db2`* file to the target database. Run them with the proper privilege to activate the license.

> **Tip:** You can check whether the database has been activated for an InfoSphere Optim Query Workload Tuner license. Open the client, connect to that database, and choose the connection for tuning configuration (Figure 3-28 on page 135). The panel shows the license status for that database (it is a blue color):
>
> ```
> InfoSphere Optim Query Workload Tuner for DB2 Linux, UNIX and
> Windows is activated on the database server.
> ```

If you do not activate the license of the target database, you can still use this no-charge feature of InfoSphere Optim Query Workload Tuner for that database. See the list of features (no-charge and priced) for InfoSphere Optim Query Workload Tuner:

http://publib.boulder.ibm.com/infocenter/dstudio/v3r2/topic/com.ibm.datatools.qrytune.relinfo.doc/topics/featuresbyproduct.html

3.4.5 User security for InfoSphere Optim Query Workload Tuner

InfoSphere Optim Query Workload Tuner is desktop-based client product. As such, there is minimal user security. However, because it is used to analyze and tune the queries on the target database, and also integrate with other product, there are required privileges for the target databases.

User privilege to activate license

To activate a tuning license against the target database, you must have the authority or privilege to run the **CREATE FUNCTION** statement on that database.

User authority and privilege to configure for tuning

To configure the database for tuning, you must have following privileges for the target database:

- You must have the authority to install a JAR file with a stored procedure on the database server. For more information, see the information center at the following address:

  http://publib.boulder.ibm.com/infocenter/db2luw/v10r5/index.jsp?topic=%2Fcom.ibm.db2.luw.apdv.routines.doc%2Fdoc%2Fr0006425.html

  (This is required for database prior to DB2 V9.7.5, while optional for DB2 v9.7.5 and follow on release)
- You must have the authorities and privileges for running the **CREATE PROCEDURE** (external) statement.
- You must have the authorities and privileges to create EXPLAIN tables on the database, if these tables are not yet created.
- You must have the authorities, privileges, or both to grant the required privileges to users who will be tuning SQL statements with InfoSphere Optim Query Workload Tuner.

User privilege to capture and tune queries

After the database is configured for query tuning, you then have proper privilege to access the target database and capture the queries, and to analyze and tune the queries. For detailed user privileges required, see the information center:

http://publib.boulder.ibm.com/infocenter/dstudio/v3r2/topic/com.ibm.datatools.qrytune.installconfig.doc/topics/authprivluw.html

User privilege to capture from InfoSphere Optim Performance Manager repository

InfoSphere Optim Query Workload Tuner can integrate with InfoSphere Optim Performance Manager to get the query runtime metrics for the monitored database. To start tuning requires a connection to the InfoSphere Optim Performance Manager repository database with proper privileges. The connection user must belong to the group that was specified in the InfoSphere Optim Performance Manager installation (as in “Console installation on AIX Server” on page 76, step 14 on page 82) to specify the connection user information.

3.4.6 Configure database for query tuning

Now that the InfoSphere Optim Query Workload Tuner client is installed, the target database has been activated for license (with scripts) or will be activated after you first connect to the database and choose **Configure Tuning** or **Start Tuning**, and also know the required privilege to configure, capture, and tune against the target database. The next step is to configure the database for query tuning. You can configure either through the InfoSphere Optim Query Workload Tuner client UI or by using the scripts to configure directly on the target database side for tuning.

Configure tuning with InfoSphere Optim Query Workload Tuner client UI

In the Data Studio (InfoSphere Optim Query Workload Tuner Client) UI, open Database Explorer and choose the database connection that you plan to configure for monitoring. (You can already have a database connection create.) Connect the database connection with user credentials, and right-click **Analyze Tune** → **Configure for Tuning** → **Advanced Configuration and Privilege Management** (Figure 3-27 on page 134) to start the configuration of query tuning.

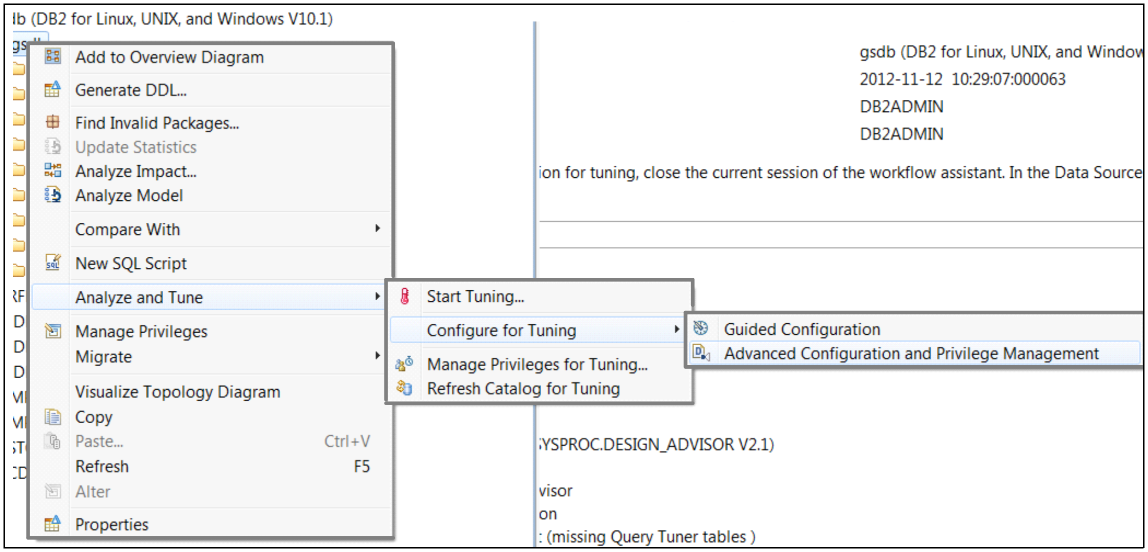

Figure 3-27 Configure database for query tuning

Figure 3-28 on page 135 shows the advanced configuration panel. Use the following steps:

1. If the database does not have the EXPLAIN tables created, or the EXPLAIN tables are out of date, you see the EXPLAIN tables section is invalid. You can use the InfoSphere Optim Query Workload Tuner UI to create the EXPLAIN tables.
2. Deploy the stored procedure for Index Advisor (this is only for DB2 prior to V9.7.5). Starting from DB2 v9.7.5, the stored procedure is built into the database system.
3. Create the Workload Control Center (WCC) object for workload tuning.
4. You can also configure the privilege (*Grant Privilege*) to users who can access (insert, delete, update, select) the WCC tables, and therefore can do the workload tuning and capture workload management.

If you do not want to do the configuration in these steps, you may also use the *Configure Database Automatically* feature to let InfoSphere Optim Query Workload Tuner configure for you using default settings, or choose the *Guided Configuration* option to configure the tuning automatically with default settings.

Note: When you create the EXPLAIN tables and WCC tables, you can also specify the table space to hold those tables. InfoSphere Optim Query Workload Tuner display a list of usable table space in that database. A good approach is to choose a table space with a page size greater than 8 KB.

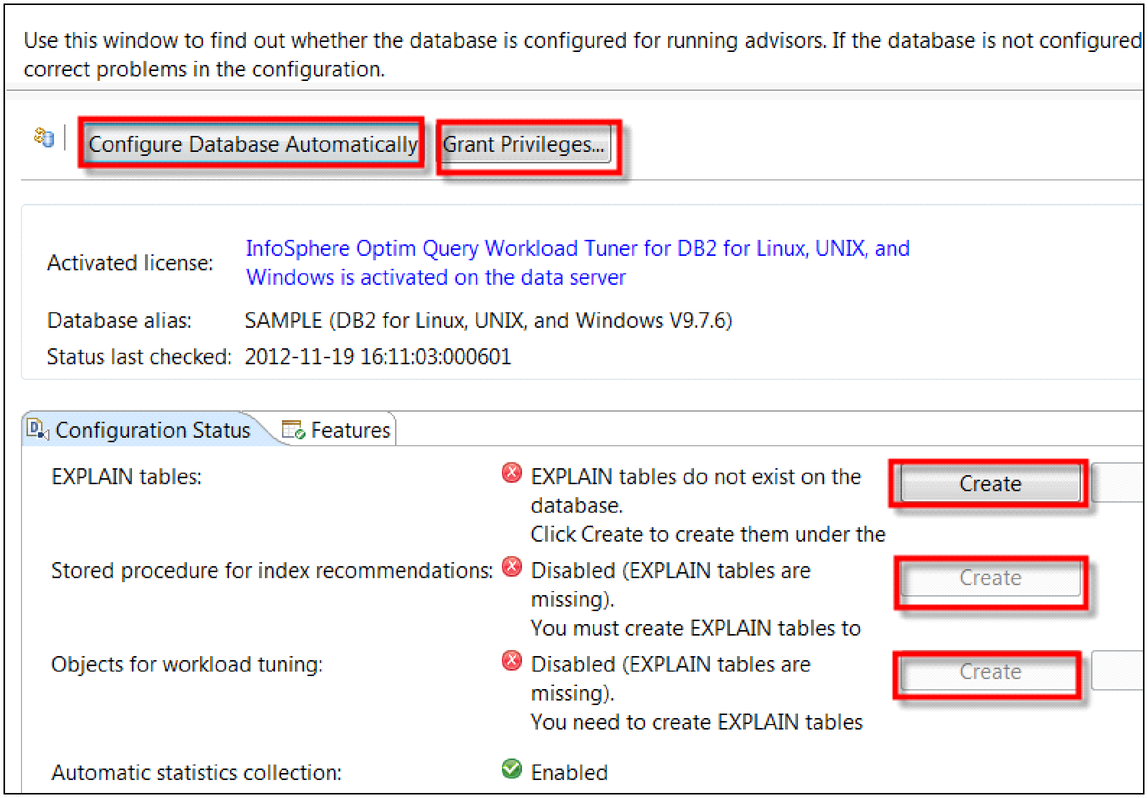

Figure 3-28 Advanced configuration for tuning

When you finish the configuration, in the same panel as Figure 3-28 on page 135, open the Feature tab to view the available features. It shows that all the tuning features are available, as in Figure 3-29.

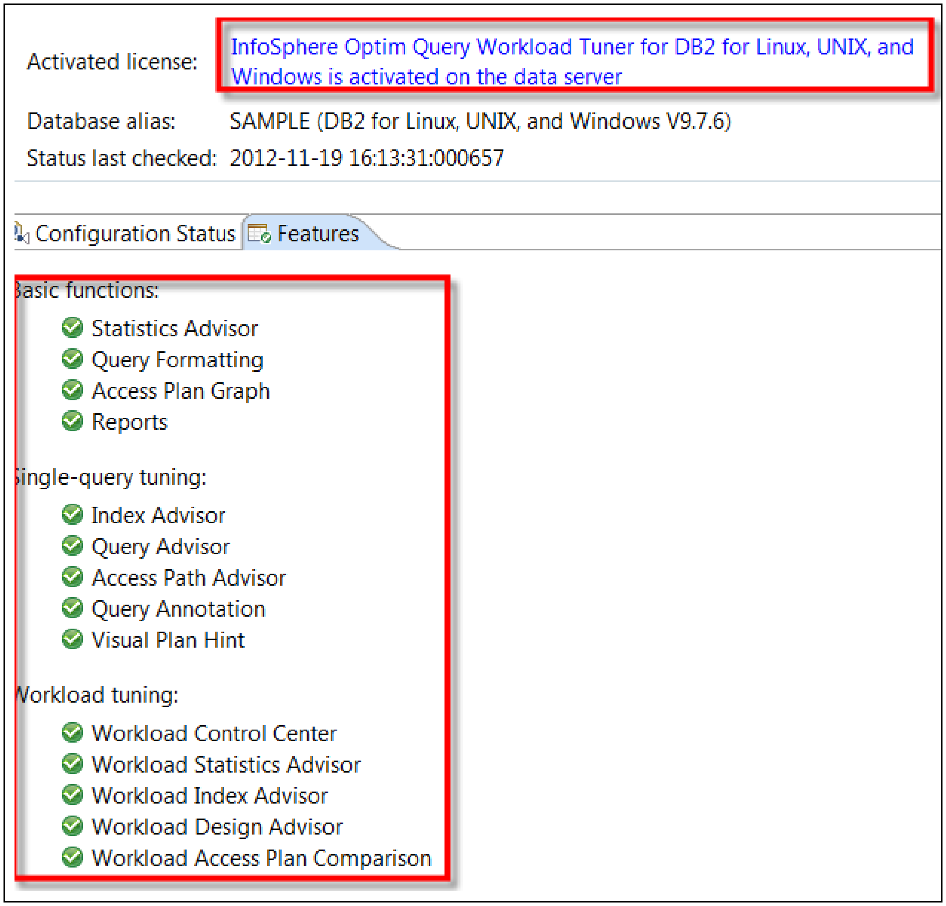

Figure 3-29 Query tuning configuration validation

Configure query tuning with configuration scripts

You can also configure the query tuning of the target database with configuration scripts provided by InfoSphere Optim Query Workload Tuner. Go to the installation directory (for example, `C:\Program Files\IBM\DS3.2.0`), then change to `QueryTunerServerConfig\all_features\LUW` location. There are two executable scripts:

- `enablement_win.bat` is for DB2 on Windows
- `enablement_LinuxUnix.sh` is for DB2 on Linux

Follow the steps described in the script file to configure tuning. For details, see the information center:

http://publib.boulder.ibm.com/infocenter/dstudio/v3r2/topic/com.ibm.datatools.qrytune.installconfig.doc/topics/ioqwtdb2luwscrpts.html

Manage the database EXPLAIN tables and WCC tables

InfoSphere Optim Query Workload Tuner will primarily use two sets of tables in the target database for query tuning: the EXPLAIN tables and WCC tables. The first step to analyze and tune a query is to explain the query, which will store the *explain* information to EXPLAIN tables; for workload tuning, InfoSphere Optim Query Workload Tuner will save all the workload information in WCC tables. The following topics can help you manage those tables so tuning can be a smoother process.

Manage the EXPLAIN tables

Consider the following information:

- The EXPLAIN tables can be created by the EXPLAIN.DDL in the database installation directory, or also can be created with InfoSphere Optim Query Workload Tuner through the configuration UI (Figure 3-28 on page 135).
- The EXPLAIN tables will be created under schema SYSTOOLS.
- When creating the EXPLAIN tables from the InfoSphere Optim Query Workload Tuner UI, you can specify a different table space for EXPLAIN tables.
- During the configuration, InfoSphere Optim Query Workload Tuner checks the format of EXPLAIN tables. If they are not current, perhaps the database was migrated from a previous version but the EXPLAIN tables are not yet migrated. Then, you can use either the **db2exmig** (migrate explain tables command) or migration DDL in InfoSphere Optim Query Workload Tuner:
  - **db2exmig** migrate explain tables command

    http://pic.dhe.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.admin.cmd.doc/doc/r0023722.html
  - Migrating EXPLAIN tables

    http://publib.boulder.ibm.com/infocenter/dstudio/v3r2/topic/com.ibm.datatools.qrytune.installconfig.doc/topics/authmigrateluw.html
- If there is too much explain data in the EXPLAIN tables, you might need to do the house keeping of the EXPLAIN tables. That is, the table space might be almost full or some of the explain information might no longer be necessary. You can delete the data in EXPLAIN_INSTANCE table by the EXPLAIN_TIME column.

For more information about the EXPLAIN tables, see the DB2 information center:

http://pic.dhe.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=%2Fcom.ibm.db2.luw.admin.perf.doc%2Fdoc%2Fc0005137.html

Manage the WCC tables

Consider the following information:

- The WCC tables are created in the target database and used to store all the workload information, including statements, runtime metrics, the explain information, advisors, and comparisons.
- All the WCC tables are created under the SYSTOOLS schema, and managed by InfoSphere Optim Query Workload Tuner.
- During the creation of WCC tables, you can specify the table space. A good approach is to choose a table space with page size greater than 8K for WCC tables.
- When a new version of WCC tables is available, InfoSphere Optim Query Workload Tuner migrates the tables automatically without user interaction.
- You can monitor the WCC tables size and corresponding table space utilization with monitoring product (InfoSphere Optim Performance Manager), and can either increase the table space size or clean up unnecessary workload from the InfoSphere Optim Query Workload Tuner UI.

3.4.7 Deploying InfoSphere Optim Query Workload Tuner in a DMZ environment

InfoSphere Optim Query Workload Tuner is a desktop-based client tool that must have a direct connection from a client desktop to the target database for query tuning and monitoring. In some cases, the user's desktop is in the office area and not allowed to access the production database in the production network. There is typically an area that can access the production database, and can also be accessed by a desktop in the office area. In this case, you can perhaps set up an InfoSphere Optim Query Workload Tuner in that area. Then, when needed, you can go to that area to do the query analysis and tuning.

Or, in the area (server) that has InfoSphere Optim Query Workload Tuner deployed, you can install remote desktop software. Then, you can use the remote desktop connection to that server from the desktop in the office area, and start the query tuning and analysis.

Figure 3-30 shows a proposed solution that deploys the InfoSphere Optim Query Workload Tuner in the area (Linux server), which can connect to the target database in production. There is also a Virtual Network Computing (VNC) server installed. Several users have been created, each with an assigned port for the VNC service. In the office area, the user must have only the assigned ID to connect to the Linux server, and start InfoSphere Optim Query Workload Tuner in the workspace with the configuration for query tuning.

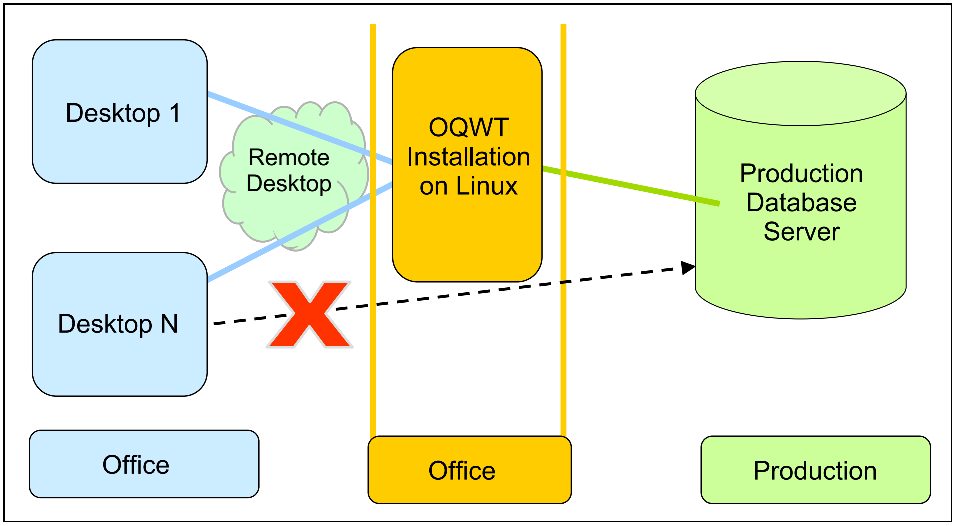

Figure 3-30 Deploy InfoSphere Optim Query Workload Tuner client in DMZ environment

If a new version of InfoSphere Optim Query Workload Tuner is necessary, update the InfoSphere Optim Query Workload Tuner in the Linux server. One user can set up the InfoSphere Optim Query Workload Tuner and related preferences, and then share the tuning preferences with other remote desktop users. Then, all of the same organization can work with the same settings. For more details, see the information center about sharing share the preferences:

http://publib.boulder.ibm.com/infocenter/dstudio/v3r2/topic/com.ibm.datatools.qrytune.project.doc/topics/expprefs.html

3.4.8 Configuration to capture from InfoSphere Optim Performance Manager repository

InfoSphere Optim Query Workload Tuner can integrate with InfoSphere Optim Performance Manger to streamline the performance management of queries. The user can identify the problem queries and applications with InfoSphere Optim Performance Manager monitoring with runtime metrics, and then go to InfoSphere Optim Query Workload Tuner to analyze and tune the performance of those problem queries. After deployment of the advisors, the user can validate the queries and applications performance again. To enable the integration with InfoSphere Performance Manager, some configuration is necessary:

1. Make sure the embedded server is started in the InfoSphere Optim Query Workload Tuner client.

   InfoSphere Optim Query Workload Tuner uses the embedded server to accept the context launch request from the InfoSphere Optim Performance Manager web UI. By default it is *auto start*, and the port for HTTP is 56788 and for HTTPS is 56789. Be sure the ports are not occupied before starting InfoSphere Optim Query Workload Tuner, which accepts only local requests. It is not necessary for these two ports to be past the firewall. You can check the status of the embedded server status on the InfoSphere Optim Query Workload Tuner client toolbar (Figure 3-31). If it is not started, you can start it by clicking the icon that is highlighted in the figure.

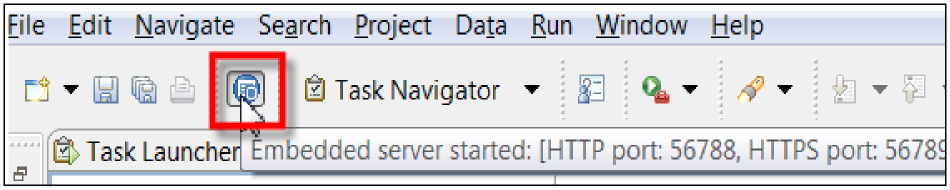

Figure 3-31 Embedded server status and icon menu

2. Configure the connection to the InfoSphere Optim Performance Manager repository database to capture.

   InfoSphere Optim Query Workload Tuner can capture the queries from the InfoSphere Optim Performance Manager repository database, which stores all the historical performance information for the monitored database. Before you can capture from the repository database, configure the InfoSphere Optim Performance Manager profiles in InfoSphere Optim Query Workload Tuner. First, create a database connection to the repository database, and then go to **Windows** and click **Preferences** → **Data Management** → **InfoSphere Optim Performance Manager Profile**, as shown in Figure 3-32 on page 141.

   Specify a name of the profile, choose the database connection you just created for the InfoSphere Optim Performance Manager repository, and click **OK** to finish. The EI access information is optional for InfoSphere Optim Query Workload Tuner integration. The user defined in the database connection to InfoSphere Optim Performance Manager repository must have the required privilege to access the performance data for the monitored database (see 3.4.5, "User security for InfoSphere Optim Query Workload Tuner" on page 132 for details). After the configuration, you can capture queries and related runtime information from the InfoSphere Optim

Performance Manager repository for the database which has been configured and collected by InfoSphere Optim Performance Manager server.

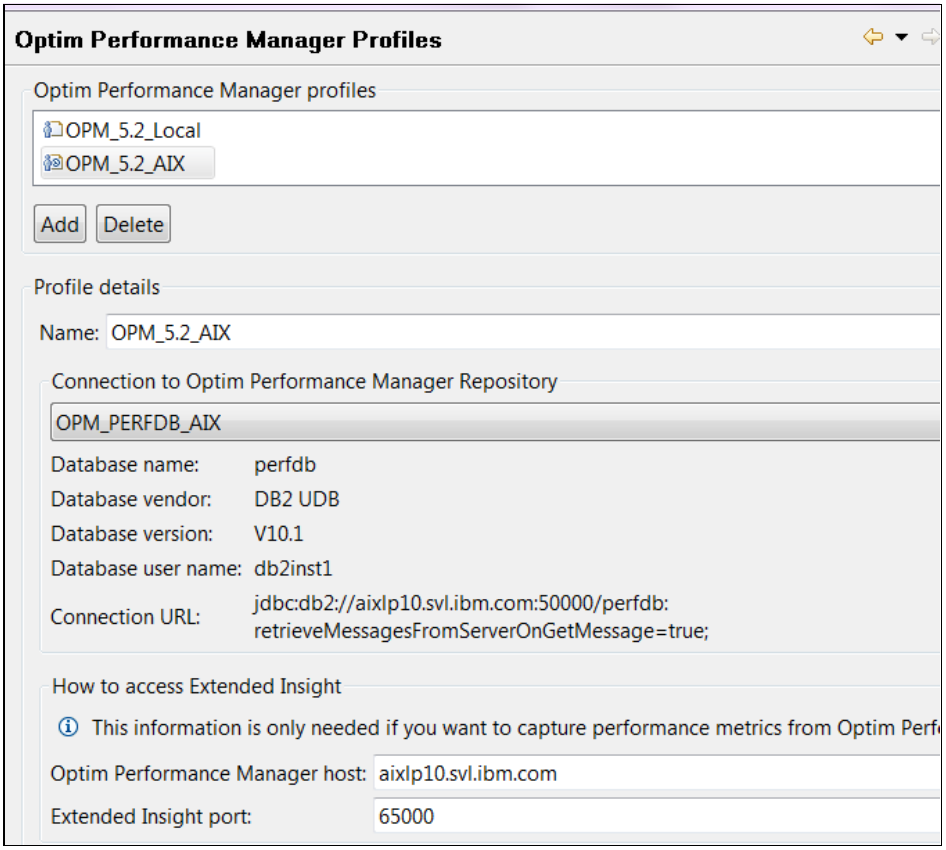

Figure 3-32 Configure the InfoSphere Optim Performance Manager profiles

3.4.9 InfoSphere Optim Query Workload Tuner client preference settings

There are several InfoSphere Optim Query Workload Tuner configurations for, as examples, capture, explain, access plans, advisors, and workloads. These configurations can be put in the preferences on the Optim Query Workload Tuner client.

Global preference

On the InfoSphere Optim Query Workload Tuner client UI, go to **Windows** → **Preferences** → **Data Management** → **Query Tuner**. A set of configuration groups is displayed for DB2 for Linux, UNIX, and Windows, and for DB2 for z/OS, because InfoSphere Optim Query Workload Tuner supports both platforms.

More details are in the information center:

http://publib.boulder.ibm.com/infocenter/dstudio/v3r2/topic/com.ibm.datatools.qrytune.project.doc/topics/glblopts.html

Local advisor preference

You can also set the preference for the advisors for current tuning activities, and you can access them in the query tuner perspective. Select **Set advisor options menu** from the **Query Tuner Workflow Assistant**.

Part 2

Using IBM products to manage performance

This part has an overview of the products available in InfoSphere Optim that can help manage performance in your environment. InfoSphere Optim Performance Manager and InfoSphere Optim Query Workload Tuner are the primary IBM products for addressing performance issues. This part also has a chapter about finding and fixing bottlenecks at the database level.

This part includes the following chapters:

- Chapter 4: This chapter has an overview of InfoSphere Optim Performance Manager. It describes individual product dashboards and reports, and how they can be used to identify, diagnose, prevent, and solve database performance problems.
- Chapter 5: This chapter introduces basic features in InfoSphere Optim Query Workload Tuner, and how to use them to assist the SQL analysis and performance tuning.
- Chapter 6: This chapter describes how you can use InfoSphere Optim Performance Manager to address bottlenecks at the database. It shows the dashboards and reports that display CPU metrics, and how you can use the CPU and related metrics to find the cause of those bottlenecks. It also shows how the Performance Expert Client can give additional insight about CPU utilization.

4

Getting to know InfoSphere Optim Performance Manager

InfoSphere Optim Performance Manager can be used to monitor and manage database and database application performance issues. Use the IBM InfoSphere Optim Performance Manager web console to isolate and analyze typical database performance problems, view a summary of the health of your databases, and drill down for more details.

You can open any dashboard in the IBM InfoSphere Optim Performance Manager web console to drill down to see specific performance information for any one database, or you can open the Health Summary page to see an overview of all of your monitored databases.

InfoSphere Optim Performance Manager uses a web interface and guided workflows to help the user navigate through the series of performance dashboards. Each dashboard surfaces collected metrics for a specific database performance category, such as memory, I/O, locking, and SQL.

With InfoSphere Optim Performance Manager, you can analyze performance data in real time and history by using state-of-the-art DB2 in-memory metrics to access the majority of the performance data on your DB2 V9.7 (or later) databases. This method of data collection provides a more basic monitoring infrastructure than the snapshot monitoring that is still used for V9.5 databases.

In addition, the use of in-memory metrics facilitates real-time data collection and sub-minute refreshes for live problem diagnosis.

InfoSphere Optim Performance Manager now uses a more efficient schema to store performance data in the repository for a DB2 V9.7 (or later) monitored database. This schema is optimized and enables improvement in retrieval times, and facilitates rolling aggregations of collected performance data to create a long-term history of performance data. Retaining long-term performance data is important for proactive tuning, trend detection, and capacity planning. InfoSphere Optim Performance Manager retains performance data at the default sampling interval (for example, 1-minute), 15-minute, 1-hour, and 1-day rolling aggregations.

With InfoSphere Optim Performance Manager V5.2, you can now use the SQL Statements dashboard to analyze statements of stored procedures. Use the controls on the dashboard to see top stored procedures by aggregated execution metrics over all executed statements and nested stored procedures, and drill down into a stored procedure to analyze the executed statements. This feature is supported for DB2 v10 FP2 for Linux, UNIX, and Windows as the monitored database.

InfoSphere Optim Performance Manager provides a set of predefined report templates that you can use to generate reports, either interactively from the web console or in batch mode from command line. You can export these reports in various formats such as PDF, PPT, or XLS. The reports provide valuable information for monitoring, trend detection or establishing baselines. Another feature that this version introduces is the ability to schedule the report generation and retain the created reports in the repository database until the retention period is reached.

Also in this version of InfoSphere Optim Performance Manager is the custom reporting function. User-defined table functions (UDFs) are now provided so that you can extract historical monitoring data from the repository database and create your own custom reports,

InfoSphere Optim Performance Manager provides capabilities to alert you early about arising early, by displaying alerts on the web console, or by sending emails or SNMP traps. A set of predefined health and performance alerts are available, but users can additionally define their own alerts either by defining SQL statements or scripts that check for certain conditions.

In addition to the health and performance alerts about problems on the monitored database, InfoSphere Optim Performance Manager creates operational alerts if it detects issues during its own operation. They are exposed on the InfoSphere Optim Performance Manager web console and alert you about warning and informational messages to help you prevent issues and operational problems that prevent InfoSphere Optim Performance Manager from monitoring the database and saving the monitoring data correctly. Examples are disk space shortages, log space full, communication problems, and DB2 errors.

You can combine any alert with an alert action that InfoSphere Optim Performance Manager executes in the form of an SQL statement or script either on the monitored database or on the repository that executes when a specific alert occurs.

The DB2 Workload Manager user interface in InfoSphere Optim Performance Manager is enhanced to enable prioritization based on the storage group of the data accessed. DB2 10 provides a WLM dispatcher feature that allows you to manage the percentages of CPU resources assigned to service classes. The DB2 Workload Manager user interface in InfoSphere Optim Performance Manager surfaces the options for dispatcher shares and limits in the WLM configuration panels. DB2 10 now supports thresholds based on statement text. Using this feature, you can set threshold limits for a specific query. This is especially useful for troubleshooting intermittent performance problems where a particular query has performance problems that manifest only occasionally. Integration between InfoSphere Optim Performance Manager Extended Edition and Tivoli Composite Application Manager provides a consolidated view of the business transactions across the enterprise, while providing comprehensive detail to help diagnose database-specific areas.

This chapter describes individual product dashboards and reports and how they can be used to identify, diagnose, prevent, and solve database performance problems. The integration of InfoSphere Optim Performance Manager with the Workload Manager tool is described in Chapter 8, “Implementing workload management” on page 329.

4.1 Getting to know the layout

this section introduces the UI features of the web console and shows you how to navigate through the dashboards, wizards, and dialogs. Subsequent sections and chapters offer specific scenarios of how to use a dashboard.

4.1.1 Task launcher

After you install InfoSphere Optim Performance Manager, you can launch the web console from any web browser while your user ID is authorized to the web console. See 2.6, “User security” on page 55.

From Windows, click **Start** → **Programs** → **InfoSphere Optim Performance Manager 5.2** → **Web Console**. The InfoSphere Optim Performance Manager login page opens (Figure 4-1 on page 148).

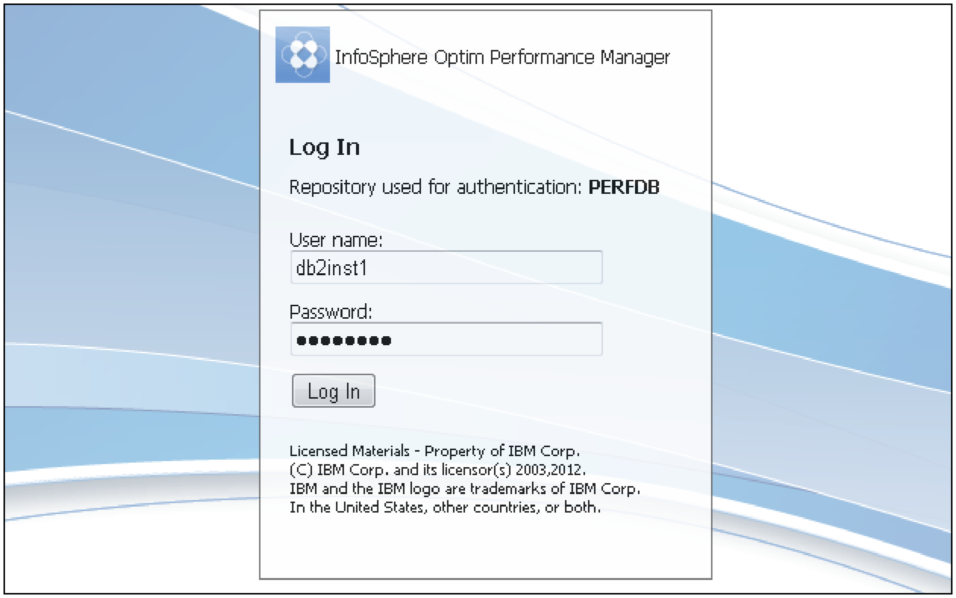

Figure 4-1 Login page

After logging in, the Task Launcher opens (Figure 4-2 on page 149). The Task Launcher describes various tasks that you can do within InfoSphere Optim Performance Manager. such as adding a database for configuration, managing privileges, and so forth.

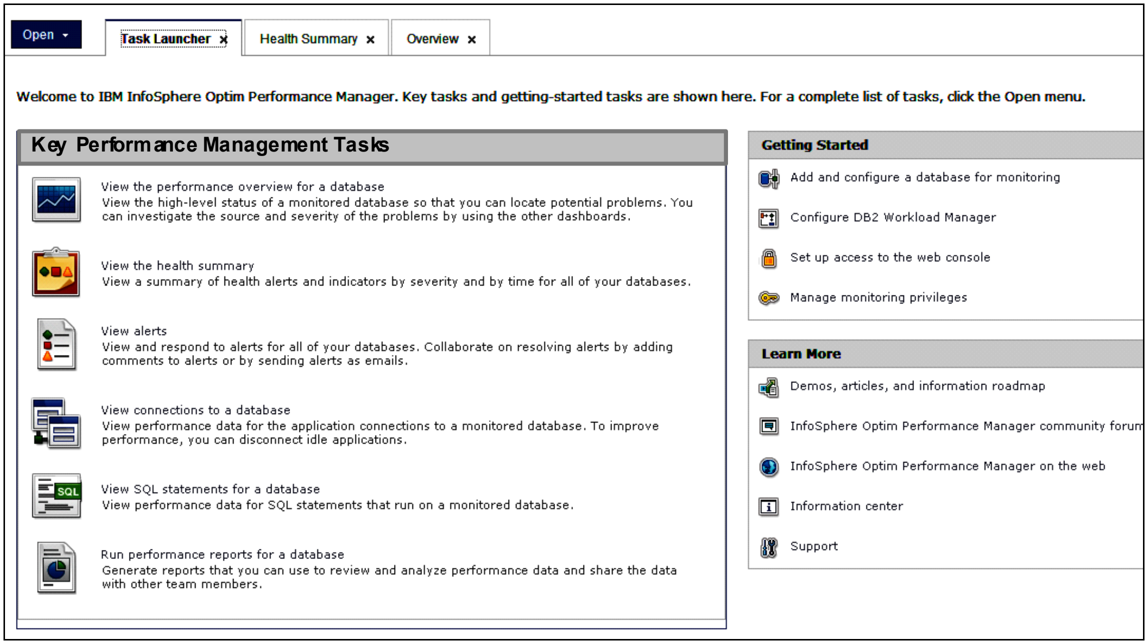

Figure 4-2 Task Launcher

The Open button at the top left corner of the Task Launcher is visible on all dashboards. You can navigate to the other dashboards from this button. If you click **Open**, you see a list of the dashboards (Figure 4-3 on page 150) in InfoSphere Optim Performance Manager, grouped by type of metrics displayed (Health and Performance), tasks (Configuration and Product Setup); at the bottom of the window is a list of administrative tasks for adding and configuring databases (Databases) and scheduling jobs and events (Job Manager and Blackout Event).

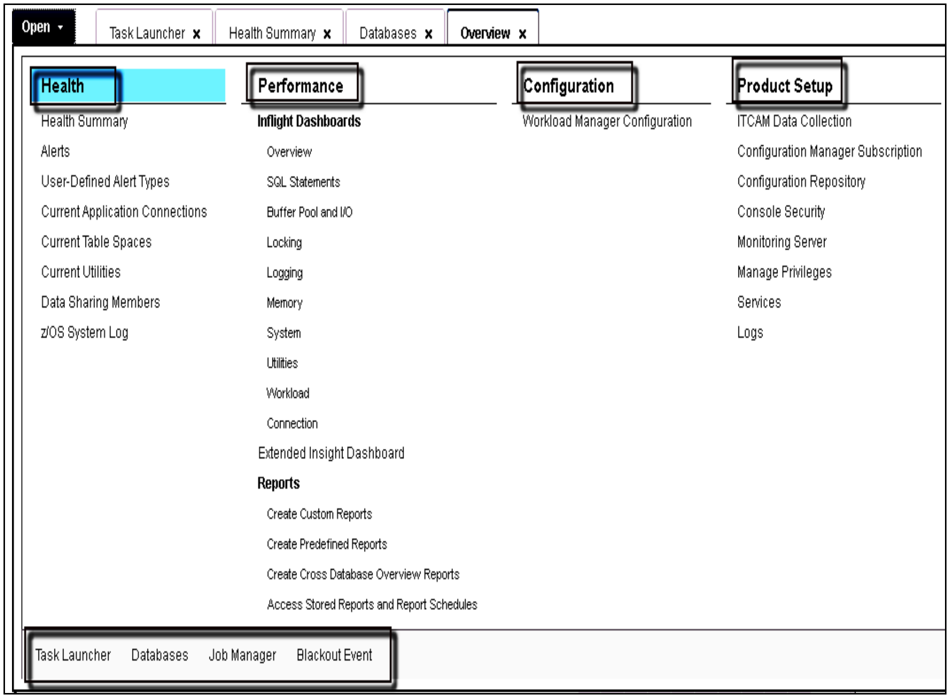

Figure 4-3 List of dashboards

4.1.2 Navigating within a dashboard

You use dashboards in the IBM InfoSphere Optim Performance Manager web console to monitor information about your databases. Most dashboards display performance data for one database for a specific period of time or *duration*. Each dashboard has a time slider that you use to control the time period for the data that is displayed.

> **Tip:** A dashboard shows performance data only if the position of the time slider on the time line contains at least two data points. If you do not see data on the dashboards, you might need to wait until enough data points are available.

The Health Summary and Alerts dashboards display cross database information. In these two dashboards, the default is to display the status of most recent 60 minutes of your databases, as shown in the left side of Figure 4-4 on page 151. You can change the duration by selecting the pull-down arrow and specifying a start and end time.

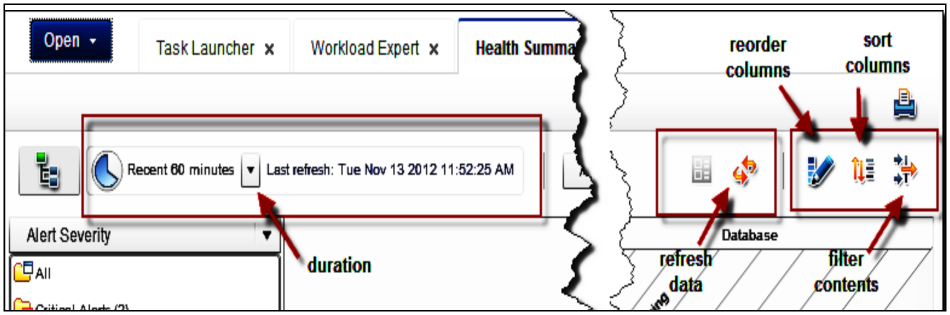

Figure 4-4 Health Summary dashboard

Inflight dashboards display metrics using a *time slider*. Section 4.1.3, "Time controls" has more information about the time slider. The default is the most recent hour or 60 minutes. However, you can use the time slider to display more or less granularity in the metrics.

On the right side of most dashboards is a series of icons for sorting, filtering, and reordering of the columns. There is also an icon for refreshing the data in the dashboard immediately. Hover the mouse over these icons for additional information.

4.1.3 Time controls

Inflight dashboards display metrics using a set of time control variables. The area above each dashboard contains these controls, and is popularly called the *time slider*. You can use the time slider to display more or less data in terms of length of time and granularity.

The major components of the time slider are highlighted in the next two figures:

- Figure 4-5 on page 152 highlights components on the left side of the slider.
- Figure 4-6 on page 152 highlights components on the left side of the slider.

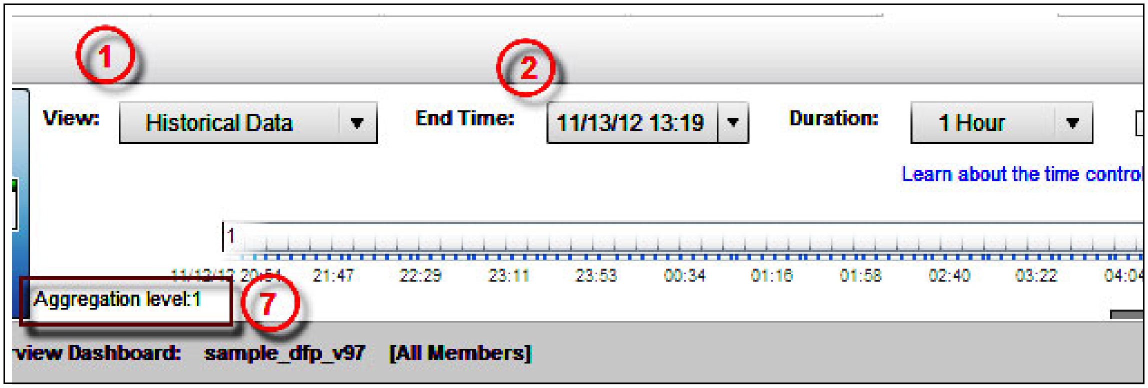

Figure 4-5 Time slider, left side

Those on the right side of the slider are highlighted in Figure 4-6 on page 152.

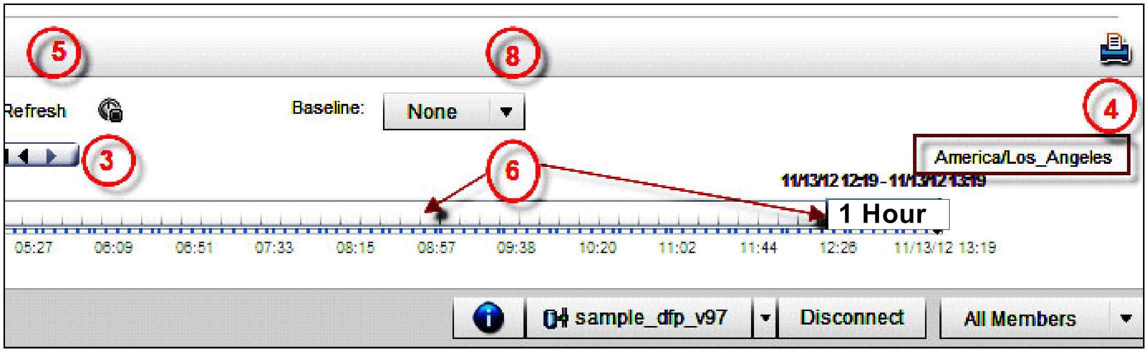

Figure 4-6 Time slider, right side

The components are described in the following list:

- View control

  Click the **down arrow** of the view controller (Figure 4-7 on page 153) to select whether you want to view the performance data in historical or real-time mode.

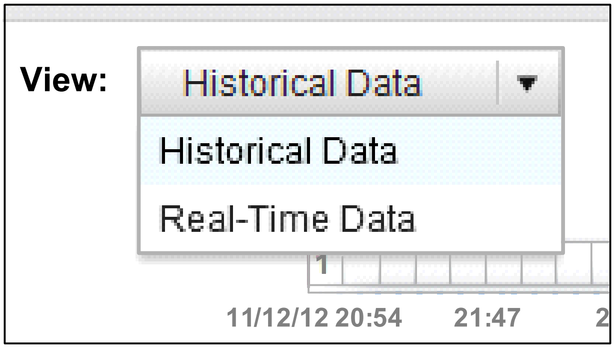

Figure 4-7 View control: Historical versus real-time mode

  In real-time mode, you can use refresh controls to either manually refresh displayed data or specify an autorefresh interval. With each refresh operation, InfoSphere Optim Performance Manager collects new data from the monitored database, and displays the data for the time frame between the previous and current refresh without persisting the data in the repository database

  Note: You must have the *CanMonitorInRealTime* privilege to display performance metrics in real time.

  In historical mode, you can position the time slider to display the performance data for a previous time frame. The length of the time line indicates the amount of historical data that is available. If more data should be available, use the Zoom Out button (described on the next page) until you can see the whole time line. The retention period, which is set when the database was configured for monitoring, controls the amount of available historical data.

- End Time

  Click the arrow to select the end time for the dashboard time slider (Figure 4-8 on page 154).

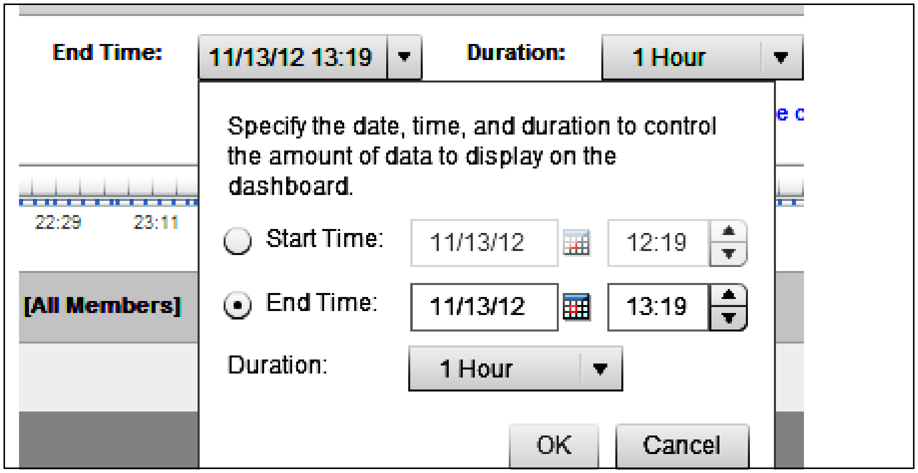

Figure 4-8 Time slider end time

The time slider moves to indicate the requested time frame; the data in the dashboard reflects that time frame.

- Zoom in or zoom out control

  Use the zoom controls to change how much of the time line is shown. For example, if the time line initially shows 60 days of data, you can zoom in so that the time line shows only 10 days of data. In this way, you can more easily manipulate the position of the time slider on the time line. Use the data point controls to move the time slider from one data point to the next. The blue lines at the bottom of the time indicate the points where data was collected. You can move from one data point to the next or previous data point.

- Timezone

  This control displays the time zone of the monitored database whose metrics are on display.

- Duration and automatic refresh

  Figure 4-9 on page 155 shows the duration and refresh controls in the middle part of the time controls area.

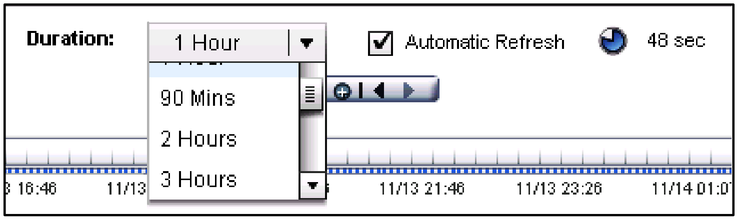

Figure 4-9 Duration, automatic refresh, refresh clock

Duration specifies how much data is shown at one time on the dashboard. The duration is reflected in the time slider (see the next control in this list).

The clock icon indicates the time that remains until the content is refreshed with the next monitoring sample. The content is refreshed based on the sampling rate that was set when the database was configured for monitoring. You can select or clear the Automatic Refresh check box to enable or disable the automatic refreshing of the data on the dashboard. Automatic refresh is useful for viewing the latest performance data.

- Time line and time slider

  The time line (Figure 4-10 on page 155) is initially expanded to show the length of time that performance data has been collected. You can use the zoom in and zoom out controls (see the control on the previous page) to change the span of the time line. Changing the span can make the time slider easier to manipulate.

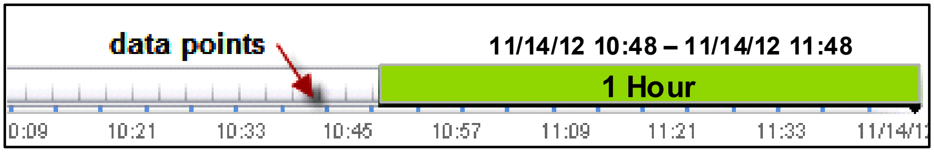

Figure 4-10 Time line showing time slider and data points

Blue data points in the time line indicate times for which monitoring data is available.

Missing lines indicate that no monitored data points are available.

The time slider shows the time interval for the data that is displayed on the dashboard. You can move this control to the left to show older data. When you do so, the color of the slider changes from green to blue. You can control the amount of data that is shown on the dashboard by increasing or decreasing the time interval with the Duration control (see page154).

- Aggregation level

  The aggregation level indicates the level of granularity of the data. The retention period of the aggregation levels, which is set when the database is configured for monitoring, controls the amount of historical data that is available. The system determines the optimal aggregation level for the selected time frame and displays the performance data at that level. Aggregation levels are as follows:

  – Level 1: Each data point represents one minute of data for Extended Insight data and the default sampling interval for inflight data.
  – Level 2: Each data point represents 15 minutes of data.
  – Level 3: Each data point represents one hour of data.
  – Level 4: Each data point represents one day of data.

- Baseline

  With this control (Figure 4-11 on page 156), you can set the time frame that is currently selected in the time control as a baseline. Typically, select a time frame that represents a normal or desirable performance, and ensure that there are few or no gaps in the data, and little or no variation in the data. Best practices for creating a baseline are described in 4.4.1, “Overview Dashboard” on page 179.

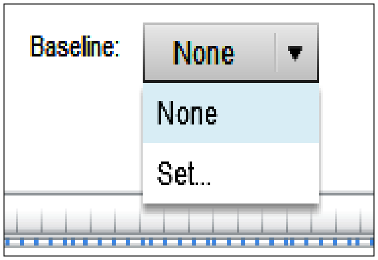

Figure 4-11 Baseline control

4.2 Health summary dashboard and alerting

This section describes various InfoSphere Optim Performance Manager dashboards, which provide information about the overall health of the monitored database.

4.2.1 Health Summary

Start by viewing the Health Summary page. Click **Open** → **Health Summary** to see an overview of the health of all monitored databases. From the Health Summary page you can determine which databases have problems, and you can drill down for more details. Because the Health Summary page is very large and detailed, only a few specific areas are shown in this book as examples of the types of information that are available.The Health Summary shows an overview of the severity and number of alerts of your data sources. The types of alerts that the Health Summary displays depend on the type of data source and the alerts that are configured for each data source.

For example, you can use the group pane (Figure 4-12 on page 157) to select the data sources to view alert severity, host, and port from the default groups. If the custom group feature is available for your product, you can also create custom groups that are shared with all users of the web console.

Figure 4-12 Health Summary page group pane

You can use the column controls (Figure 4-13 on page 158) to control which columns to display in the grid, move columns around, and filter and sort the data in the grid as you want. You can also click the **column headings** to sort the data in the grid according to any column value.

Figure 4-13 Health Summary page column control

Use the time control (Figure 4-14 on page 158) to set the time period for the health data that is displayed. You can view recent data that is refreshed at configurable intervals, or view historical data for a specific time period. For recent data, the refresh rate controls the rate at which the Health Summary page is updated. For historic data, specify the history time frame in the time zone of your web browser. The alert data is collected at a rate that you can set individually for each monitored database.

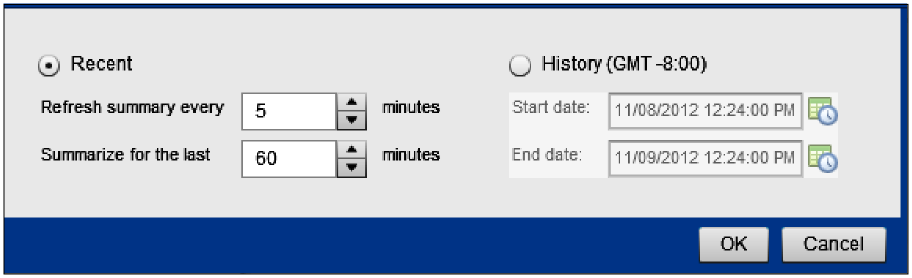

Figure 4-14 Health summary page time control

Each column in the Health Summary, as shown in Figure 4-15 on page 158, represents an alert category contributed by a data source that shares the Health Summary interface.

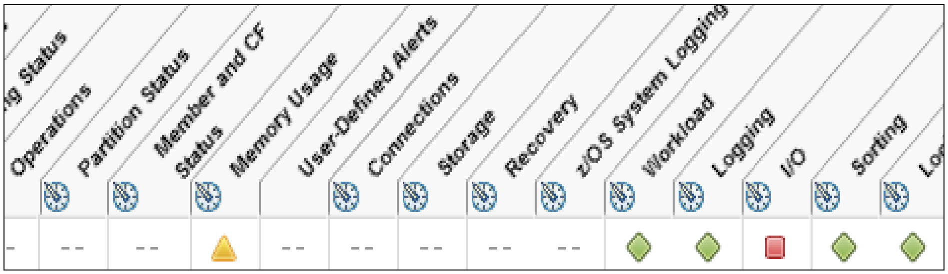

Figure 4-15 Alert categories and alert icons

To understand database performance, you monitor a set of key performance indicators for your connected DB2 data server databases. Potential areas for concern are identified by critical alerts (red squares) and warnings (yellow triangles) on the performance manager dashboards and windows.

Alert icons (Figure 4-15 on page 158) in the grid reflect alert severity for each alert category and data source. A double dash (- -) in a cell indicates that no alerts were issued in the selected time frame. Each icon represents a summation of one or more individual alerts that were encountered for the selected duration of that specific database. For example, if you had two storage alerts for a database, one critical and one warning, then the alert summary icon will identify this as critical. When you click the cell, you can drill down to the individual alerts themselves that detail the problems, including any appropriate actions you should take.

The selected time frame governs the alerts that are shown. The per-database, per-category summarizations include all of the alerts, and also highlights only those alerts that occurred in that time period. The Data Server Status category is special, because its summary represents the latest state of the database at the end of the time period.

The Health Summary and Alerts list provides an overview of alerts across all monitored databases. The Health Summary shows active alerts by category. The Alerts list also provides information about active alerts, but also provides information about all the alerts that occurred during the monitored period.

Click an alert icon in the grid to see a list of all of the alerts of that specific category for that data source. You can then select a specific alert to troubleshoot.

To better understand or resolve the cause of an alert, open the appropriate dashboard by clicking **Open Dashboard**.

4.2.2 Alert list and filtering

The Alert List page displays the most recently triggered alerts for your monitored databases. You can enable or disable alerts and change thresholds for alerts on this page, and you can configure emails or SNMP traps to be sent automatically so that you receive immediate notification if a critical problem occurs.

To display the Alert List page (Figure 4-16 on page 160 shows part of the page), from the InfoSphere Optim Performance Manager console, click **Open** → **Health** → **Alerts**. The Alert List page is very large and contains much information. Therefore, only a portion of it is shown here to help you become familiar with the report and still make it somewhat readable.

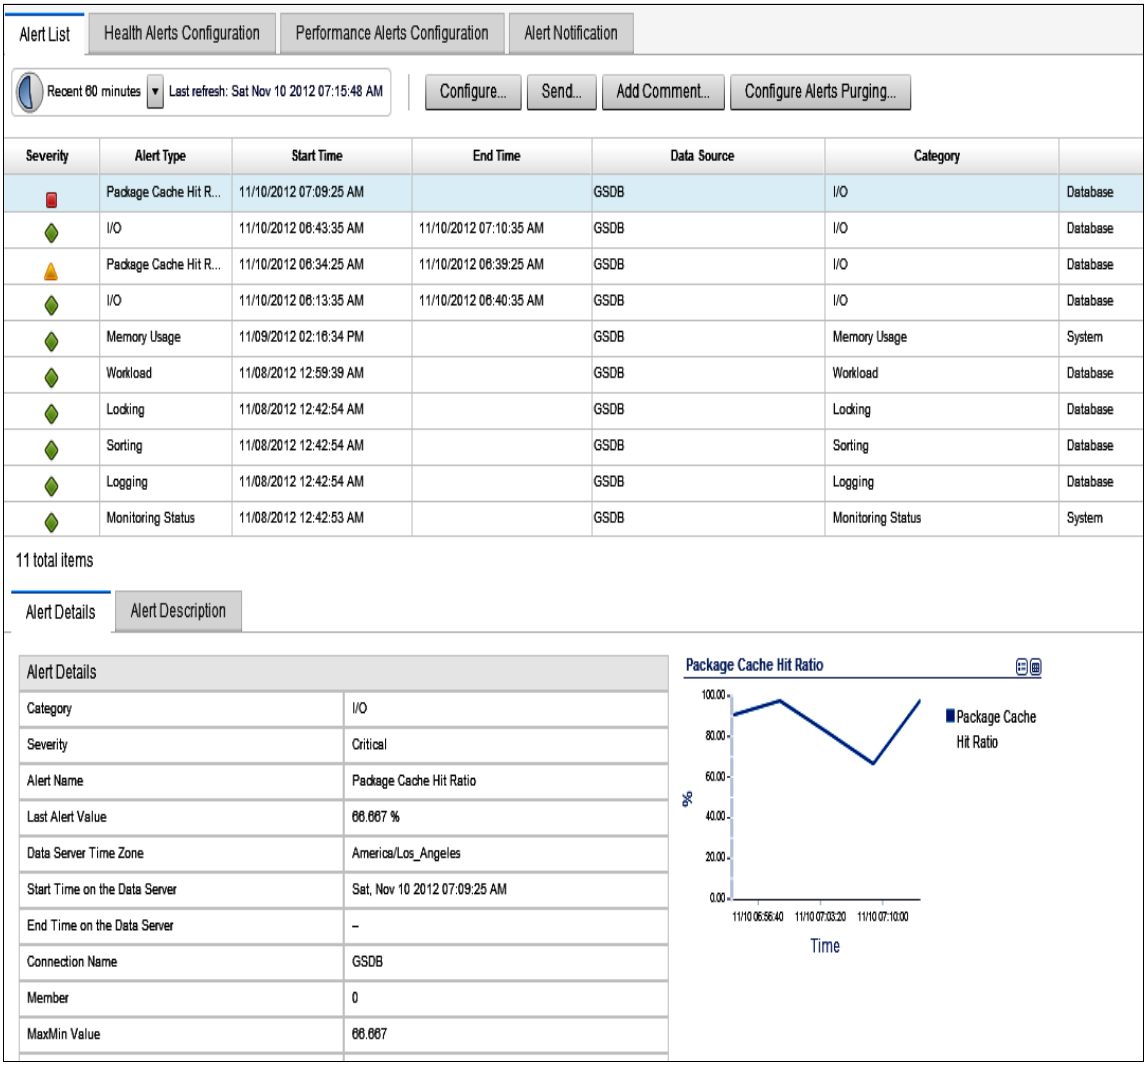

Figure 4-16 Alert List

The types of alerts that the Alert List displays depend on the type of data source and the alerts that are configured for each data source.

Alerts are enabled with default threshold values, which are set when you configure a database for monitoring. The thresholds determine when an alert is triggered. The key performance indicators are checked periodically by using a default sampling rate. If a threshold is breached, an alert is generated and an indicator is displayed on the Health Summary, the Alerts list, and the associated dashboard. You can customize the alert thresholds for your environment.

To view detailed information about an alert, click the alert in the grid (Figure 4-16 on page 160) and view the details (Figure 4-17 on page 161). This information is useful to determine the threshold-exceeded values, and the graph displays the timeline from the start to the current state of the alert.

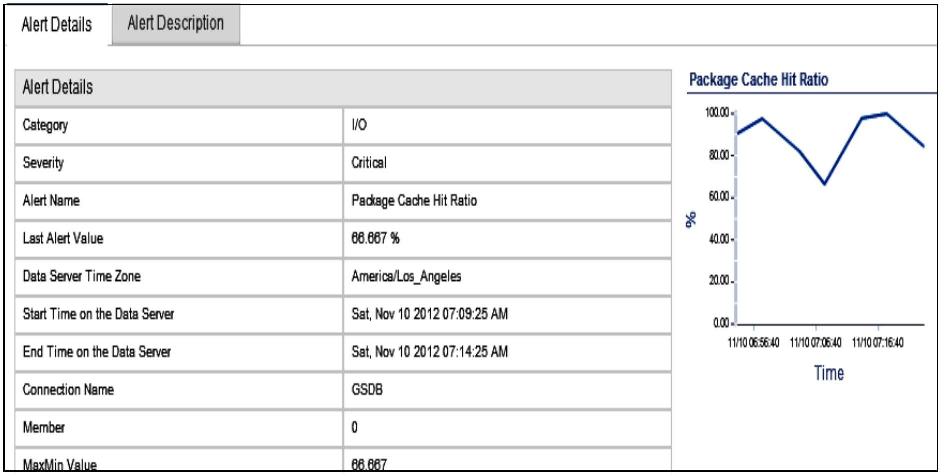

Figure 4-17 Alert Details

To view additional details and suggestions for troubleshooting the alert, from the Actions pane you can navigate to the inflight dashboards by clicking the **Navigate to** link to analyze further (Figure 4-18 on page 161).

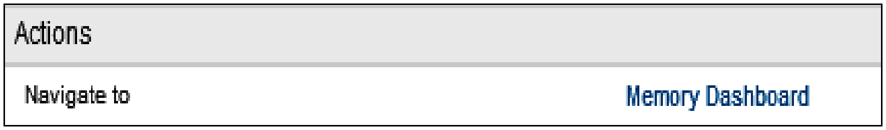

Figure 4-18 Actions

If you want to share details about an alert with system administrators or other users, you can send the alert in an email or add comments to the alert:

- To send a link to an alert in an email, select the alert and click **Send**. The email contains a link to the alert in Health Summary.

  Email: Email communication requires that the email service is configured with a valid SMTP host name and port.

- To add comments to an alert, select the alert and click **Add Comment**. Comments added to alerts are visible to all users.

InfoSphere Optim Performance Manager alerts can be classified into health and performance alert categories:

- Health alerts

  Health alerts are triggered by the Data Studio Health Monitor component of InfoSphere Optim Performance Manager. They provide database health status information, such as data server status, storage space state and utilization, and HADR state. To configure health alerts, select **Open** → **Task Manager** → **Health Alerts Configuration**. Figure 4-19 on page 162 shows a partial sample view of the Health Alerts Configuration panel.

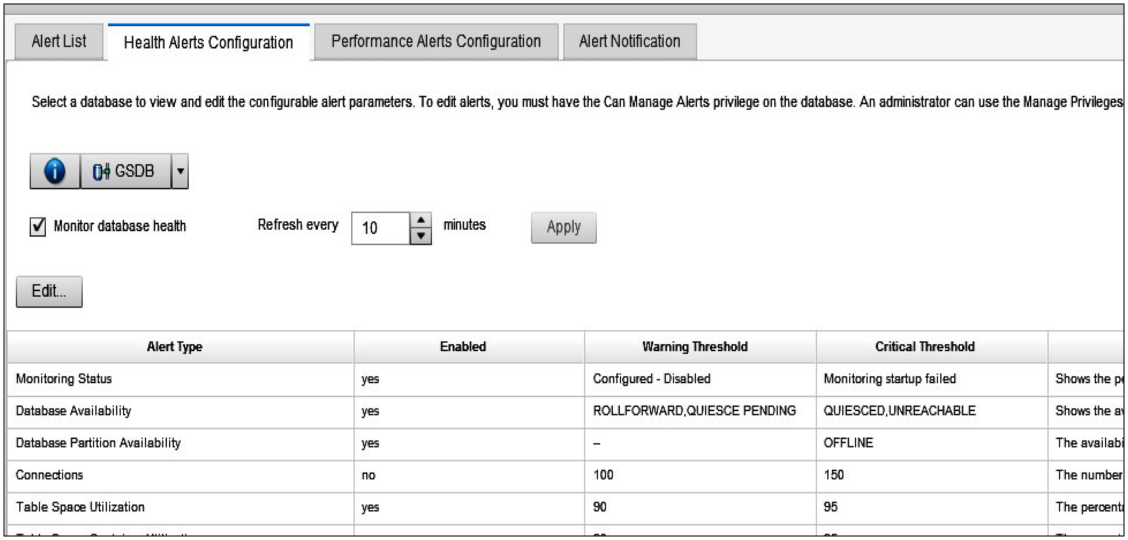

Figure 4-19 Health Alerts Configuration

> **Important:** This Health Monitoring does not require a database to be configured for monitoring. After a database connection is added in the Database page these Health Alerts take effect.

- Performance alerts

  Performance alerts represent database key performance indicators from various categories such as memory, I/O, locking, logging, and SQL. Examples of performance alerts are Rows read per fetched row, Package cache hit ratio, Currently waiting applications, and Sort overflows. To configure performance alerts, select **Open** → **Alerts** → **Performance Alert Configuration**.

> **Thresholds:** For partitioned DB2 databases, you can set thresholds to various partition roles (catalog, data, coordinator). This requires that you have assigned partition roles to individual database partitions when you have configured your database for monitoring.

4.2.3 Alert configuration

From the Alert List, you can configure the alert settings for a specific alert category and data source by selecting the alert and clicking **Configure Alerts**. This configuration is applicable only for performance alerts.

The alert configuration panel allows you to configure alert settings in two ways:

- By Alert

  You can configure individual alert settings for one or more databases. You can modify, enable, or disable setting and also warning and critical threshold levels for this particular alert. See Figure 4-20 on page 163.

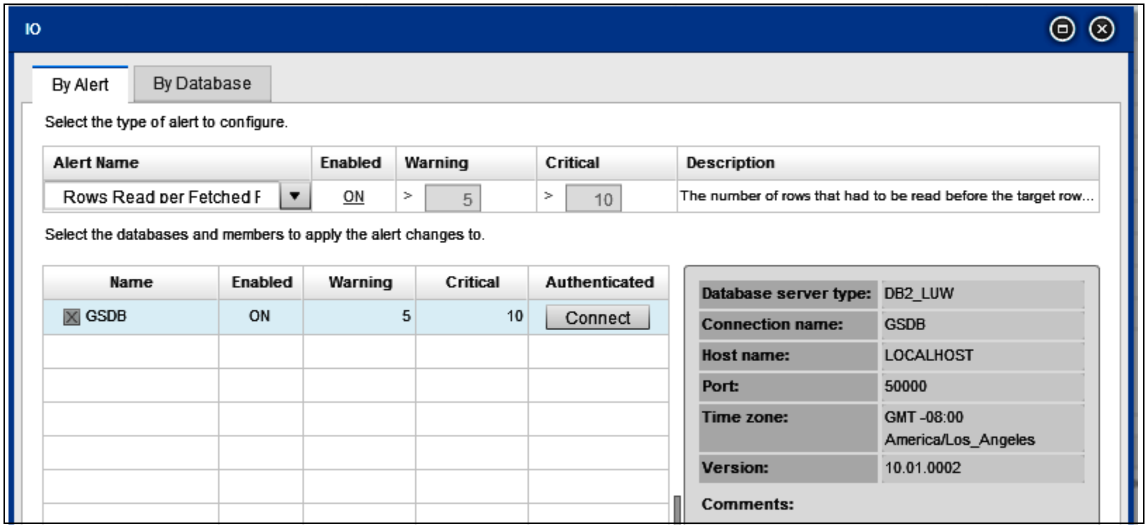

Figure 4-20 Alert configuration panel - By Alert

- By Database

  You can configure settings for any alert for a particular database. Use **Copy Settings** to replicate alert definitions to other monitored databases.

  Figure 4-21 on page 164 shows the Alert configuration panel, By Database.

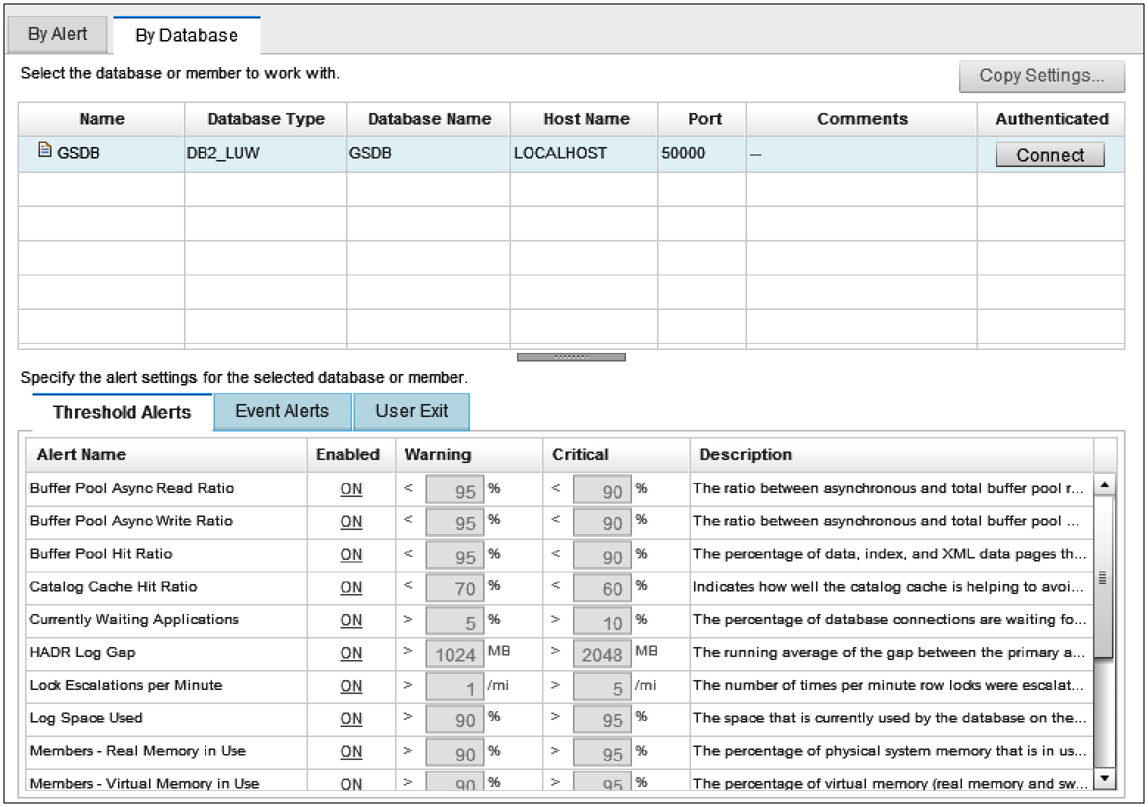

Figure 4-21 Alert configuration panel, By Database

InfoSphere Optim Performance Manager can automatically purge alerts which are no longer important. Select **Open** → **Health** → **Alerts** → **Configure Alerts Purging** to specify time interval after which alerts will be deleted. To edit these settings, you must have the *Can Manage Alerts* privilege on the database.

Figure 4-22 on page 165 shows the Configure Alerts Purging panel.

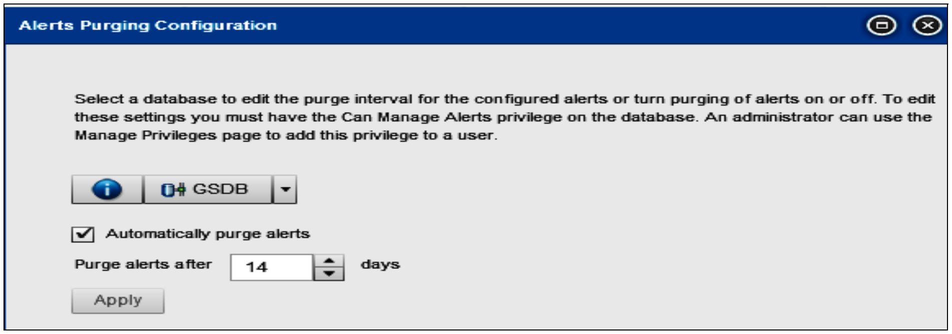

Figure 4-22 Configure Alerts Purging

4.2.4 Alert notification (email and SNMP)

After you have configured alerts for your databases, you can set up alert notifications for individual alerts per monitored database. You can also choose to subscribe to alert notifications sent through email, or sent through SNMP to another tool. With this feature, you can wait until you receive notifications about problems, or interesting events regarding your database.This allows InfoSphere Optim Performance Manager to send alert notifications to various individuals based on their responsibility for a particular monitored database or an alert category within the monitored database.

You can define alert notification parameters such as email addresses, SNMP trap generation, and alert frequency. To configure alert notifications, click **Task Manager** → **Alert Notification.** Figure 4-23 on page 166 shows the partial view of the Alert notification panel.

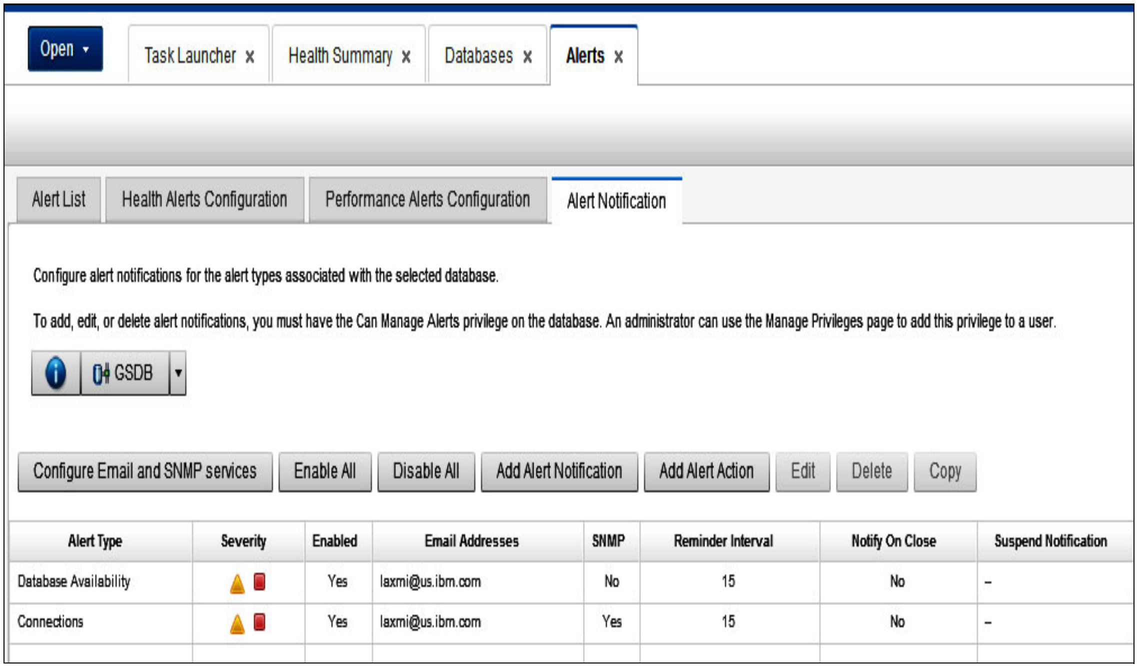

Figure 4-23 Alert Notification panel

InfoSphere Optim Performance Manager can use two alert notification methods:

- Email
- SNMP trap

To configure these notification methods, click **Configure Email and SNMP services** and provide details for each of the methods, including these items:

- Address and port number of the outbound SMTP mail server that will be used to send email notifications
- SNMP Management server that will receive the SNMP traps

After you have configured alert notification services, you can proceed to add and configure individual alert notification details. From the Alert Notification panel click **Add Alert Notification**.

Figure 4-24 on page 167 shows the panel to add new alert notification. On this panel, you can specify the following information:

- Alert severity type: Can be Warning or Critical.
- Email address: Specify email addresses of the recipients of the alert notification.

- SNMP trap generation: InfoSphere Optim Performance Manager alerts can generate SNMP traps, which can be forwarded to SNMP manager.
- Suspend Notification period: Time interval during which alert notification is disabled.

Select the **Enabled** check box to activate the alert notification for the alert.

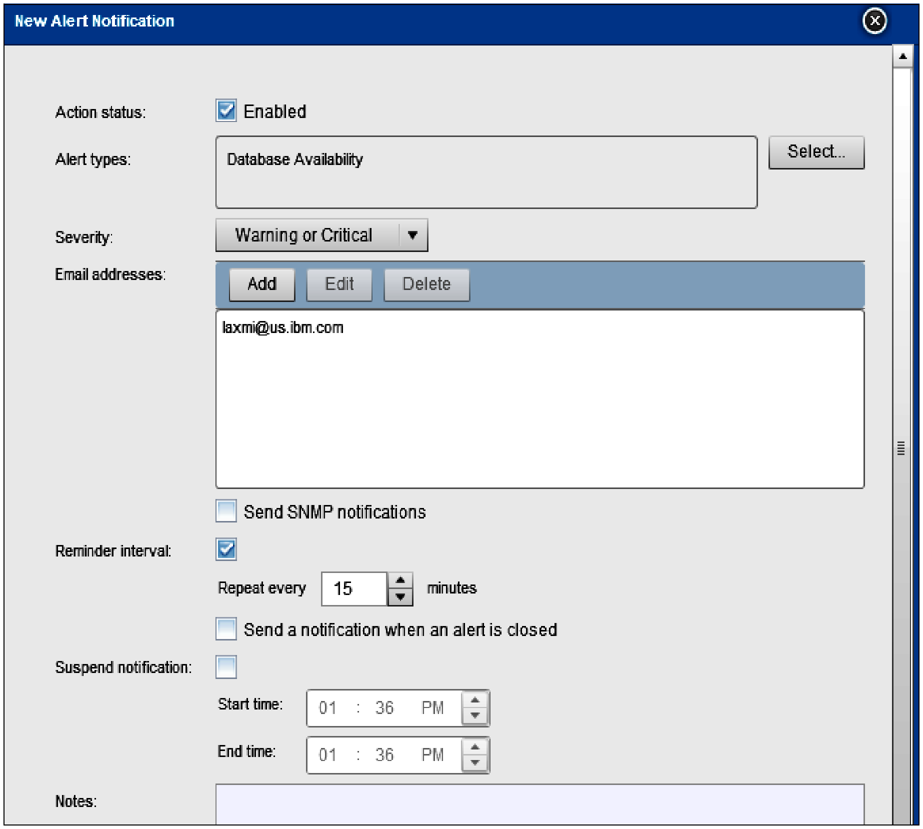

Figure 4-24 Adding new alert notification

4.2.5 Alert actions

After you configure alerts for your databases, you can set up alert actions consisting of jobs to be executed after the alert occurs. A job is a scheduled task that you want to execute as soon as the alert is encountered. To associate an alert with a job to be executed, first create the job. After a job is created you can select the job ID from the data grid. The job ID is automatically assigned to the job when you create a job. The job consists of executing a script. To run a script on a database, ensure that the user ID that runs the script has the required privileges to run the commands that are included in the script on the database. For Executable/Shell script job types, the user ID must also have permission to connect to the database server by using SSH.

To associate a job to be executed, click **Task Manager** → **Alert Notification** → **Add Alert Action.** The association opens the New Alert Action pane (Figure 4-25 on page 168).

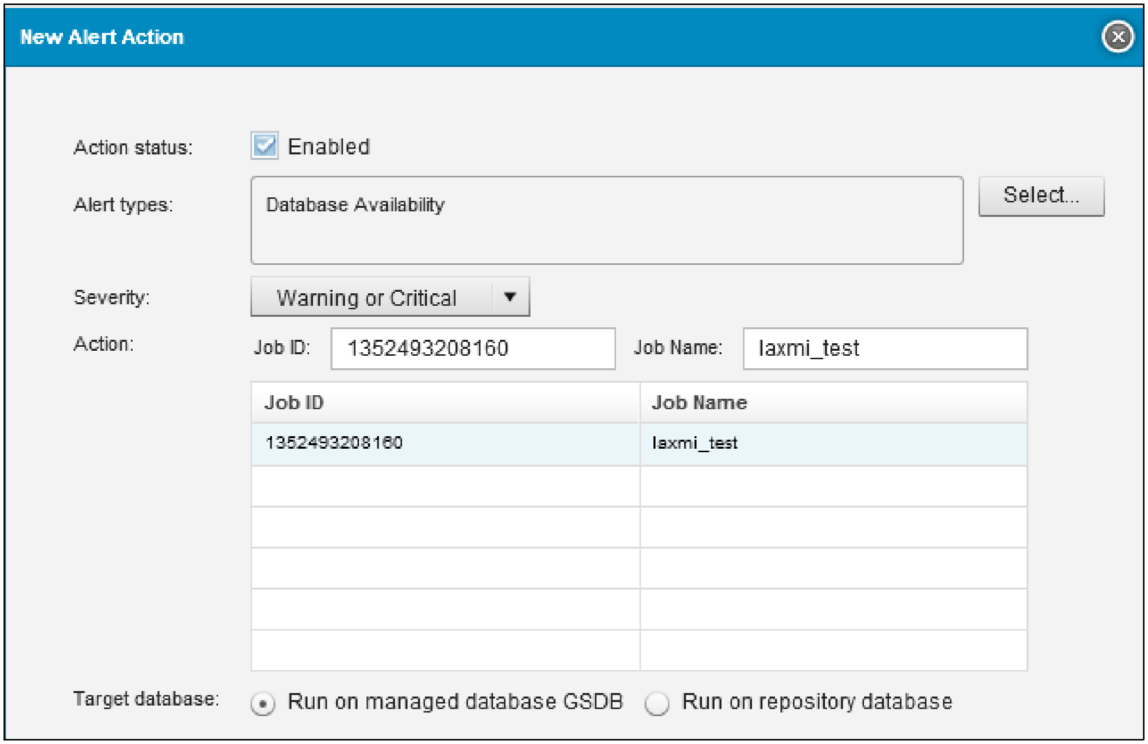

Figure 4-25 Add alert action

Provide the following details to complete the association process:

1. Select the Alert Type that you want to associate the job with.
2. Select the job you want to use. A list of existing jobs is displayed in the grid. Highlight the job on the grid that you want to add as an action to the alert notification.
3. Specify whether this job must be run on the monitored database or on the repository database.

Figure 4-26 on page 169 shows a partial view of the Alerts Notification Panel; a job has been associated with the Database Availability Alert.

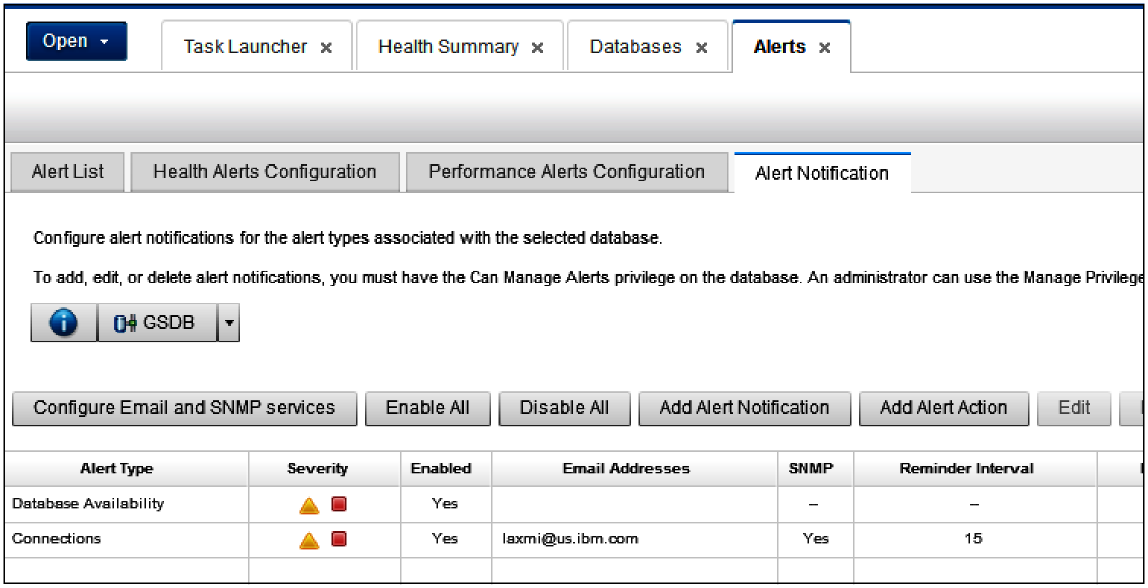

Figure 4-26 Alerts Notification pane with jobs

4.2.6 User-defined alerts

InfoSphere Optim Performance Manager includes a set of predefined alerts. However, for your requirements, they might be limited and the current criteria might not be sufficient. Therefore, user-defined alerts provide flexibility and are not limited to the alerts the tool provides. Several users have their own sets of scripts or utilities that they want to run periodically to check for specific events and act on them.

Users can create user-defined alerts by providing their own custom scripts as described in the following steps:

1. From the Open menu, click **Health** → **User-Defined Alert Types** → **Add Alert Type**.

   The Add Alert Type dialog opens (Figure 4-27 on page 170).

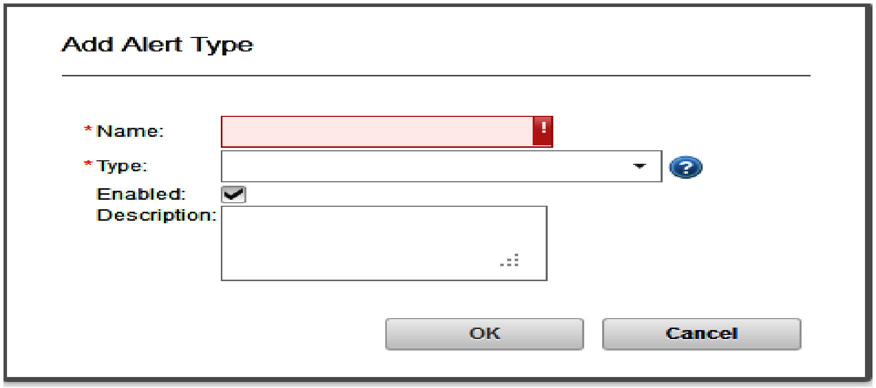

Figure 4-27 Add Alert Type

2. Enter the following information:

   – Specify a name for the user-defined alert type.

   – Click the **Type** pull-down menu to see the types. The supported types are SQL-only and Shell/Executable scripts:

     • SQL-only script: These are run on the database. To run the job, the job manager connects to the database and runs the SQL commands that are included in the job script directly on the database.

     • Executable/Shell script: These are run on the database server. To run the script, the system logs in to the database server by using SSH as the user ID that is defined in the database connection, and then runs shell commands directly on the server.

> **Important:** Specify a script that can run on the database or the database server. The script type sets the connection method to the database or database server

– The script must return an exit value to indicate the severity of the alert to be generated. The valid return values are as follows:

  - 0 = Normal condition (that is, no alert to be generated)
  - -1 = Generate a warning alert
  - -2 = Generate a critical alert

After you provide the details, click **OK**; the Script pane opens.

3. Add User-Defined Alert Types. Next, add the contents of the Alert script as shown in Figure 4-28 on page 171, and click **Save All**.

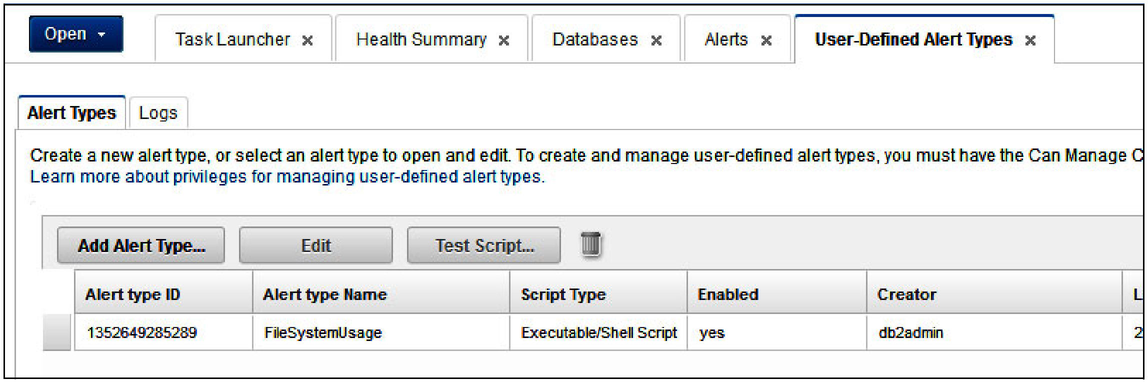

Figure 4-28 Add Script

4. Test the script. Click **Test Script**, and select one or more databases on which you want to test the script. The web console does not verify that the scripts that you enter are valid. If you choose to run the script on multiple databases, the default user ID stored in the connections profile is used to execute the script. If you choose to execute the script on a single database, you can provide the user ID and password and test connection or you can choose to use the default user ID which is the one stored in the database connection profile.

> **Important:** The web console user ID that tests the script and therefore initiates executing the script on the monitored databases or systems requires isDatabaseOwner privilege for these databases. Use the **Manage Privileges** page to configure the correct privileges. To manage the alerts, the console user ID needs Can Manage Custom Alert Types privilege issued against the Repository database.

5. Click **Test Connection** to run the script on a database or database server, or use other methods to verify that the script is correct and that it produces the expected results before you schedule the user-defined alert, shown in Figure 4-29 on page 172.

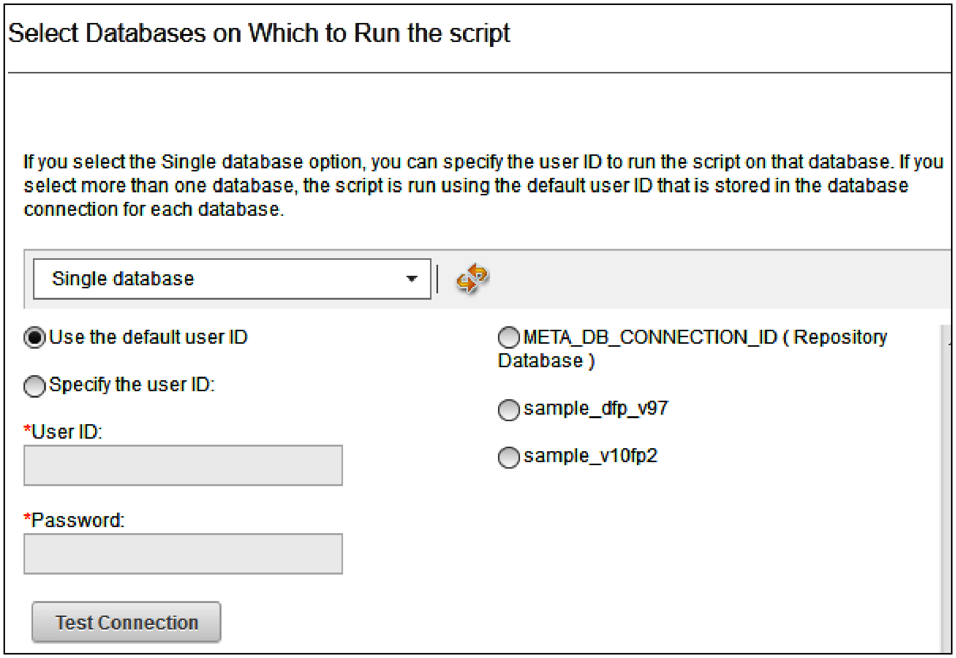

Figure 4-29 Test Script

> **Important:** To run Executable/Shell scripts on a database, ensure that the user ID that is used to run the script has permission to log in to the database server by using SSH. If you want to use DB2 commands to generate an alert, you can run the DB2 commands by using the Executable/Shell script type for the user-defined alert type. Ensure that your DB2 environment is set up and initialized before you run the script. For DB2 SQL scripts, ensure the user ID has execute privilege on the tasks.

6. View execution logs. You can view execution details in the Logs grid.You can also view the entire execution log in a browser by clicking **View log in browser**. You can rerun the script from the Logs grid and can also cancel the script. The Failed and Success records are retained for three days by default.
7. Schedule user-defined scripts. After the scripts are tested you can schedule the scripts to be executed by clicking **Schedules** → **Add Schedule**. Multiple schedules can be associated with a single script. You may select one or more

databases for the script execution by clicking the **Database** tab. Credentials from the database connection defined will be used for SSH and script execution. If you schedule a script to run on more than one database, the script runs on each database with the user ID that is stored in the database connection for that database. An example is shown in Figure 4-30 on page 173.

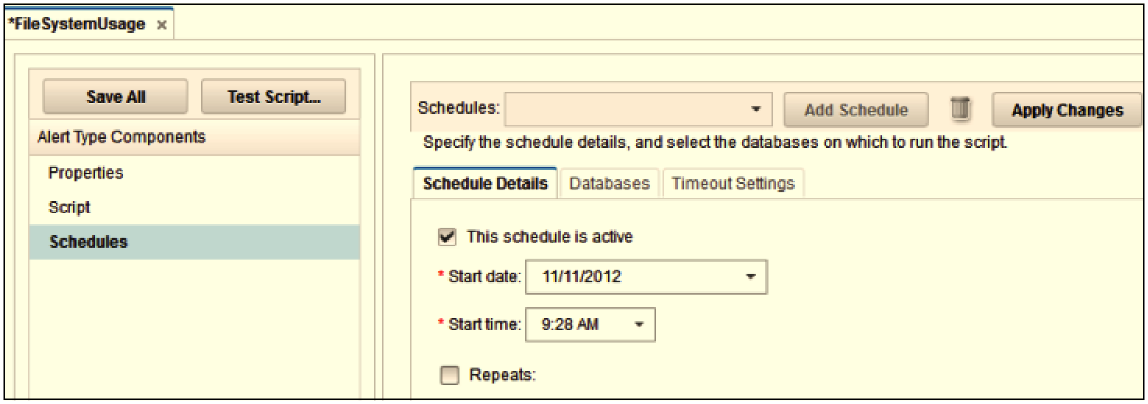

Figure 4-30 Schedule scripts

8. Use timeout settings. Figure 4-31 on page 173 shows how to choose an action to perform if the script duration exceeds a timeout period. The action can be that the script is canceled if its execution exceeds a specified duration, or the action can be to generate a specified alert (Critical, Warning, or Informational) if the execution exceeds a specified duration.

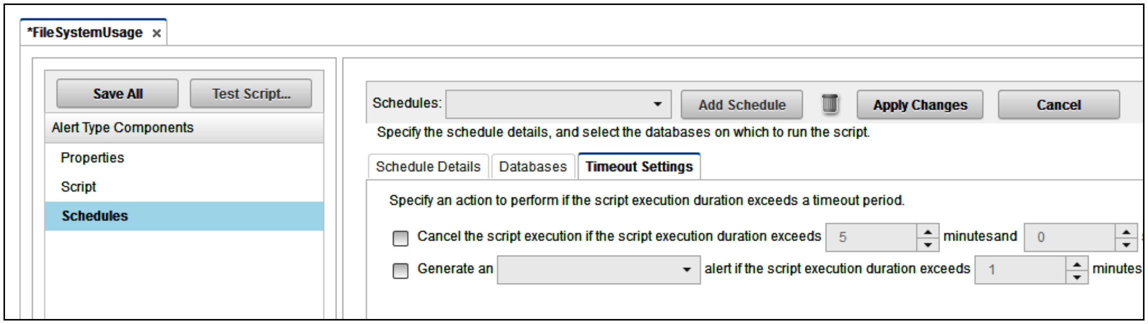

Figure 4-31 Timeout Settings

9. Consider privileges for user-defined alert types. If you enforced privileges on the Manage Privilege panel for the repository database, then InfoSphere Optim Performance Manager web console users require **Can Manage Custom Alert Type** privilege to create and manage user-defined alerts on the repository database. Otherwise, all web console users are allowed to do that.

   Additionally only web console users with isDatabaseOwner privilege for a monitored database are allowed to schedule user-defined alerts and therefore initiate the execution of scripts on the monitored database.

   To run the user-defined alert type scripts on the monitored database, the provided user ID for the monitored database requires correct permissions to perform the tasks in the script. Users can select one or more databases against which to run the script. If you select more than one database, then the user ID stored with the database connection for each database is used to run the script. SSH must be available and running on the monitored database.

10. Click **Add Schedule**; a schedule is created.

The left side of the Health Summary page is shown in Figure 4-32 on page 174.

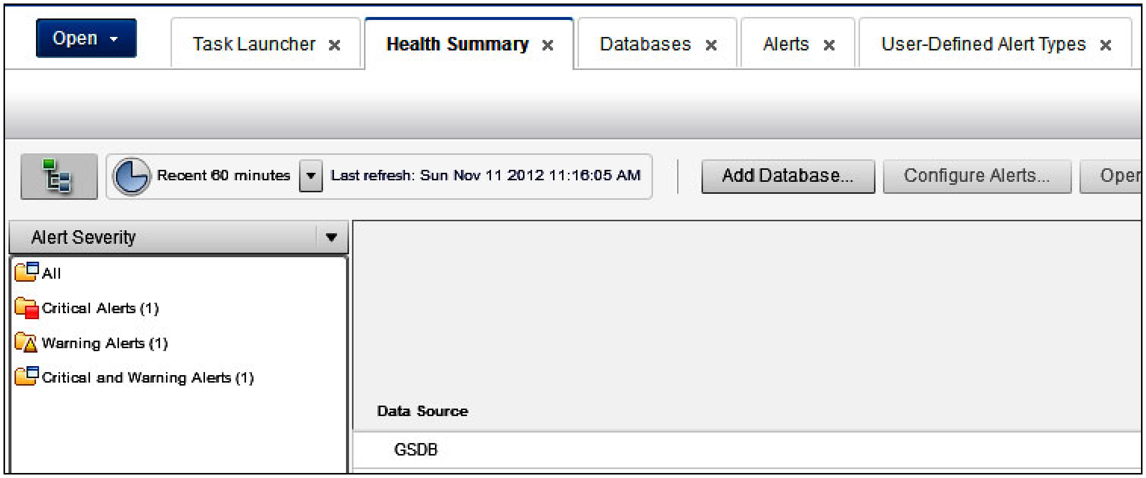

Figure 4-32 Left side of Health Summary page showing user-defined alert

You can view the user-defined alert column on the right side of the Health Summary page, shown in Figure 4-33 on page 175.

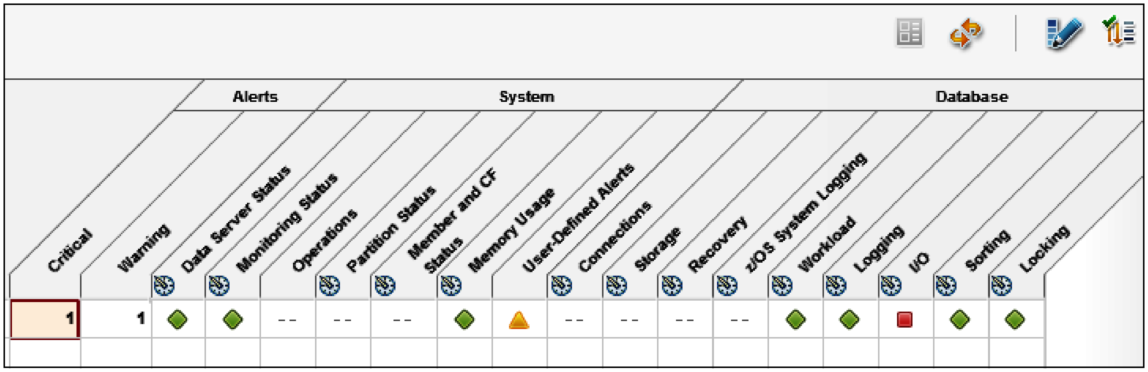

Figure 4-33 Right side of Health Summary page showing user-defined alert

List of alerts

The following sample list is of user-defined alert type scripts to handle various alerting scenarios:

- Generate an alert if file system utilization is high
- Generate an alert if amount of free space in memory is low
- Generate an alert for database backup
- Generate an alert for diagnostic records

4.2.7 Operational alerts for InfoSphere Optim Performance Manager

Operational alerts are created if InfoSphere Optim Performance Manager detects issues during its operation that might lead to incomplete or non-continuous collection of monitoring data. These issues are surfaced in the web console alert list and can therefore be detected and solved more quickly. The alert details and descriptions help you to solve the issue yourself or to give the details to the IBM support team. Operational Alerts can be classified into three kinds:

- Monitored database operation alerts
- Repository Server operation alerts
- Repository database operation alerts

The Operations column is displayed in the Health Summary page (Figure 4-34 on page 176).

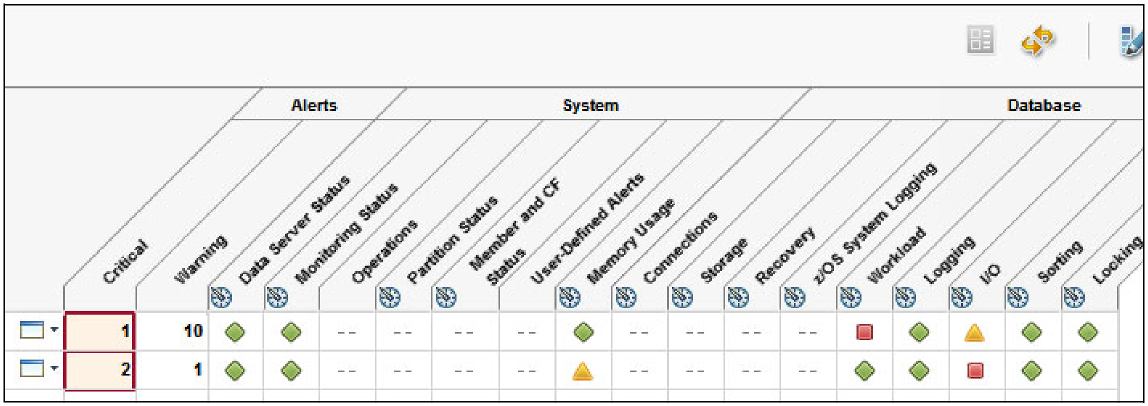

Figure 4-34 Health Summary page - Operations Alerts

4.3 Real-time monitoring and history monitoring

InfoSphere Optim Performance Manager supports real-time monitoring of your configured DB2 9.7 (or later) databases to see current data for problem determination. If you are monitoring DB2 V9.7 or later for Linux, UNIX, or Windows databases using the new in-memory metrics data collection method, you can view data collected at any point in time, whether in real-time, yesterday, last week, or last year. All dashboards can display the full range of collected monitoring data.

You can switch to real-time monitoring as highlighted on the left side of the dashboard (Figure 4-35 on page 177) to analyze current performance problems that require you to see the current activity on your database in sub-minute refresh rates. Manual refreshes and also automatic refreshes are possible on each inflight dashboard by using the time slider selection control. With each refresh, new monitoring data is collected from the monitored database and displayed but not saved in the repository database.

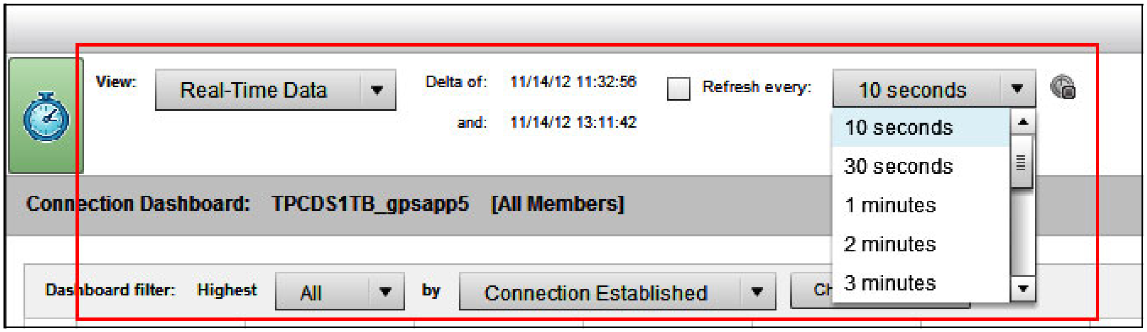

Figure 4-35 Real-time Monitoring

For real-time mode, because the monitored database is accessed additionally to the regular history collection, which increases the workload on the monitored database, the `canMonitorInRealTime` special privilege is needed for a user to receive real-time data. When you switch to real-time data by using the time selection control, then log in to the monitored database with a user having the canMonitorInRealTime privilege, then you can use the refresh controls to display real-time data. The real-time data is displayed in the following way:

- Non-cumulative metrics show the newest value.
- Cumulative metrics show the delta value between two refreshes.
- Charts show only the latest (delta) value.

InfoSphere Optim Performance Manager stores history data. You can view this data on the InfoSphere Optim Performance Manager dashboards or create reports on the long-term data, as shown in the partial display of the left side of the dashboard in Figure 4-36 on page 178. Long-term data is stored by aggregating collected data into four aggregation levels. It is a rolling aggregation mechanism. The older the data gets, the more it is aggregated. For each aggregation level, you can set a dedicated retention time. This allows you to store and analyze data, for example for multiple years in the highest level, which one-day aggregation.

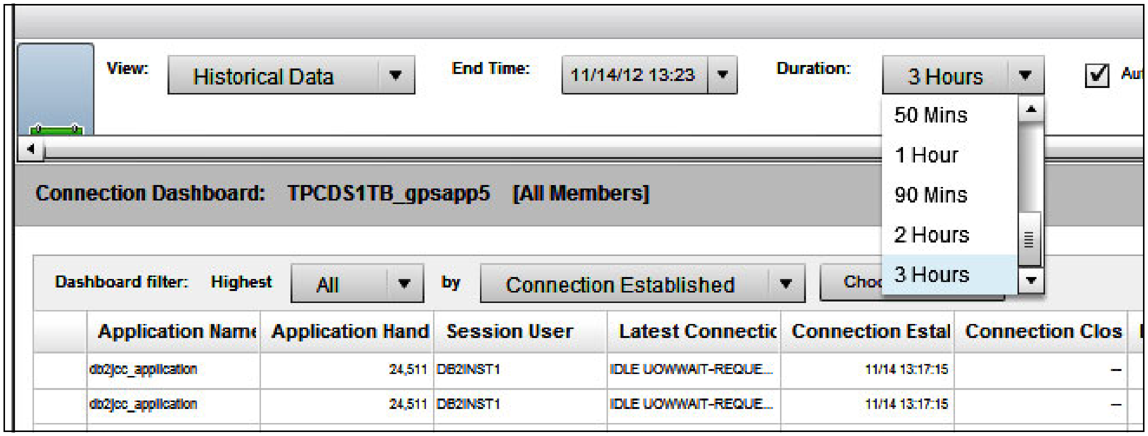

Figure 4-36 Historical view

4.4 Inflight dashboards

After initial review of the database health from the Health Summary page and the Overview Dashboard (described in 4.1, "Getting to know the layout" on page 148), you can continue to drill down for more detailed database performance information. InfoSphere Optim Performance Manager delivers this information in a series of database performance dashboards, which are grouped under the Inflight dashboards category. They provide information about a database that relates to a particular category of potential performance issues including buffer pool and I/O, locking, logging, memory, active SQL, system, utilities, and workload.

This section gives an overview of individual Inflight dashboards. To navigate to Inflight dashboards, select **Task Manager** → **Inflight dashboards**.

All inflight dashboards facilitate partition-level analysis enabling easy viewing of performance metrics across all partitions or a single partition. They show you the partition/member picker control that you find on each inflight dashboard to select the partition. By default, the dashboards display performance metrics aggregated across all partitions or members.

> **Partition sets:** For partitioned DB2 databases, the dashboards can display data for configured partitions when you use the All Members button and select the partition from the drop-down menu.

4.4.1 Overview Dashboard

The previous version of InfoSphere Optim Performance Manager had a useful overview dashboard, which showed key performance indicators (KPIs) from various other dashboards. However, the newer dashboard has many improvements, such as giving a true at-a-glance picture. With this Overview Dashboard, you can toggle between the two views, as shown in Figure 4-37 on page 179.

Figure 4-37 Overview Dashboard views

At-a-Glance View

If you are monitoring DB2 V9.7 (or later) databases using the in-memory metrics data collection method, you can use this new tab on the Overview Dashboard to see a higher level view of system activity and compare current or historical data to baseline data. On this tab you can view CPU distribution overall or within DB2, workload, plus locking, I/O, and SQL processing metrics. See Figure 4-38 on page 180 for At-a-Glance View tab on the Overview Dashboard.

> **Note:** A significant amount of information is provided in the At-a-Glance View and it cannot all be shown in Figure 4-38. Therefore only a partial view is shown, simply as an example.

Almost all metrics are depicted in bar graphs, showing the average value over the current monitoring interval. In this way, reading “at a glance” is easier than looking at a volume of numbers. By changing the interval size on the time slider, you can view either longer-term averages showing broad trends, or shorter-term values giving a fine-grained picture of performance. If no baseline is specified, you see only blue bars, no normal range bars are displayed.

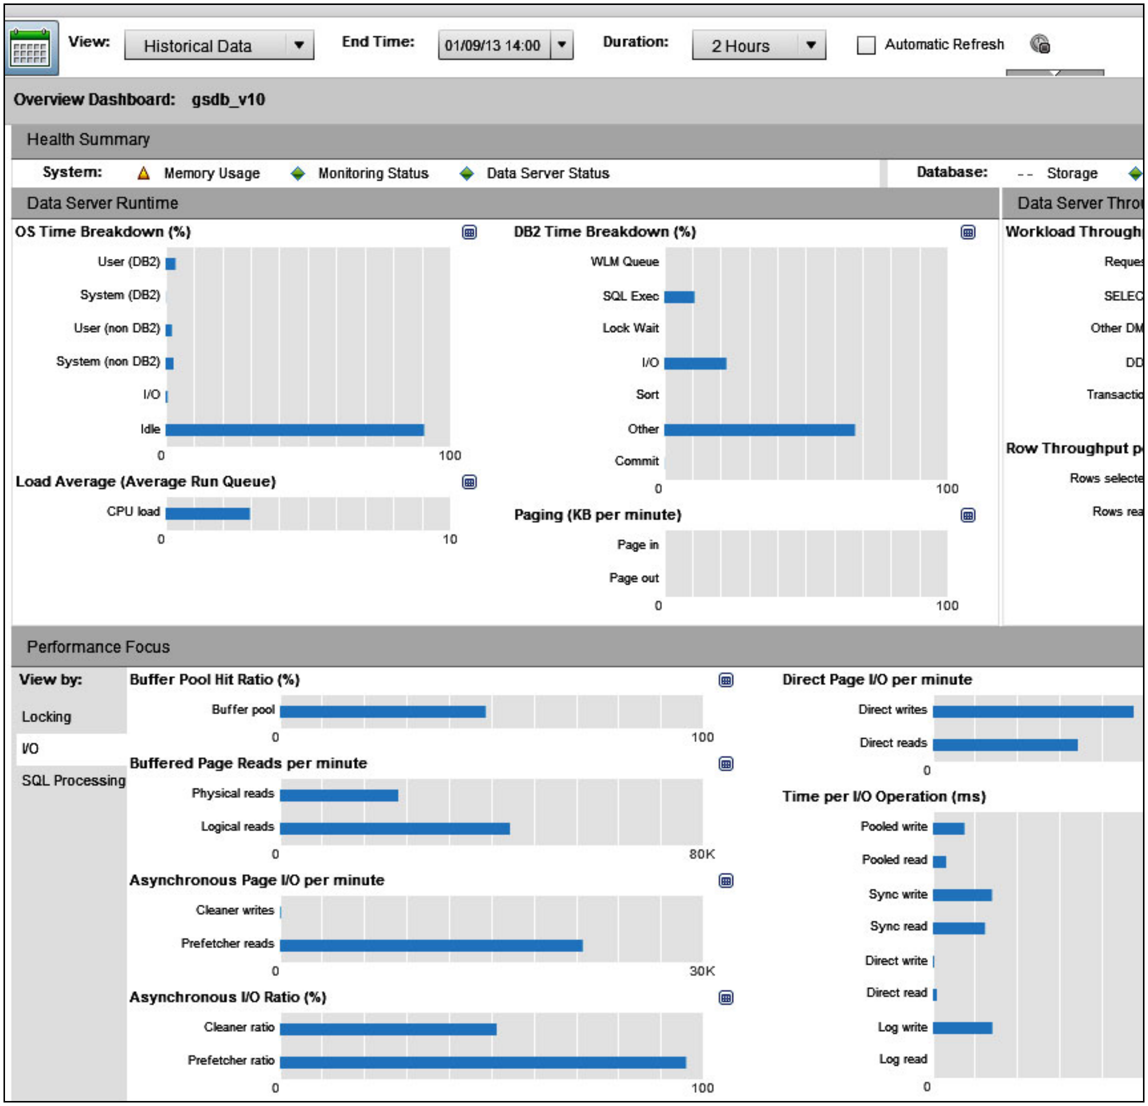

Figure 4-38 Example of the left side of the Dashboard with the "At-a-Glance View"

Use the time slider control to select the baseline time period. See item 8 in Figure 4-6 on page 152 to set a baseline. After a baseline is set at-a-glance values are automatically compared to the baseline and significant deviations are indicated.

When a baseline is specified, normal range bars are added to charts on the Performance Overview Dashboard. See Figure 4-39 on page 181.

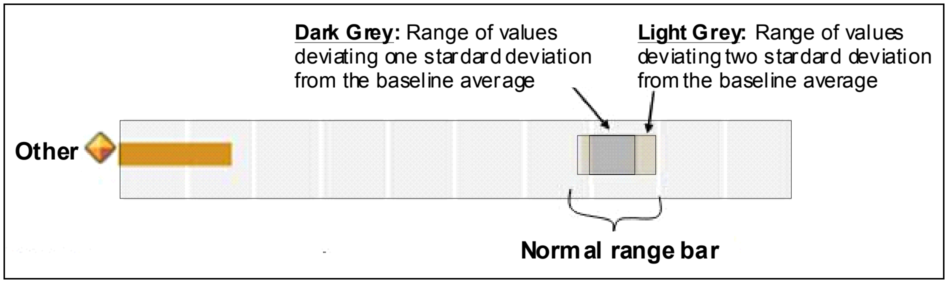

Figure 4-39 Normal range bar

If a baseline is specified, normal range bars are displayed and the length of the normal range bars represents the normal range of values for the metric (from low threshold value to high threshold value). If the current value is within this range, the current value bar will be the standard color blue. If the current value is outside of this range, the current value bar changes color to orange and the warning icon is displayed. The length of the normal range bar is determined by the following example definitions of the low and high threshold values:

- `low threshold value = max(lowest baseline value, baseline average - 2 standard deviations)`
- `high threshold value = min(highest baseline value, baseline average + 2 standard deviations)`

Figure 4-40 on page 182 shows an example of a baseline comparison overview. Some of the bars are blue and some are orange color. The orange bars indicate that the values in the interval have strayed away from a baseline that the user has set – either lower or higher. This way is convenient for determining which workload metrics have changed, and might require attention.

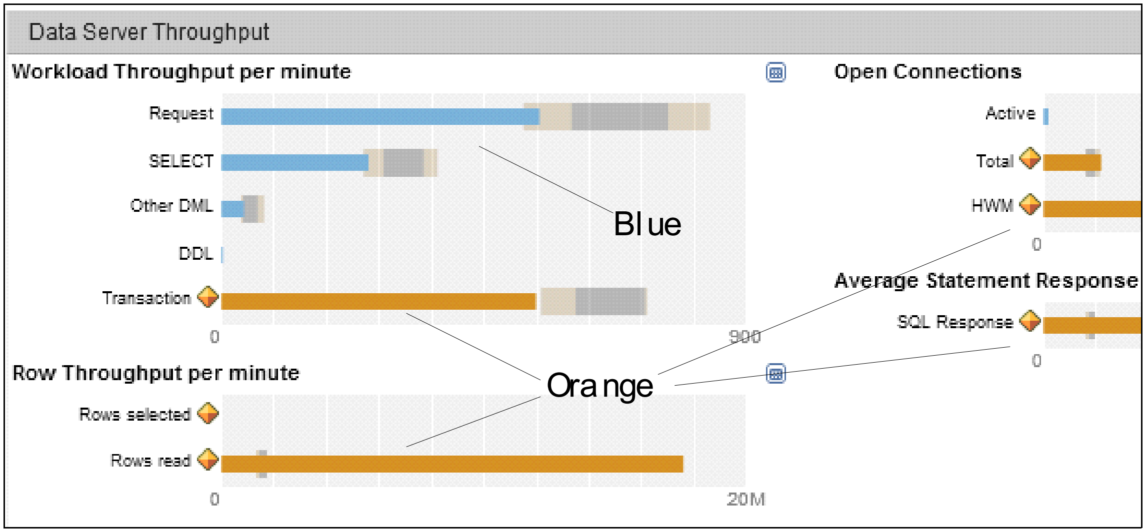

Figure 4-40 Overview baseline comparison

If a baseline is set, values that fall outside the normal range are highlighted with an icon and the color of the graph changes. Values are highlighted if they meet either of the following two conditions:

- If they are outside of two standard deviations of the baseline average
- If they are above the highest or below the lowest values that occurred during the baseline time period.

The normal range bar will be colored to represent standard deviation ranges from the baseline average. Values within plus or minus one (+/- 1) standard deviation have a darker grey color. Values within plus or minus two (+/- 2) standard deviation have a lighter grey color.

The baseline mechanism on the overview dashboard uses measurement values from an interval chosen by the user (an interval which is considered 'normal'), to compare other intervals with. After you identify and select a good baseline period, InfoSphere Optim Performance Manager will either leave the bars blue (if the values are within +/- 2 standard deviations of the average), or color them orange (and display a small diamond next to the bar) if the value is outside the 2 standard deviation range.

Hover over the graph for each performance indicator to see the value and baseline values for each metric. Extra information about that metric is displayed: average, low, and high values, standard deviation, and so forth. This is helpful in determining why a bar might have become orange; after all, it might either be higher or lower than the baseline.

You can compare the current values to the set baseline values. In the example in Figure 4-41 on page 183, the current value for the average number of transactions per minute is 538.896. The baseline average is 665.673. Calculating two standard deviations from the baseline average means that values between 547.063 and 729.323 are allowed to not fall outside the range. However, because the current value is 538.896 and is outside of the normal range, it is therefore truncated and highlighted.

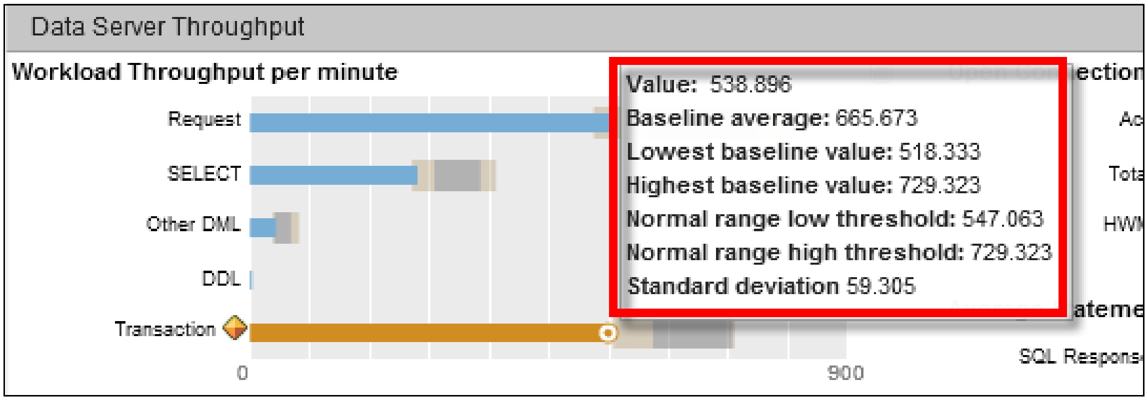

Figure 4-41 Average Statement Response time per millisecond

4.4.2 SQL Statements Dashboard

With this dashboard (Figure 4-42 on page 184) you can analyze the SQL statements that were running on your monitored database. This dashboard offers two views:

- Execution Summary view
- Top Individual Executions view

You can select an SQL statement from the list and look through the key performance metrics in the SQL Statement Details area to see if the statement negatively impacts your system. If the statement is a stored procedure call, you can drill down into the stored procedure to analyze the executed statements.

> **Note:** DB2 V10 fix pack 2 is required on the monitored database for stored procedure monitoring.

From this dashboard, you can stop queries and force applications if necessary. You can click **View Configuration Changes** to launch Configuration Manager in context so you can view relevant changes and more quickly determine whether recent changes might be the underlying cause of performance issues. Both views support tuning SQL statements using Optim query Tuner.

Execution Summary

The summary view Figure 4-42 on page 184 highlights aggregated information over all executions of a statement.

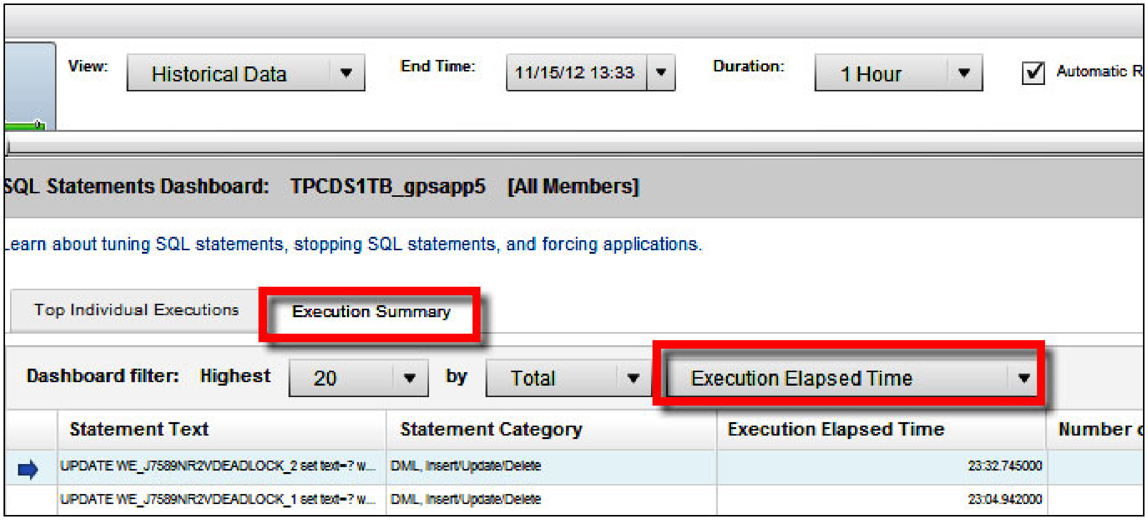

Figure 4-42 SQL Statements Dashboard, Execution Summary

This page is based on DB2's package cache and it displays a summary of all executions of a compiled statement in a grid. InfoSphere Optim Performance Manager uses MON_GET_PKG_CACHE_STMT, which is DB2's in-memory metrics function.

You can select a time frame in history to see what your top statements are within this time frame or use the real-time mode to see the newest top statements. To select your criteria and identify the top statements by this criteria use the dashboard filter controls. You can either get the top statements by total across all executions in this time frame or by average per execution value of your criteria, for example, execution time. As criteria, you can select any metrics from the displayed grid. To adjust the metrics in the grid, use the icon () to open the Select and order columns dialog. For example, if in the Overview Dashboard you detect heavy sorting activity in your database, you might want to determine which statements do heavy sorting. To do this, you can add sort processing time, sorts, and sort overflow metrics to the grid and select one of them to list the top statements.

Note: Depending on whether you selected total or average values for your top criteria in the Dashboard filter control, you might see a different set of metrics in the Select and order columns dialog. For the total values, the set of metrics is limited.

Select a statement in the grid to get more details of the statement in the lower part of the dashboard shown in different tabs (such as Overview, server execution times, row activity, locking, I/O and so on). Click the **Server Execution Times** tab (Figure 4-43 on page 185) to display graphs with detailed metrics of different wait times.

Note: The details section always shows average values by execution even if total values are displayed in the upper grid.

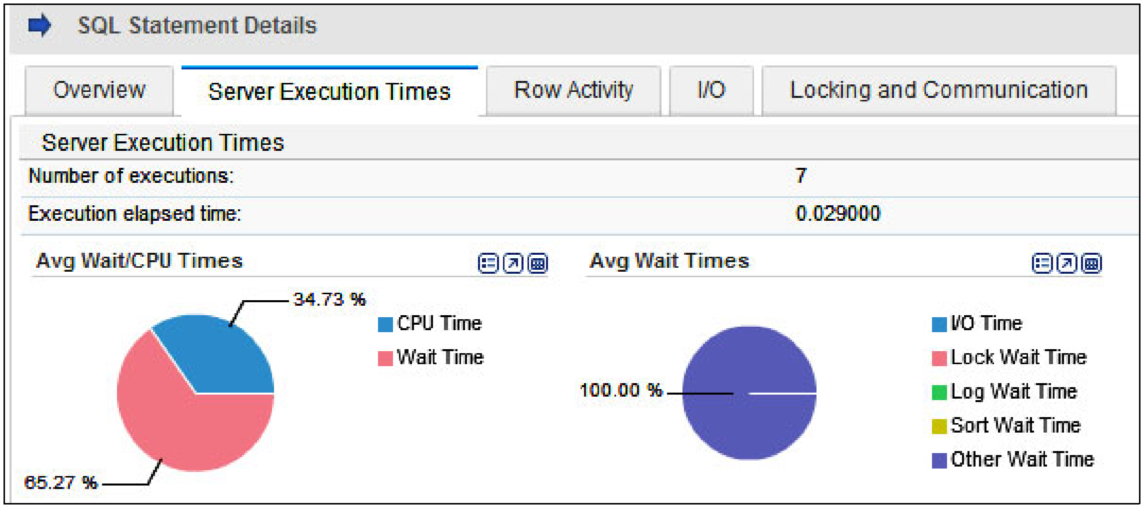

Figure 4-43 Server Execution Times

The Overview tab has an Actions pull-down menu, from which you can tune the specific statement, or switch context views by clicking either **Show the current top individual executions of this statement** or **Show top individual execution of this statement**, which switches to top execution mode either in real-time mode or in history mode and sets a filter on the statement text to see only executions of this statement and no other.

The Actions pull-down menu is also on the right side above the grid. In a multi-partition or pureScale environment this menu offers an additional action called **Show statement by Members**. In the All Members view, you can select a statement in the grid and choose this action to see and compare how this statement is performing on each partition or member. When you select this action, a filter is set on this statement and you get one row per member or partition in the grid. In this way, you can more easily compare how many rows this statement reads on each member or partition.

Tuning SQL statements

The process of tuning an SQL statement consists of analyzing the statements, and then applying the recommendations from the results of the analysis (for example, changing the statement text or creating indexes).

Check the key performance indicators in the SQL Statement Details area for statements that you might want to tune. For example, a high value of sort time might indicate that an index is missing, and a high value of rows read per fetched row might also indicate that an index is missing. If one of these two values is high and the total execution time is unusually high, then tuning these statements is advisable.

To tune the statement, start Optim Query Workload Tuner product, and then click either **Tune** or **Tune All** from one of the **Action** menus. Optim Query Workload Tuner can provide choices for indexes, materialized query tables (MQTs), or statistics to improve the execution of the SQL statements.

- Using the Tune action, a single statement is passed to InfoSphere Optim Query Workload Tuner.
- Using the **Tune All** action all statements displayed in the grid are passed to Optim Query Workload Tuner and Optim Query Workload Tuner tunes these statements as a group. This avoids making tuning decisions for one statement that might harm others.

Top Individual Executions

The Top Individual Executions view in Figure 4-44 on page 187 shows single executions of statements ordered by a specified top criteria, for example CPU time. You can use this dashboard in real-time mode to monitor current statement executions or in history mode to monitor statements executed in the specified time frame. Select a statement to get more details about the statement in the lower part of the dashboard shown in different tabs, for example, execution times, row activity, locking, I/O and so on.

If you click the **Applications** tab on the lower part of the dashboard then you get information about which connection this statement belongs, or in which service class or workload this statement is executing. From there you can see launch-in-context reports, for example, a Connection report or WLM reports about statistics, service class, or service subclass under which this statement is running.

In-context switching between the Top Individual Executions mode and SQL Summary mode is also possible from here. Select **Action → Show execution summary**, which switches to the Summary mode and sets a filter on statement text so you can see only the summary for the selected statement

As described in “Execution Summary” on page 184, you can also tune a single statement from here or select **Action → Show statement by member** in multi-partition or pureScale environments.

> **Note:** This view is mainly based on MON_GET_ACTIVITY_DETAILS, which is the DB2 in-memory metrics function, and WLM_GET_WORKLOAD_OCCURRENCE_ACTIVITIES_V97, which InfoSphere Optim Performance Manager calls in the specified sampling rate. These functions return information about activities that were submitted by a connected application and have not yet completed. As a result, InfoSphere Optim Performance Manager displays only statements that were running at the time when InfoSphere Optim Performance Manager collected the data. If InfoSphere Optim Performance Manager collects data in a one-minute sampling rate then all statements are missed that start and stop between two samples. This view is valuable for environments where you typically have long running statements, that run multiple minutes or hours, such as in business intelligence environments. For OLTP environments that usually have statements with subsecond response times, use the Execution Summary view to analyze statement execution.

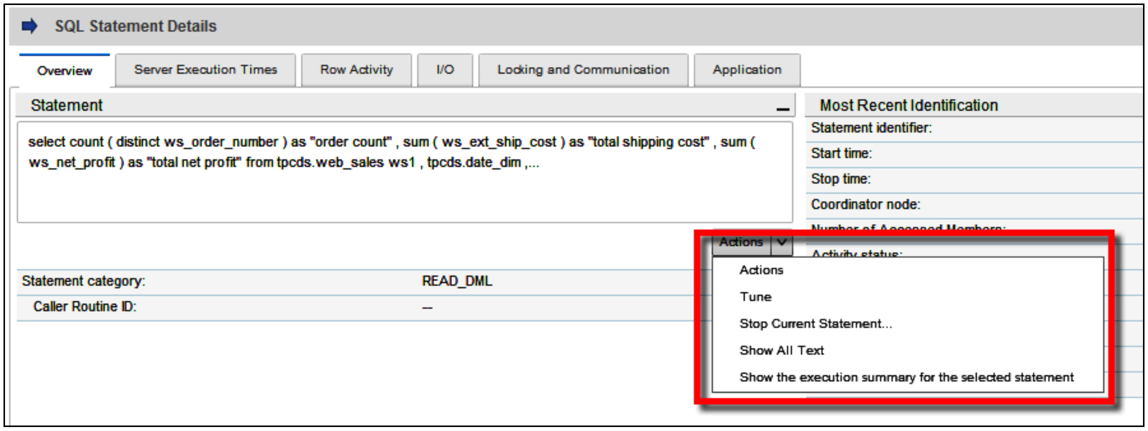

Figure 4-44 Top Individual Executions

Stop current statement

Check the key performance indicators in the SQL Statement Details area for statements that you might want to stop. For example, check the values for CPU time, sort time, and physical reads. If they are high, then other workloads might be affected. In such cases, you might want to stop the statement. Various applications might be allowed to use a large number of resources, so the decision to stop the statement also depends on the application that is executing the statement. Look at the Application section to see which application is executing the statement so that you can decide if the resource usage is unusual. If you want to release the resources to the system, you can stop the statement by clicking **Stop Current Statement**. This will roll back the current statement.

Forcing applications

Similar to stopping a currently executing SQL statement, you can also use the Applications tab from the SQL Statements Details area and force an application by clicking **Force Application**. Forcing an application cancels the database connection that is being used by the application.

> **Attention:** The impact of canceling a connection might be very large. Therefore, consider using the **Stop Current Statement** button instead to reduce the impact.

Tuning SQL statements

For information about this topic, see “Tuning SQL statements” on page 186.

4.4.3 Stored Procedure monitoring

As a new feature in InfoSphere Optim Performance Manager V5.2 you can now use the SQL Summary view on the SQL Statements Dashboard to analyze statements for stored procedures. Use the controls on the dashboard to see top stored procedures by aggregated execution metrics over all executed statements and nested stored procedures, and drill down into a stored procedure to analyze the executed statements. This feature is supported for DB2 v10 FP2 for Linux, UNIX, and Windows as the monitored database. As a further prerequisite, you must have enabled stored procedure monitoring in the SQL Statements and Connections monitoring profile, described in 3.2.2, “Adding and configuring database monitoring” on page 96.

If stored procedure monitoring is enabled, InfoSphere Optim Performance Manager uses the following DB2 in-memory metrics functions:

- MON_GET_ROUTINES to collect, for each stored procedure, the aggregated execution metrics over all executed statements
- MON_GET_ROUTINE_EXEC_LIST to get the list of all statements executed by a stored procedure

Figure 4-45 on page 189 shows stored procedures from the Execution Summary view of SQL Statements Dashboard. By selecting a stored procedure displays some general information for this stored procedure in the details section, such as the routine ID or number of executions.

Stored Procedure Information	
Routine ID of stored procedure call:	67,193
Number of Accessed Members:	1
Name:	TST10L5.ITALIAN_MEATCOURSE
Module name:	--
Specific name:	SQL121128092316204
Number of executions:	590
Number of nested stored procedures:	590
Heap use exceeded:	N
Last execution list cleanup:	--

Figure 4-45 Stored procedure information

To analyze the statements executed by a stored procedure, select a CALL statement and click **Actions** → **Show SQL Executed by the Stored Procedure**. This menu item is available only if stored procedure monitoring is enabled.

As a result, only the statements executed by the stored procedure are listed in the grid and the stored procedure name appears in the history navigator control. Each time you drill down into a stored procedure, the name is added as a further entry in the history navigator. Using the back and forward buttons in the history navigator, you can browse through stored procedures and see which statements your stored procedures are executing and how they perform. Selecting the **All Statements** entry from the history navigator always leaves the drill down view and shows you all statements in the grid.

Statements that are executed by a stored procedure are collected by InfoSphere Optim Performance Manager through the MON_GET_ROUTINE_EXEC_LIST table function. That table function returns a list of all statements (sections) executed by each procedure, external function, compiled function, compiled trigger, and anonymous block that has been invoked since the database was activated. Those metrics on the stored procedure level are then aggregated. For all execution metrics, only the sum of the metrics of the single statements are executed by the stored procedure.

Stored Procedure monitoring and Extended Insight

For an introduction of Extended Insight, see 4.5, “Extended Insight Analysis Dashboard” on page 203. If you see a stored procedure call on the Extended Insight Analysis Dashboard, this does not reflect aggregated execution metrics over all statements within the stored procedure, even if you monitor DB2 V10 fix pack 2 and have enabled stored procedure monitoring in the monitoring configuration.

To analyze aggregated stored procedure information and drill down to the CALL statements, select a CALL statement, and click **Action** → **Show the execution summary of the selected statement**. This action opens the SQL Statements Dashboard with a filter set on the routine ID and displays this CALL statement with only aggregated execution metrics. From there, you can further drill down and analyze the executed statements.

4.4.4 Buffer Pool and I/O Dashboard

Figure 4-46 on page 190 shows the left side of the Buffer Pool and I/O Dashboard; it highlights the database I/O at the buffer pool, table space, and table level.

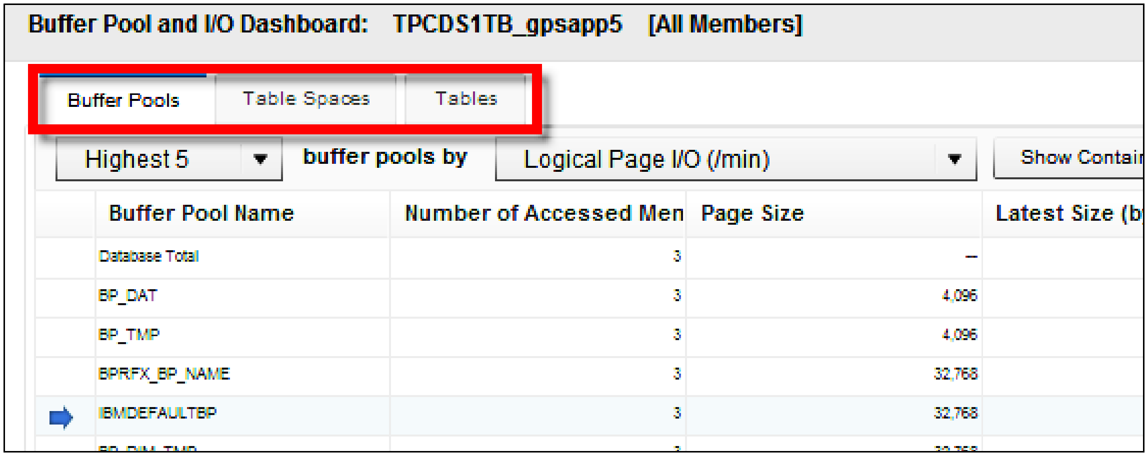

Figure 4-46 Buffer Pool and I/O Dashboard, left side

Figure 4-47 on page 191 shows the right side of the Buffer Pool and I/O Dashboard.

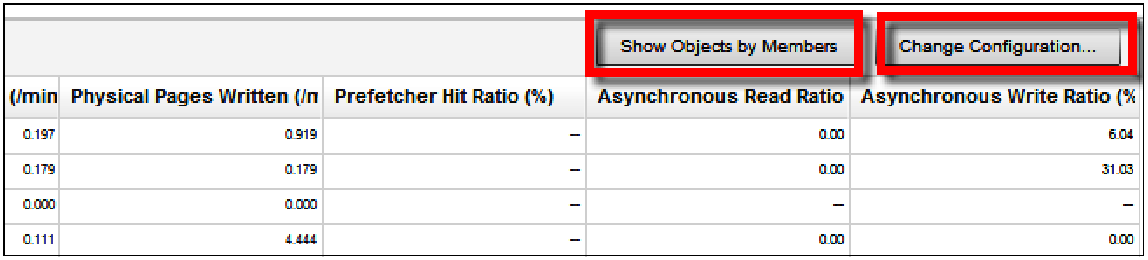

Figure 4-47 Buffer Pool and I/O Dashboard, right side

On a full view of the actual dashboard, you can see the highest and lowest buffer pools by selecting the metric that you want to learn more about. Then, you can find the buffer pools (highlighted) with low hit ratios and high activity and consider increasing their sizes to improve performance. You can also find table spaces and tables with low hit ratios and high activity and look for statements that access those tables that might need tuning.

Click an item in the grid to view details about that item in the Detailed Information area. Select a buffer pool and click **Show Contained Objects** to view the table spaces that use that buffer pool. Select a table space and click **Show Contained Objects** to view the table objects that the space contains.

Table spaces

In the Table Spaces tab, you can also find all the table spaces with, for example, low hit ratios, high activity, or high I/O times. The bottom half of the page has detailed information for a specific table space. You can view details about the I/O read and write times and get details about disk size and container, for example.

Select a table space and click **Show Contained Objects** to identify the tables that are located in the selected table space and identify “hot” tables that might be responsible for high I/O activity on that table space.

Tables

If you see hot tables in the Buffer Pool and I/O Dashboard that for example show a higher number of rows read, then you can now select such a table and drill down to the SQL statements that use that table. Figure 4-48 on page 192 shows the table view in the Buffer Pool and I/O Dashboard. Select the table with the highest number of rows read and click **Show SQL**.

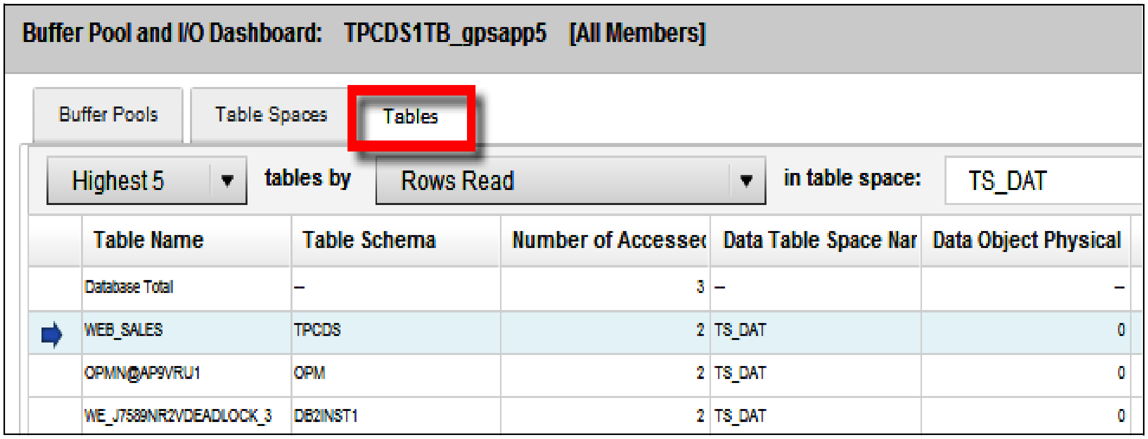

Figure 4-48 Buffer Pool and I/O Dashboard tables view

4.4.5 Locking dashboard

The Locking Dashboard (Figure 4-49 on page 193) shows the maximum wait time, block time deadlocks, timeouts, and lock wait alerts for all applications running on the database. You can use this dashboard to determine which applications have the most locking problems, and you can drill down to find the waiting or blocking connections and events.

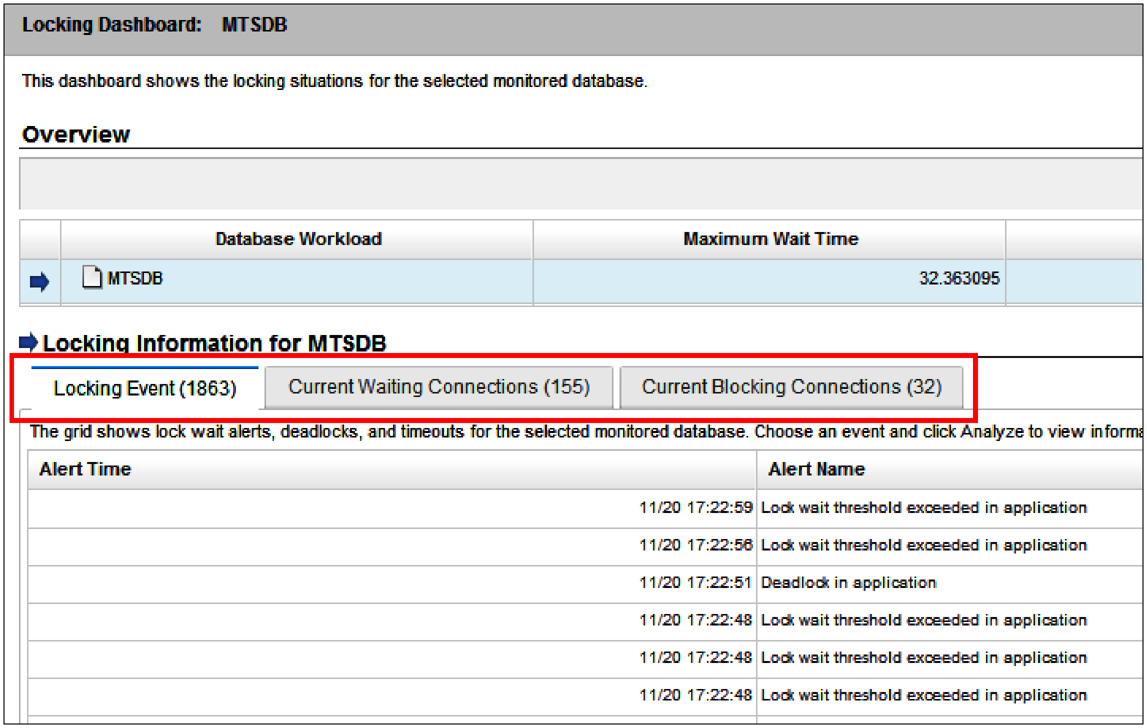

Figure 4-49 Locking dashboard

To analyze a locking event use the following steps:

1. Go to the Locking Information for the database; it displays the list of locking events alerts, current waiting connections, current blocking connections.
2. In the Locking Information grid, click the connection (waiting or blocking) or event that you want to analyze.
3. Click **Analyze**.

Figure 4-50 on page 194 shows the Analyze Locking Situation panel. This panel displays the complete lock tree for a waiting or blocking connection. Each complete set of entries in the tree includes an application that is holding a lock and the applications that are waiting because of that lock. The entries between the main entry and the leaf entry are applications that are blocking and waiting.

Each leaf entry is an application that is only waiting. If there is a deadlock event, waiting connections are also blocking connections. You can use this panel to force an application, and stop or tune the statement the application is running.

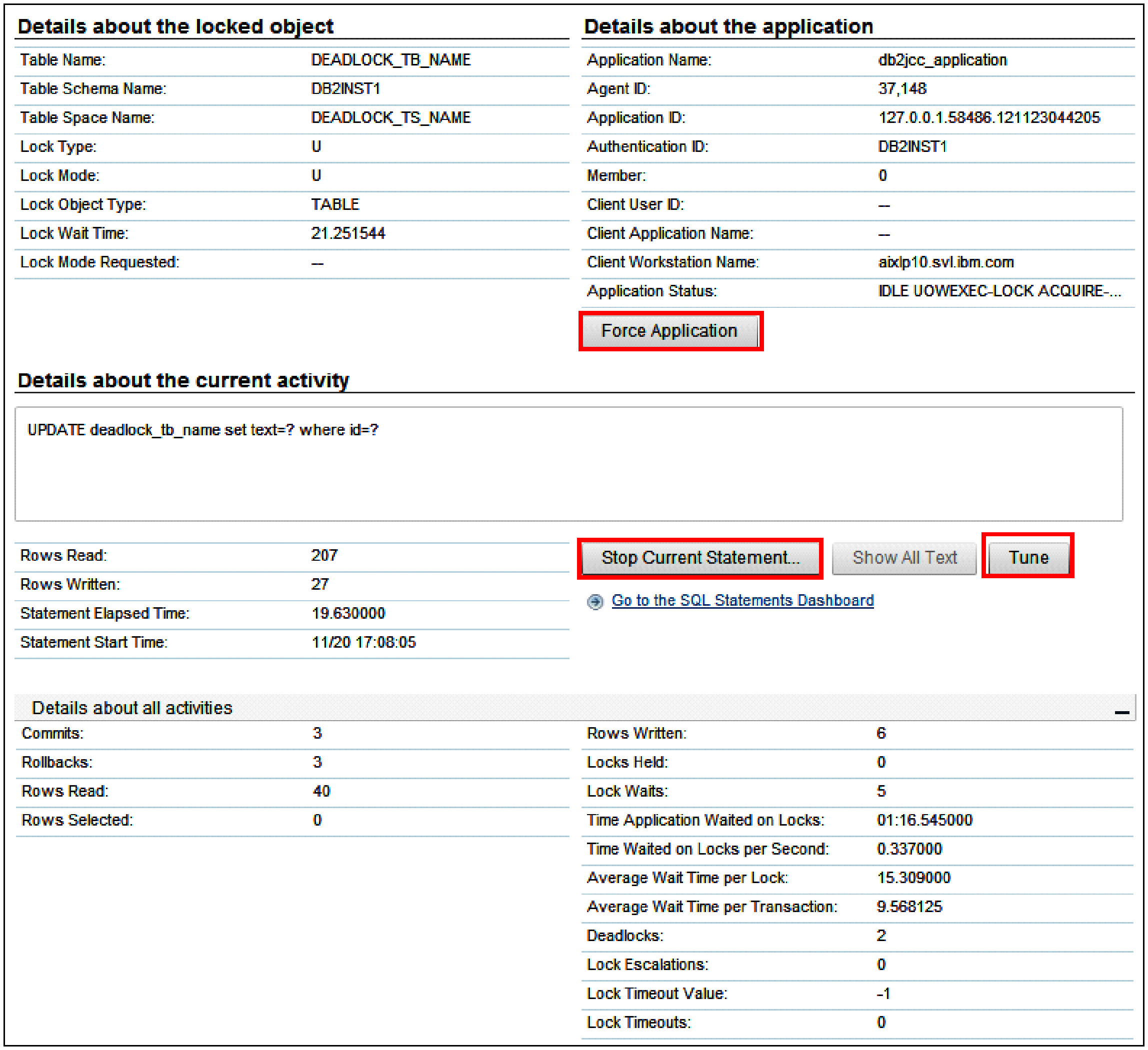

Figure 4-50 Analyze lOcking Situation panel

4.4.6 Logging Dashboard

This dashboard (Figure 4-51 on page 196) shows a partial view of the configuration and logging activity. You can use this dashboard to determine how recovery is affected and whether tuning is needed for the log files or log buffer size. High average log read/write times, which are in the Logging activity and I/O Performance section of the Logging Dashboard, can indicate whether disk configuration for the database log files can be responsible for the database performance degradation.

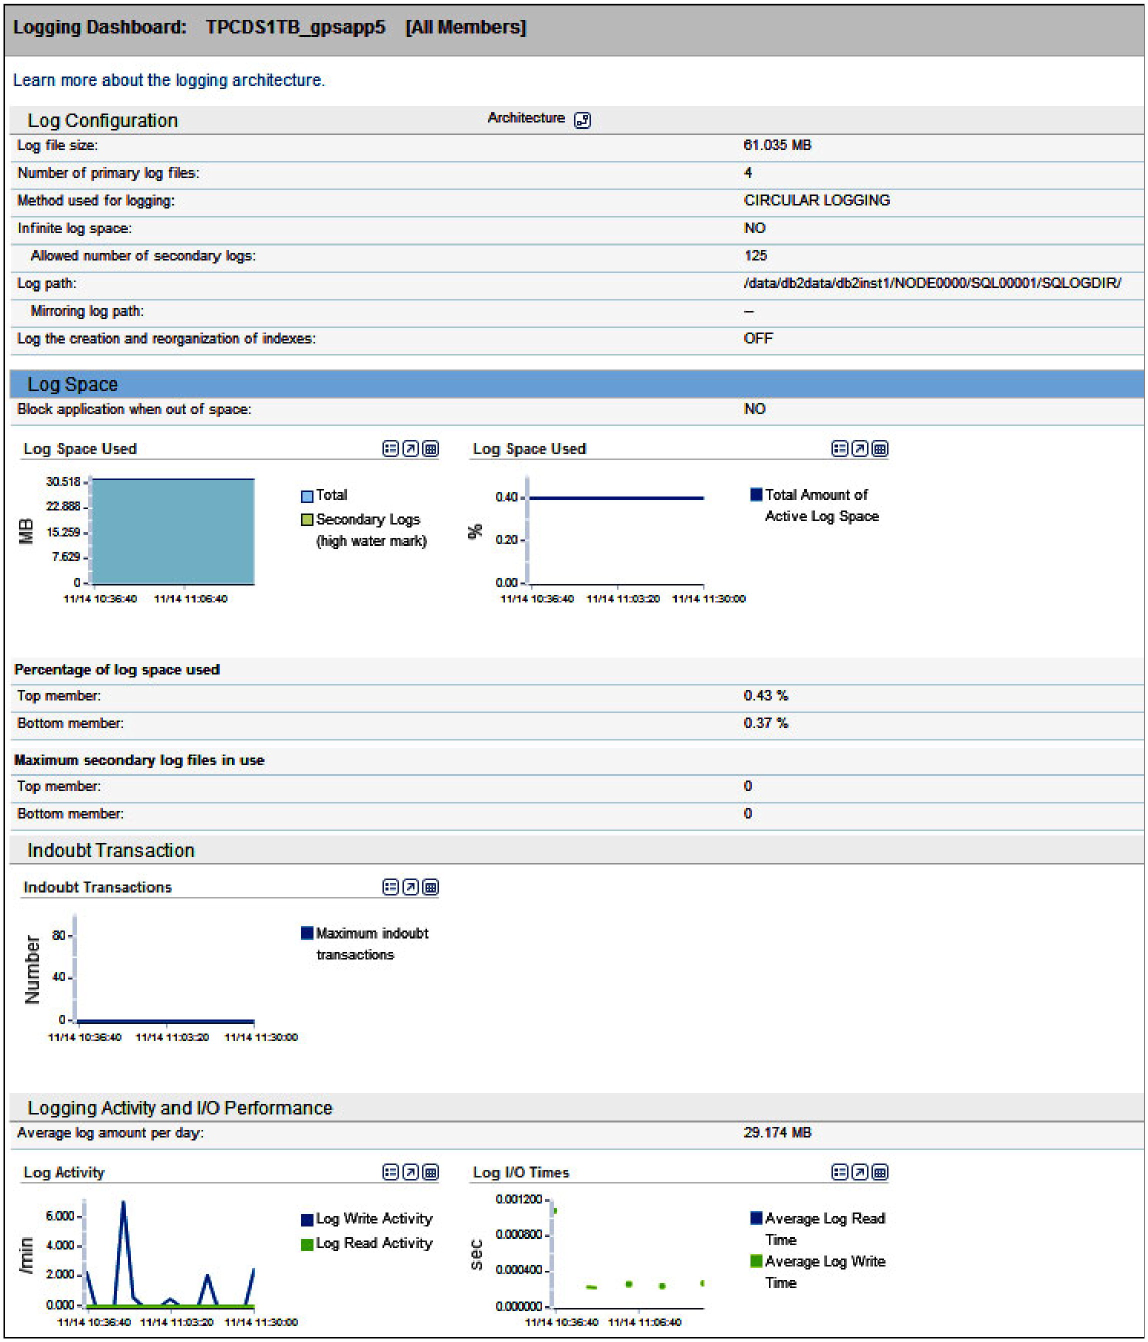

Figure 4-51 Logging dashboard

4.4.7 Memory Dashboard

This dashboard (Figure 4-52 on page 197) shows the memory consumption of the selected database. It shows memory usage by instance, by database, and by application, and shows memory that is shared between applications.

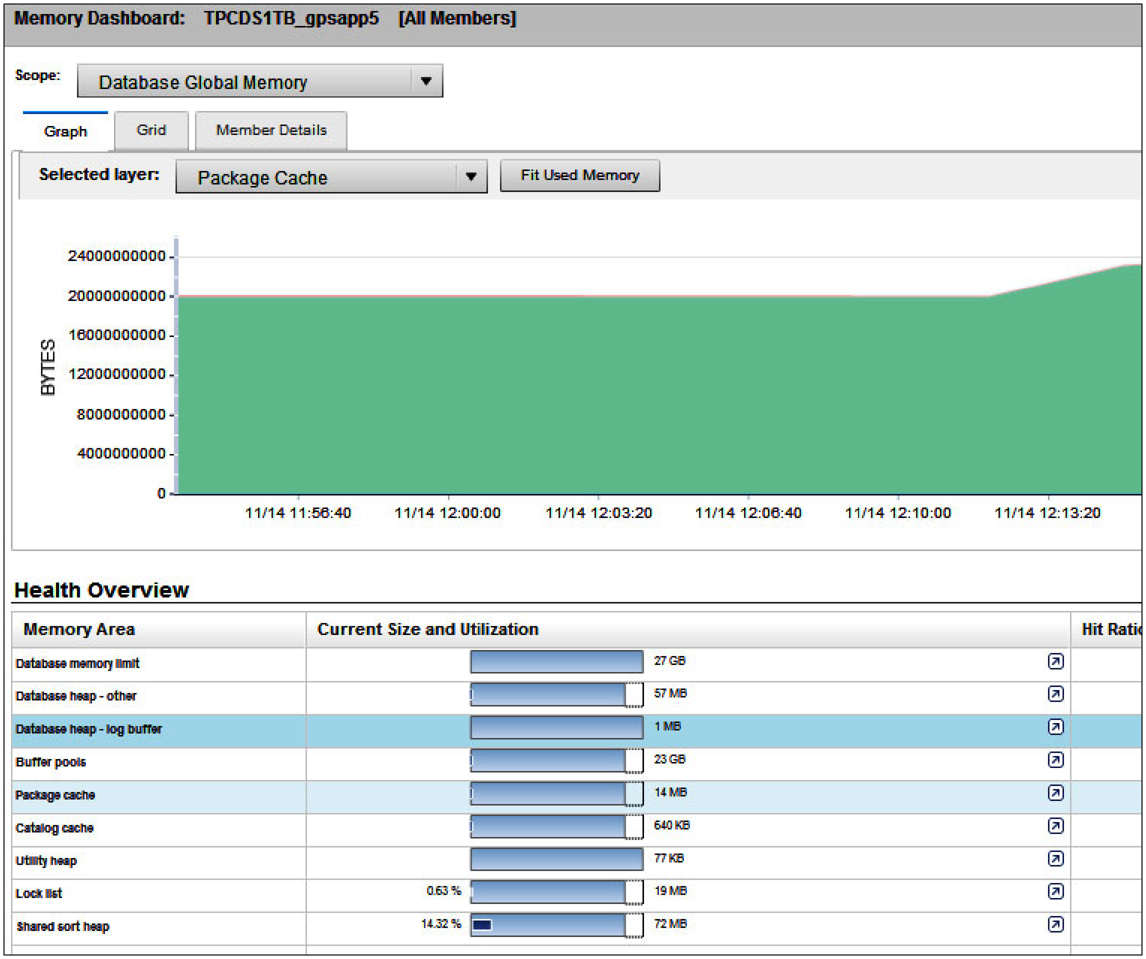

Figure 4-52 Memory dashboard

Open the Graph tab to display the statistics of the memory scope that you selected as a distribution chart. The layer (memory area) of the highest value is highlighted, by default. Use the Selected layer menu to select a specific memory area to be highlighted both in the chart and in the bottom of the Health Overview table.

In a multi-partition environment, you can open the Member details tab to compare memory area values across partitions or members. To view the memory consumption of a specific partition or member only, click the **Members** menu, and data for each of these partitions is displayed.

The bottom Health Overview table shows the health status of a database. The Health Overview table shows how the memory areas are configured, the allocated size of the areas and how much of the allocated memory is in use, and you can check the efficiency of the memory area by examining the hit ratio. By looking at this information, you can determine whether a memory area is healthy, whether you want to increase the size of the memory area, and whether you can decrease the size to release memory.

4.4.8 System dashboard

The System Dashboard (a partial view of the left side and right side is shown in Figure 4-53 on page 199) shows the system resources of the system on which the database is running. It provides information about the CPU and memory utilization. To see details about each member, click **Members**, and from the pull-down menu, select the member. The dashboard is refreshed to display information pertaining to that specific member. If you launched InfoSphere Optim Performance Manager from Tivoli Enterprise Portal Console, you can click the **Advanced System Information** link of the System dashboard, which opens the Tivoli panels that provide more detailed system information.

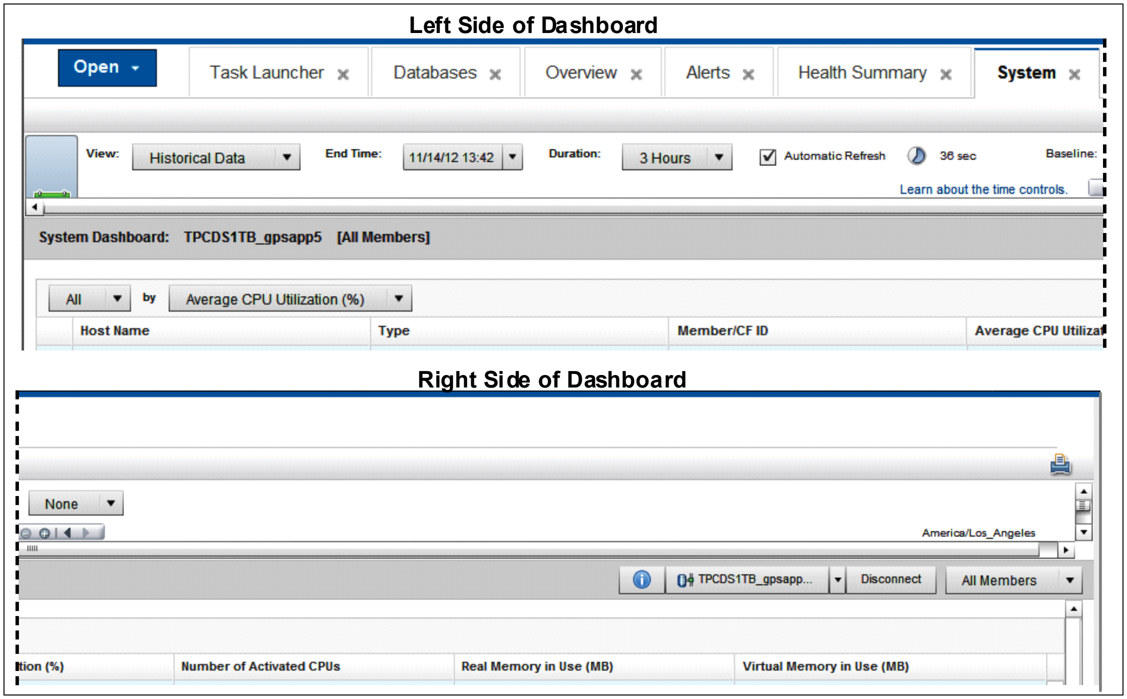

Figure 4-53 System dashboard

4.4.9 Utilities Dashboard

A partial view of the Utilities Dashboard is in Figure 4-54 on page 200. It shows the status and progress of utilities such as RUNSTATS, LOAD, and BACKUP. You can get information about utilities that are in progress and that are completed, and you can identify utilities that failed. This information can help you determine whether certain utilities must run at particular times to avoid high workloads.

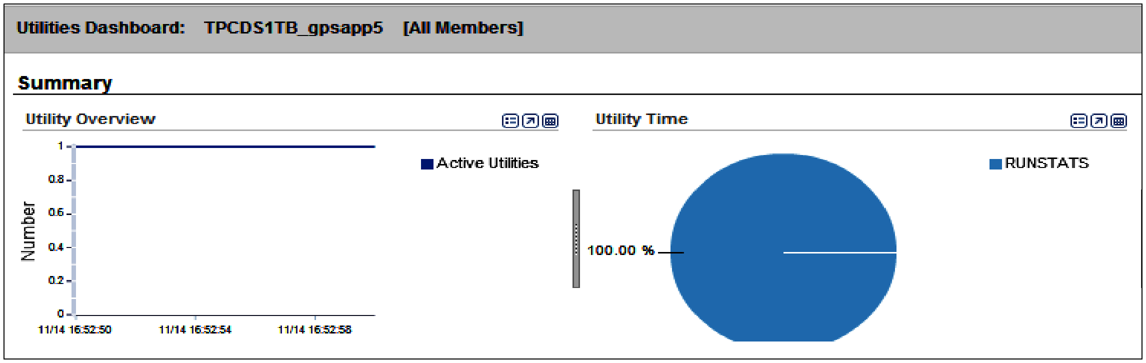

Figure 4-54 Utilities Dashboard

4.4.10 Workload Dashboard

The Workload Dashboard shows information about the workload on the database. It provides detailed information about the sort performance as seen in the partial view example of the dashboard in Figure 4-55 on page 201.

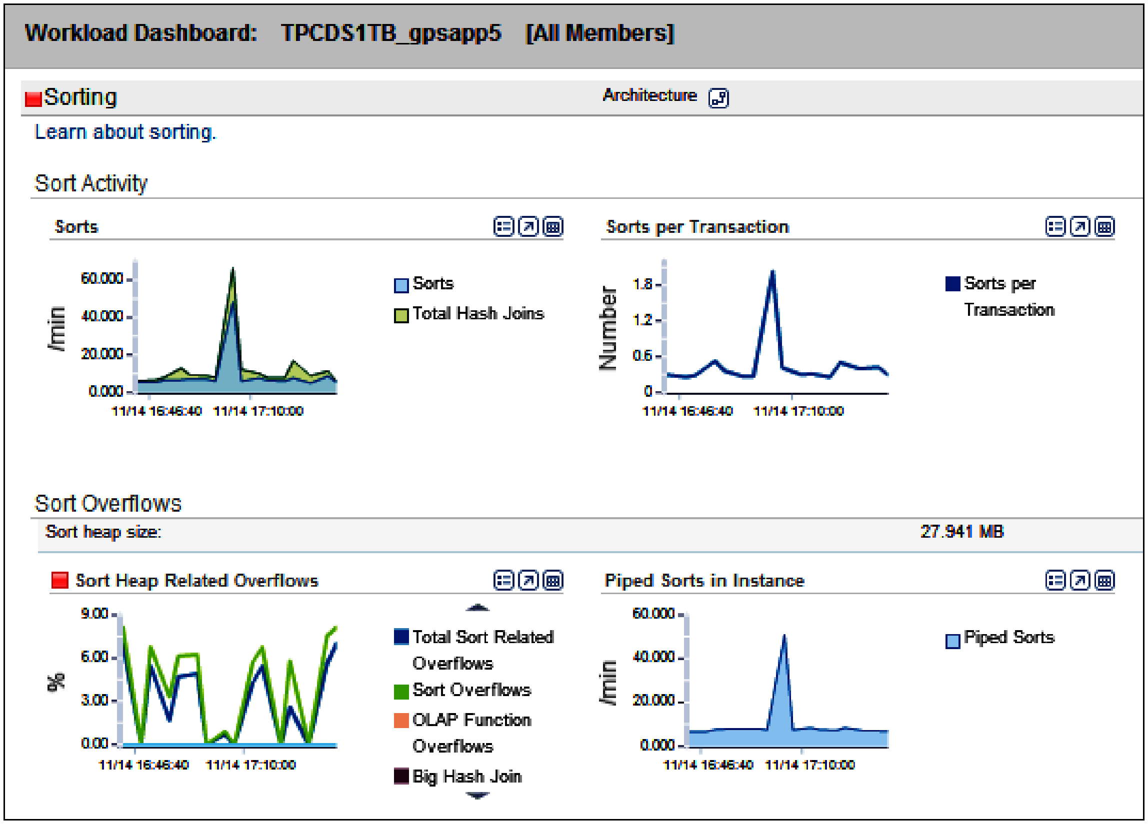

Figure 4-55 Workloads Sort Section

The throughput of transactions, SQL statements, and rows are seen in a partial view of the Workloads, Throughput and Connections dashboard in Figure 4-56 on page 202. You can use this dashboard to check the utilization of the database.

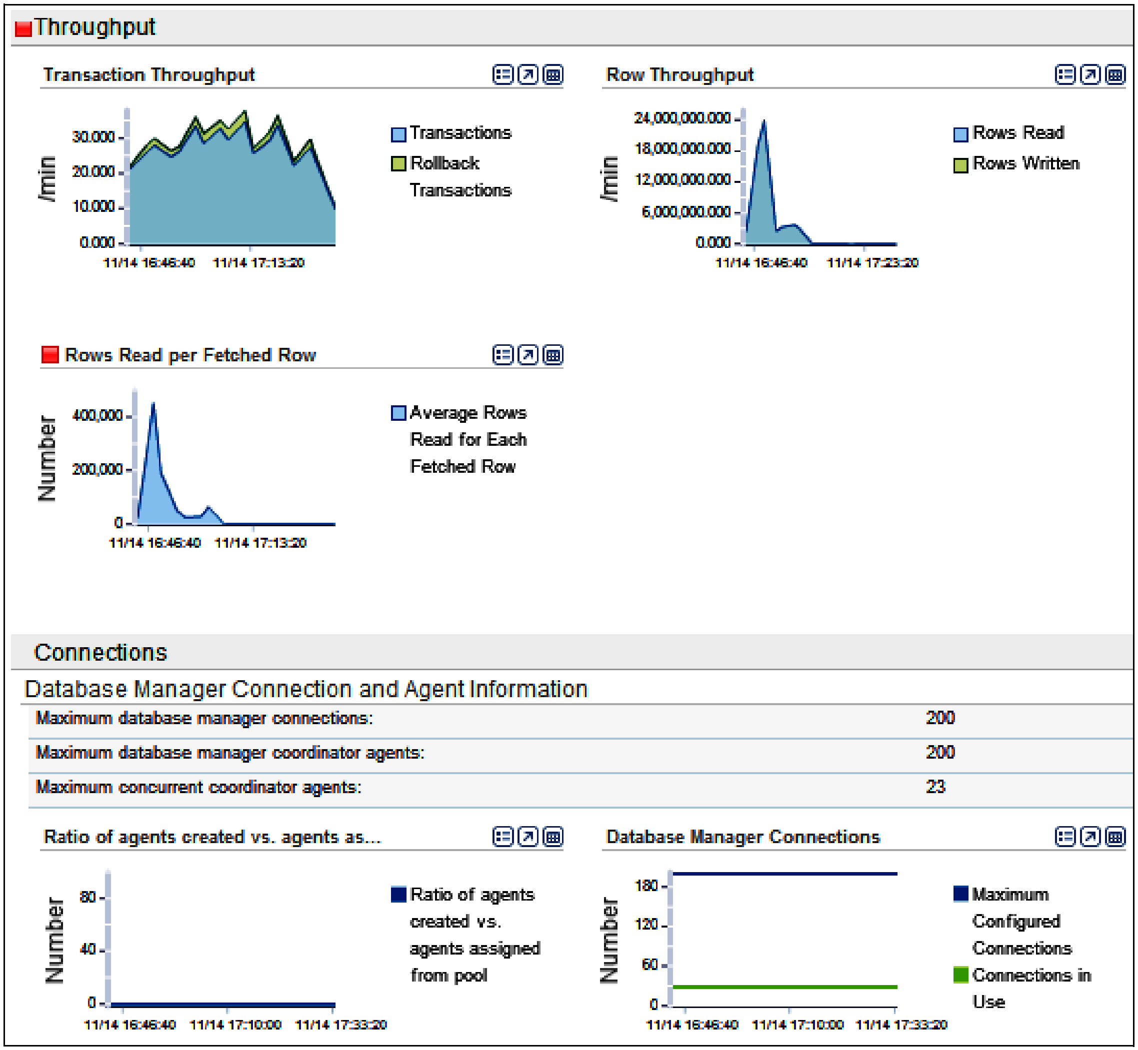

Figure 4-56 Workloads Throughput and Connections

4.4.11 Connection Dashboard

The Connections Dashboard shows the Top N current connections established (real-time mode) or the connections for a time frame in history. You can select from various Top N criteria to determine the top connections, for example by CPU time, or rows read. To list the connections by a Top N criteria you can specify filters or choose the columns to be shown in connection grid. From the dashboard, you can drill down into details of one connection or launch in-context reports, for example, the connection report. You can also force connections directly from this dashboard.

From the connection dashboard you can drill down into various details of a connection. Select a connection and use the tabs in the lower part of the dashboard to display details about the various areas, such as server times, I/O, locking, row and transaction details, application details, locking and communication and utilities.

In a multi-partition environment, to analyze how a connection is performing across members in the All Members view, select a connection and click **Show connection by members**. You can then see and compare how this connection is performing on each partition or member. When you select this action, a filter is set for this connection and you get one row per member or partition in the grid. In this way, you can compare, for example, how many rows the connection reads on each member to detect data skews.

To analyze which statements a connection executed, click **Show executed SQLs** to see a list shows the SQL statements executed by the connection and collected by InfoSphere Optim Performance Manager. The SQL dashboard is opened in the Top Execution mode filtered by the application ID of the selected connection and lists the executed SQL statements collected by InfoSphere Optim Performance Manager for this connection.

You can also force a connection from here or generate a connection report directly from this dashboard.

4.5 Extended Insight Analysis Dashboard

The InfoSphere Optim Performance Manager Extended Insight Analysis Dashboard displays end-to-end data about the entire database application system, which includes clients, application servers, data servers, and the network. Monitoring begins when you initiate a transaction, continues as that transaction is processed by each component, such as the client, network, and data server, and ends when the application finishes processing and produces the results.

InfoSphere Optim Performance Manager Extended Insight is able to identify how much time a database transaction spends in each of the software layers, such as application server, database driver, network, and database server. For monitored databases that are at DB2 9.7 Fix Pack 1 or later, InfoSphere Optim Performance Manager Extended Insight can provide further details of the transaction response times at the database server layer (examples are the average compilation time, sort processing, I/O processing, lock wait, and so on).

This information can help database administrators determine, for example, which software layer is responsible for slow database transaction response time. The end-to-end database transaction response time is defined as the time from the beginning of the first statement till the end of the commit or rollback.

The Extended Insight Analysis Dashboard consists of two panels:

- Overview panel
- Details panel

4.5.1 Extended Insight Analysis Dashboard: Overview panel

Figure 4-57 on page 204 shows a partial view of an example Extended Insight Analysis Dashboard overview panel.

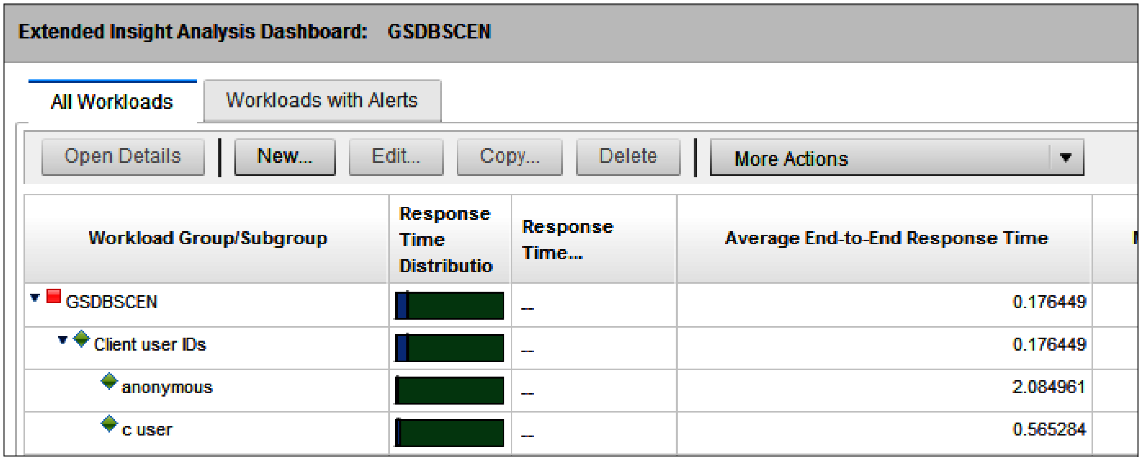

Figure 4-57 Extended Insight Analysis Dashboard overview panel

You can use the panel to view the following statistics:

- Statistics in the grid

  The Extended Insight Analysis overview grid lists statistics for the workload groups, and database that you are monitoring. InfoSphere Optim Performance Manager V5 introduces two columns:

  – Response Time Distribution shows the proportionate amount of time that was spent in each layer (server, network, and client) during the selected time period.

  – Response Time Alerts shows proportionate amount of time spent in each alert state during the selected time interval.

  The grid has the following key statistics:

  – Average end-to-end response time
  – Maximum end-to-end response time
  – Average data server time
  – Average network time
  – Average client time
  – Transactions per minute
  – Statements Failure Rate

  You can double-click a row in the grid to view response time details for a workload group, a workload subgroup, or the entire database.

- Charts

  You can view charts for a selected workload group, workload subgroup, or database by clicking the workload that is listed in the grid. Figure 4-58 on page 206 shows a partial view of the selected workload group.

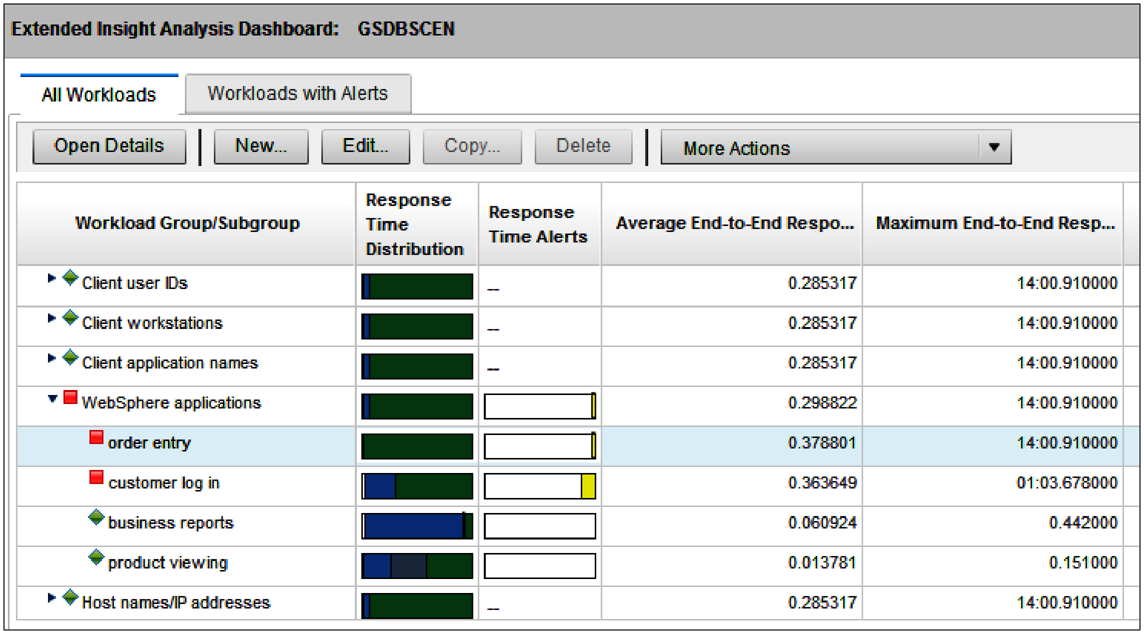

Figure 4-58 Extended Insight Analysis Dashboard overview panel

You can double-click the chart to obtain more detailed information about the transaction end-to-end response times for the selected workload group, workload subgroup, or the entire database. This detailed information also contains a response-time histogram, which groups transaction end-to-end response times into discrete elapsed time ranges. The bar for each range of values represents the percentage of transactions, which have completed within the particular range. Figure 4-59 on page 207 is an example of detailed information displayed for the selected workload group chart.

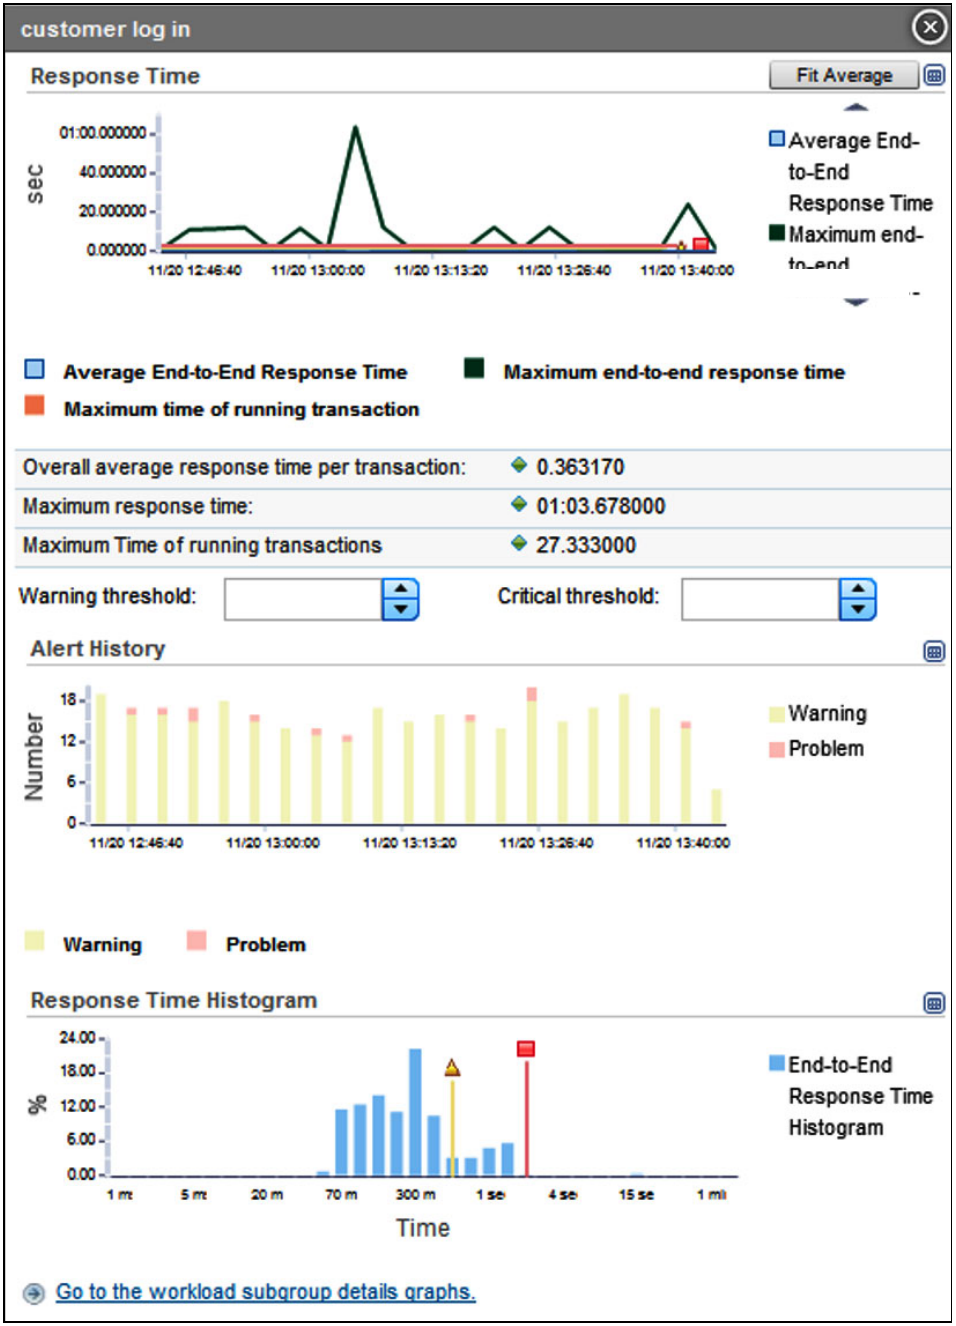

Figure 4-59 Detailed information for the selected workload group chart

4.5.2 Extended Insight Analysis Dashboard: Details panel

Use the Extended Insight Analysis Dashboard details view to locate the source of performance problems, determine how those problems affect various parts of the workload, and analyze the performance of individual SQL statements, clients, and partitions.

To view details panel, double-click the workload group or workload subgroup in the Overview panel. Figure 4-60 on page 208 shows an Extended Insight Analysis Dashboard details panel.

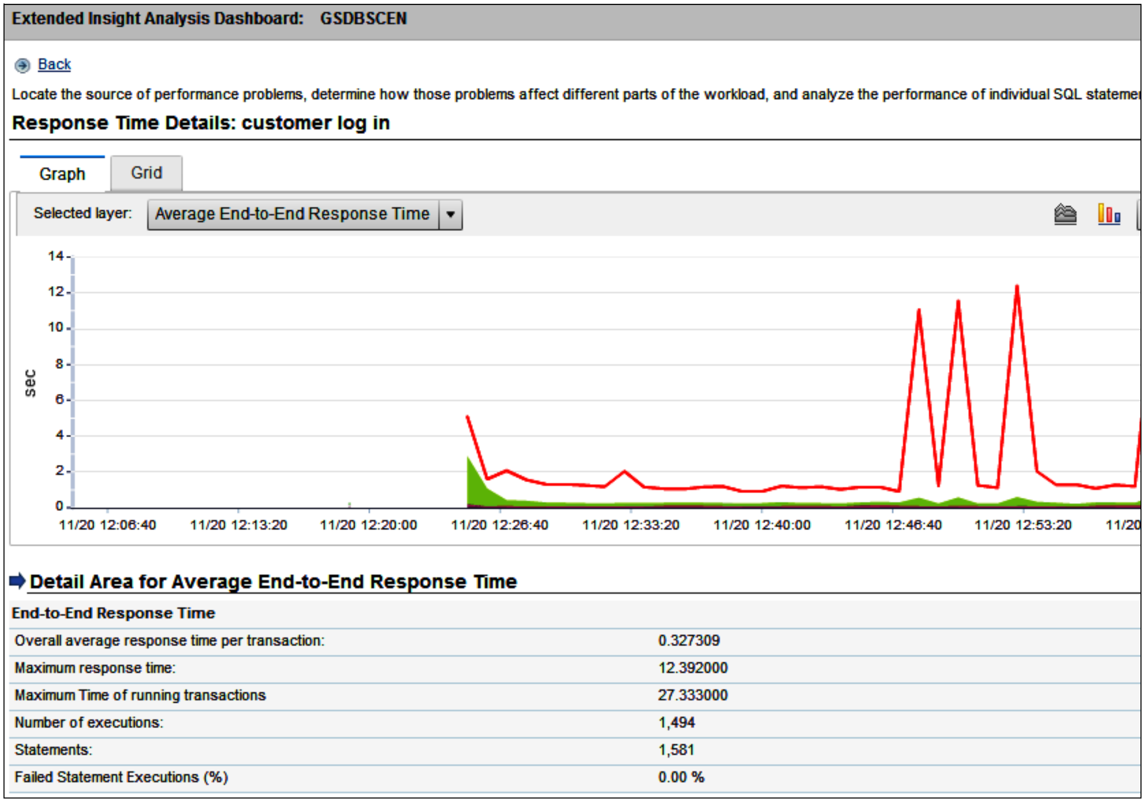

Figure 4-60 Extended Insight Analysis Dashboard details panel

The Details panel contains the following tabs:

- Graph tab
- SQL Statements tab
- Clients tab

Graph tab

The Graph tab shows response times for the components of the workload group, workload subgroup, or database for the time frame that is selected on the time slider. The graph shows time spent in each layer for each component, such as the data server, client, and network. Use the **Selected layer** drop-down list or click a layer in the graph to view details for that layer in the bottom pane of the dashboard.

Figure 4-60 on page 208 is an example of a details panel with the Graph tab for the selected layer of Average End-to-End Response time per transaction. The details area of this panel displays more information relevant to the selected layer.

SQL Statements tab

The SQL Statements tab lists the SQL statements that were run by the transactions of the workload group, workload subgroup, or database that completed their work in the time frame that is selected on the time slider. Select a statement from the list to view details for that statement in the details area of the dashboard. The details area consists of two tabs:

- DB Client Information tab

  This metrics presented in this tab show statement execution details from the client application perspective and is therefore collected on the application side. This tab shows the general SQL statement information such as statement text, end-to-end time-spent metrics. You can use this information to determine whether a statement spent most time on the application side, in the network, or on the data server side. The statements presented are statements that finished within the specified time frame. The statement metrics are aggregated over all executions.

Figure 4-61 on page 210 shows a partial view of the details panel with the SQL statement general information.

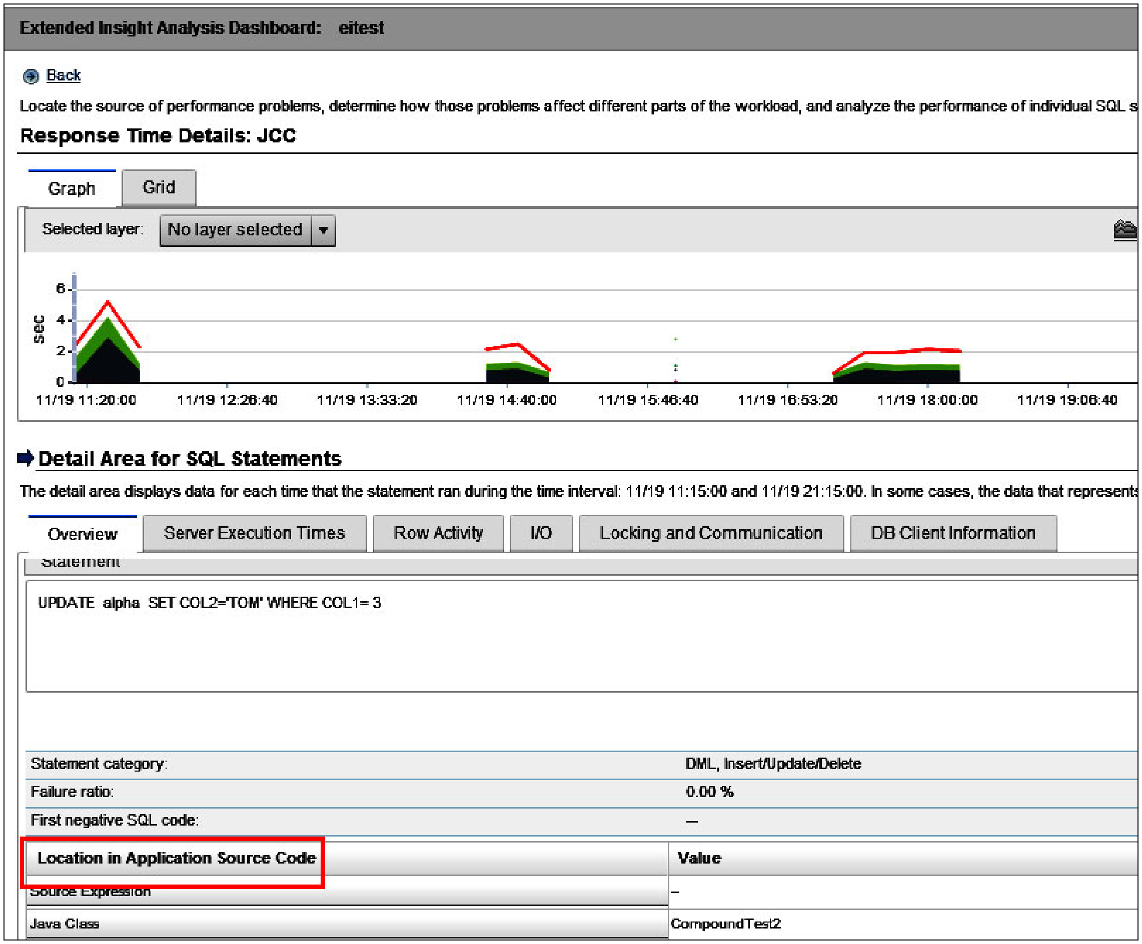

Figure 4-61 Details panel, SQL statement general information tab

- Execution Statistics tab

  The Execution Statistics tab shows the details of this SQL statement from a data server execution perspective. Figure 4-62 on page 211 shows a partial view of the average execution information across all executions of this statement in the selected time frame, such as for row and sort performance, I/O statistics, locking statistics, and data server execution times. In some cases, this average data can consist of more statement executions than executed by the workload group or subgroup that you selected on the Extended Insight overview panel.

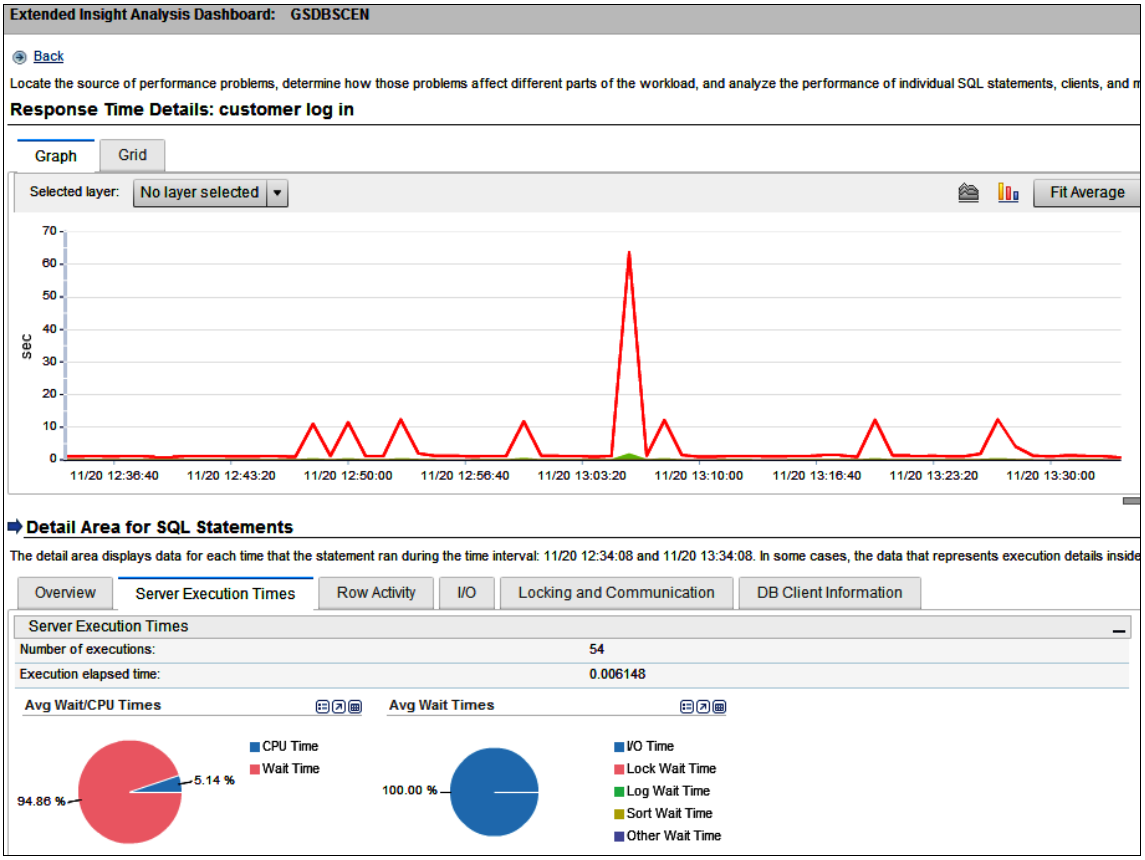

Figure 4-62 Details panel, SQL statement data server execution statistics tab

Metrics: Statement server execution details are available for DB2 V9.7 Fix Pack 1 (or later) databases. This also requires that during the InfoSphere Optim Performance Manager configuration process of the monitored database, you have selected to collect statement metrics on the data server.

Clients tab

The Clients tab (Figure 4-63 on page 212 is a partial view) lists the clients that ran the transactions for the workload group, or database for the time frame that is selected on the time slider. Select a client from the list to view the details for that client in the bottom pane of the dashboard.

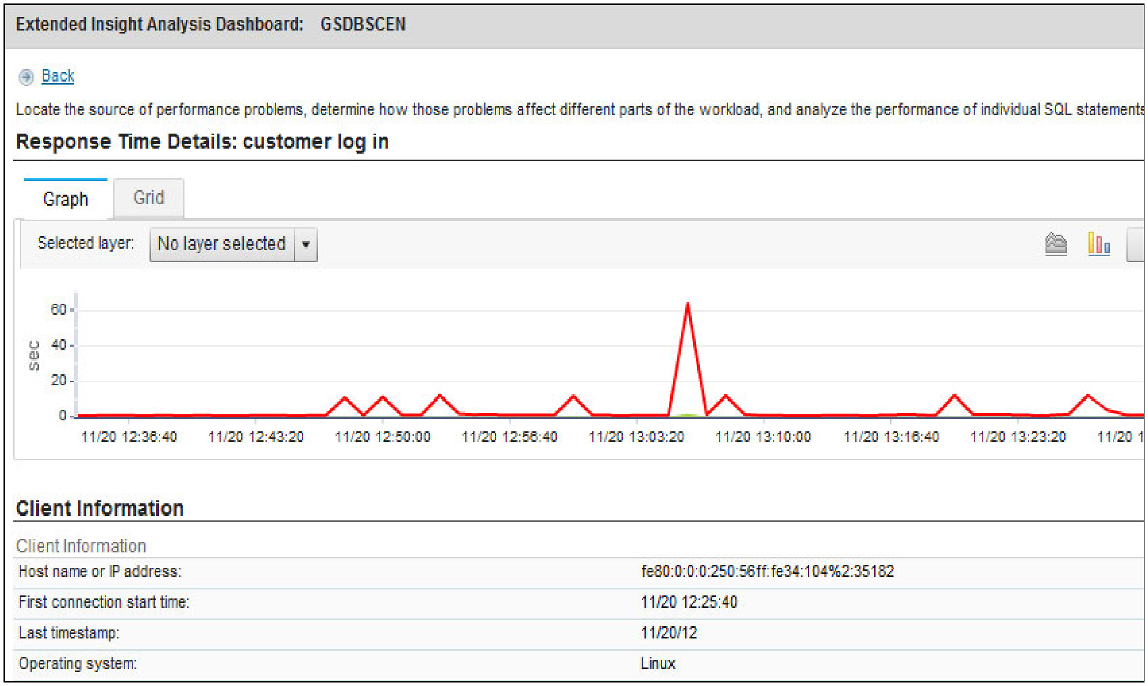

Figure 4-63 Details panel - Clients tab

4.5.3 Workload groups and subgroups

Determining where and when performance problems and bottlenecks occurred can be difficult. To isolate the source of performance problems, you can group and monitor transactions that come from specific component layers by using workload groups.

Extended Insight monitoring is based on predefined or user-defined workload groups that contain workload subgroups. The grouping of transactions into workload groups is based on the connection attributes that are set for a connection to the monitored database. The grouping can help you determine where performance problems and bottlenecks occur. You can also associate workload subgroups with alerts by setting thresholds; this will be displayed in the Workloads with Alerts tab during the selected time period.

For example, in the predefined workload group named *Client user IDs,* you get one workload subgroup for each user ID that is set in the corresponding connection attribute. In each workload subgroup of this group, all the transactions that originate from the same user ID are grouped together. The performance metrics are also aggregated within the group.

Connection attributes represent client information fields that are used at the DB2 server for determining the identity of the application or user currently using the connection. InfoSphere Optim Performance Manager users can build workload groups based on these connection attributes.

You can define values for the following connection attributes:

- Authorization ID
- Client user ID
- Client application name
- Client workstation
- Client accounting string
- IP address (DB2 Version 9.7 and later)
- Application type

Connection attributes can be set either at the database server, or at the application or application server level. For detailed information about how to set these attributes, see the DB2 information center:

http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp

You can activate or deactivate a workload group for monitoring at any time from Extended Insight Analysis Dashboard by selecting **Activate** or **Deactivate** from the **More Actions** drop-down list. When you activate a workload group for monitoring, you can view detailed data for all transactions in the group and for each of the workload subgroup in the group. Workload group definitions do not affect data collection. Workload groups can be configured for historical data collection also. Activation or deactivation process does not stop the collection of database performance data. It only affects the aggregation of the data on the dashboard.

You can create workload groups specific to your needs or use the predefined workload groups that are already activated. InfoSphere Optim Performance Manager EI provides specific support for monitoring a set of common application frameworks:

Click **More Actions** and select the workload group to activate (Figure 4-64 on page 214).

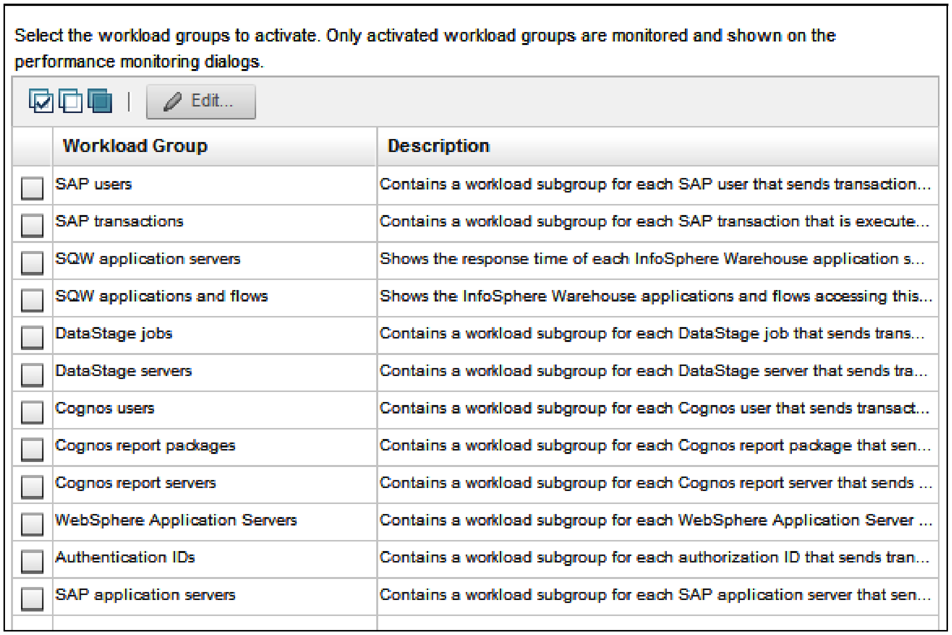

Figure 4-64 Activate predefined workload group

Consider the following examples:

- SAP: Monitoring response times per SAP user, SAP transaction, SAP source module and SAP application server
- Cognos: Monitoring response times per Cognos report, Cognos, report user, Cognos report package, and Cognos user system and Cognos server
- DataStage: Monitoring response times per DataStage job, DataStage job user, DataStage project, and DataStage server
- InfoSphere Warehouse

To create a new workload group, click **New** on the Extended Insight Analysis Dashboard.

Figure 4-65 on page 215 shows a partial view of the *Create new workload group* panel. This example will create a new workload group named global_user, which will contain database application performance data from a client user ID, `g user`. On the panel, provide a name for the new workload group name, and click **Next**.

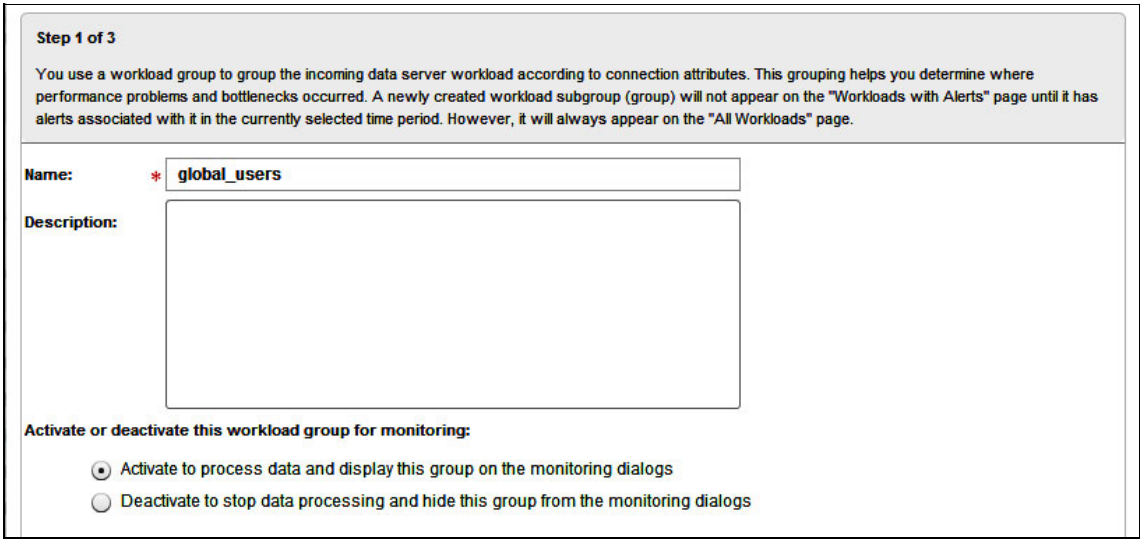

Figure 4-65 Create new workload group panel

Figure 4-66 on page 216 shows the *New workload group details* panel. In this panel, select the connection attribute for this workload group (Client User ID) and the filtering condition, which will display data only for the specific Client User ID (a user).

Figure 4-66 New workload group details

Figure 4-67 on page 217 shows the *New workload group threshold settings* panel. You can specify response time thresholds for the entire workload group or for individual workload subgroups. When the workload group is activated for monitoring, you are informed if thresholds are violated. Click **Finish**.

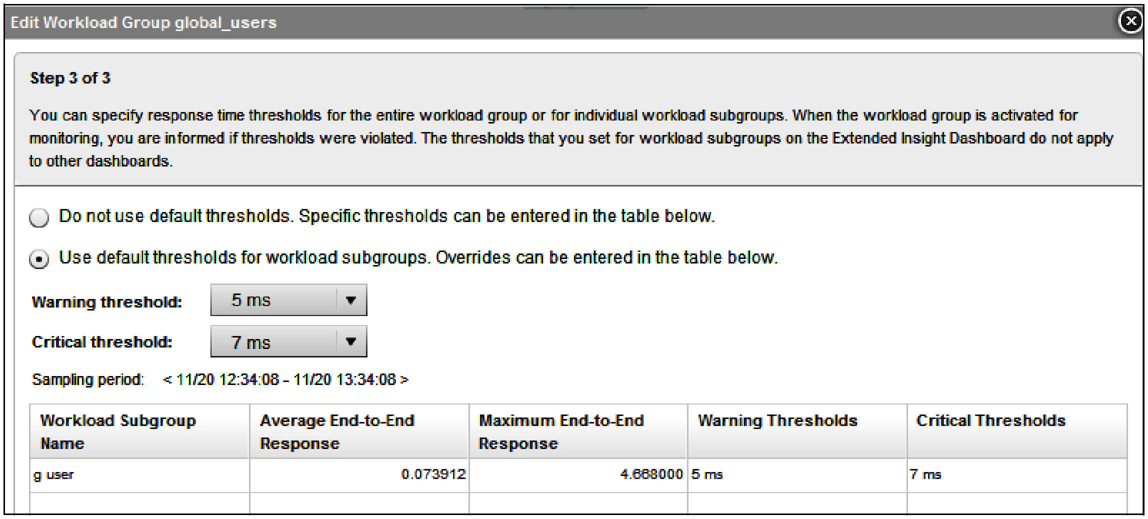

Figure 4-67 New workload group threshold settings

After the new workload group is defined and activated, you can view its database performance metrics on the Extended Insight Analysis Dashboard.

Figure 4-68 on page 218 shows a partial view of the Extended Insight Analysis Dashboard with the new workload group.

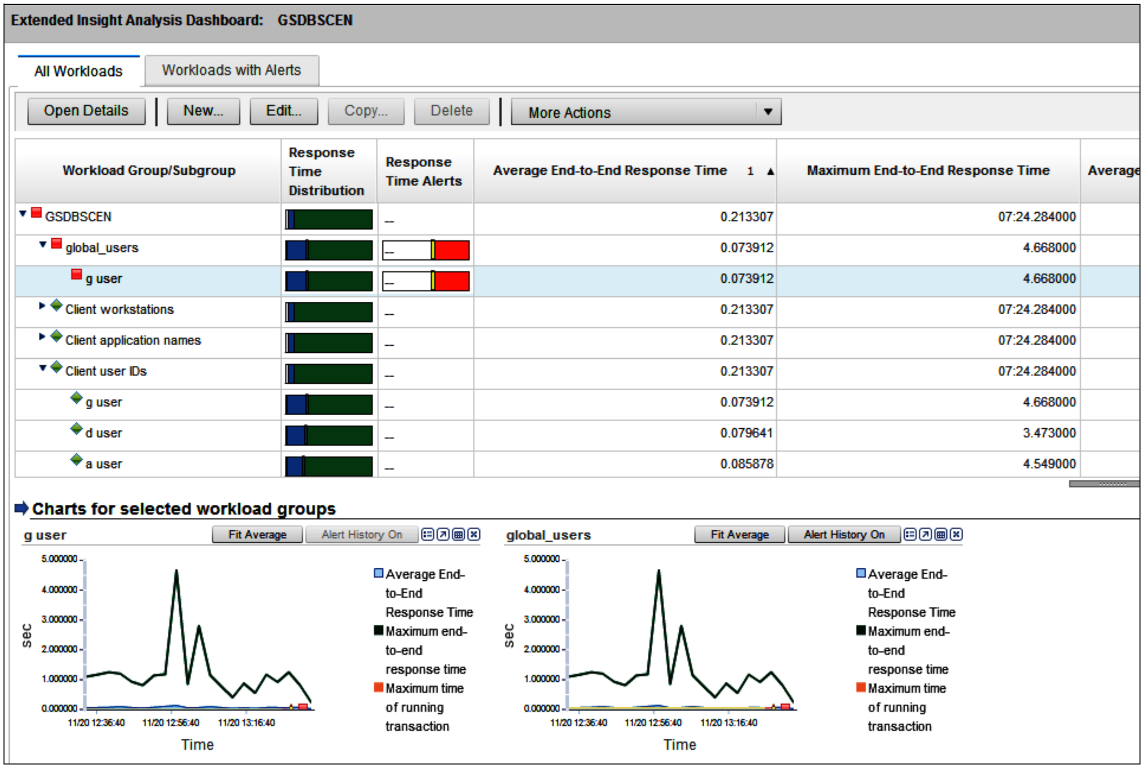

Figure 4-68 Extended Insight Analysis Dashboard with the new workload group

4.5.4 InfoSphere Optim pureQuery Runtime integration

An important step in tuning SQL is to determine the source of the SQL so it can be modified. This can be difficult, especially if the SQL was generated by a third-party framework such as Hibernate or JPA. To help identify the source code, the Extended Insight Analysis Dashboard can display the InfoSphere Optim pureQuery for Linux, UNIX, and Windows metadata, such as Java class, package, application name, method name, and source line number, as shown in Figure 4-69 on page 219.

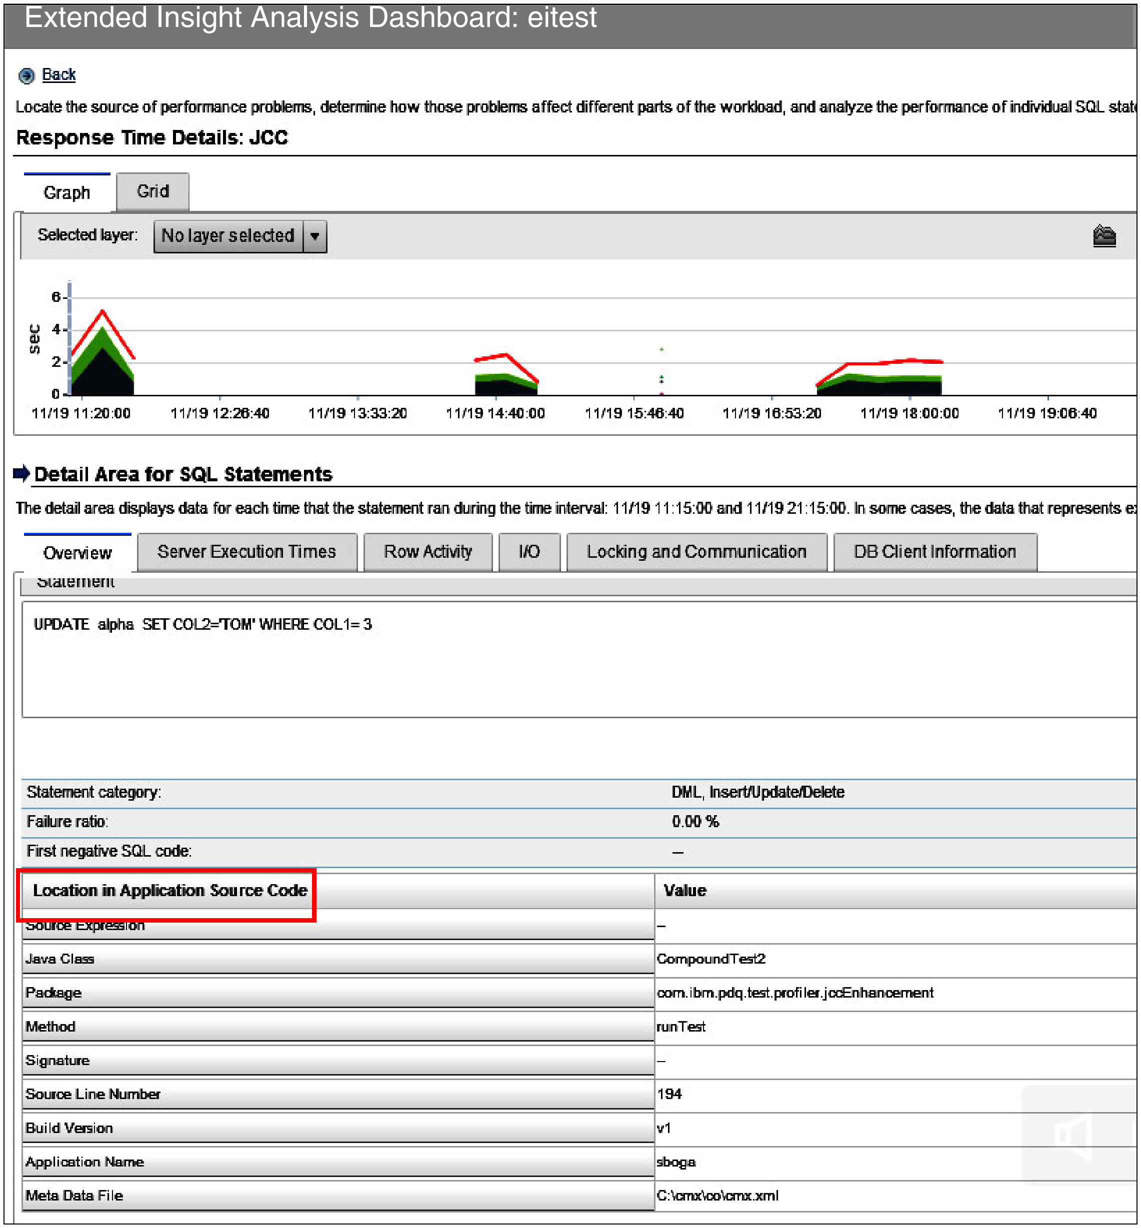

Figure 4-69 Integration of InfoSphere Optim pureQuery Runtime metadata

With this feature, the database administrator and the developer can collaborate by quickly identifying the source SQL. This feature requires a license for the InfoSphere Optim pureQuery Runtime product.

For more information about the product, see *Using Integrated Data Management To Meet Service Level Objectives*, SG24-7769.

4.6 Reporting with InfoSphere Optim Performance Manager

Because of its specific nature, information about reporting with InfoSphere Optim Performance Manager has been placed in a separate PDF document titled "Additional Material for Performance Management Using IBM InfoSphere Optim Performance Manager and Query Workload Tuner." However, be aware that although we believe it is accurate and helpful information, that material has not been subjected to the formal ITSO editing process. That material is Chapter 1 of the PDF, and titled "Reporting with InfoSphere Optim Performance Manager." To access that material, see section "Using the web material" on page 384.

5

Getting to know InfoSphere Optim Query Workload Tuner

This chapter introduces the basic features in InfoSphere Optim Query Workload Tuner, and how to use those features to assist the SQL analysis and performance tuning. It describes the workflow and best practices for how to analyze and tune SQL statements using InfoSphere Optim Query Workload Tuner. The chapter shares the overall steps to use this product and to speed the query analysis and tuning process. The chapter also examines how to find the most interesting SQL statements to tune and how to choose the tuning tools and advisors with options to get the results. Finally, it reviews the advisor results (such as the statistics and index advisors) and the use of query-related tools (such as the access plan tools, format, annotation, and reports).

5.1 Workflow for query analysis and tuning

When you do the query analysis and tuning with InfoSphere Optim Query Workload Tuner, typically four steps are involved (as Figure 5-1 shows):

1. Capture the interesting SQL statements to analyze and tune.
2. Analyze the SQL statements to understand how they look, the DB2 explain output, and how to execute the SQL statements.
3. Understand and take the proper action to resolve any performance bottlenecks of the SQL statements.
4. Validate whether the resolution works.

However, understand that the query analysis and tuning workflow is not a one-time action, but an ongoing and iterative process.

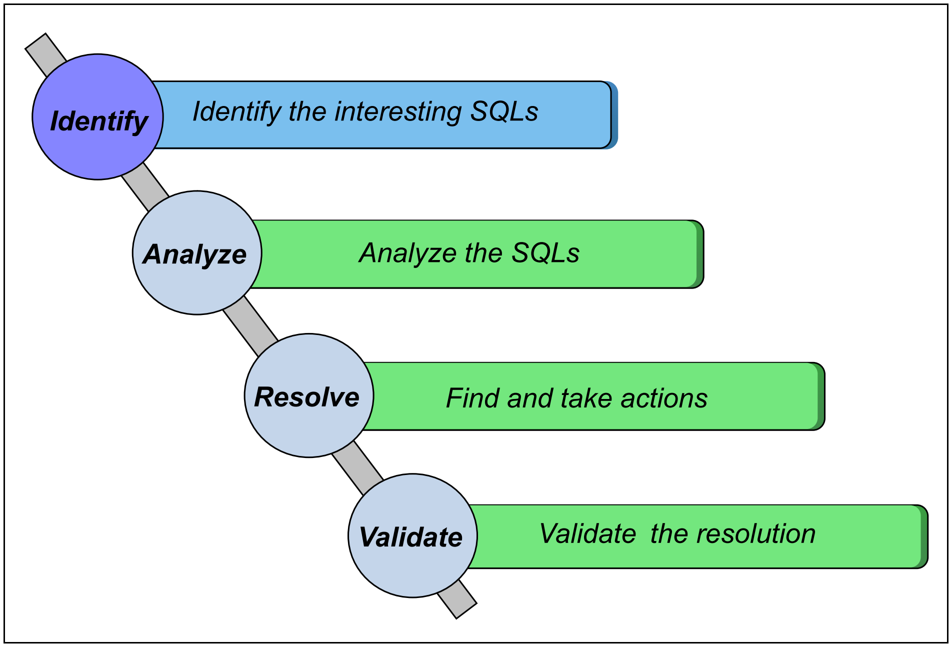

Figure 5-1 Query analysis and tuning workflow

The four steps are explained next:

1. Identity the interesting SQL statements for tuning.

   Get the interesting SQL statements to analyze and tune during the database application development, and from database performance monitoring, database catalog tables, user input, and so forth. You can identify one single SQL or a set of related SQL statements, and then navigate to InfoSphere Optim Query Workload Tuner for tuning. You can use a set of attributes to identify the SQL statements such as highest elapse time, CPU time, I/O time or sort overflow, table scans, and package names. See 5.3, "Capture SQL statements from different sources" on page 225 for further details of how to capture the interesting SQL statements.

2. Analyze the SQL statements.

   After you have the SQL statements, you can analyze them to know what types they are and how they will be explained and executed in the DB2 engine. InfoSphere Optim Query Workload Tuner provides the query format and annotation to help you learn the types of SQL statements; you can use the related access plan (the access plan graph and access plan explorer as examples) tools to know how the DB2 engine will execute the SQL.

3. Find the SQL statements and take actions to tune them.

   Find the solutions that can improve performance of the SQL statements, and take actions to implement those improvements. As examples, you can rewrite the SQL statements that do not comply with the company best practices, collect corresponding statistics so the optimizer can make more accurate choices, create indexes to avoid table scans and sorts, or apply plan hints through the optimization profile. InfoSphere Optim Query Workload Tuner provides the built-in expertise to offer the recommendation on those solutions. Then, you can review, make changes, and deploy those recommendations; you may also find the solutions by using your own expertise.

4. Validate the solution.

   You must validate the solutions that have been selected to make sure the actions really improve the performance of the SQL statements. This can be done statically and dynamically.

   – Statically, you can use the InfoSphere Optim Query Workload Tuner access plan comparison to see the changes of the estimated cost, and the access plan before and after the SQL statements tuning.

   – Dynamically, you can use the performance monitoring product, InfoSphere Optim Performance Manager, to see the runtime performance improvement of the SQL statements before and after the tuning. InfoSphere Optim Performance Manager provides an SQL baseline report to provide a detailed runtime metrics comparison.

The four steps are iterative, and on-going. If in the last step, through the validation, you find the performance improvement is not as much as expected, you might need to go through the workflow again to tune the SQL statements, and also the database application runs. As the database data grows, some SQL statements might perform worse; then, you can use this workflow again to refresh the SQL statements tuning.

InfoSphere Optim Query Workload Tuner provides a guided SQL tuning process using the workflow just described. There is a Query Tuner workflow assistant to indicate which steps you have performed, and which steps are next. The query tuning perspective (at the left-most side) has the mini-navigation bar (Figure 5-2 on page 227) with the following six steps:

1. Status
2. Capture
3. Manage
4. Invoke
5. Review
6. Compare

5.2 Entry point of InfoSphere Optim Query Workload Tuner

There are many places you can go in InfoSphere Optim Query Workload Tuner to start SQL statements tuning. As described in 1.2.6, "IBM Data Studio V3.2" on page 13, it is part of IBM Data Studio, and you can launch it inside or outside Data Studio.

Entry points of query tuning inside Data Studio are as follows:

- In Data Source Explorer, choose the connection or database, right-click, and start tuning.
- In the SQL and Routine editor, choose the SQL statements, right-click, and start tuning.
- In Data Source Explorer, drill down to view triggers or user-defined functions, right-click, and start tuning.
- In Data Source Explorer, drill down to a stored procedure, right-click, and start tuning.
- In Administration Explorer, drill down to one database, if it is connected, right-click, and start tuning.

- If you have previous tuning projects, go to InfoSphere Optim Query Workload Tuner in the data project by opening a query tuning project.
- A new Query Tuner Project can also navigate you to the query tuning perspective.

Furthermore, you can start the query tuning outside Data Studio. In the InfoSphere Optim Performance Manager web user interface (UI), you can choose one or several SQL statements, and click **Tune** to launch InfoSphere Optim Query Workload Tuner for query tuning. In this case, be sure the InfoSphere Optim Performance Manager web browser and Data Studio are in the same client, and that the Data Studio is running with an embedded server.

Launch points of InfoSphere Optim Query Workload Tuner from the InfoSphere Optim Performance Manager web UI are as follows:

- In SQL dashboard, you can choose one SQL or a set of SQL statements in the Top Individual Execution tab or the Execution Summary tab, and start tuning.
- In Extended Insight (EI) dashboard, in the details tab, you can choose a set of the top statements within one workload cluster, and start tuning.
- In the Locking Analysis dashboard, you can analyze the lock event (deadlock, lock time-out or lock wait) in the statement section, and you can click **Tune** to start tuning.

5.3 Capture SQL statements from different sources

You can capture from a lot of sources to get the SQL statements to tune, when you navigate to InfoSphere Optim Query Workload Tuner. In the IBM Query Tuning perspective mini-navigation, click **2.Capture**. You see the choice to capture from different sources. Details for each source type are at the following address:

http://pic.dhe.ibm.com/infocenter/dstudio/v3r2/index.jsp?topic=%2Fcom.ibm.datatools.qrytune.workloadtunedb2luw.doc%2Ftopics%2Flocations.html

Several special captured sources are as follows:

- InfoSphere Optim Performance Manager repository

  Be sure you have the target database monitored in InfoSphere Optim Performance Manager with Extended Insight enabled, and the corresponding InfoSphere Optim Performance Manager repository connection has been configured in InfoSphere Optim Query Workload Tuner. See 3.4.8, "Configuration to capture from InfoSphere Optim Performance Manager repository" on page 139 for details.

- Package Cache

  For DB2 v9.5, be sure the STATEMENT monitor switch is turned on. For DB2 v9.7 and later, be sure MON_ACT_METRICS is set to BASE or EXTENDED. Otherwise, you might cannot use the runtime metrics filter to determine the interesting SQL statements. In the Package Cache source, you can capture with a scheduler, and then do the recurring capture.

- Event Monitor Tables

  Make sure you create the activities event monitor, and the event monitor write to the normal tables (not to a file or un-formatted raw table). You can create the activities event monitor through scripts, or use WLM to create and trigger the monitor to capture the details of the activities. Another suggestion is to prune the data in the event monitor tables after you use it and do not need it any more.

> **Tip:** Ensure the database you will be tuning has been activated by the InfoSphere Optim Query Workload Tuner license, and has been configured for query tuning. See 3.4, "Installing and configuring IBM InfoSphere Optim Query Workload Tuner" on page 127 for the activation and configuration. Furthermore, you must have the correct privileges to do the capturing. See the following address for the required privileges:
>
> `http://ibm.co/18v8NKM`

In general, you can define a filter for most of the sources. A filter helps you determine the interesting SQL statements for tuning, and it contains the conditions (where), sort columns (order by), and display columns (select). Some sources have predefined filters (such as package cache), you can create your own filters, and edit an existed filter.

After you define the filter, InfoSphere Optim Query Workload Tuner uses the filter to capture and process from the sources, and then render the qualified statements in the Captured Statements sections. You can then select one SQL for single-query tuning, or save all the statements to a workload. A workload is a set of related SQL statements and the runtime information. You can also filter on an existing workload to create a new workload using the post-filter.

Next is an example of capturing from a package of the target database. In the Data Source Explorer, right-click the target database and select **Analyze and Tune**, then select **Start Tuning**.

The InfoSphere Optim Query Workload Tuner main UI opens (Figure 5-2). In the mini-navigation bar, click **2. Capture**. A list of the sources is displayed. Select **Package**, and on the right side, you can define a filter by clicking **New**.

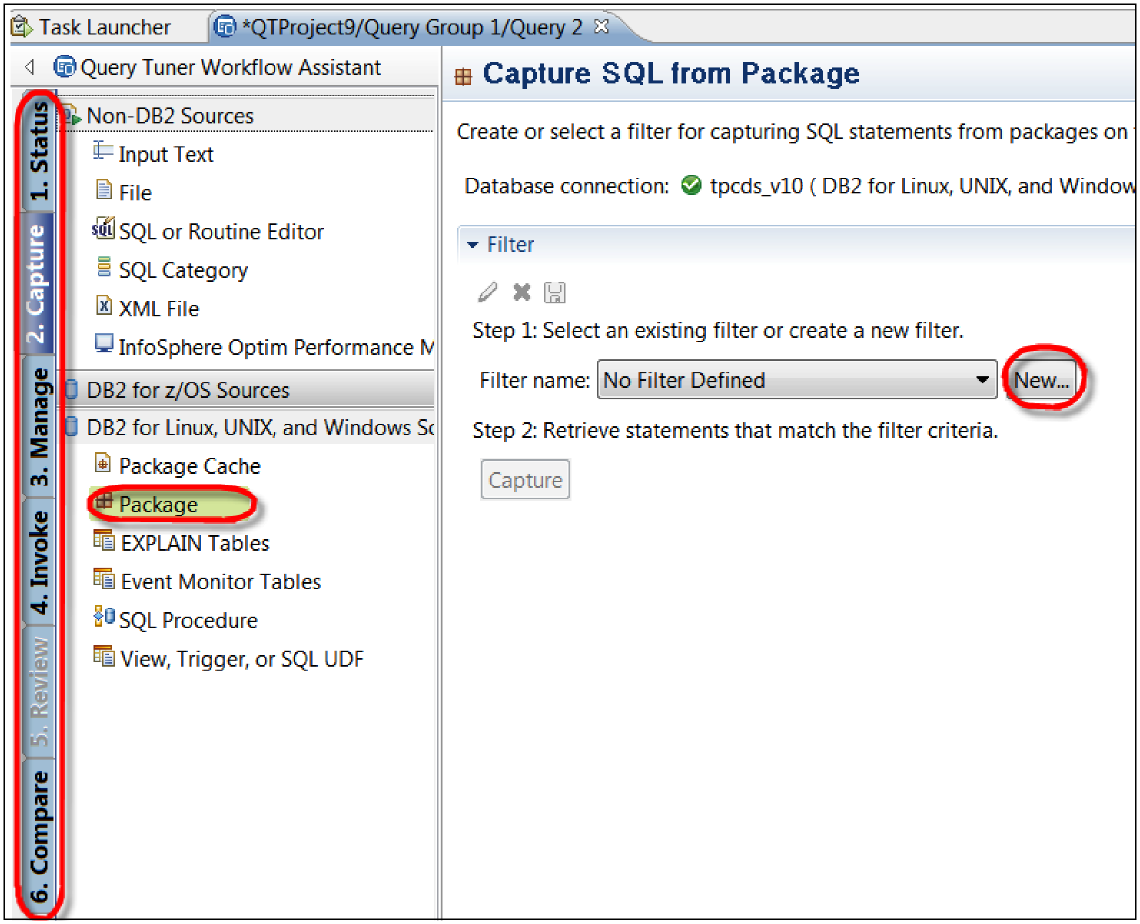

Figure 5-2 InfoSphere Optim Query Workload Tuner Main Panel and Capture from Package

You can specify the filter name, maximum number of statements, the filter condition, and filter by referenced objects (Figure 5-3). This example captures from Package (the package name includes WAPC) and also the statements that reference the TPCDS.CUSTOMER table. In the next step, you can specify sort columns and display columns you want to see in the captured statements section, save the filter, and then start the capture.

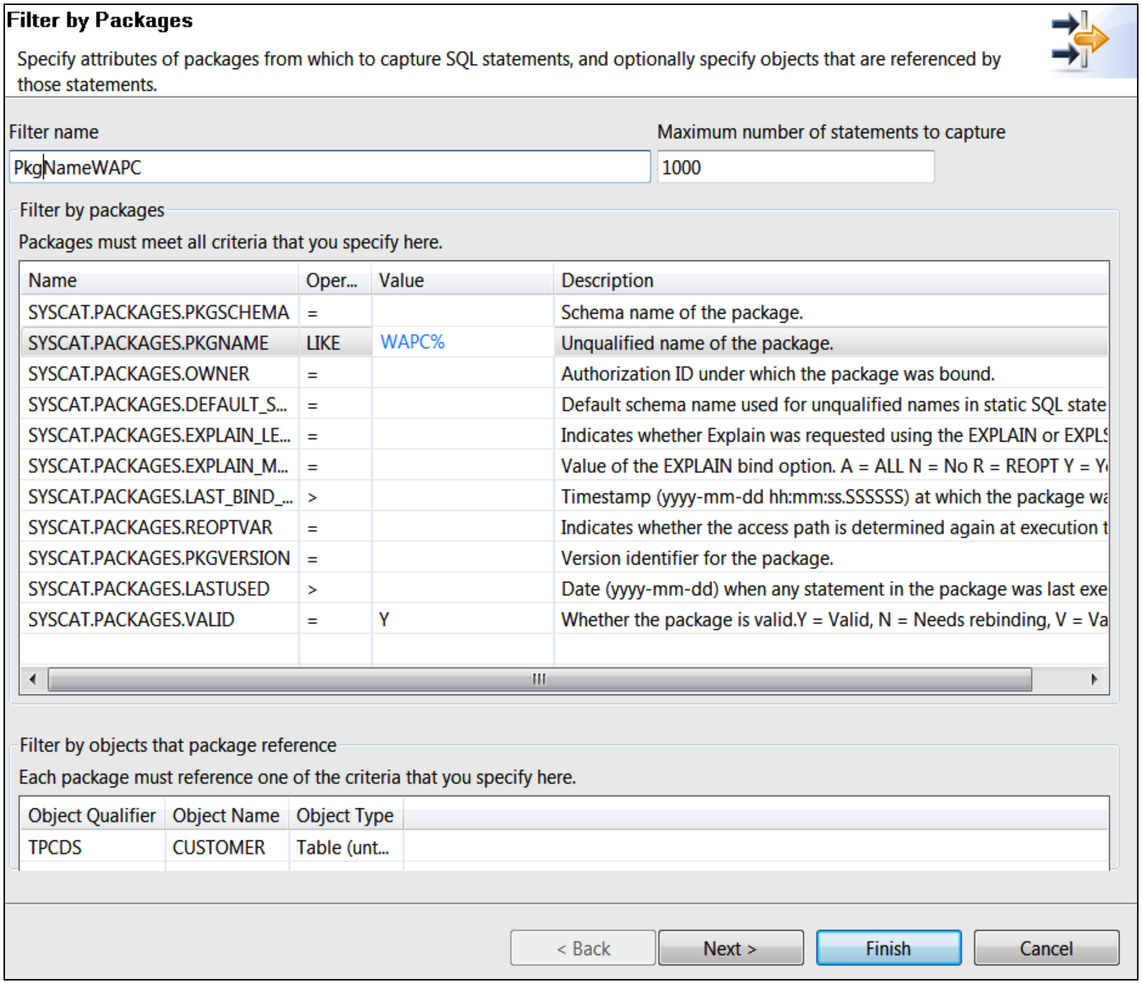

Figure 5-3 Define capture filter

When the capture finishes, you see the SQL statements in the Captured Statements section in the panel (Figure 5-4). You can review the captured statements, show the full SQL text for a specific SQL, choose one single query to start the analysis and tuning, and search in the captured SQL statements to find the interesting SQL statements, as examples. You can also save the captured statements into a workload. The InfoSphere Optim Query Workload Tuner will save the workload information to the repository on the target database side for future analysis. Save the captured statements to a workload.

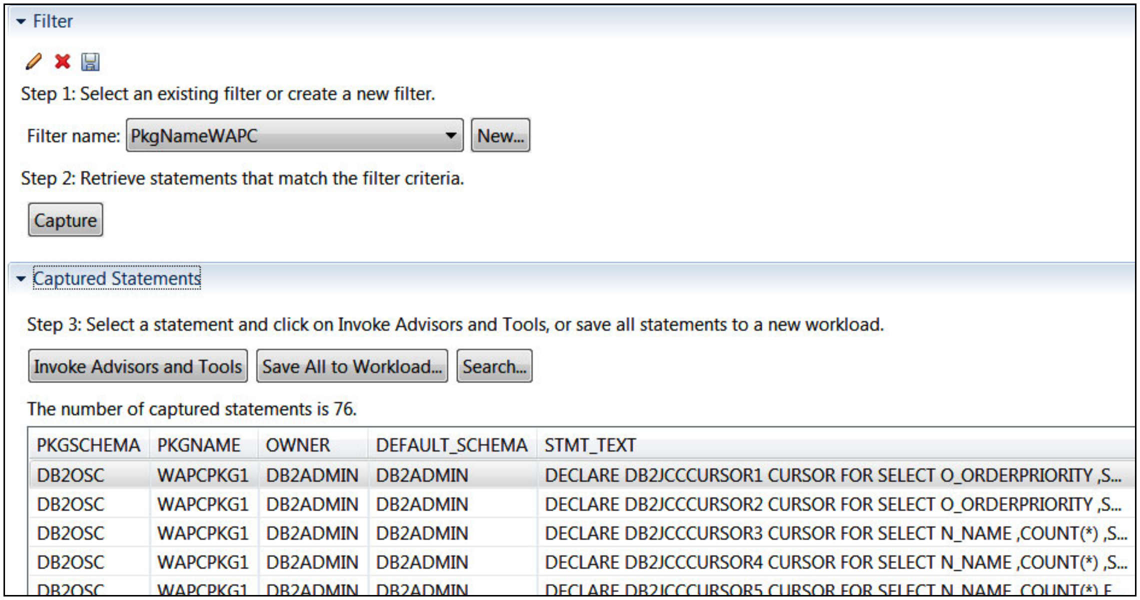

Figure 5-4 Captured Statements

After you save the captured statements to a workload, you are ready for the *Manage* step, where you can manage the workload (Figure 5-5). You can list the workloads on the target database, show the statements for each workload, invoke the advisor of the workload, and explain the workload using the current statistics information. When you drill down to one specific workload, you can refine the workload using the post filter to create a new workload using the runtime metrics or access plan operations such as table scan, index scan, and join types.

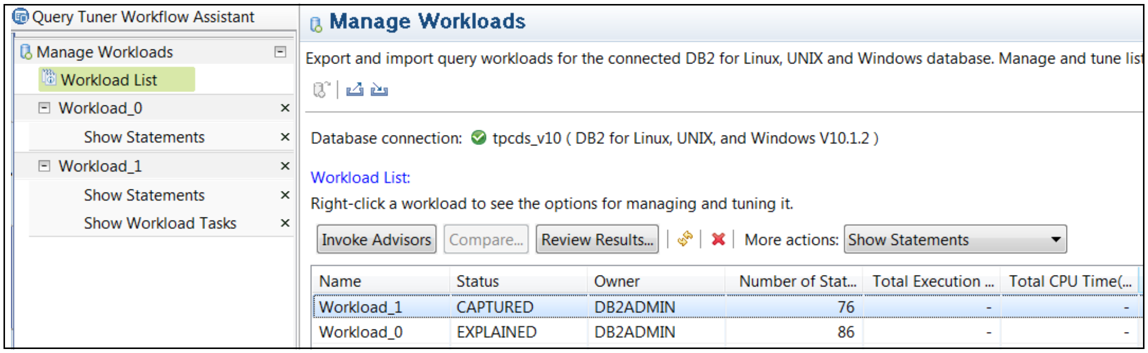

Figure 5-5 Manage Workloads

5.4 Run the tuning tools and advisors

After you get the interesting SQL statements, which are either a single SQL or a set of SQL statements (workload), the next step is to invoke query analysis advisors and tools to understand the SQL statements and the solutions to improve the performance.

For the single query, you can choose that query and click **Invoke Advisors and Tools** to go to the Run Single Query Advisors and Tools panel (Figure 5-6 on page 231). You can make several settings before running the advisors and tools, such as the SQL schema, add description, and whether to re-explain the query and related options. If the SQL is not fully qualified, InfoSphere Optim Query Workload Tuner will use the schema specified here to explain and analyze the query. If you do not re-explain, it will try to use the existing explain information to analyze the query. If InfoSphere Optim Query Workload Tuner cannot find the explain information for this SQL, then it will re-explain it. If you choose to re-explain the query, and also the related options, then it will use the explain setting and current database statistics information to explain and analyze the SQL. You can also format the SQL using the *Format SQL* feature, where InfoSphere Optim Query Workload Tuner will format the SQL to be tuned in a separate window, so you can also copy the formatted SQL text.

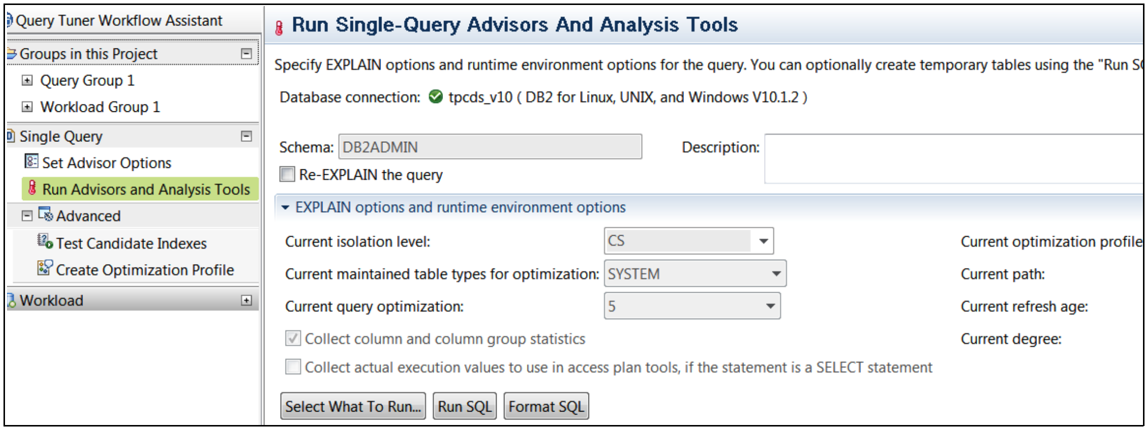

Figure 5-6 Run Single Query Advisor and Tools

Before running the advisors, you can also set the advisor options by clicking **Set Advisor Options** at the left side query tuner workflow assistant. In the advisor options, you can set both global and current preferences. If you intend to set only the options for the current running advisor, set it and then run the advisors. If you intend to set the options for all other tuning activities, then you can set the options and save to global preferences. You can also restore the setting from global preferences. As described in 3.4.9, “InfoSphere Optim Query Workload Tuner client preference settings” on page 142, you can also set the options in the product preferences settings (**Windows** → **Preferences** → **Data Management** → **Query Tuner**).

After you make all the settings, select the activities to run by clicking **Select What to Run**.

In the select activities panel (Figure 5-7), choose the activities you want to run to tune the SQL, including three sections:

- Analysis tools, to format, annotate the SQL, and generate visual presentation of the SQL
- Advisors, use InfoSphere Optim Query Workload Tuner built-in expertise to generate the recommendations on statistics, query revision, index and access path
- Reports, to give a tuning summary in HTML format for review, and share.

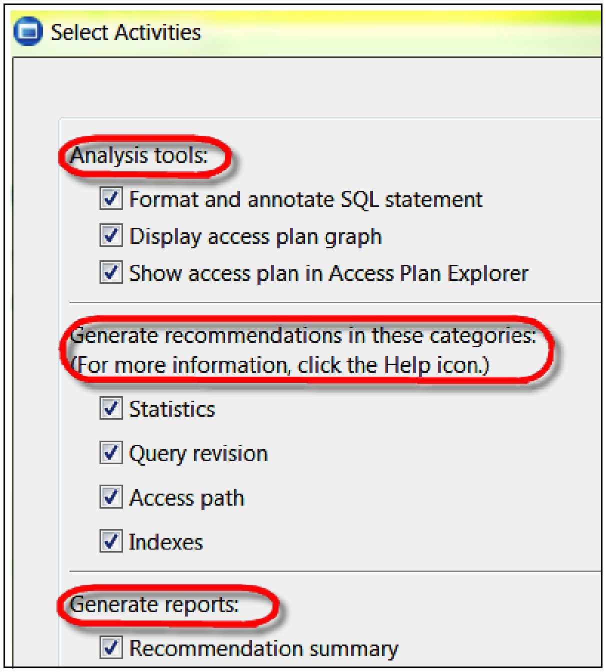

Figure 5-7 Select Activities

The more you choose, the more detailed information and recommendations you get from InfoSphere Optim Query Workload Tuner, and which requires a longer time to finish the job. You can get the explanation for each activity using the question marker at the bottom left. It embeds the information center with context information, make sure your client has Internet access to the information center, if you do not install the information center locally.

After you select activities you want to run, click **OK** to start the tuning job. A progress window shows the current progress (Figure 5-8) and which step. You can put this task into background mode, or cancel it. If you put the task into background mode, you can proceed with other tasks in the Data Studio. After the tuning is finished, InfoSphere Optim Query Workload Tuner will render the results in its panel.

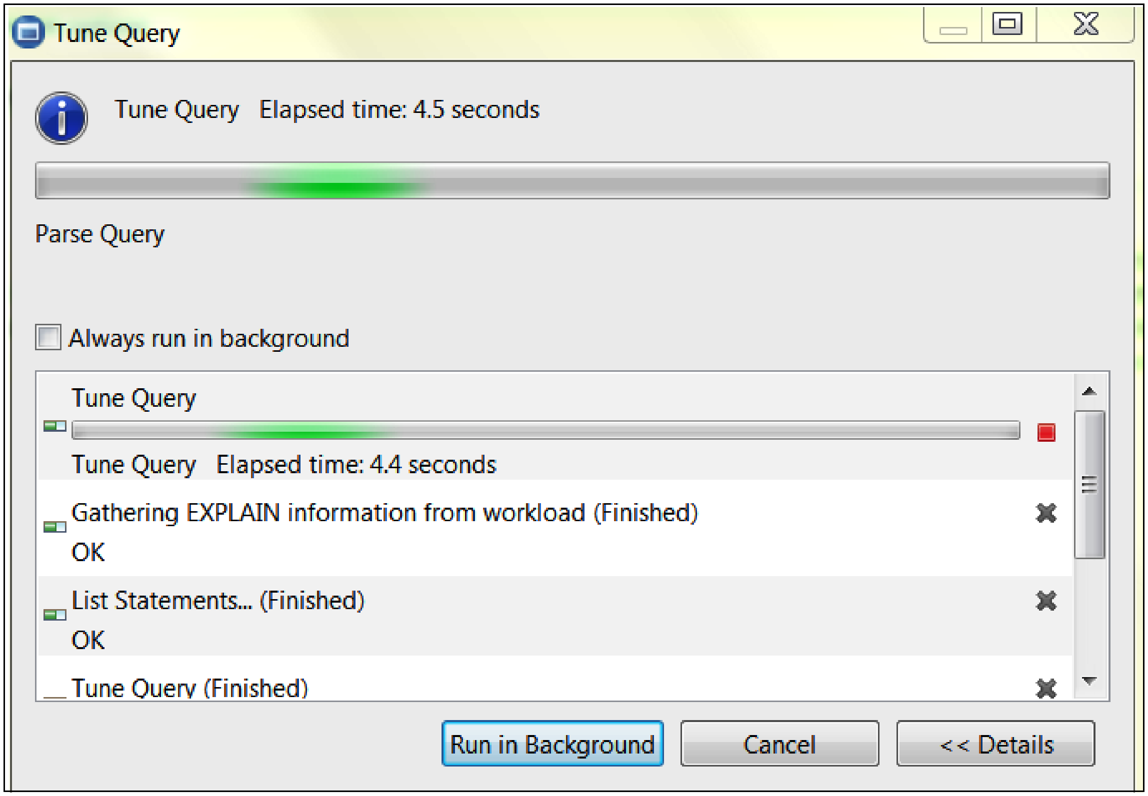

Figure 5-8 Tuning Task Progress window

When the tuning task is finished, the Single Query recommendations page opens (Figure 5-9). Use it to review the recommendations InfoSphere Optim Query Workload Tuner provides for that query, and take proper actions. By default, you see the analysis results summary, which lists the recommendations for this query; click each recommendation type in the table to view the details. You can also navigate to other tools and report from the left side.The details for types of analysis tools and advisors are described later.

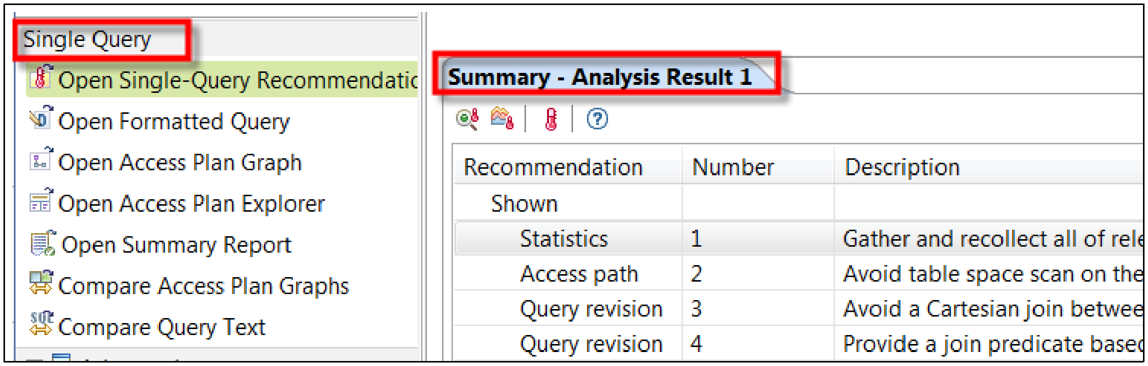

Figure 5-9 Single Query Tuning Results

Workload tuning follows a similar process. You choose an interesting workload in the manage workload step, and then invoke the run advisors. InfoSphere Optim Query Workload Tuner opens the Run Workload Advisors panel (Figure 5-10). In the panel, you can see the status of the workload, add a description of the workload, and choose whether to re-explain the workload for this tuning task. In the Workload Statements section of this panel, you can also view the statements and related runtime metrics in this workload.

Note: InfoSphere Optim Query Workload Tuner will save each set of explain information to its own repository on the target database side for each workload. If you choose not to re-explain the workload for tuning, InfoSphere Optim Query Workload Tuner will use the latest version of the explain information to analyze the workload. If there is no previous explain information, InfoSphere Optim Query Workload Tuner explains the workload before running the tuning activities.

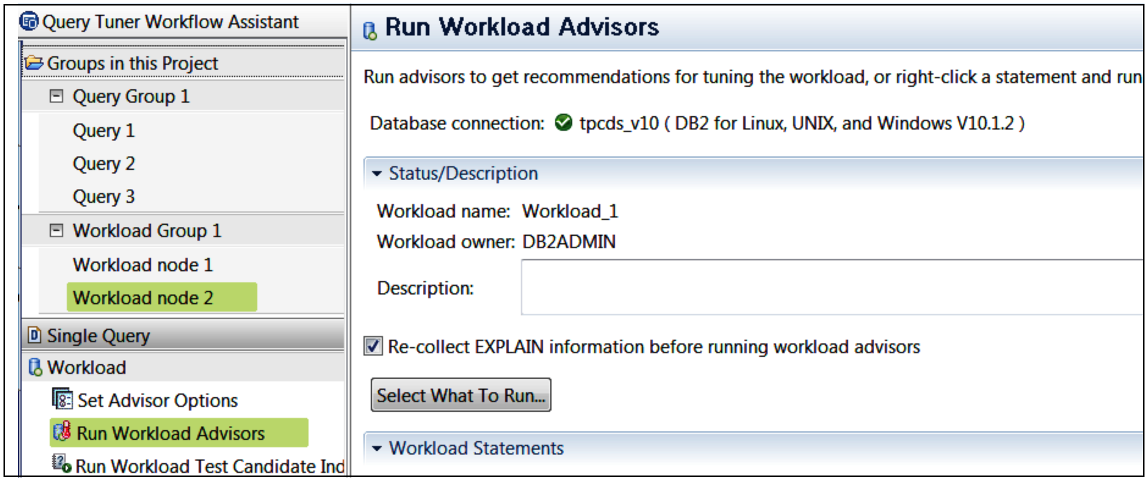

Figure 5-10 Run Workload Advisors Panel

Click **Select What To Run**. In the Select Activities dialog (Figure 5-11), choose the recommendations on statistics, statistical view, indexes, MQT, MDC, and partition key, and the tuning summary. You can also get the detailed explanation for each type of recommendation in context help by clicking the question mark (**?**) at bottom.

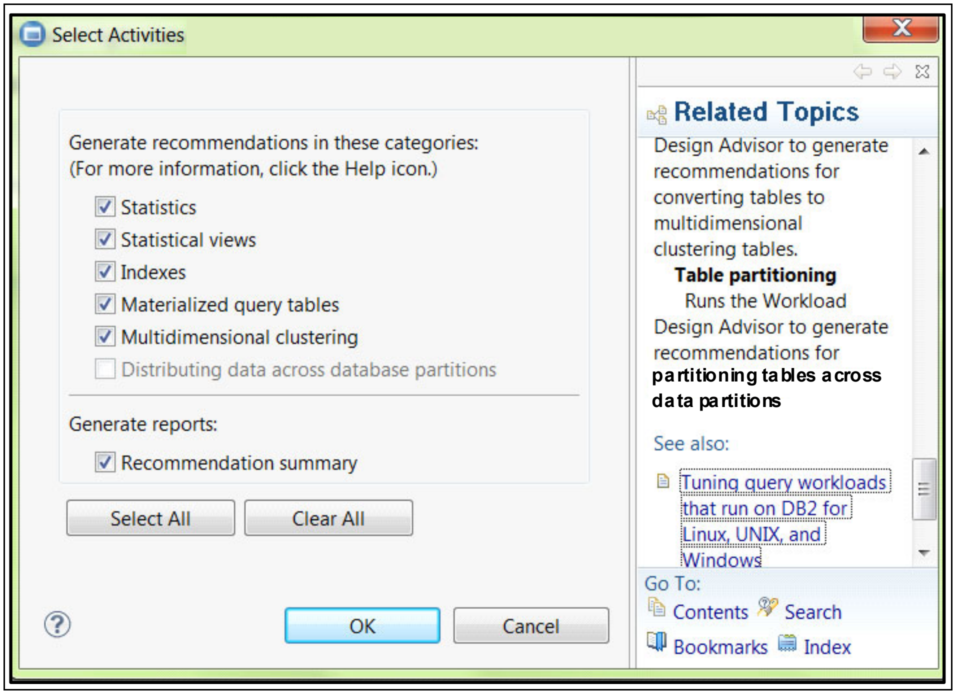

Figure 5-11 Select workload tuning activities dialog

After the selection is done, click **OK** to start the tuning task. The progress window opens, similar to single query tuning. You can put the task into the background, and the results are displayed upon completion. When the tuning task is done, you will be brought to the workload tuning result page (Figure 5-12 on page 238), which provides the options to view the tuning report and statement in the workload, and opens the recommendations. The next step is to review the tuning results and take proper actions, which are described in the next sections.

Questions and answers - Tuning:

- What is the difference between single query tuning and workload tuning?

  Single query tuning focuses on one single query and does the in-depth analysis of that query, provides the visual representation of the query access plan, provides recommendations on query rewrite, performs index creation, and collect statistics, as examples. If you find one hotspot query that is the bottleneck for the database system performance, the preferred way is to use single query tuning. If you have a set of related SQL statements for analysis, you might find tuning the SQL one by one is time consuming, and more problematic is that the recommendation, which is good for one SQL, might have a negative impact to the others, or the recommendations have some overlap. Workload tuning on a workload level provides a global optimization recommendation for the whole workload, consolidates the recommendations for each SQL, and also weighs the impact to each query based on its importance (number of executions, and elapsed time as examples). So in this case, a more suitable approach is to use the workload tuning for a set of related SQL statements, such as the SQL statements from one application, within one UOW, and belonging to one package, as examples.

- How do I choose the correct tuning type and choose activities?

  As an application developer, during the development of database applications, you can use the single query tuning to tune a specific SQL in depth. For a DBA tuning the query performance on production, a good approach is to use the workload tuning to save time and ensure the global performance. The DBA can also choose one single query that is critical, and then start the single query tuning for in-depth analysis.

  You must balance the recommendations you get from InfoSphere Optim Query Workload Tuner, the impact to the target database, and tuning elapsed time. In general, the more activities you choose, the more information and recommendations you get, the more impact is made to the target database, and the longer the time you need to see the results. For single query tuning, you can start from analysis tool and statistic advisor, and in the second round, choose all of them. For workload tuning, a good approach is to keep the statement number within 1000 to ensure performance and minimize impact to the target database. Start from statistics advisors, and next the index advisors; choose only the necessary advisors.

For each tuning task you can save it for the records and future reference. The single query tuning information is kept in the InfoSphere Optim Query Workload Tuner client side. For workload, all the information is kept in InfoSphere Optim Query Workload Tuner repository on the target database side. Go to the following address to learn more about managing a tuning project in InfoSphere Optim Query Workload Tuner:

`http://publib.boulder.ibm.com/infocenter/dstudio/v3r2/topic/com.ibm.datatools.qrytune.project.doc/topics/parent.html`

5.5 SQL tuning report

On the single tuning analysis page (Figure 5-9 on page 234), you can go to the Summary Report from the Query Tuner workflow assistant on the left side. The report (Figure 5-12) contains a summary of the recommendations from the Query Tuner advisors and tools. The tuning report is a good starting point to know the SQL and review the recommendations, and to discuss with others. You can examine the recommendations and corresponding DDL scripts, if applicable, and take appropriate actions to tune your query. You can also examine the formatted query and access plan summary, and cross reference the recommendations generated by the advisors. Use the table, column, and index information to do further analysis and tuning. It also provides the capability to save the report as HTM file, and then share with other people in the organization for the tuning report.

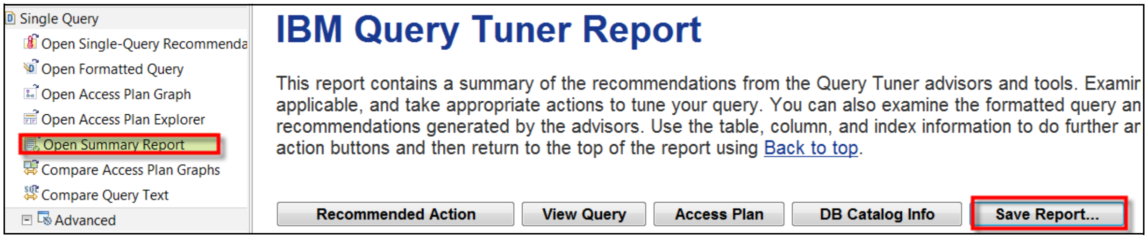

Figure 5-12 Single Query Tuner Report

As for workload tuning, when the tuning is finished, InfoSphere Optim Query Workload Tuner opens the tuning summary report panel (Figure 5-13 on page 239). The report lists recommendations from the workload advisors and information to help you understand them. Examine the recommendations and any corresponding scripts, and take appropriate actions to tune the SQL statements in the query workload. It also provides basic information of the target database and the workload been tuned. You can save the report as HTM file, and share with the others for review and discussion.

Note: Make sure you choose the **Recommendation Summary** option when selecting the tuning activities for single query (Figure 5-7 on page 232) and workload tuning (Figure 5-11 on page 236) to generate the tuning report.

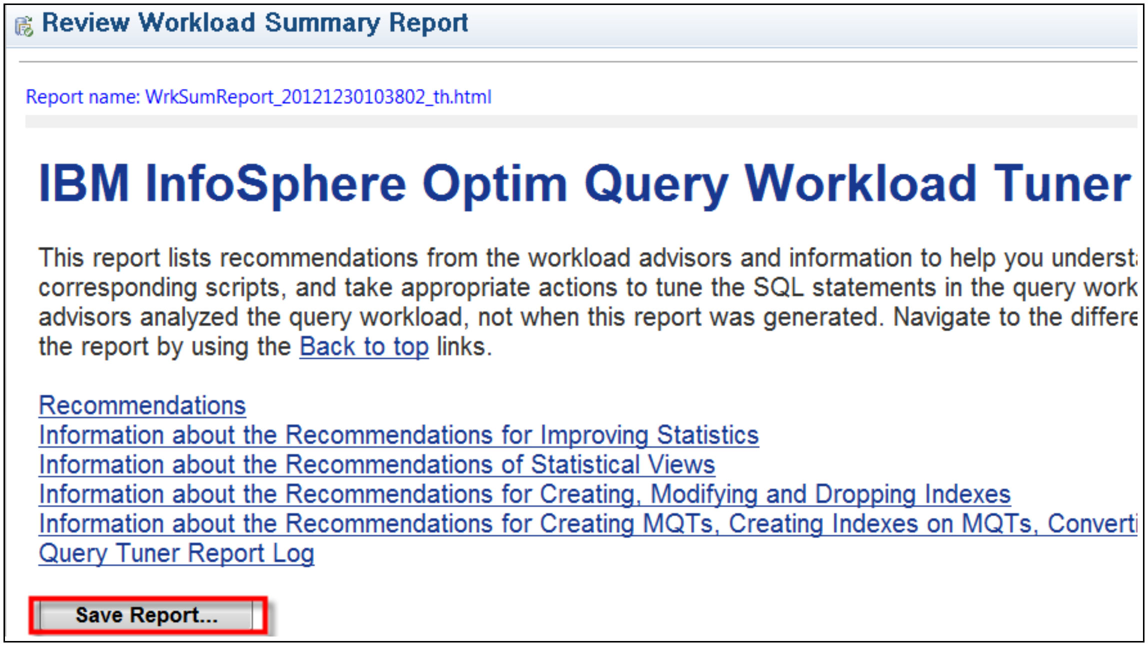

Figure 5-13 Workload tuning summary report

5.6 Query formatter and annotation

The single query analysis result has the option to open a formatted query where you see the formatted and annotated query to help you more fully understand the query. The formatted query shows a separate line for each table reference, each column reference in the SELECT clause, and each predicate. On the right side is the key statistics information for the objects referenced. When you click one row, the related rows are highlighted. You also see the missing statistics information for the referenced objects that might be candidates for statistics collection. In the additional information, you see the data skew suggestion in the example (Figure 5-14 on page 240).

Note: Choose the format and annotate the SQL statement option when you select the tuning activities (Figure 5-7 on page 232) to get the query format and annotation.

At the top of the formatted query table are options to expand or collapse the query, set the options for annotation, save the formatted query to an SQL file, print the query, or copy the formatted query to the clipboard.

Tip: To see the query format and annotation for a single query in a workload, use the following steps:

1. Right-click that single query in the workload and choose Run Single Query Analysis Tool and Advisor for the Selected Statement.
2. Choose the tools (format and annotate the query) and advisors.
3. Start the tuning task.

In the tuning results, you see the formatted query.

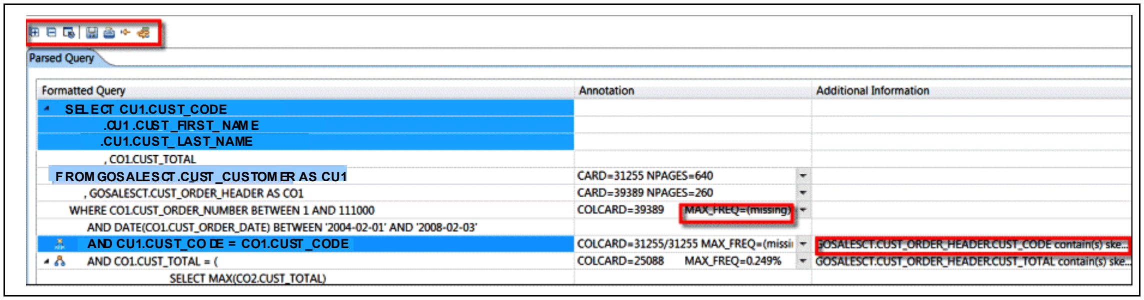

Figure 5-14 Query format and annotation

5.7 Analyze the access plan

Now, you understand *what* the query does through the query format and annotation. Next, you should also understand *how*, through the access plan, the DB2 engine explains and executes the query. InfoSphere Optim Query Workload Tuner provides the capability to render the visual representation of the query access plan (access plan graph and access plan explorer).

5.7.1 Access plan graph

You can navigate to the access plan graph panel in the tuning results section of the Query Tuner workflow assistant. The access plan graph provides a graphical view of the access plan for your query. You may open the graph in a separate window to view it if the access plan is big, as shown in Figure 5-15. In the main window, you see the graph. If an advisor issued a warning for an operation in the access plan, an icon of that advisor type is shown above the node that represents the operation in the graph. Click a node to see the details at the left side. When you hover the pointer over one node, the basic details for that node are displayed. Details might include node type, operation type, cost, estimated cardinality, and actual cardinality, and the recommendations.

> **Note:** You must select the **Display access plan graph** option when you select the tuning activities to see the access plan graph.

The top left of the panel lists more options for the access plan graph. You can export the EXPLAIN model to an XML file, which you can open and view the plan graph in the XML file. Using another InfoSphere Optim Query Workload Tuner, you can view the SQL text, zoom in, zoom out, and print the graph.

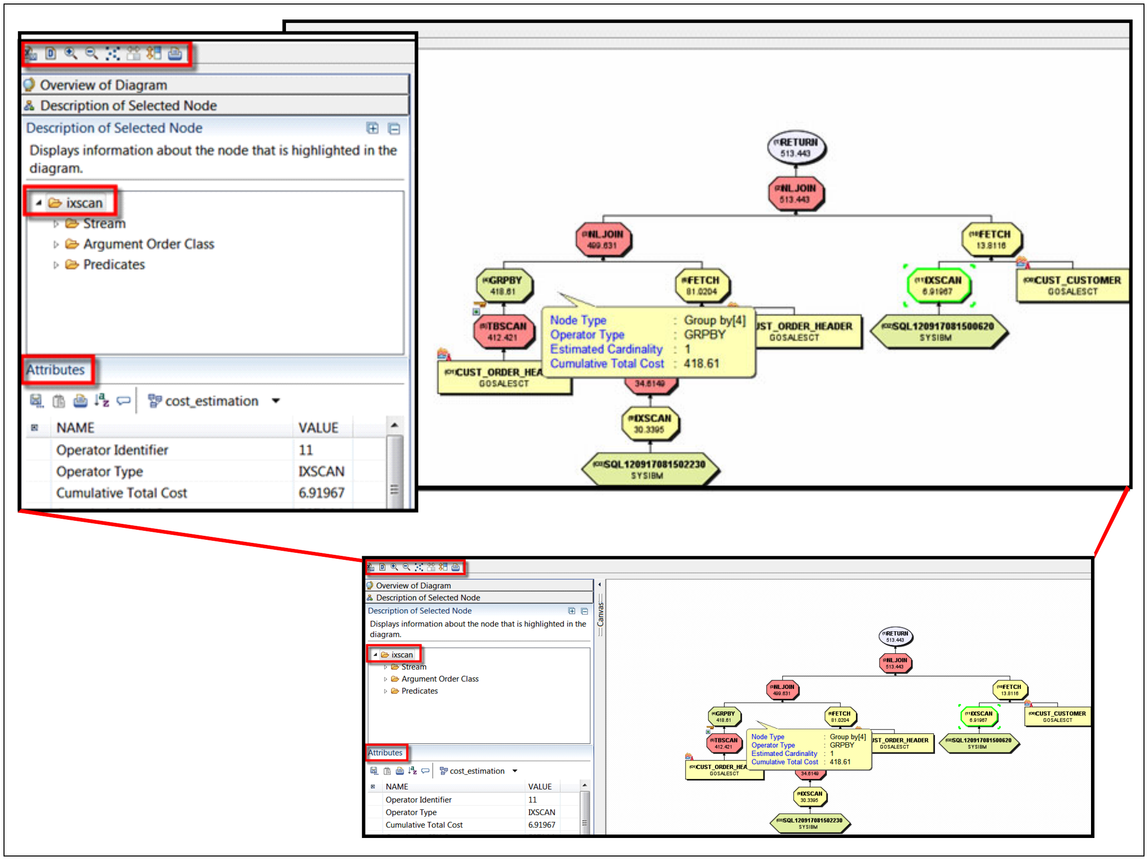

Figure 5-15   Access Plan Graph

An explanation of each node in the graph is in the information center:

http://pic.dhe.ibm.com/infocenter/dstudio/v3r2/index.jsp?topic=%2Fcom.ibm.datatools.visualexplain.data.doc%2Ftopics%2Fluwnodes.html

Tip: In the invoke query analysis tools and advisors, you can choose the **Collect actual values to use in access plan tools** option for the SELECT statement under the RE-EXPLAIN options section. InfoSphere Optim Query Workload Tuner then executes the statement and tracks the actual values for each operation. You can then see the actual cardinality in the access plan graph. A big gap between the estimated cardinality and actual cardinality typically means the related statistics are not collected or not current.

5.7.2 Access Plan Explorer

You can also view the access plan in a text representation format (Figure 5-17 on page 243), with more interactions, by clicking the **Open Access Plan Explorer** in the tuning results with Query Tuner workflow assistant. The access plan, with costs and statistics, is in a table or a hierarchical tree. You can change the view between table and tree using the menu button in the top left of the access plan area. Each row represents an operation in the access plan, and can be sorted by several metrics. In the example (Figure 5-16), it is sorted by the total cost of the operations. Meanwhile, it also provides the capability to highlight the related operations (inflow and outflow) with different colors for one specific operation.

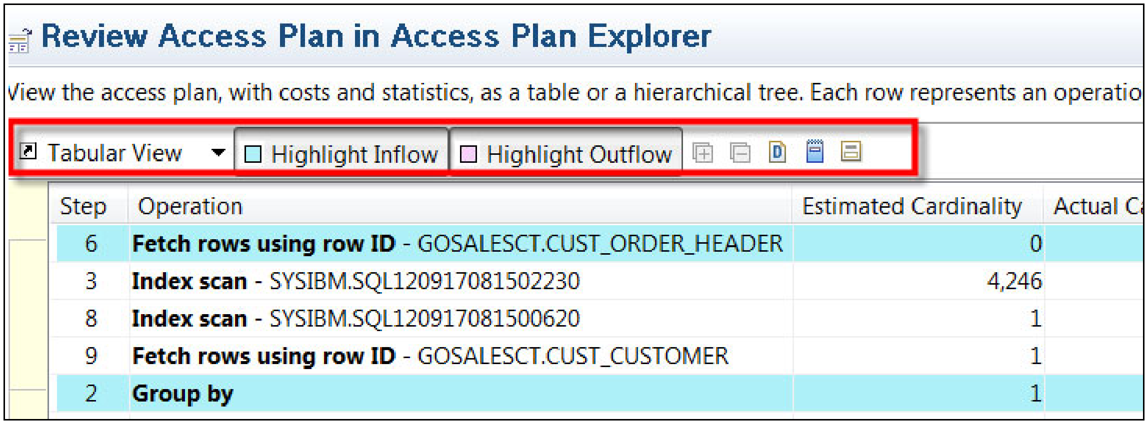

Figure 5-16   Access Plan Explorer main section

Select a row in the table and the details for that operation will be displayed in the lower section (Figure 5-17). They include basic properties, predicates, tables, index used, and also output stream information.

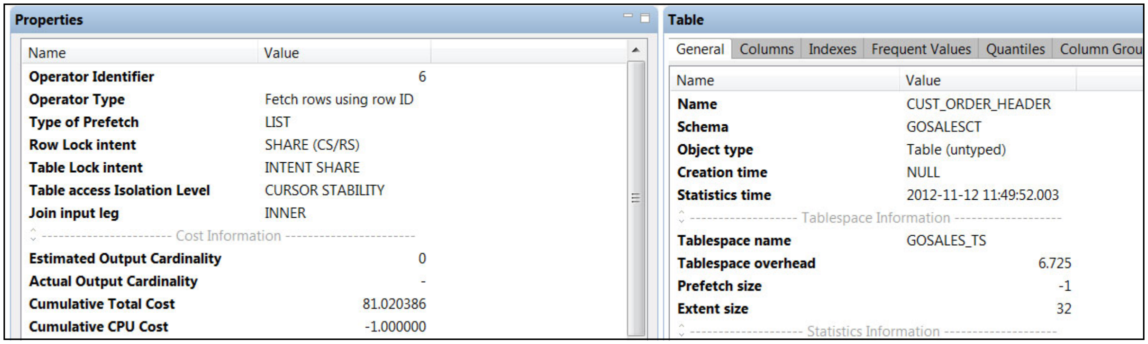

Figure 5-17   Access Plan Explorer Operation Details section

> **Note:** To view the Access Plan Explorer, choose the **Show access plan in Access Plan Explorer** option in the tuning Select Activities dialog (Figure 5-7 on page 232).

5.8 Statistics recommendations

Statistics are crucial for DB2 optimizer to choose the correct access path for the query. If the statistics information cannot reflect the current actual characteristic of the data object referenced in the query, the optimizer might fail to choose the optimized access path for that query. InfoSphere Optim Query Workload Tuner analyzes the statistics information of the object referenced in the query based on explain information, determines the statistics that should be collected, but are not collected yet, the previous collection time that exceeded a given time, or that conflicted with each other, and provides the user with the recommendations with run statistics scripts, and corresponding explanation. InfoSphere Optim Query Workload Tuner provides the advisor of statistics for both single query and workload. For workload, there is one more advisor on the statistical view. When you tune an SQL performance, start with the statistics recommendations to collect the required statistics information, and then do the remaining analysis and tuning.

5.8.1 Statistics advisor for single query

In the single query analysis results Summary tab, you see a list of all the recommendations. Choose the statistics item, and double-click to see details of the recommendations (Figure 5-18 on page 245). On the statistics recommendation page, you see the recommended RUNSTATS commands at the left side; at the right side, you can retrieve the RUNSTATS commands stored on the target database side. In the lower section is the explanation of why InfoSphere Optim Query Workload Tuner recommends the RUNSTATS commands:

- The statistics are not collected on the object that has been referenced by the query.
- The statistics are not current (you can set the statistics collection time to expire in the advisor settings).
- The statistics are conflicted in two objects.

The conflict detail section gives a detailed explanation of the conflict statistics. You can run the RUNSTATS command in place, and save them to the statistics profile for automatic statistics collection, or you can save the commands to file, and integrate to other maintenance scripts. If the database does not have a time window for RUNSTATS now, you can save the scripts as a file and run it later during the maintenance time window of the database.

More information about automatic statistics collection with a statistics profile is available in the information center:

http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.admin.perf.doc/doc/c0011392.html

> **Note:** To get the statistics recommendations, select the **Statistics** option in recommendations category in the tuning Select Activities dialog (Figure 5-7 on page 232).

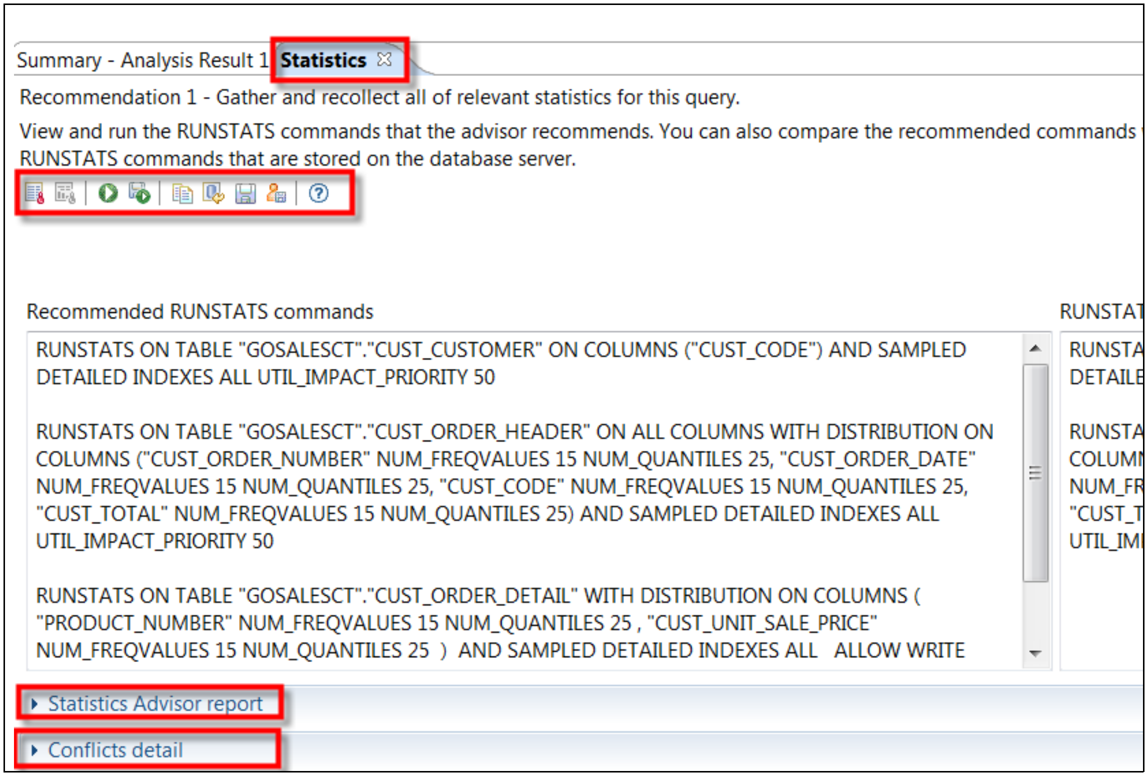

Figure 5-18 Single Query Statistics Advisor

5.8.2 Statistics advisor for workload

In the workload recommendation summary, you can choose the statistics and drill down to see the details about statistics (Figure 5-19 on page 246). It shows the recommendation in a table format. Each row is a data object (table or index) that requires RUNSTATS, including the cardinality of the object, how many SQL statements referenced this object, total cost of the operation on this object, last statistics collection time stamp, and reason why it should be collected (missing, conflict, or obsolete). You can select one row to show the details in the lower section of the selected object on the statistics information.

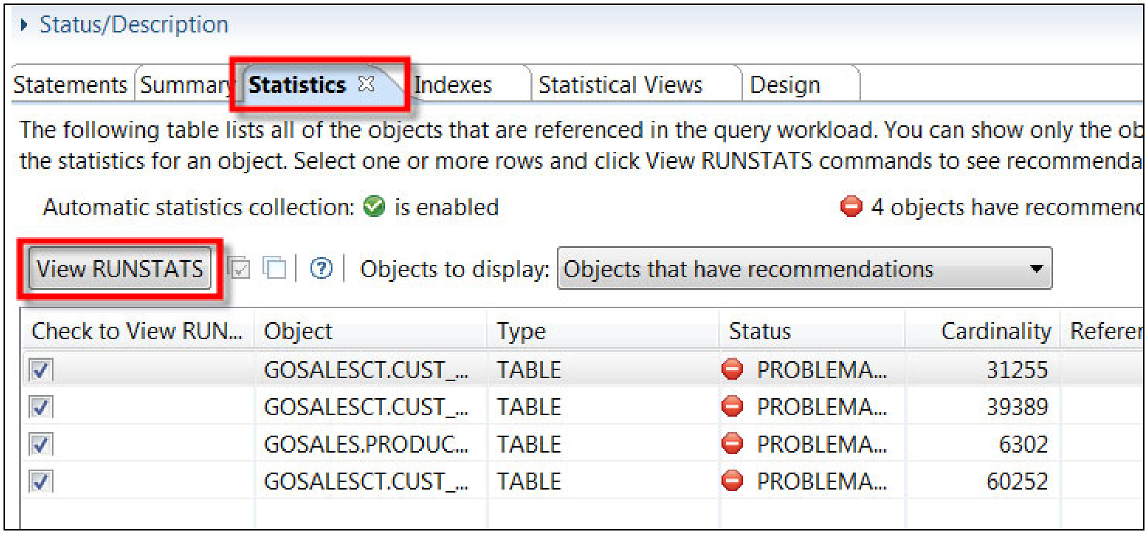

Figure 5-19 Statistics advisor for workload

Click **View RUNSTATS** to view and select the **RUNSTATS** commands for all the selected objects in the table. In the Review Recommended RUNSTATS Command window (Figure 5-20 on page 247), you can choose whether to use the commands generated by advisor, or merge the commands from advisor and from the statistics profile. You can view and run the consolidated RUNSTATS commands directly, save to a file, or save to the clipboard for future use. You can also choose to save the consolidate RUSTATS command in the statistics profile.

> **Note:** To get the statistics recommendations for workload, select the **Statistics** option in recommendations category in the tuning Select Activities dialog (Figure 5-11 on page 236).

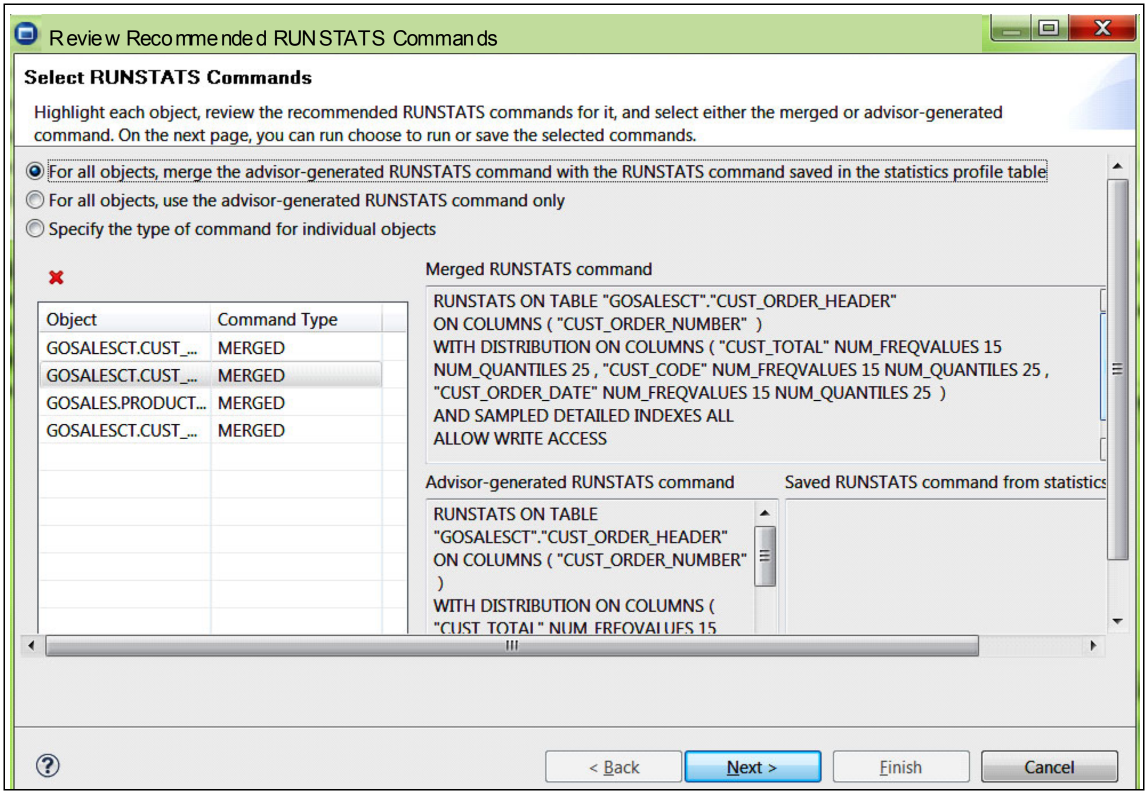

Figure 5-20 Review Recommend RUNSTATS Commands

Questions and answers - Statistics recommendations:

- Does the advisor recommend all possible features for a RUNSTATS statement?
  - Yes, it does recommend sampling, uniform and distributions statistics, and column groups on tables.
- Does the advisor recommend statistics on MQTs?
  - Yes, because the advisor uses the predicate and other interesting information from the explain, this information is generated after MQTs have been matched to the query and replaced the base table in the query.
  - Predicates, ordering, and grouping features used by the advisor will involve MQTs when they are used in query plans. This information enables the advisor to generate recommendations for RUNSTATS on MQTs and base tables.
- What information does the Statistics Advisor consider?
  - The advisor considers columns that are involved in predicates, ORDER BY, GROUP BY, and DISTINCT.
  - For column groups, it also notes local predicates that are used in the same statement to form these, so that the dependencies of column values can be accounted for by the collection of column group statistics.
- Does the advisor recommend statistics for statistical views?
  - Only if the Statistical Views Advisor was not invoked at the same time.

5.8.3 Statistical View recommendations for workload

For complex report and analytic SQL statements, usually only two major problems cause the optimizer to fail to choose the correct access plan:

- Data skew and join predicate filter miscalculation, because there is no statistics information for the join predicate or predicate with complex expression, such as the following example:

  ```
  WHERE salary+bonus > …
  ```

- Overloaded dimension in a star join

Time dimension might contain 20 years of data but the transaction table in a data warehouse keeps only the last five quarters of data. Uniform distribution is a wrong assumption after the join.

Statistical view is considered a key solution in solving these problems. Create a view on the join predicate, such as a predicate with complex expression, enable query optimization, and then run statistics on the view. Statistics on the views will be used in query planning (Example 5-1).

Example 5-1 Statistical view

```
CREATE  view sv_store_dailysales as
(SELECT  s.*
         FROM  store s, daily_sales ds
         WHERE  s.storekey = ds.storekey)
ALTER VIEW  sv_store_dailysales  ENABLE QUERY OPTIMIZATION
RUNSTATS  on table db2dba.sv_store_dailysales  WITH DISTRIBUTION
```

InfoSphere Optim Query Workload Tuner will take the query and related statistics information into consideration, and generate the recommendations on the statistical view to improve the query performance for complex query. In the recommendations tab, you see a list of the statistical views either new as proposed or existing. Choose one view and you see the statements in the workload that the view can assist the performance, and also the tables with which the view is defined. You can choose the proposed statistical views and review the scripts to generate the views, and save the scripts to file, run the scripts, and save the RUNSTATS command to statistics profile for auto-collection in the future.

For more information regarding the DB2 statistical view feature, see the DB2 information center:

http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.admin.perf.doc/doc/c0021713.html

Note: To get the statistical view recommendations, choose the statistical view option in the workload tuning activities selection (Figure 5-11 on page 236).

5.9 Indexes recommendations

The index is one of the key points for physical database design and also critical to SQL performance. The index can improve the performance of the query, which requires retrieving records from the database with predicates. In this case the index can avoid full table scans, sort, and improve efficiency of filtering, thus save both I/O time and CPU time. Also, the SQL that makes changes to the database, requires additional updates to the index, and the index requires additional storage. In addition, the index requires more effort in terms of maintenance such as RUNSTATS, REORG, and BACKUP. InfoSphere Optim Query Workload Tuner provides the capability to design the necessary indexes for the database based on the given SQL statements, consolidate indexes that are overlapped, enable the what-if analysis for creating or dropping indexes with the performance information about query performance and disk storage. Meanwhile, corresponding scripts that the user can run immediately or save to file for later use are generated

> **Note:** To get the index recommendations, select the **Indexes** option in the recommendations category on the single query tuning activities dialog (Figure 5-7 on page 232), or **Indexes** in the workload tuning activities selection (Figure 5-11 on page 236).

5.9.1 Indexes recommendations for single query

InfoSphere Optim Query Workload Tuner will examine the query, statistics information, and access plan to determine whether new indexes will help the performance of the given query. In the recommendation list, double-click the item for **Indexes**. The details of the index recommendation (Figure 5-21 on page 251) is displayed. At the beginning of the recommendations, you see the estimated performance improvement (93.17% in the example) and estimated storage (1.16 MB in the example) with the recommended indexes listed in the table below those recommendations. There, you see the details of the proposed indexes, name, key columns, estimated storage, and, if any, included columns. You can click **Test Candidate Indexes** to test how the proposed indexes will affect the access plan of the query (described in 5.9.2, "Indexes recommendations for workload" on page 251). At the bottom, you also see the existing indexes of the tables referenced by the query with the last-used time.

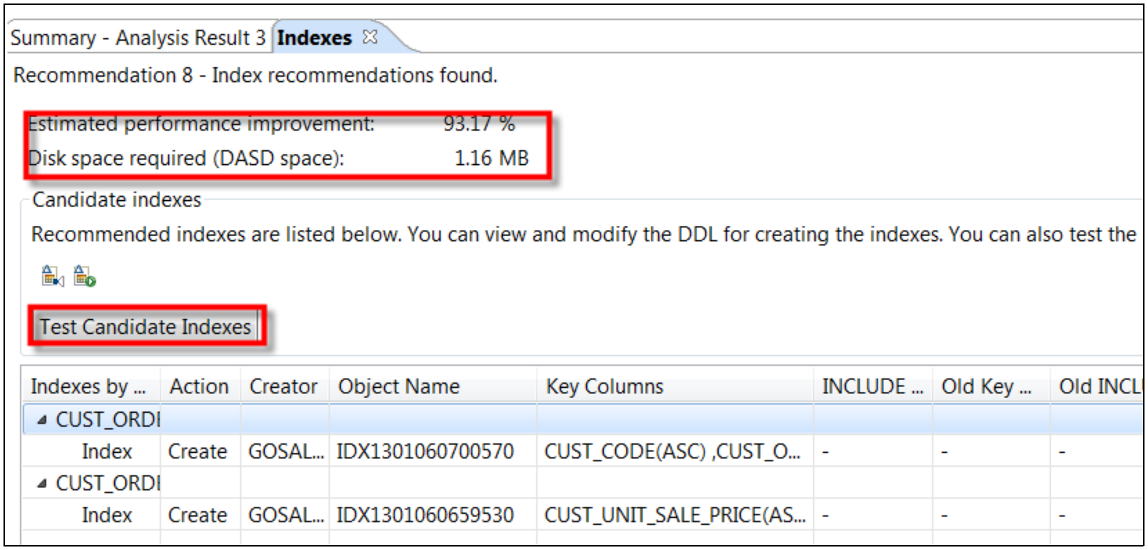

Figure 5-21 Index recommendation for single query

5.9.2 Indexes recommendations for workload

Index recommendations for a workload will evaluate the performance gains of one index on the workload level, based on the runtime information for each SQL in the workload (such as the number of executions and elapsed time).

As an example, for an index that is beneficial to SELECT statements, while having a negative impact to the statement with changes to the database records, InfoSphere Optim Query Workload Tuner will evaluate the total benefits, and determine the recommendations.

You can launch the index recommendations details from the workload recommendations summary to get the details of the proposed indexes, which can beneficial for the workload, and options to generate the DDL and test candidate indexes (Figure 5-22 on page 252). Initially, you see the overall performance gains and required storage with the proposed indexes. The next table describes the indexes recommended for each table, the cardinality of the table, number of SQL statements referenced to this table, and so forth. You can select the proposed index on the table level, based on the information presented.

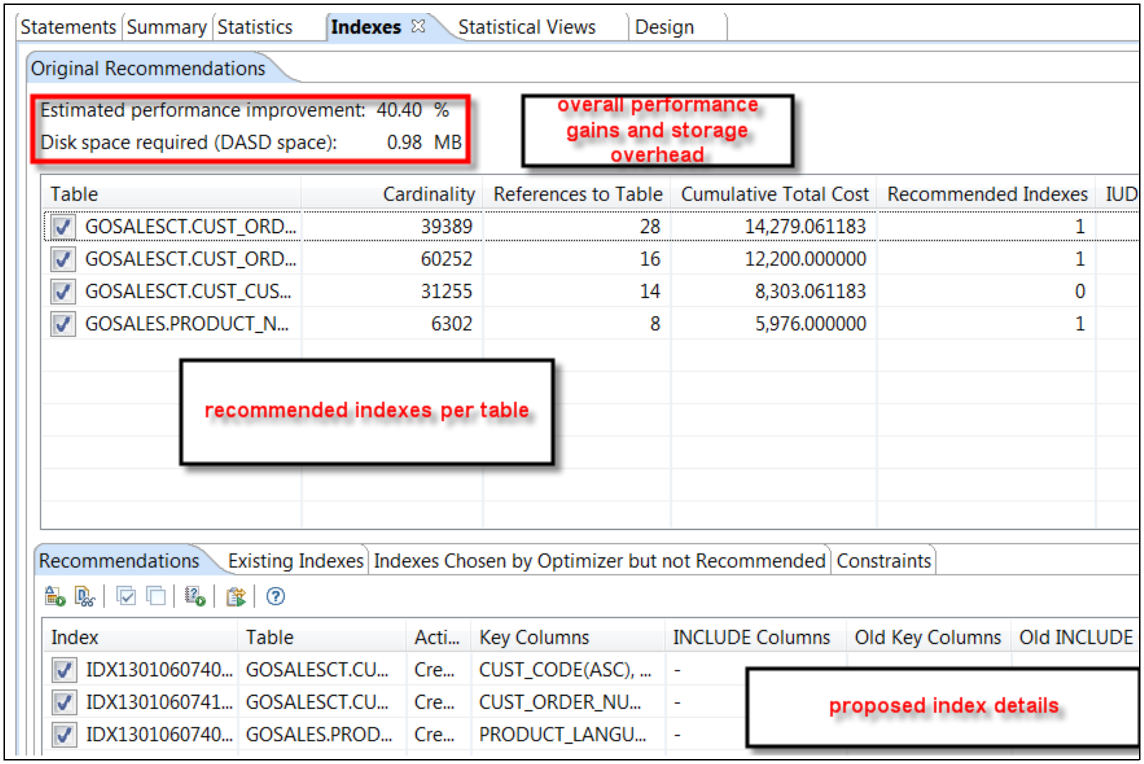

Figure 5-22 Indexes, recommendations for workload

The lower section has a series of tabbed tables and options. A recommendations tab lists the proposed index, the name, key columns, included columns, action type, average performance gains for the affected SQL statements in the workload and storage overhead of the index. The action types can be *create* and *drop*. *Create* is the new proposed index, and *drop* is the index that is recommended to drop with consolidation to avoid unnecessary indexes. You can do the following tasks:

- Select one or more indexes to examine the affected queries of those indexes, and how the indexes can improve the performance for each single query.
- Review the DDL to create the index, run it or save to file.
- Select a set of candidate indexes to test the indexes virtually to see how the candidate indexes will impact the access plans of the SQL statements in the workload, and to validate whether the indexes proposed by InfoSphere Optim Query Workload Tuner will be picked by the DB2 optimizer.
- Compare the access plan on the workload level with the index recommendations and workload without index recommendations.

In other tabs, you see the existing indexes for the tables that have been referenced in the workload, and indexes that are chosen by the optimizer but not recommended in the product, with explanation, and also the constraint to recommend the index.

5.9.3 Test Candidate Indexes

After you get the index recommendations, consider validating whether the recommended indexes will really be chosen by the optimizer to improve the query performance, or if you want to test one index to see whether the index can improve performance, and also how it will affect the access plan of the query. Moreover, you might want to check how the access plan will be changed if one existing index has been dropped. You can do the validation *physically* by creating the index or dropping the index, which requires much effort and causes a large amount of overhead to the target database. InfoSphere Optim Query Workload Tuner provides the capability to validate such changes *virtually*, which means not deploying the changes (create or drop indexes) physically to the database, but using the virtual index capabilities built into the DB2 engine to perform what-if analysis of the index changes to see the affects on the single query or workload. You can use this feature to test the index changes before deployment on both single query and workload level.

You can launch the test candidate indexes for single query from either the single query index recommendations (Figure 5-21 on page 251) or directly from the query tuner workload assistant (step **4. Invoke** for single query); under the Advanced section is the Test Candidate Indexes option. For workload level, you can launch from the workload indexes recommendations (Figure 5-23 on page 254) or from the query tuner workload assistant (step **4. Invoke** for workload); under Workload section is the Run Workload Test Candidate Indexes option.

In the test candidate indexes panel (Figure 5-23), you can manually add a candidate index in addition to the recommended indexes from the index recommendations. You can also update the index definition, such as the name, key columns, and include columns, through the icon menu on top of the indexes table. In the lower section, you see existing indexes; you can choose one or several existing indexes for a virtually drop. You can choose only the NON-UNIQUE index for virtual drop. After you finish changing all the index settings, click **Test Candidate Indexes** to start the what-if analysis. You are prompted to specify the statistics information for the candidate indexes. You can input the values if you know them, or let InfoSphere Optim Query Workload Tuner estimate the statistics information automatically. The DB2 engine will explain the single query or workload with the given virtual indexes definitions, and provide the new access plans with an estimated cost.

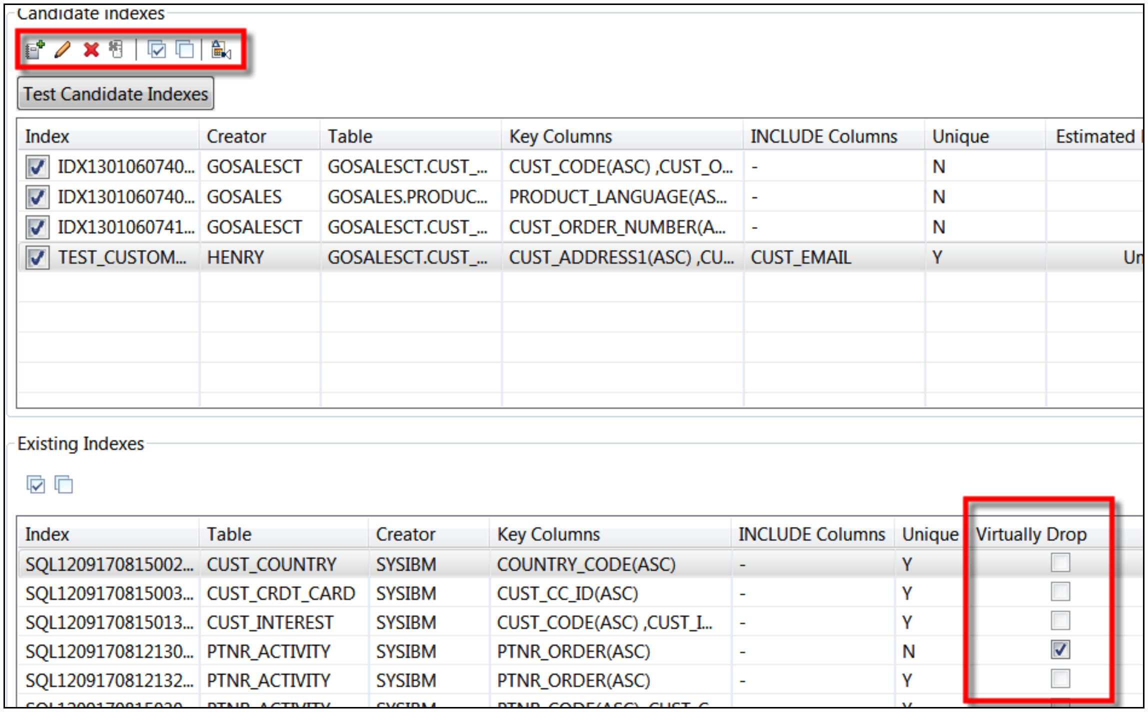

Figure 5-23 Test Candidate Index Panel

When the process is complete, the results are displayed in the Test Candidate Indexes Summary page (Figure 5-24). It shows the overall estimated performance improvement and storage overhead. For each index, you can view the impacted statements by this index. You can also compare the access plan differences of the single query or workload with and without the tested indexes to validate the performance improvement. You can select the indexes that improve the performance, and review the DDL, and run or save it to a file.

Tip: To see the effect of the actions based on recommendations such as the statistics and index advisors, you must rebind the static SQL statements in the target database using the latest statistics and index design. For dynamic SQL statements, the recommended actions will take effect when the SQL statements been compiled again.

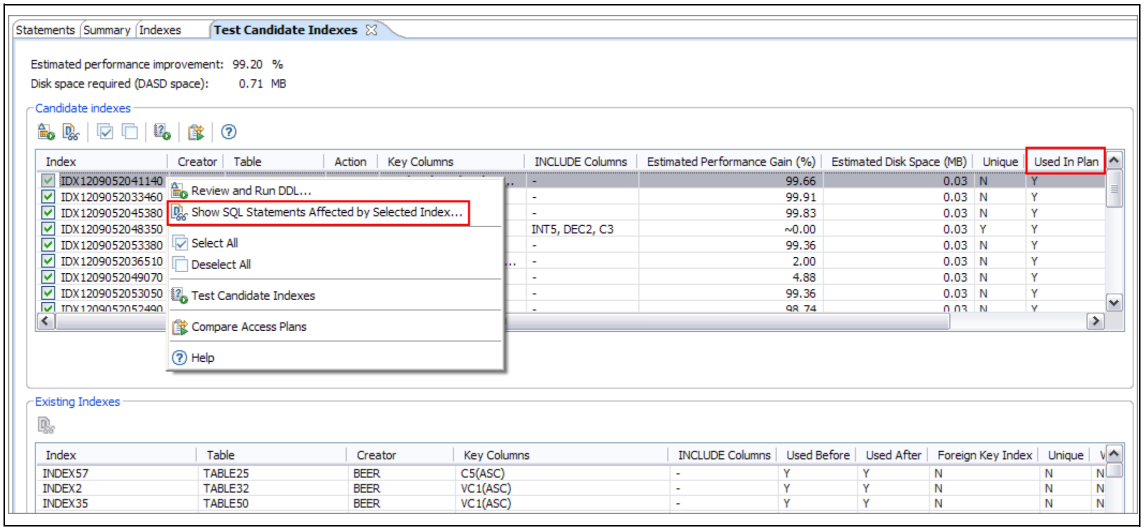

Figure 5-24 Test Candidate Indexes summary

Questions and answers - Advantages of index-related tools:

- What are the advantages of index related tools in InfoSphere Optim Query Workload Tuner compared with db2advis tool?

  The **db2advis** (DB2 design advisor) command-line tool generates recommendations on indexes, MQT, MDC, and partition keys based on a set of SQL statements. Compared with db2advis, InfoSphere Optim Query Workload Tuner related tools have the following advantages:

  - A user interface that is easier to use.

    The **db2advis** tool provides a command-line interface, with text output.

  - Consolidating the indexes recommendations.

    The **db2advis** tool may provide a set of overlapped indexes; InfoSphere Optim Query Workload Tuner consolidates the recommendations into one index.

  - Threshold based recommendations.

    InfoSphere Optim Query Workload Tuner provides the way to set the threshold on the performance gains and size of the required storage of the recommended index; the **db2advis** tool generates the recommendations for all indexes.

  - Test candidate indexes function.
  - Suggests multi-column fact key (MCFK) indexes (for zigzag join), and jump-scan support in consolidation.

    This is for DB2 v10.1, because the **db2advis** does not use the latest features in DB2 v10.1.

  - Index elimination for consolidations and virtual drop.
  - Integration with access plan comparison feature.

  InfoSphere Optim Query Workload Tuner provides the capability from index recommendation, virtual index, and access plan comparison to review and validate how the recommendations can affect the access plans.

5.10 Compare and lock down access plan

One of the key methods to validate the query performance changes during query tuning is to review the access plan changes. InfoSphere Optim Query Workload Tuner provides a GUI-based mechanism to compare the access plan between any two tuning activities on one single query, one workload, or workloads from different systems. With the access plan comparison, you can understand how the plan changes, whether it is what you expect, and whether you must lock down the access plan for one single query or a set of queries.

This section describes the access plan comparison and lock down feature in InfoSphere Optim Query Workload Tuner to assist the query performance tuning.

5.10.1 Single query access plan comparison

When you take actions based on the recommendations or experience, such as collect statistics, create an index, rewrite the query text, and update database configurations, you might want to see the new access plan of the query, and also the difference between the new one and previous one. You can launch the access plan comparison from the mini-navigation bar (by click step **6. Compare**), or in the query tuner workflow assistant in the tuning results of single query, you see the *Compare Access Plan Graphs*. You can choose any two instances of the query explain information, and then start the comparison. In the comparison panel, the differences in the access plan are highlighted in the access plan graph. In the lower section, review the difference details. InfoSphere Optim Query Workload Tuner highlights the table access and table join differences, the difference of estimated cost for each step are not highlighted. You then decide whether to use the latest access plan for this query, or revoke back using the optimization profile to lock down. The plan lock-down support in InfoSphere Optim Query Workload Tuner is described in the following sections.

5.10.2 Workload access plan comparison

In some situations, you might want to compare the access plans on a workload level for workload level query analysis and tuning, such as in the following list:

- Check whether the access plan of statements of the application will change after the application is migrated from the testing system to the production system.
- Confirm and review the statement plan changes during the migration of database from a lower version to a higher one.

- Get problem access plan change warnings after RUNSTATS or create new indexes.
- Generate scripts to lock down access plans based on an existing EXPLAIN snapshot.

InfoSphere Optim Query Workload Tuner offers the capability to compare the access plan on a workload level. You can launch the workload access plan comparison features from various locations:

- In the manage workload panel: Click **Compare** above the workload list.
- In the workload indexes recommendations panel: Use the icon menu.
- In the workload candidate test indexes results: Use the icon menu.
- In the mini-navigation bar, click step **6.Compare**: use the menu for Compare Workload Access Plan.

You can compare with either of the following options (Figure 5-25):

- Compare two EXPLAIN snapshots for one workload.
- Compare two separate workloads.

The first option is usually for validation of the tuning results of one workload based on the workload recommendations, or to lock down the access plans within the workload in the target database. If you want to see the access plan changes of a set of statements during the database movement from test system to production system, or database version upgrade, use the second option to capture the same or similar workload in the two database systems, and compare the access plans on the two workloads.

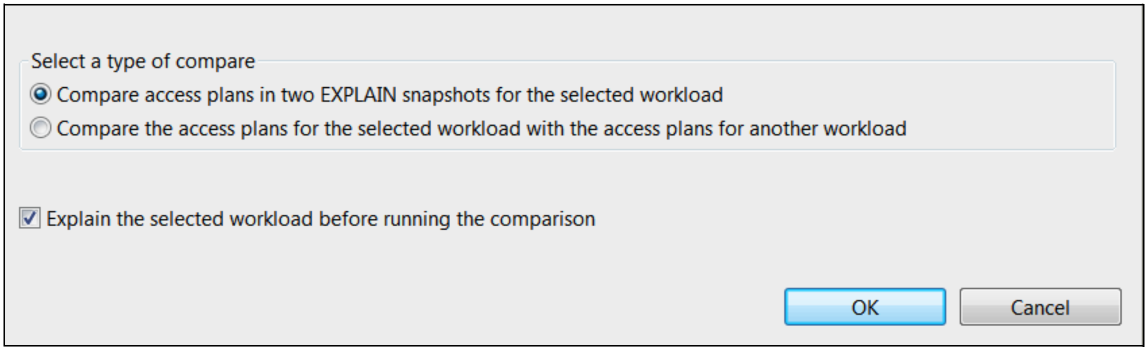

Figure 5-25 Workload access plan comparison options

InfoSphere Optim Query Workload Tuner can compare the access plan on table access, join methods, and cost for each query in the workload, and then generate the report. After the comparison finishes, you can review the results from the *comparison history*. The comparison results show basic comparison information, a summary of the comparison, and details of the comparison for each query. The differences are highlighted in colors:

- Green means improvement.
- Red indicates the regression.

Several more actions are available in the comparison results panel:

- View SQL statement comparison details.

  Right-click on a selected query to view the comparison details with the access plan, in graph mode to show the difference.

- Create an optimization hint using the existing plan.

  Right-click on a selected query to create an optimization hint for that query, either for the source or target. The GUI-based optimization profile panel opens.

- Apply, edit, or remove filter.

  You can apply a filter on the results to view only the interesting queries in the results table; the filter conditions are on the table access, join method, and cost changes.

- Lock down access plan for all SQL statements.

  Use the InfoSphere Optim Query Workload Tuner to generate optimization profiles for all the queries changed in either the source or target to lock down those queries.

- Generate an HTML comparison report

  Generate an HTML report for the comparison results, which can be shared across the organization.

- Generate a new workload

  Generate a new workload with the filter on the comparison, and the standard workload post-filters to refine a new workload of interested queries to tune.

5.10.3 Edit optimization profile for single query lock down

In some cases, the performance and access plan gets worse again after the tuning as the data characteristics change. You might need to lock down the access plan to ensure the stability of one SQL. Or, perhaps you use all the possible ways to tune the query, but it still a less optimal access plan than expected. In this case, you can use DB2 *optimization profiles* to customize the plan hint for that SQL. However, the optimization profile is not that easy to edit and is error-prone. InfoSphere Optim Query Workload Tuner, however, provides a GUI-based editor so you can create, edit, validate, and deploy the optimization profiles.

To use this feature, see the information center:

http://publib.boulder.ibm.com/infocenter/qrytuner/v3r1/index.jsp?topic=%2Fcom.ibm.datatools.qrytune.sngqry.doc%2Ftopics%2Foptprofiles.html

For information about DB2 optimization profiles, see the information center:

http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.admin.perf.doc/doc/c0024522.html

5.10.4 Workload plan lock down

When you finish the workload access plan comparison, and view the comparison result you can choose to *lock down the access plan for all statements* as depicted in Figure 5-26 on page 261. InfoSphere Optim Query Workload Tuner will generate the optimization profile for each statement which has changes on the access plan. You can also generate a subset of the workload with the filter, and then lock down the statements that show interest. In the pop-up dialog, you specify the basis of the access plan (either the target or the source), profile prefix and schema, and the target database platform to deploy the optimization profile. After the optimization profile has been generated you can view the scripts, or save the script as an SQL file, and later execute the SQL file on the target database server side to deploy the optimization profile. After you deploy the profile, rebind the packages to make the optimization profile effective for static SQL statements. For dynamic SQL statements, you might need to flush the package cache to make the SQL compile again so it will pick up the profile.

Important: The workload plan lock down supports target database version is DB2 V9.7 FP7, DB2 V10.1 FP2 and following on fix packs, versions.

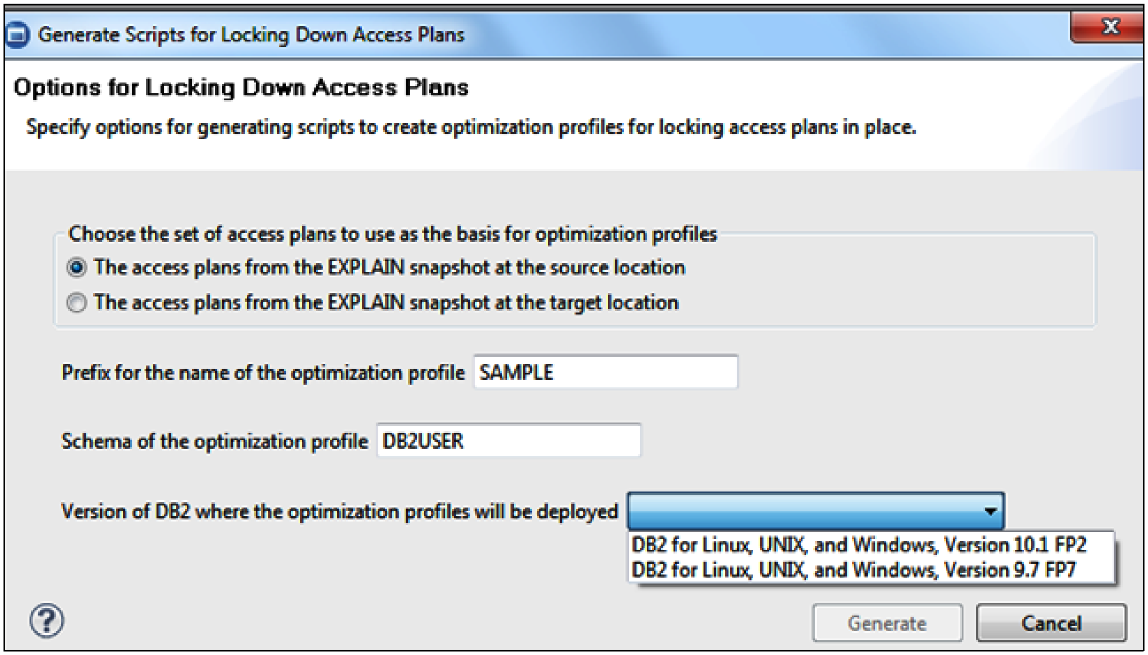

Figure 5-26 Workload plan lock down

Other useful query analysis advisors and tools are available in InfoSphere Optim Query Workload Tuner. You can access them through the InfoSphere Optim Query Workload Tuner UI. Although they are listed here, for more details about them, see the information center.

- Query Advisor: Uses the best practice of the DB2 SQL statements rule to find the part of the query that potentially does not comply with the rules, and might cause the optimizer to select a suboptimal access plan. You can follow the advisor explanations and examples to make proper changes of the SQL statements. More information is available from the information center:

  http://publib.boulder.ibm.com/infocenter/qrytuner/v3r1/index.jsp?topic=%2Fcom.ibm.datatools.qrytune.sngqry.doc%2Ftopics%2Fgenrecsqa.html

- Access Path Advisor: Examines the access plan that is chosen by the optimizer and identifies certain common access path issues. The warnings that this advisor provides can help you to understand where to look for trouble in an access plan graph or in the Access Plan Explorer. More information is available at the following address:

  http://publib.boulder.ibm.com/infocenter/qrytuner/v3r1/topic/com.ibm.datatools.qrytune.sngqry.doc/topics/genrecsapa.html

- Workload Design Advisor, including advisor on MQT (materialized query tables), multidimensional clustering (MDC), and partition key in the DPF environment. You can choose the advisors when selecting tuning activities for workload tuning.
- Compare Query Text: You can use InfoSphere Optim Query Workload Tuner to compare any two statements text with well-formatted form to see the query difference. You can access this feature through the Compare step in the workflow assistant. Make sure the two queries to be compared has been formatted through the query format tuning tool.

6

Finding and fixing database level bottlenecks

Performance refers to the way that a system behaves in response to a particular application. It can be measured in terms of system response time and resource utilization. Performance is generally affected by factors such as these factors:

- The resources that are available on your system
- How well those resources are used and shared

An important initial step when there is an unknown bottleneck is to categorize it by basic type. Does it have to do with disk, or maybe a CPU? Memory? Locking or network? These broad categories tend to cover about 95% of the performance problems you will see.

This chapter shows how InfoSphere Optim Performance Manager can help narrow the problems to some sample root causes under each type. It also describes how you can use InfoSphere Optim Performance Manager to analyze high disk and high CPU utilization on your monitored system using baselines.

Chapters 5 and 6 of *IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows*, SG24-7925, describe how to use the tool to monitor I/O and CPU by first looking at the high-level Health Summary dashboard and then drilling down through the alerts. This process is still true when dealing with alerts and thresholds. The concept of baselines was introduced in InfoSphere Optim Performance Manager, V5.1.1.

This publication concentrates on using baselines for the analysis. It starts with the Overview Dashboard and shows how to drill down to the dashboards, explain some of the metrics and discuss the reports that can assist with analyzing the performance of your database.

6.1 Disk-related bottlenecks

Because of what databases do, disk bottlenecks are probably the most common type. These odds are increased by the fact that many systems are under-provisioned in terms of disk requirements. Thus, I/O performance plays an important role in the overall stability of a database system.

I/O bottlenecks can occur for the following reasons:

- Buffer pool too small for active data set
- Increase in activity on a single disk
- Table scans resulting from insufficient indexing
- Bad plans or poor SQL
- Hardware issues

This section describes how to narrow the I/O activity seen on the operating system level down to the database object and the application consuming I/O. It shows how InfoSphere Optim Performance Manager can help you do the following tasks:

- Identify high I/O that affects your DB2 response time using alerts.
- Examine the high I/O with respect to the typical range of values as captured in the baseline.
- Drill down into problem detail and analyze the high I/O utilization on your monitored database.
- Resolve the high I/O causes and improve the performance of your applications by using expert advice.
- Prevent problems by monitoring historical trends for planning and building performance from the ground up.

6.2 Identifying and diagnosing high I/O

In the environment used for this scenario has an AIX machine with several databases, both local and remote.

6.2.1 Symptoms of high I/O utilization

InfoSphere Optim Performance Manager gives you “clues” about your I/O issues in the various dashboards. The following basic symptoms can indicate a disk bottleneck:

- High percentage of I/O wait time in the Overview Dashboard (perhaps more than 20 - 30%)
- Long I/O times in I/O dashboard (more than approximately 5 - 10 ms.)
- I/O or log alerts in Health Summary dashboard
- Low-to-mid CPU usage seen in Overview or System dashboard
- Saturated disks as seen in Tivoli (80% or more are busy)

6.2.2 Assessing database health-related problems using alerts

One way to start analyzing performance problems is to use a top-down approach that starts with the Health Summary dashboard on the InfoSphere Optim Performance Manager web console.

This dashboard gives a system-wide view of the health of all your databases. Chapter 3, “Installing and configuring Optim performance management tools” on page 69 describes how to configure a database for monitoring. Chapter 4, “Getting to know InfoSphere Optim Performance Manager” on page 145, introduces the Health Summary dashboard. The key performance indicators (KPI) for System and Database are checked periodically by using a sampling rate set during configuration of the monitored database.

These indicators appear on the Health Summary dashboard, the Alerts list, and other related dashboards:

- If a threshold is about to be reached, a warning alert is generated; the alert is represented by a yellow triangle indicator.
- When a threshold is reached or violated, a red square indicator is displayed.
- A green diamond indicator means that the metric is within the desired values.

You can click any of the red indicators to see more information about the specific alert generated, such as the time the alert happened and the severity.

By clicking the red indicator under I/O heading in the Health Summary dashboard, you get a list of the all the alerts related to I/O for this database. You can filter the alert list by clicking the **Define Filter** icon in the I/O alerts dialog. Figure 6-1 shows the filter icon and a filtered view of all the critical I/O alerts.

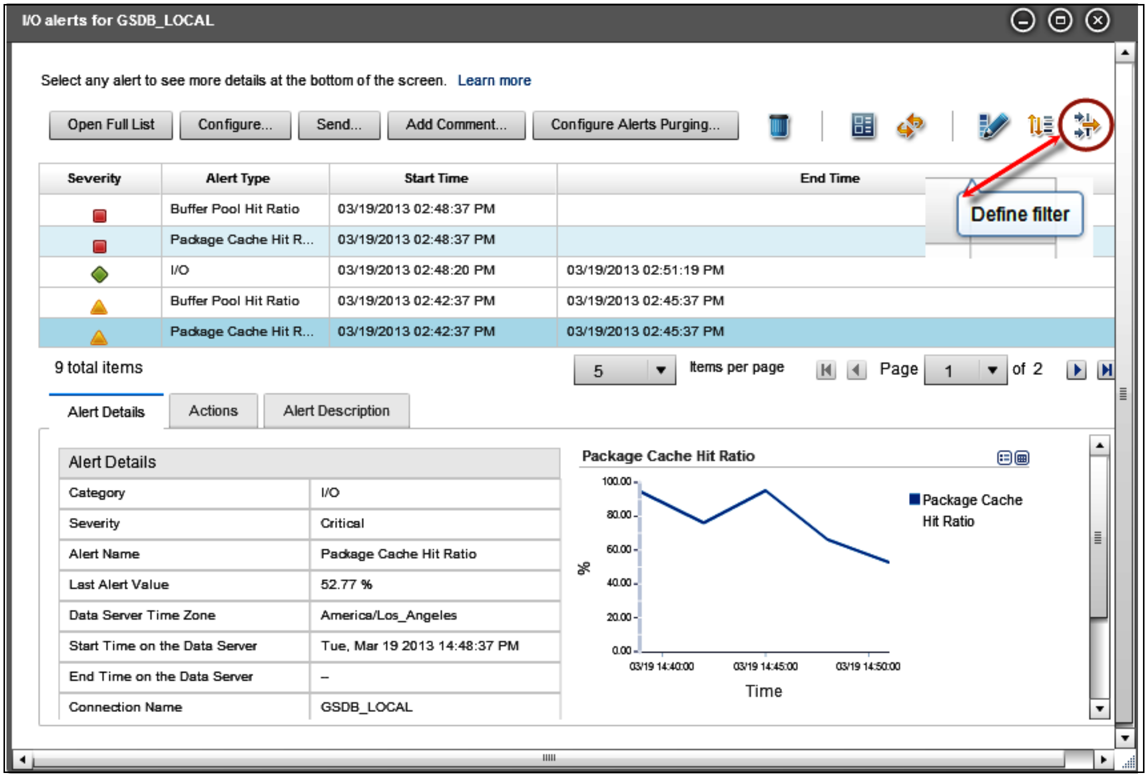

Figure 6-1 I/O alerts list launched from the Health Summary dashboard

Next, we do a survey of the metrics that can cause I/O bottlenecks in the system.

Buffer pool hit ratio

You typically see the buffer pool hit ratio alert generated several times. The buffer pool hit ratio is a good metric for monitoring I/O. A low hit ratio indicates that part of the data requested by the application was not in the buffer pool. As a result, the data had to be read from the disk rather than the buffer pool. Reading data from the disk requires more time and resources.

Clicking any one of the buffer pool hit ratio alerts opens the lower half of the window and gives you more details about this alert. Here, you see a graph that shows a steady degradation and not just a downward spike. The values are also well below what was set as a critical threshold of 90%.

An alert is reasonable here, but some systems might have lower buffer pool hit ratios than others. Therefore, consider that alert thresholds for buffer pool hit ratio and other metrics can be set to reflect the normal behavior of your system. After the metric goes outside of what is considered normal, the alert will be triggered.

Also, a low buffer pool hit ratio might result from several other possible root causes, such as a change in workload or an increase in volume of data. You can make a note to try to increase the buffer pool size, and rerun the workload. However, increasing the buffer pool size might address only one of the causes of an I/O issue.

The Actions Tab in the Alert Details (Figure 6-2) directs you to the Buffer Pool and I/O Dashboard, which is where you can analyze the problem with greater granularity.

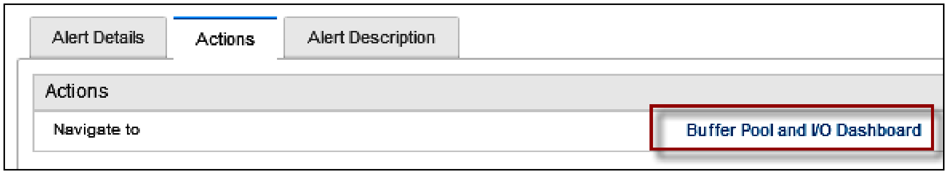

Figure 6-2 Actions tab

The Buffer Pool and I/O Dashboard (Figure 6-3) shows additional metrics that can be examined. Those metrics are described here in the context of the low buffer pool hit ratio problem.

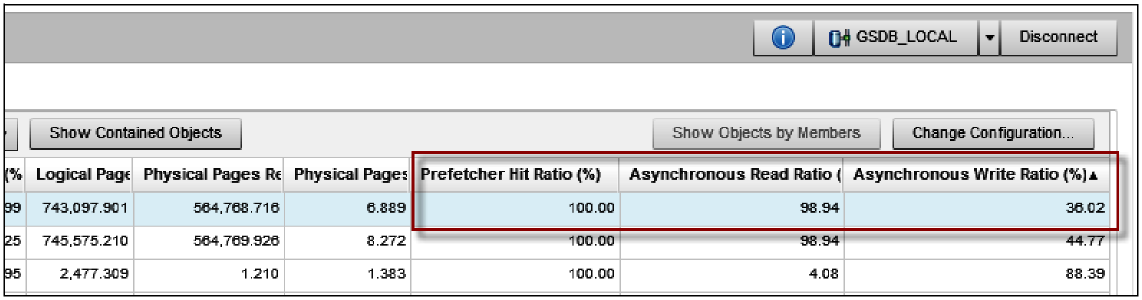

Figure 6-3 Buffer Pool and I/O Dashboard launched from Alert Details

Buffer pool asynchronous read ratio

This is the ratio between the asynchronous reads and the total number of buffer pool reads. A high ratio indicates that the majority of data requested was read by prefetchers.

Prefetching is the retrieval of data (one or more extents of pages) from disk in anticipation of their use. This can significantly improve performance in SELECT statements, which scan many rows, by reducing the time waiting for I/O to complete. Setting the PREFETCHSIZE table space tells DB2 to place that number of pages in the buffer pool in anticipation of its use. The default PREFETCHSIZE is 32.

When DB2 has to scan many pages to find result sets, DB2 uses asynchronous prefetch I/O. A high percentage of asynchronous I/O indicates that DB2 is doing a lot of scanning to find result sets. Scans may occur when indexes are missing or are sub-optimally defined, or when indexing is not practical.

A low asynchronous read ratio might indicate that no large scans are happening, or if large scans are happening, the prefetch size is too small. Figure 6-3 on page 267 shows that the ratio is above 98.94%, which is really good for asynchronous read ratio. The high value in the example indicates that the majority of the data requested was read by prefetchers.

Buffer pool asynchronous write ratio

This is the ratio between the asynchronous writes and the total buffer pool writes. A high ratio indicates that the changed data in the buffer pools is written asynchronously by the page cleaners to the disk and that the applications are not waiting. A low value might indicate that you do not have enough page cleaners, or that parameters like SOFTMAX or CHNGPGS_THRESH might need to be tuned.

The last column in Figure 6-3 on page 267 shows that the asynchronous write ratio is about 36%, which indicates that the page cleaners are not doing a good job. Before trying to fix that by tuning the system, be sure that buffer pool writes are a significant portion of the activity on the system. If the system is heavily skewed to reads, then trying to optimize a few writes is not a good use of time. You can look at other metrics to analyze the overall health of the system.

Rows Read per Fetched Row metric

Although the Rows Read per Fetched Row metric is not an I/O metric, this alert is one of the Workload alerts generated. Note that the Rows Read per Fetched Row alert occurred during the same time frame as the I/O alert, which is interesting. This alert is from the Workload category, not from the I/O category. However, it can affect I/O, so you should consider this alert here.

This alert means that DB2 is reading more rows for every fetched row than it normally does, and, therefore suggests that the access to data in the system is now more commonly using a table scan than an index, at least as compared to the normal way the system operates. This change in behavior can definitely affect buffer pool I/O. Figure 6-4 shows how high numbers of rows that are read tie into the high buffer pool activity you are seeing.

The figure shows that the Rows Read per Fetched Row exceeded the critical threshold, as evidenced by the red squares. But is this value much higher than normal? The baseline comparisons are key here, to know that there is really something wrong. For the Rows Read per Fetched Row metric, you can tie it back into the I/O problem by comparing it to a baseline. Baselines are described in the next section.

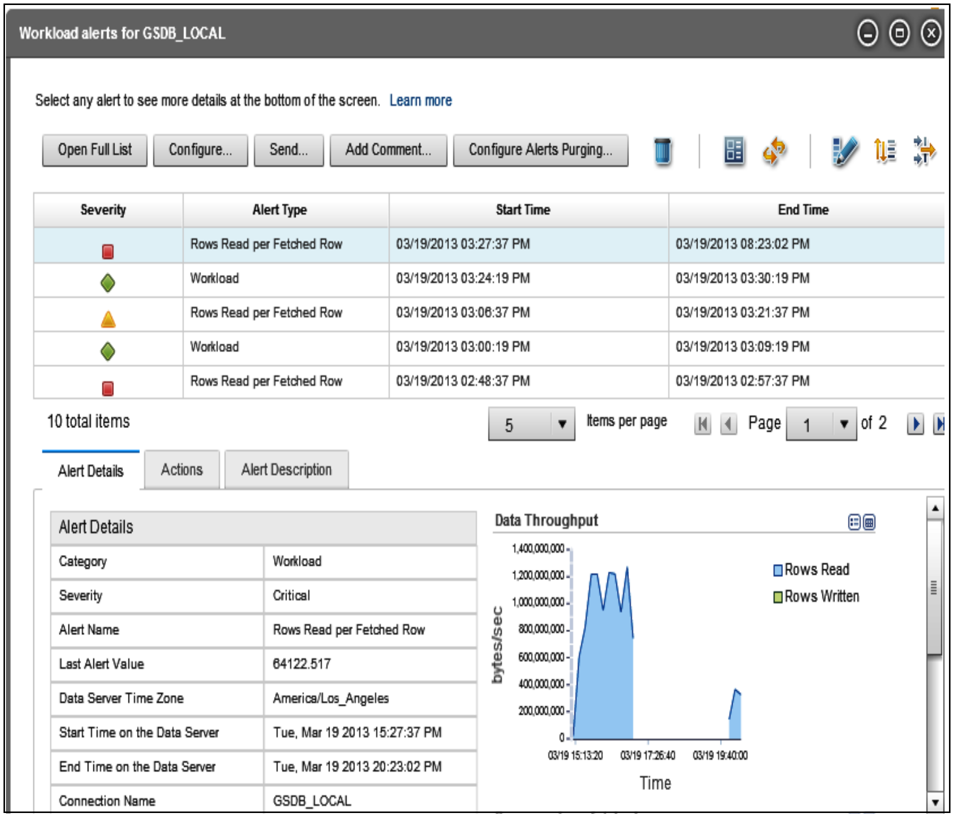

Figure 6-4 Workload alerts

You can check back on the Buffer Pool and I/O Dashboard for details about the I/O activities in the buffer pool, table spaces, and tables in the database. Figure 6-5 shows the number of logical reads and the physical reads per minute for the GOSALES_BP buffer pool, which is the buffer pool used by the application.

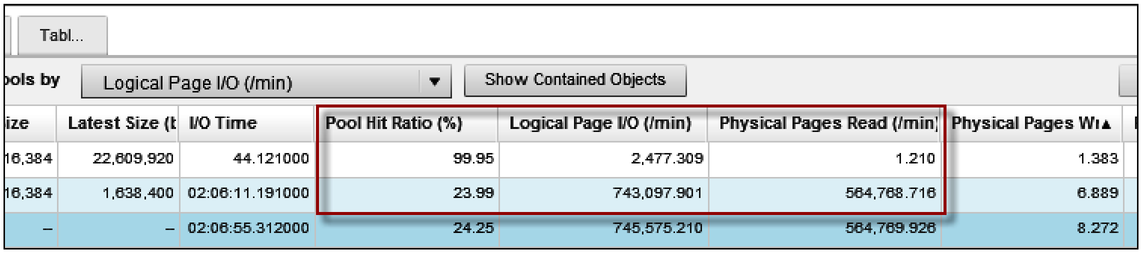

Figure 6-5 Buffer Pool and I/O Dashboard, logical and physical reads

You can see that the physical reads are lower than the logical reads, which indicates that DB2 was able to obtain the rows in the buffer pool without doing I/O. This observation is almost always true, so you cannot make any deductions from this metric.

Continue the survey and check the table spaces that uses this buffer pool by selecting the Table Spaces tab (Figure 6-6). An alternative approach to come to this panel is from the Buffer Pool and I/O Dashboard, select the buffer pool, and click **Show Contained Objects**.

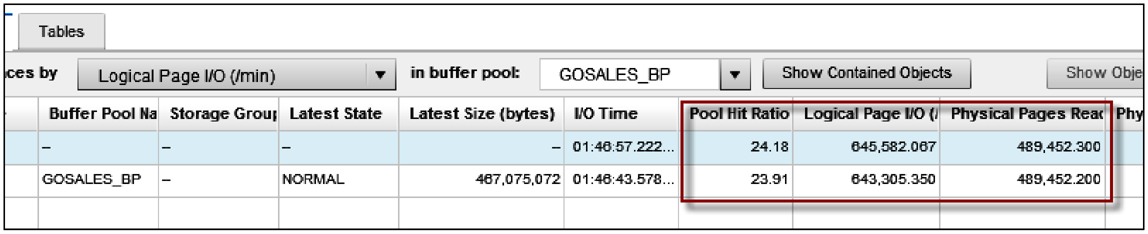

Figure 6-6 Buffer Pool and I/O Dashboard, Table Spaces tab

You can see that there is only one table space associated with this buffer pool, GOSALES_TS. The high physical and logical read values show that a lot of activity occurs in this table space, which also explains the high level of activity in the buffer pool that hosts it.

Continue the survey by looking at the activity details of the tables associated with the GOSALES_TS table space. You can either select the Tables tab (Figure 6-7) or select GOSALES_TS and click **Show Contained Objects**.

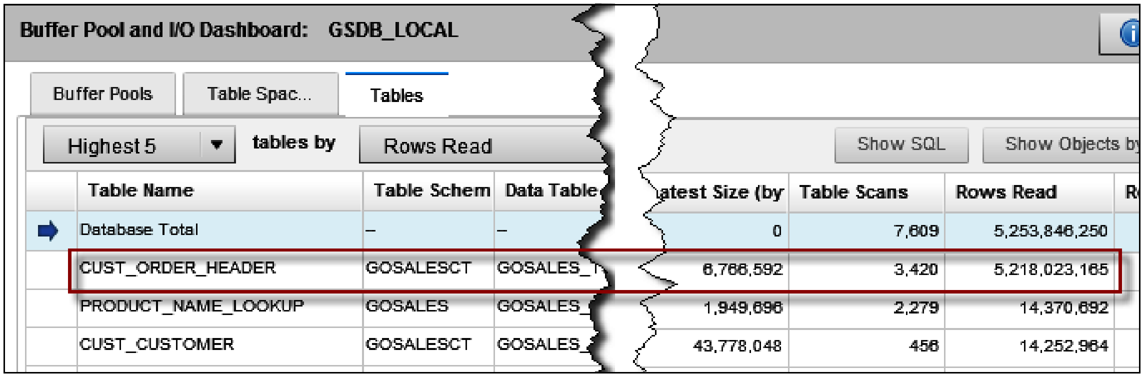

Figure 6-7 Buffer Pool and I/O Dashboard, Tables tab

You can see that one table has significantly more rows read and table scans than the others in the list, and that is the CUST_HEADER_ORDER table. This table had the most rows read, and a higher number of table scans than the larger CUST_CUSTOMER table, which ties you back to the Rows Read per Fetched Row alert you see in Figure 6-4 on page 269. Continue to analyze what SQL statements are accessing this table by selecting the table, and then clicking **Show SQL**, which transfers you to the SQL Statements dashboard. The SQL Statements dashboard is described in 6.4, “Monitoring statement execution” on page 286.

6.2.3 Diagnosing I/O problems using the Overview Dashboard

Prior to InfoSphere Optim Performance Manager V5.1, drilling down from the alert details to the other dashboards was how you typically diagnosed the problem. This is still a good strategy. Starting with InfoSphere Optim Performance Manager 5.1.1, you can also verify if the metric, in this case, the buffer pool hit ratio, is indicative of an excessively small buffer pool size, or if a certain workload had triggered the problem, by comparing your current metrics to a baseline.

Look at the Overview Dashboard. The At-a-Glance tab gives a system-wide view of the database metrics. Furthermore, you can compare the current execution metrics to a baseline here.

Chapter 4., “Getting to know InfoSphere Optim Performance Manager” on page 145 describes how to create a baseline. The choice of a baseline has a big impact on how the Overview Dashboard flags the differences, so you want the choice to be correct. On the positive side, you can always change the baseline, forward or backward, any time you want. A poor choice of baseline is one that selects a time when things were not normal or good, or is one that gives too many “false positive” results, that is, it flags differences that are not significant. The baseline should represent a time of typical or good performance.

Good advice on choosing a baseline is to be sure to get a baseline that is long enough to have some typical variability within it. That way, day-to-day fluctuations in monitored database values will not regularly exceed the thresholds for marking them out-of-bounds. But if a bottleneck develops, it will continue to display as a colored bar in the Overview Dashboard.

On the right side of the Overview are the system metrics. Examine the DB2 Time Breakdown, as shown in Figure 6-8. At a high level, an indication of time spent in DB2 shows a big spike in I/O time, with the baseline average increasing from 69% to 96%.

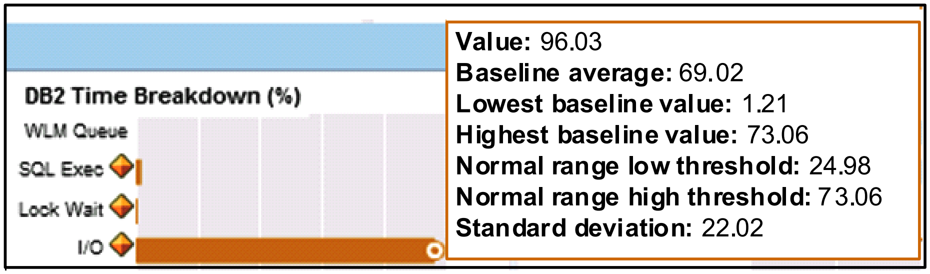

Figure 6-8 DB2 Time Breakdown

In the Performance Focus area of the Overview Dashboard, the average time for I/O operations, which is the time for synchronous reads (that is, normal, non-prefetched buffer pool reads), has increased from approximately 10 ms to over 17 ms, as shown in Figure 6-9 on page 273. This is indicates that there is a possible disk I/O bottleneck.

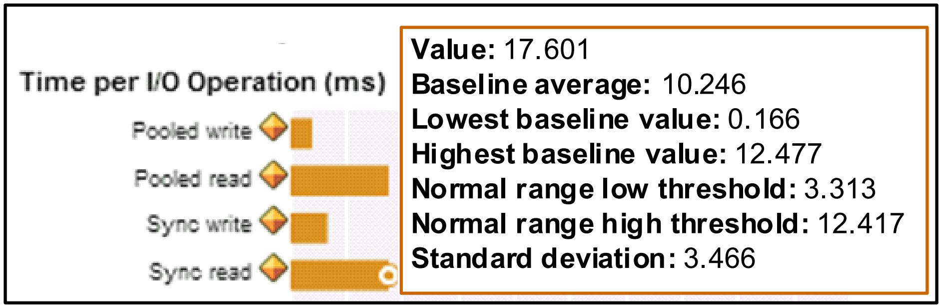

Figure 6-9 I/O time comparison to baseline

As in 6.2.2, "Assessing database health-related problems using alerts" on page 265, there is a way to drill down and diagnose this problem more closely using the Buffer Pool and I/O Dashboard. A link to this dashboard is available from the Overview Dashboard.

The Buffer Pool and I/O Dashboard gives you details about the I/O activities on the buffer pool, table spaces, and tables on the database. The Size and I/O Times tab offers useful graphs of I/O response over time, which can be used to see if this is just a spike of bad performance, or a longer lasting issue.

Figure 6-10 shows that in the time per I/O operation is higher than the baseline. In the Buffer Pool and I/O Dashboard, the average read time per I/O is also high, as shown in Figure 6-10. Does this means the disk has become slower, or maybe just busier?

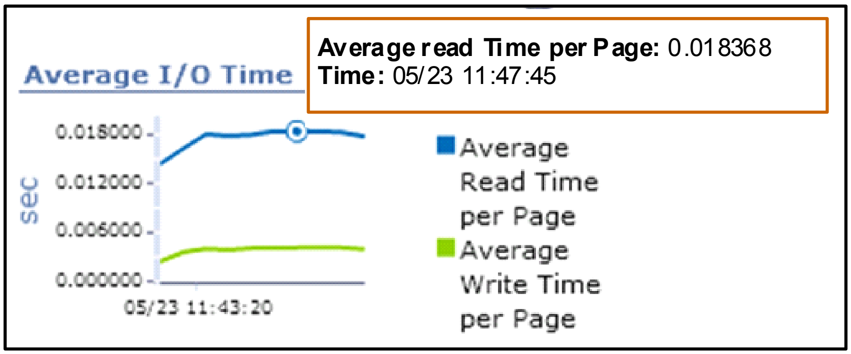

Figure 6-10 Average I/O Time

Inspect the rate at which you are doing disk reads, and then compare it with the rate at the baseline interval. If you move the time slider back to the baseline time (or perhaps open another instance of InfoSphere Optim Performance Manager at the baseline time,) you will find that the current rate of disk reads is much, much higher than at the baseline. This suggests that the longer I/O times are likely because of a much greater level of disk activity. You will certainly want to learn more about this extra activity.

6.2.4 Drilling down to the table space and tables

In the Buffer Pool and I/O Dashboard you can check the buffer pools, table spaces, and tables on the database. Open the Table Space tab to display the table spaces in the buffer pool, IBMDEFAULTBP. An alternative approach to get to this panel is from the Buffer Pool and I/O Dashboard, select the buffer pool, and click **Show Contained Objects**.

Figure 6-11 shows the Table Space tab and highlights the physical reads per minute for the table spaces contained in buffer pool, IBMDEFAULTBP,

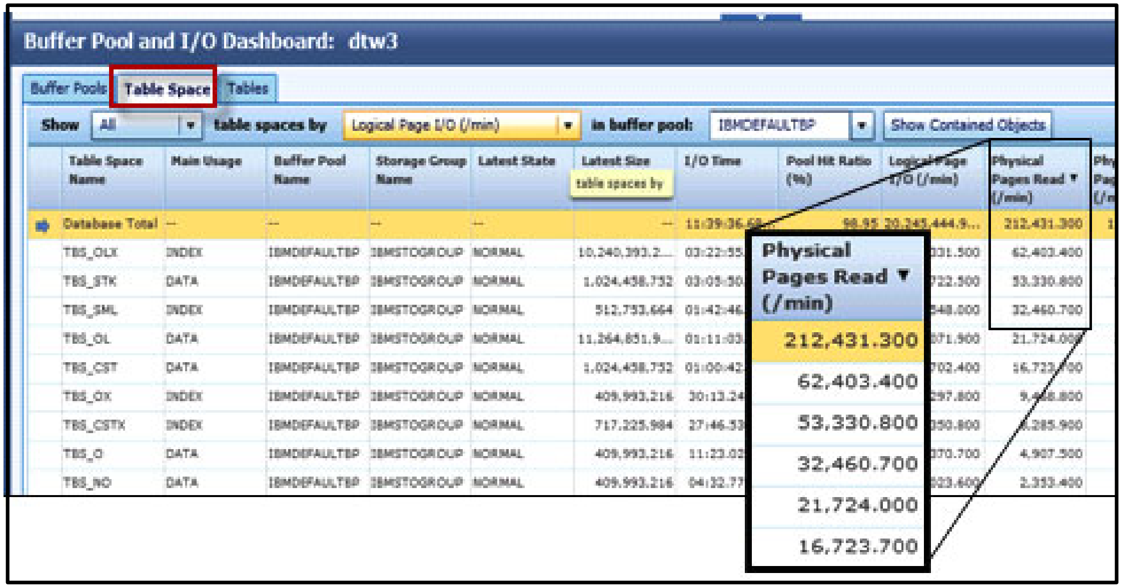

Figure 6-11 Buffer Pool and I/O Dashboard, Table Space tab

You can see that the physical reads are evenly spread over all the table spaces. No individual table space seems to be causing the increase in the physical pages read metric from the baseline in Figure 6-11.

To see the activity details of the tables within the table spaces, select the Tables tab or click **Show Contained Objects**. Figure 6-12 shows all tables for all table spaces in the IBMDEFAULTBP buffer pool.

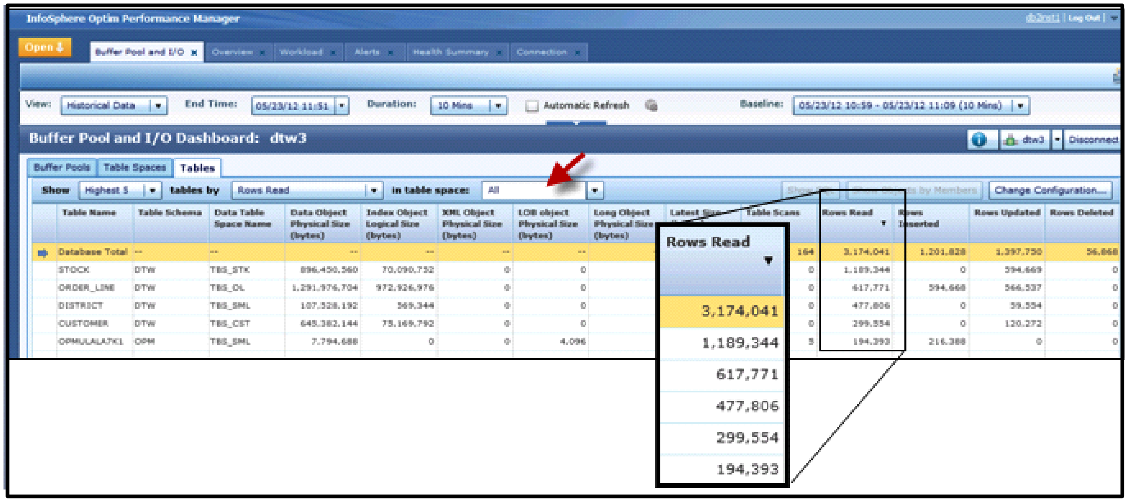

Figure 6-12 Buffer Pool and I/O Dashboard, Tables tab

You see that the STOCK table has a high number of rows read compared to the other tables. This is generally referred to as a "hot table." However, you must check if this is typical for that table or if there was a sudden spike in activity for this table. Compare this metric with the baseline, as shown in Figure 6-13.

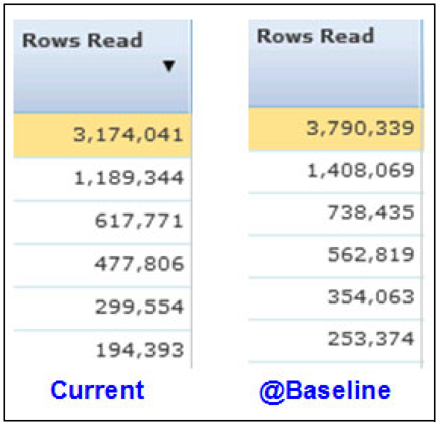

Figure 6-13 Comparison of current number of rows read to baseline for table STOCK

What is surprising is that there is not a big difference in read activity for this table. This suggests that either the database suddenly got bigger, or more likely, the buffer pool is not keeping up.

Staying on the same dashboard, you can now open the Hit Ratio tab, which gives you a graph of buffer pool hit ratio over time, and in different types of buffer pool I/O. Hover the mouse pointer over the various lines, and see that the data hit ratio is 78%. Although this value might be perfectly normal, using the same technique again, of comparing with the value at baseline time, you can see that when performance was normal, data hit ratio was slightly better at 94%. (Figure 6-14).

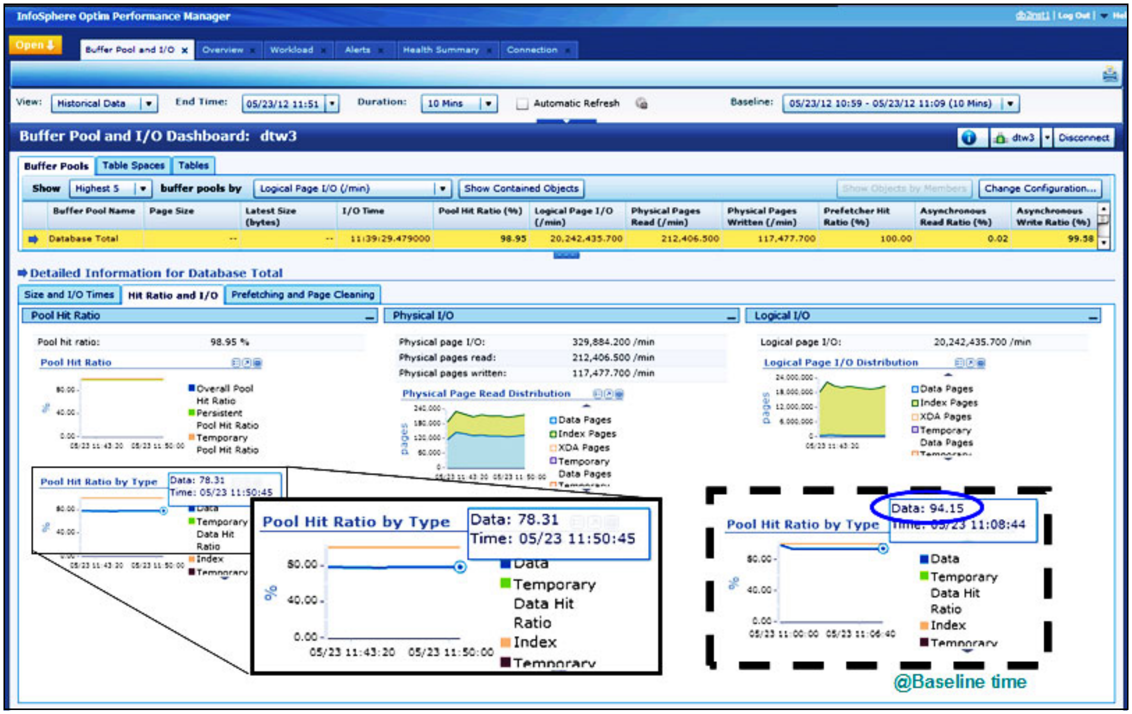

Figure 6-14 Overview Dashboard, Hit Ratio and I/O tab

What has been analyzed so far is summarized in the following list:

- From the Overview Dashboard, there is a 69% to 96% increase in percent time spent on I/O, and a 10 ms to 17 ms increase in average pool read time;
- From the Buffer Pool and I/O Dashboard, there is approximately 36,000 to 212,000 increase in BP reads per minute and 94% to 78% reduction in BP data hit ratio.
- There is no apparent table space or table hotspot.

Consider the following two possible causes:

- One or more bad SQL statement plans. This however is not likely because there is no one really active table in the table space
- Poor disk configuration. This is possible, but buffer pool configuration must be checked first.

6.2.5 Resolving high I/O problem to improve performance

IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows, SG24-7925 describes how creating an index resolves the I/O performance problem. Starting with InfoSphere Optim Performance Manager v5.1, you now have an integration point with IBM Data Studio and InfoSphere Optim Query Workload Tuner. You can perform single query tuning or workload tuning with a click of the button.

From the SQL Statements dashboard, when you select a SQL statement, the Statement Details are shown below the grid. Select **Actions** → **Tune**, below the the SQL Statement text window, to capture the query and send it to a query project in InfoSphere Optim Query Workload Tuner. Chapter 3, "Installing and configuring Optim performance management tools" on page 69 describes the installation and configuration of InfoSphere Optim Query Workload Tuner. InfoSphere Optim Query Workload Tuner must be installed and opened on the same machine where the browser to access the InfoSphere Optim Performance Manager web console is launched. You can perform single query tuning in IBM Data Studio. However, to use the index advisor or to perform workload query tuning, you must install the InfoSphere Optim Query Workload Tuner license on the monitored database. For more details about SQL query tuning, see Chapter 3, "Installing and configuring Optim performance management tools" on page 69.

The previous scenarios showed that increasing the buffer pools might resolve the high I/O problems. For this and other types of DDL changes that might benefit the monitored database, IBM Data Studio Administration Explorer allows you to generate the DDL needed to alter the size of your buffer pools. You can choose to generate the DDL, save it in a script or change plan for review, and then schedule to run the script at a later time.

6.2.6 Avoiding high I/O

In a production environment, database performance tuning is a complicated task and should be entered into for both gaining performance and avoiding future performance problems.

Optim Performance Manager, Optim Query Workload Tuner, Optim Configuration Manager, and IBM Data Studio work together to provide you with information and assistance to build performance from the ground up.

With InfoSphere Optim Performance Manager reports (predefined and custom reports), you can perform trend analysis and spot resources that might become constrained later. You can generate these reports periodically, specify the duration of time your report covers, and also keep these reports in the repository database for future reference.

Examine several of the following reports:

- Storage Group report (for DB2 v10 databases only)

  This report shows, for a specified period of time, an overview of the table spaces that are used by the storage groups. You can analyze table spaces that are frequently accessed (“hot”), less frequently accessed (“warm”), and rarely accessed (“cold”), and you can check their growth rates to plan for future disk space requirements. For each storage group, you can drill down to the Disk Space Consumption report to see details about table spaces, containers, and tables.

- Disk Space Consumption report

  This report provides an overview of the disk space usage by table space over time. You can examine details about each table space such as the size, tables contained, data skew and so on, over time.

- Table Usage report

  This report identifies hot tables or fast-growing tables that might cause disk contention or that might need reorganization.

- Database Configuration report

  Changes to your database configuration parameters, such as PREFETCHSIZE, NUM_IOCLEANERS, CHNGPGS_THRESH, and NUM_IOSERVERS, can affect the performance of your I/O and buffer pools. The Database Manager Configuration report provides an overview of the current database manager configuration and identifies the parameters that changed in a given time frame. It contains details about system management, system monitoring parameters, instance administration, capacity management, and communications.

Note: With Optim Configuration Manager, you can track database server configuration changes, and it can be set up to notify you when there are configuration changes.

When you want to implement changes to specific SQL statements identified in InfoSphere Optim Performance Manager Table Usage or Top SQL reports, you can use InfoSphere Optim Query Workload Tuner to test how your changes will affect performance. You can tune individual SQL statements that are not performing well, or tune multiple SQL statements as a workload. You can define your workload as the set of SQL statements that affect a specific table only. You can do the following tasks:

- Use the advisors to improve the query design of a single query.
- Virtually create an index and evaluate the improvements.
- Use the Workload Advisors to consolidate and optimize recommendations for the entire workload with regard to index and statistics recommendations.
- Compare access plans for single query statements or workloads.

Finally, when your application data grows, InfoSphere Optim Query Workload Tuner allows you to periodically check the health of the application to find potential problems earlier, before costly application outages. You can use the Statistics Advisor to verify whether your statistics are stale, and assess the data growth has on the performance of your workload.

6.3 CPU and memory related bottlenecks

High CPU utilization situations can slow the response time of your database system. If the CPU utilization is higher than normal or close to 100%, investigate this situation to find a reason why that is happening. The following examples are possible reasons:

- Execution of long-running or not well-tuned SQL statements
- Execution of DB2 utilities such as LOAD
- Memory problems such as under-sized DB2 memory areas or over-committed system memory

6.3.1 Diagnosing high CPU time from the Overview Dashboard

This scenario monitors the same database but with a different workload, using the baseline method to do the analysis. From the Overview Dashboard of this database, you can compare the current values to the baseline (Figure 6-15).

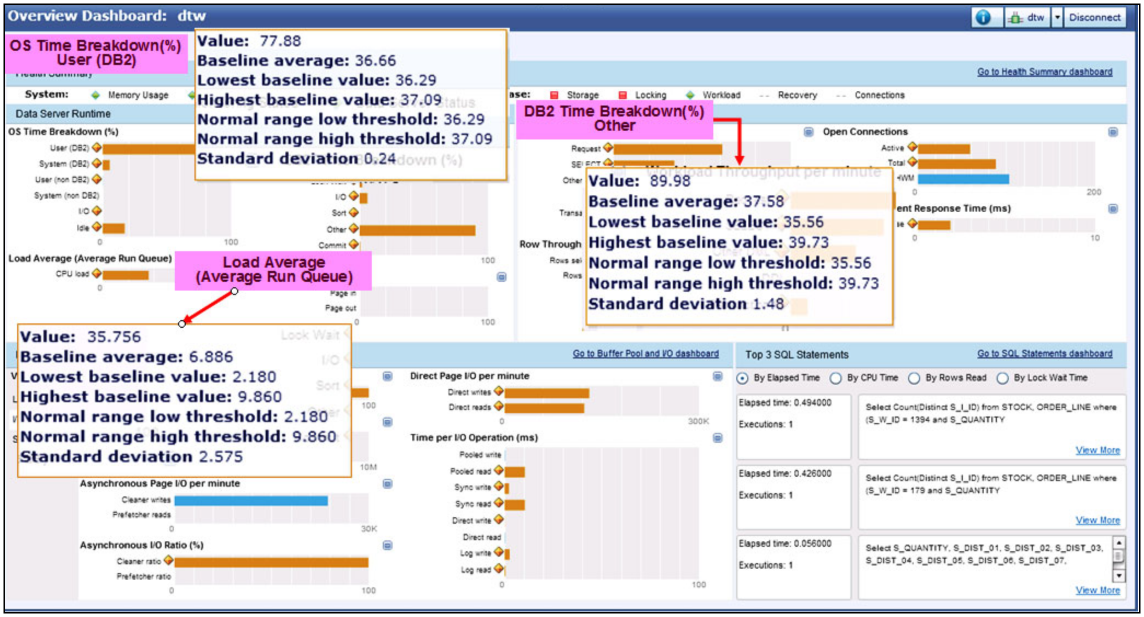

Figure 6-15 Overview Dashboard

Consider these steps:

1. Check the OS time breakdown. This metric indicates how time is spent from an OS perspective, user and system CPU time, both inside and outside of DB2, and also I/O wait and idle time. This metric shows that the percent of total CPU time for User (DB2) is currently at 78% (in the baseline it was 37%).
2. The load average (which tells you how long the run queue is – that is, the lineup of tasks for the operating system to execute), shows a current backlog or average backlog (run queue length) of about 36 threads, compared to 7 threads in the baseline (Figure 6-15). As with total CPU utilization, this is another good indicator of a busy system.
3. The Overview Dashboard also shows DB2 Breakdown, which is a breakdown of time spent inside DB2, derived from the new wait and component time metrics in DB2. These metrics are a useful way to see what DB2 is doing (or waiting for), such as executing SQL, waiting on locks, waiting on disk I/O, and sorting. There is a big shift from time spent in I/O, SQL Exec, and so forth to

“Other” (Figure 6-15 on page 280). What this represents is not yet clear from the Overview Dashboard.

Package Cache Hit Ratio

CPU saturation is frequently caused by being weighed down by many small issues rather than a few large issues. There might not be a small number of heavy statements – perhaps there are hundreds or thousands of only moderately heavy statements, but the aggregate effect is the same.

A useful place to look for clues to any sudden problem is the Alerts dashboard (Figure 6-16 on page 282). Here you can see that a Package Cache alert (a known CPU sponge) coincides with the period of high CPU. A low package cache hit ratio can indicate that much time is spent in the SQL compiler, which can lead to high CPU utilization.

Double-clicking the alert gives you much information about the alert, including a graph of package cache hit ratio. Hover the mouse over different points in time to determine the hit ratio at that time, which shows that there was a period of low package cache hit ratio. A low package cache hit ratio generally means that many SQL statements being executed are having to be compiled first, which is an expensive operation and can cause a CPU bottleneck.

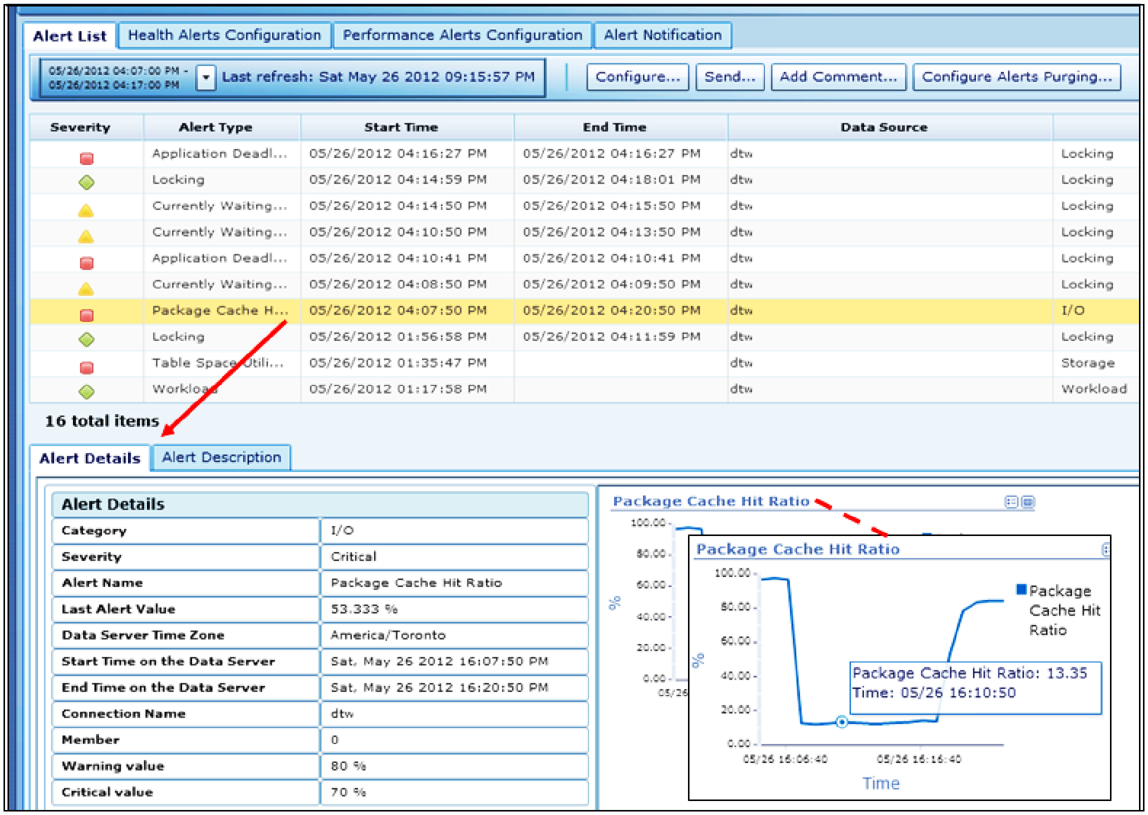

Figure 6-16 Alert List, Package Cache Hit Ratio alert

Low package cache hit ratio might also be explained by a big increase in the number of distinct SQL statements being executed. The SQL Statements dashboard is helpful, and you can choose the Execution Summary tab, to show the total metrics for all executions of each statement. You can also choose to display the statements in descending order of total CPU time. An interesting aspect, as shown in Figure 6-17 on page 283, is that the number one statement (in terms of total CPU for all executions) was executed only once.

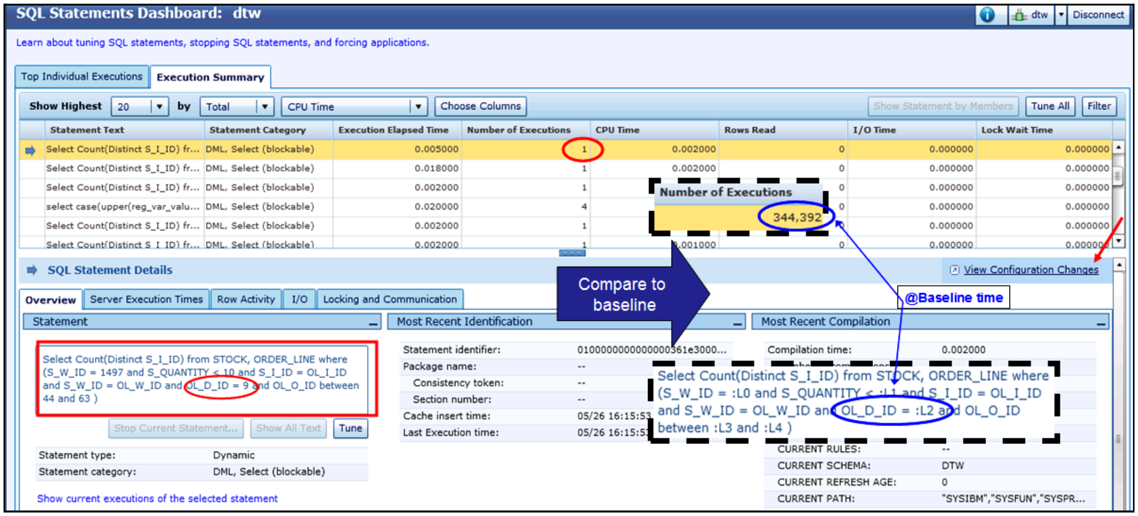

Figure 6-17 SQL Statements dashboard, Execution Summary tab

As before, a useful approach is to compare current metrics with baseline metrics. Here, you notice that the top SQL statement in the baseline was executed 344,000 times, which is much more than the current environment, where the top statement was executed once. This is particularly surprising when you consider that two top statements (baseline and current) are essentially the same, although not completely the same, because the current statement is using literal values for parameters, and the baseline statement uses host variables.

Although you can continue to diagnose this problem, if you take all the symptoms together (high CPU, low package cache hit ratio, and the particular parameter style found in the baseline period) it seems likely that the statement concentrator was being used during the baseline period, and now it is not. The statement concentrator, STMT_CONC, is a DB2 configuration setting that, when enabled, maps literal values to host variables.

You can use the Database Manager Configuration report to confirm when this setting was disabled, or click the **View Configuration Changes** link, at the bottom of the Statements grid (Figure 6-17). This link launches the InfoSphere Optim Configuration Manager. With this tool, you can view all the configuration changes in a specific time period.

6.3.2 Diagnosing memory bottlenecks

Now on to memory bottlenecks, which tend to appear in two main types (as shown in Figure 6-18). The first is more common but generally less severe. Essentially, it occurs when DB2 data structures, such as buffer pools, sort heap, and utility heaps, are undersized. This type is interesting, because it usually manifests itself as either a CPU problem (if you do not have enough memory to cache things, you might have to recalculate them), or a disk problem (if you do not have enough memory to cache things, you might have to re-read them from disk.) Figure 6-18 shows two common examples: a memory bottleneck in the package cache that shows up as a CPU issue because of extra SQL compiles, and a memory bottleneck in the buffer pool that shows up as a disk bottleneck.

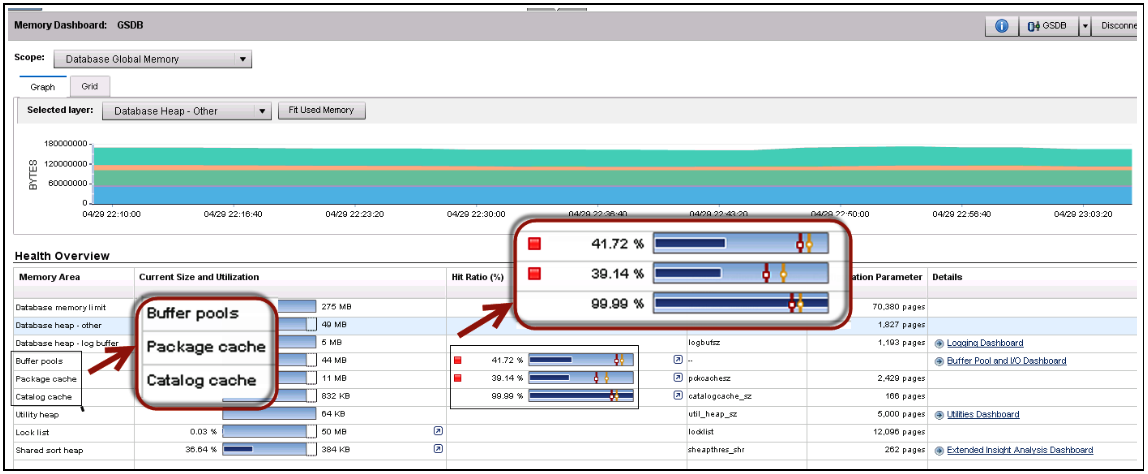

Figure 6-18 Metric sections on the Memory dashboard

A low package cache hit ratio can indicate that much time is spent in the SQL compiler, which can lead to high CPU utilization. The low package cache hit ratio coincides with a high CPU utilization. You see this high utilization for the same time frame in the System dashboard (Figure 6-19).

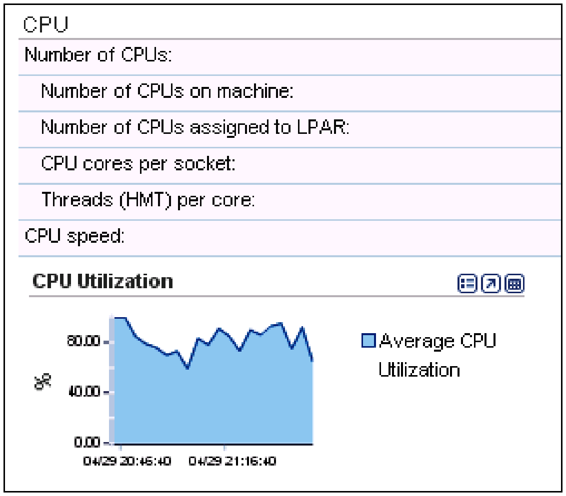

Figure 6-19 CPU utilization in the System dashboard

To investigate further what might be causing the high CPU utilization, you can check quickly what utilities are executing; perhaps they are contributing to the high CPU utilization. *IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows*, SG24-7925 describes how to monitor utilities in InfoSphere Optim Performance Manager. In 6.4, "Monitoring statement execution" on page 286 use of the Overview Dashboard and baselines for monitoring statement executions continues.

The other symptom of a memory bottleneck occurs when system memory (including DB2's memory) is over-committed. That is, the combined memory consumption of all of DB2's shared and private memory areas (plus whatever is consumed by other things on the system) are more than the system can afford, and as a result the system starts paging.

You can also monitor the paging rate (pages in and pages out) in the Overview Dashboard, in the Data Server Runtime section.

6.4 Monitoring statement execution

System-wide bottlenecks of the types described (disk, CPU, memory) are fairly common, but there are of course many cases where individual statements run slowly, in the context of a system that is running well overall.

Again, starting with the Overview Dashboard you see a reduction in overall system throughput. The baseline had a throughput of about 385,000 SQL requests per minute, and with the bottleneck, it is now reduced to about 75,000 (Figure 6-20).

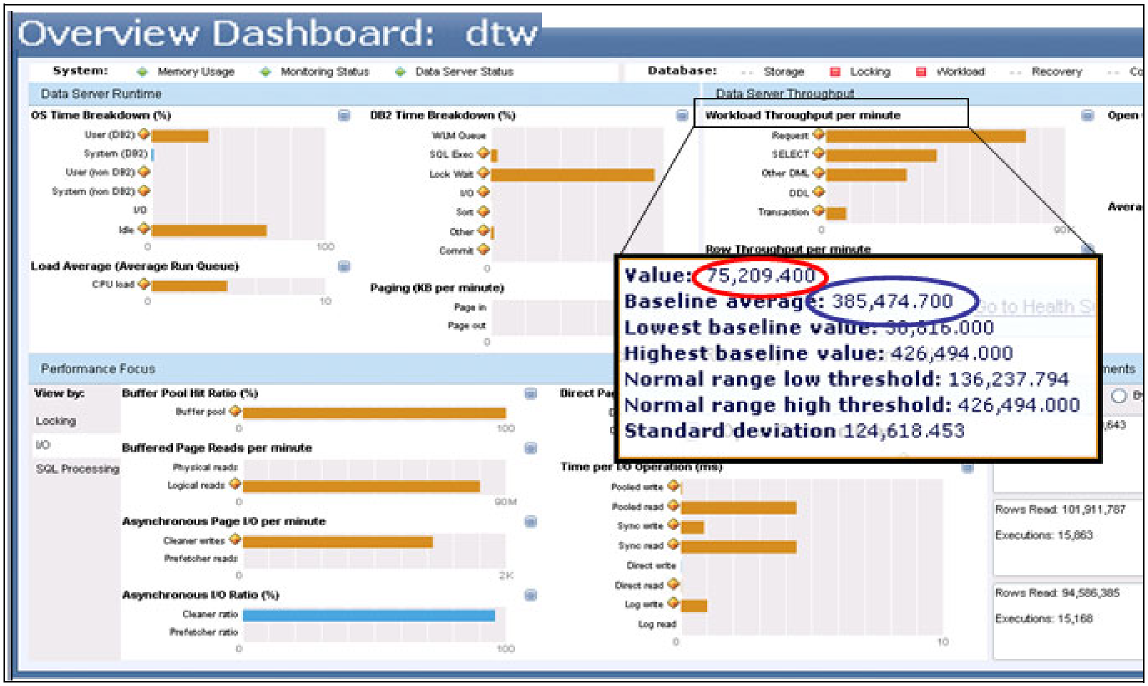

Figure 6-20 Overview Dashboard, Workload throughput per minute

Also, average statement response time, which was fast at 0.23 ms in the baseline, has increased to 58 ms (Figure 6-21).

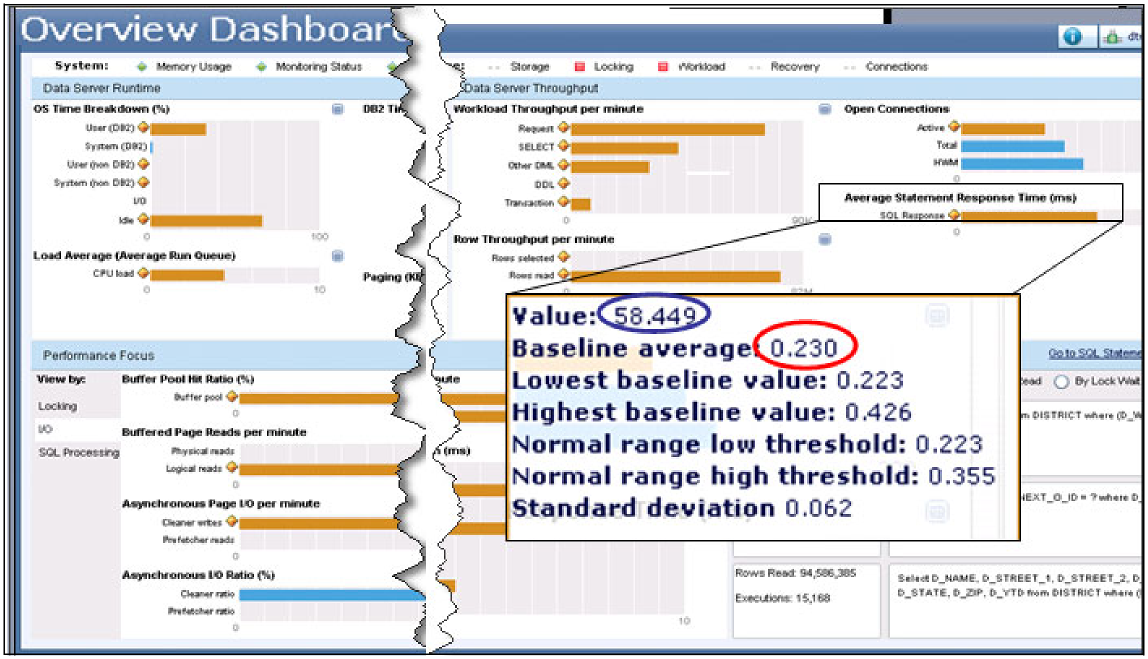

Figure 6-21 Overview Dashboard, Average Statement Response Time

In addition, the Overview Dashboard also includes CPU usage, provided by the operating system. On the baseline, it was 17% and now it is increased to 32%. This is especially interesting, because the statement throughput has decreased dramatically, yet more CPU is being used (Figure 6-22).

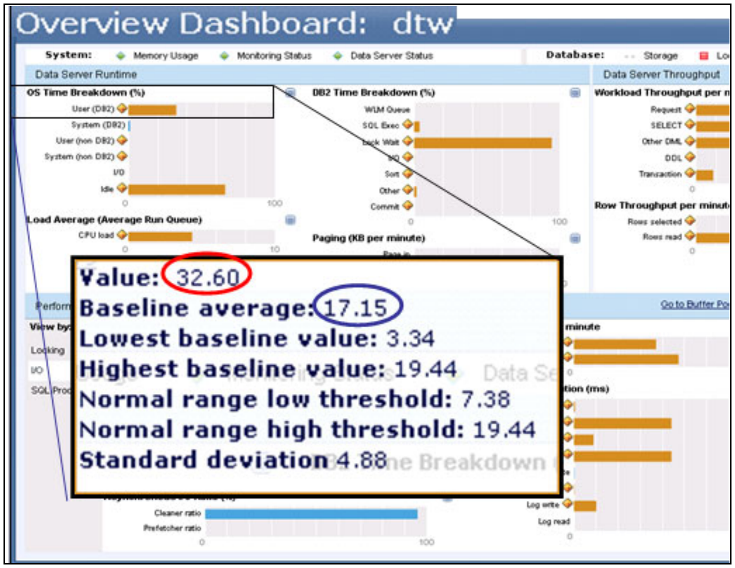

Figure 6-22 Overview Dashboard, OS Time Breakdown

The previous observations were expected (given the lower throughput observed in the application) and not surprising. However, here something might be helpful. In the breakdown of time spent in DB2, you can see that the lock wait time in the baseline was 7%, and this increased to 94% in the current measurements. (Figure 6-23 on page 289).

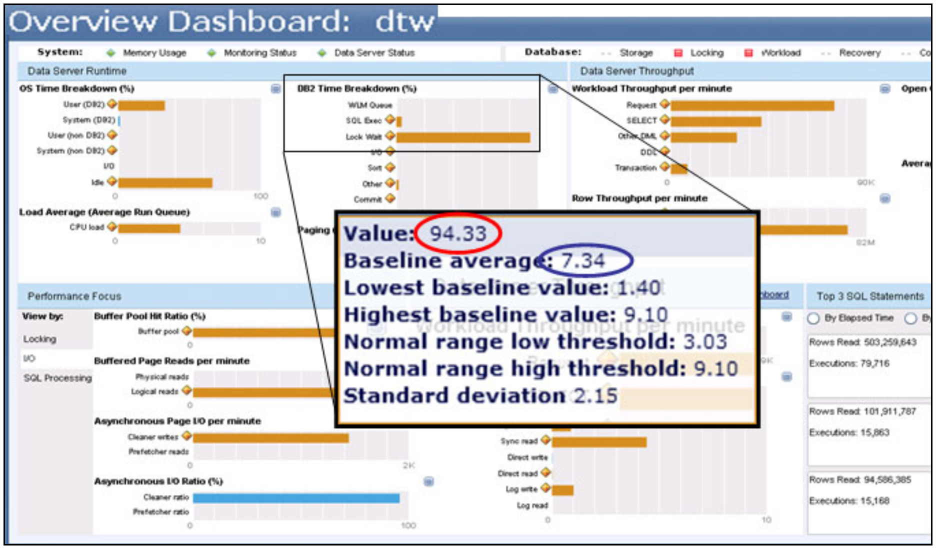

Figure 6-23 Overview Dashboard, DB2 Time breakdown

And to go along with that massive increase in lock wait time, is a very large decrease in the amount of time spent actually doing work. SQL Exec time in DB2 has dropped from 41% in the baseline, to about 4% now (Figure 6-24).

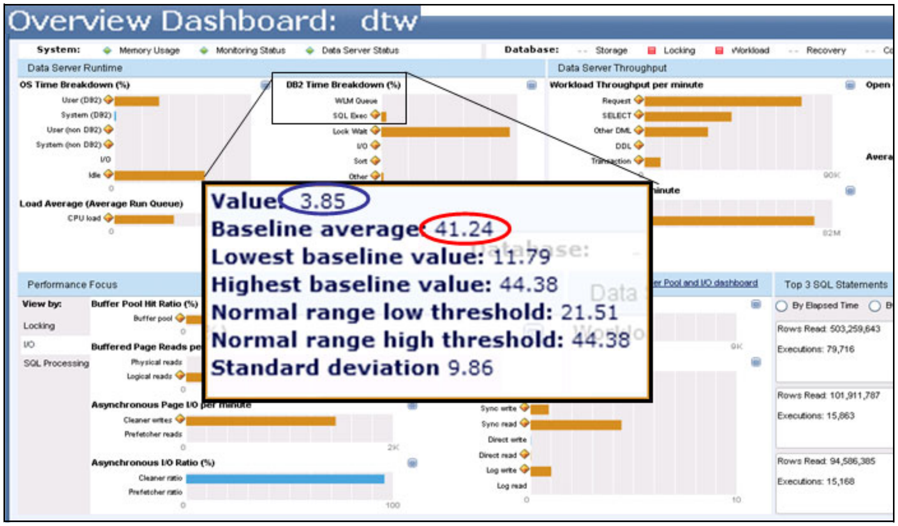

Figure 6-24 Overview Dashboard, SQL Time in DB2 Time Breakdown

So far, with this problem occurring, there are drops in statement throughput and time spent in useful work, and increases in lock wait time and CPU usage.

Another item to consider is a change in the row throughput per minute. In the baseline, there were about 296,000 rows per minute, and now with the lower overall statement throughput, there is an increase in row throughput, to about 74 million rows per minute (Figure 6-25). This increase in rows read might explain why there is an increase in CPU, although the total number of statements executed has dropped.

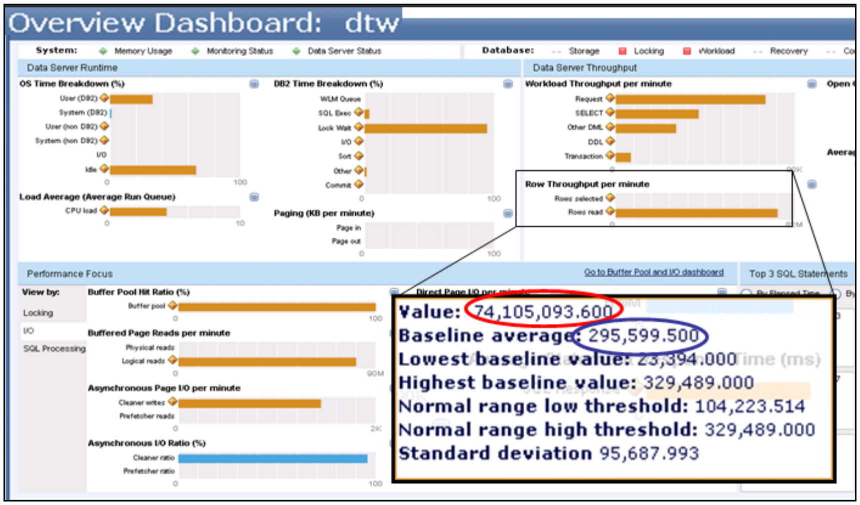

Figure 6-25 Overview Dashboard, Row Throughput per minute

The Overview Dashboard so far tells you that fewer statements are executed but more rows and higher CPU and higher locking. So far, compared to the baseline, you have seen the following changes:

- A reduction in the overall system throughput,
- An increase in the average statement response time
- An increase in CPU usage
- An increase in the lock wait time
- An increase in rows read, and
- A drop in number of statements executed

In fact, there is still one more piece of information to learn here, by looking at the Top 3 SQL Statements section, part of the Overview Dashboard shown in Figure 6-26 on page 291. In the baseline data (again, outlined as a dashed frame), you see that the top statement by rows read drove 547,000 rows reads in 547,000 executions: so, one row per execution. But the new top statement by rows read goes through 503 million rows in about 80,000 executions: much higher than one each.

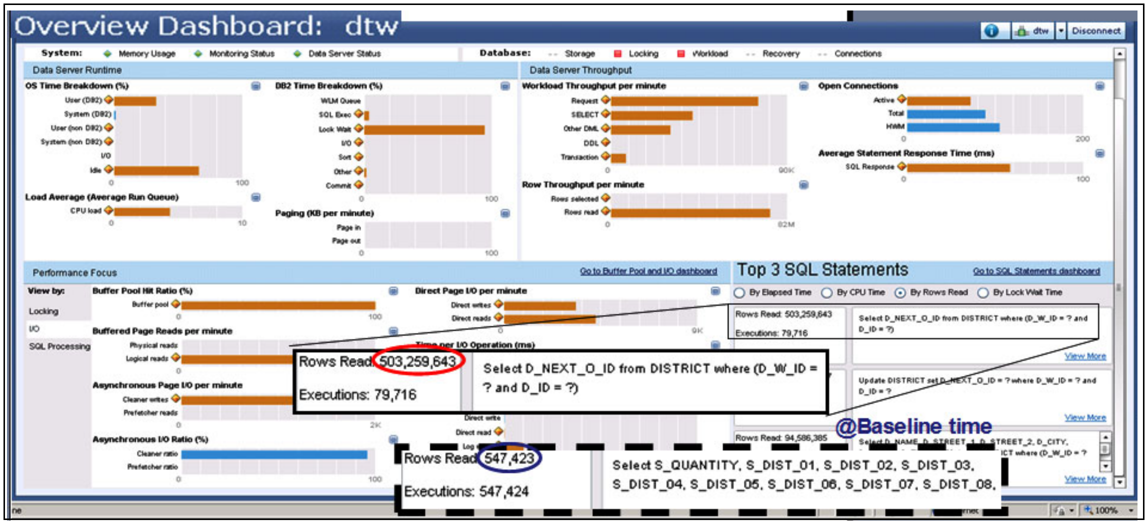

Figure 6-26 Overview Dashboard, top three SQL statements

Because an increase in rows read at the system level is factored into the high-level analysis, following this lead makes sense. Therefore, click the link that takes you to the SQL dashboard for further analysis of this statement.

6.4.1 The SQL Statements dashboard

From the Overview Dashboard, drill down to SQL Statements dashboard to look for more information about those same heavy statements, and compare current values against how things were at the baseline time.

The SQL Statements dashboard shown in Figure 6-27 on page 292, has two tabs:

- Top SQL Executions displays the top N statements that run in the specified monitoring interval.
- Execution Summary indicates the most executed SQL statements.

Chapter 4, "Getting to know InfoSphere Optim Performance Manager" on page 145 describes the metrics that are displayed in the SQL Statements dashboard and how they are used in analyzing statement execution performance.

The Execution Summary tab gives you information about each statement from the DB2 package cache, divided to show values of rows read, elapsed time, and so on, per statement execution. As seen here, each execution of the heaviest statement currently takes about 0.38 seconds (or 380 ms) and reads over 6000 rows. However, during the baseline period, it was much faster at 0.1 ms per execution and reading *one* row each time.

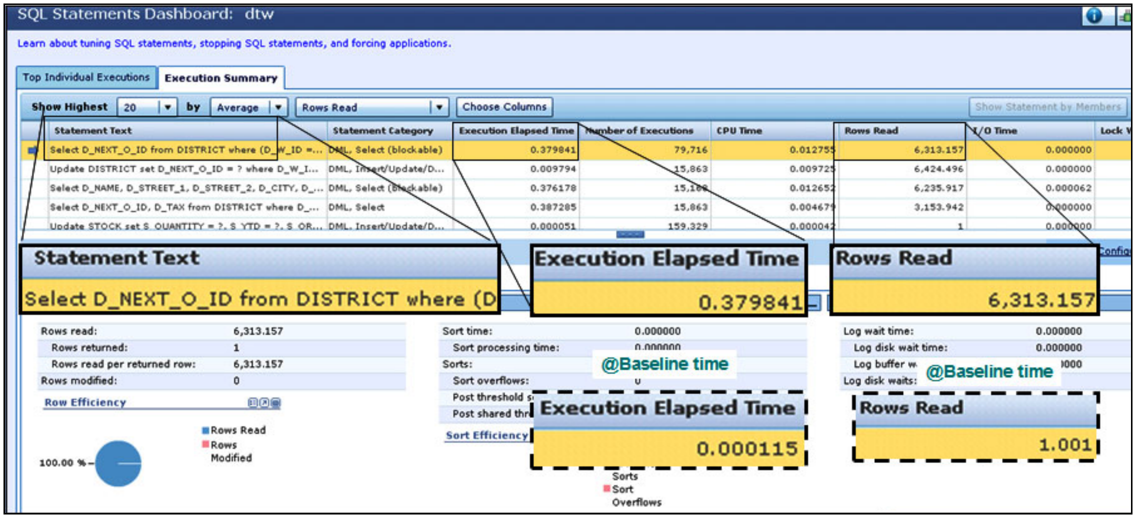

Figure 6-27 SQL Statements Dashboard, Statement Text, Execution Elapsed Time, and Rows Read

If you select the first statement in the SQL Statements Dashboard, the bottom part of the dashboard fills with details about the statement. You can perform several actions on this statement, one of which is to tune it. In 6.4.3, "Resolving high CPU-consuming statements" on page 294, you continue the analysis using InfoSphere Query Workload Tuner.

6.4.2 Using the SQL Baseline Comparison report

Starting with InfoSphere Optim Performance Manager V5.1.1, the SQL Baseline Comparison report is available. It is shown in Figure 6-28 on page 293. The report can be used to further analyze the heavy statement, and also more easily determine whether any other problematic statements are lurking. The report is particularly handy for a situation like this (a degradation in performance), because it gives you a side-by-side comparison of statement metrics between the current time (or whatever was chosen on the time slider) and a baseline time.

To generate this report, select **Reports** → **Pre-defined Reports** from the Open menu, then select **SQL Baseline Comparison report** from the pull-down list of predefined reports. In the report generation page, choose the two time periods to compare, what metric to sort on (in this case, it is rows read, to build on what is already known about the problem being seen), and whether regressions (degradations) or improvements, or both, are shown. We focus on regressions for now.

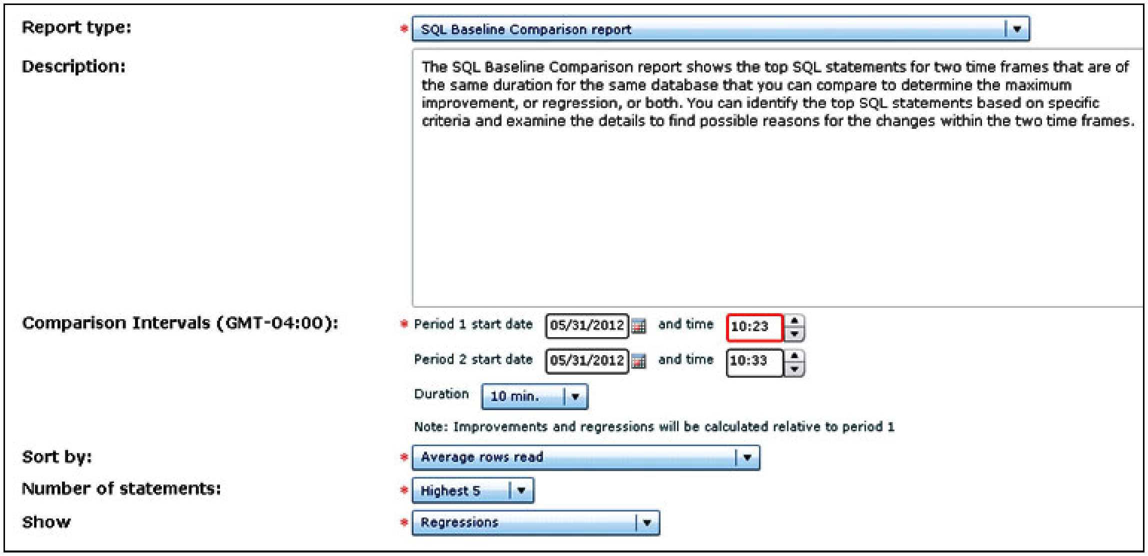

Figure 6-28 SQL Baseline report

When the report is run, it shows the top five most degraded statements, in terms of rows read. However, there are already several substantially degraded statements, not only one or two. They are displayed side-by-side (Figure 6-29 on page 294): the baseline value for average rows read, and then the report interval, followed by red bars for both absolute value and percent changes. Picking out the top two statements, you can see they have each degraded by about 18x (in fact, all of the top four statements are similarly impacted.) The leftmost column shows the statement texts; all of the top four impacted queries operate on the DISTRICT table. Such a large increase in rows read, across the board for one table, tends to suggest that these queries are not using indexes any more, which might be because of such things as bad statistics or a dropped index. InfoSphere Optim Performance Manager provides a useful way to take the next steps.

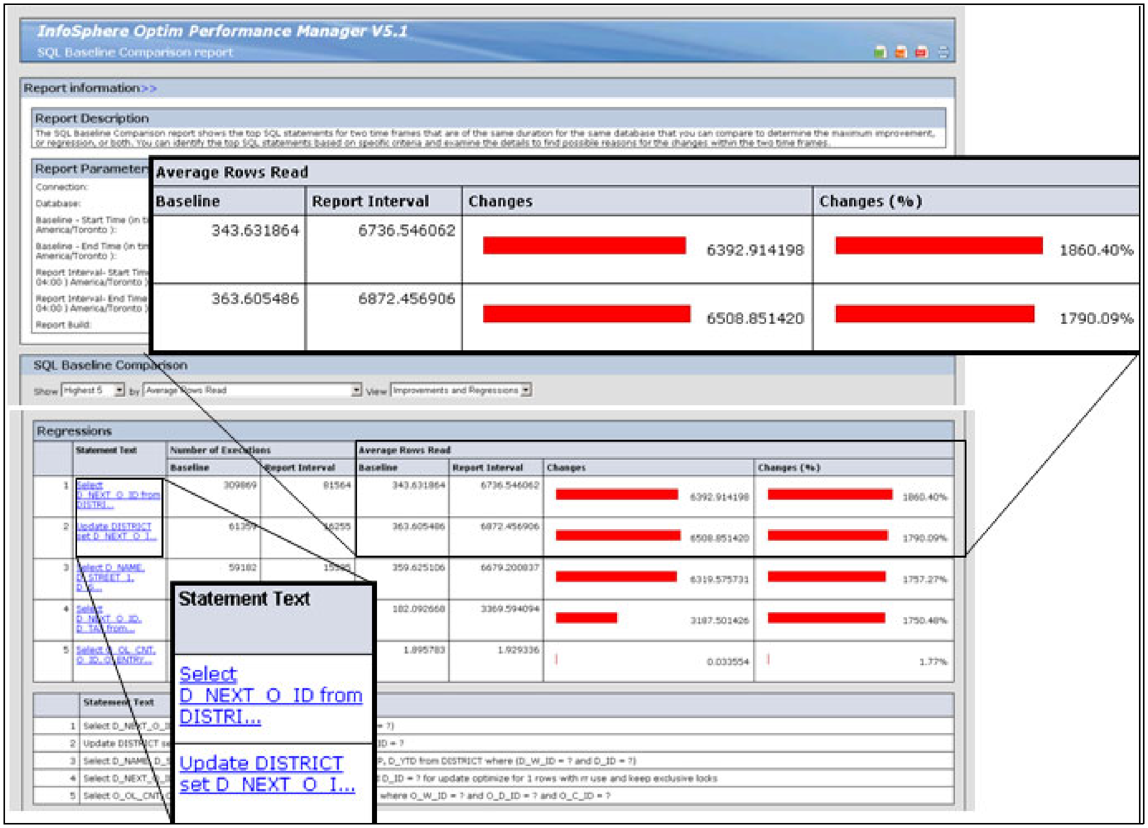

Figure 6-29 SQL Baseline report, regressions

6.4.3 Resolving high CPU-consuming statements

From the SQL Comparison report, return to the SQL Statements Dashboard, which is where you can launch InfoSphere Optim Query Workload Tuner (InfoSphere Optim Query Workload Tuner) to learn more about the problems statements. InfoSphere Optim Performance Manager V4 had the ability to do this also on a statement-by-statement basis, but InfoSphere Optim Performance Manager and InfoSphere Optim Query Workload Tuner have now been extended to be able to tune multiple statements as a group. This avoids making tuning decisions for one statement that might harm other ones.

As before, you can list the top statements by rows read, but also exploit another feature in InfoSphere Optim Performance Manager V5: the Filter button, which you can use to restrict the statements to only those that refer to the DISTRICT table. This gives a tight group of four statements to pass to InfoSphere Optim Query Workload Tuner with the Tune All button (Figure 6-30 on page 295).

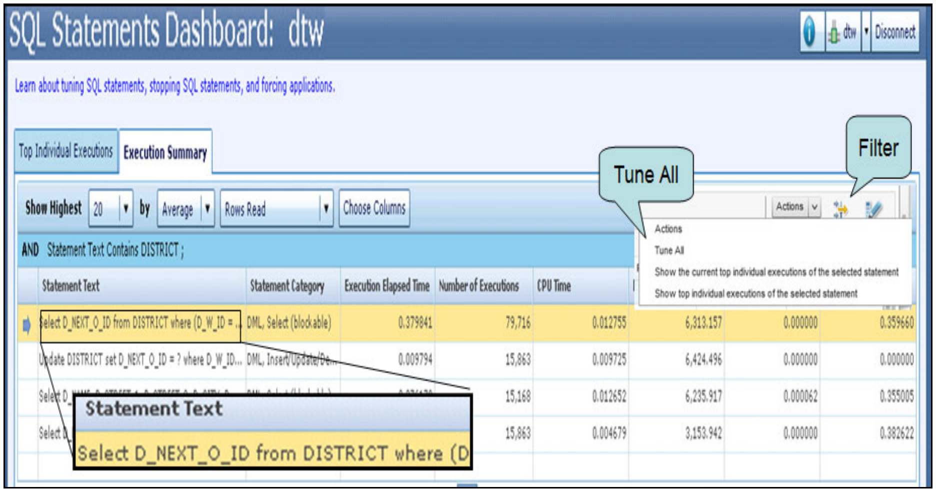

Figure 6-30 SQL Statements dashboard, Tune All

After you click **Tune** or **Tune All**, InfoSphere Optim Performance Manager activates InfoSphere Optim Query Workload Tuner and fills it with the group of statements selected.

In InfoSphere Optim Query Workload Tuner, you can select **Save All to Workload**, to save this group of statements away for later use. This lets you return to the group later, after looking at the issues of a single statement.

The analysis of the workload and the evaluation of the recommendations are described in detail in Chapter 5, “Getting to know InfoSphere Optim Query Workload Tuner” on page 221. By analyzing all the slow-running SQL statements that target a specific table, meaningful suggestions can be made.

6.5 Checking configuration changes

The database manager configuration and database configuration settings influence workload performance. You can use the Database Manager Configuration report and the Database Configuration report to check how your database manager and databases are configured and which parameters have changed over time.

In *IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows*, SG24-7925, the use of these reports to analyze the effects of changes in the database and database configuration manager parameters is explained (in the section about checking configuration changes using reports). This method is still true for InfoSphere Optim Performance Manager V5.2. Next is a description of the integration of InfoSphere Optim Performance Manager with InfoSphere Optim Configuration Manager.

6.5.1 Launching the InfoSphere Optim Configuration Manager

In the Execution Summary tab of the SQL Statements Dashboard, one SQL statement is using a lot of CPU time, which means a configuration change might have occurred. You can launch the InfoSphere Optim Configuration Manager by clicking **View Configuration Changes** (Figure 6-31).

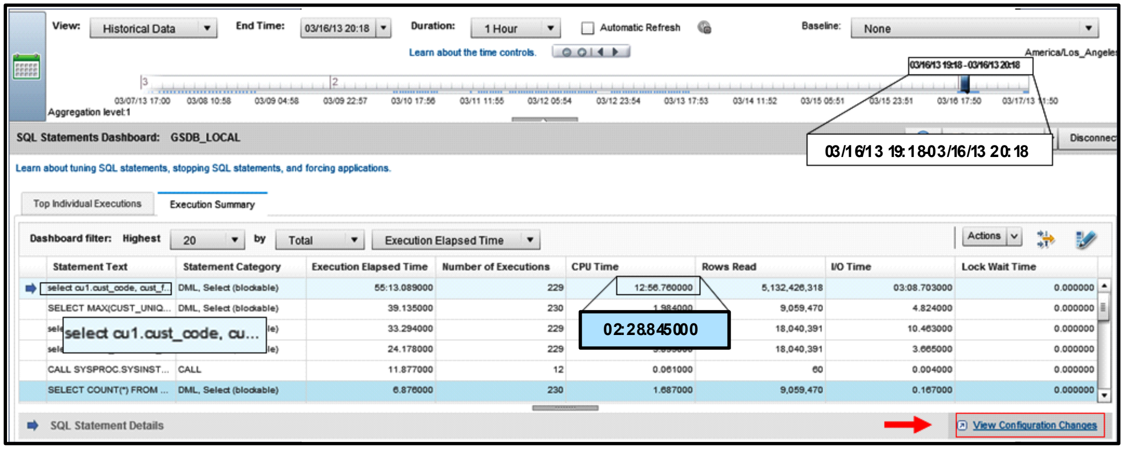

Figure 6-31 SQL Statements, View Configuration Changes

InfoSphere Optim Configuration Manager will be launched from another browser. Note the time range when the problem occurred. At the time of this writing, you must set your time frame in InfoSphere Optim Configuration Manager to include the time frame in InfoSphere Optim Performance Manager, and to see the changes that have occurred during that time.

In InfoSphere Optim Configuration Manager, set the duration to the `Last 30 days`, and select the **Connection**. InfoSphere Optim Configuration Manager gives a summary of all changes in the last 30 days for all the objects specified (Figure 6-32 on page 297).

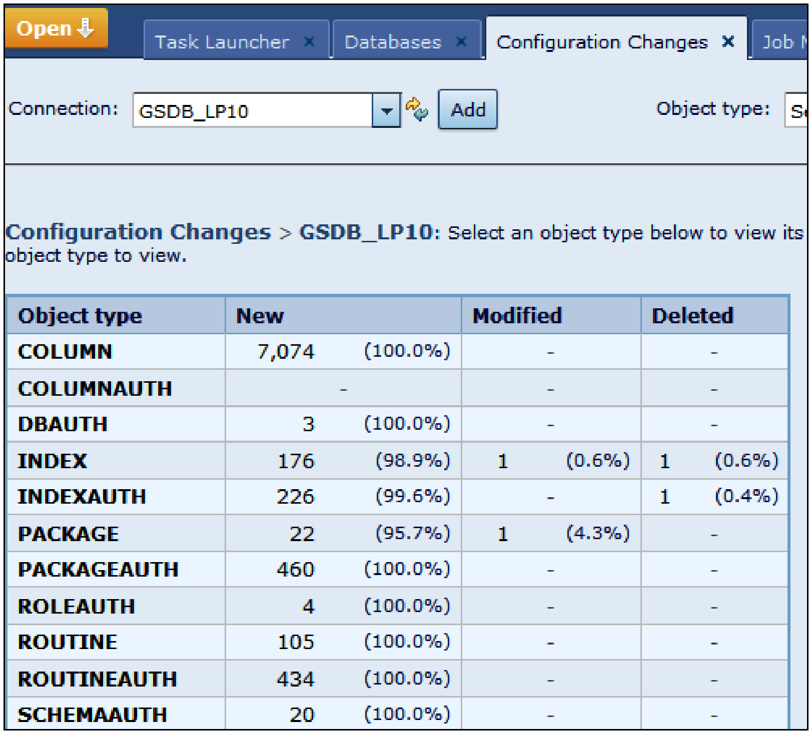

Object type	New		Modified		Deleted	
COLUMN	7,074	(100.0%)	-		-	
COLUMNAUTH	-		-		-	
DBAUTH	3	(100.0%)	-		-	
INDEX	176	(98.9%)	1	(0.6%)	1	(0.6%)
INDEXAUTH	226	(99.6%)	-		1	(0.4%)
PACKAGE	22	(95.7%)	1	(4.3%)	-	
PACKAGEAUTH	460	(100.0%)	-		-	
ROLEAUTH	4	(100.0%)	-		-	
ROUTINE	105	(100.0%)	-		-	
ROUTINEAUTH	434	(100.0%)	-		-	
SCHEMAAUTH	20	(100.0%)	-		-	

Figure 6-32 Configuration Changes, all objects

Because there were some changes in the INDEX objects, double-click this row to see more details about the index changes. You can see in Figure 6-33 on page 298 that an index on the table accessed by the high-CPU SQL statement was recently dropped.

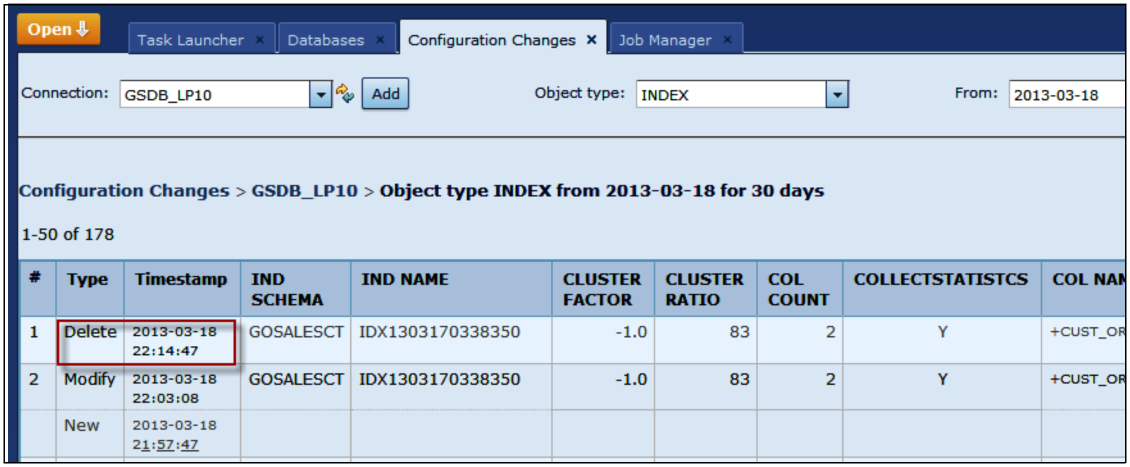

Figure 6-33 Configuration Changes, Index dropped

Part 2, "Using IBM products to manage performance" on page 143 examines more capabilities of InfoSphere Optim Configuration Manger.

Part 3

Performance management in a DB2 Distributed Environment

InfoSphere Optim Performance Manager monitors DB2 Distributed Partitioning Feature (DPF) and pureScale systems without special configuration requirements. This part has examples of the types of information that can help when developing a performance analysis on these types of systems.

This part includes the following chapter:

- Chapter 7: This chapter focuses on performance, particularly in DB2 distributed environments. The two primary distributed environments are centered around the DB2 Distributed Partitioning Facility and DB2 pureScale. The chapter describes functions and features of each and how they impact performance of each environment, and how that impact can be measured and monitored, thus enabling you to tune the environments for maximum performance.

7

Performance management: Distributed DB2 environments

InfoSphere Optim Performance Manager monitors DB2 Distributed Partitioning Feature (DPF) and pureScale systems without special configuration requirements. This chapter has examples of the types of information that can help when you develop a performance analysis on these types of systems.

7.1 Performance management in DB2 DPF environments

This section examines several special features that InfoSphere Optim Performance Manager provides for monitoring databases using the DB2 Database Partitioning Feature (DPF). All the same dashboards and reports are available as are for non-partitioned databases, but additional views are available, specifically for a DPF system.

First, tips and suggestions are offered for data collection and data viewing in InfoSphere Optim Performance Manager.

Then several scenarios show the effective use of InfoSphere Optim Performance Manager in a DPF environment.

The following scenarios are included in this section:

- Detecting activity skews
- Identifying skew at a database partition level
- Identifying most active objects

7.1.1 Data collection from DPF and viewing in web console

Chapter 2, "Planning your InfoSphere Optim Performance Manager deployment" on page 31 describes how to inform InfoSphere Optim Performance Manager what performance data to collect, how often to collect it, and how long to keep it. When monitoring a DPF system, the partitions to collect performance data on can also be specified.

When a very large database spans multiple physical machines, you may opt to collect only a sampling of data (one partition from each machine, for example) rather than gathering data from every partition. Using a sampling approach reduces overhead on the monitored database and reduces the amount of data InfoSphere Optim Performance Manager must store. However, using this approach therefore also means no performance data will be available for those partitions not monitored. That is a trade-off to be made in your own environment.

When working with performance data from a DPF system, you might want to view the performance for the whole system, or narrow some analysis to a specific partition or some subset of partitions. InfoSphere Optim Performance Manager allows for this in several ways, such as the following examples:

- Use member roles to enable easy filtering on the inflight dashboards.
- Define which partitions you want to monitor by using a member set.

Member roles and member sets

As you move through the InfoSphere Optim Performance Manager Configure Monitoring dialog for a DPF database, on step 4 you can choose to use roles and member sets.

A *member role* is a description of the type of work that partition does. The role choices are as follows:

- Catalog partition
- Data partition
- ETL partition
- Coordinator partition

The Members tab of step 4 of the monitoring configuration dialog is used to set the roles. By default, no roles are assigned to any partition. Figure 7-1 shows choosing the catalog partition role for partition 0. The other four partitions are set to the data partition role. Using roles is not required, but can be helpful later when analyzing data on the inflight dashboards.

The system used in this publication, has five partitions, across two physical machines. Machine1 houses only the catalog node and Machine2 has four logical partitions where the data tables are stored. So only two roles are defined, as shown in Figure 7-1. To choose the role, click in the role cell on the grid, which opens a drop-down list.

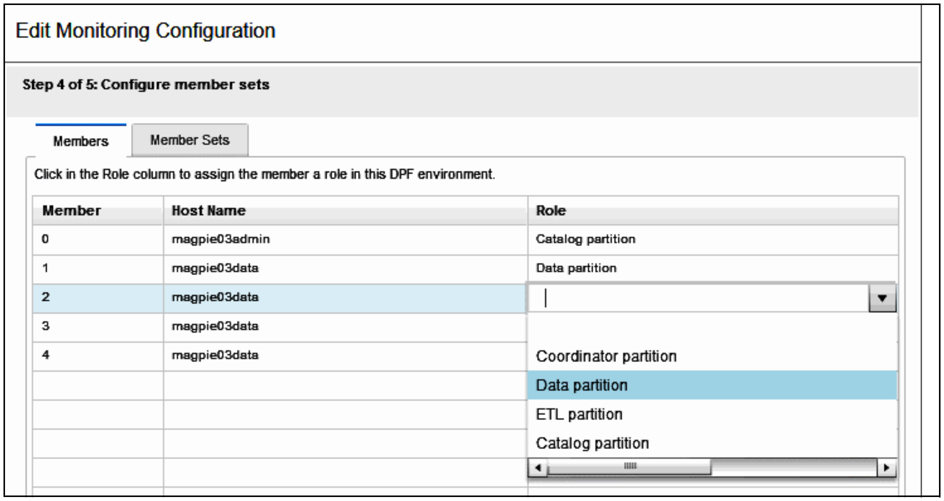

Figure 7-1 Members tab of Step4 on Configure Monitoring: Setting member roles

Another use of a named partition role is in alerts. For example, if there is much ETL activity, the metric tolerances might differ for ETL functions compared to day-to-day work on the other data partitions. You can identify one or more ETL partitions as a role and then set different alert thresholds for that role. For more information about this topic, see Chapter 4, “Getting to know InfoSphere Optim Performance Manager” on page 145.

A *member set* is a named group of members or partitions. InfoSphere Optim Performance Manager has several common groupings already defined, such as `All partitions` or `One for each machine`, as shown in Figure 7-2 on page 305.

One and only one member set can be selected as the active group that InfoSphere Optim Performance Manager will monitor. By default, InfoSphere Optim Performance Manager will monitor all partitions. If you want to monitor only one partition per machine, for example, select **One for each machine** Member Set on the Member Set tab, in step 4 of the monitoring configuration dialog. In the examples for this book, the All Partitions member set is being used.

If member roles have been defined on the Members tab in step 4 of the monitoring configuration dialog, those members will be shown on the Member Sets tab. In this example, the catalog and data partition roles have been set.

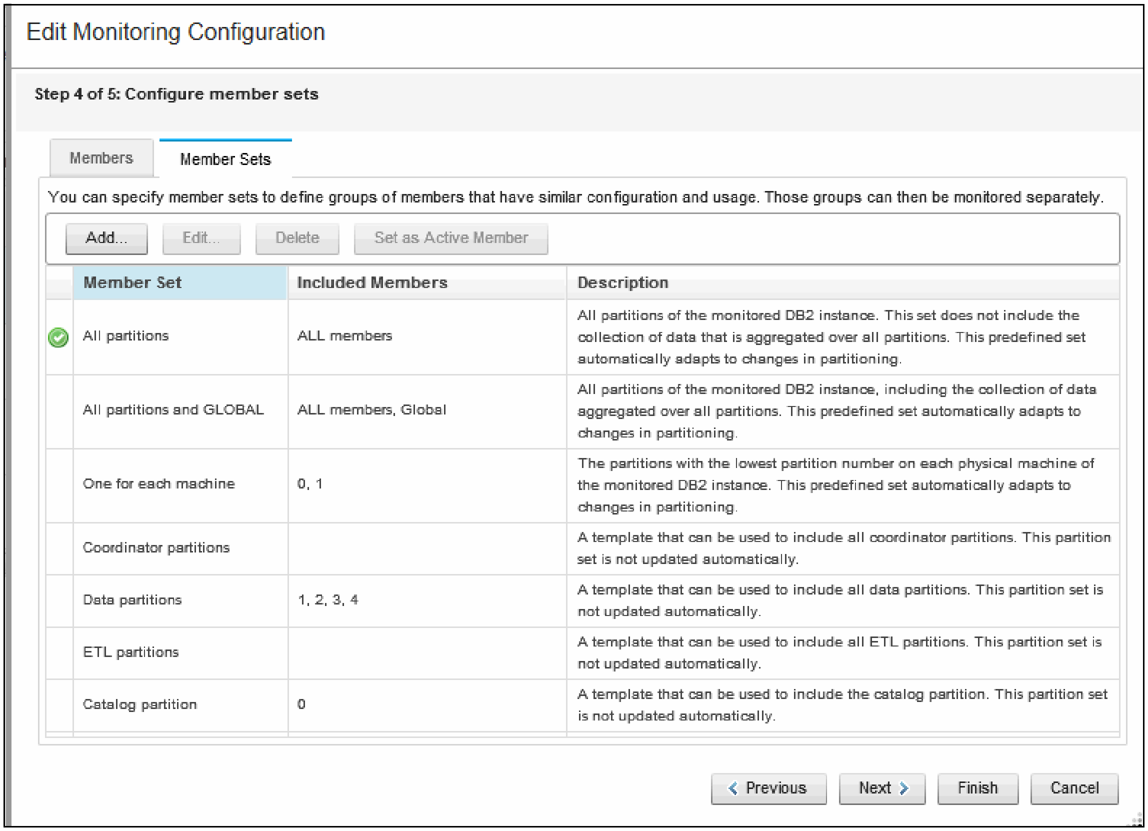

Figure 7-2 Edit Monitoring Configuration, Member Sets

The active member set controls what partition's data is collected by the InfoSphere Optim Performance Manager server.

For example if you only cared about monitoring your particular data partitions, you could select that as the active member, and InfoSphere Optim Performance Manager would only collect monitor data from those partitions.

You can create custom member sets if you want, by clicking **Add** (button not shown here) to open a dialog.

The active member set can be changed later, if you change your mind about which partitions you want to collect from. InfoSphere Optim Performance Manager will make the change after the configuration is saved. Previously collected data is not affected, however you must consider the period as you look at dashboard or reports, when more or less data is collected after the change.

Note: In previous releases of InfoSphere Optim Performance Manager, the member set was called a *partition set*. The name was changed to allow consistency in naming with pureScale systems. In general in InfoSphere Optim Performance Manager, partitions are now referred to as members, although you might continue to see references to partitions, or even possibly *nodes*, an even older term. Be assured; these terms can be used interchangeably for DPF monitoring.

Tip: A *role* is mainly used on the inflight dashboards, while the *member set* controls the performance monitoring data collection.

A good reference for optimizing your InfoSphere Optim Performance Manager data collection choices is the article on developerWorks, "Managing data warehouse performance with IBM InfoSphere Optim Performance Manager" at the following address:

http://ibm.co/1aFLIcw

7.1.2 Detecting activity skews

InfoSphere Optim Performance Manager provides a set of predefined ("canned") reports, which are interactive when viewed in the browser, or can be run offline and saved to various formats. See "Using the web material" on page 384 for more information.

The Database Connection report can be useful in a DPF environment to quickly visualize whether you have traffic arriving at unintended nodes, which might signal an activity skew.

Open the Report Dashboard, and select the Database Connection report. In this example, the same four-hour window that was being used on the regular dashboards has been chosen.

In Figure 7-3, the top 20 connections, sorted by Rows Read, are shown.

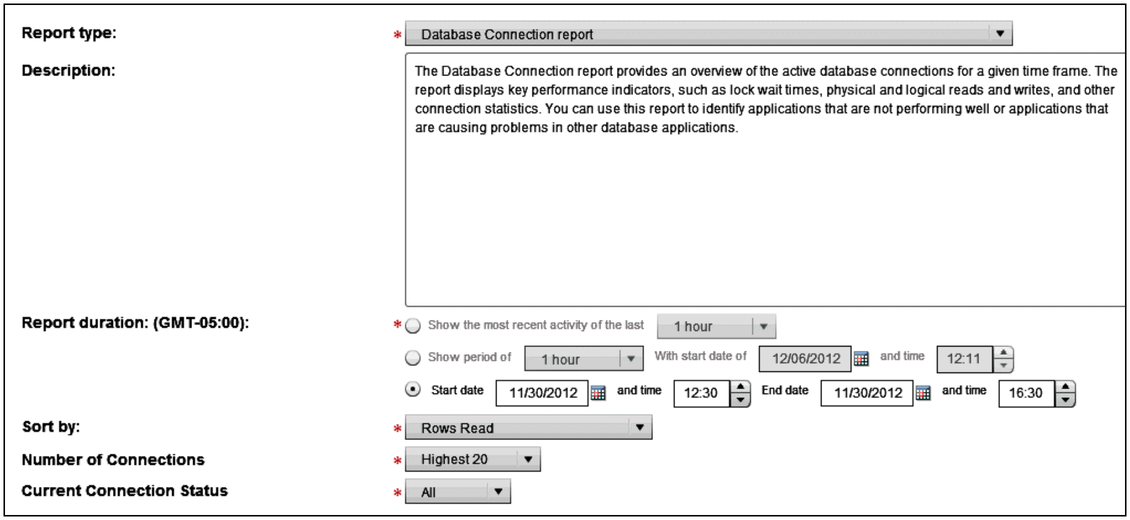

Figure 7-3 Predefined Report - Database Connection report selection criteria

Scroll to and click **Generate Report** to initiate the report. InfoSphere Optim Performance Manager will open a new browser tab for the report. The reports are interactive, maintaining the ability to drill down to more detail, as in the inflight dashboards.

Part of the resulting report is shown in Figure 7-4 on page 308. However, no connections were coming in to the other partitions during the reporting period, so only data for the catalog partition can be seen. This might be the normal state, because in this environment, there is only one coordinator node and all the workload run was simple. In a large environment with multiple coordinator nodes, or ETL work, more variation in this report might be seen.

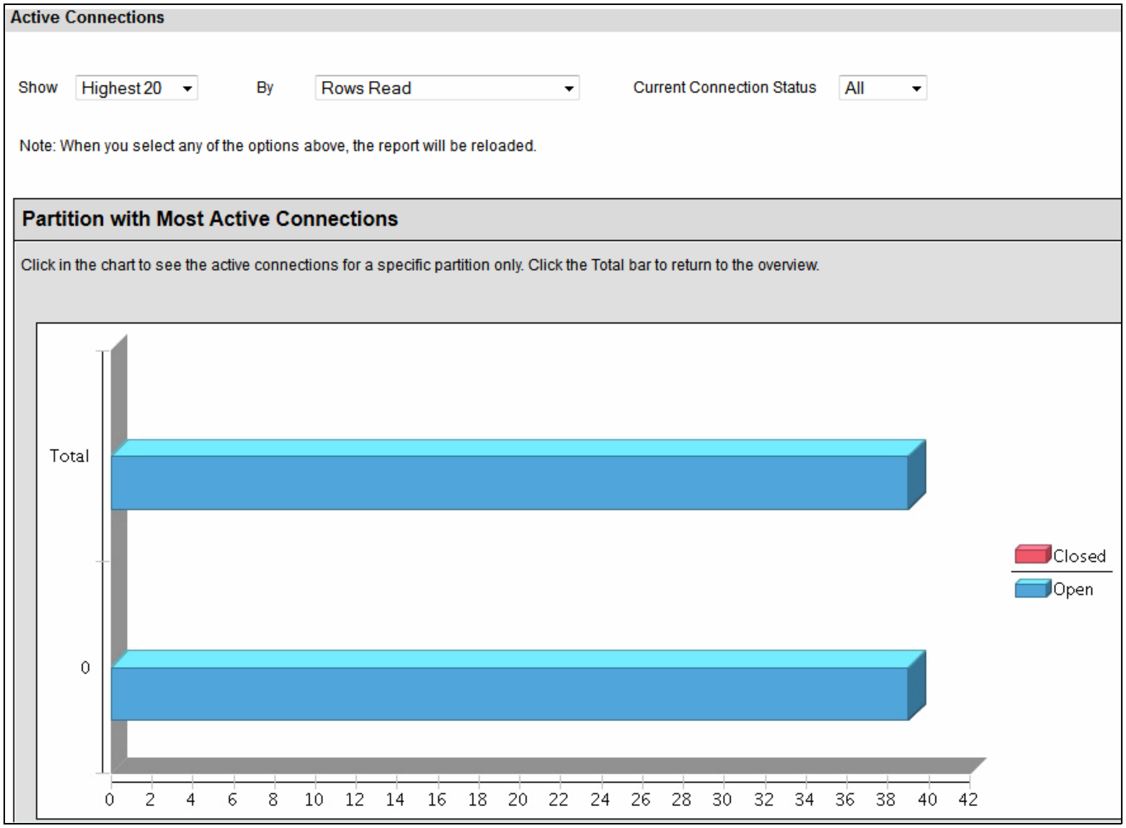

Figure 7-4 Database Connection report - DPF - no skews

Another example of Database Connection report in a DPF environment is in 4.4.11, "Connection Dashboard" on page 203.

7.1.3 Identifying skew at database partition level

On the Performance Overview Dashboard, when monitoring a DPF system, there is an additional chart named Member Skew, as shown in Figure 7-5 on page 309. The chart displays the minimum, average, and maximum count for rows read per member.

This example shows the performance data across *all* monitored members. The chart shows a large variation between minimum and maximum, but what does this mean?

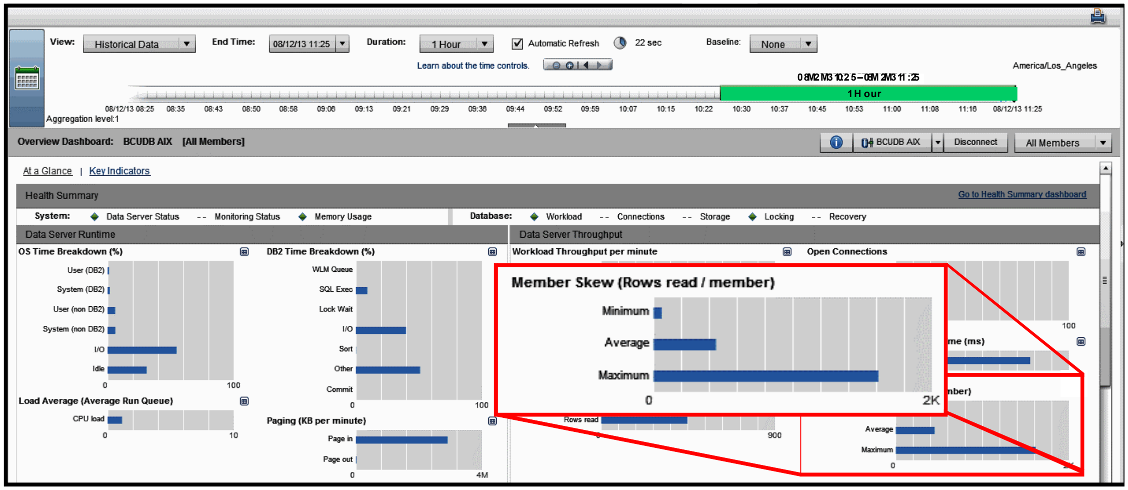

Figure 7-5 Performance Overview, DPF, Member skew

Next, look at the same metric, but for a different grouping of members. Remember certain partitions were assigned to certain roles, and in this case there is one catalog partition and four data partitions. The drop-down box near top right was used to change the view to the Data Partitions role (Figure 7-6).

Now the story is much different: Not much variation seems to exist, so the conclusion is that there is not much skew in the data partitions.

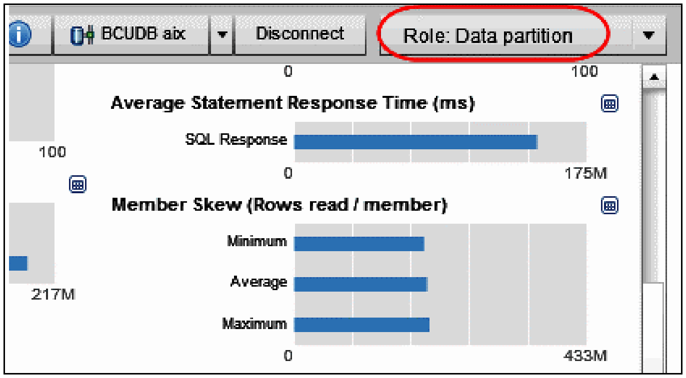

Figure 7-6 Performance Overview, DPF, Member skew for data partitions only

Why did the other view show such difference? The reason is, that in the all-member view, the catalog partition is included, and because there is not any data there, the minimum rows-read value is small compared to the data partitions.

Note: If a single partition is selected, either by using a role that only has one member, or if a specific partition is selected, the Member Skew chart is not displayed. The metric is a per-member calculation, so if there is only one member, there is nothing interesting to show because there is nothing to compare it to.

The Rows-Read metric can be seen in many places, and in several variations. Still on the Performance Overview page, the rows read per minute is shown to the left of the member skew. In Figure 7-7 on page 311, see the huge difference in scale between the rows read on catalog partition compared to the data partitions. The catalog partition has somewhere around 200,000 reads per minute, while the data partitions are more in the 200 million range.

Note: A baseline has not been set in this example, but if there were a baseline, you can visualize the deviation from norm, if some skew is creeping in. As an example, see the SG24-8111-AM.zip file listed in section "Using the web material" on page 384; Chapter 1 of the PDF in that ZIP file contains the "Disk Space Consumption report."

Most likely you will want to view only the Member Skew chart while filtered on only the data partitions.

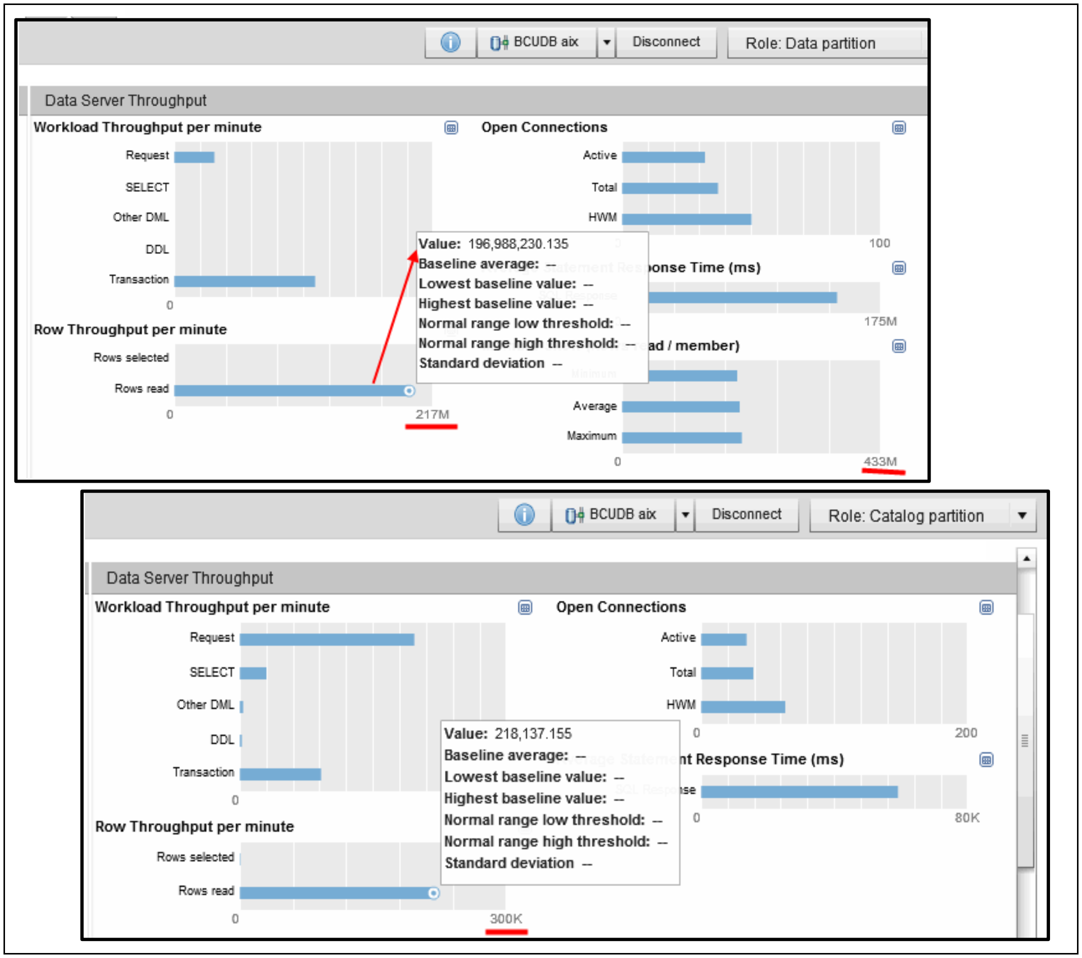

Figure 7-7 Performance Overview: Rows read throughput comparison

Workload alert

While looking at the Member Skew, notice that the Workload alert symbol is red (highlighted), so must be investigated. Switching to the Workload dashboard, as shown in Figure 7-8 on page 312 see the alert is on the Rows Read per Fetched Row metric. The current view is still the All Members group, so this is data for all monitored partitions. The chart displays the average Rows Read per Fetched Row, for the time window specified in the time slider. And because this is the All Members view, it is also averaged across all the partitions.

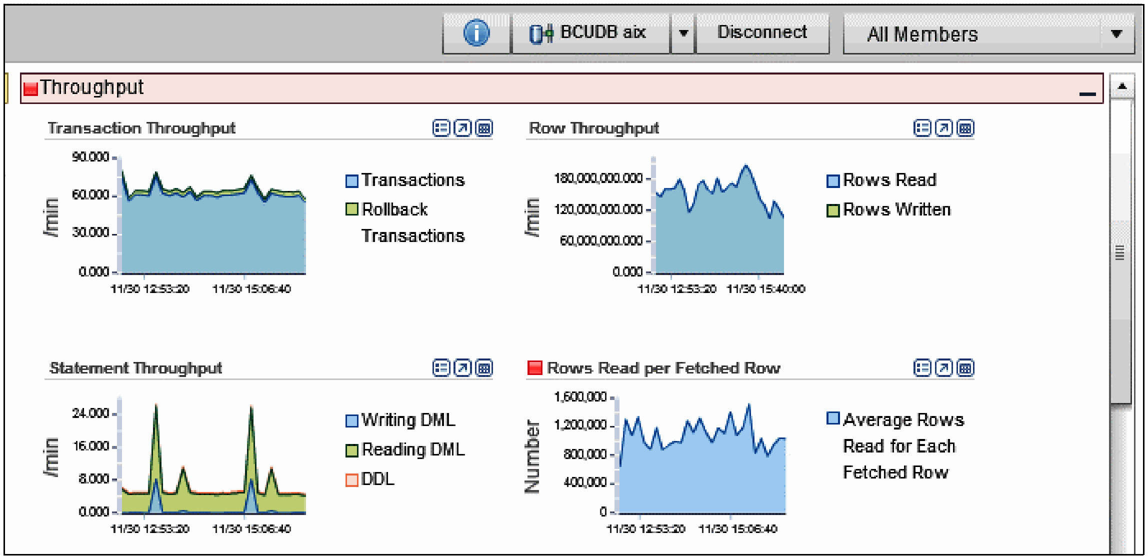

Figure 7-8 Workload dashboard, Rows Read per Fetched Row, all members

To quickly view the data for all monitored partitions, click the **Show data in grid** symbol in the upper right of the chart, and an overlay is then shown, as in Figure 7-9 on page 313. Click the **Member Details** tab to see per-partition specifics. Here, you can see that there are very large numbers for data partitions, and a small number for catalog partition. Because this is a data warehouse query-heavy system, these variations are not necessarily unusual.

Tip: Most (but not all) grid views of InfoSphere Optim Performance Manager graphs have a Member Details tab to show the metrics per monitored partition.

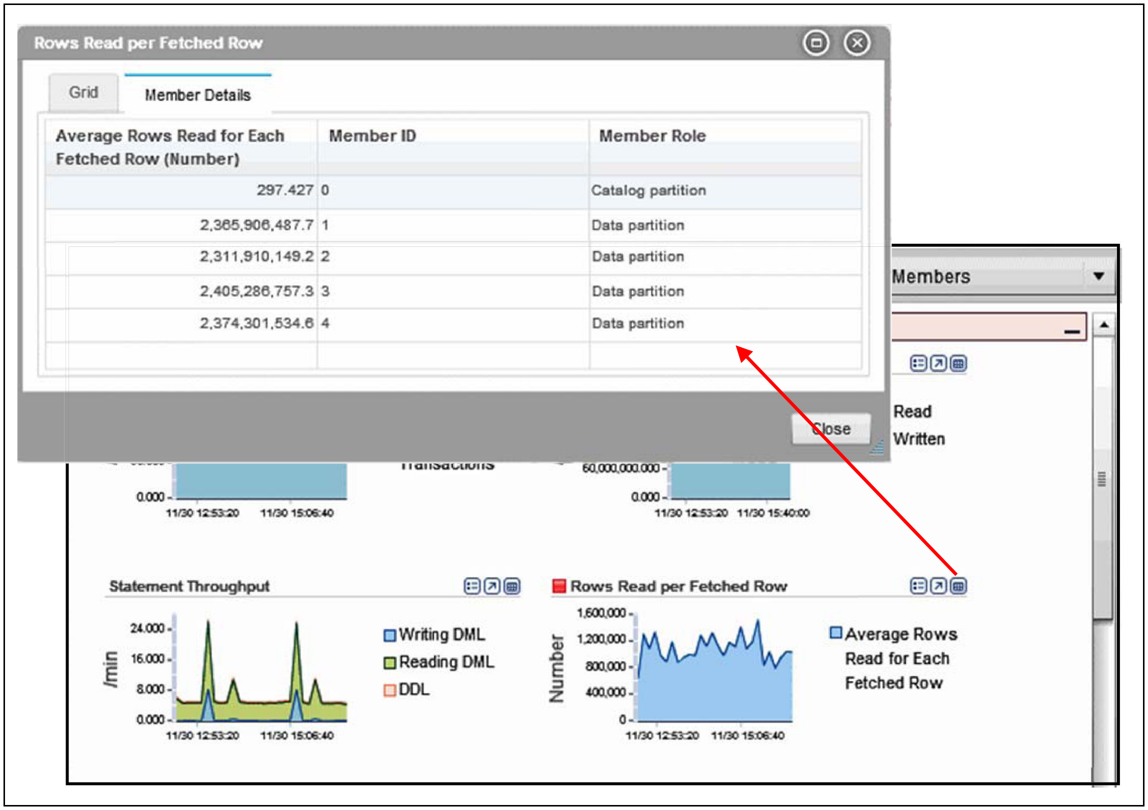

Figure 7-9 Workload dashboard, Rows Read per Fetched Row, per-member grid view

The alert definition may be set too low for this system, if reading 2 billion rows per fetched row is considered normal. (This could be a good case for some tuning exercise.)

In the configuration for Performance Alerts, you might want to set this Rows Read per Fetched Row alert to a higher value for the data partition role, to reduce the false positives, such as shown in Figure 7-10 on page 314. For this particular metric, you might want to disable the alert completely, if the information it provides is not useful. The point here is that the user can customize the alerts to best suit their own organization.

Although this book does not describe how to set the alert, Figure 7-10 shows the Alert page where a different alert threshold per role can be set.

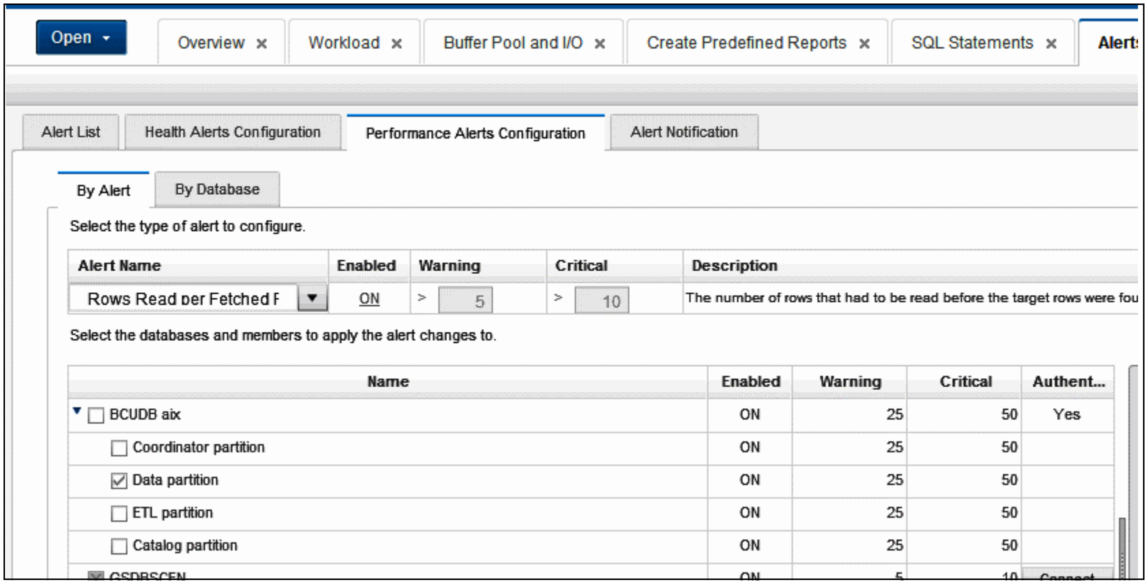

Figure 7-10 Performance Alert configuration, DPF by Role

7.1.4 Identifying most active objects

Database performance analysis is sometimes a science and sometimes an art. You might not always have a specific alert condition to handle, but might be looking only at general behavior, or exploring what performance data is available.

In this scenario, the database traffic is examined from the perspective of what is the busiest ("heavy hitter") table. You are not reacting to an alert or a complaining user. So, InfoSphere Optim Performance Manager can be used to explore the data and even uncover potential performance issues.

Identify heavy hitter table

Open the Buffer Pool and I/O Dashboard, then go to the Tables tab and view the highest 5 tables by rows read. Figure 7-11 on page 315 shows STORE_SALES as the busiest table. In this figure, the default view is being used, showing data for all monitored members.

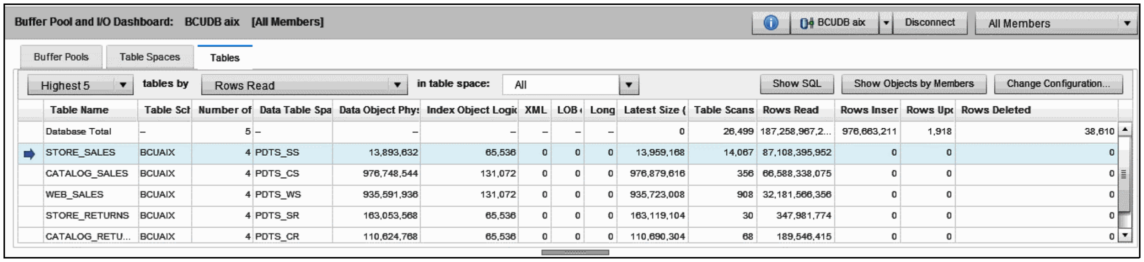

Figure 7-11 Buffer Pool and IO Dashboard, Tables tab, Highest 5 tables by Rows Read

InfoSphere Optim Performance Manager collects much performance data, and narrowing to a specific interest area can help. Therefore InfoSphere Optim Performance Manager provides rich filtering mechanisms in the web console.

Start by looking at what can be learned about the STORE_SALES table. Assume for the moment that you are not interested in the Buffer Pool or Table Space tabs, and you stay focused on the Tables tab.

As shown in Figure 7-11, the STORE_SALES has the most rows read, but it is not the largest table (Data Object Physical Size column). It also has more table scans than other tables, by a huge margin. (Yes the grid is sorted by rows read. But was also sorted by Table Scans, and STORE_SALES still are on top.)

Notice that none of the tables are doing any writing; inserts and updates are all zero. This is expected because it is a query application using the database.

View table by partition or member

Look now to see if anything is interesting about this table across the monitored partitions. Remember that in this example database, all the partitions are not being monitored. By using a member set to monitor only a subset of partitions, of course you will not see any data for the un-monitored partitions.

There are several ways to look at per-partition or per-member data:

- Use the **Show objects by members** button the Tables tab as shown in Figure 7-12 can be used.

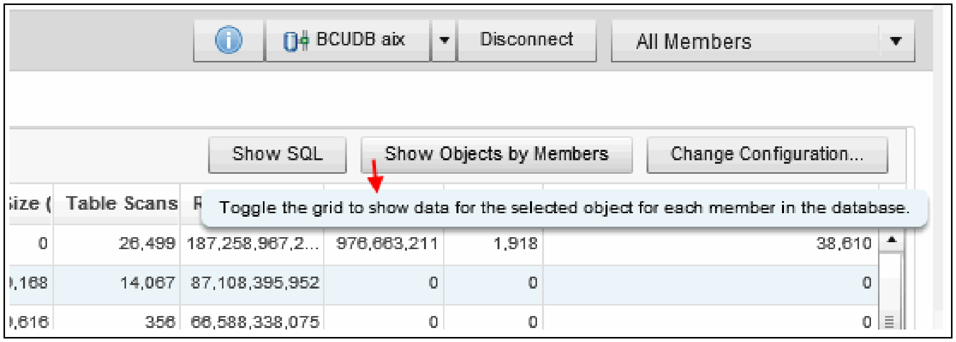

Figure 7-12 How to view tables by member

- Use the **Show by Member** button. The object grid will change to show the per-partition metrics in separate rows of the grid. In this case, the collected metrics for STORE_SALES table were of interest.

> **Tip:** When working with a grid view, click a column header to sort the list.

The resulting grid is shown in Figure 7-13 on page 317. What can be learned from this view of the data?

First notice the catalog partition is not mentioned. This is fine, because all the data for STORE_SALES is on the data partitions of course.

Next notice the index size is much smaller on member 1. Generally, the expectation is that data will be fairly even across the partitions; you try to avoid skews. This particular case might or might not be causing a problem, but you should note it for future though.

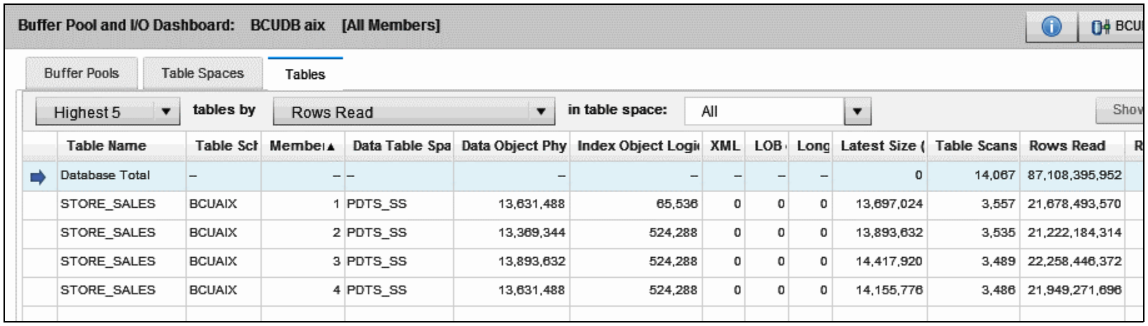

Table Name	Table Sch	Member	Data Table Spa	Data Object Phy	Index Object Logi	XML	LOB	Long	Latest Size (	Table Scans	Rows Read
Database Total	–	–	–	–	–	–	–	–	0	14,067	87,108,395,952
STORE_SALES	BCUAIX	1	PDTS_SS	13,631,488	65,536	0	0	0	13,697,024	3,557	21,678,493,570
STORE_SALES	BCUAIX	2	PDTS_SS	13,369,344	524,288	0	0	0	13,893,632	3,535	21,222,184,314
STORE_SALES	BCUAIX	3	PDTS_SS	13,893,632	524,288	0	0	0	14,417,920	3,489	22,258,446,372
STORE_SALES	BCUAIX	4	PDTS_SS	13,631,488	524,288	0	0	0	14,155,776	3,486	21,949,271,696

Figure 7-13 STORE_SALES table by member

The other numbers appear to be fairly close, so you should continue exploring.

The Show object by members view on the Buffer Pool and I/O Dashboard is key to helping identify skews for any of the physical objects: buffer pools, table spaces, and tables.

So far you can see only that STORE_SALES is the busiest table, and that it has a small index on member 1. But you do not know whether that is a problem. Examining the database objects is interesting, but you probably want to know what kind of work is being done against those objects. InfoSphere Optim Performance Manager provides a quick link from the Tables view to show the SQL that has been run against that table.

Show SQL hitting the busy table

So far you are still looking at only the STORE_SALES table. As Figure 7-14 shows, you can click **Show SQL**. InfoSphere Optim Performance Manager will launch the SQL Statements Dashboard, with a filter to show only SQLs that hit the STORE_SALES table.

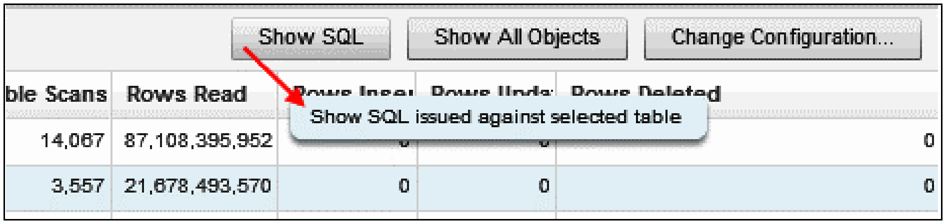

Figure 7-14 Tables view - Show SQL button

On the SQL Statements dashboard, as shown in Figure 7-15 on page 318, you now can see the top 20 statements issued against the STORE_SALES table, sorted by average elapsed time. The same time slider period is used across all inflight dashboards, so if you are looking at some historical data, you can be assured the time window will stay in context as you navigate throughout InfoSphere Optim Performance Manager.

In this example, the interest is in the slowest statement for the busiest table. But you can just as easily show the top N statements by some other metric, such as rows read, CPU, and lock wait time, by using the predefined filter buttons.

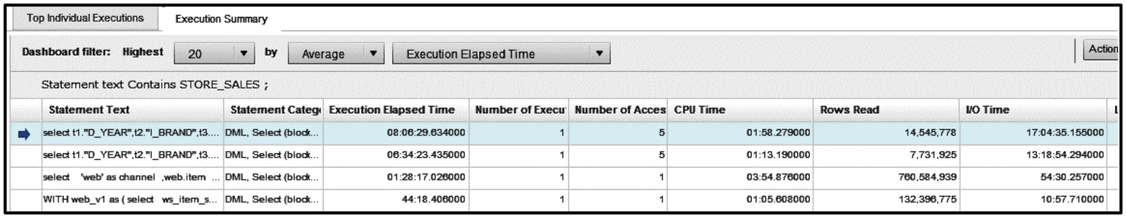

Figure 7-15 SQL Statements for specific table

On the Buffer Pool and I/O Dashboard, the STORE_SALES data was being viewed for each partition individually. But when launching to the SQL Statements dashboard that context is not maintained, so again you are looking at statements across all partitions. You can, however, view any single statement's per-partition data just as on the Buffer Pool and I/O dashboards.

On the SQL Statements Dashboard use the Actions drop-down menu, as shown in Figure 7-16, to change to a per-partition view. Select the first (slowest) statement in the list and click **Actions** → **Show Statement by Members**.

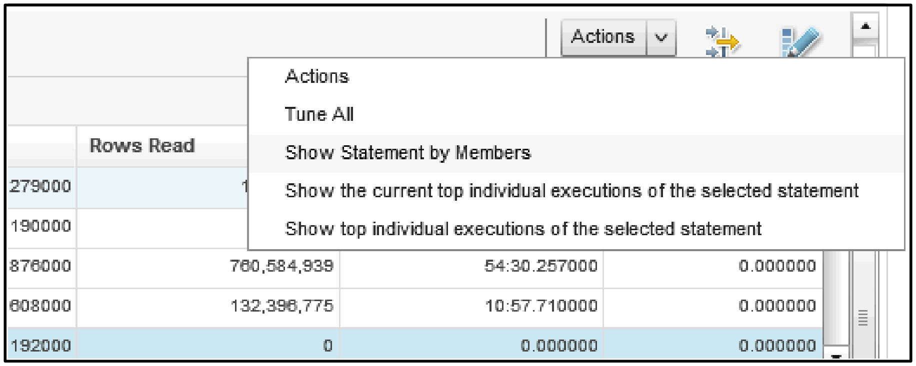

Figure 7-16 SQL Statements, Show Statement by Members

InfoSphere Optim Performance Manager will filter the grid to show only the selected statement, with a row for the performance data from each monitored partition (Figure 7-17).

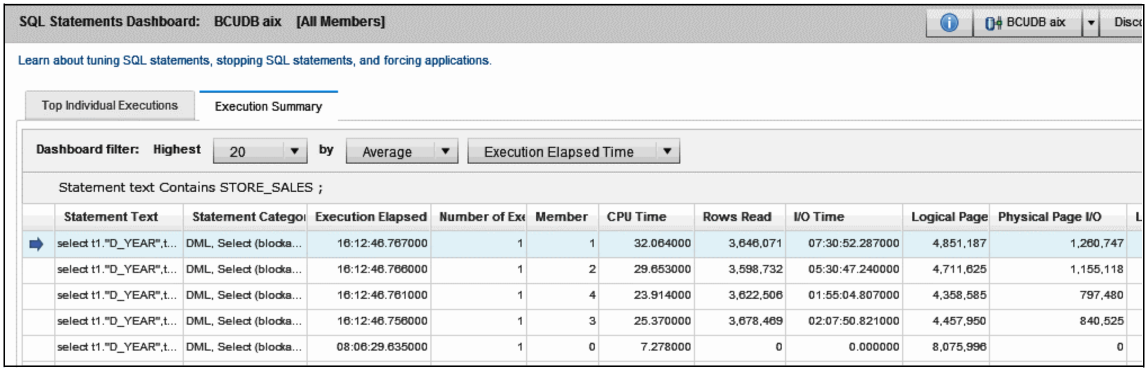

Figure 7-17   SQL Statements - single statement metrics per partition

From this view of the statement, you can look for any other skews, for example, during execution of the statement across partitions. Member 1 is showing more CPU time and more I/O time than other partitions. You already saw that member 1 had a much smaller index, so this can be a correlation.

This information alone might not be enough to reach any definite conclusions. The statement shown here might be worth tuning, and you can easily launch to the Optim Query Tuner at this point. The Optim Query Tuner is described in Chapter 5, “Getting to know InfoSphere Optim Query Workload Tuner” on page 221.

This query was run through Optim Query Tuner, which had recommendations, as shown in Figure 7-18. They were not applied because a full tuning exercise is beyond scope of this example.

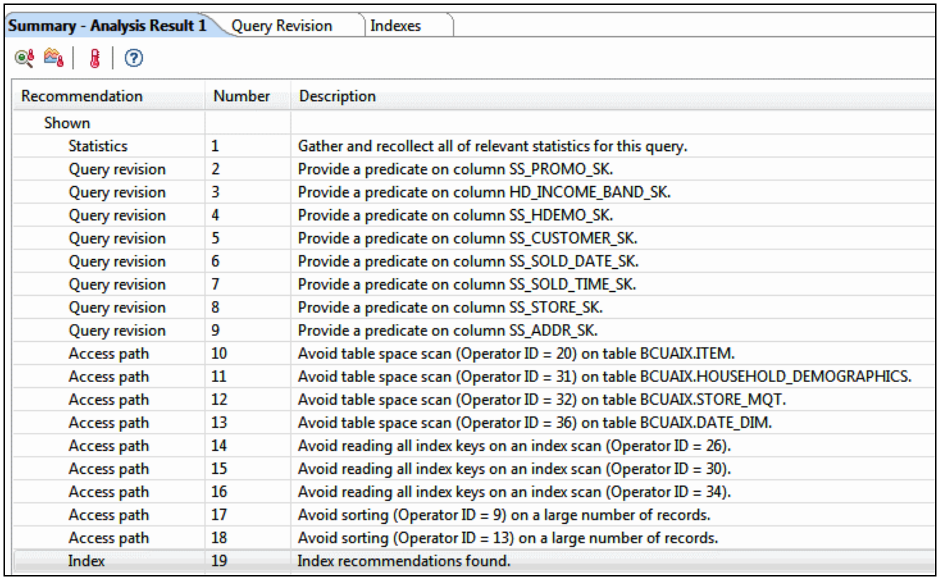

Summary - Analysis Result 1 | Query Revision | Indexes

Recommendation	Number	Description
Shown		
Statistics	1	Gather and recollect all of relevant statistics for this query.
Query revision	2	Provide a predicate on column SS_PROMO_SK.
Query revision	3	Provide a predicate on column HD_INCOME_BAND_SK.
Query revision	4	Provide a predicate on column SS_HDEMO_SK.
Query revision	5	Provide a predicate on column SS_CUSTOMER_SK.
Query revision	6	Provide a predicate on column SS_SOLD_DATE_SK.
Query revision	7	Provide a predicate on column SS_SOLD_TIME_SK.
Query revision	8	Provide a predicate on column SS_STORE_SK.
Query revision	9	Provide a predicate on column SS_ADDR_SK.
Access path	10	Avoid table space scan (Operator ID = 20) on table BCUAIX.ITEM.
Access path	11	Avoid table space scan (Operator ID = 31) on table BCUAIX.HOUSEHOLD_DEMOGRAPHICS.
Access path	12	Avoid table space scan (Operator ID = 32) on table BCUAIX.STORE_MQT.
Access path	13	Avoid table space scan (Operator ID = 36) on table BCUAIX.DATE_DIM.
Access path	14	Avoid reading all index keys on an index scan (Operator ID = 26).
Access path	15	Avoid reading all index keys on an index scan (Operator ID = 30).
Access path	16	Avoid reading all index keys on an index scan (Operator ID = 34).
Access path	17	Avoid sorting (Operator ID = 9) on a large number of records.
Access path	18	Avoid sorting (Operator ID = 13) on a large number of records.
Index	19	Index recommendations found.

Figure 7-18 Query Tuner results for slow STORE_SALES query

Using reports to examine skews

In the examples in this section, the web console is used to look at the data for STORE_SALES table. Good information can also be gotten from the interactive reports. General report usage is described in 4.6, “Reporting with InfoSphere Optim Performance Manager” on page 220.

Next a report is run for the Table Usage. Figure 7-19 shows the report selection criteria (the same 4-hour window, specifically for the top 5 tables by Rows Read.

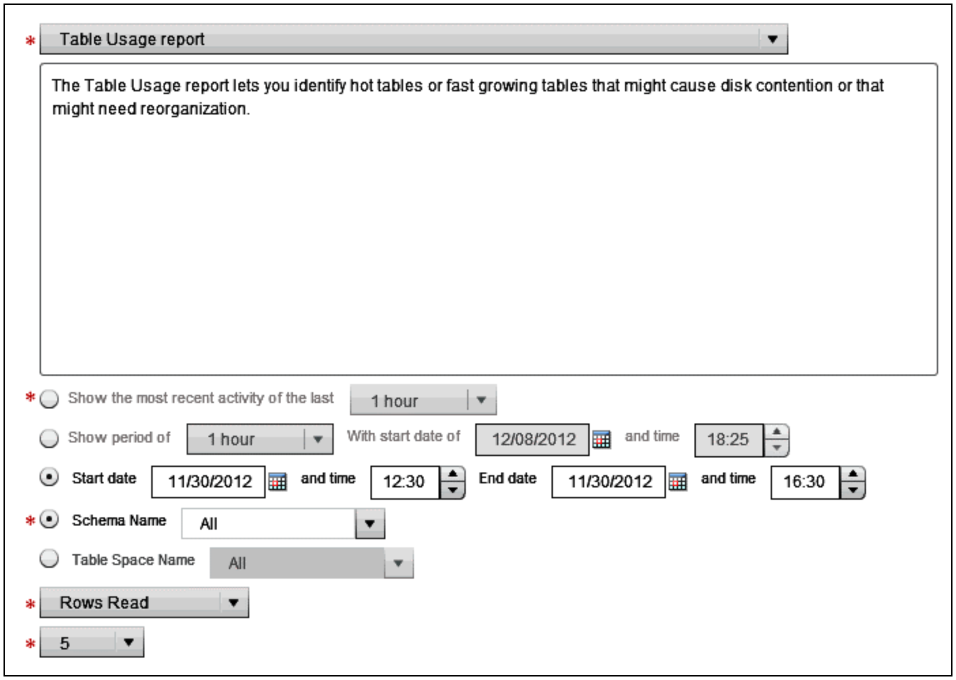

Figure 7-19 Table Usage report, selection criteria

When the report is ready, it is displayed in another browser tab. In this case, the STORE_SALES table is at the top, as expected, and shown in Figure 7-20. Notice the Number of Data Partitions column, which shows that the table is also partitioned itself. Remember this is not the same as database partitioning. Currently the data partition information (also known as *range partitions*) is not displayed in the web console inflight dashboards.

Report information>>

Report Configuration

Sort by: Rows Read   Number of Tables Top 5

Note: When you select any of the options above, the report will be reloaded.

Active Tables

Note: This information is available only when the database performance monitor is configured to collect table information. To change the configuration settings and go to Step 2: Configure monitoring profiles. Edit the I/O and disk space configuration.

Table Name	Schema Name	Type	Content Type	Number of Data Partitions	Number of Pages	▼ Rows Read
STORE_SALES	BCUAIX	USER_TABLE	Table with Index	14	1,439,208	87,108,395,952
CATALOG_SALES	BCUAIX	USER_TABLE	Table with Index	14	1,183,160	66,588,338,075
WEB_SALES	BCUAIX	USER_TABLE	Table with Index	14	569,464	32,181,566,356

Figure 7-20 Table Usage report, Summary view, STORE_SALES

When viewed in the browser, the reports are interactive, which means you can click objects to drill down to more detail. In this case the STORE_SALES table was clicked to see more about it. Another browser tab opens, displaying the new sub-report for STORE_SALES, as shown in Figure 7-21.

With this detail report, you can see how many data pages are used for each data partition across each member. If it shows skews, they can easily be found.

Table Partitions

Note: This information is available only when the database performance monitor is configured to collect table information. To change the configuration settings, go to the Manage Database Connections page, click Configure Monitoring and go to Step 2: Configure monitoring profiles. Edit the I/O and disk space configuration.

Table Name	Schema Name	Member	Data Partition ID	Type	Content Type	▼ Number of Pages	Rows Read	Rows Inserted	Rows Updated	Rows Deleted	Growth (in pages)	Growth Rate (%)
STORE_SALES	BCUAIX	3	0	USER_TABLE	Table with Index	140,752	9,376,557,362	0	0	0	0	0.00
STORE_SALES	BCUAIX	1	0	USER_TABLE	Table with Index	139,728	9,211,637,941	0	0	0	0	0.00
STORE_SALES	BCUAIX	4	0	USER_TABLE	Table with Index	138,336	9,210,132,960	0	0	0	0	0.00
STORE_SALES	BCUAIX	2	0	USER_TABLE	Table with Index	137,744	9,051,748,196	0	0	0	0	0.00
STORE_SALES	BCUAIX	3	1	USER_TABLE	Table with Index	70,368	5,157,196,088	0	0	0	0	0.00
STORE_SALES	BCUAIX	3	3	USER_TABLE	Table with Index	70,112	3,467,346,910	0	0	0	0	0.00
STORE_SALES	BCUAIX	1	1	USER_TABLE	Table with Index	70,048	4,617,188,681	0	0	0	0	0.00

Figure 7-21 Table Usage Report - STORE_SALES detail

7.2 Performance management for a DB2 pureScale environment

The pureScale feature of DB2 provides a scale-out, shared data clustered database environment. It is similar to DB2 DPF in some ways, in that it combines the power of multiple systems to enable greater overall capacity. In both DPF and pureScale, different component systems have different roles, such as coordinator partitions in DPF, and cluster caching facilities (CFs) in pureScale. However, DPF is a shared-nothing design where the database is hash-partitioned across the cluster, in pureScale there is a single database, equally accessible by all members. From a monitoring perspective, both the similarities and differences are listed in various InfoSphere Optim Performance Manager dashboards.

For pureScale monitored database, InfoSphere Optim Performance Manager presents data in a global or a per-member view, similar to the DPF examples in the previous section. Those examples are not repeated here.

When monitoring a pureScale system, additional pureScale-specific performance metrics are shown on many InfoSphere Optim Performance Manager dashboards and reports. The metrics would not appear when viewing a non-pureScale system.

The following types of information can be viewed:

- Cluster caching facility (CF) CPU and memory utilization
- Group buffer pool (GBP) hit ratio at database, connection, buffer pool, table space, and statement levels
- CF lock time-outs, lock escalations, and transaction lock wait time per database
- CF requests/time on connection or statement level
- Global Lock Manager information
- Page reclaim information
- CF configuration parameters in Database and Database Manager reports

For DB2 10, several additional metrics are available:

- Average cross invalidation (XI) time
- Number of XI requests

In addition to pureScale performance metrics visible on dashboards and reports, InfoSphere Optim Performance Manager includes ready-for-use alerts specific to pureScale, and provides predefined monitoring templates for pureScale systems.

Tip: For an excellent presentation, view the online video from January 2013, “Best practices for DB2 pureScale performance and monitoring,” which is listed at the following address:

https://ibm.biz/BdxMvG

When monitoring a DB2 pureScale instance, the Overview Dashboard provides a pureScale tab in the Performance Focus section, as shown in Figure 7-22.

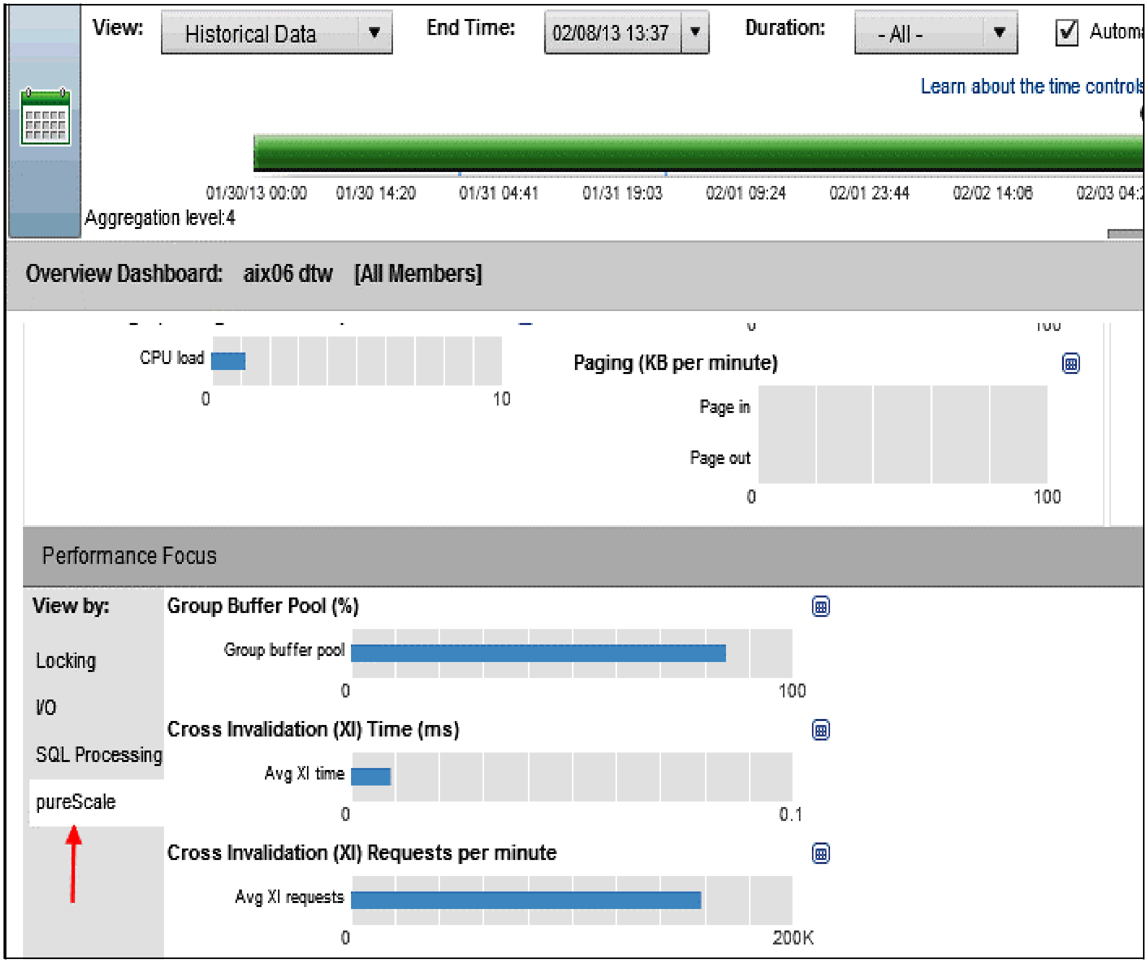

Figure 7-22 Overview Dashboard, pureScale tab

A key factor in the performance of a pureScale system is the response time of messages sent between members and the cluster caching facility (CF). If the interconnect gets congested, overall cluster performance will suffer. The best way to monitor this is with the Cross Invalidation (XI) Time metric in the pureScale tab; the XI times are shown starting with V10. This reports the amount of time required to send a certain kind of lightweight message from the CF to a member. In general, if the XI time is below approximately 15 ms (0.015 s, as reported in the InfoSphere Optim Performance Manager pureScale tab; Figure 7-23 shows 9 ms), congestion is not a problem. However if the XI response time climbs and stays above approximately 20 ms, congestion is a real possibility. In a case like this, cluster performance might benefit from additional interconnect capacity being added (which typically means adding an extra InfiniBand or Ethernet adapter to the CF).

The pureScale tab also shows Group Buffer Pool (GBP) hit ratio. If a drop in GBP hit ratio (as compared to the normal baseline value) is seen, use the Buffer Pool and I/O Dashboard to drill down through the problem.

A simple example of how the pureScale metrics are exposed is shown in Figure 7-23, where the Overview Dashboard highlights a deviation from norm for the Group Buffer Pool hit ratio.

In the case of this sample problem (where the GBP is too small), the GBP hit ratio metric shows us a huge drop from 93% to 27%. This is a place where the baseline is handy, because GBP hit ratio can be quite low even under normal circumstances. How do you know that 27% is not perfectly normal for this environment? The answer is that there is a baseline to compare against, and the baseline hit ratio is high.

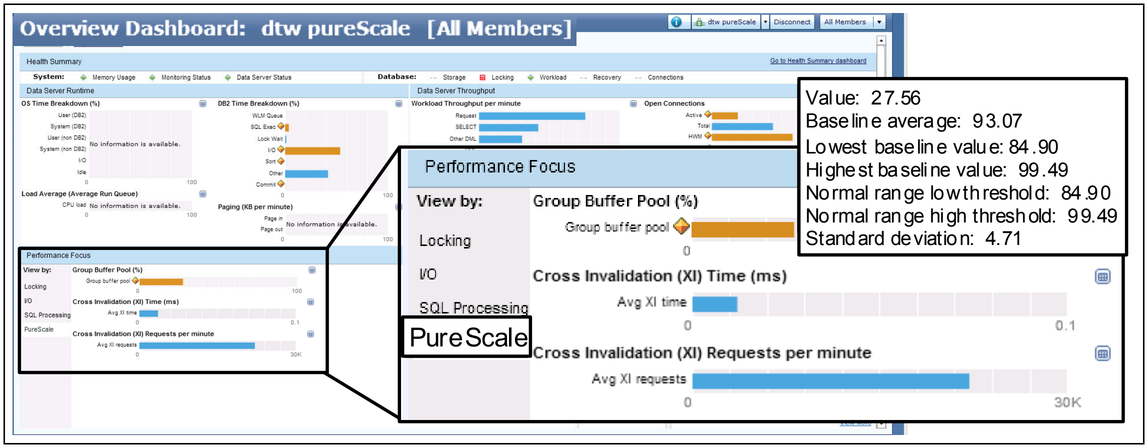

Figure 7-23 Overview Dashboard for pureScale system

A reduced hit ratio is not necessarily bad; the same is true (actually, even more true) for GBP compared to the regular DB2 buffer pool. In fact, the GBP hit ratio can actually be low in many systems. So, you must investigate further by drilling down into the Buffer Pool and I/O Dashboard, to get an idea of the physical I/O characteristics, as shown in Figure 7-24.

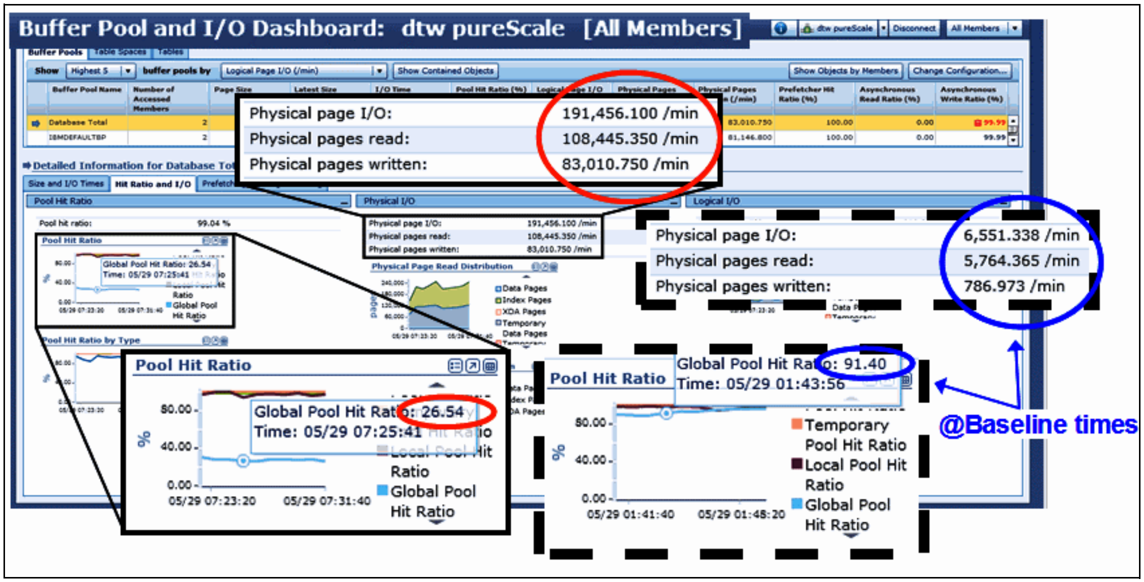

Figure 7-24 pureScale Buffer Pool and I/O dashboard

To start, verify the GBP hit ratio on the charts that show hit ratio versus time. As shown on the lower left, the hit ratio is not only spiked down, it is fairly uniformly low at 27%. Likewise, this dashboard can also show the rate of physical disk I/O, divided into pages read and written per minute.

Not only is there a large decrease in GBP hit ratio, it is correlated with a big increase in I/Os per minute; the combination can give you a fairly clear indication of a problem.

As before, comparing with values at the baseline (shown here in the boxes with the dashed outline) indicates that the hit ratio was once fine. In addition, the physical I/O at the baseline time was also significantly lower: approximately 6,500 I/Os per minute in the baseline, compared to 191,000 I/Os per minute now.

The combination of a decrease in GBP hit ratio and an increase in costly disk I/Os as compared to the baseline suggests that something near the root of the problem has been found. The ultimate cause might be poor SQL or an increase in user data. The easiest item to check (and should be checked first) is for a decrease in the size of the group buffer pool. In this example, this decrease is what produced the symptoms being seen.

Part 4

Specific topics for DB2 performance management

This part focuses on specific topics for DB2 performance management.

The following information is included in this part:

- Chapter 8: This chapter describes implementing workload management with InfoSphere Optim Performance Manager. There is an IBM DB2 for Linux, UNIX, and Windows published best practices (WLM Best Practices) document for implementing DB2 Workload Management in a Data Warehouse environment. See "Implementing DB2 workload management in a data warehouse" in the list of "Online resources" on page 386. That document provides detailed step-by-step guidance on how to help ensure the stability and predictability of your database system. By using the Workload Manager configuration tool and custom report facility provided by InfoSphere Optim Performance Manager, the steps to implement and manage a workload management configuration within DB2 are greatly simplified and the effort reduced.
- Chapter 9: InfoSphere Optim Performance Manager Extended Edition extends the capability of InfoSphere Optim Performance Manager with end-to-end database monitoring for Java technology and DB2 call level interface (CLI) applications (apps), with ready-for-use configurations for SAP, IBM WebSphere, IBM Cognos, IBM InfoSphere DataStage, and InfoSphere SQL Warehouse applications. The Extended Edition also includes integration with IBM Tivoli monitoring solutions to provide deep database insights into existing IBM Tivoli Composite Application Manager (ITCAM) application-monitoring environments.

  The chapter has an overview of these capabilities and identifies several articles and videos you can review. These articles are often updated and therefore the information might be more current than what is available in this Redbooks publication,

8

Implementing workload management

This chapter describes implementing workload management with InfoSphere Optim Performance Manager. There is an IBM DB2 for Linux, UNIX, and Windows published best practices document for implementing DB2 workload management in a data warehouse environment. See “Implementing DB2 workload management in a data warehouse” at the following location:

http://www.ibm.com/developerworks/data/bestpractices/workloadmanagement/

It has detailed guidance for how to help ensure the stability and predictability of your database system. By using the Workload Manager (WLM) configuration tool and custom report facility provided by InfoSphere Optim Performance Manager, the steps to implement and manage a workload management configuration within DB2 are greatly simplified and the effort is reduced.

WLM is a graphical user interface that can be used to configure and monitor DB2 Workload Manager V9.5 or later. It is automatically installed with InfoSphere Optim Performance Manager. The WLM is designed to help you implement and follow the published best practices to configure a DB2 workload management configuration. The WLM can automatically create a starting configuration following the best practices methodology. It can also help you customize this initial configuration for your particular environment and workload and help your more easily use additional features of DB2 Workload Manager.

The best practices presents a set of definitions representing the stages of maturity for a workload management configuration in a DB2 for Linux, UNIX, and Windows database. These stages range from stage 0 through advanced stage 3 configuration. See Figure 8-1.

Stage	Name	Objective
0	Default workload management configuration	
1	Untuned workload management configuration	Learning about your system
2	Tuned workload management configuration	Stabilizing your system
3	Advanced workload management configuration	Dealing with unique aspects

Figure 8-1 Stages of DB2 workload management configuration

This chapter documents the setup and deployment of a workload management solution using WLM following the high-level activities required to implement a stage 2 workload management configuration. See the best practices document for a full description of the various stages:

http://www.ibm.com/developerworks/data/bestpractices/workloadmanagement/

Those high-level activities include the following tasks:

1. Apply the best practices workload management configuration template to your system.
2. Collect baseline monitoring information for estimated cost distribution for SQL statements on your system.
3. Adjust the template work class definitions (and service classes if needed) to better reflect the real work executing in your environment.
4. Collect baseline monitoring information for resource consumption by service classes on your system.
5. Adjust the concurrency thresholds defined on the different service classes to better reflect the resource allocation that you want.
6. Define activity thresholds to protect your system from abnormal queries.
7. Establish a monitoring regime to ensure the ongoing fitness of your tuned configuration.

After the foundation workload management configuration is implemented, InfoSphere Optim Performance Manager also can help you create and manage more advanced configurations such as those described under the Stage 3 scenarios in the best practices document.

8.1 Implementing a stage 2 configuration

WLM is designed with the published best practices in mind. This section shows how to use it to achieve a stage 2 Workload Manager configuration.

8.1.1 Applying the best practices template configuration

The first step in the best practices methodology is to apply the default template configuration to your database. WLM can help you to easily do this. The remainder of this section shows how to accomplish this.

Starting the Workload Manager Configuration tool

To start WLM do the following steps:

1. From the InfoSphere Optim Performance Manager web console, click **Open** and select **Workload Manager Configuration** under the Configuration column (Figure 8-2).

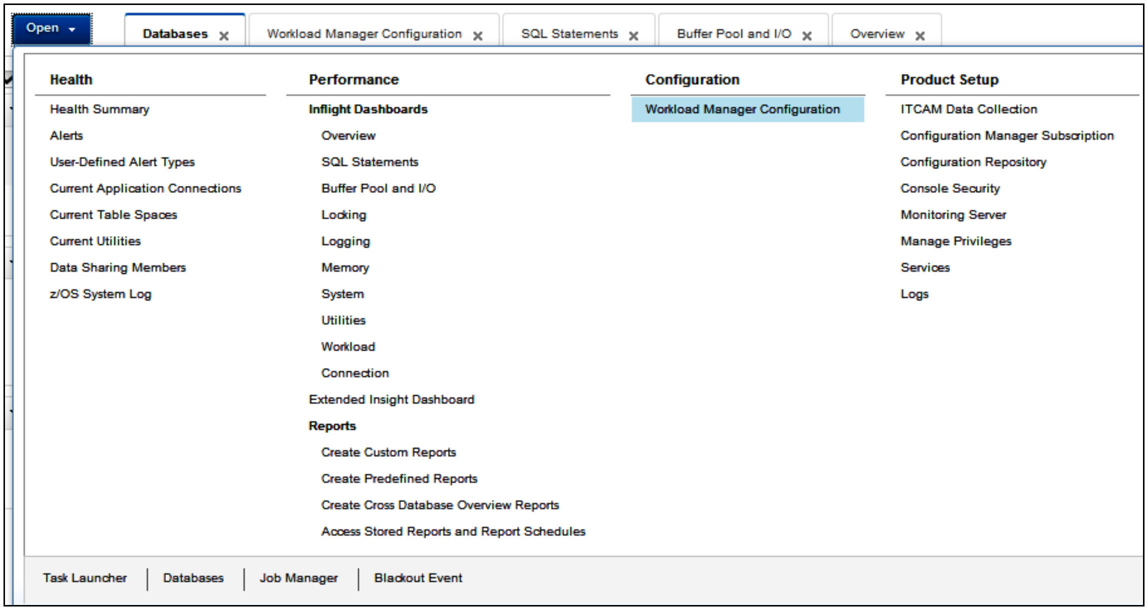

Figure 8-2 Invoking Workload Manager Configuration

2. Select the database that you want to configure and click **Connect** to connect to the database. A panel opens to show an overview of the current WLM configuration (Figure 8-3).

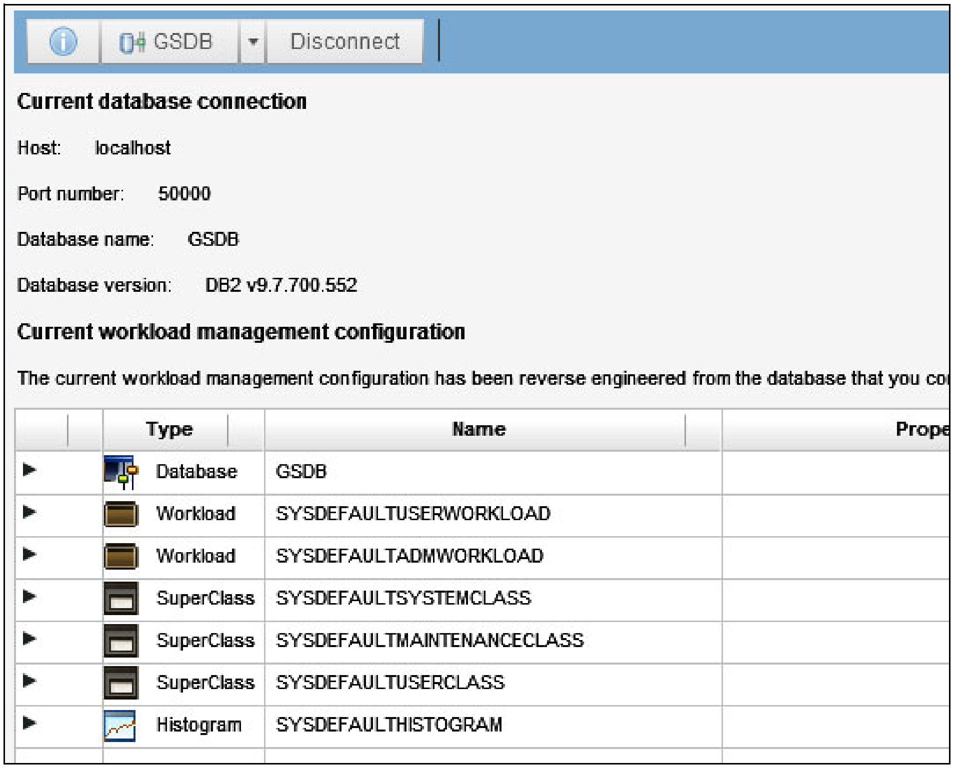

Figure 8-3 Workload Manager Configuration Wizard

This panel shows the existing WLM configuration. If you have not done any WLM configuration yet, you see only the system defined workloads and service classes (Figure 8-3). If you previously added any workloads or service classes, you see them listed here.

The purpose of this panel is to enable the recovery of changes made in a previous session that have not been deployed. When working in the WLM configuration tool, it saves a draft of the changes you have made whenever you preview SQL or explicitly request it to save a draft. If there is no saved draft when you connect, or if there have been WLM-related changes to the database that has been configured since the draft was created, you see only the existing WLM configuration from the database you are configuring, as shown in Figure 8-4 on page 333. Click **OK** and proceed to the configuration tool.

If there is a saved draft, use this panel to determine if it contains changes you want to preserve in this session. The Differences column has a symbol (!=) next to any workloads or service classes that are different in the draft, as depicted in Figure 8-4.

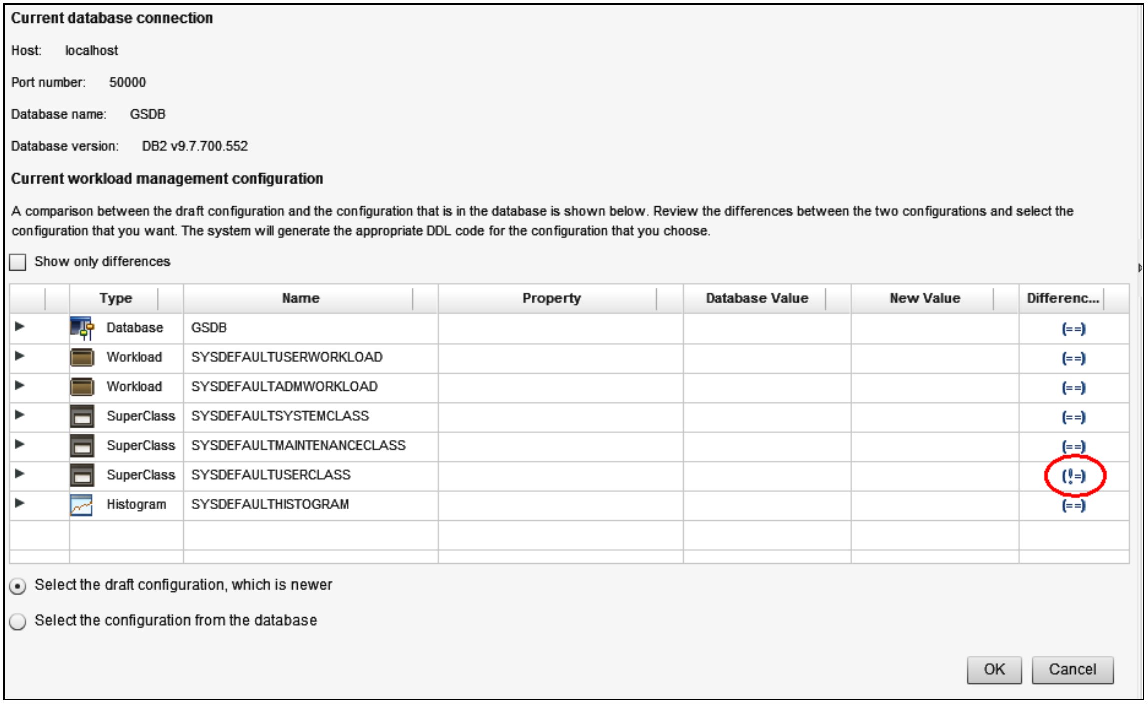

Figure 8-4 Choose to work with the saved draft, or discard it

Drill down into any workloads or service classes that are marked as changed to see what the differences are. In this example, the Collect property of the SYSDEFAULTSUBCLASS service subclass, under the service superclass SYSDEFAULTUSERCLASS is set to `NONE` in the database being configured and `Base` in the saved draft. See Figure 8-5.

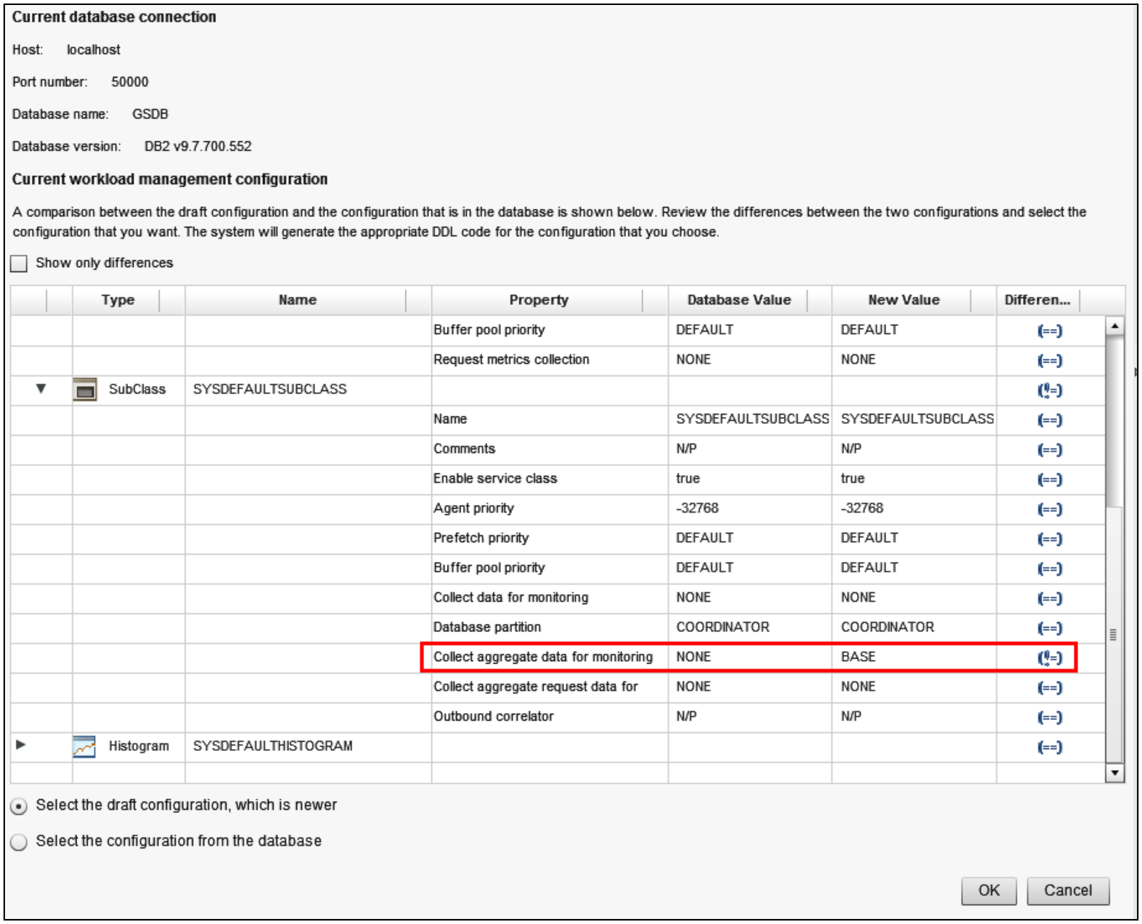

Figure 8-5 Examining the differences in the saved draft

3. After selecting the appropriate choice for working with the draft configuration or the configuration from database click **OK** to begin working with the WLM configuration.

The best practices template configuration

One of the primary objectives in the published best practices for workload management is to protect your data server from becoming overloaded by limiting the number of large queries that are allowed to run concurrently. This is accomplished by creating separate service subclasses for various sized queries. Incoming queries are mapped to the appropriate service subclass based on their estimated cost. Figure 8-6 shows this concept.

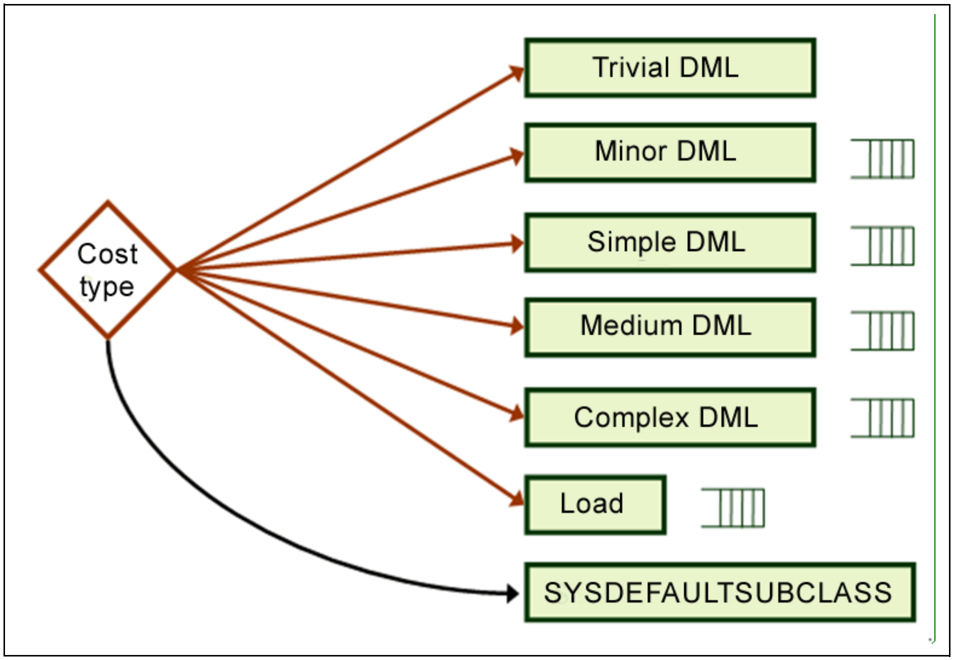

Figure 8-6 Mapping activities to service subclasses based on type or estimated cost

Whether you are explicitly following the steps outlined in the best practices, the infrastructure for mapping queries to service subclasses based on their type or estimated cost is a necessary foundation for almost any WLM solution. For this reason, the Workload Manager Configuration tool offers to create a service superclass containing this infrastructure whenever you use it to work with a database that does not already contain such a service superclass.

When you connect to a database that does not yet contain at least one service superclass with work actions to map activities to service subclasses, it will offer to create one for you. See Figure 8-7.

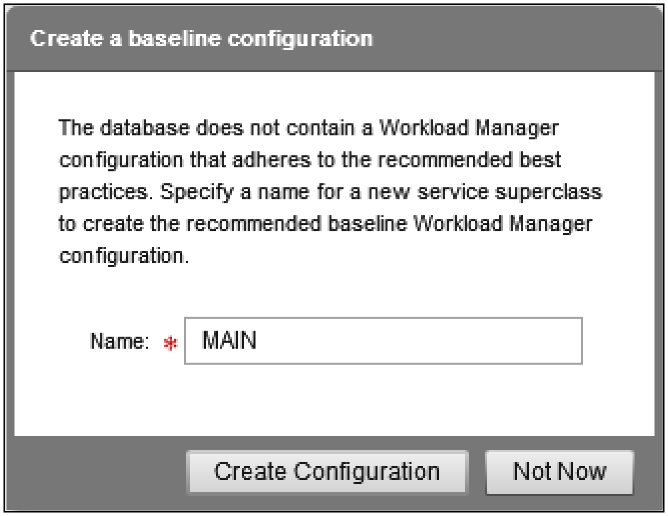

Figure 8-7 The configuration tool offers to create a template service superclass

In the template configuration generated by the Workload Manager Configuration tool, controls such as thresholds for limiting concurrency are present but disabled. This means it is safe to deploy the template configuration and gather monitoring data and there will be no change in the behavior of your database until you are ready to begin imposing limits.

Deploying your changes

Changes made in the Workload Manager Configuration tool are not deployed to the database until you explicitly do so. You can see the DDL that will implement the changes by clicking **Preview and Run SQL**. The DDL will be displayed and you can then have InfoSphere Optim Performance Manager run the DDL or you can copy and paste it to run later, manually. See Figure 8-8 and Figure 8-9 on page 337.

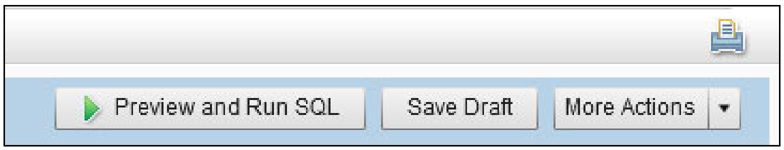

Figure 8-8 Click "Preview and Run SQL" to generate and display DDL

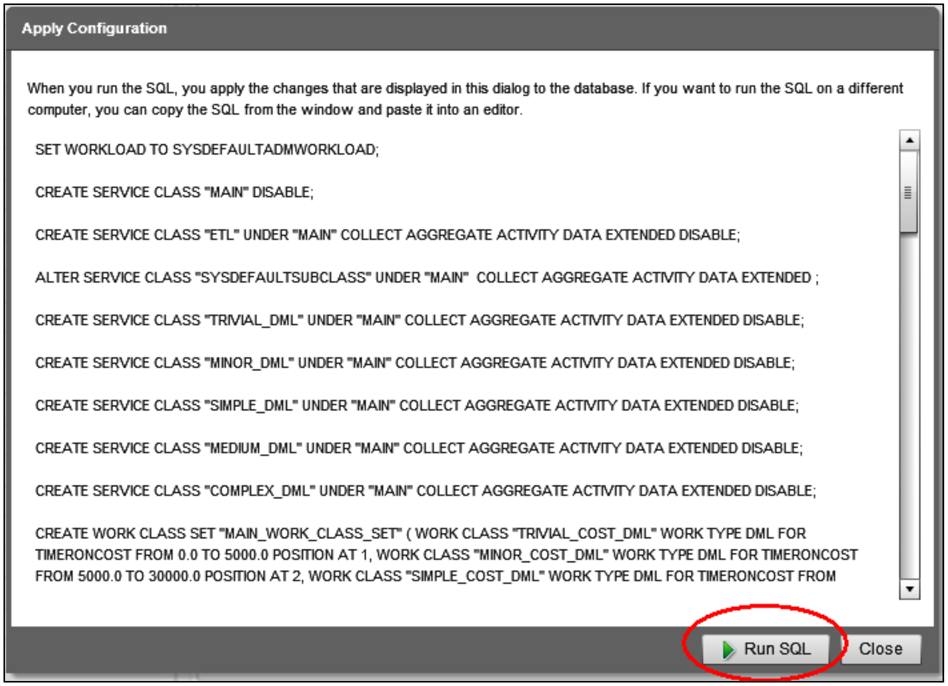

Figure 8-9 Execute the DDL directly, or copy and paste it.

Running the generated DDL from the Workload Manager Configuration tool has several benefits. One benefit is that if your changes involve dropping WLM objects that are in use, they can be dropped only after they have been disabled and all activity involving them has completed. The Workload Manager Configuration tool will automate this “wait and drain” process for you. Another benefit is that DB2 does not support multiple WLM DDL statements in a single transaction, so you cannot use transactions to automatically roll back a partially completed deployment. The Workload Manager Configuration tool will handle this for you; if something prevents the deployment from completing, it will attempt to restore your WLM configuration to its original state by generating and executing compensating DDL.

8.1.2 Collect baseline monitoring for estimated cost distribution

Now that you have configured workload management to categorize queries based on their estimated cost, you are ready to collect monitoring information. This monitoring information will be used to tune the workload management configuration to match your particular workload and environment.

Configuring InfoSphere Optim Performance Manager for collection of WLM statistics

Before you can use the monitoring graphs and reports of WLM, you must configure InfoSphere Optim Performance Manager to monitor the WLM statistics:

1. On the InfoSphere Optim Performance Manager web console, click **Databases,** select the database for which you want to collect Workload Manager metrics, and click **Configure Monitoring** → **Monitor**.
2. Step through the configure monitoring wizard and enable collection of WLM monitoring information, as shown in Figure 8-10.

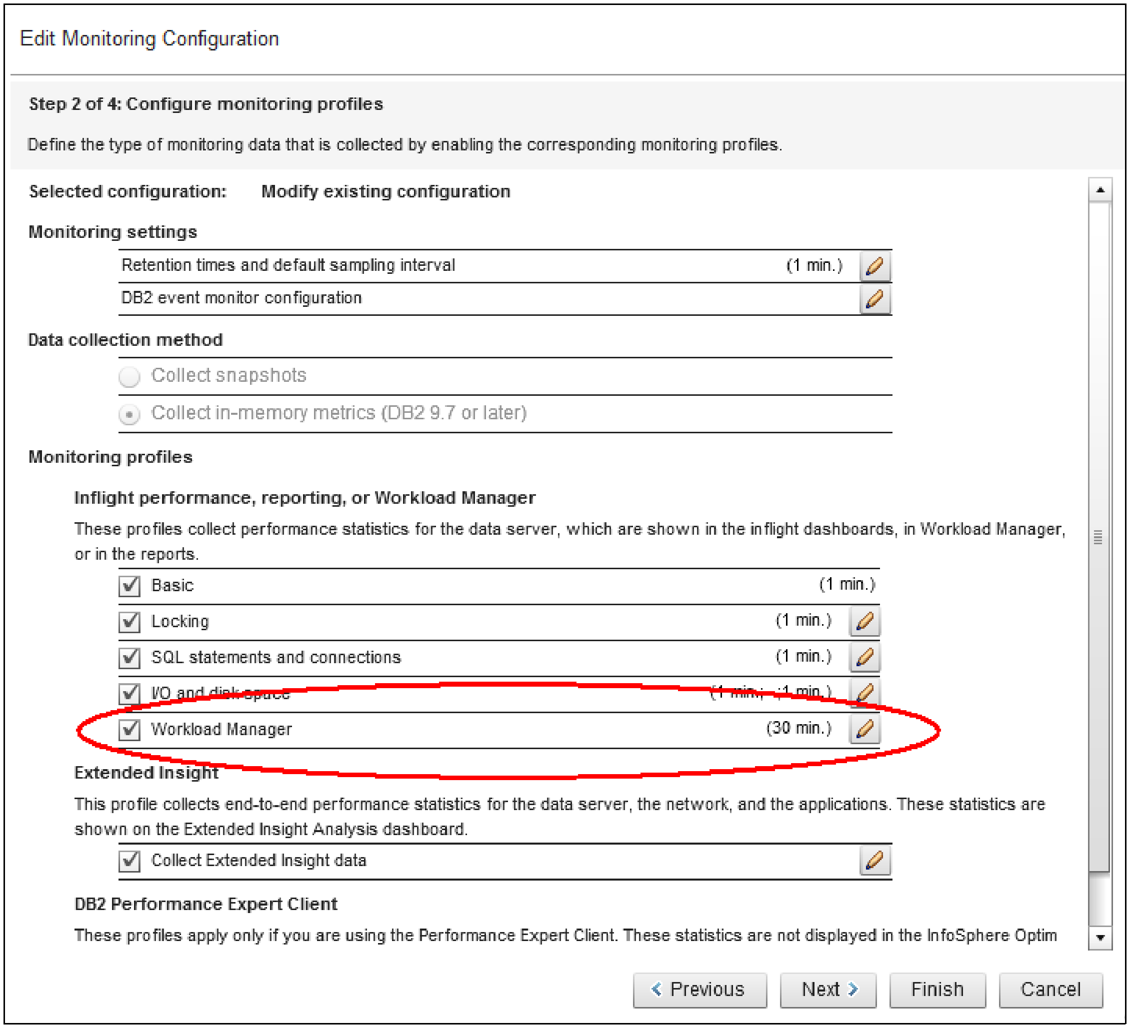

Figure 8-10 Collecting the WLM monitoring information

3. Adjust how frequently statistics are captured for Workload Manager by clicking the **pencil icon**. The recommended sampling interval for ongoing monitoring is the default value of once every 30 minutes. This provides sufficient granularity for monitoring a working WLM configuration without consuming too much space in the repository database. When first tuning the WLM configuration, a useful step is to temporarily set a more frequent value so that you can see the effects of changes to the configuration more quickly. See Figure 8-11.

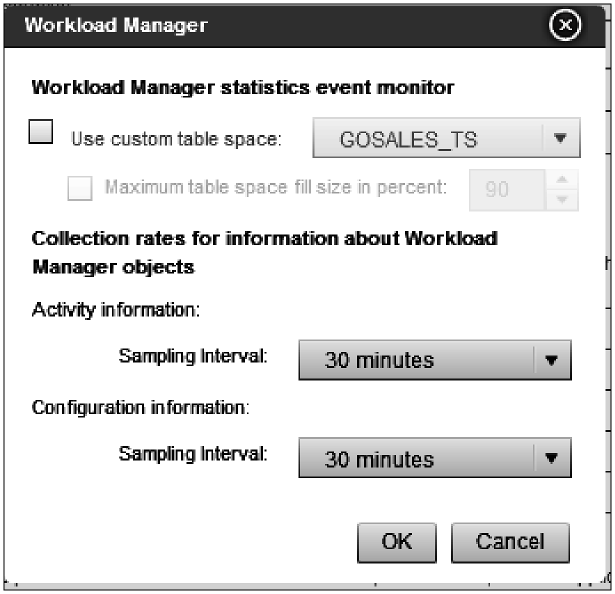

Figure 8-11 Setting the collection interval for the workload statistics event monitor

Accumulating monitoring data

After deploying the template WLM configuration and enabled monitoring, wait for representative monitoring data to accumulate. For the initial round of tuning, extensive monitoring data is not required, only enough to be somewhat representative of the normal workload for the data server.

8.1.3 Adjusting template definitions to better suit your environment

After you have monitoring data for a representative sample of the workload for your data server, the next step is to examine the distribution of estimated costs to see if you need to adjust how queries are mapped to service subclasses. The template used by the configuration tool includes five service subclasses for queries with various ranges of estimated costs. What is considered a small, medium, or large query differs for each database. Therefore, one of the first adjustments you might make is to the estimated costs that define what is mapped to each service subclass.

Evaluate the distribution of estimated cost

The objective of this step is simply to group queries into relatively homogenous groups of work. If there are no obvious imbalances in the distribution, no action is necessary. The default values for estimated costs provided in the template often work well with no adjustment.

You may merge some of these service subclasses as you proceed. The template provides more buckets than are necessary for many databases because it is easier to merge under utilized service subclasses than to split crowded ones.

Monitoring charts and data are accessed by clicking **Show charts** in the workload, service superclasses, and service subclasses tabs. In this step, examine the distribution of estimated costs for all activities. Because all activities are running in the workload SYSDEFAULTUSERWORKLOAD, look at that workload to see estimated costs for all activities in a single chart. In Figure 8-12 on page 341 tall bars (blue) have been superimposed on the right side of this figure showing the default boundaries for mapping queries to service subclasses, as supplied in the template configuration. Because this sample workload is composed of smaller queries than would be seen in a typical warehouse, all of the works fall into the first two buckets.

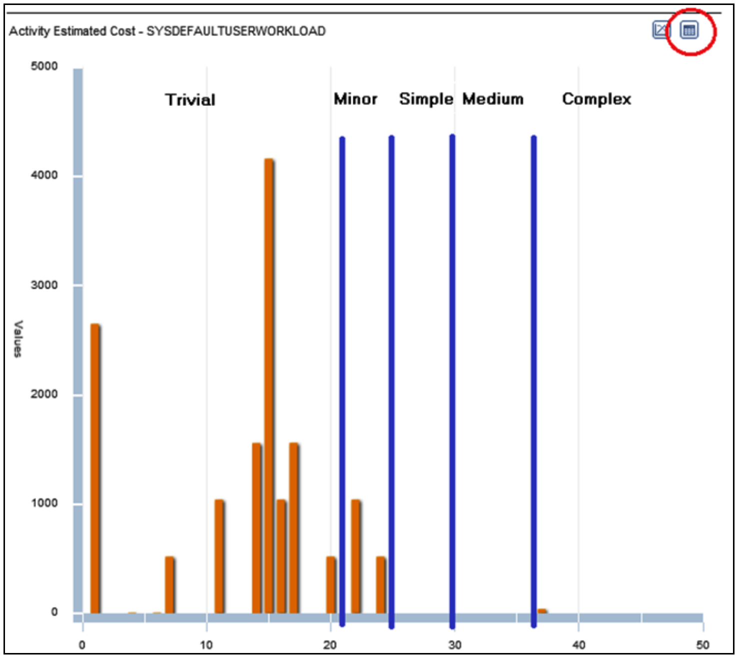

Figure 8-12 Estimated cost histograms annotated with service subclass mappings

The histograms as shown in Figure 8-12 are useful for getting a sense of the distribution. For choosing the actual numbers to adjust to, it is helpful to look at the tabular view. This can be accessed by clicking the charts icon circled (in red) in Figure 8-12.

This tabular view for same data is shown in Figure 8-13.

Activity Estimated Cost

Bin Number	Value	Bin Top	Bin Bottom
11	1040	103	68
12	0	158	103
13	0	241	158
14	1557	369	241
15	4159	562	369
16	1040	858	562
17	1560	1309	858
18	0	1997	1309
19	0	3046	1997
20	520	4647	3046

List Data

Figure 8-13 Tabular view for trivial DML service subclass

Consider that the histogram data is presented in a logarithmic scale so that the bins to the actual range of values represented by each bin are larger as you moves from left to right. See the Bin Magnitude line (Figure 8-14 on page 343), which represents the scaled bin heights, that is superimposed over the normal logarithmic representation for an illustration of this principle. In this case, although there is an apparent wide spread of data in the trivial service class range of bins, it is really a small spread, in actual cost values, and fairly homogenous. The minor range of value has a small and relatively balanced across the whole range and the other ranges have no queries present except for a few outliers at the low end of the complex range.

In summary, adjusting any of the estimated cost boundaries for mapping activities to service subclasses to distribute them more evenly is probably not worthwhile doing. If this information reflected the entire workload, then the only change to be considered is whether to collapse the simple, medium, and complex ranges into one complex range representing all three of the original ranges.

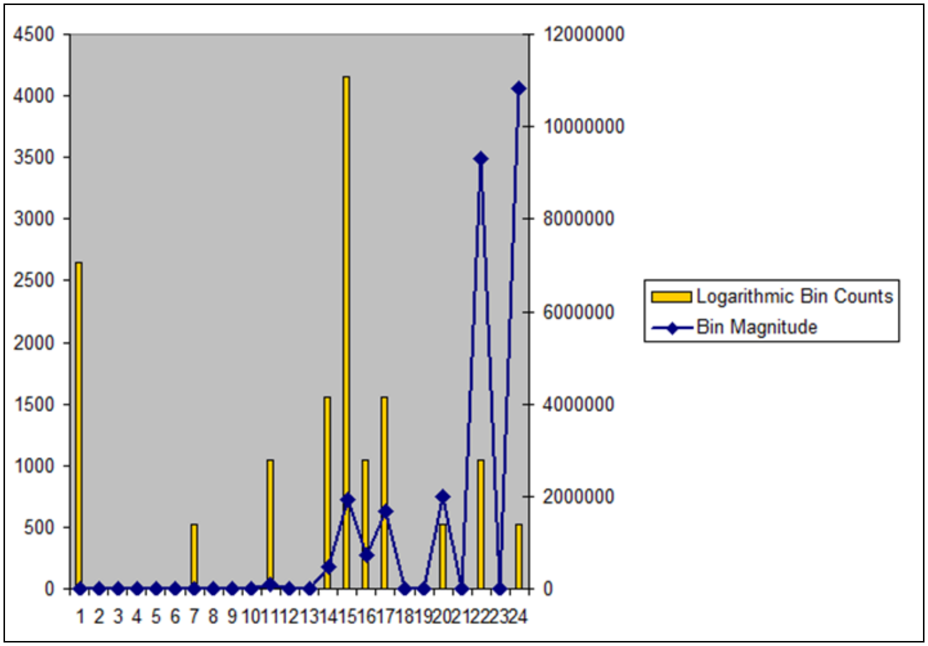

Figure 8-14 Scaled bin heights superimposed over normal logarithmic representation

Distribute queries based on estimated cost

The top level Service subclasses tab in the tooling is where you tune the template WLM configuration for your particular environment and workload. Here, you can adjust the costs that define what are considered simple, medium, or complex queries in your environment and specify the limits on how many are allowed to run concurrently.

Based on the analysis described, you need only three service subclasses. So you can delete TRIVIAL_DML and MINOR_DML, then adjust the estimated costs, as shown in Figure 8-15 on page 344.

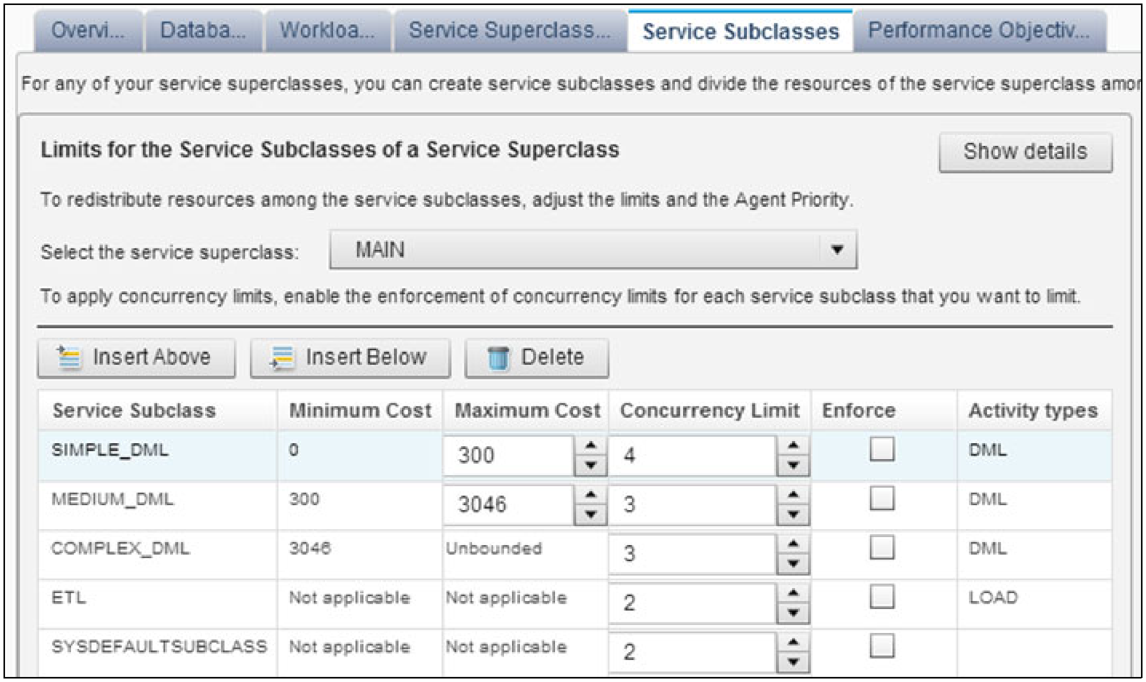

Figure 8-15 Estimated costs for distributing queries after adjustment

Enable concurrency thresholds

The WLM Best Practices document provides guidelines for determining some initial concurrency settings. These guidelines are based on a multiplier that depends on the processing power of the hardware. Workload Manager Configuration in InfoSphere Optim Performance Manager does this calculation for you but always uses a multiplier of 10x when it instantiates a template, regardless of the actual hardware it is running on. For the hardware used in this example, the WLM Best Practices guidelines suggest a multiplier of 5x. There are four CPUs, so the initial estimate of the total concurrency of this hardware is 20.

In this specific example, the service classes being used have been reduced to TRIVIAL_DML, MINOR_DML and COMPLEX_DML (which now encompasses the range of estimated cost values once represented by the original default SIMPLE_DML, MEDIUM_DML, and COMPLEX_DML ranges). In this scenario, no concurrency thresholds are currently enforced on these service classes.

Depending on the overall system utilization and the actual performance of the work in each service class compared to the objectives, you might want to define an active concurrency threshold on either or both of the following service classes to direct system resource to the work in the TRIVIAL_DML service class:

- COMPLEX_DML
- MINOR_DML

8.1.4 Collecting baseline monitoring for resource consumption

This section describes how to bring actual resources consumption by each service subclass closer to the levels you want. Specifically, you want to adjust the CPU resources based on the response times being seen versus the performance expectations for each service subclass. Resource consumption is adjusted by changing the concurrency threshold values used in each service subclass.

Before making these adjustments, collect and review baseline monitoring data on CPU consumption and response time. By default, the service classes from the template will be configured to gather the necessary monitoring data.

InfoSphere Optim Performance Manager will collect workload management statistics from DB2 and store them in the repository database where the information is available for access using the custom report facility. The custom reporting facility can be used to review it. Although many data categories are available, the following steps are of most interest when working with workload management:

- Operating system metrics and statistics

  This data is accessed through the REPORT_SYSTEM table function and provides insight into the overall system resource utilization and performance. For more information see the following location:

  http://ibm.co/18sJ4sr

- Workload Manager service class statistics

  This data represents the statistics for the work processed within each service class during the last monitoring interval. It is accessed through the REPORT_SCSTATS table function, with more information available at the following address:

  http://ibm.co/16yWAYN

- Workload Manager workload statistics

  This data represents the statistics for the work submitted by each workload during the last monitoring interval. It is accessed through the REPORT_WLSTATS table function, with more information available at the following address:

  http://ibm.co/1971ytA

- Workload Manager queue statistics

  This data represents the statistics for each concurrency threshold (all types) for work queued during the last monitoring interval. It is accessed through the REPORT_QSTATS table function, with more information available at the following address:

  http://ibm.co/18sJsa9

- Workload Manager histogram statistics

  This data contains the histogram data, collected for each monitoring interval, for the different sets of distribution information that can be collected by the different workload management objects. It is accessed through the REPORT_WLM_HISTOGRAM table function, with more information available at the following address:

  http://ibm.co/1gqeX30

- Workload Manager work class statistics

  This data represents the statistics from each monitoring interval for all the work classified by the different work class definitions as they are processed through the different work action sets in the workload management configuration. It is accessed through the REPORT_WCSTATS table function, with more information available at the following address:

  http://ibm.co/1bnDV64

Example 8-1 is a query that uses the custom reporting table functions that shows the metrics needed for collecting baseline monitoring for resource consumption.

Example 8-1 Using the custom reporting table functions

```
WITH
TOTAL_OS AS
(SELECT OPM_MIN_COLLECTION_TIMESTAMP, OPM_MAX_COLLECTION_TIMESTAMP,
MEMBER,
        CPU_USAGE_TOTAL AS OS_UTILIZATION,
        ((DB2_CPU_USER + DB2_CPU_SYSTEM * 100.00) / (CPU_USER +
CPU_SYSTEM)) AS DB2_RATIO_TO_OS
 FROM
TABLE(REPORTS_1.REPORT_SYSTEM_SHORT_TERM('2012-09-27-21.00.00.000000',
```

```
'2012-09-28-02.00.00.000000',
'HOUR',
'PER_MEMBER'))
ORDER BY OPM_MIN_COLLECTION_TIMESTAMP, OPM_MAX_COLLECTION_TIMESTAMP,
MEMBER),

TOTAL_MEMBER AS
(SELECT OPM_MIN_COLLECTION_TIMESTAMP, OPM_MAX_COLLECTION_TIMESTAMP,
MEMBER,
       TOTAL_CPU_TIME AS TOTAL_MEMBER_CPU_TIME
 FROM
TABLE(REPORTS_1.REPORT_SCSTATS_SHORT_TERM('2012-09-27-21.00.00.000000',

'2012-09-28-02.00.00.000000',
'HOUR',
'PER_MEMBER',
'DATABASE'))
ORDER BY OPM_MIN_COLLECTION_TIMESTAMP, OPM_MAX_COLLECTION_TIMESTAMP,
MEMBER),

TOTAL_SC AS
(SELECT OPM_MIN_COLLECTION_TIMESTAMP, OPM_MAX_COLLECTION_TIMESTAMP,
MEMBER,
SERVICE_SUPERCLASS_NAME, SERVICE_SUBCLASS_NAME,
TOTAL_WAIT_TIME AS TOTAL_SC_WAIT_TIME,
TOTAL_CPU_TIME AS TOTAL_SC_CPU_TIME,
COORD_ACT_LIFETIME_AVG AS AVG_SC_LIFETIME,
COORD_ACT_LIFETIME_TOP AS MAX_SC_LIFETIME,
COORD_ACT_QUEUE_TIME_AVG AS AVG_SC_QUEUE_TIME
 FROM
TABLE(REPORTS_1.REPORT_SCSTATS_SHORT_TERM('2012-09-27-21.00.00.000000',

'2012-09-28-02.00.00.000000',
                                                'HOUR',
                                                'PER_MEMBER',
                                                'SERVICESUBCLASS'))
ORDER BY OPM_MIN_COLLECTION_TIMESTAMP, OPM_MAX_COLLECTION_TIMESTAMP,
MEMBER,
         SERVICE_SUPERCLASS_NAME, SERVICE_SUBCLASS_NAME)

SELECT TOTAL_SC.OPM_MIN_COLLECTION_TIMESTAMP,
TOTAL_SC.OPM_MAX_COLLECTION_TIMESTAMP,
TOTAL_SC.MEMBER,
       TOTAL_SC.SERVICE_SUPERCLASS_NAME,
TO-TAL_SC.SERVICE_SUBCLASS_NAME,
```

```
       TOTAL_SC_WAIT_TIME,
AVG_SC_LIFETIME,
MAX_SC_LIFETIME,
AVG_SC_QUEUE_TIME,
       TOTAL_SC_CPU_TIME,
       TOTAL_MEMBER_CPU_TIME,
       (100.00 * TOTAL_SC_CPU_TIME)/TOTAL_MEMBER_CPU_TIME AS
SC_RATIO_TO_DB2,
       DB2_RATIO_TO_OS,
       OS_UTILIZATION
FROM TOTAL_SC, TOTAL_OS, TOTAL_MEMBER
WHERE TOTAL_SC.OPM_MIN_COLLECTION_TIMESTAMP =
TO-TAL_OS.OPM_MIN_COLLECTION_TIMESTAMP
  AND TOTAL_SC.OPM_MAX_COLLECTION_TIMESTAMP =
TO-TAL_OS.OPM_MAX_COLLECTION_TIMESTAMP
  AND TOTAL_SC.MEMBER = TOTAL_OS.MEMBER
  AND TOTAL_SC.OPM_MIN_COLLECTION_TIMESTAMP =
TO-TAL_MEMBER.OPM_MIN_COLLECTION_TIMESTAMP
  AND TOTAL_SC.OPM_MAX_COLLECTION_TIMESTAMP =
TO-TAL_MEMBER.OPM_MAX_COLLECTION_TIMESTAMP
  AND TOTAL_SC.MEMBER = TOTAL_MEMBER.MEMBER
ORDER BY TOTAL_SC.OPM_MIN_COLLECTION_TIMESTAMP,
         TOTAL_SC.OPM_MAX_COLLECTION_TIMESTAMP,
         TOTAL_SC.MEMBER,
         TOTAL_SC.SERVICE_SUPERCLASS_NAME,
         TOTAL_SC.SERVICE_SUBCLASS_NAME;
```

To use this, you must know which schema in the repository database corresponds to the database being configured. One way to determine the schema number is by querying the DB2PM.DATABASES table in the repository database, as shown in Example 8-2.

Example 8-2 Querying the DB2PM.DATABASES table

```
SELECT D_DB_NAME, D_I_INSTANCE_ID FROM DB2PM.DATABASES;
```

The database that was to be queried had instance ID 1, so the schema to use in the custom reporting is REPORTS_1. See the following address for more details:

http://ibm.co/1aOG4bv

8.1.5 Adjusting concurrency thresholds

Determine which service classes should get more or fewer CPU resources based on their current performance compared to any performance objectives for that set of work. If there are no specific objectives, then this exercise is intended to try to find the right balance between restricting the more complex work to maintain the responsiveness of the lighter work: try some iterative exercises of lowering or raising the concurrency thresholds to see the effect on response times.

For the machine that produced these examples, using the guidelines outlined in the Best Practices document, the estimated capacity is a concurrency of approximately 20. The suggested starting point is to allocate about 5% of the total concurrency to the COMPLEX_DML service subclass. This works out to only one concurrency ticket for COMPLEX_DML. Although it is a valid level, it is a restrictive one. Because you will be dealing with a range of work within each service class, it might be possible, depending on the current resource utilization of the system, to start with a slightly higher value to maximize the throughput of the work in this service class while still restraining it sufficiently to achieve the goals for the other service classes. With this in mind, for this example, you can choose to start with 3, and if that is not sufficient, try reducing it further to 2.

The response times for each service subclass can be seen in the service subclasses response time report, for response time percentiles after enforcement of concurrency limits. The overall concurrency for the environment can be viewed by examining the service superclass chart as seen in Concurrency for service superclass MAIN.

You can also see a breakdown of concurrency by service subclass by navigating to the Concurrency and Time in the Queue tab, under the Service Subclasses top level tab, and also the average percentage of time queries spent queued during each collection interval. Although they sometimes spend as much as 60% of their time queued, the overall response time is typically much improved. The benefits of avoiding interference from too many large queries running concurrently is so great that it outweighs the time queries queued waiting to start.

You can see a different perspective on these same performance statistics by looking at the results from the sample query, shown at the beginning of this section. Figure 8-16 on page 350 shows the selected metrics for the same set of queries run before and after enforcing concurrency limits:

SERVICE_SUBCLASS_NAME	AVG_SC_LIFETIME	MAX_SC_LIFETIME	AVG_SC_QUEUE_TIME	SC_RATIO_TO_DB2
TRIVIAL_DML	1331550	41497000	0	16.88916714
MINOR_DML	6813950	85580000	0	45.25605366
COMPLEX_DML	166886650	1095869000	0	34.82373724
ETL	0	0	0	0
SYSDEFAULTSUBCLASS	1100	120000	0	0.993526108

Figure 8-16 Chart before enforcement of Concurrency limits

Figure 8-17 shows the metrics after concurrency limits were enforced:

SERVICE_SUBCLASS_NAME	AVG_SC_LIFETIME	MAX_SC_LIFETIME	AVG_SC_QUEUE_TIME	SC_RATIO_TO_DB2
TRIVIAL_DML	1244850	23980000	0	20.99830921
MINOR_DML	4723250	42972000	0	61.08241989
COMPLEX_DML	87583000	821250000	30316900	15.08448789
ETL	0	0	0	0
SYSDEFAULTSUBCLASS	750	90000	0	1.508516142

Figure 8-17 Chart after enforcement of concurrency limits

In addition to improving performance overall, restricting concurrency on the COMPLEX_DML service subclass has freed more CPU time for use by the other service subclasses.

Response time metrics were also collected with the concurrency on COMPLEX_DML limited to 2, and also for a concurrency of 3. Figure 8-18 shows the response time for each service subclass for the various concurrency limits tried on COMPLEX_DML. Response times are number of seconds in which 90% of the queries in that service completed.

Concurrency limit on COMPLEX_DML	Response time for TRIVIAL_DML	Response time for MINOR_DML	Response time for COMPLEX_DML
(unlimited)	12.13	50.99	808.1
3	9.01	26.29	717.1
2	8.8	26.5	802.6

Figure 8-18 Response time for the different Service classes

As shown in Figure 8-18, significant improvement in response time was achieved for queries in all service subclasses by applying a concurrency limit 3 on COMPLEX_DML, compared to the response times with no concurrency limit in place. Further restricting the concurrency on COMPLEX_DML gave a smaller incremental benefit to TRIVIAL_DML and MINOR_DML, but at the cost of a disproportionate penalty for response time in COMPLEX_DML. So the choice between 2 and 3 might be based on response time goals and business priorities. If the queries running in COMPLEX_DML are much less important than those running in TRIVIAL_DML and MINOR_DML, you might prefer to maximize their response times at the expense of COMPLEX_DML.

8.1.6 Defining thresholds to protect your system

On some systems, an occasional SQL statement is disruptive. This is perhaps because it runs too long, reads too many rows, or consumes an excessive amount of temporary table space. These queries are sometimes referred to as *rogue* or *outlaw* queries. InfoSphere Optim Performance Manager enables you to establish limits to reduce the impact of such queries by creating thresholds that monitor or stop them if they exceed the limit.

Thresholds for controlling abnormal queries

When you are satisfied that the concurrency limits are set appropriately, examine the metrics collected by InfoSphere Optim Performance Manager to determine appropriate limits for each service subclass. Some thresholds that are most useful for dealing with rogue queries are as follows:

- ActivityTotalTime

  This threshold limits how long a query can run, including time spent queued waiting to start, waiting for locks, and so forth. Because the total time for activities is affected by the concurrency settings, the best approach is to have them well established before trying to set values for this threshold.

- CPUTime

  This threshold limits how much CPU time a query can consume. Because it does not include wait times, it is more indicative of how much real work a query has done.

- SQLRowsRead

  This threshold limits how many rows a query can read. It can be used to prevent large table scans, such as can happen in the case of missing indexes, missing filter predicates or accidental Cartesian products.

- SQLRowsReturned

  This threshold prevents queries from consuming excessive network bandwidth by returning too many rows to the client. It is sometimes also useful in situations to enforce policies to discourage users from extracting large blocks of data from the database.

Determine threshold limits from monitoring data

For setting each of these thresholds, corresponding monitoring metrics are collected by InfoSphere Optim Performance Manager that show the high values for each service subclass during a period of time. Example 8-3 shows a query using the customer reporting facility for viewing the values.

Example 8-3 Top Service subclass Values

```
WITH
TOTAL_SC AS
(SELECT OPM_MIN_COLLECTION_TIMESTAMP, OPM_MAX_COLLECTION_TIMESTAMP,
MEMBER,
   SERVICE_SUPERCLASS_NAME, SERVICE_SUBCLASS_NAME,
   TOTAL_WAIT_TIME AS TOTAL_SC_WAIT_TIME,
   TOTAL_CPU_TIME AS TOTAL_SC_CPU_TIME,
   COORD_ACT_LIFETIME_AVG AS AVG_SC_LIFETIME,
   COORD_ACT_LIFETIME_TOP AS MAX_SC_LIFETIME,
   COORD_ACT_QUEUE_TIME_AVG AS AVG_SC_QUEUE_TIME,
   ACT_CPU_TIME_TOP,
   ACT_ROWS_READ_TOP,
   ROWS_RETURNED_TOP,
   UOW_TOTAL_TIME_TOP
 FROM
TABLE(REPORTS_5.REPORT_SCSTATS_SHORT_TERM('2012-09-27-21.00.00.000000',

'2012-09-28-02.00.00.000000',
                                                'ALL',
                                                'PER_MEMBER',
                                                'SERVICESUB-CLASS'))
ORDER BY OPM_MIN_COLLECTION_TIMESTAMP, OPM_MAX_COLLECTION_TIMESTAMP,
MEMBER,
         SERVICE_SUPERCLASS_NAME, SERVICE_SUBCLASS_NAME)
SELECT *
FROM TOTAL_SC;
```

Consider the following information:

- For CPU Time, refer to column value ACT_CPU_TIME_TOP returned from the SQL query output that is executed.
- For SQLRowsRead, refer to column value at ACT_ROWS_READ_TOP returned from the SQL query output that is executed.
- For SQLRowsReturned, refer to column value at ROWS_RETURNED_TOP returned from the SQL query output that is executed.
- For ActivityTotalTime, look at AVG_SC_LIFETIME and MAX_SC_LIFETIME. You can also see the histograms for ActivityTotalTime in the tab labeled Service Subclass Histograms on the monitoring page of the top level Service Subclasses tab, depicted in Figure 8-19.

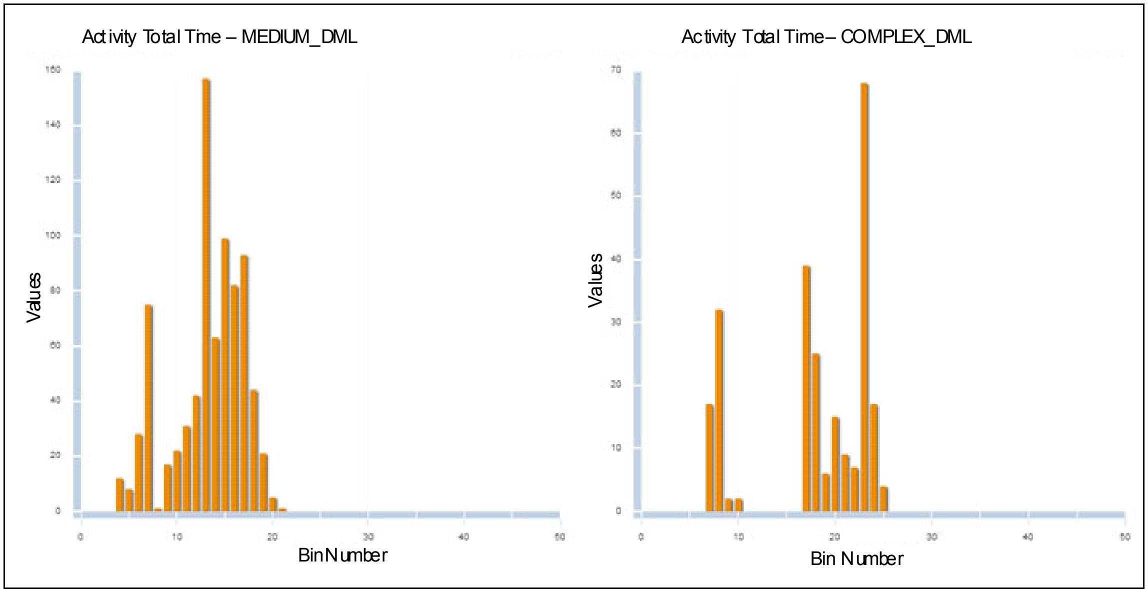

Figure 8-19 Activity Total Time Histogram from the WLM configuration tooling

Thresholds for service subclasses can be set from the Thresholds tab under the top level Service Subclasses tab. Select the **Stop the activities that exceed the limit** box if you want them to automatically be stopped, as in Figure 8-20.

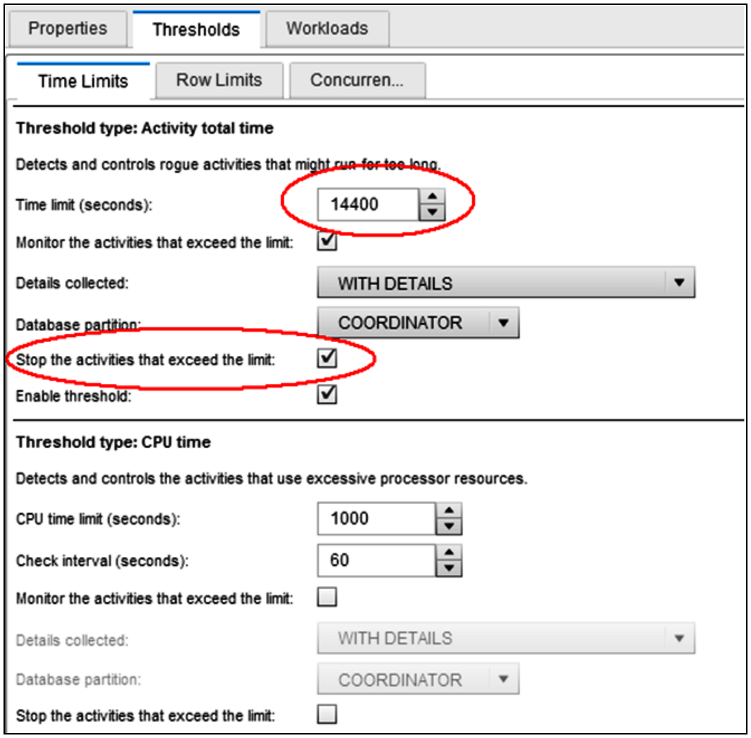

Figure 8-20 Threshold settings

8.1.7 Establishing a monitoring regime to ensure ongoing fitness

Monitoring is a critical part of maintaining a system that has reached a stable stage 2 workload management implementation to ensure that the system continues to remain over time. There are three fundamental objectives underlying the proposed monitoring regime:

- Watching for signs that the system remains in a healthy state
- Investigating and resolving any individual activities that are identified as disruptive (also referred to as rogue or outlaw queries)
- Watching for changes in resource consumption patterns

Keeping the system in a healthy state can be accomplished by monitoring the key attributes that represent a healthy, responsive system to ensure that they remain in the optimal range over time. Key attributes are as follows:

- Run queue length is less than 10.
- Overall CPU utilization is approximately 80 - 95%, with system CPU usage below 10 - 20%. The target value for system CPU usage varies by platform. For example, the target is less than 10% on Linux and less than 20% on AIX operating systems.
- Memory utilization is below 100% (that is, no paging).
- I/O waits are 10% or less.
- System workload is evenly balanced across all members (that is, no skew or uneven resource demands).

The run queue length is visible on the overview dashboard. CPU and memory metrics are also visible on the overview dashboard, and also the system dashboard.

To get I/O wait time, use the operating system tools such as **vmstat** or **nmon**. The **db2pd** tool used for monitoring and administrating DB2 databases can also be used to get a higher level overview of the setup.

The Member Details tab, available in most of the InfoSphere Optim Performance Manager dashboards, can be used to determine if data skews occur between members. You can also use the database connection report to quickly visualize if there is more than expected traffic on specific nodes in a DPF environment.

8.2 More best practices: Advanced configurations

After deployment of the WLM template configuration, use the Workload Manager Configuration tool to further customize WLM objects, service superclasses, workloads, service subclasses, and thresholds.

8.2.1 Privileged and restricted users or applications

Often you will want particular users or applications to be immune to the restrictions imposed by the template service superclass for expensive queries. The easiest way to accomplish this is to create a workload that characterizes the privileged user or application and associate it directly with a separate service class that has fewer or no restrictions.

Generally, a better approach is to create a new service subclass for the work from the privileged user, rather than routing it to one of the existing service subclasses that are part of the template.

The same mechanism can also be used to impose additional restrictions on a particular user or application. Figure 8-21 shows how this might look.

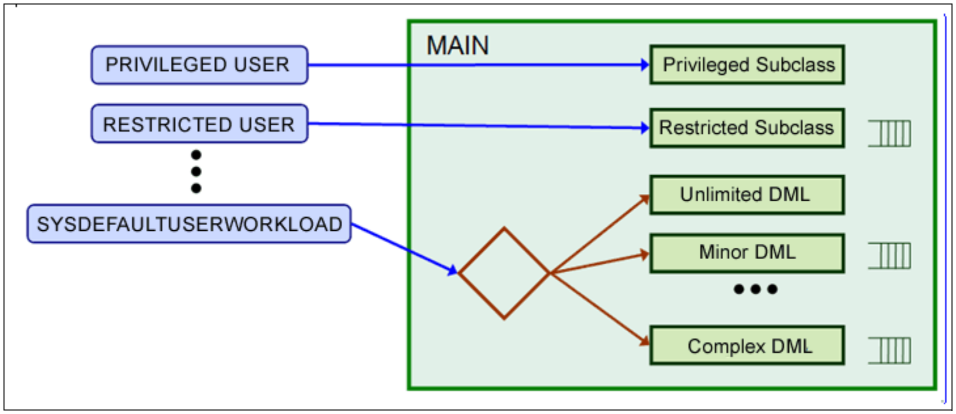

Figure 8-21 Dedicated service classes for activities from privileged and restricted users

8.2.2 Divide resources among lines of business

In most organizations, there are multiple groups of users or business units accessing the data server. A useful approach is to roughly allocate resources to each group to reduce interference. That is, you want to prevent any one group from using so much capacity that other groups are unduly impacted.

To accomplish this, use the following steps:

1. Create one or more workloads to categorize activities for each group.
2. Create a service superclass for each group, using the template to map activities to service subclasses based on their estimated cost or type.
3. Associate the workloads for each group with the service superclass where its activities will run.
4. Adjust the concurrency limits for the service classes to set the appropriate resource limits, as depicted in Figure 8-22 on page 357.

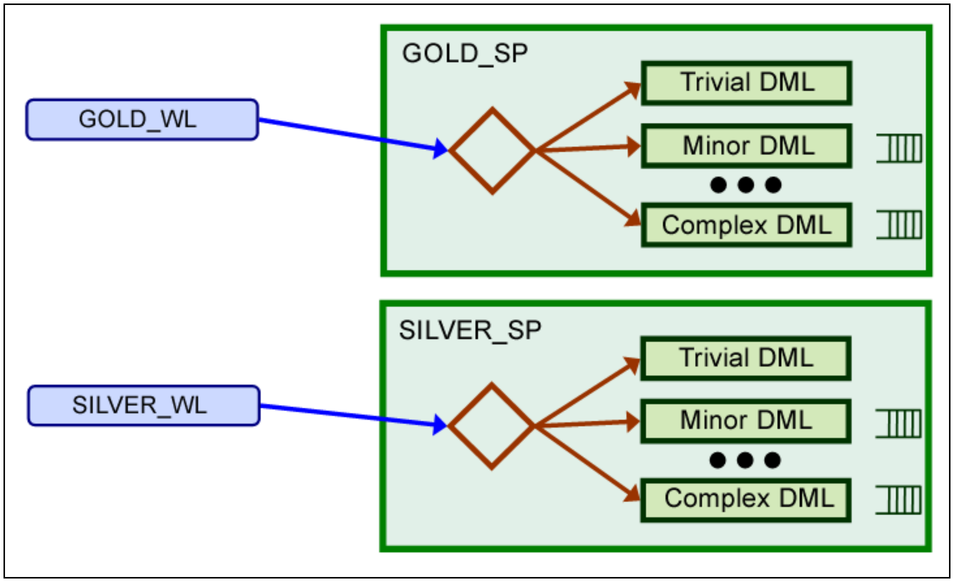

Figure 8-22 Implementing tiered services with a service superclass for each tier

If you want resources divided 60/40 between the service superclasses labeled GOLD_SP and SILVER_SP, for example, then adjust the concurrency limits in each service subclass accordingly. The final result might be similar to those shown in Figure 8-23.

Service Subclass	Concurrency limit in GOLD_SP	Concurrency limit in SILVER_SP
Minor DML	24	16
Simple DML	9	6
Medium DML	6	4
Complex DML	3	2

Figure 8-23 Concurrency adjustments

8.3 Working with Workload Manager configurations

Because of its specific nature, information about working with Workload Manager configurations has been placed in a separate PDF document titled "Additional Material for Performance Management Using IBM InfoSphere Optim Performance Manager and Query Workload Tuner." However, be aware that although we believe it is accurate and helpful information, that material has not been subjected to the formal ITSO editing process. That material is Chapter 2 of the PDF, and titled "Working with Workload Manager configurations." To access that material, see section "Using the web material" on page 384.

8.4 Autonomic performance objectives for workloads

For information about the autonomics performance objectives feature, see *IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows*, SG24-7925.

9

Monitoring packaged database application systems

InfoSphere Optim Performance Manager Extended Edition extends the capability of InfoSphere Optim Performance Manager with end-to-end database monitoring for Java technology and DB2 call level interface (CLI) applications (apps), with ready-for-use configurations for SAP, IBM WebSphere, IBM Cognos, IBM InfoSphere DataStage, and InfoSphere SQL Warehouse applications. The Extended Edition also includes integration with IBM Tivoli monitoring solutions to provide deep database insights into existing IBM Tivoli Composite Application Manager (ITCAM) application-monitoring environments.

This chapter provides an overview of these capabilities and identifies several articles and videos about these environments. These articles are often updated and hence the information might be more current than what is available in this Redbooks publication,

9.1 Monitoring SAP environments

Monitoring an SAP environment with InfoSphere Optim Performance Manager requires different configuration settings. IBM works with SAP to test the recommendations, and to obtain feedback from customers and consultants that continuously improves and shares the knowledge base. The optimization of InfoSphere Optim Performance Manager for an SAP environment includes lock monitor event levels, special DB2 monitor switches, watchdog for event monitors, and predefined system templates.

9.1.1 Configure and monitor SAP workloads with Extended Insight

The feature in InfoSphere Optim Performance Manager (Extended Insight) provides end-to-end database performance monitoring for Java technology, CLI, and .NET database applications, giving you the ability to quickly understand where your database applications are spending their time. With Extended Insight, DBAs can quickly see and understand where database applications are spending time across the IT stack from the application through the SAP application server, the database client, the database server, and the network. When a critical business application is not performing the way you expect, you can use Extended Insight to give you end-to-end visibility from the line of code that issues a database statement, to the database server. This enables you to quickly isolate and address problems by understanding and identifying where the problem is. When the problem is not in the database itself, Extended Insight can help you to more easily work with the people who are managing other parts of the infrastructure or with developers to solve the problem.

Extended Insight recognizes transactions from an SAP Server automatically and groups them by SAP application server host name, SAP user, SAP source module, or SAP transaction. For all characteristics of each attribute, you can see the breakdown of the average response time into the times spent in the SAP application server, in the database client, and in the database server. For the database server time spent, you will get a detailed breakdown into such values as average lock wait time, sort processing time, and I/O processing time. This helps to identify problematic SAP Application Server machines, specific problematic applications, or users.

The following articles and videos explain the installation and configuration of Extended Insight for monitoring SAP applications.

- "Configure and monitor SAP applications with InfoSphere Optim Performance Manager Extended Insight:"

  http://www.ibm.com/developerworks/data/library/techarticle/dm-1202sapoptimperfmgr/index.html

- "InfoSphere Optim Performance Manager in SAP environment -- Part 1:"

  http://www.channeldb2.com/video/infosphere-optim-performance-manager-in-sap-environment-part-1

- "InfoSphere Optim Performance Manager in SAP environment -- Part 2""

  http://www.channeldb2.com/video/infosphere-optim-performance-manager-in-sap-environment-part-2

9.2 Integrate with Tivoli for enterprise level end-to-end monitoring

InfoSphere Optim Performance Manager Extended Edition integrates the deep database performance insight of Optim with the broad enterprise-wide insights provided by IBM Tivoli monitoring products. This powerful combination extends transaction response-time monitoring from the database to the complete end-to-end transaction path.

Database application environments can be complex, often including multiple middleware components through which transactions can flow, including Web servers, application servers, message servers, transaction servers, and database servers.

9.2.1 InfoSphere Optim Performance Manager and ITCAM integration

The IBM Tivoli Composite Application Manager (ITCAM) for Transactions product can keep a watch over the entire end-to-end transaction path that touches many of these components. When ITCAM for Transactions detects a transaction execution problem, it can isolate the problem to individual components in the end-to-end transaction path. It can then provide a launch point for deep-dive investigation into the components.

For any transaction problems in the DB2 database component, ITCAM for Transactions can launch the Extended Insight dashboard in InfoSphere Optim Performance Manager Extended Edition in the context of the problematic database transactions. This capability enables you to use the deep database insights provided by Optim to further isolate the problem and drive it swiftly to resolution. Furthermore, Tivoli monitoring provides deeper, more extensive operating system, network, and storage information that you can access from within the system dashboard of InfoSphere Optim Performance Manager.

Detailed descriptions of how to install and configure the required components, and several usage scenarios, are available in Chapter 10, "Integration with Tivoli monitoring components" in *IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows*, SG24-7925.

The information and steps are not changed in this version of the InfoSphere Optim Performance Manager, so they are not repeated here.

Another good reference is the five-part IBM video series, "OPM-ITCAM Integration Overview" on Channel DB2:

- Part 1: An overview that introduces the capabilities and relationships between both products: IBM Tivoli Composite Application Manager (ITCAM) and IBM InfoSphere Optim Performance Manager Extended Edition;

  `http://www.channeldb2.com/video/opmitcam-integration-4`

- Part 2: Shows how high level database information is displayed in the Tivoli Enterprise Portal and the ITCAM for Transaction Topology window:

  `http://www.channeldb2.com/video/opmitcam-integration`

- Part 3: Drills into the detailed InfoSphere Optim Performance Manager database information from within the context of the Tivoli Enterprise Portal:

  `http://www.channeldb2.com/video/opmitcam-integration-1`

- Part 4: Shows you how to get detailed system information from Tivoli by launching from InfoSphere Optim Performance Manager within the Tivoli Enterprise Portal:

  `http://www.channeldb2.com/video/opmitcam-integration-2`

- Part 5: Shows you how to drill up into the end-to-end transaction context provided by ITCAM for Transactions by launching from InfoSphere Optim Performance Manager within the Tivoli Enterprise Portal:

  `http://www.channeldb2.com/video/opmitcam-integration-3`

9.2.2 Integrate InfoSphere Optim Performance Manager alerts to Tivoli event manager with SNMP

Although Tivoli ITCAM for Transactions integration is built into the Extended Edition, you might also want to send performance alerts to other monitoring solutions. You can extend InfoSphere Optim Performance Manager with a user exit to integrate DB2 database monitoring with other enterprise monitoring systems such as IBM Tivoli Netcool/OMNIbus. You can use the Simple Network Management Protocol (SNMP) to communicate database alert information to an enterprise system management software such as IBM Tivoli Netcool®. SNMP is a popular standard protocol that allows system management software to request and receive information from network devices, servers, and software.

The article, “Use Optim Performance Manager with Simple Network Management Protocol”, located at the following address, is about the integration and setup required to send SNMP traps from InfoSphere Optim Performance Manager to Tivoli’s event manager with a user exit.

http://www.ibm.com/developerworks/data/library/techarticle/dm-1007optimwithsnmp/index.html

9.3 Monitoring Cognos environments

The Extended Insight feature of InfoSphere Optim Performance Manager recognizes transactions executed from reports that are run on the Cognos Business Intelligence Server and can group them by Cognos users, report packages, and report servers. This helps to identify where the problematic SQL queries are executed, which Cognos user executed them, and which report package they are part of. With this information, the Cognos user, if needed, can improve his report performance by further tuning the queries.

9.3.1 Setting up for Cognos monitoring

Before you set up Extended Insight on a Cognos Business Intelligence server machine, ensure that the InfoSphere Optim Performance Manager Extended Edition license is activated and the Extended Insight communication ports are configured. See 3.1.2, “Activating the InfoSphere Optim Performance Manager license” on page 83 for steps on activating the license and specifying the Extended Insight ports.

Configure the databases that your Cognos report packages are using. Because Cognos is a business intelligence (BI) software, select any of the BI production templates as shown inFigure 9-1.

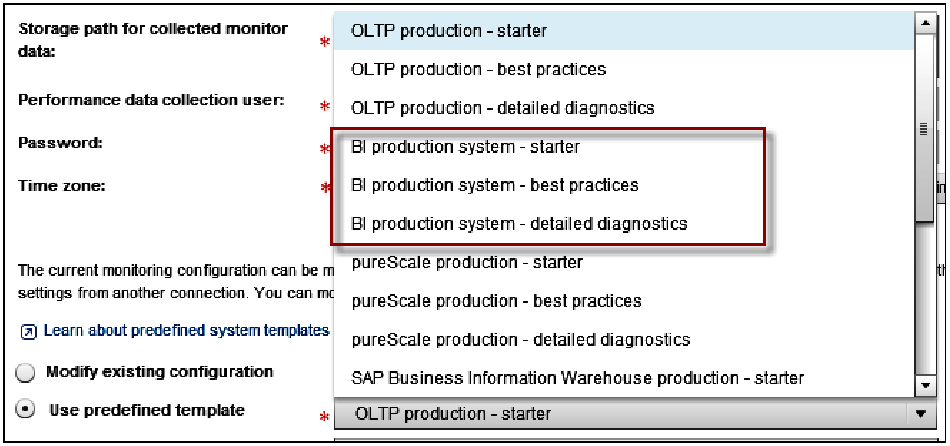

Figure 9-1 Business intelligence (BI) templates

9.3.2 Setting up the Cognos BI Server for Extended Insight

After you configure the databases, you can set up the Extended Insight client or Data Tools Runtime Client on the same machine as your Cognos BI server. Your Cognos BI server must be at version 8.4.1 or later. It should be running in compatibility or Classic Stack later.

To ensure that the workload generated by Cognos reports is correctly identified by InfoSphere Optim Performance Manager, add one new entry in the Cognos CQEConfig.xml configuration file. This file is in your <Cognos installation dir>/configuration folder. After Cognos installation, there is only a sample file CQEConfig.xml.sample that you must rename to CQEConfig.xml and open in a text editor.

Figure 9-2 shows the section that must be updated.

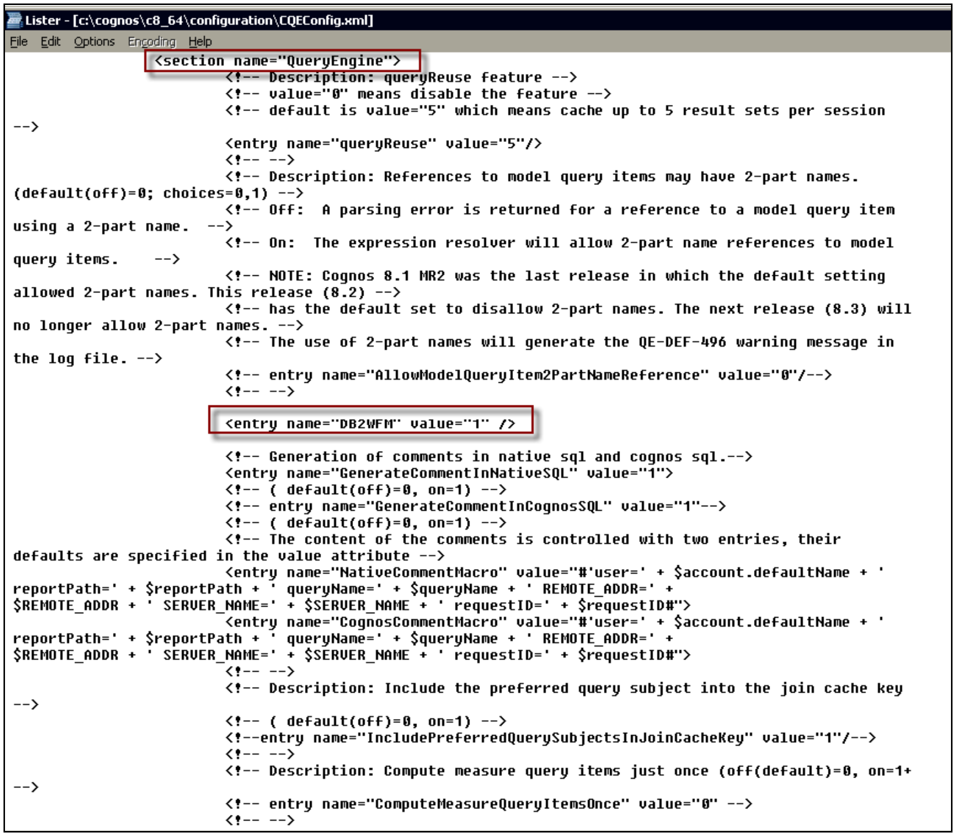


```
<section name="QueryEngine">
        <!-- Description: queryReuse feature -->
        <!-- value="0" means disable the feature -->
        <!-- default is value="5" which means cache up to 5 result sets per session
-->
        <entry name="queryReuse" value="5"/>
        <!-- -->
        <!-- Description: References to model query items may have 2-part names.
(default(off)=0; choices=0,1) -->
        <!-- Off:  A parsing error is returned for a reference to a model query item
using a 2-part name.  -->
        <!-- On:  The expression resolver will allow 2-part name references to model
query items.    -->
        <!-- NOTE: Cognos 8.1 MR2 was the last release in which the default setting
allowed 2-part names. This release (8.2) -->
        <!-- has the default set to disallow 2-part names. The next release (8.3) will
no longer allow 2-part names. -->
        <!-- The use of 2-part names will generate the QE-DEF-496 warning message in
the log file. -->
        <!-- entry name="AllowModelQueryItem2PartNameReference" value="0"/-->
        <!-- -->

        <entry name="DB2WFM" value="1" />

        <!-- Generation of comments in native sql and cognos sql.-->
        <entry name="GenerateCommentInNativeSQL" value="1">
        <!-- ( default(off)=0, on=1) -->
        <!-- entry name="GenerateCommentInCognosSQL" value="1"-->
        <!-- ( default(off)=0, on=1) -->
        <!-- The content of the comments is controlled with two entries, their
defaults are specified in the value attribute -->
        <entry name="NativeCommentMacro" value="#'user=' + $account.defaultName + '
reportPath=' + $reportPath + ' queryName=' + $queryName + ' REMOTE_ADDR=' +
$REMOTE_ADDR + ' SERVER_NAME=' + $SERVER_NAME + ' requestID=' + $requestID#">
        <entry name="CognosCommentMacro" value="#'user=' + $account.defaultName + '
reportPath=' + $reportPath + ' queryName=' + $queryName + ' REMOTE_ADDR=' +
$REMOTE_ADDR + ' SERVER_NAME=' + $SERVER_NAME + ' requestID=' + $requestID#">
        <!-- -->
        <!-- Description: Include the preferred query subject into the join cache key
-->
        <!-- ( default(off)=0, on=1) -->
        <!--entry name="IncludePreferredQuerySubjectsInJoinCacheKey" value="1"/-->
        <!-- -->
        <!-- Description: Compute measure query items just once (off(default)=0, on=1+
-->
        <!-- entry name="ComputeMeasureQueryItemsOnce" value="0" -->
        <!-- -->
```

Figure 9-2 Section in CQEConfig.xml file to be updated

Make changes to the Query Engine section of the `CQEConfig.xml` file, and restart:

1. Add the `<entry name="DB2WFM" value="1" />` entry inside the QueryEngine section.

   The `<entry name="DB2WFM" value="1" />` entry is responsible for flowing Cognos context information to InfoSphere Optim Performance Manager. Without this setting, Cognos workload is still monitored but it will not be recognized as a COGNOS application type, and predefined workload cluster groups for Cognos will not work.

2. Uncomment the following entries:

```
<entry name="GenerateCommentInNativeSQL" value="1">
<entry name="GenerateCommentInCognosSQL" value="1">
<entry name="NativeCommentMacro" value="#'user=' +
$account.defaultName + ' reportPath=' + $reportPath + ' queryName='
```

```
+ $queryName + ' REMOTE_ADDR=' + $REMOTE_ADDR + ' SERVER_NAME=' +
$SERVER_NAME + ' requestID=' + $requestID#">
<entry name="CognosCommentMacro" value="#'user=' +
$account.defaultName + ' reportPath=' + $reportPath + ' queryName='
+ $queryName + ' REMOTE_ADDR=' + $REMOTE_ADDR + ' SERVER_NAME=' +
$SERVER_NAME + ' requestID=' + $requestID#">
```

3. Save the `CQEConfig.xml` file
4. Restart the Cognos services so the changes take effect.

9.3.3 Setting up the Data Tools Runtime Client for Cognos

The Data Tools Runtime Client should be installed on the same machine as the Cognos BI server. Follow the steps in 3.3, “Installing and configuring Extended Insight Client” on page 127 for installing the product. Then launch the configuration tool.

In the Configuration Tool wizard, configure the client as an InfoSphere Optim Performance Manager client, and the clients applications as Non-Java applications as shown in Figure 9-3. Cognos BI Server uses the CLI driver connectivity in its reports.

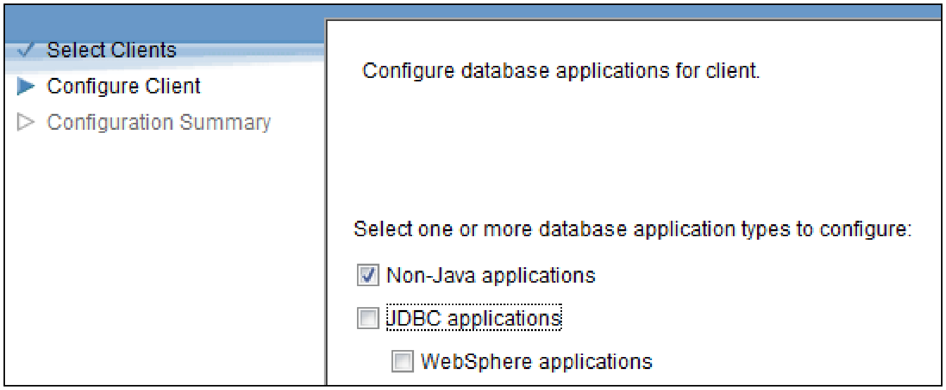

Figure 9-3 Configure Cognos application as a Non-Java application

The Configuration Tool also searches the `db2dsdriver.cfg` file, which is the configuration file for DB2 CLI driver, as shown in Figure 9-4. If it cannot be found, provide its path manually.

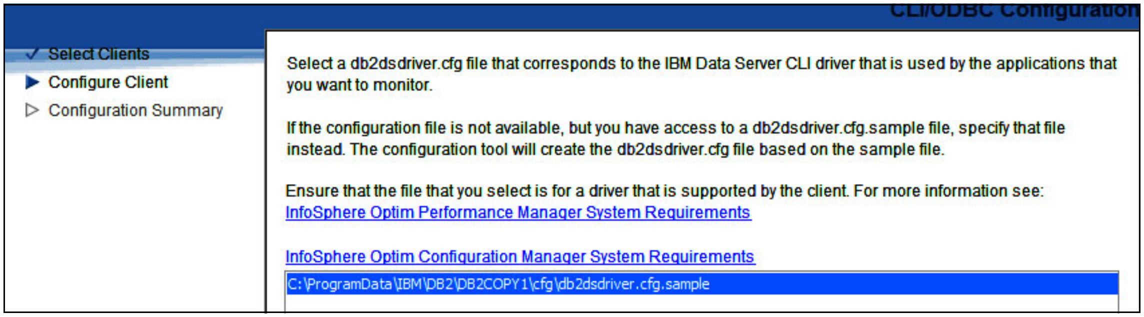

Figure 9-4 DSDriver.cfg file

The tool also looks for the appropriate CLI driver, as shown in Figure 9-5. The tool selects this driver based on your previous selection of the client type, that is, non-Java applications.

Figure 9-5 Selected CLI driver

After installation and configuration of Extended Insight on your Cognos server, restart Cognos service so that the DB2 driver configuration changes in the `db2dsdriver.cfg` file can take effect.

9.3.4 Viewing the Extended Insight monitoring data for Cognos reports

Cognos applications are listed in the Extended Insight dashboard under the default Application Types. To recognize these applications better, *activate* the Cognos workload cluster groups as shown in Figure 9-6.

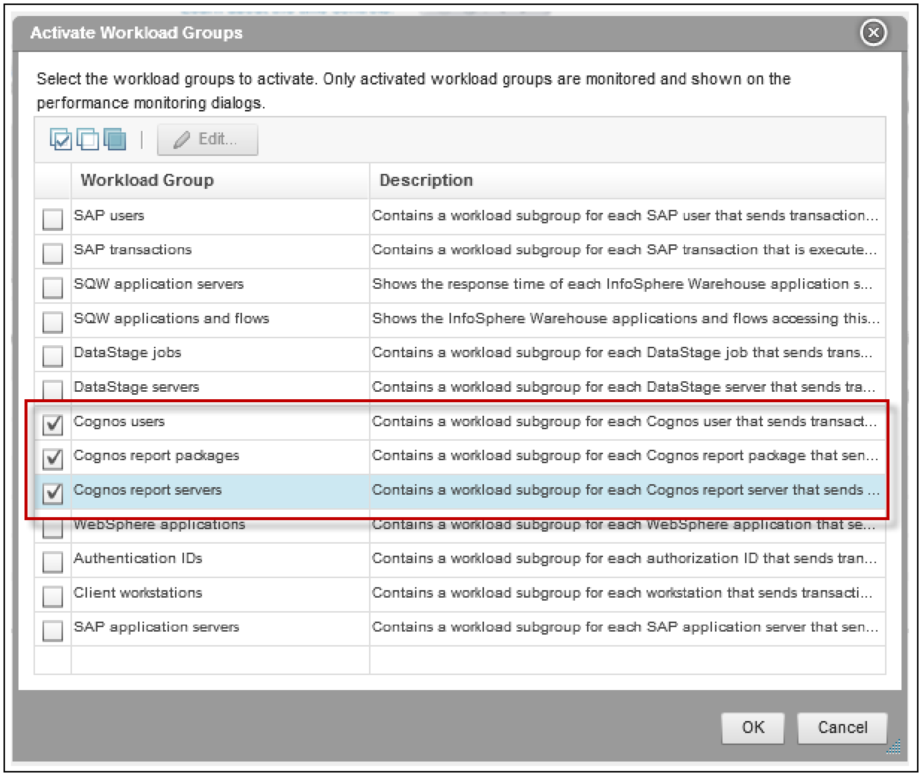

Figure 9-6 Cognos Workload Groups

When activated, the Extended Insight dashboard will listed the Cognos applications under the Cognos heading.

To monitor end-to-end response time across individual reports, you can set up a custom workload cluster group that will be filtered by the report details. In the Extended Insight overview dashboard, click **New**. The New Workload Group wizard starts. Provide the workload group name in the first page.

On the next page, select **Type of Workload Group** → **Cognos**. The clustering attributes are refreshed. Select the **Report details** connection attribute, (Figure 9-7), because this is the clustering property that contains the Cognos report name.

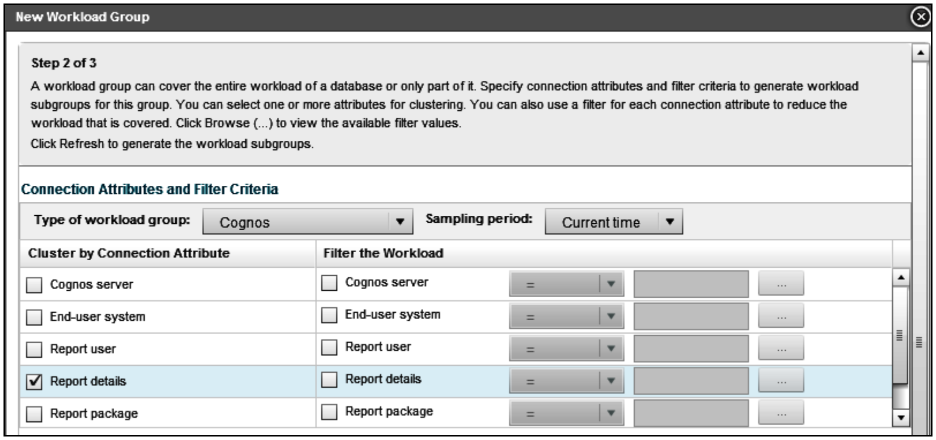

Figure 9-7 New workload group, cluster by Report Details

As 4.5, “Extended Insight Analysis Dashboard” on page 203 describes, you can view details about each application and view information such as the components of the end-to-end response time, and the top SQL statements executed by the application. You have the option to tune the SQL statement through InfoSphere Optim Performance Manager integration with Optim Query Workload Tuner.

9.4 Monitoring WebSphere Application Server applications

This section has an example of how you might use InfoSphere Optim Performance Manager Extended Insight to monitor applications running on WebSphere Application Server is shown. Also included are references to additional materials with more information.

9.4.1 Configure WebSphere Application Server for application monitoring

If your monitored application runs under WebSphere Application Server, the Data Tools Runtime Client (DTRC) must be installed at each application server, see 3.3, "Installing and configuring Extended Insight Client" on page 127. If you use WebSphere Application Server Network Deployment and are using the WebSphere Deployment Manager (DMGR) to administer your application servers, you can install the DTRC just once at the DMGR node and the appropriate JAR files will be propagated to the application servers in the group.

9.4.2 Monitor end-to-end performance for a WebSphere application

See section 8.2 in *IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows*, SG24-7925 for examples of Extended Insight monitoring for WebSphere applications.

In this section, however, we describe several applications and application environments, and other examples.

Your application environment might have many WebSphere Application Servers running the same application against a single database. Whether you are using the WebSphere Application Server Network Deployment or some other configuration, the Extended Insight feature can help with problem determination and analysis.

In the lab environment for this book, there were two stand-alone servers (not ND), both with the same WebSphere Application Server version, and both otherwise similar. One server was in San Jose, CA; another was in Raleigh, NC.

A shopping application and the database, GSDB, was on DB2 V10. Figure 9-8 on page 371 shows the overall end-to-end response times for the WebSphere applications. In this example, various performance thresholds were set for this group, and, as you can see, some are showing warning levels.

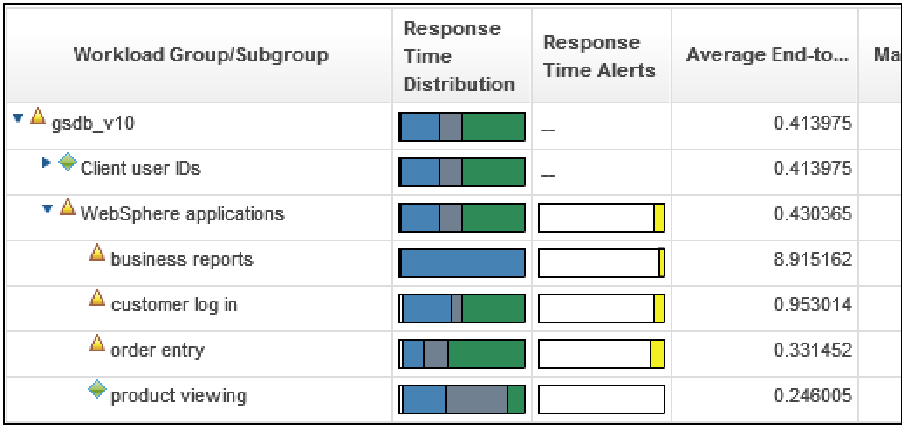

Workload Group/Subgroup	Response Time Distribution	Response Time Alerts	Average End-to...	Ma
gsdb_v10		–	0.413975	
Client user IDs		–	0.413975	
WebSphere applications			0.430365	
business reports			8.915162	
customer log in			0.953014	
order entry			0.331452	
product viewing			0.246005	

Figure 9-8 Extended Insight Summary, WebSphere applications

One way to analyze the information is to drill down to the individual application, and see the various statements and client information. However, for this example, two custom Workload Groups were set up. The two groups are defined to show transactions from each of the two servers separately. Perhaps you can imagine this in your own organization, where you might have multiple regional server farms running applications for the same database, and you can easily set up various Extended Insight Workload Groups to track the regions, for example.

Note: For more information and detailed example of creating Workload Groups (formerly called Workload Cluster Groups), read about "Understanding workload clusters" in *Optim Performance Manager for DB2 for Linux, UNIX and Windows*, SG24-7925.

The "San Jose apps" workload group setup is shown in Figure 9-9 on page 372. The workload group type was set as WebSphere, and clustered (grouped) on the WebSphere Application Server Application Type. Also included was a filter where only the server in San Jose was selected. A group for the Raleigh server was created using the same steps. There were no threshold values set for these groups in this example, but you could easily set them and even have different thresholds for different servers.

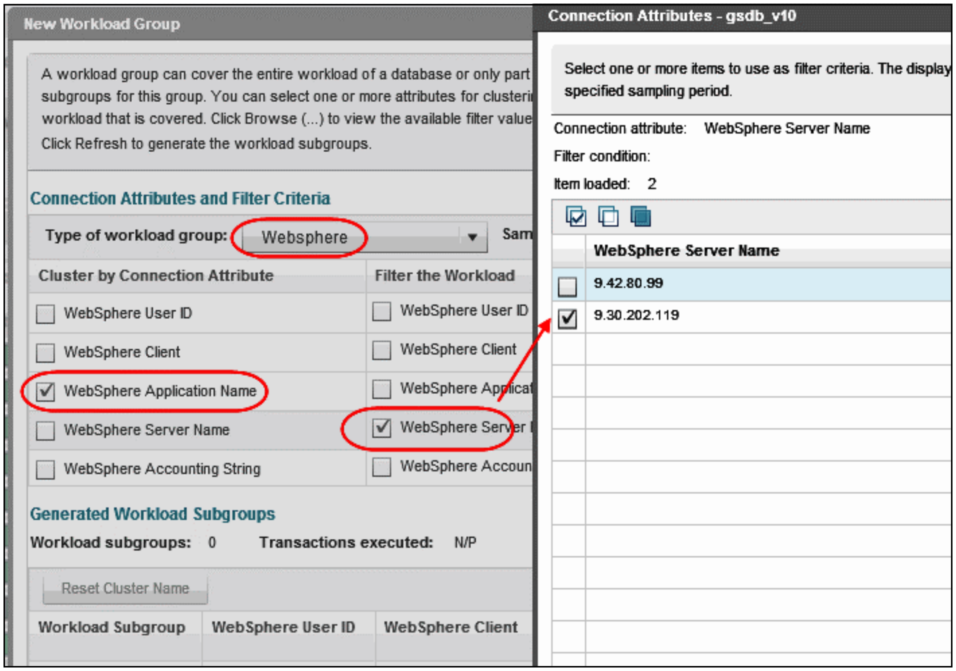

Figure 9-9 Creating custom Workload Group for San Jose server

The resulting Extended Insight page now shows both the San Jose and the Raleigh groups, with the corresponding applications that are running on each server, as shown in Figure 9-10.

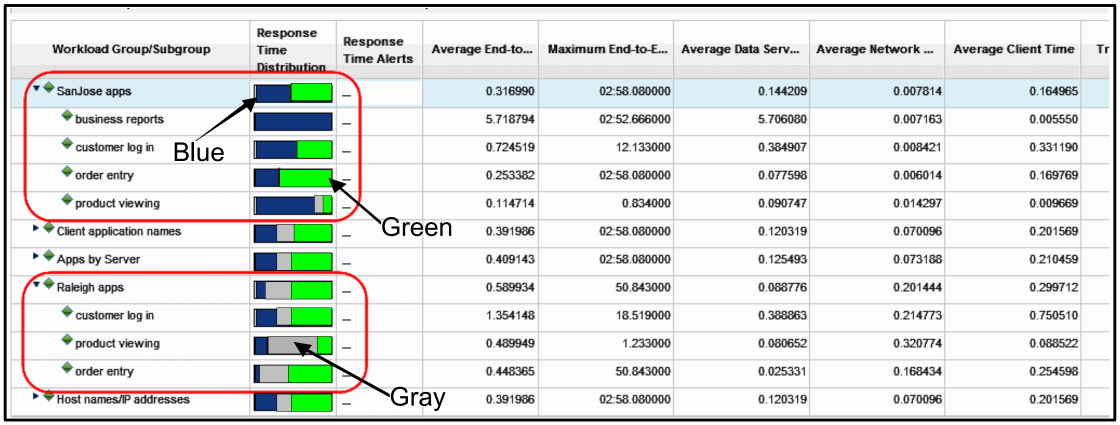

Workload Group/Subgroup	Response Time Distribution	Response Time Alerts	Average End-to...	Maximum End-to-E...	Average Data Serv...	Average Network ...	Average Client Time	Tr
SanJose apps		–	0.316990	02:58.080000	0.144209	0.007814	0.164965	
business reports		–	5.718794	02:52.666000	5.706080	0.007163	0.005550	
customer log in		–	0.724519	12.133000	0.384907	0.008421	0.331190	
order entry		–	0.253382	02:58.080000	0.077598	0.006014	0.169769	
product viewing		–	0.114714	0.834000	0.090747	0.014297	0.009669	
Client application names		–	0.391986	02:58.080000	0.120319	0.070096	0.201569	
Apps by Server		–	0.409143	02:58.080000	0.125493	0.073188	0.210459	
Raleigh apps		–	0.589934	50.843000	0.088776	0.201444	0.299712	
customer log in		–	1.354148	18.519000	0.388863	0.214773	0.750510	
product viewing		–	0.489949	1.233000	0.080652	0.320774	0.088522	
order entry		–	0.448365	50.843000	0.025331	0.168434	0.254598	
Host names/IP addresses		–	0.391986	02:58.080000	0.120319	0.070096	0.201569	

Figure 9-10 Extended Insight, Regional app servers

Various observations can be made about the overview page:

- The San Jose server has one more application running than Raleigh: the Business Reports app. This is not unexpected, because the main office for the sample company is in San Jose, where the data research people are located. Their traffic is being routed to the San Jose server only. If Business Reports transactions were showing up in the Raleigh group, you might want to investigate because a configuration issue might be incorrectly routing the transactions.
- The Raleigh server shows consistently longer network times (the gray part of the colored bar chart) than San Jose. Again this might be somewhat expected knowing that the application database, GSDB, is located in San Jose, physically near the San Jose WebSphere Application Server, while traffic from the Raleigh server has much longer distance to travel.
- Maybe network issues need to be investigated, or perhaps you might need to tune an application, if the application is too chatty. Neither of these might be what a DBA is expected to resolve, but by using InfoSphere Optim Performance Manager with Extended Insight, the DBA can help direct the problem to the correct support team.
- Without even drilling down into details, you can get benefits from the Extended Insight summary view when using custom Workload Groups.
- Another observation about the applications is that the Business Reports group has significantly longer response times than the shopping applications. Again, this is not unexpected because the application is doing reporting, rather than OLTP-type transactions. However, because both currently share the same database for shoppers and report-running, the best approach is to keep the reports running as smoothly as possible.

A

Performance management differences: DB2 V9.7 or later and DB2 V9.5 databases

The previous chapters describe functions and usage scenarios of InfoSphere Optim Performance Manager and Optim Workload Query Tuner for monitoring and tuning DB2 V9.7 or later databases.

This appendix describes the architectural and functional differences if you use InfoSphere Optim Performance Manager or InfoSphere Optim Workload Query Tuner to manage performance of DB2 V9.5 databases.

A.1 InfoSphere Optim Performance Manager differences for monitoring DB2 V9.7 or later and DB2 V9.5 databases

This appendix provides an overview of the new features of InfoSphere Optim Performance Manager V5 for monitoring DB2 V9.5 database and the available user interfaces for monitoring DB2 V9.5 databases. It describes the architectural and functional differences for monitoring DB2 V9.7 or later an DB2 V9.5 databases.

A.1.1 Overview of monitoring DB2 V9.5 databases with InfoSphere Optim Performance Manager

Although the focus of InfoSphere Optim Performance Manager V5 is on delivery of new functions and architectural enhancements for monitoring DB2 V9.7 and DB2 10 databases, it offers substantial value to DB2 V9.5 deployments. In addition to the monitoring features that were already available in InfoSphere Optim Performance Manager V4 there are a few key features added in various InfoSphere Optim Performance Manager V5 deliveries that also support monitored databases of DB2 V9.5:

- Inflight dashboard enhancements

  The connection dashboard is completely revised and, for DPF databases, partition-level details are uniformly available across all inflight dashboards.

- Reporting enhancements

  The new Performance Overview and Table usage reports are also available for DB2 V9.5 databases. Command-line reporting and the new report scheduling and report retention features can also be used for DB2 V9.5 databases.

- Alerting enhancements

  Various predefined alerts are introduced that are also checked for DB2 V9.5 databases, for example lock escalations per minute. In addition, user-defined alerts, alert actions, and blackout events can be defined for DB2 V9.5 databases.

From an architectural perspective, InfoSphere Optim Performance Manager V5 uses the same monitoring infrastructure as InfoSphere Optim Performance Manager V4 to collect monitoring data from DB2 V9.5 databases. See *IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows*, SG24-7925 for an overview of the architecture and collection method. It also describes a set of user scenarios to monitor for example I/O utilization, CPU and memory usage or to analyze locking problems. These scenarios provide valuable information how to use InfoSphere Optim Performance Manager for monitoring DB2 V9.5 databases.

Performance Expert Client is available as an additional user interface to the InfoSphere Optim Performance Manager web interface to perform performance monitoring tasks for DB2 V9.5 databases. Performance Expert Client is the original client user interface of the previous DB2 Performance Expert product and is still available in the InfoSphere Optim Performance Manager product.

Note: You cannot use Performance Expert Client for DB2 V9.7 or later databases using the in-memory metrics collection method because of the architectural differences of in-memory metrics and snapshot collection.

Performance Expert Client provides a few useful monitoring functions that are not available in the InfoSphere Optim Performance Manager web interface for DB2 V9.5 databases:

- Real-time monitoring

  Most panels on Performance Expert Client allow you to click a refresh button to get and display actual snapshot data on the panel.

- Display of top dynamic SQL statements

  This panel lists dynamic SQL statements collected from the dynamic SQL snapshot. You can sort the statements by various performance metrics to get the top statements, or you can filter them or launch Visual Explain for single statements. The InfoSphere Optim Performance Manager web interface provides the Dynamic SQL report for DB2 V9.5 database, which is based on the same data collected from the dynamic SQL snapshot. But because this report is limited in the number of statements it displays and also limited in filter capabilities and ad-hoc sorting it is valuable to use Performance Expert Client additionally to analyze dynamic SQL statements

- Long term performance analysis through Performance Warehouse

  For DB2 V9.5 databases, InfoSphere Optim Performance Manager aggregates collected data into separate long term tables. No retention period exists for data in these tables. The data can be accessed through the Performance Warehouse in Performance Expert Client through queries or reporting functions and provide useful information for long-term trend analysis or capacity planning.

- Display of operating system metrics based on CIM integration

  If you set up a CIM server on the monitored system, InfoSphere Optim Performance Manager can collect operating system data from the monitored system such as CPU utilization, memory usage, file system and process information. This information is displayed on dedicated panels on Performance Expert Client.

A.1.2 Architectural differences

The main architectural differences between monitoring DB2 V9.7 (or later) and DB2 V9.5 database are in the method to collect data and how the collected data is stored.

For monitoring DB2 V9.7 or DB2 10 databases, InfoSphere Optim Performance Manager mainly collects the lightweight DB2 in-memory metrics from the monitored databases in periodic intervals by establishing a connection and executing SQL statements that call the appropriate DB2 in-memory metrics monitoring user-defined functions (UDFs) such as MON_GET_CONNECTIONS. The collected data is stored in the repository database in various fact and dimension tables. The data model behind these tables follows a star schema and therefore avoids data redundancy, for example a buffer pool name or a SQL statement text is only stored once in a dimension table and the fact tables store only the periodically collected performance metrics for that buffer pool or SQL statement. Each fact table exists four times, one table for each aggregation level, and the aggregation level number is appended to the table name to distinguish the tables. The InfoSphere Optim Performance Manager repository server in background aggregates the collected data into the next higher aggregation level and stores the aggregated data in the table for the next higher aggregation level. Additionally it also deletes the data from these tables after the retention time is reached.

For monitoring DB2 V9.5 databases, InfoSphere Optim Performance Manager mainly collects DB2 snapshot data from the monitored databases in periodic intervals by attaching to the DB2 instance, retrieving snapshot data using DB2 APIs and detaching. The collected data is stored in the repository database in various short-term history tables. After the retention time is reached, InfoSphere Optim Performance Manager deletes the data from these tables. The InfoSphere Optim Performance Manager repository server in the background aggregates the collected data into a set of long term history tables having the same table name but residing in a different schema. No automatic deletion takes places from these tables.

In addition, InfoSphere Optim Performance Manager optionally collects data from event monitors if enabled in the monitoring configuration. For monitoring DB2 V9.7 (or later) databases, the lock and statistic event monitor are created and, if Extended Insight is configured, also the package cache and transaction event monitors are created. For monitoring DB2 V9.5 databases, the deadlock and statistic event monitors are created. If you use Performance Expert Client, you can also start statement or activity event monitors on demand from the Performance Warehouse function.

A.1.3 Inflight dashboard differences

The main dashboard differences for DB2 V9.7 (or later) and DB2 V9.5 databases exist in the number of available views on dashboards, in the time controls for browsing through collected data, and in the set of performance metrics that are displayed on the dashboards.

The following inflight dashboard views are not available if you monitor DB2 V9.5 databases:

- At-a-Glance View on the Overview Dashboard including Baseline settings
- SQL Summary view on the SQL dashboard to analyze top dynamic and static statements from package cache

  For DB2 V9.5, the SQL Dashboard is called Active SQL dashboard. The Active SQL dashboard provides only one view showing the top executing statements that were running at the time when InfoSphere Optim Performance Manager collected the snapshot data. The purpose of the Active SQL dashboard is the same as the Top Individual Execution view available on the SQL Dashboard for DB2 V9.7 (or later) databases.

  As an alternative to the missing SQL Summary view either use the Dynamic SQL report from the InfoSphere Optim Performance Manager web interface or use the Dynamic SQL Statements panel from Performance Expert Client. Both have the limitation that only dynamic statements are displayed.

The following time controls are not available on inflight dashboards if you monitor DB2 V9.5 databases

- Real-time mode

  Only historical data that is already collected is displayed for DB2 V9.5 databases. For real-time monitoring, use Performance Expert Client.
- Display of aggregated data with aggregation level indicator

  Only short-term history data is displayed for DB2 V9.5 databases that have a retention time between multiple days and one or two weeks. For displaying aggregated long-term history data, use Performance Warehouse of Performance Expert Client

The differences in the set of performance metrics that are displayed on the dashboards for DB2 V9.7 (or later) and DB2 V9.5 databases mainly consist of a new set of time-spent monitor elements including wait and component times that were introduced in the DB2 in-memory metrics for DB2 V9.7 or later. These time-spent monitor elements are mainly exposed on the Connection dashboard, SQL Dashboard, and the At-a-Glance View of the Overview Dashboard.

A.1.4 Reporting differences

The main reporting differences for DB2 V9.7 (or later) and DB2 V9.5 databases exist in the number of available reports and in the custom reporting feature. The custom reporting feature to extract collected data from the repository database using predefined UDFs is available only for DB2 V9.7 (or later) databases that use the in-memory metrics collection method.

The following predefined reports are available for DB2 V9.5 databases:

- Database Configuration
- Database Manager Configuration
- Database Connection
- Performance Overview
- Disk Space Consumption
- Table Usage
- Dynamic SQL statement
- Workload Manager Overview
- Workload Manager Service Subclass
- Workload Manager Service Superclass
- Workload Manager Work Action Set
- Workload Manager Workload

For DB2 V9.7 (or later) databases, the following reports are also available:

- SQL Baseline Comparison
- Storage Group report (DB2 V10)
- Top SQL Package
- Top SQL (Dynamic and Static)

  The Dynamic SQL Statement report is not available; it is replaced by the Top SQL report.

A.2 Optim Workload Query Tuner differences for DB2 V9.7 or later and DB2 V9.5 databases

Optim Query Workload Tuner supports both DB2 V9.5, and DB2 V9.7 (and later). It primarily uses the features inside the database engine, which get improved in each newer version. So some features are available only for DB2 V9.7 and later, or more information can be used if the tuned database is DB2 V9.7 and later.

Several new query optimization improvements are in DB2 v10.1 such as the zigzag join, jump scan, and others. This section describes differences of the supported features in Optim Query Workload Tuner for different database levels from three perspectives.

A.2.1 Captured source difference

Differences in captured sources are as follows:

- Capture from package cache

  InfoSphere Optim Query Workload Tuner can capture from package cache for both DB2 V9.5 and DB2 V9.7 (and later); for DB2 V9.5, only the dynamic statements can be captured with limited metrics, because DB2 V9.5 engine does not provide any interfaces to get the runtime information for static statements. The filter condition set is quite small when captured from DB2 V9.5 because the DB2 engine exposes only a small set of metrics of the dynamic statements. Furthermore, if you want to get the runtime information of the statements, you must turn on the STATEMENT monitor switch. For DB2 V9.7 (and later), DB2 engine provides both dynamic and static statements in package cache. Included is runtime information and much more metrics. This enables the user to capture both dynamic and static statements in package cache, and use more filter conditions to get the interesting queries.

- Capture from event monitor tables

  User can capture from activity event monitor tables for statements to tune for both DB2 V9.5 and V9.7 (or later). For V9.5, there is no explain information of the captured statements, so you must explain again using the current statistics information to get the explain information; for V9.7 and later, the explained information can be retrieved directly during capture.

- Capture from InfoSphere Optim Performance Manager performance repository

  This is supported only for DB2 V9.7 and later. If your target database is DB2 V9.5, you do not see such an option in the capture source list in the product.

A.2.2 Query analysis options difference

A feature in Optim Query Workload Tuner is to use the section actuals information to compare the actual returned rows and estimated returned rows in access plan. This is available only for DB2 V9.7 and later. When tuning a SQL against DB2 V9.7 and later, in the Run Advisors and Analysis Tools page, EXPLAIN options section, you can see the option (Figure A-1 on page 382). For

DB2 V9.5, this feature is unavailable because it uses the section actual feature in DB2 V9.7.

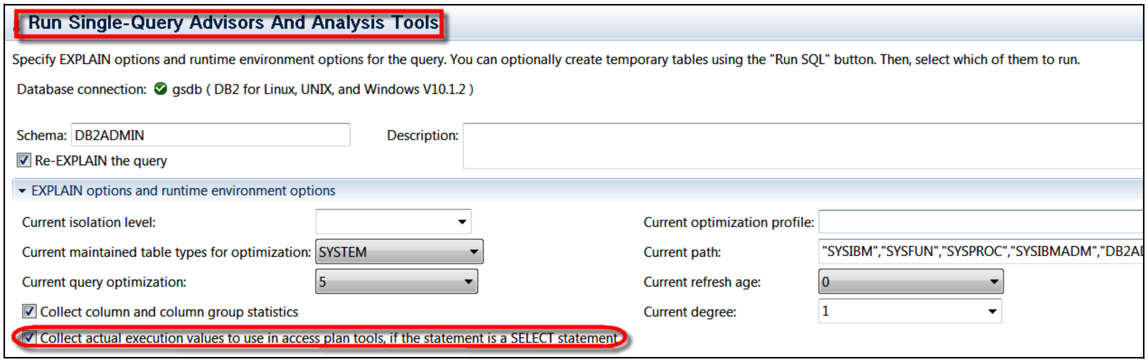

Figure A-1 Leverage section actuals

For details regarding DB2 section features, see the information center:

http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.admin.perf.doc/doc/c0056362.html

A.2.3 Access plan and tools difference

The plans and tools differences are as follows:

- Edit Optimization Profile for plan hint

  This feature is used to create or edit a query optimization profile for the plan hint. It is available only for database that is DB2 V9.7 fix pack 2 and later. The details of this feature is described in Chapter 5, “Getting to know InfoSphere Optim Query Workload Tuner” on page 221.

- Workload access plan lock down

  The workload plan lock down feature can help generate the optimization profile to lock down the access plan of the interested queries. It is available only for DB2 V9.7 fix pack 7 and follow-on fix packs, and DB2 V10.1 fix pack 2 and follow-on fix packs. The details of this feature are described in Chapter 5, “Getting to know InfoSphere Optim Query Workload Tuner” on page 221.

- New access plan operation type in DB2 v10.1

  The new SQL features in DB2 V10.1, such as zigzag join, index jump scan, and others, are also supported in Optim Query Workload Tuner. In Access Plan Graph, Access Plan Explorer shows the new operation type, and in the post filter of workload, provides an option to filter the workload on the new operation type.

B

Additional material

This book refers to additional material that can be downloaded from the Internet as described in the following sections.

Locating the web material

The web material associated with this book is available in softcopy on the Internet from the IBM Redbooks Web server. Point your web browser at:

`ftp://www.redbooks.ibm.com/redbooks/SG248111`

Alternatively, you can go to the IBM Redbooks website at:

ibm.com/redbooks

Select the **Additional materials** and open the directory that corresponds with the IBM Redbooks form number, SG248111.

Using the web material

The additional Web material that accompanies this book includes the following files:

File name	*Description*
SG24-8111-AM.zip	This ZIP file contains a PDF with Additional Material relevant to Redbooks publication SG24-8111. There are two chapters in the PDF. Chapter 1 is titled “Reporting with InfoSphere Optim Performance Manager” (which is relevant to Chapter 4) and Chapter 2 is titled “Working with Workload Manager configurations” (which is relevant to Chapter 8).

Downloading and extracting the Web material

Create a subdirectory (folder) on your workstation, and extract the contents of the Web material .zip file into this folder.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows, SG24-7925*
- *Using Integrated Data Management To Meet Service Level Objectives*, SG24-7769

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- *Database Monitoring Guide and Reference*, SC27-2458
- *Troubleshooting and Tuning Database Performance*, SC27-2461

Online resources

These websites are also relevant as further information sources:

- "Performance management and optimization solutions for IBM DB2 for Linux, UNIX and Windows"

  http://public.dhe.ibm.com/common/ssi/ecm/en/imb14129usen/IMB14129USEN.PDF

- Implementing DB2 workload management in a data warehouse

  http://www.ibm.com/developerworks/data/bestpractices/workloadmanagement/

- Integrated Data Management Information Center:

  http://publib.boulder.ibm.com/infocenter/idm/v2r2/topic/com.ibm.dstudio.nav.doc/topics/welcome.html

- DB2 Information Center:

  http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/

- Information Management:

  http://www.ibm.com/software/data/

- DB2 developerWorks:

  http://www.ibm.com/developerworks/db2/

- IBM Tivoli Composite Application Manager for Transactions 7.2.0.2:

  http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/topic/com.ibm.itcamt.doc_7.2.0.2/ic-homepage.html

- IBM Tivoli Composite Application Manager for Application Diagnostics 7.1.0.2:

  http://publib.boulder.ibm.com/infocenter/tivihelp/v24r1/topic/com.ibm.itcamfad.doc_7101/ic-homepage.html

- IBM Tivoli Monitoring (latest version):

  http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.2fp2/welcome.htm

Help from IBM

IBM Support and downloads:

ibm.com/support

IBM Global Services:

ibm.com/services

Performance Management: Using IBM InfoSphere Optim Performance Manager and Query Workload Tuner

(0.5” spine)
0.475”<->0.875”
250 <-> 459 pages

Performance Management

Using IBM InfoSphere Optim Performance Manager and Query Workload Tuner

Prevent problems before they impact the business

Align monitoring and business objectives

Improve query workload performance

This IBM Redbooks publication describes the architecture and components of IBM InfoSphere Optim Performance Manager Extended Edition. Intended for DBAs and those involved in systems performance, it provides information for installation, configuration, and deployment. InfoSphere Optim Performance Manager delivers a new paradigm used to monitor and manage database and database application performance issues. It describes product dashboards and reports and provides scenarios for how they can be used to identify, diagnose, prevent, and resolve database performance problems.

IBM InfoSphere Optim Query Workload Tuner facilitates query and query workload analysis and provides expert recommendations for improving query and query workload performance. Use InfoSphere Optim Performance Manager to identify slow running queries, top CPU consumers, or query workloads needing performance improvements and seamlessly transfer them to InfoSphere Optim Query Workload Tuner for analysis and recommendations.This is done using query formatting annotated with relevant statistics, access plan graphical or hierarchical views, and access plan analysis. It further provides recommendations for improving query structure, statistics collection, and indexes including generated command syntax and rationale for the recommendations.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-8111-00 ISBN 0738438456