Infrastructure Solutions: SAP IS Automotive on an IBM Platform

Case study of a real world scenario

Make-to-Order manufacturing data in a complex industry

Benchmark testing in an IBM System z and System p environment

Abbas Birjandi
Eric Romain
Hervé Sabrié
Nicole Pompejus
Till Poertner
Valery Charles P Desmoulin
Hans HPJ Westheim

Redbooks

ibm.com/redbooks

International Technical Support Organization

Infrastructure Solutions: SAP IS Automotive on an IBM Platform

December 2008

SG24-7540-00

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

First Edition (December 2008)

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 give 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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.

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.

Trademarks

Preface

This IBM® Redbooks® publication describes a proof of concept project conducted by a team of IT specialists from IBM, SAP®, and an automotive industry client. The project involved developing a supply chain management implementation for complex make-to-order repetitive manufacturing for configurable products, specifically the production of a variety of automobiles destined for an international market.

The project was based on the current and projected requirements of the automotive OEM. The performance testing described here was undertaken because the volume of data to be processed surpasses that of other SAP implementations of the make-to-order scenario currently in production. The purpose of the tests was to ensure that the application could be restructured to meet the requirements of this data volume and that the IBM infrastructure would be capable of processing this load in the tight time windows specified by the customer's business processes.

The customer selected a set of SAP functional performance tests and defined run time targets for these key performance indicators. The modules tested were Rapid Planning Matrix, Materials Requirements Planning, Backflush, and Materials Ledger. Actual data was provided from the customer environment.

The test environment included IBM System z® with DB2® for the SAP central services and databases, System p® (AIX®) for the SAP application servers and liveCache, and a DS8300 Turbo with 13 TB capacity for the data. The tests were performed with mySAP™ ERP 2005 with IS Automotive activated; the workload was simulated through specific BladeCenter® injectors.

IBM and SAP performance experts ran the benchmark tests and made some configuration adjustments. The result was an implementation that met the customer's requirements. These tests will influence major customer automotive architecture decisions.

The team that wrote this book

This book was produced by a team of specialists from around the world working in the International Technical Support Organization (ITSO), at the Product Solution Support Center in Montpellier France.

Abbas Birjandi is an IBM certified professional with more than 35 years of experience in the computer industry, including almost 15 years with IBM France. He holds a Ph.D. in Computer Science. His career has included research and faculty positions at various universities, systems architecture, project leadership, and management positions. Prior to joining the ITSO Abbas was EMEA Project Manager for the Early Support Program. For the last few years he has worked on projects related to IBM educational activities.

Eric Romain is a Certified I/T Specialist leading IM/DB2 activities within the PSSC Customer Support Center in Montpellier, France. He has worked with DB2 for z/OS® since version 1, and also has 9 years of experience with SAP. Eric's areas of expertise include DB2, DB2 DataSharing and SAP; he has filled roles in education, installation, performance management, and migrations. Prior to joining IBM in 2000, he worked for a large distribution company as technical leader for DB2 on z/OS in charge of operations for production systems. As DB2 for z/OS expert, Eric is currently involved in running customer benchmarks as part of the European System z benchmark center. He has worked on projects like Dynamic Pricing Engine Application, SAP BW migration using DB2 for z/OS as Database Server and Linux® on System z as Application Servers, as well as on large banking applications.

Hervé Sabrié is an IBM certified professional with more than 18 years of experience in the computing industry. He graduated from the Toulouse University in 1988 with a DEA (postgraduate certificate) in Fundamental Physics and a master in Micro-electronics. Hervé joined IBM France and subsequently PSSC in 1994 and has held positions ranging from IT specialist to project leader. He was previously the manager of the European Design Center, leading teams of IT architects and IT Specialists running design workshops to leverage the IBM on demand strategy, and he led the European System p Benchmark Center during the 1990s. He is currently managing large and complex projects as part of the PSSC Montpellier Benchmarking Center.

Nicole Pompejus is solution manager for the Industry Solution Automotive and part of the Product and Technology Group of SAP headquarters in Walldorf, Germany. Nicole started her professional career with an SAP partner consulting company, working for several years in international projects as a consultant to the automotive industry. She joined SAP in 2002 and is responsible for Automotive Demand management and Manufacturing. Nicole holds a diploma degree from HTW Saarbrücken (Hochschule für Technik und Wirtschaft des Saarlandes) as well as a diploma from ESAM Paris (Ecole Supérieure d'Administration et de Management), equated by the French Government as BAC+ 4.

Till Poertner joined SAP AG in 1998 and is solution manager for Industry Solution Automotive. He is part of the Product and Technology Group of SAP headquarters, based in Walldorf (Germany). Till holds a Master of Science degree in Computer Science from the University of Bonn. He started his SAP career as a developer for the Automotive Industry Solution, moving up to the global Roll Out team and supporting various Automotive customer projects in implementing the SAP industry solution. He now works as solution manager and is responsible for the Automotive High-volume Make-to-Order Solution.

Valery Charles P Desmoulin is an IT Solution Architect, dedicated to the design of SAP related architectures. He has been with the SPGIT S&D organization since February 2006. Before joining S&D, he was a Senior Consultant for GBS in SPGIT, dealing with complex SAP implementation projects, and Architecture and IT Strategy projects, leading the technology related areas of the projects. He has over 12 years experience applying IT in a variety of sectors, including Banking, Telco, Public, and Industry sectors.

Hans HPJ Westheim is a Senior IT Specialist leading a team of SAP Basis Consultants, zOS System Programmers and DB2-DBAs. This team is responsible for maintaining SAP systems, and perform SAP-related projects for a number of customers who run their SAP solutions on zSeries®. Hans has worked for over 24 years in IT, and during the last 12 years of his career he has specialized in zOS/DB2 in relation with SAP. One of the projects he worked on was to build a solid, high available infrastructure for a large utilities company. This involved maintaining a large number (60+) of DB2/zOS SAP databases of various sizes (50 Gb to 4 Tb) in an efficient and cost-effective way.

A very special thanks to the following people for their contributions to this book and project:

Christoph Pertzsch, SAP AG, Germany

Elke Menninger, SAP AG, Germany

Mathias Uhlmann, SAP AG, Germany

Hans Baessler, SAP AG, Germany

Dong Liu, SAP AG, Germany

Jens Bekker, SAP AG, Germany

Ivan Sesar, SAP AG, Germany

Nick W. Scherer, SAP AG, Germany

Franck Almarcha, PSSC, IBM France

Richard Corrihons, PSSC, IBM France

Thierry Chenillot, PSSC, IBM France

Robert Nicolas, PSSC, IBM France

Gordon Watson, PSSC, IBM France

Chulhee Cho, Technical Sales Support, IBM Korea

Yougrok Kwon, Technical Sales Support, IBM Korea

Hyun Baek, Technical Sales Support, IBM Korea

Steve Bergmann , IBM SAP International Competency Centre, IBM Germany

International Technical Support Organization, Poughkeepsie Center

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

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 e-mail 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

Part 1

Project overview

In this part we provide an introduction to the overall project environment based on SAP and IBM solutions. In addition we present an in-depth look at the architecture used and a brief overview of the test results.

This part includes the following:

1

SAP and IBM solutions in the automotive environment

The SAP automotive environment is a suite of software modules used to manage *Make-in-Time* manufacturing and tailored to the automotive industry. In this chapter we describe in general terms the features of the SAP solution that support the unique aspects of automotive manufacturing. We also describe the specific circumstances of the OEM with whom SAP and IBM undertook this project, and explain how the SAP and IBM solution satisfies the requirements of this particular organization.

The topics covered are:

1.1 Introduction

IBM, SAP, and an automotive industry client joined forces in developing an integrated information system to address the ever growing complexities and demands of running an international operation covering all steps of automobile production.

The project presented in this book is a *Supply Chain Management* implementation based on the current and projected requirements of the automotive OEM. The volume of data to be processed surpasses that of other SAP implementations of the Make-to-Order scenario currently in production.

IBM and SAP therefore proposed this proof of concept and performance test to ensure that the application could be restructured to meet the requirements of this data volume and that the IBM infrastructure would be capable of processing this load in the tight time windows specified by the customer's business processes.

This book covers stress testing of selected components from the SAP IS Automotive solution, together with target results referred to as key performance indicators (KPIs). The scope of these KPIs is described fully in later sections. The required systems configuration and infrastructure was defined based on this scope.

1.2 Objectives of the PoC test

This Proof of Concept project was intended, as a benchmark, to help the customer in selecting the most appropriate platform for their future production system. The main objectives were the following:

- Provide a test environment to demonstrate how the System z platform adds value as a database server for SAP.
- Provide the client with the basis for a production platform decision by executing tests and measuring results against defined key performance indicators (KPIs).
- Demonstrate results to the client.

The targeted KPIs focus on the business and functional aspects of the scenario, and demonstrate the pure performance of the SAP modules selected. No *operational* KPIs were included because the client already is confident that the System z platform can handle the non-functional requirements such as availability, manageability, and scalability without the necessity to have them demonstrated within this zPoC.

Four KPIs were identified as the most significant indicators for evaluating the overall appropriateness and efficacy of the SAP and IBM solution. Table 1-1 lists the KPIs and their target execution times.

Table 1-1 KPIs and their expected execution targets

KPI	Target execution time
Rapid Planning Matrix (RPM)	30 minutes
Material Requirement Planning (MRP)	30 minutes
Backflush	30 minutes
Material Ledger (ML)	4 hours

IBM and SAP provided recommendations about tuning parameters based on the various tests. When KPIs were not originally met, the project team researched the inhibitors and adjusted the runs accordingly, ultimately modifying the system to achieve, or even exceed, the targets established by the client.

1.3 The SAP solution (and how it fits with the IBM solution)

During this benchmark testing the team focused on the OEM *Make-to-Order* scenario portion of the SAP automotive solution. In this section we provide an overview of the scenario and relationships among the components.

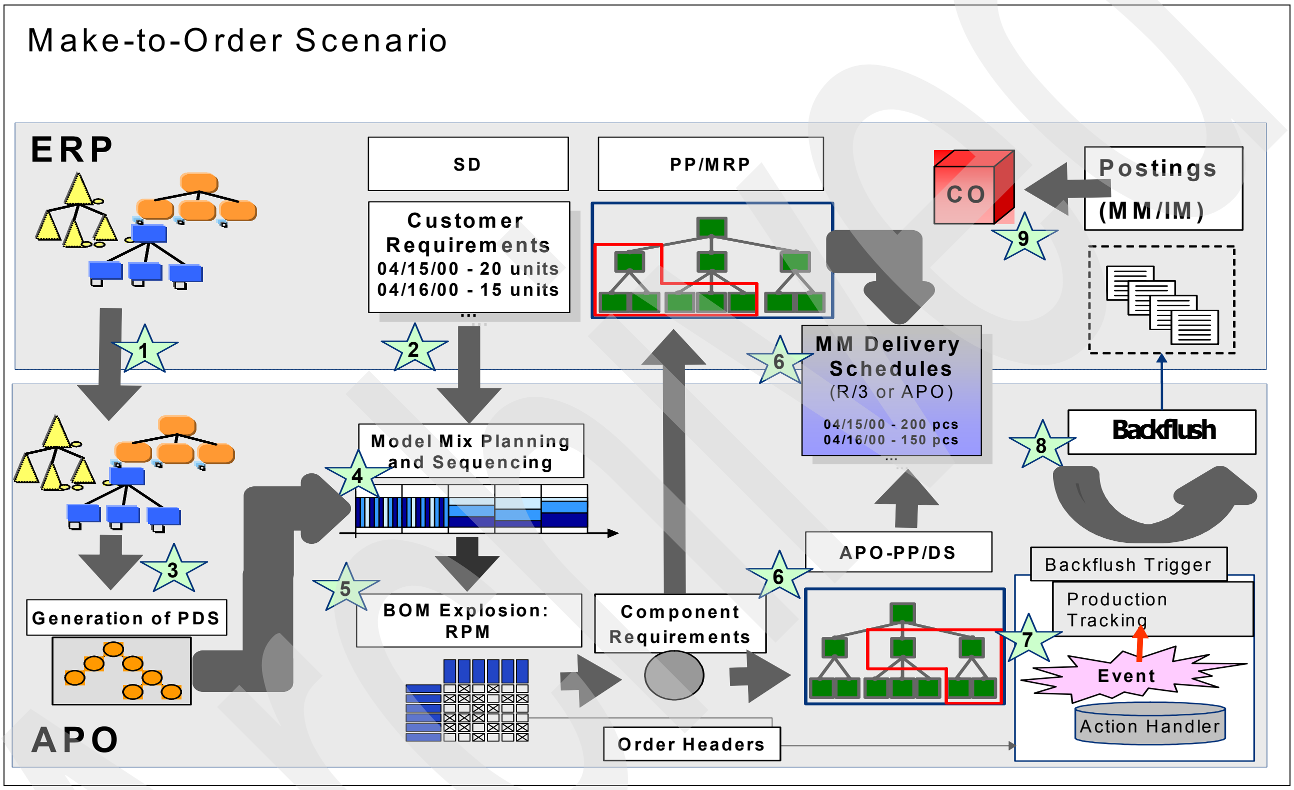

Figure 1-1 Overview of the Make-to-Order scenario

This process can be used for *make-to-stock* production and for *make-to-order* repetitive manufacturing for configurable products.

The basis of this scenario is the use of *integrated Product and Process Engineering* (*iPPE*). iPPE enables you to collect all the data for an entire product life cycle in one integrated model. It is particularly suited to highly variant products.

IPPE consists of a bill of material called Product-Variant-Structure, the Process Structure, and the Factory Layout, as well as relationships between all of these objects.

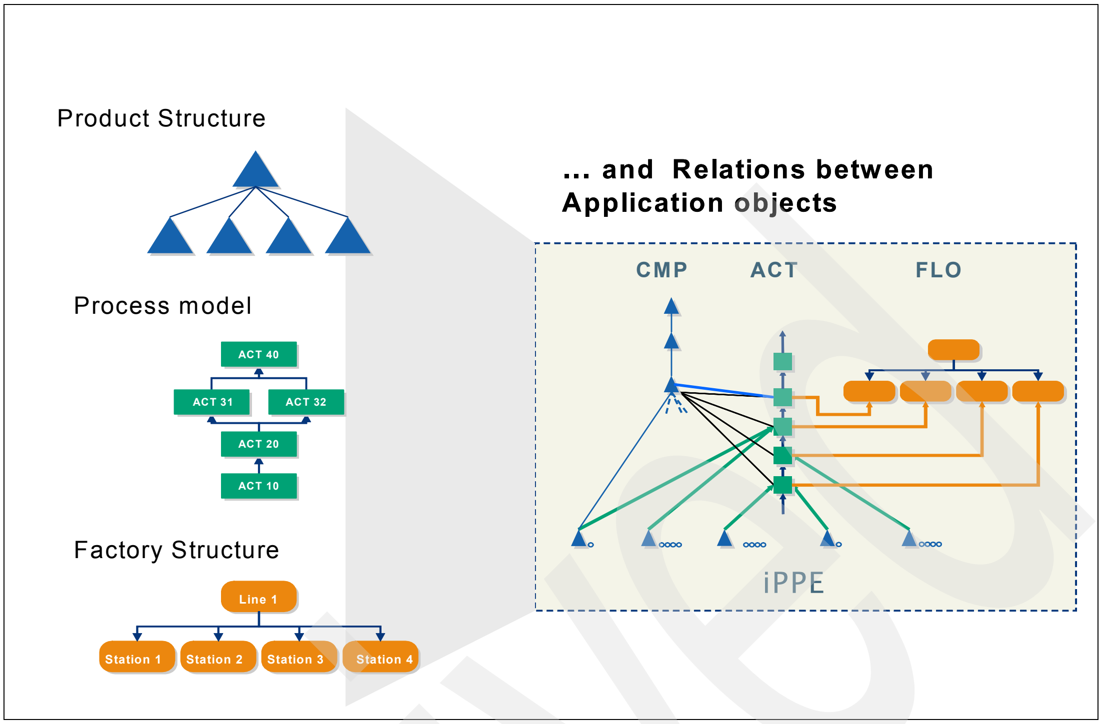

Figure 1-2 Overview of iPPE structures

Product-Variant-Structure

A distinctive characteristic of the automotive industry is that one single car model has lots of possible feature variations. Building up a specific bill of material for any possible feature combination would result in an unmanageable amount of data. Figure 1-3 demonstrates how fast the number of combinations grows for just eight features with only a few variations per feature.

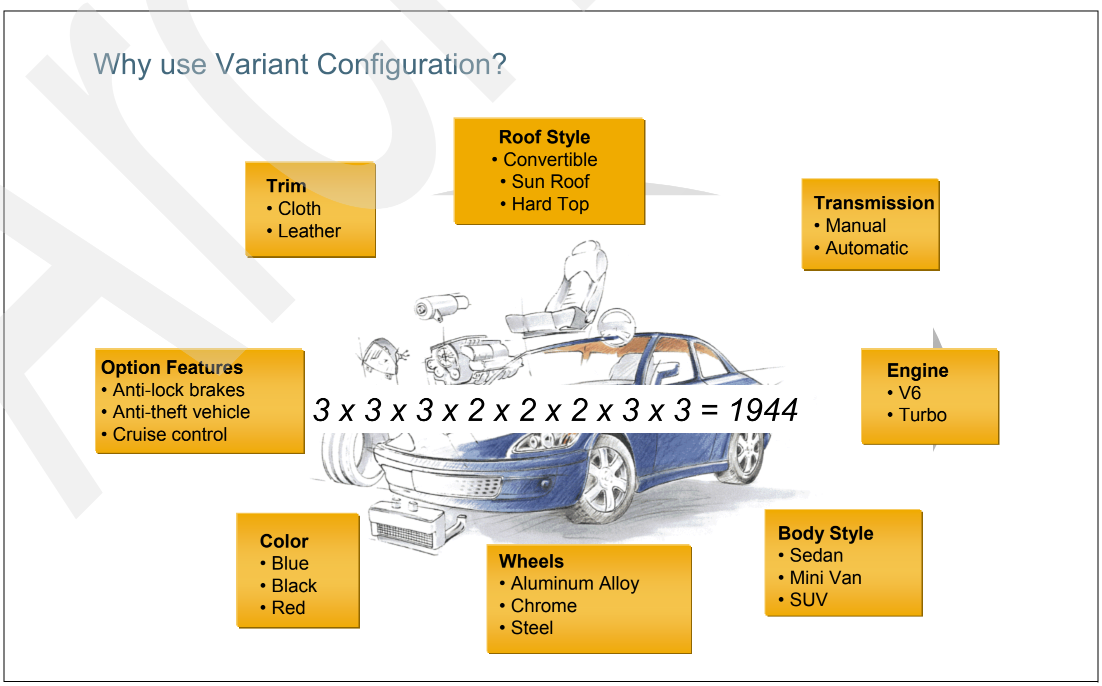

Figure 1-3 Why variant configuration is needed

The solution to this problem is the iPPE *Product-Variant-Structure (PVS)*. The PVS is a Super Bill of Material (BOM) that contains all possible components of a car in one single BOM with one header material, which can be thought of as the car's model name. A customer orders the header material, the model, and selects the features he would like to have (like engine, air conditioning, and so forth). The specific components of the customer car are then selected from the Super BOM with the help of object dependencies. You can also define so-called "Car variants," which refer to the header material and represent a few fully specified cars that, for example, would be stocked or semi-specified cars with a few very important features predefined (such as engine, shape, and country), and other features that can still be selected by the customer (such as air-conditioning, color, sports package, and so forth). This model allows us to handle a huge number of possible BOMs in one single BOM and therefore in a very intelligent and slim way.

The other components of iPPE are the *Process Structure*, where all the activities and times are contained, and the *Factory Layout*, which represents the real factory and therefore the production flow.

You can represent a complete production model consisting of BOM, routing, and line design in one single model.

iPPE also supports the various change statuses of your master data using Engineering Change Management in the Discrete Industries System.

The iPPE Make-to-Order process starts with the creation of master data in the Discrete Industries System. The master data is then transferred to the SCM system, where production planning takes place.

The SCM system enables you to plan configurable products of a high order volume very quickly with the help of the *Rapid Planning Matrix*. The various change statuses are taken into account during the planning process.

Next the *backflush* is done and transferred to the Discrete Industries System for further execution, such as postings in Material Management, the delivery processes, and the financial postings.

Figure 1-4 on page 8, along with the discussion that follows the figure, describes the planning process for repetitive manufacturing for configurable and non-configurable products with APO and a DI System.

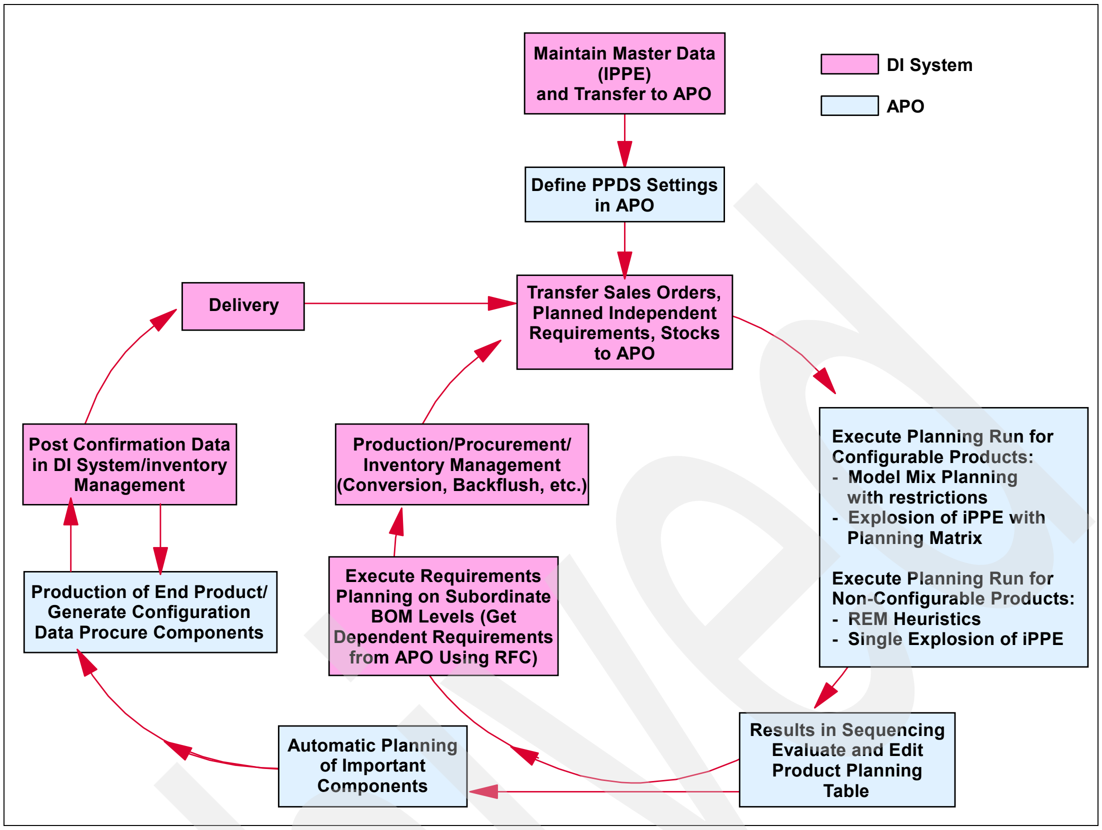

Figure 1-4 The production planning process in the DI System and APO

1. Maintain all the necessary master data in the DI System and then transfer this data to SCM using the SCM Core Interface.
2. In SAP SCM, you generate a production data structure (PDS) from iPPE data. This production data structure is used as the basis for planning in SAP SCM.
3. Maintain the PP/DS (Production Planning and Detailed Scheduling) settings in SCM, such as the factory calendar, propagation range, and so on.
4. Transfer the necessary transaction data (for example, sales orders, planned independent requirements, stocks) to create a consistent planning situation between the two systems.
5. Execute a planning run in APO for the product to be planned.

   To plan configurable products, you can carry out the planning run in *Model Mix Planning*.

   In Model Mix Planning, depending on your settings, an APO planned order with order quantity 1 is created for every sales order. In Model Mix Planning, the planned order is assigned to a line resource on which final assembly is to be carried out, and the order sequence is determined with specification of the start and finish dates of the planned orders. This APO planned order contains no component list if, for the product to be planned, you set that explosion is to be carried out using the planning matrix. In the last step, the planning run in Model Mix Planning explodes the product variant structure and creates a planning matrix. The planning matrix determines the requirement quantities and dates of the required components. This process is very fast.

   As an alternative to the procedure described here, you can also carry out a single explosion of iPPE after Model Mix Planning if you set single explosion of the iPPE data for the product to be planned, or you do not use Model Mix Planning and carry out a planning run in Production Planning. In this planning run, SCM planned orders are created for the product to be planned and a single explosion of iPPE is carried out to determine the components required for the finished product (if you set single explosion of iPPE for the product to be planned).

If necessary, you can also define further heuristics for planning the lower-level BOM levels, which means you can plan all BOM levels in one planning run. However, planning is only executed for the components that you have defined to be further planned in APO. In this case, SCM also creates suitable planned orders for the dependent requirements. You can also execute this planning step in a separate planning run as described in step 7.

6. Evaluate the planning results.

   In the short-term planning period, you can use the sequence schedule to evaluate the planning results of Model Mix Planning (for configurable products). In the sequence schedule, you can display and change the sequence of the dispatched planned orders. You can also use the product planning table. To evaluate the planning results of the non-configurable products, use either the Product Planning Table or the Product View.

7. Less important components.

   Less important components (B and C parts) should not be further planned in SCM, but in the Discrete Industries System instead. When planning these parts, you do not usually require any planning functions other than those in this system.

8. You continue to plan the important components produced in-house or procured externally (A parts), such as the engine, in SCM.

   – Carry out a planning run in SCM for these components. The planning run creates SCM planned orders to cover the dependent requirements determined in SCM. In the planning run a single explosion of the iPPE data is carried out for components produced in-house to determine the dependent requirements for the lower-level components.

   – Evaluate the planning results for the finished product and for the components which you continued to plan in SCM and, if necessary, process the procurement elements created in SCM. To do this, you can use the Product Planning Table or the Product View, for example.

   – Now, you can transfer the SCM planned orders for the components to the DI System, where they are produced or procured (conversion, backflush, and so on).

   – You can process the SCM planned orders or SCM purchase requirements for the externally procured components using collaborative procurement directly in SCM and convert them into purchase orders. Collaborative procurement provides you with support in determining the supply source, delivery schedules, and so forth.

9. You plan and control the production of the finished product (final assembly) in SCM, with the help of the product planning table and sequencing (only for configurable products). You also create the production backflush in SCM. The backflushing data is transferred to the DI System. Here, the goods receipt of the finished product and, at the same time, the goods issue of the components are posted automatically (backflush). However, you can also accumulate the goods issues and post them later. The updated component stock levels are transferred from the DI System to SCM. The dependent requirements are updated in SCM simultaneously. That is, the component requirements are reduced.

Additional details about the parts of the process that are relevant for the performance test are provided in the following sections.

Rapid Planning Matrix

The RPM, or Rapid Planning Matrix, is a product-related database that is stored in SAP liveCache for SAP SCM. It provides two matrices for planning: a component variant matrix for calculating the required components, and an activity matrix for calculating the required activities.

The RPM is the basis of a new form of production planning in SAP APO. It is used for products with many variants and large numbers of orders to quickly determine requirement quantities and dates for the components of a product. It also lets users specify whether the software needs to generate an activity matrix during the planning run.

Overview of the RPM

MRP with the RPM in SAP APO is based on the master data from iPPE. Any changes made to master data, including those made to the PVS, the line resource, or the order sequence, are automatically transferred to the RPM.

The planning matrix includes the following:

- Data from iPPE (PVS, line, and process structures)
- Data for the line resource, such as task time or shift sequence
- Data from orders planned for the product in SAP APO, such as the start and end dates for the order
- Information on which components are needed for which order, determined by exploding the PVS and taking into account object dependencies in iPPE and the validity data maintained in ECM

The following activities require data from the RPM:

- Determination of the component requirements for each order and each defined period for further planning in SAP SCM and SAP ERP
- Generation of order BOMs for planned order management
- Generation of order BOMs for backflush at reporting points
- Transfer of data from the matrices as the basis for sequencing
- Forecasting the delivery schedule, creating a just-in-time (JIT) delivery schedule, issuing a sequenced JIT call

The following example illustrates the type of bottleneck problems that can occur in MRP and the advantages of planning with an RPM: A typical maximum BOM for a car has 40,000 components, including all historical changes. An exploded customer order consists of 2,000 components of the maximum BOM selected, which results in storage of 60 GB of dependent requirements in the MRP database in SAP ERP in relational form. This data must be read and written during an MRP run. The order explosion must consider all historical changes to a component variant, resulting in 1.2 billion checks of effectivity parameters and, in the worst case, 1.2 billion checks of selection conditions. The alternative is to model this data in iPPE and use SAP SCM, which stores data in RAM instead of on a hard drive and uses an optimized form for the data that requires very little memory and allows for object oriented access. The matrix representation in this case allows for simultaneous explosion of all orders instead of explosion on an order-by-order basis.

Discrete Industries Backflush

The *DI Backflush* component is used in an environment with a high volume of orders with small order quantities; often only one piece flow as described in this scenario. The backflush function was developed specifically for this “high volume scenario,” which is particularly common in the automotive as well as the high-tech industries.

The backflush can be carried out very quickly. This is due to both the new link to the SAP SCM system, where the creation of the backflush is separated from the processing of the backflush data, and to the increased separation of the Logistics and Controlling functions.

The backflush is created in the SAP SCM system. The subsequent processes, that is, the posting of goods movements and production activities, are carried out separately in the ERP system.

A backflush can only be carried out for a reporting point that has been defined in iPPE. Therefore, the product structure of the materials to be backflushed must be maintained in iPPE. A direct link exists to Periodical Product Cost Controlling via the product cost collector. The following functions are included in the production backflush:

- The backflush can be executed for make-to-order repetitive manufacturing and make-to-stock repetitive manufacturing, as well as for shop-floor-oriented manufacturing; in this scenario only make-to-order repetitive manufacturing is considered.
- Backflush at reporting points is possible in MTO production.
- To speed up the processing in the ERP system, synchronous and asynchronous goods movements are posted separately.
- The goods receipt in the MTO procedure is evaluated on the basis of the material consumed, the labor executed, and the overhead. This significantly improves performance and enables the system to determine the goods receipt value.
- WIP (work in process) can be calculated in Controlling.
- The production backflush also supports iPPE products with phantom assemblies or by-products.

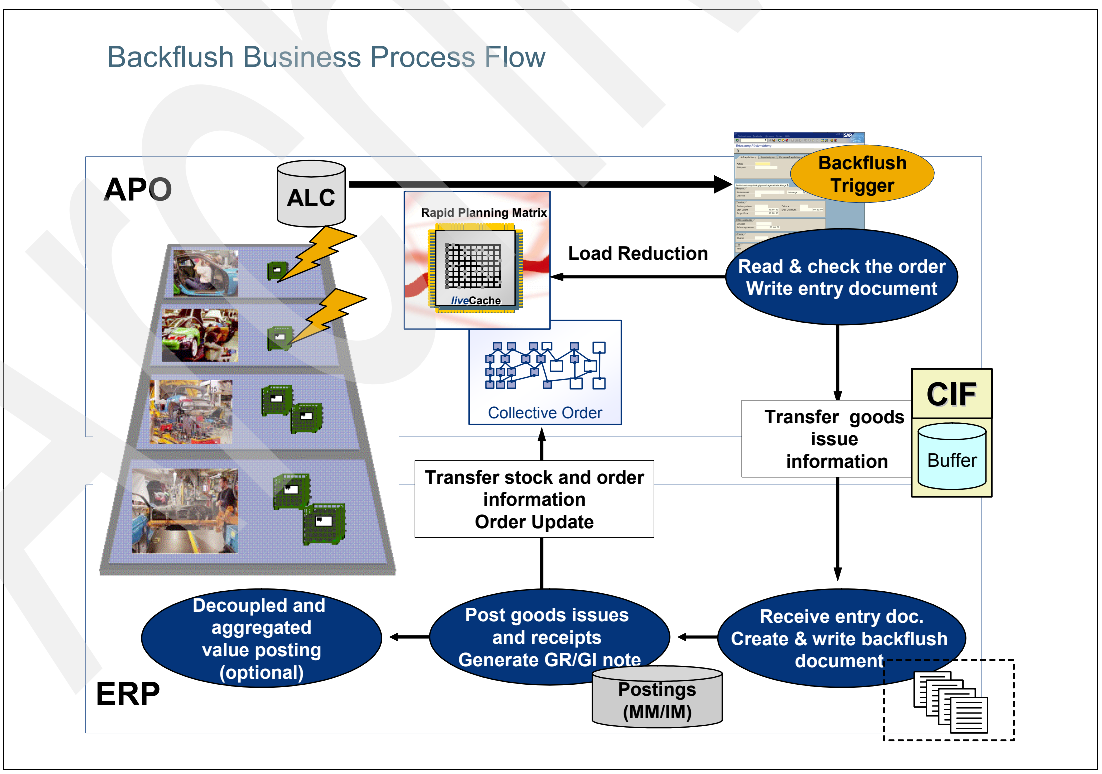

Figure 1-5 Backflush process

With the help of an SCM backflush profile, flexible access to the backflush transaction is possible, as well as the possibility of triggering a backflush via the Action Handler. While processing the SCM orders, the capacities are reduced immediately and the information for inventory management is transferred to the ERP system.

In the ERP system several goods movements take place:

- All goods receipts and goods issues of synchronous materials are posted immediately.
- Goods issues of asynchronous materials are periodically aggregated and posted.
- The processes can be distributed to several servers.
- Failed goods movements end up in the corresponding postprocessing with cumulative and individual records, including those for goods receipts as well.

To perform the actual backflush, several date-specific prerequisites must be maintained in the ERP system. The iPPE data must be maintained and a production version must exist because it is used to determine the cost collector for cost settlements. Additionally, a REM profile must be entered in the material master (for example, movement types, batch and stock determination) and the components with synchronous goods issue as well as these with asynchronous goods issue must be determined. Classes and characteristics must be created, as well as sales orders with assigned characteristic values. Finally, a product cost collector must be created and set as “SCM cost collector” in the first backflush.

The data from iPPE and the production versions in the ERP system are transferred to their correspondents in the SCM system via the CIF interface. The data from the material master is transferred to the product master in the SAP SCM system; additionally, planned orders are created from transmitted sales orders.

The orders in the RPM in SAP SCM are planned production orders. They have a pegging context to a sales order (MTO) or forecast order (MTS). There is no quantity differentiation within the matrix: exactly one column is created per order. Planned production orders must always be split to the quantity of one if the RPM is used; otherwise, the complete planned production order must be backflushed.

Backflushing in the SCM system is always carried out for a certain reporting point. The reporting point structures can be changed during production. One order always refers to exactly one reporting point structure. The definition of reporting points is possible in Line Design at the beginning and/or end of a line segment. Unlike lines, activities contain no reporting point definition. Reversing an action is only possible in a document-neutral way by using the opposite movement direction during posting.

The valuations in Controlling (WIP and GR cost estimate) are based on the actual postings.

When an order has passed a reporting point the matrices are updated to reflect the inputs used for the order. While consumed requirements are shown as “0” or “+”, open requirements are shown as “X”.

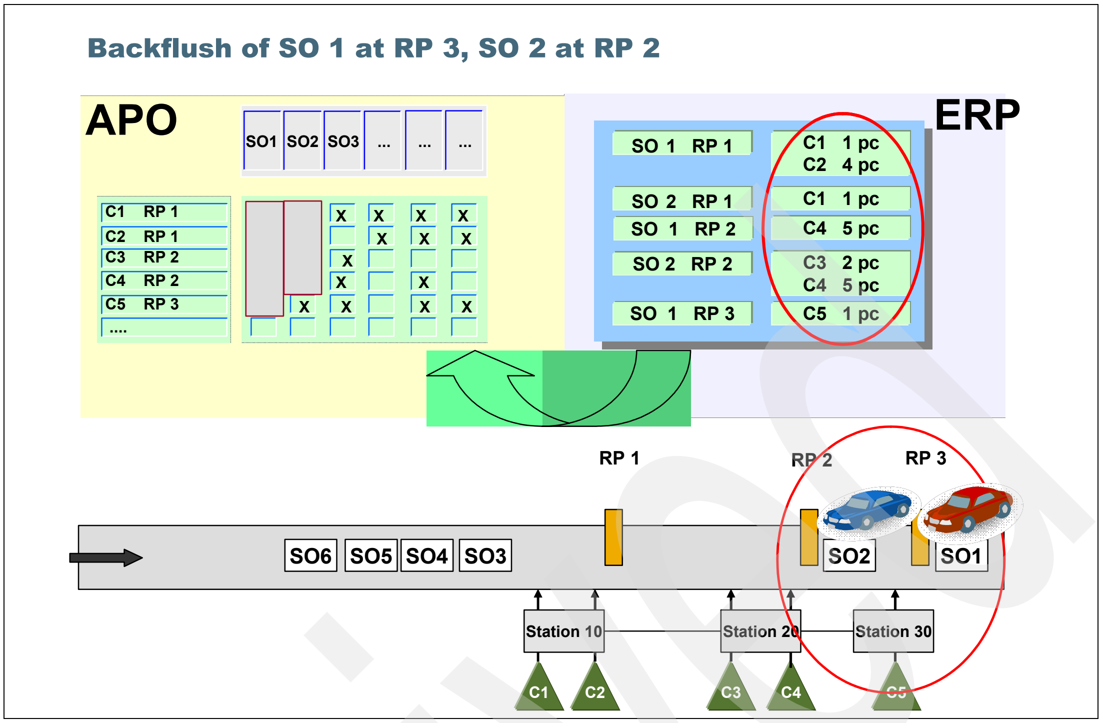

Figure 1-6 Backflush and Rapid Planning Matrix

The backflush data is sent from SAP SCM to mySAP ERP for posting in FI and CO. Because the vehicle is still in production at this time, component costs are allocated to WIP. The material price is maintained in the mySAP ERP system. Upon goods receipt, the value of the components is transferred from WIP to Finished Goods (synchronous posting). This is done by posting a credit to the WIP account. In case of scrap, adjusting entries need to be made to reduce the WIP account and increase the Scrap account.

Backflush process flow (Backflush in two steps)

Production backflushes, which are sent from an APO system to an ERP system, contain components and activities to be posted. Components can be posted either synchronously or asynchronously, whereas activities are always posted asynchronously (transaction PPCGO).

Synchronously means that the goods movement is posted immediately on transferring the backflushing data from the SCM system to the ERP system. Goods movement or production activity that is not posted immediately on transfer is posted asynchronously. The backflushing data from the SCM system to the ERP system is initially recorded in a buffer on the database. Processing and posting the data is triggered later by a report.

The indicator "synchronous posting GI" in the SCM product master determines how a material is posted. Goods issues are always posted immediately (that is, synchronously) irrespective of the indicator in the product master.

During the backflush process in a high volume production process, the number of cost collectors increases significantly, which could lead to a deterioration in performance. This occurs because material documents must be created for each cost collector, even if the components are the same. (This step is necessary for the postings in FI/CO.) To avoid this deterioration in performance, the backflush can be executed in two steps.

This processing has the advantage that the logistical data remains consistent even if the FI/CO data is only posted once a day. In the case of the two-step production backflush, the requirements are adjusted in SCM after the first step.

The two-step production backflush is illustrated in Figure 1-7 and works in the following way:

- The system reads all open backflushes and corresponding components. The quantities of components are grouped according to posting date. The system posts the components, without reference to a product cost collector, to a non-MRP storage location. The production backflushes are updated. The goods movements are posted. At the same time, the product cost collectors and the information about the components are stored in a new table. The backflush data is updated in the APO system.
- The components are grouped according to posting date and product cost collectors. The goods issues are posted per product cost collector (for example, once a day). The data in CO/FI is updated.
- If errors occur in either of the steps, postprocessing of the production backflushes is possible with the transaction COGI.

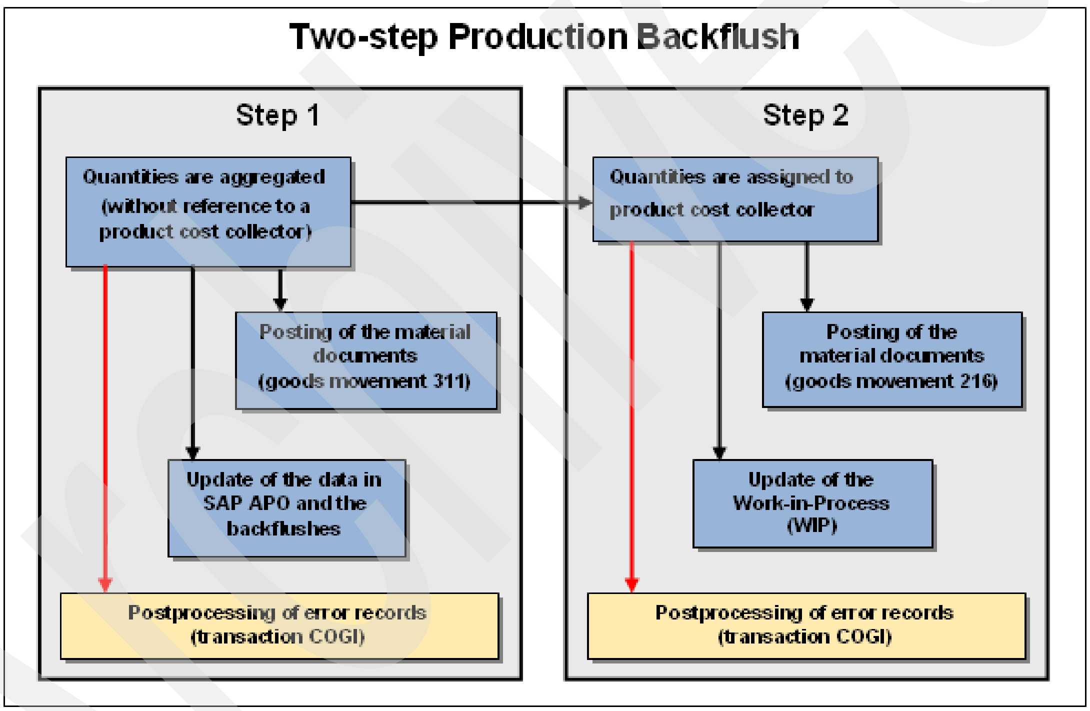

Figure 1-7 The two-step backflush

The advantage of the two-step production backflush is that the logistical data is already updated before the data in CO/FI is posted. The goods movements are aggregated before they are posted, thereby improving the performance of the system.

For example: Assume there are 100 assemblies, each made up of 500 components, that are produced every 30 minutes. The posting results are quite different depending on the backflush.

Using the one-step production backflush the system creates 100 material documents every 30 minutes, with 500 items each. This leads in total to a posting of 50,000 line items, and in two hours 200,000 line items are posted.

In contrast, in the two-step production backflush only one material document with 500 line items is created every 30 minutes. After two hours the goods issue is posted for each product cost collector from the non-MRP relevant storage location. This results in a posting of 50,000 line times - 52,000 line items in total compared to the 200,000 line items in the one-step production backflush.

1.4 The customer project

The customer for whom the performance test was carried out is a large automotive manufacturer from the Asia-Pacific region with a sales network extending to 193 countries.

In 2004, the customer celebrated the export of 10 million cars. In order to be closer to consumers, he has established a manufacturing presence in the USA, China, India, and Europe.

The customer has recently selected SAP as the core of his IT system strategy. When fully implemented, the objective of the IT solution is to manage his automotive production, which is targeted to produce 5 million vehicles per year by 2010.

A multi-year SAP implementation project plan and approach has been defined with the following key aspects:

- ERP implemented in local Asian plants when technically proven in the smaller facilities.
- Reutilize fully the experience of resources who were involved in the implementation projects of the smaller facilities.
- Minimize implementation risks with a phased implementation.

The test objects included in the performance tests were the following:

- RPM (Rapid Planning Matrix)
- MRP (Materials Requirements Planning)
- Backflush
- Material Ledger

Figure 1-8 on page 16 provides an overview of the customer’s process.

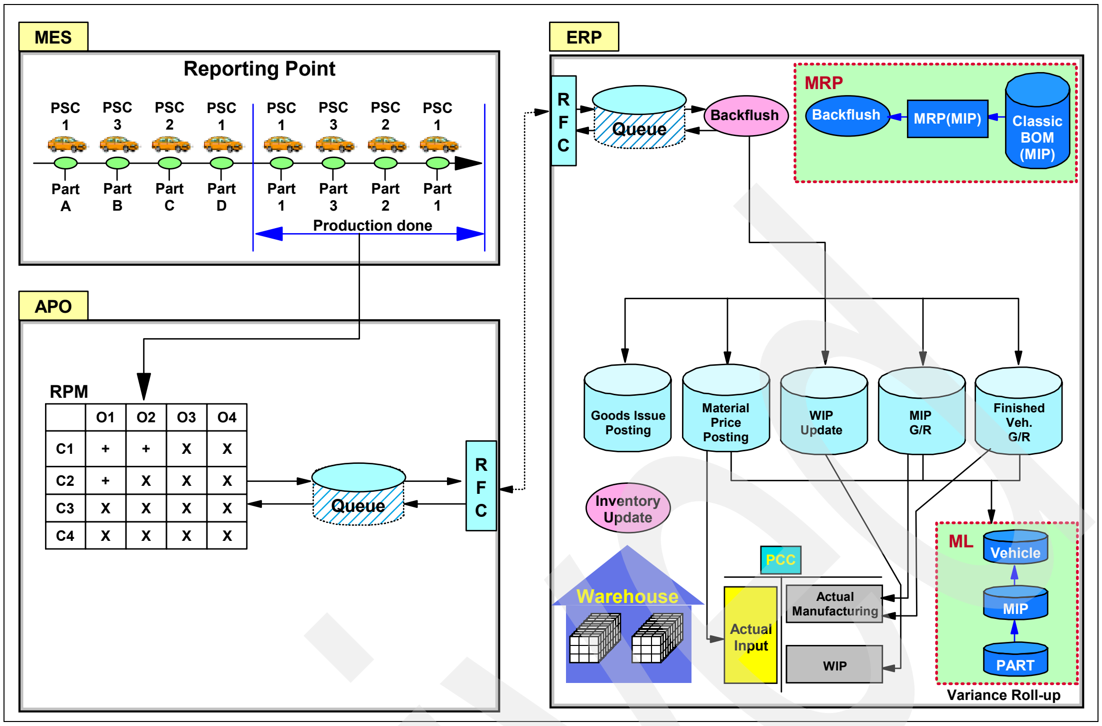

Figure 1-8 SAP Operational model

Customer specifics

The customer currently uses fully specified BOMs. Because this model is not robust enough to accommodate the customer's future growth plans, it was decided to use iPPE to handle the anticipated large volume of data. In the first project stage, however, the plan is not to use the full capacity of iPPE to reduce the data amount, but to create for every fully specified BOM a header variant in the iPPE Access Node as entry points to the BOM.

The customer's Product Variant Structure will be the three-level BOM shown in Figure 1-9 on page 17.

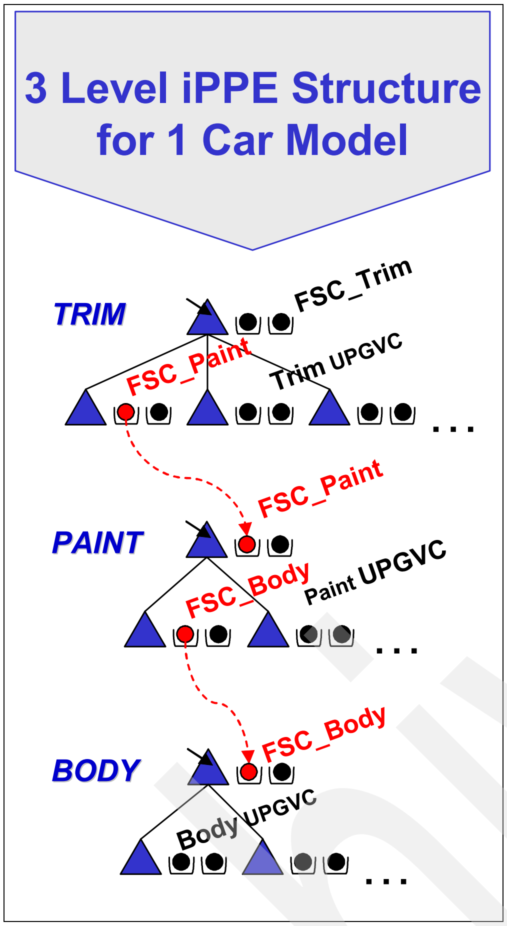

Figure 1-9 Customer PVS

Production will take place as shown in Figure 1-10.

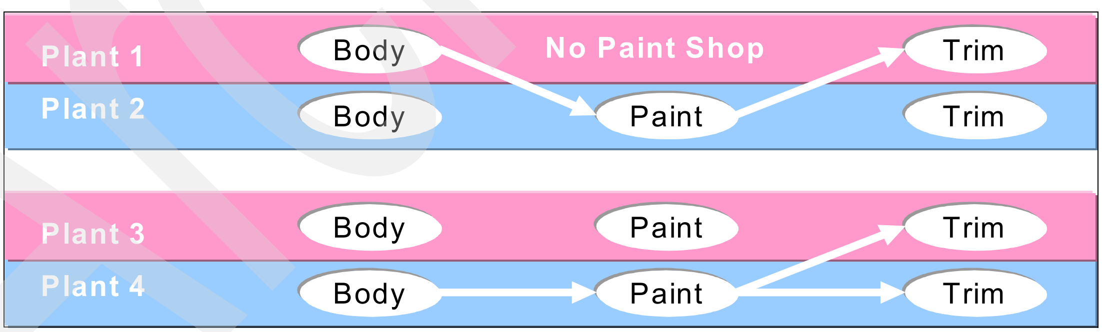

Figure 1-10 Customer production

In addition, the customer has the following business requirements:

1. They need to manage inventory for Painted Body and Body.
2. They need to have “Controlling” for each shop.
3. They need to move materials between plants.

2

Architecture definition

This chapter describes the architecture of the system we used to perform benchmark testing. The architecture was designed to satisfy the complexity and magnitude of the customer's business, and therefore the proposed SAP solution.

This chapter includes the following topics:

2.1 Architecture overview

This section provides an overview of the architecture we designed and implemented to run the performance tests.

2.1.1 SAP solution logical components overview

The SAP solution components are shown in Figure 2-1.

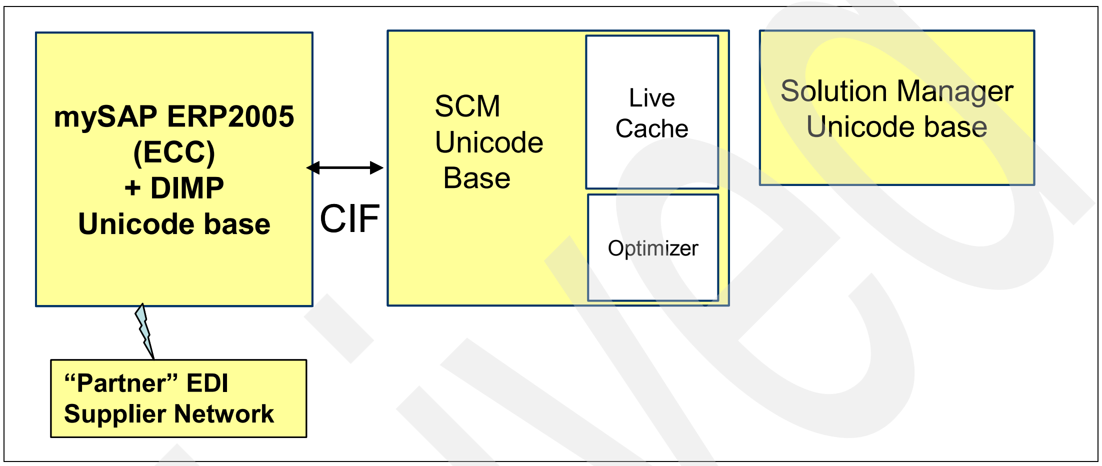

Figure 2-1 SAP solution logical Components

The components are:

- ECC: ERP Central Component that replaced the SAP R/3® Enterprise core and provides essential ERP module functions. The ECC version we used is 6.0.
- DIMP: Discrete Industries & Mill Products, the Industry Solution selectable in ECC which is required for automotive enterprises[1].
- SCM: Supply Chain Management supports planning, execution and coordination of the supply chain network.
- APO: Advanced Planner & Optimizer is a component of SCM that is used for the planning and optimization of supply chain processes at a strategic, tactical, and operational planning level.
- liveCache: In-memory database backed by disk and used to manage complex data structures and data flows in APO.
- CIF: Core InterFace is a standardized interface between SCM/APO and various releases of the ERP system. SCM and APO are linked to the ERP system, which provides the master data. The planning results are subsequently re transferred to the ERP system for further processing and execution.
- The Solution Manager component provides support for SAP implementation, operation, and optimization.
- Unicode provides a unique number for every character, regardless of the platform, program, and language. As of 1st January 2007, all SAP releases are Unicode (except upgrade cases). For further details see: `http://www.unicode.org`

[1] As of mySAP ERP 2005 the Industry Solution components are now integrated into the core Enterprise Resource Planning (ERP) package.

2.1.2 Non-functional requirements for architecture definition

The initial aim of the zPOC is to test the reliability and the performance of the previously explained functional processes, namely:

- Production planning
- Backflush
- Materials requirements planning
- Materials ledger

In this case we have to consider that the customer's manufacturing capacity represents large volumes for these processes. From these business needs we established the following non-functional requirements.

The infrastructure must be:

- Highly available
- Scalable
- Reliable
- Secure
- Highly performing
- Flexible
- Easy to manage

2.1.3 Architectural decisions

These non-functional requirements led us to the following platforms:

- The SAP ECC and SCM databases run on IBM System z9® S38 because:
  - IBM System z offers the most highly available, scalable, reliable, and secure technology platform for key SAP databases and Central Services.
  - A subsequent choice of IBM System z for the database environment could facilitate:
    - Migration from existing mainframe applications to SAP
    - Interfaces to other non-SAP applications
  - The IBM System z9 S38 offers more processing capacity than will actually be required for the zPoC; however, actual utilization will be reported to define the production requirement.
- The SAP application servers run on IBM System p because IBM System p for application servers provides the performance and throughput, flexibility and manageability required to handle a broad mix of SAP workloads.
- The SAP liveCache component runs on System p (although it is a database, MaxDB in this case) because it is not supported on System z by SAP. The liveCache runs on a dedicated System p server.

Using a mix of IBM System z and IBM System p as the infrastructure has several advantages:

- The flexibility of both platforms allows resources to be easily assigned or reassigned during the tests, which is especially valuable where the workload and application behavior cannot be accurately predicted, such as for materials ledger.

- ► The performance and scalability of both platforms provide the best opportunity to meet or exceed the customer's requirements, unless the application structure and processing flow are the main inhibitors.
- ► The combination of these two platforms in an appropriate configuration would also address customer non-functional requirements in an actual deployment.

2.1.4 Operational model - Physical architecture

In this section we provide an overview of the physical architecture supporting the SAP components. Figure 2-2 shows the servers, disk subsystem, network configuration, and SAN configuration.

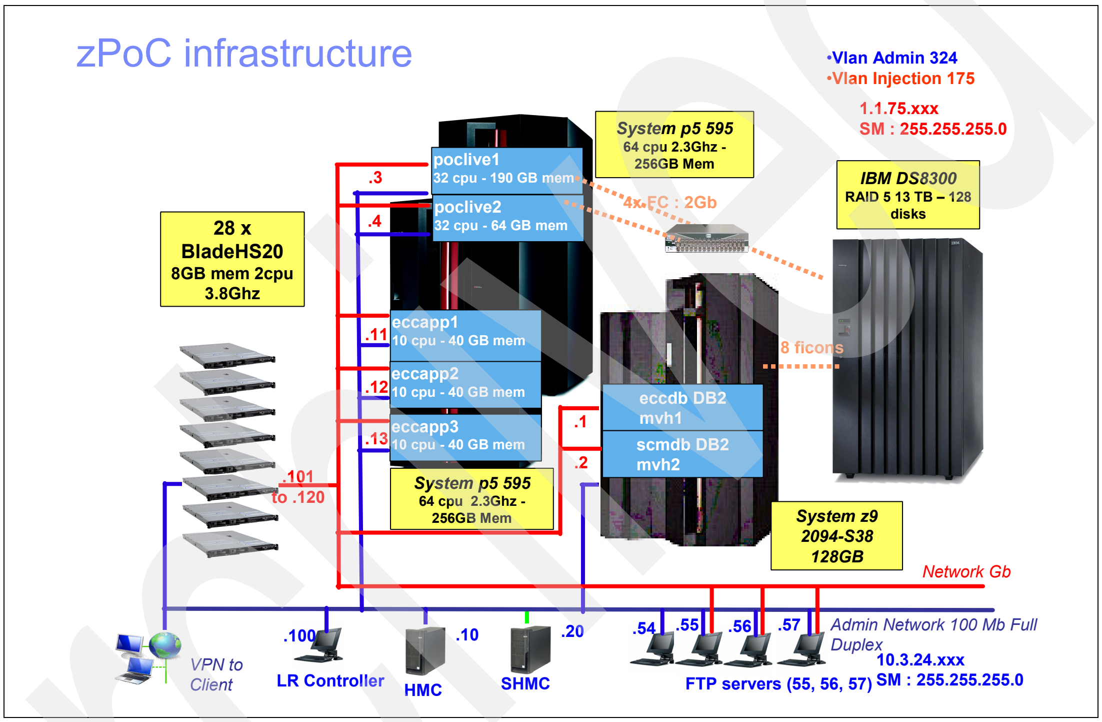

Figure 2-2 zPOC infrastructure

This architecture includes:

- ► Servers:
  - – 2 System p5® 595 POWER5+™ 2.3 GHZ with 64 cores and 256 GB RAM each
  - – 1 System z9 2094-S38 (2094-738) with 128 GB as Database servers
  - – 28 injector machines (blade technology) for workload simulation
- ► Disk subsystem:
  - – 1 DS8300 Turbo with 13TB capacity
- ► Network:
  - – 3 VLANs are defined:
    - • Administration network - VLAN 324 at 100 Mbps
    - • Injection network - VLAN 175 at 1 Gbps
    - • User network: IBM IP address / local IP address at 1 Gbps
  - – All the servers used for this benchmark are connected to private VLANs. The IBM Montpellier Team set up a VPN connection with port-forwarding to allow the customer

to connect remotely to the systems. An external user can connect to those systems with a proper authentication to the VPN concentrator located in IBM Montpellier. Once a user is authenticated through a Web site, a Cisco Applet on the user's local machine listens locally on a specific port and redirects the service to the target server with the right port. Two VLAN Networks are defined on a Cisco switch #7609 for the following purposes: one front end VLAN Network for Load Runner injectors communication to the online application, and a back end VLAN Network for db2 intercommunications.

– All the ports are on the same Cisco Ethernet switch dedicated to the zPoC. There is no trunk used between switches except for a switch that concentrates the Load Runner injectors' network.
– For storage attachment, we use eight FICONS to link the system z platform to the storage unit and four fiber cards of 2 Gigabits between system p and the storage unit.
– System p servers are connected to the storage subsystem through a SAN. Each server has 4 HBAs at 4 GBs connected to a switch. Load balancing is done by MPIO (multi-path I/O) and SDDPCM (Subsystem Device Driver Path Control Module).
– System z has 8 Ficon channels connected to the storage subsystem.

Bear in mind that this architecture has been defined for benchmark purposes, taking into account that many hardware resources are available. The architecture would be different for a real production environment because of the need to also consider high availability, resources optimization (virtualization, micro-partitioning), access connections, security, and so forth.

2.1.5 Operating systems and software levels

Table 2-1 provides a quick view of the basic software components and levels used for these tests. Notice that we used the latest software versions available at the time the project started.

Table 2-1 Software components

Name	Version/Level	Updates	Purpose	Notes
z/OS	V1.8	Put0703+	System z Operating System	
AIX 5L™	V5.3	TL5 SP6	System p UNIX® Operating System at latest level	ML = Maintenance level ITL = Technical Level
DB2	V8.1	Put0703+	DB2 on z/OS	
DB2 Connect™	8	FixPack 12	Database connection from AIX to z/OS	SAP Special Build 16563
mySAP ERP 2005	ECC 6.0		ERP Central Components	Customer licenses
mySAP SCM	V5.0		Supply Chain Management includes Advanced Planner & Optimizer (APO) component	Customer licenses
Windows® XP	SP2 + latest fixes		OS for Injection controllers	
Windows 2000 Advanced Server	SP4 + latest fixes		OS for Injection generators	
Mercury Loadrunner	8.1.4		Dialogue injection software	Customer licenses

Name	Version/Level	Updates	Purpose	Notes
Tivoli® OMEGAMON® XE for DB2 Performance Expert	3.1		DB2 Performance	
MaxDB 7.6.00	7.6.02		SAP Live Cache 7.6.1	Build 014-123-152-175

2.2 Solution components configuration

In this section we provide details about the infrastructure configuration that was set up for this benchmark.

2.2.1 SAP systems landscape

In 2.1, "Architecture overview" on page 20, we had a first look at the overall solution and technical infrastructure set up for this benchmark. In this section we describe how SAP systems fit into the infrastructure, as shown in Figure 2-3.

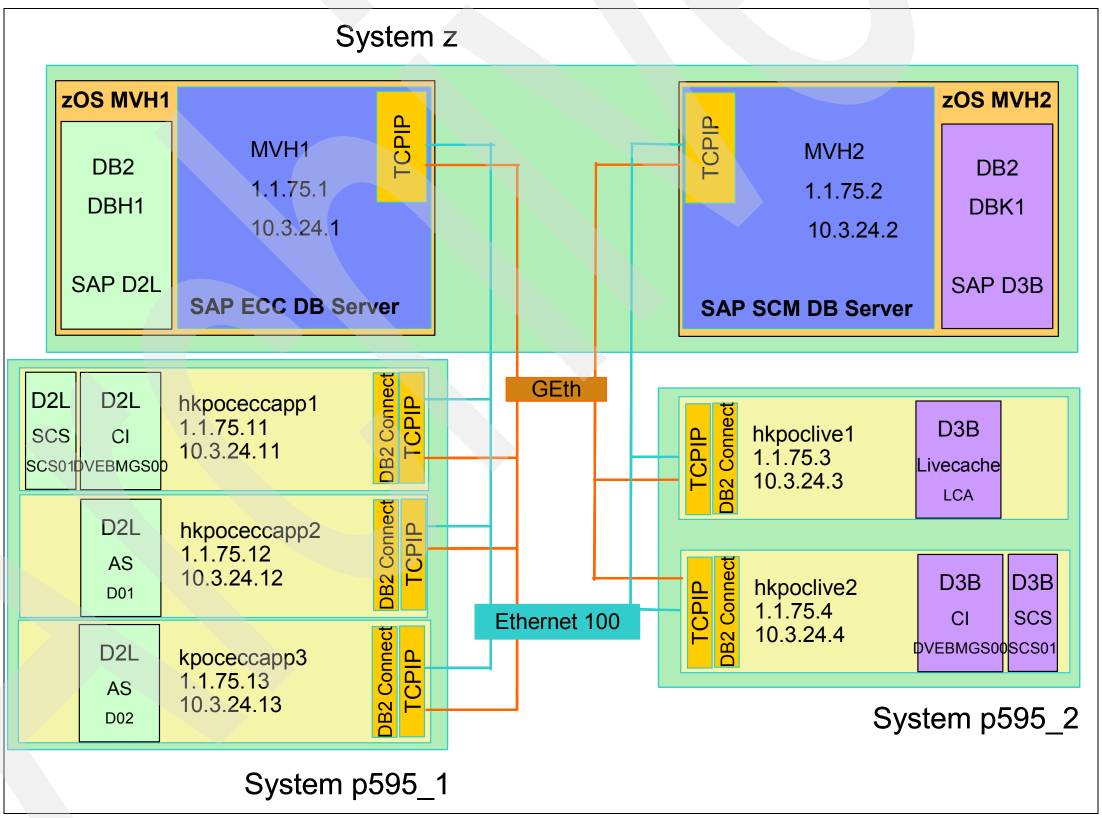

Figure 2-3 SAP systems landscape

On System z9 we defined 2 LPARs:

- MVH1 runs the SAP ECC database server.
- MVH2 runs the SAP SCM database server.

We used two System p595 servers:

- System p595_1, shown on the lower left side of Figure 2-3, contains 3 LPARs:
  - hkpoceccapp1 running SAP ECC central services and central instance DVEBMGS00
  - hkpoceccapp2 running SAP ECC dialog instance D01
  - hkpoceccapp3 running SAP ECC dialog instance D02
- System p595_2, on the lower right side of Figure 2-3, contains 2 LPARs:
  - hkpoclive1 running SAP liveCache instance LCA
  - hkpoclive2 running SAP SCM central system

Now that we have shown which SAP systems and instances run on which hardware resources, we next describe the hardware resource configuration in more detail.

2.2.2 IBM System z configuration

The first three KPIs (MRP, RMP, and BF) were tested with a different configuration than the configuration used for testing the last KPI, ML. The original configuration was an S38 with 36 CPs and 2 zAAP processors.

After a microcode change the configuration was changed to an S38 with 20 CPs and 18 zIIP processors.

As shown in Figure 2-2 on page 22, two LPARs are defined:

- MVH1 for hosting the ECC-database
- MVH2 for hosting the SCM-database

Figure 2-4 and Figure 2-5 show the z9 HMC panels where the original LPAR settings were specified in preparation for testing of the first three KPIs. Both LPARs are defined identically: 8 CPs and 24 Gb Storage. Thus, during these tests 20 CPs, 2 zAAPs, and 80 Gb storage were not used at all.

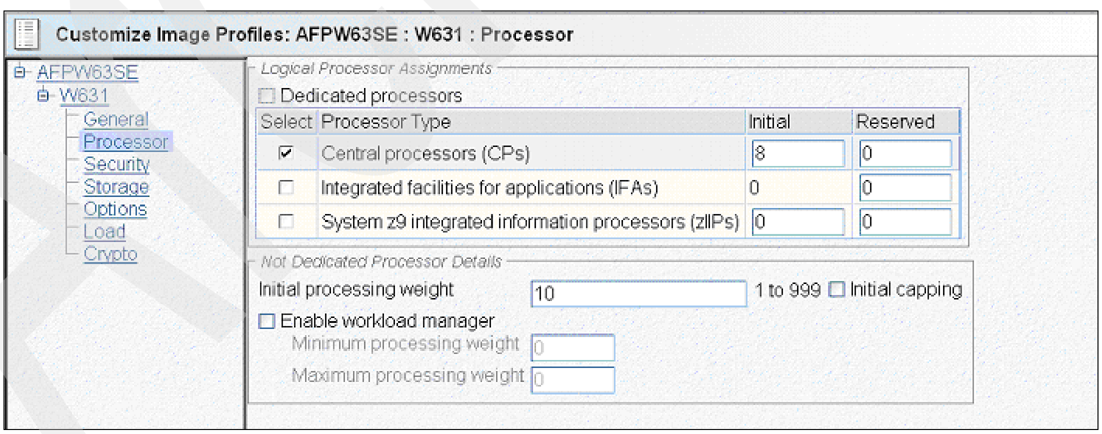

Figure 2-4 z9 HMC Processor info original

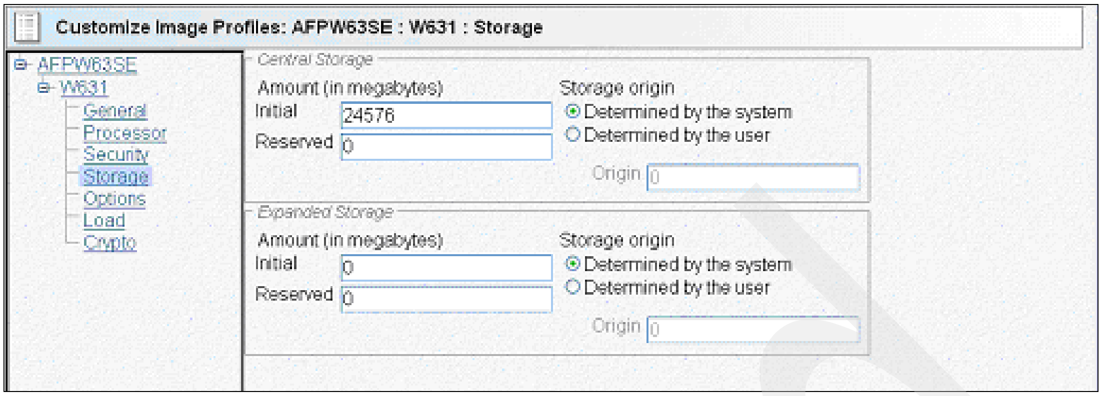

Figure 2-5 z9 HMC Storage info original

Figure 2-6 shows the z9 HMC panel with the modified LPAR settings used during testing of the ML KPI. Both LPARs are defined identically: 10 CPs, 9 zIIPs, and 24 Gb Storage, which is not changed. (Note: Only the initial numbers of CPs and zIIPs are available after an IPL. The reserved processors can be configured online as needed.) During those tests we have more flexibility to use the available processors. The 80 Gb unconfigured storage is still not used.

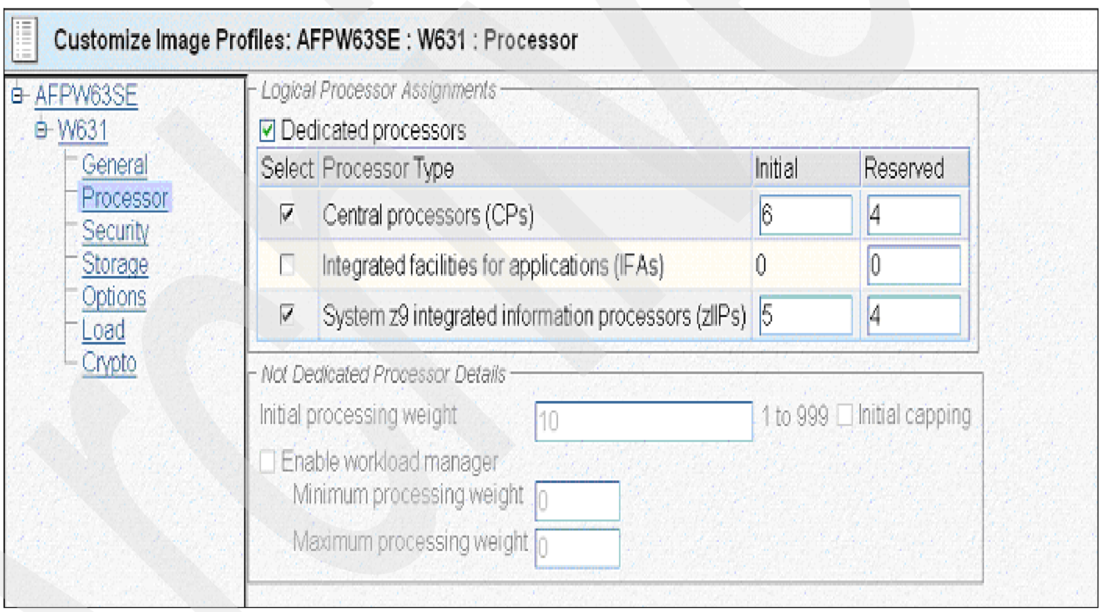

Figure 2-6 z9 Processor info changed

On the Channel Path Activity panel shown in Figure 2-7 on page 27, the relevant channels are:

- Channel 20 is a shared OSA Express channel used for connecting the z9 to the administrators network (TSO, RMF™ traffic). Both LPARs are using this channel.
- Channel 30 is a dedicated OSA Express Channel used for connecting the z9 to the P595-machines and is used as a non routed database VLAN. (On MVH2 there is also a dedicated OSA Express channel 31 defined.)
- Channels 40 through D1 are shared FICON® channels used for connecting the z9 to the DS8000™ storage environment. Both LPARs are using those channels.

```
                         RMF V1R8   Channel Path Activity           Line 1 of 11
Command ===>                                                   Scroll ===> CSR

Samples: 100     System: MVH1  Date: 10/16/07  Time: 14.11.20  Range: 100   Sec

 Channel Path        Utilization(%)    Read(B/s) Write(B/s)    MSG  MSG Send Recv
ID No  G  Type   S    Part  Tot  Bus    Part  Tot  Part  Tot   Rate Size Fail Fail

20        OSD    Y     0.0  0.0  0.0     21   21    2K   2K
30        OSD         13.2 13.3  1.2     5M   5M    4M   4M
40     4  FC_S   Y     3.0  3.1  0.9     2M   2M    3M   3M
41     4  FC_S   Y     3.1  3.2  0.9     2M   2M    3M   3M
50     4  FC_S   Y     3.1  3.2  0.9     2M   2M    3M   3M
51     4  FC_S   Y     3.0  3.1  0.9     2M   2M    3M   3M
C0     4  FC_S   Y     3.1  3.1  0.9     2M   2M    3M   3M
C1     4  FC_S   Y     3.0  3.0  0.8     2M   2M    3M   3M
D0     4  FC_S   Y     3.0  3.0  0.9     2M   2M    3M   3M
D1     4  FC_S   Y     3.1  3.1  0.9     2M   2M    3M   3M
```

Figure 2-7 z9 Channel paths

Figure 2-8 shows the Cache Summary displaying all DASD control units. All volumes that are relevant to our testing are randomly striped over these control units.

```
                        RMF V1R8   Cache Summary    - PORTINH1      Line 1 of 21
Command ===>                                                   Scroll ===> CSR

Samples: 100     Systems: 1     Date: 10/16/07  Time: 14.11.20  Range: 100   Sec
                               CDate: 10/16/07 CTime: 14.11.16 CRange: 100   Sec

SSID  CUID Type-Mod  Size  I/O   Hit  Hit   -- Miss ---  Read  Seq  Async  Off
                           Rate   %   Rate  Total Stage   %    Rate Rate   Rate

A0FF  A000 2107-921   30G   0.0  100   0.0    0.0   0.0   100   0.0   0.0   0.0
C000  C000 2107-932   61G  25.3  100  25.3    0.0   0.0  72.9   0.0   8.2   0.0
C100  C101 2107-932   61G 113.4 96.0 108.9    4.6   4.6  80.2   4.3  18.2   0.0
C200  C221 2107-932   61G  33.0 91.3  30.2    2.9   2.6  61.4   0.0  15.4   0.0
C300  C301 2107-932   61G 209.7 80.9 169.6   40.1   1.0  58.5   0.0 250.1   0.0
C400  C402 2107-932   61G  72.6 98.4  71.4    1.2   1.1  55.2   0.6  46.1   0.0
C500  C501 2107-932   61G 309.8 91.0 281.8   28.0  25.8  57.8   0.0  62.9   0.0
C600  C602 2107-932   61G   0.8  100   0.8    0.0   0.0   100   0.0   0.8   0.0
C700  C702 2107-932   61G   0.8 83.3   0.7    0.1   0.1   100   0.0   0.2   0.0
C800  C801 2107-932   61G   1.1  100   1.1    0.0   0.0  75.0   0.0   0.4   0.0
C900  C902 2107-932   61G   0.8 82.1   0.7    0.1   0.1   100   0.0   1.5   0.0
CA00  CA21 2107-932   61G   0.8  100   0.8    0.0   0.0   100   0.0   0.8   0.0
CB00  CB21 2107-932   61G 366.5 95.2 348.9   17.6  10.3  46.8   2.3  95.8   0.0
CC00  CC01 2107-932   61G 878.5 91.9 807.2   71.3  66.7  46.4   8.3 204.5   0.0
CD00  CD21 2107-932   61G  80.6 98.4  79.3    1.3   0.4  81.8  13.2  30.5   0.0
CE00  CE01 2107-932   61G   0.0  0.0   0.0    0.0   0.0   0.0   0.0   0.0   0.0
CF00  CF22 2107-932   61G   0.0  0.0   0.0    0.0   0.0   0.0   0.0   0.0   0.0
D000  D002 2107-932   61G   0.0  0.0   0.0    0.0   0.0   0.0   0.0   0.0   0.0
D100  D122 2107-932   61G   0.0  0.0   0.0    0.0   0.0   0.0   0.0   0.0   0.0
D200  D222 2107-932   61G   0.0  0.0   0.0    0.0   0.0   0.0   0.0   0.0   0.0
D300  D321 2107-932   61G   0.0  0.0   0.0    0.0   0.0   0.0   0.0   0.0   0.0
```

Figure 2-8 Cache Summary

Among the volumes defined are those for:

- The operating system.
- DB2 (catalog, directory, active logs, workdatabase, and tablespaces). These volumes are all defined on the CUs C000 through CF00. Those CUs are the only ones considered relevant during the tests.
- DB2 archiving.
- Flashcopys.

Table 2-2 and Table 2-3 on page 28 define the DASD layouts of the ECC and SCM databases respectively.

Table 2-2 ECC DASD layout

	VOLSER	HLQ
Cat/Dir	DBH0D1 DBH0D2	DSNDBH0
Workdatabase	DBH1W1 DBH1W2 DBH1W3 DBH1W4 DBH1W5	DSNDBH0
Active logs	DBH1L1 DBH1L2 DBH1L3	DSNLGH0
SAP spaces	HK1000 through HK1117	DB2CLNT1

Table 2-3 SCM DASD Layout

	VOLSER	HLQ
Cat/Dir	DBK0D1 DBK0D2	DSNDBK0
Workdatabase	DBK1W1 DBK1W2 DBK1W3 DBK1W4 DBK1W5	DSNDBK0
Active logs	DBK1L1 DBK1L2 DBK1L3	DSNLGK0
SAP spaces	HK2000 through HK2069	DB2CLNT2

2.2.3 IBM System p configuration

As you can see in the HMC panel in Figure 2-9 on page 29, the following two System p595s are used for this benchmark:

- MCP17M-595-SN51263CF
- MCP17M-595-SN512644F

These are shown as System p595_2 and System p595_1 in Figure 2-3 on page 24.

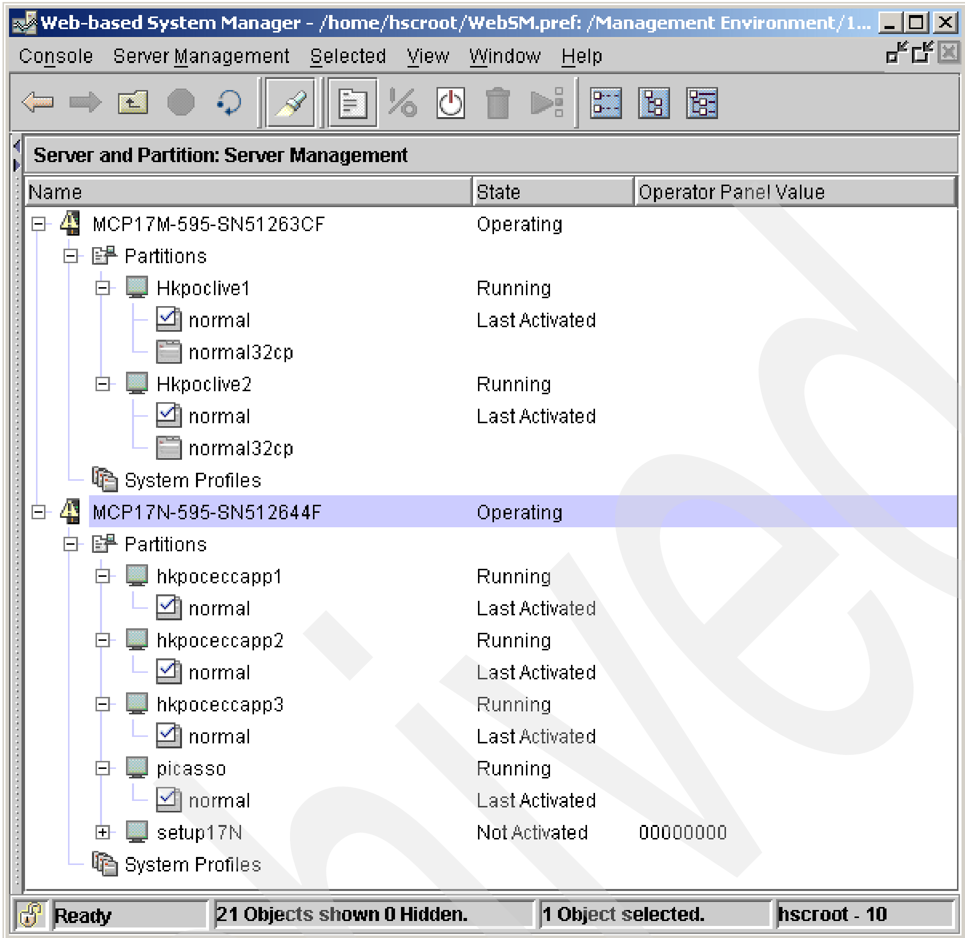

Figure 2-9 HMC view

2.2.4 MCP17M-595-SN512644F and MCP17M-595-SN51263CF configuration

Table 2-4 summarizes the logical partitioning of the two System p machines and their application throughout the benchmark testing. For complete details and illustrations of the configuration process, refer to 2.3, "MCP17M-595-SN512644F and MCP17M-595-SN51263CF configuration details" on page 31.

Table 2-4 Logical partitioning summary

LPAR	#CPU	RAM (GB)	#FC adapters	#Ethernet adapters	Used for
MCP17M-595-SN5152633CF					
hkpoceccapp1	10	40	0	2	ECC central instance
hkpoceccapp2	10	40	0	2	ECC dialog instance
hkpoceccapp3	10	40	0	2	ECC dialog instance
MCP17M-595-SN512644F					
hkpoclive1	32	196	4	2	liveCache
hkpoclive2	32	56	0	2	SCM central instance

The System p configuration was used for the MRP, RPM, and Backflush KPIs. For the Material Ledger KPI, the configuration was changed by the Montpellier team to address the type of workload executed. The DS8300 has four Power5 cores and 128 GB memory. The system is not running in LPAR mode. This means that the same resources (cache, processors, disks) are used simultaneous by the z and the P environments.

2.2.5 Storage subsystem configuration

One storage subsystem DS8300 Turbo was used for this benchmark. The DS8300 has four Power5 cores and 128 GB memory. The box is not running in LPAR mode. This means that the same resources (cache, processors, disks) are used simultaneously by the z and the P environments. DS8300 contains 128 disks, grouped in 16 blocks of 8 disks. Each disk has 146 GB capacity and runs at 15 Krpm.

The configuration details are in 2.4, “Storage configuration” on page 61.

The storage dedicated for the pSeries® for the SAP applications servers was composed of six ranks of eight disks in RAID 5:

- This storage box was connected to the SAN with four fiber cards (4 Gb/s).
- The pSeries applications servers were connected to the SAN with two fibers (4 Gb/s).
- For each application server, six LUNs of 200 GB were created on separate Raid5 arrays.

2.2.6 SAP systems configuration

The settings used within SAP are described in the chapters related to the KPI results.

2.2.7 Load runner

To simulate the online activity workload, Mercury LoadRunner was used to generate a workload scenario against the system.

The main characteristics of a scenario are the number of virtual users (Vusers), the script replay mode (loop or one shot), and the duration of the loop.

The LoadRunner infrastructure is composed of a master workstation serving as a controller and 28 blade servers actually generating the load (injectors). Mercury LoadRunner is used on both the generators and the controller.

2.2.8 Network

The network used was described in 2.1.4, “Operational model - Physical architecture” on page 22. As described there, three VLANs were set up and used in this benchmark for server communication, administration, and VPN to customer.

2.3 MCP17M-595-SN512644F and MCP17M-595-SN51263CF configuration details

As you can see in Figure 2-10, this System p595 server is configured with 64 cores at 2.3 GHz, as stated in 2.1.4, “Operational model - Physical architecture” on page 22.

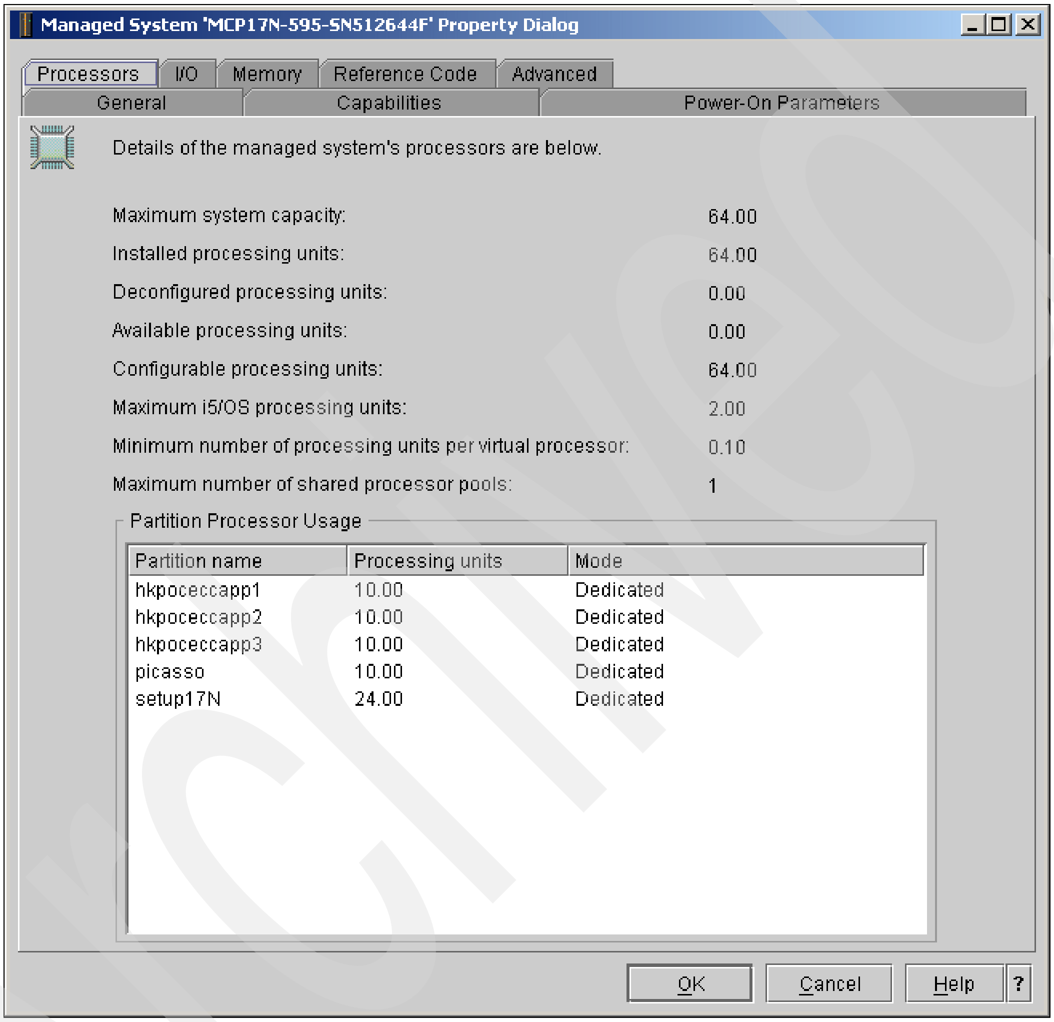

Figure 2-10 MCP17M-595-SN512644F CPU configuration

This server is configured with 256 GB RAM, as shown in Figure 2-11.

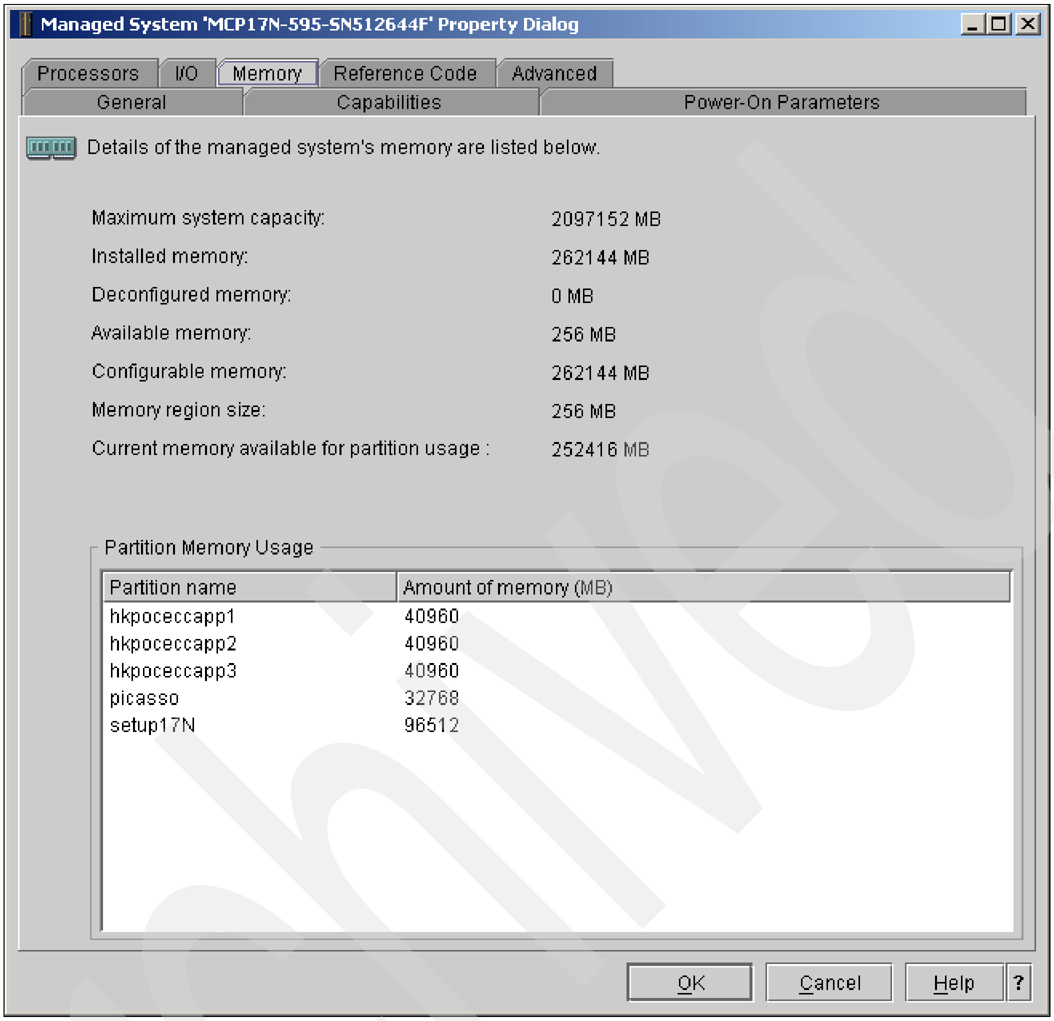

Figure 2-11 MCP17M-595-SN512644F memory configuration

Note that:

- No virtualization is used on this server; all resources are assigned to capped LPARs.
- The LPARs called "picasso" and "setup17N" are not relevant for this project.

LPAR hkpoceccapp1 configuration

The SAP ECC central instance uses an LPAR with 10 CPUs and 40 GB RAM, and it has the following I/O capability assigned:

- One SCSI controller with the local disks
- One PCI slot with one Ethernet Fiber card with two ports at 1 Gbps
- One PCI slot with one Ethernet 10/100/1000 card with two ports

Figure 2-12 on page 33 and Figure 2-13 on page 34 display these partition properties.

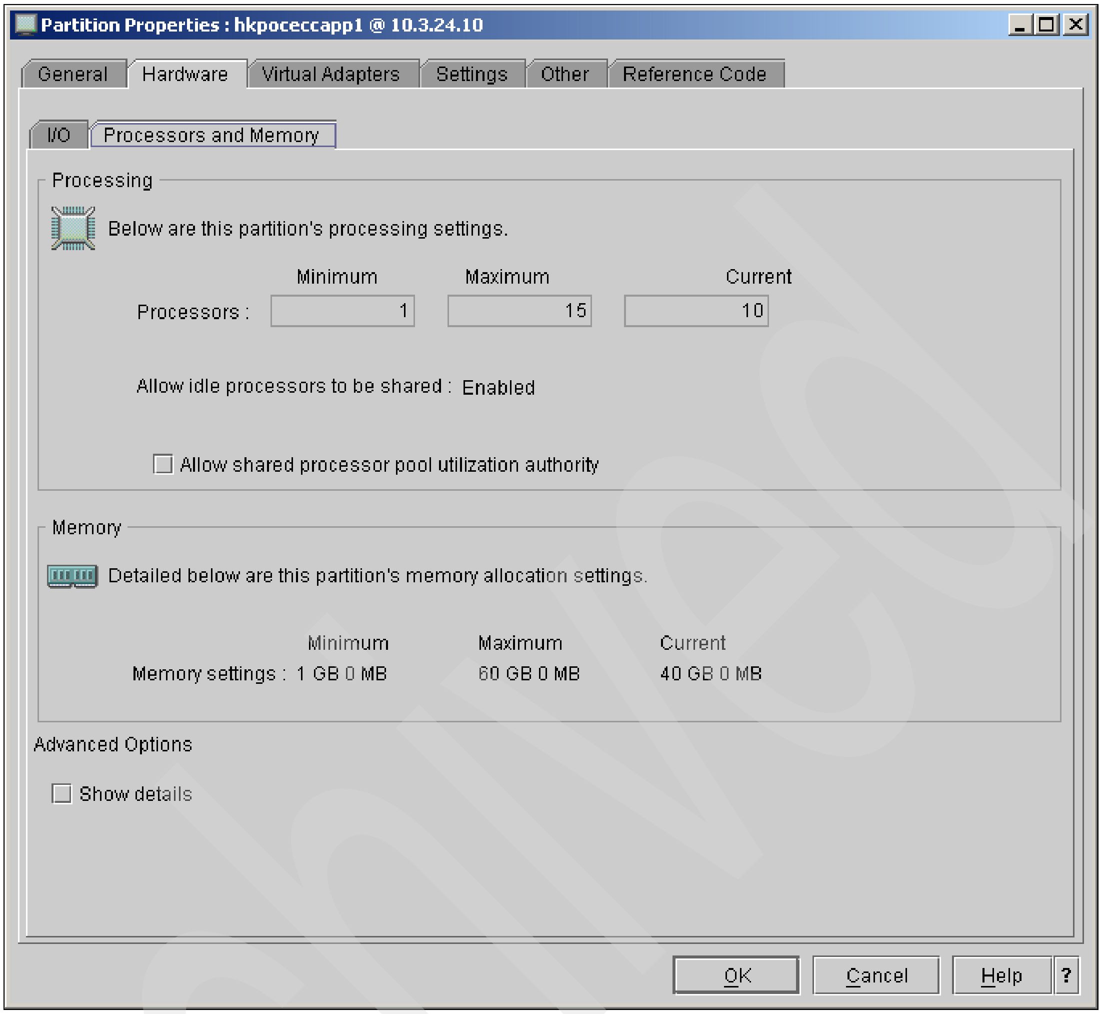

Figure 2-12 hkpoceccapp1 CPU and memory configuration

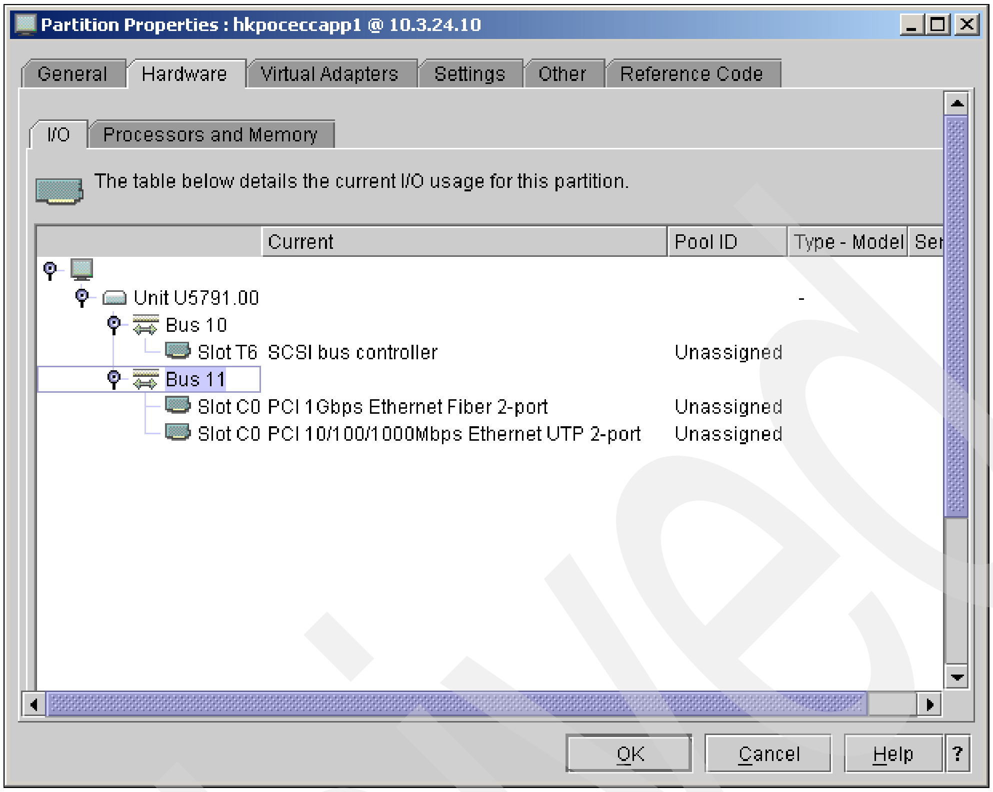

Figure 2-13 hkpoceccapp1 I/O configuration

For the AIX configuration, four internal SCSI disks are available. Each disk has 143 GB capacity. On this server only two disks are used, as shown in Figure 2-14. In this figure you can see also that two volume groups are defined on this LPAR, rootvg and d2lvg, for the SAP application components.

```
{hkpoceccapp1:root}/ -> lspv
hdisk0          00c2644fae867d01                    rootvg          active
hdisk1          00c2644fbb217f06                    rootvg          active
hdisk2          00c2644f851e52b4                    d2lvg           active
hdisk3          00c2644fbaf6a167                    d2lvg           active
{hkpoceccapp1:root}/ ->
```

Figure 2-14 hkpoceccapp1 assigned disks

Figure 2-15 displays the characteristics of one of the disks. We specified hdisk0 for this example.

```
{hkpoceccapp1:root}/ -> lspv hdisk0
PHYSICAL VOLUME:    hdisk0                   VOLUME GROUP:     rootvg
PV IDENTIFIER:      00c2644fae867d01 VG IDENTIFIER     00c2644f00004c0000000112ae869980
PV STATE:           active
STALE PARTITIONS:   0                        ALLOCATABLE:      yes
PP SIZE:            256 megabyte(s)          LOGICAL VOLUMES:  14
TOTAL PPs:          546 (139776 megabytes)   VG DESCRIPTORS:   2
FREE PPs:           46 (11776 megabytes)     HOT SPARE:        no
USED PPs:           500 (128000 megabytes)   MAX REQUEST:      256 kilobytes
FREE DISTRIBUTION:  00..00..00..00..46
USED DISTRIBUTION:  110..109..109..109..63
{hkpoceccapp1:root}/ ->
```

Figure 2-15 hdisk0 characteristics

The volume group rootvg has the characteristics shown in Figure 2-16.

```
{hkpoceccapp1:root}/ -> lsvg rootvg
VOLUME GROUP:       rootvg                   VG IDENTIFIER:  00c2644f00004c0000000112ae869980
VG STATE:           active                   PP SIZE:        256 megabyte(s)
VG PERMISSION:      read/write               TOTAL PPs:      1092 (279552 megabytes)
MAX LVs:            256                      FREE PPs:       92 (23552 megabytes)
LVs:                15                       USED PPs:       1000 (256000 megabytes)
OPEN LVs:           14                       QUORUM:         1
TOTAL PVs:          2                        VG DESCRIPTORS: 3
STALE PVs:          0                        STALE PPs:      0
ACTIVE PVs:         2                        AUTO ON:        yes
MAX PPs per VG:     32512
MAX PPs per PV:     1016                     MAX PVs:        32
LTG size (Dynamic): 256 kilobyte(s)          AUTO SYNC:      no
HOT SPARE:          no                       BB POLICY:      relocatable
{hkpoceccapp1:root}/ ->
```

Figure 2-16 rootvg characteristics

Several logical volumes have been defined on this rootvg volume group, as shown in Figure 2-17.

```
{hkpoceccapp1:root}/ -> lsvg -l rootvg
rootvg:
LV NAME             TYPE       LPs   PPs   PVs  LV STATE      MOUNT POINT
hd5                 boot       1     2     2    closed/syncd  N/A
hd6                 paging     64    128   2    open/syncd    N/A
hd8                 jfs2log    1     2     2    open/syncd    N/A
hd4                 jfs2       8     16    2    open/syncd    /
hd2                 jfs2       13    26    2    open/syncd    /usr
hd9var              jfs2       8     16    2    open/syncd    /var
hd3                 jfs2       8     16    2    open/syncd    /tmp
hd1                 jfs2       1     2     2    open/syncd    /home
hd10opt             jfs2       2     4     2    open/syncd    /opt
lg_dumplv           sysdump    8     16    2    open/syncd    N/A
fslv00              jfs2       1     2     2    open/syncd    /XmRec
fslv01              jfs2       1     2     2    open/syncd    /tmp/m2
lvsapd2l            jfs2       256   512   2    open/syncd    /sap/D2L
paging00            paging     128   128   1    open/syncd    N/A
paging01            paging     128   128   1    open/syncd    N/A
{hkpoceccapp1:root}/ ->
```

Figure 2-17 rootvg logical volumes

The volume group d2lvg has the characteristics shown in Figure 2-18.

```
{hkpoceccapp1:root}/ -> lsvg d2lvg
VOLUME GROUP:       d2lvg                    VG IDENTIFIER:  00c2644f00004c0000000115851e
VG STATE:           active                   PP SIZE:        128 megabyte(s)
VG PERMISSION:      read/write               TOTAL PPs:      2186 (279808 megabytes)
MAX LVs:            256                      FREE PPs:       986 (126208 megabytes)
LVs:                2                        USED PPs:       1200 (153600 megabytes)
OPEN LVs:           2                        QUORUM:         2
TOTAL PVs:          2                        VG DESCRIPTORS: 3
STALE PVs:          0                        STALE PPs:      0
ACTIVE PVs:         2                        AUTO ON:        yes
MAX PPs per VG:     32768                    MAX PVs:        1024
LTG size (Dynamic): 256 kilobyte(s)          AUTO SYNC:      no
HOT SPARE:          no                       BB POLICY:      relocatable
{hkpoceccapp1:root}/ ->
```

Figure 2-18 d2lvg characteristics

Several logical volumes have been defined on this d2lvg volume group, as shown in Figure 2-19.

```
{hkpoceccapp1:root}/ -> lsvg -l d2lvg
d2lvg:
LV NAME             TYPE       LPs     PPs     PVs  LV STATE      MOUNT POINT
d2llv               jfs2       400     400     2    open/syncd    /usr/sap/D2L/DVEBMGS00
sapmntlv            jfs2       800     800     2    open/syncd    /sapmnt/D2Lfranck
{hkpoceccapp1:root}/ -> _
```

Figure 2-19   d2lvg logical volumes

The logical volume d2llv is defined to host the core SAP application, so for performance purposes, the Montpellier team made the decision to stripe it, as shown in Figure 2-20.

```
{hkpoceccapp1:root}/ -> lslv d2llv
LOGICAL VOLUME:     d2llv                  VOLUME GROUP:   d2lvg
LV IDENTIFIER:      00c2644f00004c0000000115851e552e.1 PERMISSION:     read/write
VG STATE:           active/complete        LV STATE:       opened/syncd
TYPE:               jfs2                   WRITE VERIFY:   off
MAX LPs:            9999                   PP SIZE:        128 megabyte(s)
COPIES:             1                      SCHED POLICY:   striped
LPs:                400                    PPs:            400
STALE PPs:          0                      BB POLICY:      relocatable
INTER-POLICY:       maximum                RELOCATABLE:    no
INTRA-POLICY:       middle                 UPPER BOUND:    4
MOUNT POINT:        /usr/sap/D2L/DVEBMGS00 LABEL:          /usr/sap/D2L/DVEBMGS00
MIRROR WRITE CONSISTENCY: on/ACTIVE
EACH LP COPY ON A SEPARATE PV ?: yes (superstrict)
Serialize IO ?:     NO
STRIPE WIDTH:       4
STRIPE SIZE:        1m
DEVICESUBTYPE : DS_LVZ

{hkpoceccapp1:root}/ ->
```

Figure 2-20   d2llv characteristics

Mounted file systems are the necessary ones for the operating system and for the SAP dialog instance, as you can see in Figure 2-21. Further installation information about SAP dialogue instances is available on the Web at `http://www.service.sap.com/instguides` and `http://www.service.sap.com/notes`. Note that the NFS mounted file systems are not relevant to this project.

```
{hkpoceccapp1:root}/ -> mount
  node       mounted        mounted over    vfs       date        options
-------- --------------- --------------- ------ ------------ ---------------
         /dev/hd4        /               jfs2   Sep 27 13:29 rw,log=/dev/hd8
         /dev/hd2        /usr            jfs2   Sep 27 13:29 rw,log=/dev/hd8
         /dev/hd9var     /var            jfs2   Sep 27 13:29 rw,log=/dev/hd8
         /dev/hd3        /tmp            jfs2   Sep 27 13:30 rw,log=/dev/hd8
         /dev/hd1        /home           jfs2   Sep 27 13:30 rw,log=/dev/hd8
         /proc           /proc           procfs Sep 27 13:30 rw
         /dev/hd10opt    /opt            jfs2   Sep 27 13:30 rw,log=/dev/hd8
         /dev/fslv00     /XmRec          jfs2   Sep 27 13:30 rw,log=/dev/hd8
         /dev/fslv01     /tmp/m2         jfs2   Sep 27 13:30 rw,log=/dev/hd8
         /dev/lvsapd2l   /sap/D2L        jfs2   Sep 27 13:30 rw,log=/dev/hd8
hkpoclive1 /tmp/FRANCK     /mnt            nfs3   Sep 27 13:30 bg,hard,intr,sec=sys,rw
hkpoclive2 /sapmnt/D3B/exe /mnt            nfs3   Sep 27 13:38
{hkpoceccapp1:root}/ ->
```

Figure 2-21   Mounted file systems

Disk space used by those file systems is shown in Figure 2-22.

```
{hkpoceccapp1:root}/ -> df -g
Filesystem    GB blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4           2.00      1.89    6%    12880     3% /
/dev/hd2           3.25      0.13   97%    54066    60% /usr
/dev/hd9var        2.00      1.98    1%      537     1% /var
/dev/hd3           2.00      1.65   18%      802     1% /tmp
/dev/hd1           0.25      0.20   22%      322     1% /home
/proc                 -         -    -         -     -  /proc
/dev/hd10opt       0.50      0.26   49%     6269    10% /opt
/dev/fslv00        0.25      0.25    1%        4     1% /XmRec
/dev/fslv01        0.25      0.24    3%       30     1% /tmp/m2
/dev/lvsapd2l     64.00     36.33   44%    53225     1% /sap/D2L
hkpoclive1:/tmp/FRANCK     64.00     46.44   28%    79573     1% /mnt
hkpoclive2:/sapmnt/D3B/exe     64.00     46.44   28%    79573     1% /mnt
{hkpoceccapp1:root}/ ->
```

Figure 2-22 File systems disk usage

This LPAR exports two file systems (kernel and transport) used across the SAP landscape, as shown in Figure 2-23.

```
{hkpoceccapp1:root}/ -> exportfs -v
/sap/D2L/trans  -sec=sys:krb5p:krb5i:krb5:dh,rw,root=hkpoceccapp2:hkpoceccapp3:hkpoclive1:hkpoclive2
/sap/D2L/sapmnt -sec=sys:krb5p:krb5i:krb5:dh,rw,root=hkpoceccapp2:hkpoceccapp3
{hkpoceccapp1:root}/ ->
```

Figure 2-23 Exported file systems

Note that the transport file system (/sap/D2L/trans) is exported to the whole SAP system landscape (ECC and SCM), whereas the kernel file system (/sapmnt/D2L/sapmnt) is only exported to the other ECC dialog instances.

The paging area is defined as shown in Figure 2-24.

```
{hkpoceccapp1:root}/ -> lsps -a
Page Space      Physical Volume   Volume Group    Size %Used Active  Auto  Type
paging01        hdisk1            rootvg       32768MB     1   yes   yes    lv
paging00        hdisk0            rootvg       32768MB     1   yes   yes    lv
hd6             hdisk0            rootvg       16384MB     1   yes   yes    lv
{hkpoceccapp1:root}/ -> lsps -a
```

Figure 2-24 Paging area definition

Regarding the network configuration, we have two adapters defined (plus loopback) with two IP addresses corresponding to the VLANs we described previously, as shown in Figure 2-25 on page 38.

```
{hkpoceccapp1:root}/sapmnt/D2L/exe -> ifconfig -a
en0: flags=5e080863,c0<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,CHECKSUM_OFFL
        inet 1.1.75.11 netmask 0xffffff00 broadcast 1.1.75.255
         tcp_sendspace 131072 tcp_recvspace 65536
en2: flags=5e080863,c0<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,CHECKSUM_OFFL
        inet 10.3.24.11 netmask 0xffffff00 broadcast 10.3.24.255
         tcp_sendspace 131072 tcp_recvspace 65536
lo0: flags=e08084b<UP,BROADCAST,LOOPBACK,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT>
        inet 127.0.0.1 netmask 0xff000000 broadcast 127.255.255.255
        inet6 ::1/0
         tcp_sendspace 131072 tcp_recvspace 131072 rfc1323 1
{hkpoceccapp1:root}/sapmnt/D2L/exe ->
```

Figure 2-25 TCP/IP configuration

We have now completely described the configuration of the LPAR hkpoceccapp1. This configuration is standard for running an SAP instance.

LPAR hkpoceccapp2 configuration

This LPAR holds the SAP ECC dialog instance. This LPAR has 10 CPUs and 40 GB RAM assigned, as shown in Figure 2-26.

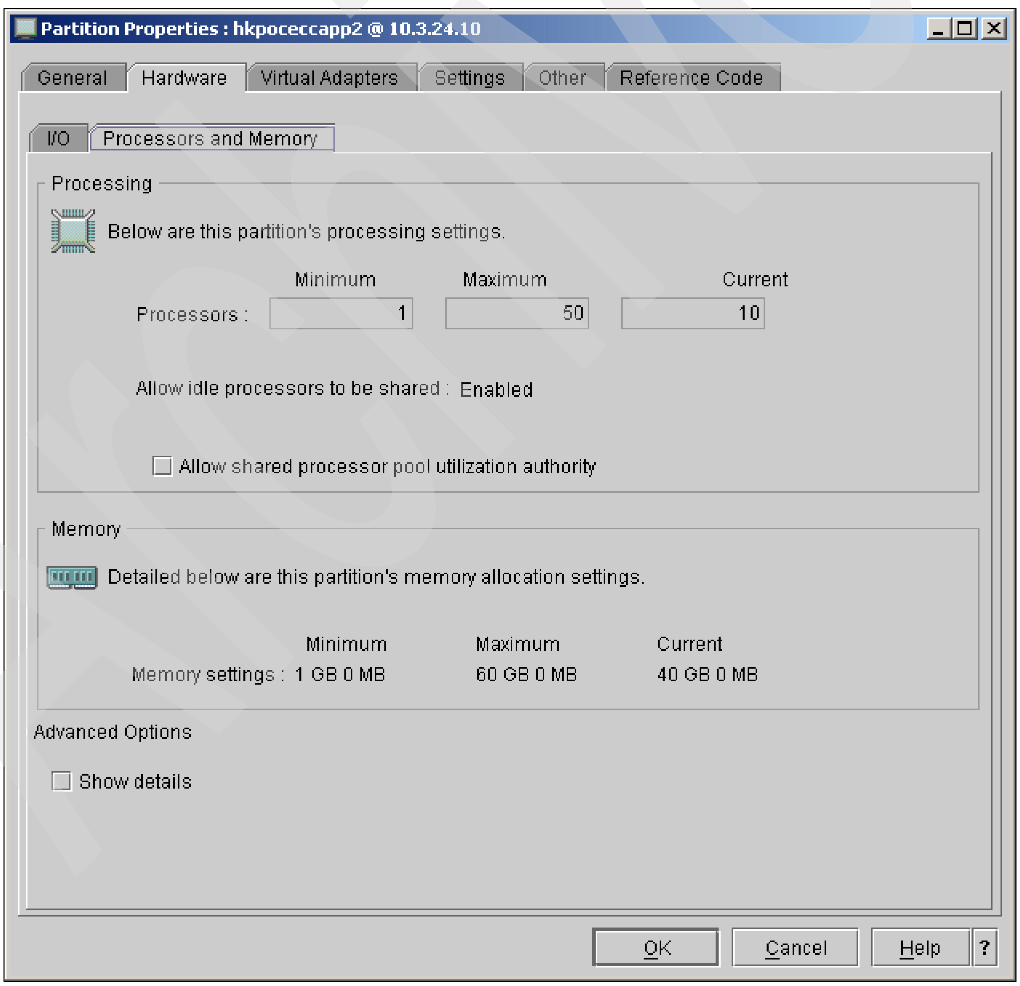

Figure 2-26 hkpoceccapp2 CPU and memory configuration

As shown in Figure 2-27, this LPAR has the following I/O usage:

- One SCSI controller with the local disks
- One PCI slot with one Ethernet Fiber card with two ports at 1 Gbps
- One PCI slot with one Ethernet 10/100/1000 card with two ports

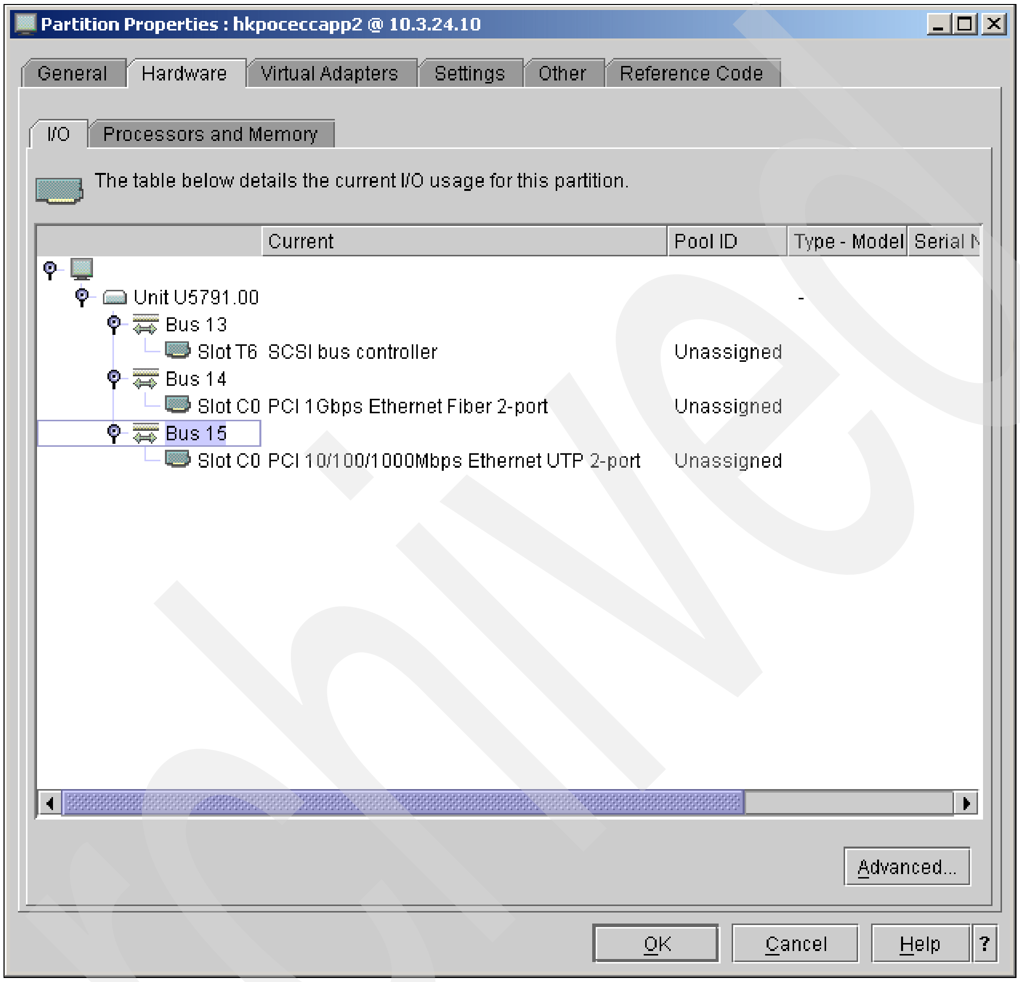

Figure 2-27 hkpoceccapp2 I/O configuration

For the AIX configuration level, four internal SCSI disks are available, each with a capacity of 143 GB. On this server all disks are used as shown in Figure 2-28. This figure also shows that the only volume group defined in this LPAR is rootvg.

```
{hkpoceccapp2:root}/ -> lspv
hdisk0          00c2644fbccfe786                    rootvg          active
hdisk1          00c2644fbccfe9f2                    rootvg          active
hdisk2          00c2644fb28f9729                    rootvg          active
hdisk3          00c2644f6f27133a                    rootvg          active
{hkpoceccapp2:root}/ ->
```

Figure 2-28 hkpoceccapp2 assigned disks

Figure 2-29 displays the characteristics of one of the disks. We specified hdisk0 for this example.

```
{hkpoceccapp2:root}/ -> lspv hdisk0
PHYSICAL VOLUME:     hdisk0                   VOLUME GROUP:     rootvg
PV IDENTIFIER:       00c2644fbccfe786 VG IDENTIFIER     00c2644f00004c0000000112b28fa1bd
PV STATE:            active
STALE PARTITIONS:    0                        ALLOCATABLE:      yes
PP SIZE:             256 megabyte(s)          LOGICAL VOLUMES:  14
TOTAL PPs:           546 (139776 megabytes)   VG DESCRIPTORS:   1
FREE PPs:            46 (11776 megabytes)     HOT SPARE:        no
USED PPs:            500 (128000 megabytes)   MAX REQUEST:      256 kilobytes
FREE DISTRIBUTION:   16..00..00..00..30
USED DISTRIBUTION:   94..109..109..109..79
{hkpoceccapp2:root}/ ->
```

Figure 2-29 hdisk0 characteristics

The volume group rootvg has the characteristics shown in Figure 2-30.

```
{hkpoceccapp2:root}/ -> lsvg rootvg
VOLUME GROUP:       rootvg                   VG IDENTIFIER:  00c2644f00004c0000000112b28fa1bd
VG STATE:           active                   PP SIZE:        256 megabyte(s)
VG PERMISSION:      read/write               TOTAL PPs:      2184 (559104 megabytes)
MAX LVs:            256                      FREE PPs:       672 (172032 megabytes)
LVs:                16                       USED PPs:       1512 (387072 megabytes)
OPEN LVs:           15                       QUORUM:         1
TOTAL PVs:          4                        VG DESCRIPTORS: 4
STALE PVs:          0                        STALE PPs:      0
ACTIVE PVs:         4                        AUTO ON:        no
MAX PPs per VG:     32512
MAX PPs per PV:     1016                     MAX PVs:        32
LTG size (Dynamic): 256 kilobyte(s)          AUTO SYNC:      no
HOT SPARE:          no                       BB POLICY:      relocatable
{hkpoceccapp2:root}/ ->
```

Figure 2-30 rootvg characteristics

Several logical volumes have been defined on this rootvg volume group, as shown in Figure 2-31.

```
{hkpoceccapp2:root}/ -> lsvg -l rootvg
rootvg:
LV NAME             TYPE       LPs    PPs    PVs  LV STATE      MOUNT POINT
hd5                 boot       1      2      2    closed/syncd  N/A
hd6                 paging     64     128    2    open/syncd    N/A
hd8                 jfs2log    1      2      2    open/syncd    N/A
hd4                 jfs2       8      16     2    open/syncd    /
hd2                 jfs2       13     26     2    open/syncd    /usr
hd9var              jfs2       8      16     2    open/syncd    /var
hd3                 jfs2       8      16     2    open/syncd    /tmp
hd1                 jfs2       1      2      2    open/syncd    /home
hd10opt             jfs2       2      4      2    open/syncd    /opt
lg_dumplv           sysdump    8      16     2    open/syncd    N/A
fslv00              jfs2       1      2      2    open/syncd    /XmRec
fslv01              jfs2       1      2      2    open/syncd    /tmp/m2
lvsapd21            jfs2       256    512    2    open/syncd    /sap/D2L
paging00            paging     128    128    1    open/syncd    N/A
paging01            paging     128    128    1    open/syncd    N/A
paging02            paging     256    512    2    open/syncd    N/A
{hkpoceccapp2:root}/ ->
```

Figure 2-31 rootvg logical volumes

Mounted file systems are the necessary ones for the operating system and for the SAP dialog instance, as shown in Figure 2-32.

```
(hkpoceccapp2:root)/ -> mount
  node       mounted        mounted over    vfs       date        options
-------- ---------------  ---------------  ------ ------------ ---------------
         /dev/hd4         /                jfs2   Oct 08 10:25 rw,log=/dev/hd8
         /dev/hd2         /usr             jfs2   Oct 08 10:25 rw,log=/dev/hd8
         /dev/hd9var      /var             jfs2   Oct 08 10:25 rw,log=/dev/hd8
         /dev/hd3         /tmp             jfs2   Oct 08 10:25 rw,log=/dev/hd8
         /dev/hd1         /home            jfs2   Oct 08 10:25 rw,log=/dev/hd8
         /proc            /proc            procfs Oct 08 10:25 rw
         /dev/hd10opt     /opt             jfs2   Oct 08 10:25 rw,log=/dev/hd8
         /dev/fslv00      /XmRec           jfs2   Oct 08 10:25 rw,log=/dev/hd8
         /dev/fslv01      /tmp/m2          jfs2   Oct 08 10:25 rw,log=/dev/hd8
         /dev/lvsapd2l    /sap/D2L         jfs2   Oct 08 10:25 rw,log=/dev/hd8
hkpoceccapp1 /sap/D2L/sapmnt  /sapmnt/D2L      nfs3   Oct 08 10:26 bg,hard,intr,sec=sys:dh:krb5:krb5i
hkpoceccapp1 /sap/D2L/trans   /usr/sap/trans   nfs3   Oct 08 10:26 bg,hard,intr,sec=sys:dh:krb5:krb5i
hkpoclive1 /tmp/FRANCK      /mnt             nfs3   Oct 08 10:26 bg,hard,intr,sec=sys,rw
(hkpoceccapp2:root)/ ->
```

Figure 2-32   Mounted file systems

In this LPAR, we mount the SAP kernel and the transport file system from LPAR hkpoceccapp1. Other NFS mounted file systems are not relevant for this project.

Disk space used by the file systems is shown in Figure 2-33.

```
(hkpoceccapp2:root)/ -> df -g
Filesystem    GB blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4           2.00      1.94    4%    12847     3% /
/dev/hd2           3.25      0.13   96%    53894    59% /usr
/dev/hd9var        2.00      1.98    1%      528     1% /var
/dev/hd3           2.00      1.34   33%      394     1% /tmp
/dev/hd1           0.25      0.19   23%      269     1% /home
/proc                 -         -    -         -     -  /proc
/dev/hd10opt       0.50      0.26   49%     6269    10% /opt
/dev/fslv00        0.25      0.25    1%        5     1% /XmRec
/dev/fslv01        0.25      0.24    3%       30     1% /tmp/m2
/dev/lvsapd2l     64.00     55.16   14%    23883     1% /sap/D2L
hkpoceccapp1:/sap/D2L/sapmnt     64.00     29.00   55%    54605     1% /sapmnt/D2L
hkpoceccapp1:/sap/D2L/trans     64.00     29.00   55%    54605     1% /usr/sap/trans
hkpoclive1:/tmp/FRANCK      5.50      0.66   88%     3348     3% /mnt
(hkpoceccapp2:root)/ ->
```

Figure 2-33   File systems disk usage

The paging area is defined as shown in Figure 2-34.

```
(hkpoceccapp2:root)/ -> lsps -a
Page Space      Physical Volume   Volume Group    Size %Used Active  Auto  Type
paging02        hdisk2            rootvg       65536MB     1    yes   yes    lv
paging01        hdisk1            rootvg       32768MB     1    yes   yes    lv
paging00        hdisk0            rootvg       32768MB     1    yes   yes    lv
hd6             hdisk0            rootvg       16384MB     1    yes   yes    lv
(hkpoceccapp2:root)/ ->
```

Figure 2-34   Paging area definition

Regarding the network configuration, we have two adapters defined (plus rollback) with two IP addresses corresponding to the VLANs we described previously, as shown in Figure 2-35.

```
(hkpoceccapp2:root)/ -> ifconfig -a
en0: flags=5e080863,c0<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,CHECKSUM_OFFL(
        inet 1.1.75.12 netmask 0xffffff00 broadcast 1.1.75.255
         tcp_sendspace 131072 tcp_recvspace 65536
en2: flags=5e080863,c0<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,CHECKSUM_OFFL(
        inet 10.3.24.12 netmask 0xffffff00 broadcast 10.3.24.255
         tcp_sendspace 131072 tcp_recvspace 65536
lo0: flags=e08084b<UP,BROADCAST,LOOPBACK,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT>
        inet 127.0.0.1 netmask 0xff000000 broadcast 127.255.255.255
        inet6 ::1/0
         tcp_sendspace 131072 tcp_recvspace 131072 rfc1323 1
(hkpoceccapp2:root)/ ->
```

Figure 2-35 TCP/IP configuration

We have now completely described the configuration of the LPAR hkpoceccapp2. This configuration is standard for running an SAP instance.

LPAR hkpoceccapp3 configuration

This LPAR holds the SAP ECC dialog instance. This LPAR has 10 CPUs and 40 GB RAM assigned, as shown in Figure 2-36.

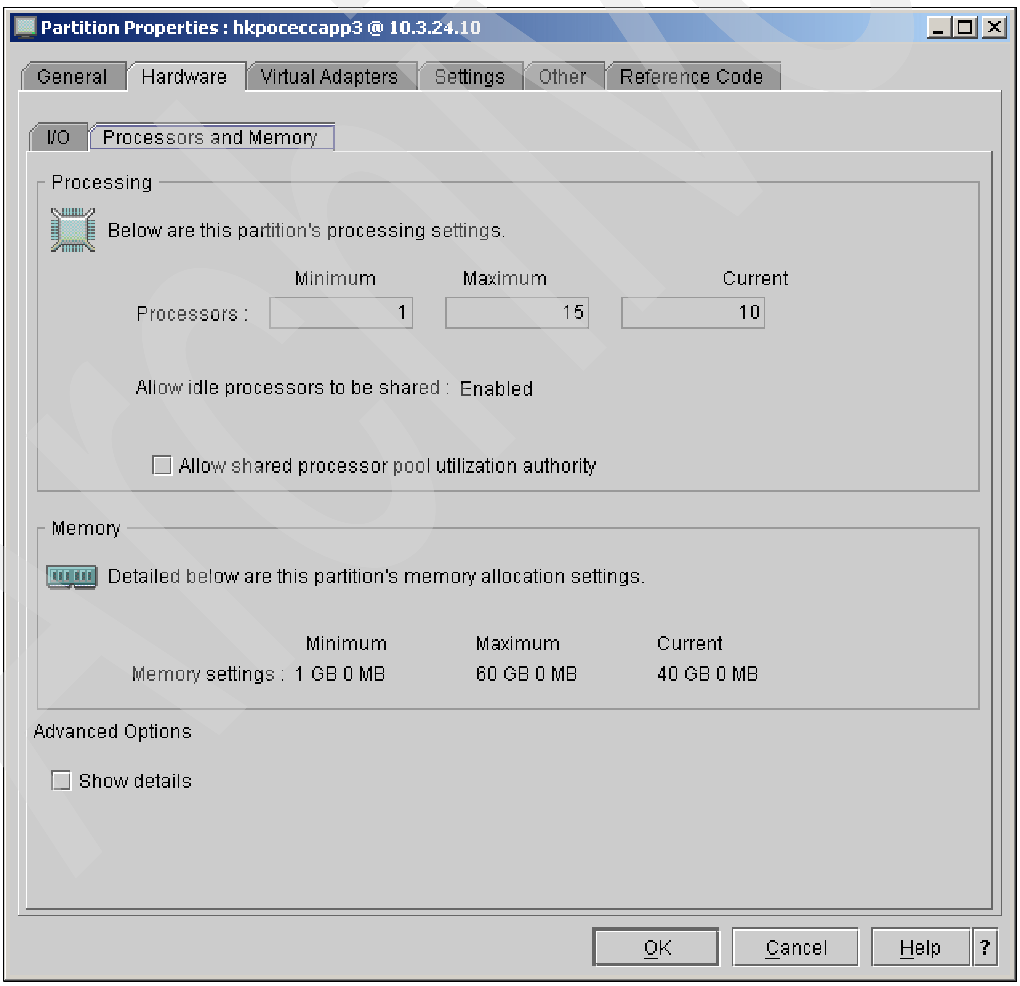

Figure 2-36 hkpoceccapp3 CPU and memory configuration

As shown in Figure 2-37, this LPAR has the following I/O assignments:

- 1 SCSI controller with the local disks
- 1 PCI slot with 1 Ethernet Fiber card with 2 ports at 1 Gbps
- 1 PCI slot with 1 Ethernet 10/100/1000 card with 2 ports

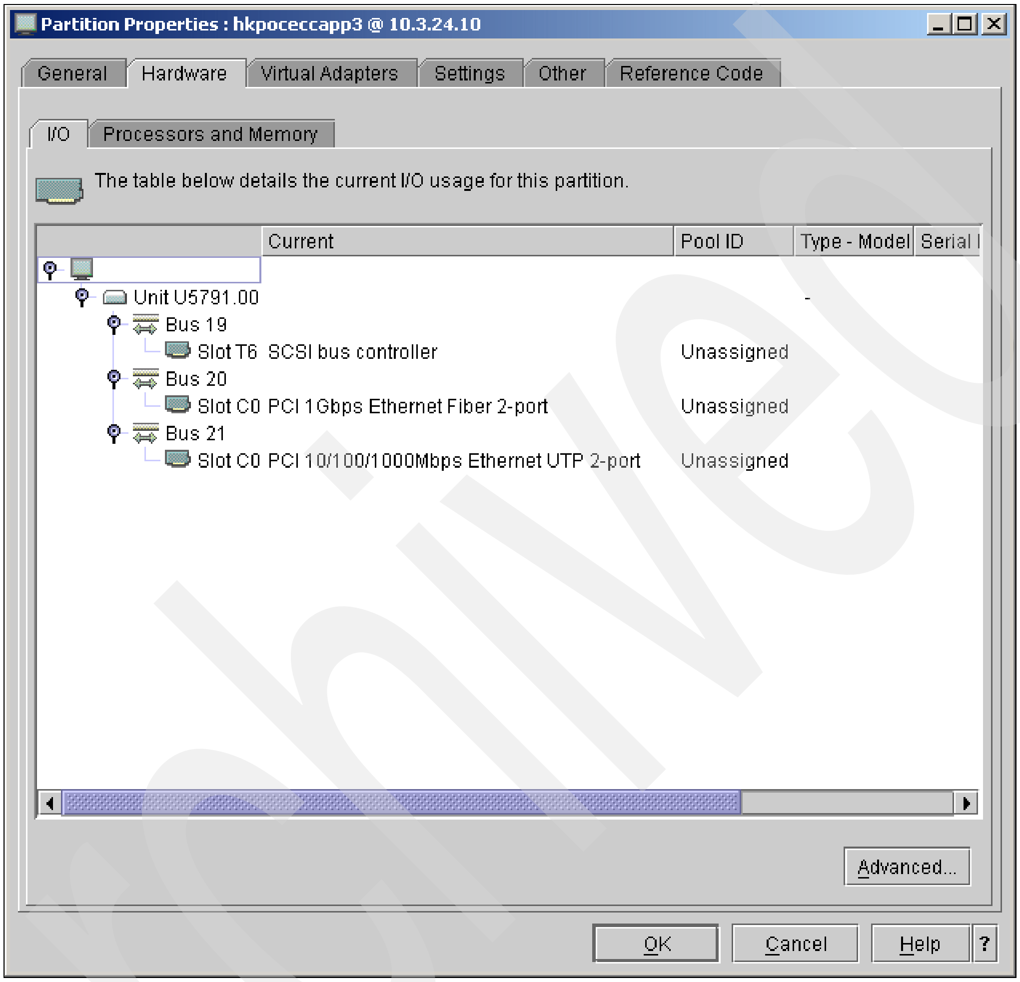

Figure 2-37 hkpoceccapp3 I/O configuration

At AIX configuration level, four internal SCSI 143 GB disks are available. Each disk has 143 GB capacity. On this server only two disks are used, as shown in Figure 2-38. This figure also shows that the only volume group defined in this LPAR is rootvg.

```
{hkpoceccapp3:root}/ -> lspv
hdisk0          00c2644fbcd557da                    rootvg          active
hdisk1          00c2644fbcd55a46                    rootvg          active
hdisk2          none                                None
hdisk4          00c2644fb2ba2b54                    None
{hkpoceccapp3:root}/ ->
```

Figure 2-38 hkpoceccapp3 assigned disks

Figure 2-39 displays the characteristics of one of the disks. We specified hdisk0 for this example.

```
{hkpoceccapp3:root}/ -> lspv hdisk0
PHYSICAL VOLUME:     hdisk0                   VOLUME GROUP:     rootvg
PV IDENTIFIER:       00c2644fbcd557da VG IDENTIFIER     00c2644f00004c0000000112b2ba452b
PV STATE:            active
STALE PARTITIONS:    0                        ALLOCATABLE:      yes
PP SIZE:             256 megabyte(s)          LOGICAL VOLUMES:  14
TOTAL PPs:           546 (139776 megabytes)   VG DESCRIPTORS:   2
FREE PPs:            46 (11776 megabytes)     HOT SPARE:        no
USED PPs:            500 (128000 megabytes)   MAX REQUEST:      256 kilobytes
FREE DISTRIBUTION:   00..00..00..00..46
USED DISTRIBUTION:   110..109..109..109..63
{hkpoceccapp3:root}/ ->
```

Figure 2-39 hdisk0 characteristics

The volume group rootvg has the characteristics shown in Figure 2-40.

```
{hkpoceccapp3:root}/ -> lsvg rootvg
VOLUME GROUP:        rootvg                   VG IDENTIFIER:  00c2644f00004c0000000112b2ba452b
VG STATE:            active                   PP SIZE:        256 megabyte(s)
VG PERMISSION:       read/write               TOTAL PPs:      1092 (279552 megabytes)
MAX LVs:             256                      FREE PPs:       92 (23552 megabytes)
LVs:                 15                       USED PPs:       1000 (256000 megabytes)
OPEN LVs:            14                       QUORUM:         1
TOTAL PVs:           2                        VG DESCRIPTORS: 3
STALE PVs:           0                        STALE PPs:      0
ACTIVE PVs:          2                        AUTO ON:        no
MAX PPs per VG:      32512
MAX PPs per PV:      1016                     MAX PVs:        32
LTG size (Dynamic): 256 kilobyte(s)           AUTO SYNC:      no
HOT SPARE:           no                       BB POLICY:      relocatable
{hkpoceccapp3:root}/ ->
```

Figure 2-40 rootvg characteristics

Several logical volumes have been defined on this rootvg volume group, as shown in Figure 2-41.

```
{hkpoceccapp3:root}/ -> lsvg -l rootvg
rootvg:
LV NAME             TYPE       LPs    PPs    PVs  LV STATE       MOUNT POINT
hd5                 boot       1      2      2    closed/syncd   N/A
hd6                 paging     64     128    2    open/syncd     N/A
hd8                 jfs2log    1      2      2    open/syncd     N/A
hd4                 jfs2       8      16     2    open/syncd     /
hd2                 jfs2       13     26     2    open/syncd     /usr
hd9var              jfs2       8      16     2    open/syncd     /var
hd3                 jfs2       8      16     2    open/syncd     /tmp
hd1                 jfs2       1      2      2    open/syncd     /home
hd10opt             jfs2       2      4      2    open/syncd     /opt
lg_dumplv           sysdump    8      16     2    open/syncd     N/A
fslv00              jfs2       1      2      2    open/syncd     /XmRec
fslv01              jfs2       1      2      2    open/syncd     /tmp/m2
lvsapd21            jfs2       256    512    2    open/syncd     /sap/D2L
paging00            paging     128    128    1    open/syncd     N/A
paging01            paging     128    128    1    open/syncd     N/A
{hkpoceccapp3:root}/ ->
```

Figure 2-41 rootvg logical volumes

Mounted file systems are the necessary ones for the operating system and for the SAP dialog instance, as shown in Figure 2-42.

```
(hkpoceccapp3:root)/ -> mount
  node       mounted        mounted over    vfs       date        options
-------- ---------------  ---------------  ------ ------------ ---------------
         /dev/hd4         /                jfs2   Oct 08 10:22 rw,log=/dev/hd8
         /dev/hd2         /usr             jfs2   Oct 08 10:22 rw,log=/dev/hd8
         /dev/hd9var      /var             jfs2   Oct 08 10:22 rw,log=/dev/hd8
         /dev/hd3         /tmp             jfs2   Oct 08 10:22 rw,log=/dev/hd8
         /dev/hd1         /home            jfs2   Oct 08 10:23 rw,log=/dev/hd8
         /proc            /proc            procfs Oct 08 10:23 rw
         /dev/hd10opt     /opt             jfs2   Oct 08 10:23 rw,log=/dev/hd8
         /dev/fslv00      /XmRec           jfs2   Oct 08 10:23 rw,log=/dev/hd8
         /dev/fslv01      /tmp/m2          jfs2   Oct 08 10:23 rw,log=/dev/hd8
         /dev/lvsapd21    /sap/D2L         jfs2   Oct 08 10:23 rw,log=/dev/hd8
hkpoceccapp1 /sap/D2L/sapmnt  /sapmnt/D2L      nfs3   Oct 08 10:23 bg,hard,intr,sec=sys:dh:krb5:krb5i
hkpoceccapp1 /sap/D2L/trans   /usr/sap/trans   nfs3   Oct 08 10:23 bg,hard,intr,sec=sys:dh:krb5:krb5i
hkpoclive1 /tmp/FRANCK      /mnt             nfs3   Oct 08 10:23 bg,hard,intr,sec=sys,rw
(hkpoceccapp3:root)/ ->
```

Figure 2-42   Mounted file systems

In this LPAR, we mount the SAP kernel and the transport file system from LPAR hkpoceccapp1. Other NFS mounted file systems are not relevant for this project.

Disk space used by those file systems is shown in Figure 2-43.

```
(hkpoceccapp3:root)/ -> df -g
Filesystem    GB blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4           2.00      1.94    3%    12837     3% /
/dev/hd2           3.25      0.13   96%    53883    59% /usr
/dev/hd9var        2.00      1.98    1%      522     1% /var
/dev/hd3           2.00      1.27   37%      505     1% /tmp
/dev/hd1           0.25      0.20   22%      268     1% /home
/proc                 -         -     -        -     -  /proc
/dev/hd10opt       0.50      0.26   49%     6269    10% /opt
/dev/fslv00        0.25      0.25    1%        4     1% /XmRec
/dev/fslv01        0.25      0.24    3%       30     1% /tmp/m2
/dev/lvsapd21     64.00     52.41   19%    23942     1% /sap/D2L
hkpoceccapp1:/sap/D2L/sapmnt     64.00     29.99   54%    54604     1% /sapmnt/D2L
hkpoceccapp1:/sap/D2L/trans     64.00     29.99   54%    54604     1% /usr/sap/trans
hkpoclive1:/tmp/FRANCK      5.50      0.66   88%     3348     3% /mnt
(hkpoceccapp3:root)/ ->
```

Figure 2-43   File systems disk usage

The paging area is defined as shown in Figure 2-44.

```
(hkpoceccapp3:root)/ -> lsps -a
Page Space      Physical Volume   Volume Group    Size %Used Active  Auto  Type
paging01        hdisk1            rootvg       32768MB     1   yes   yes    lv
paging00        hdisk0            rootvg       32768MB     1   yes   yes    lv
hd6             hdisk0            rootvg       16384MB     1   yes   yes    lv
(hkpoceccapp3:root)/ ->
```

Figure 2-44   Paging area definition

Regarding the network configuration, we have two adapters defined (plus loopback) with two IP addresses corresponding to the VLANs we described previously, as shown in Figure 2-45.

```
(hkpoceccapp3:root)/ -> ifconfig -a
en0: flags=5e080863,c0<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,CHECKSUM_OFFL(
        inet 1.1.75.13 netmask 0xffffff00 broadcast 1.1.75.255
         tcp_sendspace 131072 tcp_recvspace 65536
en2: flags=5e080863,c0<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,CHECKSUM_OFFL(
        inet 10.3.24.13 netmask 0xffffff00 broadcast 10.3.24.255
         tcp_sendspace 131072 tcp_recvspace 65536
lo0: flags=e08084b<UP,BROADCAST,LOOPBACK,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT>
        inet 127.0.0.1 netmask 0xff000000 broadcast 127.255.255.255
        inet6 ::1/0
         tcp_sendspace 131072 tcp_recvspace 131072 rfc1323 1
(hkpoceccapp3:root)/ ->
```

Figure 2-45 TCP/IP configuration

We have now completely described the configuration of the LPAR hkpoceccapp3. This configuration is standard for running an SAP dialog instance.

MCP17M-595-SN51263CF configuration

As shown in Figure 2-46, this System p595 server is configured with 64 cores at 2.3 GHz, as mentioned in 2.1.4, “Operational model - Physical architecture” on page 22.

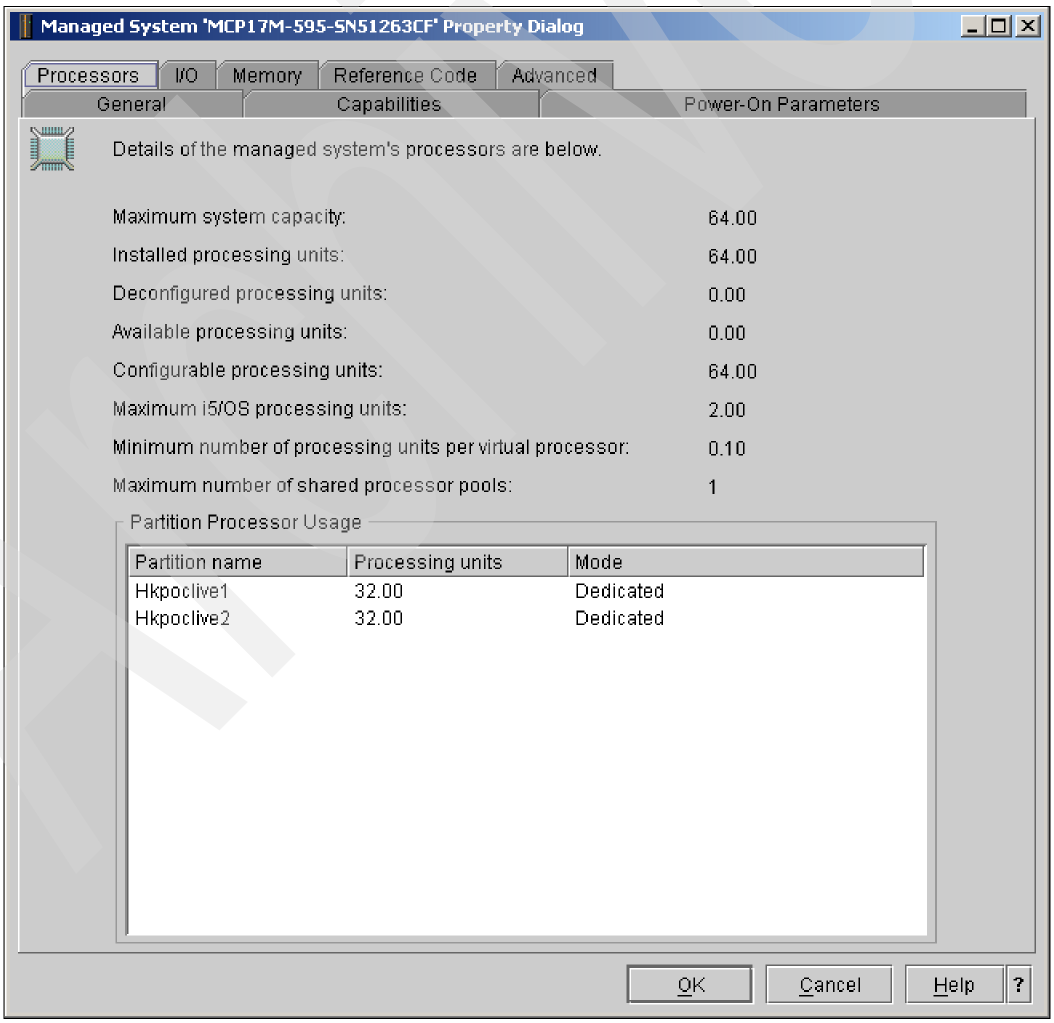

Figure 2-46 MCP17M-595-SN51263CF CPU configuration

Regarding memory, this server is configured with 256 GB RAM, as shown in Figure 2-47.

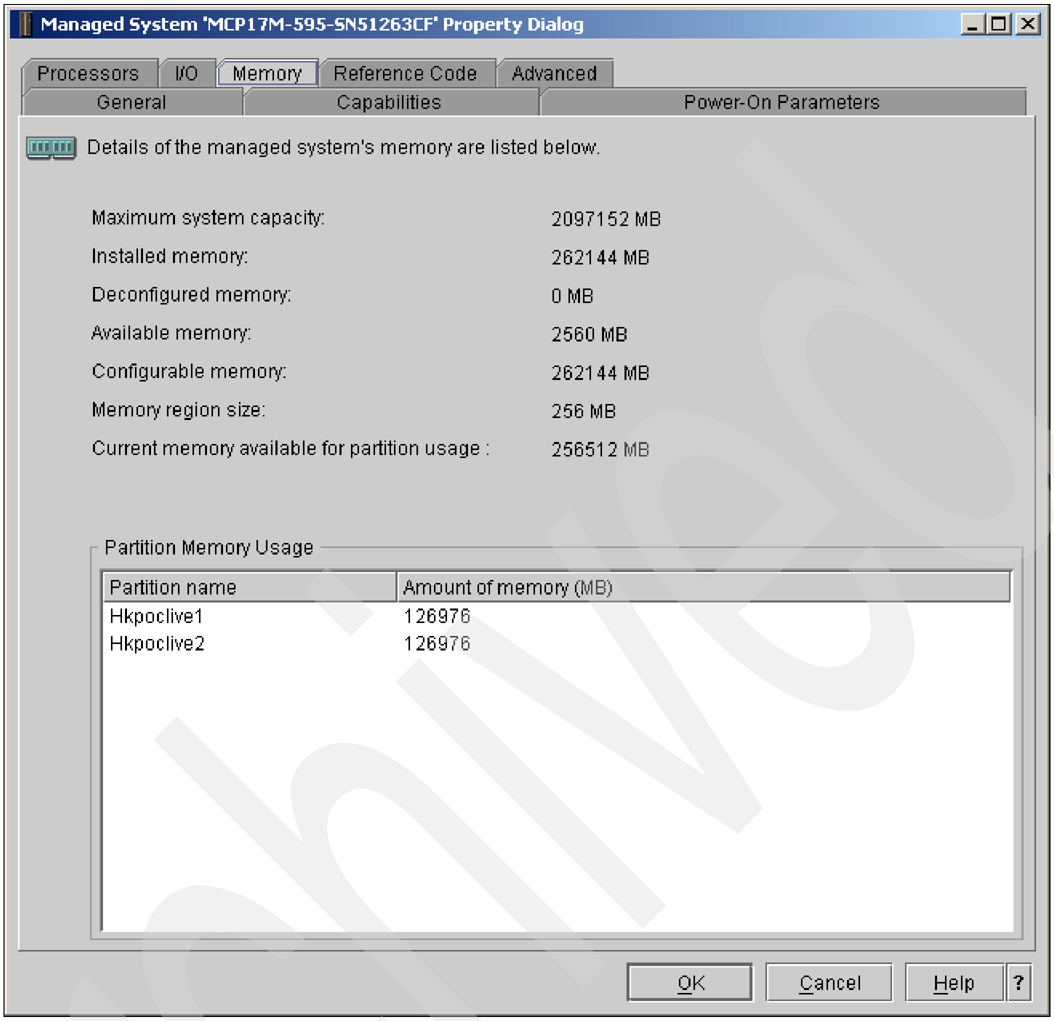

Figure 2-47 MCP17M-595-SN51263CF memory configuration

No virtualization is used on this server; all resources are assigned to capped LPARs.

LPAR hkpoclive1 configuration

This LPAR holds the SAP liveCache instance. This liveCache instance consists of a MaxDB database instance. This LPAR has 32 CPUs and 124 GB RAM assigned, as shown in Figure 2-48.

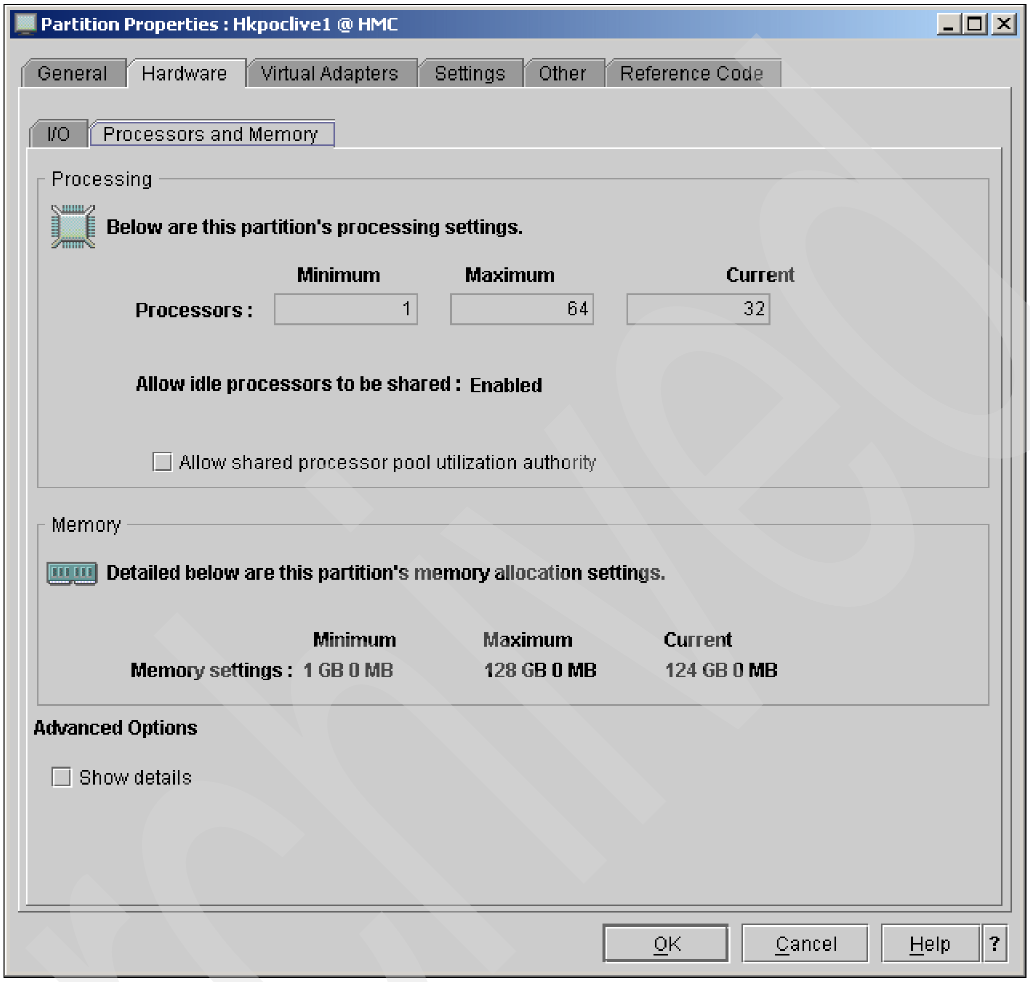

Figure 2-48 hkpoclive1 CPU and memory configuration

As shown in Figure 2-49, this LPAR has the following I/O assignments:

- Three SCSI controllers with the local disks
- Three PCI slots with Ethernet Fiber cards with two ports at 1Gbps or Ethernet 10/100/1000 card with two ports
- Four slots with HBAs

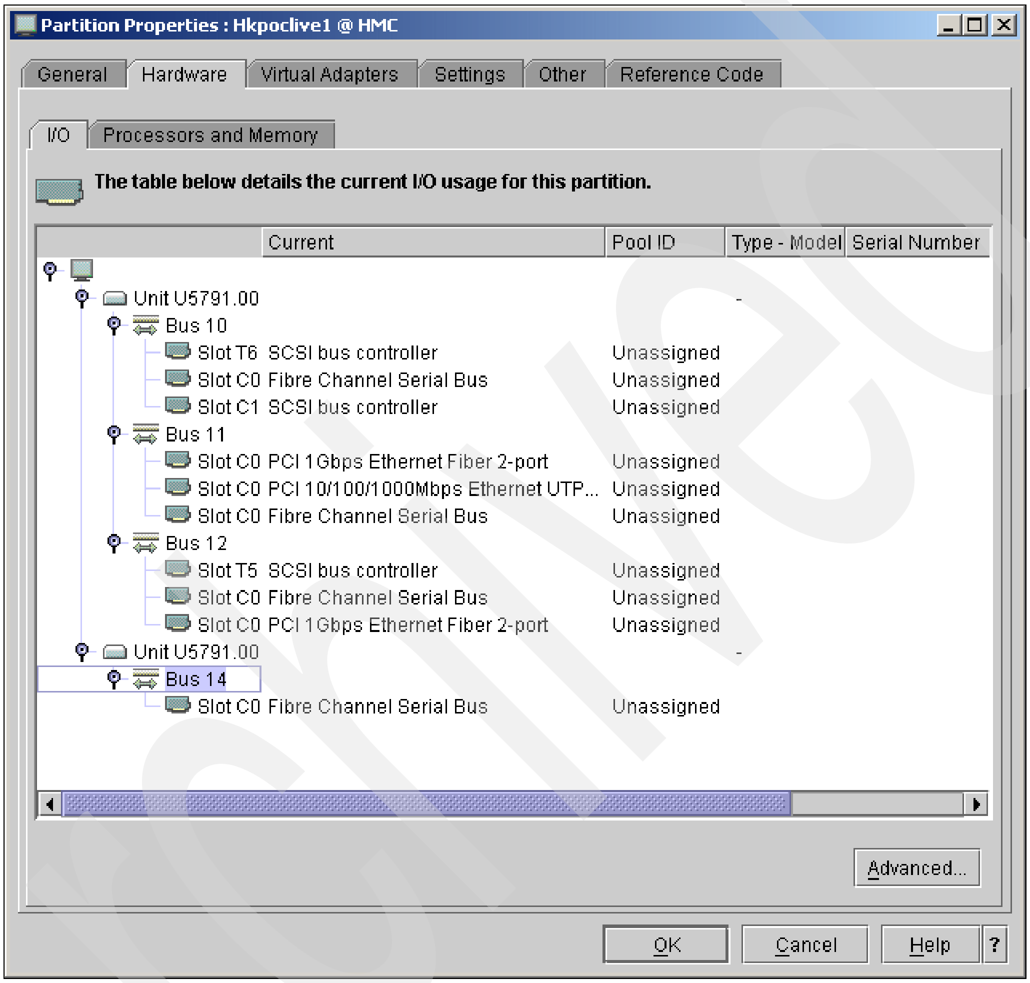

Figure 2-49 hkpoclive1 I/O configuration

At AIX configuration level, eight internal SCSI 143 GB disks are available. Each disk has 143 GB capacity. This LPAR has some DS8300 disks assigned as well. What we see in Figure 2-50 are not physical disks, but AIX recognizes them as such. We discuss the storage configuration details in 2.4, “Storage configuration” on page 61. This LPAR is the only one with storage-attached disks in the System p server environment.

```
{hkpoclive1:root}/ -> lsdev -Cc disk
hdisk0  Available 0C-08-00-8,0  16 Bit LVD SCSI Disk Drive
hdisk1  Available 0C-08-00-9,0  16 Bit LVD SCSI Disk Drive
hdisk2  Available 0C-08-00-10,0 16 Bit LVD SCSI Disk Drive
hdisk3  Available 0C-08-00-11,0 16 Bit LVD SCSI Disk Drive
hdisk4  Available 0K-08-00-8,0  16 Bit LVD SCSI Disk Drive
hdisk5  Available 0K-08-00-9,0  16 Bit LVD SCSI Disk Drive
hdisk6  Available 0K-08-00-10,0 16 Bit LVD SCSI Disk Drive
hdisk7  Available 0K-08-00-11,0 16 Bit LVD SCSI Disk Drive
hdisk8  Available 0D-08-02      IBM MPIO FC 2107
hdisk9  Available 0D-08-02      IBM MPIO FC 2107
hdisk10 Available 0L-08-02      IBM MPIO FC 2107
hdisk11 Available 0E-08-02      IBM MPIO FC 2107
hdisk12 Available 0E-08-02      IBM MPIO FC 2107
hdisk13 Available 0E-08-02      IBM MPIO FC 2107
hdisk14 Available 0D-08-02      IBM MPIO FC 2107
hdisk15 Available 0D-08-02      IBM MPIO FC 2107
hdisk16 Available 0D-08-02      IBM MPIO FC 2107
hdisk17 Available 0D-08-02      IBM MPIO FC 2107
{hkpoclive1:root}/ ->
```

Figure 2-50   hkpoclive1 disk configuration

It is also interesting to see which disks are assigned to which volume groups. This is shown in Figure 2-51.

```
{hkpoclive1:root}/ -> lspv
hdisk1          00c263cfb455ae7c                    rootvg          active
hdisk2          00c263cfb438ebdc                    sdblcpvg1       active
hdisk3          00c263cfb438ee63                    sdblcpvg1       active
hdisk0          00c263cf9f25d154                    rootvg          active
hdisk4          00c263cfb438f0f8                    sdblcpvg1       active
hdisk5          00c263cfb438f37f                    sdblcpvg1       active
hdisk6          00c263cfb460e2df                    sdblcpvg1       active
hdisk7          00c263cfb7d2cf06                    sdblcpvg1       active
hdisk8          00c263cfbb8665f5                    sdblcpvg        active
hdisk9          00c263cfbb86688c                    sdblcpvg        active
hdisk10         00c263cfbb866b25                    sdblcpvg        active
hdisk11         00c263cfbb866db8                    sdblcpvg        active
hdisk12         00c263cfbb867045                    sdblcpvg        active
hdisk13         00c263cfbb8672d1                    sdblcpvg        active
hdisk14         00c263cf7eb251d2                    psscsaplogvg    active
hdisk15         00c263cf7eb25463                    psscsaplogvg    active
hdisk16         00c263cf7eb256f2                    psscsaplogvg    active
hdisk17         00c263cf7eb2597f                    psscsaplogvg    active
{hkpoclive1:root}/ ->
```

Figure 2-51   Disk assigned to volume groups

Figure 2-52 on page 51 displays the characteristics of one of the disks. We specified hdisk0 for this example.

```
{hkpoclive1:root}/ -> lspv hdisk0
PHYSICAL VOLUME:     hdisk0                   VOLUME GROUP:     rootvg
PV IDENTIFIER:       00c263cf9f25d154 VG IDENTIFIER     00c263cf00004c00000001129f25f3bb
PV STATE:            active
STALE PARTITIONS:    0                        ALLOCATABLE:      yes
PP SIZE:             256 megabyte(s)          LOGICAL VOLUMES:  13
TOTAL PPs:           546 (139776 megabytes)   VG DESCRIPTORS:   2
FREE PPs:            12 (3072 megabytes)      HOT SPARE:        no
USED PPs:            534 (136704 megabytes)   MAX REQUEST:      256 kilobytes
FREE DISTRIBUTION:   00..00..00..00..12
USED DISTRIBUTION:   110..109..109..109..97
{hkpoclive1:root}/ ->
```

Figure 2-52 hdisk0 characteristics

Other physical devices available include Ethernet cards and HBAs. As you can see in Figure 2-53, this LPAR has four FC adapters available.

```
{hkpoclive1:root}/ -> lsdev -C |grep fcs
fcs0          Available 0D-08         FC Adapter
fcs1          Available 0E-08         FC Adapter
fcs2          Available 0L-08         FC Adapter
fcs3          Available 0U-08         FC Adapter
{hkpoclive1:root}/ ->
```

Figure 2-53 hkpoclive1 FC adapters

Regarding Ethernet adapters, hkpoclive1 has four 2-port GB adapters and two 2-port 10/100/1000 Base T adapters, as shown on Figure 2-54.

```
{hkpoclive1:root}/ -> lsdev -C |grep ent
ent0          Available 0G-08         2-Port Gigabit Ethernet-SX PCI-X Adapter (14108802)
ent1          Available 0G-09         2-Port Gigabit Ethernet-SX PCI-X Adapter (14108802)
ent2          Defined   0N-08         2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent3          Defined   0N-09         2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent4          Defined   0H-08         2-Port Gigabit Ethernet-SX PCI-X Adapter (14108802)
ent5          Defined   0H-09         2-Port Gigabit Ethernet-SX PCI-X Adapter (14108802)
ent6          Available 0I-08         2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent7          Available 0I-09         2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902)
ent8          Available 0J-08         2-Port Gigabit Ethernet-SX PCI-X Adapter (14108802)
ent9          Available 0J-09         2-Port Gigabit Ethernet-SX PCI-X Adapter (14108802)
{hkpoclive1:root}/ ->
```

Figure 2-54 hkpoclive1 Ethernet adapters

Regarding the network configuration, only two adapters are defined (plus loopback) with two IP addresses corresponding to the VLANs we described previously, as shown in Figure 2-55.

```
{hkpoclive1:root}/ -> ifconfig -a
en0: flags=5e080863,c0<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,CHECKSUM_OFFL
        inet 1.1.75.3 netmask 0xffffff00 broadcast 1.1.75.255
         tcp_sendspace 131072 tcp_recvspace 65536
en8: flags=5e080863,c0<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,CHECKSUM_OFFL
        inet 10.3.24.3 netmask 0xffffff00 broadcast 10.3.24.255
         tcp_sendspace 131072 tcp_recvspace 65536
lo0: flags=e08084b<UP,BROADCAST,LOOPBACK,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT>
        inet 127.0.0.1 netmask 0xff000000 broadcast 127.255.255.255
        inet6 ::1/0
         tcp_sendspace 131072 tcp_recvspace 131072 rfc1323 1
{hkpoclive1:root}/ ->
```

Figure 2-55 TCP/IP configuration

As shown in Figure 2-56, the following volume groups exist on this LPAR:

- rootvg
- sdblcpvg1: Not used in this configuration
- sdblcpvg: Volume group created for all SAP-related file systems, except logs
- psscsaplogvg: Volume group dedicated to logging structure

```
{hkpoclive1:root}/ -> lsvg
rootvg
sdblcpvg1
sdblcpvg
psscsaplogvg
{hkpoclive1:root}/ ->
```

Figure 2-56 hkpoclive1 volume groups

The volume group rootvg is on internal disks and has the characteristics shown in Figure 2-57.

```
{hkpoclive1:root}/ -> lsvg rootvg
VOLUME GROUP:       rootvg                   VG IDENTIFIER:  00c263cf00004c00000001129f25f3bb
VG STATE:           active                   PP SIZE:        256 megabyte(s)
VG PERMISSION:      read/write               TOTAL PPs:      1092 (279552 megabytes)
MAX LVs:            256                      FREE PPs:       24 (6144 megabytes)
LVs:                13                       USED PPs:       1068 (273408 megabytes)
OPEN LVs:           11                       QUORUM:         1
TOTAL PVs:          2                        VG DESCRIPTORS: 3
STALE PVs:          0                        STALE PPs:      0
ACTIVE PVs:         2                        AUTO ON:        yes
MAX PPs per VG:     32512
MAX PPs per PV:     1016                     MAX PVs:        32
LTG size (Dynamic): 256 kilobyte(s)          AUTO SYNC:      no
HOT SPARE:          no                       BB POLICY:      relocatable
{hkpoclive1:root}/ ->
```

Figure 2-57 rootvg characteristics

Several logical volumes have been defined on this rootvg volume group, as shown in Figure 2-58.

```
{hkpoclive1:root}/ -> lsvg -l rootvg
rootvg:
LV NAME             TYPE       LPs   PPs   PVs  LV STATE      MOUNT POINT
hd5                 boot       1     2     2    closed/syncd  N/A
hd6                 paging     8     16    2    open/syncd    N/A
hd8                 jfs2log    1     2     2    open/syncd    N/A
hd4                 jfs2       8     16    2    open/syncd    /
hd2                 jfs2       17    34    2    open/syncd    /usr
hd9var              jfs2       8     16    2    open/syncd    /var
hd3                 jfs2       22    44    2    open/syncd    /tmp
hd1                 jfs2       1     2     2    open/syncd    /home
hd10opt             jfs2       2     4     2    open/syncd    /opt
lg_dumplv           sysdump    8     16    2    closed/syncd  N/A
fslv00              jfs2       1     2     2    open/syncd    /XmRec
fslv01              jfs2       1     2     2    open/syncd    /tmp/m2
lvsapd3b            jfs2       456   912   2    open/syncd    /sap/D3B
{hkpoclive1:root}/ ->
```

Figure 2-58 rootvg logical volumes

The volume group sdblcpvg on the storage subsystem has the characteristics shown in Figure 2-59.

```
{hkpoclive1:root}/ -> lsvg sdblcpvg
VOLUME GROUP:       sdblcpvg                 VG IDENTIFIER:  00c263cf00004c0000000113bb867f61
VG STATE:           active                   PP SIZE:        16 megabyte(s)
VG PERMISSION:      read/write               TOTAL PPs:      57570 (921120 megabytes)
MAX LVs:            256                      FREE PPs:       0 (0 megabytes)
LVs:                7                        USED PPs:       57570 (921120 megabytes)
OPEN LVs:           7                        QUORUM:         4
TOTAL PVs:          6                        VG DESCRIPTORS: 6
STALE PVs:          0                        STALE PPs:      0
ACTIVE PVs:         6                        AUTO ON:        yes
MAX PPs per VG:     65536                    MAX PVs:        1024
LTG size (Dynamic): 256 kilobyte(s)          AUTO SYNC:      no
HOT SPARE:          no                       BB POLICY:      relocatable
{hkpoclive1:root}/ ->
```

Figure 2-59 sapdblcpvg characteristics

Several logical volumes have been defined on this sdblcpvg volume group, as shown in Figure 2-60.

```
{hkpoclive1:root}/ -> lsvg -l sdblcpvg
sdblcpvg:
LV NAME             TYPE       LPs   PPs   PVs  LV STATE      MOUNT POINT
jfs2log_log         jfs2log    4     4     1    open/syncd    N/A
jfs2log_data        jfs2log    4     4     1    open/syncd    N/A
sapdb_data1         jfs2       11648 11648 6    open/syncd    /sapdb/LCP/sapdata1
sapdb_data2         jfs2       11648 11648 6    open/syncd    /sapdb/LCP/sapdata2
sapdb_data3         jfs2       11648 11648 6    open/syncd    /sapdb/LCP/sapdata3
sapdb_data4         jfs2       11648 11648 6    open/syncd    /sapdb/LCP/sapdata4
sapdb_log           jfs2       10970 10970 6    open/syncd    /sapdb/LCP/saplog
{hkpoclive1:root}/ ->
```

Figure 2-60 sapdblcpvg logical volumes

The volume group psscsaplogvg has the characteristics shown in Figure 2-61.

```
{hkpoclive1:root}/ -> lsvg psscsaplogvg
VOLUME GROUP:       psscsaplogvg             VG IDENTIFIER:  00c263cf00004c00000001157ebf832b
VG STATE:           active                   PP SIZE:        32 megabyte(s)
VG PERMISSION:      read/write               TOTAL PPs:      7540 (241280 megabytes)
MAX LVs:            256                      FREE PPs:       3 (96 megabytes)
LVs:                2                        USED PPs:       7537 (241184 megabytes)
OPEN LVs:           2                        QUORUM:         3
TOTAL PVs:          4                        VG DESCRIPTORS: 4
STALE PVs:          0                        STALE PPs:      0
ACTIVE PVs:         4                        AUTO ON:        yes
MAX PPs per VG:     32768                    MAX PVs:        1024
LTG size (Dynamic): 256 kilobyte(s)          AUTO SYNC:      no
HOT SPARE:          no                       BB POLICY:      relocatable
{hkpoclive1:root}/ ->
```

Figure 2-61 psscsaplogvg characteristics

Several logical volumes have been defined on this psscsaplogvg volume group, as shown in Figure 2-62 on page 54.

```
{hkpoclive1:root}/ -> lsvg -l psscsaplogvg
psscsaplogvg:
LV NAME             TYPE       LPs   PPs   PVs  LV STATE      MOUNT POINT
psscsaploglv        jfs2       7536  7536  4    open/syncd    /psscsaplog
loglv00             jfs2log    1     1     1    open/syncd    N/A
{hkpoclive1:root}/ ->
```

Figure 2-62   psscsaplogvg logical volumes

Mounted file systems are the necessary ones for the operating system and for the SAP liveCache instance, as shown in Figure 2-63. Further installation information about SAP instances can be found on the Web at http://www.service.sap.com/instguides and http://www.service.sap.com/notes. Note that the picasso NFS mounted file system is used for performance statistics collection and is not directly related to this project.

```
{hkpoclive1:root}/ -> mount
  node       mounted        mounted over    vfs       date        options
-------- ---------------  ---------------  ------ ------------ ---------------
         /dev/hd4         /                jfs2   Oct 08 16:06 rw,log=/dev/hd8
         /dev/hd2         /usr             jfs2   Oct 08 16:06 rw,log=/dev/hd8
         /dev/hd9var      /var             jfs2   Oct 08 16:06 rw,log=/dev/hd8
         /dev/hd3         /tmp             jfs2   Oct 08 16:06 rw,log=/dev/hd8
         /dev/hd1         /home            jfs2   Oct 08 16:07 rw,log=/dev/hd8
         /proc            /proc            procfs Oct 08 16:07 rw
         /dev/hd10opt     /opt             jfs2   Oct 08 16:07 rw,log=/dev/hd8
         /dev/fslv00      /XmRec           jfs2   Oct 08 16:07 rw,log=/dev/hd8
         /dev/fslv01      /tmp/m2          jfs2   Oct 08 16:07 rw,log=/dev/hd8
         /dev/lvsapd3b    /sap/D3B         jfs2   Oct 08 16:07 rw,log=/dev/hd8
         /dev/sapdb_data1 /sapdb/LCP/sapdata1 jfs2   Oct 08 16:07 rw,cio,log=/dev/jfs2log_log
         /dev/sapdb_data2 /sapdb/LCP/sapdata2 jfs2   Oct 08 16:07 rw,cio,log=/dev/jfs2log_data
         /dev/sapdb_data3 /sapdb/LCP/sapdata3 jfs2   Oct 08 16:07 rw,cio,log=/dev/jfs2log_data
         /dev/sapdb_data4 /sapdb/LCP/sapdata4 jfs2   Oct 08 16:07 rw,cio,log=/dev/jfs2log_data
         /dev/sapdb_log   /sapdb/LCP/saplog jfs2   Oct 08 16:07 rw,cio,log=/dev/jfs2log_data
         /dev/psscsaploglv /psscsaplog      jfs2   Oct 08 16:07 rw,log=/dev/loglv00
hkpoceccapp1 /sap/D2L/trans   /usr/sap/trans    nfs3   Oct 08 16:07 bg,hard,intr,sec=sys:dh:krb5:krb5i
hkpoclive2 /sap/D3B/sapmnt   /sapmnt/D3B      nfs3   Oct 08 16:10 bg,hard,intr,sec=sys:dh:krb5:krb5i:k
picasso-onn /Results          /Results         nfs3   Oct 08 16:14 bg,hard,intr,sec=sys,rw
{hkpoclive1:root}/ ->
```

Figure 2-63   Mounted file systems

Disk space used by those file systems is shown in Figure 2-64.

```
{hkpoclive1:root}/ -> df -g
Filesystem    GB blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4           2.00      1.89    6%    12955     3% /
/dev/hd2           4.25      1.13   74%    53908    17% /usr
/dev/hd9var        2.00      1.94    4%      585     1% /var
/dev/hd3           5.50      0.66   88%     3347     3% /tmp
/dev/hd1           0.25      0.20   19%      381     1% /home
/proc                 -         -    -         -     -  /proc
/dev/hd10opt       0.50      0.23   54%     6313    11% /opt
/dev/fslv00        0.25      0.25    1%        4     1% /XmRec
/dev/fslv01        0.25      0.24    3%       31     1% /tmp/m2
/dev/lvsapd3b    114.00    110.56    4%    26453     1% /sap/D3B
/dev/sapdb_data1     182.00     84.75   54%        7     1% /sapdb/LCP/sapdata1
/dev/sapdb_data2     182.00     61.42   67%      226     1% /sapdb/LCP/sapdata2
/dev/sapdb_data3     182.00    132.43   28%        6     1% /sapdb/LCP/sapdata3
/dev/sapdb_data4     182.00     86.60   53%        6     1% /sapdb/LCP/sapdata4
/dev/sapdb_log     171.41     73.19   58%      112     1% /sapdb/LCP/saplog
/dev/psscsaploglv    235.50    139.13   41%      263     1% /psscsaplog
hkpoceccapp1:/sap/D2L/trans      64.00     31.79   51%    54767     1% /usr/sap/trans
hkpoclive2:/sap/D3B/sapmnt      64.00     31.04   52%    83493     2% /sapmnt/D3B
picasso-onn:/Results    136.25    126.10    8%    11409     1% /Results
{hkpoclive1:root}/ ->
```

Figure 2-64   File systems disk usage

The paging area is defined as shown in Figure 2-65.

```
{hkpoclive1:root}/ -> lsps -a
Page Space      Physical Volume   Volume Group    Size %Used Active  Auto  Type
hd6             hdisk0            rootvg        2048MB     8   yes   yes    lv
{hkpoclive1:root}/ ->
```

Figure 2-65 Paging area definition

We have now completely described the configuration of the LPAR hkpoclive1. This configuration is standard for running an SAP liveCache instance.

LPAR hkpoclive2 configuration

This LPAR holds the SAP SCM central system. This LPAR has 32 CPUs and 124 GB RAM assigned, as shown in Figure 2-66.

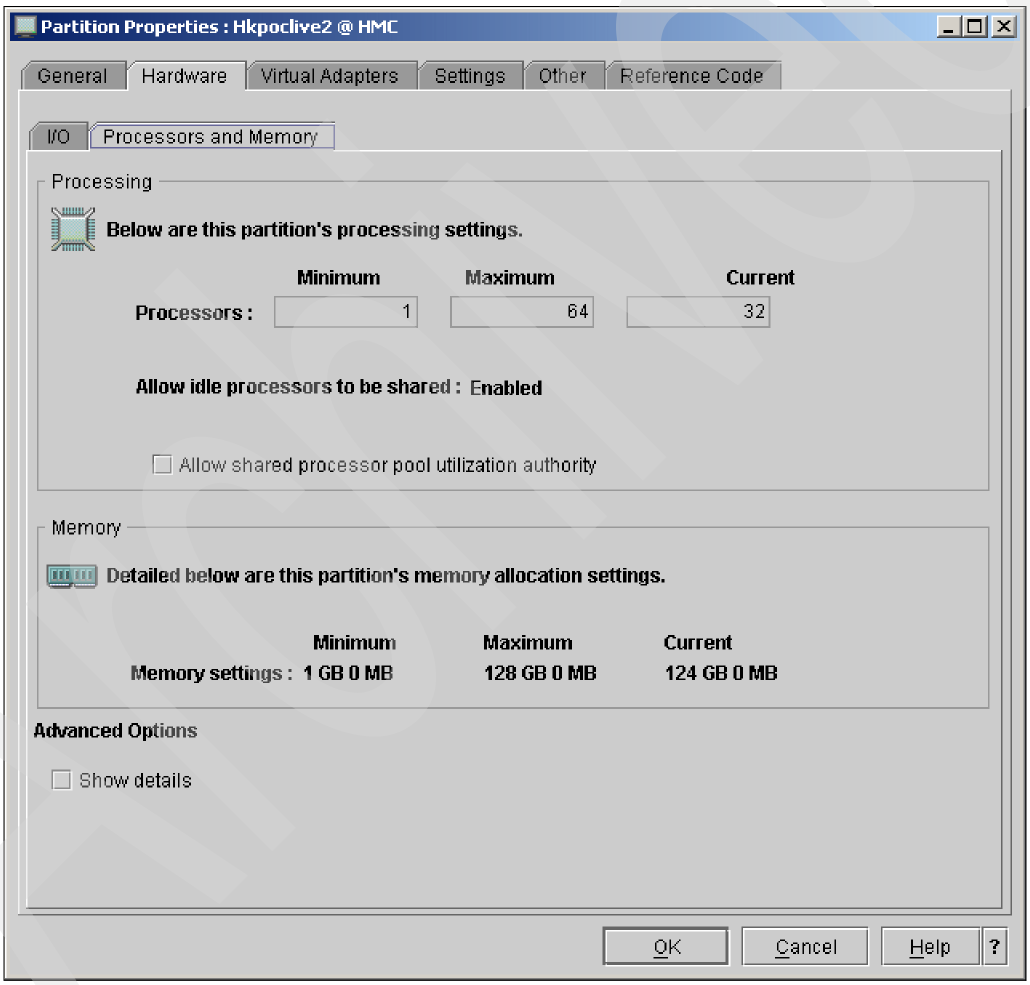

Figure 2-66 hkpoclive2 CPU and memory configuration

As shown in Figure 2-67, this LPAR has:

- One SCSI controller with the local disks
- Two PCI slots with Ethernet Fiber cards with two ports at 1 Gbps or Ethernet 10/100/1000 card with two ports
- Four slots with HBAs

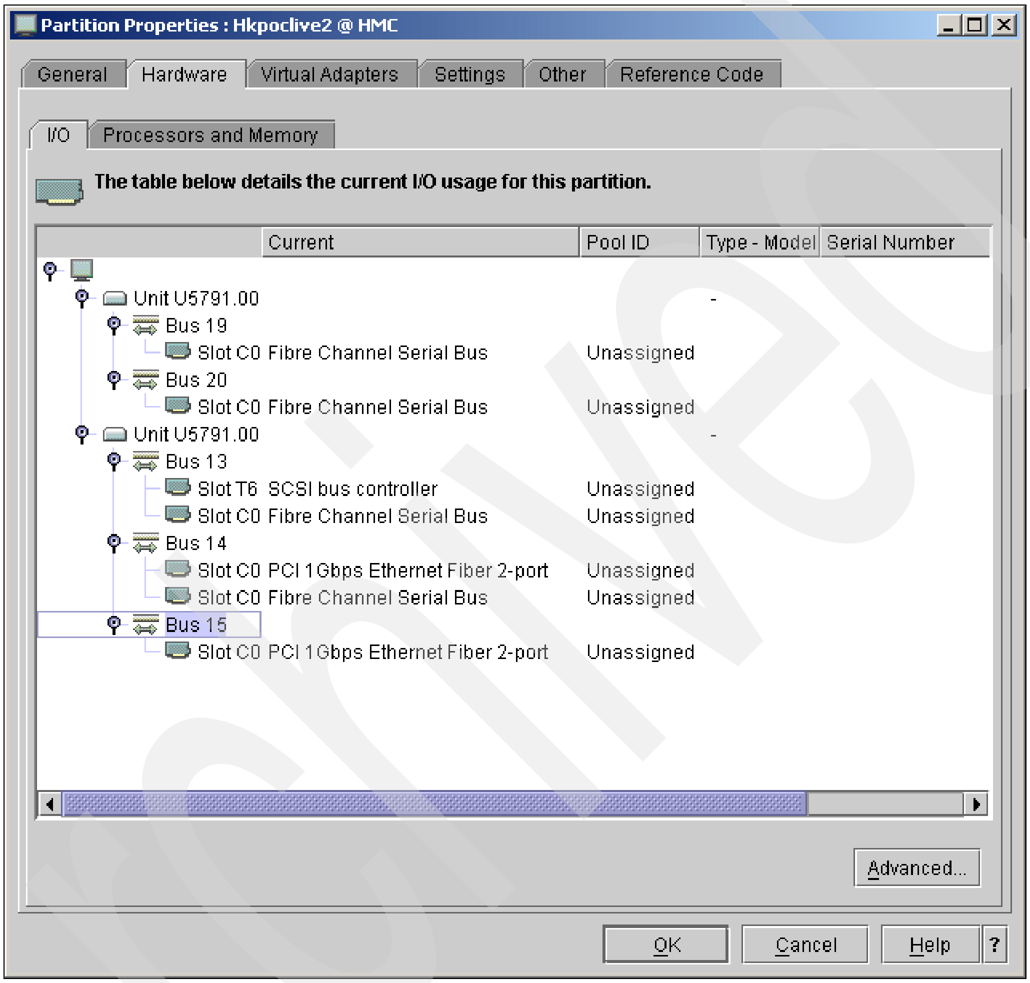

Figure 2-67 hkpoclive2 I/O configuration

At AIX configuration level, four internal SCSI 143 GB disks are available. Each disk has 143 GB capacity. Although it has FC adapters, this LPAR does not use disks from the storage subsystem.

```
{hkpoclive2:root}/ -> lsdev -Cc disk
hdisk0 Available 0M-08-00-8,0  16 Bit LVD SCSI Disk Drive
hdisk1 Available 0M-08-00-9,0  16 Bit LVD SCSI Disk Drive
hdisk2 Available 0M-08-00-10,0 16 Bit LVD SCSI Disk Drive
hdisk3 Available 0M-08-00-11,0 16 Bit LVD SCSI Disk Drive
{hkpoclive2:root}/ ->
```

Figure 2-68 hkpoclive2 disk configuration

It is also interesting to see which disk are assigned to which volume groups, as shown in Figure 2-69.

```
{hkpoclive2:root}/ -> lspv
hdisk1          00c263cfbadb6d0a                    rootvg          active
hdisk0          00c263cfb293fe1a                    rootvg          active
hdisk2          00c263cfbadb6f7e                    d3bvg           active
hdisk3          00c263cff73a0feb                    d3bvg           active
{hkpoclive2:root}/ ->
```

Figure 2-69 disk to volume group mapping

Figure 2-70 displays the characteristics of one of the disks. We specified hdisk0 for this example.

```
{hkpoclive2:root}/ -> lspv hdisk0
PHYSICAL VOLUME:    hdisk0                   VOLUME GROUP:     rootvg
PV IDENTIFIER:      00c263cfb293fe1a VG IDENTIFIER     00c263cf00004c0000000112fafd6a16
PV STATE:           active
STALE PARTITIONS:   0                        ALLOCATABLE:      yes
PP SIZE:            256 megabyte(s)          LOGICAL VOLUMES:  13
TOTAL PPs:          546 (139776 megabytes)   VG DESCRIPTORS:   1
FREE PPs:           247 (63232 megabytes)    HOT SPARE:        no
USED PPs:           299 (76544 megabytes)    MAX REQUEST:      256 kilobytes
FREE DISTRIBUTION:  109..91..00..00..47
USED DISTRIBUTION:  01..18..109..109..62
{hkpoclive2:root}/ ->
```

Figure 2-70 fdisk0 characteristics

Other physical devices available include Ethernet cards and HBAs. As shown in Figure 2-71, this LPAR has four FC adapters available.

```
{hkpoclive2:root}/ -> lsdev -C |grep fcs
fcs0       Available 00-08         FC Adapter
fcs1       Available 0P-08         FC Adapter
fcs2       Available 0R-08         FC Adapter
fcs3       Available 0S-08         FC Adapter
{hkpoclive2:root}/ ->
```

Figure 2-71 hkpoclive2 FC adapters

Regarding Ethernet adapters, hkpoclive2 has four 2-port GB adapters, as shown on Figure 2-72.

```
hkpoclive2:root}/ -> lsdev -C |grep ent
nt0       Available 0H-08        2-Port Gigabit Ethernet-SX PCI-X Adapter (14108802)
nt1       Available 0H-09        2-Port Gigabit Ethernet-SX PCI-X Adapter (14108802)
nt2       Available 0Q-08        2-Port Gigabit Ethernet-SX PCI-X Adapter (14108802)
nt3       Available 0Q-09        2-Port Gigabit Ethernet-SX PCI-X Adapter (14108802)
hkpoclive2:root}/ ->
```

Figure 2-72 hkpoclive2 Ethernet adapters

But regarding the network configuration, only two adapters are defined (plus loopback) with 2 IP addresses corresponding to the VLANs we described previously, as shown in Figure 2-73

```
{hkpoclive2:root}/ -> ifconfig -a
en0: flags=5e080863,c0<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,CHECKSUM_OFFL
        inet 1.1.75.4 netmask 0xffffff00 broadcast 1.1.75.255
         tcp_sendspace 131072 tcp_recvspace 65536
en2: flags=5e080863,c0<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,CHECKSUM_OFFL
        inet 10.3.24.4 netmask 0xffffff00 broadcast 10.3.24.255
         tcp_sendspace 131072 tcp_recvspace 65536
lo0: flags=e08084b<UP,BROADCAST,LOOPBACK,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT>
        inet 127.0.0.1 netmask 0xff000000 broadcast 127.255.255.255
        inet6 ::1/0
         tcp_sendspace 131072 tcp_recvspace 131072 rfc1323 1
{hkpoclive2:root}/ ->
```

Figure 2-73 TCP/IP configuration

As shown in Figure 2-74, several volume groups exist on this LPAR:

- rootvg
- old_rootvg: Not used in this configuration, varied off
- db3vg: Volume group used for SAP instance

```
{hkpoclive2:root}/ -> lsvg
old_rootvg
rootvg
d3bvg
{hkpoclive2:root}/ ->
```

Figure 2-74 hkpoclive2 volume groups

The volume group rootvg has the characteristics shown in Figure 2-75.

```
{hkpoclive2:root}/ -> lsvg rootvg
VOLUME GROUP:       rootvg                   VG IDENTIFIER:  00c263cf00004c0000000112fafd6a16
VG STATE:           active                   PP SIZE:        256 megabyte(s)
VG PERMISSION:      read/write               TOTAL PPs:      1092 (279552 megabytes)
MAX LVs:            256                      FREE PPs:       494 (126464 megabytes)
LVs:                13                       USED PPs:       598 (153088 megabytes)
OPEN LVs:           12                       QUORUM:         1
TOTAL PVs:          2                        VG DESCRIPTORS: 3
STALE PVs:          0                        STALE PPs:      0
ACTIVE PVs:         2                        AUTO ON:        no
MAX PPs per VG:     32512
MAX PPs per PV:     1016                     MAX PVs:        32
LTG size (Dynamic): 256 kilobyte(s)          AUTO SYNC:      no
HOT SPARE:          no                       BB POLICY:      relocatable
{hkpoclive2:root}/ ->
```

Figure 2-75 trootvg characteristics

Several logical volumes have been defined on this rootvg volume group, as shown in Figure 2-76.

```
{hkpoclive2:root}/ -> lsvg -l rootvg
rootvg:
LV NAME             TYPE       LPs   PPs   PVs  LV STATE      MOUNT POINT
hd5                 boot       1     2     2    closed/syncd  N/A
hd6                 paging     8     16    2    open/syncd    N/A
hd8                 jfs2log    1     2     2    open/syncd    N/A
hd4                 jfs2       1     2     2    open/syncd    /
hd2                 jfs2       13    26    2    open/syncd    /usr
hd9var              jfs2       1     2     2    open/syncd    /var
hd3                 jfs2       5     10    2    open/syncd    /tmp
hd1                 jfs2       1     2     2    open/syncd    /home
hd10opt             jfs2       2     4     2    open/syncd    /opt
lg_dumplv           sysdump    8     16    2    open/syncd    N/A
fslv00              jfs2       1     2     2    open/syncd    /XmRec
fslv01              jfs2       1     2     2    open/syncd    /tmp/m2
lvsapd3b            jfs2       256   512   2    open/syncd    /sap/D3B
{hkpoclive2:root}/ ->
```

Figure 2-76 rootvg logical volumes

The volume group d3bvg has the characteristics shown in Figure 2-77.

```
{hkpoclive2:root}/ -> lsvg d3bvg
VOLUME GROUP:       d3bvg                    VG IDENTIFIER:   00c263cf00004c00000001158e24
VG STATE:           active                   PP SIZE:         128 megabyte(s)
VG PERMISSION:      read/write               TOTAL PPs:       2186 (279808 megabytes)
MAX LVs:            256                      FREE PPs:        1786 (228608 megabytes)
LVs:                1                        USED PPs:        400 (51200 megabytes)
OPEN LVs:           1                        QUORUM:          2
TOTAL PVs:          2                        VG DESCRIPTORS:  3
STALE PVs:          0                        STALE PPs:       0
ACTIVE PVs:         2                        AUTO ON:         yes
MAX PPs per VG:     32768                    MAX PVs:         1024
LTG size (Dynamic): 256 kilobyte(s)          AUTO SYNC:       no
HOT SPARE:          no                       BB POLICY:       relocatable
{hkpoclive2:root}/ ->
```

Figure 2-77 D3bvg characteristics

Several logical volumes have been defined on this d3bvg volume group, as shown in Figure 2-78.

```
{hkpoclive2:root}/ -> lsvg -l d3bvg
d3bvg:
LV NAME             TYPE       LPs   PPs   PVs  LV STATE      MOUNT POINT
d3blv               jfs2       400   400   2    open/syncd    /usr/sap/D3B/DVEBMGS00
{hkpoclive2:root}/ ->
```

Figure 2-78 d3bvg logical volumes

Mounted file systems are the necessary ones for the operating system and for the SAP SCM central instance, as shown in Figure 2-79 on page 60. Further installation information about the SAP SCM central instance is on the Web at `http://www.service.sap.com/instguides` and `http://www.service.sap.com/notes`. Note that the picasso and hkpoclive1 NFS mounted file systems are used for performance statistics collection and are not directly related to this project.

```
(hkpoclive2:root)/ -> mount
  node       mounted        mounted over    vfs       date        options
-------- --------------- --------------- ------ ------------ ---------------
         /dev/hd4         /               jfs2   Oct 08 16:08 rw,log=/dev/hd8
         /dev/hd2         /usr            jfs2   Oct 08 16:08 rw,log=/dev/hd8
         /dev/hd9var      /var            jfs2   Oct 08 16:08 rw,log=/dev/hd8
         /dev/hd3         /tmp            jfs2   Oct 08 16:08 rw,log=/dev/hd8
         /dev/hd1         /home           jfs2   Oct 08 16:09 rw,log=/dev/hd8
         /proc            /proc           procfs Oct 08 16:09 rw
         /dev/hd10opt     /opt            jfs2   Oct 08 16:09 rw,log=/dev/hd8
         /dev/fslv00      /XmRec          jfs2   Oct 08 16:09 rw,log=/dev/hd8
         /dev/fslv01      /tmp/m2         jfs2   Oct 08 16:09 rw,log=/dev/hd8
         /dev/lvsapd3b    /sap/D3B        jfs2   Oct 08 16:09 rw,log=/dev/hd8
hkpoceccapp1 /sap/D2L/trans   /usr/sap/trans   nfs3   Oct 08 16:09 rw,bg,hard,intr,sec=sys:dh:krb5:kr
hkpoclive1 /tmp/sapact      /tmp/sapact      nfs3   Oct 08 16:09 rw,bg,soft,intr,sec=sys
picasso-onn /Results         /Results         nfs3   Oct 08 16:15 bg,hard,intr,sec=sys,rw
(hkpoclive2:root)/ ->
```

Figure 2-79   Mounted file systems

Disk space used by those file systems is shown in Figure 2-80.

```
(hkpoclive2:root)/ -> df -g
Filesystem    GB blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4           0.25      0.16   36%    12907    26% /
/dev/hd2           3.25      0.13   96%    53920    59% /usr
/dev/hd9var        0.25      0.22   12%      566     2% /var
/dev/hd3           1.25      0.13   90%      793     3% /tmp
/dev/hd1           0.25      0.20   20%      298     1% /home
/proc                 -         -    -         -     -  /proc
/dev/hd10opt       0.50      0.26   49%     6269    10% /opt
/dev/fslv00        0.25      0.25    1%        4     1% /XmRec
/dev/fslv01        0.25      0.24    3%       30     1% /tmp/m2
/dev/lvsapd3b     64.00     31.04   52%    83526     2% /sap/D3B
hkpoceccapp1:/sap/D2L/trans      64.00     31.78   51%    54770     1% /usr/sap/trans
hkpoclive1:/tmp/sapact       5.50      0.66   88%     3347     3% /tmp/sapact
picasso-onn:/Results    136.25    126.10    8%    11409     1% /Results
(hkpoclive2:root)/ ->
```

Figure 2-80   File systems disk usage

The paging area is defined as shown in Figure 2-81.

```
(hkpoclive2:root)/ -> lsps -a
Page Space      Physical Volume   Volume Group    Size %Used Active  Auto  Type
paging00        hdisk2            psvg         32000MB     1   yes   yes    lv
hd6             hdisk1            rootvg        2048MB     2   yes   yes    lv
(hkpoclive2:root)/ ->
```

Figure 2-81   Paging area definition

We have now completely described the configuration of the LPAR hkpoclive2. This configuration is standard for running an SAP SCM central instance.

Summary

In the previous sections we have described the System p configuration that was used for MRP, RPM, and Backflush KPIs. The configuration was changed by the Montpellier team before we performed the Material Ledger KPI. These configurations are summarized in the following two tables.

Table 2-5 Initial System p configuration (used for MRP, RPM, and Backflush KPIs)

LPAR	CPU	RAM (GB)	FC adapters	Ethernet adapters	Purpose
hkpoceccapp1	10	40	0	2	ECC central instance
hkpoceccapp2	10	40	0	2	ECC dialog instance
hkpoceccapp3	10	40	0	2	ECC dialog instance
hkpoclive1	32	196	4	2	liveCache
hkpoclive2	32	56	0	2	SCM central instance

Table 2-6 Modified configuration (used for ML KPI)

LPAR	CPU	RAM (GB)	FC adapters	Ethernet adapters	Purpose
hkpoceccapp1	10	60	0	2	ECC central instance
hkpoceccapp2	10	60	0	2	ECC dialog instance
hkpoceccapp3	10	60	0	2	ECC dialog instance
hkpoclive1	32	128	4	2	liveCache
hkpoclive2	32	128	0	2	SCM central instance

2.4 Storage configuration

In this project we used one storage subsystem DS8300 Turbo. This section describes the configuration of the storage subsystem for this benchmark. For further details about DS8000, refer to `http://www-03.ibm.com/systems/storage/disk/ds8000/index.html` or to *IBM System Storage DS8000 Series: Architecture and Implementation*, SG24-6786.

The DS8300 has four Power5 cores and 128 GB memory. The box is not running in LPAR mode. This means that the same resources (cache, processors, and disks) are used simultaneous by the z and P environments. Figure 2-82 displays details about the storage subsystem. Two CPUs have two cores each, and each CPU has 64 GB RAM.

```
dscli showsu -fullid   IBM.2107-75M1040 -cfg
C:\Progra~1\IBM\dscli\profile\W10.profile
Date/Time: October 11, 2007 10:44:56 AM CEST IBM DSCLI Version: 5.2.410.333 DS:
IBM.2107-75M1040
Name             -
desc             -
ID               IBM.2107-75M1040
Model            932
WWNN             5005076306FFF8C3
config           Undefined
pw state         On
pw mode          Remote Manual
reqpm            Remote Manual
Processor Memory 128.0 GB
MTS              IBM.2421-75M1040
```

Figure 2-82   DS8300 characteristics

The DS8300 has eight HBAs at 4 Gbps; each HBA has four ports. The details are shown in Figure 2-83. This figure displays the ID of each HBA, and also the ID of each port.

```
dscli lshba  -l IBM.2107-75M1041  -cfg C:\Progra~1\IBM\dscli\profile\W10.profile
Date/Time: October 11, 2007 10:45:26 AM CEST IBM DSCLI Version: 5.2.410.333 DS:
IBM.2107-75M1041
ID                           State  loc                      FC  interfID
==========================================================================================
IBM.1300-001-34150/R-1-P1-C1 Online U1300.001.RJ34150-P1-C1 N/A 0x0200,0x0201,0x0202,0x0203
IBM.1300-001-34150/R-1-P1-C4 Online U1300.001.RJ34150-P1-C4 N/A 0x0230,0x0231,0x0232,0x0233
IBM.1300-001-34151/R-1-P1-C1 Online U1300.001.RJ34151-P1-C1 N/A 0x0000,0x0001,0x0002,0x0003
IBM.1300-001-34151/R-1-P1-C4 Online U1300.001.RJ34151-P1-C4 N/A 0x0030,0x0031,0x0032,0x0033
IBM.1300-001-34258/R-1-P1-C1 Online U1300.001.RJ34258-P1-C1 N/A 0x0300,0x0301,0x0302,0x0303
IBM.1300-001-34258/R-1-P1-C4 Online U1300.001.RJ34258-P1-C4 N/A 0x0330,0x0331,0x0332,0x0333
IBM.1300-001-34259/R-1-P1-C1 Online U1300.001.RJ34259-P1-C1 N/A 0x0100,0x0101,0x0102,0x0103
IBM.1300-001-34259/R-1-P1-C4 Online U1300.001.RJ34259-P1-C4 N/A 0x0130,0x0131,0x0132,0x0133
```

Figure 2-83   HBA configuration

Figure 2-84 displays the details for each port. We actually have to support two protocols: FICON for System z and FCP for System p. In the figure you can see the mapping between HBA ports and the protocols they support.

```
dscli lsioport  -l  -cfg C:\Progra~1\IBM\dscli\profile\W10.profile
Date/Time: October 11, 2007 10:45:09 AM CEST IBM DSCLI Version: 5.2.410.333 DS:
IBM.2107-75M1041
ID    WWPN             State  Type             topo     portgrp Speed
=======================================================================
I0000 50050763060000C3 Online Fibre Channel-SW FICON    0       4 Gb/s
I0001 50050763060040C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0002 50050763060080C3 Online Fibre Channel-SW SCSI-FCP 0       4 Gb/s
I0003 500507630600C0C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0030 50050763060300C3 Online Fibre Channel-SW FICON    0       4 Gb/s
I0031 50050763060340C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0032 50050763060380C3 Online Fibre Channel-SW SCSI-FCP 0       4 Gb/s
I0033 500507630603C0C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0100 50050763060800C3 Online Fibre Channel-SW FICON    0       4 Gb/s
I0101 50050763060840C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0102 50050763060880C3 Online Fibre Channel-SW SCSI-FCP 0       4 Gb/s
I0103 500507630608C0C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0130 50050763060B00C3 Online Fibre Channel-SW FICON    0       4 Gb/s
I0131 50050763060B40C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0132 50050763060B80C3 Online Fibre Channel-SW SCSI-FCP 0       4 Gb/s
I0133 50050763060BC0C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0200 50050763061000C3 Online Fibre Channel-SW FICON    0       4 Gb/s
I0201 50050763061040C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0202 50050763061080C3 Online Fibre Channel-SW SCSI-FCP 0       4 Gb/s
I0203 500507630610C0C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0230 50050763061300C3 Online Fibre Channel-SW FICON    0       4 Gb/s
I0231 50050763061340C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0232 50050763061380C3 Online Fibre Channel-SW SCSI-FCP 0       4 Gb/s
I0233 500507630613C0C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0300 50050763061800C3 Online Fibre Channel-SW FICON    0       4 Gb/s
I0301 50050763061840C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0302 50050763061880C3 Online Fibre Channel-SW SCSI-FCP 0       4 Gb/s
I0303 500507630618C0C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0330 50050763061B00C3 Online Fibre Channel-SW FICON    0       4 Gb/s
I0331 50050763061B40C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
I0332 50050763061B80C3 Online Fibre Channel-SW SCSI-FCP 0       4 Gb/s
I0333 50050763061BC0C3 Online Fibre Channel-SW FC-AL    0       4 Gb/s
```

Figure 2-84 HBA ports details

Regarding internal access to disk connectivity, this configuration uses four HBAs with four ports at 4 Gb/s, as shown in Figure 2-85.

```
scli lsda  -l IBM.2107-75M1041  -cfg C:\Progra~1\IBM\dscli\profile\W10.profile
Date/Time: October 11, 2007 10:45:32 AM CEST IBM DSCLI Version: 5.2.410.333 DS: IBM.2107-
75M1041
ID                            State  loc                    FC Server DA pair interfs
==============================================================================IBM.1300-001-
34150/R-1-P1-C3 Online U1300.001.RJ34150-P1-C3 -  0      2       0x0230,0x0231,0x0232,0x0233
IBM.1300-001-34151/R-1-P1-C3 Online U1300.001.RJ34151-P1-C3 -  0      0
0x0030,0x0031,0x0032,0x0033
IBM.1300-001-34258/R-1-P1-C6 Online U1300.001.RJ34258-P1-C6 -  1      2
0x0360,0x0361,0x0362,0x0363
IBM.1300-001-34259/R-1-P1-C6 Online U1300.001.RJ34259-P1-C6 -  1      0
0x0160,0x0161,0x0162,0x0163
```

Figure 2-85 Internal connectivity details

The internal protocol used is FCP.

DS8300 is made of 128 disks, grouped in 16 blocks of 8 disks each. Each disk has 146 GB capacity and runs at 15 Krpm. All disks are the same, so we show the characteristics of only one disk in Figure 2-86.

```
dscli lsddm  -l IBM.2107-75M1041  -cfg C:\Progra~1\IBM\dscli\profile\W10.profile
Date/Time: October 11, 2007 10:45:16 AM CEST IBM DSCLI Version: 5.2.410.333 DS: IBM.2107-
75M1041
ID                       Model      loc                     firmwarelevel DA Pair
DualLoop dkcap (10^9B) diskrpm dkinf dkrate


(Gb/sec) dkuse          arsite Position State  diskclass
================================================================================



==================================================
IBM.2107-D01-47351/R1-P1-D1  S5BE146146 U2107.D01.RJ47351-P1-D1  940D         0
2         146.0   15000 FC-AL 2.2


  array member   S3           6 Normal ENT
```

Figure 2-86 Disk example

In this figure we can very easily see that we have one 146 GB disk at 15 Krpm connected through FCP.

Array sites are groups of eight physical disks. As we explained before, we have 16 array sites, as shown in Figure 2-87.

```
dscli lsarraysite  -cfg C:\Progra~1\IBM\dscli\profile\W10.profile
Date/Time: October 11, 2007 10:46:48 AM CEST IBM DSCLI Version: 5.2.410.333 DS:
IBM.2107-75M1041
arsite DA Pair dkcap (10^9B) State    Array
===========================================
S1     0               146.0 Assigned A0
S2     0               146.0 Assigned A1
S3     0               146.0 Assigned A2
S4     0               146.0 Assigned A3
S5     0               146.0 Assigned A4
S6     0               146.0 Assigned A5
S7     0               146.0 Assigned A6
S8     0               146.0 Assigned A7
S9     2               146.0 Assigned A8
S10    2               146.0 Assigned A9
S11    2               146.0 Assigned A10
S12    2               146.0 Assigned A11
S13    2               146.0 Assigned A12
S14    2               146.0 Assigned A13
S15    2               146.0 Assigned A14
S16    2               146.0 Assigned A15
```

Figure 2-87 Array sites list

These array sites are formatted into arrays. In our case all arrays are RAID5, as shown in Figure 2-88.

```
dscli lsarray -l  -cfg C:\Progra~1\IBM\dscli\profile\W10.profile
Date/Time: October 11, 2007 10:46:55 AM CEST IBM DSCLI Version: 5.2.410.333 DS:
IBM.2107-75M1041
Array State    Data   RAIDtype  arsite Rank DA Pair DDMcap (10^9B) diskclass
============================================================================
A0    Assigned Normal 5 (6+P+S) S1     R0   0                146.0 ENT
A1    Assigned Normal 5 (6+P+S) S2     R1   0                146.0 ENT
A2    Assigned Normal 5 (6+P+S) S3     R2   0                146.0 ENT
A3    Assigned Normal 5 (6+P+S) S4     R3   0                146.0 ENT
A4    Assigned Normal 5 (7+P)   S5     R4   0                146.0 ENT
A5    Assigned Normal 5 (7+P)   S6     R5   0                146.0 ENT
A6    Assigned Normal 5 (7+P)   S7     R6   0                146.0 ENT
A7    Assigned Normal 5 (7+P)   S8     R7   0                146.0 ENT
A8    Assigned Normal 5 (6+P+S) S9     R8   2                146.0 ENT
A9    Assigned Normal 5 (6+P+S) S10    R9   2                146.0 ENT
A10   Assigned Normal 5 (6+P+S) S11    R10  2                146.0 ENT
A11   Assigned Normal 5 (6+P+S) S12    R11  2                146.0 ENT
A12   Assigned Normal 5 (7+P)   S13    R12  2                146.0 ENT
A13   Assigned Normal 5 (7+P)   S14    R13  2                146.0 ENT
A14   Assigned Normal 5 (7+P)   S15    R14  2                146.0 ENT
A15   Assigned Normal 5 (7+P)   S16    R15  2                146.0 ENT
```

Figure 2-88 Arrays

We can easily see parity and spare disks, and also the correspondence between array sites and arrays.

Next we can view the ranks, which are arrays formatted to one system or the other as shown in Figure 2-89.

```
dscli lsrank -l  -cfg C:\Progra~1\IBM\dscli\profile\W10.profile
Date/Time: October 11, 2007 10:47:02 AM CEST IBM DSCLI Version: 5.2.410.333 DS:
IBM.2107-75M1041
ID Group State datastate Array RAIDtype extpoolID extpoolnam stgtype exts   usedexts
===================================================================================
R0   0   Normal  Normal   A0      5    P0     extp_P0       ckd      873     873
R1   1   Normal  Normal   A1      5    P1     extp_P1       ckd      873     873
R2   0   Normal  Normal   A2      5    P0     extp_P0       ckd      873     873
R3   1   Normal  Normal   A3      5    P1     extp_P1       ckd      873     873
R4   0   Normal  Normal   A4      5    P0     extp_P0       ckd     1018     999
R5   1   Normal  Normal   A5      5    P1     extp_P1       ckd     1018     990
R6   0   Normal  Normal   A6      5    P0     extp_P0       ckd     1018     990
R7   1   Normal  Normal   A7      5    P1     extp_P1       ckd     1018     999
R8   0   Normal  Normal   A8      5    P0     extp_P0       ckd      873     873
R9   1   Normal  Normal   A9      5    P1     extp_P1       ckd      873     873
R10  0   Normal  Normal   A10     5    P2     ForPserie     fb       779     750
R11  1   Normal  Normal   A11     5    P3     ForPserie     fb       779     750
R12  0   Normal  Normal   A12     5    P4     ForPserie     fb       909     809
R13  1   Normal  Normal   A13     5    P5     ForPserie     fb       909     809
R14  0   Normal  Normal   A14     5    P6     ForPserie     fb       909     809
R15  1   Normal  Normal   A15     5    P7     ForPserie     fb       909     809
```

Figure 2-89 Ranks

Column `stgtype` indicates that ranks R0 to R9 are formatted for System z (ckd) and R10 to R15 are formatted for System p (fb).

Column `ID` shows that, for System p, each rank is assigned to an extent pool. In our project, extent pools are defined as shown in Figure 2-90.

```
dscli lsextpool -l   -cfg C:\Progra~1\IBM\dscli\profile\W10.profile
Date/Time: October 11, 2007 10:47:10 AM CEST IBM DSCLI Version: 5.2.410.333 DS:
IBM.2107-75M1041
Name ID stgtype rankgrp status availstor (2^30B) %allocated avail reserved numvols numranks

extp_P0   P0 ckd      0  below      41       98      47     0     512      5
extp_P1   P1 ckd      1  below      41       98      47     0     512      5
ForPserie P2 fb       0  below      29       96      29     0       5      1
ForPserie P3 fb       1  below      29       96      29     0       5      1
ForPserie P4 fb       0  below     100       88     100     0       6      1
ForPserie P5 fb       1  below     100       88     100     0       6      1
ForPserie P6 fb       0  below     100       88     100     0       6      1
ForPserie P7 fb       1  below     100       88     100     0       6      1
```

Figure 2-90 Extent pools

As shown, we defined 2 extent pools for System z (P0 and P1) and 6 for System p (P2, P3, P4, P5, P6, and P7).

2.4.1 System z extent pools

This section applies to storage configuration for System z only.

On the System z extend pools (P0 and P1) we defined 16 control units/ssids, 8 on extend pool P0 and 8 on extend pool P1 (see Figure 2-91).

```
dscli lslcu -l   -cfg C:\Progra~1\IBM\dscli\profile\HKMC_MOP.profile
Date/Time: October 16, 2007 4:14:53 PM CEST IBM DSCLI Version: 5.2.420.682 DS: IBM.2107-75M1041
ID Group addrgrp confgvols subsys conbasetype
=============================================
00      0 0            256 0xC000 3990-6
01      1 0            256 0xC100 3990-6
02      0 0            256 0xC200 3990-6
03      1 0            256 0xC300 3990-6
04      0 0            256 0xC400 3990-6
05      1 0            256 0xC500 3990-6
06      0 0            256 0xC600 3990-6
07      1 0            256 0xC700 3990-6
08      0 0            256 0xC800 3990-6
09      1 0            256 0xC900 3990-6
0A      0 0            256 0xCA00 3990-6
0B      1 0            256 0xCB00 3990-6
0C      0 0            256 0xCC00 3990-6
0D      1 0            256 0xCD00 3990-6
0E      0 0            256 0xCE00 3990-6
0F      1 0            256 0xCF00 3990-6
```

Figure 2-91 CKD CU/SSIDs

Each control unit has a maximum of 256 addresses for defining volumes. Of these, 64 are defined as base volumes 3390 model-9. The remaining addresses are defined as Parallel Access Volumes (PAV) aliases. Every base volume has 3 aliases assigned. This distribution of PAVs is not relevant anymore because HyperPAV is used. This means that all defined aliases are gathered into one pool of aliases per CU/SSID. Those aliases can be used as additional PAVs on any base volume when needed.

Figure 2-92 shows the volumes defined for CU/SSID 00.

```
dscli lsckdvol -l   -cfg C:\Progra~1\IBM\dscli\profile\HKMC_MOP.profile
Date/Time: October 16, 2007 4:15:08 PM CEST IBM DSCLI Version: 5.2.420.682 DS: IBM.2107-75M1041
Name ID   accstate datastate configstate deviceMTM volser datatype voltype   orgbvols extpool sam      cap (cyl) cap (10^9B) cap (2^30B)
======================================================================================================================================
-    0000 Online   Normal    Normal      3390-9    H1SYP1 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    0001 Online   Normal    Normal      3390-9    H1SYS1 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    0002 Online   Normal    Normal      3390-9    H1PAG1 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    0003 Online   Normal    Normal      3390-9    H1SMF1 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    0004 Online   Normal    Normal      3390-9    H1LOC1 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    0005 Online   Normal    Normal      3390-9    H1JES2 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    0006 Online   Normal    Normal      3390-9    H1W001 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    0007 Online   Normal    Normal      3390-9    H1OS01 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    0008 Online   Normal    Normal      3390-9    H1OS02 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    0009 Online   Normal    Normal      3390-9    H1OU01 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    000A Online   Normal    Normal      3390-9    H1TSO1 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    000B Online   Normal    Normal      3390-9    H1TSO2 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    000C Online   Normal    Normal      3390-9    H1W002 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    000D Online   Normal    Normal      3390-9    H1CICS 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    000E Online   Normal    Normal      3390-9    H1DB81 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    000F Online   Normal    Normal      3390-9    XXC00F 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    ....
-    ....
-    003C Online   Normal    Normal      3390-9    F1C03C 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    003D Online   Normal    Normal      3390-9    F1C03D 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    003E Online   Normal    Normal      3390-9    F1C03E 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    003F Online   Normal    Normal      3390-9    F1C03F 3390     CKD Base  -        P0      Standard     10017         8.5         7.9
-    0040 -        -         -           -         F1C03F -        CKD Alias 003F     -       Standard         0         0.0         0.0
-    0041 -        -         -           -         F1C03F -        CKD Alias 003F     -       Standard         0         0.0         0.0
-    0042 -        -         -           -         F1C03F -        CKD Alias 003F     -       Standard         0         0.0         0.0
-    ....
-    ....
-    00FA -        -         -           -         H1SYS1 -        CKD Alias 0001     -       Standard         0         0.0         0.0
-    00FB -        -         -           -         H1SYS1 -        CKD Alias 0001     -       Standard         0         0.0         0.0
-    00FC -        -         -           -         H1SYS1 -        CKD Alias 0001     -       Standard         0         0.0         0.0
-    00FD -        -         -           -         H1SYP1 -        CKD Alias 0000     -       Standard         0         0.0         0.0
-    00FE -        -         -           -         H1SYP1 -        CKD Alias 0000     -       Standard         0         0.0         0.0
-    00FF -        -         -           -         H1SYP1 -        CKD Alias 0000     -       Standard         0         0.0         0.0
```

Figure 2-92 CKD Volumes on CU/SSID

2.4.2 System p extent pools

This section applies to storage configuration for System p only.

On our six extent pools we created LUNs. There are several possibilities for creating LUNs depending on various criteria; in our case the LUNs were defined as shown in Figure 2-93 on page 68.

```
dscli lsfbvol -l  -cfg C:\Progra~1\IBM\dscli\profile\W10.profile
Date/Time: October 11, 2007 10:49:46 AM CEST IBM DSCLI Version: 5.2.410.333 DS: IBM.2107-75M1041
Name           ID   accstate datastate configstate deviceMTM datatype extpool sam      captype cap
(2^30B) cap (10^9B) cap (blocks) volgrp


hk1_P2_01_1000 1000 Online   Normal    Normal      2107-900  FB 512   P2      Standard      DS
150.0          -    314572800 V0
hk2_P2_01_1001 1001 Online   Normal    Normal      2107-900  FB 512   P2      Standard      DS
150.0          -    314572800 V1
hk1_P3_01_1001 1002 Online   Normal    Normal      2107-900  FB 512   P4      Standard      DS
150.0          -    314572800 V0
hk2_P3_01_1004 1003 Online   Normal    Normal      2107-900  FB 512   P4      Standard      DS
150.0          -    314572800 V1
hk1_P4_01_1001 1004 Online   Normal    Normal      2107-900  FB 512   P6      Standard      DS
150.0          -    314572800 V0
hk2_P4_01_1005 1005 Online   Normal    Normal      2107-900  FB 512   P6      Standard      DS
150.0          -    314572800 V1
hk1_P5_01_1002 1100 Online   Normal    Normal      2107-900  FB 512   P3      Standard      DS
150.0          -    314572800 V0
hk2_P5_01_1005 1101 Online   Normal    Normal      2107-900  FB 512   P3      Standard      DS
150.0          -    314572800 V1
hk1_P6_01_1002 1102 Online   Normal    Normal      2107-900  FB 512   P5      Standard      DS
150.0          -    314572800 V0
hk2_P6_01_1006 1103 Online   Normal    Normal      2107-900  FB 512   P5      Standard      DS
150.0          -    314572800 V1
hk1_P7_01_1003 1104 Online   Normal    Normal      2107-900  FB 512   P7      Standard      DS
150.0          -    314572800 V0
hk2_P7_01_1006 1105 Online   Normal    Normal      2107-900  FB 512   P7      Standard      DS
150.0          -    314572800 V1
hk1_SAPLOG_P4  1200 Online   Normal    Normal      2107-900  FB 512   P4      Standard      DS
59.0           -    123731968 V0
hk1_SAPLOG_P6  1201 Online   Normal    Normal      2107-900  FB 512   P6      Standard      DS
59.0           -    123731968 V0
hk1_SAPLOG_P5  1300 Online   Normal    Normal      2107-900  FB 512   P5      Standard      DS
59.0           -    123731968 V0
hk1_SAPLOG_P7  1301 Online   Normal    Normal      2107-900  FB 512   P7      Standard      DS
59.0           -    123731968 V0
hk1_P2_50_5000 5000 Online   Normal    Normal      2107-900  FB 512   P2      Standard      DS
150.0          -    314572800 V2
hk1_P4_50_5001 5001 Online   Normal    Normal      2107-900  FB 512   P4     Standard      DS
150.0          -    314572800 V2
hk1_P6_50_5002 5002 Online   Normal    Normal      2107-900  FB 512   P6      Standard      DS
150.0          -    314572800 V2
```

Figure 2-93 LUNs definition

As shown in the figure, all LUNs are defined with 150 GB, except LUNs that will be used for MaxDB logging purposes, which are defined with 59 Gb.

The LUNs are grouped into volume groups, and the volumes groups are presented to System p LPARs according to a path. Figure 2-94 on page 69 shows the volume groups.

> **Note:** We are not dealing with UNIX volume groups but with DS8300; the concept has nothing to do with UNIX although the name is the same.

```
dscli lsvolgrp -l   -cfg C:\Progra~1\IBM\dscli\profile\W10.profile
Date/Time: October 11, 2007 10:49:59 AM CEST IBM DSCLI Version: 5.2.410.333 DS:
IBM.2107-75M1041
Name                ID  Type
=======================================
hkpoclive1          V0  SCSI Mask
hkpoclive1_FC4      V1  SCSI Mask
hkpoclive1_FC       V2  SCSI Mask
hkpoclive1_FC2      V3  SCSI Mask
hkpoclive1_FC3      V4  SCSI Mask
All CKD             V10 FICON/ESCON All
All Fixed Block-512 V20 SCSI All
All Fixed Block-520 V30 OS400 All
```

Figure 2-94   DS8300 Volume groups

Volume group form V0 is presented to the System p LPARs. In Figure 2-95 on page 69 we can see how volume groups are presented to LPARs. Volume groups V1 to V4 are used for FlashCopy®. Other volume groups do not apply to System p.

```
dscli lshostconnect -l   -cfg C:\Progra~1\IBM\dscli\profile\W10.profile
Date/Time: October 11, 2007 10:50:06 AM CEST IBM DSCLI Version: 5.2.410.333 DS: IBM.2107-75M1041
Name         ID   WWPN             HostType LBS addrDiscovery Profile          portgrp volgrpID
atchtopo ESSIOport                 speed   desc
==========================================================================================================
hkpoclive1   0000 10000000C95FAA28 pSeries  512 reportLUN     IBM pSeries - AIX     0 V0      SCSI-FCP
I0332,I0232,I0032,I0132,I0202,I0002,I0102,I0302 Unknown -

hkpoclive1   0001 10000000C95FB7A9 pSeries  512 reportLUN     IBM pSeries - AIX     1 V0      SCSI-FCP
I0332,I0232,I0032,I0132,I0202,I0002,I0102,I0302 Unknown -

hkpoclive1   0002 10000000C95FB4C7 pSeries  512 reportLUN     IBM pSeries - AIX     2 V0      SCSI-FCP
I0332,I0232,I0032,I0132,I0202,I0002,I0102,I0302 Unknown -

hkpoclive1   0003 10000000C95FB726 pSeries  512 reportLUN     IBM pSeries - AIX     3 V0      SCSI-FCP
I0332,I0232,I0032,I0132,I0202,I0002,I0102,I0302 Unknown -

hkpoclive2   0004 10000000C95FB839 pSeries  512 reportLUN     IBM pSeries - AIX     0 -       SCSI-FCP
all                Unknown -

hkpoclive2   0005 10000000C95FB7F9 pSeries  512 reportLUN     IBM pSeries - AIX     1 -       SCSI-FCP
all                Unknown -

hkpoclive2   0006 10000000C95FB6C8 pSeries  512 reportLUN     IBM pSeries - AIX     2 -       SCSI-FCP
all                Unknown -

hkpoclive2   0007 10000000C95FB5D9 pSeries  512 reportLUN     IBM pSeries - AIX     3 -       SCSI-FCP
all                Unknown -
```

Figure 2-95   LUN masking

As you can see, volume group V0 is presented to LPAR `hkpoclive1`, so only `hkpoclive1` is able to see DS8300 FCP ports and therefore make use of the volume groups.

3

Results overview

In this chapter we present an overview of the test cases performed during this test campaign and provide selected run logs as well. Part 2 of this book provides more details and analysis of some specific tests related to the four key performance indicators: RPM, MRP, Backflush and Materials Ledger.

3.1 RPM test

During the RPM test process we investigated two ways of running RPM:

- Using three separate jobs for three levels
- Using one job with one step for three levels

We also tested a running process with overlap timing between the different steps. Run 7 and Run 10 are detailed in Chapter 4, “Rapid Planning Matrix” on page 79.

Conduction					SAP values				
Run_ID	**Date**	**Job / Step**	**start**	**end**	**duration in min.**	**total duration in**	**total duration in min.**	Status	Comments
RUN0001	06.09.2007	RPM Trim	18:15	18:30	15	15	16		3 separate jobs for 3 levels. 2 jobs were canceled because the log file size was full.
		RPM Paint	-	-	-				
		RPM Body	-	-	-				
RUN0002	06.09.2007	RPM Trim	19:00	19:34	4	*16*	*16*		3 separate jobs for 3 levels. Log file size has been incresed. The result has **NOT** included the run time of time series creation.
		RPM Paint	19:19	19:23	4				
		RPM Body	19:30	19:38	8				
RUN0003	07.09.2007	RPM Trim	10:50	10:56	6	*13*	*13*		3 separate Jobs for 3 levels. Number of parallel processes have been chenged from 40 to 21. The result has **NOT** included the run time of time series creation.
		RPM Paint	11:03	11:07	4				
		RPM Body	11:20	11:23	3				
RUN0004	07.09.2007	RPM Trim	11:45	11:54	9	28	28		1 job with 4 steps. 3 steps for 3 levels and last step for timeseries creatation.
		RPM Paint	11:54	11:56	2				
		RPM Body	11:56	11:58	2				
		PM Time Seri	11:58	12:13	15				
RUN0005	07.09.2007	RPM All	14:06	14:22	16	**16**	16		1 job with 1 step for 3 levels. Number of parallel processes have been chenged from 21 to 66.
RUN0006	07.09.2007	RPM All	15:44	16:46	62	62	62		1 job with 1 step for 3 levels. The RPM was blocked for 40 min. because the REORG of RUN0005 was still running.
RUN0007	07.09.2007	RPM All	17:11	17:26	15	**15**	**15**		Same setting as RUN0005
RUN008	10.09.2007	RPM All	13:21	17:00	-				Job for investigating RPM fatal errors. The job was canceled because of SQL lock of DB2.
RUN009	10.09.2206	PPDS Trim	22:25	22:40	15	44	59		Overlap process with overwriting the exsiting plan. orders and separate creating of time series
	11.09.2207	RPM Trim	10:20	10:31	11				
		PPDS Paint	10:20	10:35	14				
		RPM Paint	11:00	11:03	3				
		PPDS Body	11:00	11:13	13				
		RPM Body	11:55	12:11	17				
RUN0010	11.09.2007		16:10	17:04	22				
		PPDS Trim	16:10	16:13	3	19	22		Overlap process with reuser of the plan. orders and direct creating of time series
		RPM Trim	16:24	16:37	13				
		PPDS Paint	16:25	16:26	1				
		RPM Paint	16:50	16:55	5				
		PPDS Body	16:50	16:51	1				
		RPM Body	17:00	17:05	5				

Figure 3-1 RPM run log

3.2 MRP test

The MRP test covers two steps: Vehicle and Made In Plant (MIP). Most of the tuning effects are based on the optimization done at the application level, the number of parallel processes used, and the optimization done at the database level. Run 14 is detailed in Chapter 5, “Material requirements planning” on page 105.

Conduction Run_ID	Date	Job / Step	start	end	SAP Value in min.	Total in min	Status	Comments
RUN0001	12.09.2007	MRP for Vehicle	10:00	10:31	31	128		One job with 2 steps (Vehicle and MIP)
		MRP for MIP	10:31	12:08	97			
RUN0002	12.09.2007	MRP for Vehicle	19:00		-	-		Two separate jobs for Vehicle and MIP, the result is not valid because of invalid data.
		MRP for MIP		21:08	-			
RUN0003	13.09.2007	MRP for Vehicle	16:15	16:44	29	121		One job with 2 steps (Vehicle and MIP), job started after applying SAP Notes for low level code.
		MRP for MIP	16:44	18:16	92			
RUN0004	14.09.2007	MRP for Vehicle	16:25	16:55	30	121		One job with 2 steps (Vehicle and MIP), job started after applying SAP Notes for low level code and DB patch
		MRP for MIP	16:55	18:26	91			
RUN0005	14.09.2007	MRP for Vehicle	18:55	19:13	18	67		One job with 2 steps (Vehicle and MIP), job started withou restore and any changes in system. This is **NOT** a initial MRP run.
		MRP for MIP	19:13	20:02	49			
RUN0006	14.09.2007	MRP for Vehicle	20:55	21:09	14	63		One job with 2 steps (Vehicle and MIP), job started with DB RunStat and changes for process but without restore. This is **NOT** a initial MRP run.
		MRP for MIP	21:09	21:58	49			
RUN0007	17.09.2007	MRP for Vehicle	9:50		-	-		One job with 2 steps (Vehicle and MIP), job was canceled because of hanging process in Report SAPSSMY1 which was probably caused by an
		MRP for MIP		11:18	-			
RUN0008	17.09.2007	MRP for Vehicle and MIP	15:50		-	-		We identified during the run that the modification in system has only minor improvement for performance. Budi MD_MRP_RUN_PARALLEL has not
				17:15	-			
RUN0009	17.09.2007	MRP for Vehicle	20:08	22:34	94	94		MRP for Vehicle in 9 parallel jobs for 9 plants. The job for HV21 was started
		MRP for MIP	22:51	23:28	37			MRP for MIP in one job.
RUN0010	18.09.2007	MRP for Vehicle	10:15	10:40	25	128		One job with 2 steps (Vehicle and MIP), job started after patitioning in tabel PLAF and RESB.
		MRP for MIP	10:40	12:23	103			
RUN0011	18.09.2008	MRP for Vehicle	17:43	18:06	23	39		One job with 2 steps (Vehicle and MIP), job started after patitioning in tabel PLAF & RESB, modification in BADI.
		MRP for MIP	18:06	18:22	16			
RUN0012	19.09.2007	MRP for Vehicle	16:01	16:21	20	39		One job with 2 steps (Vehicle and MIP), job started with 20% of planning horizon.
		MRP for MIP	16:21	16:40	19			
RUN0013	19.09.2007	MRP for Vehicle	19:25	19:41	16	26		One job with 2 steps (Vehicle and MIP), job started with 20% of planning horizon.
		MRP for MIP	19:41	19:51	10			
RUN0014	20.09.2007	MRP for Vehicle	20:47	21:03	16	26		One job with 2 steps (Vehicle and MIP), job started with 20% of planning horizon.
		MRP for MIP	21:03	21:13	10			
RUN0015	21.09.2007	MRP for Vehicle	12:38	13:12	34	34		We stared 9 parallel jobs for Vehicle and 20 parallel jobs for MIP.
		MRP for MIP	13:15	13:34	19			
RUN0016	21.09.2007	MRP for Vehicle and MIP	16:00	16:35	35	35		One job with 2 steps (Vehicle and MIP). Job started with 85 parallel processes and increased number range in table RESB and PLAF.

Figure 3-2 MRP run log

3.3 Backflush test

The execution of the backflush test occurred with two different operating processes, either as a *single-step process* or as a *two-step process*. Our best results were achieved on Runs 10 and 14. Chapter 6, “Discrete industries production confirmation: Backflush” on page 131 presents the details.

Conduction Run_ID	Date	Job / Step	start	end	SAP values duration in min	Total duration in min	Status	Comments
RUN0001	24.09.2007	Single-step Backflush	18:30	19:19	-	-		started with 40 parallel processes. Job was canceled because of locks
RUN0002	25.09.2007	Single-step Backflush	10:50	12:52	122	122		started without parallel processes.
RUN0003	25.09.2007	Single-step Backflush	13:50	14:22	-	-		started with 2 parallel processes. Job was canceled because of locks coursed by parallel processes.
RUN0004	25.09.2007	1st step of two-step Backflush	17:15	18:09	52	52		(PP04). Job started without parallel processes.
RUN0005	25.09.2007 26.09.2007	1st step of two-step Backflush	20:30	21:13	43	396		11 jobs for 1st step of two-step Backflush and each one with one hour's data volume (PP06 - PP15). Job started without parallel processes. Jobs with user PP06 (production 6 AM - 7 AM) and PP07 (production 7 AM - 8 AM) finished in short time because there is not production between 6 AM and 8 AM.
		1st step of two-step Backflush	21:13	21:14	1			
		1st step of two-step Backflush	21:14	21:14	1			
		1st step of two-step Backflush	21:14	21:58	44			
		1st step of two-step Backflush	21:58	22:42	44			
		1st step of two-step Backflush	22:42	23:27	45			
		1st step of two-step Backflush	23:27	0:12	45			
		1st step of two-step Backflush	0:12	0:55	43			
		1st step of two-step Backflush	0:55	1:39	44			
		1st step of two-step Backflush	1:39	2:22	43			
		1st step of two-step Backflush	2:22	3:05	43			
RUN0006	26.09.2007		10:25		921	921		One job with one step for 10 hours' data volume (50% of daily volume)
	27.09.2007	2nd step of two-step Backflush		1:46				
RUN0007	26.09.2007	1st step of two-step Backflush **with** activities posting	14:50	15:01	11	64		One job with 2 steps for one hour's data volume (PP16). The activities have been posted with 4 parallel processes in first step. The rest items have been posted without parallel processes.
		1st step of two-step Backflush **without** activities posting	15:01	15:54	53			
RUN0008	26.09.2007	1st step of two-step Backflush	18:50	19:33	43	43		(PP17). Job started with 5000 max. material items in package and without parallel processes
RUN0009	27.09.2007	Modified Single-step Backflush	11:16	11:52	36	36		(PP20). This is not standard KPI because of project scope coding modification. 10 parallel processes were configured.
RUN0010	27.09.2007	Modified Single-step Backflush	16:00	16:26	26	27		One job for modified Single-step Backflush with one hour's data volume (PP21). This is not standard KPI because of project scope coding
		Backflush	16:27	16:28	1			
RUN0011	27.09.2007	1st step of two-step Backflush	21:07	21:41	34	34		(PP22). Job started with 10 parallel processes and 5000 max. material items in package. In parallel we started also the payroll and load runner.
RUN0012	27.09.2007	1st step of two-step Backflush	22:05	22:39	34	34		(PP23). Job started with 40 parallel processes and 5000 max. material items in package. At this moment, the payroll job and load runner which started in last run were still running.
RUN0013	27.09.2007		23:20		500	500		One job with one step for 4 hours' data volume. At this moment, the payroll job was still running.
	28.09.2007	2nd step of two-step Backflush		7:40				
RUN0014	28.09.2007	Modified Single-step Backflush	14:45	15:16	31	31		(PP24). This is not standard KPI because of project scope coding modification. 80 parallel processes were configured. In parallel we started also the payroll and load runner.
RUN0015	28.09.2007	Modified Single-step Backflush	15:50	16:22	32	32		(PP01). This is not standard KPI because of project scope coding modification. 80 parallel processes were configured. At this moment, the payroll job and load runner which started in last run were still running.

Figure 3-3 BF run log

3.4 ML test

The ML test is divided into multiple steps corresponding to a required process for generating documents on calculated data. Complete details about Run 10 are presented in Chapter 7, “Materials ledger” on page 161. The KPI was achieved by working at the code level and at the database optimization level.

RUN0001		RUN01 SELECT MATERIALS	21:03:30	21:06:09	159	668.47		Test run type CASE1 for ML with 60 parallel processes, paket size 50, FI number range buffered (main memory, bufsize 2000)
		RUN01 DETERMINE COSTING SEQ.	21:06:09	3:02:15	21366			
		RUN01 SNGL-LVL PRICE DETERM.	3:02:16	5:19:27	8231			
		RUN01 MULTILEVEL PRICE DETER	5:19:27	5:31:57	750			
		RUN01 CLOSING ENTRY	5:31:57	8:11:59	9602			
RUN0002	16.10.2007 17.10.2007	RUN02 SELECT MATERIALS	11:58:26	12:01:33	187	256.03		Test run type CASE1 for ML with 60 parallel processes, paket size 50, FI number range buffered (main memory, bufsize 2000)
		RUN02 DETERMINE COSTING SEQ.	13:52:15	14:04:00	712			
		RUN02 SNGL-LVL PRICE DETERM.	14:10:39	14:23:12	753			
		RUN02 MULTILEVEL PRICE DETER	17:31:08	21:19:38	13710			
RUN0003	18.10.2007	RUN03 SELECT MATERIALS	10:05:18	10:08:03	165	212.50		Test run type CASE1 for ML with 60 parallel processes, paket size 50, FI number range buffered (main memory, bufsize 2000) - modified coding! for Closing step
		RUN03B DETERMINE COSTING SEQ.	10:48:06	10:55:10	424			
		RUN03 SNGL-LVL PRICE DETERM.	11:15:32	11:28:09	757			
		RUN03 MULTILEVEL PRICE DETER	11:52:48	13:58:14	7526			
		RUN03 CLOSING ENTRY	14:10:02	15:14:40	3878			
RUN0004	18.10.2007	RUN04 CLOSING ENTRY NO BUFFER	0.8364	0.9687	11431	190.52		Client requested test run type CASE1 for ML ONLY Closing step step with 60 parallel processes, paket size 1, FI number range NOT buffered - modified coding! for Closing step
RUN0005	19.10.2007	RUN05 SELECT MATERIALS	14:54:01	14:56:47	166	246.10		Test run type CASE1 for ML with 60 parallel processes, paket size 50, FI number range buffered (main memory, bufsize 2000) - modified coding! for Closing step
		RUN05 DETERMINE COSTING SEQ.	14:56:47	15:04:23	456			
		RUN05 SNGL-LVL PRICE DETERM.	15:04:23	15:19:04	881			
		RUN05 MULTILEVEL PRICE DETER	15:19:04	17:31:39	7955			
		RUN05 CLOSING ENTRY	17:31:39	19:00:07	5308			
RUN0006	19.10.2007	RUN06 MULTILEVEL PRICE DETER	0.88885	0.92891	3461	57.68		Test run type CASE1 for ML with 60 parallel processes, paket size 50, FI number range buffered (main memory, bufsize 2000) - modified coding! for Closing step
RUN0007	19.10.2007	RUN07 SELECT MATERIALS	22:23:48	22:27:50	242	42.93		Test run type CASE1 for ML with 60 parallel processes, paket size 50, FI number range buffered (main memory, bufsize 2000) - modified coding! for Closing step
		RUN07 DETERMINE COSTING SEQ.	22:28:28	22:35:28	420			
		RUN07 SNGL-LVL PRICE DETERM.	22:37:59	22:55:10	1031			
		RUN07 MULTILEVEL PRICE DETER	23:00:46	23:15:29	883			
RUN0008	22.10.2007	RUN08 MULTILEVEL PRICE DETER	9:37:22	11:23:21	6359			Test run type CASE1 for ML with 60 parallel processes, paket size 50, FI number range buffered (main memory, bufsize 2000) - modified coding! for Closing step
		RUN08B MULTILEVEL PRICE DETER	11:30:16	13:08:03	5867			
RUN0009	22.10.2007	RUN09 SELECT MATERIALS	20:34:48	20:37:50	182	198.73		Test run type CASE1 for ML with 60 parallel processes, paket size 50, FI number range buffered (main memory, bufsize 2000) - modified coding! for Closing step
		RUN09 DETERMINE COSTING SEQ.	20:39:45	20:47:14	449			
		RUN09 SNGL-LVL PRICE DETERM.	20:51:04	21:08:16	1032			
		RUN09 MULTILEVEL PRICE DETER	21:15:28	22:22:48	4040			
		RUN09 CLOSING ENTRY	22:22:48	0:06:29	6221			
RUN0010	23.10.2007	RUN10 SELECT MATERIALS	15:41:41	15:44:33	172	217.70		Test run type CASE1 for ML with 60 parallel processes, paket size 50, FI number range buffered (main memory, bufsize 2000) - modified coding! for Closing step & Loadrunner users
		RUN10 DETERMINE COSTING SEQ.	15:44:32	15:52:12	460			
		RUN10 SNGL-LVL PRICE DETERM.	15:52:12	16:08:20	968			
		RUN10 MULTILEVEL PRICE DETER	16:08:20	17:21:33	4393			
		RUN10 CLOSING ENTRY	17:21:34	19:19:23	7069			
RUN0011	24.10.2007	RUN11 SELECT MATERIALS	11:56:14	11:58:56	162	196.20		Test run type CASE3 for ML with 60 parallel processes, paket size 50, FI number range buffered (main memory, bufsize 2000) - modified coding! for Closing step & Loadrunner users
		RUN11 DETERMINE COSTING SEQ.	11:58:55	12:06:17	442			
		RUN11 SNGL-LVL PRICE DETERM.	12:06:17	12:23:11	1014			
		RUN11 MULTILEVEL PRICE DETER	12:23:11	13:19:09	3358			
		RUN11 CLOSING ENTRY	13:19:09	15:12:25	6796			
		RUN11 CLOSING REVERSAL	16:10:57	17:47:54	5817			
RUN0012	24.10.2007	RUN12 CLOSING ENTRY	17:56:28	19:29:53	5605	93.42		Test run type CASE3 for ML with 60 parallel processes, paket size 50, FI number range buffered (parallel buffering, bufsize 2000) - modified coding! for Closing step
RUN0013	24.10.2007	RUN13 CLOSING REVERSAL	20:22:03	21:30:35	4112	73.05		Test run type CASE3 for ML with 60 parallel processes, paket size 50, FI number range buffered (parallel buffering, bufsize 2000) - modified coding! for Closing step
		RUN13 CLOSING ENTRY	21:36:28	22:49:31	4383			
RUN0014	25.10.2007	RUN14 SETTLEMENT	9:44:17	9:56:46	749	12.48		Test run Settlement with 40 parallel processes
RUN0015	25.10.2007	RUN15 SELECT MATERIALS	14:11:42	14:14:31	169	25.28		Test run type CASE1 for ML with 60 parallel processes, paket size 50, FI number range buffered (parallel buffering, bufsize 2000) - modified coding! for Closing step & Loadrunner
		RUN15 DETERMINE COSTING SEQ.	14:14:30	14:22:01	451			
		RUN15 SNGL-LVL PRICE DETERM.	14:22:01	14:36:58	897			
		RUN15 MULTILEVEL PRICE DETER	14:36:58					
		RUN16 SELECT MATERIALS	17:23:17	17:27:53	276			Test run type CASE1 for ML with 60

Figure 3-4 ML run log

Part 2

Test details

In this part we provide details about how we configured our test environment, conducted the benchmark tests, and tuned the four major application scenarios, as well as the non-functional requirement tests, for optimum performance.

This part contains the following:

4

Rapid Planning Matrix

The RPM scenario takes place in the APO system. The purpose of this test is to verify that the RPM job scenarios can be completed within the required time frame—30 minutes—on the IBM system z and p configuration. This test also provides an indication of the optimal settings for those parameters that are most important for performance.

Ten runs were performed for this KPI. Run 7 and Run 10 are discussed in detail here. These particular runs were chosen for discussion because they use two different methods of handling the RPM execution process.

The topics included in this chapter are:

4.1 RPM test overview

This section provides an overview of the RPM test scenarios, which are based on iPPE and Rapid Planning Matrix Vehicle, along with data generation, parameters set up, test execution steps, the test deliverable, target time, and reset mechanism.

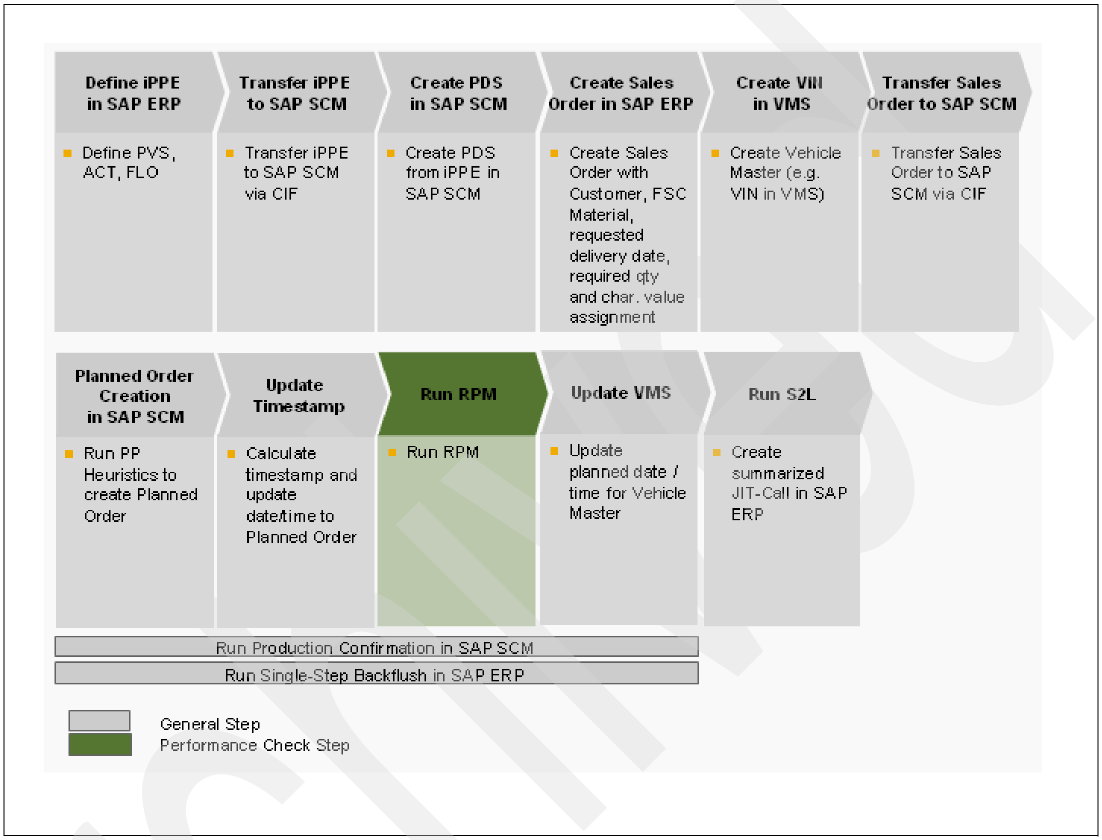

Figure 4-1 Test scenario 1a - RPM

4.1.1 Purpose of the test

The RPM scenario takes place in the APO system. The purpose of the test is to execute the RPM jobs in 30 minutes or less on the IBM System z and System p configuration. The test also provides an indication of the optimal settings for those parameters that have the greatest effect on performance.

4.1.2 Master data and volumes

The master data to consider for this test includes the iPPE structures and material master. For the scope of the performance test it is assumed that the configurable material is defined at *FSC* (full spec code = model index level and year).

The customer uses a three-level iPPE structure: Finished product, Body in paint, and Body in white. There will be orders at each of these three levels, and on each level there will be 1800 material variants for each model.

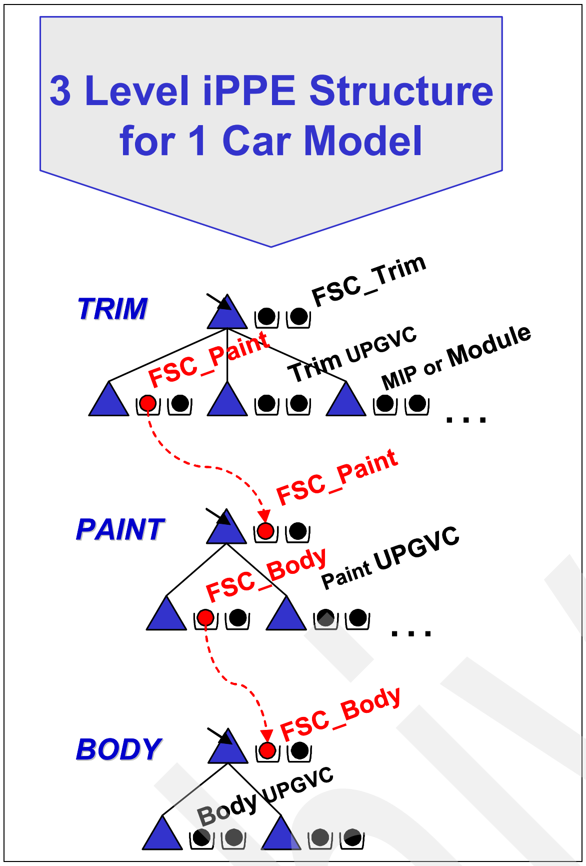

Figure 4-2 Customer PVS Structure

Table 4-1 shows the quantities and types of data generated in zPOC for this scenario.

Table 4-1 Generated data type and quantities

Item	Number of data fields (rounded)
Number of vehicles produced per month	187500 (7500/workday)
Number of orders per configurable material	8929
Configurable material (model/year)	21*3
Characteristics on FSC	Max. 300, average 200 per vehicle
Material variants on finished product level	1800*21= 37800
Material variants on body in white level	1800*21= 37800
Material variants on painted body level	1800*21= 37800
Number of components per vehicle	1200
Number of components per finished product order	1000
Number of components per painted body order	20
Number of components per body in white order	200
iPPE structures	21*3
Material master records	2087400

4.1.3 Master data

The customer provided the following data, which is required for the RPM execution:

- 23 plants:
  - 7 vehicle plants
  - 6 engine plants
  - 3 T/M plants
  - 2 raw material plants
  - 1 seat plant
  - 4 additional plants, like sales, knock down (KD), and so forth
- Material master: 630000 (for trim line)
- iPPE: 21 access nodes with 3 levels, routings for each of the 21 model code/year and level, line design for each plant and level
- 21 models (iPPE access nodes) are equally distributed in the 7 vehicle plants
- Production versions
- Cost collectors with preliminary cost estimation and standard cost estimation for each FSC
- PDS of all 63 structures (21 models * 3 levels = 63 PDS)
- Customer master
- Sales orders for two months (187500 per month)

Table 4-2 summarizes the parameter settings for the execution of this scenario.

Table 4-2 Settings

Parameter	Settings
Transaction code	/SAPAPO/CDPSB0 - production planning run
Time profile	SAP001 (planning period -2 to +3 months)
Propagation range	SAPALL
Heuristic	ZAP_MRP_001 (Comp. acc. LLevl Code)
Add-on program	Upload sales order

4.1.4 Test execution

We performed the following steps for this scenario:

1. Run Heuristic for finished product material in all plants in parallel
2. Run RPM for finished product material in all plants in parallel
3. Run Heuristic for painted body material in all plants in parallel
4. Run RPM for painted body material in all plants in parallel
5. Run Heuristic for body in white material in all plants in parallel
6. Run RPM for body in white material in all plants in parallel
7. Time series generation

4.1.5 Test deliverable

The following results were generated:

- Time series generation: Confirm that generating the time series after the BOM explosion gives better performance then generating it at the same time.
- Test series to determine optimal degree of parallelism: Measure different degrees of parallelism, specifically 10, 20, 30, 40, 50, 60, 70, and 80 parallel processes used for the RPM call.

4.1.6 Test reset mechanism

The test is reset by shutting down the servers and going back to a backup in the state directly before executing the test.

4.1.7 Target time

The target time for the test is 30 minutes.

4.1.8 Activities

For the verification of the runtime we created a job that included the steps listed in 4.1.4, “Test execution” on page 82.

We ran the tests in two different ways:

1. Run all the steps in sequence, referred to as *Scenario 1* or *Run0007*.
2. Run the test with cascading steps to simulate a more realistic situation, referred to as *Scenario 2* or *Run0010*.

4.2 RPM Run0007 (Scenario 1)

Before the actual test started we generated Planned Orders to satisfy the demand coming from customer orders. We generated the Planned Orders for all three levels of the Bill of Material (finished product, painted body, and body in white). These Planned Orders were then used during the RPM Call as input elements.

We ran the RPM Call in the background with the help of the SAP job engine.

In the job definition you can specify that the job should use parallel processing, but this feature does not work for the RPM Call because each Job can use only one single session process.

This means that we need to create separate jobs to be able to use parallel processing. To avoid a bottleneck in the live cache we created seven jobs per plant for each level that can run in parallel (see Figure 4-3 on page 84).

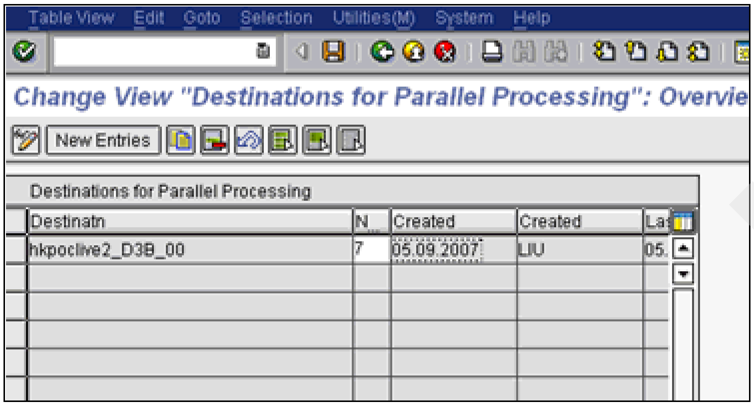

Figure 4-3 Destinations for Parallel Processing

Figure 4-4 shows the variant we created for background execution for the three levels of the Bill of Material.

Figure 4-4 Variant for RPM Call Trim

For the RPM calls of the body in white and the painted body the product selection varies (Body: M*_07_KHP, Paint: M*_07_KHB) but the other settings remain the same.

The settings for the production planning run are shown in Figure 4-5.

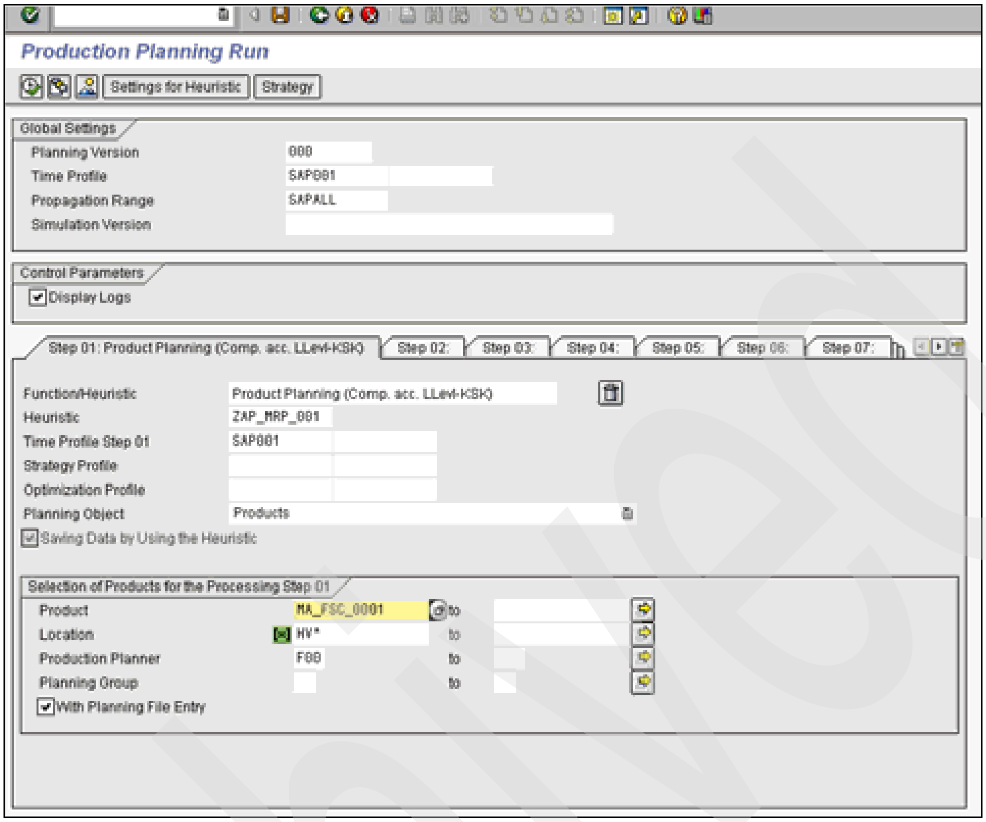

Figure 4-5 Production planning run

We changed a number of the parameters for performance reasons. The changes are summarized in Table 4-3 and discussed in detail in the steps that follow the table.

Table 4-3 Scenario 1 parameter changes

ID	Parameter	Old value	New value
01	Overwrite mode for the log area	unselected	selected
02	Create a log backup every __ minutes	300 minutes	180 minutes
03	Automatic log backup (After adjustments for Parameter 01 and 02)	Off	On
04	liveCache Database Parameter "SHAREDSQL_COMMANDCACHESIZE"	262144	2621440
05	Log files size	LOG001 with 1 GB LOG002 with 7 GB	LOG001 with 1 GB LOG002 with 7 GB LOG003 with 30 GB LOG004 with 30 GB
06	Number of parallel processing	40	66

1. We increased the log files sizes in live cache to avoid the automatic cancellation of jobs.
2. We changed the log writing mode to allow it to overwrite. By default, the log area can be overwritten only if the appropriate log entries have been backed up. You must therefore make regular log backups.
3. Due to limited disk space, and also to ensure optimal performance and throughput, we decided to change the configuration of the log area so that the logs can be overwritten (Figure 4-6). The log area is overwritten cyclically, even if one does not perform any log backups.

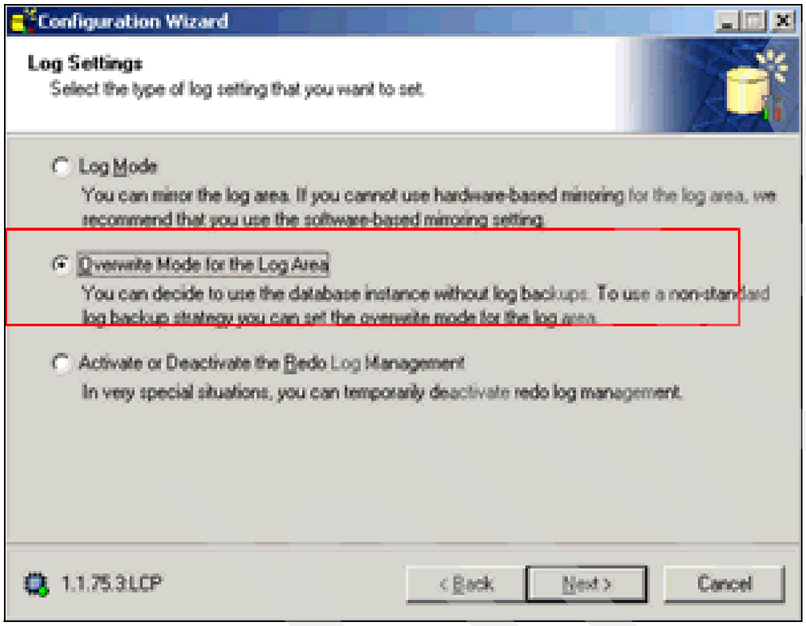

Figure 4-6 Parameter 01: Log settings

4. We decided that the log backup should be performed every 180 minutes to recycle storage space, as shown in Figure 4-7.

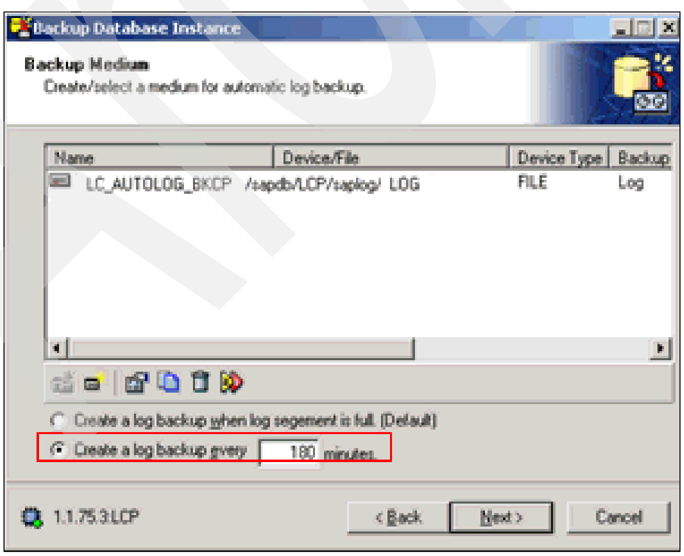

Figure 4-7 Parameter 02: Backup medium

5. We checked, with transaction LC10, that the Automatic Log backup was switched on (Figure 4-8).

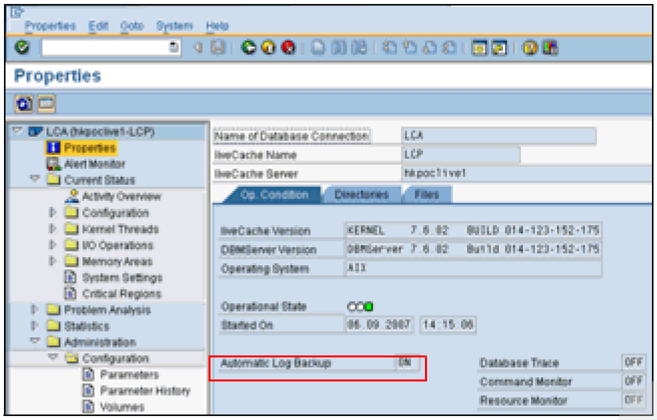

Figure 4-8 Parameter 03: Automatic log backup

6. We increased the parameter "SHAREDSQL_COMMANDCACHESIZE" to allow additional caching for the SharedSQL statements (Figure 4-9).

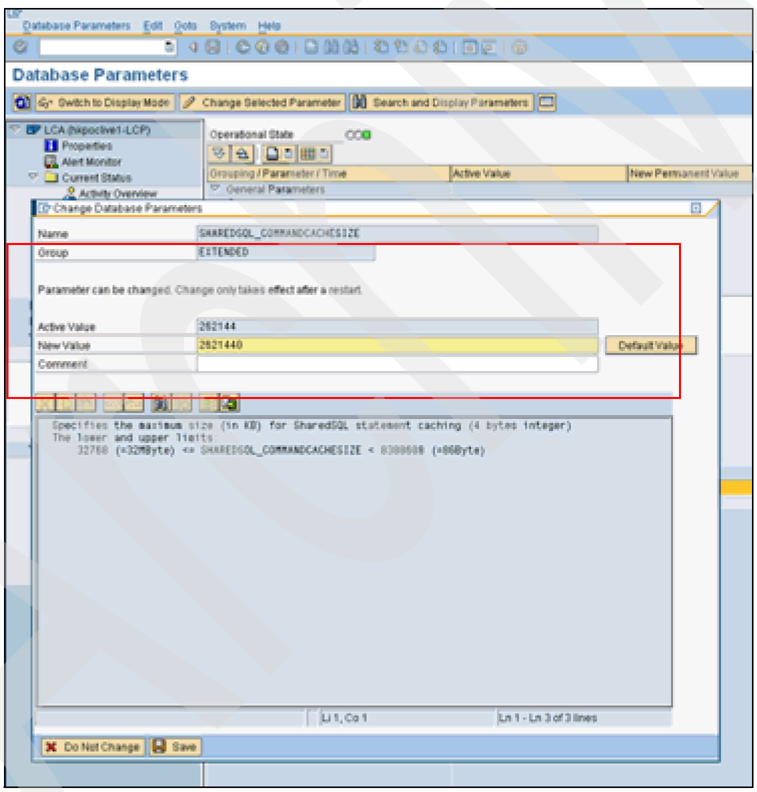

Figure 4-9 Database parameters

7. With the help of the Database Manager tool we checked that the new log areas had been configured and also that the size of the new log areas was correct (Figure 4-10).

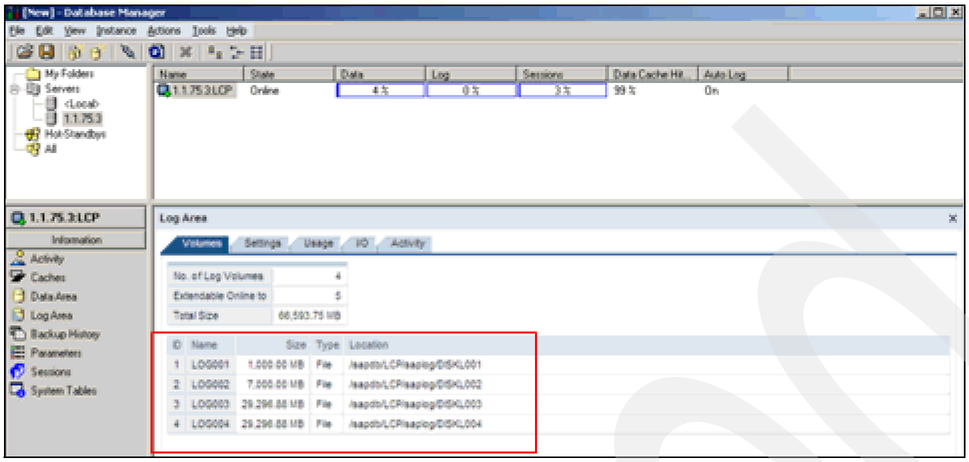

Figure 4-10 Database Manager tool

4.3 RPM Run0010 (Scenario 2)

In terms of day-to-day activity at the client site, Run 10 corresponds most closely to reality by using the most suitable process with overlap between the sub-tasks.

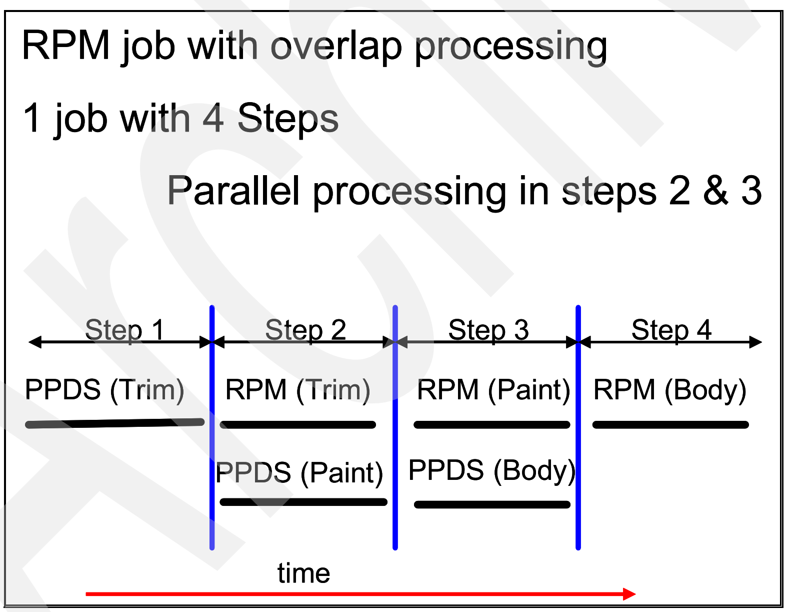

Figure 4-11 RPM Jobs with overlap processing

SAP application layer

In this scenario we ran the steps of the RPM Call in an cascading mode to test a more realistic situation. We performed the overlap process by specifying the reuse mode as "`Use suitable receipt elements`" in Heuristic ZAP_PP_002.

The cascading process steps are as follows:

1. Create Planned Orders for the customer orders on the finished product level.
2. Do an RPM Call of the finished product level to determine the dependent requirements for the body in white.
3. Use the result of the RPM Call for the finished product to do a planned order creation for the painted body.
4. Do an RPM Call for the painted body.
5. Use the result of the RPM call for the painted body to do a planned order creation for the body in white.
6. Do an RPM Call for the body in white.

General setup

Table 4-4 provides an overview of the SAP parameters that were changed in profile settings for the system. Details about each parameter follow the table.

Table 4-4 SAP parameter setup

ID	Parameter	Old value	New value
01	timestamp_duration_add	1800	60
02	Number of parallel sessions	40	23
03	Heuristic setting of ZAP_PP_002 Reuse mode	Delete unfixed receipts	Use suitable receipt elements
04	RPM time series	Do not generate	Generate directly
05	Live Cache parameter SHAREDSQL	Yes	No

Parameter 01: timestamp_duration_add

During the RPM Call process various time series are generated in the background. The time series generation is blocked if a reorg process is already running in the background and postponed to half an hour later even if such a long time is not needed. We modified the coding and set the parameter timestamp_duration_add to only one minute.

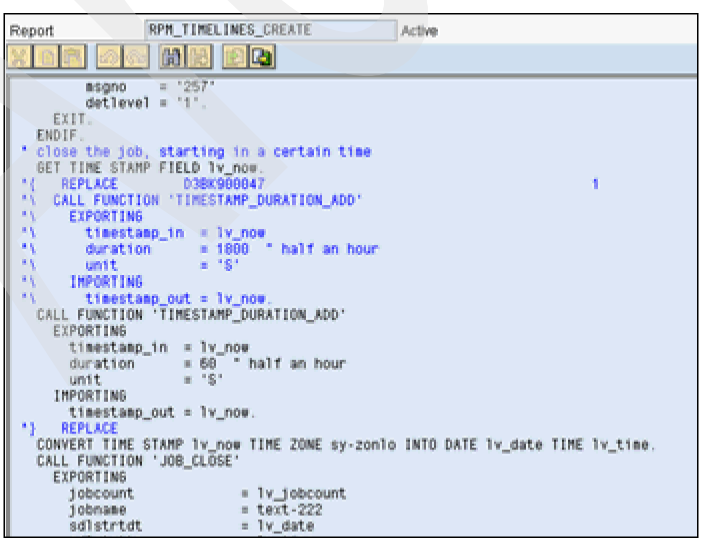


```
        msgno    = '257'
        detlevel = '1'.
    EXIT.
  ENDIF.
* close the job, starting in a certain time
  GET TIME STAMP FIELD lv_now.
*{   REPLACE        D38K900047                                        1
*\  CALL FUNCTION 'TIMESTAMP_DURATION_ADD'
*\    EXPORTING
*\      timestamp_in  = lv_now
*\      duration      = 1800  " half an hour
*\      unit          = 'S'
*\    IMPORTING
*\      timestamp_out = lv_now.
  CALL FUNCTION 'TIMESTAMP_DURATION_ADD'
    EXPORTING
      timestamp_in  = lv_now
      duration      = 60  " half an hour
      unit          = 'S'
    IMPORTING
      timestamp_out = lv_now.
*}   REPLACE
  CONVERT TIME STAMP lv_now TIME ZONE sy-zonlo INTO DATE lv_date TIME lv_time.
  CALL FUNCTION 'JOB_CLOSE'
    EXPORTING
      jobcount              = lv_jobcount
      jobname               = text-222
      sdlstrtdt             = lv_date
```

Figure 4-12 RPM_timelines_create

Parameter 02

The number of parallel processes is reduced from 40 to 23. The reason for this change is that a maximum of only 23 processes is needed in each step: 21 processes for 21 matrixes in each level and 2 processes for PP/DS.

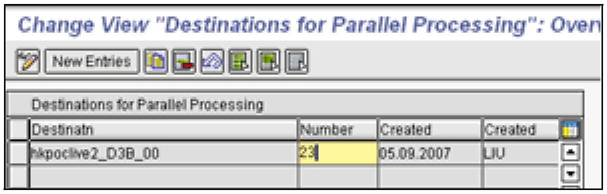

Figure 4-13 Destinations for parallel processing

Parameter 03

ZAP_PP_02 is a modified copy of the standard heuristic SAP_PP_02. SAP_PP_02 is available in every SAP system as the standard.

In the heuristic settings for ZAP_PP_02 we changed the value from “delete unfixed receipts” to “use suitable receipt elements” because this mode will be used in the customer’s real production mode. This setting means that we do not generate new planned orders for the amount of already existing ones (for example, from a previous run) to satisfy the demand.

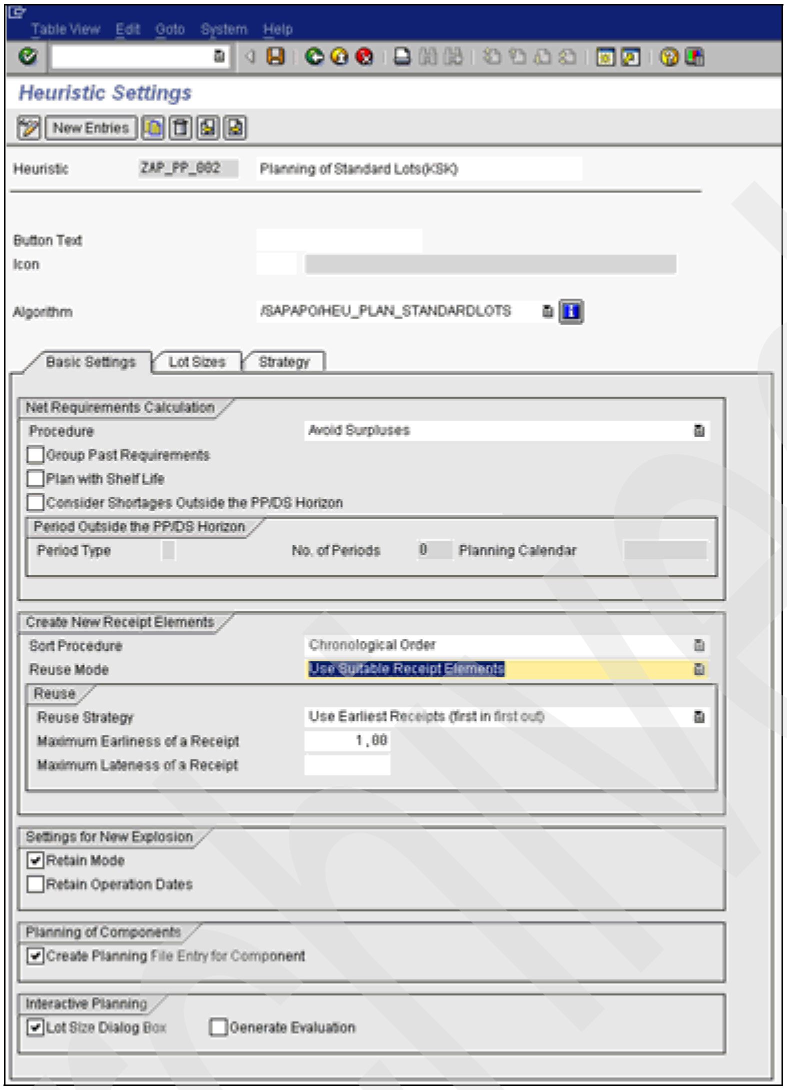

Figure 4-14 Heuristic settings

Parameter 004

For the RPM Call we left the parameter "RPM time series - generate directly" because this is done in parallel with the RPM Call process and does not use performance from the RPM Call. The REORG job after RPM might block the time series generation; by choosing this option we also make sure to eliminate such a possibility.

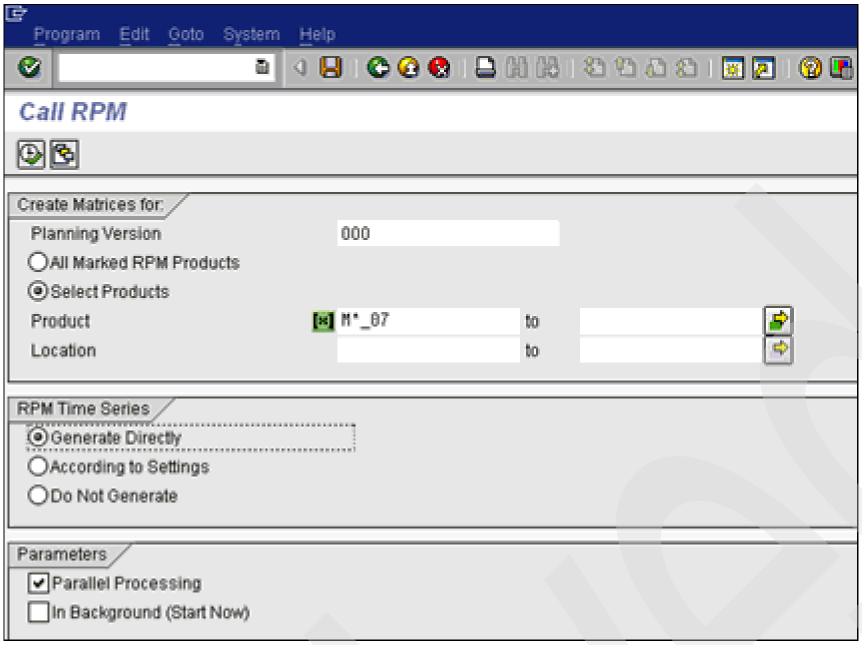

Figure 4-15 RPM Call Settings

4.4 Results overview

This section presents an overview of the test runs; details are provided in subsequent sections.

Scenario 1

Run_ID	Date	Job / Step	Start	End	Duration in min.
RUN0007	07.09.2007	RPM All	17:11	17:26	15

Figure 4-16 Run 7 results

Scenario 2

Run_ID	Date	Job / Step	Start	End	Duration in min.	RPM duration in min.	Total duration in min.
RUN0010	11.09.2207	PPDS Trim	16:10	16:13	3	19	22
	11.09.2207	RPM Trim	16:25	16:36	11		
		PPDS Paint	16:25	16:26	1		
		RPM Paint	16:50	16:54	4		
		PPDS Body	16:50	16:51	1		
		RPM Body	17:00	17:04	4		

Figure 4-17 Run 10 general results

Based on the given customer data (375000 fully configured vehicle orders) we achieved a run time of 15 minutes for scenario 1 and 22 minutes for scenario 2.

In both scenarios the processes are highly scalable and the KPI of 30 minutes requested by the customer is greatly improved upon.

The second approach is widely used by our real-world customers and is therefore the scenario recommended by SAP.

4.5 Analysis of System p infrastructure layer test results

In this section we present the system activity summary for the five System p server LPARs during the run. The whole run began at 16:10 and finished at 17:05. The KPI itself covers only the RPM steps for “Trim”, “Paint”, and “Body”, which lasted 22 minutes.

4.5.1 CPU details

In the following graphs we present the detailed CPU activity on LPARs hkpoclive1 and hkpoclive2. There is no CPU load on the other partitions because they are not used by the RPM process.

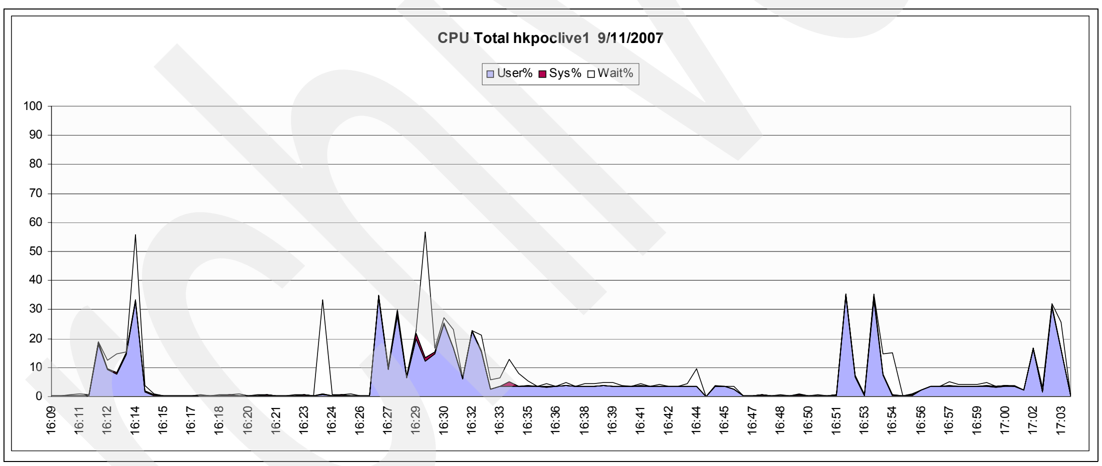

Figure 4-18 CPU details for hkpoclive1

In Figure 4-18 it is interesting to note that the CPU usage peak at 17:20 is due to a CPU I/O wait. At this particular moment the wait% is up to 63%. In 4.5.3, “Disk details” on page 95 we study the disk behavior to get an explanation of that point. There is almost no AIX kernel CPU usage; almost all is used by the liveCache itself.

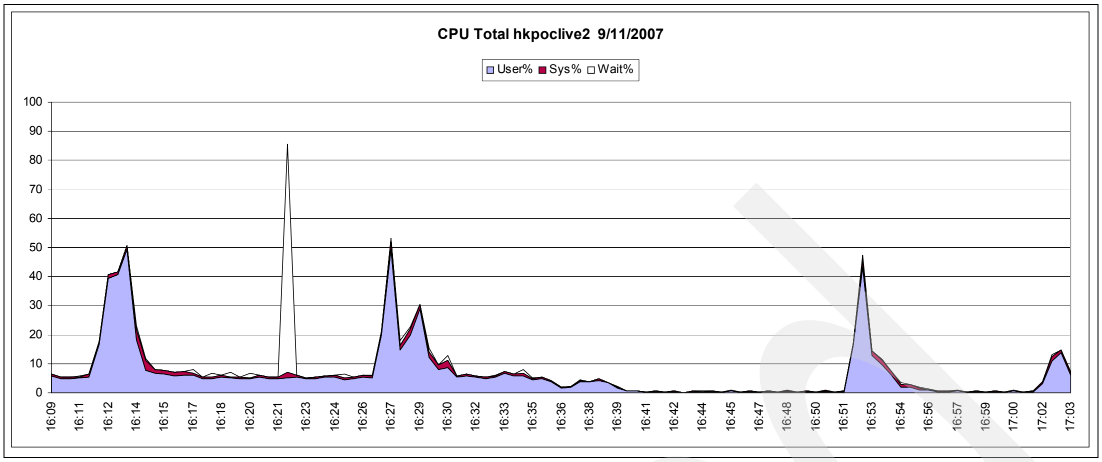

Figure 4-19 CPU details for hkpoclive2

For LPAR hkpoclive2, Figure 4-19 does not show any unexpected behavior. CPU usage peaks at 40% and there is almost no wait%. Looking at the CPU distribution, we would see a homogenous CPU load distribution in both cases.

4.5.2 Memory details

In this section we present the memory details at the LPAR level.

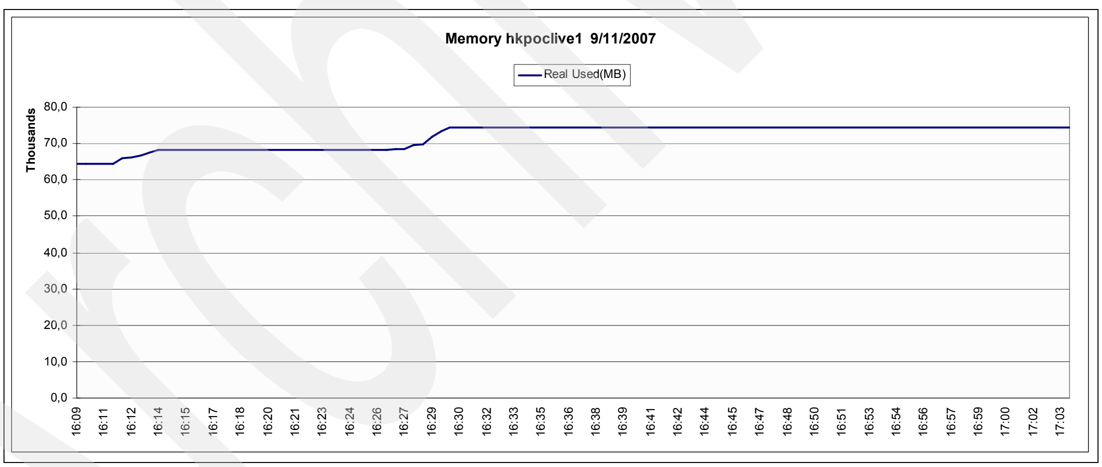

Figure 4-20 Free memory on hkpoclive1

At the moment of the run, hkpoclive1 had 196 GB RAM configured. As shown in Figure 4-20, hkpoclive1 used memory and available memory is about 135 GB during the whole run; the liveCache needed some more memory (15 MB). The system has no problem on that side.

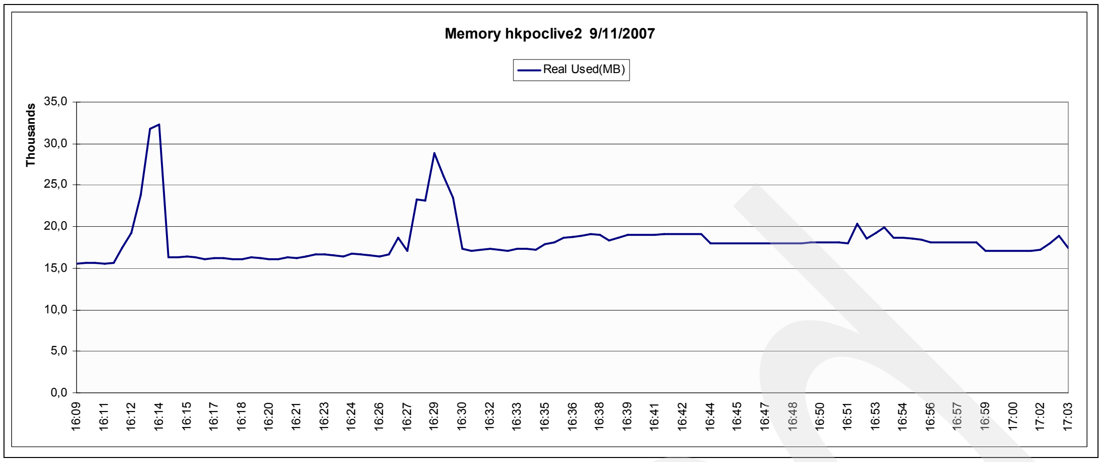

Figure 4-21 Free memory on hkpoclive2

During this run, the SCM SAP processes on hkpoclive2 needed some additional memory, as show in Figure 4-21. At the moment of the run hkpoclive2 had 56 GB RAM configured, but we can see that free memory is never less to 26 GB, which is about 50% of the configured memory.

4.5.3 Disk details

Disk transfers

Disk transfers on hkpoclive1 are shown in the next two figures.

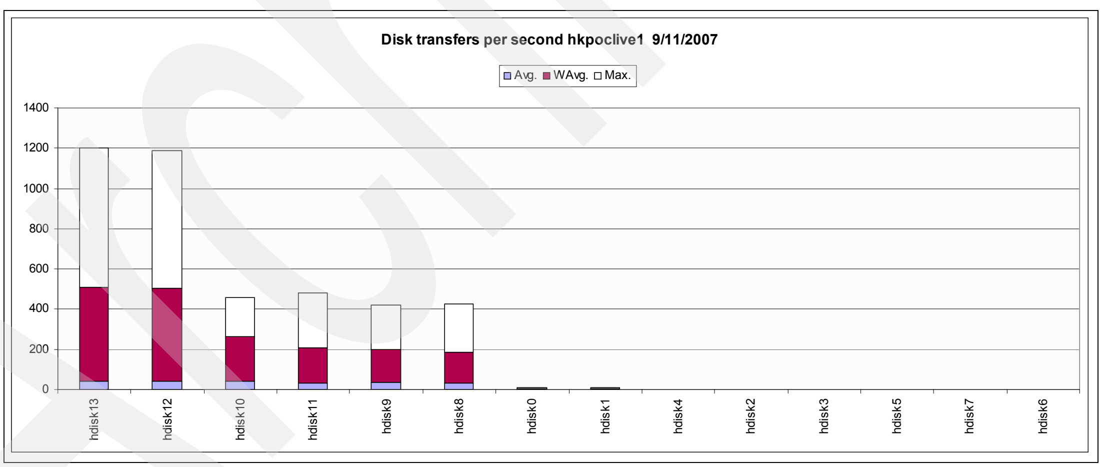

Figure 4-22 Disk transfers on hkpoclive1

In Figure 4-22, disk transfers on hkpoclive1, we can see that hdisk12 and hdisk 13 have a peak rate of up to 1200 I/O per second. Remember that they are LUNs defined on the storage subsystem, as explained in 2.2.5, “Storage subsystem configuration” on page 30. That is the reason we see so much I/O per second. Hdisk0 and hdisk1 hold rootvg, so we do not expect much activity there.

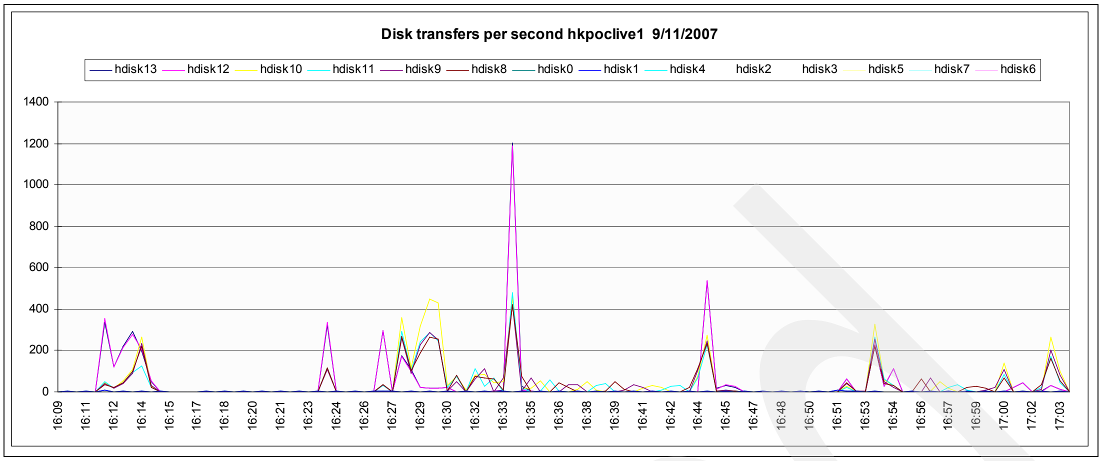

Figure 4-23 disks transfers on time frame

The interesting data in Figure 4-23 shows that there are peaks in disk activity at 16:33 and 16:44, corresponding to CPU peaks on hkpovlive2, which is logical and normal considering the application behavior.

The next two figures show disk activity on hkpoclive2.

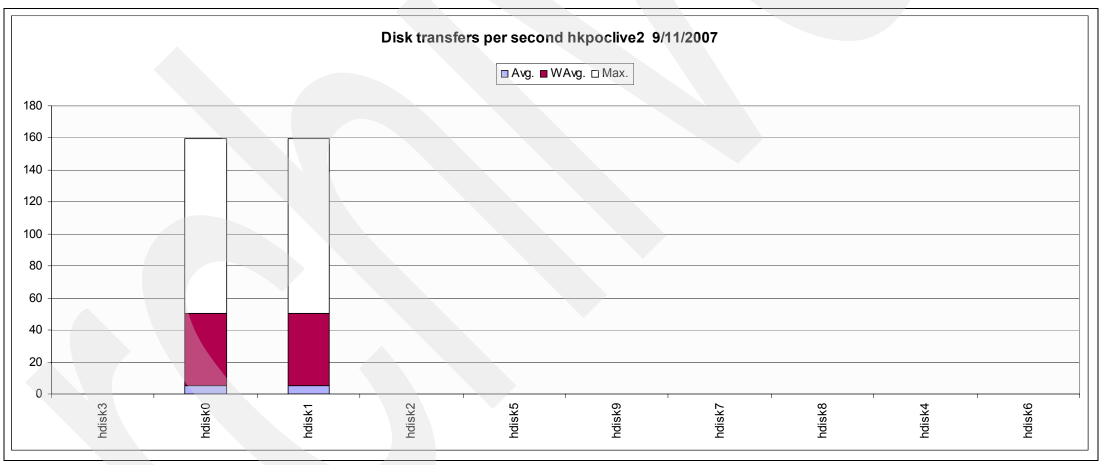

Figure 4-24 Disk transfers hkpoclive2

Regarding disk transfers on hkpoclive2, we can see Figure 4-24 that hdisk0 and hdisk1 (SAP application and operating system on the volume group rootvg) have peaks that extend almost their maximum capacity in I/O per second. We probably have a contention there.

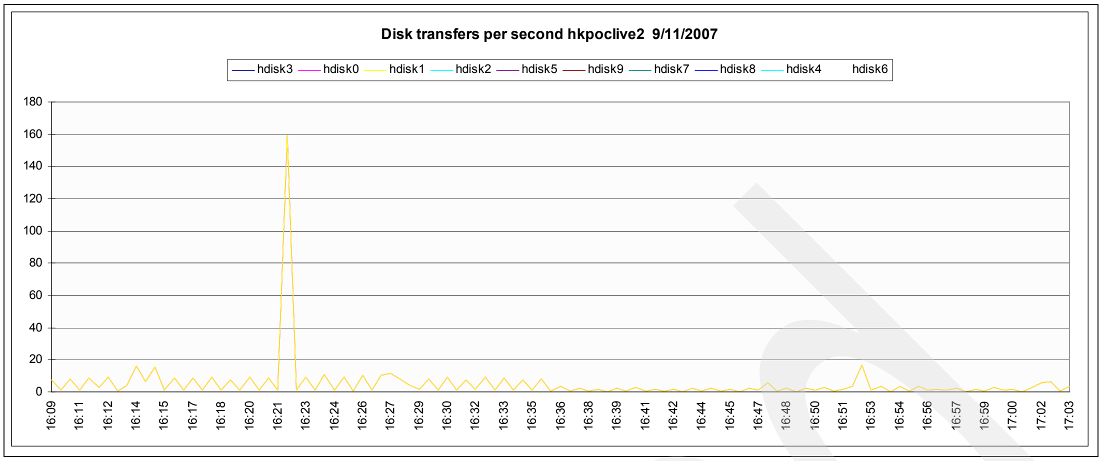

Figure 4-25 Disk transfers on time frame

Figure 4-25 shows a concrete peak in disk activity at 16:21, corresponding with the CPU peak. Here we first note that this peak does not affect the run because it is not sustained, but we do have a small contention there. The cause of this contention is that SAP writes on rootvg at the same time as the sustained activity, so the corrective action we adopted is to create a separate volume group on a separate disk, striping the logical volume, and thereby we eliminated the contention.

Regarding disk analysis, everything was OK during the run, and on a purely technical level the Montpellier team fixed the contention. We now check the I/O adapters.

I/O adapters

Regarding I/O adapters on hkpoclive1, Figure 4-26 shows that the load is well balanced.

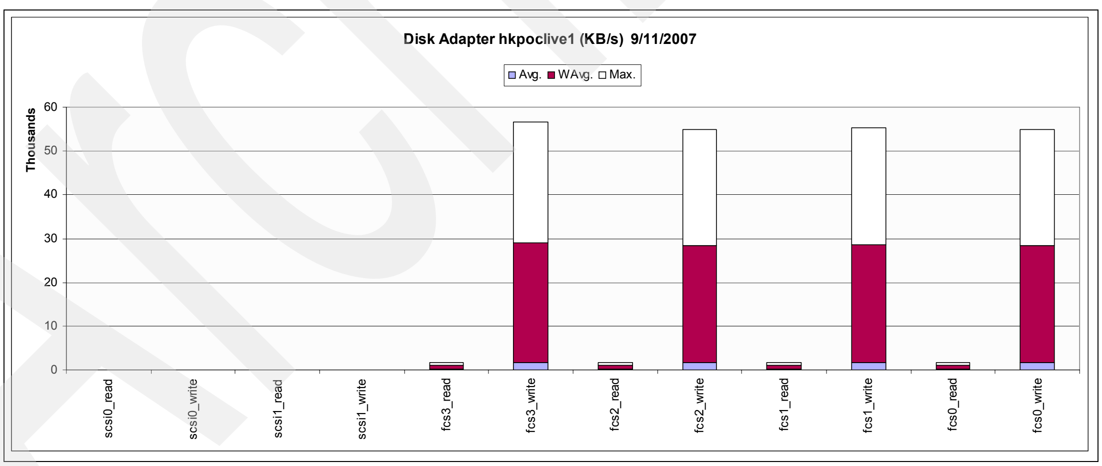

Figure 4-26 I/O adapters load

4.5.4 Network details

This section concerns the TCP/IP network usage profile.

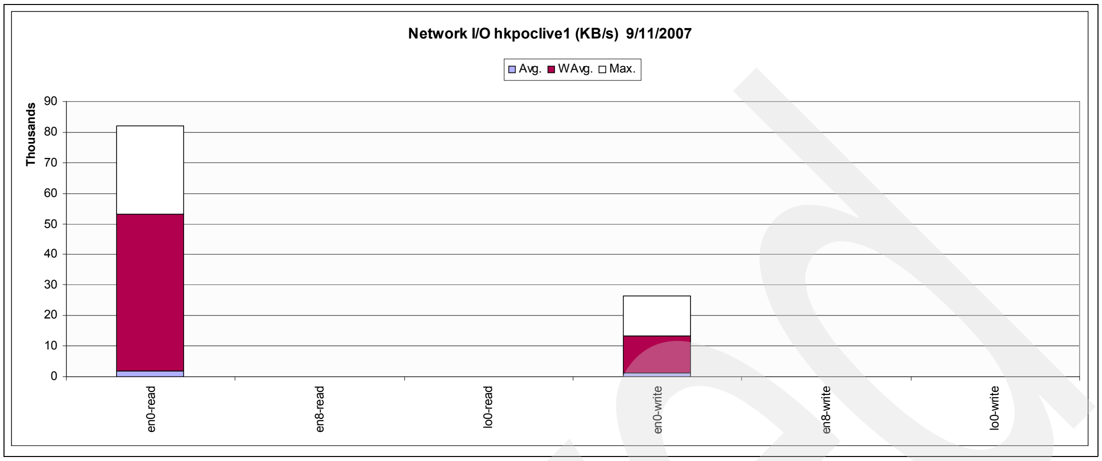

Figure 4-27 hkpoclive1 network I/O per adapter

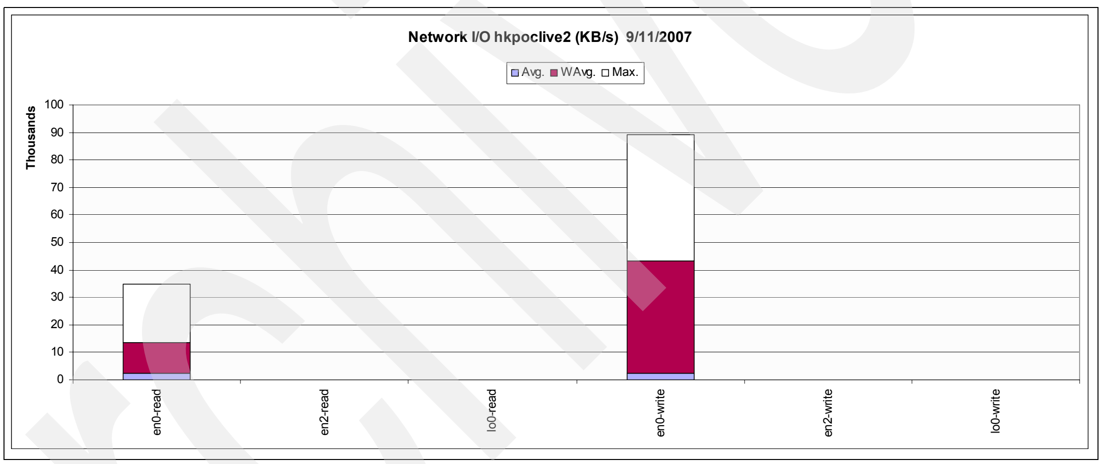

Figure 4-28 hkpoclive2 network I/O per adapter

Figure 4-28 and Figure 4-29 show the network load by adapter. We notice that the highest load is on en0 adapters, which is logical because these adapters are connected to the 1 Gbps network for server communications.

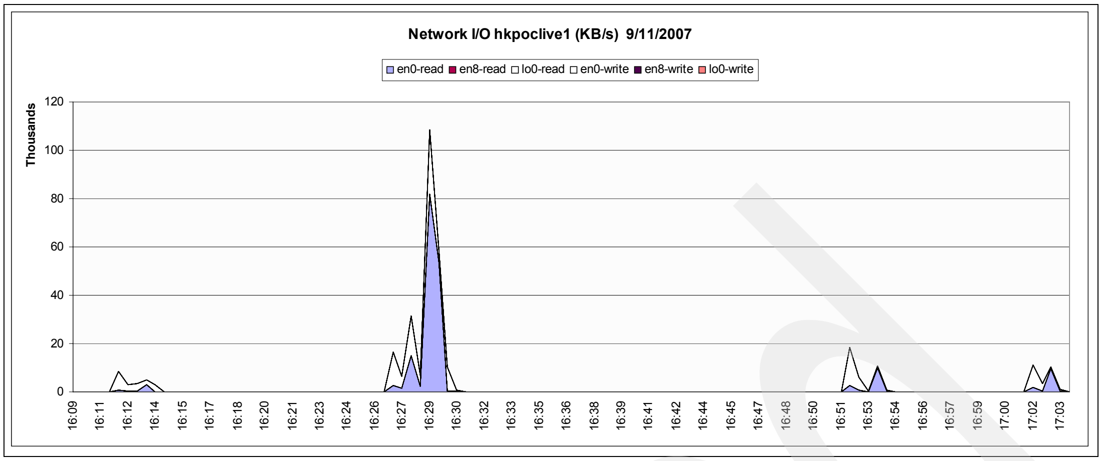

Figure 4-29 hkpoclive1 network I/O

We can also see that reads and writes are inverted in hkpoclive1 and hkpoclive2, which is also logical because these two servers have a very strong bidirectional communication.

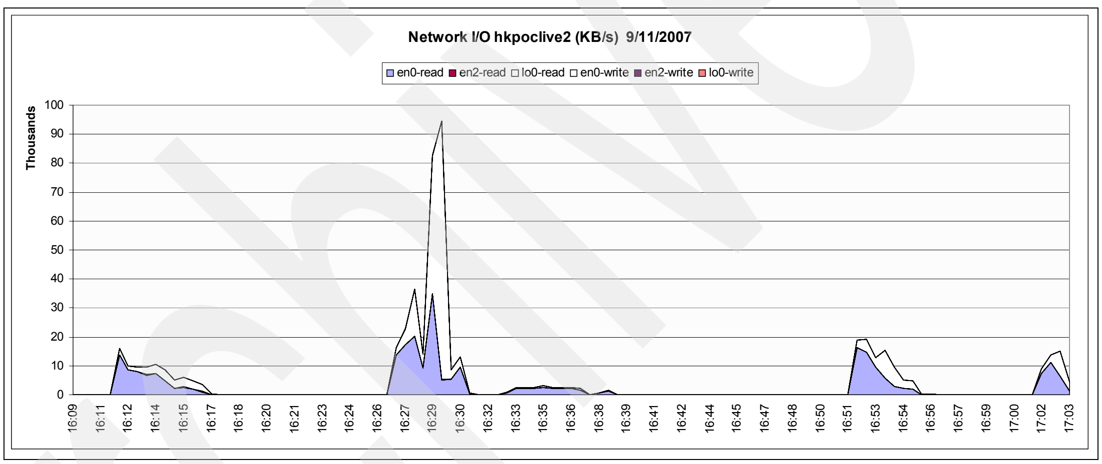

Figure 4-30 hkpoclive2 network I/O

No network errors were detected, and overall there was no contention due to the network architecture.

4.5.5 Top processes

In this section we determine what the top CPU and memory consumers are.

High CPU use

As Figure 4-31 shows, the top demands on the CPU during the run are labeled as kernel. These are the liveCache processes.

Figure 4-32 shows the top CPU demand processes are disp+work (SAP processes) and db2sysc (DB2 connect, connection manager between DB2/Z and SAP).

The CPU is primarily used by core SAP or DB2 processing.

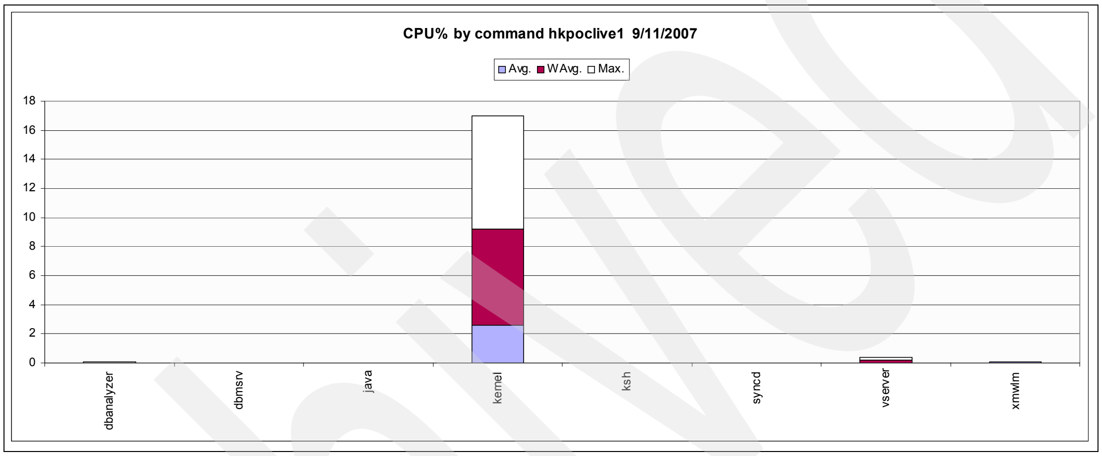

Figure 4-31 hkpoclive1 top CPU

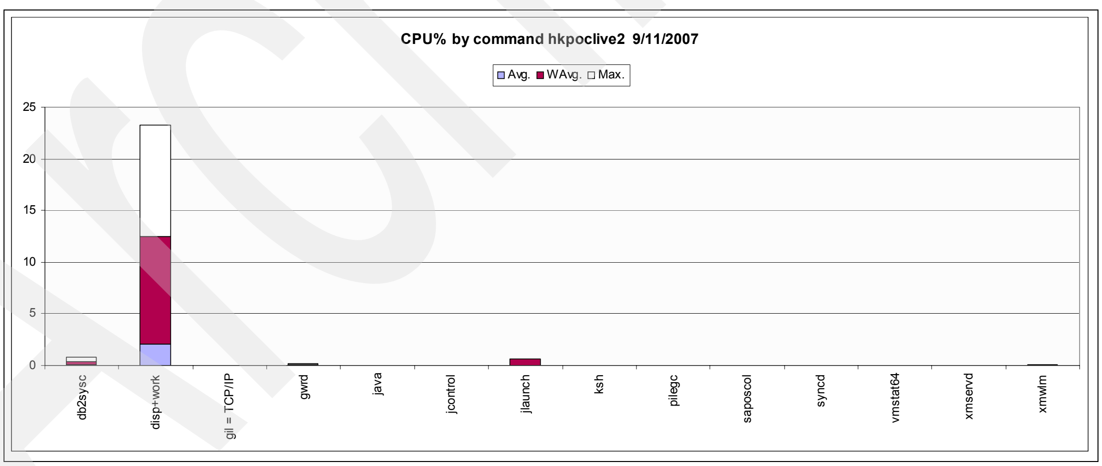

Figure 4-32 hkpoclive2 top CPU

High memory use

The next two figures deal with memory use.

In Figure 4-33 we see that during the run, the highest memory use is labeled as kernel. These are the liveCache processes.

Figure 4-34 shows the highest memory use per process was by disp+work (SAP processes) and db2sysc (DB2 connect, connection manager between DB2/Z and SAP).

The same conclusion can be made for memory use as for CPU use, namely that most of the memory consumed is due to the core SAP or DB2 processes.

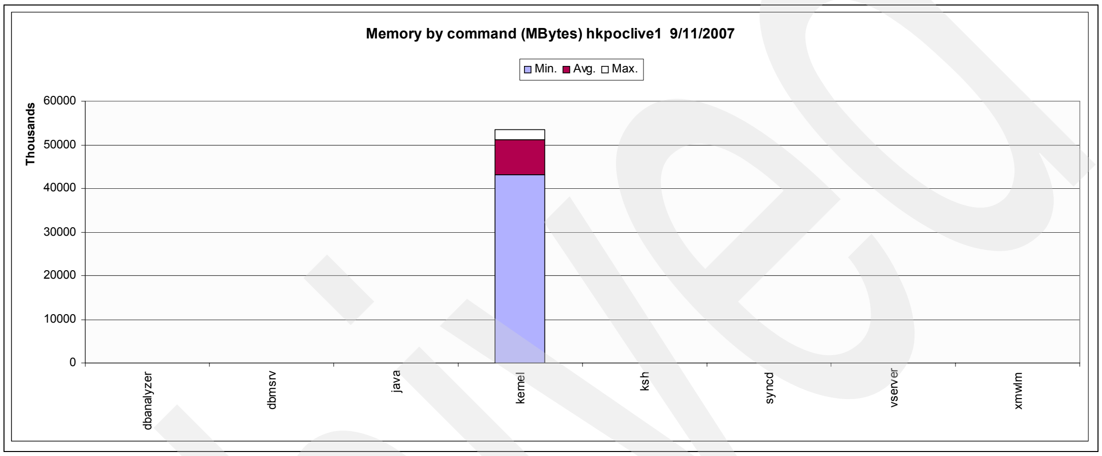

Figure 4-33 hkpoclive1 top memory

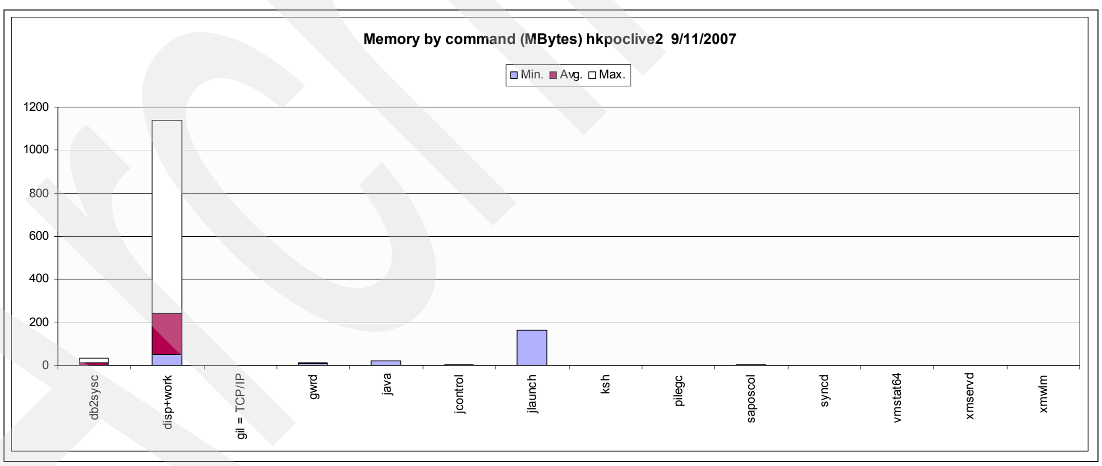

Figure 4-34 hkpoclive2 top memory

4.5.6 System p architecture conclusions

At the infrastructure level, System p servers perfectly responded to run 10 needs. Absolutely no negative impact was experienced at this level during the run.

4.6 Analysis of System z infrastructure layer test results

From System p performance analysis we saw that there is no activity on the ECC LPAR. As a consequence there is no need to analyze the activity on the MVH1 database, so we limit ourselves to only reporting about the SCM database (MVH2). The run consisted of four parts: PPDS-Trim, RPM-Trim & PPDS-Paint, RPM-Paint & PPDS Body, and RPM-Body. The following results were obtained:

1. PPDS-Trim began at 16:10 and finished at 16:13, lasting 3 minutes.
2. RPM-Trim & PPDS-Paint began at 16:25 and finished at 16:36, lasting 11 minutes.
3. PRM-Paint & PPDS-Body began at 16:50 and finished at 16:54, lasting 4 minutes.
4. RPM-Body began at 17.00 and finished at 17.04, lasting 4 minutes.

Figure 4-35 shows the LPAR summary for the three RPM steps (2, 3, and 4) for run 10.

The CPU utilization and the DASD rate for LPAR MVH2 is low. The peak CPU usage is 18% of eight CPs. The load shown gives no reason for further investigation or preparation of more detailed reports. There are no bottlenecks on MVH2.

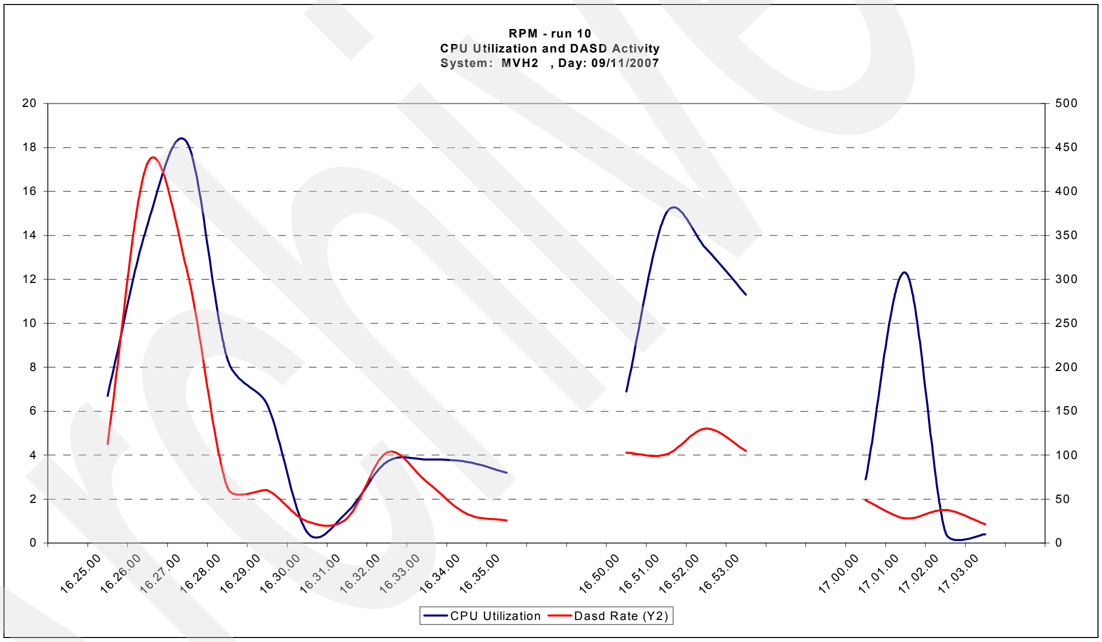

Figure 4-35 LPAR overview RPM run 10

Table 4-5 summarizes the time observed for DB2 processing using DB2 Accounting reports. All times are given in seconds and are averages based on the number of accounting records collected during the interval.

Table 4-5 DB2 times RPM run 10

RPM run 10	STEP 2	STEP 3	STEP 4
	in second	in second	in second
APPL (CL.1) Elapsed	0.018728	0.013775	0.014525
DB2 (CL.2) Elapsed	0.002969	0.001160	0.001567
DB2 (CL.2) CPU	0.001396	0.000755	0.000810
DB2 (CL.2) Suspended	0.001536	0.000393	0.000739
DB2 (CL.2) Not Accounted	0.000037	0.000013	0.000018
Accounted records #	400	150	150

The average elapsed time spent in DB2 is between 1 and 3 milliseconds. In step 2 and step 4 half of this time is spent in suspension. More than 96% of the suspension time in all three steps is because of synchronous write I/O on the DB2 logs.

Table 4-6 shows that the DB2 activity is very low and the number of SQL statements executed per second is few.

Table 4-6 SQL activity (Step 2)

RPM run 10	SQL activity per second
SELECT	0.00
INSERT	0.38
UPDATE	0.40
DELETE	0.05
PREPARE	0.22
DESCRIBE	0.02
DESCRIBE TABLE	0.00
OPEN	1.97
CLOSE	0.18
FETCH	1.97

Because the KPI is achieved and the DB2 activity is very low, no more investigation has been done.

4.7 Conclusions: Results and recommendations

The runs RUN0005 and RUN0007 (Scenario 1) are considered as valid test runs for the RPM KPI; the best run time achieved was 15 minutes. RUN0010 (Scenario 2) is also a valid test, which uses overlap mode to best simulate the client's daily activity.

Table 4-7 Run results

Run_ID	Date	Job / Step	Start	End	Duration in min.
RUN0005	07.09.2007	RPM All	14:06	14:22	16
RUN0007	07.09.2007	RPM All	17:11	17:26	15
RUN0010	11.09.2007	3 jobs	16:10	17:04	22

Results

During the performance test there were different ways to run the RPM because of the client's specific three-level iPPE (Trim, Paint, and Body).

The possible runtime strategies are:

- Parallelization of all 63 matrixes (all planned orders at the three levels are created beforehand).
- Level by level (planned orders are created level by level, RPM runs level by level, sequential jobs for the three levels).
- Overlap scenario (planned order trim for one model, RPM trim, planned order paint for one model, RPM paint, planned order body for one model, RPM body, job chain for models in parallel).

Whatever the strategy used, the KPI for RPM is met. Which concept the client uses will depend on creation of the planned order and updating of the production times (timestamp update) from the legacy system. This will be decided during the implementation project.

Regarding the infrastructure, most of the workload occurs on the SCM application and liveCache Database. The System z infrastructure is running well below capacity. Regarding System p we can see that liveCache is reading and writing on disks, and at the CPU level we can appreciate that this is a sequential process, as shown in Figure 4-18 on page 93. In fact, only one thread is really used.

There are two ways to improve the performance for this KPI:

- Change the (SAP) process internally to be able to achieve more parallelization and use more threads.
- Use a more powerful CPU (Power6 for instance) in the System p environment.

Recommendations

The complete RPM process consists of three parts: RPM, create time series, REORG_SOFT process.

After generating the time series the result is visible to the user. The REORG_SOFT process runs only in the background and is not included in the RPM runtime. During the performance test we had some problems with blocking of time series. For further discussion about the problem of creating the time series during REORG_SOFT processing, refer to the OSS-message 4319927 and consider the recommendations.

The creation of time series could run as a separate job, but this program cannot run in parallel. The recommendation is to run the RPM and creation of time series together, because then the time series are generated also in parallel.

5

Material requirements planning

This chapter describes our experiences with benchmark testing of the MRP module.

It includes the following topics:

5.1 MRP test overview

This test involves running the entire material requirement planning module, excluding the pre-step RPM, which was tested separately. The complete planning process entails several MRP runs. The first run is for all parts in all plants except for the press plant; the second run is for the press plant only.

The test objective is that, for a one month production planning period, material requirements planning should be completed within 30 minutes.

The expected outputs from the planning run are: dependent requirements, planned orders, and purchase requisitions.

Infrastructure layer

The objective of these test runs at the IBM infrastructure level is to demonstrate that the architecture setup is totally adequate to fulfill the KPI. In this case we have to achieve a target run time of 30 minutes.

SAP application layer

A process called “80/20” was implemented to perform the MRP scenarios. This process is based on two step execution; the first step generates 80% of the data volume and the second step generates 20% of the data volume, with the tables already filled with 80% of the data volume. The 20% step is more closely related to the client’s business needs.

Table 5-1 summarize the process type (80% or 20%) used during the MRP scenarios.

Table 5-1 Process types used in the MRP scenarios

Runid	Process type
run0004 or run4	80%
run0010 or run10	80%
run0011 or run11	80%
run0013 or run13	20%
run0014 or run14	20%

The 80/20 setup is used in this run. Unlike previous runs (that is, run0012), the flag RESDT is reset for all entries in DBVM and the flags GSAEN (NETCH) and AKKEN (NETPL) are set for all entries in DBVM. Additionally, the partitioning of RESB and PLAF was adjusted further from run0013.

MRP process

Figure 5-1 and Figure 5-2 on page 107 locate the MRP within the customer’s overall manufacturing process IT system.

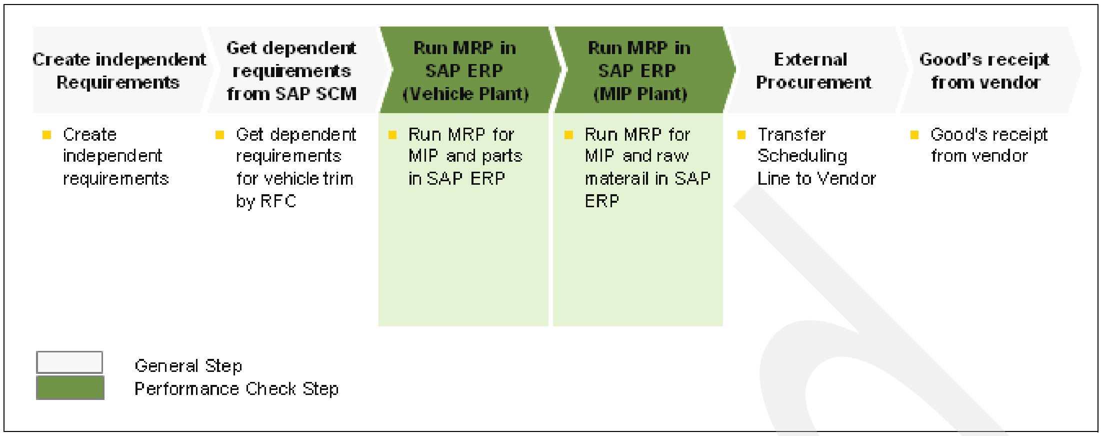

Figure 5-1 MRP-Process I

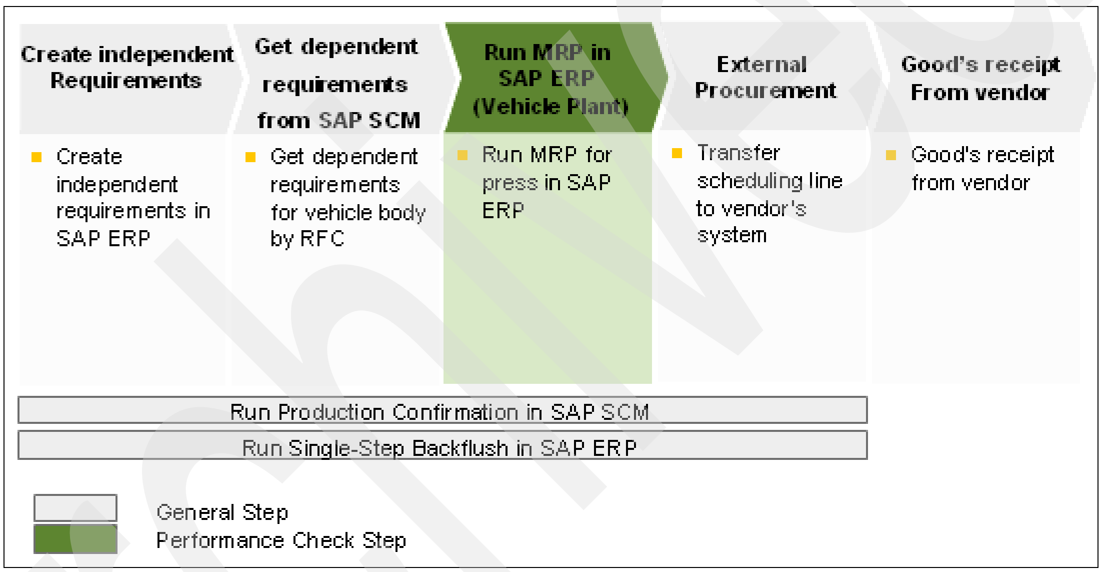

Figure 5-2 MRP Process II

5.1.1 Data required for the MRP execution

The following data was provided by the customer:

- Plants: 9 vehicle plants and 20 component plants
- Material master with MRP type "A2" (which means, get dependent requirement from an external system, in this case the SAP SCM system) or "PD" (planned with MRP)
- Classic bills of material with multilevel structures
- Customer Master
- Info Records
- Dependent requirements coming from the rapid planning matrix in SAP SCM.

Press plant run in SAP ERP

We carried out the run for the press plant in SAP ERP in parallel with the other MRP run.

Configuring the system

For each scenario we performed several test runs to determine the optimal settings to achieve the best performance and to reach the KPI.

In this chapter we describe the test runs that delivered results that best meet the customer's requirements.

5.2 Test execution

We used the same environment as for the RPM test case and used a backup of the results of the Rapid Planning Matrix test.

After every run we reset the systems by shutting down the servers and going back to a backup in the state directly before executing the test.

5.2.1 Initial settings

We set up one job with two steps. The first job step is for the vehicle plants while the second one is for the component plants. The settings for these steps are shown in the next two figures.

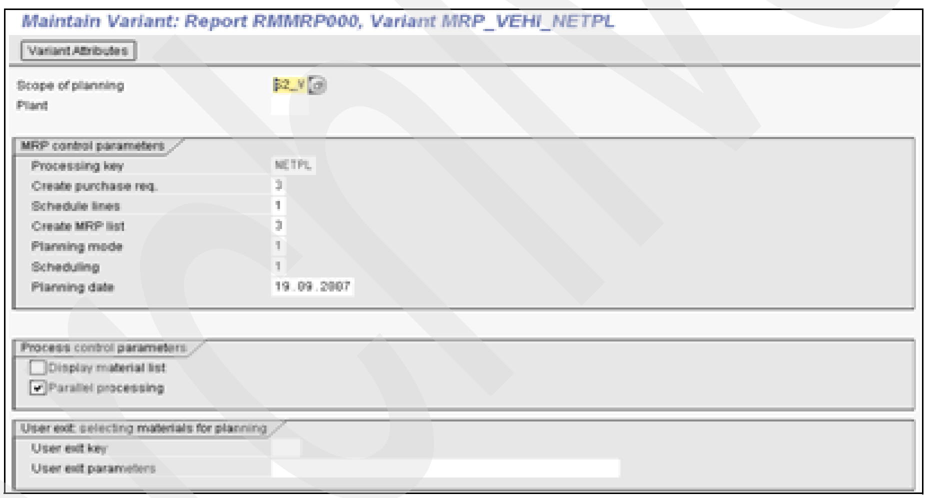

Figure 5-3 Settings for vehicle plants

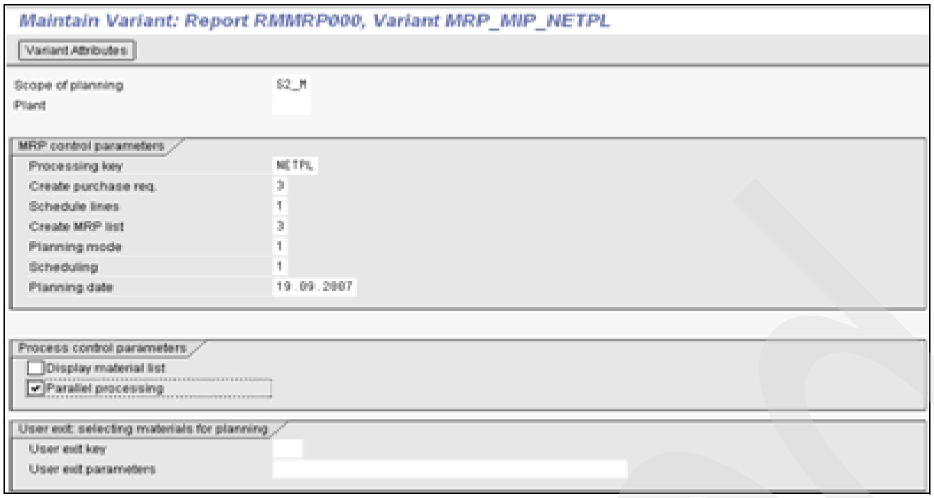

Figure 5-4 Settings for component plants

The *processing key* NETPL stands for the planning procedure that does net change planning in the planning horizon. With that processing key the system only plans materials that have undergone a change relevant to MRP within the planning horizon.

The *planning mode* controls how the system deals with procurement proposals (Planned Orders, Purchase Requisitions, Scheduling Agreement Lines) from the last planning run, which are not yet firmed.

According to the customer a daily change of 20% of the planning data is expected. At the same time, these daily MRP runs do not start on empty tables. We therefore performed test runs in which only 20% of the planning data was generated in order to compare with the defined KPI. We also set the flag RESDT in DBVM to reuse existing procurement proposals from the last planning run, which are not yet firmed, because this is the standard procedure in the company's planning run.

For these tests we first ran the two step job to generate the initial 80% of the data. The second run then simulated the daily run, in which 20% of the planning data was generated; this is the run we used to calculate the KPI.

In transaction OMDX of SAP customizing we set the planning horizon to 14 days.

During the runs we identified the following areas for improvement:

- Partitioning of tables PLAF for planned orders and RESB for reservation of materials is needed as input for the finished product.
- BADI implementation for the parallelization of the MRP runs to influence the packet size is needed.

We therefore applied the settings described in the next two sections.

5.2.2 Partitioning of tables PLAF and RESB

Table 5-2 Initial settings

Parameter	Value
Partitioning of RESB and PLAF	Partitioning of RESB and PLAF by MANDT, MATNR: '500','G' '500','MA' '500','MB' '500','MC' '500','MD' '500','ME' '500','MF' '500','MG' '500','MH' '500','MI' '500','MJ' '500','MK' '500','ML' '500','MM' '500','MN' '500','MO' '500','MP' '500','MQ' '500','MR' '500','MS' '500','MT' '500','MU' '500','N' '500','T' '500','Z'
BADI MD_MRP_RUN_PARALLEL for packet size 1	Active; Settings like screen shot
Planning horizon	13 for all plants
Flag for Reset in table DBVM	NETPL and NETCH flag set for all entries in DBVM, RESDT not set for all entries in DBVM

We worked with three destinations for parallel processing with 15, 25, and 25 sessions, for a total of 65 processes in parallel. During the test we found that we did not gain performance improvement when increasing the number of parallel processes from 65 without applying other optimization measures, for example, to increase the insert performance on the database.

For the partitioning of the tables RESB and PLAF we decided to use the material number object and we created 25 partitions for both tables. Unfortunately the given naming of the material numbers does not allow further partitioning.

After creation of index statistics the expensive reads on table RESB no longer dominated the run time. For the inserts to table RESB we observed a substantial improvement in the average run time for inserting one row. This average insert time originally was 13.4 ms/row and we were able to reduce it to an average insert time of 4.2 ms/row, which means an improvement of nearly 70%. This improvement originates from the partitioning of RESB in DB2.

During earlier runs we observed that the material packets were getting smaller during the planning run for the component plants. This was caused by longer run times for the individual

MRP processes. At the same time, the insert performance into RESB was getting slower as more insert requests would head for the same RESB partition on DB2 and the insert performance was worsened due to the decreasing packet size. This vicious cycle continued until we reduced the packet size to values as small as 2. Such situations need to be avoided, meaning the packet size has to be determined independent of the actual run time of MRP processes.

5.2.3 BADI implementation for influencing packet size

The MRP calculates the standard packet size by doubling or halving the last packet size if the run time of the MRP work processes is below or above a given run-time window. The initial packet size is 1. The standard packet size calculation does not take the load distribution or the efficiency of parallelization into account. Therefore it can be concluded that the partitioning itself is not a stable optimization option if the material master data is named as in the test data.

Note on the material numbers of the test data: The first two digits of the material number describe the car model. Each combination of first two digits occurs only in one vehicle plant. In the current case, the partitioning of RESB and PLAF by the first two digits of the material number requires the processing of relatively large packets to avoid conflicting inserts to the same partition as much as possible. This can only be assured by fixing the packet size with the available BADI MD_MRP_RUN_PARALLEL as shown in Figure 5-5 and Table 5-5 on page 111.

Method IF_EX_MD_MRP_RUN_PARALLEL~SET_PACKAGE_SIZE Active

```
METHOD if_ex_md_mrp_run_parallel~set_package_size.

  IF im_plwrk CS 'HV'.
    CASE im_disst.
      WHEN '000'. ch_tsize = 50.
      WHEN '001'. ch_tsize = 100.
      WHEN '002'. ch_tsize = 10.
      WHEN '003'. ch_tsize = 5.
      WHEN '999'. ch_tsize = 200.
      WHEN OTHERS. ch_tsize = 1.
    ENDCASE.
  ELSE.
    CASE im_disst.
      WHEN '001'.
        IF im_plwrk CS 'HE'.
          ch_tsize = 5.
        ELSEIF im_plwrk = 'HH11'.
          ch_tsize = 15.
        ELSE.
          ch_tsize = 1.
        ENDIF.
      WHEN '999'. ch_tsize = 100.
      WHEN OTHERS. ch_tsize = 1.
    ENDCASE.
  ENDIF.

ENDMETHOD.
```

Figure 5-5 Implementation of BADI MRP_Run_Parallel - programming

Table 5-3 Overview programming settings for BADI MRP_Run_Parallel

Plant	Low-level code	Packet size
HV*	000	50
HV*	001	100
HV*	002	10
HV*	003	5
HV*	999	200
HV*	other	1
HE*	001	5
HH11	001	15
Not HV*, HE*, HH11	001	1
Not HV*	999	100
Not HV*	Not 001 or 999	1

During the test we ran one job per vehicle plant to parallelize as much as possible the MRP runs for these plants and to use the total efficiency of the BADI MRP_Run_Parallel, but we observed that a packet size of one for the vehicle plants is very inefficient. In the vehicle plants we plan dependent requirements coming from the Rapid Planning Matrix in SAP SCM.

5.3 Results overview

Table 5-4 presents an overview of our test results.

Table 5-4 Test results

Parameter	1st step (Vehicle plants)	2nd step (MIP plants)
Scope of planning	S2_V	S2_M
Processing key	NEUPL	NEUPL
Planning mode	1	1
No. of materials	572727	39974
Runtime [s]	942	584
Runtime / material [ms]	1.64	14.61
Total runtime [min]	25.4	
Total runtime [% of KPI]	84.8	

With the performance measures we took, we achieved a total run time of 25.4 minutes. This represents 84.8% of the KPI of 30 minutes that was requested by the customer, meaning that we surpassed the customer's expectation.

5.4 Recommendations

All optimization measures applied during the volume test were technical optimization measures. Further reduction potential exists on both the technical and the business levels. This section describes some of the suggested optimization techniques.

5.4.1 Partitioning

Because of the existing data structure, partitioning on DB2 is limited to distinguishing by the first two digits of the object. Therefore, it was only possible to use the material number (MATNR) for partitioning of tables RESB and PLAF. We could utilize the fact that the first two digits allow a separation by model, meaning the planned orders and dependent requirements belonging to one model are written to the same partition. However, the MRP is planning (within one low-level code) material by material and plant by plant (if a scope of planning is used). In the current setup, the efficiency of the partitioning is thus strongly influenced by the material packet size of the MRP run.

The partitioning can be more efficient if you assign material numbers in a way that the partitioning will also separate the planned orders and dependent requirements that are generated for materials that belong to the same model.

For example, instead of choosing the material numbers as M<letter for car model>_<identifier of material type>_<serial number> the numbers could be chosen to M<letter for car model>_<serial number>_<identifier of material type>. This would allow us to include also the serial number in the partitioning key and thus to use more partitions.

5.4.2 Packet size

The packet size can be varied to find a setting that comes closer to achieving optimum performance. In general, a smaller packet size means that the workload can be distributed more evenly over the available parallel processes, but a smaller packet size also means that the parallelization overhead (RFC communication, functionality that takes place for every packet, and so forth) has a higher share on the total computational effort. This became evident during our tests when the read accesses to table CIF_IMOD ruined the performance of the vehicle plant planning when packet size 1 was used.

In the tests, the packet size was chosen to result in approximately 10 packets per MRP process. Further tests with the packet size might be done to fine tune this.

5.4.3 Number of parallel processes

A very important parameter for the MRP run is the number of processes. This parameter was only rarely varied during the tests because the STAD records showed that the time spent on the DB was in all cases dominating the response time. Furthermore, the accesses on the DB were blocking each other, which made partitioning of RESB and PLAF on DB2 such an important optimization measure.

For the implementation at hand, the customer should check the amount of planning data that is actually generated and read during an MRP run. The number of parallel processes can be increased if the following conditions occur:

- CPU and memory resources are available on application servers as well as on the database server.
- The database performance is not determined by wait times, and so forth. This means the average access times per row are nearly unchanged if the number of parallel processes is increased.

The SAP recommendation is to set the maximum number of MRP processes to the number of available CPUs.

5.4.4 Recommended MRP variants

The duration of the MRP run is considerably influenced by the execution parameters. By optimizing the use of these parameters, you can reduce the run time by up to 90%. The pertinent parameters are discussed in this section.

Table 5-5 specifies the execution parameter variants that we recommend for background MRP runs for cases when you do not want to set a planning horizon to improve the MRP performance.

Table 5-5 Recommended MRP variant 1

Period	Processing key	Planning mode	MRP list	Scheduling	Parallel processing
Daily	NETCH	1	2 or 3	1	☑
start/exceptional	NEUPL	3	2 or 3	1	☑

Table 5-6 specifies variants of execution parameters for background MRP runs for cases when you want to set a planning horizon to improve the MRP performance. We recommend these settings for optimal performance and planning consistency.

Table 5-6 Recommended MRP variant 2

Period	Processing key	Planning mode	MRP list	Scheduling	Parallel processing
daily	NETPL	1	2 or 3	1	☑
weekly	NETCH	1	2 or 3	1	☑
start/exceptional	NEUPL	3	2 or 3	1	☑

Processing key

The "Regenerative planning" (NEUPL) option is performance intensive because all MRP materials are proposed for planning. However, to ensure data consistency when using the scheduling function (with a limited scheduling horizon), you should execute materials planning with NEUPL and Planning mode option 2 using longer time intervals. With the "Net change planning" (NETCH) option, only MRP materials for which the delivery date or the quantity has changed since the last planning run are proposed for planning. With the "Net change planning in the planning horizon" (NETPL) option, only materials for which the delivery date or quantity has changed within the planning horizon are proposed for planning. When and how often each processing key option has to be executed can be determined from the previous tables, depending on how the planning horizon is used.

Planning mode

The *planning mode* has a significant influence on performance, as described in SAP Note 78867. This indicator controls how the order proposals (planned orders, purchase requisitions, and delivery schedules) from the last planning run are handled. Executing materials planning with Planning mode option 3 is very performance intensive and unnecessary during normal planning with NETCH or NETPL. For optimum performance during normal production, use option 1, "Reactivate Planning Dates," because an explosion of the bill of material and rescheduling does not occur. With option 2, "BOM is re-exploded," the bills of material are re-exploded but rescheduling does not occur.

Only use planning mode 3 to update existing planned orders if work center data or customizing tables have been changed. For example, if you change the processing time for purchasing in the plant Customizing settings of material requirements planning, it only affects existing order proposals if you run material requirements planning with planning mode 3. Also note that no planning file entries are written when work center data or Customizing tables are changed.

Create MRP list

This parameter specifies whether to create a list of all planned materials processed by the MRP run. Option 1, "Create MRP List," is performance intensive. We recommend that you use this option only if you require all MRP lists. For optimum performance, use option 3, "No MRP list." You can also use option 2, "Depending on exception messages," to create MRP lists only for those materials that generate exception messages in normal operation.

Scheduling key

This parameter specifies whether only the basic dates are determined for newly created planned orders or whether lead time scheduling is also performed. Option 2, “Lead time scheduling and capacity planning,” is performance intensive. For optimum performance, execute option 1, “Basic dates will be determined.”

Parallel processing

The MRP is run in parallel if this indicator is set and a valid definition of parallel processing exists.

Planning horizon

To speed up your MRP run, you can reduce the plant-level planning horizon to a value that is slightly above the replenishment lead time of most of your materials and use MRP with the execution parameter NETPL. To ensure that no material is left unplanned, create an MRP group with a planning horizon that is long enough for your materials with the longest replenishment lead times.

As noted previously, the MRP has to be run with execution parameter NETCH from time to time if planning horizons are used.

Bulk materials

Consider configuring all materials for which Planned Orders are not necessary as bulk materials.

Good examples of bulk materials are low-price raw materials that are most frequently found in the bills of material and not configured as batch materials. They are usually available in large quantities at the work center (for example, washers or grease) and are not taken into account in the net requirements calculation. They are planned on the basis of consumption. You can improve system performance by configuring materials as bulk material because Planned Orders are not created for bulk materials.

A material that appears in several bills of material may cause severe performance bottlenecks because multiple Planned Orders are generated system wide. However, dependent requirements are still created for bulk materials if they are not excluded from planning.

You can set the bulk material indicator in the material master record. If a material is used as a bulk material only in specific cases, set the indicator in the BOM item.

Note that since bulk materials are planned on the basis of consumption, you have to define a costing procedure.

MRP lot sizes

All material records used in the volume test had lot size “EX” configured. Consider using different lot sizes (for example, weekly lot sizes) to generate fewer planned orders and, accordingly, fewer dependent requirements. Additionally, lot sizing procedures with different granularities can be defined.

The lot-sizing procedure determines how many receipt elements the MRP run creates for a material. Three different periods can be defined for a lot-sizing procedure in transaction OMI4. In the first period the lot size has to be exact, while the lot size can be freely defined for subsequent short- and long-term periods. Only one lot size is used if neither the period of days for lot-for-lot order quantity nor the long-term lot size have been maintained.

Consider this example: A lot-sizing procedure can be defined for which exact lot sizes are used for the first two weeks (14 calendar days). Then a three-day lot size (short-term lot sizing procedure: daily; number of periods for grouping of requirements: 3) is used for eight weeks. After these eight weeks, detailed planning of this kind is not required and so a two-week lot size (lot-sizing procedure: weekly; number of periods for grouping of requirements: 2) is used.

If no long term lot-sizing procedure exists but the end of the short-term period has been specified, no receipt elements are created by the MRP run after the short-term period.

5.5 Analysis of System p infrastructure layer test results

We first present the system activity summary for the five System p server LPARs during the run. This run began at 20:47 and finished at 21:13, lasting 26 minutes.

5.5.1 LPAR activity summary

The following five figures (Figure 5-6 through Figure 5-10) present the system activity summary for the System p server LPARs.

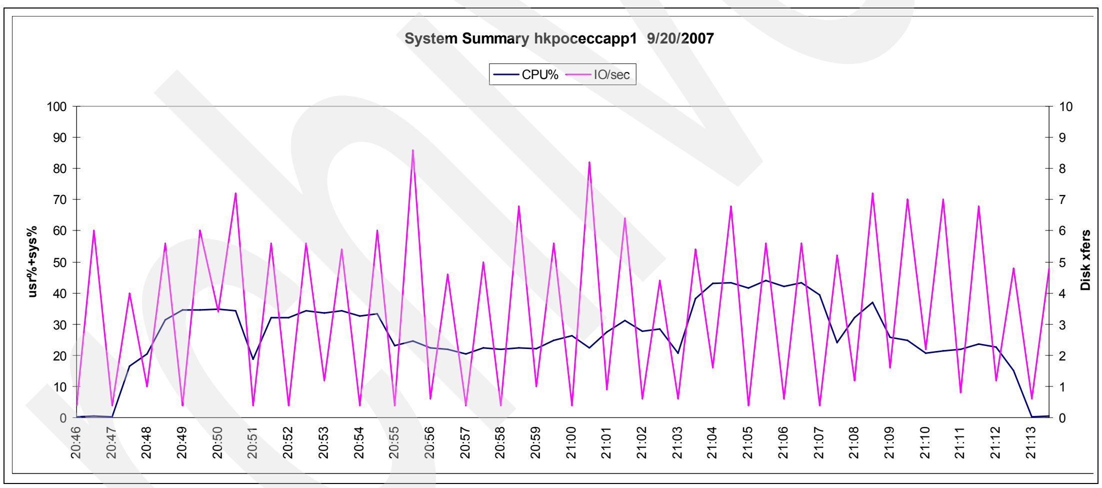

Figure 5-6 Activity summary for hkpoceccapp1

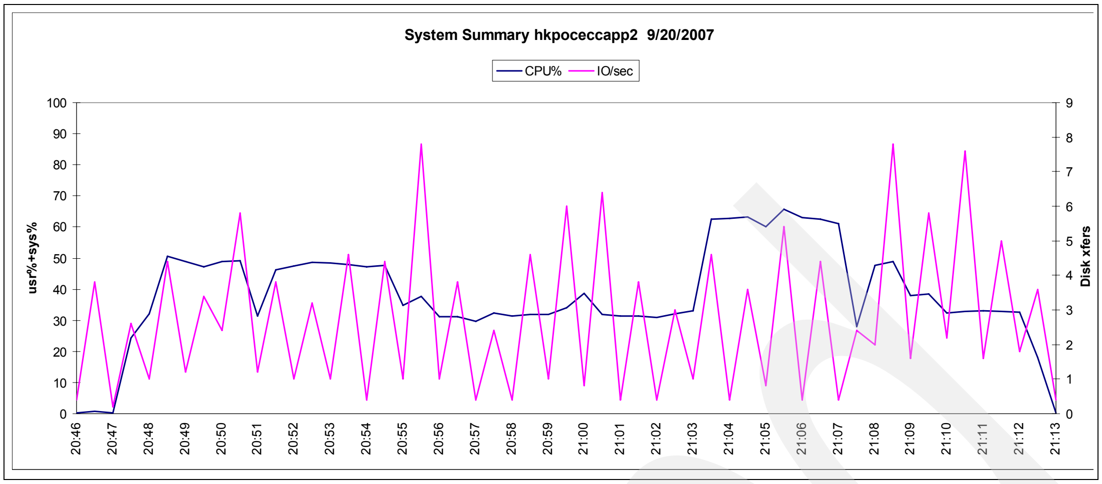

Figure 5-7 Activity summary for hkpoceccapp2

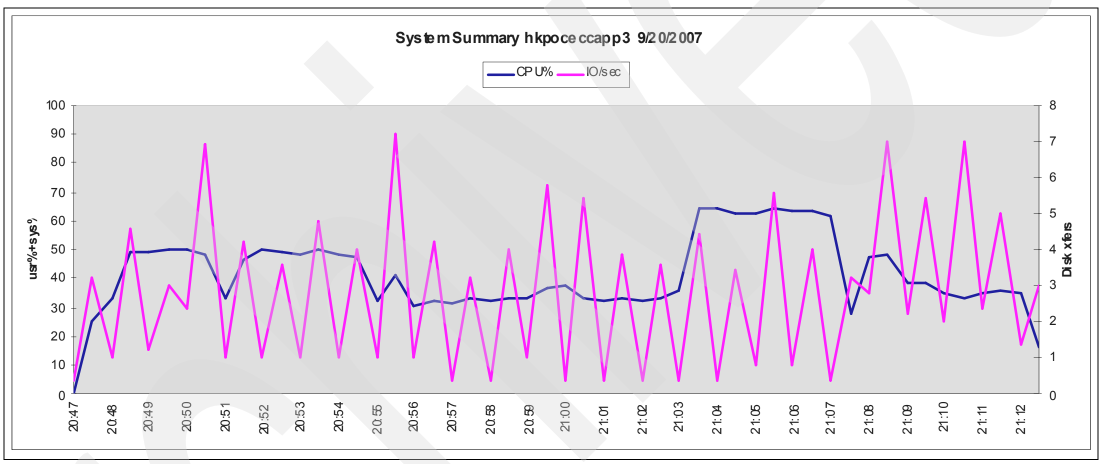

Figure 5-8 Activity summary for hkpoceccapp3

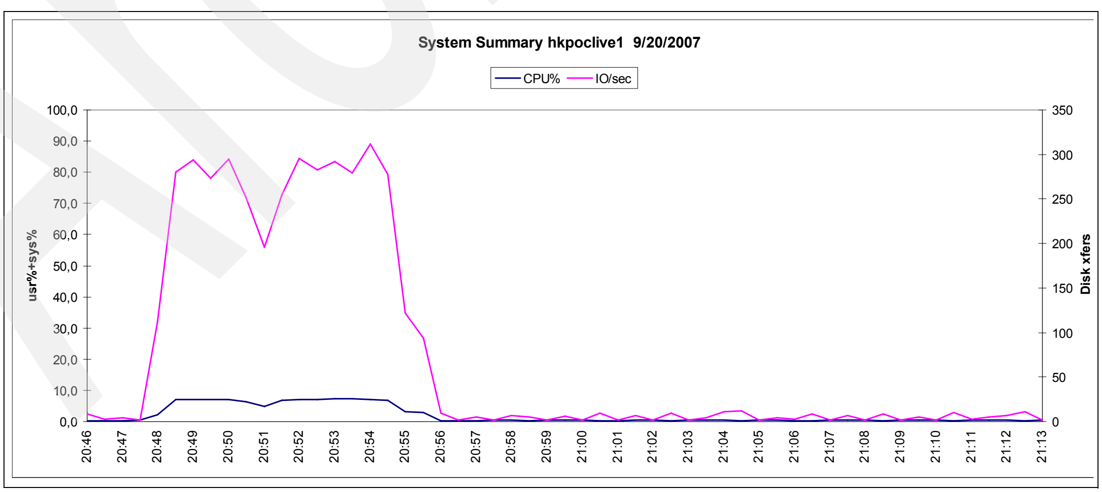

Figure 5-9 Activity summary for hkpoclive1

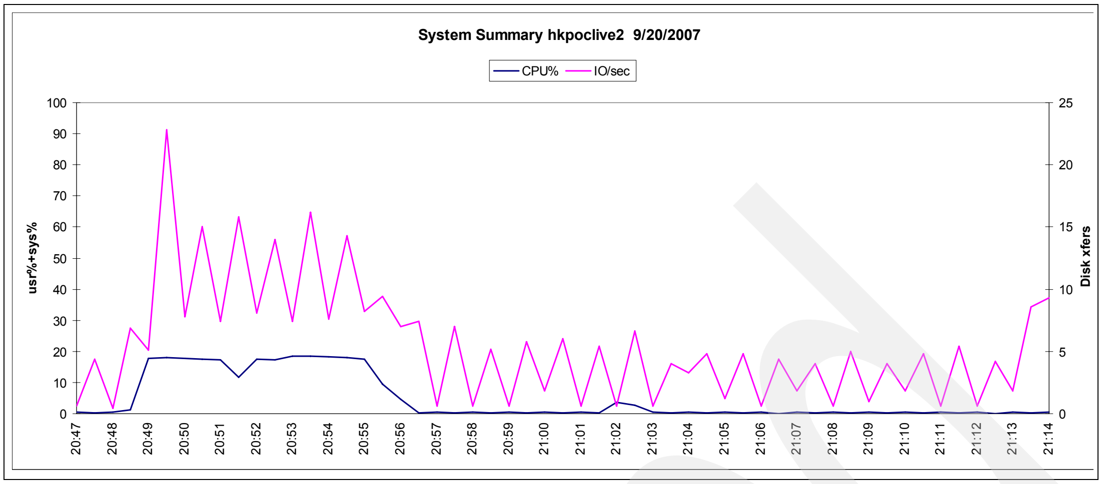

Figure 5-10 Activity summary for hkpoclive2

Regarding the ECC system LPARs (hkpoceccapp1, shown in Figure 5-6 on page 116; app2 or app3 would show the same behavior), we see that the ECC system has significant activity during this run. Hkpoceccapp1 has an average usage of 23.6%, with peaks of 41.5%. Disk I/O is close to 0 in this LPAR.

Hkpoceccapp2 and hkpoceccapp3 show similar behavior during the run, so the activity graph is similar for both. Hkpoceccapp2 has an average CPU consumption of 35.3%, with peaks of 62.7%. I/O per second is not significant. Hkpoceccapp3 as an average CPU consumption of 36.1%, with peaks of 61.4%, and again, I/O per second is not significant.

In other words, on these 3 LPARs, there is a sustained activity but the systems are performing perfectly.

Regarding the liveCache instance, we can see in Figure 5-9 on page 117 that the CPU activity average is 1.9%, with peaks of 6.9%. Disk activity is not relevant.

Finally, regarding the SCM instance, Figure 5-10 on page 118 shows an average CPU consumption of 4.4%, with peaks at 17.7%. But if we focus on the interval from 20:49 until 20:56, the CPU usage average is 17%, so it seems that LPAR is involved only in the first step of the run. Disk activity is not relevant.

No LPAR had relevant CPU wait time, nor did any LPAR use paging area. This leads to the observation that the systems are responding to the run perfectly.

Note: These charts do not consider CPU wait percentage.

5.5.2 Memory details

About 11 GB is used on average for app1, and 8.5 GB for app2 and app3. Figure 5-11 shows memory usage for app1 only.

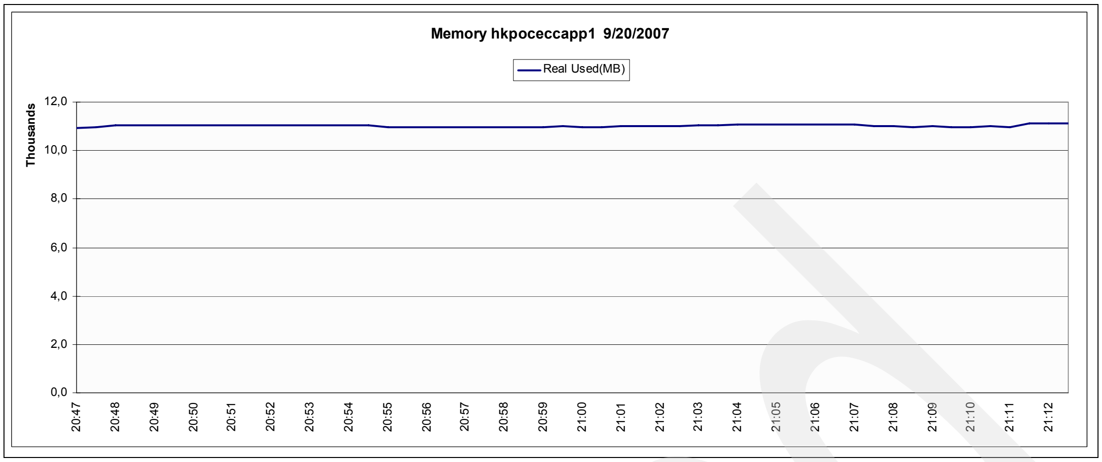

Figure 5-11 Memory hkpoceccapp1

5.5.3 CPU details

The next four figures (Figure 5-12 through Figure 5-15) show the CPU usage per process. The highest usage is by disp+work (SAP processes) and db2sysc (DB2 connect, connection manager between DB2/z and SAP).

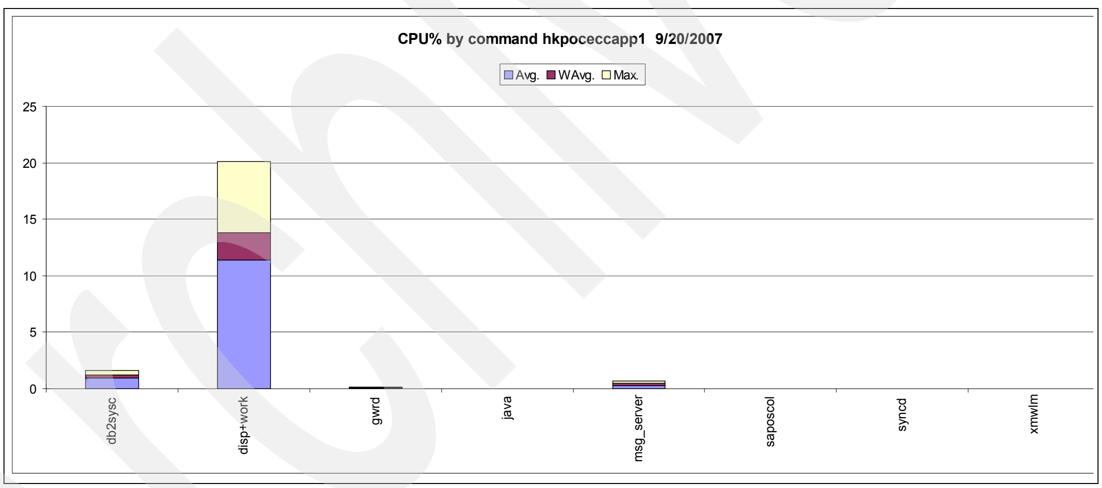

Figure 5-12 hkpoceccapp1 top CPU

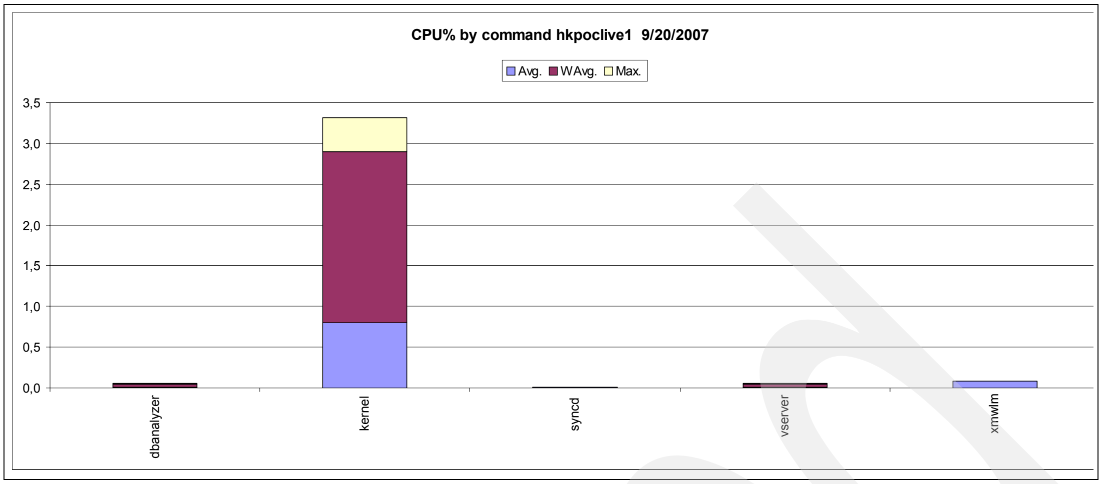

Figure 5-13 hkpoclive1 top CPU

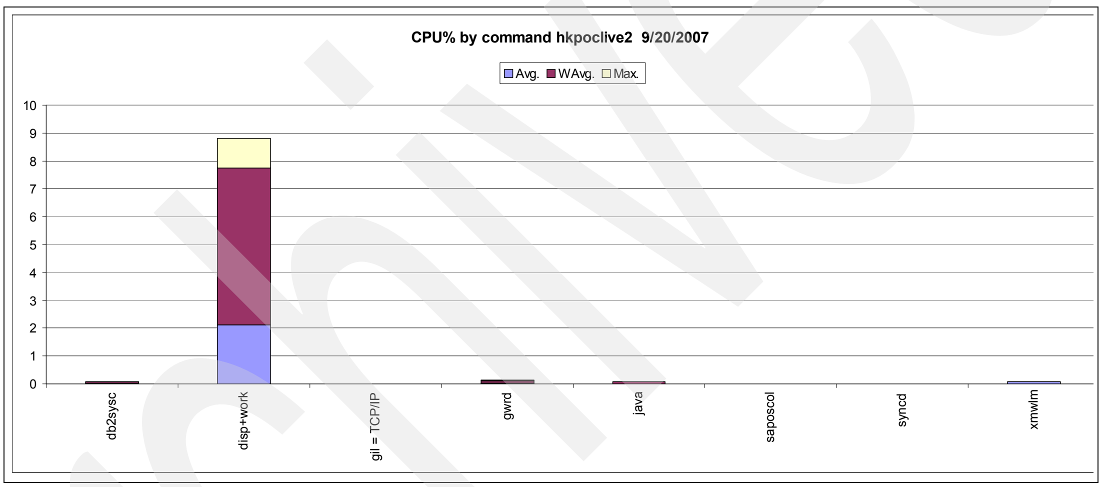

Figure 5-14 hkpoclive2 top CPU

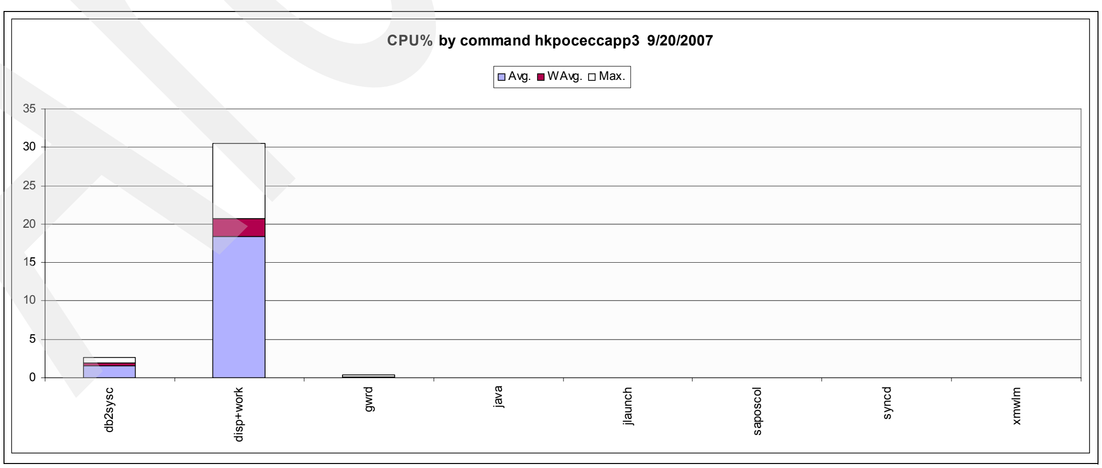

Figure 5-15 hkpoclive3 top CPU

Top memory use

Observing the behaviors of hkpoclive1 and hkpoclive2, we noticed that top memory consuming processes are the core SAP processes, and we do not have any other processes consuming memory. This behavior is similar to that observed for top CPU use.

5.5.4 System p architecture conclusion

Based on run 14, we can say that at the infrastructure level System p servers perfectly responded to customer requirements. Absolutely nothing impacted negatively on the run at this level.

5.6 Analysis of System z infrastructure layer test results

In this section we present the system activity summary for the two System z LPARs during the run. This run began at 20:47 and finished at 21:13, lasting 26 minutes.

Figure 5-16 presents an overview of system activity during run 14.

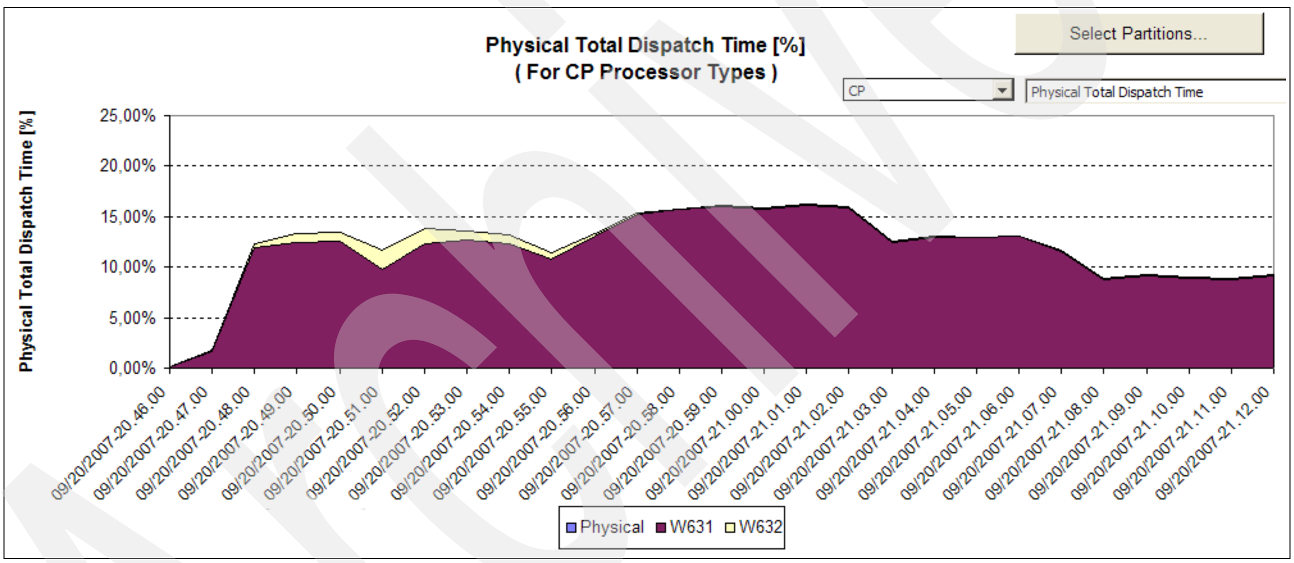

Figure 5-16 LPAR overview MRP run 14

During the test there is a substantial load on MVH1 and during the first period a minor load on MVH2. Remember the LPAR overview shows values related to all available CPs (of which there are 36), so a value of 15% would mean that an LPAR used 5.4 CPs.

5.6.1 Activity on MVH1

Figure 5-17 shows the LPAR summary for MVH1.

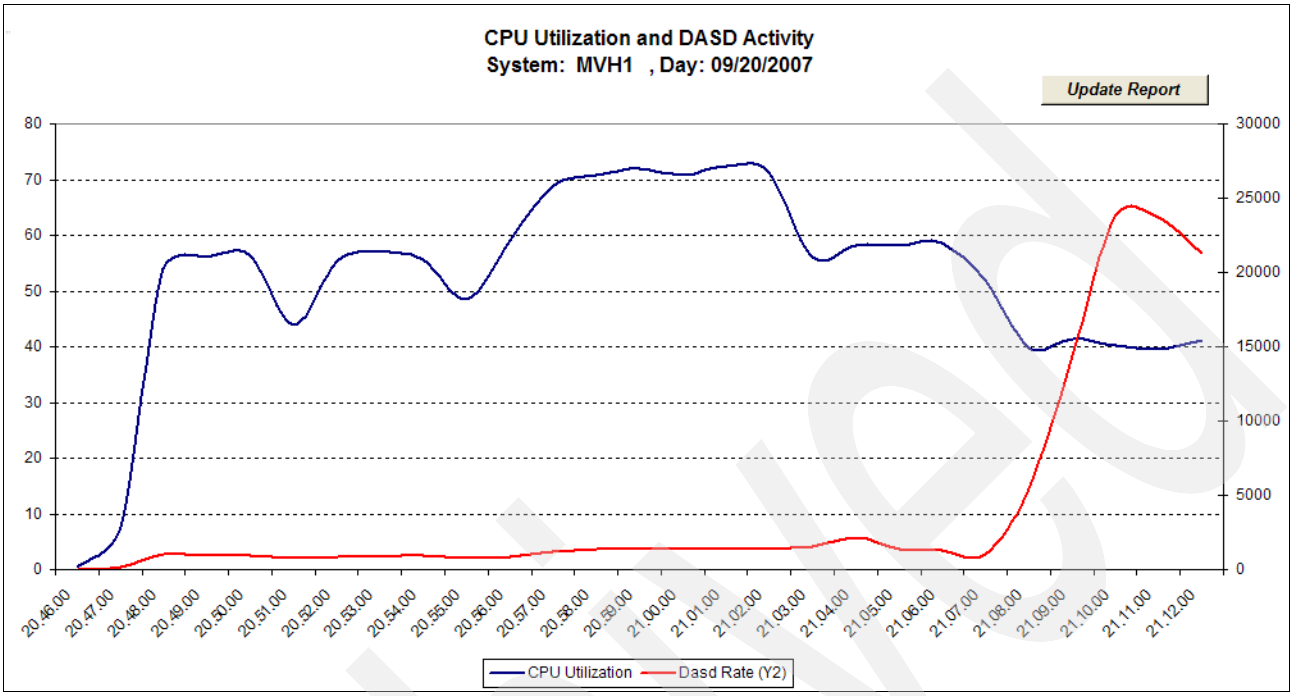

Figure 5-17 LPAR Summary MVH1 MRP run 14

The CPU utilization and the DASD rate (at the end of the run) is high. Also interesting is the different behavior during the run: During the first period we see high CPU usage and low DASD rates; at the end of the run we see lower CPU usage and very high DASD rates.

Each MRP run is divided into two steps:

Step 1: MRP for Vehicle plants called (VEH)

Step 2: MRP for MIP called (MIP)

To gain a better understanding of how DB2 for z/OS helped in achieving the KPI goal for MRP we have to use the results achieved before run 14, focussing on the implementation done at the database level.

In this discussion we present the results of three runs in addition to run 14; run 4, run 10, and run 11.

Table 5-7 on page 123 summarizes the results of these four runs at the DB2 level using DB2 accounting and statistics reports. Later on we identify the contention and how it was relieved.

Table 5-7 Table summary - DB2 Accounting and Statistics reports data

	RUN 4		RUN 10		RUN 11		RUN 14	
	VEH	MIP	VEH	MIP	VEH	MIP	VEH	MIP
Run Duration (minutes)	30	91	25	103	23	16	16	10
DB2 Numbers								
Accounted records	710213	40599	739380	78365	763496	44873	697702	59416
Commits	710254	40603	739415	78367	763537	44873	697735	59419
DB2 Times (in seconds)								
DB2 Elapsed time	0.056641	4.835897	0.012480	2.957493	0.014644	0.608366	0.007819	0.212623
DB2 CPU time	0.004805	0.096222	0.004257	0.072680	0.004339	0.057383	0.004130	0.027902
DB2 Suspend time	0.051069	4.743812	0.007263	2.895991	0.009170	0.544275	0.002713	0.179848
DB2 Not accounted t	0.000767		0.000960	N/C	0.001135	0.006708	0.000976	0.004873
Synchronous I/O	0.001597	0.001257	0.002305	0.000343	0.001843	0.002156	0.001537	0.001116
DB2 Suspended time (in seconds)								
Page Latch	0.034309	4.696032	0.001280	2.821926	0.002595	0.497727	0.000176	0.009394
Other suspension	0.016760	0.047780	0.005983	0.074065	0.006575	0.046548	0.002537	0.170454
SQL Activity (DML) - Total								
SELECT	183	956	158	123	191	23	0	48
INSERT	3 833 732	1 789 602	3 855 199	1 805 489	3 879 082	1 712 340	721 410	356 555
UPDATE	571 963	650 956	601 019	686 570	629 713	650 270	585 406	845 772
DELETE	687 300	61 488	625 642	122 360	717 937	24 049	548 207	55 228
DESCRIBE	2 693	2 091	2 530	1 956	3 064	934	448	231
DESC.TBL	0	0	0	0	0	0	0	0
PREPARE	24 001	22 190	24 523	18 142	31 860	4 828	4 844	2 375
OPEN	17 439 563	3 452 958	17 005 079	3 893 770	17 998 799	3 088 859	17 242 480	4 416 748
FETCH	17 441 300	3 580 040	17 006 842	4 021 237	18 003 933	3 188 431	17 247 114	4 674 519
CLOSE	9 594	4 311	14 483	4 884	16 266	314	1 093	193
SQL Activity (DML) - per second								
SELECT	0.12	0.18	0.09	0.02	0.12	0.01	0.00	0.16
INSERT	2 547.22	330.71	3 198.17	297.71	3 251.35	1 671.42	964.29	739.95
UPDATE	218.75	119.07	289.57	112.72	468.17	715.84	600.89	1 955.59
DELETE	242.32	6.62	281.36	14.89	456.13	17.10	543.19	46.40
DESCRIBE	1.55	0.39	2.10	0.33	1.39	0.61	0.52	0.18
DESC.TBL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
PREPARE	14.20	4.11	19.07	3.01	19.99	3.08	5.41	2.33
OPEN	6 918.02	595.37	8 050.44	600.71	13 400.00	3 113.49	17 700.00	7 837.16
FETCH	6 918.58	618.60	8 052.32	621.06	13 400.00	3 125.17	17 700.00	8 262.07
CLOSE	5.40	0.79	11.88	0.81	7.92	0.27	1.56	0.08

Runs 4, 10, and 11 are full MRP runs and are similar in terms of SQL statements executed (Select, Insert, Update, and Delete) and so are valid runs with which to make comparisons. Run 14 is the 20% part of an 80/20 percent process and is more closely related to the client business needs.

Figure 5-18 on page 124 shows the number of SQL statement executed during the four runs and indicates that runs 4, 10, and 11 are comparable runs.

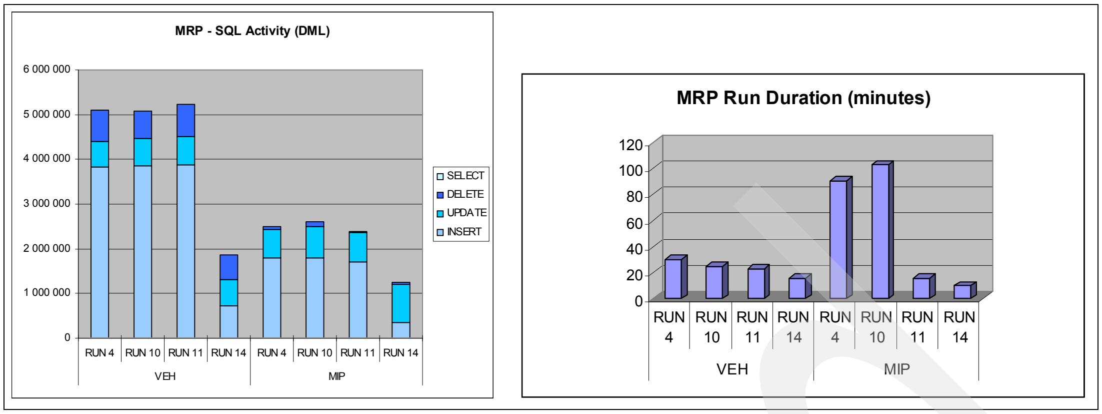

Figure 5-18 SQL Activity during MRP runs

The SQL traces in Figure 5-19 show the typical situation in the second step (planning for MIP plants) during execution of run 4. The biggest share of the total run time is spent on inserts to table RESB (via view MDSB).

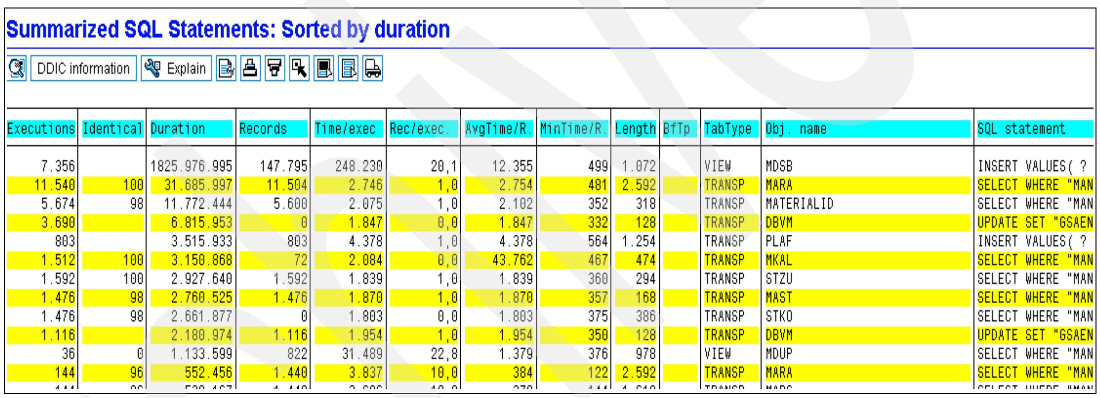

Summarized SQL Statements: Sorted by duration

Executions	Identical	Duration	Records	Time/exec	Rec/exec.	AvgTime/R.	MinTime/R.	Length	BfTp	TabType	Obj. name	SQL statement
7.356		1825.976.995	147.795	248.230	20,1	12.355	499	1.072		VIEW	MDSB	INSERT VALUES(?
11.540	100	31.685.997	11.504	2.746	1,0	2.754	481	2.592		TRANSP	MARA	SELECT WHERE "MAN
5.674	98	11.772.444	5.600	2.075	1,0	2.102	352	318		TRANSP	MATERIALID	SELECT WHERE "MAN
3.690		6.815.953	0	1.847	0,0	1.847	332	128		TRANSP	DBVM	UPDATE SET "GSAEN
803		3.515.933	803	4.378	1,0	4.378	564	1.254		TRANSP	PLAF	INSERT VALUES(?
1.512	100	3.150.868	72	2.084	0,0	43.762	467	474		TRANSP	MKAL	SELECT WHERE "MAN
1.592	100	2.927.640	1.592	1.839	1,0	1.839	360	294		TRANSP	STZU	SELECT WHERE "MAN
1.476	98	2.760.525	1.476	1.870	1,0	1.870	357	168		TRANSP	MAST	SELECT WHERE "MAN
1.476	98	2.661.877	0	1.803	0,0	1.803	375	386		TRANSP	STKO	SELECT WHERE "MAN
1.116		2.180.974	1.116	1.954	1,0	1.954	350	128		TRANSP	DBVM	UPDATE SET "GSAEN
36	0	1.133.599	822	31.489	22,8	1.379	376	978		VIEW	MDUP	SELECT WHERE "MAN
144	96	552.456	1.440	3.837	10,0	384	122	2.592		TRANSP	MARA	SELECT WHERE "MAN

Figure 5-19 SQLtrace planning for MIP plants

To determine why the average execution time of the Insert SQL statement on RESB table is so high, we have to take a look at the observed time inside DB2. The graphs in Figure 5-20 on page 125 show the time spent in DB2 divided by CPU time (processing of SQL statement) and Suspension time (waiting time). The numbers are provided by DB2 Accounting reports. The graphs in Figure 5-20 are separated into VEH and MIP because of the different scaling units. All the times are averages based on more than 700000 accounting records for VEH and 40000 for MIP.

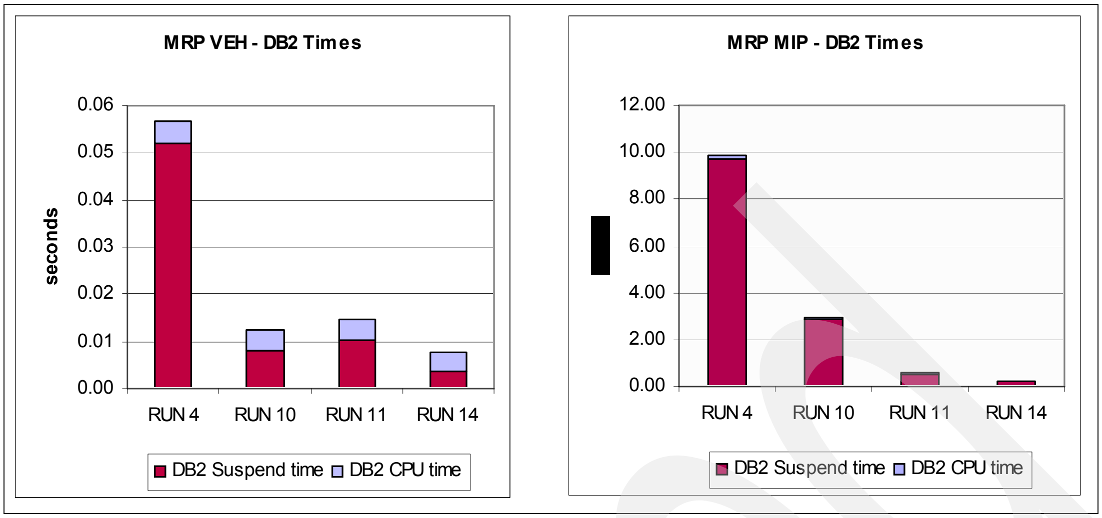

Figure 5-20 MRP VEH/MIP DB2 Times

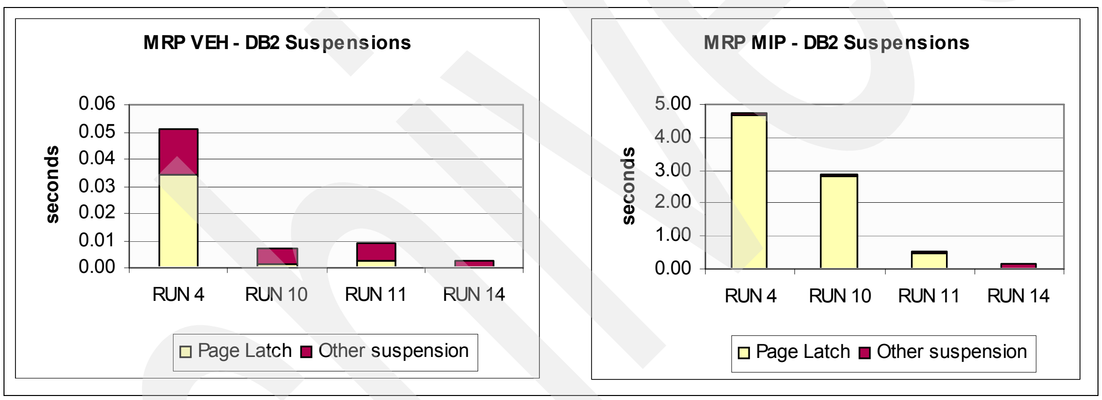

Figure 5-21 MRP VEH/MIP DB2 Suspensions

As shown in Figure 5-21, the major suspension type is *Page Latch*. A Page Latch is a DB2 internal mechanism used to serialize writes into a page; this kind of suspension is observed with massive parallel insert tasks. A Lock/latch trace was collected and showed us that the main Page Latch contentions were on space map pages.

In the following paragraphs we described the tuning activities we executed to relieve the contention from run 4 to run 11. Two tables were mainly concerned: RESB and PLAF. Before starting run 4, RESB and PLAF tables were compressed. The compression ratio was 80%

- Between run 4 and run 10 the RESB and PLAF tables were partitioned using MANDT and MATNR columns to spread the insert activity onto distinct partitions. The partitioning key used was:
  - '500','G'
  - '500','MA'
  - '500','MB'
  - '500','MC'
  - '500','MD'
  - '500','ME'
  - '500','MF'
  - '500','MG'
  - '500','MH'

- '500','MI'
- '500','MJ'
- '500','MK'
- '500','ML'
- '500','MM'
- '500','MN'
- '500','MO'
- '500','MP'
- '500','MQ'
- '500','MR'
- '500','MS'
- '500','MT'
- '500','MU'
- '500','N'
- '500','T'
- '500','Z'

- The RESB and PLAF table spaces were isolated in a distinct buffer pool (BP10) and the indexes were isolated on two buffer pools (BP11 and BP12).
- The previous graphs show that this partitioning relieved a part of the contention, but the suspension was still too high during the MIP part of the run.
- Before run 11 RESB and PLAF tables were recreated with the same partitioning criteria, buffer pool isolation, and some special options at the Tablespace level:
  - MEMBER CLUSTER
  - FREEPAGE = 0
  - PCTFREE = 0
- All the partitions were loaded with data to build the compression dictionaries and the statistics were collected.
- As you can see from the previous graphs, these changes reduced the Page Latch contention by 940 percent, from 4.7 seconds on run 4 to 0.5 seconds on run 11.
- More improvements could be achieved by using a different material numbering scheme. The first two digits of the material number identify the car model. Each combination of first two digits occurs only in one vehicle plant. We could spread the data for one plant onto more partitions using a different material numbering plan.

5.6.2 Activity on MVH2

Figure 5-22 shows the LPAR summary for MVH2.

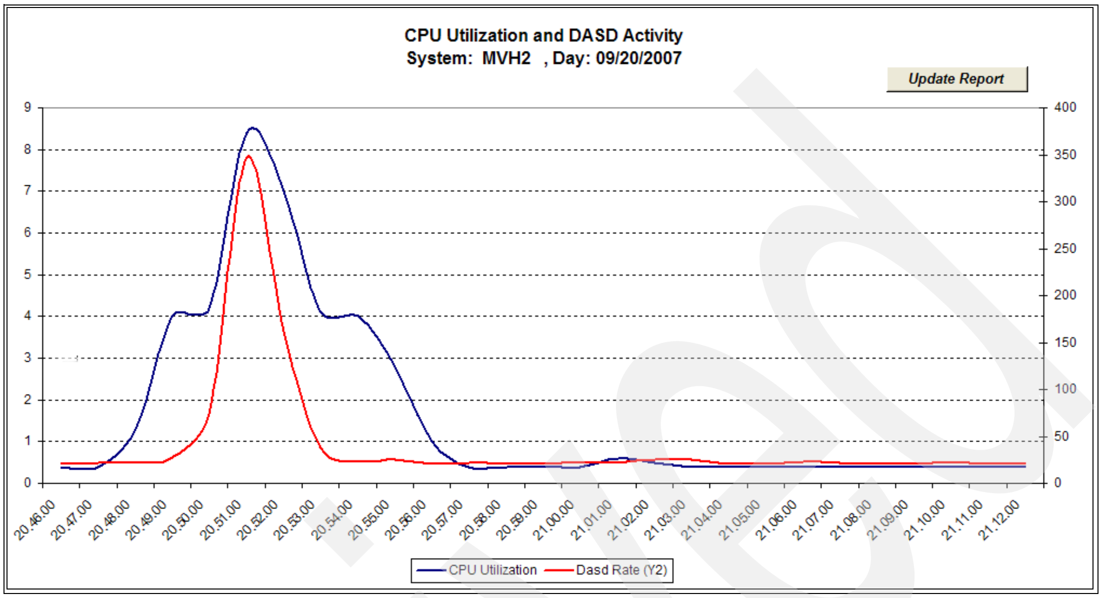

Figure 5-22 LPAR Summary MVH2 MRP run 14

The CPU utilization and the DASD rate for LPAR MVH2 are both low. There is no need for further investigation or preparation of more detailed reports for this LPAR.

5.7 Conclusions and further recommendations

This section presents some final thoughts about the MRP tests. The recommendations that follow are provided *in addition* to those presented in 5.4, “Recommendations” on page 112.

5.7.1 Valid test runs for KPI

Run 0013 and run 0014 are considered valid test runs for KPI MRP, with each test achieving a run time of 26 minutes (Table 5-8).

Table 5-8 Run durations

Run_ID	Date	Job / Step	Start	End	Duration in min.
RUN0013	19.09.2007	MRP for Vehicle and MIP	19:25	19:51	**26**
RUN0014	20.09.2007	MRP for Vehicle and MIP	20:47	21:13	**26**

The joint effort from SAP and Z-experts produced an impressive improvement in response time for achieving the KPI. The DB2 partitioning functionality together with SAP customization made it possible to parallelize the process in such a way that the KPI was achieved and the underlying hardware infrastructure was used effectively.

5.7.2 Further recommendations

See 5.4, “Recommendations” on page 112 for information about performance tuning. This section provides suggestions for some additional performance tuning techniques.

Rebuild of low-level codes

The test data contains some inconsistencies in the low-level codes between tables MARA and DBVM and the BOM. These inconsistencies were corrected using the Z-Reports provided in SAP note 140573 and report RMMDVM10 (transaction MDRE) as described in SAP note 186216. Report RMMDVM10 is valid also if table DBVM is used for the planning file.

The Z-Reports of note 140573 do not decrease the low-level codes in table MARA. The procedure is as follows:

1. A list of BOMs is created.
2. For each BOM in the list, check whether the materials in the BOM have a lower low-level code than the header material. The low-level codes of the materials in the BOM are increased if an inconsistency is found.

Report RMMDVM10 assures that the entries in table DBVM are consistent with the material master data. This report has to be run plant by plant.

This procedure was applied from the third through the ninth run. Before the ninth run a new flashcopy was made and the changes to the low-level codes were made permanent.

The low-level codes in the system range from 000 to 050. This does not pose a problem as long as the order of low-level codes reflects the structure of the BOMs. To rebuild the low-level codes the following procedure was proposed by the responsible developer of SAP note 140573:

1. Execute report ZZDISTINIT (refer to Appendix B, “Profiles and configuration files” on page 277).
2. Execute report ZZKLAHDISGAP from SAP note 140573.
3. Execute reports ZZDIS502, ZZDIS503, ZZDIS504, and ZZDISK01 of SAP notes 140573 and 84153.
4. Ensure that the report ZZDISK01 is executed for class type 200 as well. The default value would be 300; this must be changed to 200 once the execution for 300 is over.

Note: Ensure that all these reports are executed and tested in the test/development system first. No other transactions/reports like MRP or costing which use the low-level codes should be executed until the above described procedure is finished.

We tested the procedure described before run 15. Without retesting it is not possible to judge whether the low-level code rebuild is responsible for the changes in the planning results for the MIP plants.

The reset flag in table DBVM

In the tests the flag DBVM-RESDT was set for more than 75% of all entries in DBVM. The purpose of this flag is to trigger a regenerative planning (planning mode 3 - Delete and recreate planning data) for the respective materials. The flag DBVM-RESDT overrides the planning mode setting in the variant used for the MRP run. This flag is set when changes to master data occurred that make a regenerative planning necessary.

The client should carefully test when this flag is set and whether it is planned to execute the responsible step during the daily business. It is recommended to avoid changes to master data during daily business operation.

Note: The flag DBVM-RESDT should never be reset manually or by a Z-Report in a production environment

Further optimization of the MRP run

All optimization measures applied during the volume test were technical optimization measures. Further reduction potential exists on the technical level as well as on the business level. Business-related optimization options are focused on reducing the amount of planning that is actually needed. For details, see 5.4, "Recommendations" on page 112.

Useful SAP notes

The SAP notes listed in Table 5-9 contain general information about the MRP run and its performance.

Table 5-9 SAP notes

Note number	Description
23278	MD01, MDBT: Performance during MRP run
78867	MD01: Documentation on the planning mode
204517	MRP: Performance problems - typical causes
440016	No longer any data in MDTB
483672	Performance and bulk material settings
550441	FAQ: Exception messages in material requirements planning
550568	FAQ: MRP run (MD01, MD02, MD03, MDBT...)
553746	FAQ: Planning file entries in material requirements planning
568593	FAQ: Parallel planning run (MD01, MDBT and so on)

6

Discrete industries production confirmation: Backflush

Backflush is the process of consuming components from inventory as production is reported. This chapter describes our experiences with backflush testing.

It includes the following topics:

6.1 Backflush test overview

The backflush scenario takes place in the ERP system. There are two alternatives: Single-step or two-step backflush. The test will give an indication of optimal settings for those alternatives that are most important for performance.

The purpose of the test is to execute the backflush jobs (PPCGO*) in a maximum of 30 minutes. From an infrastructure perspective, the test also is intended to demonstrate that the architecture set up is adequate to fulfill the KPI.

Description

The customer wants to execute the backflush periodically every hour. That means that the run time should remain under 30 minutes for the one-step backflush or for the first step of the two-step backflush.

Parallelization

The transaction PPCGO (calls program ppc_go_conf in dialogue mode) cannot be executed in parallel. There is a lock in the transaction that blocks a second run while another instance of the program is still running.

The program ppc_go_conf (one-step backflush) itself is able to be run in parallel batch processes.

6.1.1 Scope of the tests

This section locates the backflush tests within the customer's overall process.

Alternative one: One-step backflush

For the one-step backflush the jobs run in parallel. The material document is on the cost collector level, which means on the material variant level. The expected material documents number approximately ~600*3.

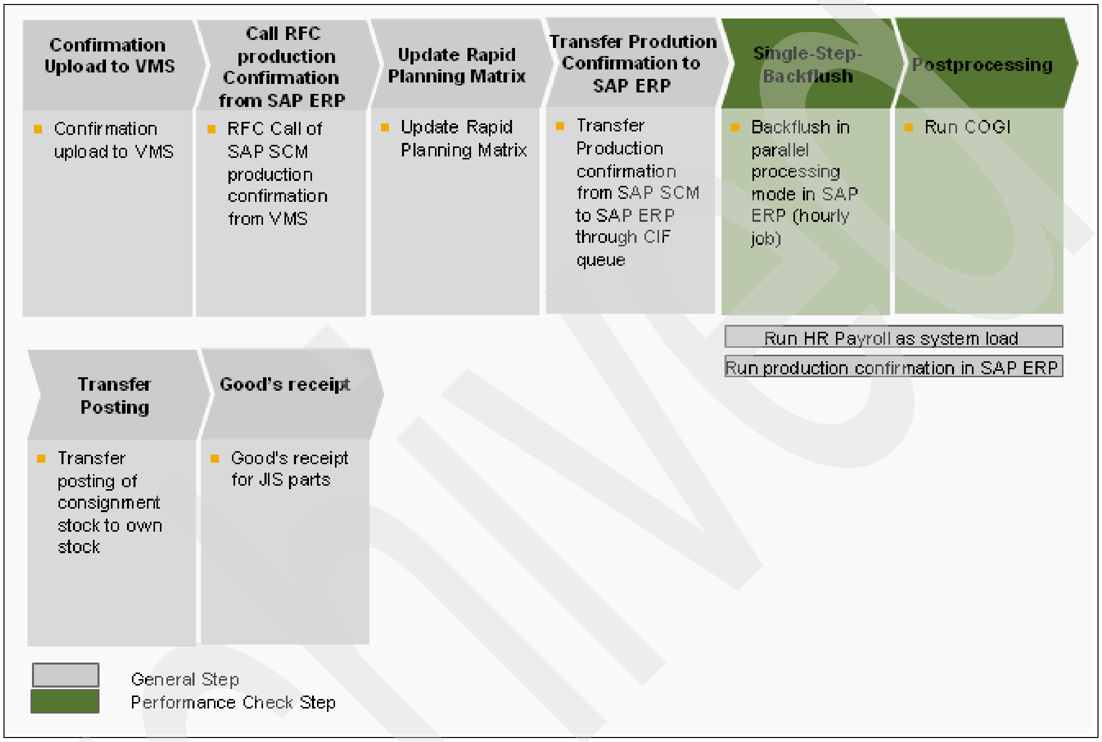

Figure 6-1 Single-step backflush

Alternative two: Two-step backflush

In the two-step backflush scenario the job for the first step cannot be parallelized. The accumulation in the material document is on the material level. The expected line items in the material documents number approximately 210000 (21 models *10.000 materials). The first step must be processed 22 times with different users. After this, the job for the second step can run. The program will run without parallelization. The number of created material documents should be a percentage of 113400 materials (1800*21*3).

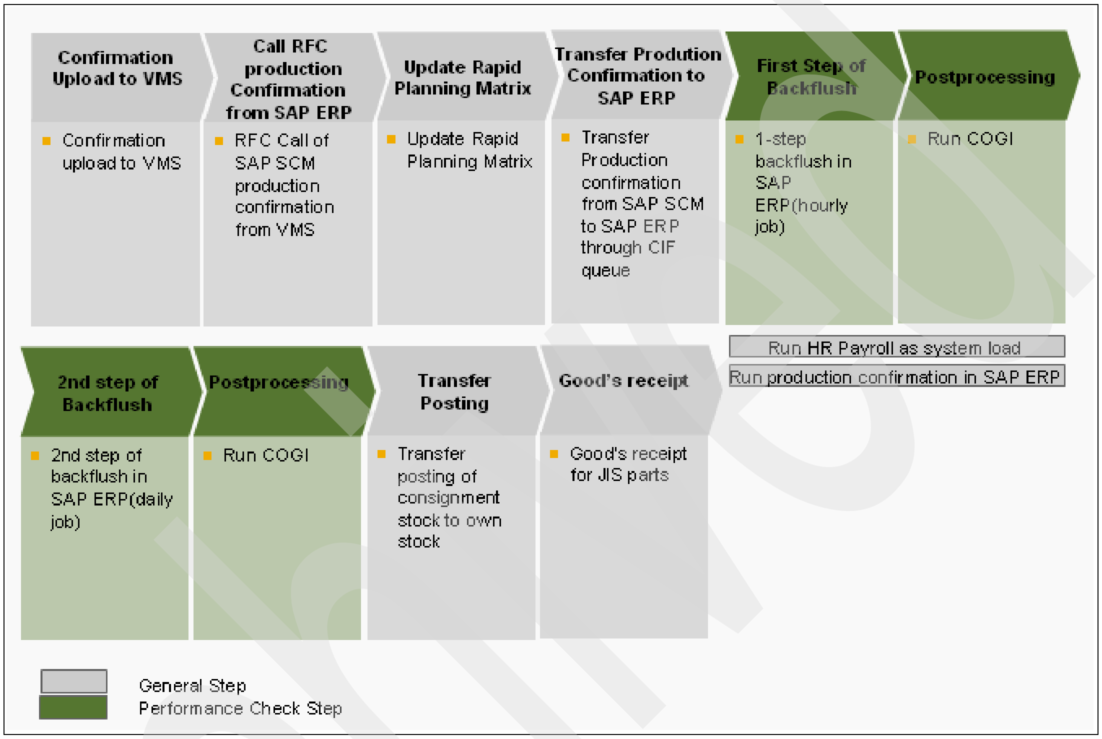

Figure 6-2 Two-step backflush

Scope

The scenario covers only the steps in the SAP ERP system. The time consumed during data transfer via the Core Interface and the backflush posting in the SAP SCM system as a prerequisite step for the MRP run are *not* in the scope of the performance test. The single- and the two-step backflush are alternatives for the test.

6.1.2 Data required for the backflush execution

The following data required for the backflush execution was provided by the customer, or was derived by the test team where noted:

- Customizing for the Discrete Industries-Backflush-Profile without movement types.
- Customizing for non-MRP-relevant storage locations.
- Creation of material master for parts in non-MRP-relevant storage locations.
- Customizing for negative stock.
- OSS note for two-step backflush (only CIF-transfer of MRP relevant stock).
- Material master for FSC without DI-Profile.
- Cost collector with cost estimation.

- To simulate the synchronous backflush trigger from an MES-system we created an Excel® file with confirmations for 1 hour (7500) for all three levels. These confirmations should represent approximately 600*3 different material variants with different checkpoints, considering the production progress (Work in process).
- 22 Excel files each time, with confirmation for 1 hour (7500) for all three levels with different reporting points and different users, considering the production progress (Work in process).
- Add-on program for uploading the Excel file for generating the confirmations in SAP SCM.
- Planned Orders on the three levels (from previous test) and backup on this stage.
- SAP SCM confirmations, Core Interface transfer to the SAP ERP.

Table 6-1 gives an indication of the magnitude of the data handled in this backflush process.

Table 6-1 Backflush data

Item	Number of data fields
Number of vehicles produced per day	7500
Number of backflush reporting point triggers per 1 hour	7500
Total number of backflush reporting points	18
Number of cost collectors per 1 hour	about 600
Material Master	about 2000000

6.1.3 Approach and prerequisites for the test execution

The transaction PPCGO is run in a background job. The KPI relevant time we calculate is the total run time of the job. A total of 7500 Planned Orders for approximately 600*3 different material variants with different checkpoints are backflushed. During each test run there is a workload of confirmations coming from APO.

Deliverables

The following test runs were completed:

- One test run is used to compare the performance of a one-step backflush to a two-step backflush.
- Test series to determine the optimal degree of parallelism
  Measure different degrees of parallelism by testing with 10, 20, 30, 40, 50, 60, 70, 80 parallel processes used for the single-step backflush.
- Test series, using the optimal degree of parallelism determined in the previous test, to determine material block options. The parameters tested are:
  - Material block: Without late locking
  - Material block: Late locking, 10 second wait
  - Material block: Late locking, 30 second wait
  - Material block: Late locking, 60 second wait

Test reset mechanism

The test is reset by shutting down the servers and reverting to a backup in the state of having the confirmations in the SAP ERP and not posted with PPCGO*.

Exit status of the environment after the test

After the test, the posting of all parts is done in the SAP ERP system. The application log (transaction SLG1) holds information on the backflush execution time. The inbound queues of the SAP ERP system are empty (except for usual idle system activity) and the confirmations are carried out (check in PPCSHOW).

Run time goal

The maximum run time requested by the customer is 30 minutes.

We carried out a number runs for the two alternative scenarios. This document describes in detail two of the runs we performed for the one-step backflush scenario.

6.2 Test execution

From the various test runs we carried out we describe here the best two runs of the one-step backflush. In the results area we include also the evaluation of the two-step backflush.

One-step backflush: First run

First we ran the one-step backflush without parallelization. We carried out the one-step backflush with the settings shown in Figure 6-3.

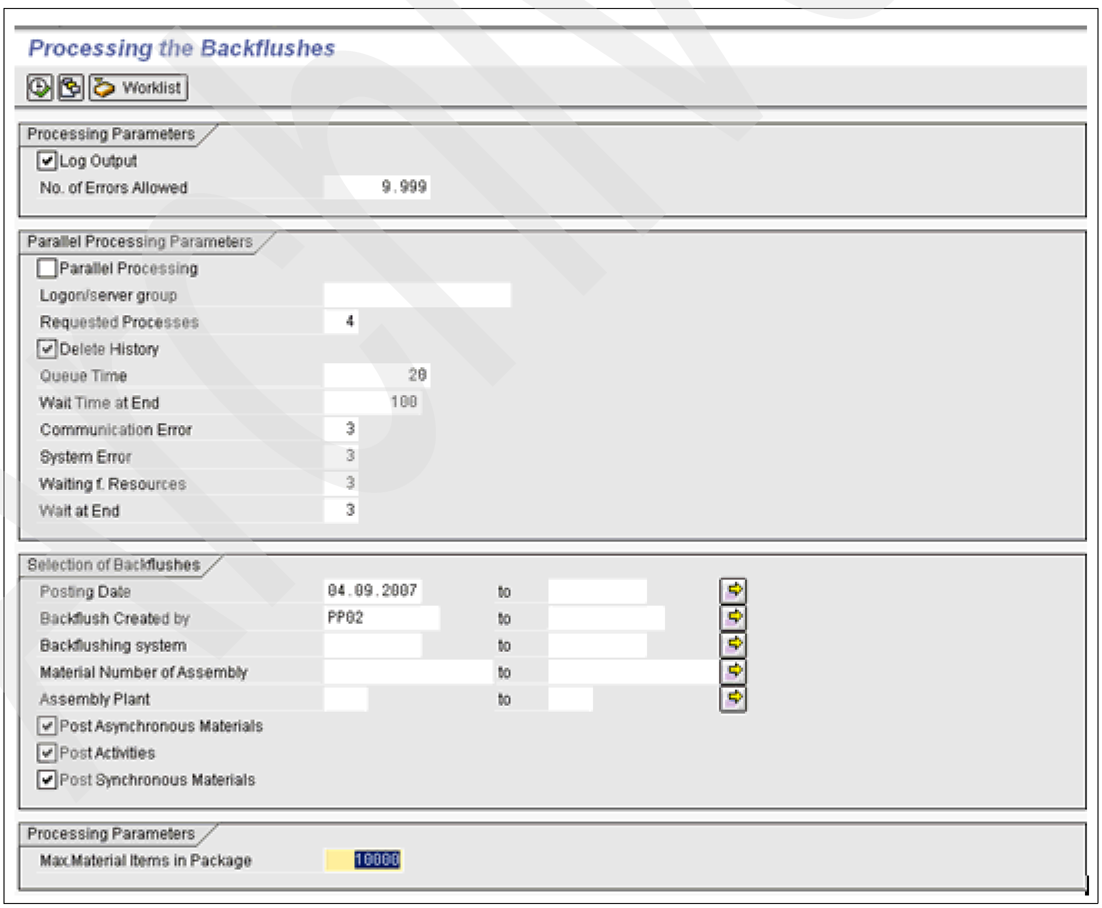

Figure 6-3 Backflush settings

Figure 6-4 displays a sample STAD record of one of the update processes. It indicates that most of the time was spent on the database. The average times per row are very good and indicate little tuning potential. Results of the first test are shown in Table 6-2.

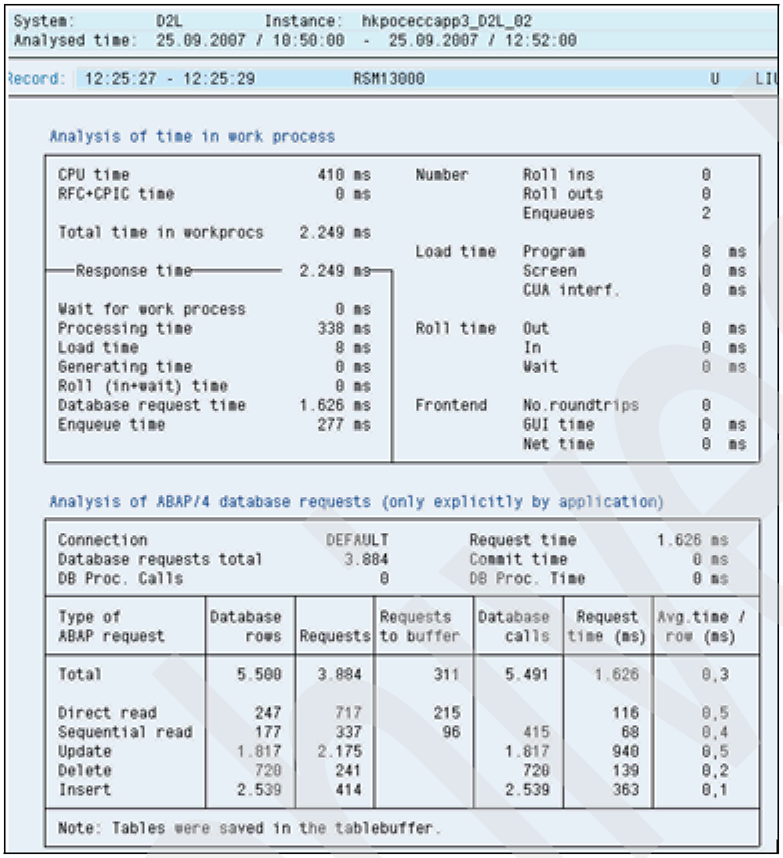

Figure 6-4 Sample STAD record

Table 6-2 Result of the one-step backflush:

Materials successfully posted	Incorrect material items	Successfully posted service items	Service Items without errors	Run time [min]
90563	0	11858	0	121.5

At this point very little tuning potential—except parallelization—is evident. Parallelization, however, does not work because the individual parallel processes would block each other because common parts lead to conflicts during the posting of goods movements.

One-step backflush: Second run

Obviously no tuning is possible for the one-step backflush, but the customer will not be able to work with such a long runtime. Therefore, we decided to modify the backflush code to enable a parallelization of the processes. We made this modification only to show that the long backflush runtime is not a show stopper for the customer.

The procedure we used to make the modification follows.

We copied the report PPC_CONF_GO (calling it ZPPC_CONF_GO) and the function group PPC1PR into the customer name space and modify the coding in the following way:

- The form routine that spawns the parallel processes now checks whether an LUW has common parts with currently running LUWs before it starts a new process with this LUW.
- This form routine waits until at least one job has come back before it tries to start additional LUWs.
- To avoid problems in the communication with the enqueue process we used a server group in this test that utilizes only the central instance and we defined a second ENQ work process on the central instance (SAP note 1227773). We did this because no communication with the enqueue process is necessary if a process is run on the central instance.

Such a modification can only improve the performance if it is actually possible to find a significant number of Logical Units of Work (LUWs) with component lists that have no common parts so that no locking problems can occur. The parallel processing of the backflush in one step will, in any case, not work if there are common parts in all or most of the backflushes.

We started the job with one hour's backflush volume and 80 parallel processes. In parallel we started the payroll (10 jobs in parallel) and the load runner.

We created the backflush job using the report ZPPC_CONF_GO with the settings shown in Figure 6-5.

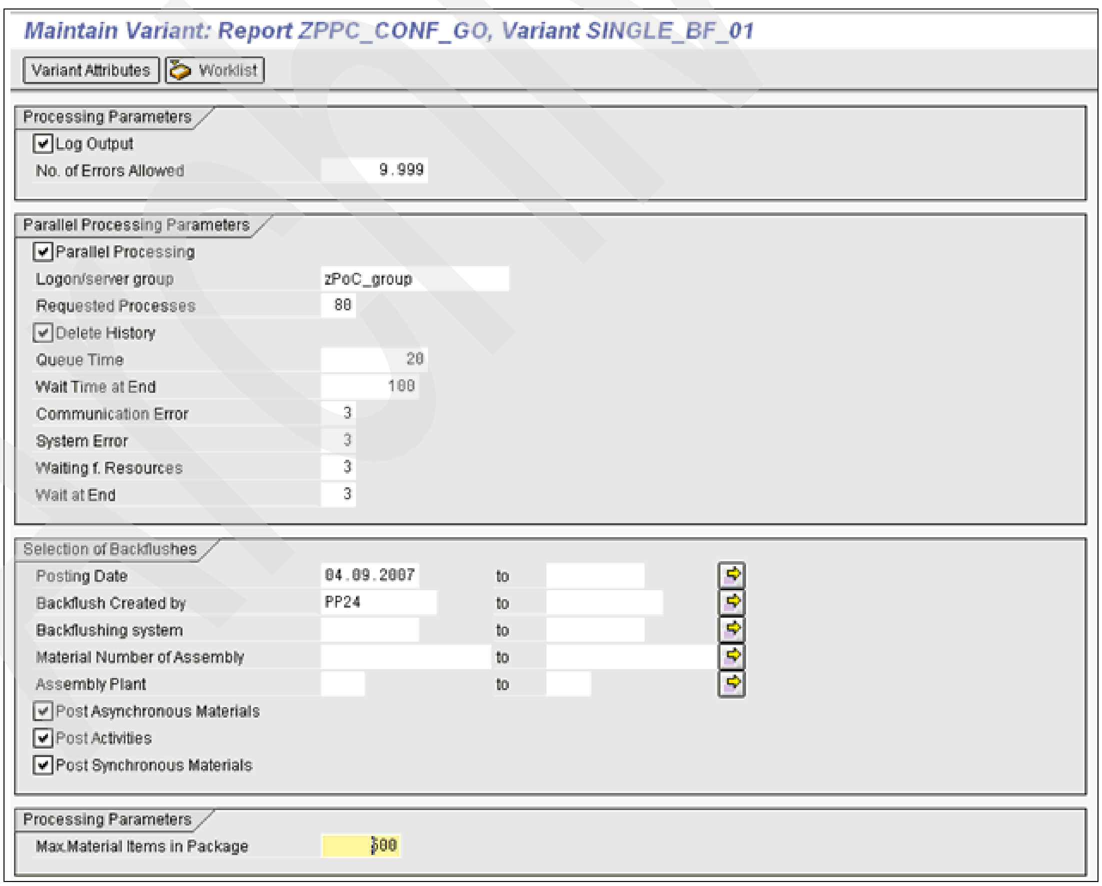

Figure 6-5 Variant for BF Job

Setup for the payroll process

We started the program RPCSC000 in the ABAP™ workbench in parallel with the backflush run, with the settings shown in Figure 6-6.

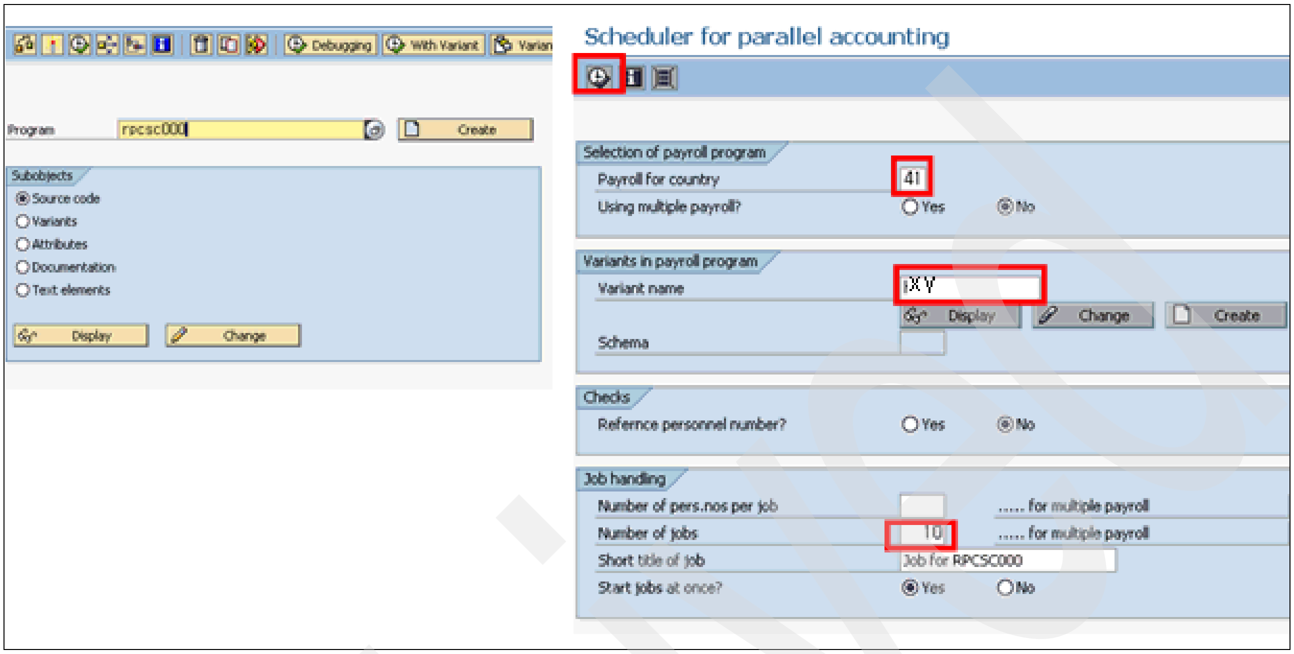

Figure 6-6 Setup for the payroll process

We observed that 59 enqueue requests were waiting in the queue. This was a typical situation observed during the run, and indicates that problems with setting and releasing enqueues existed either because of the enqueue performance or because of the number of requests.

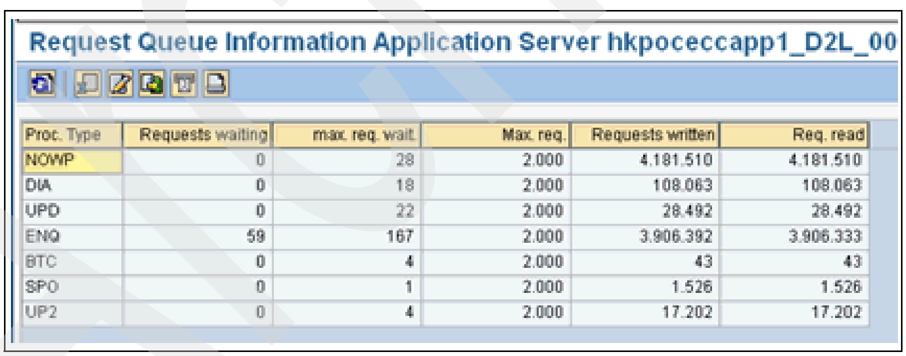

Request Queue Information Application Server hkpoceccapp1_D2L_00

Proc. Type	Requests waiting	max. req. wait	Max. req.	Requests written	Req. read
NOWP	0	28	2.000	4.181.510	4.181.510
DIA	0	18	2.000	108.063	108.063
UPD	0	22	2.000	28.492	28.492
ENQ	59	167	2.000	3.906.392	3.906.333
BTC	0	4	2.000	43	43
SPO	0	1	2.000	1.526	1.526
UP2	0	4	2.000	17.202	17.202

Figure 6-7 Enqueue requests

We regularly observed the situation shown in Figure 6-8 on page 140, namely, a large number of processes performing enqueue operations.

Global Work Process Overview

CPU | Debugging | Long <-> short names | Select process | Settings

Sort: Server

Name	No.	Type	PID	Status	Reason	Sem	Start	Error	CPU	Time	User	Report	Action	Table
hkpoceccapp1_D2L_00	0	DIA	1257564	running		26	Yes			9	SAPSYS	SAPLSENA		
hkpoceccapp1_D2L_00	5	DIA	831550	running		26	Yes			2	PP24	SAPLSENT		
hkpoceccapp1_D2L_00	30	UPD	1069260	running			Yes			9	PP24	SAPLGLIU	Update	COFIT
hkpoceccapp1_D2L_00	31	UPD	987296	running		26	Yes			7	PP24	SAPLSENA		
hkpoceccapp1_D2L_00	32	UPD	385442	running			Yes			7	PP24	SAPLGLIU	Update	COFIT
hkpoceccapp1_D2L_00	33	UPD	405986	running		26	Yes			7	PP24	SAPLSENA		
hkpoceccapp1_D2L_00	35	ENQ	1011884	running		26	Yes				SAPSYS			
hkpoceccapp1_D2L_00	36	ENQ	1015982	running		*	Yes			2	SAPSYS			
hkpoceccapp1_D2L_00	37	BTC	389554	running		26	Yes			1144	LIU	SAPLSENA		
hkpoceccapp1_D2L_00	38	BTC	1024180	running		26	Yes			1144	LIU	SAPLSENA		
hkpoceccapp1_D2L_00	39	BTC	1028278	running		26	Yes			1143	LIU	SAPLSENA		
hkpoceccapp2_D2L_01	0	UPD	377016	running			Yes			9	PP24		Update	COFIT
hkpoceccapp2_D2L_01	2	DIA	1470972	stopped	ENQ		Yes			5	LR_MM037			
hkpoceccapp2_D2L_01	4	UPD	638978	stopped	ENQ		Yes			7	PP24			
hkpoceccapp2_D2L_01	5	UPD	524310	stopped	ENQ		Yes			14	PP24			
hkpoceccapp2_D2L_01	6	DIA	1339874	stopped	ENQ		Yes			3	LR_MM055			
hkpoceccapp2_D2L_01	7	DIA	389210	stopped	ENQ		Yes			2	LR_MM026			
hkpoceccapp2_D2L_01	8	DIA	704636	stopped	ENQ		Yes			1	LR_MM050			
hkpoceccapp2_D2L_01	9	DIA	1544634	stopped	ENQ		Yes			1	LR_MM051			
hkpoceccapp2_D2L_01	10	DIA	1049014	stopped	ENQ		Yes				LR_MM056			
hkpoceccapp2_D2L_01	11	DIA	372860	stopped	ENQ		Yes				LR_MM040			
hkpoceccapp2_D2L_01	12	DIA	598020	stopped	ENQ		Yes			72	PP24	SAPLSENA		
hkpoceccapp2_D2L_01	13	DIA	569500	stopped	ENQ		Yes			16	PP24			
hkpoceccapp2_D2L_01	14	DIA	606398	stopped	ENQ		Yes			33	PP24	SAPLSENA		
hkpoceccapp2_D2L_01	15	DIA	1249540	stopped	ENQ		Yes				LR_MM042			
hkpoceccapp2_D2L_01	16	DIA	512156	stopped	ENQ		Yes			65	PP24	SAPLSENA		
hkpoceccapp2_D2L_01	18	DIA	790612	stopped	ENQ		Yes			5	LR_PP002			
hkpoceccapp2_D2L_01	19	DIA	520404	stopped	ENQ		Yes			5	LR_PP005			
hkpoceccapp2_D2L_01	20	DIA	295064	stopped	ENQ		Yes			5	LR_PP001			
hkpoceccapp2_D2L_01	22	DIA	303278	stopped	ENQ		Yes			17	PP24			
hkpoceccapp2_D2L_01	25	DIA	1360164	running			Yes				LR_FI040			
hkpoceccapp2_D2L_01	26	DIA	733342	stopped	ENQ		Yes			60	PP24	SAPLKCOA		
hkpoceccapp2_D2L_01	28	DIA	786518	stopped	ENQ		Yes			5	LR_PP006			
hkpoceccapp2_D2L_01	29	DIA	647346	stopped	ENQ		Yes				LR_MM049			
hkpoceccapp2_D2L_01	37	DIA	839836	stopped	ENQ		Yes			5	LR_SD008			
hkpoceccapp2_D2L_01	38	DIA	729226	stopped	ENQ		Yes			2	LR_MM007			
hkpoceccapp2_D2L_01	41	UPD	725118	stopped	ENQ		Yes			8	PP24			
hkpoceccapp2_D2L_01	45	BTC	610486	stopped	CPIC		Yes			17	PP24			
hkpoceccapp2_D2L_01	46	BTC	1565026	stopped	ENQ		Yes			1144	LIU			
hkpoceccapp2_D2L_01	47	BTC	1044910	stopped	ENQ		Yes			1144	LIU	SAPLRPHI		
hkpoceccapp2_D2L_01	48	BTC	1409426	stopped	ENQ		Yes			1143	LIU			
hkpoceccapp2_D2L_01	49	BTC	1397084	stopped	ENQ		Yes			1142	LIU			
hkpoceccapp3_D2L_02	1	UPD	639170	stopped	ENQ		Yes			9	PP24	SAPLSENA		
hkpoceccapp3_D2L_02	2	DIA	1614100	stopped	ENQ		Yes			3	LR_MM025	SAPLSENA		
hkpoceccapp3_D2L_02	3	UPD	495712	running			Yes			15	PP24	SAPLGLIU	Update	COFIT
hkpoceccapp3_D2L_02	4	UPD	491550	stopped	ENQ		Yes			18	PP24	SAPLORFC		
hkpoceccapp3_D2L_02	5	UPD	831750	stopped	ENQ		Yes			15	PP24	SAPLSENA		
hkpoceccapp3_D2L_02	6	DIA	950732	stopped	CPIC		Yes			17	PP24	SAPLSENT		

Figure 6-8 Global Work process overview

Nevertheless, the enqueues were held only for a short time, as Figure 6-8 shows for some enqueues.

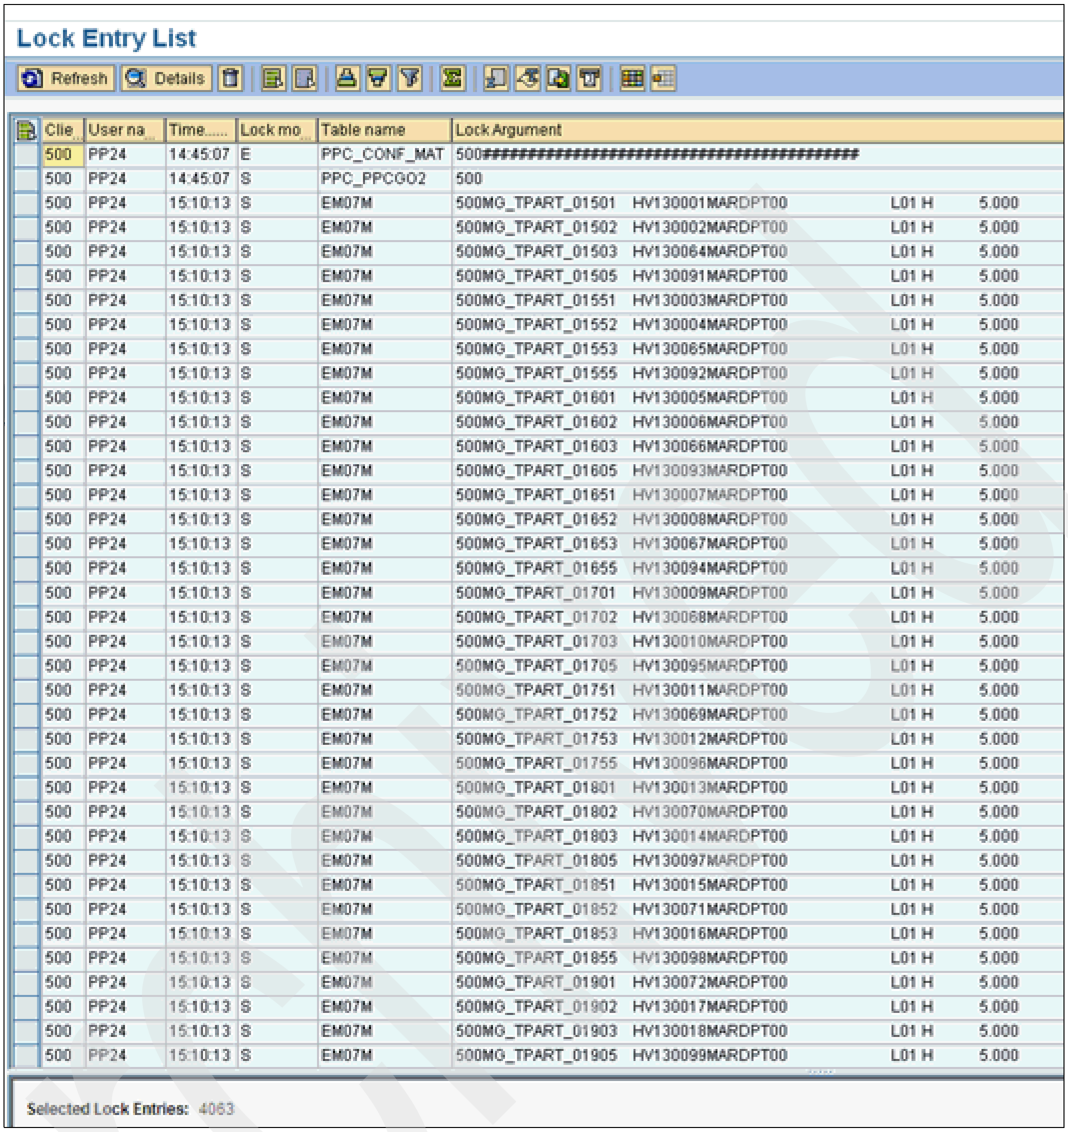

Lock Entry List

Refresh | Details

Clie	User na	Time	Lock mo	Table name	Lock Argument
500	PP24	14:45:07	E	PPC_CONF_MAT	500##
500	PP24	14:45:07	S	PPC_PPCGO2	500
500	PP24	15:10:13	S	EM07M	500MG_TPART_01501 HV130001MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01502 HV130002MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01503 HV130064MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01505 HV130091MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01551 HV130003MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01552 HV130004MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01553 HV130065MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01555 HV130092MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01601 HV130005MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01602 HV130006MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01603 HV130066MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01605 HV130093MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01651 HV130007MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01652 HV130008MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01653 HV130067MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01655 HV130094MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01701 HV130009MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01702 HV130068MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01703 HV130010MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01705 HV130095MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01751 HV130011MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01752 HV130069MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01753 HV130012MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01755 HV130096MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01801 HV130013MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01802 HV130070MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01803 HV130014MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01805 HV130097MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01851 HV130015MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01852 HV130071MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01853 HV130016MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01855 HV130098MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01901 HV130072MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01902 HV130017MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01903 HV130018MARDPT00 L01 H 5.000
500	PP24	15:10:13	S	EM07M	500MG_TPART_01905 HV130099MARDPT00 L01 H 5.000

Selected Lock Entries: 4063

Figure 6-9 Lock Entry list

ZPPC_CONF_GO is a copy of PPC_CONF_GO but with the customer modification (parallelization). "Z" or "Y" in SAP usually means that this is customer-owned coding.

To improve the enqueue performance the central instance should have at least as many dialog work processes as the number of parallel processes that are started by report ZPPC_CONF_GO.

The following SQL trace (Figure 6-10 on page 142) shows that all database operations are fast.

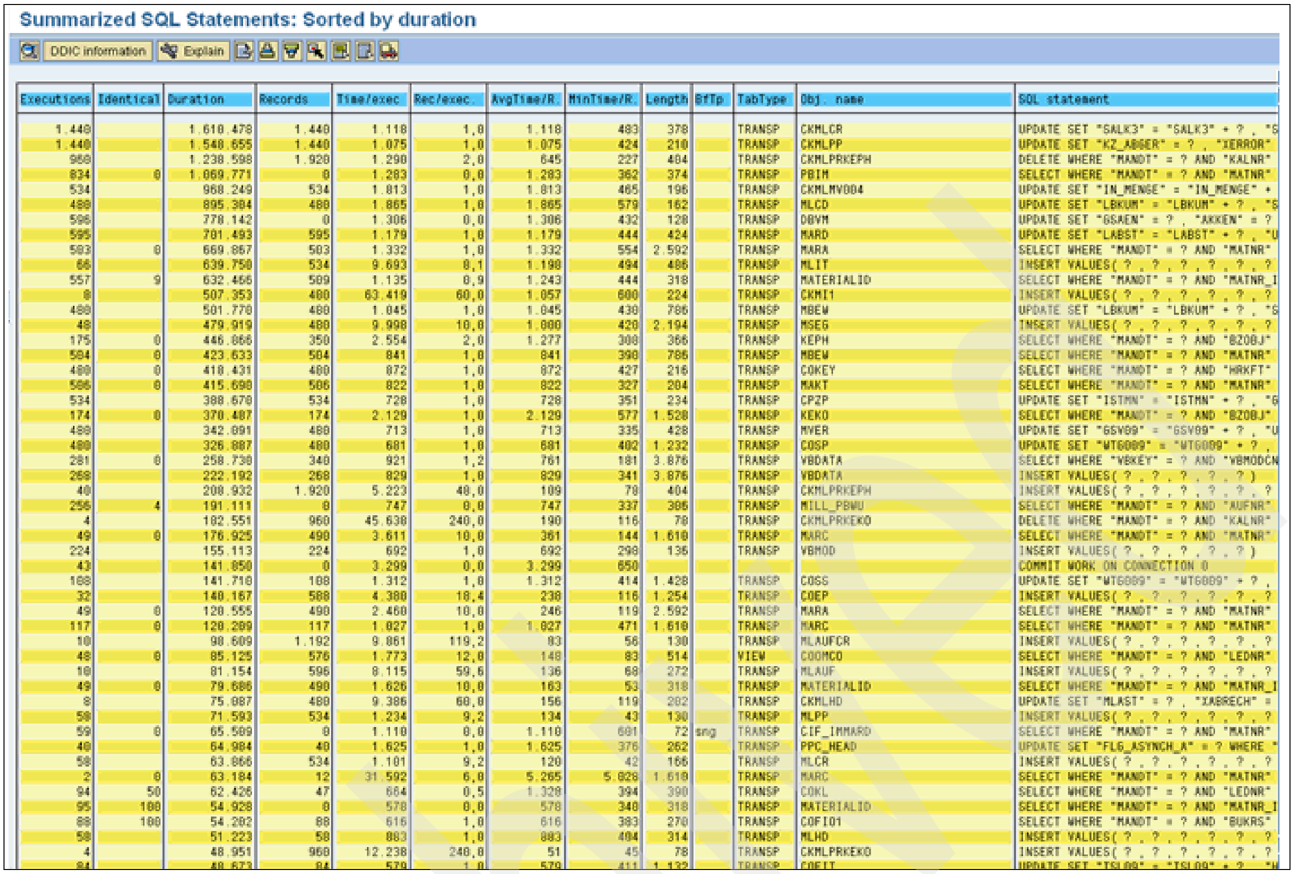

Summarized SQL Statements: Sorted by duration

Executions	Identical	Duration	Records	Time/exec	Rec/exec.	AvgTime/R.	MinTime/R.	Length	BfTp	TabType	Obj. name	SQL statement
1.440		1.610.478	1.440	1.118	1,0	1.118	483	378		TRANSP	CKMLCR	UPDATE SET "SALK3" = "SALK3" + ? , "S
1.440		1.548.655	1.440	1.075	1,0	1.075	424	210		TRANSP	CKMLPP	UPDATE SET "KZ_ABGER" = ? , "XERROR"
960		1.238.598	1.920	1.290	2,0	645	227	404		TRANSP	CKMLPRKEPH	DELETE WHERE "MANDT" = ? AND "KALNR"
834	0	1.069.771	0	1.283	0,0	1.283	362	374		TRANSP	PBIM	SELECT WHERE "MANDT" = ? AND "MATNR"
534		968.249	534	1.813	1,0	1.813	465	196		TRANSP	CKMLMV004	UPDATE SET "IN_MENGE" = "IN_MENGE" +
480		895.384	480	1.865	1,0	1.865	579	162		TRANSP	MLCD	UPDATE SET "LBKUM" = "LBKUM" + ? , "S
596		778.142	0	1.306	0,0	1.306	432	128		TRANSP	DBVM	UPDATE SET "GSAEN" = ? , "AKKEN" = ?
595		701.493	595	1.179	1,0	1.179	444	424		TRANSP	MARD	UPDATE SET "LABST" = "LABST" + ? , "U
503	0	669.867	503	1.332	1,0	1.332	554	2.592		TRANSP	MARA	SELECT WHERE "MANDT" = ? AND "MATNR"
66		639.750	534	9.693	8,1	1.198	494	486		TRANSP	MLIT	INSERT VALUES(? , ? , ? , ? , ? , ?
557	9	632.466	509	1.135	0,9	1.243	444	318		TRANSP	MATERIALID	SELECT WHERE "MANDT" = ? AND "MATNR_I
8		507.353	480	63.419	60,0	1.057	600	224		TRANSP	CKMI1	INSERT VALUES(? , ? , ? , ? , ? , ?
480		501.770	480	1.045	1,0	1.045	430	786		TRANSP	MBEW	UPDATE SET "LBKUM" = "LBKUM" + ? , "S
48		479.919	480	9.998	10,0	1.000	420	2.194		TRANSP	MSEG	INSERT VALUES(? , ? , ? , ? , ? , ?
175	0	446.866	350	2.554	2,0	1.277	308	366		TRANSP	KEPH	SELECT WHERE "MANDT" = ? AND "BZOBJ"
504	0	423.633	504	841	1,0	841	398	786		TRANSP	MBEW	SELECT WHERE "MANDT" = ? AND "MATNR"
480	0	418.431	480	872	1,0	872	427	216		TRANSP	COKEY	SELECT WHERE "MANDT" = ? AND "HRKFT"
506	0	415.690	506	822	1,0	822	327	204		TRANSP	MAKT	SELECT WHERE "MANDT" = ? AND "MATNR"
534		388.670	534	728	1,0	728	351	234		TRANSP	CP2P	UPDATE SET "ISTMN" = "ISTMN" + ? , "G
174	0	370.487	174	2.129	1,0	2.129	577	1.528		TRANSP	KEKO	SELECT WHERE "MANDT" = ? AND "BZOBJ"
480		342.091	480	713	1,0	713	335	428		TRANSP	MVER	UPDATE SET "GSV09" = "GSV09" + ? , "U
480		326.887	480	681	1,0	681	402	1.232		TRANSP	COSP	UPDATE SET "WTG009" = "WTG009" + ? ,
281	0	258.730	340	921	1,2	761	181	3.876		TRANSP	VBDATA	SELECT WHERE "VBKEY" = ? AND "VBMODCN
268		222.192	268	829	1,0	829	341	3.876		TRANSP	VBDATA	INSERT VALUES(? , ? , ? , ? , ?)
40		208.932	1.920	5.223	48,0	109	78	404		TRANSP	CKMLPRKEPH	INSERT VALUES(? , ? , ? , ? , ? , ?
256	4	191.111	0	747	0,0	747	337	306		TRANSP	MILL_PBWU	SELECT WHERE "MANDT" = ? AND "AUFNR"
4		182.551	960	45.638	240,0	190	116	78		TRANSP	CKMLPRKEKO	DELETE WHERE "MANDT" = ? AND "KALNR"
49	0	176.925	490	3.611	10,0	361	144	1.610		TRANSP	MARC	SELECT WHERE "MANDT" = ? AND "MATNR"
224		155.113	224	692	1,0	692	298	136		TRANSP	VBMOD	INSERT VALUES(? , ? , ? , ? , ?)
43		141.850	0	3.299	0,0	3.299	650					COMMIT WORK ON CONNECTION 0
108		141.710	108	1.312	1,0	1.312	414	1.428		TRANSP	COSS	UPDATE SET "WTG009" = "WTG009" + ? ,
32		140.167	588	4.380	18,4	238	116	1.254		TRANSP	COEP	INSERT VALUES(? , ? , ? , ? , ? , ?
49	0	120.555	490	2.460	10,0	246	119	2.592		TRANSP	MARA	SELECT WHERE "MANDT" = ? AND "MATNR"
117	0	120.209	117	1.027	1,0	1.027	471	1.610		TRANSP	MARC	SELECT WHERE "MANDT" = ? AND "MATNR"
10		98.609	1.192	9.861	119,2	83	56	130		TRANSP	MLAUFCR	INSERT VALUES(? , ? , ? , ? , ? , ?
48	0	85.125	576	1.773	12,0	148	83	514		VIEW	COOMCO	SELECT WHERE "MANDT" = ? AND "LEDNR"
10		81.154	596	8.115	59,6	136	68	272		TRANSP	MLAUF	INSERT VALUES(? , ? , ? , ? , ? , ?
49	0	79.686	490	1.626	10,0	163	53	318		TRANSP	MATERIALID	SELECT WHERE "MANDT" = ? AND "MATNR_I
8		75.087	480	9.386	60,0	156	119	202		TRANSP	CKMLHD	UPDATE SET "MLAST" = ? , "XABRECH" =
58		71.593	534	1.234	9,2	134	43	130		TRANSP	MLPP	INSERT VALUES(? , ? , ? , ? , ? , ?
59	0	65.509	0	1.110	0,0	1.110	601	72	sng	TRANSP	CIF_IMMARD	SELECT WHERE "MANDT" = ? AND "MATNR"
40		64.984	40	1.625	1,0	1.625	376	262		TRANSP	PPC_HEAD	UPDATE SET "FLG_ASYNCH_A" = ? WHERE "
58		63.866	534	1.101	9,2	120	42	166		TRANSP	MLCR	INSERT VALUES(? , ? , ? , ? , ? , ?
2	0	63.184	12	31.592	6,0	5.265	5.028	1.610		TRANSP	MARC	SELECT WHERE "MANDT" = ? AND "MATNR"
94	50	62.426	47	664	0,5	1.328	394	390		TRANSP	COKL	SELECT WHERE "MANDT" = ? AND "LEDNR"
95	100	54.928	0	578	0,0	578	348	318		TRANSP	MATERIALID	SELECT WHERE "MANDT" = ? AND "MATNR_I
88	100	54.202	88	616	1,0	616	383	270		TRANSP	COFI01	SELECT WHERE "MANDT" = ? AND "BUKRS"
58		51.223	58	883	1,0	883	404	314		TRANSP	MLHD	INSERT VALUES(? , ? , ? , ? , ? , ?
4		48.951	960	12.238	240,0	51	45	78		TRANSP	CKMLPRKEKO	INSERT VALUES(? , ? , ? , ? , ? , ?
84		48.673	84	579	1,0	579	411	1.132		TRANSP	COFIT	UPDATE SET "TSL09" = "TSL09" + ? , "H

Figure 6-10 Summarized SQL statements

During this run, we never observed that 80 parallel processes are used. This indicates that the work packages (LUWs) cannot be run with such a high parallelization without causing either wait times when trying to get material enqueues or failed material movements (that is, COGI records).

A decrease in the total run time can be expected if the enqueue performance is improved.

With the run time of the modified one-step backflush (27.3 minutes) we achieved a result that exceeded the customer's expectations.

6.3 Results overview

In this test, the one-step backflush (PPChGO) and the first step of the two-step backflush (also PPCGO) were compared against each other and against the KPI of 30 minutes. The second step of the two-step backflush (PPCGO2) was evaluated as far as possible. However, due to the tight schedule we are not able to test this second step with a daily data volume.

One-step backflush results

Report: PPC_CONF_GO

The best run time we achieved was 122 minutes.

The one-step backflush cannot be parallelized with the present data because the component lists of the work packages contain too many common parts. This leads to a situation in which parallel processes are requesting enqueues (SAP locks) for the components in their work package. At some point, the request is rejected because another process is already holding the enqueue. The standard is to retry every second; the number of retries is specified in the customization. This mechanism becomes critical when the component lists contain the same components but in different order. In the debugger, we observed exactly this behavior.

Without parallelization, the KPI of 30 minutes cannot be met. We can conclude that it is impossible with the standard one-step backflush to process the hourly backflush volume of the customer (as given in the test data) within one hour.

Results for the first step of the two-step backflush

Report: PPC_CONF_GO

The best run time we achieved was 34 minutes.

We were not able to optimize the parameters of the first step of the two-step backflush in such a way that a run time below 30 minutes was possible. However, it is clearly possible to process an hourly backflush volume of the customer (as given in the test data) within one hour. Furthermore, parallelization is possible (but not implemented in the standard) and can be realized by customer adoption of the programs.

The first step of the two-step backflush will not exhibit worse performance if the number of common parts between the car models is increased. Therefore, SAP recommends that the customer use the two-step backflush.

Results for the second step of the two-step backflush

Report: PPC_STEP2_GO

No run time was available for this test.

We could not prepare and perform testing of the second step of the two-step backflush with the daily data volume because of our tight schedule. The tests that could be performed are the following:

- A 10 hour volume that takes 15.4 hours
- A 4 hour volume that takes 8.3 hours.

It is obvious that the run time does not scale linearly with the data volume. That means, by increasing the data volume by a factor of 2.5 (10 hours instead of 4 hours) we observed only an increase in the runtime by a factor of 1.84.

Nevertheless, after the test it was not clear whether the standard PPCGO2 is able to perform a daily data volume of the customer (as given in the test data) within one day. Furthermore, the customer requires much faster processing. This can be achieved only by parallelizing PPCGO2. This parallelization has been discussed between the customer and SAP on several occasions.

Results for the modified single-step backflush

Report: ZPPC_CONF_GO

The best run time we achieved for this test was 27.3 minutes.

After discussions with the customer and SAP development, we created a modification of PPC_CONF_GO (and function group PPC1PR) to assess the customer's claim that one-step backflush should work for their data. The modification has not yet been tested thoroughly enough to allow it to be used in a production environment.

The modified one-step backflush avoids the enqueue problems that were observed for the standard one-step backflush by checking the component lists before starting new parallel tasks. It is therefore possible to measure whether the number of common parts in the present test data is small enough to achieve a sufficiently high parallelization without locking problems.

The results of our tests are the following:

- It is possible to avoid enqueue problems by modifying the one-step backflush.
- The customer's data (as given in the test) allow for a relatively high parallelization initially. After some minutes, we observed usually fewer than 10 parallel processes.
- It is possible to reach the KPI with the one-step backflush if the logic of starting the parallel tasks is changed.
- The modified one-step backflush is only slightly faster than the first step of the two-step backflush.

The last point deserves special attention. The run time of the one-step backflush is not stable if the bills of material are changed in such a way that more common parts exist between the different models. We anticipate that the global trend in the automotive industry of having an increasing number of common parts between the different model families will also be reflected in the customer's data structures at some point. The modified version of the single-step backflush will run into problems if this happens, meaning the run time of the backflushing process (including the reprocessing in transaction COGI) will increase.

We see two options for handling such a case:

- Allow for a small overlap in the component list
- Run only work packages with overlapping components

In the first case, a (possibly large) increase in the number of reprocessing records will be observed. The time for processing an hourly volume will then be increased primarily by the serial processing with transaction COGI.

In the second case, fewer parallel processes will be started in order to avoid material enqueue problems. In the extreme case, there will be no parallelization. The run time of the single-step backflush itself will increase in this case.

To summarize this briefly: The modified one-step backflush will have worse performance if more common parts are used. The performance deterioration could be so extreme that the run time of the unparallelized standard one-step backflush is observed. The customer would then not be able to handle an hourly backflush volume within one hour.

Common parts between model families are a standard means in the automotive industry to save costs and to increase the efficiency of production. It is therefore very likely that the present test data does not represent the customer's data structure during the whole lifetime of the planned SAP system.

Taking into account this risk and evaluating the run time difference observed between the standard (unparallelized) first step of the two-step backflush and the modified (highly parallelized) one-step backflush, our conclusion is the following:

> It is not feasible for the customer to use the one-step backflush even in a changed version that avoids material enqueue problems. The comparison between the gain in run time (in

comparison to the two-step backflush) and the risk for the future operations of the customer clearly shows that the risks inherent in the one-step backflush are too high.

Recommended backflush variant

Based on the argument presented in the previous section we recommend that the customer use the two-step backflush. Both steps of the two-step backflush can be parallelized, as has been discussed with the customer.

6.4 Analysis of System p infrastructure layer test results

The objective of the test runs at the IBM infrastructure level is to demonstrate that the architecture setup is totally adequate to fulfill the KPI. In this case we have to achieve this KPI in 30 minutes.

In this section we present the system activity summary for the five System p servers LPARs during the run. This run began at 14:45 and finished at 15:16, lasted 31minutes.

6.4.1 LPARs activity summary

The following five figures (Figure 6-11 through Figure 6-15) present the system activity summary for the System p server LPARs, based upon the nmon results.

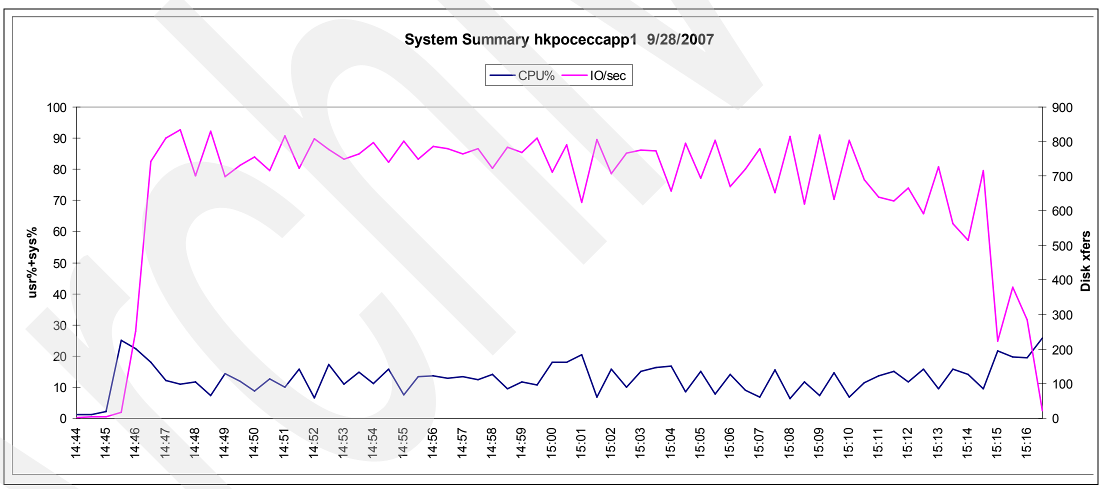

Figure 6-11 Activity summary for hkpoceccapp1

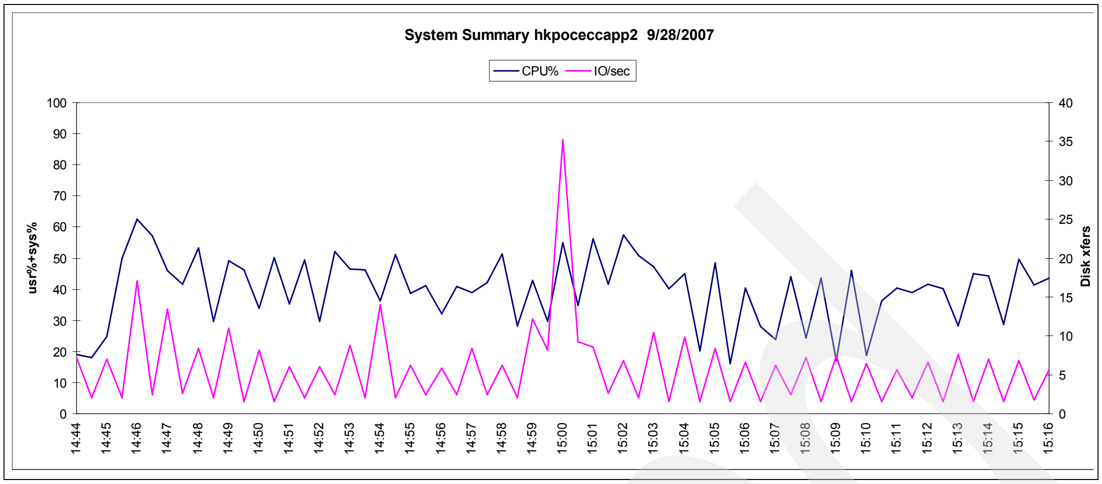

Figure 6-12 Activity summary for hkpoceccapp2

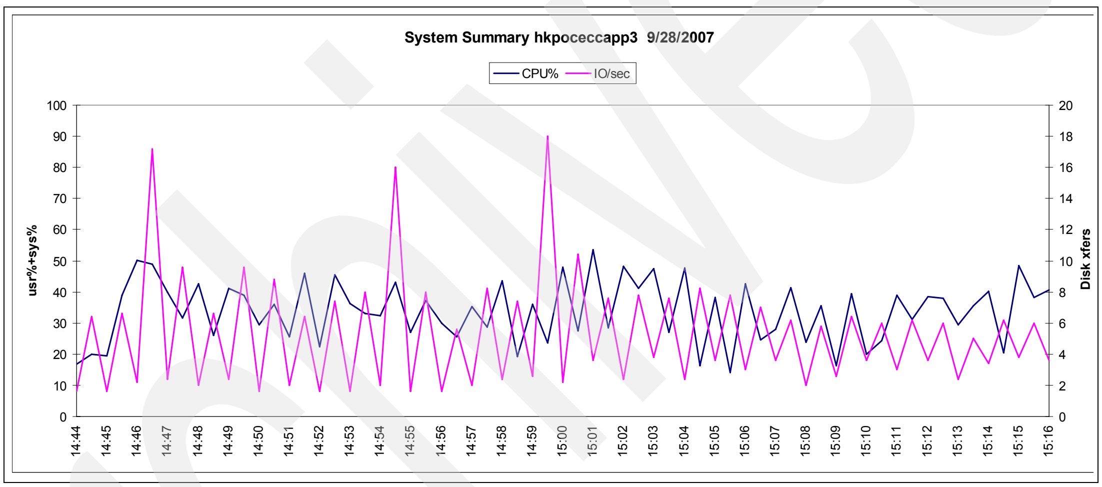

Figure 6-13 Activity summary for hkpoceccapp3

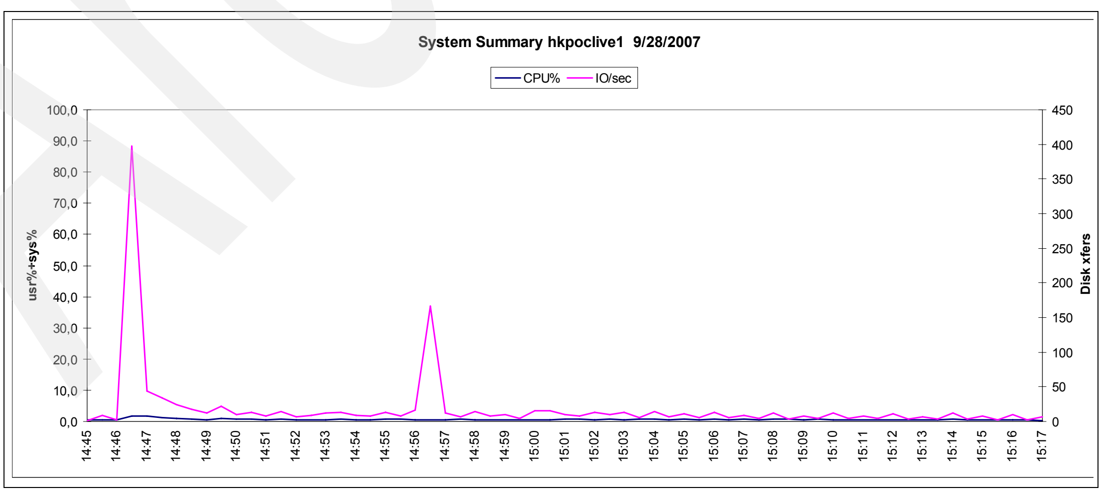

Figure 6-14 Activity summary for hkpoclive1

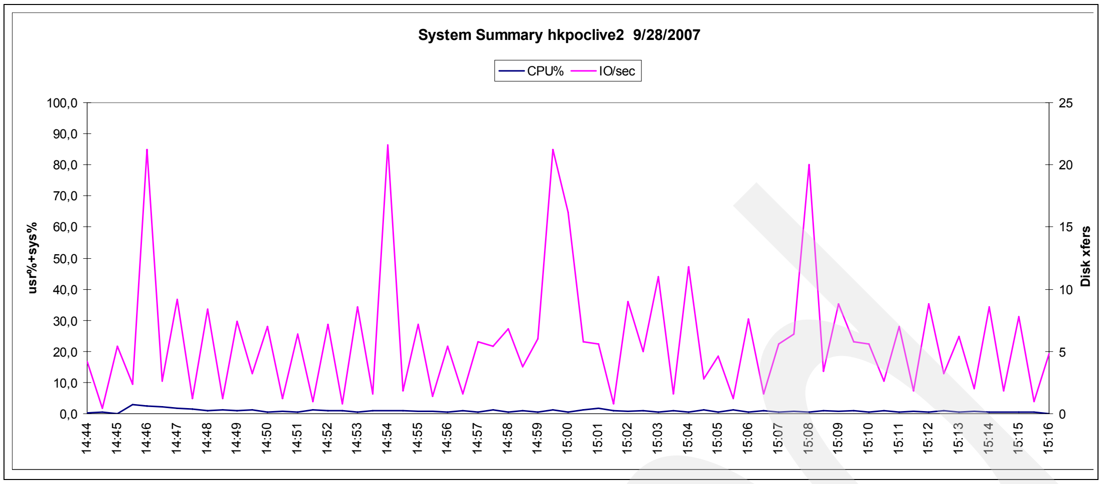

Figure 6-15 Activity summary for hkpoclive2

Regarding the ECC system LPARs (hkpoceccapp1, hkpoceccapp2, and hkpoceccapp3), we see that the ECC system has significant activity on this run. Hkpoceccapp1 has an average usage of 11%, with peaks up to 24%, *not* including CPU wait%. Looking at it, the CPU wait% average is 27% and the CPU wait% peak is 63.4%. We consider this more carefully in the next paragraph. Disk I/Os are very high on this LPAR, averaging 662 I/Os per second, and with a maximum of 833 I/Os per second, so we probably have a disk contention on this LPAR. We address that point in 6.4.4, "Disk details" on page 152.

Hkpoceccapp2 and hkpoceccapp3 have similar behavior during the run. Hkpoceccapp2 has an average CPU consumption of 36.6%, with peaks up to 57.6%. Nevertheless, we can see that there is almost no CPU wait time. I/Os per second (20) are not significant. Hkpoceccapp3 has an average CPU consumption of 30.6%, with peaks of 49.8%, but again we can see that there is almost no CPU wait time, and I/Os per second (20) are not significant. In summary, there is activity on these two LPARs, but they seem to be performing perfectly.

Regarding the SCM instance, the graph shows an average CPU and disk usage close to 0, so we can conclude that this instance has very little influence on the run, and therefore we did not dig any deeper into the details.

Finally, regarding the liveCache instance, the CPU activity is very close to 0. We noticed peak I/Os at the beginning of the run (16:46) of 400 I/O per second, then we had a second peak at 14:56 of 200 I/Os per second. Considering that the liveCache technically is a database, this behavior is not surprising. Having a look at the disk details shows liveCache writes at these moments, but the activity is so small that we did not dig deeper into the details.

> **Note:** These summary graphs do not consider CPU wait percentage.

6.4.2 CPU details

Global CPU study

The next three figures (Figure 6-16 through Figure 6-18) show the detailed CPU activity on ECC system LPARs.

In Figure 6-16, it is interesting to note that the majority of the hkpoceccapp1 CPU time is spent waiting for an I/O. We have no way to know if this is normal for the application, but the most probable answer would be negative. We plan to study in detail the disk behavior for this LPAR.

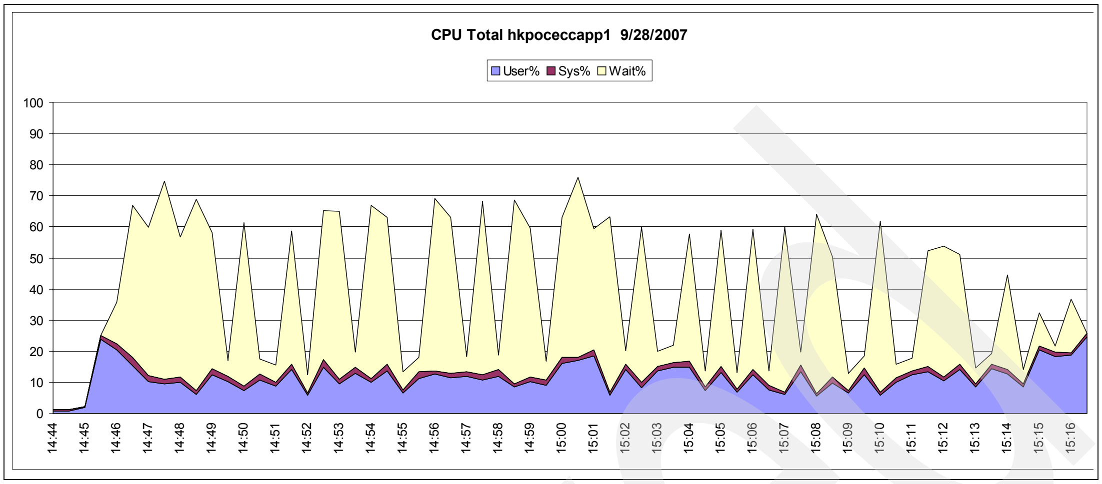

Figure 6-16 CPU details for hkpoceccapp1

Figure 6-17 shows no unexpected behavior for LPAR hkpoceccapp2. CPU usage is never higher than 65% and there is almost no wait time.

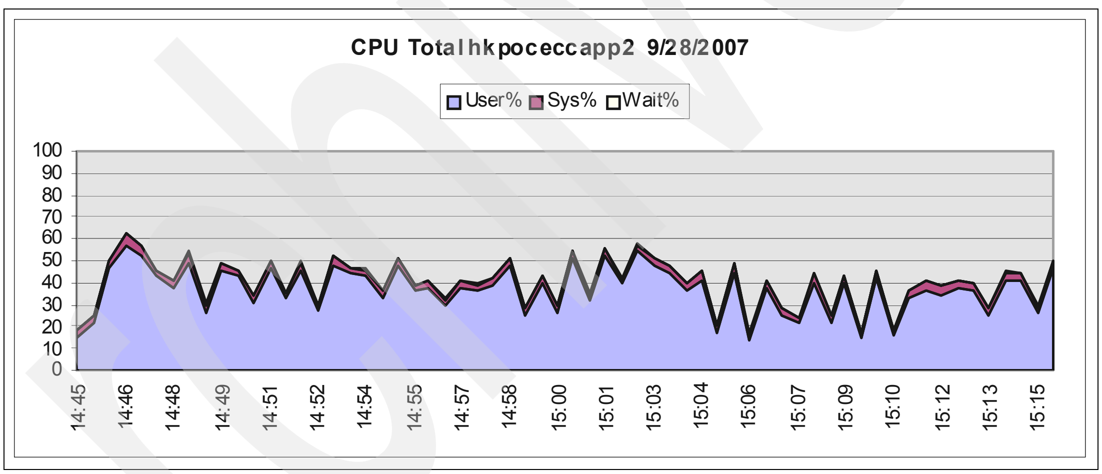

Figure 6-17 CPU details for hkpoceccapp2

Regarding LPAR hkpoceccapp3, the behavior is very similar to hkpoceccapp2 as we can see in Figure 6-18. CPU usage never exceeds 55% and there is almost no wait time.

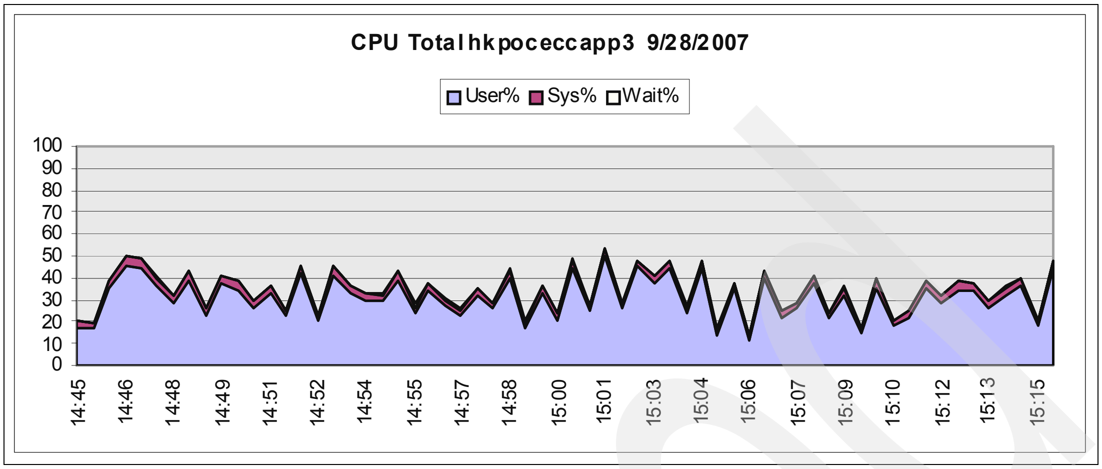

Figure 6-18 CPU details for hkpoceccapp3

Unitary CPU study

As shown in Figure 6-19, the hkpoceccapp1 CPU usage by processor is balanced, and we notice that CPU17 and CPU19 are less used on this LPAR and that they have no wait time. This figure also shows that SMT is not used, so on the whole run, physical CPUs are sufficient to carry the load. We can see that the CPU is absolutely not overloaded in general, so we only need to lower the wait% in order to improve the overall performance.

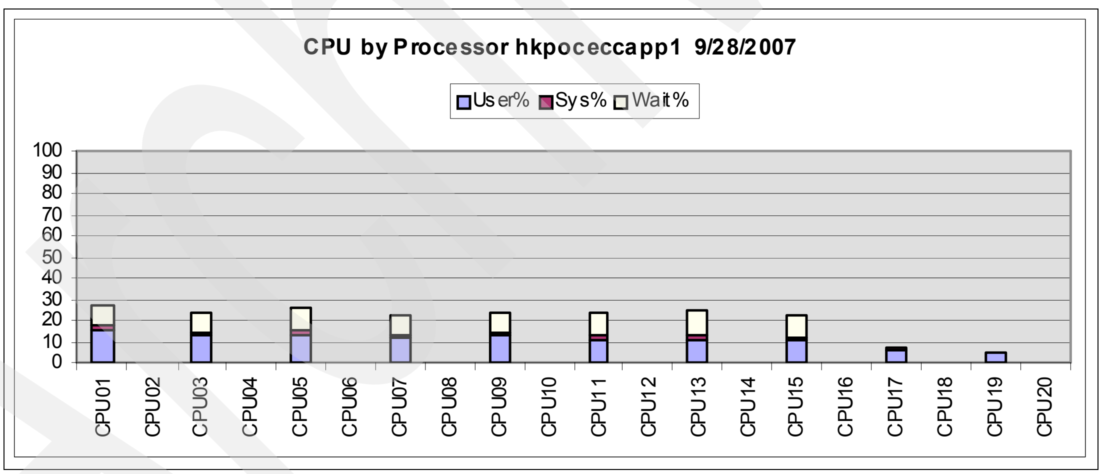

Figure 6-19 hkpoceccapp1 CPU by processor

Regarding LPAR hkpoceccapp2, Figure 6-20 shows an almost homogeneous load distribution between the CPUs (only CPU03 has a different behavior), but SMT is used for all CPUs, so the application is using all benefit from parallelization. Comparing with run 02, a lot of improvement has been realized by the functional team. These same comments apply to hkpoceccapp3, which has a very similar behavior, as shown in Figure 6-21.

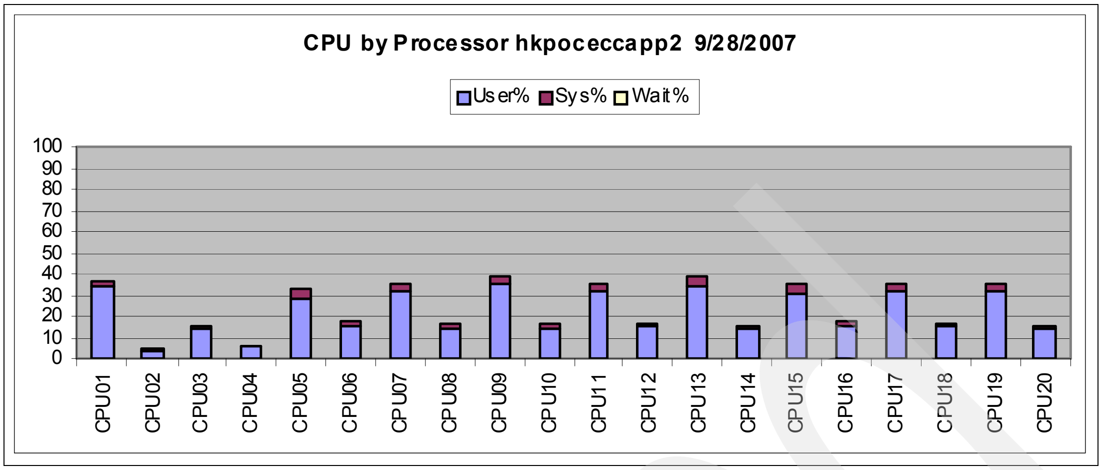

Figure 6-20 hkpoceccapp2 CPU by processor

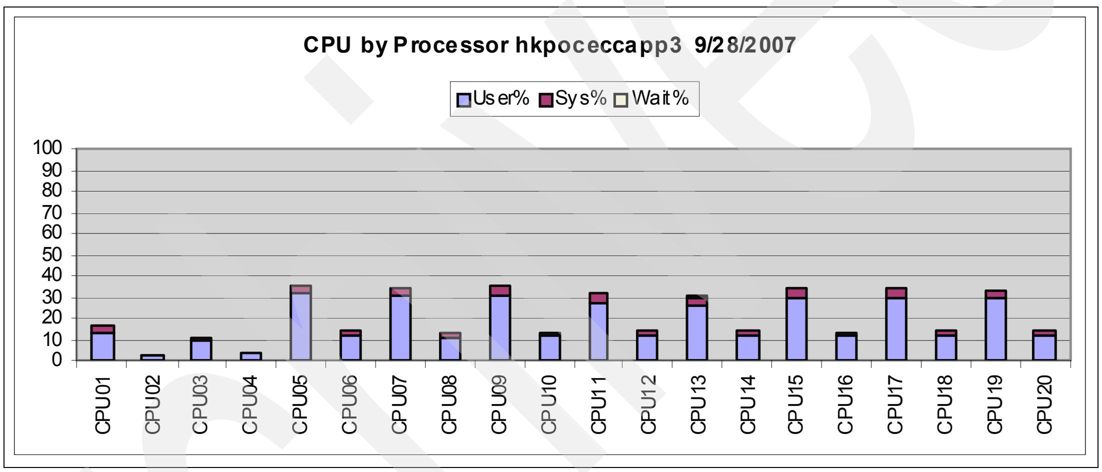

Figure 6-21 hkpoceccapp3 CPU by processor

6.4.3 Memory details

The next three figures (Figure 6-22 through Figure 6-24 on page 151) present the memory details at the LPAR level.

As you can see from these three graphs, the three LPARs have a very similar behavior regarding memory. At the time of the run, each LPAR had 40 GB memory configured. We note that no LPAR consumed more than 15 GB during the run, and no LPAR had paging during the run. Therefore, regarding memory for run 14, everything was fine.

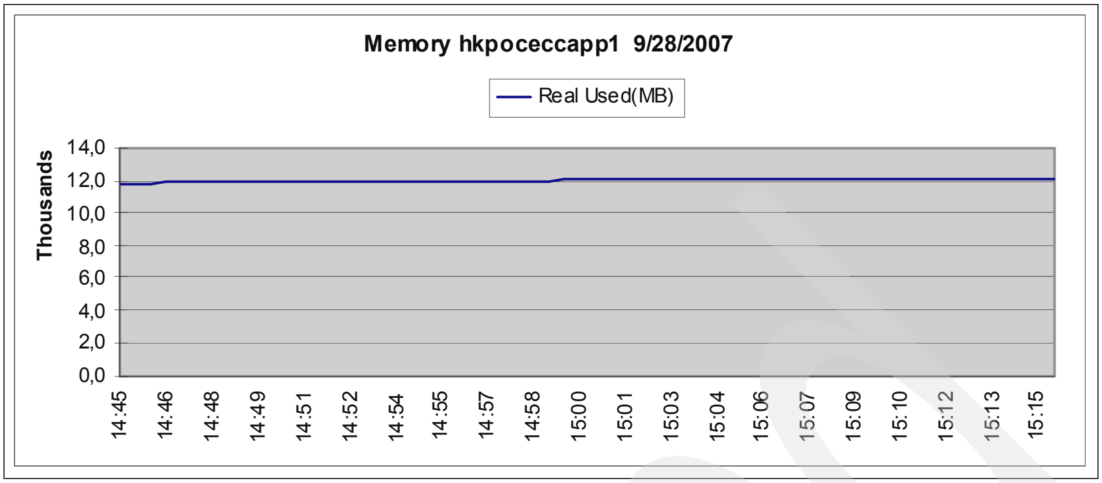

Figure 6-22 hkpoceccapp1 used memory

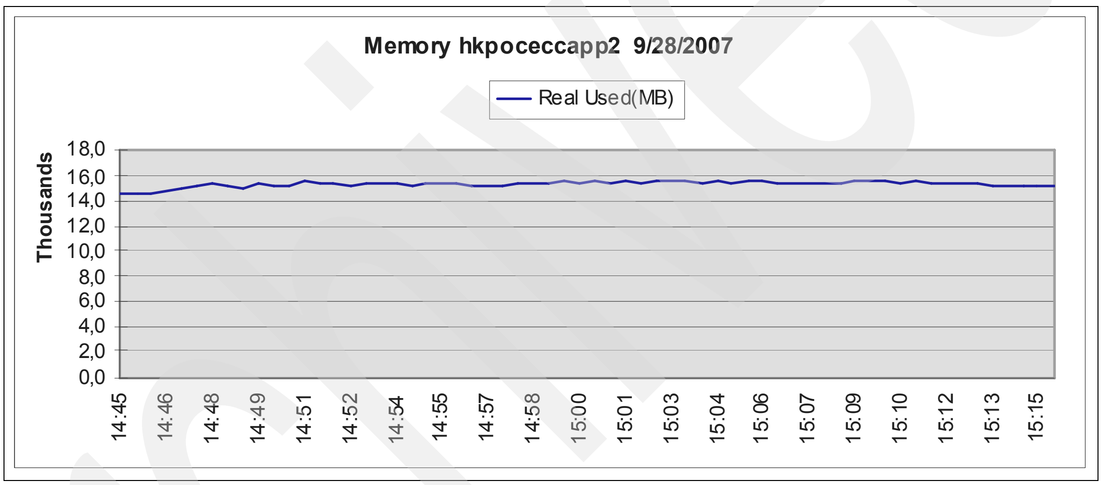

Figure 6-23 hkpoceccapp2 used memory

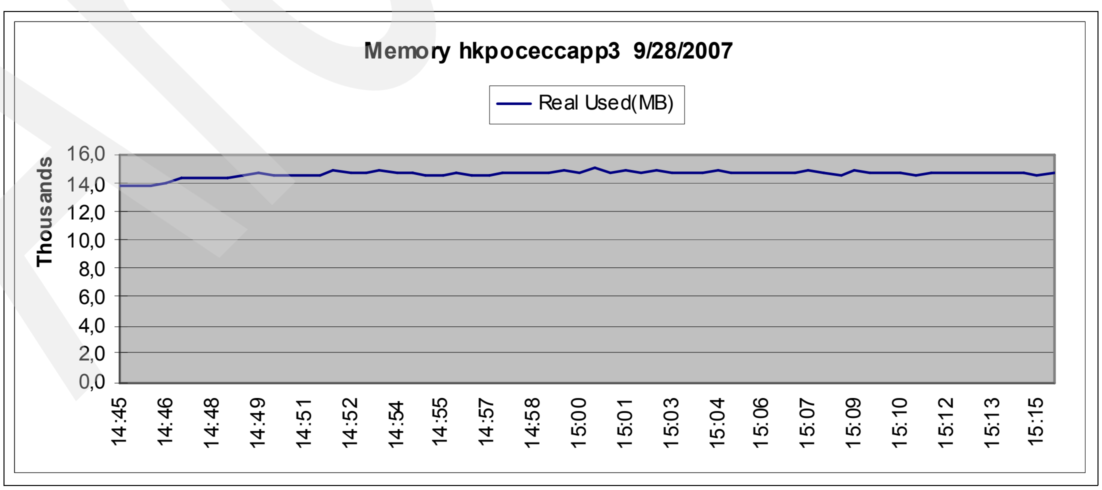

Figure 6-24 hkpoceccapp3 used memory

6.4.4 Disk details

As stated earlier, hkpoceccapp1 shows an unexpected behavior, with a great deal of wait time at the CPU level. In this section we consider hkpoceccapp1 LPAR disk details in an effort to improve this situation.

Disk activity

Disk activity on hkpoclive1 is shown in Figure 6-25 and Figure 6-26.

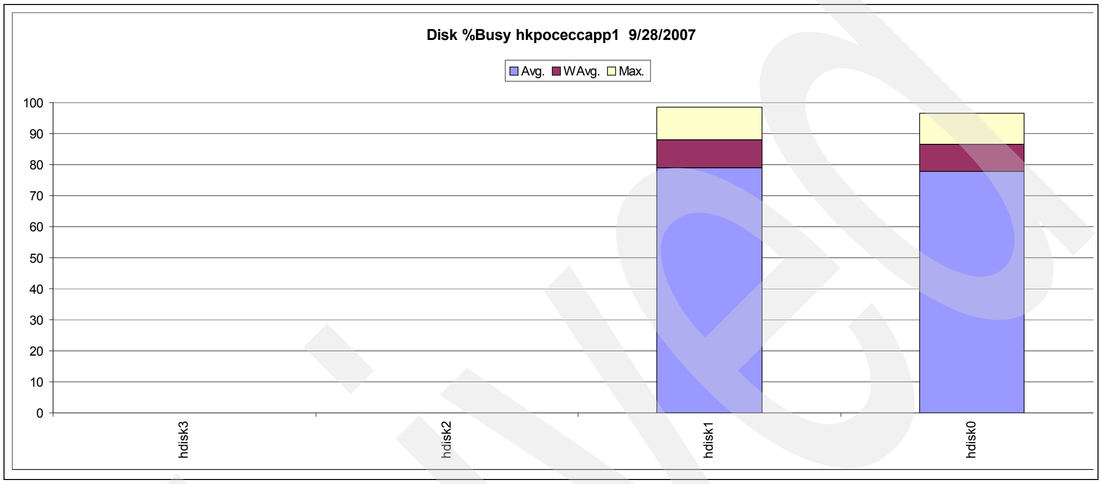

Figure 6-25 Disk activity hkpoceccapp1

In Figure 6-25 we see hdisk0 and hdisk1 have an average activity of 80% and a maximum activity of almost 100%. Because these two disks are part of a mirror link, the activity is almost the same. The “write” activity is performed on both disks, while the “read” activity is performed from either one disk or the other.

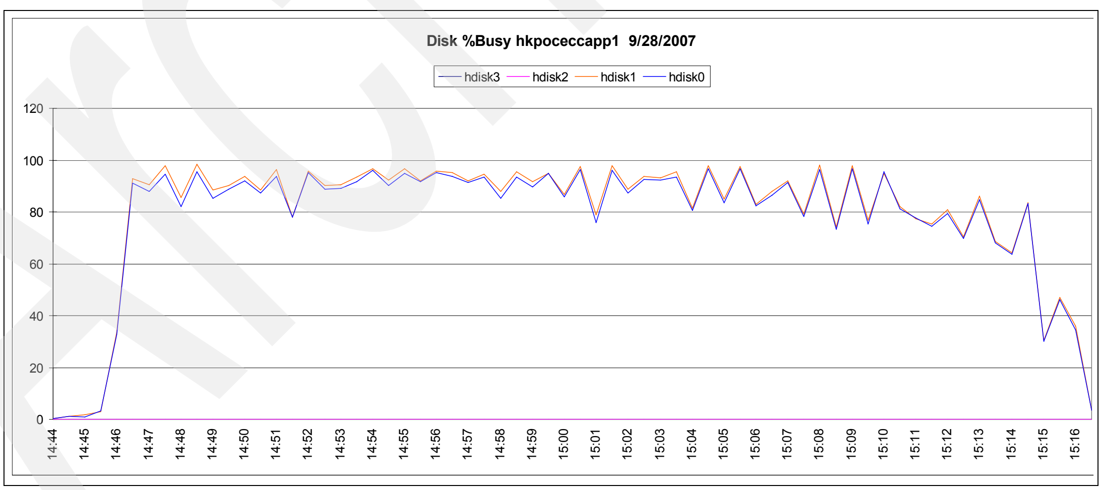

Figure 6-26 Disk activity on time frame

In Figure 6-26 we see that during the entire run we have activity of about 90% at the disk level, so this is the cause of the high wait%. Next we study whether this activity is read or write, and then which processes are causing it.

Disk reads

The next two figures give us statistics about disk reads per disk and within the time frame.

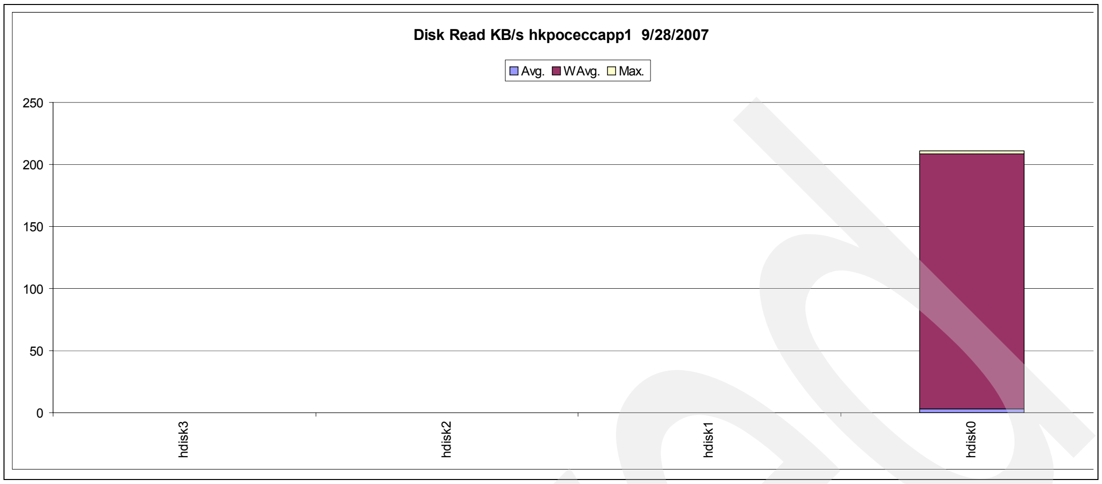

Figure 6-27 hkpoceccapp1 disk reads

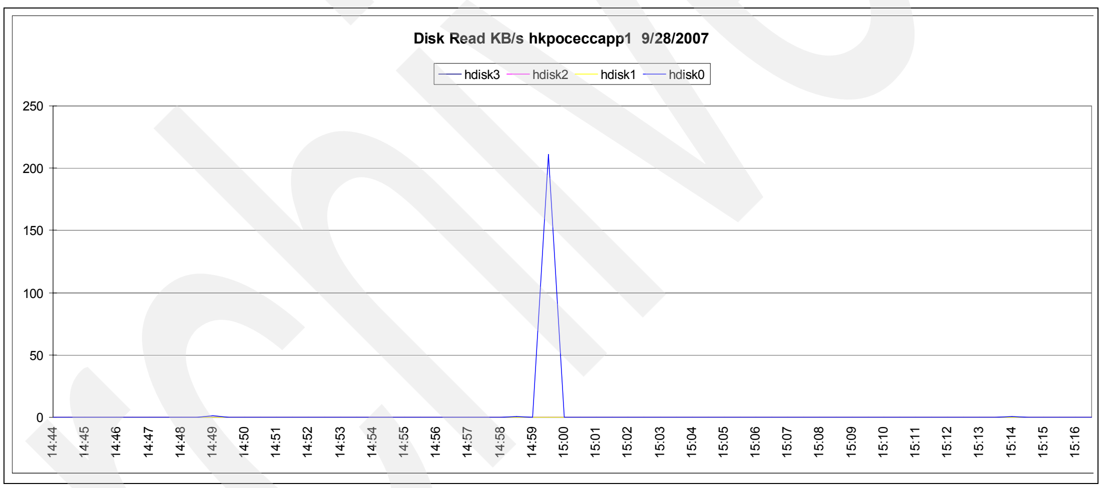

Figure 6-28 hkpoceccapp1 disk reads on time frame

As Figure 6-27 and Figure 6-28 show, there are almost no disk reads (and those there are occur only at 14:59).

Disk writes

Figure 6-29 and Figure 6-30 confirm that all the disk contention is due to disk writes. During the whole run a sustained disk writing activity exists.

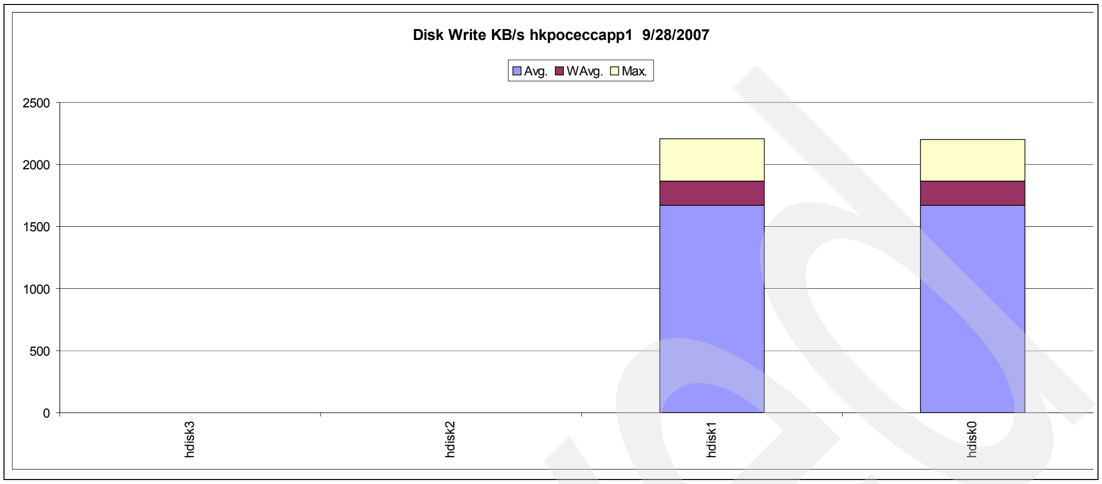

Figure 6-29 hkpoceccapp1 disk writes

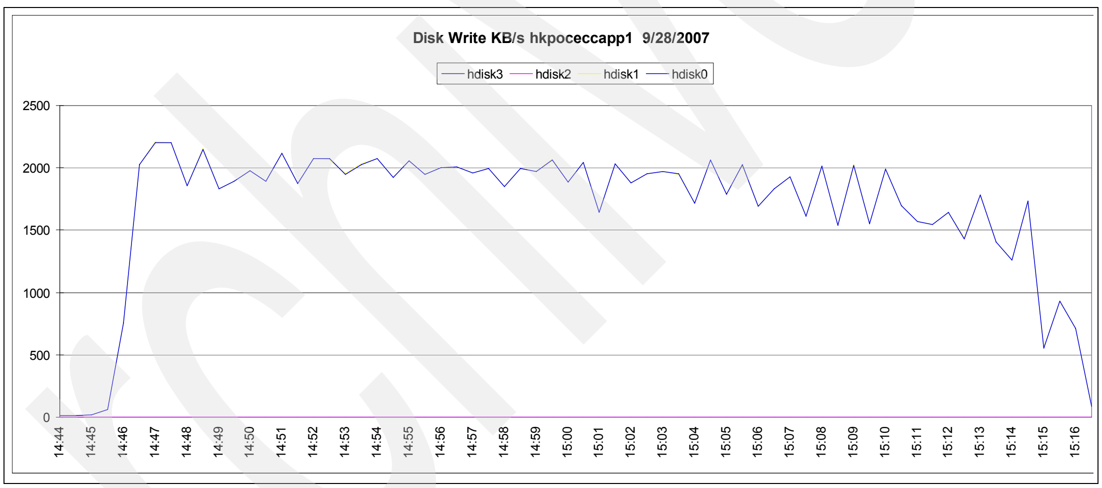

Figure 6-30 hkpoceccapp1 disk writes on time frame

Top disk using processes

Figure 6-31 shows that the top disk consumer is the SAP message server. It is typically writing logs and this is not a usual behavior.

Therefore, the recommendation to improve the performance is to add some disks and put the SAP application server in a different volume group, not rootvg. The Montpellier team took this action. From a pure benchmark run perspective, this disk contention did not affect the results of the run because a lot of CPU resources were available.

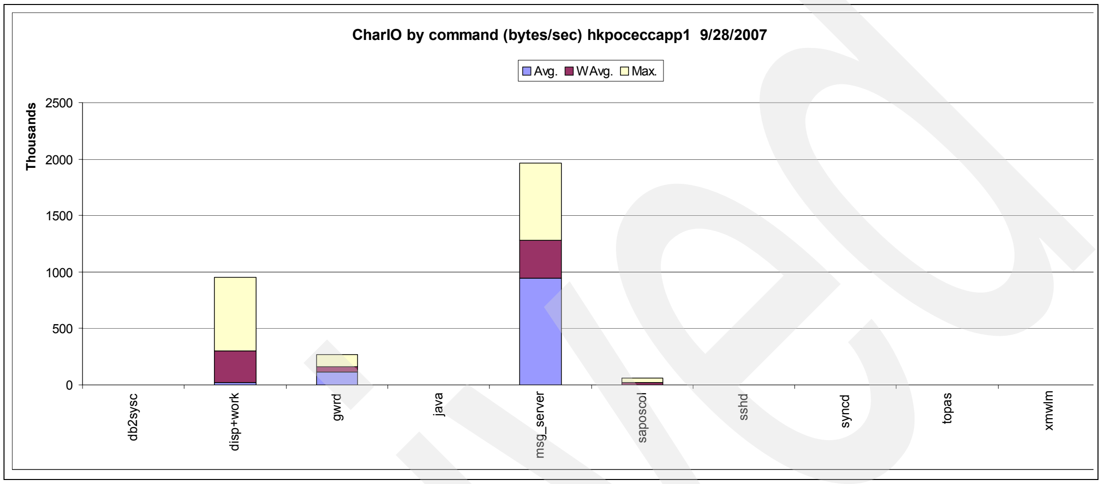

Figure 6-31 Top disk consumers

6.4.5 Network details

Figure 6-32 shows the network load by adapter. We notice that the highest load is on en0 adapters, which is logical because these adapters are connected to the 1 Gbps network for server communications.

The three LPARs have a similar behavior, which is logical because they run the same application, mainly reading from the database and writing.

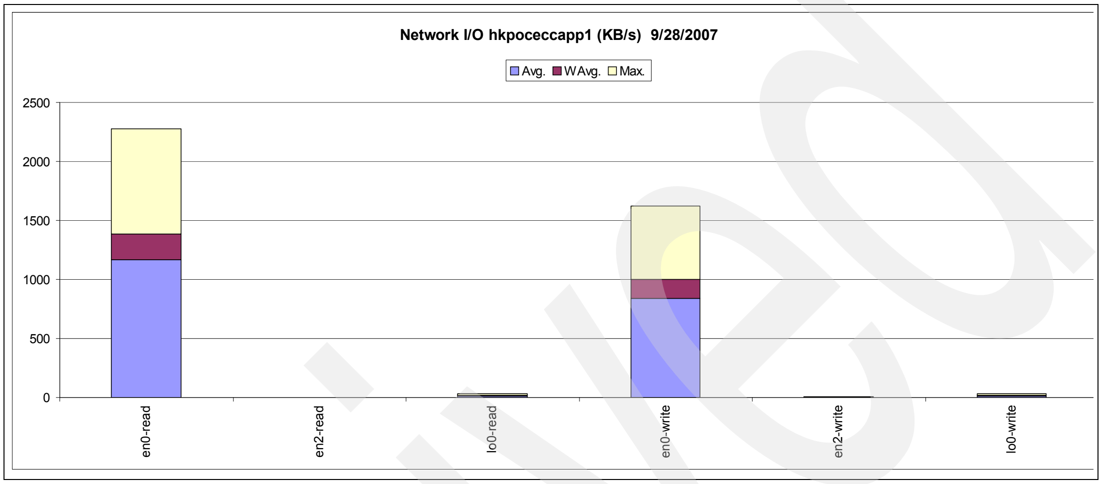

Figure 6-32 hkpoceccapp1 network I/O per adapter

The cumulative network activity for the three LPARs is similar and no contention was produced. Figure 6-33 shows this behavior for hkpoceccapp1. For this run we can conclude that the network is well sized.

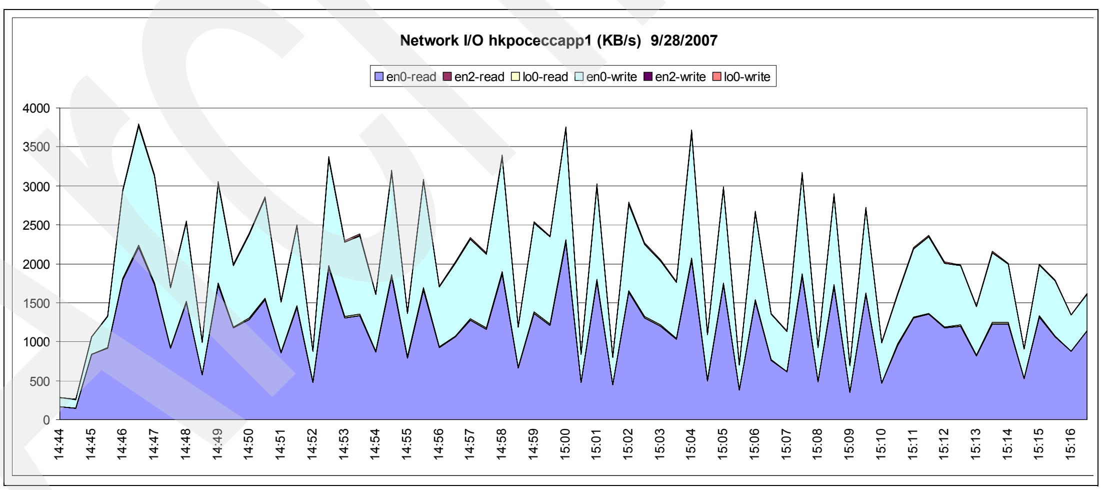

Figure 6-33 hkpoceccapp1 network I/O

6.4.6 Top processes

This section discusses the top CPU and memory consumers.

Top CPU

Figure 6-34 shows that the top processes are disp+work (SAP processes) and db2sysc (DB2 connect, connection manager between DB2/Z and SAP).

On the three application servers (hkpoceccapp1, hkpoceccapp2, hkpoceccapp3) the main contributor to the CPU consumption is the SAP work process activity, which is the expected behavior.

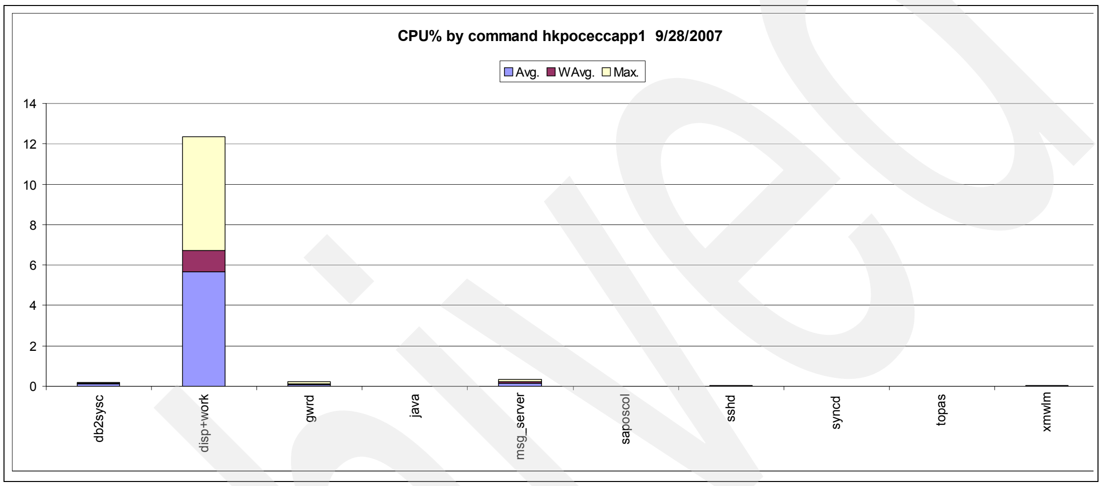

Figure 6-34 hkpoceccapp1 top CPU

Top memory

The three application servers (hkpoceccapp1, hkpoceccapp2, hkpoceccapp3) are again the top processes in terms of memory usage, specifically by disp+work SAP processes and db2sysc (DB2 connect, connection manager between DB2 for z/OS and SAP).

6.4.7 System p architecture conclusions

At the infrastructure level, System p servers responded perfectly to run 14 demands. Absolutely nothing had a negative impact on the run at this level except the disk contention on hkpoceccapp1, which was solved by the Montpellier team.

We performed the following corrective steps:

1. We added two disks and created a new striped volume group on these disks.
2. We created a logical volume reserved for SAP application file systems.
3. We withdrew the SAP elements from rootvg.

This solution eliminated the disk contention.

6.5 Analysis of System z infrastructure layer test results

Unfortunately some SMF data was lost during the tests, so there is no SMF data available for BF run 14. Instead we present data for BF run 10 here. The behavior of the z infrastructure during this test was comparable to the behavior during run 14. The major difference between the two runs was the number of parallel process used: 80 for run 14 and 40 for run 10.

The run began at 16:00 and finished at 16:27, lasting 27 minutes.

Figure 6-35 shows the LPAR summary. The CPU utilization and the DASD rate for LPAR MVH1 was considerably higher compared to BF run 2. However, the values shown are still relatively low for this powerful infrastructure, so there was no need for further investigation on possible bottlenecks in the z environment.

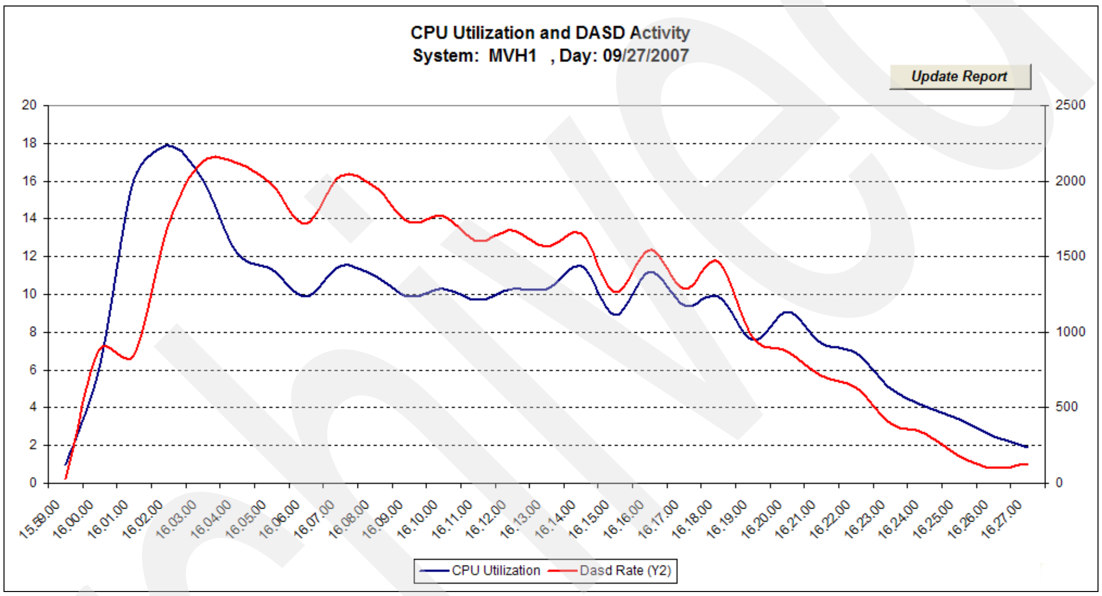

Figure 6-35 CPU Utilization and DASD Activity

Figure 6-36 shows the time repartition between the application and the database. The results are average time based on 28602 accounting records collected during the run interval. As you can see, only 9% of the time is spent in DB2.

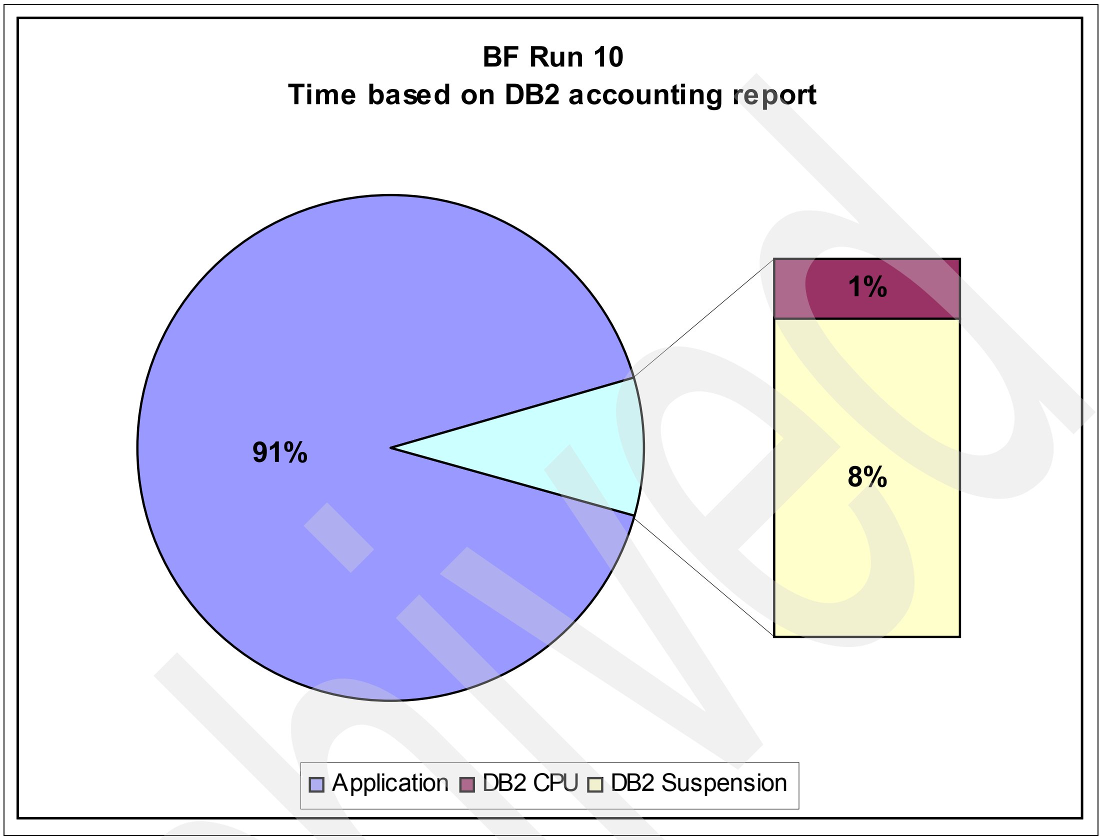

Figure 6-36 DB2 accounting report - Time repartition

The next graph, Figure 6-37, shows the processor usage and the workload eligible for zIIP processing for one minute during the CPU peak usage between 16:02 and 16:03.

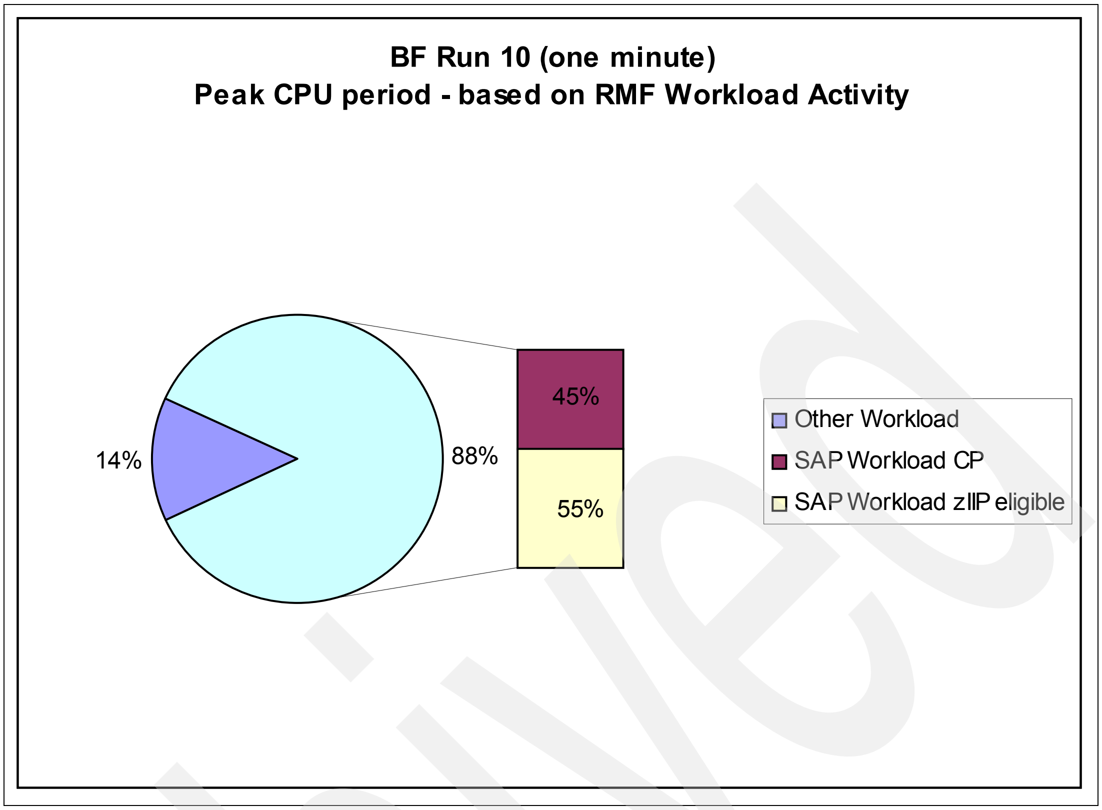

Figure 6-37 One minute RMF Workload Activity - CPU usage

During this one minute of peak CPU usage we used 1.4 CP, of which 1.2 was used by the SAP activity on DB2. Of this SAP activity on DB2, 55% is eligible to be processed by a zIIP processor. This percentage of zIIP processor use is estimated because no processors of this kind were configured on the LPAR at that time. During the Materials Ledger test some zIIP processors were enabled. We include more detailed results in the next chapter.

6.6 Conclusions

In this test, the single-step automotive backflush and the first step of the two-step automotive backflush were compared against each other and against the KPI of 30 minutes. The second step of the two-step backflush was evaluated as far as possible. Yet, the tight schedule made it impossible to test this second step with a daily data volume.

During this test the bottleneck was obviously still not on the hardware. The use of the System p and System z infrastructure is better compared to the initial BF tests, that is, BF run002. However, the values still show no bottlenecks on the System p and System z hardware.

6.6.1 Recommended Backflush variant

Following the discussion in “Results for the modified single-step backflush” on page 143, we recommended to the client that they use the two-step backflush. Both steps of the two-step backflush can be parallelized as discussed between the client and SAP on several occasions.

7

Materials ledger

This chapter describes our experiences testing the materials ledger (ML) performance. We performed 13 runs for this KPI. We describe run 10 in detail here because it achieved the KPI performance objective and because it was composed of the batch process along with an online users load created with Loadrunner.

The topics included in this chapter are:

7.1 ML test overview

Because ML post closing generates documents on calculated data, processing time increases in proportion to the number of accounting documents and the number of line items in each document. Based on optimization of SAP running in other subsidiaries, the customer believes the KPI can be met. The current legacy application's equivalent time window is only four hours, so the client has set this four hour duration as the target for the ML KPI.

The objective of these test runs at the IBM infrastructure level is to demonstrate that the architecture set up is totally adequate to fulfill the KPI in the allotted four hours.

Description

Figure 7-1 locates the material ledger within the client's overall IT process; Figure 7-2 on page 163provides a detailed view of the ML process.

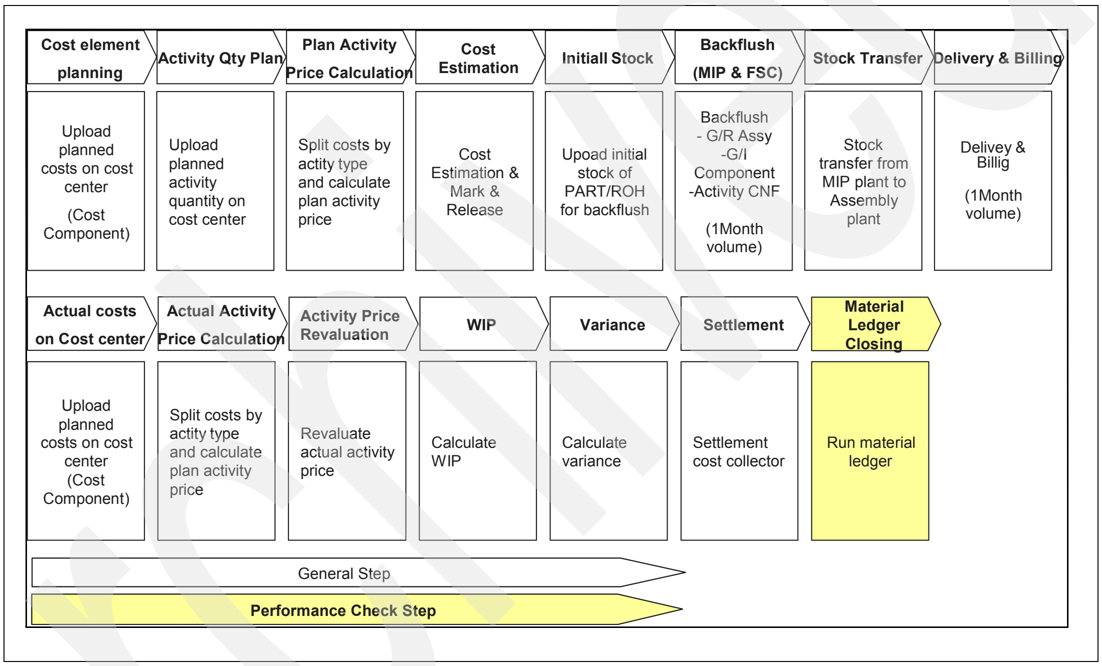

Figure 7-1 ML process within overall process

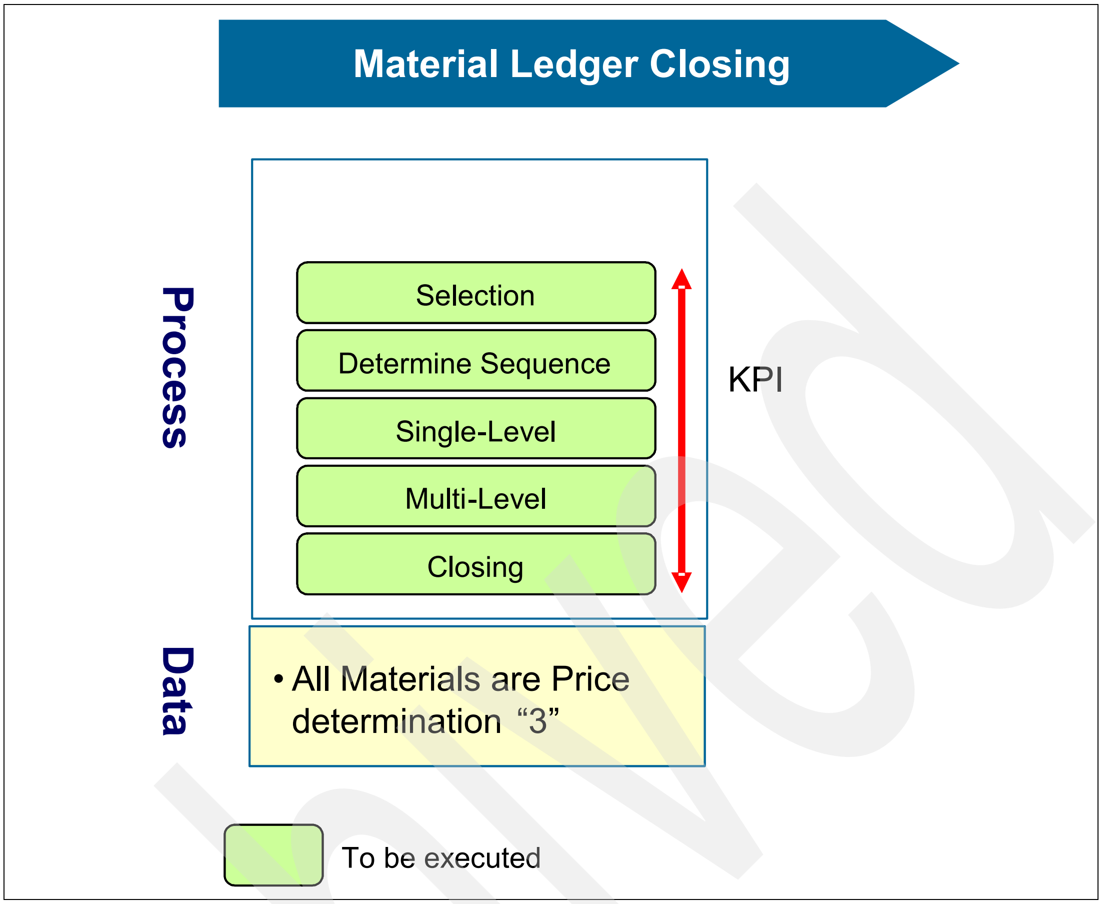

Figure 7-2 ML process: Detailed view

Purpose of the test

The purpose of the test is to determine whether the material ledger computation run can perform on the IBM System z and System p configuration in the required time window. The test also gives an indication of optimal settings for the parameters that are most important for performance.

Test input business volumes:

- Number of models: 21*3
- Number of orders: 185750
- Number of variants: 113400 (1800 variants x 21*3 models)

Test conditions

Our test environment conformed to the following conditions:

- We continued testing in the environment used for RPM and BF.
- Loadrunner background dialogue activity was run.
- One month of materials management data was created with repeated RPM and backflush runs.
- We assume that no time stamping problem occurs when uploading data.

The expected output of the test run is a Valuation Information report.

Test preparation

To measure ML performance, monthly transaction data should be prepared in the following areas of logistics. After generating variance in product cost collection, ML settlement should be carried out.

- Standard cost data generation for all materials to generate standard cost for vehicle, MIP, and end-items.
- Sales closing, which includes normal sales, returned goods, and so forth.
- Production closing, which includes production data such as consumption in production, re-production, and so forth.
- Purchase closing, to have purchase price variance in real situations.
- Cost closing, to close overhead cost data (labor cost, expense, depreciation, and so forth.
- Product cost collection, to generate cost variance in each production step for vehicle/MIP. Some examples are defective goods, variance in inventory due diligence, and so forth.

7.2 Test execution

This test run includes the whole process chain, including the load runner.

Setup

Table 7-1 identifies the SAP parameters that were changed in profile settings for system. The 11 figures that follow (Figure 7-3 on page 166 through Figure 7-13 on page 170) show the dialog windows where the parameters were set.

Table 7-1 Changed SAP Parameters

No	Flow step / Program	Parameter	Last used Value	New Value
001	Selection			
002	Single-Level Price Determination	Number of Tasks No. of Materials in ML Doc.	60 50	60 50
003	Multilevel Price Determination	Packet Size Number of Tasks	10 45	10 40
004	Modification for Multilevel Price Determination	ABAP coding select on a table Include: LCKMLMVRUNF03	`SELECT kalnr kalst INTO CORRESPONDING FIELDS OF TABLE st_kalst_tbl FROM ckmlmv011 WHERE laufid = gd_mgv_runcontrol-laufid. SORT st_kalst_tbl BY kalnr.`	New Coding from From SAP Development: `SELECT SINGLE kalst INTO CORRESPONDING FIELDS OF ls_kalst_tbl FROM ckmlmv011 WHERE laufid = gd_mgv_runcontrol-laufid AND kalnr = v_mat_kalnr`
005	Post Closing	Number of tasks No. of Materials in ML Doc,	60 1	60 1
005	Modification for Post Closing	TPART,ROH	500	250
006	Entry SPLIT-FI-DOC	CKMLVADMIN via TA SE16 to activate the document split according to SAP Note 539452	active	active
007	new index of table CKMLV003	Index z01	-	Index z01
008	Change coding for Load Runner test according to note 51789	RSRZLQVAL	ANSWER_TIME_WEIGHT_FACTOR TYPE I VALUE 5, USER_AMOUNT_WEIGHT_FACTOR TYPE I VALUE 1,	ANSWER_TIME_WEIGHT_FACTOR TYPE I VALUE 1, USER_AMOUNT_WEIGHT_FACTOR TYPE I VALUE 5,
009	Extend RFC server group zPoC_group	Include server hkpoceccapp1	Server 2 & 3	Server 1, 2 & 3
010	Update Catalog for SAPD2L.CKMLMV011~0		NLEVEL = 4 , NPAGE = 10000	NLEVEL = 2 , NPAGE = 800
011	Extend SMLG logon group	Include server hkpoceccapp1	Server 2 & 3	Server 1, 2 & 3

Parameter 001: Selection

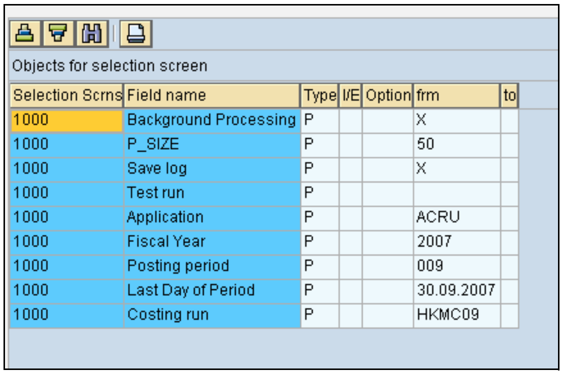

Selection Scrns	Field name	Type	I/E	Option	frm	to
1000	Background Processing	P			X	
1000	P_SIZE	P			50	
1000	Save log	P			X	
1000	Test run	P				
1000	Application	P			ACRU	
1000	Fiscal Year	P			2007	
1000	Posting period	P			009	
1000	Last Day of Period	P			30.09.2007	
1000	Costing run	P			HKMC09	

Figure 7-3 Parameter 001

Parameter 002: Determine sequence

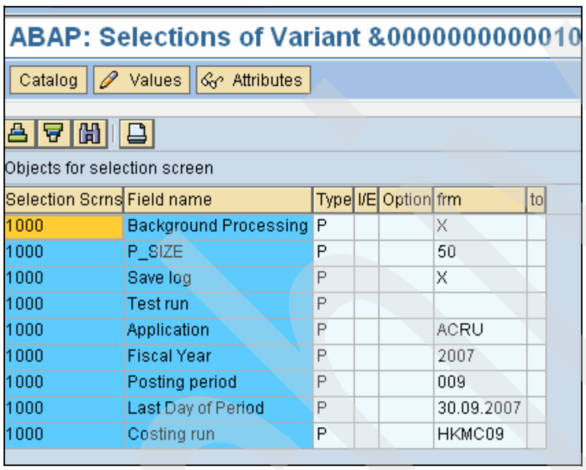

Selection Scrns	Field name	Type	I/E	Option	frm	to
1000	Background Processing	P			X	
1000	P_SIZE	P			50	
1000	Save log	P			X	
1000	Test run	P				
1000	Application	P			ACRU	
1000	Fiscal Year	P			2007	
1000	Posting period	P			009	
1000	Last Day of Period	P			30.09.2007	
1000	Costing run	P			HKMC09	

Figure 7-4 Parameter 002

Parameter 003: Single-level price determination

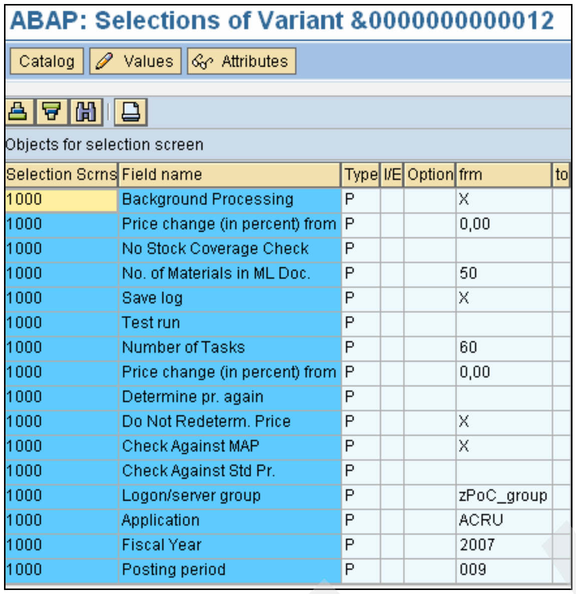

ABAP: Selections of Variant &0000000000012

Objects for selection screen

Selection Scrns	Field name	Type	I/E	Option	frm	to
1000	Background Processing	P			X	
1000	Price change (in percent) from	P			0,00	
1000	No Stock Coverage Check	P				
1000	No. of Materials in ML Doc.	P			50	
1000	Save log	P			X	
1000	Test run	P				
1000	Number of Tasks	P			60	
1000	Price change (in percent) from	P			0,00	
1000	Determine pr. again	P				
1000	Do Not Redeterm. Price	P			X	
1000	Check Against MAP	P			X	
1000	Check Against Std Pr.	P				
1000	Logon/server group	P			zPoC_group	
1000	Application	P			ACRU	
1000	Fiscal Year	P			2007	
1000	Posting period	P			009	

Figure 7-5 Parameter 003

Parameter 004: Multilevel price determination

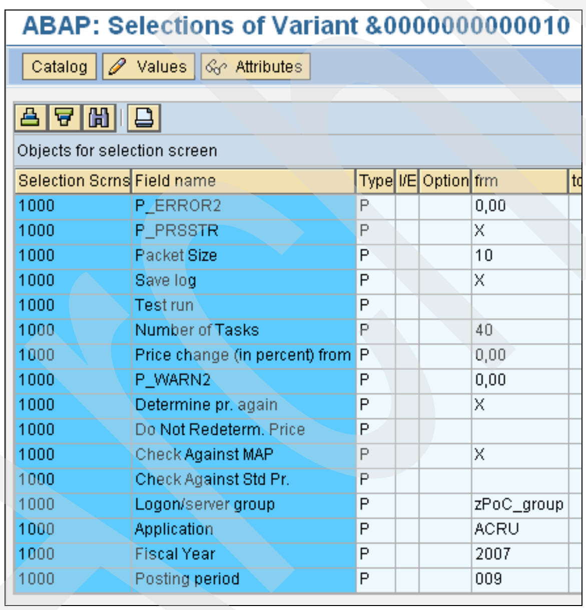

ABAP: Selections of Variant &0000000000010

Objects for selection screen

Selection Scrns	Field name	Type	I/E	Option	frm	to
1000	P_ERROR2	P			0,00	
1000	P_PRSSTR	P			X	
1000	Packet Size	P			10	
1000	Save log	P			X	
1000	Test run	P				
1000	Number of Tasks	P			40	
1000	Price change (in percent) from	P			0,00	
1000	P_WARN2	P			0,00	
1000	Determine pr. again	P			X	
1000	Do Not Redeterm. Price	P				
1000	Check Against MAP	P			X	
1000	Check Against Std Pr.	P				
1000	Logon/server group	P			zPoC_group	
1000	Application	P			ACRU	
1000	Fiscal Year	P			2007	
1000	Posting period	P			009	

Figure 7-6 Parameter 004

Parameter 005: Post closing

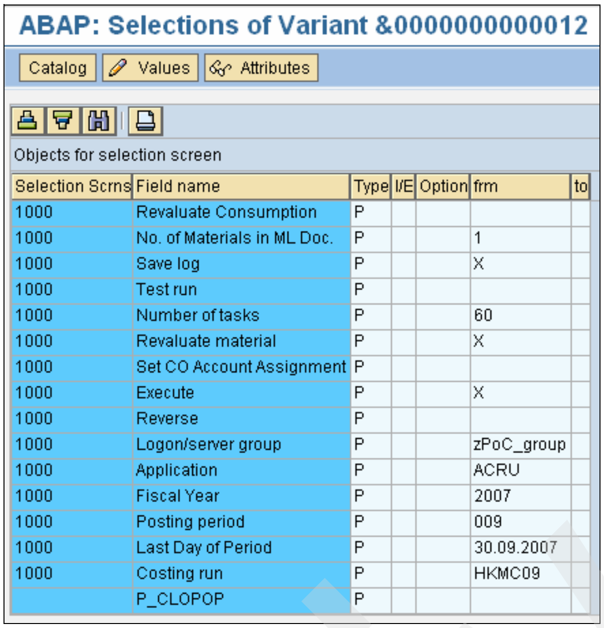

ABAP: Selections of Variant &0000000000012

Selection Scrns	Field name	Type	I/E	Option	frm	to
1000	Revaluate Consumption	P				
1000	No. of Materials in ML Doc.	P			1	
1000	Save log	P			X	
1000	Test run	P				
1000	Number of tasks	P			60	
1000	Revaluate material	P			X	
1000	Set CO Account Assignment	P				
1000	Execute	P			X	
1000	Reverse	P				
1000	Logon/server group	P			zPoC_group	
1000	Application	P			ACRU	
1000	Fiscal Year	P			2007	
1000	Posting period	P			009	
1000	Last Day of Period	P			30.09.2007	
1000	Costing run	P			HKMC09	
	P_CLOPOP	P				

Figure 7-7 Parameter 005

Parameter 006: Entry SPLIT-FI-DOC

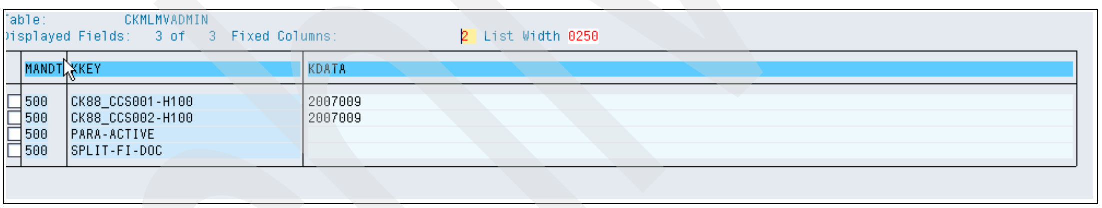

MANDT	XKEY	KDATA
500	CK88_CCS001-H100	2007009
500	CK88_CCS002-H100	2007009
500	PARA-ACTIVE	
500	SPLIT-FI-DOC	

Figure 7-8 Parameter 006

Parameter 007: New index of table CKMLV003

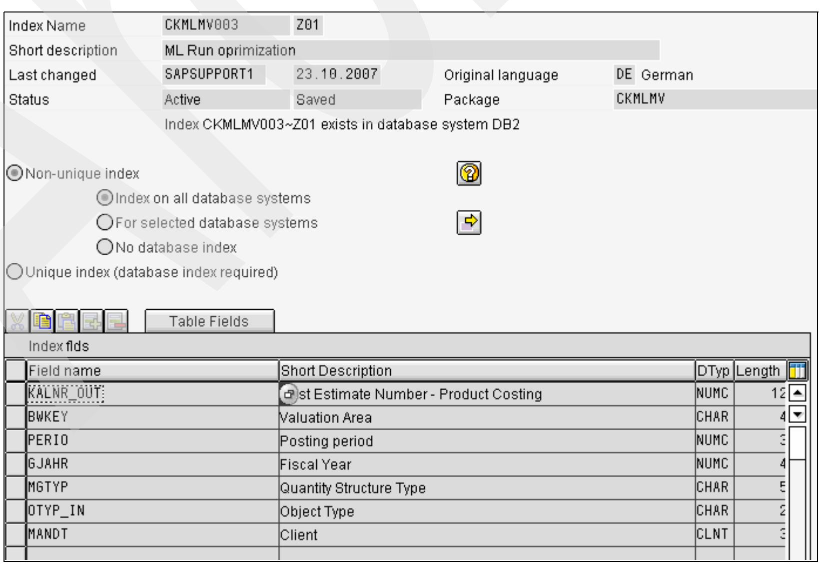

Field name	Short Description	DTyp	Length
KALNR_OUT	Cost Estimate Number - Product Costing	NUMC	12
BWKEY	Valuation Area	CHAR	4
PERIO	Posting period	NUMC	3
GJAHR	Fiscal Year	NUMC	4
MGTYP	Quantity Structure Type	CHAR	5
OTYP_IN	Object Type	CHAR	2
MANDT	Client	CLNT	3

Figure 7-9 Parameter 007

Parameter 008: Change coding, note 51789

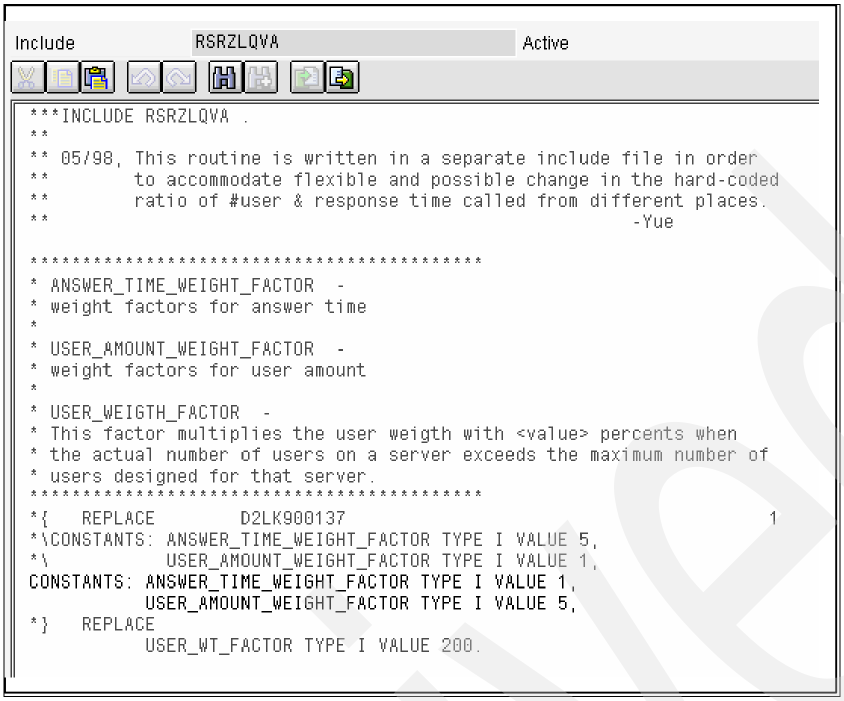

```
Include          RSRZLQVA                        Active

***INCLUDE RSRZLQVA .
**
** 05/98, This routine is written in a separate include file in order
**        to accommodate flexible and possible change in the hard-coded
**        ratio of #user & response time called from different places.
**                                                              -Yue

********************************************
* ANSWER_TIME_WEIGHT_FACTOR  -
* weight factors for answer time
*
* USER_AMOUNT_WEIGHT_FACTOR  -
* weight factors for user amount
*
* USER_WEIGTH_FACTOR  -
* This factor multiplies the user weigth with <value> percents when
* the actual number of users on a server exceeds the maximum number of
* users designed for that server.
********************************************
*{   REPLACE        D2LK900137                                          1
*\CONSTANTS: ANSWER_TIME_WEIGHT_FACTOR TYPE I VALUE 5,
*\           USER_AMOUNT_WEIGHT_FACTOR TYPE I VALUE 1,
CONSTANTS: ANSWER_TIME_WEIGHT_FACTOR TYPE I VALUE 1,
           USER_AMOUNT_WEIGHT_FACTOR TYPE I VALUE 5,
*}   REPLACE
           USER_WT_FACTOR TYPE I VALUE 200.
```

Figure 7-10   Parameter 008

Parameter 009: Change RFC server group

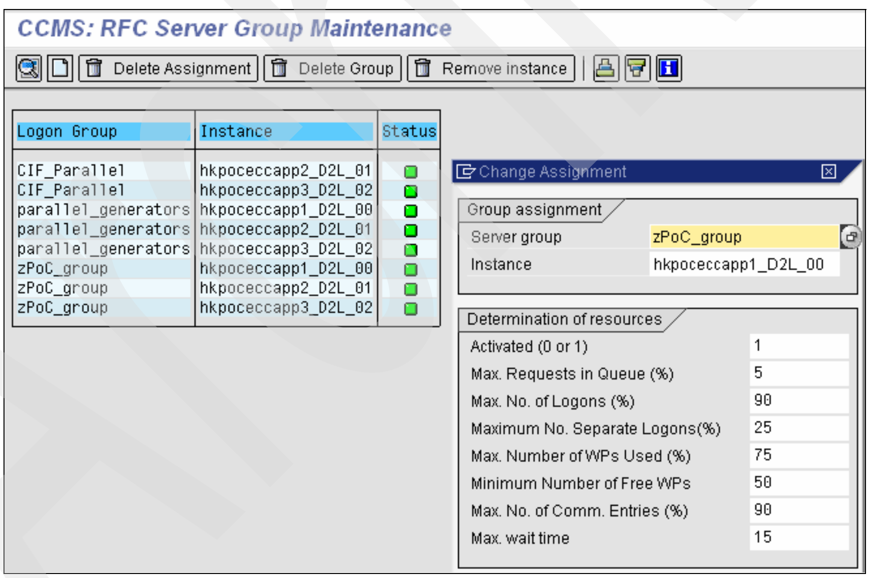

Figure 7-11   Parameter 009

Parameter 010: Update DB2 catalog

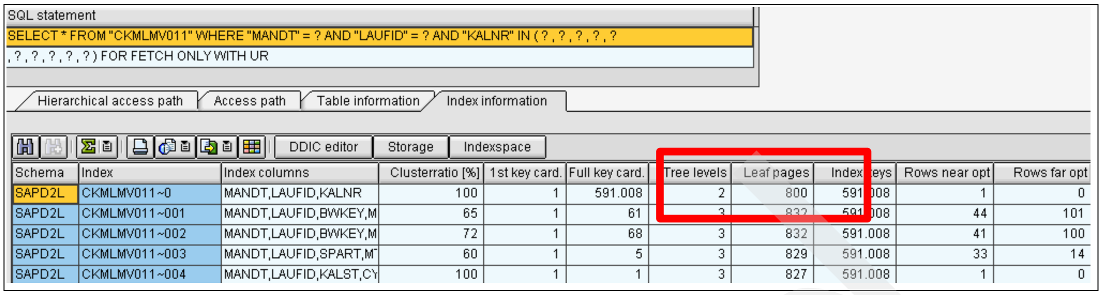

Figure 7-12 Parameter 010

Parameter 011: Change SMLG group zPoC_ECC

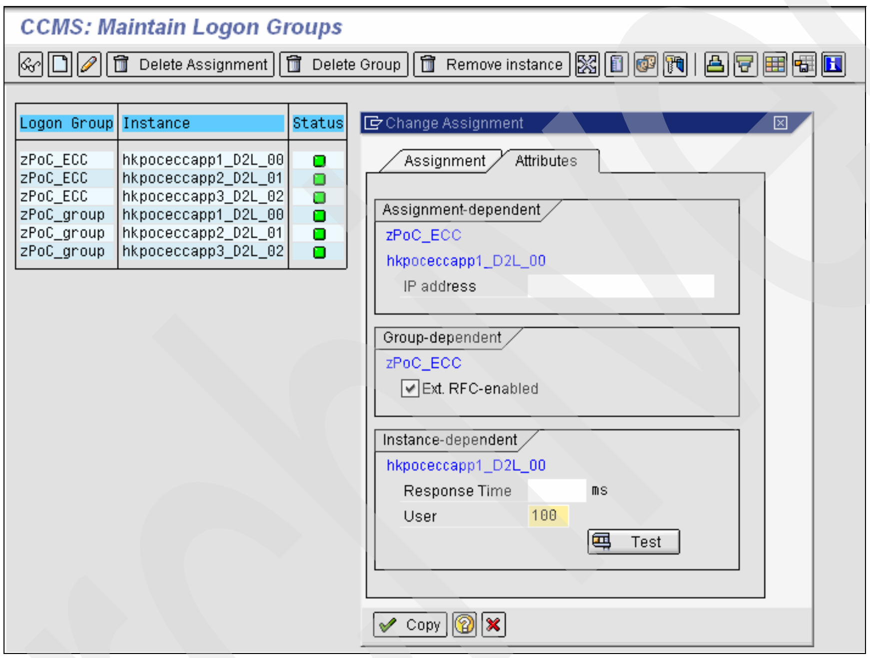

Figure 7-13 Parameter 011

7.3 Results overview

The next five figures show the first ML calculation results and the overview of run 10 when ML was executed along with Loadrunner.

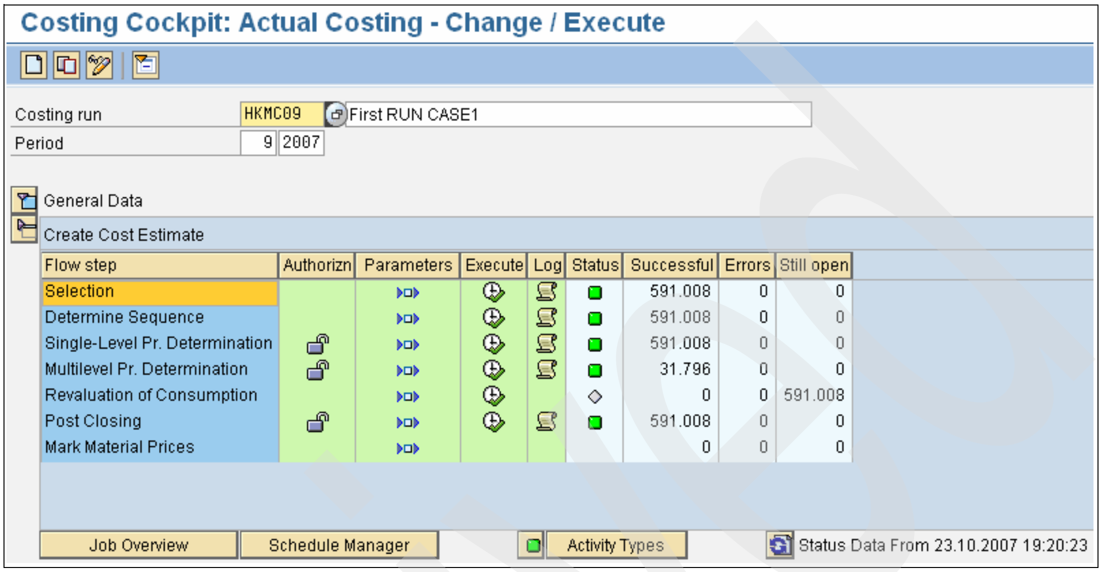

Figure 7-14 ML calculation result

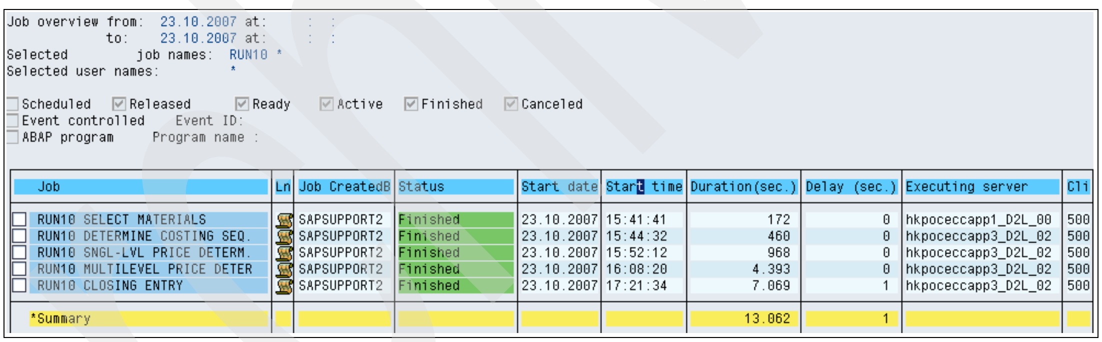

Figure 7-15 Overview of run10 (1)

The run time was about 3.6 hours, so the KPI of 4 hours was achieved even with the load runner. Compared to run 09, the run time was increased by about 0.3 hours.

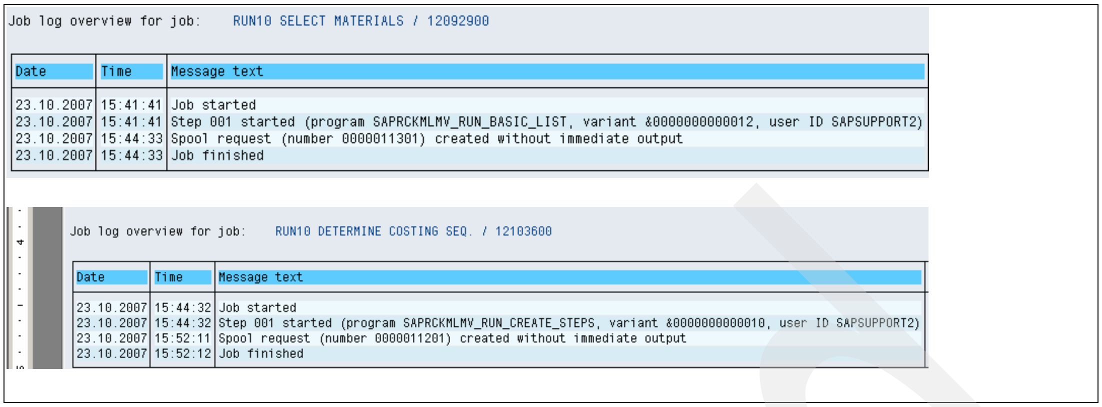

Job log overview for job: RUN10 SELECT MATERIALS / 12092900

Date	Time	Message text
23.10.2007	15:41:41	Job started
23.10.2007	15:41:41	Step 001 started (program SAPRCKMLMV_RUN_BASIC_LIST, variant &0000000000012, user ID SAPSUPPORT2)
23.10.2007	15:44:33	Spool request (number 0000011301) created without immediate output
23.10.2007	15:44:33	Job finished

Job log overview for job: RUN10 DETERMINE COSTING SEQ. / 12103600

Date	Time	Message text
23.10.2007	15:44:32	Job started
23.10.2007	15:44:32	Step 001 started (program SAPRCKMLMV_RUN_CREATE_STEPS, variant &0000000000010, user ID SAPSUPPORT2)
23.10.2007	15:52:11	Spool request (number 0000011201) created without immediate output
23.10.2007	15:52:12	Job finished

Figure 7-16 Log overview of run 10 (2)

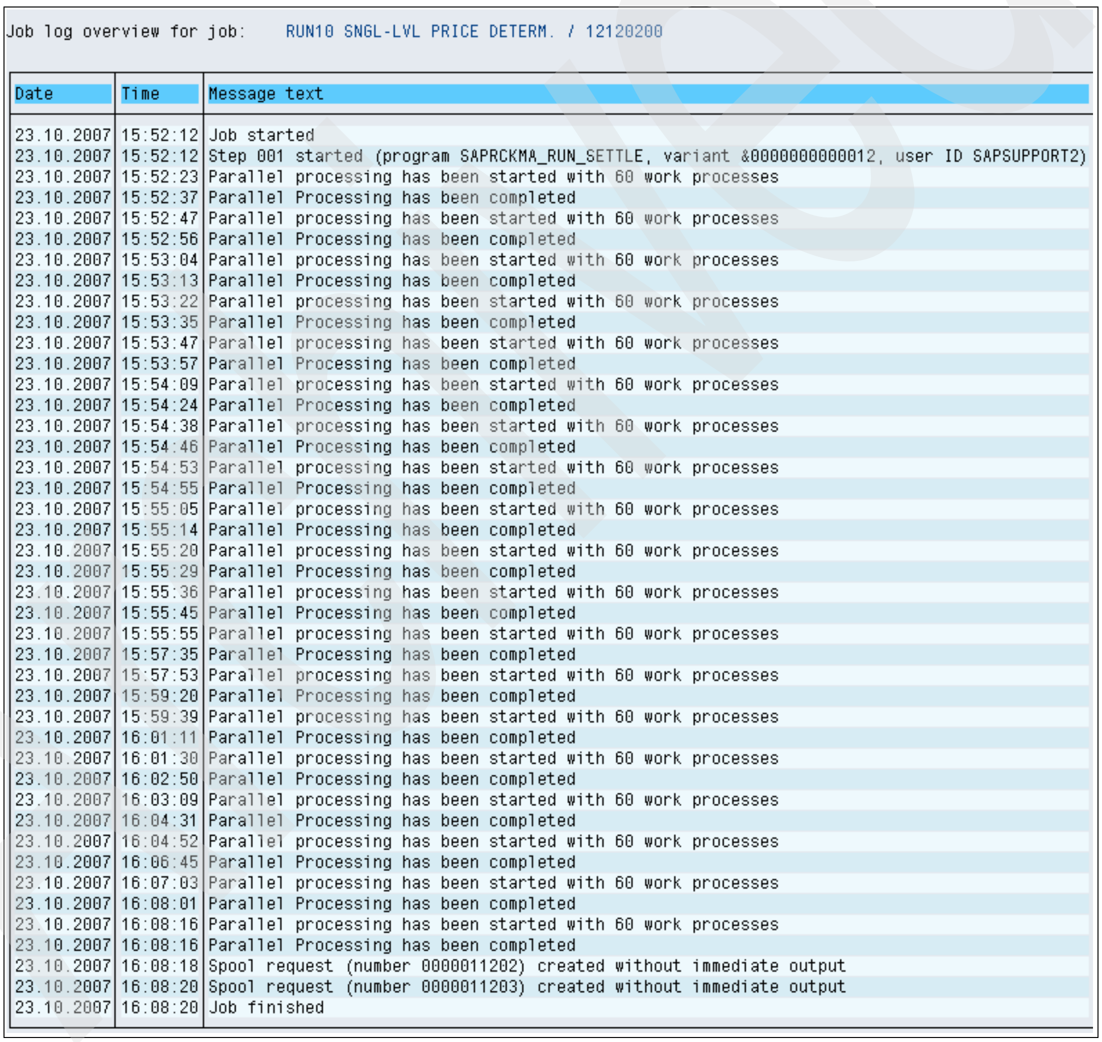

Job log overview for job: RUN10 SNGL-LVL PRICE DETERM. / 12120200

Date	Time	Message text
23.10.2007	15:52:12	Job started
23.10.2007	15:52:12	Step 001 started (program SAPRCKMA_RUN_SETTLE, variant &0000000000012, user ID SAPSUPPORT2)
23.10.2007	15:52:23	Parallel processing has been started with 60 work processes
23.10.2007	15:52:37	Parallel Processing has been completed
23.10.2007	15:52:47	Parallel processing has been started with 60 work processes
23.10.2007	15:52:56	Parallel Processing has been completed
23.10.2007	15:53:04	Parallel processing has been started with 60 work processes
23.10.2007	15:53:13	Parallel Processing has been completed
23.10.2007	15:53:22	Parallel processing has been started with 60 work processes
23.10.2007	15:53:35	Parallel Processing has been completed
23.10.2007	15:53:47	Parallel processing has been started with 60 work processes
23.10.2007	15:53:57	Parallel Processing has been completed
23.10.2007	15:54:09	Parallel processing has been started with 60 work processes
23.10.2007	15:54:24	Parallel Processing has been completed
23.10.2007	15:54:38	Parallel processing has been started with 60 work processes
23.10.2007	15:54:46	Parallel Processing has been completed
23.10.2007	15:54:53	Parallel processing has been started with 60 work processes
23.10.2007	15:54:55	Parallel Processing has been completed
23.10.2007	15:55:05	Parallel processing has been started with 60 work processes
23.10.2007	15:55:14	Parallel Processing has been completed
23.10.2007	15:55:20	Parallel processing has been started with 60 work processes
23.10.2007	15:55:29	Parallel Processing has been completed
23.10.2007	15:55:36	Parallel processing has been started with 60 work processes
23.10.2007	15:55:45	Parallel Processing has been completed
23.10.2007	15:55:55	Parallel processing has been started with 60 work processes
23.10.2007	15:57:35	Parallel Processing has been completed
23.10.2007	15:57:53	Parallel processing has been started with 60 work processes
23.10.2007	15:59:20	Parallel Processing has been completed
23.10.2007	15:59:39	Parallel processing has been started with 60 work processes
23.10.2007	16:01:11	Parallel Processing has been completed
23.10.2007	16:01:30	Parallel processing has been started with 60 work processes
23.10.2007	16:02:50	Parallel Processing has been completed
23.10.2007	16:03:09	Parallel processing has been started with 60 work processes
23.10.2007	16:04:31	Parallel Processing has been completed
23.10.2007	16:04:52	Parallel processing has been started with 60 work processes
23.10.2007	16:06:45	Parallel Processing has been completed
23.10.2007	16:07:03	Parallel processing has been started with 60 work processes
23.10.2007	16:08:01	Parallel Processing has been completed
23.10.2007	16:08:16	Parallel processing has been started with 60 work processes
23.10.2007	16:08:16	Parallel Processing has been completed
23.10.2007	16:08:18	Spool request (number 0000011202) created without immediate output
23.10.2007	16:08:20	Spool request (number 0000011203) created without immediate output
23.10.2007	16:08:20	Job finished

Figure 7-17 Log overview of run 10 (3)

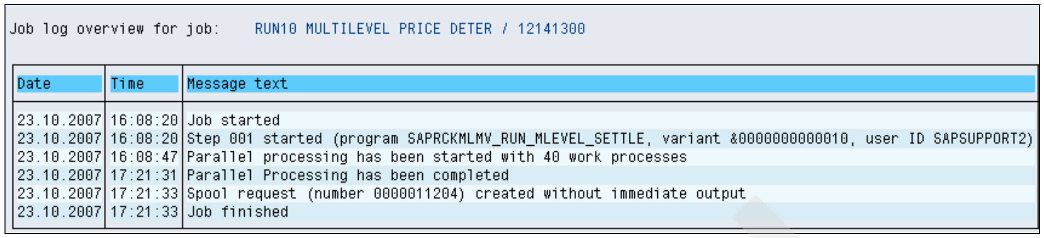

Job log overview for job: RUN10 MULTILEVEL PRICE DETER / 12141300

Date	Time	Message text
23.10.2007	16:08:20	Job started
23.10.2007	16:08:20	Step 001 started (program SAPRCKMLMV_RUN_MLEVEL_SETTLE, variant &0000000000010, user ID SAPSUPPORT2)
23.10.2007	16:08:47	Parallel processing has been started with 40 work processes
23.10.2007	17:21:31	Parallel Processing has been completed
23.10.2007	17:21:33	Spool request (number 0000011204) created without immediate output
23.10.2007	17:21:33	Job finished

Figure 7-18 Log overview of run 10 (4)

Spool

Figure 7-19 shows a spool request in system D2L.

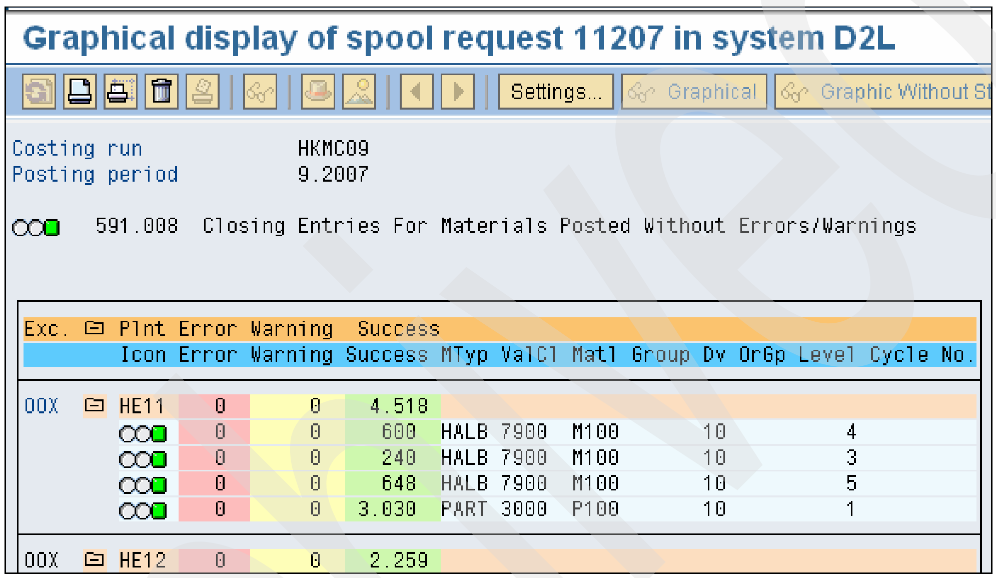

Graphical display of spool request 11207 in system D2L

Settings... Graphical Graphic Without St

Costing run HKMC09
Posting period 9.2007

591.008 Closing Entries For Materials Posted Without Errors/Warnings

Exc.	Plnt	Error	Warning	Success						
	Icon	Error	Warning	Success	MTyp	ValCl	Matl Group	Dv	OrGp Level	Cycle No.
OOX	HE11	0	0	4.518						
		0	0	600	HALB	7900	M100		10	4
		0	0	240	HALB	7900	M100		10	3
		0	0	648	HALB	7900	M100		10	5
		0	0	3.030	PART	3000	P100		10	1
OOX	HE12	0	0	2.259						

Figure 7-19 Spool

7.3.1 Observations

All steps ran without errors and within the KPI requirement of four hours. It should be noted that during the run the load runner was also running.

Posted documents are shown in Figure 7-20.

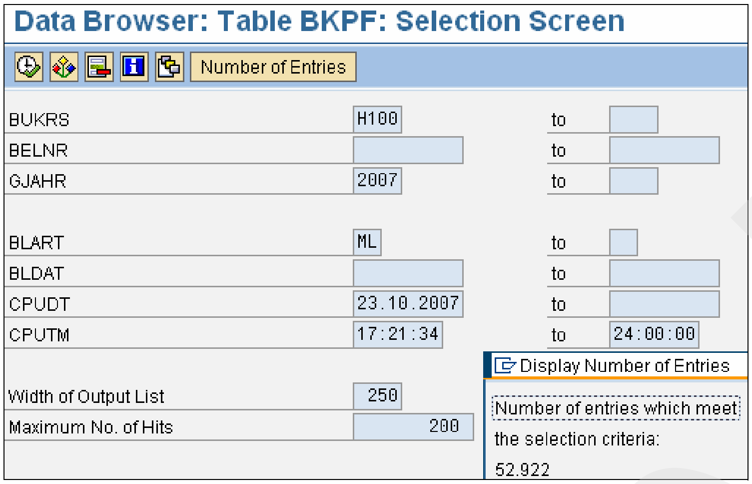

Figure 7-20 Table BKPF

In this run 1600 more documents were posted than in run 09 because of the change in package size for ROH and PART from 500 to 250.

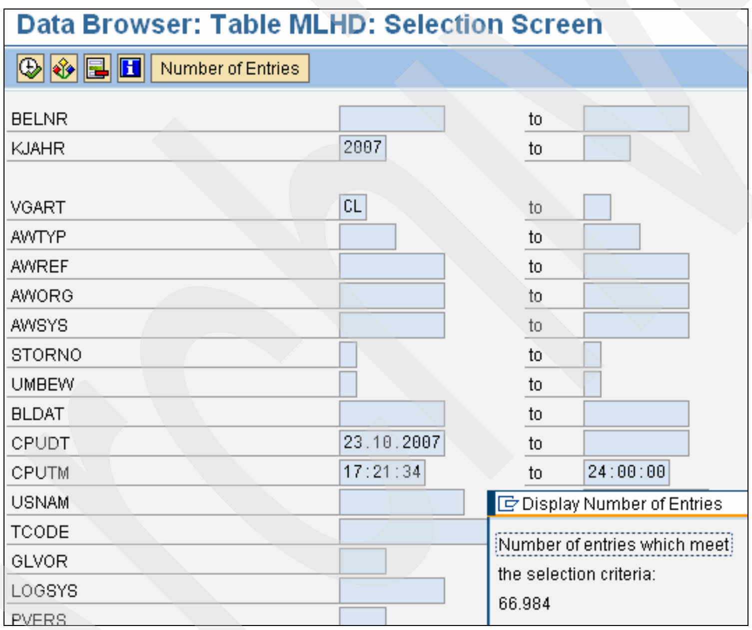

Figure 7-21 Table MLHD

7.4 Analysis of System p infrastructure layer test results

We first present the system activity summary for the five System p server LPARs during the run. This run began at 15:41 and finished at 19:19, lasting 217 minutes.

7.4.1 LPAR activity summary

In this section we discuss the system activity summary based upon the nmon results presented in the following five graphs (Figure 7-22 through Figure 7-26).

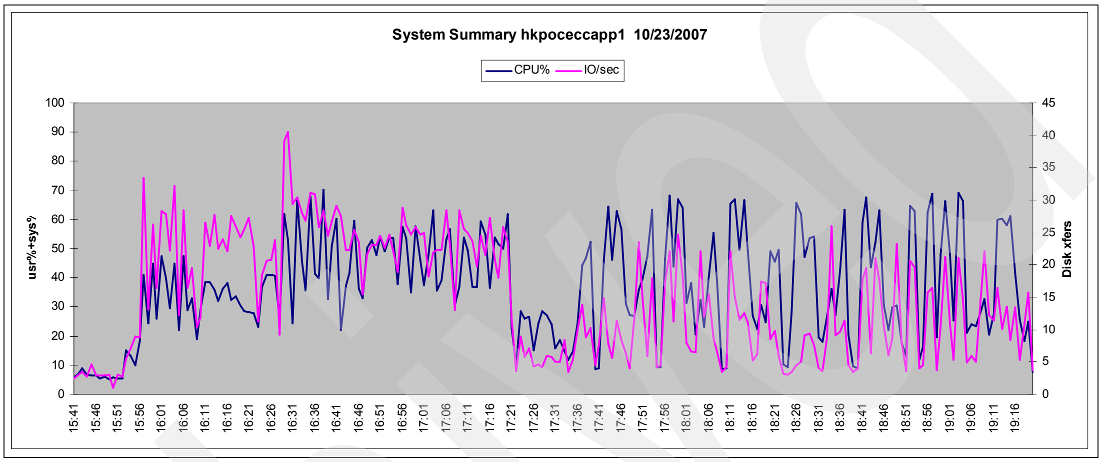

Figure 7-22 Activity summary for hkpoceccapp1

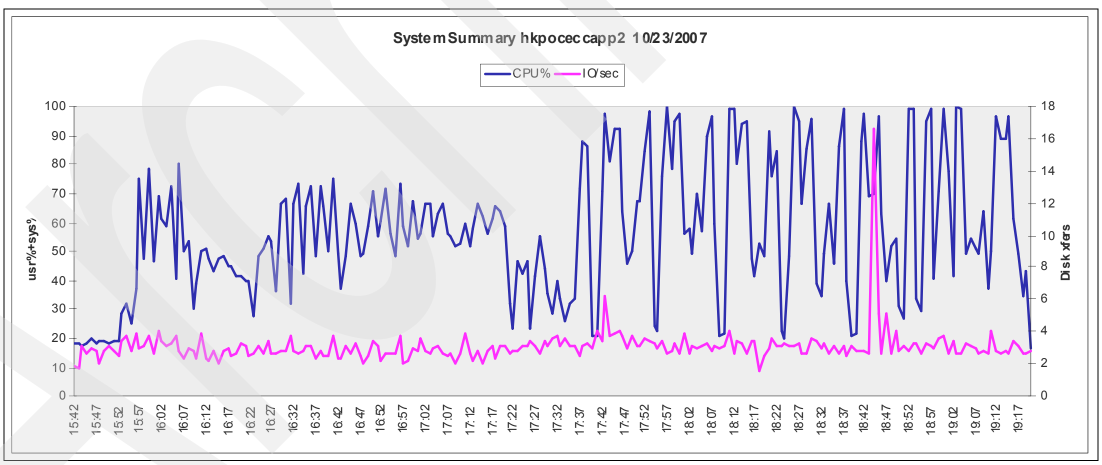

Figure 7-23 Activity summary for hkpoceccapp2

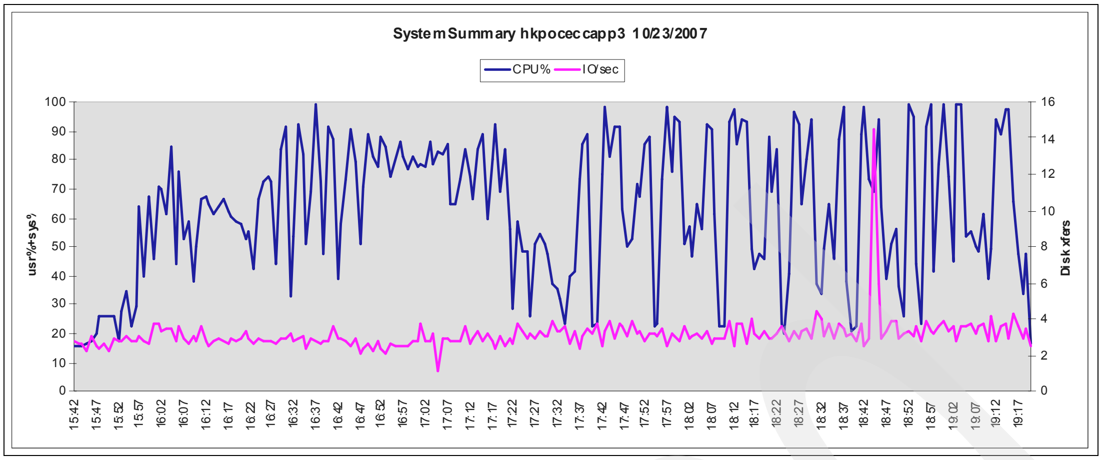

Figure 7-24 Activity summary for hkpoceccapp3

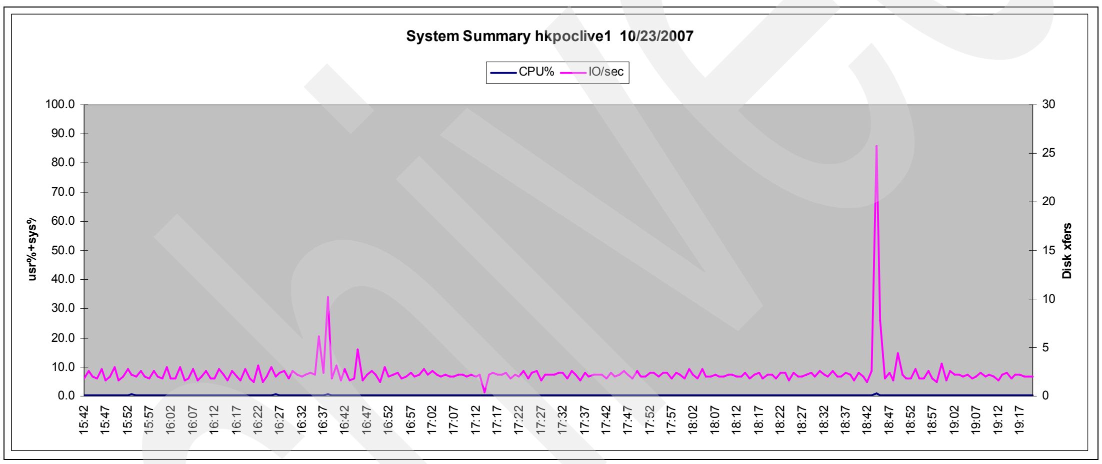

Figure 7-25 Activity summary for hkpoclive1

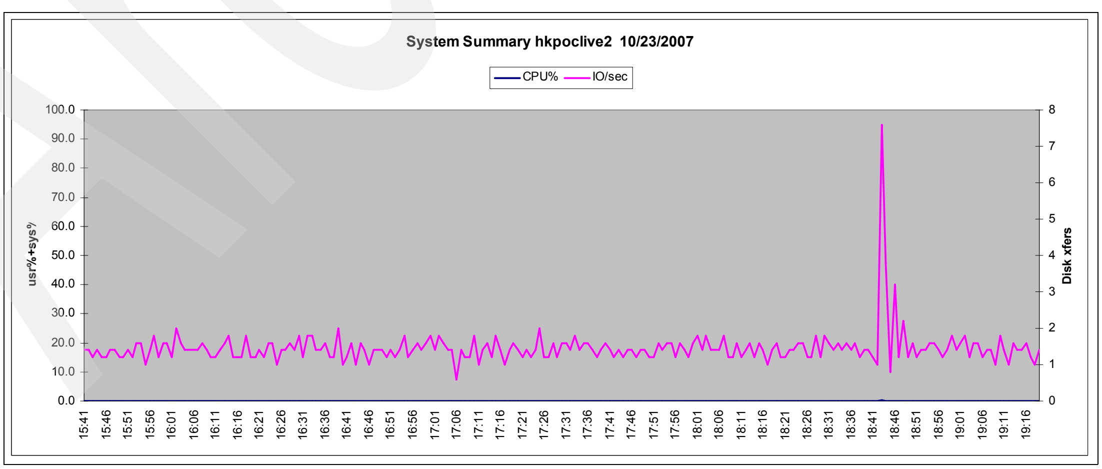

Figure 7-26 Activity summary for hkpoclive2

Regarding the ECC system LPARs, we see that the ECC system has very significant activity on this run. Hkpoceccapp1 has an average usage of 37%, with peaks of 70%. Disk I/Os are very low on this LPAR and are not significant.

Hkpoceccapp2 and hkpoceccapp3 exhibit similar behavior during the run. Hkpoceccapp2 has an average CPU consumption of 57%, with peaks of 99.9%. Nevertheless, we can see that there is almost no CPU wait time, and I/Os per second are not significant. Hkpoceccapp3 has an average CPU consumption of 62.6%, with peaks of 99.2%. However, once again, we can see that there is almost no CPU wait time, and I/Os per second are not significant. Therefore, on these two LPARs, there is a huge amount of activity but the systems seem to be performing perfectly.

Regarding the SCM instance, Figure 7-26 shows an average CPU and disk usage close to 0, so we conclude that this instance has very little influence on the run, and therefore do not dig any deeper into details. We can conclude that this LPAR does not interfere in this run.

Finally, regarding the liveCache instance, we can see in Figure 7-25 on page 176 that the CPU and disk activity are very close to 0, so again we can conclude that this LPAR does not interfere in this run.

> **Note:** These summary graphs do not consider CPU wait percentage.

7.4.2 CPU details

Global CPU study

In the next three graphs we present the detailed CPU activity on ECC system LPARs.

As shown in the Figure 7-27, it is interesting to note that there is absolutely no wait time during the run, so the application is taking perfect advantage of the available resources.

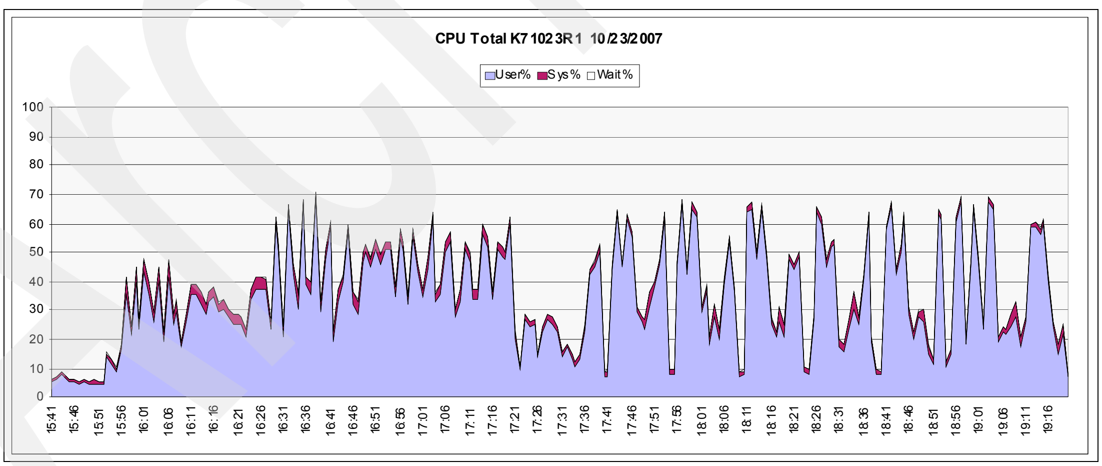

Figure 7-27 hkpoceccapp1 CPU details

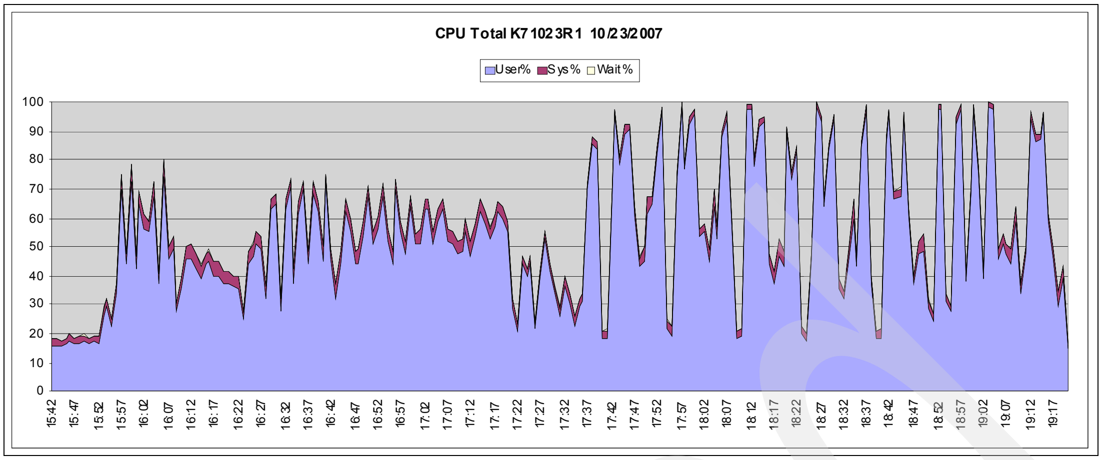

Figure 7-28 hkpoceccapp2 CPU details

On LPAR hkpoceccapp2, no unexpected behavior is seen. CPU usage peaks at 99.9% and there are no wait times. The usage peaks are not sustained, which is probably due to the behavior of the application itself.

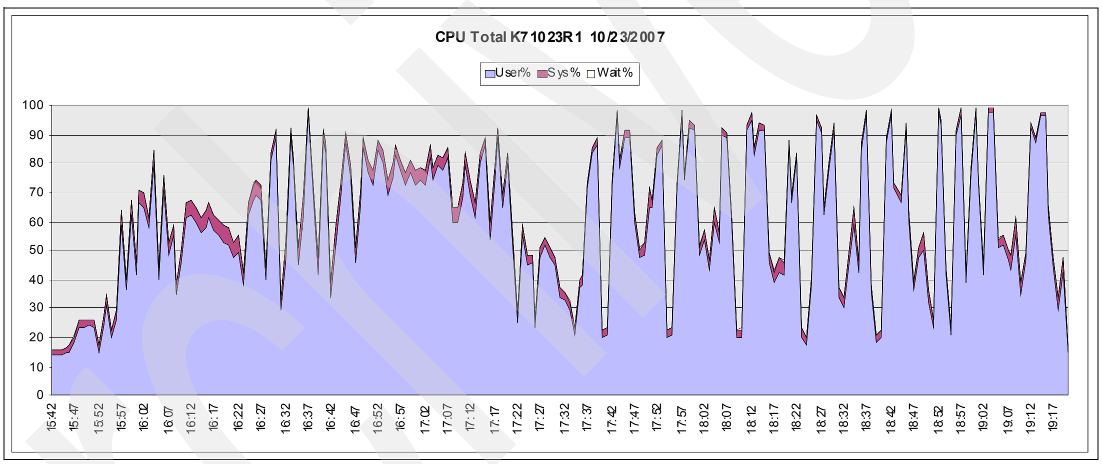

Figure 7-29 hkpoceccapp3 CPU details

Regarding LPAR hkpoceccapp3, the behavior is very similar to that shown in hkpoceccapp2 during the later part of the run (from 17:22 until the end) as you can see in Figure 7-29. CPU usage peaks at 99% and there is no wait time. During the first part of the run (from the beginning until 17:22), we can see that the CPU usage on this LPAR is higher that on LPAR hkpoceccapp2, although it follows the same pattern. This is probably due to the behavior of the application itself.

Unitary CPU study

The next three graphs (Figure 7-30 through Figure 7-32) show the unitary CPU and thread usage.

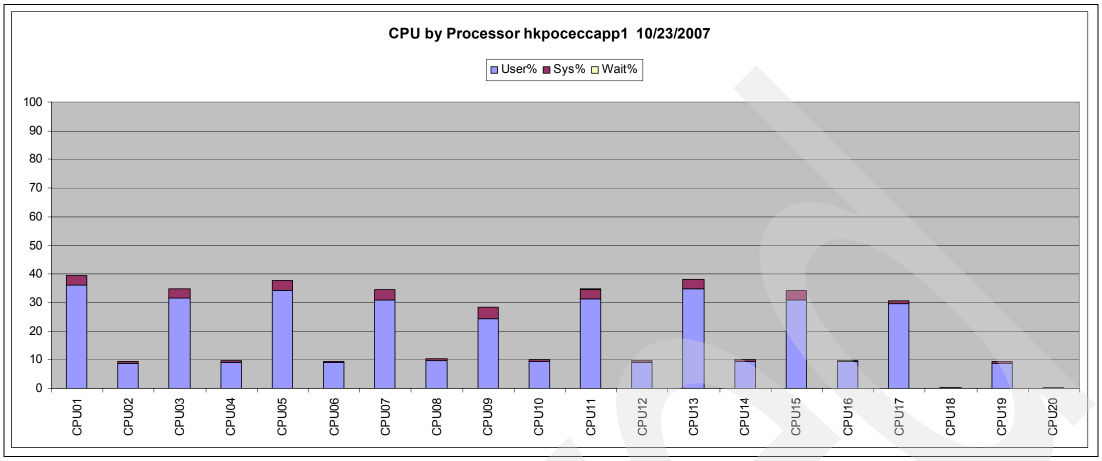

Figure 7-30 hkpoceccapp1 CPU by processor

As shown in Figure 7-30, the hkpoceccapp1 CPU usage by processor is balanced. We notice that CPU19 is less used on this LPAR. This figure also shows that SMT is used for almost all CPUs, so physical CPUs are not sufficient to balance the load. As we can see, the CPU is absolutely not overloaded in general on this LPAR, although there is a significant load.

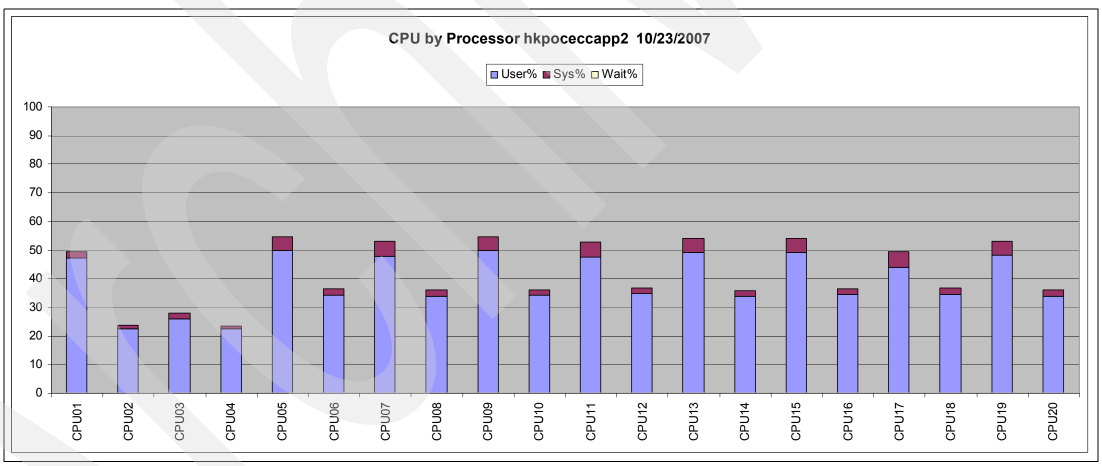

Figure 7-31 hkpoceccapp2 CPU by processor

Regarding LPAR hkpoceccapp2, Figure 7-31 shows an almost homogeneous load distribution between the CPUs (only CPU03 has a different behavior), but SMT is used for all CPUs, so the application is using all benefit from parallelization and the load consumes almost all available CPU resources. These same comments apply to hkpoceccapp3, which has a very similar behavior, as shown on Figure 7-32 on page 180.

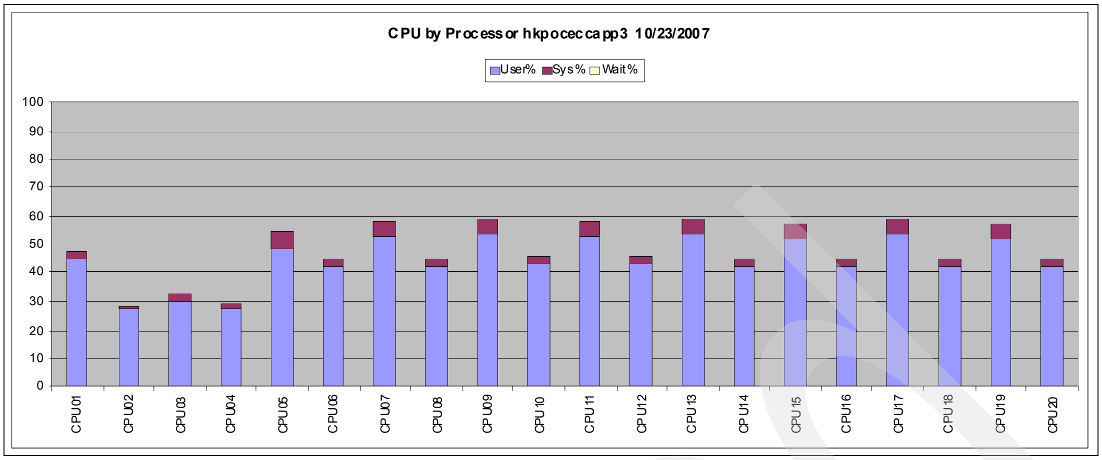

Figure 7-32 hkpoceccapp3 CPU by processor

In this case (hkpoceccapp2 and hkpoceccapp3 LPARs), increasing the number of CPUs might improve the performance of the run.

7.4.3 Memory details

Figure 7-33 shows the memory details at the LPAR level.

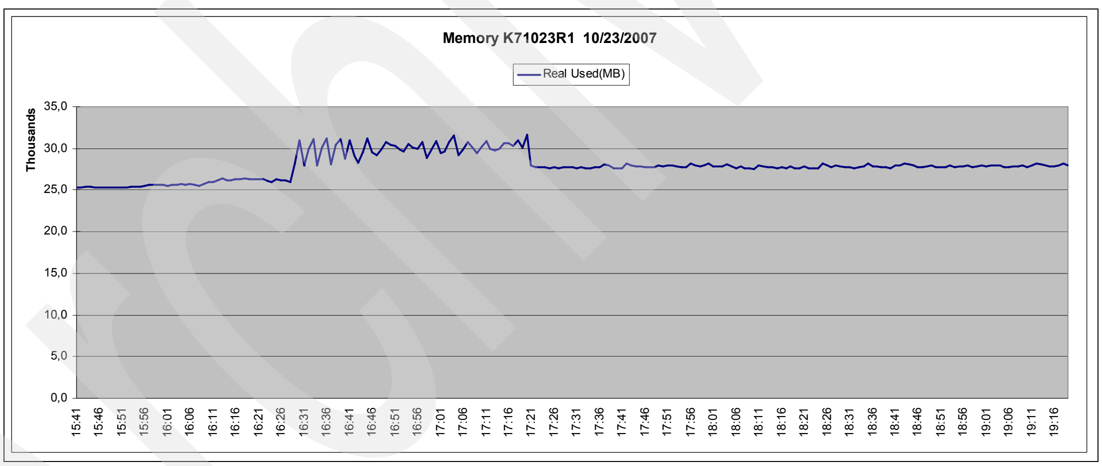

Figure 7-33 hkpoceccapp1 used memory

The three LPARs hkpoceccapp1, hkpoceccapp2, and hkpoceccapp3 exhibit similar behavior regarding memory; Figure 7-33 shows hkpoceccapp1 memory behavior as an example. At the time of the run, each LPAR had 60 GB memory. We observed that no LPAR consumed more than 25 GB during the run, and no LPAR had paging during the run. So, regarding memory for run 10, everything was fine.

7.4.4 Disk details

In this case there was no disk activity in any system.

7.4.5 Network details

In this section we examine the TCP/IP network usage profile.

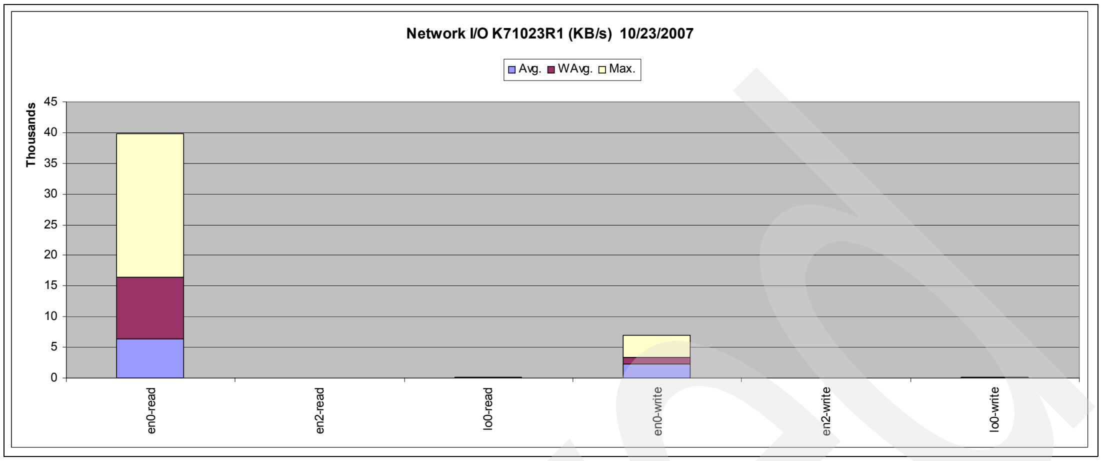

Figure 7-34 hkpoceccapp1 network I/O per adapter

Figure 7-34 shows network load by adapter. We notice that the highest load is on en0 adapters, which is logical because these adapters are connected to the 1 Gbps network for server communications.

The three LPARs exhibit similar behavior, which is logical because they run the same application, mainly reading from the database and writing some data. Figure 7-34 shows hkpoceccapp1 behavior.

The cumulative network activity for hkpoceccapp1 is shown in Figure 7-35. We saw similar behavior for hkpoceccapp2 and hkpoceccapp3.

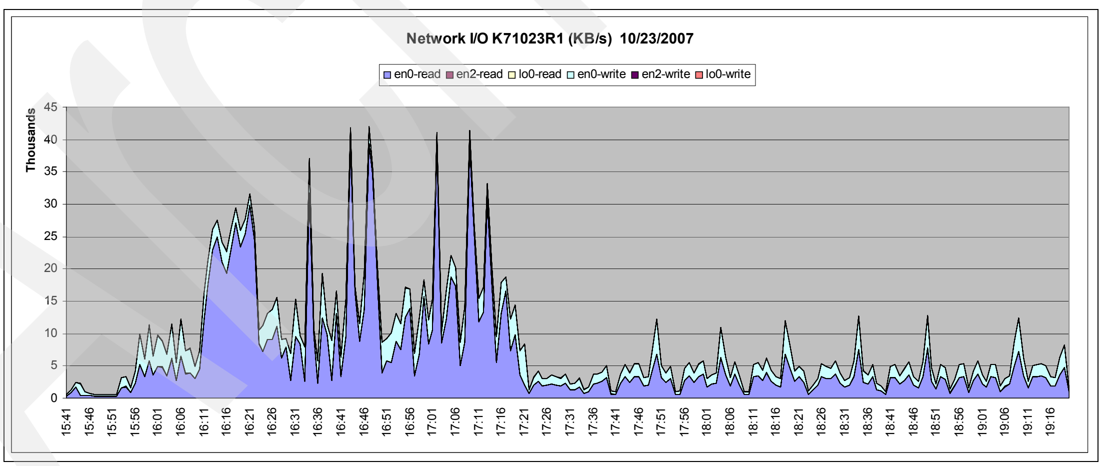

Figure 7-35 hkpoceccapp1 network I/O

We saw no contention during the run, so we can conclude that the network is well sized. It is pertinent to mention that we did not experience any network errors at all, and the servers behaved the same way.

7.4.6 Top processes

In this section we analyze the top CPU and memory consumers.

Top CPU

As Figure 7-36 shows, hkpoceccapp1 top processes are disp+work (SAP processes) and db2sysc (DB2 connect, connection manager between DB2 for z/OS and SAP). The same patterns were exhibited for hkpoceccapp2 and hkpoceccapp3, which indicates that SAP necessary core processes are the main consumers of the CPU.

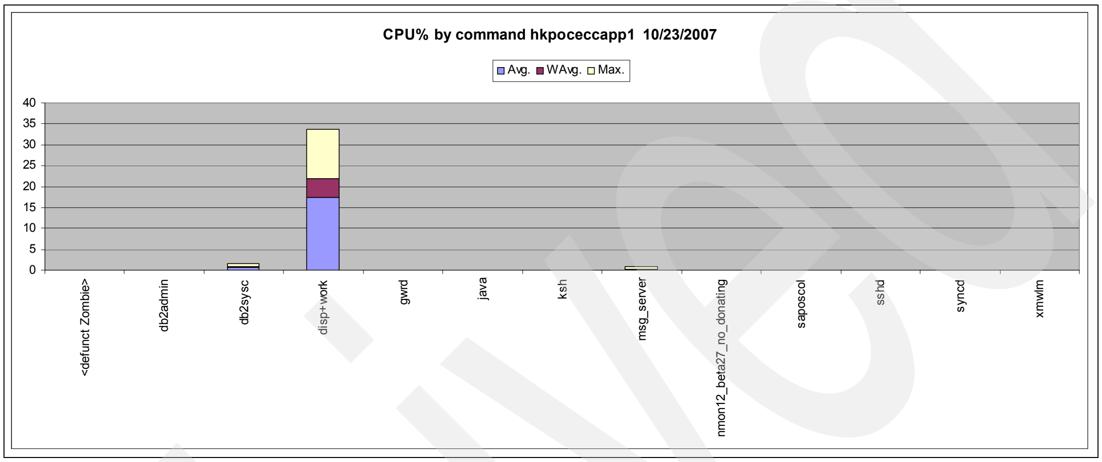

Figure 7-36 hkpoceccapp1 top CPU

Top memory

We found the same behavior regarding top memory consumers, meaning that disp+work SAP processes and db2sysc (DB2 connect, connection manager between DB2 on z and SAP) are the top memory consumers.

7.4.7 System p architecture conclusions

At the infrastructure level, System p servers responded perfectly to run 10 demands. Absolutely nothing had a negative impact on the run at this level. We notice that on LPARs hkpoceccapp2 and hkpoceccapp3, CPU usage peaks at 100%, so implementing more CPUs might improve performance. However, this test fulfilled the KPI requirements.

7.5 Analysis of System z infrastructure layer test results

This run began at 15:41 and finished at 19:19, lasting 217 minutes. Figure 7-37 shows the LPAR activity summary during the test interval.

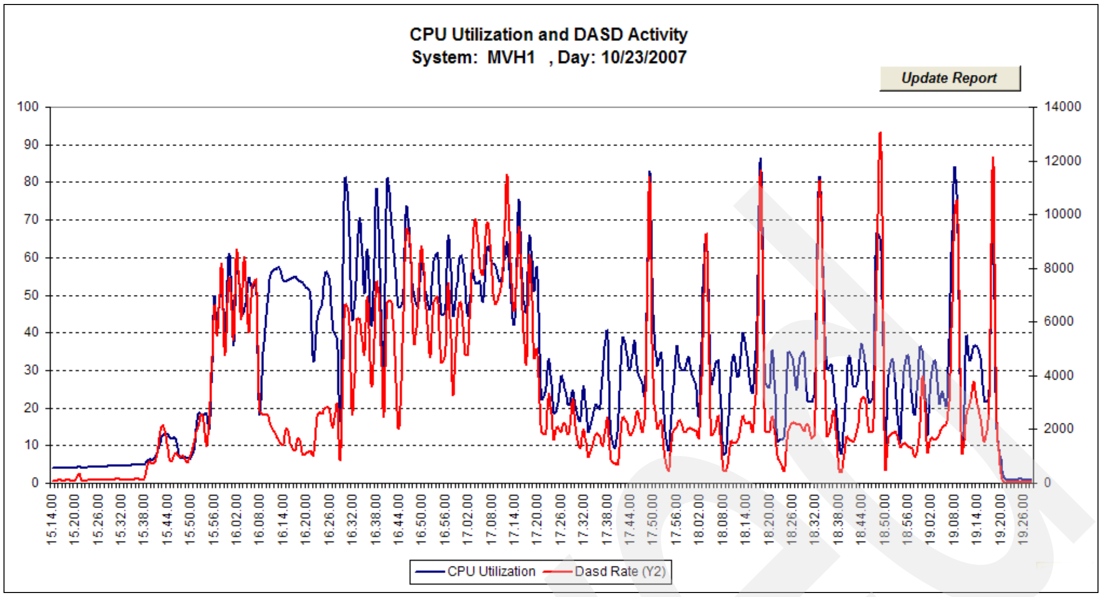

Figure 7-37 LPAR summary MVH1 ML run 10

During the test there was substantial CPU utilization and I/O activity.

In the following pages we will present more detailed reports on:

- – CPU use
- – OSA adapter
- – Channel path activity
- – Cache activity
- – DASD activity

7.5.1 CPU use

Figure 7-38 on page 184 shows CPU use during the test.

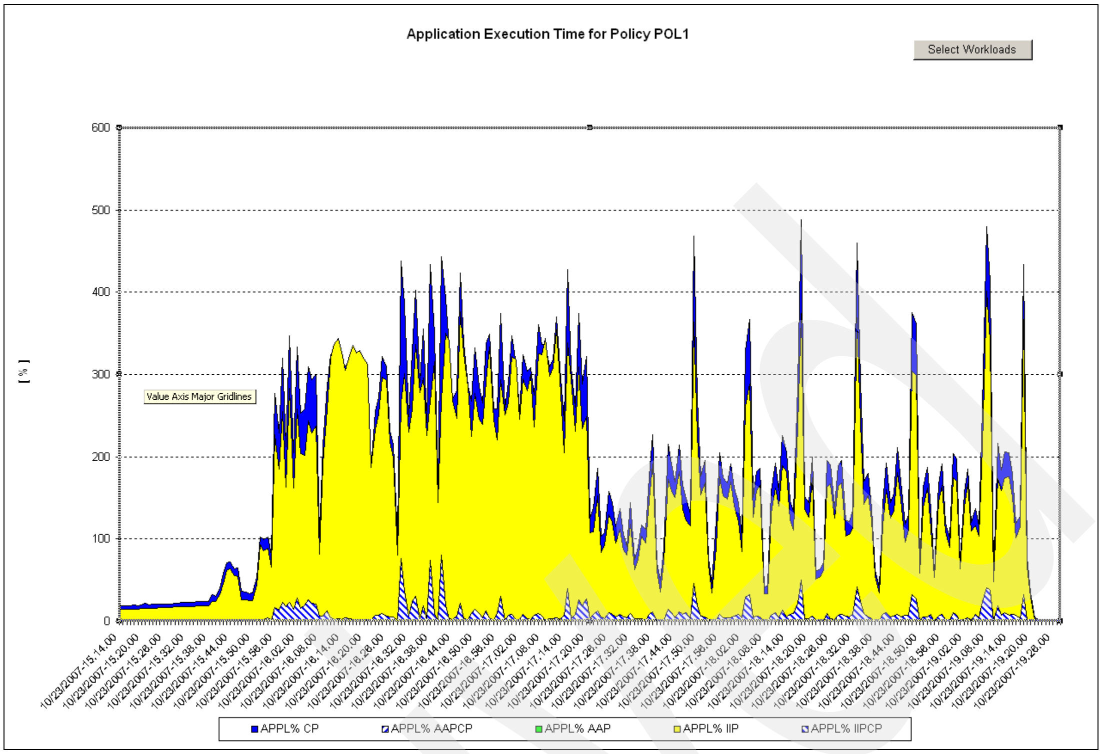

Figure 7-38 CP-zIIP activity MVH1 ML run 10

This graph shows the CP/zIIP activity during the test. Additionally, details are shown about load that ran on CPs but that could have run on zIIP. Obviously, at some peek moments all the available zIIPs were occupied. The usage of the z9 is substantial, but the box is capable of handling the load very well. Figure 7-39 shows additional information on the capture ratio. The capture ratio indicates how much of the used resources are spent on application processing. The graph shows that during the busy periods more than 90% of the used resources are spent on application processing. This is considered a good value.

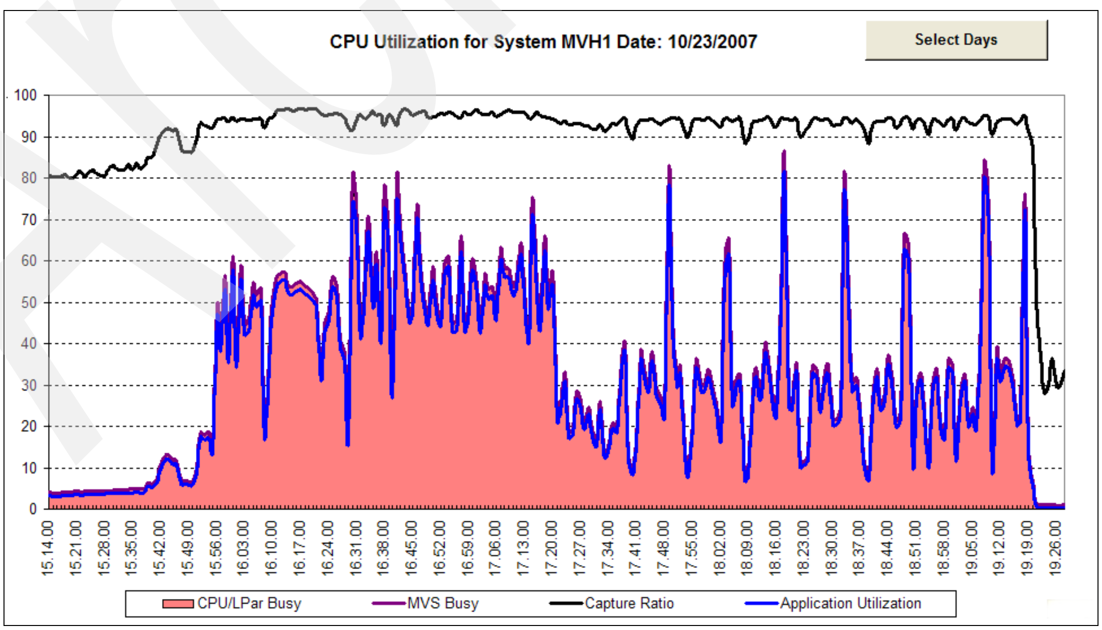

Figure 7-39 Capture ratio during MVH1 ML run 10

7.5.2 OSA adapter

The OSA adapter use shown in Figure 7-40 is for communicating to the SAP application servers. All the SQL requests come in through this adapter and all the results go out through this adapter.

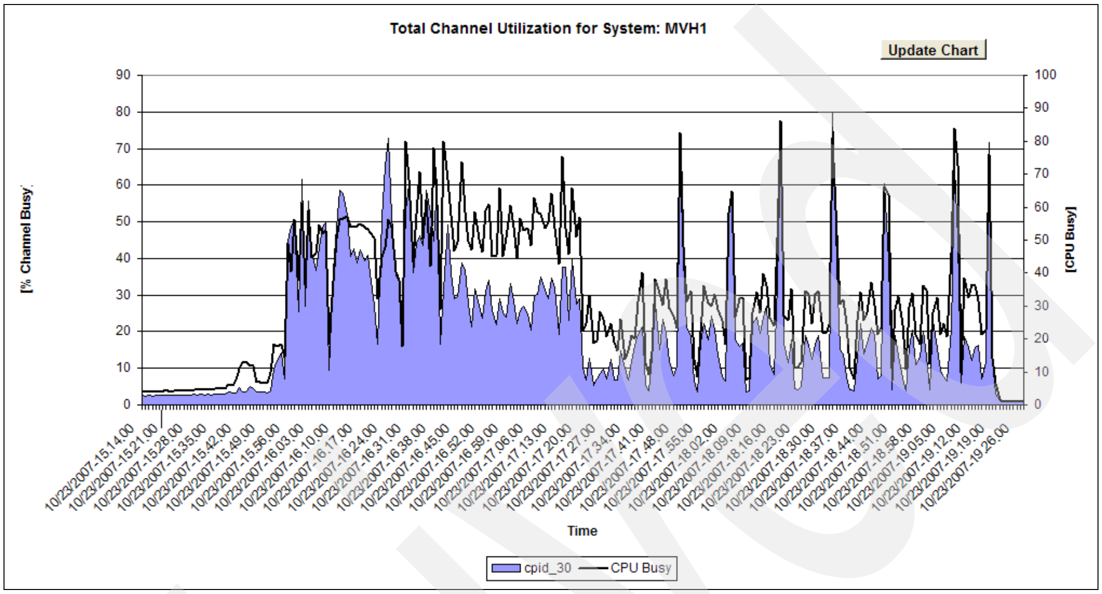

Figure 7-40 OSA adapter usage MVH1 ML run 10

The graph shows substantial activity on the OSA adapter. The adapter was capable of handling the load very well during the test. In a "normal" production situation there is always more than one OSA adapter configured for redundancy reasons. With TCP/IP on z/OS there are methods available to configure the adapters in such a way that they are used simultaneously and provide even more bandwidth.

7.5.3 Channel path activity

Figure 7-41 shows the activity on the FICON-Express2 channels attached to the DS8300 storage box used during the test.

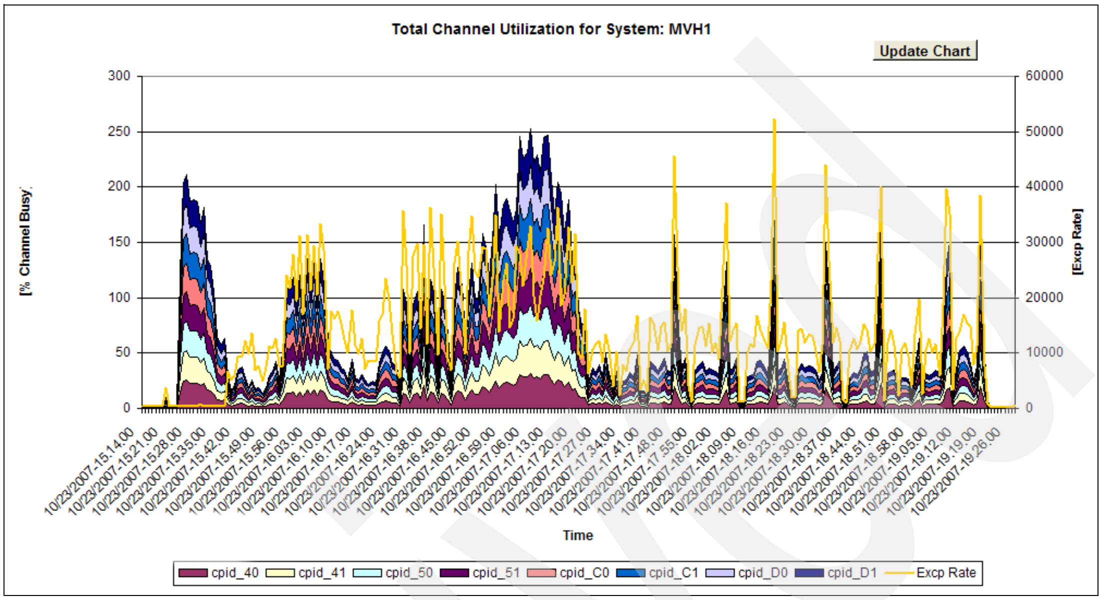

Figure 7-41 FICON Activity ML run 10

The FICON Express2 channels handle the load very well. The load is balanced over the available channels. At peek moments we see over 50000 Excp's; individual channel utilization is around 20%. The available FICON Express2 channels can easily handle much more additional load.

7.5.4 Cache activity

The next three figures show activity on the DASD cache of the DS8300 storage box used during the run. Figure 7-42 shows the I/O rate for all relevant control units.

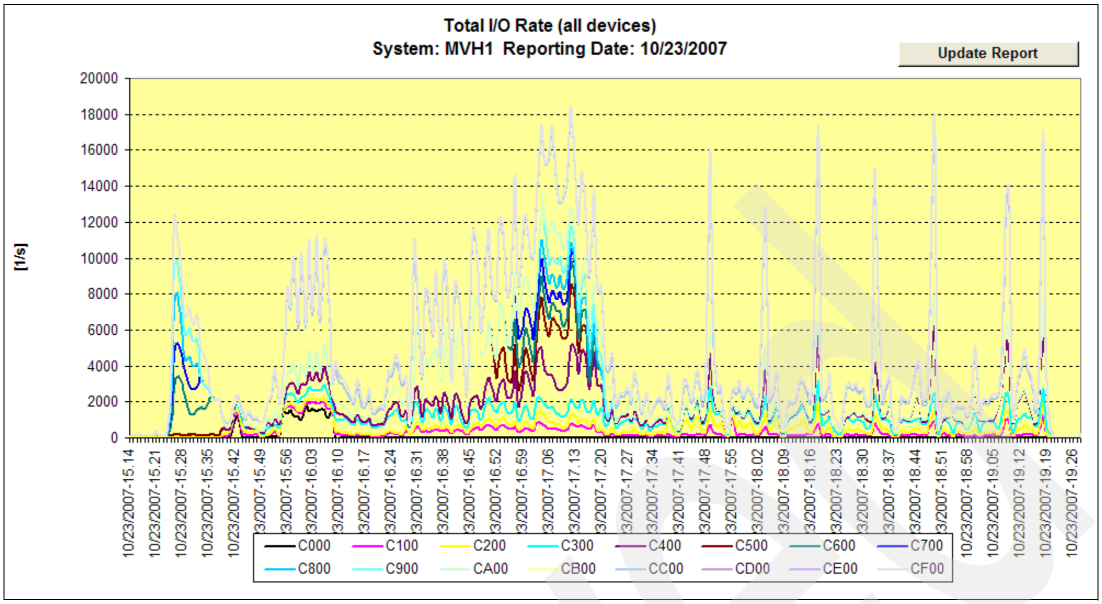

Figure 7-42 Total I/O rate ML run 10

Figure 7-43 shows the Cache Rate.

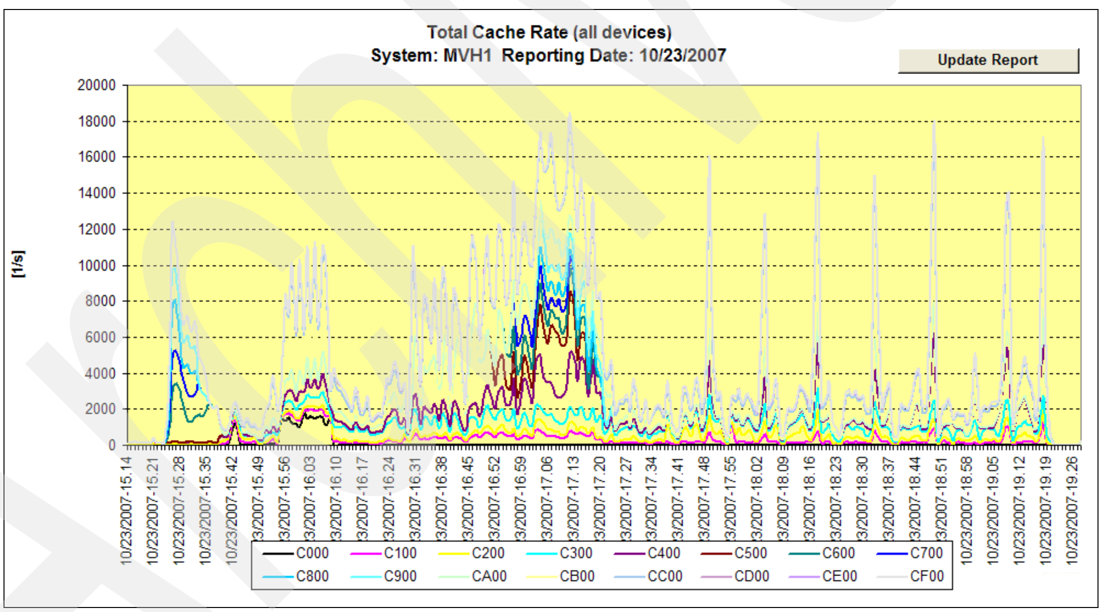

Figure 7-43 Total cache rate ML run 10

As you can see, both graphs are almost identical. This means that a high number of the I/Os are cache hits. Nevertheless, we took a closer look at the busiest control unit, CC00, as shown in Figure 7-44.

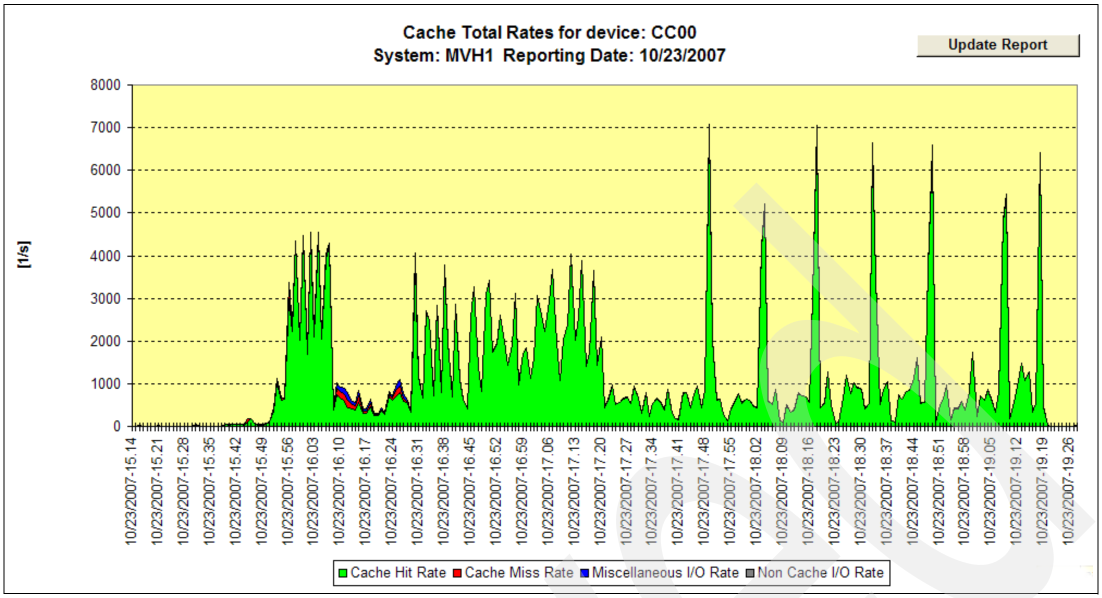

Figure 7-44 Cache Rates for CC00

This again shows that the cache is handling the load very well. Almost 100% of the I/Os are cache hits. The cache is obviously not overcommitted.

7.5.5 DASD activity

Figure 7-45 shows the I/O intensity for the volumes containing the DB2 catalog and directory (DBH0D1 and DBHOD2), the workfile database (DBH1W1 through DBH1W5) and the DB2 active logs (DBH1L1, DBH1L2, and DBH1L3).

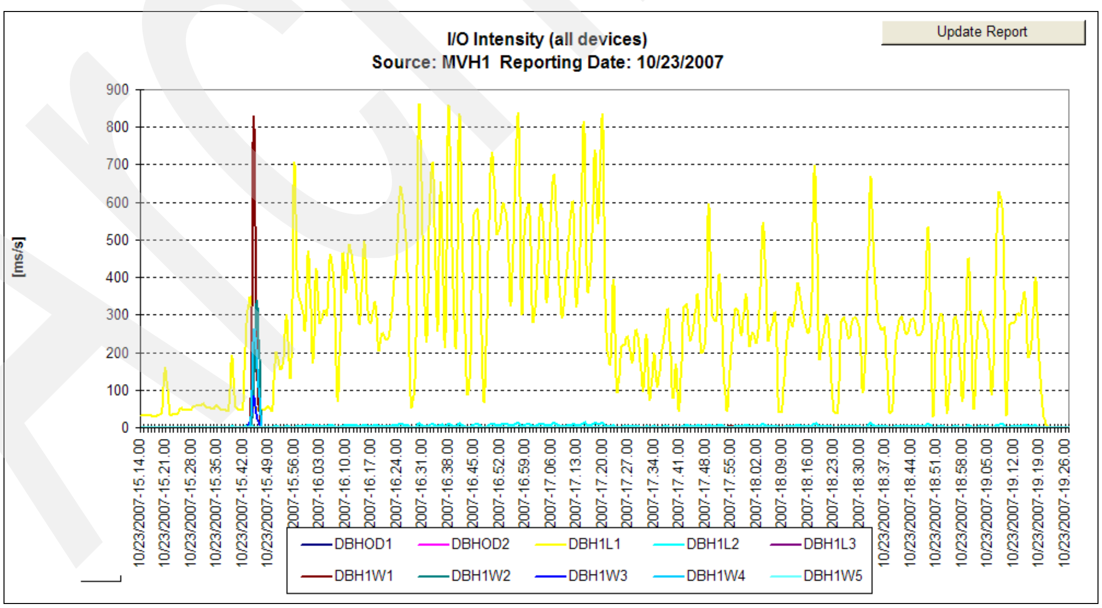

Figure 7-45 DASD activity 0 ML run10

The graph shows only activity on the DBH1L1 volume. This behavior was unexpected because the logs are supposed to be striped across L1, L2, and L3. Analysis showed that the striping did not work because a wrong SMS class was assigned during creation of the logs.

Nevertheless, the logs not being striped did not create a bottleneck. Figure 7-46 shows details for volume DBH1L1.

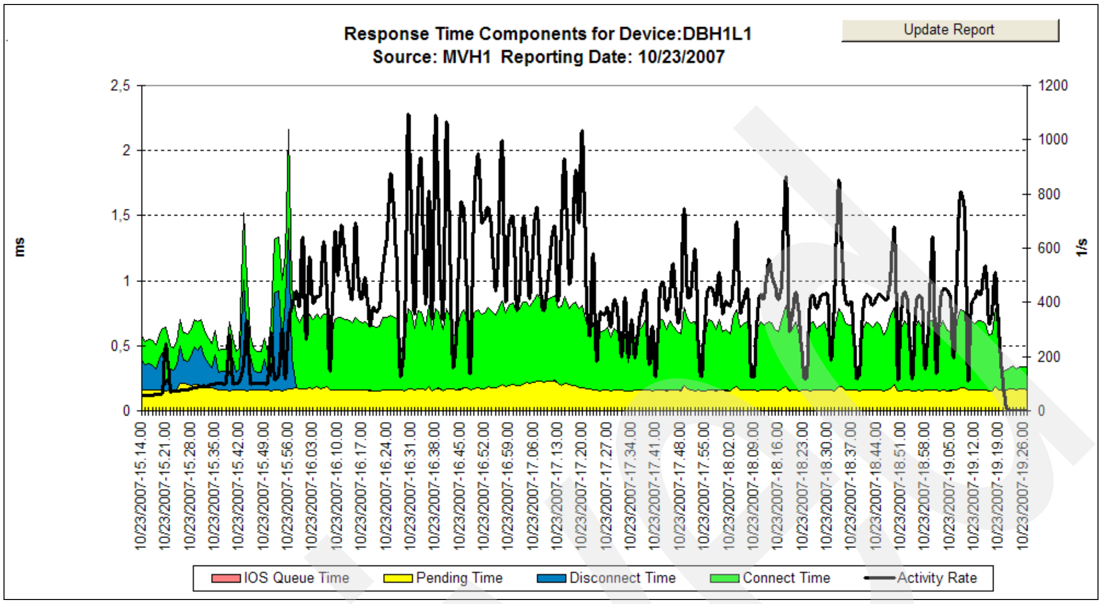

Figure 7-46 DASD activity DBH1L1 ML run 10

During the test the response time was always below 1 ms, which is good.

The following six graphs (Figure 7-47 through Figure 7-52) show details about the volumes (HK1000 through HK1117) containing the SAP ECC database. We take a closer look at the volumes with I/O intensity higher than 1000.

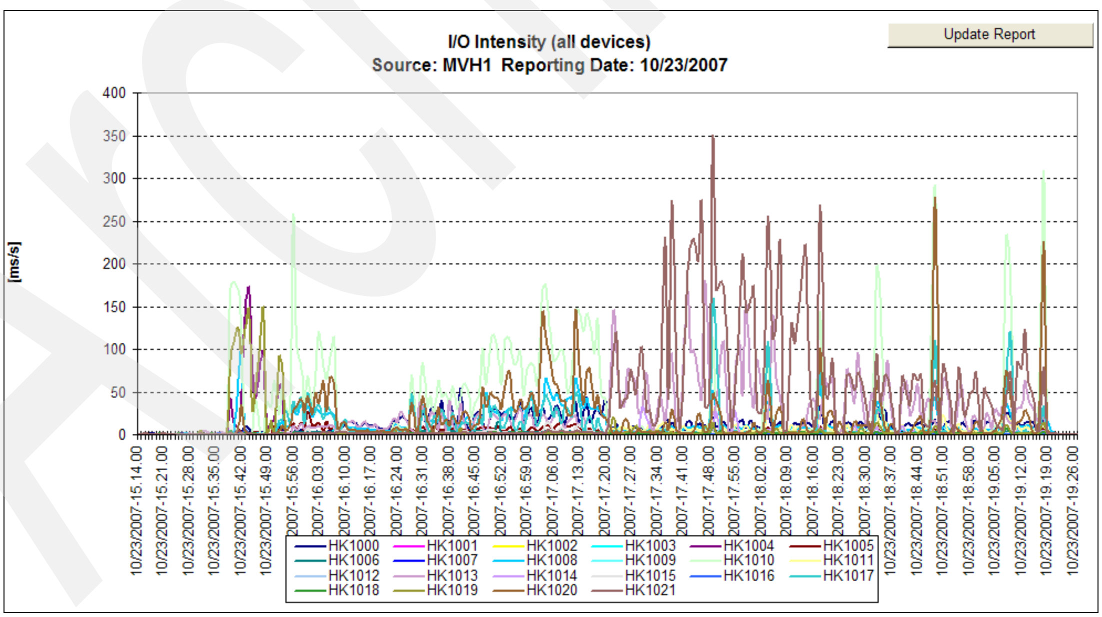

Figure 7-47 DASD activity 1 ML run10

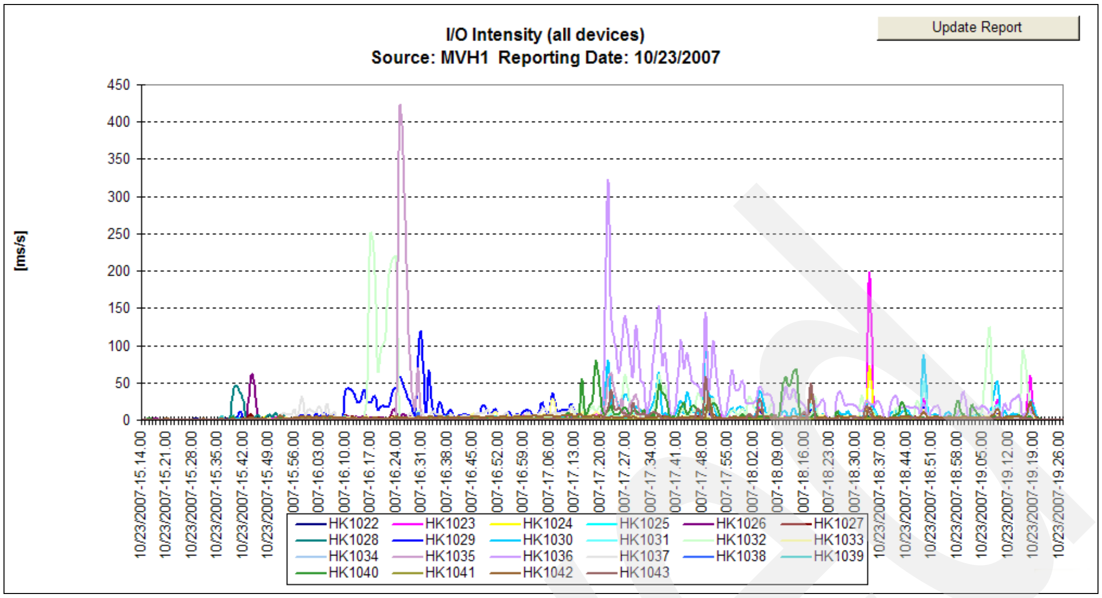

Figure 7-48 DASD activity 2 ML run10

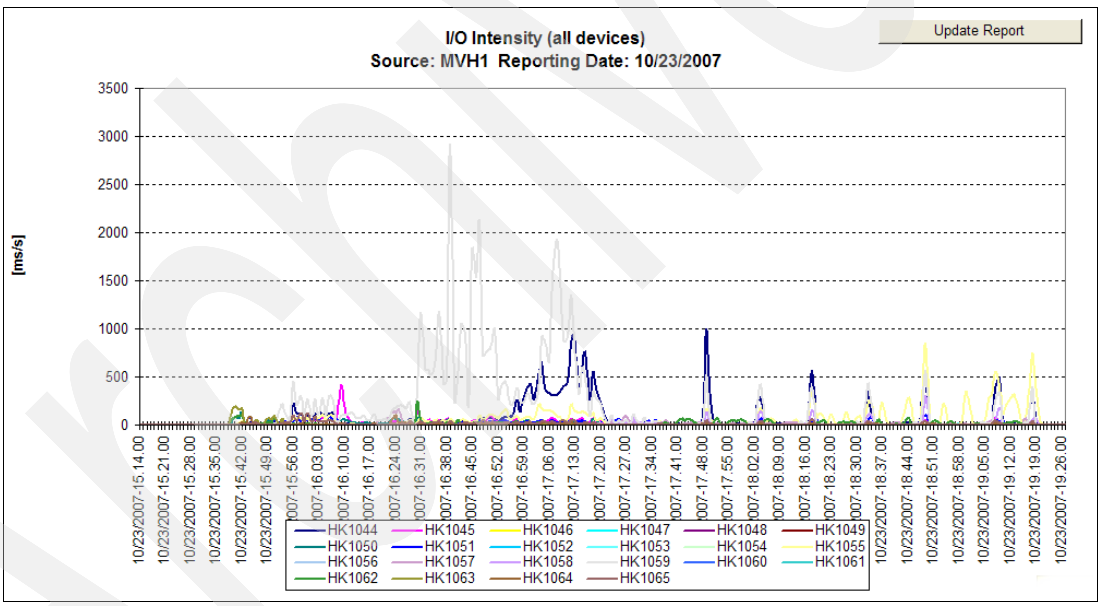

Figure 7-49 DASD activity 3 ML run10

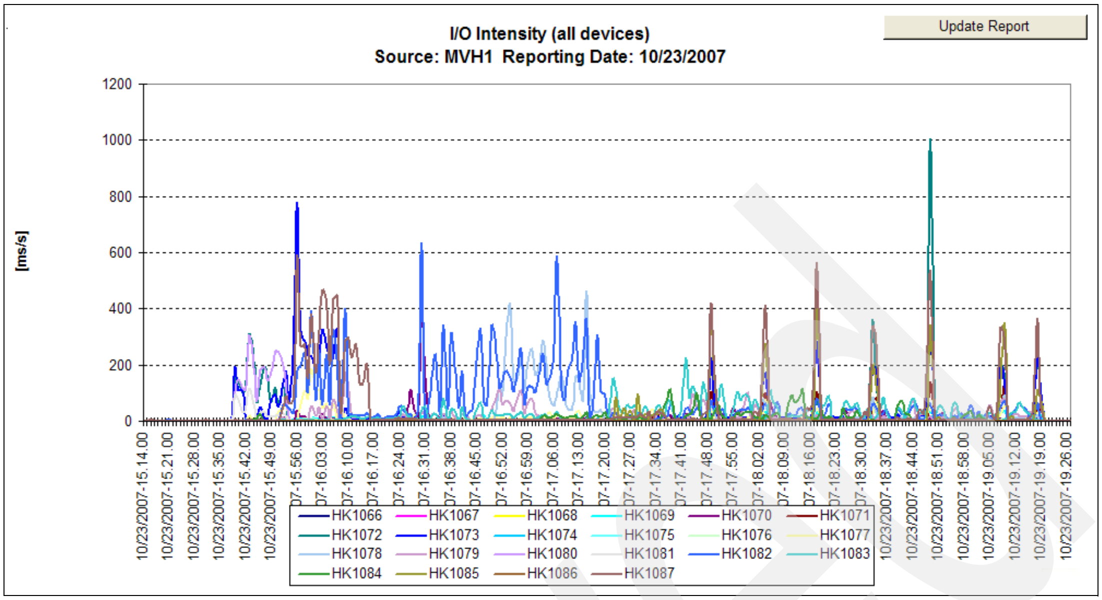

Figure 7-50 DASD activity 4 ML run10

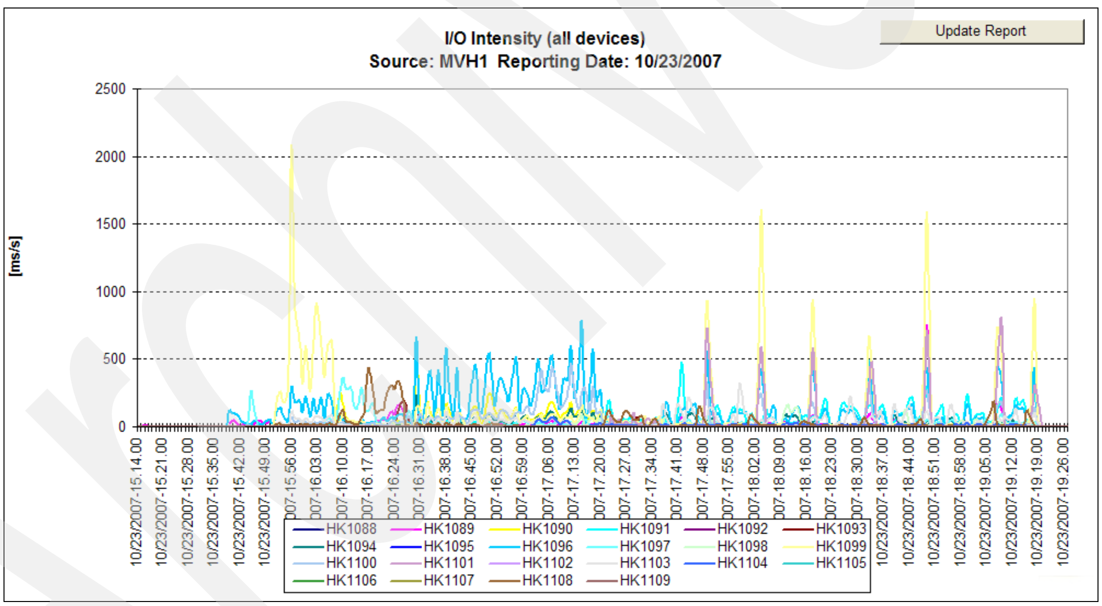

Figure 7-51 DASD activity 5 ML run10

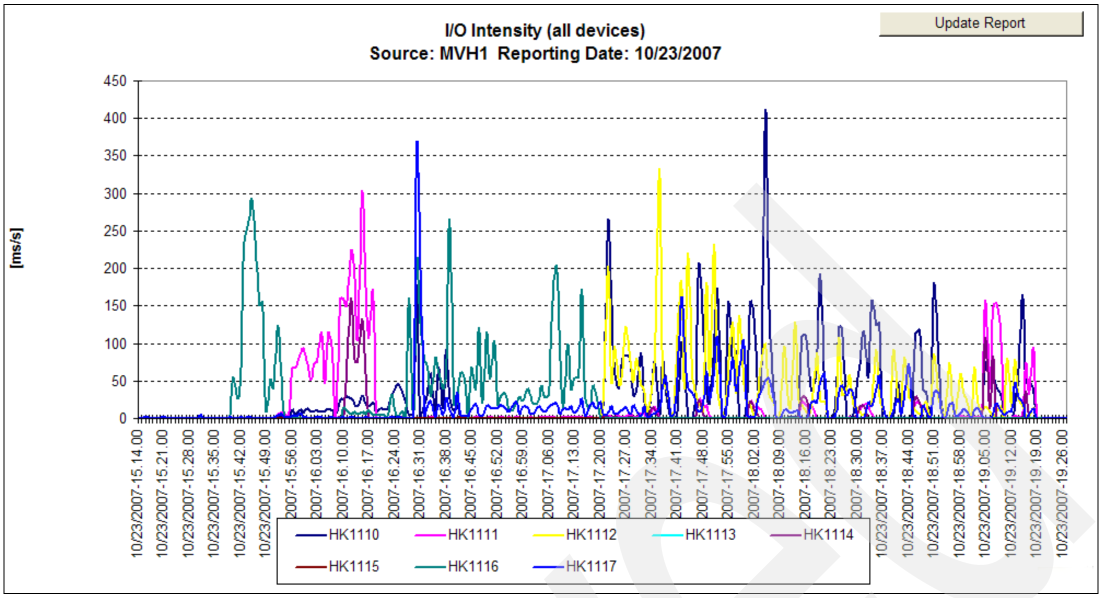

Figure 7-52 DASD activity 6 ML run10

The only volumes showing an I/O intensity higher than 1000 were HK1059 and HK1099, so we took an even closer look at them.

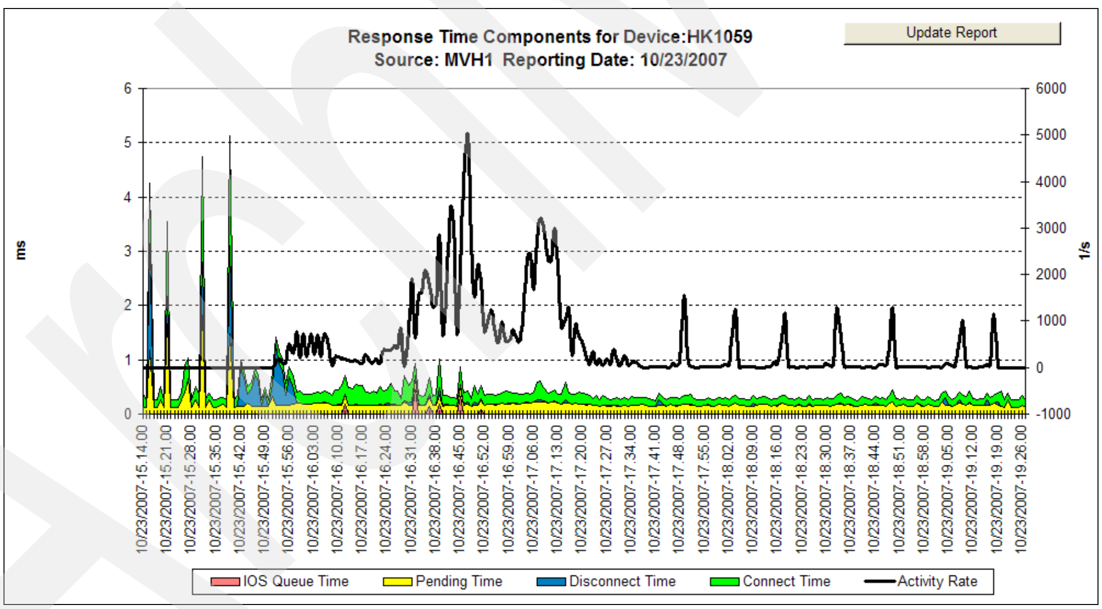

Figure 7-53 Activity HK1059 ML run10

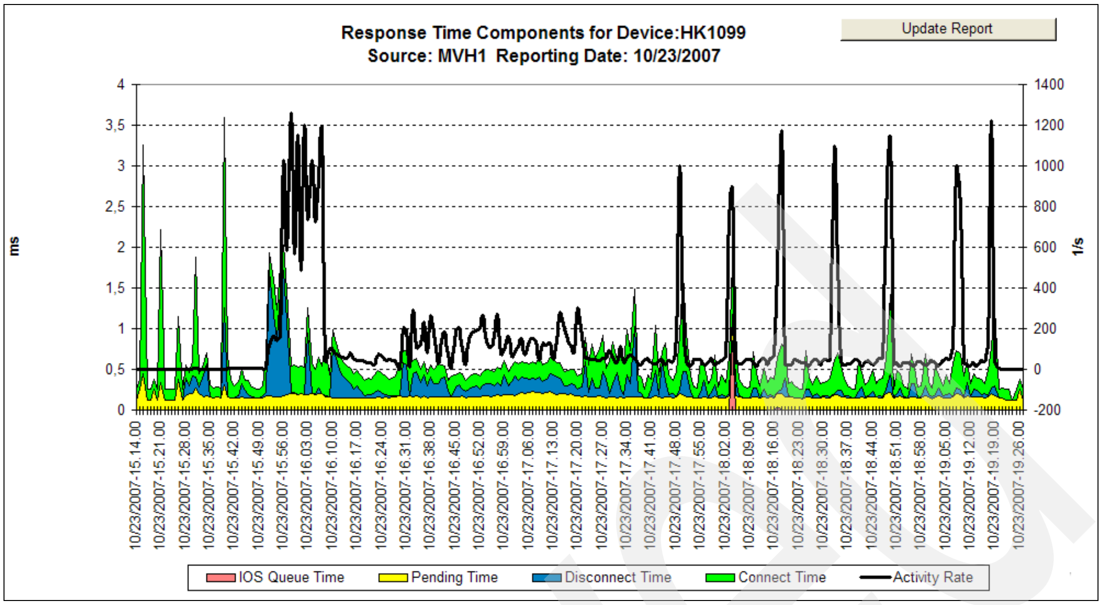

Figure 7-54 DASD activity HK1099 ML run10

Both volumes' response times were always below 1 ms, which is good.

7.5.6 DB2 response time

In this section we describe the observed DB2 response time during run 10. All reported times are averages based on 1426352 accounting records.

We performed several tuning activities for the ML runs. Specifically, we:

- Updated DB2 catalog tables related to statistics on CKMLMV011~0 because of a bad access path
- Changed the tree level from 4 to 2
- Changed the leafpages to 800
- Created new index CKMLMV003~Z01 on columns: KALNR_OUT, BWKEY, PERIO, GJAHR, MGTYP, OTYP_IN, and MANDT
- Updated statistics on CKMLMV003~Z01

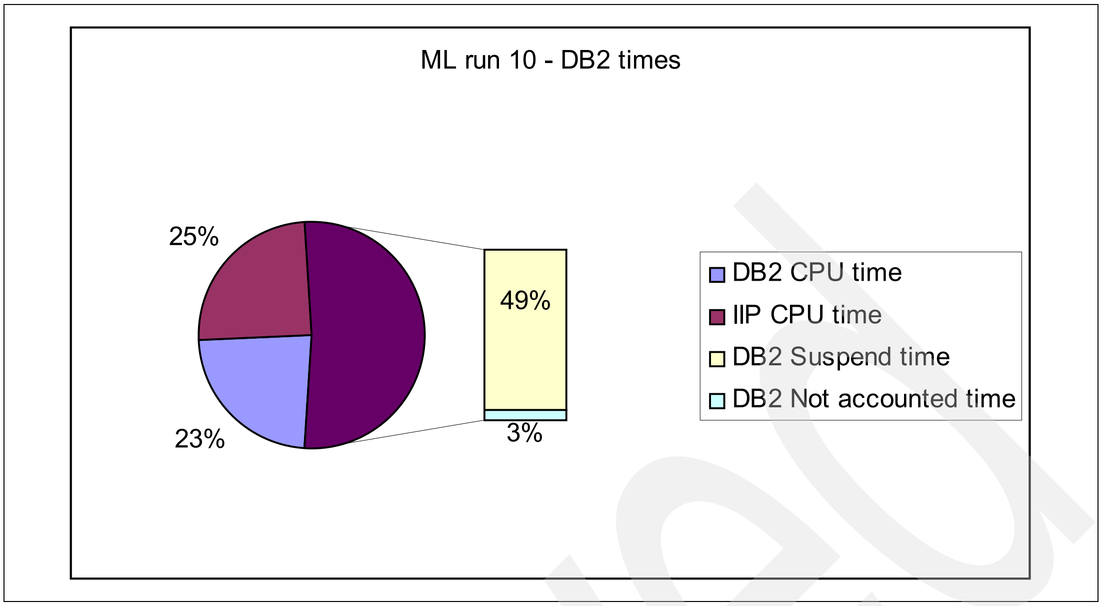

Figure 7-55 ML run 10 DB2

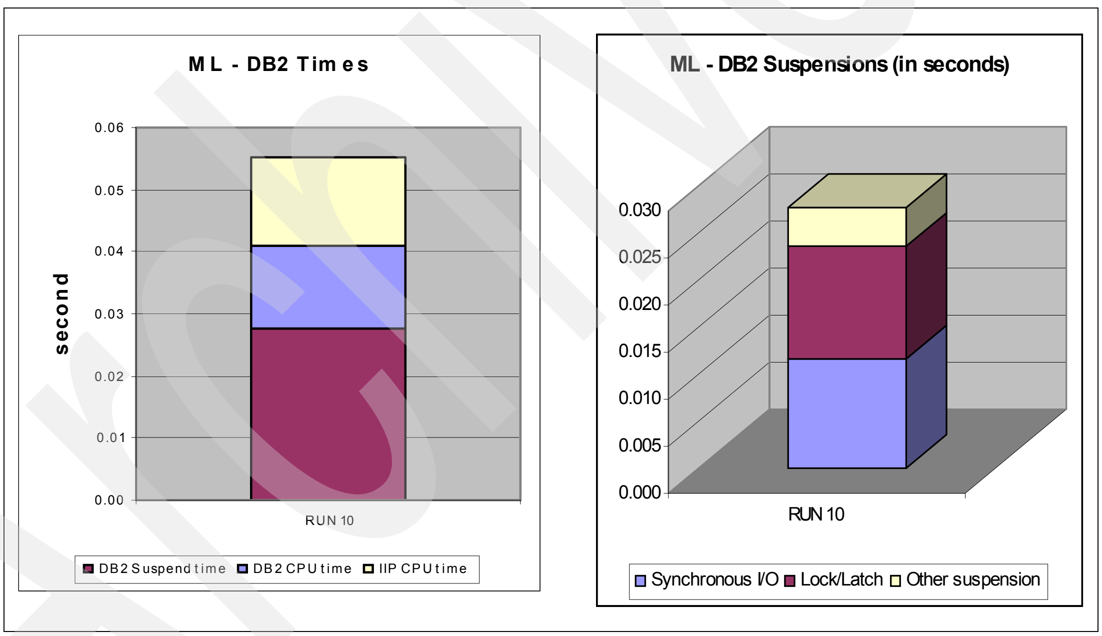

Figure 7-56 ML run 10 DB2 activities

If we look in detail at the type of suspensions that occurred, there are two ways of improvement:

- Reduce the number of synchronous I/Os
- Reduce the number of Lock/Latches

These two contentions represent 85% of the total contention time. We did not investigate beyond this stage due to the fact that we achieved our KPI objectives.

7.6 Conclusions and recommendations

As expected, ML activity is reflected on the ECC environment. During the tests of the so-called “Case1” (running the steps of the material ledger closing in a sequence), we found several optimization possibilities. Finally, we were able to reach the KPI target of four hours. The closing consists of 5 steps:

1. Selection
2. Determine sequence
3. Single level price determination
4. Multilevel price determination
5. Post closing

The runtimes of steps 1 (Selection) and 2 (Determine sequence) were only a few minutes, so we focused more on the remaining three steps.

The following parameter settings were changed, optimized, and tested:

- Degree of parallel processing.
- Package_size (number of materials per ML document).
- Number range buffering.
- Code modification for step 4 (Multilevel price determination). The modification concerns the selection to table CKMLMV011, which should be optimized.
- Code modification for step 5 (Post closing). The modification concerns dynamic value for package_size depending on the costing level.

In the test environment we determined the optimal settings were those shown in Table 7-2.

Table 7-2 Optimal settings

Flow step / Program	Parameter	Last used Value
Single level price determination	Number of parallel processes Package_size	60 50
Multi level price determination	Number of parallel processes Package_size	40 10
Post Closing	Number of parallel processes Package_size	60 1

Additional modifications made included:

- Buffer the number range for the object RF_BELEG. The test results in the use of parallel buffering with a buffer size of 2000.
- The modification for the Multilevel price determination step was handled in the OSS message 5528707 2007. SAP Note 1113247 was already created and should be applied.
- The modification for the Post closing step is still handled in the OSS message 900575 2007.

Additional recommendations:

- Implement the document split according to SAP Note 539452.
- Create a new Secondary index for the table CKMLMV003.

The ML run showed a demanding load on the System z Infrastructure which was handled very well. On average three CPs and three zIIPs were used. Sometimes we saw short spikes where 5 CPs and 4 zIIPs were used. The only potential bottleneck that lies ahead is the OSA adapter. Should load increase, we suspect this is the first component that would need adjustment. (Solution: Have an extra OSA adapter and configure it to be used simultaneously with the current one. This would be like having an extra driving lane in each direction on the highway.) The FICON Express2 channels can handle much more load. It is also possible to configure additional channels. The current DS8000 can handle much more additional load. This box is currently configured with 128 GB of storage, which can be doubled to 256 GB. In other words, the entire configuration still has a great deal of upward scalability potential.

Regarding System p, all resources are used on hkpoceccapp2 and hkpoceccapp3. Hkpoceccapp1 has free CPU resources, so to improve the results these free resources should be able to be used to balance the load better at the SAP level. On the other hand, more CPU resources on hkpoceccapp2 and hkpoceccapp3 would probably improve the performance as well, for instance, more CPUs or more powerful CPUs.

7.6.1 OSS notes

Implement the following notes (and corrections from the notes) from OSS in the SAP system after the upgrade (ECC6, Patch level SAP_APPL, 600, SAPKH60009).

System performance

- 51789: Poor user distribution in logon distribution

Functional corrections

- 980475
- 1025058
- 1057851

8

Compression/decompression test with MRP

This chapter describes an additional MRP test we performed. The objective of this test, which was not in the original plan, was to compare MRP runs in a DB2 compressed versus a DB2 uncompressed environment to evaluate the impact of compression in terms of both performance and disk space savings.

The topics included in this chapter are:

8.1 Test overview

Table 8-1 shows the steps we performed to test MRP in compressed and uncompressed environments.

Table 8-1 Test steps

1. Nov. 20th (contains MRP data) Level2 Backup Restore	
2. Clearing of main table related to MRP	
a	**Reorg with statistics PLAF,EBAN,MDTB,MDKP,KBED + MARA, MARC**
b	**Runstats Tablespaces/Index listed in MRP Note: MARA, ….. CECUSFT**
c	**MRP Run after PIR deletion**
d	**RESB/PLAF Table size check (with SAP transaction – DB02)**
e	**RESB/PLAF DB Reorganization with ‘keepdictionary’ reorg parameter ‘on’**
3. Master Backup creation for Compression Test (Flashcopy)	
4. MRP run in Compression status	
a	**MRP run after new PIR entry**
b	**Monitor dynamic statement cache to determine the top 10 most frequently accessed DB2 tables**
c	**RESB/PLAF Table size check with SAP transaction – DB02)**
d	**MRP Run Time check**
5. Master Backup Restore (Flashback)	
6. MRP run in Decompression status	
a	**Decompression of top 10 tables identified during step 4b (includes RESB & PLAF) including REORG+statistics**
b	**MRP run after new PIR entry**
c	**RESB/PLAF Table size check using SAP transaction DB02**
d	**MRP Run Time check**

Figure 8-1 compares the run times achieved; the same MRP run is three times faster using DB2 compression than not using it.

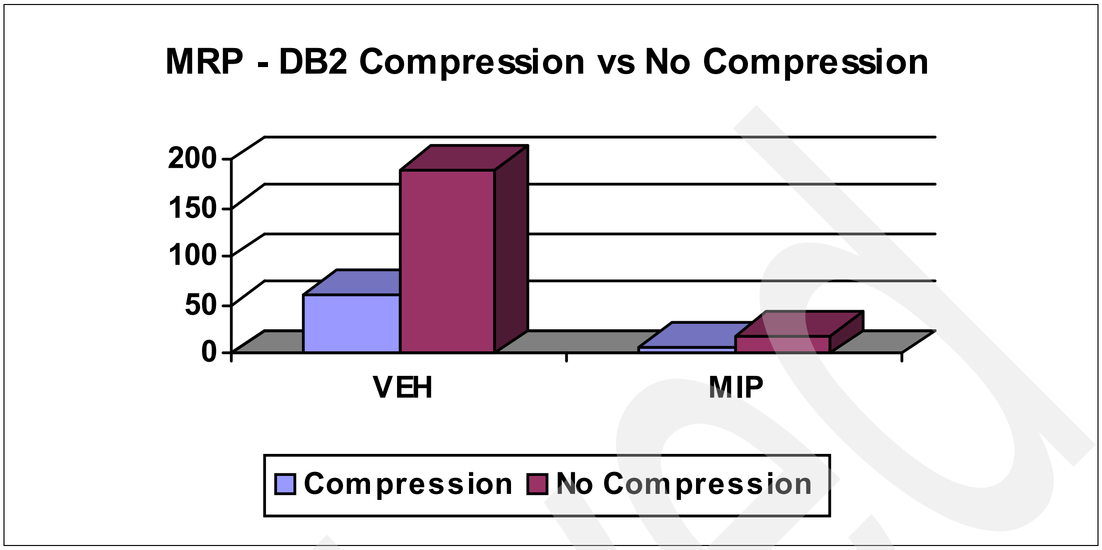

Figure 8-1 MRP Run duration in minutes

To execute the MRP run without compression we turned off the compression on the following ten tables that are used most by the MRP application:

- RESB
- PLAF
- MARA
- DBVM
- MATERIALID
- MARC
- MKAL
- MAST
- STKO
- STZU

8.2 Analysis of System p infrastructure layer test results

The following two graphs (Figure 8-3 and Figure 8-2 on page 200) compare CPU usage on application server 1 during the compressed and the uncompressed tests. The uncompressed DB2 test consumed only 15% of CPU resources during the MRP run, and it was completed approximately 3 times faster.

Figure 8-2 shows CPU usage on the SAP application server during the DB2 uncompressed run; with an average use of 15% of 20 processors during 3 hours, we can consider that the processor usage reduced to one hour is equivalent to 9 processors (3 processors x 3 hours).

Figure 8-3 shows CPU usage on the SAP application server during the DB2 compressed run; with an average use of 45% of 20 processors during one hour, we can consider that the processor usage reduced to one hour is equivalent to 9 processors.

Comparing these two runs, the overall processor usage on the SAP application server is the same, but it took longer during the DB2 uncompressed run. In an SAP process, the application server waits for database access to process data. If the time needed to retrieve the data in the database increases, you will have more processes waiting on the application server and therefore less processor utilization. The increase in database processing time during the DB2 uncompressed run is described further in 8.3.2, “DB2 times” on page 202.

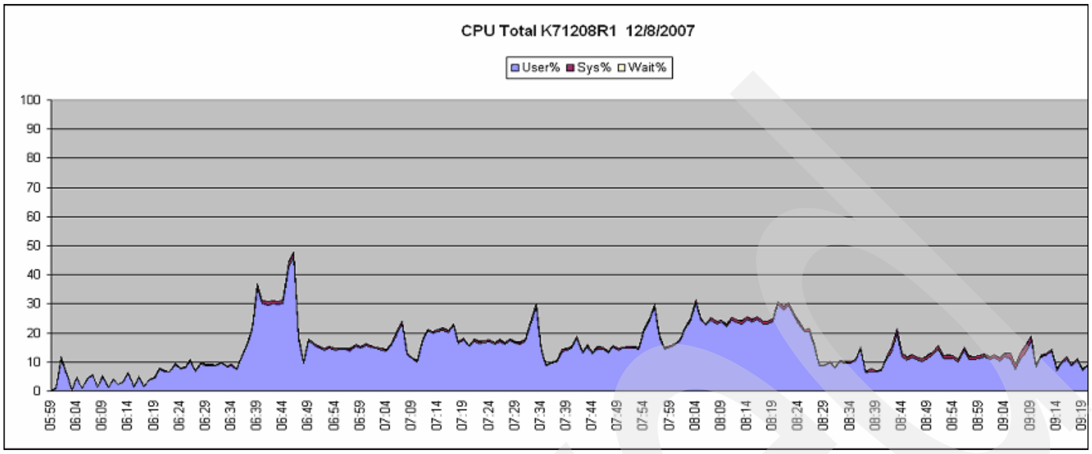

Figure 8-2 CPU summary - hkpoceccapp1 - uncompressed DB2

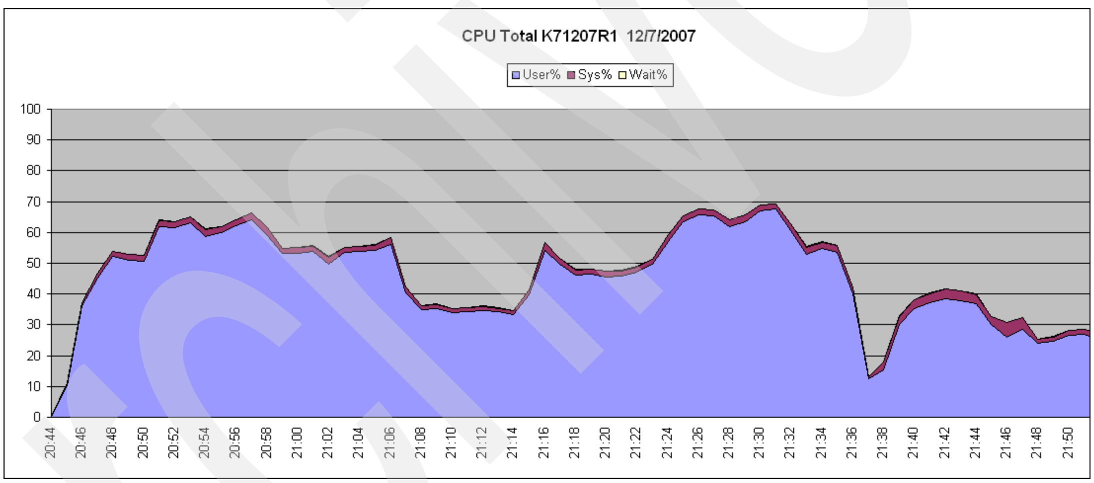

Figure 8-3 CPU summary - hkpoceccapp1 - compressed DB2

Memory use

The next two graphs (Figure 8-4 and Figure 8-5 on page 201) compare memory usage in the two environments. It is easy to see that in either cases, compressed or uncompressed, the application server uses the same amount of memory.

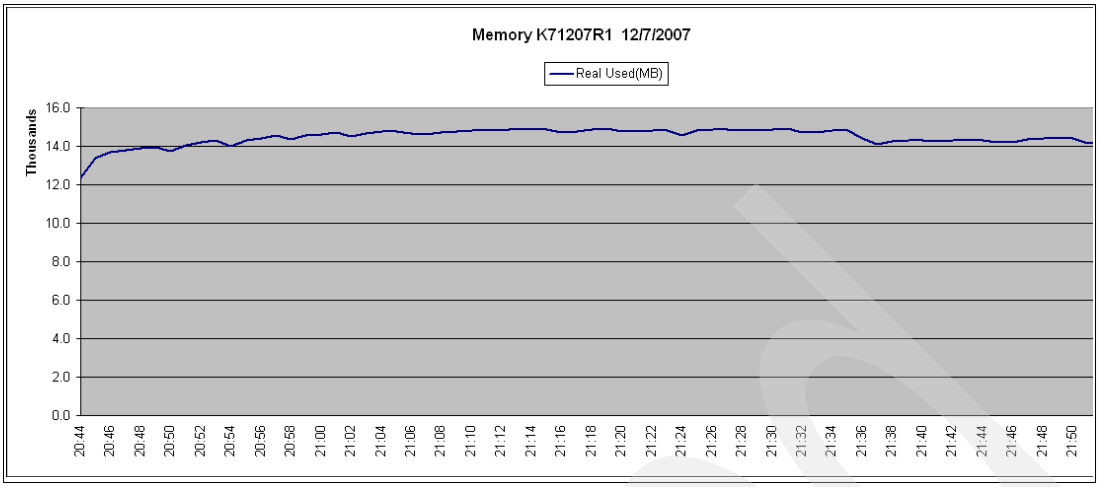

Figure 8-4 Memory summary on hkpoceccapp1 - compressed DB2

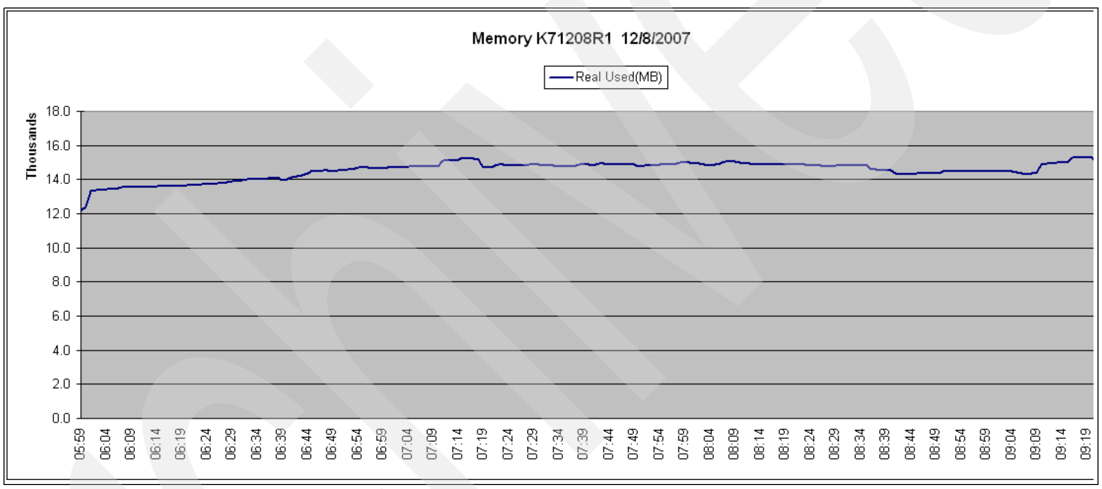

Figure 8-5 Memory summary on hkpoceccapp1 - uncompressed DB2

8.3 System z infrastructure layer and DB2 behavior

8.3.1 System z infrastructure

As shown in the next two figures, during the MRP with DB2 compression run we used 29859 CPU seconds, compared with 39683 CPU seconds used during the run without DB2 compression. In other words, the MRP run with DB2 compression used 32% less CPU resources.

As these figures also show, the average number of I/Os per second in the uncompressed environment is 22000, compared with 2500 on the run with DB2 compression.

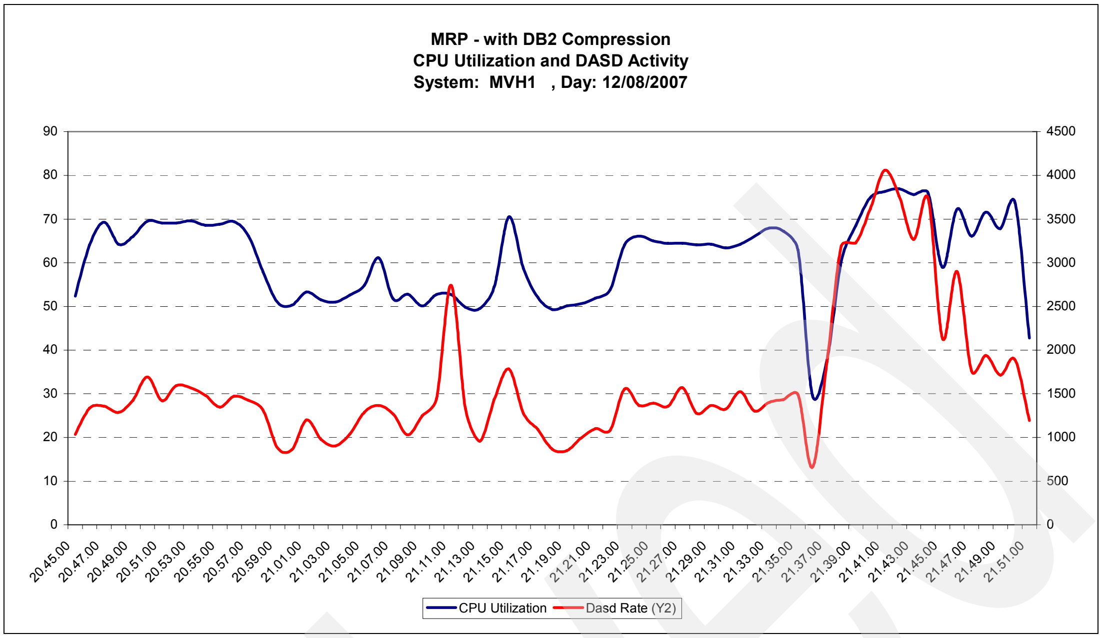

Figure 8-6 MRP with DB2 compression - CPU usage

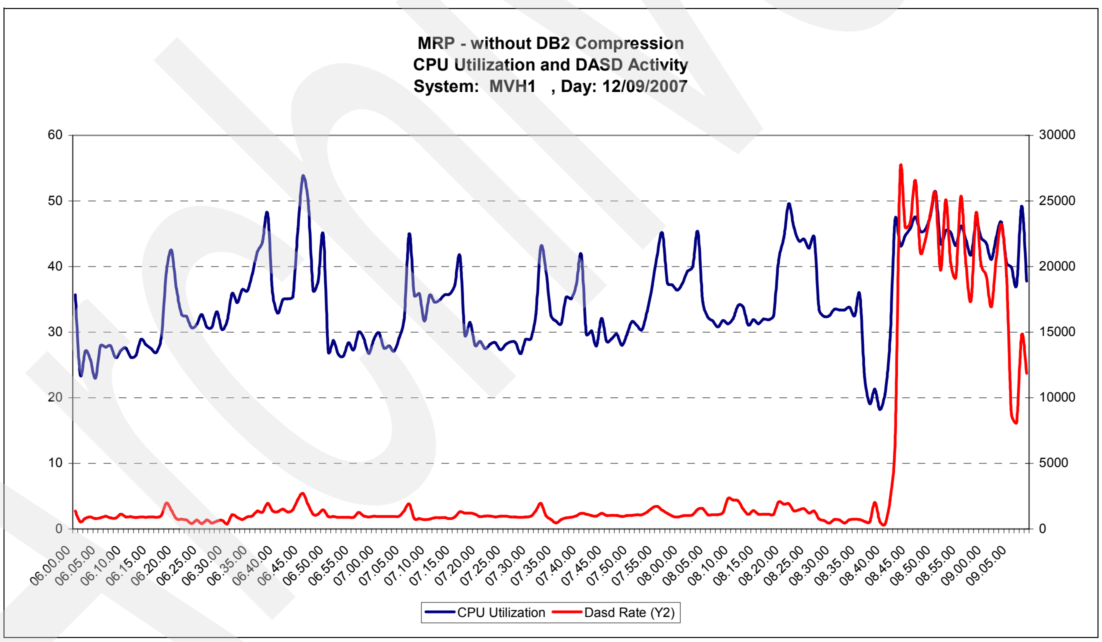

Figure 8-7 MRP without DB2 compression - CPU usage

8.3.2 DB2 times

The times reported are extracted from DB2 accounting reports executed after each run. Each MRP run is separated between VEH and MIP. The times are average seconds based on more than 400000 accounting records for both the VEH runs and more than 40000 accounting records for both the MIP runs.

Figure 8-8 displays the will report the number of executed SQL statements. The numbers are similar, and we can conclude that the MRP with DB2 compression run is comparable to the MRP without DB2 compression run.

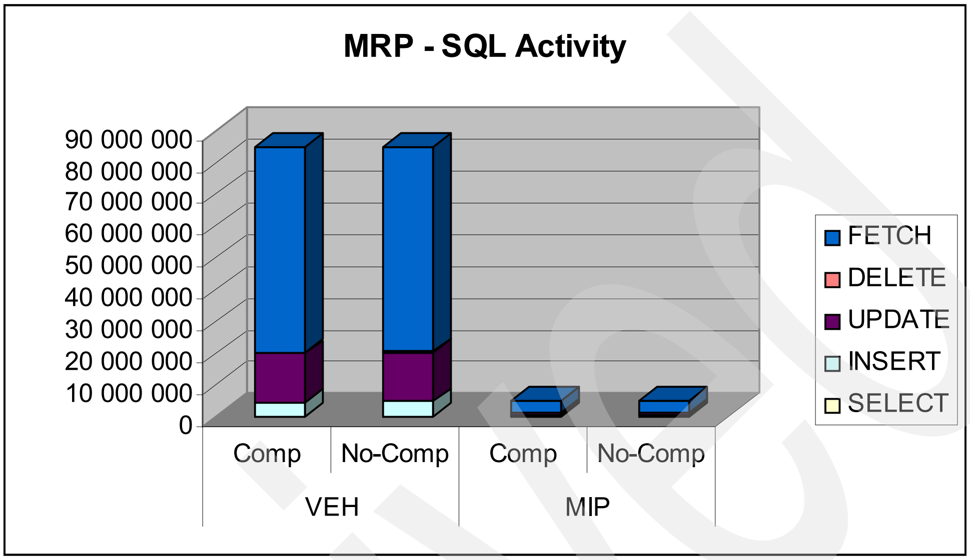

Figure 8-8 SQL activity during MRP runs

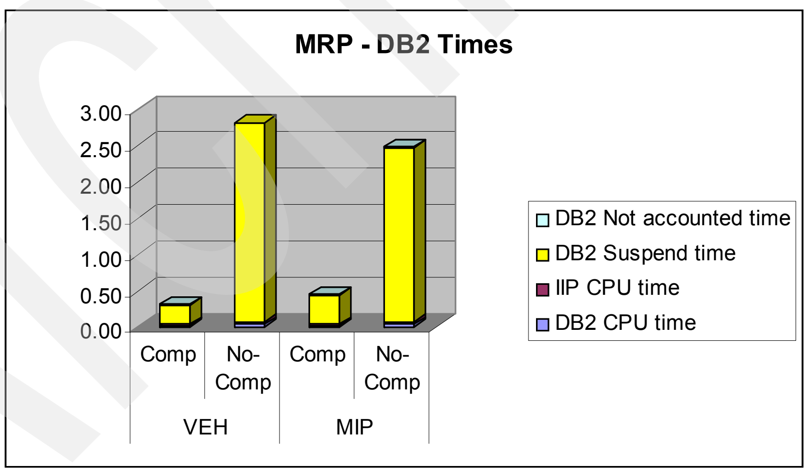

Figure 8-9 DB2 Class 2 times in seconds

The run duration for the MRP without DB2 compression is affected by the time spent in suspension. This suspension time is more than 2 seconds in both the VEH and MIP runs.

Figure 8-10 on page 204 shows detail about the suspension times. the main reason for suspension is Page Latch for both VEH and MIP runs.

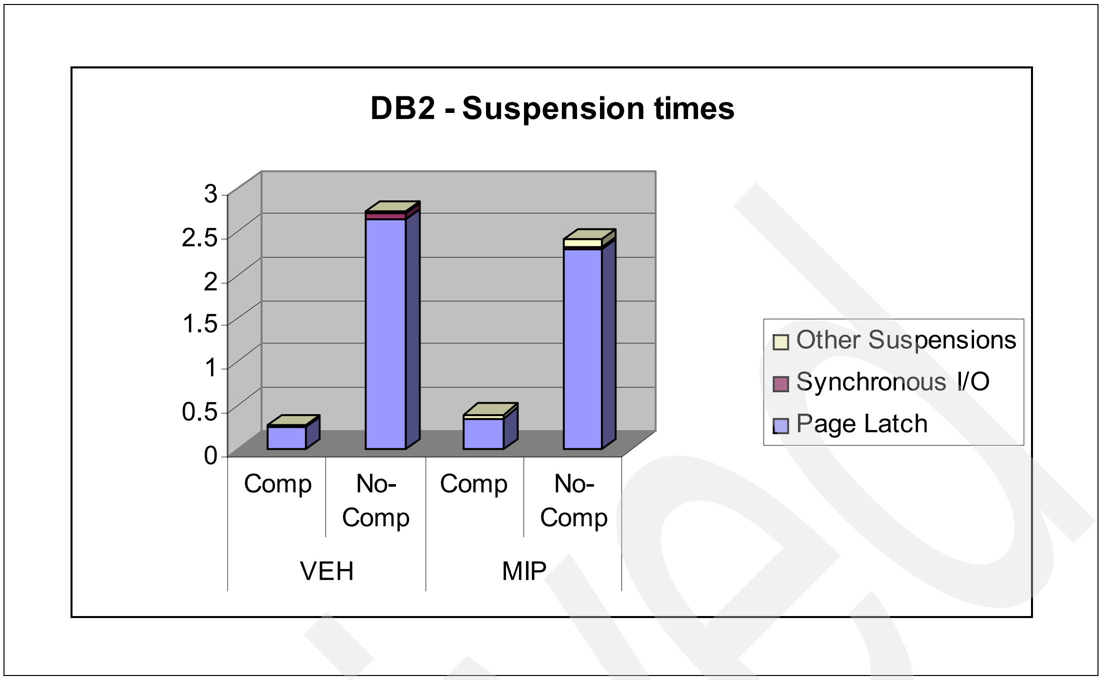

Figure 8-10 DB2 Class 3 Suspension times

8.3.3 DB2 tablespace layout

The 10 tables that were used most during the MRP run were uncompressed and reorganized, and the statistics were collected before starting the run without compression.

Table 8-2 Table layouts with and without DB2 compression

	Comp After MRP run				No-Comp After MRP run		
Table name	#Rows	Avg Rowlgth	Size in KB	%Compr	#Rows	Avg Rowlgth	Size in KB
RESB	70 363 249	253	18 585 380	80	70 363 544	1251	93 052 036
PLAF	6 704 775	163	1 844 728	83	6 706 599	1090	13 863 276
MARA	2 337 161	180	533 188	85	2 337 161	1289	4 675 276
DBVM	1 715 122	57	182 160	58	1 715 122	140	301 372
MATERIALID	2 337 161	58	170 404	59	2 337 161	144	421 592
MARC	2 348 520	193	57 612	83	2 348 523	1205	4 698 008
MKAL	521 387	73	47 428	78	521 387	343	229 512
MAST	328 847	41	16 860	72	328 847	154	64 324
STKO	202 847	54	13 772	78	202 847	264	67 636
STZU	202 847	46	11 808	77	202 847	208	54 112

Table 8-2 shows that the compression ratio is 80% on the biggest tables. For the RESB table, at the end of the MRP run the table size is 18.5 GB with DB2 compression and 93 GB without DB2 compression. In other words, the RESB table is 5 times larger without DB2 compression.

In summary, on just 10 tables the storage saving is nearly 100 GB. This is illustrated in Figure 8-11 on page 205.

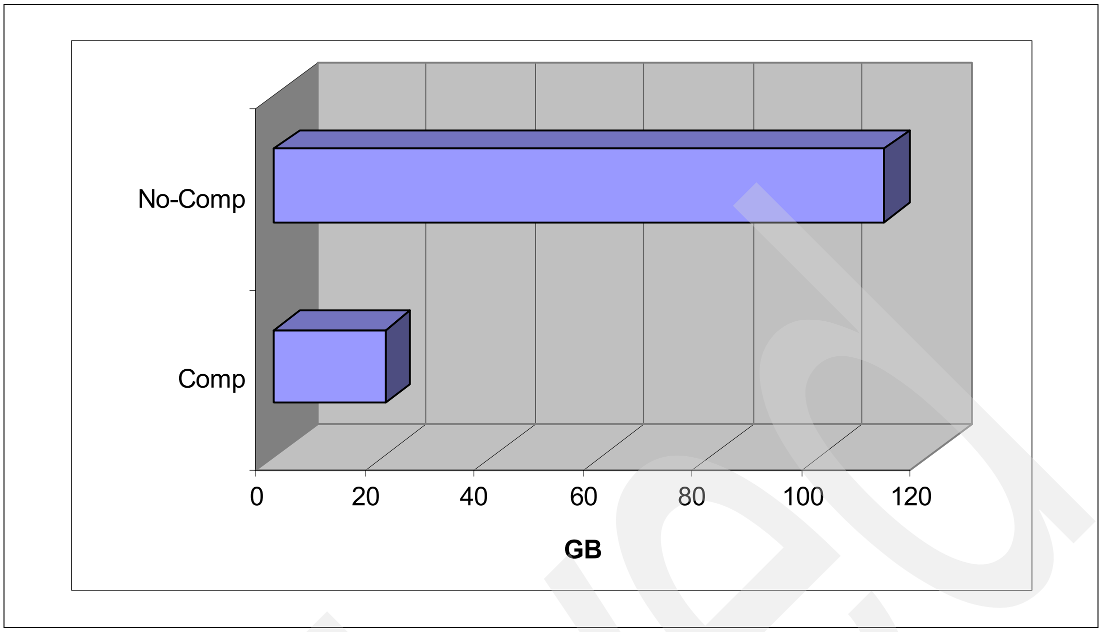

Figure 8-11 Storage used by the top 10 tables

8.3.4 DB2 administration

The impact of DB2 compression was also visible when we used utilities for database management. To demonstrate this impact two utilities were executed after each MRP run:

– Reorganization of one partition on the RESB table

– Image copy of one partition on the RESB table

As shown in Figure 8-12, the run time for each utility without DB2 compression is twice that of the same utility with DB2 compression.

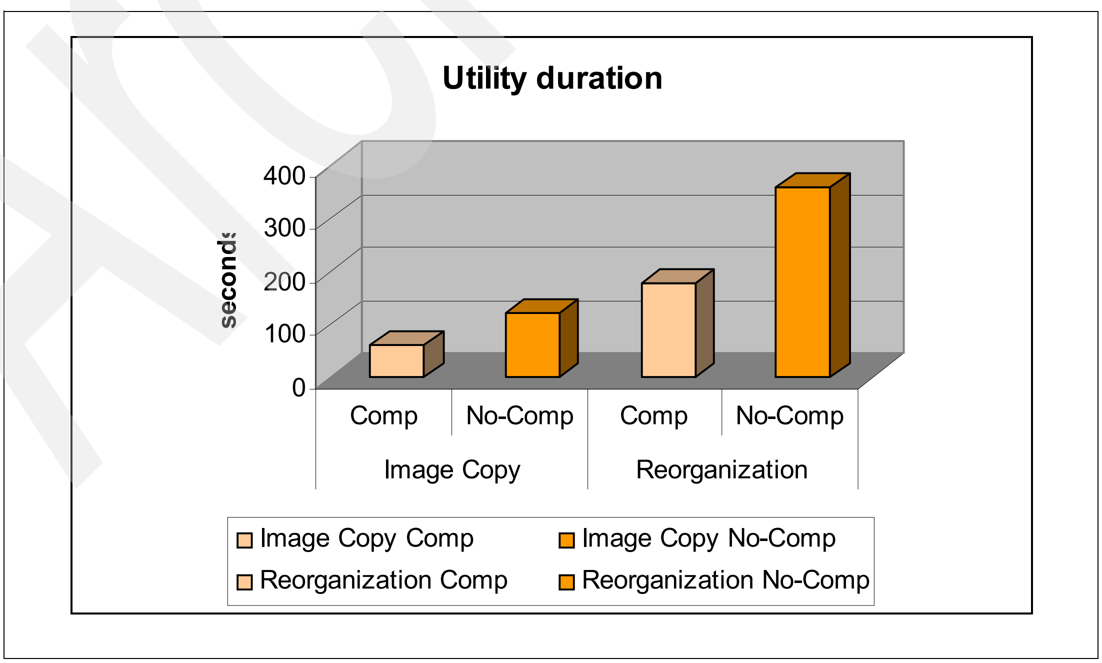

Figure 8-12 DB2 impact

8.4 Conclusions: Results and recommendations

The test demonstrated the savings provided by DB2 compression in the following areas:

- Run time of MRP
- CPU usage
- Storage needed
- Database administration

We recommend using DB2 compression, especially with the Unicode SAP environment. For a little overhead we could reduce the run time for our SAP application (by 3 time during this test) and even reduce the CPU consumption by minimizing the number of I/Os and the DB2 contention.

Without DB2 compression, it would have been impossible to achieve the performance results described in the previous KPI.

9

Non-functional requirements

This chapter presents a discussion of the non-functional requirements expressed by the customer. It is concerned with understanding the customer's expectations regarding the behavior of the solution for a defined workload rather than the business function the solution provides.

The topics covered in this chapter are:

9.1 Overview

The non-functional requirements for a business system address those aspects of the system that, while not directly affecting the business functionality of the system as seen by the users, can have a profound impact on how that business system is accepted by both the users and the people responsible for supporting that system.

The definition of these requirements is an essential factor in developing a total customer solution that viably delivers functionality; however, detail here is limited to the early stages of an implementation project. It is usual to refine these during the “blueprint” phase because functional choices will determine the critical areas to be measured.

Benchmark or Proof-of-Concept tests can be used to verify that the implementation meets these requirements or to indicate what corrective action is necessary. Ideally a series of tests should be planned that map to the development schedule and grow in complexity.

Purpose

Non-functional requirements (NFRs) are used primarily to drive the operational aspects of the architecture (platform and network infrastructure, placement, and so forth), in other words, to address major operational and technical areas of the system in order to ensure the robustness of the system.

The NFR definitions are used:

- To define requirements and constraints on the IT system.
- As a basis for early system sizing and estimates of cost.
- To assess the viability of the proposed IT system.
- To drive design of the operational models. Non-functional requirements are frequently the most important determining factor of the architecture.
- As an input to component design in order to meet NFRs such as performance.

The performance- and scalability-related NFRs provide the baseline against which the subsequent performance measurement and tuning activities are scoped and executed.

9.2 Key non-functional requirements

The requirements of the system to be delivered must be understood in each of the areas identified in this section in order to facilitate the design and development of the operational model, that is, the computers, networks, and other platforms on which the application will execute and by which it is managed.

These are key for the execution of benchmark tests. They also affect the design of technical and application components. For example, service level requirements may imply application component performance requirements.

9.2.1 Performance and capacity

Performance and capacity definitions measure characteristics such as throughput, response time, utilization, and static volumetric requirements.

Response time requirements relate to the times used to complete specific business processes, batch or interactive, within the target business system.

- For batch, this could be the elapsed time to complete the overnight bill run.
- For interactive transactions this could be either:
  - The end-to-end response time associated with a specific single user-system interaction, for example, the time between a user selecting the “process button” and the result being displayed to the user.
  - The elapsed time to complete a specific business process or transaction, for example, the elapsed time to complete the order-entry business process.

To set a value for elapsed time, determine the expected time for the execution of business processes by different user populations. Then, establish acceptable variations between these user populations.

It is recommended that response time “bands” be defined. The assignment of business transactions to bands should depend on their frequency of use and processing complexity.

Throughput requirements relate to the ability of the business system to execute a given number of business- or system-related processes within a given unit of time, in other words, the workload. For example:

- The number of account balance inquiries processed per day
- The number of new orders processed per day
- The number of telephone call records processed per nightly bill run

This activity involves an understanding of the frequency of invocation of each of the business processes specified within the process model (for example, the number of orders per day) along with an understanding of the size of each process (for example, the average number of items per order).

The approach taken depends on the number of processes defined within the overall business system. In practice, many business systems have so many processes that it may be necessary to limit a detailed examination of the volumetric to a subset. This is typically a combination of high volume processes (for example, the top 20% most frequently executed processes) plus those that we know will place significant demands on one or more system components. The remaining processes are then managed by making an agreed upon set of working assumptions.

Utilization requirements relate to the maximum acceptable loading of the nodes on which the business system is to be implemented when running the defined workload. Examples are that the network bandwidth utilization must not exceed 20%, or, the database server will be no more than 60% utilized.

Static volumetric requirements relate to the volumetric for the data entities that exist within the target system that, although relatively static, are likely to have a significant effect on the overall sizing of the target system. Examples are the user community makeup and input/output devices and the number of business system users by type, different user locations, customers, customer accounts, and products.

The combination of these four sets of performance sub-requirements (throughput, response time, utilization, and static volumetric) is required to properly define the performance baseline to which the system must be designed. For a benchmark test we must clearly state the conditions under which the tests will be run, for example, the specific workload mix, any background load, and so forth.

In summary, the system must be designed to meet the agreed upon response time requirements, while supporting the defined workload mapped against the given static volumetric baseline, on a system platform that does not exceed the stated utilization.

9.2.2 Scalability

Scalability is the ability to expand the system architecture to accommodate more users, more processes, more transactions, more data, and additional systems and services as the requirements change and the system evolves to meet the needs of the business in the future. The existing systems are extended as far as possible without necessarily having to replace them.

This requirement directly affects the architecture as well as the selection of hardware and system software components.

The solution must allow the hardware and the deployed software services and components (both application related and technical) to be scaled horizontally as well as vertically. Horizontal scaling involves replicating the same functionality across additional nodes; vertical scaling involves the same functionality to be executed across bigger and more powerful nodes.

Scalability definitions measure volumes of users and data the system will support. This can be expressed as a profile over time:

Stage 1 …. Stage 2 …. Stage x …...

9.2.3 Availability, recoverability, reliability

The availability expectations of a system relate to how many hours in the day, days per week, and weeks per year the application is going to be available to its users, and how quickly they should be able to recover from failures. Since the system constitutes Software (including applications), Hardware, and Network components, this requirement extends to all three of them.

Hardware availability, recoverability, and reliability definitions measure system up-time. For example, this can be specified in terms of mean time between failures (MTBF).

9.2.4 Disaster recovery

In the event of a disaster like a flood, tornado, and so forth, the entire complex where the system is hosted may become completely inoperable or inaccessible. Mission-critical systems should have a plan in place to recover from such disasters within a reasonable amount of time.

The solution implementing the various processes must be integrated with the existing enterprise disaster recovery plans. The processes must be analyzed to understand the criticality of each process to the business, the impact of loss to the business in case of non-availability of the process, the maximum tolerance on any down time. Based on this analysis, appropriate disaster procedures must be developed and plans made.

As a part of the disaster recovery, electronic backups of data and procedures must be maintained at a disaster recovery location and be retrievable within the appropriate time frame for system function restoration. In case of high criticality, real-time mirroring to a mirror site must also be considered.

9.2.5 Flexibility

This encompasses maintainability, portability, and extensibility, meaning the ease of extending the architecture to include new business functions and technologies in the future.

Functional

As customer's business needs grow and evolve, the solution must be extensible in order to address future functionality and changes without having to be completely rewritten.

Specifically, the solution must allow deployed processes to be easily modified and redeployed, and also allow new processes to be easily deployed without impact to the overall solution in a major way. The solution must be flexible enough to incorporate new services without significant changes.

Technical

An on-demand business may grow by mergers, so the solution must be flexible to allow new systems on varying platforms to be incorporated. Also, the solution must be flexible enough to modify the existing processes to incorporate new systems.

Additionally, the architecture must be innovative and flexible enough to accommodate related technological changes that could be leveraged in the future.

The solution must allow additional channels of information delivery to be supported in the future with no significant changes to the architecture.

9.3 Other non-functional requirements

Although included here for completeness, the non-functional requirements in this section are generally not included in benchmark testing.

9.3.1 Security

Definition and monitoring of effectiveness in meeting the security requirements of the system, for example to avoid financial harm in accounting systems, is important.

9.3.2 Manageability

In general, this refers to system management, but also the solution must allow for the service implementations to be dynamically chosen based on changing needs. It should be easy to define and implement policies for service selection based on desired qualities of service (QOS).

9.3.3 Environmental

This includes, for example, safety, meaning the ability of the system to not harm users, including physical harm, as in medical systems or other forms of harm.

9.3.4 Usability

Usability measures characteristics such as aesthetics and consistency in the user interface. Aesthetics is the artistic, visual quality of the user interface. Consistency is the constant use of mechanisms employed in the user interface. This applies both within the system and with other systems.

9.3.5 Accessibility

This is a measure of the ease with which different facets of the system are exercised, for example, by persons with disabilities who need assistive technologies.

9.3.6 Data integrity

Data integrity includes considerations such as currency of information (valid and up to date), locality of updating, data retention standards, and so forth. It also includes transactionality, which is defined as the system's ability to deliver application functionality in a predictable and reliable manner. The main transactionality requirement is based on the principle that no transaction failure shall result in irrecoverable data loss.

9.4 Constraints

These are additional "rules" the system must conform to or satisfy that will impact the chosen solution and may also need to be included in benchmark tests. Some examples are:

- The business constraints which the system must satisfy:
  - Geographical locations
  - Internal organization and other factors
  - Government regulations
  - Economic and other external factors
- The technical standards the system must satisfy:
  - Technical accessibility standards (mobile)
  - Customer standards and policies
  - The technical choices which constrain the system
  - Existing hardware, current platforms and equipment
  - Any software givens (packages, middleware, legacy, and so forth)
  - Any data givens (interfaces and access to customer data on specific existing systems)
  - Database management system to be used
  - Network configurations

9.5 Examples

A system either satisfies its functional requirements and constraints or it doesn't; the choice is binary. In contrast, non-functional requirements are characterized by a more continuous scale of measurements.

Non-functional requirements are expressed as attributes. Attributes describe a scale on which some property of the system can, in principle, be measured. The requirement may be that the value is above or below some target threshold or series of thresholds.

9.5.1 Example: Response time requirements

The following specific notation is commonly used when discussing the performance NFRs:

- Peak Day to Average Day ratio: The ratio of the peak day workload to the average day workload.
- Peak Hour Scaling Factor: The ratio of the workload over the notional peak hour compared with the workload averaged over the online day.
- Average Response Time: Defined as the sum of all sampled response times for a given interaction, divided by the number of measurements within the sample. Usually this is a measure of the typical response times that a user will see.
- Percentile Response Times (x%): The time within which x% of all responses complete.

Table 9-1 provides a sample set of generic response time bands for the most significant user-system interaction within the classes of transaction present in the target business system.

Table 9-1 End-to-end response time

	Target Average End-to-End Response Time (seconds)		
Frequency of Use	**High Frequency** (> 100 times/day)	**Medium Frequency** (>10 times/day but < 100 times/day)	**Low Frequency** (< 10 times/day)
Simple Transaction	*1 – 2*	*2 - 3*	*3 - 4*
Medium Transaction	*3 - 5*	*4 - 7*	*5 - 10*
Complex Transaction	*6 - 10*	*8 - 15*	*11 - 20*
Very Complex Transaction	*11 - 20*	*16 - 30*	*21 - 40*
Long-running	*> 20*	*> 30*	*> 40*

In this example, a number of different target response time bands were defined to reflect the range of processing complexity that existed within the target business system. Effort should be concentrated on delivering the fastest response to those processes used most frequently by the business system users.

9.5.2 Example: Throughput

The following specific notation is commonly used when discussing throughput NFRs:

- Peak Day to Average Day ratio: The ratio of the peak day workload to the average day workload.
- Peak Hour Scaling Factor: The ratio of the workload over the notional peak hour compared with the workload averaged over the online day.

If the workload is generated by users or processes at a number of different geographical locations, then the overall process workload should be broken out by location, as shown in this example.

Table 9-2 provides a sample set of business process volumetric requirements, corresponding to the Design Point workload (in this case, the peak day) for our target business system.

Table 9-2 Design Point workloads

Business process	**Design Point volumetric (Number per *peak* day)**	
	Site A	**Site B**
Update page through GUI	4,000	6,000
Update page through APG	10,000	35,000
Update page through PCI client	35,000	10,000
Re-initialize page through GUI	200	300
Re-initialize page through APG	0	0
Re-initialize page through PCI client	1,750	500
Create new page through GUI	0	1
Review page through GUI (including accessable alternatives)	40,000	60,000
Retrieve page from archive	40	80
Update pagination	1	1
Broadcast page to air	50,000	50,126

9.5.3 Example: Availability

Availability is frequently an important SLR. Table 9-3 gives an example of an availability specification. Availability requirements typically vary by use case; each row represents a collection of use cases with common availability requirements.

Table 9-3 Availability specification

Use Cases	Service hours	Special characteristics	Impact of loss: Critical/ High/Medium/Low				Fall back plan	Avail-ability Req.	Number of outages accept-able	Recovery Require-ment
			A few sec.	1-5 min	½-1 hour	3-4 hour		VH/H/M		
U1 [*mgmt stats*.] U3 [*maintain prices*]	08:00-18:00	Over time, impact will become more serious	-	L	L	M		H	1/month	Next day
U7 [*Set up account*] U8 [*Account Inquiry*] U9 [*Process Order*]	07:00-22:00	Peak period (end Feb to mid-April) needs highest availability	M	H (C for phone staff: U2 & U7)	H (C for phone staff: U2 & U7)	C	Users moved if possible Switch to backup facility	VH	2/year (major) 1/month (minor)	3-4 hrs to get back either normal or reduced service ½ hour recovery
U6 [*bulk mailing*] (batch)	24:00-07:00	Plan and download several days in advance	-	-	-	L		H	1/month	Next night
S5 [*fax*] S6 [*EDI*] S7 [*Email*]	almost 24 hours		-	L	M	M		VH	2/year (major) 1/month (minor)	Next day

10

Conclusions

This chapter summarizes the benchmark results and the lessons that we took away from the project. Topics included are:

- 10.1, “Test results” on page 218
- 10.2, “Lessons learned” on page 220

10.1 Test results

In this benchmark project we performed tests to achieve KPI on four functional modules: RPM, BF, MRP and ML, using "real" customer data and a high-volume make-to-order scenario, with an SAP APO configuration.

All KPIs were achieved according to the customer requirements and we can also highlight that the hardware layer demonstrated its strength to achieve these KPIs.

First, we demonstrated that the task of exploding 375,000 fully configured orders for a three-level bill of material (RPM) is possible within a small time window. The resulting run time was 15 minutes for the simple scenario and 22 minutes for the realistic scenario.

Second, with the observed performance measurements, we reached a total run time of just 25.4 minutes for the R/3 MRP run, which represents 84.8 percent of the KPI of 30 minutes that was requested by the customer.

Third, with an adjustment of the one-step backflush, we achieved a run time of 27.3 minutes, and we found out that with an adjustment of the two-step backflush, we will be able to achieve an even better result. Unfortunately, the limited time for this project prevented us from testing this in detail.

Finally, we achieved an execution time of 3 hours 16 minutes for running the material ledger job that calculates the actual cost of each material and reflects it in logistics flow for the next month, based on the accumulated information for the current month of activity.

As displayed in Table 10-1, it is interesting to notice the major improvement achieved between the initial runs and the final ones.

Table 10-1 KPI summary

KPI Summary	RPM	MRP	Backflush	Material Ledger
1st test	**59 min**	**128 min**	**122 min**	**11 h 08**
Target	**30 min**	**30 min**	**30 min**	**4 hours**
Best Achieved	**15 min**	**26 min**	**27 min**	**3 h 16**

In some cases, we improved the results by more than 400 percent. Assuming that all the changes applied to our benchmark configuration are fully applicable to a production system, it clearly shows that a personalized approach and an optimized tuning are required to minimize the impact on the client business activities.

The graph in Figure 10-1summarizes the database size growth for the ECC environment. During this project period the data size grew from 0.35TB to 0.66TB during scenario 1 activities and up to 1.1TB when running ML tests with scenario 2. Most of the associated tablespaces are compressed, but there may be some remaining which are not compressed because we did not run the reorganization of the whole DB.

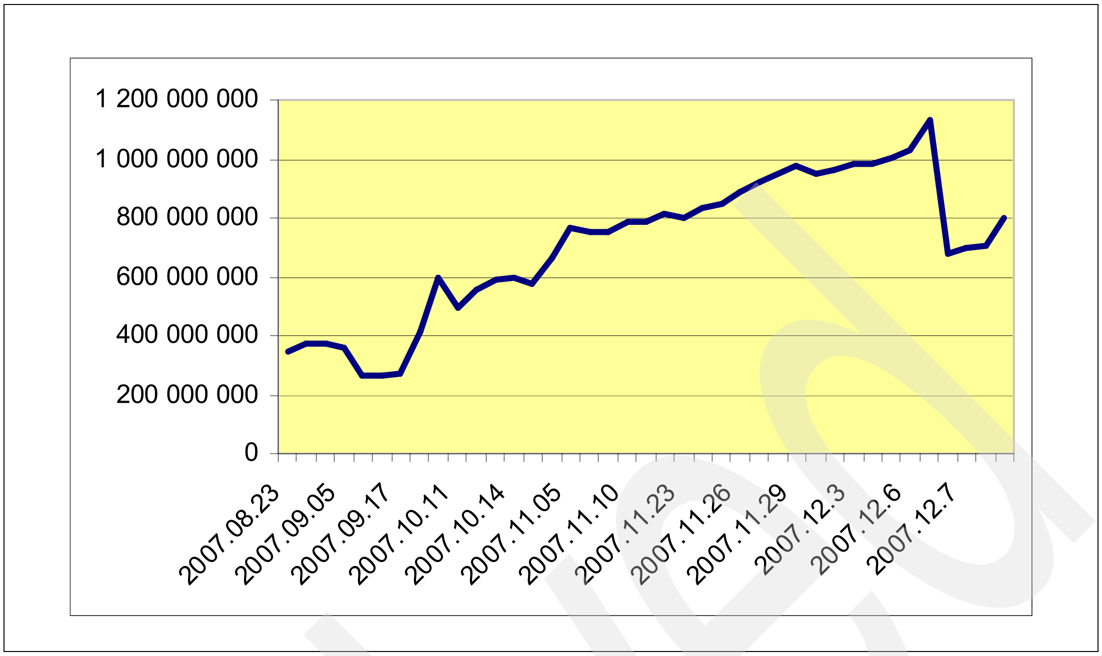

Figure 10-1 ECC database size growth

This growth must be carefully monitored in a production environment to anticipate storage space requirements, and to keep control of the production data growth.

For SAP this was a technical proof of concept, a proof of the strength and scalability of the solution. For IBM this was a demonstration that its hardware and software capabilities can easily accommodate the taxing requirements of an automotive client without any issues.

For our client it was a demonstration of:

- The joint capability of SAP and IBM.
- The fact that the proposed design will work within their defined parameters, especially concerning performance.
- The technologies that were used can easily adapt to future needs, thus growing with the customer.

Finally, as a summary, Table 10-2 on page 220 presents the key resource utilization parameters for various test runs that complied with the KPIs.

Table 10-2 Resource utilization and hit patio

CPU Utilization

CPU per system			RPM run10		MRP run14		Backflush run14		Material Ledger run9	
			Ave (%)	Max (%)	Ave (%)	Max (%)	Ave (%)	Max (%)	Ave (%)	Max (%)
ECC	AP #1	10	0,2	1,3	26,8	44,1	12,6	25	31,1	81,3
	AP #2	10			39	65,8	39,8	62,5	48,3	100
	AP #3	10			39,9	64,7	33,8	53,6	54,8	99,7
	DB	8C 6C-5I	n.a.	n.a.	54 (91,91)	71,9	9,3 (*) (85,18)	17,9 (*)	34 (95,73)	87,8
APO	AP	32	7,1	52,7	4,8	18,5	0,9	2,9	0,1	0,3
	DB	8C 6C-5I	6,0 (99,6)	18,2	1,7 (97,86)	8,4	n.a.	n.a.	n.a.	n.a.
	Live cache	32	5,5	35	2,1	7,4	0,7	1,6	0,4	1,8

(*) run10 used

Memory Utilization

Memory per system			RPM run10		MRP run14		Backflush run14		Material Ledger run9	
			Ave GB	Max GB	Ave GB	Max GB	Ave GB	Max GB	Ave GB	Max GB
ECC	AP #1	40/60	8	8	10	10	11	11	14	19
	AP #2	40/60			7	7	13	14	14	21
	AP #3	40/60			7	7	13	14	14	19
	DB	24	5-6	5-6	5-6	5-6	5-6	5-6	5-6	5-6
APO	AP	58	17	35	17	17	16	16	58	58
	DB	24	5-6	5-6	5-6	5-6	5-6	5-6	5-6	5-6
	Live cache	196	68	78	60	60	60	60	22	22

10.2 Lessons learned

Through the execution of this proof of concept, we proved the strength and scalability of the solution. We also demonstrated that the proposed architecture, with the application server on System p and database server on System z, is the optimum configuration to fulfill the client's requirements in terms of performance.

The SAP tuning features applied and the code optimization developed were possible thanks to the underlying infrastructure based on DB2. We cannot guarantee that similar performance results can be achieved on a full UNIX configuration, using the same data structure and similar scenarios.

Based on the specificity of the client's data structure and volume, and considering the necessary link between the new SAP production environment and the existing legacy systems, we strongly recommend using the architecture tested in this PoC as the baseline for the future production system.

Furthermore, the team work among the client, SAP, and IBM was crucial in providing the needed skills in order to:

- Understand the business context and its associated business processes
- Convert the needs into application steps
- Optimize the executed tasks on the underlying infrastructure
- Tune the database

These factors were all important to the success of this benchmark test project.

Part 3

Appendixes

In this part we provide details for setting up and monitoring various configurations. The topics included are:

A

Benchmark reporting

In this appendix we describe the reports and reporting tools used during the test.

The reports can be divided into four categories:

- SAP reports
- AIX reports
- z/OS reports
- DB2 reports

The topics presented in this appendix are:

- A.1, “SAP reports” on page 224
- A.2, “Transactions for technical monitoring” on page 228
- A.3, “AIX monitoring tools” on page 240
- A.5, “zOS and DB2 reports” on page 259

A.1 SAP reports

This section describes the tools and procedures SAP usually applies when doing performance tests and which were also applied to the performance tests for this specific customer.

A.1.1 Monitoring

Procedure

The execution phase of a performance test consists of several steps. Once the data and system have been provided, the load generation begins. A single-user test is followed by a multi-user test consisting of several steps until the maximum load has been reached. The load generation is then followed by an analysis phase.

When monitoring the test you must distinguish between two separate monitoring concepts, both of which are generally used and complement each other:

- Real-time monitoring

  Real-time monitoring means that the running load test is monitored using appropriate monitoring transactions. In addition to monitoring load tests using the SAP CCMS performance monitors, the currently available load test tools also provide functions for real-time monitoring and for the collection of performance data. The available options differ from tool to tool. For example, SAP LoadRunner by Mercury provides the correlation of response times, CPU utilization, script errors, and the test progress.

- Retroactive analysis

  A retroactive analysis automatically evaluates collected information after the end of the test execution. This also includes the (manual) correlation of values from the execution log of the load test tool with the values collected in the monitoring process.

Real-time monitoring fulfills an important function in allowing you to interfere with the test process at an early stage in order to directly obtain first-hand information. Moreover, you can use the information provided by the current monitoring process to configure the data collection for the subsequent, more detailed analysis. The retroactive analysis is used to more exactly narrow down performance problems and to provide general information on the system behavior. Table A-1 provides a list of transactions for real-time monitoring and for retroactive analysis.

Table A-1 Transactions for real-time monitoring and retroactive analysis

Transaction for the current monitoring process	Transaction for the retroactive analysis
AL08 - List of all users logged on	SM21 - Display system log
OS07 - Operating System Monitor	ST04 - Database monitor
SM04 - User List	ST05 - Performance analysis
SM12 - Lock entry list	ST22 - ABAP Runtime Error
SM66 - Global work process overview	STAD - Display statistical records
	OS07 (via OS Collector)

Consider this example: The run-time monitors of a load test tool indicate unusually long response times. The following actions are taken already during the running load test generation:

- Use of Transaction STAD

  The analysis shows that the database time represents a large portion of the total run time.

- Use of the operating system analysis of the database server via Transaction ST06 in order to check the hardware utilization of the database server (CPU, memory)

  The analysis returns a high degree of CPU utilization of the database server.

Once the load generation has finished, the following measures are taken for a more exact analysis:

- Use of an SQL trace

  An SQL trace is used to analyze the database accesses that have been identified as limiting the performance at the SQL level. In this way you can obtain information on the cause of the performance problem.

- Database analysis

  The status of the database is checked using ST04.

- Program analysis

  A more detailed analysis of the database access using Transaction ST05, based on the traces that have been previously created. In addition, you can retroactively create traces of individual transactions. This process does not consider the load profile for the individual trace, but it still plays a very important role in the detailed analysis of individual transactions.

- Evaluation of the trace files

  This is done to identify an expensive SQL statement or a missing index on a table.

As a classical client-server application, the SAP system consists of three layers in the SAP GUI for Windows environment (Figure A-1). Each of those layers could contain the cause for the long response times we observed: the client, the application server, or the database server.

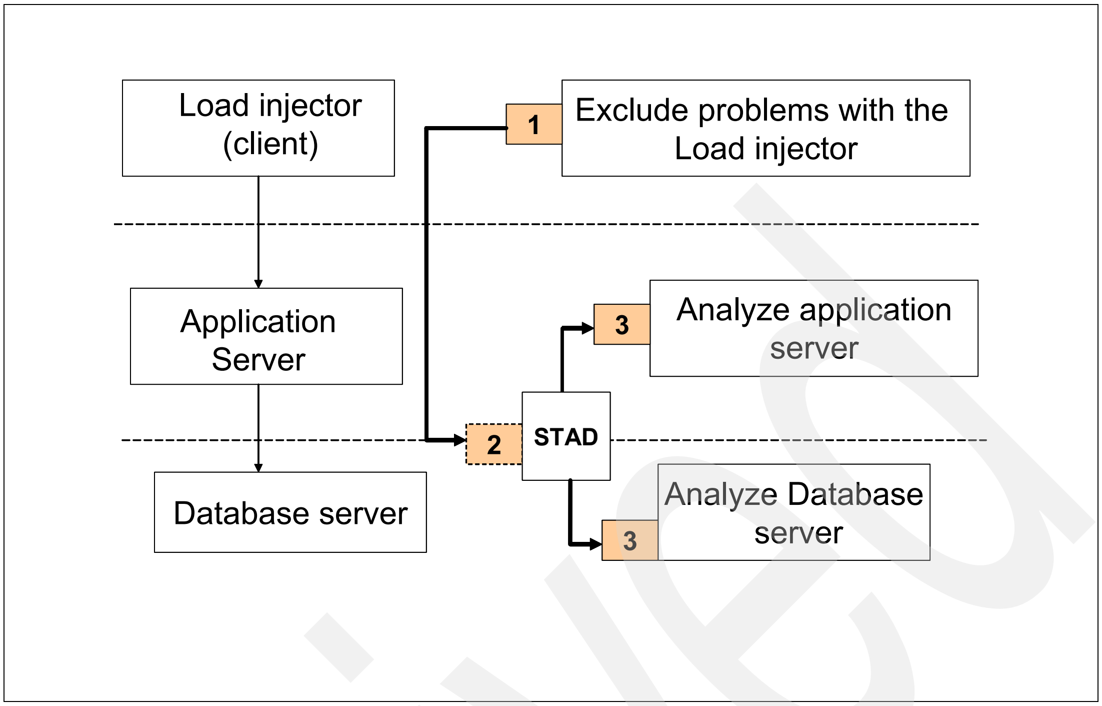

Figure A-1 Client-server layer

The first step is to rule out the load injectors as the source of the problem. Because each front end processes only the requests of a single user, the SAP GUI is responsible for performance problems in the least of all cases. For this reason, the load test employs load injectors instead of real front ends. These load injectors process several SAP GUI instances simultaneously. In order to guarantee a correct execution of the performance tests, you must make sure that the load injectors themselves will not become the performance bottleneck. For this reason you should first check the utilization of the machines that are used as load injectors. To do that, you must observe the CPU utilization of the load injectors. Single CPU peaks of 100% give no cause for concern; however, plateaus at 100% indicate the overload of an injector. In this case you must reduce the number of virtual users on the injector. In cases of doubt, you must use additional load injectors.

If you were able to rule out the load injectors as sources of error, you should use the next step to find out whether the response time is primarily used within the application server or by the database. The most efficient way to do that is to use the SAP Workload Monitor (see A.2.4, “STAD: Transaction analysis” on page 231). This transaction enables you to analyze the response time of a transaction separated by database response time and processing time. If the response times of the database exceed an acceptable limit, you should continue with an analysis of the database.

If the response times of the database are acceptable, you can rule out the database server as the performance-limiting component. The next step then consists of analyzing the application server.

If, however, you have identified the database server as the performance-limiting component, you can use a number of additional options for diagnosis. First you should use the database monitor ST04 and, if possible, the operating system monitor OS07 to check whether the hardware of the database server is overloaded. The relevant parameters in this context are primarily the CPU utilization, memory usage, and hard disk activity (see A.2.8, “OS07/ST06: System monitor” on page 235). You can use Transaction SM66 to find out whether only one transaction or all transactions within the determined load profile are affected.

This results in the decision matrix shown in Table A-2.

Table A-2 Matrix for analyzing the database server

Observation	Symptom Database server overloaded	Symptom Database server not overloaded
Only one transaction is affected	Assumption: Inefficient or expensive SQL programming R consult with SQL developer	Assumption: Problem with DB locks, indices, or IO problem R consult with database expert
All transactions are affected	Assumption: Insufficient hardware capacity of the database server R consult with database and Basis expert	Assumption: IO problem or incorrect configuration of the database server R consult with database expert

The matrix presented here contains typical suspicious factors and the related recommendations for taking action. Although this matrix does not provide enough information to locate the cause of the problem or to remove it entirely, it is often sufficient to identify the right person to contact when a problem occurs.

The analysis of the application server should start with a check for a possible overload of the application server. To do that, you must use the operating system monitor OS07. Then you should again try to find out whether the long response times are a general problem of the system; that is, whether all transactions specified in the load profile are affected, or whether it is only one transaction that causes the problem. To do that, you can use Transaction SM66.

Table A-3 Matrix for Analyzing the Application Server

Observation	Symptom Application server overloaded	Symptom Application server not overloaded
Only one transaction is affected	Assumption: Inefficient or poorly performing ABAP programming R consult with ABAP developer	Assumption: Problem with locks queues, and the like R consult with development
All transactions are affected	Assumption: Insufficient hardware capacity of the application server R consult with Basis employee	Many problems possible; for example, network problems, locks, incorrectly configured work processes, or another configuration problem R further investigation required; consult with Basis and IT employees

If you use the operating system monitor OS07 and find out that the application servers are overloaded due to a high CPU utilization or insufficient memory, the next actions to be taken are easy. If the problem involves only one transaction, it is very likely that this transaction has been programmed inefficiently. In this case the possible symptoms would be that the largest part of the CPU time or of the memory is occupied by a single transaction. If that is the case, you should involve the development team. However, if the system load is equally generated by several transactions, you can usually solve the problem by adding more hardware in the application servers. Here you should involve employees from SAP Basis and members of the project management team to indicate and discuss an extension of the hardware.

If, however, the analysis using OS07 does not indicate any OS-side overload of the application servers, the next steps to be taken are somewhat more difficult. If only one transaction shows longer response times, it may be possible that locks are the cause of the problem. To verify this, you should use Transaction SM12. If the assumption is correct, you should involve the development team.

If the performance of all transactions in the load profile is unsatisfactory, there can be several causes that you should investigate sequentially.

Network diagnosis

You can either use appropriate network diagnosis tools or reduce the available bandwidth in order to prove an overload of the network. To do that, you should consult with an expert from the SAP Basis team.

A number of work processes that have been configured too small typically causes long dispatcher queues for waiting transactions. This can be uncovered using Transaction SM51. Another possible cause can be found in an incorrect size of the various buffers in the SAP system, in other words, the buffers are too small. Use Transaction ST02 to check the buffers. If you suspect a configuration problem in the application server to be the cause of the trouble, you should definitely consult with an expert from the SAP Basis team.

Conclusion

The most demanding step within a performance test consists of narrowing down performance problems and then finding out their causes. To do this, you should use the diagnostic tools provided by your SAP systems.

The first step to take when a performance problem occurs is to narrow down the components that limit the performance. You must determine whether it is the application, the database, the hardware, or another component that causes the long response times.

Once you have identified the performance-limiting component, you must further isolate the possible cause within the component and finally involve an expert to locate, analyze, and eliminate the cause.

A.2 Transactions for technical monitoring

This section provides an overview of the transactions introduced in the previous methodological section. The sequence of the descriptions presented here is based on the alphabetical order of the technical transaction names.

A.2.1 AL08: List of all users logged on

Transaction AL08 enables you to obtain an overview regarding the distribution of users across the available instances (see Figure A-2). You can, for example, assess the effectiveness of a load balancing process.

Transaction AL08 provides the following options:

- Display of the number of test users already logged on
- Check whether test users are still logged on but are no longer active, for example between two test runs

- Identification of unwanted users in the system that represent a risk because they distort the test results and must therefore be avoided

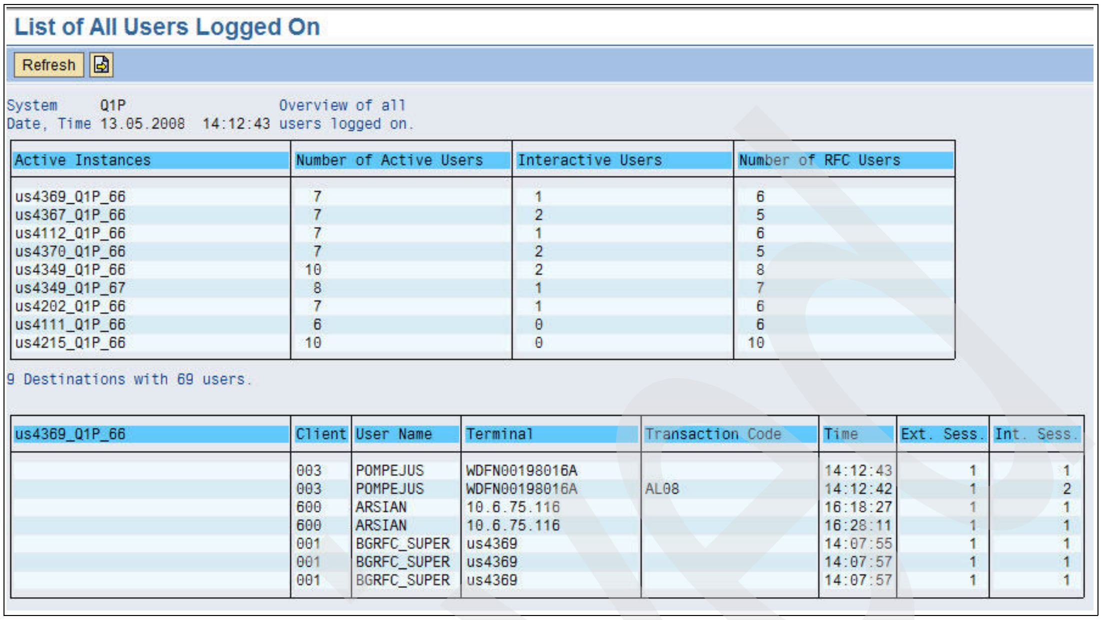
List of All Users Logged On

Refresh

System Q1P Overview of all
Date, Time 13.05.2008 14:12:43 users logged on.

Active Instances	Number of Active Users	Interactive Users	Number of RFC Users
us4369_Q1P_66	7	1	6
us4367_Q1P_66	7	2	5
us4112_Q1P_66	7	1	6
us4370_Q1P_66	7	2	5
us4349_Q1P_66	10	2	8
us4349_Q1P_67	8	1	7
us4202_Q1P_66	7	1	6
us4111_Q1P_66	6	0	6
us4215_Q1P_66	10	0	10

9 Destinations with 69 users.

us4369_Q1P_66	Client	User Name	Terminal	Transaction Code	Time	Ext. Sess.	Int. Sess.
	003	POMPEJUS	WDFN00198016A		14:12:43	1	1
	003	POMPEJUS	WDFN00198016A	AL08	14:12:42	1	2
	600	ARSIAN	10.6.75.116		16:18:27	1	1
	600	ARSIAN	10.6.75.116		16:28:11	1	1
	001	BGRFC_SUPER	us4369		14:07:55	1	1
	001	BGRFC_SUPER	us4369		14:07:57	1	1
	001	BGRFC_SUPER	us4369		14:07:57	1	1

Figure A-2 Users logged on list

In a configuration, the dialog users should not work on the central instance, but on the dialog instances. Transaction AL08 enables you to check during a load test whether the load balancing process regulates this as it should. It is advisable to update Transaction AL08 regularly during the test, and also to check the transaction before the test begins. At this point you should note that depending on the configuration being used, a load balancing process can itself become a problem during the load test. Usually, this happens if the period of time that elapses between the logon of new users is significantly shorter than the time interval based on which the load balancer checks the user distribution.

Note: If the application servers are suddenly overloaded, and this cannot be explained on the basis of the test run up until that point, you should verify whether users foreign to the performance test are active in the system. Unfortunately, it frequently happens in practice that in spite of early announcements of the performance test the system is still used for other purposes during the test run. The result is that the performance test becomes distorted and useless.

A.2.2 SM12: Lock entry list

If you suspect locks in the application server to be the cause of the problems, you should use Transaction SM12 to check the locks that have been set. If your suspicion turns out to be true, you should consult with the development team.

Transaction SM12 provides an overview of active locks in the system (see Figure A-3). For each lock the following information is displayed:

- Client and user that set the lock
- Start time of the transaction that set the lock

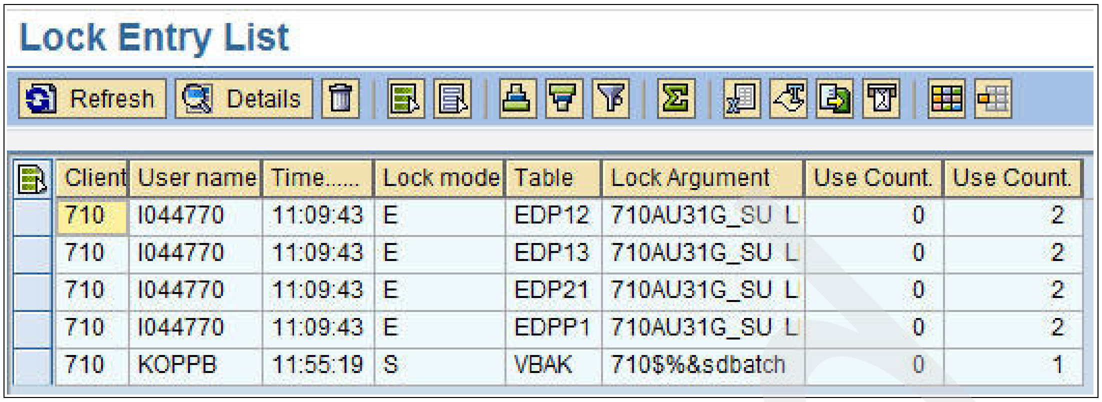

	Client	User name	Time......	Lock mode	Table	Lock Argument	Use Count.	Use Count.
	710	I044770	11:09:43	E	EDP12	710AU31G_SU L	0	2
	710	I044770	11:09:43	E	EDP13	710AU31G_SU L	0	2
	710	I044770	11:09:43	E	EDP21	710AU31G_SU L	0	2
	710	I044770	11:09:43	E	EDPP1	710AU31G_SU L	0	2
	710	KOPPB	11:55:19	S	VBAK	710$%&sdbatch	0	1

Figure A-3 Overview of active locks in the Lock Entry List

You should check the lock entry list whenever you think that performance problems are caused by locks that have not been deactivated. This is the case if a specific transaction has long response times in the application server whose hardware is not overloaded.

In that case you must check the locks using Transaction SM12. If the locks are indeed the cause of the problem, the locks are probably set on the tables required by the slow process. In this case you should involve the development team.

A.2.3 SM66: Global work process overview

The global work process overview provides you with all the relevant information on the active work processes (Figure A-4). You can use this information to collect first hints on problems with locks and in order to check which transactions are affected by performance problems.

The global work process overview provides the following options:

- Monitoring the load of the work processes for all active instances across the entire system
- Identifying locks in the database
- Determining transactions that cause long system response times

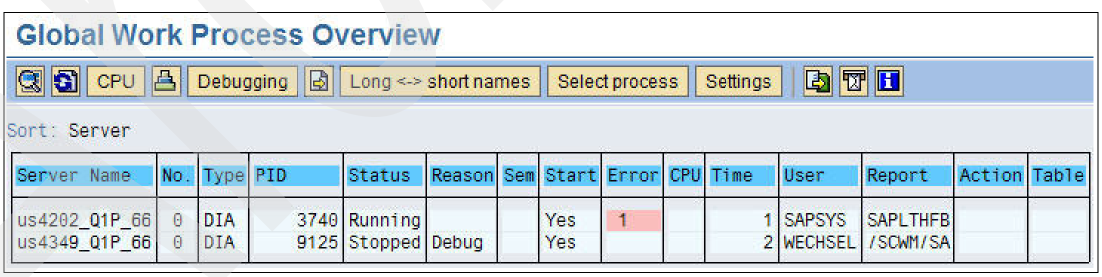

Server Name	No.	Type	PID	Status	Reason	Sem	Start	Error	CPU	Time	User	Report	Action	Table
us4202_Q1P_66	0	DIA	3740	Running			Yes	1		1	SAPSYS	SAPLTHFB		
us4349_Q1P_66	0	DIA	9125	Stopped	Debug		Yes			2	WECHSEL	/SCWM/SA		

Figure A-4 Global work process overview

The following information is displayed in the global work process overview:

- The type of a process

  Dialog process: DIA
  Update: UPD
  enqueue: ENQ
  batch: BTC
  spool: SPO
  V2 update: UP2

- The status of the individual application servers

  Here, the hold (wait) status is of particular interest. If a process is in “hold” status, the reason is specified in the subsequent row.

- The CPU time and the execution time

  Both times provide you with information about how long the process has already been busy with its current task. You can use this information to determine whether an identified performance problem affects only one or several programs.

A.2.4 STAD: Transaction analysis

Identifying problems

In most cases, calling the transaction analysis (Transaction STAD) represents the first step towards narrowing down a performance problem. You can use the transaction analysis to identify programs with long run times. In addition, Transaction STAD also subdivides the total response time into its subcomponents, CPU time and database time, and can therefore be used to narrow down long response times on the application servers or the database server.

Programs with long response times

Apart from assigning a long response time to an application or database server, you can also use Transaction STAD to obtain an overview of the programs affected by a performance problem (see Figure A-5). To do that, check the program names (Name column) of the programs with long response times. This way you can get information as to whether the performance bottlenecks are caused by one or several programs.

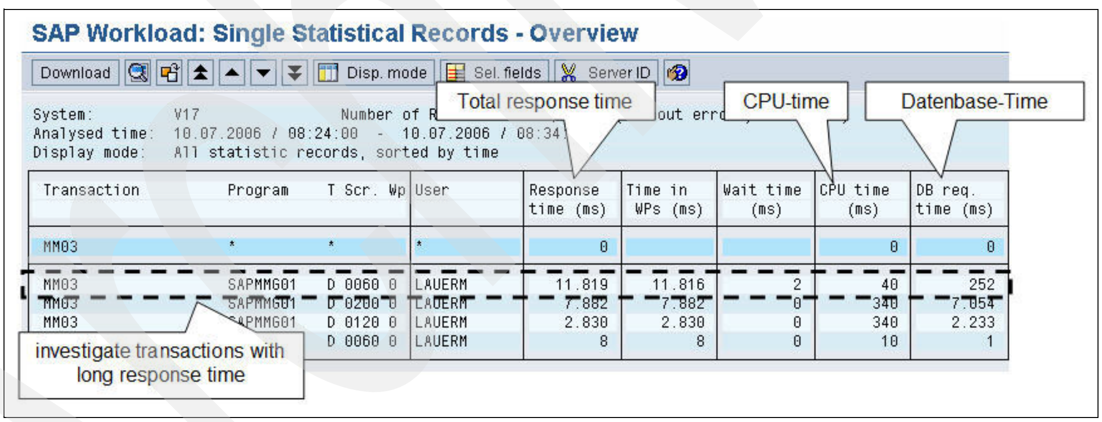

Figure A-5 Evaluating the response time of a transaction

A.2.5 ST02: Overview of SAP buffers

Transaction ST02 displays an overview of the buffers of the application server (see Figure A-6 and Figure A-7). It is important that the buffers are large enough to allow a reloading of objects without deleting existing objects, and those objects that are still needed from the buffer. If that is the case, the number of reloads from the database increases, which in turn affects system performance.

Limits

In this context, two limits must be observed. One way to measure the effectiveness of the buffer sizing is to determine the hit ratio. In general, the hit ratio should be at least 98%.

Swaps should not occur. An exception is the program buffer, in which swaps can be accepted.

System: dsni02_V17_00 Tune summary
Date & time of snapshot: 27.02.2006 13:52:31 Startup: 25.02.2006 15:51:44

Buffer	Hitratio [%]	Allocated [kB]	Free space [kB]	Free space [%]	Dir. size Entries	Swaps	Database accesses
Nametab (NTAB)							
Table definition	99,63	7.593	5.826	89,82	35.332	0	4.111
Field description	99,41	54.546	42.187	79,56	35.332	0	4.207
Short NTAB	99,45	5.890	5.105	96,27	8.833	0	872
Initial records	99,80	11.193	9.943	93,75	8.833	0	2.123
Program	99,87	500.000	187.721	38,74	125.000	0	15.060
CUA	99,94	10.000	8.186	93,37	5.000	0	127
Screen	99,85	20.000	15.991	80,43	2.000	0	413
Calendar	100,00	488	369	77,20	200	0	146
OTR	100,00	4.096	3.603	100,00	2.000	0	0
Tables							
Generic key	99,85	60.000	40.103	68,95	10.000	0	4.275
Single record	95,72	20.000	13.399	67,35	500	0	10.776
Export/import	88,42	4.096	3.198	88,76	2.000	0	0
Exp./Imp. SHM	81,48	4.096	3.577	99,28	2.000	0	0

Figure A-6 Overview using ST02 (Lower part)

System: dsni02_V17_00 Tune summary
Date & time of snapshot: 10.01.2006 12:37:22 Startup: 07.01.2006 15:24:58

SAP memory	Current use [%]	Current use [kB]	Max. use [kB]	In memory [kB]	On disk [kB]	SAP cursor cache	Hitrat [%]
Roll area	0,17	452	1.256	262.144	0	IDs	99,9
Paging area	0,03	84	28.440	128.000	134.144	Statements	94,0
Extended Memory	38,22	74.752	167.936	195.584			
Heap Memory		0	0				

Call statistics 794 tables buffered	Hitratio [%]	ABAP/4 Requests	Processor Fails	Database Total calls	Database AvgTime[ms]	Database Rows affected
Select single	99,40	1.161.763	404.934	24.303	0,000	756.829
Select	75,17	721.615	0	434.387	0,000	409.803
Insert		5.224	3.719	5.579	0,000	11.966
Update		495	375	556	0,000	1.548
Delete		9.148	5.769	9.221	0,000	11.471
Total	87,29	1.898.245	414.797	474.046		1.191.617

Figure A-7 Buffer overview using ST02 (Upper part)

Use Transaction ST02 if you suspect an incorrect configuration of the buffer sizes to be the cause of the problem. If that is true, consult with an SAP Basis expert to optimize the buffer configuration.

A.2.6 ST04N: Database overview

Transaction ST04N (previously ST04) provides an overview of the database status. ST04N is a database-specific transaction (Figure A-8).

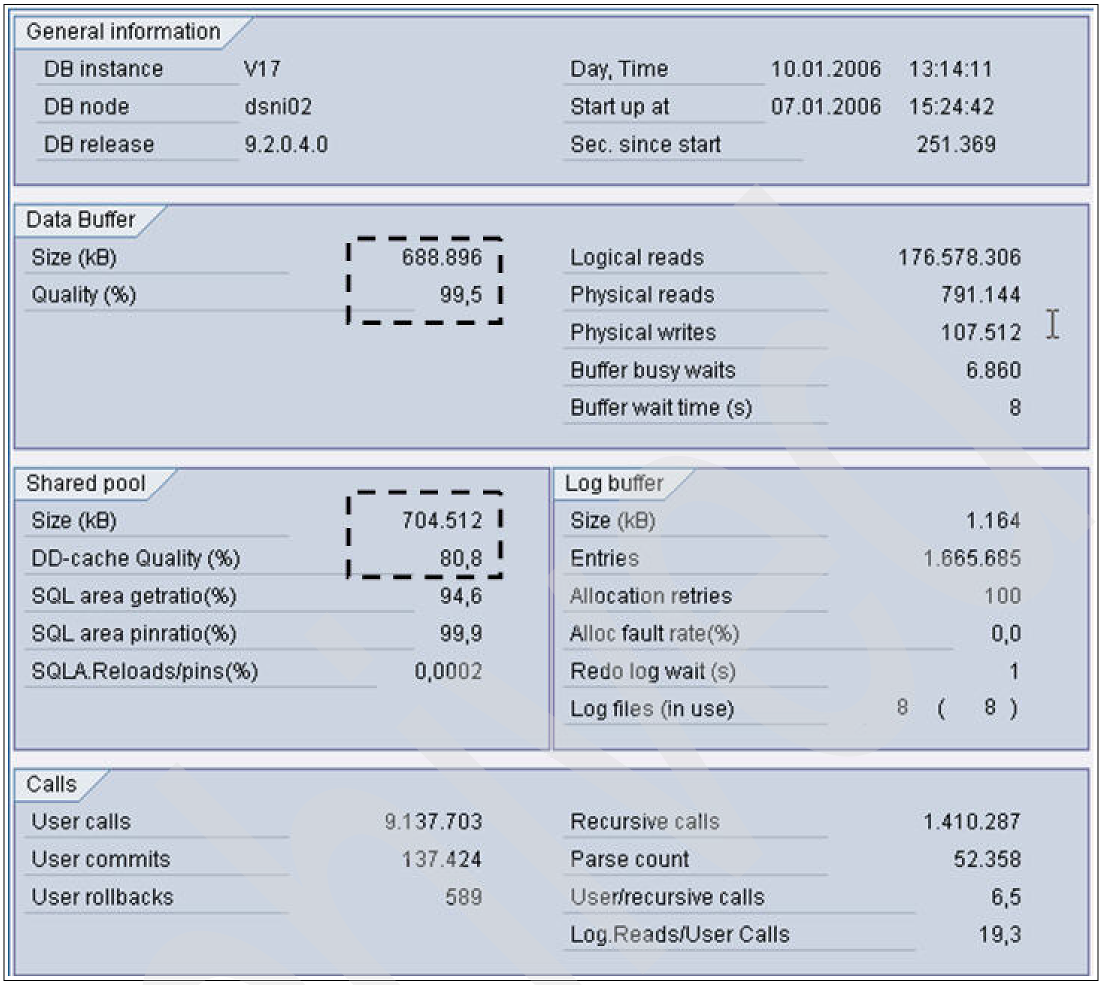

Figure A-8 Database overview using ST04N

Vendor-dependent values

The database overview (Transaction ST04N) offers a general overview of the status of the database. Here, the details on the cache size and quality of the data buffers and of the shared pool are of specific interest. If these values range below the reference values recommended by the vendor, consult with a database expert to adjust the configuration. Because these values depend on the database being used, we cannot provide reference values here.

A.2.7 ST05: SQL trace analysis

The SQL trace analysis (ST05) enables you to evaluate the communication between the application server and database server. In particular you can use the option to directly jump to the generating ABAP code to clarify whether an SQL statement with a long run time was created by a custom development. This question represents an important step towards preparing the optimization of the database access by the development team. Thus transaction ST05 is a valuable utility that allows you to carry out first steps in the analysis of a performance problem on the database server.

Process

An SQL Trace process consists of the following four steps (Figure A-9):

1. Start SQL Trace using ST05

2. -n.Generate load
3. -n.Stop SQL Trace
4. .-n.Analyze the trace

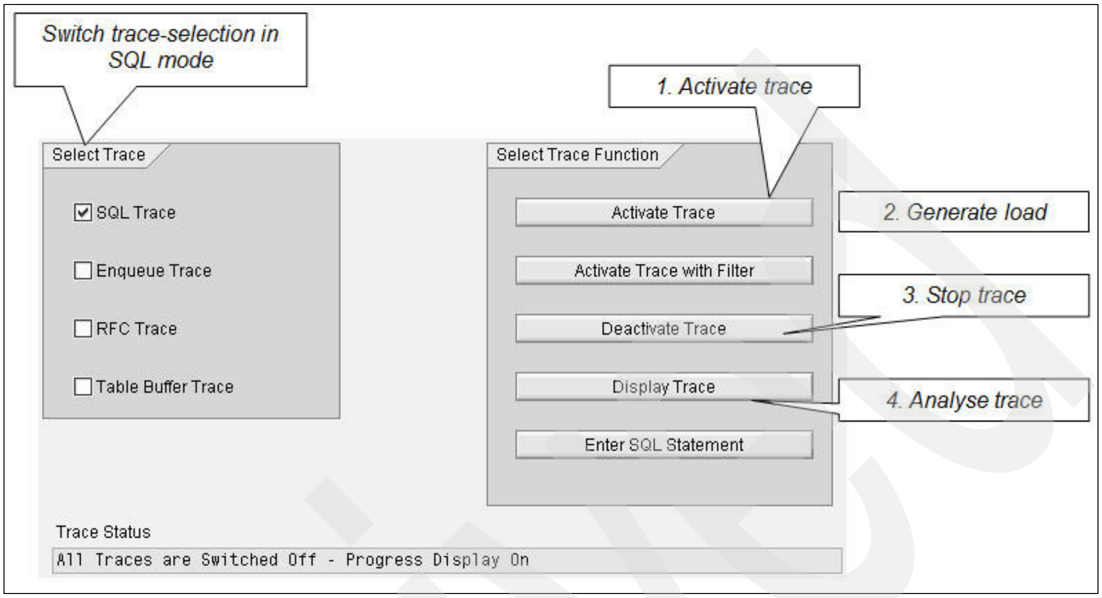

Figure A-9 Menu for the trace selection (ST05)

The analysis of the SQL trace provides a lot of interesting information. First of all you can view the response times of the individual queries in the Duration column, which enables you to identify particularly slow queries. The overview of individual response times usually helps you to answer the question whether or not a performance problem exists within the database.

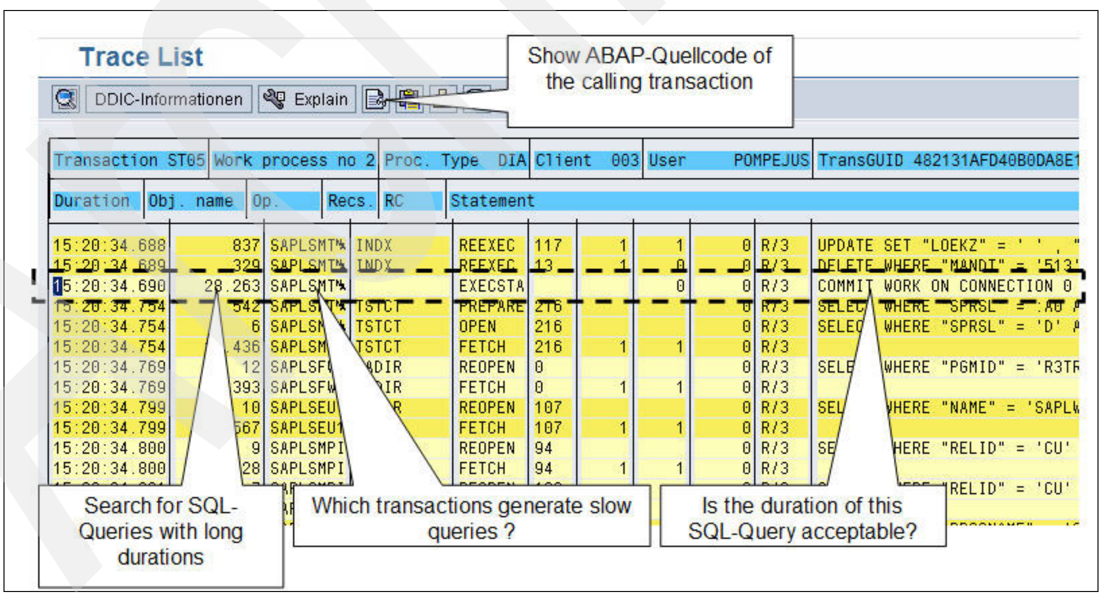

Figure A-10 Analysis of Database Response Times

If that is the case, the analysis of additional information in the SQL trace (Figure A-10) can then provide a first hint about the location of the problem. The Program column displays the name of the ABAP object that generated the SQL call. Use this to determine whether all slow queries are caused by the same transaction or if they are distributed across several transactions.

Jump to the source code

For further analysis you can jump directly from a selected query into the ABAP source code. This is particularly useful if you want to check whether the query in question was generated by custom development or by a component of the standard SAP version.

A.2.8 OS07/ST06: System monitor

Transactions OS07/ST06 query OS-specific performance parameters without using generic tools of the operating system. You can use transaction OS07 on the application or database server to check whether the hardware of the respective system is overloaded. The most important indicators in this context are the CPU utilization, memory usage, and hard disk activity.

Start with transaction OS07

In the context of a performance test, transaction OS07 represents a good starting point for the analysis of performance problems. The transaction enables you to monitor an instance from the point of view of the operating system. For this purpose, you should first select an instance that you want to monitor as shown in Figure A-11.

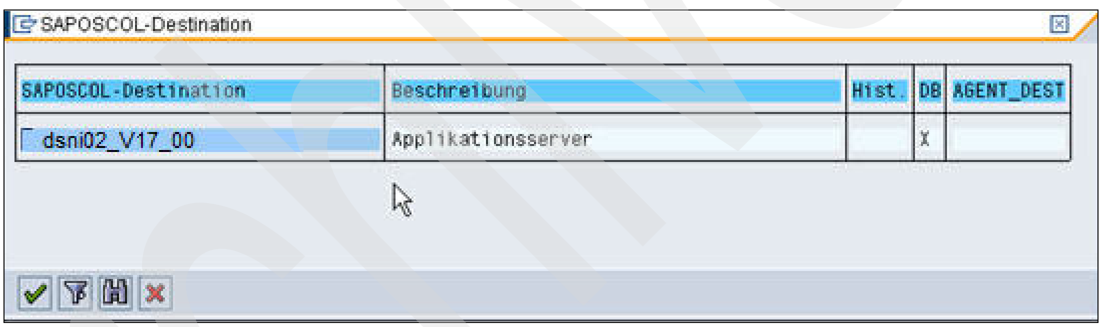

Figure A-11 Selecting the application server

The system then displays the OS monitor (transaction code ST06) that provides information on resources made available by the operating system (see Figure A-12). These resources comprise the following:

- Virtual memory
- Physical memory
- CPU time
- Storage space in the file systems of the hard disks

The OS monitor helps you identify bottlenecks in these resources. For this purpose, it displays a number of performance indicators:

- Average load and utilization of the CPU
- Memory utilization
- Hard disk utilization
- Network activity

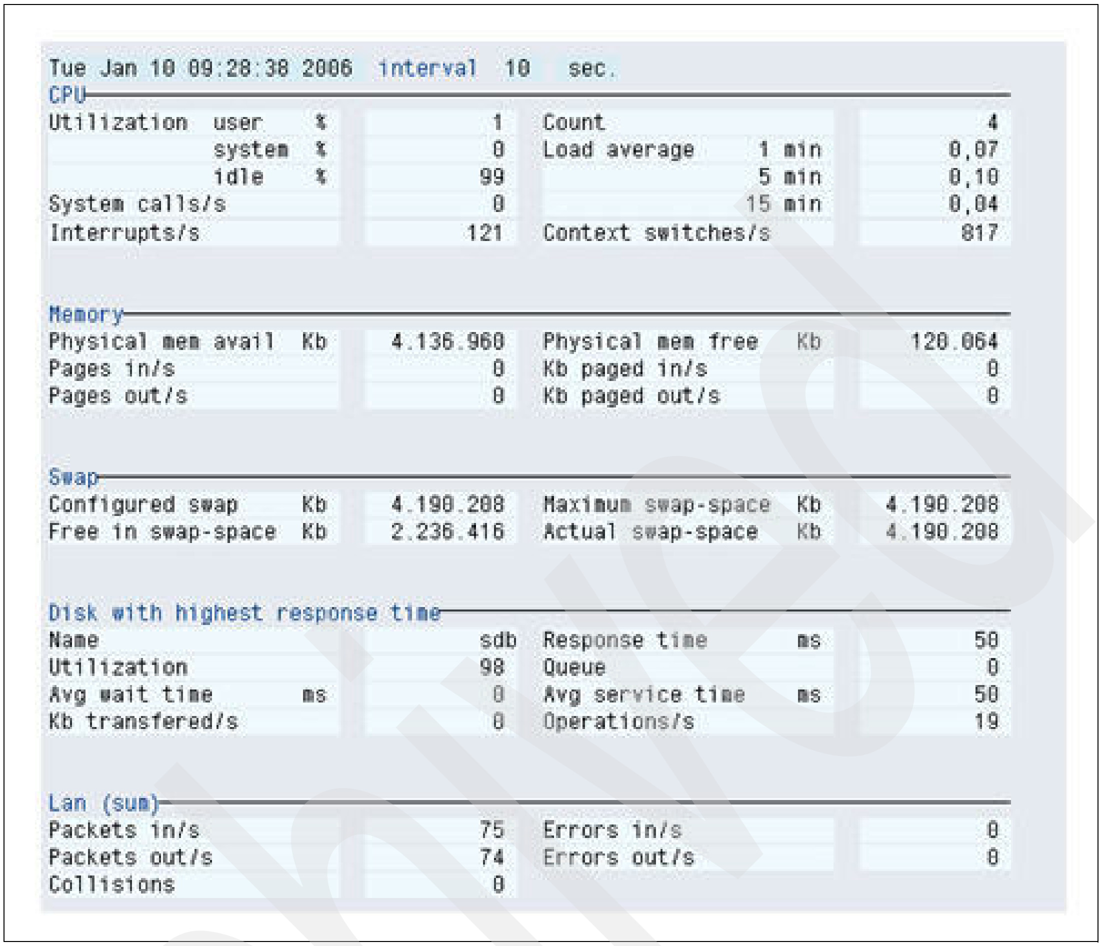

Tue Jan 10 09:28:38 2006 interval 10 sec.

CPU

Utilization user %	1	Count	4
system %	0	Load average 1 min	0,07
idle %	99	5 min	0,10
System calls/s	0	15 min	0,04
Interrupts/s	121	Context switches/s	817

Memory

Physical mem avail Kb	4.136.960	Physical mem free Kb	120.064
Pages in/s	0	Kb paged in/s	0
Pages out/s	0	Kb paged out/s	0

Swap

Configured swap Kb	4.190.208	Maximum swap-space Kb	4.190.208
Free in swap-space Kb	2.236.416	Actual swap-space Kb	4.190.208

Disk with highest response time

Name	sdb	Response time ms	50
Utilization	98	Queue	0
Avg wait time ms	0	Avg service time ms	50
Kb transfered/s	0	Operations/s	19

Lan (sum)

Packets in/s	75	Errors in/s	0
Packets out/s	74	Errors out/s	0
Collisions	0		

Figure A-12 Querying The Status of the OS Using Transaction OS07

You should compare the displayed values with the values you have seen in your own experience. The CPU utilization (CPU Utilization user) should not constantly range above 60 to 70% per CPU. If the degree of utilization is higher, it is very likely that the application server is overloaded. You should then try to identify the process that causes the load.

The swap memory on the hard disk should not be overtaxed. If the system shows a high activity of virtual memory management (Memory Pages in/s on Windows NT®, Memory Pages out/s on UNIX systems), the application server has insufficient memory. As a rule of thumb, it is recommended that not more than 20% of the physical main memory should be swapped per hour. Thus, in a system with 4 GB of memory, swap rates of up to 200 KB/second are acceptable. But in general you can say the smaller the swap rate, the better.

The utilization of the slowest hard disk (Disk with highest response time value in the Response time group) should range below 50%. If that is not the case, the hard disk may be overloaded.

The network should not receive any erroneous packets (LAN Errors in/s and LAN Errors out/s should be 0). If it does, there can be a problem in the area of the network hardware. In this case, you should consult with your network administrators to further narrow down and analyze the cause of the problem. The same goes for packet collisions (LAN Collisions greater than 0).

A.2.9 SM04: User list

Transaction SM04 (User List) provides an overview of logged-on users for an instance as well as of their cumulative memory usage (Figure A-13). The transaction also allows you to monitor the memory usage of a user in real time. This way you can easily identify the originator of a memory bottleneck.

Modes per user

In addition, you can use a terminal identification to find out how many modes are currently in use by a single user, and where exactly in the SAP system that user has logged on. Similar to transaction AL08, you can also identify the point in time at which the last end user interaction with the SAP system occurred.

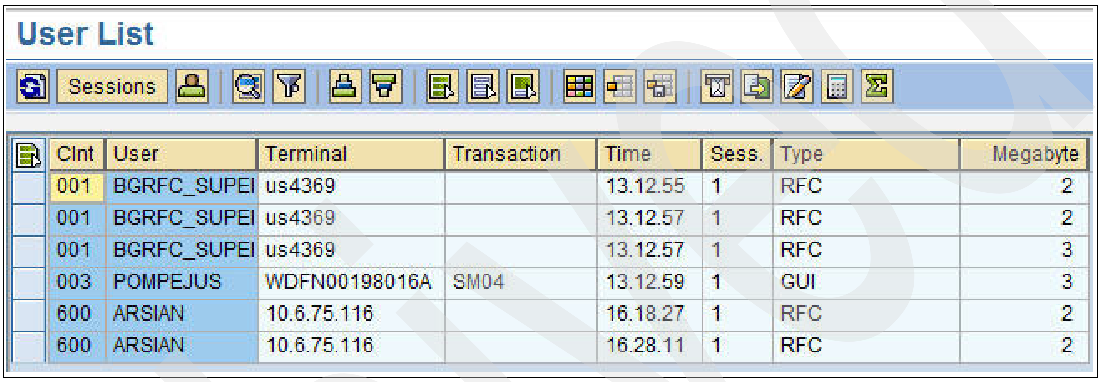

Clnt	User	Terminal	Transaction	Time	Sess.	Type	Megabyte
001	BGRFC_SUPEI	us4369		13.12.55	1	RFC	2
001	BGRFC_SUPEI	us4369		13.12.57	1	RFC	2
001	BGRFC_SUPEI	us4369		13.12.57	1	RFC	3
003	POMPEJUS	WDFN00198016A	SM04	13.12.59	1	GUI	3
600	ARSIAN	10.6.75.116		16.18.27	1	RFC	2
600	ARSIAN	10.6.75.116		16.28.11	1	RFC	2

Figure A-13 Viewing the User List via Transaction SM04

If the front end being used is UNIX-based, the terminal name matches the display variable of the front-end process. In the case of a Windows or OS/2® front end, the terminal name matches the name of the machine on which the front end was started.

> **Note:** When setting up the test scenarios, make sure that the test users are not arbitrarily distributed across processes or transactions to be tested. Instead, bundle them into groups of users. This facilitates the retroactive restriction of performance problems to individual transactions or programs.

A.2.10 SM21: System log

Transaction SM21 (System Log) allows you to analyze system or error messages during or after a test run (see Figure A-15).

You can use the input fields in the Selection group to define a filter (see Figure A-14). The most frequently used filters restrict the period in which a message occurred (From date/time and To date/time). Alternative filter criteria refer to the user, the transaction code, and the SAP process. In order for a change to the filter settings to take effect, you must use the Reread System Log view.

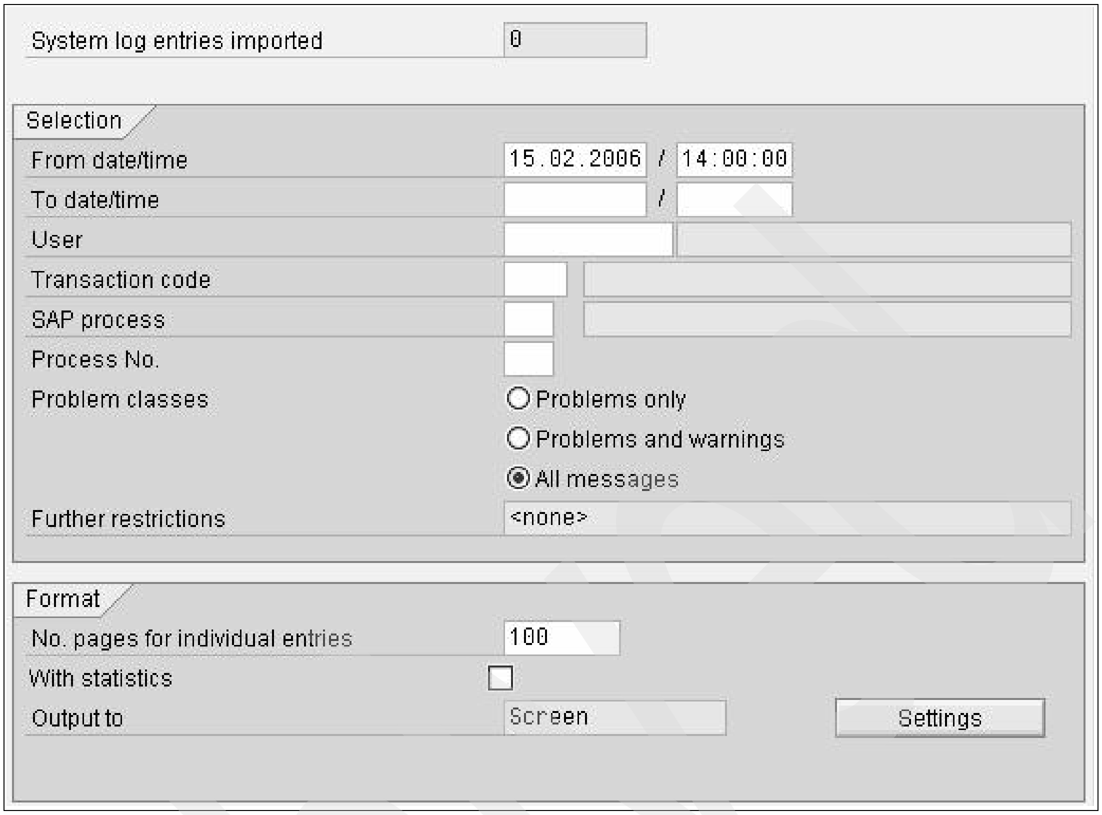

Figure A-14   Setting filter criteria in transaction SM21

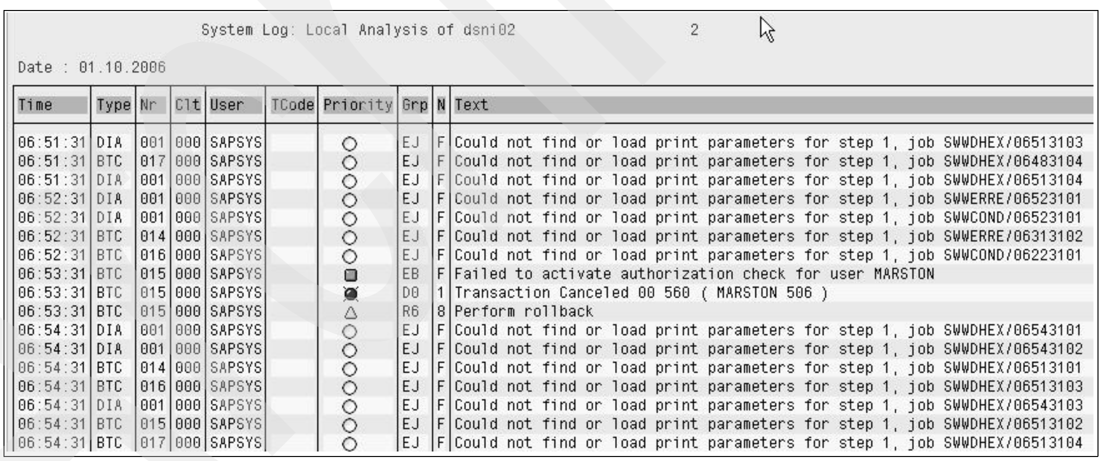
System Log: Local Analysis of dsni02   2

Date : 01.10.2006

Time	Type	Nr	Clt	User	TCode	Priority	Grp	N	Text
06:51:31	DIA	001	000	SAPSYS			EJ	F	Could not find or load print parameters for step 1, job SWWDHEX/06513103
06:51:31	BTC	017	000	SAPSYS			EJ	F	Could not find or load print parameters for step 1, job SWWDHEX/06483104
06:51:31	DIA	001	000	SAPSYS			EJ	F	Could not find or load print parameters for step 1, job SWWDHEX/06513104
06:52:31	DIA	001	000	SAPSYS			EJ	F	Could not find or load print parameters for step 1, job SWWERRE/06523101
06:52:31	DIA	001	000	SAPSYS			EJ	F	Could not find or load print parameters for step 1, job SWWCOND/06523101
06:52:31	BTC	014	000	SAPSYS			EJ	F	Could not find or load print parameters for step 1, job SWWERRE/06313102
06:52:31	BTC	016	000	SAPSYS			EJ	F	Could not find or load print parameters for step 1, job SWWCOND/06223101
06:53:31	BTC	015	000	SAPSYS			EB	F	Failed to activate authorization check for user MARSTON
06:53:31	BTC	015	000	SAPSYS			D0	1	Transaction Canceled 00 560 (MARSTON 506)
06:53:31	BTC	015	000	SAPSYS			R6	8	Perform rollback
06:54:31	DIA	001	000	SAPSYS			EJ	F	Could not find or load print parameters for step 1, job SWWDHEX/06543101
06:54:31	DIA	001	000	SAPSYS			EJ	F	Could not find or load print parameters for step 1, job SWWDHEX/06543102
06:54:31	BTC	014	000	SAPSYS			EJ	F	Could not find or load print parameters for step 1, job SWWDHEX/06513101
06:54:31	BTC	016	000	SAPSYS			EJ	F	Could not find or load print parameters for step 1, job SWWDHEX/06513103
06:54:31	DIA	001	000	SAPSYS			EJ	F	Could not find or load print parameters for step 1, job SWWDHEX/06543103
06:54:31	BTC	015	000	SAPSYS			EJ	F	Could not find or load print parameters for step 1, job SWWDHEX/06513102
06:54:31	BTC	017	000	SAPSYS			EJ	F	Could not find or load print parameters for step 1, job SWWDHEX/06513104

Figure A-15   Display of the message list in the system log

A.2.11 ST22: ABAP run time error

Causes of shortdumps

Transaction ST22 enables you to analyze shortdumps that have occurred (see Figure A-16). Possible causes of shortdumps are the following:

- ABAP exceptions
- Time-out dumps for dialog processes that have been running longer than the permitted maximum response time

You can use transaction ST22 to identify and display the section of the programming code at which the exception occurred. The transaction provides analysis and correction recommendations, which are, however, rather generic. For this reason it is advisable to search for additional information in SAP Service Marketplace.

A.2.12 LC10: liveCache monitoring

liveCache is unique to SAP SCM and represents a separate database.

Transaction LC10 is the entry point for the liveCache assistant and offers monitoring functions and tools in the following areas:

- Performance
- OMS locks
- SQL locks
- SQL performance
- Messages
- Logs
- Tables/Views/Synonyms
- DB procedures

A.2.13 AANA: Transaction for the statistical evaluation of database tables

Use this transaction to analyze how entries for a table have been distributed to selected fields. Table analysis counts the table entries and assigns the number of entries found to the selected field values (such as organizational units or periods).

Due to the usage of Native-SQL, transaction TAANA is to be preferred to evaluations with transaction SE16.

A.2.14 Additional transactions

- ST12 - Trace analysis
- SE30 - Runtime analysis
  - Measurement in dialog status
  - Measurement of external session
  - Planning a measurement
  - Selection of measurement restrictions
  - Analyzing measurement results
- SM37 - Job overview

  This transaction gives an overview about all the scheduled programs to run in the background and also the history of them. This is to identify if some background action took parts of the performance. It also gives the run times of the scheduled programs.

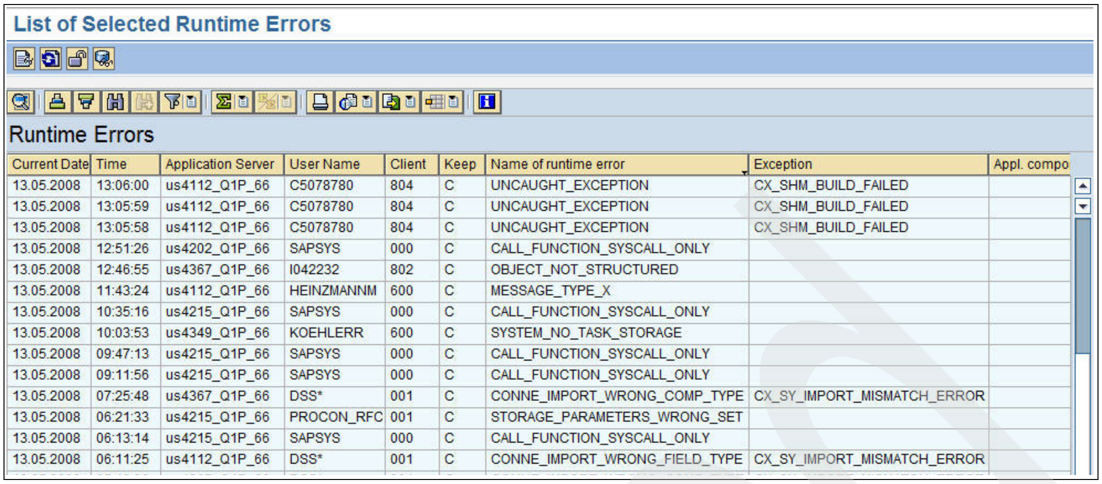

List of Selected Runtime Errors

Runtime Errors

Current Date	Time	Application Server	User Name	Client	Keep	Name of runtime error	Exception	Appl. compo
13.05.2008	13:06:00	us4112_Q1P_66	C5078780	804	C	UNCAUGHT_EXCEPTION	CX_SHM_BUILD_FAILED	
13.05.2008	13:05:59	us4112_Q1P_66	C5078780	804	C	UNCAUGHT_EXCEPTION	CX_SHM_BUILD_FAILED	
13.05.2008	13:05:58	us4112_Q1P_66	C5078780	804	C	UNCAUGHT_EXCEPTION	CX_SHM_BUILD_FAILED	
13.05.2008	12:51:26	us4202_Q1P_66	SAPSYS	000	C	CALL_FUNCTION_SYSCALL_ONLY		
13.05.2008	12:46:55	us4367_Q1P_66	I042232	802	C	OBJECT_NOT_STRUCTURED		
13.05.2008	11:43:24	us4112_Q1P_66	HEINZMANNM	600	C	MESSAGE_TYPE_X		
13.05.2008	10:35:16	us4215_Q1P_66	SAPSYS	000	C	CALL_FUNCTION_SYSCALL_ONLY		
13.05.2008	10:03:53	us4349_Q1P_66	KOEHLERR	600	C	SYSTEM_NO_TASK_STORAGE		
13.05.2008	09:47:13	us4215_Q1P_66	SAPSYS	000	C	CALL_FUNCTION_SYSCALL_ONLY		
13.05.2008	09:11:56	us4215_Q1P_66	SAPSYS	000	C	CALL_FUNCTION_SYSCALL_ONLY		
13.05.2008	07:25:48	us4367_Q1P_66	DSS*	001	C	CONNE_IMPORT_WRONG_COMP_TYPE	CX_SY_IMPORT_MISMATCH_ERROR	
13.05.2008	06:21:33	us4215_Q1P_66	PROCON_RFC	001	C	STORAGE_PARAMETERS_WRONG_SET		
13.05.2008	06:13:14	us4215_Q1P_66	SAPSYS	000	C	CALL_FUNCTION_SYSCALL_ONLY		
13.05.2008	06:11:25	us4112_Q1P_66	DSS*	001	C	CONNE_IMPORT_WRONG_FIELD_TYPE	CX_SY_IMPORT_MISMATCH_ERROR	

Figure A-16 Overview of ABAP runtime errors in transaction ST22

A.2.15 Summary

The monitoring tools contained in the SAP CCMS provide you with the information needed to narrow down, locate, and analyze performance problems. You can use these tools during a test run to monitor the load behavior of the system. In addition, the tools provide monitors for retroactive analysis. No other tool provides more information on the status of your SAP system. All currently available monitoring and load test tools access a subset of this information.

You can find more details under help.sap.com (search for “ABAP-Analysis Tools”) or in the book *Testing SAP Solutions* (http://www.sap-press.de/katalog/buecher/titel/gp/titelID-1408).

A.3 AIX monitoring tools

In this section we identify the tools that we used in this benchmark project to measure performance on System p and AIX.

What was monitored on the System p side during the runs is mainly CPU, memory consumption, and I/O usage.

We used several monitoring tools, first for online monitoring purposes, second for reporting purposes.

A.3.1 Online monitoring

For online monitoring we basically use the xmperf graphical monitor. Online monitoring on System p is focused on CPU usage, and in case of abnormal behavior of the system, memory and I/O are analyzed.

xmperf graphical monitor

The xmperf graphical monitor is used for real-time system monitoring. Recordings and snapshots from the graphical monitor are also used to track load distribution across LPARs and over the multiple components on System p.

A great deal of documentation about this tool in general is available on the Internet; detailed documentation can be found at
http://publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/perftool/prfusrgd/ch02body.htm.

As an example, Figure A-17 shows CPU usage for the five LPARs we used in this project.

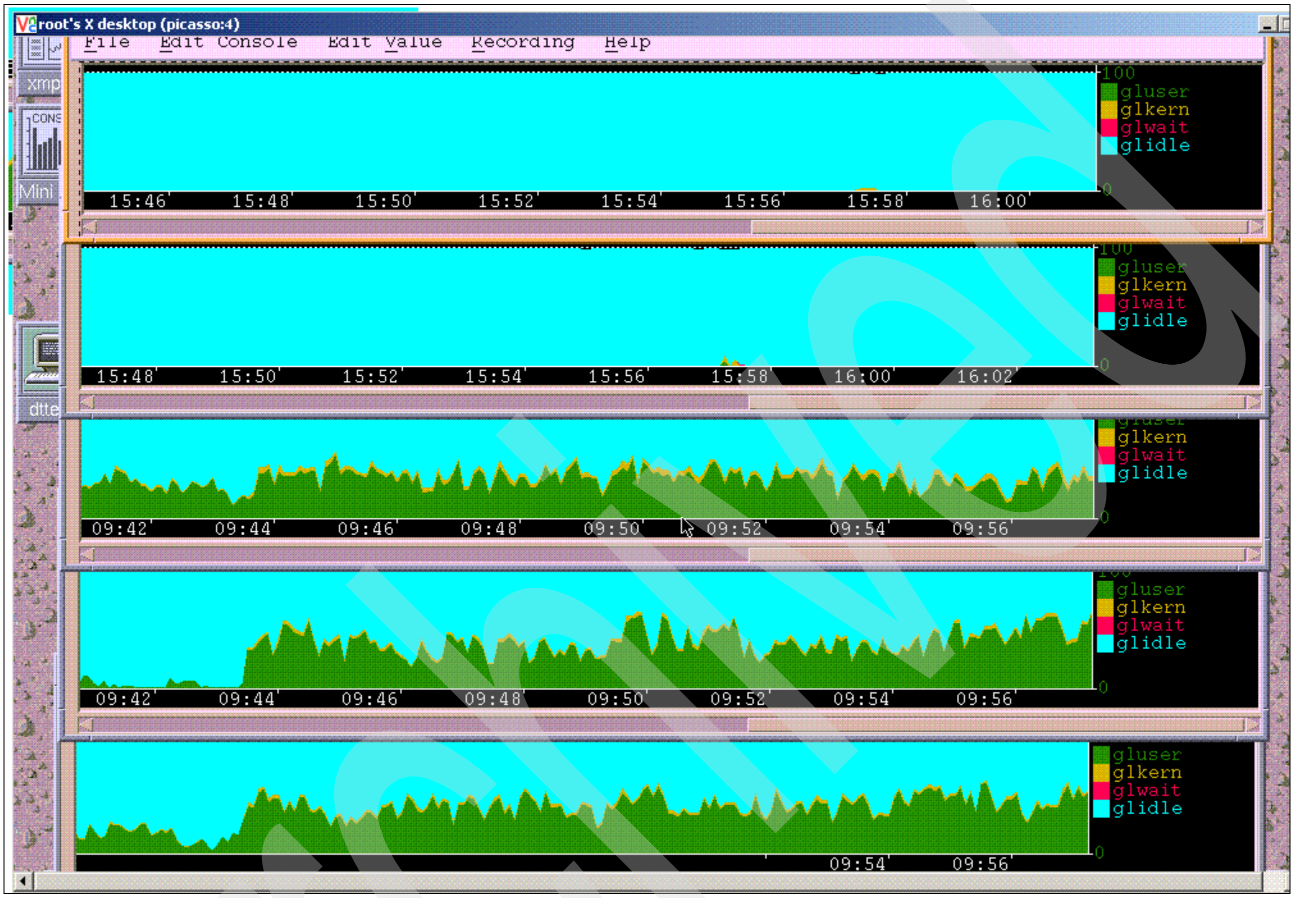

Figure A-17 xmperf CPU monitoring

Here we can easily see that CPU consumption of ECC applications servers is about 50%, sustained. This indicates that the load is well balanced between the three application servers.

A.3.2 Operating system tools

In case we need to monitor any other components of the System p5 servers regarding performance, we can either configure new monitors on xmperf or use operating system tools. In this case the second solution is chosen because it is usually quicker. Tools and commands we used included topas, vmstat, iostat, and nmon. Details are provided in the sections that follow.

topas

This command is used to see the top “resources consumer” processes in the system. For instance, if we call command **topas** on application server hkpoceccapp2, we see the data displayed as shown in Figure A-18.

```
Topas Monitor for host:    hkpoceccapp2          EVENTS/QUEUES    FILE/TTY
Fri Oct 12 11:14:04 2007   Interval:  2          Cswitch   22044  Readch  554.4K1
                                                 Syscall   41738  Writech1372.5K7
Kernel    2.7   |#                            |  Reads       182  Rawin        0
User     69.9   |####################         |  Writes     3095  Ttyout     416
Wait      0.0   |#                            |  Forks         1  Igets        0
Idle     27.4   |########                     |  Execs         3  Namei      649
                                                 Runqueue    4.0  Dirblk       0
Network  KBPS   I-Pack  O-Pack   KB-In  KB-Out   Waitqueue   0.0
en0    6461.9   4545.5  4275.5  4030.0  2431.9
lo0      38.8    473.0   473.0    19.4    19.4   PAGING           MEMORY
en2       0.8      4.0     1.0     0.4     0.5   Faults     5275  Real,MB   61440
                                                 Steals        0  % Comp     16.5
Disk    Busy%     KBPS     TPS KB-Read KB-Writ   PgspIn        0  % Noncomp   3.3
hdisk1    0.5      4.0     1.0     0.0     4.0   PgspOut       0  % Client    3.3
hdisk0    0.5      4.0     1.0     0.0     4.0   PageIn        0
hdisk2    0.0      0.0     0.0     0.0     0.0   PageOut       0  PAGING SPACE
hdisk3    0.0      0.0     0.0     0.0     0.0   Sios          0  Size,MB  147456
                                                                  % Used      0.5
Name            PID  CPU%  PgSp Owne             NFS (calls/sec)  % Free     99.4
disp+wor    1343814   9.0  23.0 d2ladm           ServerV2       0
disp+wor    1278262   7.2  24.8 d2ladm           ClientV2       0   Press:
disp+wor     987504   6.8  23.2 d2ladm           ServerV3       0   "h" for help
disp+wor     889140   6.6  22.6 d2ladm           ClientV3       0   "q" to quit
disp+wor     975208   6.2  26.5 d2ladm
disp+wor     946518   4.0  23.3 d2ladm
disp+wor    1139004   3.9  24.6 d2ladm
disp+wor     397320   3.6  23.9 d2ladm
disp+wor     971110   3.3  23.7 d2ladm
disp+wor     676068   3.2  22.6 d2ladm
disp+wor     938322   3.1  23.3 d2ladm
disp+wor     905536   3.1  22.5 d2ladm
disp+wor     913730   2.9  23.9 d2ladm
disp+wor     942420   2.8  24.3 d2ladm
gwrd        1220978   1.6   6.2 d2ladm
disp+wor     880942   1.3  13.5 d2ladm
disp+wor    1442292   0.5   2.2 d2ladm
db2sysc     1294624   0.2   1.3 d2ladm
db2sysc     1110458   0.2   1.3 d2ladm
db2sysc     1069474   0.1   1.5 d2ladm
```

Figure A-18 topas

At a glance we can see total CPU consumption, network usage, disk usage, top CPU consumers, memory, paging, and so forth.

Option **h** for **topas** means “help,” so if we press **h**, we can see all the possibilities for performance reporting, as shown in Figure A-19.

```
                              Topas Help
One-character commands:
  a - Show all the variable subsections being monitored. Pressing the
      the 'a' key always returns topas to the main inital display.
  c - Pressing the 'c' key repeatedly toggles the CPU subsection
      between the cumulative report, off, and a list of busiest CPUs.
  d - Pressing the 'd' key repeatedly toggles the disk subsection between
      busiest disks list, off, and total disk activity for the system.
  f - Moving the cursor over a WLM class and pressing 'f' shows the list of
      top processes in the class on the bottom of the screen(WLM Display Only)
  n - Pressing the 'n' key repeatedly toggles the network interfaces subsection
      between busiest interfaces list, off, and total network activity.
  p - Pressing the 'p' key toggles the hot processes subsection on and off.
  P - Toggle to the Full Screen Process Display
  q - Quit the program
  r - Refresh the screen
  w - Pressing the 'w' key toggles the WorkLoad Management(WLM)
      classes subsection on and off.
  W - Toggle to the Full Screen WLM Class Display
  L - Toggle to the Logical Partition/CPU Display
  D - Toggle to the Disk Display
  C - Toggle to the Cross-LPAR (CEC) Panel

WLM,Disk,Process displays, the network,disk,cpu,and wlm sections may be sorted
using the ARROW keys and/or TAB. A highlighted column header shows sorted-by
      Use one of the one-character commands to return to monitoring.
```

Figure A-19 topas options

As an example, choosing **c** (CPU) option returns a listing as shown in Figure A-20 on page 244.

The listing shows the CPU usage and CPU consumer processes from highest to the lowest.

```
Topas Monitor for host:    hkpoceccapp2        EVENTS/QUEUES    FILE/TTY
Fri Oct 12 11:17:24 2007   Interval:  2        Cswitch   29524  Readch      555
                                               Syscall   64052  Writech3160.8K
CPU      User%   Kern%   Wait%   Idle%         Reads         4  Rawin         0
cpu2      88.5     0.6     0.0    10.9         Writes     2954  Ttyout      551
cpu15     71.9     1.8     0.0    26.3         Forks         0  Igets         0
cpu14     61.2     4.4     0.0    34.4         Execs         0  Namei        16
cpu5      61.1     1.9     0.0    36.9         Runqueue    5.0  Dirblk        0
cpu19     60.0     2.0     0.0    38.0         Waitqueue   0.0
cpu18     59.1     4.8     0.0    36.1
cpu4      58.4     5.6     0.0    36.0         PAGING           MEMORY
cpu10     57.8     5.1     0.0    37.1         Faults     2193  Real,MB   61440
cpu16     57.7     5.1     0.0    37.2         Steals        0  % Comp     16.5
cpu11     56.1     2.0     0.0    41.9         PgspIn        0  % Noncomp   3.3
cpu12     55.2     4.8     0.0    40.1         PgspOut       0  % Client    3.3
cpu8      54.6     8.6     0.0    36.8         PageIn        0
cpu6      54.4     5.2     0.0    40.4         PageOut       0  PAGING SPACE
cpu17     48.7     1.5     0.0    49.8         Sios          0  Size,MB  147456
cpu9      35.4     1.3     0.0    63.4                          % Used      0.5
cpu0      22.7     1.5     0.0    75.8         NFS (calls/sec)  % Free     99.4
cpu7       6.5    17.0     0.0    76.5         ServerV2       0
cpu13      0.0     0.7     0.0    99.3         ClientV2       0   Press:
cpu3       0.0     0.8     0.0    99.2         ServerV3       0   "h" for help
cpu1       0.0     0.7     0.0    99.3         ClientV3       0   "q" to quit

Network  KBPS   I-Pack  O-Pack   KB-In  KB-Out
Total    9.8K   8213.0  8149.5  5669.1  4346.1

Name            PID  CPU%  PgSp Owner
disp+wor     880942   7.6  13.5 d2ladm
disp+wor     913730   4.8  23.9 d2ladm
disp+wor     942420   4.8  24.3 d2ladm
disp+wor     971110   4.7  23.7 d2ladm
disp+wor     938322   4.7  23.3 d2ladm
disp+wor     975208   4.6  26.5 d2ladm
disp+wor     987504   4.4  23.2 d2ladm
disp+wor     676068   4.4  22.6 d2ladm
disp+wor     889140   4.3  22.6 d2ladm
disp+wor    1343814   3.4  23.0 d2ladm
disp+wor     397320   3.4  23.9 d2ladm
disp+wor    1278262   3.3  24.8 d2ladm
disp+wor    1442292   1.6   2.2 d2ladm
gwrd        1220978   0.9   6.2 d2ladm
disp+wor     905536   0.2  22.5 d2ladm
db2sysc      774302   0.2   2.1 d2ladm
db2sysc      819458   0.2   1.3 d2ladm
db2sysc      745482   0.2   1.3 d2ladm
db2sysc     1110458   0.1   1.3 d2ladm
db2sysc     1294624   0.1   1.3 d2ladm
```

Figure A-20 topas c

vmstat

vmstat is a useful command to see the LPAR resources consumptions at a glance, as shown in Figure A-21.

```
{hkpoceccapp2:root}/ -> vmstat 2

System configuration: lcpu=20 mem=61440MB

kthr    memory              page              faults        cpu
----- ----------- ------------------------ ------------ -----------
 r  b   avm   fre  re  pi  po  fr   sr  cy  in   sy  cs us sy id wa
10  0 2539469 12716107   0   0   0   0    0   0 367 4909 2367 79  0 21  0
11  0 2562494 12693082   0   0   0   0    0   0 675 3130 3098 74  0 26  0
 9  0 2582733 12672843   0   0   0   0    0   0 1236 7114 7382 70  1 29  0
 8  0 2601943 12653633   0   0   0   0    0   0 1826 15341 14245 76  1 23  0
 9  0 2622071 12633505   0   0   0   0    0   0 2426 19608 18526 72  1 27  0
 9  0 2631346 12624230   0   0   0   0    0   0 2030 17040 14042 77  1 22  0
11  0 2635870 12619706   0   0   0   0    0   0 2190 15787 15429 72  1 27  0
 8  0 2640028 12615548   0   0   0   0    0   0 3489 26341 24523 70  2 28  0
10  0 2642851 12612725   0   0   0   0    0   0 2844 21032 19484 75  1 24  0
 8  0 2644627 12610949   0   0   0   0    0   0 3261 25276 22901 71  1 27  0
 6  0 2647828 12607748   0   0   0   0    0   0 3485 28768 25545 50  2 48  0
 5  0 2650026 12605550   0   0   0   0    0   0 3657 26675 26013 25  2 74  0
 2  0 2653660 12601916   0   0   0   0    0   0 3402 25500 24993 13  2 86  0
 1  0 2656479 12599097   0   0   0   0    0   0 3089 23623 22944  6  1 92  0
 0  0 2659248 12596328   0   0   0   0    0   0 3125 23546 22774  6  1 93  0
 1  0 2662648 12592928   0   0   0   0    0   0 3203 27221 25119  7  2 91  0
 1  0 2659572 12596004   0   0   0   0    0   0 3038 26536 26098  6  2 93  0
 1  0 2661384 12594192   0   0   0   0    0   0 2870 24666 24784  5  2 94  0
 0  0 2662668 12592908   0   0   0   0    0   0 2573 22425 22910  7  2 92  0
 0  0 2662750 12592826   0   0   0   0    0   0 2353 21295 20607  3  2 95  0
kthr    memory              page              faults        cpu
```

Figure A-21 vmstat

In this example we see that while the system is busy, everything is fine. This is indicated by the fact that there are no kernel threads in wait queue and no paging.

vmstat does not provide any data on disk and network usage.

iostat

iostat provides disk usage statistics, as shown in Figure A-22.

```
{hkpoceccapp2:root}/ -> iostat 2

System configuration: lcpu=20 drives=4 paths=4 vdisks=0

tty:      tin         tout    avg-cpu: % user % sys % idle % iowait
          0.0         29.5              19.9    1.4   78.7      0.0

Disks:         % tm_act     Kbps      tps    Kb_read   Kb_wrtn
hdisk2            0.0       0.0       0.0          0         0
hdisk1            0.0       0.0       0.0          0         0
hdisk0            0.0       0.0       0.0          0         0
hdisk3            0.0       0.0       0.0          0         0

tty:      tin         tout    avg-cpu: % user % sys % idle % iowait
          0.0        218.6              22.3    1.8   75.9      0.0

Disks:         % tm_act     Kbps      tps    Kb_read   Kb_wrtn
hdisk2            0.0       0.0       0.0          0         0
hdisk1            0.0       0.0       0.0          0         0
hdisk0            0.0       0.0       0.0          0         0
hdisk3            0.0       0.0       0.0          0         0

tty:      tin         tout    avg-cpu: % user % sys % idle % iowait
          0.0        228.4              37.6    2.9   59.4      0.0

Disks:         % tm_act     Kbps      tps    Kb_read   Kb_wrtn
hdisk2            0.0       0.0       0.0          0         0
hdisk1            2.0     260.4       4.0          0       521
hdisk0            2.0     260.4       4.0          0       521
hdisk3            0.0       0.0       0.0          0         0
{hkpoceccapp2:root}/ ->
```

Figure A-22 iostat 2

As these read and write statistics indicate, in this case everything is OK; we only have on those disks operating system and swap space.

Obviously, this list of AIX tools and commands is not complete, but represents those most used in this project for performance monitoring purposes.

A.4 nmon

Nmon is another useful tool used in this project to monitor System p. It is a free tool that can be downloaded from http://www.ibm.com/developerworks/aix/library/au-analyze_aix/

Nmon first of all gives details about the system configuration, as shown in Figure A-23.

```
------------------------------
N    N  M    M   OOOO   N    N   For online help type: h
NN   N  MM   MM  O    O  NN   N   For command line option help:
N N  N  M MM M  O    O  N N  N      quick-hint  nmon -?
N  N N  M    M  O    O  N  N N    full-details  nmon -h
N   NN  M    M  O    O  N   NN   To start nmon the same way every time?
N    N  M    M   OOOO   N    N    set NMON ksh variable, for example:
------------------------------    export NMON=cmt
     Version v12beta28 for AIX53
                       20 - CPUs currently
                       20 - CPUs configured
                     2147 - MHz CPU clock rate
            PowerPC_POWER5 - Processor
                    64 bit - Hardware
                    64 bit - Kernel
                   Dynamic - Logical Partition
             5.3.0.56 ML05 - AIX Kernel Version
               hkpoceccapp2 - Hostname
```

Figure A-23   nmon

All nmon options can be displayed by pressing **h**. These options are shown in Figure A-24.

```
-HELP---------key=statistics which toggle on/off---------------------------
    h = This help information - use h to remove this screen
    r = Resources CPU/cache/AIX etc. p = Partitions (LPAR) Stats
    t = Top Process Stats 1=basic  2=CPU-Use 3=CPU 4=Size 5=Disk-I/O
        u = shows command arguments (hit u again to refresh)
    c = CPU by processor              #=PURR for p5 MicroLPAR
    C = CPU can show up to 128 CPUs   l = longer term CPU averages
    m = Memory and Paging stats       k = Kernel Internal stats
    n = Network stats                 j = JFS Usage Stats
    d = Disk I/O Graphs D=Stats       o = Disks %Busy Map
    a = Adapter I/O (not on AIX 5.1)  e = ESS vpath Stats
    A = Async I/O Servers             g = Disk Groups (see cmdline -g)
    v = VerboseChecks-OK/Warn/Danger  N = NFS stats
    b = black & white mode            w = view AIX wait processes
    U = WLM Classes (Top Section)     W = WLM Section (S=SubClasses)
    V = Volume Group disk Stats       P = Paging Space Stats
    --- controls ---
    + and - = double or half the screen refresh time
    q = quit                        space = refresh screen now
    . = Minimum Mode =display only busy disks and processes
    0 = reset peak counts to zero (peak = ">")

    nmon version v12beta28  - written by Nigel Griffiths, nag@uk.ibm.com
```

Figure A-24   nmon options

For instance, we can look at CPU consumption with nmon, as shown Figure A-25.

```
-nmon-----------l=LongTerm-CPU------Host=hkpoceccapp2----Refresh=2 secs----11:31.26---
 CPU-Utilisation-Small-View ---------------------------------------------------------
                                  0-----------25-----------50-----------75-----------100
 CPU User%  Sys% Wait% Idle%|            |            |            |            |
   0   0.0   0.0   0.0 100.0|                                                   >
   1   0.0   0.0   0.0 100.0|  >                                                |
   2   0.0   0.0   0.0 100.0|                                   >               |
   3   0.0   0.0   0.0 100.0|  >                                                |
   4  36.0   6.0   0.0  58.0|UUUUUUUUUUUUUUUUUUsss                              >
   5  24.2   1.4   0.0  74.4|UUUUUUUUUUUU                 >                     |
   6  35.5   6.5   0.0  58.0|UUUUUUUUUUUUUUUUUsss                               >
   7  29.0   1.0   0.0  70.0|UUUUUUUUUUUUUU >                                   |
   8  36.5   4.0   0.0  59.5|UUUUUUUUUUUUUUUUUUss                              >|
   9   0.0   0.0   0.0 100.0|                              >                    |
  10  60.0   6.0   0.0  34.0|UUUUUUUUUUUUUUUUUUUUUUUUUUUUUUsss                 >|
  11  23.1   1.0   0.0  76.0|UUUUUUUUUUU                              >         |
  12  35.0   7.0   0.0  58.0|UUUUUUUUUUUUUUUUUsss                              >|
  13   5.4   0.0   0.0  94.6|UU                           >                     |
  14  43.5   5.5   0.0  51.0|UUUUUUUUUUUUUUUUUUUUUss                            >
  15   0.0   0.0   0.0 100.0|                         >                         |
  16  37.5   5.5   0.0  57.0|UUUUUUUUUUUUUUUUUUss                      >        |
  17  13.0   1.5   0.0  85.5|UUUUUU                                             >
  18  60.5   5.0   0.0  34.5|UUUUUUUUUUUUUUUUUUUUUUUUUUUUUUss                   >
  19   0.0   0.0   0.0 100.0|                     >                             |
 Physical Averages          +------------|------------|------------|------------+
 All  41.0   3.9   0.0  55.1|>UUUUUUUUUUUUUUUUUUUs                              |
                            +------------|------------|------------|------------+
```

Figure A-25 CPU usage

As another example, paging details are shown in Figure A-26.

```
-nmon-----------t=Top-Processes-----Host=hkpoceccapp2----Refresh=2 secs----11:29.04---------
 Paging-Space ------------------------------------------------------------------------------
    Volume-Group PagingSpace-Name Type LPs   MB    Used IOpending
          rootvg              hd6  LV   64 16384    0%     0    Active    Auto
          rootvg         paging00  LV  128 32768    0%     0    Active    Auto
          rootvg         paging01  LV  128 32768    0%     0    Active    Auto
          rootvg         paging02  LV  256 65536    0%     0    Active    Auto
```

Figure A-26 paging details

We do not show all details, but nmon can give detailed information about all activity at AIX level.

A.4.1 Reporting and documentation

Nmon has the great advantage of being able to run in background, logging the information collected in a file. Information to be collected is tunable, as are the collecting periods and intervals.

For the run performance reporting on System p, we used nmon as the collecting tool and nmon analyzer as the analysis and reporting tool. Nmon analyzer is a free tool and can be downloaded at:

http://www.ibm.com/developerworks/aix/library/au-nmon_analyser/index.html.

Nmon analyzer generates and displays information in an easy to understand graphical format that enables you to analyze system performance at a glance.

The main performance components considered there are CPU, memory, and I/O. Several examples of nmon analyzer graphical reports are provided in this section.

One element we always study is the system summary as shown in Figure A-27.

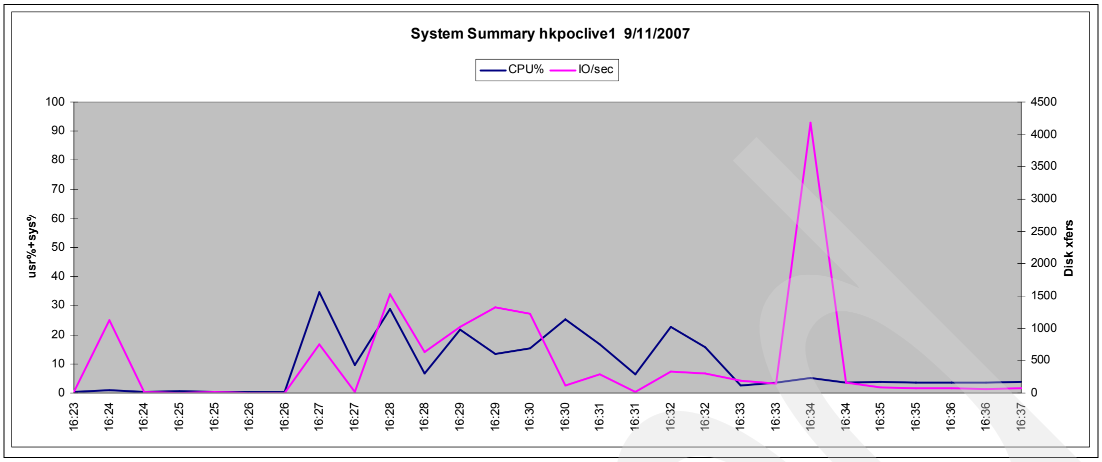

Figure A-27 system performance summary

On this graph, we can see at a glance the percentage of CPU consumption and the I/O per second at the disk level. This gives us an overview of the system behavior during the run. The time scale at the bottom is always adjusted to fit the run it is documenting.

The next graph shows a high-level view of CPU consumption. As shown in Figure A-28, we can distinguish user used CPU, system used CPU, and wait percentage throughout the runs.

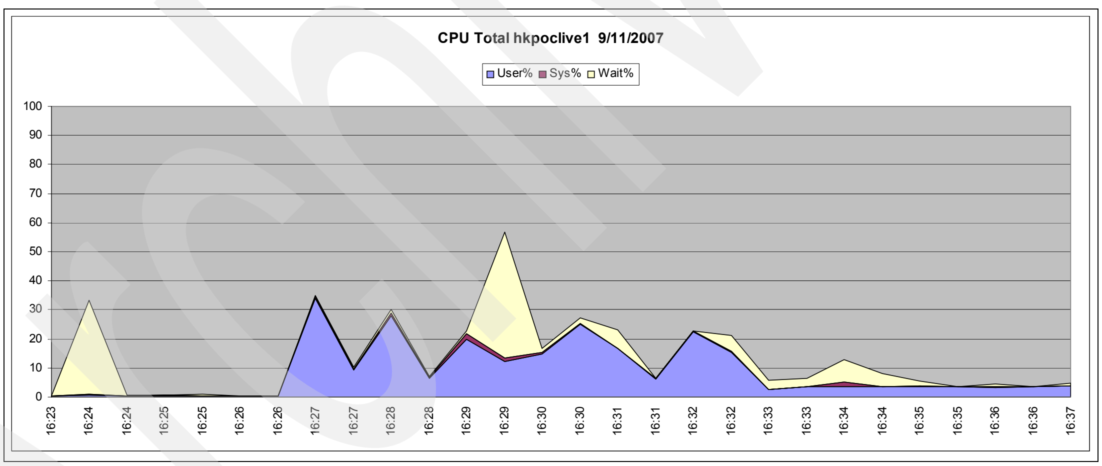

Figure A-28 CPU Total

The next interesting element to identify possible bottlenecks and performance issues is the CPU usage spread by CPU, where we can see if the load is well balanced between all CPUs. This is charted in Figure A-29 on page 250.

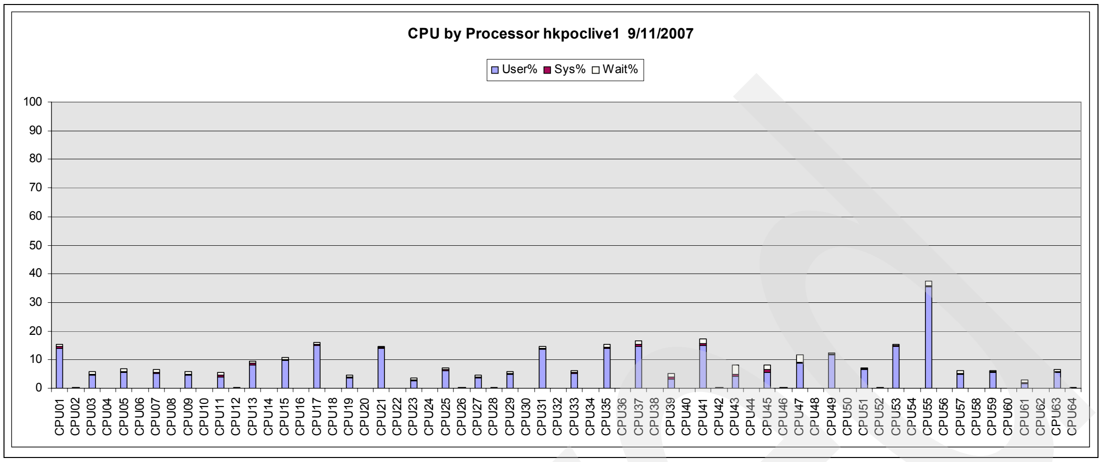

Figure A-29 CPU details

In this particular case we can notice that the application running on this LPAR doesn't use SMT (simultaneous multi threading) because it only uses one thread of the two available per CPU. This is a good sign because AIX balances first on available physical CPUs before using threads defined.

The next element we look at is the memory usage, as shown in Figure A-30.

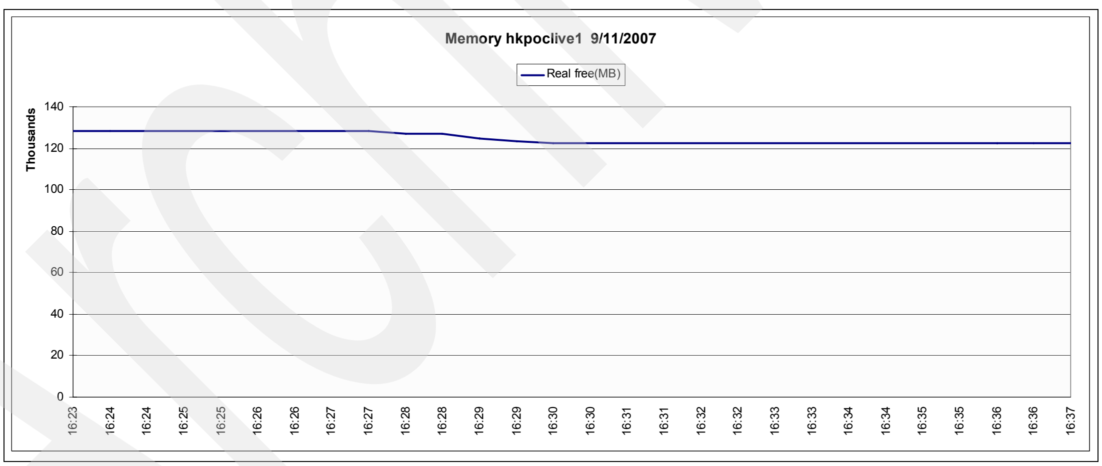

Figure A-30 Memory usage

SAP applications usually use memory according to their needs. In this example, we see stable usage of the available memory.

Disk I/O is also interesting data to look carefully at. Some contention can happen there, so it is always important to monitor disks in detail. In Figure A-31 we can see the I/Os per second, disk reads, and disk writes.

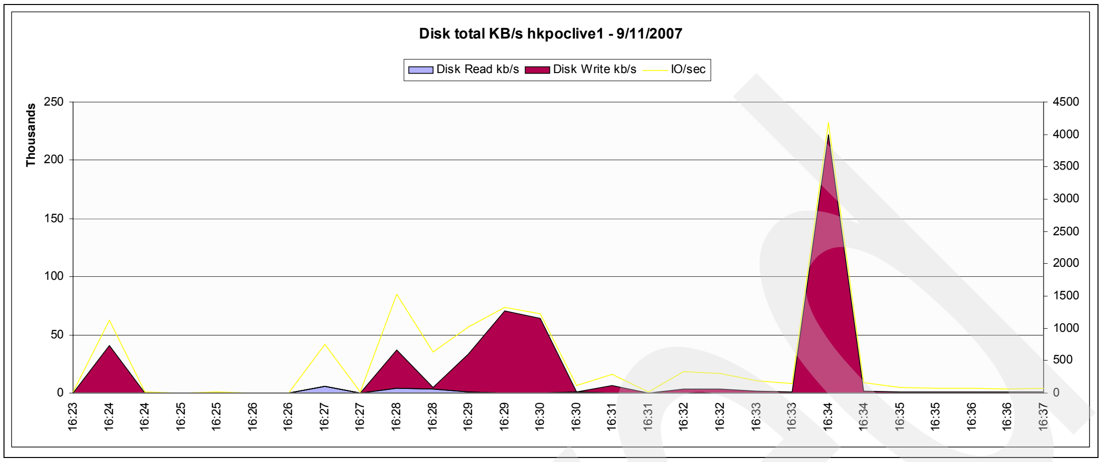

Figure A-31 Disk statistics

In this particular example we have some peaks during a few minutes and a maximum of 415 I/Os per second, which is very few for this LPAR. Everything is OK at the disk level on this system.

It is also interesting to see activity by disk during the run interval, as shown in Figure A-32 and Figure A-33 on page 252. Hdisk0 and hdisk1 have the same activity because they are mirrored. These two disks are internal and hold the operating system. The other used disks (hdisk8 to hdisk13) are logical units presented by the storage subsystem. We see that they have a very small average activity (less that 5% busy) and peak activity is about 50% for hdisk9.

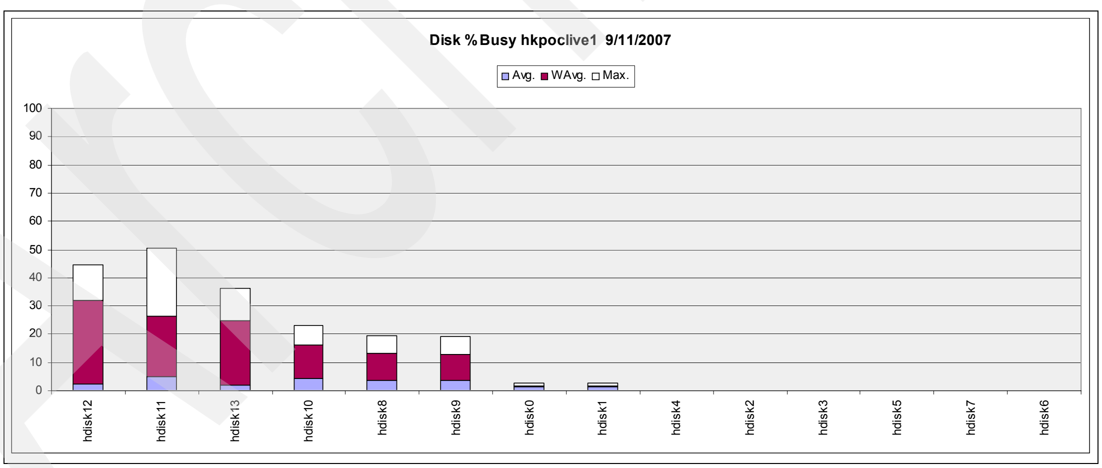

Figure A-32 Disk activity report

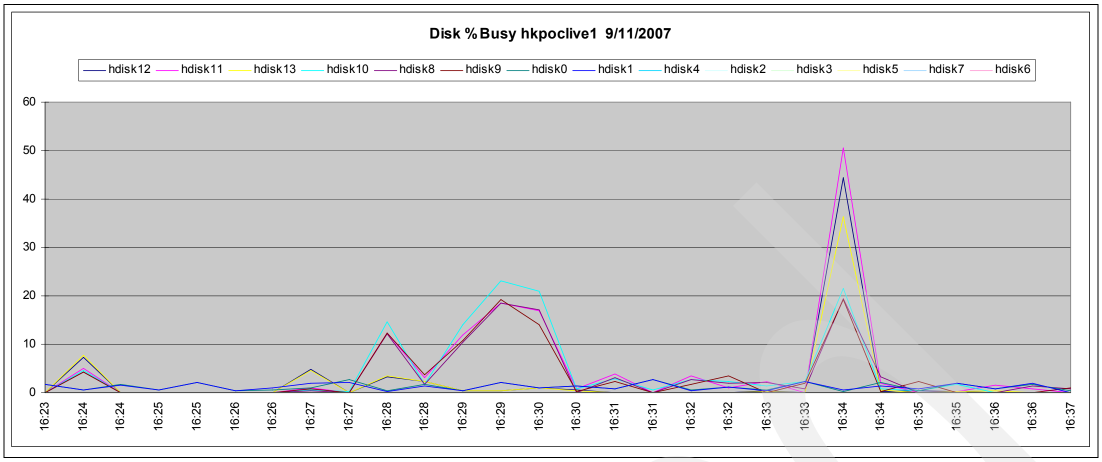

Figure A-33 Disk activity during the period

If now we observe the disk activity in the time scale of the run, we can see that the peak load of 50% was reached only during one minute, not as a sustained load.

We can display additional information such as I/Os per second per disk, which is also interesting, as shown in Figure A-34 and Figure A-35 on page 253. Again we can see very low activity averages and peak loads that are not very high.

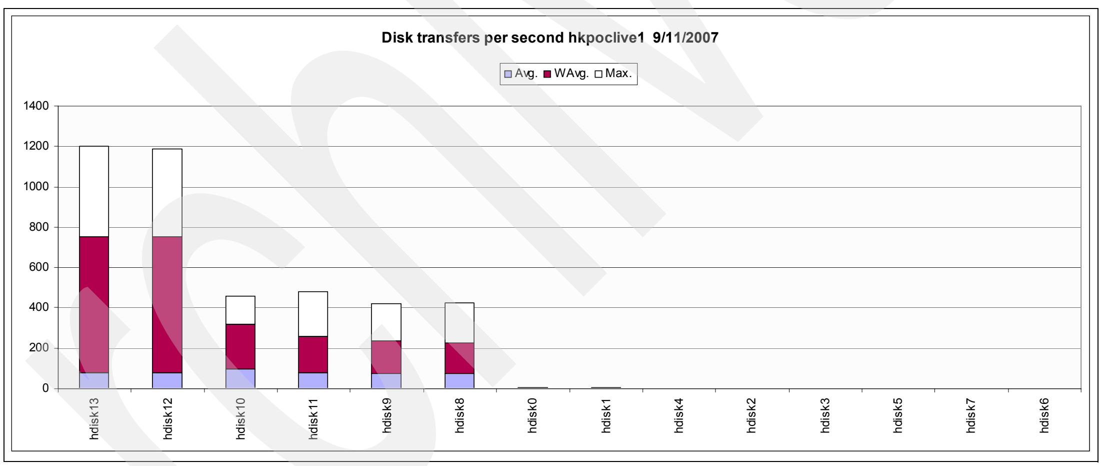

Figure A-34 I/O per second per disk

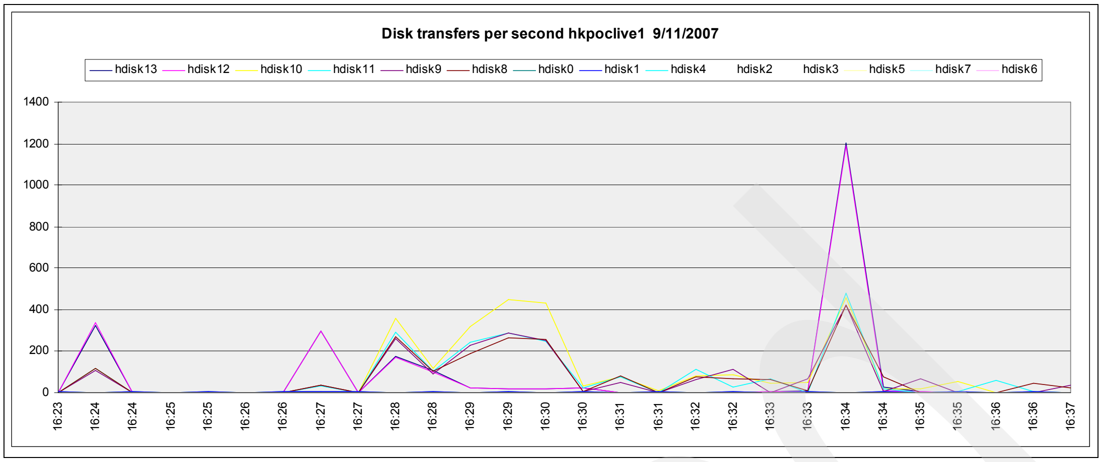

Figure A-35 I/O per sec per disk on time frame

Thus we can conclude that everything is fine on the disk side. There are no contentions and no performance issues related to the disks.

It also might be interesting to have a look at the adapters themselves in some cases. We can obviously analyze read and write contention at the adapter level, as shown in Figure A-36 and Figure A-37 on page 254. We can see that SCSI adapters have no activity; they hold operating systems and swap so this is a good sign, and we see that FC adapters have the same usage profile. It is expected because they access the same disk subsystem areas. We notice that there are few reads and some writes, which is not surprising given the run functional profile.

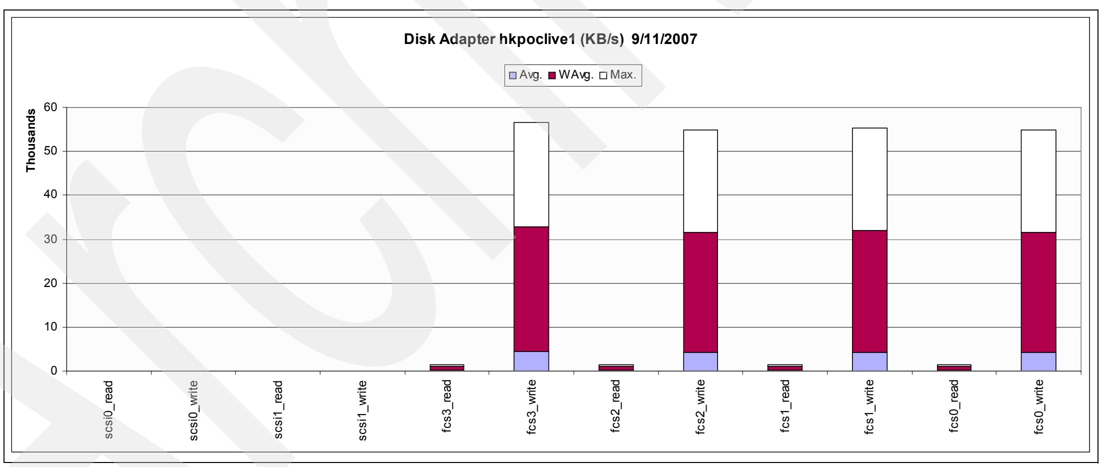

Figure A-36 Adapters load

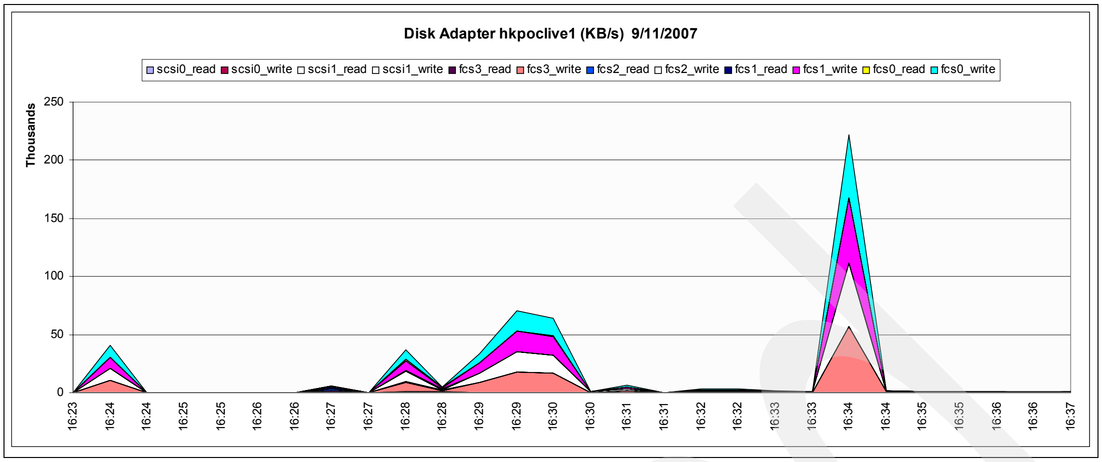

Figure A-37 Adapters load on time frame

If we check the time scale, we see the same few small peaks that were observed with the disk profile study, which is quite logical.

Therefore, we can say that on the adapter side everything is running fine.

Now we take a look at the network usage profile. In Figure A-38 we have a clear picture of the total network reads and writes. There are more reads, and we see a peak at 80 KB/s, which is not much in that case, but we have to study the capacity of the network adapter to be sure there is no bottleneck.

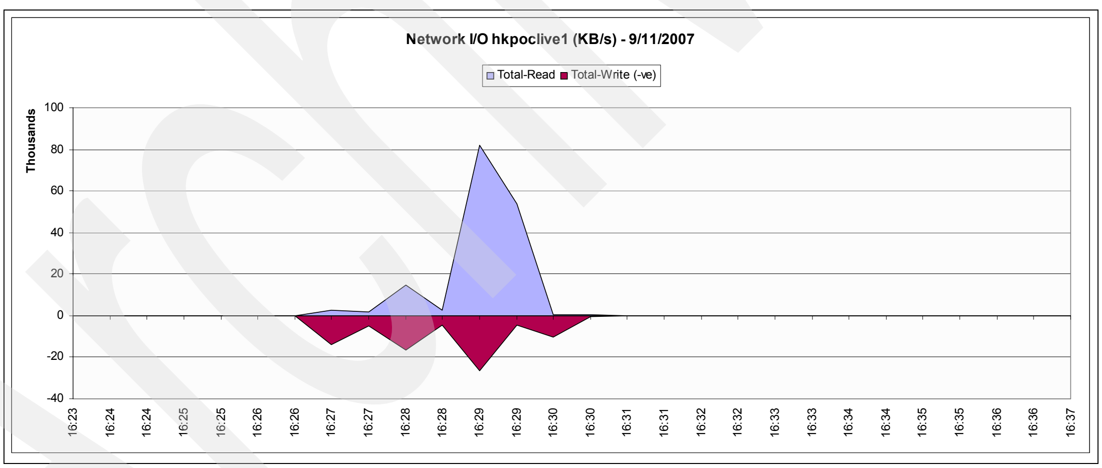

Figure A-38 Network profile

So to have more detail, we study the network load per adapter, as shown in Figure A-39.

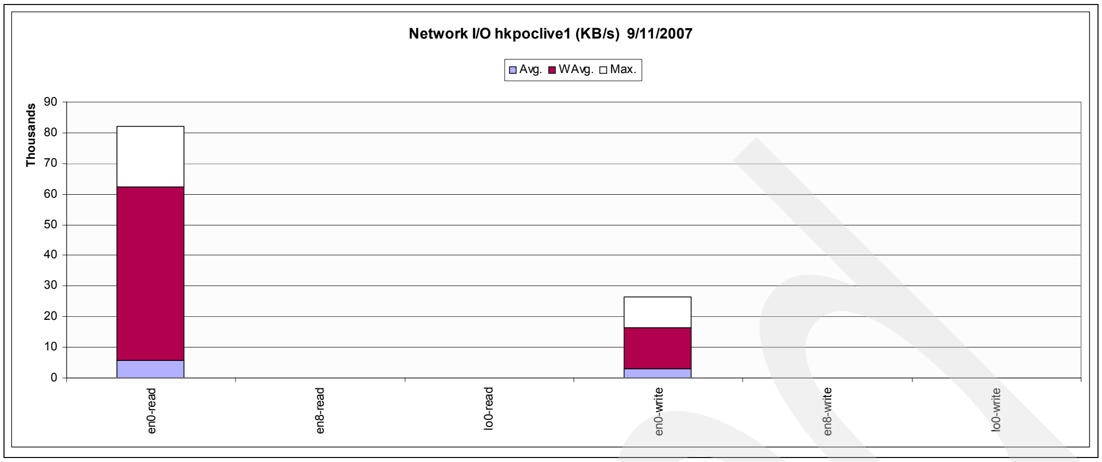

Figure A-39 Network usage by adapter

For data transfer between all systems in the SAP landscape we only use one adapter (en0), which is a 1 Gb adapter. We can observe a little more detail regarding the adapter usage on the run time frame shown in Figure A-40.

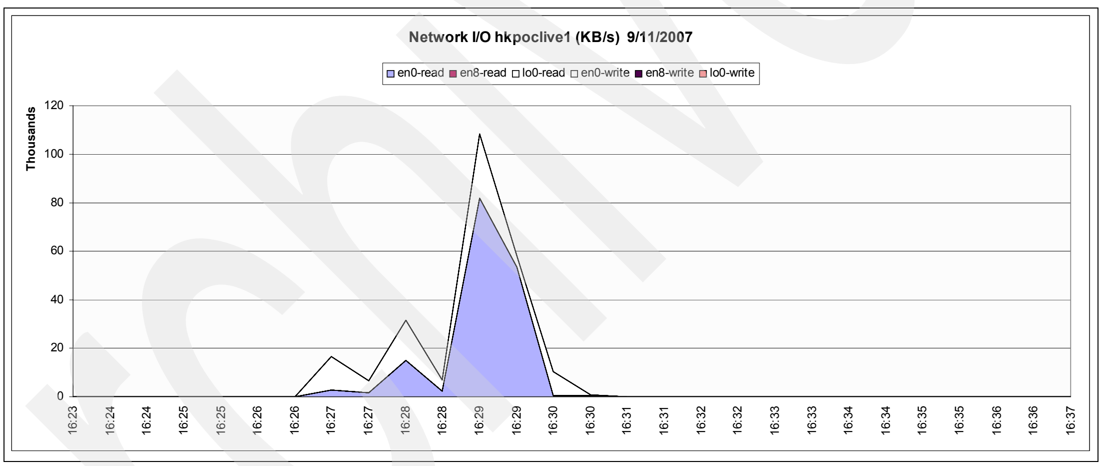

Figure A-40 Adapter usage on time frame

In this graph it is interesting to note that for reads and writes we have a peak of 110 MB/s, not sustained, and the capacity of the adapter is 1 Gbps, so this peak almost reaches the capacity of the adapter.

We can also check the network errors, as shown in Figure A-41.

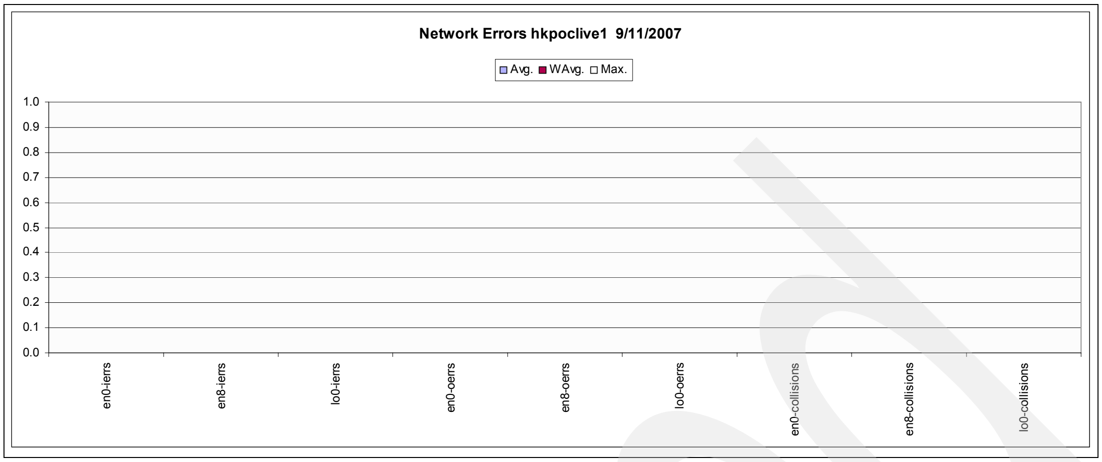

Figure A-41 Net errors

This graph is empty, so we had no network errors during the run.

Another interesting piece of information we gather from nmon is the Kbps per adapter. We can also distinguish between read and write, as shown in Figure A-42.

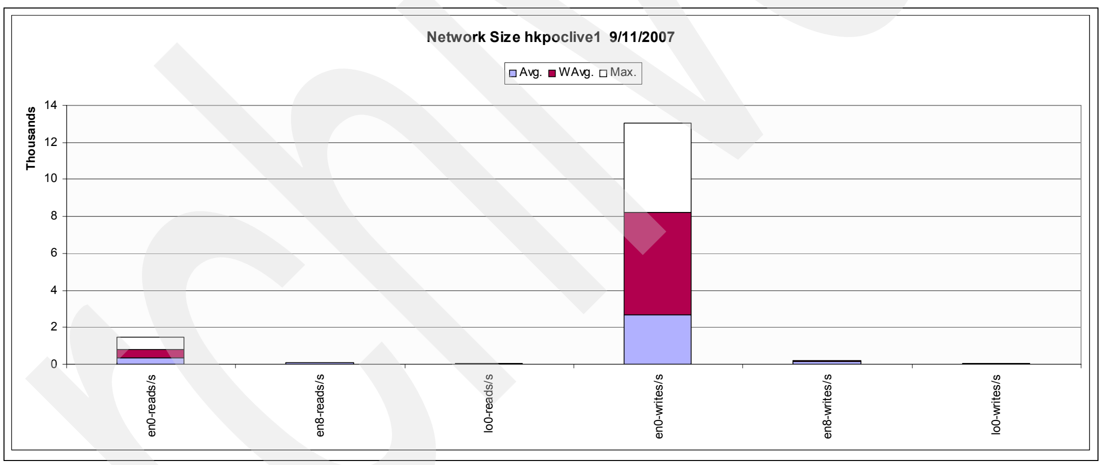

Figure A-42 read and write size

It is also interesting to take a look at this information plotted on the run time frame as shown in Figure A-43.

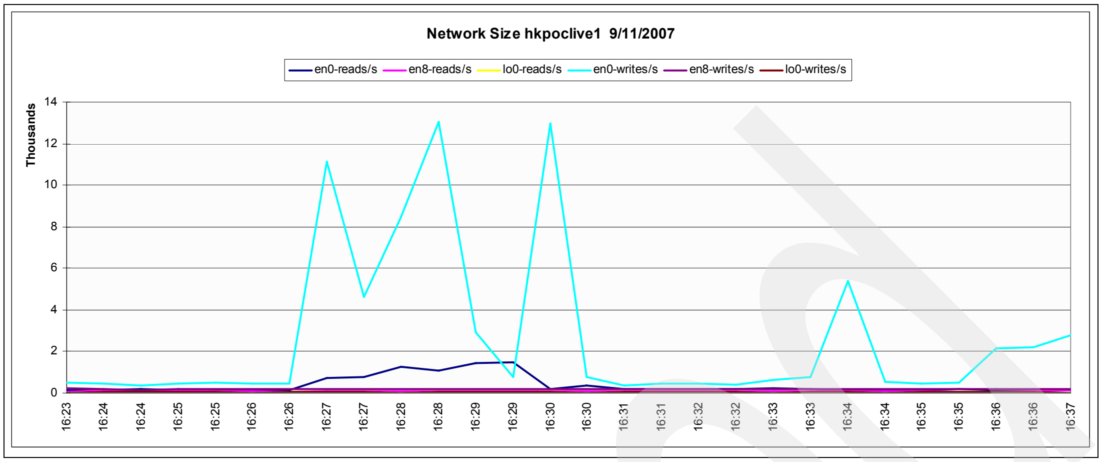

Figure A-43 Size on time frame

Figure A-44 shows the absence of paging activity, indicating the absence of any problems.

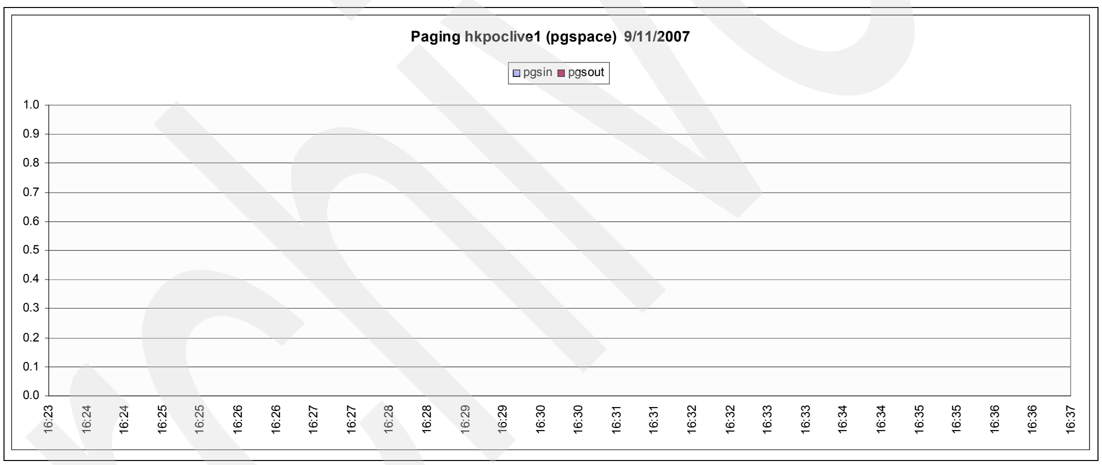

Figure A-44 paging

No paging occurred, so there were no problems in that respect.

Top CPU consuming processes are shown in Figure A-45.

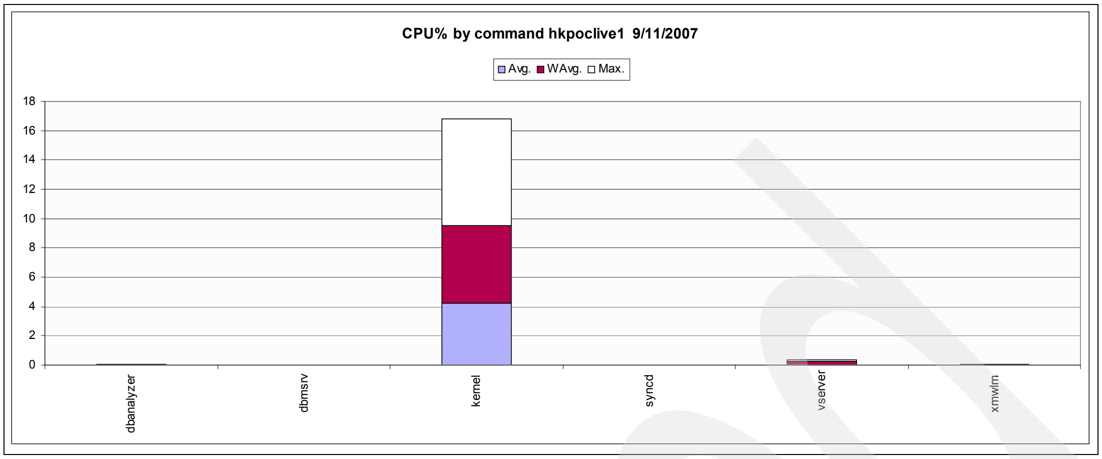

Figure A-45 Top CPU consumers

liveCache is the top consumer (under kernel), then we have a small use from the vserver process (communication manager between liveCache and the SAP SCM ABAP stack).

Regarding memory, Figure A-46 shows that the main memory consumer is liveCache, with about 50 GB during the run. There is not much difference between average and maximum used memory because liveCache has quite static memory usage.

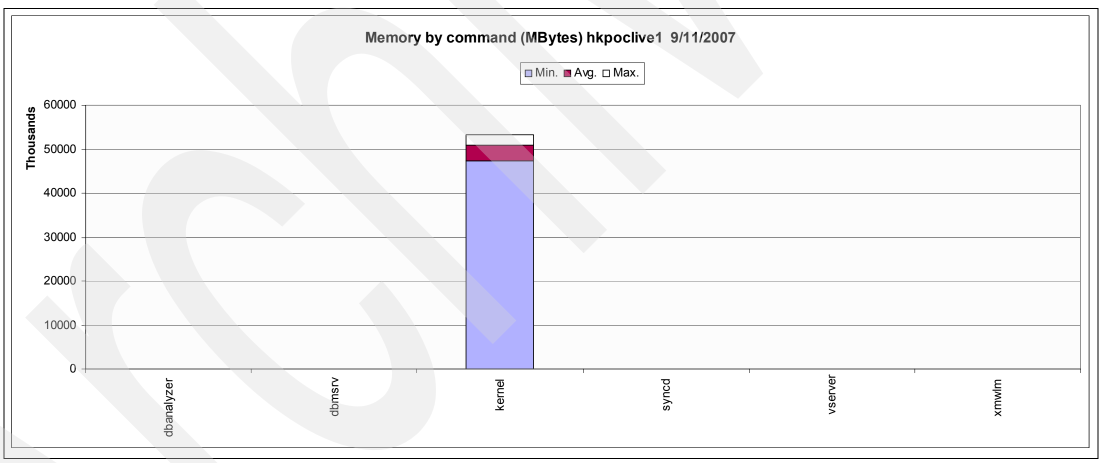

Figure A-46 Memory top consumers

This information provides us with a clear idea of the system activities during the runs.

A.5 zOS and DB2 reports

All zOS reports created during the tests and shown in this book are based on SMF recording and post processing of SMF records. This method is sufficient for the zPOC since all KPIs are batch oriented. If KPIs had been defined for real-time transactions (for example, Transaction X should be finished in x seconds) additional monitoring tools would have been needed.

The SMF interval was reduced to 1 minute (the default value is 30 minutes). The result is that every minute all SMF recorded information is externalized to disk. All SMF data created during the zPOC is kept. This setup provides us with the needed granularity for creating meaningful reports. It also gives us ability to create the reports, create additional, more detailed reports, and so on.

After a test for one of the KPIs, all relevant SMF data is cut from the whole lot of recorded SMF data. To do this the startdate/time and the enddate/time of the test run is used. This subset of SMF data is used to create a set of reports and do a preliminary analysis. Based on this information (and of course real time monitoring during the tests), decision and modifications are made for subsequent test runs.

The tools used for creating the reports are:

- RMF Postprocessor
- RMF Spreadsheet reporter

A.5.1 RMF Postprocessor

The RMF Postprocessor is used to create some traditional character-based reports. The reports are created for the entire period the job ran. Figure A-47 is an example of the parameters used for the RMF Postprocessor job.

```
DATE(07269,07269)
RTOD(1450,1501)
STOD(1450,1501)
DINTV(0011)
REPORTS(CPU)
REPORTS(CHAN)
REPORTS(IOQ)
REPORTS(PAGING)
REPORTS(DEVICE(DASD))
REPORTS(CACHE)
SYSRPTS(WLMGL(SCLASS,SCPER,POLICY,RCLASS))
SYSOUT(G)
```

Figure A-47 Control Statement for RMF postprocessing

In this example a KPI was tested which started and ended on 09/25/2007. The start time was 14:50 and the end time was 15:01, so the interval is 11 minutes long. The following reports were created:

- RMF summary report
- CPU activity report
- Channel Path activity report
- I/O queuing activity report

- Cache subsystem activity report
- DASD activity report
- Paging activity report
- Workload activity report

Examples and brief descriptions of the reports follow. For a more detailed discussion refer to the RMF manuals.

RMF summary report

This report is created because no SUMMARY statement is specified for the RMF postprocessor. Not specifying this statement defaults to SUMMARY(INT,TOT). The report gives a good high-level insight into what was happening on the LPAR during the test, as shown in Figure A-48.

```
                                          R M F    S U M M A R Y    R E P O R T
                                                                                                                     PAGE   001
            z/OS   V1R8                SYSTEM ID MVH1             START 09/26/2007-14.50.00  INTERVAL 00.00.59
                                       RPT VERSION V1R8 RMF       END   09/26/2007-15.01.00  CYCLE 1.000 SECONDS

NUMBER OF INTERVALS 11              TOTAL LENGTH OF INTERVALS 00.10.53
DATE   TIME      INT    CPU    DASD  DASD    JOB    JOB    TSO    TSO    STC    STC   ASCH   ASCH   OMVS   OMVS SWAP DEMAND
MM/DD HH.MM.SS MM.SS   BUSY   RESP  RATE    MAX    AVE    MAX    AVE    MAX    AVE    MAX    AVE    MAX    AVE RATE PAGING

09/26 14.50.00 00.59   14.8    0.6 410.3      0      0      3      2     90     90      0      0      7      5 0.00   0.03
09/26 14.51.00 00.59   14.9    0.5 346.2      0      0      3      2     90     90      0      0      5      5 0.00   0.03
09/26 14.52.00 00.59   16.1    0.8 145.1      0      0      2      2     90     90      0      0      7      5 0.00   0.00
09/26 14.53.00 01.00   15.4    0.6 194.9      0      0      2      2     90     90      0      0      7      5 0.00   0.00

09/26 14.54.00 00.59   16.1    0.6 207.0      0      0      2      2     90     90      0      0      7      5 0.00   0.00
09/26 14.55.00 01.00   15.3    0.5 182.0      0      0      2      2     90     90      0      0      5      5 0.00   0.00
09/26 14.56.00 00.59   15.7    0.5 210.4      0      0      2      2     90     90      0      0      6      5 0.00   0.00
09/26 14.57.00 01.00   15.4    0.5 223.0      0      0      2      2     90     90      0      0      5      5 0.00   0.00

09/26 14.58.00 00.59   15.4    0.5 226.7      0      0      2      2     90     90      0      0      7      5 0.00   0.00
09/26 14.59.00 01.00   16.2    0.6 520.6      1      1      2      2     91     90      0      0      7      5 0.00   0.00
09/26 15.00.00 00.59   15.9    0.7 328.9      0      0      2      2     91     91      0      0      5      5 0.00   0.00

TOTAL/AVERAGE          15.5    0.6 272.3      1      0      3      2     91     90      0      0      7      5 0.00   0.01
```

Figure A-48 RMF summary report

Interesting information is in the following columns:

CPUBUSY: To determine whether we had high CPU activity during the test. A separate, more detailed CPU activity report is also created to further investigate when there is a need to do so. In this case CPU activity is approximately 15% so there is no CPU bottleneck.

DASDRESP/DASDRATE: To determine whether we had good DASD response rates or high DASD activity rates. Response rates shown should be below 1.0 ms. Higher rates need further investigation. High DASD rates (values over x-thousand) might also be a reason for investigation. The reports Channel Path activity, I/O queuing activity, Cache subsystem activity, and DASD activity report are helpful. In this case response rates are good and the DASD rates are low, so there is no need for further investigation.

DEMAND PAGING: To determine whether we had high paging activity. Values over x-hundreds might be a reason to investigate what tasks are paging. The report Paging Activity will be helpful for this. Obviously in this case the system is not paging and there is no need for further investigation.

CPU activity report

This report is created because REPORTS(CPU) is specified. The report will be helpful for further investigation in situations with high CPU usage. The report itself consists of two parts, one part with details on the CPU activity, and one part containing Partition data.

The part containing the details on the CPU activity is shown in Figure A-49

```
                                        C P U   A C T I V I T Y
                                                                                                    PAGE    1
             z/OS V1R8               SYSTEM ID MVH1            START 09/26/2007-14.50.00  INTERVAL 000.10.59
                                     RPT VERSION V1R8 RMF      END   09/26/2007-15.01.00  CYCLE 1.000 SECONDS
CPU  2094   MODEL  736   H/W MODEL  S38

---CPU---  ONLINE TIME   LPAR BUSY       MVS BUSY      CPU SERIAL  I/O TOTAL         % I/O INTERRUPTS
NUM  TYPE  PERCENTAGE    TIME PERC       TIME PERC     NUMBER      INTERRUPT RATE    HANDLED VIA TPI
 0   CP    100.00        15.62           15.61         0181DF      78.27             22.01
 1   CP    100.00        17.60           17.59         0181DF      86.36             22.18
 2   CP    100.00        16.83           16.82         0181DF      87.90             21.59
 3   CP    100.00        17.88           17.87         0181DF      93.93             22.35
 4   CP    100.00        12.62           12.61         0181DF      75.88             20.98
 5   CP    100.00        13.45           13.44         0181DF      76.82             21.79
 6   CP    100.00        15.16           15.15         0181DF      79.04             20.84
 7   CP    100.00        15.22           15.21         0181DF      308.1              4.62
CP   TOTAL/AVERAGE       15.55           15.54                     886.3             15.76


SYSTEM ADDRESS SPACE ANALYSIS           SAMPLES =    660

NUMBER OF ADDRESS SPACES

        -------------------- QUEUE TYPES --------------------   --------- ADDRESS SPACE TYPES ----------
                    IN      OUT      OUT    LOGICAL  LOGICAL       BATCH     STC      TSO     ASCH    OMVS
              IN    READY   READY    WAIT   OUT RDY  OUT WAIT

MIN           53        1       0       0         0        38          0       87        2       0       5
MAX           59        4       0       0         0        45          1       91        3       0       7
AVG         54.6      2.5     0.0     0.0       0.0      42.5        0.1     90.0      2.0     0.0     5.0

DISTRIBUTION OF IN-READY QUEUE
 NUMBER OF              0    10   20   30   40   50   60   70   80   90   100
 ADDR SPACES    (%)     |....|....|....|....|....|....|....|....|....|....|
<=  N          100.0    >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
 =  N +  1       0.0
 =  N +  2       0.0
 =  N +  3       0.0
<=  N +  5       0.0
<=  N + 10       0.0
<=  N + 15       0.0
<=  N + 20       0.0
<=  N + 30       0.0
<=  N + 40       0.0
<=  N + 60       0.0
<=  N + 80       0.0
>   N + 80       0.0

N = NUMBER OF PROCESSORS ONLINE (8.0 ON AVG)
```

Figure A-49   CPU activity

Interesting information is the distribution of the IN-READY QUEUE. This gives an insight into how many tasks are waiting to use a processor. Obviously, this system had no tasks waiting for a processor during the tests. The number of tasks in the in-ready queue was always lower than the number of processors N.

Figure A-50 is an example of the partition data report.

```
                                     P A R T I T I O N   D A T A   R E P O R T
                                                                                                    PAGE    2
             z/OS V1R8               SYSTEM ID MVH1            START 09/26/2007-14.50.00  INTERVAL 000.10.59
                                     RPT VERSION V1R8 RMF      END   09/26/2007-15.01.00  CYCLE 1.000 SECONDS

MVS PARTITION NAME                     W631        NUMBER OF PHYSICAL PROCESSORS          38                GROUP NAME       N/A
IMAGE CAPACITY                          402                  CP                           36                     LIMIT       N/A
NUMBER OF CONFIGURED PARTITIONS           9                  AAP                           2
WAIT COMPLETION                          NO                  IFL                           0
DISPATCH INTERVAL                   DYNAMIC                  ICF                           0
                                                             IIP                           0

--------- PARTITION DATA -----------------  -- LOGICAL PARTITION PROCESSOR DATA --   -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --

                 ----MSU----  -CAPPING--  PROCESSOR-  ----DISPATCH TIME DATA----   LOGICAL PROCESSORS  --- PHYSICAL PROCESSORS ---
NAME       S   WGT  DEF    ACT  DEF   WLM%  NUM   TYPE   EFFECTIVE       TOTAL       EFFECTIVE    TOTAL  LPAR MGMT  EFFECTIVE  TOTAL

W631       A    10    0     63  NO     0.0    8   CP    00.13.40.098  00.13.40.935      15.53     15.55      0.00       3.45   3.46
W632       A    10    0      2  NO     0.0    8   CP    00.00.19.978  00.00.20.964       0.38      0.40      0.00       0.08   0.09
*PHYSICAL*                                                            00.00.04.225                          0.02              0.02
                                                        ------------  ------------                        ------     ------ ------
  TOTAL                                                 00.14.00.076  00.14.06.125                          0.03       3.54   3.56

W633       D
W634       D
W635       D
W636       D
W637       D
W638       D
W639       D
```

Figure A-50   Partition data report

This part of the CPU activity reports gives a good high-level insight of the entire CEC running in LPAR mode. As stated before, for the tests of the first three KPIs, two LPARs, each with a maximum of 8 CPs are defined. Some KPI tests use the LPARs simultaneously. This report is very helpful to see where most of the action is happening and to see for which LPAR detailed reports are needed: W631 or W632 or maybe both.

Channel Path activity report

This report is created because REPORTS(CHANNEL) is specified. It shows an average usage for all channels available. High DASD response rates (above 1 ms) might indicate channel contention. Figure A-51 is an example of this report.

```
                                             C H A N N E L   P A T H   A C T I V I T Y
                                                                                                                        PAGE    1
            z/OS V1R8                SYSTEM ID MVH1            START 09/26/2007-14.50.00  INTERVAL 000.10.59
                                     RPT VERSION V1R8 RMF      END   09/26/2007-15.01.00  CYCLE 1.000 SECONDS

  IODF = 80   CR-DATE: 08/21/2007   CR-TIME: 11.04.46   ACT: POR         MODE: LPAR       CPMF: EXTENDED MODE          CSSID: 0

-----------------------------------------------------------------------------------------------------------------------------------
                                              DETAILS FOR ALL CHANNELS
-----------------------------------------------------------------------------------------------------------------------------------
    CHANNEL PATH     UTILIZATION(%)     READ(MB/SEC) WRITE(MB/SEC)     CHANNEL PATH     UTILIZATION(%)     READ(MB/SEC) WRITE(MB/SEC)
  ID TYPE   G SHR  PART  TOTAL    BUS   PART  TOTAL   PART  TOTAL   ID TYPE   G SHR  PART  TOTAL    BUS   PART  TOTAL   PART  TOTAL

  20 OSD      Y    0.01  0.01   0.00   0.00  0.00   0.00  0.00   C1 FC_S   4  Y   0.64  0.71   0.14   0.55  0.58   0.21  0.22
  30 OSD           0.95  0.96   0.00   0.56  0.56   0.73  0.73   D0 FC_S   4  Y   0.64  0.71   0.15   0.55  0.59   0.21  0.22
  40 FC_S   4 Y    0.64  0.72   0.15   0.54  0.59   0.22  0.22   D1 FC_S   4  Y   0.64  0.71   0.15   0.55  0.60   0.21  0.22
  41 FC_S   4 Y    0.64  0.70   0.14   0.53  0.57   0.22  0.22
  50 FC_S   4 Y    0.64  0.71   0.14   0.54  0.58   0.21  0.22
  51 FC_S   4 Y    0.65  0.72   0.15   0.57  0.62   0.22  0.22
  70 CNC_?    Y    0.00  0.00
  C0 FC_S   4 Y    0.64  0.71   0.14   0.53  0.57   0.21  0.22

    CHANNEL PATH     WRITE(B/SEC)     MESSAGE RATE     MESSAGE SIZE   SEND FAIL    RECEIVE FAIL
  ID TYPE   G SHR     PART   TOTAL     PART   TOTAL     PART   TOTAL     PART       PART   TOTAL
  EC IQD      Y          0       0        0       0  ------- -------        0          0       0
  ED IQD      Y          0       0        0       0  ------- -------        0          0       0
  EE IQD      Y          0       0        0       0  ------- -------        0          0       0
  EF IQD      Y          0       0        0       0  ------- -------        0          0       0
```

Figure A-51 Channel Path activity report

Interesting information is the utilization on the important channel paths. In this case the important channels are 40, 41, 50, 51, C0, C1, D0, D1. Those are the FICON channels used for connecting to the DASD. High utilization would also mean high DASD activity or Cache subsystem activity or both. There are separate reports for those items.

Also important can be channel 30. This is the OSA-Express2 gigabit adapter for connecting the SAP application servers. High utilization on this channel path could indicate a network bottleneck.

The other channel paths are irrelevant.

DASD activity report

This report is created because REPORTS(DEVICE(DASD)) is specified. Figure A-52 is an example.

```
                                 D I R E C T   A C C E S S   D E V I C E   A C T I V I T Y
                                                                                                                  PAGE    1
          z/OS V1R8              SYSTEM ID MVH1           START 09/26/2007-14.50.00  INTERVAL 000.10.59
                                 RPT VERSION V1R8 RMF     END   09/26/2007-15.01.00  CYCLE 1.000 SECONDS

TOTAL SAMPLES =    660   IODF = 80   CR-DATE: 08/21/2007   CR-TIME: 11.04.46     ACT: POR
                                     DEVICE   AVG  AVG   AVG  AVG   AVG  AVG  AVG    %      %     %    AVG    %      %
STORAGE  DEV  DEVICE   VOLUME PAV  LCU  ACTIVITY RESP IOSQ  CMR  DB    PEND DISC CONN   DEV    DEV   DEV  NUMBER ANY    MT
 GROUP   NUM  TYPE     SERIAL           RATE     TIME TIME  DLY  DLY   TIME TIME TIME   CONN   UTIL  RESV ALLOC  ALLOC  PEND

SGDBOCAT C100 33909    DBHOD1   4  000F   1.305  0.6  0.0   0.0  0.0   0.2  0.1  0.4   0.01   0.01  0.0  51.1   100.0  0.0
SGDBOWRK C101 33909    DBH1W1   4  000F   0.156  0.4  0.0   0.0  0.0   0.1  0.0  0.2   0.00   0.00  0.0   3.0   100.0  0.0
SGHKMC1  C102 33909    HK1000   5  000F   5.846  0.4  0.0   0.0  0.0   0.1  0.1  0.2   0.02   0.03  0.0  36.0   100.0  0.0
SGHKMC1  C103 33909    HK1001   4  000F   5.402  0.4  0.0   0.0  0.0   0.2  0.0  0.2   0.03   0.03  0.0  25.0   100.0  0.0
SGHKMC1  C104 33909    HK1002   5  000F   0.770  0.3  0.0   0.0  0.0   0.2  0.0  0.1   0.00   0.00  0.0  33.0   100.0  0.0
SGHKMC1  C105 33909    HK1003   4  000F   0.344  1.5  0.0   0.0  0.0   0.2  0.9  0.4   0.00   0.01  0.0  39.0   100.0  0.0
SGHKMC1  C106 33909    HK1004   5  000F   5.017  0.7  0.0   0.0  0.0   0.1  0.1  0.4   0.04   0.05  0.0  27.0   100.0  0.0
SGHKMC1  C107 33909    HK1005   4  000F   4.812  1.4  0.0   0.0  0.0   0.2  1.0  0.2   0.03   0.15  0.0  65.0   100.0  0.0
SGHKMC1  C108 33909    HK1006   4  000F   0.068  0.4  0.0   0.0  0.0   0.2  0.0  0.2   0.00   0.00  0.0  29.0   100.0  0.0
SGHKMC1  C109 33909    HK1007   4  000F   2.580  0.4  0.0   0.0  0.0   0.1  0.1  0.2   0.01   0.02  0.0  34.0   100.0  0.0
SGHKMC1  C10A 33909    HK1008   6  000F   1.059  0.8  0.0   0.0  0.0   0.2  0.1  0.6   0.01   0.01  0.0  36.0   100.0  0.0
SGHKMC1  C10B 33909    HK1009   4  000F  20.571  0.4  0.0   0.0  0.0   0.1  0.0  0.3   0.14   0.14  0.0  39.0   100.0  0.0
SGHKMC1  C10C 33909    HK1010   4  000F   2.329  1.0  0.0   0.0  0.0   0.2  0.4  0.4   0.02   0.05  0.0  45.0   100.0  0.0
SGHKMC1  C10D 33909    HK1011   7  000F   1.479  0.4  0.0   0.0  0.0   0.1  0.0  0.2   0.00   0.00  0.0  41.0   100.0  0.0
SGHKMC1  C10E 33909    HK1012   4  000F   3.874  0.4  0.0   0.0  0.0   0.1  0.1  0.2   0.02   0.02  0.0  37.0   100.0  0.0
SGHKMC1  C10F 33909    HK1013   4  000F   0.392  1.4  0.0   0.1  0.0   0.3  0.1  1.0   0.01   0.01  0.0  38.0   100.0  0.0
|...............................
devices C110 thru C12F not shown
...............................
         C130 33909    F2C130   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C131 33909    F2C131   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C132 33909    F2C132   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C133 33909    F2C133   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C134 33909    F2C134   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C135 33909    F2C135   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C136 33909    F2C136   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C137 33909    F2C137   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C138 33909    F2C138   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C139 33909    F2C139   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C13A 33909    F2C13A   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C13B 33909    F2C13B   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C13C 33909    F2C13C   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C13D 33909    F2C13D   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C13E 33909    F2C13E   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0
         C13F 33909    F2C13F   4  000F   0.000  0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0   0.0   100.0  0.0

                        LCU        000F  75.469  0.6  0.0   0.0  0.0   0.2  0.1  0.4   0.01   0.01  0.0   578   100.0  0.0
```

Figure A-52 Direct Access Device Activity

The report shows a line for each connected DASD devices. With this report the devices showing high activity rates can be identified, and more importantly, so can devices with high average response rates. Average response rates below 1.0 msecs are considered good rates.

After determining which devices show high rates, the reason for the high rates should be investigated. Because during the first tests dynamic PAV / HiperPAV was used (During test for the first KPIs dynamic PAV is used. Later on HiperPAV, the better version of PAV, is switched on) the values for IOS-queuing should be negligible.

High response rates can be caused by poor cache hit ratios for the DASD cache. High activity rates could indicate inefficient SQLs causing table/index-scan or bufferpools with low hit ratios. High average pending times should be investigated. This can be done by using the I/O queuing activity report.

I/O queuing activity report

This report is created because REPORTS(IOQ) is specified. Figure A-53 shows a example.

```
                                    I/O  Q U E U I N G   A C T I V I T Y
                                                                                                        PAGE    1
           z/OS V1R8              SYSTEM ID MVH1           START 09/26/2007-14.50.00  INTERVAL 000.10.59
                                  RPT VERSION V1R8 RMF     END   09/26/2007-15.01.00  CYCLE 1.000 SECONDS
TOTAL SAMPLES =    660   IODF = 80   CR-DATE: 08/21/2007   CR-TIME: 11.04.46     ACT: POR
      - INITIATIVE QUEUE -  ------- IOP UTILIZATION -------   -- % I/O REQUESTS RETRIED --   -------- RETRIES / SSCH ---------
 IOP      ACTIVITY   AVG Q    % IOP   I/O START    INTERRUPT         CP   DP   CU   DV            CP     DP     CU     DV
            RATE     LNGTH    BUSY      RATE         RATE      ALL BUSY BUSY BUSY BUSY      ALL  BUSY   BUSY   BUSY   BUSY

 00        73.792     0.00    0.08     73.792       73.984     0.0  0.0  0.0  0.0  0.0     0.00  0.00   0.00   0.00   0.00
 01         0.000     0.00    0.00      0.000        0.000    **** **** **** **** ****    ***** *****  *****  *****  *****
 02         0.000     0.00    0.00      0.000        0.000    **** **** **** **** ****    ***** *****  *****  *****  *****
 03        73.730     0.00    0.08     73.730       73.825     0.0  0.0  0.0  0.0  0.0     0.00  0.00   0.00   0.00   0.00
 04         0.000     0.00    0.00      0.000        0.000    **** **** **** **** ****    ***** *****  *****  *****  *****
 05        73.545     0.00    0.06     73.545       74.028     0.0  0.0  0.0  0.0  0.0     0.00  0.00   0.00   0.00   0.00
 06         0.261     0.00    0.34      0.261      616.172     4.4  4.4  0.0  0.0  0.0     0.05  0.05   0.00   0.00   0.00
 07        74.892     0.00    0.07     74.892       76.789     0.0  0.0  0.0  0.0  0.0     0.00  0.00   0.00   0.00   0.00

 SYS      296.218     0.00    0.08    296.218      914.797     0.0  0.0  0.0  0.0  0.0     0.00  0.00   0.00   0.00   0.00

 LCU    CU     DCM GROUP   CHAN     CHPID    % DP   % CU   CUB   CMR  CONTENTION      Q   CSS    HPAV
              MIN MAX DEF  PATHS    TAKEN    BUSY   BUSY   DLY   DLY     RATE     LNGTH   DLY  WAIT  MAX

 000D   A000               50       0.115    0.00   0.00   0.0   0.0
                           D1       0.068    0.00   0.00   0.0   0.0
                            *       0.183    0.00   0.00   0.0   0.0     0.000     0.00   0.1

 000E   C010               40       3.544    0.00   0.00   0.0   0.0
                           41       2.953    0.00   0.00   0.0   0.0
                           50       2.938    0.00   0.00   0.0   0.0
                           51       2.994    0.00   0.00   0.0   0.0
                           C0       2.923    0.00   0.00   0.0   0.0
                           C1       2.999    0.00   0.00   0.0   0.0
                           D0       2.927    0.00   0.00   0.0   0.0
                           D1       3.039    0.00   0.00   0.0   0.0
                            *      24.317    0.00   0.00   0.0   0.0     0.000     0.00   0.1
...........................
cu C011 thru C019 not shown
...........................
 0018   C01A               40       0.009    0.00   0.00   0.0   0.0
                           41       0.017    0.00   0.00   0.0   0.0
                           50       0.015    0.00   0.00   0.0   0.0
                           51       0.018    0.00   0.00   0.0   0.0
                           C0       0.021    0.00   0.00   0.0   0.0
                           C1       0.024    0.00   0.00   0.0   0.0
                           D0       0.012    0.00   0.00   0.0   0.0
                           D1       0.017    0.00   0.00   0.0   0.0
                            *       0.133    0.00   0.00   0.0   0.0     0.000     0.00   0.1
```

Figure A-53   I/O queuing activity

The information in this report can be helpful to further investigate high pending times in the DASD activity report. A common solution to avoid pending times is to change the distribution of the data over the available control units.

Cache subsystem activity report

This report is created because REPORTS(CACHE) is specified.

The Cache Subsystem Activity report provides cache statistics on a subsystem basis as well as on a detailed device-level basis. Therefore the reports consists of two parts. Figure A-54 shows the part with the cache statistics on the subsystem level.

```
                                   C A C H E   S U B S Y S T E M   A C T I V I T Y
                                                                                                                        PAGE    1
            z/OS V1R8                SYSTEM ID MVH1            START 09/26/2007-14.50.00  INTERVAL 000.09.59
                                     RPT VERSION V1R8 RMF      END   09/26/2007-15.00.00

 SUBSYSTEM  2107-01     CU-ID  C101     SSID C100     CDATE  09/26/2007     CTIME 14.50.00      CINT  00.09.57
 TYPE-MODEL 2107-932    MANUF  IBM      PLANT  75     SERIAL 0000000M1041

 ------------------------------------------------------------------------------------------------------------------------------------
                                                         CACHE SUBSYSTEM STATUS
 ------------------------------------------------------------------------------------------------------------------------------------

 SUBSYSTEM STORAGE            NON-VOLATILE STORAGE         STATUS

 CONFIGURED       62816M      CONFIGURED      2048.0M      CACHING                   - ACTIVE
 AVAILABLE        55898M      PINNED              0.0      NON-VOLATILE STORAGE      - ACTIVE
 PINNED              0.0                                   CACHE FAST WRITE          - ACTIVE
 OFFLINE             0.0                                   IML DEVICE AVAILABLE      - YES

 ------------------------------------------------------------------------------------------------------------------------------------
                                                        CACHE SUBSYSTEM OVERVIEW
 ------------------------------------------------------------------------------------------------------------------------------------

 TOTAL I/O     51806   CACHE I/O     51806   CACHE OFFLINE        0
 TOTAL H/R     0.987   CACHE H/R     0.987

 CACHE I/O     -------------READ I/O REQUESTS-------------    ----------------------WRITE I/O REQUESTS----------------------        %
 REQUESTS        COUNT      RATE       HITS      RATE     H/R      COUNT      RATE       FAST      RATE       HITS      RATE     H/R      READ

 NORMAL          28249      47.3      27595      46.2   0.977       5814       9.7       5814       9.7       5814       9.7   1.000      82.9
 SEQUENTIAL      17501      29.3      17478      29.3   0.999        242       0.4        242       0.4        242       0.4   1.000      98.6
 CFW DATA            0       0.0          0       0.0     N/A          0       0.0          0       0.0          0       0.0     N/A       N/A
 TOTAL           45750      76.6      45073      75.5   0.985       6056      10.1       6056      10.1       6056      10.1   1.000      88.3

-----------------------CACHE MISSES-----------------------            ------------MISC------------          ------NON-CACHE I/O-----
 REQUESTS        READ     RATE    WRITE     RATE   TRACKS     RATE                            COUNT   RATE                      COUNT   RATE
                                                                       DFW BYPASS                 0    0.0          ICL               0    0.0
 NORMAL           654      1.1        0      0.0      823      1.4      CFW BYPASS                 0    0.0          BYPASS            0    0.0
 SEQUENTIAL        23      0.0        0      0.0    23586     39.5      DFW INHIBIT                0    0.0          TOTAL             0    0.0
 CFW DATA           0      0.0        0      0.0                        ASYNC (TRKS)            3574    6.0

 TOTAL            677     RATE      1.1

 ---CKD STATISTICS---      ---RECORD CACHING---                         ----HOST ADAPTER ACTIVITY---       --------DISK ACTIVITY-------
                                                                                       BYTES   BYTES                 RESP   BYTES   BYTES
 WRITE              591     READ MISSES        0                                        /REQ    /SEC                 TIME    /REQ    /SEC
 WRITE HITS         589     WRITE PROM      3216                        READ           37.9K    2.9M       READ     1.414   56.6K    2.3M
                                                                        WRITE           8.0K   81.0K       WRITE   26.350   96.4K  279.9K
```

Figure A-54 Cache Subsystem Activity (part 1)

Figure A-55 on page 266 shows the part with the device-level statistics.

```
                                      C A C H E   S U B S Y S T E M   A C T I V I T Y
                                                                                                                      PAGE    2
            z/OS V1R8                SYSTEM ID MVH1            START 09/26/2007-14.50.00  INTERVAL 000.09.59
                                     RPT VERSION V1R8 RMF      END   09/26/2007-15.00.00

SUBSYSTEM  2107-01     CU-ID  C101     SSID C100     CDATE  09/26/2007     CTIME 14.50.00      CINT  00.09.57
TYPE-MODEL 2107-932    MANUF  IBM      PLANT  75     SERIAL 0000000M1041

------------------------------------------------------------------------------------------------------------------------------------
                                                         CACHE SUBSYSTEM DEVICE OVERVIEW
------------------------------------------------------------------------------------------------------------------------------------

VOLUME     DEV  XTNT     %    I/O   ---CACHE HIT RATE--   ----------DASD I/O RATE----------    ASYNC   TOTAL    READ   WRITE      %
SERIAL     NUM  POOL   I/O   RATE    READ    DFW    CFW   STAGE  DFWBP    ICL    BYP  OTHER     RATE     H/R     H/R     H/R   READ

*ALL                 100.0   86.8    75.5   10.1    0.0     1.1    0.0    0.0    0.0    0.0      6.0   0.987   0.985   1.000   88.3
*CACHE-OFF             0.0    0.0
*CACHE               100.0   86.8    75.5   10.1    0.0     1.1    0.0    0.0    0.0    0.0      6.0   0.987   0.985   1.000   88.3
DBHOD1    C100  0001   1.8    1.5     1.4    0.2    0.0     0.0    0.0    0.0    0.0    0.0      0.0   0.998   0.998   1.000   88.0
DBH1W1    C101  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
HK1000    C102  0001   6.3    5.5     4.9    0.5    0.0     0.0    0.0    0.0    0.0    0.0      0.4   0.993   0.992   1.000   90.0
HK1001    C103  0001   6.5    5.6     4.0    1.6    0.0     0.0    0.0    0.0    0.0    0.0      1.7   0.993   0.991   1.000   72.2
HK1002    C104  0001   1.0    0.8     0.8    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0   0.998   0.998   1.000   97.2
HK1003    C105  0001   0.3    0.3     0.2    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0   0.843   0.841   1.000   98.8
HK1004    C106  0001   5.7    4.9     4.8    0.1    0.0     0.1    0.0    0.0    0.0    0.0      0.0   0.987   0.987   1.000   98.8
HK1005    C107  0001   5.5    4.8     4.0    0.0    0.0     0.8    0.0    0.0    0.0    0.0      0.0   0.836   0.835   1.000   99.3
HK1006    C108  0001   0.1    0.1     0.1    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0   1.000   1.000   1.000   75.0
HK1007    C109  0001   3.7    3.2     3.2    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0   0.990   0.990   1.000   99.2
HK1008    C10A  0001   2.2    1.9     1.0    0.9    0.0     0.0    0.0    0.0    0.0    0.0      0.8   0.999   0.998   1.000   53.4
HK1009    C10B  0001  35.7   31.0    24.6    6.4    0.0     0.0    0.0    0.0    0.0    0.0      2.5   1.000   1.000   1.000   79.3
HK1010    C10C  0001   2.3    2.0     1.8    0.1    0.0     0.1    0.0    0.0    0.0    0.0      0.1   0.966   0.963   1.000   92.8
HK1011    C10D  0001   2.0    1.7     1.7    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0   0.998   0.998   1.000   99.7
HK1012    C10E  0001   4.9    4.2     4.2    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0   0.995   0.995   1.000   99.5
HK1013    C10F  0001   0.4    0.3     0.1    0.2    0.0     0.0    0.0    0.0    0.0    0.0      0.4   0.944   0.861   1.000   40.5
................................
devices C110 thru C12F not shown
................................
F2C130    C130  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C131    C131  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C132    C132  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C133    C133  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C134    C134  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C135    C135  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C136    C136  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C137    C137  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C138    C138  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C139    C139  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C13A    C13A  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C13B    C13B  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C13C    C13C  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C13D    C13D  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C13E    C13E  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
F2C13F    C13F  0001   0.0    0.0     0.0    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0     N/A     N/A     N/A    N/A
```

Figure A-55 Cache Subsystem Activity (part 2)

This report provides insight into the details concerning the cache and determine the optimum size of the cache, the amount of needed non-volatile storage (NVS).

Paging activity report

This report is created because REPORTS(PAGING) is specified. Figure A-56 shows the first page of this report.

```
                                          P A G I N G   A C T I V I T Y
                                                                                                                  PAGE    1
          z/OS V1R8            SYSTEM ID MVH1          START 09/26/2007-14.50.00  INTERVAL 000.10.59
                               RPT VERSION V1R8 RMF     END   09/26/2007-15.01.00  CYCLE 1.000 SECONDS

OPT = IEAOPT00    MODE = ESAME              CENTRAL STORAGE PAGING RATES - IN PAGES PER SECOND
-----------------------------------------------------------------------------------------------------------------------------------

                          PAGE IN                                   PAGE OUT
                 ------------------------------------       ----------------------------
                                 NON              %                                     %
                         NON     SWAP,            OF                                    OF
                         SWAP,   NON     TOTAL   TOTL                NON     TOTAL     TOTL
CATEGORY         SWAP    BLOCK   BLOCK   RATE    SUM        SWAP     SWAP    RATE      SUM

PAGEABLE SYSTEM
AREAS (NON VIO)
     LPA                  0.00    0.01     0.01  100

     CSA                  0.00    0.00     0.00    0                 0.00    0.00       0
                         ------  ------   ------ ----              ------  ------    ----
     SUM                  0.00    0.01     0.01  100                 0.00    0.00       0

ADDRESS SPACES
     HIPERSPACE           0.00             0.00    0                 0.00    0.00       0

     VIO                  0.00             0.00    0                 0.00    0.00       0

     NON VIO      0.00    0.00    0.00     0.00    0        0.00     0.00    0.00       0
                 ------  ------  ------   ------ ----     ------   ------  ------    ----
     SUM          0.00    0.00    0.00     0.00    0        0.00     0.00    0.00       0

TOTAL SYSTEM
     HIPERSPACE           0.00             0.00    0                 0.00    0.00       0

     VIO                  0.00             0.00    0                 0.00    0.00       0

     NON VIO      0.00    0.00    0.01     0.01  100        0.00     0.00    0.00       0
                 ------  ------  ------   ------ ----     ------   ------  ------    ----
     SUM          0.00    0.00    0.01     0.01  100        0.00     0.00    0.00       0

     SHARED                       0.00     0.00                      0.00    0.00

PAGE MOVEMENT WITHIN CENTRAL STORAGE      39.18             PAGE MOVEMENT TIME %      0.0
AVERAGE NUMBER OF PAGES PER BLOCK          0.0
BLOCKS PER SECOND                          0.00
PAGE-IN EVENTS (PAGE FAULT RATE)           0.01
```

Figure A-56 Paging Activity

The report can be used to determine whether there is sufficient central storage. Obviously this is the case because the report shows negligible paging activity.

Workload Activity report

This report is created because SYSRPTS(WLMGL(SCLASS,SCPER,POLICY,RCLASS)) is specified. In general the report shows information about the defined versus the achieved goals. The information can be used to determine what workloads are not meeting their goals. The information will be meaningfull in a constrained situation, meaning WLM has to determine where the resources should be used to achieve the defined goals.

Figure A-57 shows an example.

```
                              W O R K L O A D   A C T I V I T Y
                                                                                              PAGE  10
     z/OS V1R8              SYSPLEX PORTINH1          START 09/26/2007-14.50.00 INTERVAL 000.10.59   MODE = GOAL
                            RPT VERSION V1R8 RMF      END   09/26/2007-15.00.59

                                    POLICY ACTIVATION DATE/TIME 09/13/2007 11.24.57

------------------------------------------------------------------------------------------------------------ SERVICE CLASS PERIODS

REPORT BY: POLICY=POL1       WORKLOAD=OMVS       SERVICE CLASS=VEL50_O3   RESOURCE GROUP=*NONE      PERIOD=1 IMPORTANCE=3
                                                 CRITICAL     =NONE

TRANSACTIONS   TRANS-TIME HHH.MM.SS.TTT  --DASD I/O--  ---SERVICE----  SERVICE TIMES  ---APPL %---  PAGE-IN RATES    ---STORAGE---
AVG      3.05  ACTUAL                91  SSCHRT   2.4  IOC      47175  CPU       2.0  CP      0.34  SINGLE      0.0  AVG    5470.99
MPL      3.05  EXECUTION             90  RESP     0.6  CPU      49171  SRB       0.0  AAPCP   0.00  BLOCK       0.0  TOT   16668.59
ENDED     130  QUEUED                 0  CONN     0.5  MSO      37627  RCT       0.2  IIPCP   0.00  SHARED      0.0  CEN   16668.59
END/S    0.20  R/S AFFIN              0  DISC     0.0  SRB        323  IIT       0.0                HSP         0.0  EXP       0.00
#SWAPS    139  INELIGIBLE             0  Q+PEND   0.1  TOT     134296  HST       0.0  AAP      N/A  HSP MISS    0.0
EXCTD       0  CONVERSION             0  IOSQ     0.0  /SEC       203  AAP       N/A  IIP      N/A  EXP SNGL    0.0  SHR     222.70
AVG ENC  0.00  STD DEV               29                                IIP       N/A                EXP BLK     0.0
REM ENC  0.00                                          ABSRPTN    67                                EXP SHR     0.0
MS ENC   0.00                                          TRX SERV   67

GOAL: EXECUTION VELOCITY 50.0%     VELOCITY MIGRATION:   I/O MGMT   100%     INIT MGMT  100%

         RESPONSE TIME EX   PERF  AVG   ------ USING% ----- ------------ EXECUTION DELAYS % ------------ ---DLY%-- -CRYPTO%-    %
SYSTEM                 VEL% INDX ADRSP   CPU  AAP  IIP  I/O  TOT                                         UNKN IDLE  USG  DLY QUIE

MVH1        --N/A--     100  0.5   5.0   0.1  N/A  N/A  0.0  0.0                                         40.8 59.2  0.0  0.0  0.0

------------------------------------------------------------------------------------------------------------ SERVICE CLASS(ES)

REPORT BY: POLICY=POL1       WORKLOAD=OMVS       SERVICE CLASS=VEL50_O3   RESOURCE GROUP=*NONE
                                                 CRITICAL     =NONE
                                                 DESCRIPTION  =vel 50 importance 3 for OMVS

TRANSACTIONS   TRANS-TIME HHH.MM.SS.TTT  --DASD I/O--  ---SERVICE----  SERVICE TIMES  ---APPL %---  PAGE-IN RATES    ---STORAGE---
AVG      3.05  ACTUAL                91  SSCHRT   2.4  IOC      47175  CPU       2.0  CP      0.34  SINGLE      0.0  AVG    5470.99
MPL      3.05  EXECUTION             90  RESP     0.6  CPU      49171  SRB       0.0  AAPCP   0.00  BLOCK       0.0  TOT   16668.59
ENDED     130  QUEUED                 0  CONN     0.5  MSO      37627  RCT       0.2  IIPCP   0.00  SHARED      0.0  CEN   16668.59
END/S    0.20  R/S AFFIN              0  DISC     0.0  SRB        323  IIT       0.0                HSP         0.0  EXP       0.00
#SWAPS    139  INELIGIBLE             0  Q+PEND   0.1  TOT     134296  HST       0.0  AAP      N/A  HSP MISS    0.0
EXCTD       0  CONVERSION             0  IOSQ     0.0  /SEC       203  AAP       N/A  IIP      N/A  EXP SNGL    0.0  SHR     222.70
AVG ENC  0.00  STD DEV               29                                IIP       N/A                EXP BLK     0.0
REM ENC  0.00                                          ABSRPTN    67                                EXP SHR     0.0
MS ENC   0.00                                          TRX SERV   67

----------------------------------------------SERVICE CLASSES BEING SERVED--------------------------------------------------------
SAPHIGH
```

Figure A-57 Workload activity

This reports shows information about service class VEL50_03. For this service class the goal execution velocity of 50.0% is defined in WLM. All work in this service class is meeting its goals.

A.5.2 OMEGAMON XE for DB2 Performance Expert (V3)

The batch reporting utility DB2PM from OMEGAMON XE for DB2 Performance Expert (V3) is used to create some relevant character-based reports. As stated previously, the recorded SMF records during the tests of the KPIs are culled from the entire recorded SMF data. This subset of SMF data is used as input for DB2PM. An example of the parameters used as input for DB2PM is shown in Figure A-57.

```
GLOBAL
  TIMEZONE(-2)
SYSPARMS TRACE DDNAME(SYTRCDD1)

ACCOUNTING
      REPORT
      LAYOUT(LONG)
      INCLUDE(SUBSYSTEMID(DBH1))

STATISTICS
      REPORT
      LAYOUT(LONG)
      INCLUDE(SUBSYSTEMID(DBH1))
EXEC
```

Figure A-58 Control statement for OMEGAMON XE, DB2PM

The following reports are created:

- System parameters report
- Accounting report (long)
- Statistics report (long)

In the following sections a brief description of each report is given, along with an explanation of the purpose it served during the tests. For a more detailed discussion refer to the OMEGAMON XE for DB2 Performance expert manuals.

System parameters report

This report is created because the statement SYSPARMS TRACE DDNAME(SYTRCDD1) is specified for the DB2PM utility. The report lists all the DB2 parameters that were in effect during the reporting interval. An example is in Figure A-59.

```
   LOCATION: HKMCDBH1                          OMEGAMON XE FOR DB2 PERFORMANCE EXPERT (V3)            PAGE: 1-1
      GROUP: N/P                                  SYSTEM PARAMETERS REPORT
     MEMBER: N/P
  SUBSYSTEM: DBH1                                                                              ACTUAL FROM: 09/26/07 14:50:12.25
DB2 VERSION: V8

 MVS PARMLIB UPDATE PARAMETERS (DSNTIPM)                           IRLM INSTALLATION PARAMETERS (DSNTIPI)
 ---------------------------------------                           --------------------------------------
 SUBSYSTEM DEFAULT (SSID)..................................DBH1    IRLM SUBSYSTEM NAME (IRLMSID)..............................DJH1
 SUPPRESS SOFT ERRORS (SUPERRS).............................YES    IRLM RESOURCE TIMEOUT IN SECONDS (IRLMRWT)..................600
                                                                   IRLM AUTOMATIC START (IRLMAUT).............................YES
 STORAGE SIZES INSTALLATION PARAMETERS (DSNTIPC,DSNTIPE)           IRLM START PROCEDURE NAME (IRLMPRC)...................DBH1IRLM
 -------------------------------------------------------           SECONDS DB2 WILL WAIT FOR IRLM START (IRLMSWT)..............300
 MAX NO OF USERS CONCURRENTLY RUNNING IN DB2 (CTHREAD).....1,024   UTILITY TIMEOUT FACTOR (UTIMOUT)..............................3
 MAX NO OF TSO CONNECTIONS (IDFORE)..........................100   U LOCK FOR REPEATABLE READ OR READ STABILITY (RRULOCK).......NO
 MAX NO OF BATCH CONNECTIONS (IDBACK)........................200   X LOCK FOR SEARCHED UPDATE/DELETE (XLKUPDLT)................YES
 MAX NO OF REMOTE CONNECTIONS (CONDBAT)...................30,000   IMS/BMP TIMEOUT FACTOR (BMPTOUT)..............................4
 MAX NO OF CONCURRENT REMOTE ACTIVE CONNECTIONS (MAXDBAT)....512   IMS/DLI TIMEOUT FACTOR (DLITOUT)..............................6
 MAXIMUM SIZE OF EDM POOL IN BYTES (EDMPOOL).........155,889,664   WAIT FOR RETAINED LOCKS (RETLWAIT)............................1
 MAXIMUM SIZE OF SORT POOL IN BYTES (SRTPOOL).........28,672,000   ENABLE DB CHECKING...........................................NO
 MAXIMUM SIZE OF RID POOL IN BYTES (MAXRBLK).........102,400,000
 3990 CACHE (SEQCACH).......................................SEQ    IRLM PROCESSING PARAMETERS
 UTILITY CACHE OPTION (SEQPRES).............................YES    --------------------------
 .... lines excluded ....
 .... lines excluded ....
 .... lines excluded ....

 TIMESTAMP              09/26/07 14:50:12.25          VIRTUAL POOL TYPE                                       N/A
 BUFFER POOL ID                        BP8K0          VIRTUAL POOL SIZE                                     10000
                                                      HIPERPOOL SIZE                                          N/A
                                                      VIRTUAL POOL SEQUENTIAL THRESHOLD                        40
                                                      HIPERPOOL SEQUENTIAL THRESHOLD                          N/A
                                                      HORIZONTAL DEFERRED WRITE THRESHOLD                      50
                                                      VERTICAL DEFERRED WRITE THRESHOLD (PERCENTAGE)           30
                                                      VERTICAL DEFERRED WRITE THRESHOLD (BUFFERS)               0
                                                      VIRTUAL POOL PARALLEL SEQUENTIAL THRESHOLD               50
                                                      ASSISTING PARALLEL SEQUENTIAL THRESHOLD                   0
                                                      PGFIX ATTRIBUTE                                         YES
                                                      PAGE STEAL METHOD                                       LRU

 TIMESTAMP              09/26/07 14:50:12.25          VIRTUAL POOL TYPE                                       N/A
 BUFFER POOL ID                       BP16K0          VIRTUAL POOL SIZE                                      5000
                                                      HIPERPOOL SIZE                                          N/A
                                                      VIRTUAL POOL SEQUENTIAL THRESHOLD                        80
                                                      HIPERPOOL SEQUENTIAL THRESHOLD                          N/A
                                                      HORIZONTAL DEFERRED WRITE THRESHOLD                      30
                                                      VERTICAL DEFERRED WRITE THRESHOLD (PERCENTAGE)            5
                                                      VERTICAL DEFERRED WRITE THRESHOLD (BUFFERS)               0
                                                      VIRTUAL POOL PARALLEL SEQUENTIAL THRESHOLD               50
                                                      ASSISTING PARALLEL SEQUENTIAL THRESHOLD                   0
                                                      PGFIX ATTRIBUTE                                         YES
                                                      PAGE STEAL METHOD                                       LRU
```

Figure A-59 DB2PM Sysparms report

During the tests of the KPIs no system parameters were changed.

Accounting report (long)

This report is created because the statement ACCOUNTING REPORT LAYOUT(LONG) INCLUDE(SUBSYSTEMID(DBH1)) is specified for the DB2PM utility. The report lists a wealth of information about the DB2 activity during the KPI test, as shown in the sample report in Figure A-60.

```
   LOCATION: HKMCDBH1                OMEGAMON XE FOR DB2 PERFORMANCE EXPERT (V3)                 PAGE: 1-1
      GROUP: N/P                              ACCOUNTING REPORT - LONG                REQUESTED FROM: NOT SPECIFIED
     MEMBER: N/P                                                                                 TO: NOT SPECIFIED
  SUBSYSTEM: DBH1                              ORDER: PRIMAUTH-PLANNAME                INTERVAL FROM: 09/26/07 14:50:16.97
DB2 VERSION: V8                                    SCOPE: MEMBER                                  TO: 09/26/07 15:00:59.93

PRIMAUTH: SAPD2LC  PLANNAME: DISTSERV

ELAPSED TIME DISTRIBUTION                                              CLASS 2 TIME DISTRIBUTION
----------------------------------------------------------------       ----------------------------------------------------------------
APPL   |========================================> 83%                  CPU     |===========> 23%
DB2    |==> 4%                                                         NOTACC  |=> 2%
SUSP   |======> 13%                                                    SUSP    |====================================> 75%

AVERAGE       APPL(CL.1)  DB2 (CL.2)  IFI (CL.5)    CLASS 3 SUSPENSIONS    AVERAGE TIME  AV.EVENT    HIGHLIGHTS
------------  ----------  ----------  ----------    --------------------   ------------  --------    --------------------------
ELAPSED TIME    0.396590    0.066851         N/P    LOCK/LATCH(DB2+IRLM)       0.005514      0.83    #OCCURRENCES    :      6185
 NONNESTED      0.396211    0.066807         N/A    SYNCHRON. I/O              0.041950     28.92    #ALLIEDS        :         3
 STORED PROC    0.000378    0.000044         N/A     DATABASE I/O              0.038012     27.97    #ALLIEDS DISTRIB:         0
 UDF            0.000000    0.000000         N/A     LOG WRITE I/O             0.003938      0.95    #DBATS          :      1235
 TRIGGER        0.000000    0.000000         N/A    OTHER READ I/O             0.002549      0.72    #DBATS DISTRIB. :         0
                                                    OTHER WRTE I/O             0.000006      0.00    #NO PROGRAM DATA:      1238
CP CPU TIME     0.020684    0.015466         N/P    SER.TASK SWTCH             0.000297      0.00    #NORMAL TERMINAT:      1238
 AGENT          0.020684    0.015466         N/A     UPDATE COMMIT             0.000000      0.00    #ABNORMAL TERMIN:         0
  NONNESTED     0.020609    0.015443         N/P     OPEN/CLOSE                0.000005      0.00    #CP/X PARALLEL. :      1238
  STORED PRC    0.000076    0.000023         N/A     SYSLGRNG REC              0.000013      0.00    #IO PARALLELISM :         0
  UDF           0.000000    0.000000         N/A     EXT/DEL/DEF               0.000263      0.00    #INCREMENT. BIND:         0
  TRIGGER       0.000000    0.000000         N/A     OTHER SERVICE             0.000016      0.00    #COMMITS        :      6185
 PAR.TASKS      0.000000    0.000000         N/A    ARC.LOG(QUIES)             0.000000      0.00    #ROLLBACKS      :         0
                                                    ARC.LOG READ               0.000000      0.00    #SVPT REQUESTS  :         0
 IIPCP CPU      0.011315         N/A         N/A    DRAIN LOCK                 0.000000      0.00    #SVPT RELEASE   :         0
                                                    CLAIM RELEASE              0.000000      0.00    #SVPT ROLLBACK  :         0
IIP CPU TIME    0.000000    0.000000         N/A    PAGE LATCH                 0.000019      0.05    MAX SQL CASC LVL:         1
                                                    NOTIFY MSGS                0.000000      0.00    UPDATE/COMMIT   :     51.08
SUSPEND TIME    0.000100    0.050335         N/A    GLOBAL CONTENTION          0.000000      0.00    SYNCH I/O AVG.  :  0.001451
 AGENT               N/A    0.050335         N/A    COMMIT PH1 WRITE I/O       0.000000      0.00
 PAR.TASKS           N/A    0.000000         N/A    ASYNCH CF REQUESTS         0.000000      0.00
 STORED PROC    0.000100         N/A         N/A    TOTAL CLASS 3              0.050335     30.53
 UDF            0.000000         N/A         N/A

NOT ACCOUNT.         N/A    0.001050         N/A
DB2 ENT/EXIT         N/A         N/P         N/A
EN/EX-STPROC         N/A        1.31         N/A
EN/EX-UDF            N/A        0.00         N/A
DCAPT.DESCR.         N/A         N/A         N/P
LOG EXTRACT.         N/A         N/A         N/P
.... lines excluded ....
.... lines excluded ....
.... lines excluded ....
TOTAL  BPOOL ACTIVITY   AVERAGE     TOTAL
---------------------  --------  --------
BPOOL HIT RATIO (%)       91.75       N/A
GETPAGES                 856.35   5296521
BUFFER UPDATES           188.56   1166213
SYNCHRONOUS WRITE          0.18      1127
SYNCHRONOUS READ          27.79    171859
SEQ. PREFETCH REQS         0.30      1831
```

Figure A-60   DB2PM Accounting report

A particularly interesting part of this report is the distribution of the CLASS 2 TIME DISTRIBUTION. The following discussion is from a document called "IBM SAP Technical Brief/Tuning SAP/DB2/zSeries" from Mark Gordon:

> DB2 accounting data can be used to determine where time is spent processing in DB2. Time is gathered on each DB2 thread, and is broken down into "class 1", "class 2", and "class 3" time.
>
> – Class 3 suspend time is when the thread is suspended, and waiting for a DB2 action (commit, I/O, row lock, etc.).
>
> – Class 2 Elapsed time is when DB2 is processing a request – it contains the Class 3 delay time, as well as CPU time in DB2 (class 2 CPU time).
>
> – Class 1 CPU is time that a thread is processing SQL from SAP – it contains class 2 CPU, plus time processing outside DB2
>
> – Class 1 elapsed time is the time a thread is allocated. This is not a useful performance metric with SAP.
>
> – Not attributed time = *Class 2 elapsed time – Class 2 CPU time – Class 3 suspension time*. It is the time that is left when all the time that DB2 can account for is removed from DB2 elapsed time. Not attributed time happens when DB2 considers a thread to be run-able, but the thread is not running. This might be caused by a CPU overload on the DB server, a paging problem on the DB server, and so forth.

Key indicators in DB2 times

Suspend in DB2 high (*class 3 / Class 2 elapsed*): When this is high (for example, over 75 to 80%), DB2 execution is often blocked while DB2 waits for events such as I/O, locks, and so forth.

Not attributed (or Other) high: When this is high (over 20 to 25%), DB2 execution is blocked due to an operating system issue such as CPU overload, workload prioritization, paging, and so forth.

CPU time high (*Class 2 CPU / Class 2 elapsed*): If this value is high (over 60 to 70%), there may be problems with inefficient SQL, such as table scans on moderate sized memory resident tables. It may be a sign of a well-tuned system (high hit rates, short suspend times), though in general, it is unusual to see a system with Class 2 CPU greater than Class 2 Elapsed.

Length of individual suspensions: Long average duration for I/O suspension, other read I/O, other write I/O, and commit can be indicators of I/O performance problems.

Statistics report (long)

This report is created because the statement STATISTICS REPORT LAYOUT(LONG) INCLUDE(SUBSYSTEMID(DBH1)) is specified for the DB2PM utility. The report lists a wealth of statistical information about the DB2 activity during the KPI test, as shown in the sample report in Figure A-61.

```
   LOCATION: HKMCDBH1                  OMEGAMON XE FOR DB2 PERFORMANCE EXPERT (V3)                         PAGE: 1-1
      GROUP: N/P                                STATISTICS REPORT - LONG                          REQUESTED FROM: NOT SPECIFIED
     MEMBER: N/P                                                                                            TO: NOT SPECIFIED
  SUBSYSTEM: DBH1                                                                                  INTERVAL FROM: 09/26/07 14:50:12.25
DB2 VERSION: V8                                   SCOPE: MEMBER          |                                  TO: 09/26/07 15:00:02.60

---- HIGHLIGHTS ----------------------------------------------------------------------------------------------------------
INTERVAL START   : 09/26/07 14:50:12.25    SAMPLING START: 09/26/07 14:50:12.25    TOTAL THREADS       :     0.00
INTERVAL END     : 09/26/07 15:00:02.60    SAMPLING END  : 09/26/07 15:00:02.60    TOTAL COMMITS       :  5511.00
INTERVAL ELAPSED:             9:50.351665   OUTAGE ELAPSED:              0.000000   DATA SHARING MEMBER:      N/A

SQL DML                       QUANTITY  /SECOND  /THREAD  /COMMIT   SQL DCL                       QUANTITY  /SECOND  /THREAD  /COMMIT
---------------------------   --------  -------  -------  -------   ---------------------------   --------  -------  -------  -------
SELECT                           46.00     0.08      N/C     0.01   LOCK TABLE                        0.00     0.00      N/C     0.00
INSERT                          137.9K   233.54      N/C    25.02   GRANT                             0.00     0.00      N/C     0.00
UPDATE                          132.6K   224.69      N/C    24.07   REVOKE                            0.00     0.00      N/C     0.00
DELETE                        22799.00    38.62      N/C     4.14   SET HOST VARIABLE                 0.00     0.00      N/C     0.00
                                                                    SET CURRENT SQLID                 0.00     0.00      N/C     0.00
PREPARE                        2445.00     4.14      N/C     0.44   SET CURRENT DEGREE                0.00     0.00      N/C     0.00
DESCRIBE                        287.00     0.49      N/C     0.05   SET CURRENT RULES                 0.00     0.00      N/C     0.00
DESCRIBE TABLE                    0.00     0.00      N/C     0.00   SET CURRENT PATH                  0.00     0.00      N/C     0.00
OPEN                            318.4K   539.32      N/C    57.77   SET CURRENT PRECISION             0.00     0.00      N/C     0.00
CLOSE                           254.00     0.43      N/C     0.05
FETCH                           318.9K   540.26      N/C    57.87   CONNECT TYPE 1                    0.00     0.00      N/C     0.00
                                                                    CONNECT TYPE 2                    2.00     0.00      N/C     0.00
TOTAL                           933.7K  1581.57      N/C   169.42   RELEASE                           0.00     0.00      N/C     0.00
                                                                    SET CONNECTION                    0.00     0.00      N/C     0.00

                                                                    ASSOCIATE LOCATORS                0.00     0.00      N/C     0.00
                                                                    ALLOCATE CURSOR                   0.00     0.00      N/C     0.00

                                                                    HOLD LOCATOR                      0.00     0.00      N/C     0.00
                                                                    FREE LOCATOR                      0.00     0.00      N/C     0.00

                                                                    TOTAL                             2.00     0.00      N/C     0.00

STORED PROCEDURES             QUANTITY  /SECOND  /THREAD  /COMMIT   TRIGGERS                      QUANTITY  /SECOND  /THREAD  /COMMIT
---------------------------   --------  -------  -------  -------   ---------------------------   --------  -------  -------  -------
CALL STATEMENT EXECUTED          19.00     0.03      N/C     0.00   STATEMENT TRIGGER ACTIVATED       0.00     0.00      N/C     0.00
PROCEDURE ABENDED                 0.00     0.00      N/C     0.00   ROW TRIGGER ACTIVATED             0.00     0.00      N/C     0.00
.... lines excluded ....
.... lines excluded ....
.... lines excluded ....
 PREPARE REQUEST                   0.00      0.00
LAST AGENT REQUEST                0.00      0.00
TWO PHASE COMMIT REQUEST          0.00      0.00
TWO PHASE BACKOUT REQUEST         0.00      0.00
FORGET RESPONSES                  0.00      0.00
COMMIT RESPONSES                  0.00      0.00
BACKOUT RESPONSES                 0.00      0.00
THREAD INDOUBT-REM.L.COORD.       0.00
COMMITS DONE-REM.LOC.COORD.       0.00
BACKOUTS DONE-REM.L.COORD.        0.00

STATISTICS REPORT COMPLETE
```

Figure A-61   Accounting report (long)

A.5.2.1 RMF Spreadsheet reporter

The RMF Spreadsheet reporter can be considered as the workstation front end to the RMF postprocessor. The tool is Java™-based and contains functionality (macros) to process RMF Postprocessor Reports (shown in Figure A-62) and RMF Postprocessor Overview Records (shown in Figure A-63 on page 273).

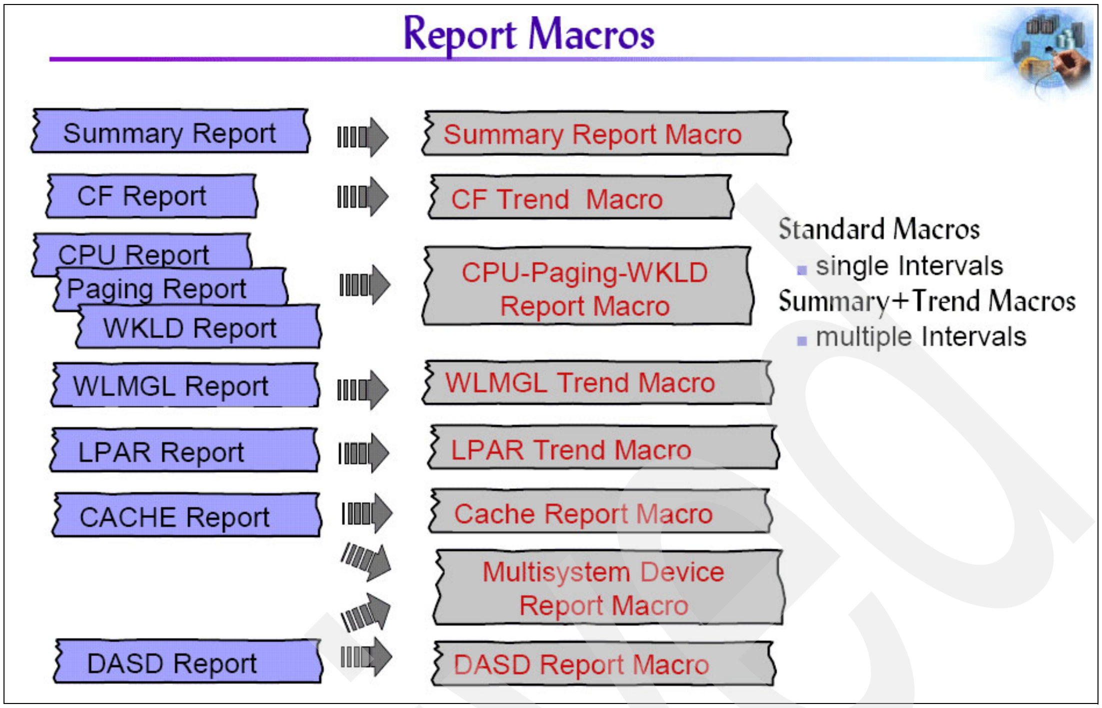

Figure A-62 RMF Spreadsheet report macro's

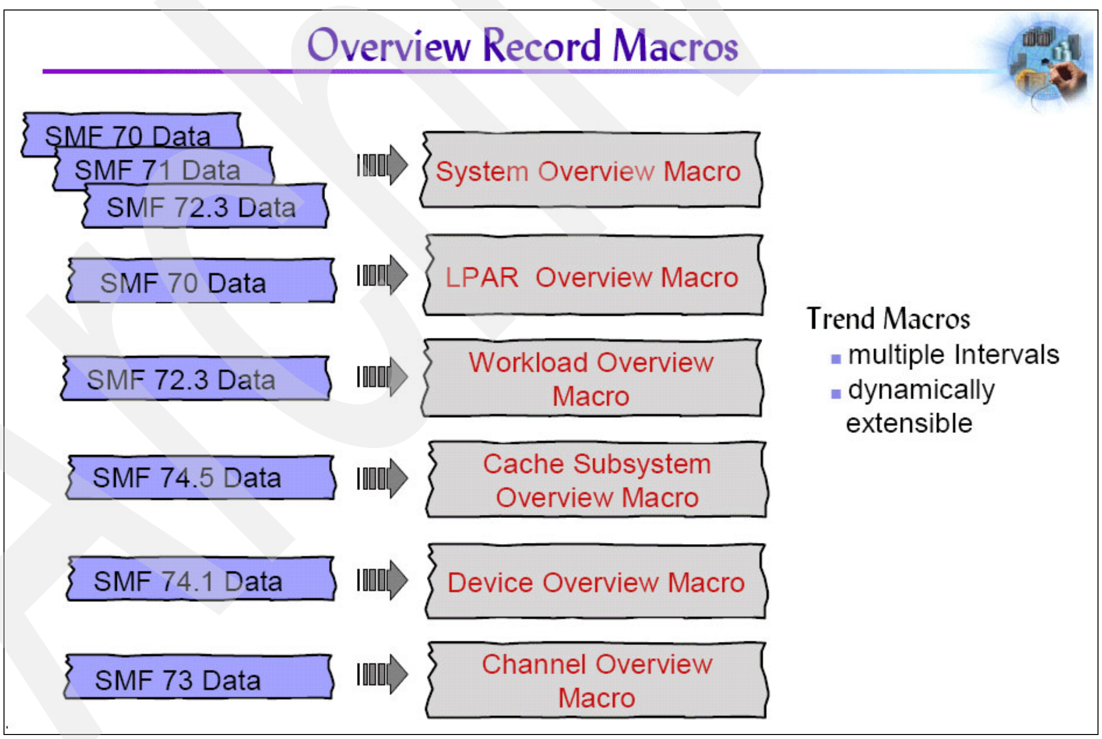

Figure A-63 RMF Spreadsheet overview macro's

These macros convert the RMF Postprocessor Reports and Records into so called "working sets." Those working sets are the input for the Spreadsheets. Some report macros are limited when reporting multiple intervals. This multiple interval (remember that every minute, an SMF interval record is created) reporting is used extensively during the zPOC.

The macros used are:

- Summary Report Macro
- The LPAR Trend Macro
- All overview Macros

Following are examples from the reports.

Figure A-64 displays information from the RMF summary report shown previously in a graphical way. The RMF spreadsheet reporter made this an easy task.

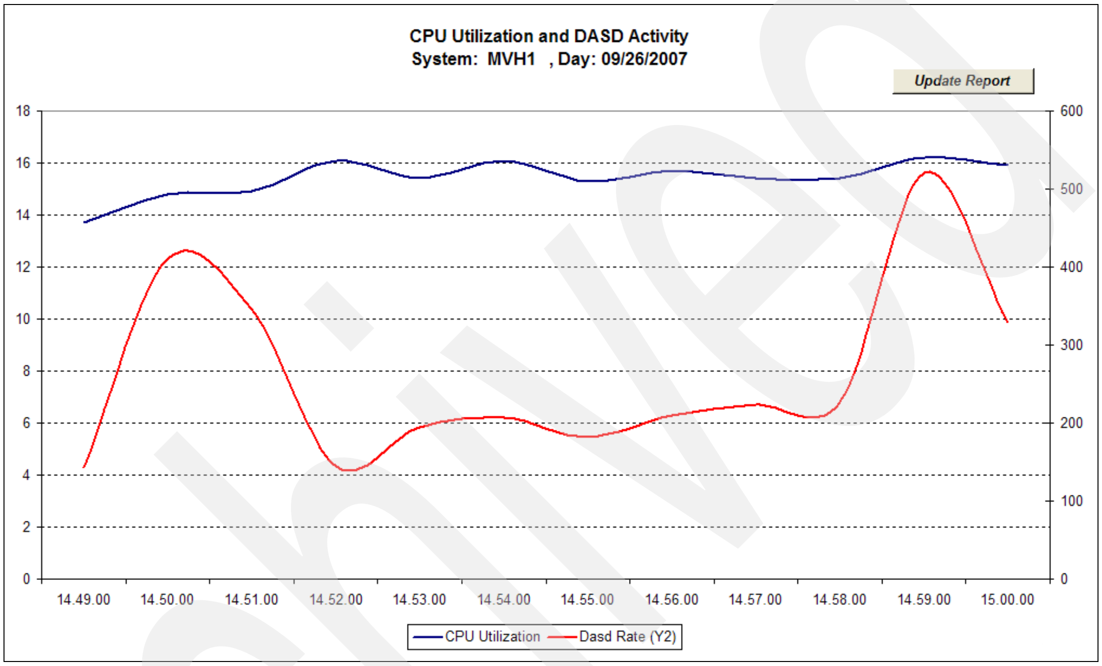

Figure A-64 LPAR Summary example

Figure A-65 shows information from the RMF CPU activity report (partition data) discussed previously, now in a graphical way. In this case the RMF spreadsheet reporter combined information from 12 CPU activity reports into one graph. Doing this without the RMF spreadsheet reporter would be a quite time consuming activity.

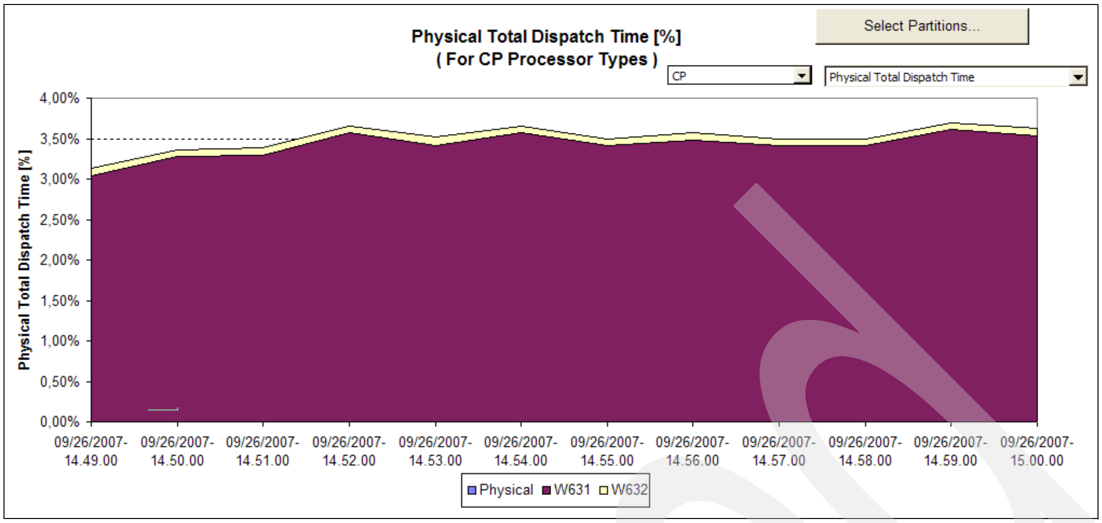

Figure A-65 LPAR overview examples

Figure A-66 is based on an overview record macro. It shows information about the channels. The same information is presented in the RMF Postprocessor - Channel Path activity report. This report is only able to show this information for one interval. The graphs based on overview records are very powerful in presenting this kind of information as a trend during a particular period.

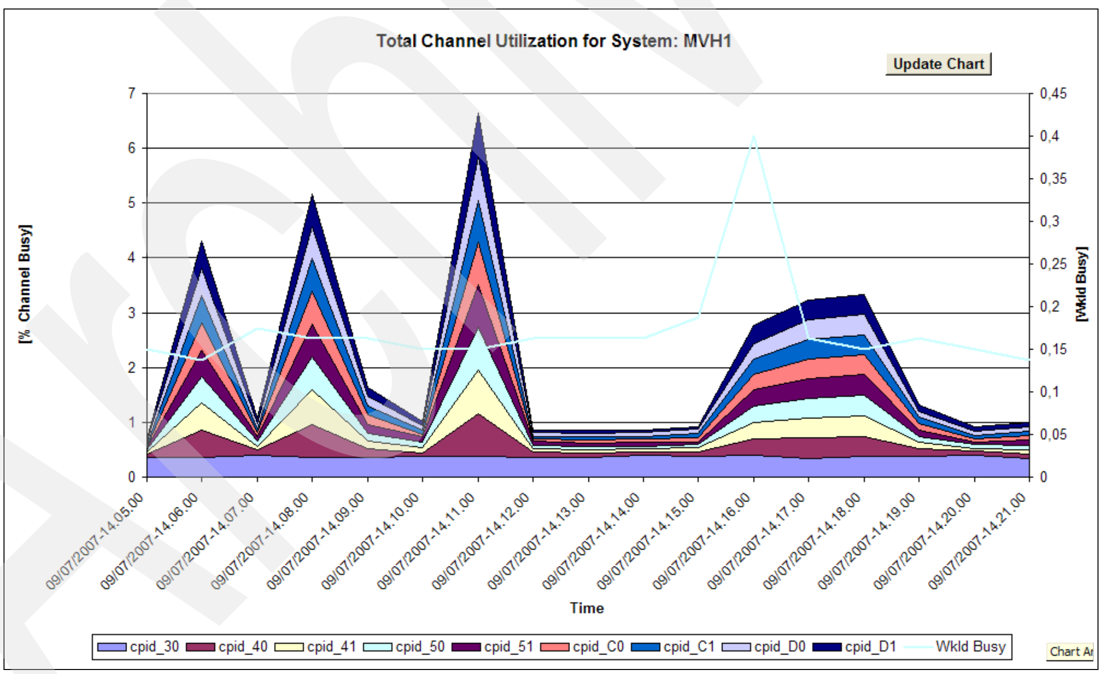

Figure A-66 Channel Utilization example

B

Profiles and configuration files

This appendix provides the profiles and configuration files that were used throughout this benchmark.

The files included are:

B.1 SAP D2L profile for Application Server hkpoceccapp1

```
#.******************************************************************************
#.*                                                                            *
#.*       Instance profile D2L_DVEBMGS00_HKPOCECCAPP1                          *
#.*                                                                            *
#.*       Version                 = 000036                                     *
#.*       Generated by user = FRANCK                                           *
#.*       Generated on = 20.07.1205 , 19:10:29                                 *
#.*                                                                            *
#.******************************************************************************
#parameter created                        by: SAPSUPPORT1  01.10.2007 14:03:13
stat/tabrec = 10
SAPSYSTEMNAME = D2L
SAPSYSTEM = 00
INSTANCE_NAME = DVEBMGS00
DIR_CT_RUN = $(DIR_EXE_ROOT)/run
DIR_EXECUTABLE = $(DIR_INSTANCE)/exe
jstartup/trimming_properties = off
jstartup/protocol = on
jstartup/vm/home = /usr/java14_64
jstartup/max_caches = 500
jstartup/release = 700
jstartup/instance_properties =
$(jstartup/j2ee_properties):$(jstartup/sdm_properties)
j2ee/dbdriver =
/usr/opt/db2_08_01//java//db2jcc.jar:/usr/opt/db2_08_01//java//db2jcc_license_cisu
z.jar
# PHYS_MEMSIZE = 512
exe/saposcol = $(DIR_CT_RUN)/saposcol
exe/icmbnd = $(DIR_CT_RUN)/icmbnd
rdisp/j2ee_start_control = 1
rdisp/j2ee_start = 1
rdisp/j2ee_libpath = $(DIR_EXECUTABLE)
exe/j2ee = $(DIR_EXECUTABLE)/jcontrol$(FT_EXE)
rdisp/j2ee_timeout = 600
rdisp/frfc_fallback = on
icm/HTTP/j2ee_0 = PREFIX=/,HOST=localhost,CONN=0-500,PORT=5$$00
icm/server_port_0 = PROT=HTTP,PORT=80$$
ms/server_port_0 = PROT=HTTP,PORT=81$$
#VMC related OSS notes 854170
#parameter created                        by: IBMKOREA     2007/07/12 07:14:39
vmcj/option/maxJavaHeap = 200M
#VMC related OSS notes 854170
#parameter created                        by: IBMKOREA     2007/07/12 07:13:46
vmcj/option/ps = 512M
# VMC related parameter OSS notes 854170
#parameter created                        by: IBMKOREA     2007/07/12 07:12:32
vmcj/enable = on
#-----------------------------------------------------------------------
# SAP Messaging Service parameters are set in the DEFAULT.PFL
#-----------------------------------------------------------------------
dbms/type = db2
```

```
dbs/db2/schema = SAPD2L
dbs/db2/use_drda = 1
dbs/db2/ssid = DBH1
dbs/db2/hosttcp = mvh1
dbs/db2/user = SAPD2LC
rsdb/max_blocking_factor = 10
rsdb/max_in_blocking_factor = 10
rsdb/min_blocking_factor = 3
rsdb/min_in_blocking_factor = 3
rsdb/prefer_fix_blocking = 0
rsdb/prefer_union_all = 1
rsdb/prefer_in_itab_opt = 1

#-------------------------------------------------------------
# Jcontrol: Migrated Profile Parameter
#      create at Mon Jul  2 23:17:10 2007
#-------------------------------------------------------------
j2ee/instance_id = ID0047712
#-------------------------------------------------------------

### Work Processes ###
rdisp/wp_no_dia = 75
rdisp/wp_no_btc = 10
#old_value: 1
changed: SAP* 27.09.2007 12:00:28
rdisp/wp_no_enq = 2
rdisp/wp_no_vb = 5
rdisp/wp_no_vb2 = 4
rdisp/wp_no_spo = 1

### R3 Buffer ###
enque/table_size = 409600
#old_value: 800000
changed: SAP* 22.11.2007 15:47:24
abap/buffersize = 1600000
#old_value: 500
changed: SAP* 22.11.2007 15:48:30
abap/shared_objects_size_MB = 1000
rsdb/cua/buffersize = 50000
rsdb/esm/buffersize_kb = 10000
rsdb/ntab/entrycount = 50000
rsdb/ntab/ftabsize = 100000
rsdb/ntab/irbdsize = 100000
rsdb/ntab/sntabsize = 100000
#old_value: 50000
changed: SAP* 06.09.2007 15:36:45
rtbb/buffer_length = 60000
rtbb/max_tables = 1000
sap/bufdir_entries = 20000
zcsa/calendar_area = 2000000
zcsa/db_max_buftab = 40000
zcsa/presentation_buffer_area = 50000000
zcsa/table_buffer_area = 160000000
#old_value: 2000000
changed: SAP* 06.09.2007 15:40:39
```

```
ztta/short_area = 4000000
ztta/max_memreq_MB = 2048
#old_value: 16000
changed: SAP* 06.09.2007 15:42:28
ztta/parameter_area = 32000
rdisp/max_arq = 2000
rspo/spool_id/max_number = 990000
rspo/spool_id/loopbreak = 990000

### ATP-Server and Exp/Imp buffer ###
#old_value: 500000
changed: FRANCK 15.11.2007 14:19:49
rsdb/obj/buffersize = 1000000
rsdb/obj/large_object_size = 8192
#old_value: 13000
changed: SAP* 06.09.2007 15:37:32
#old_value: 250000
changed: FRANCK 15.11.2007 14:20:46
rsdb/obj/max_objects = 500000

### Roll-Memory ###
ztta/roll_area = 6500000
ztta/roll_extension = 32000000000
ztta/roll_first = 1

### Roll-Buffer and Extended Memory (shared) ###
em/blocksize_KB = 16384
em/global_area_MB = 3072
em/initial_size_MB = 32768

rdisp/PG_MAXFS = 500000
#old_value: 65536
changed: FRANCK 22.10.2007 17:38:59
rdisp/PG_SHM = 500000
#old_value: 1000
changed: SAP* 06.09.2007 15:39:37
rdisp/PG_LOCAL = 150
rdisp/ROLL_MAXFS = 250000
rdisp/ROLL_SHM = 65536

### Heap Memory (local) ###
abap/heap_area_dia = 2000000000
abap/heap_area_nondia = 2000000000
abap/heaplimit = 20000000
abap/heap_area_total = 4000000000

# AIX: change to alternate Mem Mgmt 01.09.2006 (SAP Note #789477)
ES/TABLE = SHM_SEGS
#old_value: 4096
changed: SAP* 06.09.2007 15:41:20
ES/SHM_SEG_SIZE = 4096
EM/TOTAL_SIZE_MB = 32768
#old_value: 5
changed: SAP* 06.09.2007 15:41:40
ES/SHM_PROC_SEG_COUNT = 8
```

```
#old_value: 4
changed: SAP* 06.09.2007 15:41:50
ES/SHM_MAX_PRIV_SEGS = 7
ES/SHM_USER_COUNT = 6144
ES/SHM_SEG_COUNT = 6144
ES/SHM_SEGS_VERSION = 2
#-------------------------------------------------------------------
#Client server communication
#----------------------------------------------------------------
#old_value: 10000
changed: SAP* 25.09.2007 10:45:07
rdisp/tm_max_no = 9000
#old_value: 3200
changed: SAP* 25.09.2007 10:44:44
rdisp/wp_ca_blk_no = 9051
#old_value: 2000
changed: SAP* 25.09.2007 10:45:35
rdisp/appc_ca_blk_no = 9051
rdisp/max_comm_entries = 2000
rdisp/elem_per_queue = 2000

gw/max_conn = 2000
gw/cpic_timeout = 300
gw/max_overflow_size = 100000000
gw/max_sys = 2000
gw/max_shm_req = 400

rsts/ccc/cachesize = 6000000
sapgui/user_scripting = TRUE
rdisp/wp_auto_restart = 86400
rdisp/noptime = 30000
#----------------------------------------------------------------------
#----------------------------------------------------------------------
#Values proposed by SAP for shared memory pool sizes
#----------------------------------------------------------------------
#parameter created                          by: SAP*         06.09.2007 15:43:06
#old_value: 954000000
changed: SAP* 06.09.2007 16:03:15
#ipc/shm_psize_10 = 1400000000
ipc/shm_psize_10 = 1912000000
#parameter created                          by: SAP*         06.09.2007 15:43:06
ipc/shm_psize_14 = 0
#parameter created                          by: SAP*         06.09.2007 15:43:06
ipc/shm_psize_18 = 0
#parameter created                          by: SAP*         06.09.2007 15:43:06
ipc/shm_psize_19 = 0
#parameter created                          by: SAP*         06.09.2007 15:43:06
ipc/shm_psize_30 = -10
#parameter created                          by: SAP*         06.09.2007 15:43:06
ipc/shm_psize_40 = 448000000
#parameter created                          by: SAP*         06.09.2007 15:43:06
ipc/shm_psize_41 = 0
#parameter created                          by: SAP*         06.09.2007 15:43:06
ipc/shm_psize_51 = -10
#parameter created                          by: SAP*         06.09.2007 15:43:06
```

```
ipc/shm_psize_52 = -10
#parameter created                      by: SAP*         06.09.2007 15:43:06
ipc/shm_psize_54 = -10
#parameter created                      by: SAP*         06.09.2007 15:43:06
ipc/shm_psize_55 = -10
#parameter created                      by: SAP*         06.09.2007 15:43:06
ipc/shm_psize_57 = -10
#parameter created                      by: SAP*         06.09.2007 15:43:06
ipc/shm_psize_58 = -10
```

B.2 SAP D2L profile for application servers hkpoceccapp2

```
#.*************************************************************************
#.*                                                                       *
#.*       Instance profile D2L_D01_HKPOCECCAPP2                           *
#.*                                                                       *
#.*       Version                 = 000023                                *
#.*       Generated by user = FRANCK                                      *
#.*       Generated on = 22.11.2007 , 11:48:31                            *
#.*                                                                       *
#.*************************************************************************
#parameter created                      by: SAPSUPPORT1  01.10.2007 14:00:42
stat/tabrec = 10
SAPSYSTEMNAME = D2L
SAPSYSTEM = 01
INSTANCE_NAME = D01
DIR_CT_RUN = $(DIR_EXE_ROOT)/run
DIR_EXECUTABLE = $(DIR_INSTANCE)/exe
jstartup/trimming_properties = off
jstartup/protocol = on
jstartup/vm/home = /usr/java14_64
jstartup/max_caches = 500
jstartup/release = 700
jstartup/instance_properties = $(jstartup/j2ee_properties)
j2ee/dbdriver =
/usr/opt/db2_08_01//java//db2jcc.jar:/usr/opt/db2_08_01//java//db2jcc_license_cisu
z.jar
# PHYS_MEMSIZE = 512
exe/saposcol = $(DIR_CT_RUN)/saposcol
exe/icmbnd = $(DIR_CT_RUN)/icmbnd
rdisp/j2ee_start_control = 1
rdisp/j2ee_start = 1
rdisp/j2ee_libpath = $(DIR_EXECUTABLE)
exe/j2ee = $(DIR_EXECUTABLE)/jcontrol$(FT_EXE)
rdisp/j2ee_timeout = 600
rdisp/frfc_fallback = on
icm/HTTP/j2ee_0 = PREFIX=/,HOST=localhost,CONN=0-500,PORT=5$$00
icm/server_port_0 = PROT=HTTP,PORT=80$$
#VMC related OSS notes 854170
#parameter created                      by: IBMKOREA     2007/07/12 07:14:39
vmcj/option/maxJavaHeap = 200M
#VMC related OSS notes 854170
```

```
#parameter created                     by: IBMKOREA     2007/07/12 07:13:46
vmcj/option/ps = 512M
# VMC related parameter OSS notes 854170
#parameter created                     by: IBMKOREA     2007/07/12 07:12:32
vmcj/enable = on

dbms/type = db2
dbs/db2/schema = SAPD2L
dbs/db2/use_drda = 1
dbs/db2/ssid = DBH1
dbs/db2/hosttcp = mvh1
dbs/db2/user = SAPD2LC
rsdb/max_blocking_factor = 10
rsdb/max_in_blocking_factor = 10
rsdb/min_blocking_factor = 3
rsdb/min_in_blocking_factor = 3
rsdb/prefer_fix_blocking = 0
rsdb/prefer_union_all = 1
rsdb/prefer_in_itab_opt = 1

#-------------------------------------------------------------
# Jcontrol: Migrated Profile Parameter
#      create at Tue Jul  3 12:07:11 2007
#-------------------------------------------------------------
j2ee/instance_id = ID0138524
#-------------------------------------------------------------

### Work Processes ###
rdisp/wp_no_dia = 75
rdisp/wp_no_btc = 10
rdisp/wp_no_vb = 5
rdisp/wp_no_vb2 = 4
rdisp/wp_no_spo = 1

### R3 Buffer ###
abap/buffersize = 1500000
abap/shared_objects_size_MB = 500
rsdb/cua/buffersize = 50000
rsdb/esm/buffersize_kb = 10000
rsdb/ntab/entrycount = 50000
rsdb/ntab/ftabsize = 100000
rsdb/ntab/irbdsize = 100000
rsdb/ntab/sntabsize = 100000
#old_value: 50000
changed: SAP* 06.09.2007 15:46:20
rtbb/buffer_length = 60000
rtbb/max_tables = 1000
sap/bufdir_entries = 20000
zcsa/calendar_area = 2000000
zcsa/db_max_buftab = 40000
zcsa/presentation_buffer_area = 50000000
zcsa/table_buffer_area = 160000000
#old_value: 2000000
changed: SAP* 06.09.2007 15:47:30
ztta/short_area = 4000000
```

```
ztta/max_memreq_MB = 2048
ztta/parameter_area = 16000
rdisp/max_arq = 2000
rspo/spool_id/max_number = 990000
rspo/spool_id/loopbreak = 990000

### ATP-Server and Exp/Imp buffer ###
#old_value: 500000
changed: FRANCK 15.11.2007 14:37:36
rsdb/obj/buffersize = 1000000
#old_value: 40000
changed by eric 14.11.2007
rsdb/obj/large_object_size = 8192
#old_value: 13000
changed: SAP* 06.09.2007 15:46:45
#old_value: 250000
changed: FRANCK 15.11.2007 14:35:58
rsdb/obj/max_objects = 500000
#old_value: 20000

### Roll-Memory ###
ztta/roll_area = 6500000
ztta/roll_extension = 32000000000
ztta/roll_first = 1

### Roll-Buffer and Extended Memory (shared) ###
em/blocksize_KB = 16384
em/global_area_MB = 3072
em/initial_size_MB = 32768

rdisp/PG_MAXFS = 500000
#old_value: 65536
changed: FRANCK 22.10.2007 17:39:54
rdisp/PG_SHM = 500000
#old_value: 2000
changed: SAP* 06.09.2007 15:47:08
rdisp/PG_LOCAL = 150
rdisp/ROLL_MAXFS = 250000
rdisp/ROLL_SHM = 65536

### Heap Memory (local) ###
abap/heap_area_dia = 2000000000
abap/heap_area_nondia = 2000000000
abap/heaplimit = 20000000
abap/heap_area_total = 4000000000

# AIX: change to alternate Mem Mgmt 01.09.2006 (SAP Note #789477)
ES/TABLE = SHM_SEGS
#old_value: 4096
changed: SAP* 06.09.2007 15:47:54
ES/SHM_SEG_SIZE = 4096
EM/TOTAL_SIZE_MB = 32768
#old_value: 5
changed: SAP* 06.09.2007 15:48:09
ES/SHM_PROC_SEG_COUNT = 8
```

```
#old_value: 4
changed: SAP* 06.09.2007 15:48:23
ES/SHM_MAX_PRIV_SEGS = 7
ES/SHM_USER_COUNT = 6144
ES/SHM_SEG_COUNT = 6144
ES/SHM_SEGS_VERSION = 2
#------------------------------------------------------------------
#Client server communication
#---------------------------------------------------------------
rdisp/tm_max_no = 10000
rdisp/wp_ca_blk_no = 3200
rdisp/appc_ca_blk_no = 2000
rdisp/max_comm_entries = 2000
rdisp/elem_per_queue = 2000

gw/max_conn = 2000
gw/cpic_timeout = 300
gw/max_overflow_size = 100000000
gw/max_sys = 2000
gw/max_shm_req = 400

rsts/ccc/cachesize = 6000000
sapgui/user_scripting = TRUE
rdisp/wp_auto_restart = 86400
rdisp/noptime = 30000
#---------------------------------------------------------------------
#Values proposed by SAP for shared memory pool sizes
#---------------------------------------------------------------------
#parameter created                          by: FRANCK       15.11.2007 14:41:49
ipc/shm_psize_10 = 1520000000
#parameter created                          by: FRANCK       15.11.2007 14:41:49
ipc/shm_psize_14 = 0
#parameter created                          by: FRANCK       15.11.2007 14:41:49
ipc/shm_psize_18 = 0
#parameter created                          by: FRANCK       15.11.2007 14:41:49
ipc/shm_psize_19 = 0
#parameter created                          by: FRANCK       15.11.2007 14:41:49
ipc/shm_psize_30 = -10
#parameter created                          by: FRANCK       15.11.2007 14:41:49
ipc/shm_psize_40 = 448000000
#parameter created                          by: FRANCK       15.11.2007 14:41:49
ipc/shm_psize_41 = 0
#parameter created                          by: FRANCK       15.11.2007 14:41:49
ipc/shm_psize_51 = -10
#parameter created                          by: FRANCK       15.11.2007 14:41:49
ipc/shm_psize_52 = -10
#parameter created                          by: FRANCK       15.11.2007 14:41:49
ipc/shm_psize_54 = -10
#parameter created                          by: FRANCK       15.11.2007 14:41:49
ipc/shm_psize_55 = -10
#parameter created                          by: FRANCK       15.11.2007 14:41:49
ipc/shm_psize_57 = -10
#parameter created                          by: FRANCK       15.11.2007 14:41:49
ipc/shm_psize_58 = -10
```

B.3 SAP D2L profile for application servers hkpoceccapp3

```
#.******************************************************************************
#.*                                                                            *
#.*       Instance profile D2L_D02_HKPOCECCAPP3                                *
#.*                                                                            *
#.*       Version                 = 000026                                     *
#.*       Generated by user = FRANCK                                           *
#.*       Generated on = 22.11.2007 , 11:48:31                                 *
#.*                                                                            *
#.******************************************************************************
SAPSYSTEMNAME = D2L
SAPSYSTEM = 02
INSTANCE_NAME = D02
DIR_CT_RUN = $(DIR_EXE_ROOT)/run
DIR_EXECUTABLE = $(DIR_INSTANCE)/exe
jstartup/trimming_properties = off
jstartup/protocol = on
jstartup/vm/home = /usr/java14_64
jstartup/max_caches = 500
jstartup/release = 700
jstartup/instance_properties = $(jstartup/j2ee_properties)
j2ee/dbdriver =
/usr/opt/db2_08_01//java//db2jcc.jar:/usr/opt/db2_08_01//java//db2jcc_license_cisu
z.jar
# PHYS_MEMSIZE = 512
exe/saposcol = $(DIR_CT_RUN)/saposcol
exe/icmbnd = $(DIR_CT_RUN)/icmbnd
rdisp/j2ee_start_control = 1
rdisp/j2ee_start = 1
rdisp/j2ee_libpath = $(DIR_EXECUTABLE)
exe/j2ee = $(DIR_EXECUTABLE)/jcontrol$(FT_EXE)
rdisp/j2ee_timeout = 600
rdisp/frfc_fallback = on
icm/HTTP/j2ee_0 = PREFIX=/,HOST=localhost,CONN=0-500,PORT=5$$00
icm/server_port_0 = PROT=HTTP,PORT=80$$
#VMC related OSS notes 854170
#parameter created                            by: IBMKOREA     2007/07/12 07:14:39
vmcj/option/maxJavaHeap = 200M
#VMC related OSS notes 854170
#parameter created                            by: IBMKOREA     2007/07/12 07:13:46
vmcj/option/ps = 512M
# VMC related parameter OSS notes 854170
#parameter created                            by: IBMKOREA     2007/07/12 07:12:32
vmcj/enable = on

dbms/type = db2
dbs/db2/schema = SAPD2L
dbs/db2/use_drda = 1
dbs/db2/ssid = DBH1
dbs/db2/hosttcp = mvh1
dbs/db2/user = SAPD2LC
rsdb/max_blocking_factor = 10
```

```
rsdb/max_in_blocking_factor = 10
rsdb/min_blocking_factor = 3
rsdb/min_in_blocking_factor = 3
rsdb/prefer_fix_blocking = 0
rsdb/prefer_union_all = 1
rsdb/prefer_in_itab_opt = 1

#-------------------------------------------------------------
# Jcontrol: Migrated Profile Parameter
#      create at Tue Jul  3 12:07:11 2007
#-------------------------------------------------------------
j2ee/instance_id = ID0138524
#-------------------------------------------------------------

### Work Processes ###
rdisp/wp_no_dia = 75
rdisp/wp_no_btc = 10
rdisp/wp_no_vb = 5
rdisp/wp_no_vb2 = 4
rdisp/wp_no_spo = 1

### R3 Buffer ###
abap/buffersize = 1500000
abap/shared_objects_size_MB = 500
rsdb/cua/buffersize = 50000
rsdb/esm/buffersize_kb = 10000
rsdb/ntab/entrycount = 50000
rsdb/ntab/ftabsize = 100000
rsdb/ntab/irbdsize = 100000
rsdb/ntab/sntabsize = 100000
rtbb/buffer_length = 50000
rtbb/max_tables = 1000
sap/bufdir_entries = 20000
zcsa/calendar_area = 2000000
zcsa/db_max_buftab = 40000
zcsa/presentation_buffer_area = 50000000
zcsa/table_buffer_area = 160000000
ztta/short_area = 2000000
ztta/max_memreq_MB = 2048
#parameter created                          by: SAPSUPPORT1  01.10.2007 14:05:01
stat/tabrec = 10
#old_value: 16000
changed: SAP* 06.09.2007 15:49:26
ztta/parameter_area = 32000
rdisp/max_arq = 2000
rspo/spool_id/max_number = 990000
rspo/spool_id/loopbreak = 990000

### ATP-Server and Exp/Imp buffer ###
#rsdb/obj/buffersize = 500000
rsdb/obj/buffersize = 1000000
rsdb/obj/large_object_size = 8192
#rsdb/obj/max_objects = 250000
rsdb/obj/max_objects = 500000
```

```
### Roll-Memory ###
ztta/roll_area = 6500000
ztta/roll_extension = 32000000000
ztta/roll_first = 1

### Roll-Buffer and Extended Memory (shared) ###
em/blocksize_KB = 16384
em/global_area_MB = 3072
em/initial_size_MB = 32768

rdisp/PG_MAXFS = 500000
#old_value: 65536
changed: FRANCK 22.10.2007 17:41:30
rdisp/PG_SHM = 500000
rdisp/PG_LOCAL = 2000
rdisp/ROLL_MAXFS = 250000
rdisp/ROLL_SHM = 65536

### Heap Memory (local) ###
abap/heap_area_dia = 2000000000
abap/heap_area_nondia = 2000000000
abap/heaplimit = 20000000
abap/heap_area_total = 4000000000

# AIX: change to alternate Mem Mgmt 01.09.2006 (SAP Note #789477)
ES/TABLE = SHM_SEGS
ES/SHM_SEG_SIZE = 4096
EM/TOTAL_SIZE_MB = 32768
ES/SHM_PROC_SEG_COUNT = 5
ES/SHM_MAX_PRIV_SEGS = 4
ES/SHM_USER_COUNT = 6144
ES/SHM_SEG_COUNT = 6144
ES/SHM_SEGS_VERSION = 2
#------------------------------------------------------------------
#Client server communication
#---------------------------------------------------------------
rdisp/tm_max_no = 10000
#old_value: 3200
changed: SAP* 25.09.2007 10:52:02
rdisp/wp_ca_blk_no = 9070
#old_value: 2000
changed: SAP* 25.09.2007 10:52:20
rdisp/appc_ca_blk_no = 9070
rdisp/max_comm_entries = 2000
rdisp/elem_per_queue = 2000

gw/max_conn = 2000
gw/cpic_timeout = 300
gw/max_overflow_size = 100000000
gw/max_sys = 2000
gw/max_shm_req = 400

rsts/ccc/cachesize = 6000000
sapgui/user_scripting = TRUE
rdisp/wp_auto_restart = 86400
```

```
rdisp/noptime = 30000
#-----------------------------------------------------------------------
#-----------------------------------------------------------------------
#Values proposed by SAP for shared memory pool sizes
#-----------------------------------------------------------------------
#parameter created                            by: SAP*         06.09.2007 15:49:38
#old_value: 190000000
changed: SAP* 06.09.2007 16:04:41
ipc/shm_psize_10 = 1510000000
#parameter created                            by: SAP*         06.09.2007 15:49:38
ipc/shm_psize_14 = 0
#parameter created                            by: SAP*         06.09.2007 15:49:38
ipc/shm_psize_18 = 0
#parameter created                            by: SAP*         06.09.2007 15:49:38
ipc/shm_psize_19 = 0
#parameter created                            by: SAP*         06.09.2007 15:49:38
ipc/shm_psize_30 = -10
#parameter created                            by: SAP*         06.09.2007 15:49:38
ipc/shm_psize_40 = 448000000
#parameter created                            by: SAP*         06.09.2007 15:49:38
ipc/shm_psize_41 = 0
#parameter created                            by: SAP*         06.09.2007 15:49:38
ipc/shm_psize_51 = -10
#parameter created                            by: SAP*         06.09.2007 15:49:38
ipc/shm_psize_52 = -10
#parameter created                            by: SAP*         06.09.2007 15:49:38
ipc/shm_psize_54 = -10
#parameter created                            by: SAP*         06.09.2007 15:49:38
ipc/shm_psize_55 = -10
#parameter created                            by: SAP*         06.09.2007 15:49:38
ipc/shm_psize_57 = -10
#parameter created                            by: SAP*         06.09.2007 15:49:38
ipc/shm_psize_58 = -10
```

B.4 DB2 for z/OS Configuration

```
MVS PARMLIB UPDATE PARAMETERS (DSNTIPM)
---------------------------------------
SUBSYSTEM DEFAULT (SSID)..................................DBH1
SUPPRESS SOFT ERRORS (SUPERRS).............................YES

STORAGE SIZES INSTALLATION PARAMETERS (DSNTIPC,DSNTIPE)
-------------------------------------------------------
MAX NO OF USERS CONCURRENTLY RUNNING IN DB2 (CTHREAD).....1,024
MAX NO OF TSO CONNECTIONS (IDFORE)..........................100
MAX NO OF BATCH CONNECTIONS (IDBACK)........................200
MAX NO OF REMOTE CONNECTIONS (CONDBAT)...................30,000
MAX NO OF CONCURRENT REMOTE ACTIVE CONNECTIONS (MAXDBAT)....512
MAXIMUM SIZE OF EDM POOL IN BYTES (EDMPOOL).........155,889,664
MAXIMUM SIZE OF SORT POOL IN BYTES (SRTPOOL).........28,672,000
MAXIMUM SIZE OF RID POOL IN BYTES (MAXRBLK).........102,400,000
3990 CACHE (SEQCACH)........................................SEQ
UTILITY CACHE OPTION (SEQPRES)..............................YES
MAXIMUM KEPT DYNAMIC STATEMENTS (MAXKEEPD)................8,000
CONTRACT THREAD STORAGE (CONTSTOR)..........................YES
EDM POOL DATA SPACE SIZE IN KB (EDMDSPAC)...................N/A
MAX EDM POOL DATA SPACE SIZE IN KB (EDMDSMAX)...............N/A
MANAGE THREAD STORAGE (MINSTOR).............................YES
EDM STATEMENT CACHE SIZE IN KB (EDMSTMTC)...............475,737
EDM DBD CACHE SIZE IN KB (EDMDBDC)......................475,737
LONG-RUNNING READER (LRDRTHLD)...............................10
MAX NO OF DATA SETS CONCURRENTLY IN USE (DSMAX)..........49,800
PAD INDEX BY DEFAULT (PADIX).................................NO

IRLM INSTALLATION PARAMETERS (DSNTIPI)
--------------------------------------
IRLM SUBSYSTEM NAME (IRLMSID).............................DJH1
IRLM RESOURCE TIMEOUT IN SECONDS (IRLMRWT)..................600
IRLM AUTOMATIC START (IRLMAUT)..............................YES
IRLM START PROCEDURE NAME (IRLMPRC)....................DBH1IRLM
SECONDS DB2 WILL WAIT FOR IRLM START (IRLMSWT)..............300
UTILITY TIMEOUT FACTOR (UTIMOUT)..............................3
U LOCK FOR REPEATABLE READ OR READ STABILITY (RRULOCK).......NO
X LOCK FOR SEARCHED UPDATE/DELETE (XLKUPDLT)................YES
IMS/BMP TIMEOUT FACTOR (BMPTOUT)..............................4
IMS/DLI TIMEOUT FACTOR (DLITOUT)..............................6
WAIT FOR RETAINED LOCKS (RETLWAIT)............................1
ENABLE DB CHECKING...........................................NO

IRLM PROCESSING PARAMETERS
--------------------------
PC SPECIFIED................................................N/A
WAIT TIME FOR LOCAL DEADLOCK..............................5,000
NUMBER OF LOCAL CYCLES PER GLOBAL CYCLE.......................1
TIMEOUT INTERVAL............................................600
IRLM MAXIMUM CSA USAGE ALLOWED................................0
Z/OS LOCK TABLE HASH ENTRIES..................................0
```

```
PENDING NUMBER OF HASH ENTRIES................................0
Z/OS LOCK TABLE LIST ENTRIES..................................0

TRACING, CHECKPOINT & PSEUDO-CLOSE PARAMETERS (DSNTIPN)
-------------------------------------------------------
START AUDIT TRACE (AUDITST)..................................NO
START GLOBAL TRACE (TRACSTR).................................NO
TRACE TABLE SIZE IN 4K BYTES (TRACTBL).......................16
START SMF ACCOUNTING (SMFACCT)............................1,2,3
START SMF STATISTICS (SMFSTAT)..............................1,3
STATISTICS TIME INTERVAL IN MINUTES (STATIME).................5
SYNCHRONIZATION INTERVAL WITHIN THE HOUR (SYNCVAL)............0
ONLINE DATASET STATISTICS TIME INTERVAL IN MIN.(DSSTIME)......5
START MONITOR TRACE (MON).....................................1
MONITOR BUFFER SIZE IN BYTES (MONSIZE)..................500,000
UNICODE IFCIDS (UIFCIDS).....................................NO
DDF/RRSAF ACCUM (ACCUMACC)....................................5
AGGREGATION FIELDS (ACCUMUID).................................0

ARCHIVE LOG INSTALLATION PARAMETERS (DSNTIPA)
---------------------------------------------
CATALOG ARCHIVE DATASETS (CATALOG)..........................YES
COPY1 ARCHIVE LOG DEVICE TYPE (UNIT)........................3390
COPY2 ARCHIVE LOG DEVICE TYPE (UNIT2)...................'BLANK'
SPACE ALLOCATION METHOD (ALCUNIT)......................CYLINDER
PRIMARY SPACE ALLOCATION (PRIQTY).........................2,000
SECONDARY SPACE ALLOCATION (SECQTY)...........................100
ARCHIVE LOG BLOCK SIZE IN BYTES (BLKSIZE)................24,576
MAXIMUM READ TAPE UNITS (MAXRTU)..............................2
TAPE UNIT DEALLOCATION PERIOD (DEALLCT).................0000:00
MAX NUMBER OF DATASETS RECORDED IN BSDS (MAXARCH).........1,000
FIRST ARCHIVE COPY MASS STG GROUP NAME...................'NONE'
SECOND ARCHIVE COPY MASS STG GROUP NAME..................'NONE'
DAYS TO RETAIN ARCHIVE LOG DATA SETS (ARCRETN)............9,999
ISSUE WTOR BEFORE MOUNT FOR ARCHIVE VOLUME (ARCWTOR).........NO
COMPACT DATA (COMPACT).......................................NO
QUIESCE PERIOD (QUIESCE)......................................5
SINGLE VOLUME (SVOLARC)......................................NO

DISTRIBUTED DATA FACILITY PANEL 2 (DSNTIP5)
-------------------------------------------
TCP/IP ALREADY VERIFIED (TCPALVER)...........................NO
TCP/IP KEEPALIVE (TCPKPALV).................................120
EXTRA BLOCKS REQ (EXTRAREQ).................................100
EXTRA BLOCKS SRV (EXTRASRV).................................100
POOL THREAD TIMEOUT (POOLINAC)..............................120
DATABASE PROTOCOL (DBPROTCL)...............................DRDA
HOP SITE AUTHORIZATION (HOPAUTH).......................PKGOWNER

PROTECTION INSTALLATION PARAMETERS (DSNTIPP)
--------------------------------------------
DB2 AUTHORIZATION ENABLED (AUTH)............................YES
SYSTEM ADMINISTRATOR 1 AUTHORIZATION ID (SYSADM)........RICHARD
SYSTEM ADMINISTRATOR 2 AUTHORIZATION ID (SYSADM2)........DB2GRP
SYSTEM OPERATOR 1 AUTHORIZATION ID (SYSOPR1)............RICHARD
```

```
SYSTEM OPERATOR 2 AUTHORIZATION ID (SYSOPR2).............DB2GRP
DEFAULT (UNKNOWN) USER AUTHORIZATION ID (DEFLTID).......IBMUSER
RESOURCE LIMIT TABLE CREATOR AUTH ID (RLFAUTH)...........SYSIBM
ARCHIVE LOG RACF PROTECTION (PROTECT)........................NO
PACKAGE AUTHORIZATION CACHE SIZE (CACHEPAC)...........1,048,576
PLAN AUTHORIZATION CACHE SIZE (AUTHCACH)..................3,072
ROUTINE AUTHORIZATION CACHE SIZE (CACHERAC)...........1,048,576
BIND NEW PACKAGE (BINDNV).............................BINDADD
DBA CREATE VIEW (DBACRVW)...................................YES

DATA DEFINITION CONTROL SUPPORT (DSNTIPZ)
-----------------------------------------
INSTALL DD CONTROL (RGFINSTL)................................NO
CONTROL ALL APPLICATIONS (RGFDEDPL)..........................NO
REQUIRE FULL NAMES (RGFFULLQ)...............................YES
UNREGISTERED DDL DEFAULT (RGFDEFLT).....................ACCEPT
REGISTER TABLE OWNER (RGFCOLID)........................DSNRGCOL
DDL REGISTRATION DATABASE NAME (RGFDBNAM)..............DSNRGFDB
APPL REGISTRATION TABLE NAME (RGFNMPRT).......DSN_REGISTER_APPL
OBJECT REGISTRATION TABLE NAME (RGFNMORT).....DSN_REGISTER_OBJT
ESCAPE CHARACTER (RGFESCP)................................X'40'

DEFINE GROUP OR MEMBER (DSNTIPK)
--------------------------------
GROUP NAME (GRPNAME).....................................DSNCAT
MEMBER NAME (MEMBNAME).....................................DSN1
PARALLELISM ASSISTANT (ASSIST)..............................YES
PARALLELISM COORDINATOR (COORDNTR)..........................YES

DISTRIBUTED DATA FACILITY PANEL 1 (DSNTIPR)
-------------------------------------------
DDF STARTUP OPTION (DDF)...................................AUTO
RLST ACCESS ERROR (RLFERRD)............................NOLIMIT
RESYNCHRONIZATION INTERVAL IN MINUTES (RESYNC)................2
DBAT STATUS (CMTSTAT)..................................INACTIVE
IDLE THREAD TIMEOUT INTERVAL (IDTHTOIN).......................0
EXTENDED SECURITY (EXTSEC)..................................YES
MAX TYPE 1 INACTIVE THREADS (MAXTYPE1)........................0

LOCK ESCALATION PARAMETERS (DSNTIPJ)
------------------------------------
MAX PAGE OR ROW LOCKS PER TABLE SPACE (NUMLKTS)...............0
MAX PAGE OR ROW LOCKS PER USER (NUMLKUS)..............2,097,152

LOG INSTALLATION PARAMETERS (DSNTIPL,DSNTIPH)
---------------------------------------------
OUTPUT BUFFER SIZE IN K BYTES (OUTBUFF)...................4,000
DBM1 STORAGE FOR FAST LOG (LOGAPSTG)........................100
CHECKPOINT FREQUENCY (CHKFREQ)...............................10
UR CHECK FREQUENCY (URCHKTH)..................................1
UR LOG RECORD WRITTEN THRESHOLD IN KB (URLGWTH)..............97
LIMIT BACKOUT (LBACKOUT)....................................NO
BACKOUT DURATION (BACKODUR)...................................5
PSEUDO-CLOSE FREQUENCY (PCLOSEN).............................15
PSEUDO-CLOSE TIMER (PCLOSET).................................25
```

```
CHECKPOINTS BETWEEN LEVEL ID UPDATES (DLDFREQ)................5
WRITE THRESHOLD ON NO OF FILLED BUFFER (WRTHRSH)............N/A
NUMBER OF ACTIVE LOG COPIES (TWOACTV).........................1
NUMBER OF ARCHIVE LOG COPIES (TWOARCH)........................1
COPY 1 PREFIX (ARCPFX1)....................DSNARHO.DBH1.ARCLG1
COPY 2 PREFIX (ARCPFX2)....................DSNARHO.DBH1.ARCLG2
TIMESTAMP ARCHIVE LOG DATA SETS (TSTAMP).................... NO

APPLICATION PROGRAMMING DEFAULTS PANEL 1 (DSNTIPF)
   ---------------------------------------------
EBCDIC SBCS CCSID (SCCSID)...................................37
EBCDIC MBCS CCSID (MCCSID)..................................N/P
EBCDIC GBCS CCSID (GCCSID)..................................N/P
ASCII SBCS CCSID (ASCCSID)..................................819
ASCII MBCS CCSID (AMCCSID)..................................N/P
ASCII GBCS CCSID (AGCCSID)..................................N/P
UNICODE SBCS CCSID (USCCSID)................................367
UNICODE MBCS CCSID (UMCCSID)..............................1,208
UNICODE GBCS CCSID (UGCCSID)..............................1,200
DECIMAL POINT OPTION (DECIMAL)...........................PERIOD
DEFAULT CHARSET (CHARSET)..............................ALPHANUM
DEFAULT DELIMITER (DELIM)...............................DEFAULT
DEFAULT ENCODING SCHEME (ENSCHEME).......................EBCDIC
DEFAULT HOST LANGUAGE (DEFLANG).........................IBMCOB
DEFAULT MIXED GRAPHIC (MIXED).................................NO
LOCALE LC_TYPE (LC_CTYPE)...............................'BLANK'
DEFAULT SQL DELIMITER (SQLDELI).........................DEFAULT
DIST SQL STRING DELIMITER (DSQLDELI)......................APOST
APPLICATON ENCODING (APPENSCH)...........................EBCDIC

APPLICATION PROGRAMMING DEFAULTS PANEL 2 (DSNTIP4)
-----------------------------------------------------
MINIMUM DIVIDE SCALE (DECDIV3)................................NO
DECIMAL ARITHMETIC (DECARTH)..................................15
USE FOR DYNAMIC RULES (DYNRULS)..............................YES
STATIC DESCRIBE (DESCSTAT)...................................YES
DATE FORMAT (DATE)...........................................ISO
TIME FORMAT (TIME)...........................................ISO
LOCAL DATE LENGTH (DATELEN)..................................N/A
LOCAL TIME LENGTH (TIMELEN)..................................N/A
STD SQL LANGUAGE (STDSQL)......................................NO
PAD NULL-TERMINATED (PADNTSTR).................................NO

OPERATOR FUNCTIONS INSTALLATION PARAMETERS (DSNTIPO)
---------------------------------------------------
WTO ROUTE CODES (ROUTCDE)......................................1
RESOURCE LIMIT FACILITY AUTOMATIC START (RLF).................NO
RESOURCE LIMIT SPECIFICATION TABLE SUFFIX (RLFTBL)...........01
RESOURCE LIMIT SPEC TABLE ERROR ACTION (RLFERR).........NOLIMIT
AUTO BIND (ABIND)............................................YES
ALLOW EXPLAIN AT AUTOBIND (ABEXP)..............................NO
DPROP SUPPORT (EDPROP).........................................NO
SITE TYPE (SITETYP)....................................LOCALSITE
TRACKER SITE (TRKRSITE)........................................NO
READ COPY2 ARCHIVE (ARC2FRST)..................................NO
```

```
STATISTICS ROLLUP (STATROLL)................................YES
STATISTICS HISTORY (STATHIST)..............................NONE
REAL TIME STATS (STATSINT)...................................30

SIZES PANEL 2 (DSNTIP7)
-----------------------
USER LOB VALUE STORAGE (LOBVALA)......................1,000,000
SYSTEM LOB VALUE STORAGE (LOBVALS).......................50,000
MAXIMUM NUMBER OF LE TOKENS (LEMAX)..........................20
TABLE SPACE ALLOCATION IN KB (TSQTY)..........................0
INDEX SPACE ALLOCATION IN KB (IXQTY)..........................0
VARY DS CONTROL INTERVAL (DSVCI).............................YES
OPTIMIZE EXTENT SIZING (MGEXTSZ).............................YES

PERFORMANCE AND OPTIMIZATION (DSNTIP8)
--------------------------------------
CURRENT DEGREE (CDSSRDEF).....................................1
CACHE DYNAMIC SQL (CACHEDYN)................................YES
OPTIMIZATION HINTS ALLOWED (OPTHINTS)........................NO
VARCHAR FROM INDEX (RETVLCFK)................................NO
RELEASE CURSOR HOLD LOCKS (RELCURHL)........................YES
MAX DEGREE OF PARALLELISM (PARAMDEG)..........................0
UPDATE PART KEY COLS (PARTKEYU).............................YES
EDM BEST FIT (EDMBFIT).......................................NO
IMMEDWRITE OVERRIDE FLAG (IMMEDWRI)..........................NO
EVALUATE UNCOMMITTED (EVALUNC)..............................YES
CURRENT REFRESH AGE (REFSHAGE)................................0
CURRENT MAINT TYPE (MAINTYPE)............................SYSTEM
STAR JOIN ENABLING (STARJOIN) ...............................2
STAR JOIN MAX POOL (SJMXPOOL)...............................128

BUFFER POOL PARAMETERS (DSNTIP1)
--------------------------------
DEFAULT BUFFER POOL FOR USER DATA (TBSBPOOL)................BP2
DEFAULT BUFFER POOL FOR USER INDEXES (IDXBPOOL).............BP3

OTHER SYSTEM PARAMETERS
-----------------------
DUAL BSDS MODE (TWOBSDS)....................................YES
ROLL UP PARALLEL TASK ACCOUNTING (PTASKROL).................YES
NO. PAGES SMALL TABLE THRESHOLD (NPGTHRSH)...................10
SMS DATACLASS NAME FOR TS (SMSDCFL).........................N/P
SMS DATACLASS NAME FOR IS (SMSDCIX).........................N/P
COMPATIBILITY OPTION (COMPAT)...............................N/A
OFFLOAD OPTION (OFFLOAD)....................................YES
SU CONVERSION FACTOR........................................626
OUTER JOIN PERFORMANCE ENHANCEMENTS (OJPERFEH)..............YES
MINIMUM DIVIDE SCALE (MINDVSCL)............................NONE
STAR JOIN THRESHOLD (SJTABLES)................................4
MVS ENVIRONMENT..............................................XA
ONLINE SYSTEM PARM USER ID MONITOR..........................N/P
ONLINE SYSTEM PARM CORREL ID MONITOR........................N/P
ONLINE SYSTEM PARM TIME CHANGED.............................N/P
ONLINE SYSTEM PARM TYPE...................................X'00'
DB2-SUPPLIED DECP INDICATOR...............................X'D5'
```

B.5 DBH1: Buffer pool definition

Name	Usage	VPSIZE	VPSEQT	DWQT	VDWQT	PGFIX
BP0	DB2 Catalog & Directory	10 000	80	30	5	YES
BP1	Sort	50 000	100	80	10	YES
BP2	SAP Tablespaces	120 000	80	30	5	YES
BP3	SAP Indexes	180 000	80	30	5	YES
BP4	VB tables	10 000	10	70	50	YES
BP10	RESB & PLAF Tablespaces	100 000	80	30	10	YES
BP11	RESB & PLAF Indexes	100 000	80	30	10	YES
BP12	RESB & PLAF Indexes	100 000	80	30	10	YES
BP40	SAP LOB Tablespaces	8 000	80	30	5	YES
BP32K	DB2 Catalog & Directory	10 000	80	30	5	YES
BP8K0	DB2 Catalog & Directory	10 000	40	50	30	YES
BP16K0	DB2 Catalog & Directory	5 000	80	30	5	YES

Figure B-1 DBH1 - Buffer pool definition

B.6 DBK1: Buffer pool definition

Name	Usage	VPSIZE	VPSEQT	DWQT	VDWQT	PGFIX
BP0	DB2 Catalog & Directory	10 000	80	30	5	YES
BP1	Sort	10 000	100	80	10	YES
BP2	SAP Tablespaces	60 000	80	30	5	YES
BP3	SAP Indexes	90 000	80	30	5	YES
BP4	VB tables	10 000	10	70	50	YES
BP40	SAP LOB Tablespaces	15 000	80	30	5	YES
BP32K	DB2 Catalog & Directory	15 000	80	30	5	YES
BP8K0	DB2 Catalog & Directory	16 000	40	50	30	YES
BP16K0	DB2 Catalog & Directory	16 000	80	30	5	YES

Figure B-2 DBK1 - Buffer pool definition

B.7 z/OS change log

```
09/07  10h00 Update DSNTIJUZ for DBH1 CTHREAD=1024 (500) MAXDBAT=512
       10h15 FlashCopy Level 1 - 80 volumes in 20 mn
             Creation userid IBMKOREA in D2L(500) and D3B(300)

12/07  17h30 FlashCopy Level 1 - 80 volumes in 20 mn
22/08  17h00 FlashCopy Level 1 Volumes DB2 MVH1 et MVH2
23/08  17h00 FlashCopy Level 1 Volumes systeme + DB2 MVH1 et MVH2
       18h15 Dump sur 3592 de 192 volumes de FlashCopy sur VOLSER=CLNT01
             Dump termine a 21h35,
             Sysout dans ERIC.CLNT.DUMPVOL.D230807.OUTPUT
24/08  10h00 Mise en place FlashCopy Level 2 pour les volumes DB
       14h30 Copy du backup du 23/08 depuis le Level 1 sur le Level2
       15h45 FlashCopy Level 1 Volumes DB2 MVH1 et MVH2
       16h00 Apply maintenance DB2 V8 PUT0705 + Hiper
28/08  22h30 FlashCopy Level 2 Volumes DB2 MVH1 et MVH2
29/08  10h00 Dump sur 3592 de 160 vol de FlashCopy Lvl 2 sur VOL=CLNT02
             Dump termine a 12h43,
             Sysout dans ERIC.CLNT.DUMPVOL.D280807.OUTPUT
30/08  18h00 FlashBack from Level 2 MVH1, MVH2 et Livecache
06/09  07h30 Online Reorg TS + IX - DBH1 & DBK1
       11h00 DBH1: Reorg RESB + changes
       11h00 Runstats - DBH1 & DBK1
       16h00 EarlyWatch changes implementation on DBH1 & DBK1
       16h30 EarlyWatch changes implementation on AS parameters
       17h20 FlashCopy Level 1 Volumes DB2 + SYS MVH1 et MVH2
       17h45 MODIF: Increase number of DIA wp on hkpoclive2
             D3B_DVEMGS00_hkpoclive2 - rdisp/wp_no_dia=70  (60)
             MODIF: SMLG add hkpoclive2 in logon group zPOC
10/09  10h00 D3B: installation ST-A/PI (SAPKITAA5B)
       10h30 D2L: installation ST-A/PI (SAPKITAR7B)
       18h00 FlashBack from Level 1 MVH1, MVH2 et Livecache
       18h15 ADD volumes DBK1W3, DBK1W4, DBK1W5 on MVH2
       18h25 ADD 6 workfiles on DBK1 DSNDB07
11/09  10h00 Add WLM definitions for DDF Workload on MVH1 & MVH2
             SAPHIGH - Vel = 55 - Imp = 2
             SAPMED  - Vel = 45 - Imp = 3
             SAPLOW  - Vel = 25 - Imp = 4
       10h00 MODIF: Livecache MAXLOCKS = 4000000 (1000000 before)
                              SHAREDSQL= NO (YES before)
                              CACHE_SIZE = 7500000 (3500000 before)
       11h00 PROBLEM: Note 977359 - DB2-z/OS: Slow Array Deletes with
                       Download DB2 Connect V8 FP12 Special Build 16563
       14h30 MODIF: Installation FP12_Build_16563 on hkpoclive1 & 2
                    hkpoceccapp1 , 2 & 3
                 Procedure :
                    gunzip -dv aix_64_fp12_specialbuild16563.tar.tgz
                    tar -xvf aix_64_fp12_specialbuild16563.tar.tar
                    cd special_16563
                    ./installFixPak -y
                    /usr/opt/db2_08_01/instance/db2iupdt -e
       18h50 FlashCopy Level 2 Volumes DB2 MVH1 et MVH2
```

```
       19h30 Dump sur 3592 de 160 vol de FlashCopy Lvl 1 sur VOL=CLNT03
12/09  10h00 Start of MRP runs
             MRP RUN00001
       16h40 FlashBack from Level 2 MVH1, MVH2 & Livecache
       17h00 Data changes on tables MARA & DBVM with CLNT queries
             ERIC.CLNT.CNTL($09QUER1)
             ERIC.CLNT.CNTL($09QUER2)
              PB: Duplicate Keys on DBVM indexes
               CLNT: Drop primary key + Drop unique index
                     re-create index as non unique + Rebuild
             ERIC.CLNT.CNTL($09QUER3)
       19h00 MRP RUN00002
13/09  11h00 FlashBack from Level 2 MVH1, MVH2 & Livecache
17/09  19h20 FlashCopy Level 2 Volumes DB2 MVH1 et MVH2
18/09  09h30 FlashBack from Level 2 MVH1, MVH2 & Livecache
       10h00 Partitioning tables RESB & PLAF to 25 partitions on MATNR
       15h30 FlashBack from Level 2 MVH1, MVH2 & Livecache
       16h00 Partitioning tables RESB & PLAF to 25 partitions on
             MANDT , MATNR
             RESB table created MEMBER CLUSTER PCTFREE=FP=0
             LOAD tables RESB & PLAF with previous datas to build the
             compression dictionary + Runstats
             Dummy load on RESB table  with KEEPDICTIONARY REUSE
             Load initial on PLAF table  with KEEPDICTIONARY REUSE
19/09  07h30 FlashBack from Level 2 MVH1, MVH2 & Livecache
       08h00 Partitioning tables RESB & PLAF to 25 partitions on
             MANDT , MATNR
             RESB table created MEMBER CLUSTER PCTFREE=FP=0
             Index RESB~M changed to Cluster and partitioned
             LOAD tables RESB & PLAF with previous data  to build the
             compression dictionary + Runstats
             Dummy load on RESB table  with KEEPDICTIONARY REUSE
             Initial load on PLAF table  with KEEPDICTIONARY REUSE
       09h20 FlashCopy Level 2 Volumes DB2 MVH1
             RUNs MRP + BF data preparation
       17h30 FlashBack from Level 2 MVH1, MVH2 & Livecache
             Installation CTS EDHK900694 in D2L
20/09  17h30 FlashBack from Level 2 MVH1, MVH2 & Livecache
             Installation CTS EDHK900766 in D2L
             Installation CTS ADHK900098 in D3B
21/09  09h30 FlashBack from Level 2 MVH1, MVH2 & Livecache
       09h50 Alter TS RESBX, PLAFX, DBVM & KBVM with TRACKMOD NO
       15h30 FlashBack from Level 2 MVH1, MVH2 & Livecache
       17h15 FlashBack from Level 2 MVH1, MVH2 & Livecache
       17h30 z01 index in DB2 on MANDT, ACCID on table SAPAPO/PEGKEY
             Installation CTS EDHK900694 in D2L
             Installation CTS EDHK900766 in D2L
             Installation CTS ADHK900098 in D3B
             Start BackFlush data generation
24/09  17h15 FlashCopy Level 1 MVH1, MVH2 & Livecache
             Installation CTS KDOK6000668 in D3B & D2L
             Installation Note 1051445    in D3B & D2L
01/10  08h00 FlashBack from Level 2 MVH1, MVH2 & Livecache
       08h00 to 15h10 Lots of modif done by eric...
       17h10 FlashCopy Level 1 MVH1, MVH2 & Livecache
```

```
02/10  10h00 Applied patch 904 on live cache
       16h50 Turned HYPERPAV ON on MVH1 and MVH2 (not really needed,
             but should not hurt)
03/10  09h15 Installation CTS EDHK900790 on D2L
04/10  17h00 FlashBack Level 1 + activate archiving on DBH1 & DBK1
             Installation CTS EDHK900790 on D2L
08/10  10h55 FlashCopy Level 2
       13h00 Added 32 disks to HK1 pool
10/10  07h00 Installation CTS EDHK900798 on D2L
       11h45 FlashCopy Level 2 (with 102 disk for HK1 pool)
       20h20 added 16 volumes to HK1 pool
11/10  11h00 reorg A140X998.ACCTIT due to 64GBytes
             reorg A140X997.COEP
             reorg A140X992.MSEG
             reorg A140X998.BALDAT
             reorg A140X998.ACCTCR
             reorg A140XAAX.PPCXCONF
14/10  11h00 FlashCopy Level 2 (with 118 disk for HK1 pool)
15/10  17h00 reorg 21 TS on ECC
       19h40 FlashCopy Level 2 (with 118 disk for HK1 pool)
       19h45 change hardware add zIIP and Storage
16/10  11h15 1182016 rows in CKMLMV011 after the run
             591008 rows before the run
       11h25 FlashBack Level 2 (with 118 disk for HK1 pool)
       11h40 Deactivate CKMLMV011~004 index on CKMLMV011 table
             runstats on CKMLMV011 table
18/10  08h30 FlashBack Level 2 (with 118 disk for HK1 pool)
       09h10 Runstats on CKMLMV011 table with COLGROUP on ~0 indexes
             job in ERIC.CLNT.CNTL($40RUNST)
19/10  12h30 FlashBack Level 2 (with 118 disk for HK1 pool) only MVH1
22/10  19h00 FlashBack Level 2 (with 118 disk for HK1 pool) only MVH1
             Update DB2 Catalog statistics for CKMLMV011
             job in ERIC.CLNT.CNTL($60UPDCT)
             DBH1 set sysparm load(DBH1PAR1)   changed archive PQTY
23/10  10h35 FlashBack Level 2 (with 118 disk for HK1 pool) only MVH1
       10h40 Update DB2 Catalog statistics for CKMLMV011
             job in ERIC.CLNT.CNTL($60UPDCT)
       14h00 Copy LCU C600-CA3F to 1600-1A3F
             Copy LCU C010-C01F to 1010-101F
             IPL MVH2 on new DS8K Load adr = 1010
             New Flascopy Level 1 volumes for MVH2 are Yxxxxx
             New Flascopy Level 2 volumes for MVH2 are Zxxxxx
             All Flashcopy job on MVH2 are changed.
       20h40 FlashBack Level 2 (with 118 disk for HK1 pool) only MVH1
             Update DB2 Catalog statistics for CKMLMV011
             job in ERIC.CLNT.CNTL($60UPDCT)
             DBH1 set sysparm load(DBH1PAR1)   changed archive PQTY
24/10  23h45 FlashBack Level 2 (with 118 disk for HK1 pool) only MVH1
             Update DB2 Catalog statistics for CKMLMV011
             job in ERIC.CLNT.CNTL($60UPDCT)
             DBH1 set sysparm load(DBH1PAR1)   changed archive PQTY
25/10  12h00 FlashBack Level 2 (with 118 disk for HK1 pool) only MVH1
             Update DB2 Catalog statistics for CKMLMV011
             job in ERIC.CLNT.CNTL($60UPDCT)
             DBH1 set sysparm load(DBH1PAR1)   changed archive PQTY
```

```
27/10  14h00 Swap DS8K
             Copy LCU C000-C00F to 1000-100F
             Copy LCU C100-C11F to 1100-111F
             Copy LCU C200-C21F to 1200-121F
             Copy LCU C300-C31F to 1300-131F
             Copy LCU C400-C41F to 1400-141F
             Copy LCU C500-C51F to 1500-151F
             Copy LCU CA00-CA1F to 1A00-1A1F
             COPY LCU CB00-CB1F TO 1B00-1B1F
             COPY LCU CC00-CC1F TO 1C00-1C1F
             COPY LCU CD00-CD1F TO 1D00-1D1F
             COPY LCU CE10-CE1F TO 1E10-1E1F
             COPY LCU CF10-CF1F TO 1F10-1F1F
             IPL MVH1 on new DS8K Load adr = 1000
             New Flascopy Level 1 volumes for MVH1 are Yxxxxx
             New Flascopy Level 2 volumes for MVH1 are Zxxxxx
             All Flashcopy job on MVH1 are changed.
       17h00 FlashCopy Level 1 (with 118 disk for HK1 pool)
30/10  10h00 Add 16 volumes to HK1 pool  HK1118-HK1133
             All Flashcopy jobs updated on MVH1
             RESB repartitioning - 47 partitions
             RESB recreated with DSSIZE=32Gb
             partition dictionnary build part 27-32
       12h00 FlashCopy Level 2 (with 134 disk for HK1 pool)
       16h00 D2L Installation new SAP license
             License parameters saved in /home/d2ladm/D2L_License.txt
             on hkpoceccapp1
31/10  08h00 Import of Client 510 successfully finished
       08h30 SCC7 (Post Import processing) successfully finished
       09h00 FlashCopy Level 2 (with 134 disk for DBH1 pool)
       09h05 D3B Installation new SAP license
             License parameters saved in /home/d3badm/D3B_License.txt
             on hkpoclive2
       09h05 FlashCopy Level 1 (with 80 disk for DBK1 pool)
       10h00 Reorganization of PLAF table to build compression on new
             partitions. The compression ratio is 83%.
             Runstats on PLAF table to collect new statistics.
             Reorganization of RESB table to build compression on new
             partitions. The compression ratio is 80%.
             Runstats on RESB table to collect new statistics.
             Runstats on DB2 Objects needing new statistics
             Update Data for transaction DB02
       14h30 FlashCopy Level 1 (with 134 disk for DBH1 pool)
       >>>>> FlashCopy Level 1 is the starting point  of scenario II
             for both ECC & APO
05/11  12h45 FlashCopy Level 2  !!!!!!ECC ONLY!!!!!!
07/11  19h56 FlashBack Level 2  !!!!!!ECC ONLY!!!!!!
             added 16 volumes to CLNT1 pool
             reorg RESB table to build compression dictionaries which
             were 'lost' during DDL changes
08/11  00h10 FlashCopy Level 2  !!!!!!ECC ONLY!!!!!! (150 disk for DBH1)
12/11  09h30 FlashCopy Level 1  ECC+APO
       10h30 FlashBack Level 2  !!!!!!ECC ONLY!!!!!! (150 disk for DBH1)
13/11  07h30 Reorg ECC tables DBVM , ZPPT_REM_VEHICLE , ZPPT_REM_RP
       07h40 Runstats on objects needing new statistics
```

```
      08h30 FlashCopy Level 2  ECC+APO+Livecache (disk 5xxxx)
14/11  20h00 Profile Instance parameter change
             hkpoceccapp1:
             /rsdb/obj/buffersize=500000  before=40000  (KB)
             /rsdb/obj/max_objects=250000 before=20000
             /ipc/shm_psize_10=1400000000 before=954000000
             hkpoceccapp2:
             /rsdb/obj/buffersize=500000  before=40000  (KB)
             /rsdb/obj/max_objects=250000 before=20000
             /ipc/shm_psize_10=1008000000 before=534000000
             hkpoceccapp3:
             /rsdb/obj/buffersize=500000  before=40000  (KB)
             /rsdb/obj/max_objects=250000 before=20000
             /ipc/shm_psize_10=998000000  before=190000000
15/11  10h00 MVH1: CPU configuration change
             10 CP   before=8
              5 zIIP before=0
             hkpoceccapp1-3: CPU configuration change
             15 CP   before=10
       15h50 ALTER TS RESBX TRACKMOD NO
       16h00 Lock user FPMD01 on client 500
16/11  08h00 Unload RESB table by partitions (duration=15mn)
             Rows UNLOADED=115500619
16/11  15h20 FlashBack Level 2  ECC+APO+Livecache!!!
       15h40 Re-creation of RESB table with DSSIZE=4GB (32GB before)
             alter RESB table with TRACKMOD NO
             LOAD RESB table to build dictionaries (duration=02h12)
                   + Collect Statistics
             LOAD Dummy RESB table Part27-47       (duration=00h08)
       16h00 APO CTS transport installation
             ADHK900104
             ADHK900071
             ADHK900124
       22h40 Re-creation of RESB table with DSSIZE=4GB (32GB before)
             Change table partitioning to 153 (47 before)
             alter RESB table with TRACKMOD NO
             LOAD RESB table to build dictionaries (duration=02h12)
                   + Collect Statistics
             LOAD Dummy RESB table Part27-47       (duration=00h08)
19/11  11h00 FlashBack Level 2  ECC
       11h20 ECC CTS transport installation
             EDHK900906
20/11  15h00 REORG + Statistics on RESB part 27:47
       17h10 FlashCopy Level 2  !!!!!!ECC ONLY!!!!!! (150 disk for DBH1)
22/11  11h20 FlashBack Level 1  !!!!!!ECC ONLY!!!!!! (150 disk for DBH1)
       11h30 p595 LPAR changes
             hkpoceccapp1-3 20 CPU  (15 before)
             Application server instance profile changes
             DIA = 75    - BTC = 10
       11h30 ECC CTS transport installation
             EDHK900906 , EDHK900952
       12h30 REORG ECC Tablespaces :
             A140X998.ACCTIT , A140X99M.S033 , PR30X999.VBDATA
             A140X997.COSP   , A140X997.COSS , A130X998.CKIT
             A120X996.MKPF   , A140X992.MSEG , A100X998.BLPK
```

```
             A000XAA4.BLPP
       16h00 Application server instance profile changes
             abap/buffersize=1600000   (800000 before)   hkpoceccapp1
             abap/buffersize=1600000   (1500000 before)  hkpoceccapp2
             abap/buffersize=1600000   (1500000 before)  hkpoceccapp3
             abap/shared_object_size_MB=1000 (500 before) app1-3
26/11  19h49 SMS SG SGCLNT1: disable D000 volumes
             ICKDSF additional volumes D000/D020/D030
             FC D000 --> FC1 D020/FC2 D030
             SMS SG SGCLNT1: enable D000 volumes
             Update FC1,FC2,FB1,FB2 JCL in
             SUPPORT.MVH1.FLASHCPY.CLNT.CNTL
28/11  16h10 SMS SG SGCLNT1: disable D100 volumes
             ICKDSF additional volumes D100/D120/D130
             FC D100 --> FC1 D120/FC2 D130
             SMS SG SGCLNT1: enable D100 volumes
             Update FC1,FC2,FB1,FB2 JCL in
             SUPPORT.MVH1.FLASHCPY.CLNT.CNTL
28/11  17h00 Runstats TABLESPACE A110XAAA.AFFW due to bad response time
29/11  09h36 ABEND
        TITLE=DBH1,ABND=0C4-00000004,U=SAPD2LC,M=(N),C=810.LOCN=1.1.75.1
        LOC=DSNXGRDS.DSNXOMD +12B4
        see APAR PK57217 (target date 02/26/2008)
        Diagnostics:
        -SYSM.DUMP.MVH1.D071129.T093608.S00004
        -ROOT.SDSF (po ds with SPOOL ds DBM1,MSTR,IRLM,DIST,SYSLOG)
        This abend followed a situation in which many SAP transactions
        timed out on the following resource:
        NAME A140X997.COSS    .X'0278BC' '.X'09'
29/11  10h31 DBH1 stopped after new transactions abended in DB2 with
             DB2 reason code 00E20016
29/11  10h31 DBH1 started
29/11  13:00 Runstats / Reorg TS/IX via SAP DBA Cockpit Calendar
29/11  19:00 Change calendar date from 29/11 to 1/12
==================== Date change from 29/11 to 1 /12 ===================
01/12  20:07 Runstats A140X997.CKMLMV00
01/12  20:24 Reorg    CKMLMV00 via DBA Cockpit
01/12  20:24 Reorg    CKMLMV01 via DBA Cockpit
02/12  09h10 Runstats on CKMLMV011 table with COLGROUP on ~0 indexes
             job in ERIC.CLNT.CNTL($40RUNST)
             Update DB2 Catalog statistics for CKMLMV011
             job in ERIC.CLNT.CNTL($60UPDCT)
       12.15 rdisp/PG_MAXFS => (500000) 750000 run ok
             rdisp/PG_SHM   => (500000) 750000 run ok
       13.30 rdisp/PG_MAXFS => (750000) 875000 for new sap start
             rdisp/PG_SHM   => (750000) 875000 for new sap start
06/12  23h00 SMS SG SGDB2SAV: add D010-D01F,D110-D11F volumes
             added 32 additional volumes to provide sufficient
             workspace for DB2 Reorg activities
             SMS SG SGCLNT1: enable D000 volumes
             Update FC1,FC2,FB1,FB2 JCL in
             SUPPORT.MVH1.FLASHCPY.CLNT.CNTL
07/12  08h26 SMS SG SGCLNT1: disable D200 volumes
             ICKDSF additional volumes D200/D220/D230
             FC D200 --> FC1 D220/FC2 D230
```

```
      SMS SG SGCLNT1: enable D200 volumes
      Update FC1,FC2,FB1,FB2 JCL in
      SUPPORT.MVH1.FLASHCPY.CLNT.CNTL
11H35 Reorg EBAN
11H37 Reorg MDTB
11H40 Reorg MDKP
11H57 Reorg KBED
11H59 Reorg MARA
12h05 Reorg MARC
12h13 Reorg MAKT
12h18 Runstats tables:
      SAPD2L.CABN
      SAPD2L.CABNT
      SAPD2L.CAWN
      SAPD2L.CAWNT
      SAPD2L.CABN
      SAPD2L.AUSP
      SAPD2L.INOB
      SAPD2L.KLAH
      SAPD2L.KSML
      SAPD2L.KSSK
      SAPD2L.KSML
      SAPD2L.CUCO
      SAPD2L.CUEX
      SAPD2L.CUKB
      SAPD2L.CUKN
      SAPD2L.CUOB
      SAPD2L.CUXREF
      SAPD2L.IBIB
      SAPD2L.IBIN
      SAPD2L.IBST
      SAPD2L.IBSYMBOL
12h42 Reorg PLAF
13h33 Runstats tables:
       SAPD2L.CUVTAB_ADM
       SAPD2L.CUVTAB_FLD
       SAPD2L.CUVTAB_GR
       SAPD2L.CUVTAB_GRT
       SAPD2L.CUVTAB_IND
       SAPD2L.CUVTAB_ST
       SAPD2L.CUVTAB_STT
       SAPD2L.CUVTAB_TX
       SAPD2L.CUVTAB_VALC
       SAPD2L.CUVTAB_VALN
       SAPD2L.CECUFM
       SAPD2L.CECUSD
       SAPD2L.CECUSDT
       SAPD2L.CECUSF
       SAPD2L.CECUSFT
16h50 Reorg RESB, partitions 27:47
17h05 Update RESB related Catalog stats of partitions 27:47
      ERIC.DB2.CNTL(UPDTRES)
17h13 Reorg PLAF
18h15 FlashCopy Level 2 Volumes DB2 MVH1
23h02 Runstats PLAF
```

```
08/12
       01h00 Runstats DBVM,RESBX,PLAF
       01H40 UNLOAD Realtime statistics:
             DB2SAV.DSNDB04.INDE1EOU.D2007342.T004151.U/P
             DB2SAV.DSNDB04.INDEXSPA.D2007342.T004123.U/P
             DB2SAV.DSNRTSDB.DSNRTSTS.D2007342.T004033.U/P
             DB2SAV.DSNDB04.TABL1VXU.D2007342.T003911.U/P
             DB2SAV.DSNDB04.TABLESPA.D2007342.T003710.U/P
             DB2SAV.DSNDB04.TABLESPA.D2007342.T003656.U/P
             DB2SAV.DSNRTSDB.DSNRTSTS.D2007342.T003942.U
       02h00 BSDS Print, job BSDSPR4A
       02h02 FlashBack from Level 2 MVH1
       10H27 Unload Realtime Statistics
             DB2SAV.DSNRTSDB.DSNRTSTS.D2007342.T092817.U
             DB2SAV.DSNRTSDB.DSNRTSTS.D2007342.T092743.P
       11h10 FlashBack from Level 2 MVH1
```

B.8 Report ZZDISTINIT source code

```
Report zzdisstinit.
* This report initializes the low-level code in table MARA
data: ini_disst like mara-disst.

Update mara
       Set disst = ini_disst.
```

Abbreviations and acronyms

ABAP Programming Language of the SAP System

ABAP Programming Language of the SAP System

APO Advanced Planning and Optimization; part of SAP SCM; used for the planning and optimization of supply chain processes at a strategic, tactical, and operational planning level.

APO Advanced Planning and Optimization; part of SAP SCM; used for the planning and optimization of supply chain processes at a strategic, tactical, and operational planning level.

BOM Bill of Material

BOM Bill of Material

CIF Core Interface; SAP standard interface to connect ERP and SCM

CIF Core Interface; SAP standard interface to connect ERP and SCM

DBVM SAP table; Planning file entry, MRP Area

DI Discrete Industries; extension to the ERP system

DI Discrete Industries; includes Automotive

DI Discrete Industries; extension to the ERP system

ECC Enterprise Core Component; SAP ERP system

ECM Engineering Change Management

ECM Engineering Change Management

ERP Enterprise Resource Planning

ERP Enterprise Resource Planning

FSC Full specified code

FSC Full specified code

GR Goods receipt

GR Goods receipt

Info record Element in SAP ERP to keep detail information about relation between customer and vendor

iPPE Integrated Product and Process Engineering

iPPE Integrated Product and Process Engineering

LC Live Cache; Powerful in-memory database backed by disk to manage complex data structures & data flows in APO

LC Live Cache; Powerful in-memory database backed by disk to manage complex data structures & data flows in APO

LUW Logical Unit of Work

LUW Logical Unit of Work

MDSB Dependent requirements view on Reservations (RESBI)

MIP Made in Plant - in house component production

MM Materials Management

MM Materials Management

MMP Model Mix Planning

MMP Model Mix Planning

MRP Material Requirements Planning

MRP Material Requirements Planning

MTO Make to Order

MTO Make to Order

MTS Make to Stock

MTS Make to Stock

OEM Original Equipment Manufacturer

OEM Original Equipment Manufacturer

PLAF SAP-ERP table that contains the planned orders

PLM Product Life cycle Management

PLM Product Life cycle Management

PP/DS Production Planning and Detailed Scheduling; part of APO

PP/DS Production Planning and Detailed Scheduling; part of APO

PVS Product Variant Structure; iPPE-BOM

PVS Product Variant Structure; iPPE-BOM

RESB SAP-ERP table containing the reservations for the materials used during the production planning

RESDT Field of table DBVM - Reset Procurement Proposals

RPM Rapid Planning Matrix; fast explosion engine of SAP SCM

RPM Rapid Planning Matrix; fast explosion engine of SAP SCM

SCM Supply Chain Management

SCM Supply Chain Management

STAD Transaction in SAP for business transaction analysis of SAP workload

WIP Work in Process

WIP Work in Process

Glossary

LPAR On IBM System p the Power5 hardware supports logical partitioning functionality in which a single large server can be divided into a number of “logical” machines. Under AIX5.2, this functionality supports the division of the machine along physical boundaries only. The physical resources such as CPU, disks, and hardware adapters are assigned to an LPAR. Under AIX5.3, hardware virtualization is introduced and multiple LPARs can share physical resources. From AIX5.2, dynamic reconfiguration is supported which allows resources to be removed or added to the active LPARs so that the LPAR can be reconfigured during production.

LUN An LUN (Logical Unit Number) is an indivisible unit presented by a storage device to its host. LUNs are assigned to each disk drive in an array so the host can address and access the data on those devices. This is a very important concept in understanding the relationship between storage devices and the piece of the storage device that can be used by a typical host.

NMON The nmon tool is an cross-platform tool used for monitoring system and IO behavior over time. It provides a data recording which can then be processed into graphs based on time slices. This tool provides the basis for an in-depth analysis of the system behavior with views to multiple concurrent metrics.

SAP Interface The SAP application servers: an instance consists of an SAP dispatcher and its work processes. The work processes are the processes carrying out the SAP application work on behalf of online users, or batch jobs. SAP instances are the basis of application server scaling. The application server instances can be located anywhere in the landscape as long as they have a backbone network connection to the database and the central instance.

SAP liveCache SAP liveCache technology is an object-based enhancement of the MaxDB database system that was developed to manage complex objects (for example, in logistical solutions such as SAP SCM/APO). In solutions of this type, large volumes of data must be permanently available and modifiable. You can use SAP liveCache technology to represent data structures and data flows (such as networks and relationships) more easily and effectively.
In contrast to MaxDB, with an optimally configured SAP liveCache database instance, all data that must be accessible to the database system is located in the main memory.

SAPCI SAP Central Instance (CI) is used to differentiate the SAP instance which contains the global components which link all the instances of an SAP system. These are the message server, and the logical lock server (enqueue). This instance is often kept separate from “production” instances and protected by HACMP™ along with the database.

WLM (AIX) This is an AIX kernel level tool that is used for prioritizing workloads. It provides the framework in which to identify the processes belonging to a particular workload, group them together, and set their resource allocations. This tool is often used to group processes of a specific characteristic into component groups for monitoring purposes. WLM can run both actively and passively. In passive mode, it is merely observing the utilization of the defined workload components. Both XMPERF and NMON are WLM aware and can monitor WLM component groups.

WLM (z/OS) The z/OS Work Load Manager component dynamically allocates or redistributes server resources such as CPU, I/O and memory across a set of workloads based on user-defined goals and their resource demand within a z/OS image. WLM can function across multiple images of z/OS, Linux or VM sharing a System z processor. WLM also assists routing components and products to dynamically direct work requests associated with a multi-system workload to run on a z/OS image within a Parallel Sysplex® that has sufficient server resources to meet customer-defined goals.

XMPERF A graphical performance monitor on AIX which can access all the OS level performance metrics and present them in a highly configurable and intuitive graphics display. This tool is a client/server architecture which can monitor multiple systems simultaneously and display them together in a combined console. This tool is used to monitor the interaction of applications distributed across more than one server.

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 publications

For information about ordering these publications, see “How to get Redbooks publications” on page 310. Note that some of the documents referenced here may be available in softcopy only.

- *Infrastructure Solutions: Design, Manage, and Optimize a 20 TB SAP NetWeaver Business Intelligence Data Warehouse,* SG24-7289-00
- *Infrastructure Solutions: Design, Manage, and Optimize a 60 TB SAP NetWeaver Business Intelligence Data Warehouse*, SG24-7385-00
- *Enhancing SAP by Using DB2 9 for z/OS*, SG24-7239-00

Other publications

These publications are also relevant as further information sources:

- *SAP on DB2 for z/OS and OS/390: High Availability and Performance Monitoring with Data Sharing,* SG24-6950-00

Online resources

These Web sites are also relevant as further information sources:

- Unicode

  `http://www.unicode.org`
- SAP dialogue instances

  `http://www.service.sap.com/instguides`
  `http://www.service.sap.com/notes`
- SAP installation

  `http://www.service.sap.com/instguides`
- DS8000

  `http://www-03.ibm.com/systems/storage/disk/ds8000/index.html`
- xmperf graphical monitor

  `http://publib16.boulder.ibm.com/doc_link/en_US/a_doc_lib/perftool/prfusrgd/ch02body.html`

How to get Redbooks publications

You can search for, view, or download Redbooks documents, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy books, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Infrastructure Solutions: SAP IS Automotive on an IBM Platform

(0.5" spine)
0.475"<->0.873"
250 <-> 459 pages

Infrastructure Solutions: SAP IS Automotive on an IBM Platform

IBM®

Redbooks®

Case study of a real world scenario

Make-to-Order manufacturing data in a complex industry

Benchmark testing in an IBM System z and System p environment

This IBM Redbooks publication describes a proof of concept project conducted by a team of IT specialists from IBM, SAP, and an automotive industry client. The project involved developing a supply chain management implementation for complex make-to-order repetitive manufacturing for configurable products, specifically the production of a variety of automobiles destined for an international market.

The project was based on the current and projected requirements of the automotive OEM. The performance testing described here was undertaken because the volume of data to be processed surpasses that of other SAP implementations of the make-to-order scenario currently in production. The purpose of the tests was to ensure that the application could be restructured to meet the requirements of this data volume and that the IBM infrastructure would be capable of processing this load in the tight time windows specified by the customer's business processes.

The customer selected a set of SAP functional performance tests and defined run time targets for these key performance indicators. The modules tested were Rapid Planning Matrix, Materials Requirements Planning, Backflush, and Materials Ledger. Actual data was provided from the customer environment.

The test environment included IBM System z with DB2 for the SAP central services and databases, System p (AIX) for the SAP application servers and liveCache, and a DS8300 Turbo with 13 TB capacity for the data. The tests were performed with mySAP ERP 2005 with IS Automotive activated; the workload was simulated through specific BladeCenter injectors.

IBM and SAP performance experts ran the benchmark tests and made some configuration adjustments. The result was an implementation that met the customer's requirements. These tests will influence major customer automotive architecture decisions.

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

SG24-7540-00 ISBN 0738431680