



IBM Customer Tuning Experience

RDBMS / Application Tuning Using An ESS

Martin S. Carangelo, Ralf Schmidt-Dannert

**IBM's Advanced Technical Support for the Americas
Foster City, California**

**Version: 1.1
Status: November 2001**



General..... 3
 Edition Notice (August 2001)..... 3
 Scope..... 3
 Acknowledgments 3
 Feedback 3
 Copyright 2001 IBM Corporation. All Rights Reserved. 3
 Introduction..... 4
 Background:..... 4
 Overview of Customer Environment:..... 4
 RS/6000 & AIX Environment 5
 ESS Storage Server Analysis 6
 Overview: 6
 Initial ESS Observations From ATS ESS Team: 6
 SSD Recommendations:..... 7
 Additional ESS Recommendations: 9
 Database Analysis 11
 Database Analysis 12
 Considerations: 12
 Observations: 12
 Recommendations: 13
 General RDBMS Recommendations: 13
 Application Analysis..... 14
 Observations: 14
 Recommendations: 14
 General RDBMS Based Application Recommendations:..... 14
 Summary 15
 Additional Comments:..... 15
 Attachment A – Environment Specific Recommendations: 16
 DB2 Specific Recommendations: 16
 PeopleSoft Specific Recommendations: 19
 Attachment B – Sample Tool Output: 20

General

Edition Notice (August 2001)

This is the first edition of this document

Scope

This document is written for IBM personnel involved with the design and layout of IBM's ESS I/O Sub-system for ERP solutions, using utilizing a RDBMS database. It is provided only as a customer experience that might provide insight and assistance into other potential design efforts.

Acknowledgments

The authors want to graciously acknowledge the work of Jim Dilley, who took the time to review, critique, and contribute to the contents of this document.

Feedback

Please send feedback to: jdilley@us.ibm.com

Copyright 2001 IBM Corporation. All Rights Reserved.

Neither this documentation, nor any part of it, may be copied or reproduced in any form or by any means or translated into another language, without the prior consent of the IBM Corporation.

IBM makes no warranties or representations with respect to the content hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. IBM assumes no responsibility for any errors that may appear in this document. The information contained in this document is subject to change without any notice. IBM reserves the right to make any such changes without obligation to notify any person of such revision or changes. IBM makes no commitment to keep the information contained herein up to date.

PeopleSoft is a registered trademark of the PeopleSoft Corporation
IBM, DB2, and AIX are registered trademarks of the International Business Machines Corporation

Introduction

Background:

A customer called upon IBM in March of 2001 for tuning and performance assistance. IBM was asked if they would perform an analysis of a RDBMS based application implemented using an IBM ESS I/O Subsystem. The customer indicated that they were fast approaching a go live date that was soon to be followed by the production of additional application solutions. They were concerned about the performance of the IBM delivered architecture based upon their testing results and end-user trials. Attempts were made to remotely look into possible problems with the system. However, it was determined that an onsite visit with the customer would be best for all parties involved. This document represents the analysis process and resulting recommendations made to this customer for setting up their ESS system for optimal performance and availability, based on their environment.

Overview of Customer Environment:

IBM put together a plan to analyze the customer's environment based on a four-part approach. The investigation of the system would take place from the following points of view:

1. RS/6000 & AIX
2. ESS Storage Server
3. DB2
4. PeopleSoft (initially Financials with HRMS soon to follow)

It was felt best to analyze the system as outlined above, in order to make sure all aspects of the solution were taken into consideration; this was actually a pretty traditional approach. This approach starts with the base foundation and then works up through the layers until the top layer, the application, is reached.

Once onsite, this outlined plan was reviewed with the customer and the analysis was begun. Below is a discussion of each piece of this process along with findings and a review of the resulting recommendations. This document will also report the previously proposed recommendations from remote analysis, and where appropriate, those recommendations that were executed while onsite with the customer.

RS/6000 & AIX Environment

Since IBM has had previous contacts with this customer regarding several performance questions, the review of the RS/6000 and AIX configuration did not take long. In those previous discussions, IBM had reviewed the configurations and also proposed various changes to the AIX configuration. It was noted that these recommendations were indeed completed and remained in affect on the system. Examples of the AIX configuration recommendations were the minservers and maxservers values for AIO servers parameters of the AIX configuration. Also, AIX maintenance levels were checked to assure proper levels were applied to our system (this is always an essential step whenever approaching this type of a complex environment). Due to this fact, and time constraints involved, it was felt that the AIX and RS/6000 configurations were not an issue at this point. The team decided to move on to other components of the analysis. The OS level recommendations were performed prior to the customer visit and are therefore outside the scope of this document.

ESS Storage Server Analysis

Overview:

There were numerous recommendations associated with the reconfiguration of the ESS device. These recommendations deal with the creation and maintenance of storage objects that reside on the ESS. These recommendations were generated from the analysis that is documented in this paper.

Working with the IBM Advanced Technical Support (ATS) ESS Storage Group in San Jose California, the required information for performing a complete analysis of the ESS Server was identified. New AIX commands/utilities were received from the ESS ATS Team and installed to aid in the collection of this required information. These commands/utilities were installed on each of the five RS/6000 servers that were attached to the ESS. General information was also collected with regard to the overall configurations of each of the RS/6000 systems attached to the ESS. These commands were run on each of the attached systems. The results were then packaged and sent to the team in San Jose for additional analysis.

To see samples of the output from each of these commands, please see Attachment B.

Initial ESS Observations From ATS ESS Team:

After receiving the information from the local team, the ATS ESS Team analyzed the information and responded with the following observations and recommendations:

1. Physical cabling of the ESS was incorrect. Cables initiating from the five servers were not plugged into the ESS in a way that would allow the achievement of maximum performance and throughput. Additionally, these connections were not correctly configured to allow the HACMP fail-over process to function properly.
2. The internal cabling, within the ESS, was not setup for optimal performance and throughput.
3. The ESS was not at the proper level of Microcode. Neither was the proper level of SDD software installed on the servers. The result of this was that SDD would not function properly to balance the activity down each SCSI channel. This would also effect the operation of the HACMP fail-over process, causing it to not work properly.
4. It was finally determined that the storage allocation done on the ESS was not performed properly, or rather was not done in an appropriate order to allow for the optimal performance and throughput for the ESS device. The sizing of LUNS for a single attached system and the placement of volume groups and logical volumes did not follow the standard guidelines that have been established for optimal performance of the ESS device.

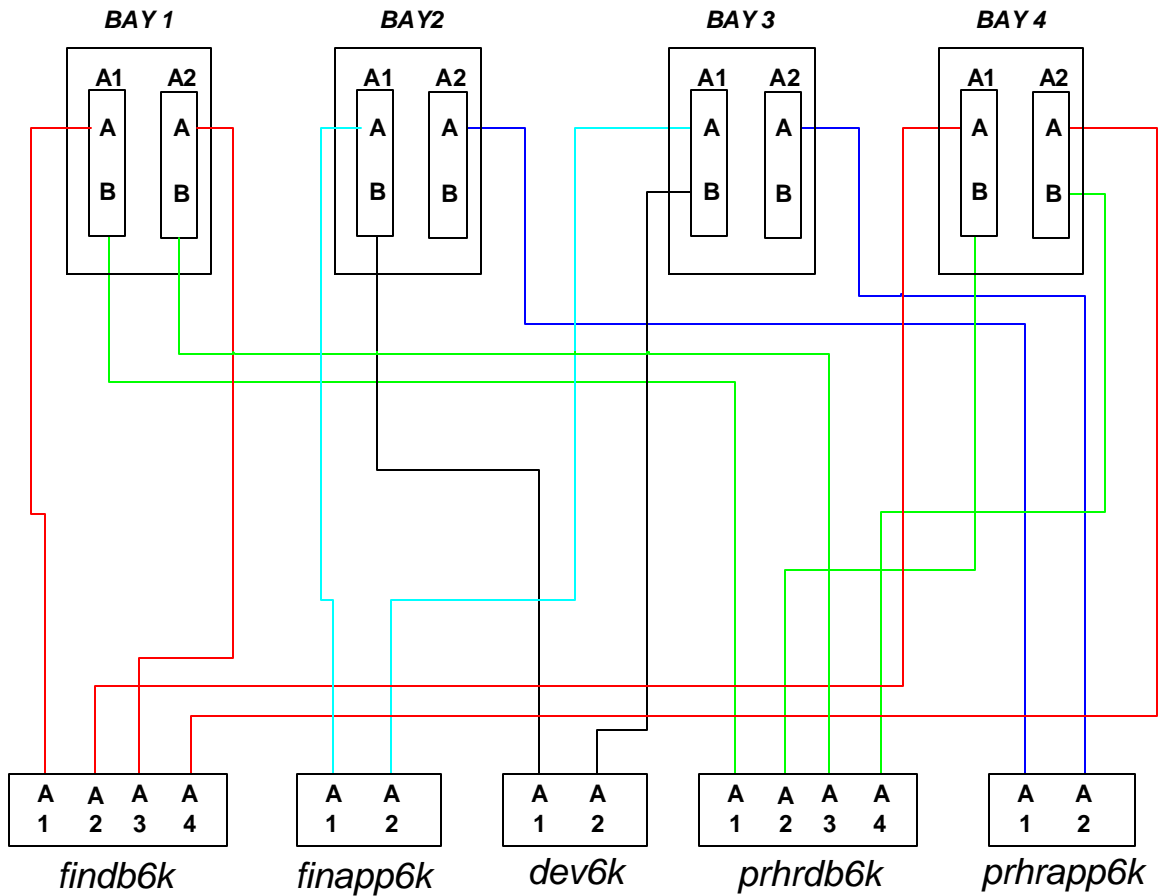
ESS ATS Team Recommendations:

From these observations and the resulting and subsequent analysis, the ESS ATS team put forth the following recommendations.

1. Obtain the proper level of Microcode and SDD software for a customer using HACMP. This customer's current level of micro-code was G3 but needed to be at level G3+2, which was the current level at that time. The latest microcode and driver levels for the ESS can be found at the following WEB site: <http://ssddom01.storage.ibm.com/techsup/swtechsup.nsf/support/sddupdates>
2. The CD containing the new micro-code would also contain the correct release of the SDD software for this environment. It was determined that this was the first action that needed to be accomplished in order to address all the identified issues. Obtaining the correct micro-code and software levels should always be a high priority in this environment.
3. It was recommended that the customer perform a complete reconfiguration of their ESS device. This reconfiguration should be done as if the device was newly delivered and being setup for the first time. The reason for this level of reconfiguration was simple. SDD Engineering felt that the failure to use standard methods of hardware setup, configuration, and allocation of the storage subsystem created a questionable base foundation. They felt it was best to start from a solid and known configuration. Therefore, this "tear it down and build it back up" approach would move the customer back to standards and give them a much better foundation for achieving the optimal performance and throughput from their ESS device.

The following graphic shows the final recommended cabling to support the customer requirements.

ESS - Server Connection Diagram



Additional ESS Recommendations:

To determine an optimal ESS sub-system layout, the following questions were discussed with the customer and should be part of any analysis for this type of environment:

1. How many systems are connected to the ESS?
2. Is shared access to disk space required?
3. How much disk space will be required for each of these systems now and what is the future growth expectation?
4. Are ESS Copy Services required, if so, which?


Based on the results of these questions, the following additional recommendations were made to the customer:

1. Each of the 6 ESS ranks (18.2GB disks) should be split up into 13 LUNs, 8GB each.
2. All volume groups created after the reconfiguration will be required to have a multiple of 6 LUNs, one from each ESS rank, assigned to them at all times. This approach will spread IO activity across the entire ESS device and result in optimal performance and throughput.
3. To support the spreading of IO over the entire ESS device volume group striping with a physical partition (PP) size of 2MB should be used. Striping on the logical volume level should not be used.
4. All logical volumes created must be sized based on a multiple of 12mb. This is because of the requirement to spread IO across the entire ESS device to achieve optimal performance and throughput. The 12MB figure is derived from the AIX physical partition size of 2MB per LUN and that there are 6 ranks in the ESS device. This means that it is important to allocate a minimum of 6 times 2MB, one PP from each of the LUNs on each of the arrays to achieve the ultimate spread of the IO across the entire ESS sub-system. Logical volumes may also be any multiple of 12MB and this will also follow the performance configurations methodology.
5. Each system attached to the ESS, which houses a database, should have a specific number of LUNs based on projected need and growth. This number of LUNs has to be further divided into a number of physical locations on the ESS device. Each system must have its placement of the LUNs calculated to allow for optimal performance. This has to be based on whether it is a production system or a test system. The creation of logical volumes, and the actual placement of these logical volumes, must also follow the developed methodology. For example, a logical volume created for the financials production environment must be placed on its assigned LUNs for the server that it resides on in order to be considered in compliance and capable of achieving optimal performance. This means that storage reserved for a particular system should not be used to support another system, no mixing.


Several LUNs will be unassigned and available for future additional disk space requirements. These LUNs should be placed in a fashion, that the assignment of them to a specific application domain follows the methodology described above.

The following graphics show the proposed distribution of assigned disk space to the different customer application domains:

Layout and Usage of LUNs Within One LSS



Distribution of Data Over All 6 LSSs



Database Analysis

Considerations:

It is important to note that hardware alone, and reconfigurations to that hardware, will generally not allow customers to achieve the optimal performance of their environment for any given hardware solution. Careful consideration must be paid to both the database and application layers, as part of the overall environment, in order to achieve optimal performance out of any implemented solution. It is in this analysis phase that IBM looked into the database layout and the configurations within the application environment.

The physical database layout, particularly the tablespace layout, is an important analysis element. Simply put, this layer of the system sits on top of the ESS storage device. It was determined that issues existed with the design, size, and placement of several database objects. These issues, and their ramifications, are of equal or greater value towards the overall system performance than that of just the ESS hardware and any potential issues associated with such a device.

Observations:

Here are some high level observations made of this customer's environment:

1. This customer had numerous database configuration parameters, which were not set for optimal database operation. These settings were causing the overall operation of the system to be performing at a less than optimal level. Database specific recommendations were made to correct these settings and are noted in Attachment A of this document.
2. There were many more tablespaces defined for this application, than are normally configured as determined by the associated application's design and architecture. After analysis, it was discovered that further separation and segregation of DB objects were still being implemented. This separation represented a deviation from the application vendor's delivered design.
When further investigated it was determined that this separation methodology was thought to represent a performance advantage. It was being done purely from database theory for laying out the database objects and did not take into consideration the underlying hardware. Older database theory says that you must not mix certain types of objects with one another or grave performance consequences may occur. This was true in the days of non-RAID systems but is not necessarily true for newer breeds of hardware technologies like the ESS.
3. This customer used no striping for the database. Even with a certified and optimally configured hardware platform below it, the database layer can be setup to further enhance or hinder the overall system performance by how sub-system

layers are utilized. This should be a point of focus when evaluating performance issues. It is important to ensure that the various striping techniques implemented at the different levels of the total solution work together for an optimally performing solution.

Recommendations:

From the observations and analysis of the database layer, the following recommendations were developed:

General RDBMS Recommendations:

1. Verify that the database configurations parameters are correct for the environment. Understand how these specific parameters may have a negative impact. The potential improvement to be gained with parameter changes may be small or large. In this regard, Attachment A contains these recommendations that are specific to DB2 for this customer environment. Each database needs to be evaluated separately.
2. This recommendation is largely based on the availability of database management tools to the DBA team. If appropriate tools are available we recommend that the customer use raw logical volumes to house their database. The advantage in using raw logical volumes is increased performance and faster recovery. With the availability of management tools, there are no real disadvantages. Tools for the DBA and system administrators will allow for ease of management and reduce the chance of errors when dealing with raw logical volumes. However, if these management tools are not available, it is recommended that the customer does NOT use raw logical volumes; instead AIX file systems should be used. This will cause some loss of the performance advantage of raw logical volumes but AIX file systems will allow easier management, reporting, and planning of the environment.
3. As previously mentioned, older storage configurations require a different mindset for database layout than the newer hardware technologies like the ESS. There are no additional advantages gained by further separating indexes or data from the organization typically delivered by the application vendor in this I/O environment. It was therefore recommended to eliminate the additional tablespaces that were created for older styles of database management and put objects back in their original locations. As a benefit, this will make management easier by placing objects where they are expected. Thus, this will make support and maintenance easier.
4. Always verify current maintenance releases have been installed on all RDBMS systems that access the ESS. This statement applies to the entire hardware and software environment.

Application Analysis

Observations:

The final area that was analyzed was the overall PeopleSoft applications environment. Here, such elements as PeopleTools, Tuxedo, the Process Scheduler, the applications servers and SQR were analyzed. The recommendations coming out of this analysis are listed below. Please note that the file system components of the PeopleSoft architecture were not analyzed but its design and use in the overall architecture of the system was discussed.

Recommendations:

A lot of the analysis and recommendations already discussed could have been placed in this category but were not for organizational purposes. Note that specific PeopleSoft recommendations are in Attachment A.

General RDBMS Based Application Recommendations:

These are general comments that apply to database applications implemented using an ESS I/O Subsystem:

1. Most ERP applications have “local” code associated with many of their various elements. Local meaning that it is not stored in the database. A classic example would be the executables for the application server layer, log files, startup scripts, etc. These should not be stored on the ESS. To put it another way, any existing file system that contains other than database objects should be defined on local disk.
2. Validate the integrity of the database in terms of the application. Many applications provide tools that can be used to “audit” the integrity of the database against what is expected from the application. As an example, a missing index will not be an operational problem to the application, but will be an overall performance problem to the system that is making database calls that were based on the expectation that the index exists correctly, and this query represents the optimal path to the required data.

Summary

At the start of this evaluation, the customer was I/O constrained during their tests. After implementing the changes recommended in this paper, they were CPU bound. Additional evaluation by the customer, after this evaluation, uncovered several indexes that were missing or did not match what was expected out of the calling application program. The correction of these discrepancies ultimately fixed the perceived problem initially blamed on the hardware.

Throughout the many performance problems that this organization has dealt with over the past several years, there is a common thread. That thread is the correcting of hardware configurations, OS configurations, and DB configurations, can only build the required base for solving the more common and typical problems. Generally, it is application design or the associated database design that provide the largest benefits to performance problems in this type of complex environment.

Additional Comments:

One final recommendation, and it applies to all areas, is the need for tools to aid in the management of this environment. Tools to help the administration staff to address the day-to-day monitoring, maintenance, and planning issues associated with these types of complex hardware and software implementations. For both the UNIX and the database side, these tools can assist with and address the most common database and UNIX System requirements. These include database and UNIX administration tools to service and support operations, performance management and other utilities. Additionally, recovery and replication management tools, that also support archival, image copy, and point in time recovery, should be considered. An example of tools in this category are offered by vendors such as IBM, with Tivoli and DB2 management tools, and Quest software, with its Quest Central for DB2.

Attachment A – Environment Specific Recommendations:

This is a discussion of customer specific parameter changes. They are provided here for reference and the potential insight it might provide to other RDBMS solutions.

DB2 Specific Recommendations:

1. Besides the two DB2 instances for the production environments there should be a single DB2 Instance for each of the product groups within the DEVELOPMENT environment. Within these instances, multiple databases should be created to support the various application systems. Then, under these instances, there should be placed all of the databases associated with those environments. For example, under the financials DEVELOPMENT instance would be a TEST database, a development database, a reporting database, and a QA database. The same for HRMS. This makes much better use of the DB2 memory management structures and also simplifies overall database management. A major benefit of this architecture is that it will use the limited amount of physical memory on the development hardware more efficiently. The allocation of a single large buffer pool, as opposed to the many smaller ones that would have been required by the use of one DB2 instance per database, will function better and optimize memory usage.
2. Delivered application objects should be created within DMS tablespaces. Currently there is a mix of SMS and DMS type tablespaces. SMS tablespaces are not bad, but should only be used for static or low growth DB objects. This is because of the system overhead incurred when having to grow a SMS DB object. The use of DMS type objects will reduce this overhead and improve overall system performance.
3. DB2 Catalog and Temp areas should reside in SMS type tablespaces.
4. Combine the multiple defined buffer pools defined inside one database into one large buffer pool, which is then used by all tablespaces. Buffer pools ultimately tie back to physical memory in the server hardware. This customer has 6GB of memory and was not utilizing it properly for the database layer of the systems. The intention is to let DB2 manage and distribute buffer space to tablespaces as needed. This is only possible if it is feeding off a single repository. Separating memory into individual buffer pools implies that some memory (ultimately the physical server memory) will lie idle for low activity tablespaces thus not utilizing the capacity of the server to its fullest. During further stress testing and production, the usage of the buffer pool should be monitored closely for tables / tablespaces using too many blocks in the buffer pool, resulting in forcing active data blocks out of the buffer pool. If this is observed, these tables / tablespaces should be assigned their own buffer pool, as long as physical memory is available.

5. Modify the following configuration settings. The values given are a starting point and need to be revisited, particularly as load is applied to the system. The values shown should be applied to the two production instances only. The configuration of the DB2 databases on the development / test system need to be investigated along the same lines and their parameter adjusted accordingly.

Set the DB2 environment variables from the shell level using the “db2set” command.

DB2 environment	<u><i>New Value</i></u>	Description
DB2_PARALLEL_IO	*	Parallelism to use when reading data from disk
DB2MEMMAXFREE	8000000	Specifies the amount of free memory that is retained by each DB2 UDB agent
DB2_STRIPED_CONTAINERS	ON	Ensure, that the extents allocated in a tablespace are aligned to the stripes in the RAID array.

Changes recommended for the database manager configuration:

DBM Parameter	<u><i>New Value</i></u>	Description
SHEAPTHRES	> (3 x SORTHEAP)	Sort heap threshold
DIAGLEVEL	3	Amount of information shown in DB log file
INTRA_PARALLEL	NO	Enable intra-partition parallelism
MAX_QUERYDEGREE	1	Maximum query degree of parallelism

Changes recommended for database configuration:

DB Parameter	<u><i>New Value</i></u>	Description
DBHEAP	20000	Database heap
CATALOGCACHE_SZ	4096	Catalog cache size
LOGBUFSZ	2560	Log buffer size
UTIL_HEAP_SZ	10000	
BUFFPAGE	320,000	Buffer pool size
SORTHEAP	4096	Sort heap size. The system needs to be watched closely for sorts going to the temporary tablespace. If such sorts are detected, this parameter should be increased, as long as physical memory is available.
APPLHEAPSZ	10000	Application heap size
PCKCACHESZ	4096	Package cache size
CHNGPGS_THRESH	40	Changed pages threshold for asynchronous page cleaners

SOFTMAX	500	Recovery range and soft checkpoint interval
DFT_DEGREE	ANY	Default degree for intra-partition parallelism
MAX_QUERYDEGREE	1	
NUM_IOCLEANERS	20	Number of asynchronous page cleaners
NUM_IOSERVERS	40	Number of I/O servers
AVG_APPLS	7	Average number of active applications
DFT_PREFETCH_SZ	<multiple of 8> and extent size used	Default prefetch size
DFT_EXTENT_SZ	<multiple of 8>	Default extent size of table spaces

6. Verify the ESS location of tablespaces:
Select TBSPACE, TBSPACEID, TBSPACETYPE, EXTENTSIZE, PREFETCHSIZE, PAGESIZE, NGNAME, BUFFERPOOLID from syscat.tablespaces.
Once located, if the tablespace does not conform to the new ESS placement and allocation methodology it must be deleted and re-created using one of the methods discussed in the parent document. Do this for all tablespaces in all of the databases residing on the ESS. Continue to abide by the developed storage allocation methodology. It will be a labor-intensive task in the beginning but will require minimal efforts to maintain and benefits will continue.
7. Use a tablespace extent size of 8 for tablespaces containing the indexes (4 KB block size assumed), to match the ESS stripe size of 32KB.
8. For large tablespaces, particularly those that PeopleSoft delivers with the word LARGE in the tablespace name, increase the tablespace extent size to minimize I/O. This addresses the idea of further improving the overall performance of the system through additional DB2 optimization. It is recommend using a one-megabyte extent size (also the stripe size) for high-growth, frequently updated, large DB2 tablespaces. For example, the PS_JRNL_LINES table of the financial product could greatly benefit from an increased extent size as applied to the tablespace GLLARGE. Ensure that the specified prefetch size is set accordingly as well.
9. DB2_STRIPED_CONTAINERS should be set to ON before the database or tablespace restore. This ensures, that the extents in the tablespace are aligned to the stripes in the RAID-5 arrays on the ESS.
10. Verify the size of the lock list and check for lock escalations. It should be kept in mind, that it is sometimes better to create a table lock than to create a huge number of rowlocks.

PeopleSoft Specific Recommendations:

The comments above, applied to a PeopleSoft environment, translate to the following recommendations:

1. Tuxedo and SQR code should be moved off of the ESS device and be placed on internal drives located on the servers. This code seldom changes and really belongs as part of the systems group of software. Placing it on the ESS makes it more difficult to manage while not providing any performance benefits.
2. The DBA's and the PeopleSoft Systems administrator should frequently run the two delivered utilities for this application, DDDAUDIT and SYSAUDIT. These utilities should be run regularly for each system that has a PeopleSoft database.

Attachment B – Sample Tool Output:

Here are partial samples for each of the types of information collected on one of the servers and subsequently sent to San Jose.

finddb_ls2105.out:

Hostname	VG	hdisk	Location	LUN SN	S	Connection	Size	LSS	Vol	Rank
findb6k	None	hdisk2	14-08-00-1,0	10112973	Y	R1-B1-H1-ZA				
findb6k	None	hdisk3	14-08-00-1,1	10312973	Y	R1-B1-H1-ZA	10.2	17	3	1100
findb6k	None	hdisk4	14-08-00-1,2	10412973	Y	R1-B1-H1-ZA	10.2	17	4	1100
findb6k	None	hdisk5	14-08-00-1,3	10512973	Y	R1-B1-H1-ZA				
findb6k	None	hdisk6	14-08-00-1,4	10912973	Y	R1-B1-H1-ZA	5.2	17	9	1100
findb6k	None	hdisk7	14-08-00-2,0	20712973	Y	R1-B1-H1-ZA	3.2	18	7	1200
findb6k	None	hdisk8	14-08-00-2,1	20812973	Y	R1-B1-H1-ZA	4.2	18	8	1200
findb6k	None	hdisk9	14-08-00-2,2	20112973	Y	R1-B1-H1-ZA				
findb6k	None	hdisk10	14-08-00-2,3	20D12973	Y	R1-B1-H1-ZA	5.2	18	13	1200
findb6k	None	hdisk11	14-08-00-2,4	20E12973	Y	R1-B1-H1-ZA	5.2	18	14	1200

finddb_lsess.out:

Disk	Location	LUN SN	Type	Size	LSS	Vol	Rank	S	Connection
hdisk3	14-08-00-1,1	10312973	IBM 2105-F20	10.2	17	3	1100	11	Y
hdisk4	14-08-00-1,2	10412973	IBM 2105-F20	10.2	17	4	1100	11	Y
hdisk6	14-08-00-1,4	10912973	IBM 2105-F20	5.2	17	9	1100	11	Y
hdisk7	14-08-00-2,0	20712973	IBM 2105-F20	3.2	18	7	1200	02	Y
hdisk8	14-08-00-2,1	20812973	IBM 2105-F20	4.2	18	8	1200	02	Y
hdisk10	14-08-00-2,3	20D12973	IBM 2105-F20	5.2	18	13	1200	02	Y
hdisk11	14-08-00-2,4	20E12973	IBM 2105-F20	5.2	18	14	1200	02	Y
hdisk13	14-08-00-3,1	40812973	IBM 2105-F20	3.2	20	8	1400	03	Y
hdisk14	14-08-00-3,2	40912973	IBM 2105-F20	4.2	20	9	1400	03	Y
hdisk15	14-08-00-3,3	40A12973	IBM 2105-F20	4.2	20	10	1400	03	Y

finddb_lsdd.out:

Hostname	VG	vpath	hdisk	Location	LUN SN	S	Connection	Size	LSS	Vol	Rank
findb6k	None	vpath0	hdisk2	14-08-00-1,0	10112973	Y	R1-B1-H1-ZA				
findb6k	None	vpath0	hdisk64	17-08-00-6,0	10112973	Y	R1-B3-H1-ZA				
findb6k	None	vpath0	hdisk94	2A-08-00-5,1	10112973	Y	R1-B2-H1-ZA				
findb6k	None	vpath0	hdisk122	2D-08-00-4,1	10112973	Y	R1-B4-H1-ZA				
findb6k	None	vpath1	hdisk3	14-08-00-1,1	10312973	Y	R1-B1-H1-ZA	10.2	17	3	1100
findb6k	None	vpath1	hdisk66	17-08-00-6,2	10312973	Y	R1-B3-H1-ZA	10.2	17	3	1100
findb6k	None	vpath1	hdisk93	2A-08-00-5,0	10312973	Y	R1-B2-H1-ZA	10.2	17	3	1100
findb6k	None	vpath1	hdisk123	2D-08-00-4,2	10312973	Y	R1-B4-H1-ZA	10.2	17	3	1100
findb6k	None	vpath2	hdisk4	14-08-00-1,2	10412973	Y	R1-B1-H1-ZA	10.2	17	4	1100
findb6k	None	vpath2	hdisk67	17-08-00-6,3	10412973	Y	R1-B3-H1-ZA	10.2	17	4	1100
findb6k	None	vpath2	hdisk95	2A-08-00-5,2	10412973	Y	R1-B2-H1-ZA	10.2	17	4	1100
findb6k	None	vpath2	hdisk121	2D-08-00-4,0	10412973	Y	R1-B4-H1-ZA	10.2	17	4	1100
findb6k	findbvg1	vpath5	hdisk7	14-08-00-2,0	20712973	Y	R1-B1-H1-ZA	3.2	18	7	1200
findb6k	findbvg1	vpath5	hdisk46	17-08-00-3,1	20712973	Y	R1-B3-H1-ZA	3.2	18	7	1200
findb6k	findbvg1	vpath5	hdisk80	2A-08-00-3,0	20712973	Y	R1-B2-H1-ZA	3.2	18	7	1200
findb6k	findbvg1	vpath5	hdisk131	2D-08-00-6,1	20712973	Y	R1-B4-H1-ZA	3.2	18	7	1200