

1. Record Nr.	UNINA9910828466903321
Titolo	IBM eserver pSeries sizing and capacity planning : a practical guide // [G. Benton Gibbs ... et al.]
Pubbl/distr/stampa	[S.l.], : IBM, International Technical Support Organization, c2004
Edizione	[1st ed.]
Descrizione fisica	xviii, 546 p. : ill
Collana	IBM redbooks
Altri autori (Persone)	GibbsG. Benton
Disciplina	004.2/4
Soggetti	Computer capacity - Planning IBM computers
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	The "e" in "eserver" is printed as the symbol for electronic. "March 2004."
Nota di bibliografia	Includes bibliographical references and index.
Nota di contenuto	Front cover -- Contents -- Notices -- Trademarks -- Preface -- The team that wrote this redbook -- Become a published author -- Comments welcome -- Part 1 Introduction to sizing and capacity planning -- Chapter 1. Overview, concepts, and approach -- 1.1 Definitions of common terms -- 1.2 Concepts -- 1.2.1 Required knowledge and experience -- 1.2.2 Sizing with capacity planning -- 1.2.3 The sizing problem -- 1.2.4 Sizing inputs -- 1.2.5 Sizing outputs -- 1.2.6 Who performs sizing -- 1.3 Sizing and resizing process -- 1.3.1 System design and requirements -- 1.3.2 Sizing model -- 1.3.3 Hardware requirements -- 1.3.4 Building block choices -- 1.3.5 eConfig for the price -- 1.3.6 Sales, purchase, install, and production -- 1.3.7 Gathering performance data -- 1.3.8 Performance tuning -- 1.3.9 Estimated or measured growth -- 1.3.10 Capacity planning -- 1.3.11 Resizing model -- 1.3.12 rPerf reliance -- 1.4 Weighing sizing components -- 1.4.1 Memory size -- 1.4.2 Disk type and number -- 1.4.3 Adapters for disk, tape and network -- 1.4.4 Software -- 1.4.5 Summary -- 1.5 The importance of the right amount of information -- 1.5.1 Brain overload -- 1.5.2 Summary -- 1.6 A practical sizing method -- 1.6.1 Segmentation -- 1.7 Performance theory -- 1.8 General rules of thumb for RDBMS memory -- 1.8.1 Application resident set -- 1.8.2 RDBMS data and file system cache -- 1.8.3 RDBMS utilization rules of thumb -- 1.8.4 Utilization -- 1.8.5 RDBMS raw data

to disk rules of thumb -- 1.8.6 RDBMS disk use rules of thumb -- 1.9 The performance saturation curve -- 1.10 Successive approximation and sizing levels -- 1.11 Plagiarism -- 1.12 Triangulation -- 1.12.1 A triangulation story -- 1.13 Common sizing mistakes -- 1.13.1 Sizing report outline -- 1.13.2 A sizing story -- 1.14 The eConfig configurator -- 1.14.1 Configurator test -- 1.15 Cost-based sizing method.

1.16 High availability and disaster recovery -- 1.17 Capacity Upgrade on Demand -- 1.18 Sizing for Linux on pSeries -- Part 2 Components involved in sizing and capacity planning -- Chapter 2. Hardware components -- 2.1 Performance methodology -- 2.2 Overview of pSeries systems -- 2.2.1 Autonomic computing -- 2.2.2 e-business on demand -- 2.2.3 Reliability, availability, and serviceability features -- 2.2.4 Capacity Upgrade on Demand -- 2.3 pSeries processors -- 2.3.1 Processor descriptions -- 2.3.2 RISC/CISC concepts -- 2.3.3 Superscalar architecture: Pipelines and parallelisms -- 2.3.4 32-bit versus 64-bit computing -- 2.3.5 Performance of processors -- 2.3.6 Processor evolution -- 2.4 Memory -- 2.4.1 Memory hierarchy -- 2.4.2 Locality concept -- 2.4.3 Caches -- 2.4.4 Memory cycles -- 2.4.5 Virtual memory concepts -- 2.4.6 Memory affinity -- 2.4.7 Large page support -- 2.5 Input/output -- 2.5.1 Peripheral Component Interconnect -- 2.5.2 PCI-X -- 2.6 Storage architectures -- 2.6.1 Direct access storage -- 2.6.2 Storage area networks -- 2.6.3 Network-attached storage -- 2.6.4 RAID -- 2.6.5 IBM TotalStorage Enterprise Storage Server -- 2.6.6 IBM TotalStorage Fibre Array Storage Technology -- 2.6.7 IBM 7133 Serial Disk System -- 2.6.8 IBM TotalStorage Expandable Storage Plus 320 -- 2.6.9 The IBM TotalStorage Network Attached Storage -- 2.7 Additional hardware considerations -- 2.7.1 Multiprocessor configurations -- 2.7.2 NUMA -- 2.7.3 Logical partitioning -- 2.7.4 Dynamic logical partitioning (5.2.0) -- 2.7.5 Dynamic CPU sparing and CPU Guard (5.2.0) -- 2.7.6 UE-Gard (5.2.0) -- Chapter 3. Software components -- 3.1 AIX -- 3.1.1 History of AIX -- 3.1.2 AIX kernel -- 3.1.3 Modes of operation (execution modes) -- 3.1.4 AIX 5L kernel subsystems -- 3.1.5 Multitasking and multithreading support -- 3.1.6 64-bit kernel -- 3.2 Workload Manager -- 3.2.1 Classes. 3.2.2 Tiers -- 3.2.3 Class attributes -- 3.3 Linux -- 3.3.1 Linux for pSeries -- 3.3.2 Linux and AIX -- 3.3.3 Logical partitioning -- 3.3.4 Other related information and links -- Chapter 4. Benchmarks -- 4.1 Introduction to benchmarks -- 4.2 OLTP benchmarks -- 4.2.1 TPC-C benchmark -- 4.3 Business intelligence benchmarks -- 4.3.1 TPC-H benchmark -- 4.4 e-business benchmarks -- 4.4.1 TPC-W benchmark -- 4.4.2 SPEC JBB2000 benchmark -- 4.4.3 SPECweb99 benchmark -- 4.5 High Performance Computing benchmarks -- 4.5.1 SPEC CPU2000 benchmark -- 4.5.2 LINPACK benchmark -- 4.6 ISV benchmarks -- 4.6.1 SAP Standard Application benchmarks -- 4.6.2 Oracle Applications Standard benchmark -- 4.6.3 Siebel platform sizing and performance program benchmark -- 4.7 Relative performance -- Part 3 Sizing pSeries systems -- Chapter 5. General sizing -- 5.1 Where to locate the Balanced System Guideline -- 5.2 Six golden sizing principles -- 5.2.1 Correct processor configuration -- 5.2.2 Balanced systems -- 5.2.3 CPU magic number calculations -- 5.2.4 Estimating CPU power -- 5.2.5 Estimating memory sizing -- 5.2.6 Estimating disk sizing -- 5.3 The Balanced System Guideline overview -- 5.3.1 Problems with sizing -- 5.3.2 Assumptions: Prerequisites for using the spreadsheet -- 5.3.3 Spreadsheets: Pros and cons -- 5.3.4 The Balanced System Guideline sections -- 5.4 The Balanced System Guideline details -- 5.4.1 Introduction sheet -- 5.4.2 Performance and

balanced systems sheets -- 5.4.3 Balanced system examples -- 5.4.4 LPAR sheet -- 5.4.5 pSeries costs -- 5.4.6 Price-based sizing -- 5.4.7 Sizing new systems -- 5.4.8 Sizing CPU and RAM sheet -- 5.4.9 Sizing and planning disks sheet -- 5.4.10 Sizing Results sheet -- 5.4.11 Calibration sheet -- 5.4.12 Calibrating a new workload example: SAP, DB2, pSeries 650 -- 5.5 Resizing existing systems for upgrades -- 5.5.1 Assumptions. 5.5.2 ResizeCPU sheet -- 5.5.3 ResizeRAM sheet -- 5.5.4 ResizeDisk sheet -- 5.5.5 ResizeDiskUse sheet -- 5.5.6 Modeling to add new workloads -- 5.6 Balanced System Guideline and sizing levels -- 5.6.1 Sizing for Level 2: 'Ball park' or rough estimates -- 5.6.2 RDBMS server sizer for level 3: Consider opinion -- 5.6.3 Sizing for Level 4: Sizing from measured data -- 5.7 Business intelligence sizing -- 5.7.1 Business intelligence golden rules -- 5.7.2 Business intelligence sizing approaches -- 5.7.3 Business intelligence sample configurations -- 5.8 Disk and stripe sizing -- 5.8.1 Disk sizing -- 5.8.2 Stripe sizing -- 5.9 pSeries 670 and 690 RIO-2 I/O Sizing Tool -- 5.9.1 Notes and assumptions -- 5.9.2 Readme sheet -- 5.9.3 Adapters sheet -- 5.9.4 Results sheet -- 5.9.5 p670/p690 errors sheet -- 5.9.6 RIO-2 loops sheet -- 5.10 Review and summary -- Chapter 6. Application-specific sizing -- 6.1 IBM applications -- 6.1.1 DB2 -- 6.1.2 Lotus Domino -- 6.1.3 Tivoli Storage Manager -- 6.1.4 WebSphere -- 6.2 ISV applications -- 6.2.1 eSizings@us.ibm.com sizing support -- 6.2.2 Quick e-sizing guides -- 6.3 IBM Eserver Sizing Guide -- 6.4 Network File System sizing -- 6.4.1 Functionality -- 6.4.2 Cache management on an NFS client -- 6.4.3 Performance considerations -- 6.4.4 Method and sizing factors -- Part 4 Capacity planning -- Chapter 7. AIX tools for data gathering -- 7.1 AIX standard tools -- 7.1.1 The vmstat command -- 7.1.2 The iostat command -- 7.1.3 The sar command -- 7.1.4 The svmon command -- 7.1.5 The ps command -- 7.1.6 The ipcs command -- 7.1.7 The topas command -- 7.2 Performance Toolbox -- 7.3 AIX Workload Manager -- 7.3.1 Configuring AIX Workload Manager -- 7.3.2 System capacity and sizing for workload management -- 7.3.3 The wlmstat command -- 7.4 Performance Management Services for AIX -- 7.4.1 Architecture -- 7.4.2 Utilization. 7.4.3 Comparison, correlation, forecast -- 7.4.4 PM/AIX usage -- 7.4.5 Data collection -- 7.4.6 Thresholds -- 7.4.7 SRM reports -- 7.4.8 Executive reports -- 7.4.9 Capacity reports -- 7.4.10 Workload specific reports -- 7.4.11 Application response metric reports -- 7.4.12 System analysis and forecast with PM/AIX -- Chapter 8. Features and tools for capacity planning -- 8.1 Performance Toolbox -- 8.1.1 Tool utilization strategy -- 8.1.2 azizo -- 8.1.3 xmtrend -- 8.1.4 jazizo -- 8.1.5 wlmperf -- 8.2 Workload Manager -- 8.2.1 Typical UNIX system capacity sizing -- 8.2.2 Server consolidation considerations -- 8.2.3 System capacity sizing for workload management -- 8.2.4 Conclusion -- 8.3 Dynamic LPAR and CUoD -- 8.3.1 Configuration alternative -- 8.3.2 DLPAR benefit -- 8.3.3 Partitioning misconceptions -- 8.3.4 Example situations using LPAR -- 8.3.5 DLPAR sizing considerations -- 8.3.6 DLPAR and applications -- 8.3.7 CUoD advantage: Pay as you grow -- 8.3.8 Workload Manager versus DLPAR -- 8.3.9 Capacity planning for DLPAR -- 8.3.10 DLPAR examples -- 8.4 IBM Insight tools -- 8.4.1 IBM Insight for SAP R/3 overview -- 8.4.2 IBM Insight for Oracle database -- Part 5 Appendixes -- Appendix A. Sanity check before upgrading -- Identifying the workloads -- Setting objectives -- Identifying critical resources -- Minimizing critical-resource requirements -- Using the appropriate resource -- Reducing the requirement for the critical resource -- Structuring for parallel use of resources -- Reflecting priorities in resource allocation -- Repeating

the tuning steps -- Applying additional resources -- Appendix B.  
Sample for CPU resource usage calculation -- Abbreviations and  
acronyms -- Related publications -- IBM Redbooks -- Other resources  
-- Online resources -- How to get IBM Redbooks -- Help from IBM --  
Index -- Back cover.

---