

E&P Data Management Innovations

Presenter: Yogi Schulz

Superior Application Performance



**Achieving
Superior Application
Performance
amidst
Skyrocketing
Data Volumes**



PPDM™

Yogi Schulz



April 2004

Copyright © 2004 by Corvelle Management Consultants



- ❑ **President of Corvelle Management Consultants**
- ❑ **Information technology related management consulting**
- ❑ **Project management and systems development**
- ❑ **Computing Canada & Calgary Herald columnist**
- ❑ **PPDM Association board member**
- ❑ **Industry presenter:**
 - Project World - 4 years
 - CIPS Informatics - 7 years
 - PMI - Information Systems SIG
 - Convergence - 4 years
 - PPDM Association - several years



Presentation Outline

Superior Application Performance




- ❑ Origin of skyrocketing data volumes
- ❑ Large data volume issues
- ❑ Achieving superior application performance
- ❑ Recommendations
- ❑ Questions & Answers



Yogi Schulz

Corvelle



Origin of Skyrocketing Data Volumes

Application Performance



- ❑ **Average data center – storage in gigabytes:**

- 1997 1,000 2002 35,000

- ❑ **Seismic traces – number per square kilometer:**

- 1997 72,000 2002 120,000

- ❑ **Well logs – megabytes per trip:**

- 1990 LIS 50 2000 DLIS 100

- ❑ **Visualization – storage in gigabytes:**

- 1997 none 2002 significant



Yogi Schulz




Large Data Volume Issues

Superior Application Performance



- ❑ **Sluggish online response times**
- ❑ **Long wait times to copy large datasets**
- ❑ **Consequences:**
 - Frustration
 - Irritation
 - Loss of productivity
 - Elongated cycle times



Dysfunctional solution:

Just throw hardware at the problem

Yogi Schulz

Corvelle

Achieving Superior Application Performance

Superior Application Performance



- Server
- Database
- Disk
- Data
- Network
- Applications
- Workstation
- Training
- Operating System
- Thin-client computing architecture
- Operations



Holistic approach will deliver the most benefits

Yogi Schulz

Corvelle

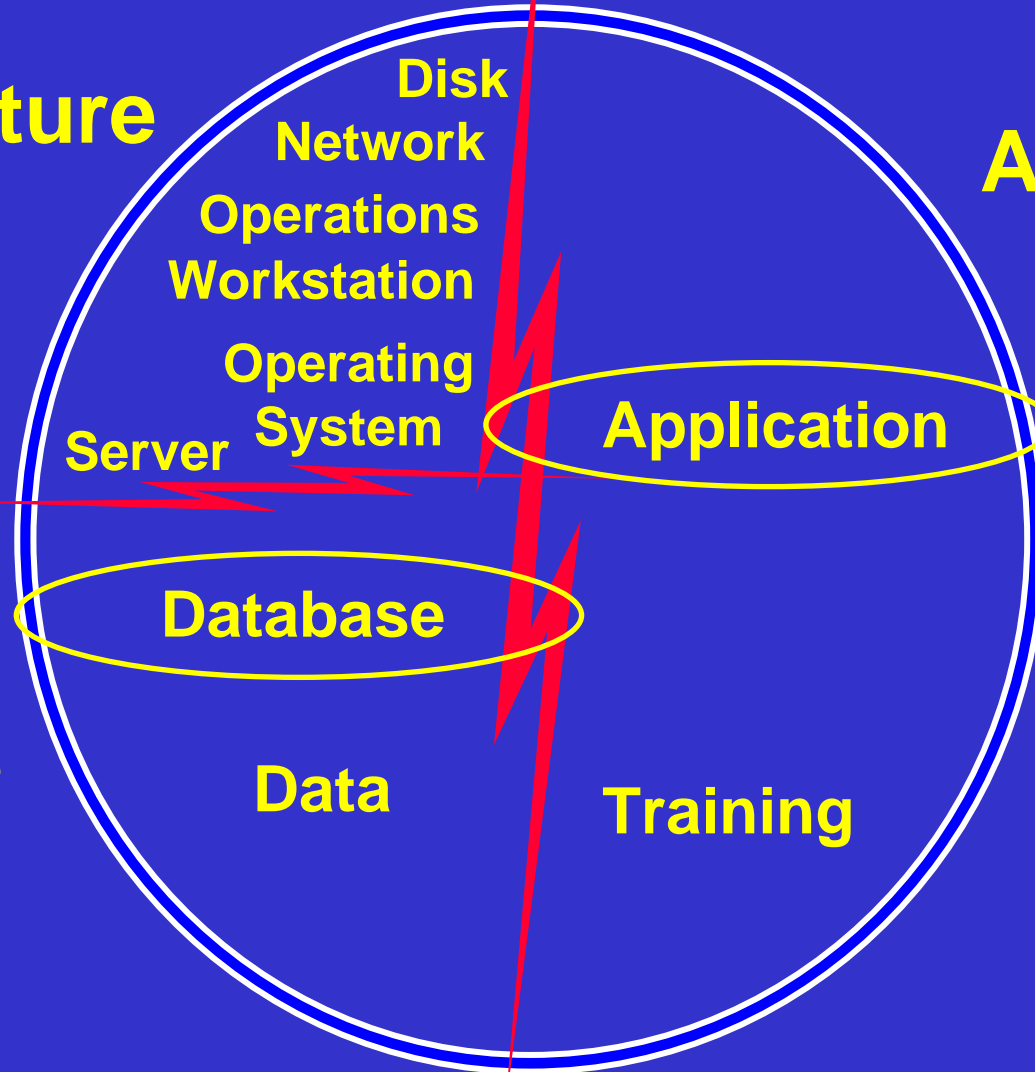
Relative Gains from Performance Improvement Opportunities

Superior Application Performance



Infrastructure

Application



Yogi Schulz



Superior Performance Server

Superior Application Performance



- ❑ **Symmetrical Multi-Processor (SMP)**
 - Not single processor
- ❑ **64-bit addressing**
 - Not 32-bit
- ❑ **Large memory**
 - Not small memory; no page faults
- ❑ **Shared memory**
 - Not processor-dedicated memory
- ❑ **Multiple I/O adapters**
 - Not too few

Yogi Schulz



Superior Performance Disk - 1

Superior Application Performance



❑ **Disk:**

- Acquire sufficient disk capacity
- De-fragment regularly

❑ **SAN/NAS:**

- Use high performance read-ahead/caching disks; not CAS
- Implement FibreChannel connection; not Ethernet

❑ **Configure disks to:**

- Maximize access paths to data
- Minimize contention for access to individual devices

Superior Performance Disk - 2

Superior Application Performance



❑ Implement RAID 5

- Maximize disk availability

❑ Consider Hierarchical Storage Management (HSM) implementation

- SAN/NAS disk
- Cheap disk
- Near-line tape
- Storage rack tape

Increasing performance

Decreasing cost

❑ Separate NFS and database datasets

- Reduce contention arising from different access patterns

Yogi Schulz

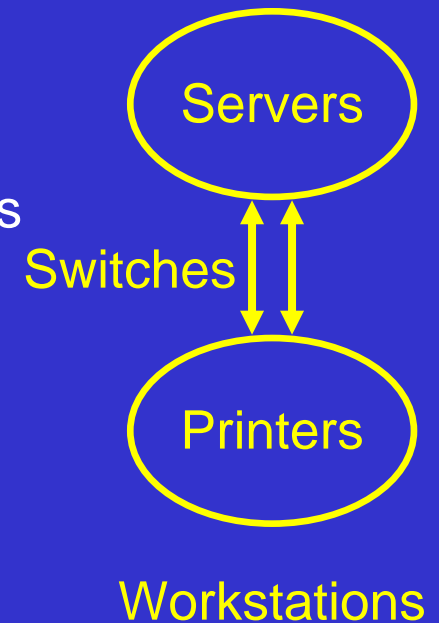


Superior Performance Network

Superior Application Performance



- ❑ **Ethernet to the workstations:**
 - Gigabit; not 100 mbit; not 10 mbit
 - Switched; not shared; not hubs
- ❑ **Multiple paths in the internal network**
 - Minimize contention on key network segments
- ❑ **Switches on server network segment**
 - Not hubs
- ❑ **Isolate workstations from servers**
 - No workstation traffic through server NICs
- ❑ **Locate printers with workstations**
 - No printing traffic through server NICs



Superior Performance Workstation

Superior Application Performance



- ❑ **Install sufficient memory**
 - Avoid paging to improve performance
- ❑ **Install sufficient graphics power & memory**
 - Minimize re-draw times
- ❑ **Install large monitor(s)**
 - Maximize physical volume that can be displayed
- ❑ **Cleanup extraneous data on disk**
 - Maximize I/O performance
- ❑ **Consider multi-processor**
 - Improve application performance

Superior Performance Operating System

Superior Application Performance



- ❑ **Turn on asynchronous I/O**
 - Reduce I/O wait

- ❑ **Allocate sufficient priority and memory to database tasks**
 - Swapping out a database task dramatically reduces performance

- ❑ **Allocate sufficient paging disk on internal drives**
 - Not on network or SAN/NAS drives

Superior Performance Operations

Superior Application Performance



- ❑ **Monitor computing infrastructure proactively**
 - Identify/correct performance problems
- ❑ **Manage with change control discipline**
 - Maximize robustness/availability
- ❑ **Define/test database recovery strategy**
- ❑ **Define/test:**
 - Disaster recovery
 - Business continuity plan

Yogi Schulz



❑ **Conduct capacity planning**

Superior Performance Database

Superior Application Performance



- ❑ **Tune databases for performance**
 - See reference slides for details
- ❑ **Monitor database operation continuously**
 - As data volumes grow, operational issues emerge
- ❑ **Consider what data to store in databases**
 - Are the benefits of manageability, multi-application access, security, auditability attractive?
- ❑ **Consider what data to store in a file system**
 - Do the benefits of faster access outweigh database benefits?

Yogi Schulz



Superior Performance Data

Superior Application Performance



- ❑ **Clean up data:**
 - Replace nulls with correct data
 - Correct edit exceptions
 - Provide foreign key values
 - Standardize reference values

- ❑ **Structure data for performance:**
 - Implement PPDM
 - Create flattened exports

- ❑ **Eliminate non-project data from datasets;
examples include:**
 - Grid, culture, Production time series, Well logs

Superior Performance Applications - 1

Superior Application Performance



- ❑ **Follow software package vendor recommendations for installation and configuration**
 - Populate properties files
 - Acquire more than the minimum hardware configuration
- ❑ **Actively monitor software package execution**
 - Monitor memory and I/O activity
 - Review log files for errors
- ❑ **Provide software package vendor with feedback on performance**
 - Be specific about function exhibiting inadequate performance

Superior Performance Applications - 2

Superior Application Performance



- ❑ **Upgrade to current version/install bug fixes**
 - Maximize your benefits from vendor investments

- ❑ **Participate in software package vendor end-user forums**
 - Collaborate with your peers

- ❑ **Encourage/nag software package vendor to improve performance**
 - Provide your supporting experience data

Superior Performance Training

Superior Application Performance



- ❑ **Describe strategies for minimizing execution of high resource consumption functions**
 - Follow a defined workflow
 - Turn off auto re-calc/auto re-fresh/auto re-draw

- ❑ **Encourage/reward reporting of performance issues/experiences**
 - Data issues
 - Network contention
 - Server/workstation hangs
 - Puzzling application behavior

Thin-client Computing Architecture

Superior Application Performance



- ❑ **Move application processing from workstations to an application server**
 - Install Citrix MetaFrame, Microsoft Terminal Server or SGI VAN
 - Deliver applications through a Web browser

- ❑ **Send mostly pixels and a few commands across the network**
 - Send no data

Benefits:

- Dramatically reduce network traffic
- Improve end-user performance
- Significantly reduce price of workstations
- Improve computing infrastructure utilization

Yogi Schulz



Conclusions

Superior Application Performance



Superior application performance is achieved through implementing a large number of improvements in:

- Infrastructure
- Database
- Application

Holistic approach will deliver the most benefits

Yogi Schulz



Recommendations

Superior Application Performance



- ❑ **Review all aspects of your computing environment**
- ❑ **Make prudent investments to achieve superior application performance**



Operate a happier, more productive exploration computing environment

Yogi Schulz

Corvelle

Questions & Answers

Superior Application Performance



Dogbert
will explain it
to you

Yogi Schulz



Achieving Superior Application Performance amidst Skyrocketing Data Volumes

Superior Application Performance



1800, 250 - 6th Ave. S.W.
Calgary, Alberta Canada T2P 3H7
Phone/Fax: (403) 249-5255
E-mail: YogiSchulz@corvelle.com
Web: www.corvelle.com

President of Corvelle Management Consultants

Information technology related management consulting

Project management and systems development

Computing Canada & Calgary Herald columnist

PPDM Association board member

Yogi Schulz



Superior Performance Tuning Database - 1

Superior Application Performance



- ❑ **Select a larger block size**
 - Retrieve more data on each I/O request

- ❑ **Set multi-block read count to a multiple of the block size**
 - Retrieve more data on each I/O request

- ❑ **Allocate sufficient SGA size**
 - Bring as much data into memory as possible

- ❑ **Execute Commits frequently**
 - Minimize size and number of rollback segments

- ❑ **Monitor application plans**
 - Un-optimized SQL produces poor performance

Superior Performance Tuning Database - 2

Superior Application Performance



- ❑ **Monitor index usage**
 - Poor index choices create poor performance
- ❑ **Initiate sufficient database writers**
 - Minimize waits on writes
- ❑ **Dedicate disk volume to redo log**
 - Minimize disk access contention on writes
- ❑ **Set Initrans, pctfree high enough**
 - Minimize allocation of secondary extents
- ❑ **Avoid replication/standby database**
 - Minimize interfering I/O activity

Bibliography - 1

Superior Application Performance



- ❑ **Brocade "What's Your SAN Plan Guide"**
 - www.brocade.com/products/WYSP/guide_welcome.jsp
- ❑ **Cisco - Network Management System: Best Practices White Paper**
 - www.bitpipe.com/detail/RES/1070907380_696.html?src=googleb
- ❑ **Database Performance Tuning on AIX**
 - publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg245511.html?Open
- ❑ **DataVera – data quality enhancement tools**
 - www.datavera.com/datavera
- ❑ **Fibre Channel Industry Association (FCIA)**
 - www.fibrechannel.org
- ❑ **Fibre Channel Overview**
 - hsi.web.cern.ch/HSI/fcs/spec/overview.htm
- ❑ **How to Improve Network Performance**
 - www.newchanneltech.com/whitepapers/3com_net_prf.pdf
- ❑ **HP Performance Tuning Framework**
 - www.hp.com/workstations/software/framework

Yogi Schulz



Bibliography - 2

Superior Application Performance



- ❑ **IBM eServer pSeries and IBM RS/6000 Performance Report**
 - www-1.ibm.com/servers/eserver/pseries/hardware/system_perf.pdf
- ❑ **Improve B2B Application Performance with Gigabit Server Network Adapter**
 - www1.us.dell.com/content/topics/global.aspx/power/en/ps1q01_broadcom?c=us&cs=555&l=en&s=biz
- ❑ **Intel® Itanium® 2 processor**
 - www.intel.com/products/server/processors/server/itanium2/
- ❑ **NetScout - nGenius® Performance Management System**
 - www.netscout.com
- ❑ **RAID 5: Independent Data disks with distributed parity blocks**
 - www.acnc.com/04_01_05.html

Bibliography - 3

Superior Application Performance



- ❑ **SGI Altix Returns to SC2003 After Breakout First Year**
 - Acclaimed Shared-Memory Architecture, 64-bit Linux, and Itanium 2 Combine to Create World's Fastest Linux Supercomputer
 - www.sgi.com/newsroom/press_releases/2003/november/altix_sc2003.html

- ❑ **SUN Performance**
 - www.sun.com/sun-on-net/performance.html

- ❑ **Three Ways to Improve Network Performance**
 - Written by Maria Piech - Published in Unisphere Magazine, October 2002
 - www.somix.com/files/02-10-01_3ways-improve_netperf.html

- ❑ **WebLogic Server Performance Tuning Guide**
 - www.weblogic.com/docs51/admindocs/tuning.html