The Retek Data Warehouse Performance Benchmark

September, 1999

Table of Contents

Executive Summary

Data Warehousing for the Retail industry involves very large volumes of data, high user populations and analytical complexity putting significant strain on the entire Decision Support System (hardware, software and applications). To meet the demands of high growth and complex business requirements, it is critical for the system to process transactions and execute reports in an acceptable timeframe.

Retek Information Systems performed a Test-to-Scale Performance Benchmark of their Retek DataWarehouse (RDW) product in August 1999 at the HP Performance Center in Cupertino, California. The purpose of the Benchmark was two-fold:

- Determine the performance and scalability of RDW using massive data volumes and analytical complexity that is typical of Retek's clients
- Produce a white paper describing the details of the Benchmark that can be subsequently used for Marketing and Capacity Planning

The benchmark was performed using commercially available hardware and software. A 1.1 Terabyte Oracle 8i database was built using dimensional data that was representative of a large retailer. Under simulated but real-world conditions, daily sales and inventory transactions were executed against the RDW Application. To simulate a large user population, thousands of queries with varying degrees of complexity & frequency were run against the test database.

The benchmark also tested different processor configurations under varying transaction loads for the batch environment. These tests were performed to test scalability and provide some insight into capacity planning for future Retek customers.

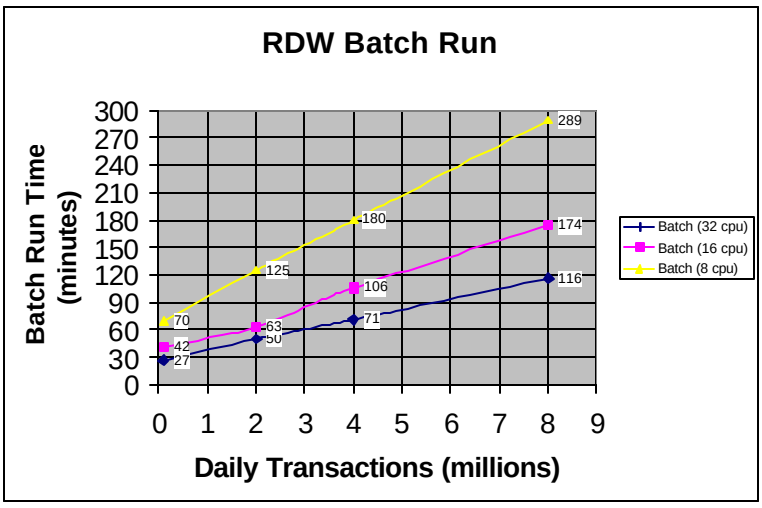

The results of the benchmark proved the scalability of the entire system. In summary, it was possible to process 2 million transactions in less than a 1 hour batch window; for the front-end reports, it was possible to achieve a report throughput rate of 50 reports per minute, allowing concurrent support of hundreds or thousands of users.

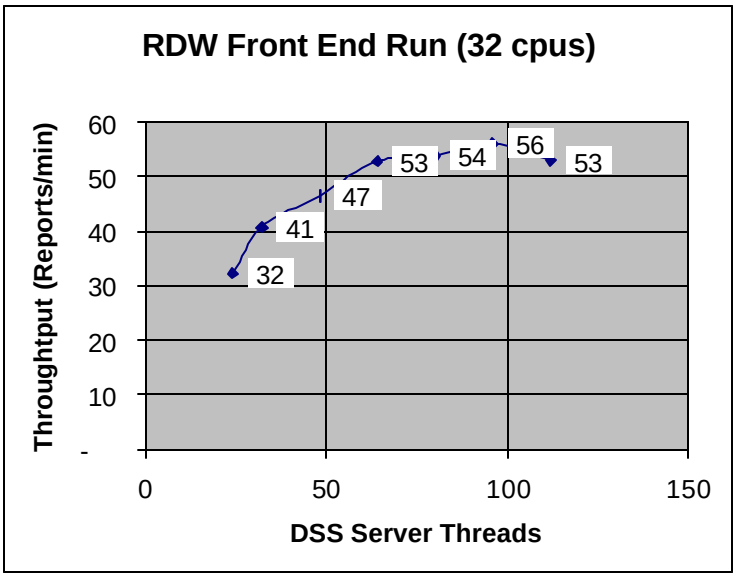

The technical environment consisted of the following configuration:

- Hewlett Packard V2500 Server, 32 cpus, 16 GB memory, UX-11
- Hewlett Packard FC60 Storage Array, 3 Terabytes physical disk
- (2) Hewlett Packard NetServers, NT 4.0
- Oracle 8.1.5.0
- Retek RDW 3.0
- MicroStrategy DSS Server 5.6
- MicroStrategy DSS Broadcaster 5.5

Finally, the tests have demonstrated that the hardware and software platform will scale and meet the demands of large retailer with massive data volumes, analytical complexity and large user populations.

The performance improvements that resulted from the benchmark study will be engineered into RDW product code and distributed to Retek's client base.

Objective

The primary purpose of the Retek Data Warehouse (RDW) Benchmark was to determine the performance and scalability of the RDW application using massive data volumes, very large batch transaction rates, high user populations and analytical complexity that is typical of Retek's clients.

As part of conducting the large-scale benchmark, many performance tuning opportunities will be identified. The tuning opportunities will be evaluated and incorporated into the base RDW application.

Another objective of the benchmark was to make the results available to Retek clients, partners, and prospects. This information will demonstrate to Retek's business partners that the RDW application can scale to massive data volumes and high numbers of concurrent users. Further, the detail information can be used to conduct some high-level capacity planning and set performance expectations for a small, medium and large retailer.

Finally, Retek RDW marketing collateral will be published describing the highlights of the benchmark. This document will advertise the partners used in the benchmark and will be used with potential customers.

Goals

The goals of the benchmark are to perform a simulation of the various functional areas of RDW for a typical retail customer. These performance goals can be quantified as:

- 1 TB RDW Database Environment. Creation and building of a 1+ Terabyte RDW environment to perform a simulation of a large retailer and over 1 year of historical data.
- Batch Load of 2 million sales, 2.2 million inventory transactions in a 4 hour window. Determination of an optimal batch window (better than 4 hours) that provides maximum throughput in the minimum possible time.
- Measure front-end report performance using varying report complexity/frequency and high degree of concurrent users. Retek will focus on optimizing the performance of a realistic mix of reports in the following ranges of complexity: 65% simple, 25% medium, and 10% complex. Complexity is determined by the number of queries and temporary tables created by the report SQL and by the volume of data and dimensional relationships involved.
- Determination of the number of concurrent users that can be supported by the RDW application. Retek will focus on emulating the realistic behavior of a user population of 750 users.
- Comparison of Batch Load times under different CPU configurations. Determine batch scalability under hardware configurations of 8 cpu, 16 cpu and 32 cpu.
- Performance Measurement of Batch Load time under varying transaction load. Batch load times of 2million, 4 million and 8 million sales transactions will be conducted. The purpose of this goal is to prove scalability with increasing transactional volume.

Technical Environment

Technical Architecture

The technical environment consisted of the following components:

- Hewlett-Packard V2500, 32 cpus, 16 GB memory (Benchmark Server)
- Hewlett-Packard K570, 6 cpus (Development Server)
- HP UX-11 Operating System
- HP FC60 Storage Array, 3 TB physical disk
- Oracle 8.1.5.0 (32 bit)
- Retek RDW 3.0
- HP NetServers, 2 cpus (2), NT 4.0
- MicroStrategy DSS Server 5.6
- MicroStrategy DSS Broadcaster 5.5

A high-level technical schematic is presented below:

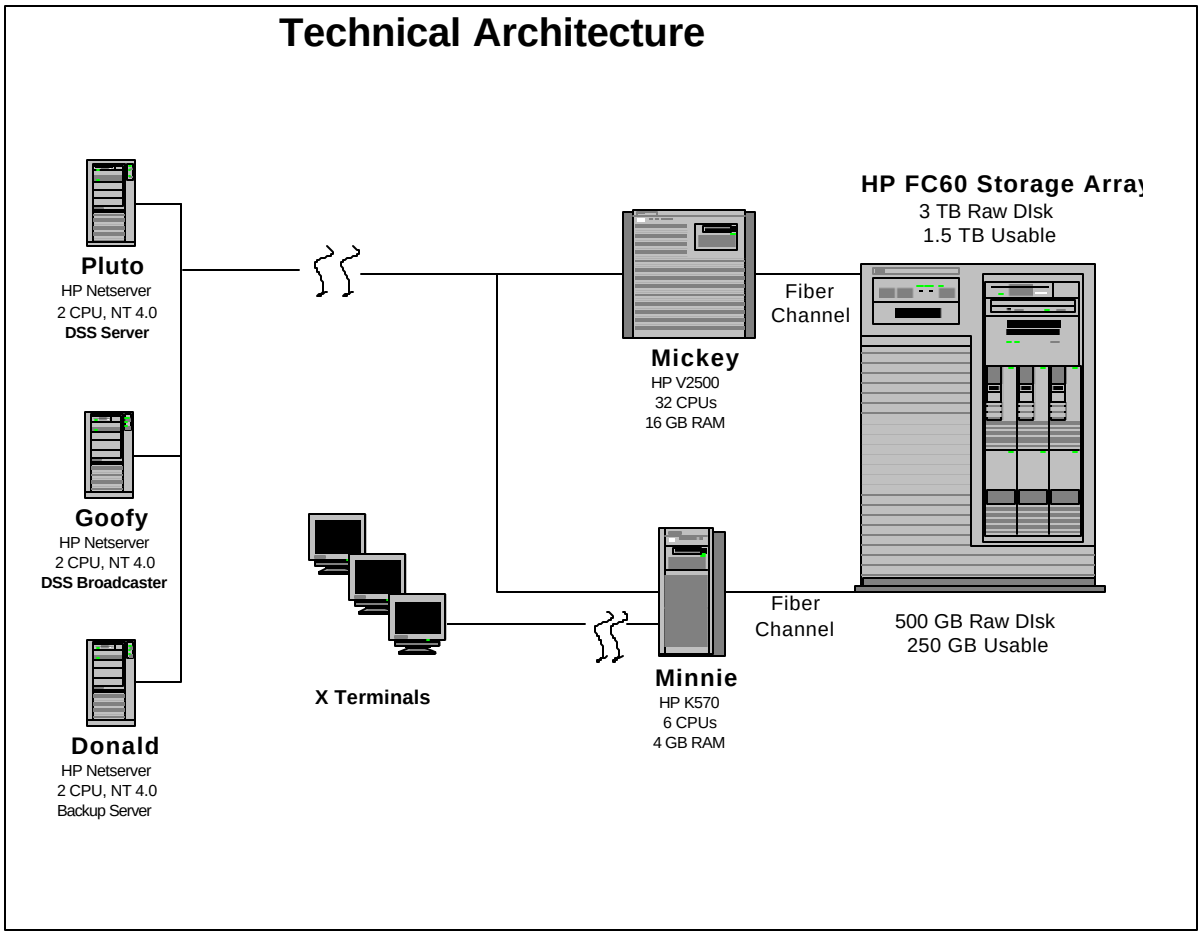

Unix Configuration

The following Unix configuration & software features were used for the Benchmark

Kernel Parameters

The following UX-11 parameters were changed:

Parameter	Value
dbc_max_pct	10
dbc_min_pct	10
maxfiles_lim	2048
maxfiles	2048
STRMSGSZ	65535
num_tachyon_adapters	5
max_fcp_reqs	512
nstrpty	60
maxdsiz	0x7b03a000
maxswapchunks	12288
maxuprc	2048
maxusers	400
maxvgs	32
nfile	8192
nflocks	1024
npty	256
semmap	512
semmni	1024
semmns	2048
semmnu	1024
semume	512
shmmax	0x40000000
shmmni	512
shmseg	512
swapmem_on	0
unlockable_mem	4096

Asynchronous I/O

Asynchronous I/O on HPUX is supported using raw (character) devices. The disk driver for Async I/O must be configured using the SAM tool. Async I/O allows multiple reads and writes to proceed in parallel, thus reducing I/O wait time and improving performance of the Oracle Server.

Async I/O was enabled for the entire benchmark period. The device `/dev/async` was owned by '`oracle`'.

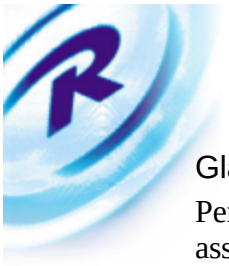

GlancePlus Pack

Performance monitoring, data collection and reporting are critical components that can assist the benchmark process. The tools allowed detail monitoring of the system and were crucial in identifying and solving tuning issues. It also enabled the team to refine and optimize the RDW batch load process.

The 'glance' tool allowed system monitoring while the benchmark tests were running. In addition, the 'scope' utility was also running on the system, collecting performance data on a continuous basis. Finally, 'pv' or Perf View was used to analyze and report system behavior after the benchmark tests were executed.

Disk Architecture

Disk configuration is a critical component of data warehouse performance. The objective was to get maximum performance, scalability, availability and flexibility of the disk storage solution. For the benchmark, the HP FiberChannel FC60 Storage Array was used. The disk was configured in Raid 0/1 format (mirror/stripe) for maximum performance. The array was configured with the following specifications:

Specification	Value
Connectivity	Fiber Channel (2)
Raid Format	Raid-01 (stripe/mirror)
Spindle Size	18 GB
Spindles/LUN	5 (+5 for mirroring)
Size per LUN	86 GB
Number of LUNs	17
Volume Group Size	1.4 TB
Logical Volume Size (raw)	6 GB
Stripe Element Size	64k
Asynchronous I/O	true

The disk connectivity to the host is shown below:

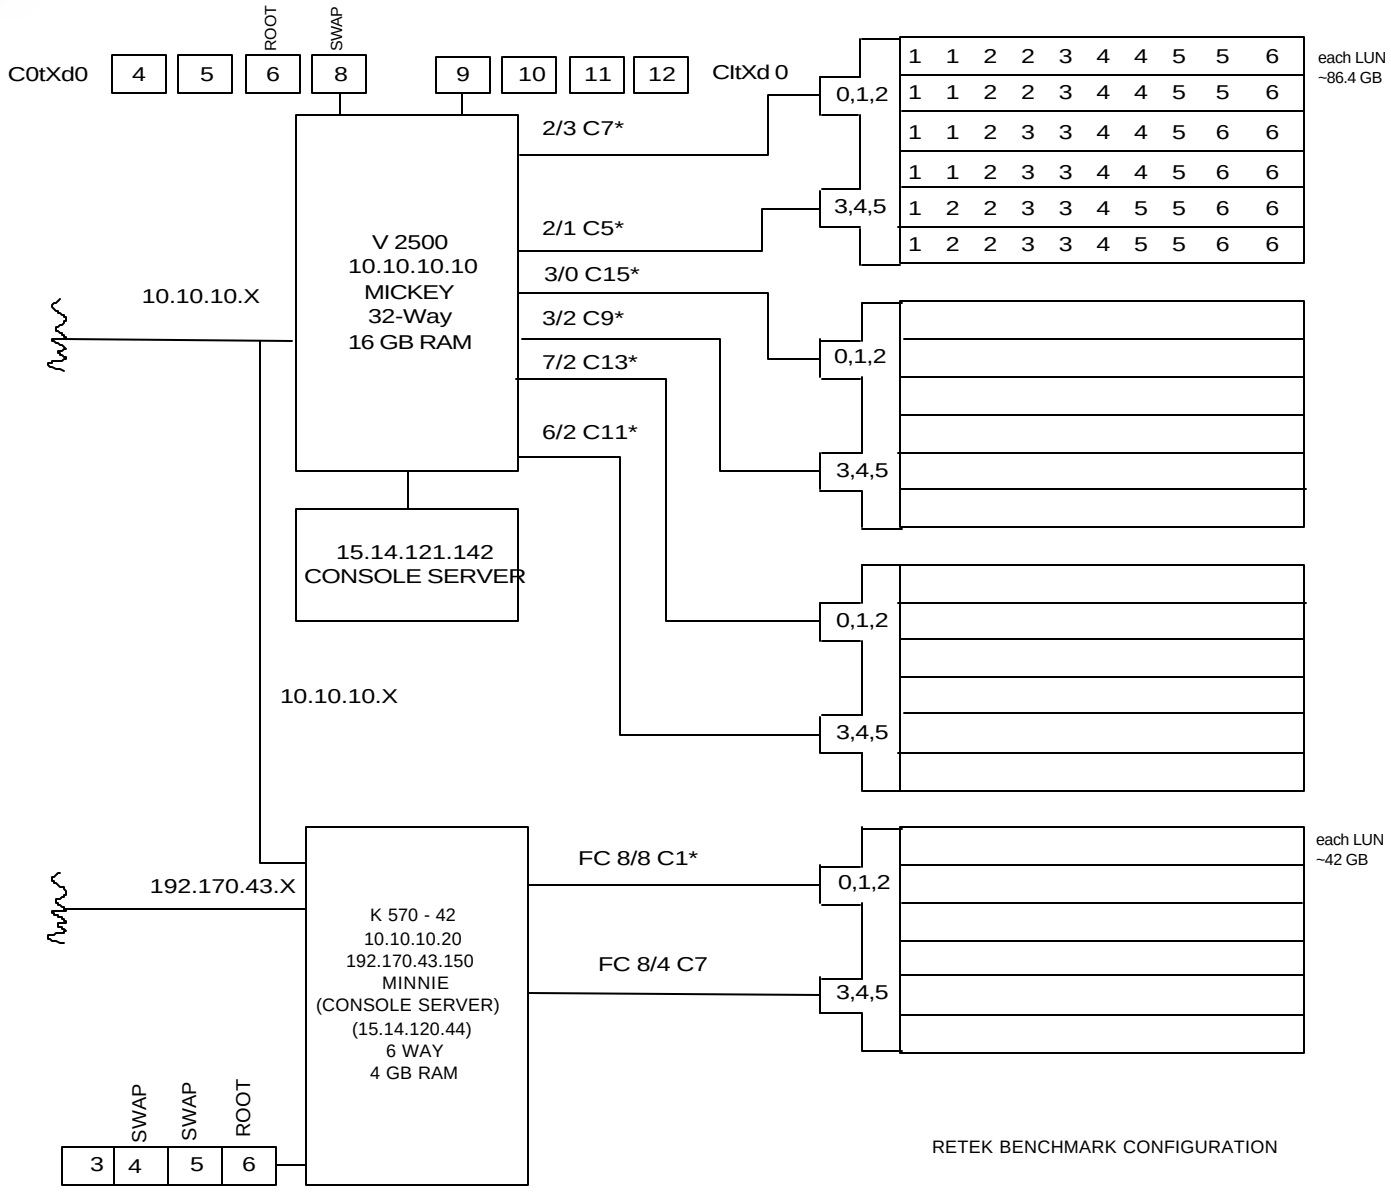

Volume Groups

Based on several discussions with HP engineering staff, we decided to use a single but large volume group for the main Oracle datafiles. Since each LUN was striped (with 5 disks) at the hardware level and the logical volume would provide further striping at the O/S level, we were attaining even load balancing *across and down* each LUN. This configuration was based on HP's internal testing for TPC-D benchmarks. The VG configuration is shown below:

```
--- Volume groups ---
VG Name                     /dev/vg00
VG Write Access             read/write
VG Status                   available
Max LV                      255
Cur LV                      8
Open LV                     8
Max PV                      16
Cur PV                      2
Act PV                      2
Max PE per PV               4350
VGDA                        4
PE Size (Mbytes)            4
Total PE                    8680
Alloc PE                    5196
Free PE                     3484
Total PVG                   0
Total Spare PVs             0
Total Spare PVs in use      0

VG Name                     /dev/vg01
VG Write Access             read/write
VG Status                   available
Max LV                      255
Cur LV                      4
Open LV                     4
Max PV                      16
Cur PV                      2
Act PV                      2
Max PE per PV               4342
VGDA                        4
PE Size (Mbytes)            4
Total PE                    8680
Alloc PE                    8192
Free PE                     488
Total PVG                   0
Total Spare PVs             0
Total Spare PVs in use      0

VG Name                     /dev/vg03
VG Write Access             read/write
VG Status                   available
Max LV                      255
Cur LV                      14
Open LV                     14
Max PV                      1
Cur PV                      1
```

```
Act PV                      1
Max PE per PV               21658
VGDA                        2
PE Size (Mbytes)            4
Total PE                    21655
Alloc PE                    21042
Free PE                     613
Total PVG                   0
Total Spare PVs             0
Total Spare PVs in use      0

VG Name                     /dev/vg04
VG Write Access             read/write
VG Status                   available
Max LV                      255
Cur LV                      243
Open LV                     243
Max PV                      17
Cur PV                      17
Act PV                      17
Max PE per PV               21658
VGDA                        34
PE Size (Mbytes)            4
Total PE                    368118
Alloc PE                    367659
Free PE                     459
Total PVG                   0
Total Spare PVs             0
Total Spare PVs in use      0
```

Logical Volumes

Using the HP hardware/software package, the maximum number of logical volumes that can be created from a single volume group is 256. Due to this limitation, the LV size was determined to be 6 GB each, allowing us to grow to 1.5 TB, certainly exceeding the available disk. Each LV was striped over 17 LUNs with a 64k stripe element size. A sample LV is shown below:

```
--- Logical volumes ---
LV Name                     /dev/vg04/spare99
VG Name                     /dev/vg04
LV Permission               read/write
LV Status                   available/syncd
Mirror copies               0
Consistency Recovery        MWC
Schedule                    striped
LV Size (Mbytes)            6052
Current LE                  1513
Allocated PE                1513
Stripes                     17
Stripe Size (Kbytes)        64
Bad block                   on
Allocation                  strict
IO Timeout (Seconds)        default
```

Oracle Database Configuration

For the Benchmark, the Oracle RDBMS used was Oracle 8.1.5.0 in 32-bit mode. The 64-bit version was not available at the time of the Performance Benchmark. The RDBMS was used with the Parallel Query and Partitioning options. The entire database was build on striped raw partitions to get better performance.

More information on the Oracle database is provided in the Database Details section.

MicroStrategy Configuration

MicroStrategy products DSS Broadcaster 5.5 and DSS Server 5.6 were used to generate and execute front-end reports. Server and Broadcaster were installed and configured on 2 dual CPU NT machines and were connected via a 100 BaseT Ethernet network.

The following .DSS file was used for Broadcaster

```
****************************************************
****
*         Initialization file for the DSS Agent
*
****************************************************
****

[applications]
Excel=D:\PROGRAM FILES\MICROSOFT
OFFICE\OFFICE\EXCEL.EXE
Word=D:\PROGRAM FILES\MICROSOFT
OFFICE\OFFICE\WINWORD.EXE
Powerpoint=D:\PROGRAM FILES\MICROSOFT
OFFICE\OFFICE\POWERPNT.EXE
RPTpath=..\reports\
MSMail=0
MSAccess=D:\PROGRAM FILES\MICROSOFT
OFFICE\OFFICE\MSACCESS.EXE

[Powerpoint]
classname=PPApplicationClass

[databases]
dssagent=ODBC;DSN=MICKEYMETA;DBQ=rdw1;UID=RDW30_DSS_SY
S;PWD=oggle
dssagenttype=1
warehouse=ODBC;DSN=MICKEY;DBQ=rdw1;UID=RDW30_DSS_SYS;P
WD=oggle
warehousetype=1

DateFormat="dd-mmm-yy"
DataMining=0
MetricJoinType=
SubQueryType=2
StarJoin=0
DistributedDataBase=0
WhereClauseOrder=1
CreateTempView=0
TempTablePrefix="ET_"
TempTableSpace=""
TempTableStarSchema=1
DeleteTempTablePrompt=1
DropTempTableLater=0
CreateTempView=0
ODBCWorkAround=0
NoCountCompound=0
IssueCartesianWarning=1
DrivingTable=0
ArithmeticCheck=1
NullFunction=NVL(???,0)
NullIndicator= "NA","??"
ZeroFunction=""
ConvertLength=""
ServerDataMart=
ServerDataMartType=
TableStatement="ANALYZE TABLE ??? ESTIMATE STATISTICS"
TempTableAnalyze=0
TempTableOption="NOLOGGING"
TempTableType=0
FilteredTransformation=1

[system]
ApplicationTitle=Retek Data Warehouse v3.0 Volume Test
ApplicationIcon=\AGENT.ICO
IntersectString="^"
UnionString="+"
NotString="-"
CtGridGovern=1
MetricParser=OFF
EISAS=3
CHAR171=171
CHAR187=187
CHAR160=160
MailSystem=0
Metadata Prefix=
System User=RDW30_DSS_SYS
Group User=
SplashForm40=
SplashLogo40=

[Caching]
UseLocalCache=0
LocalCachePath=C:\DSSTOOLS\DSSAGENT\..\cache\
UseSharedCache=0
SharedCachePath=C:\DSSTOOLS\DSSAGENT\..\cache\
ScheduleType=1
Units=0
Period=1
StartDate=01/01/1997 11:00
MDCacheLocation=C:\DSSTOOLS\DSSAGENT\..\cache\
LoadMDCache=0
WriteMDCache=0
MetricCacheLocation=C:\DSSTOOLS\DSSAGENT\..\cache\
LoadMetricCache=0

[preferences]
UserID=RDW30_DSS_SYS
PWD=rdw
project=RDW 3.0 Volume Project
projectname=RDW 3.0 Volume Project on Mickey
ReportTop=4410
ReportLeft=0
ReportWidth=7680
ReportHeight=4755
FilterTop=4410
FilterLeft=7680
FilterWidth=7680
FilterHeight=4710
AgentTop=6915
AgentLeft=0
AgentWidth=2400
AgentHeight=360
CriteriaTop=0
CriteriaLeft=7680
CriteriaWidth=7680
CriteriaHeight=4410
RObjectTop=0
RObjectLeft=0
RObjectWidth=7680
RObjectHeight=4410
AgentType=0
ScreenSettings=1
Language=English
AgentList=0
FilterList=1
ReportList=1
CriteriaList=1
RObjectList=1
```

```
ReportWindowState=0
FilterWindowState=0
AgentWindowState=1
CriteriaWindowState=0
RObjectWindowState=0
OpenAgent=0
NewExcelSheet=1
DefEmptyFilter=1
GridAdjustment=0
RCEThreshold=60
OpenHierarchy=1
EngineProps=0
AgentAnimation=0
ReportCostEstimator=1
GridFormat=1
AutoSizeReports=0
DefaultOperator=0
ComponentsFolders=1
ReadPropForAllMet=0
TopRowHeader=1
ProjectDesc=Volume Test Project of RDW 3.0
ScreenX=15360
ScreenY=11520
OverrideDrillDown=True
ODBCQueryTimeOut=0
AutoMaximize=0
FilterBuilder=1
MetricEditor=1
CustomGroupEditor=1
HelpPath=dssagent.hlp
CacheAutoprompts=0
ExpireCaches=1
AllowDrillWithNoSelection=0
ShowAllDetails=0
DatamartCustomTemplate=0
DatamartPrePostSQL=0

[Filter Details]
FontName=Arial
FontSize=8.25
FontStyle=000000
ShowFltrDet=1
Forecolor=0
Backcolor=12632256

[DSS Objects]
FontName=MS Sans Serif
FontSize=8.25

[Interface]
CriteriaWindow=1
FiltersWindow=1
ReportsWindow=1
AgentsWindow=1
ReportObjectsWindow=1
DefaultHierarchies=1
Drilling=2
DrillGoverning=1
Sorting=1
Printing=1
ExporttoExcel=1
ExporttoWord=1
ExporttoMail=0
Scrolling=1
Paging=1
Help=1
Trash=0
FilterOperator=1
Lightning=1
GraphMode=1
GridMode=1
MapMode=0
AlertMode=1
EngineMode=1
DesignMode=1
ShowStatusBar=1
ShowPrompt=1
ViewMenu=1
ShowConsolidations=0
ShowQualities=1
NewObjectFolder=1
NewObject=1
ReportWriterMode=0
Refresh=1
ViewMode=0
SendObject=1
ObjectSelector=0
DatamartWizard=0
ReportWizard=0

[tools]
HoldQueryExe=0
** SQLLog commented  to improve Broadcaster **
** sqllog=C:\TEMP\load.TXT **

[DSS Server]
ServerMode=1
ServerName=MINNIE
SwitchTiers=1
Protocol=ncacn_ip_tcp
Endpoint=1620
Address=10.10.10.30
DMString=

[System Catalog]
Catalog SQL=select table_name, column_name, data_type,
   data_length from all_tab_columns where owner =
'RDW30_ADM_DM'
Catalog SQL2=

[User Messages]
MsgString=1110111011010110

[Objects Toolbar]
DockRank=0
DockRankSeq=0
DockStatus=0
FloatingWidthInBtns=5
FloatingLeft=0
FloatingTop=0

[Actions ToolBar]
DockRank=0
DockRankSeq=1
DockStatus=0
FloatingWidthInBtns=5
FloatingLeft=0
FloatingTop=0

[Navigation ToolBar]
DockRank=0
DockRankSeq=2
DockStatus=0
FloatingWidthInBtns=4
FloatingLeft=0
FloatingTop=0

[Objects]
Filter=aa PB Filters (TPOP) Review
Template=aa PB Packet (Renewed Interest)
Agent=No object selected
Report=No object selected
FilterFolder=
TemplateFolder=
AgentFolder=No Action
ReportFolder=No Action
[DBExport]
UseCustomAlias=0
AppendToTable=0
UseCompoundMetric=1
```

Report Query Profile

Approach

For the RDW Performance Benchmark a set of standard reports were established. This set of reports or 'Query Profile' was a representative sample of the actual reports available within RDW. These reports also represented a realistic mix of reports that would be run in a production data warehouse environment.

The candidate reports within RDW were ranked according to three measures: *Complexity, Frequency, and Importance.* The ranking within each measure were assigned ascending according to intensity within that measure.

Complexity

Each report got a Complexity ranking where 1=Simple, 2=Medium, and 3=Complex. Complexity was defined based on the SQL that the report generated. Simple queries performed a single `select` against the warehouse. Medium queries generated 2-10 `selects` and up to 10 temp tables. High Complexity reports generated more than 10 `selects` and/or more that 10 temp tables.

Frequency

Frequency ranking was assigned based on the expected frequency of execution of the report. The specific rankings was defined as 1=Occasionally, 2=Weekly, 3=Daily.

Importance

Importance ranking was assigned based on the relative importance of the report, or value to the business. The specific rankings were defined as 1=Low, 2=Medium, 3=High.

Query Profile

The following table represents the percentage breakdowns used to define a 'realistic' mix of reports across Complexity and Frequency:

	Occasionally	Weekly	Daily	GT	%	Target
Low		16	20	36	**60%**	60%
Medium	4	7	7	18	**30%**	30%
High	2	4		6	**10%**	10%
GT	6	27	27	60		
%	**10%**	**45%**	**45%**		**100%**	
Target	10%	45%	45%			

Service Volumes

Since a large number of reports were required for the benchmark, DSS Broadcaster was used to generate service volumes that represented the Frequency/Complexity (F/C) as specified above. Additionally, weighting for Daily (D), Weekly (W), or Occasional (O) usage was applied to the Complexity of Low (L), Medium (M), High (H). The following table represents the report suite representing a total of 2,970 individual reports:

F/C Scale	Total	Percentage	Service Volume	Extension	Suite
LD	20	33.3%	360	5	1800
LW	16	26.7%	288	1	288
MD	7	11.7%	126	5	630
MW	7	11.7%	126	1	126
MO	4	6.7%	72	0.5	36
HW	4	6.7%	72	1	72
HO	2	3.3%	36	0.5	18
Totals	**60**	**100.0%**	**1080**		**2970**

The following table represents a Report Makeup of the Query Profile:

PB Report Name	Complex	Freq	Dimension	Metrics
Dept,Net Transfer Value(LD)	1-Low	3-Daily	Division/Department	Net Transfer Value
Org,Sales Value by Type(LD)	1-Low	3-Daily	Organization	Sales Value
Prod,Sales Value by Type(LD)	1-Low	3-Daily	Product	Sales Value
Time,Sales Value by Type(LD)	1-Low	3-Daily	Time	Sales Value
Loc,SKU,Weeks,VAT(LW)	1-Low	2-Weekly	Location/SKU by Month/Week	Sales Value/VAT
Loc,Subclass,Weeks,VAT(LW)	1-Low	2-Weekly	Location/Subclass by Month/Week	Sales Value/VAT
SKU,Week,Price(LW)	1-Low	2-Weekly		
Org,EOH Value By Type(MD)	2-Med	3-Daily	Organization	EOH Value
Org,Markdown Value By Type(MD)	2-Med	3-Daily	Organization	Markdown
Prod,EOH Value By Type(MD)	2-Med	3-Daily	Product	EOH Value
Prod,Markdown Value By Type(MD)	2-Med	3-Daily	Product	Markdown
Time,EOH Value By Type(MD)	2-Med	3-Daily	Time	EOH Value
Time,Markdown Value By Type(MD)	2-Med	3-Daily	Time	Markdown
Class, Sales, Profit(MW)	2-Med	2-Weekly	Division/Department/Class	Sales,Promo/Profit
Org,Sales,Profit Value TY vs LY(MW)	2-Med	2-Weekly	Organization	Sales/Profit
Prod,Sales,Profit Value TY vs LY(MW)	2-Med	2-Weekly	Product	Sales/Profit
Time,Sales,Profit Value TY vs LY(MW)	2-Med	2-Weekly	Time	Sales/Profit
Loc,Division,LW,Sales Value vs Chain(MO)	2-Med	1-Occasionally	Division/Location	Sales Value
Region,Division,LW,Sls Value vs	2-Med	1-	Division/Region	Sales Value

PB Report Name	Complex	Freq	Dimension	Metrics
Chain(MO)		Occasionally		
Subclass Ranking(HW)	3-High	2-Weekly		
Loc,Division,LW,Sales(HW)	3-High	2-Weekly	Area/Region/Location/Division	Sales Value & Forward WOS EOH Value
Org, LW Sales Value TD TY/LY(HW)	3-High	2-Weekly	Organization	Sales Markdown
Org,Markdown Value vs Fcst,Plan,LY(HW)	3-High	2-Weekly	Organization	Sales Value
Org,Sales Value vs Fcst,Plan,LY(HW)	3-High	2-Weekly	Organization	Sales Value
Prod, LW Sales Value TD TY/LY(HW)	3-High	2-Weekly	Product	Sales Value
Prod,Markdown Value vs Fcst,Plan,LY(HW)	3-High	2-Weekly	Product	Sales Markdown
Prod,Sales Value vs Fcst,Plan,LY(HW)	3-High	2-Weekly	Product	Sales Value
Time,Markdown Value vs Fcst,Plan,LY(HW)	3-High	2-Weekly	Time	Sales Markdown
Time,Sales Value vs Fcst,Plan,LY(HW)	3-High	2-Weekly	Time	Sales Value
Dept,Loc,LW,Sales(HO)	3-High	1-Occasionally	Div/Dept/Area/Region/Loc	Sales Value & Forward WOS EOH Value
Loc, Week, Sales Uplift(HO)	3-High	1-Occasionally	Location by Week	Sales Units/Value & Forward WOS EOH Value
Loc,Div,Dept,LW,Sales(HO)	3-High	1-Occasionally	Area/Region/Loc/Div/Dept	Sales Value & Forward WOS EOH Value
Loc,End of Season On Hand Fcst,LW(HO)	3-High	1-Occasionally	Region/Location	EOH Value
Loc,Subclass,LW,Product Analysis(HO)	3-High	1-Occasionally	Loc/Dept/Class/Subclass	Sales Value
Loc,Subclass,LW,Sales,Markdown(HO)	3-High	1-Occasionally	Loc/Dept/Class/Subclass	Sales Value
Org,Sales,Profit,MD Value(HO)	3-High	1-Occasionally	Organization	Sales/Markdown Value
Prod,Sales,Profit,MD Value(HO)	3-High	1-Occasionally	Product	Sales/Markdown Value
Region,Division,LW,Sales(HO)	3-High	1-Occasionally	Chain/Area/Region by Division	Sales Value
Region,Division,LW,Stock(HO)	3-High	1-Occasionally	Chain/Area/Region by Division	EOH Value
Time,Sales,Profit,MD Value(HO)	3-High	1-Occasionally	Time	Sales/Markdown Value
SKU,Week,Average Price(HO)	3-High	1-Occasionally		
SKU,Sales,Cost,Profit(HO)	3-High	1-Occasionally		

Sample SQL for Low, Medium and High Complexity is show in the Appendix.

Functional Scope
Retek Data Warehouse has several functional areas and it was important to limit the benchmark to subject areas that were most commonly used by the end users. The functional areas that were benchmarked are:]

- Sales
- Inventory
- Inventory Movement
- Pricing
- Cost

The functional areas that were not tested as part of this benchmarking exercise will be volume tested at a later point in time. Also the Data Warehouse Interface (DWI) programs were not executed since they primarily run in the Retek Merchandising System (RMS) environment; these will also be tested at a future time.

Dimensional Values
In order to represent a sizable retailer, the following dimensional values were established:

Dimension	Rows
SKUs	1,634,098
Sub-class	31,755
Class	16,455
Locations	125

The Dimensions represent a real "hardline" retailer with 125 stores.

Fact Generation

Retek currently has programs that will generate sample RMS data and RDW dimension data. However, there has not historically been an easy way to automatically generate the data in the RDW fact tables.

For the benchmark project, Retek created a new Fact Generation utility to perform efficient bulk loading of data into the RDW fact tables. The solution is data driven and uses 2 distinct types of input column data: primary keys that reference the dimension tables; and general (fact) columns of interest. Each of these are represented by a different table which serves as the base for the FactGen utility processing. The FactGen utility was designed to account for columns which are mutually exclusive, random population of fact columns (i.e. columns can be populated randomly 1/3 of the time, etc…), NULL values where appropriate, etc.

There were several specific goals that were necessary for FACTGEN to be considered successful. They included:

- Having the generated data simulate real data. The data that was generated needed to be "intelligently random". That is, it needed be a reasonable attempt at real data.
- The generated data must have data integrity. The primary and foreign keys must be accurate.
- The data generation process must be fast and scaleable. Due to large table sizes, the goal was several million rows per minute. This could only be accomplished by being able to insert into multiple partitions of the same table at the same time.

Benchmark Results

Overall, the Benchmark was very successful both in batch (data load) and front-end (report query) processing. The results far exceeded the initial goals that were set as targets for application benchmark.

Batch Results

Batch processing assumed that the data was already extracted from an external system (like RMS). The data was loaded into staging areas (AT) and then moved into the 1st level DataMart areas. Three levels of aggregations were performed, as provided in the base RDW product (sku-loc-week, sbc-loc-day, sbc-loc-week). The batch data flow is represented as below:

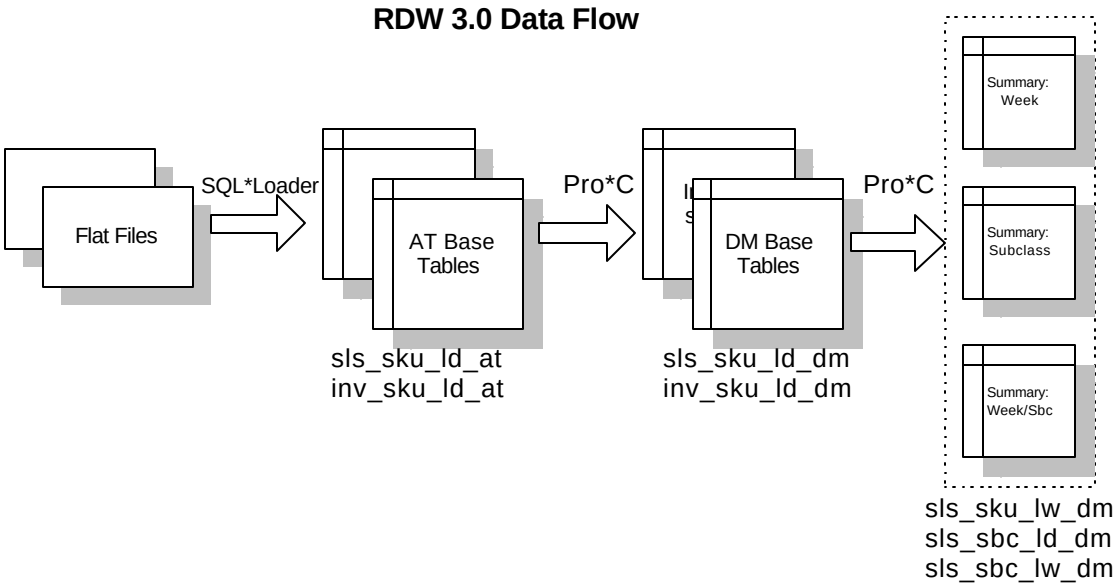

The batch results for the "end-to-end" cycle, including sales, inventory, inventory movement, pricing, and cost can be best represented by the following graph:

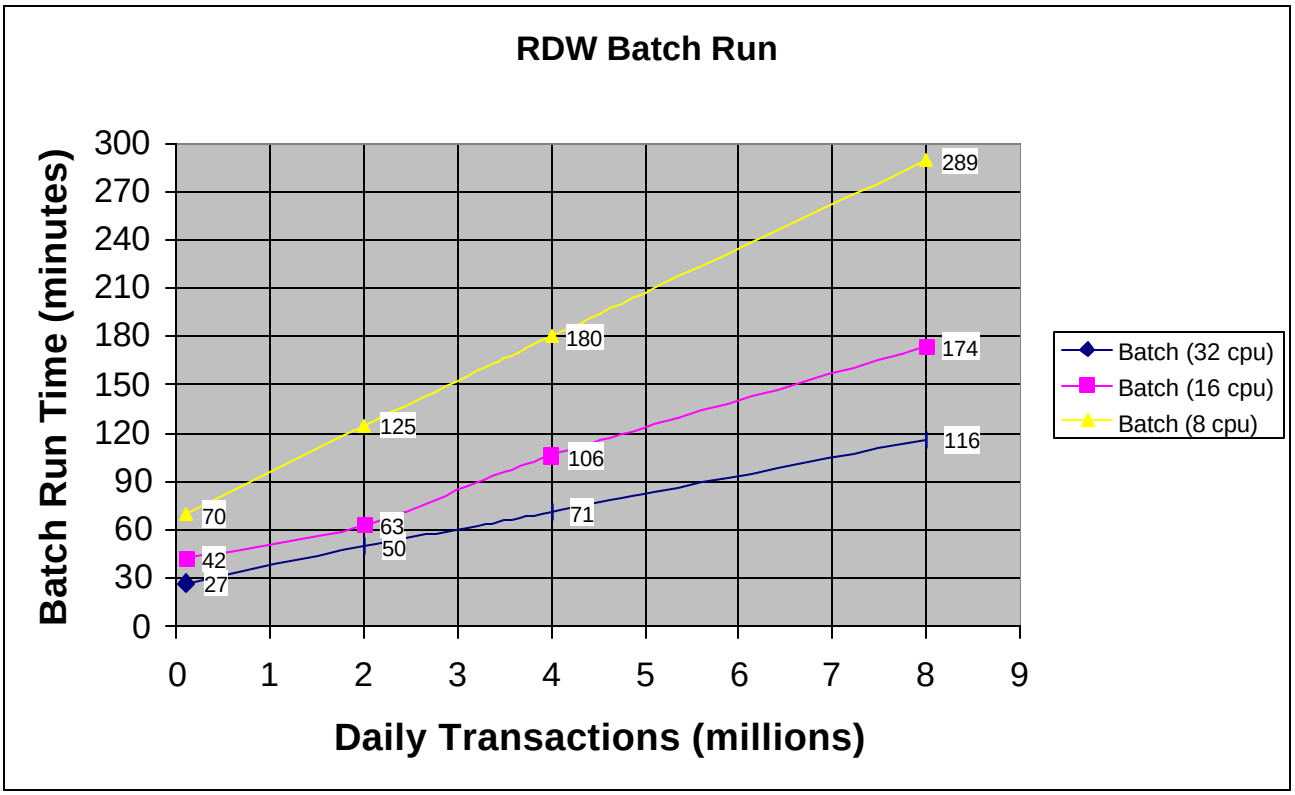

The 'Daily Transactions' indicate the daily sales volume, inventory changes, inventory movement, price changes and cost changes. A Daily Transaction of 2 million would indicate the following volumes in each functional area:

Function	Volume
Sales	2,2029,500
Inventory	2,250,000
Inventory Movement	225,000
Pricing	123,000
Cost	125,000

Daily Transaction for 4 million and 8 million would factor the numbers above by 2 times and 4 times respectively.

As can be seen from the graph, the scalability is near linear at all transaction rates (2, 4, & 8 million) and at all CPU configurations (8, 16 & 32 CPU). RDW could process 8 million transactions given the dataflow as presented above in under 5 hours with a 32 cpu configuration.

The CPU utilization for the 2 million / 32 cpu configuration is shown below:

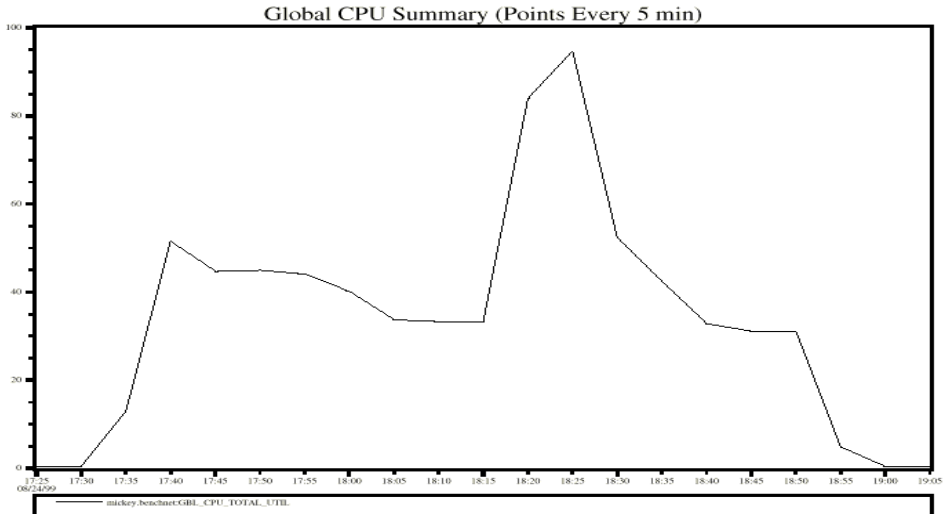

As indicated in the graph, the CPUs are about 50% used for the 1st half of the batch cycle, later approaching 90% and then back to about 40%. There is enough application level contention that beyond a certain point, all CPUs cannot be fully utilized. The peak of 90% represents a certain sort that was performed in memory.

The Disk utilization for the 2 million / 32 cpu configuration is shown below:

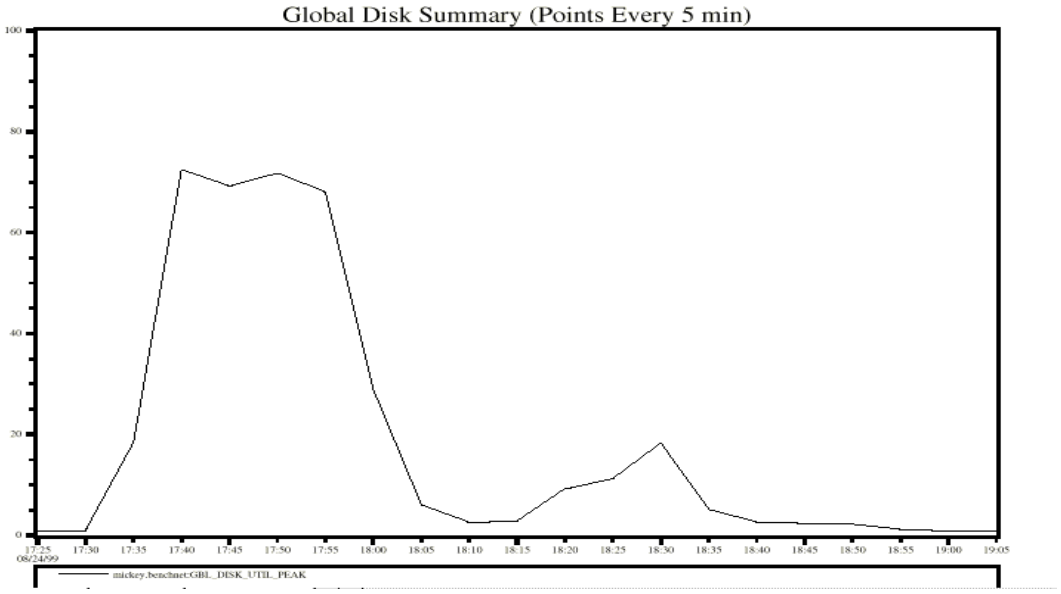

The Disk utilization was nearly 80% for the beginning of the batch cycle and later most of the information was cached in memory thus reducing the access to physical disk. The configuration of the disk with striping/mirroring allowed us to achieve very high I/O rates.

Front-End Results

The front-end tests were run using DSS Broadcaster and DSS Server. Reports were queued up from Broadcaster to Server and then the threads to the database were opened. The Report Suite was run using a varying number of threads to the database. The tests were also run in a 600 GB and 1 TB environment.

The results of the Report Suite in the 1 TB environment are shown below:

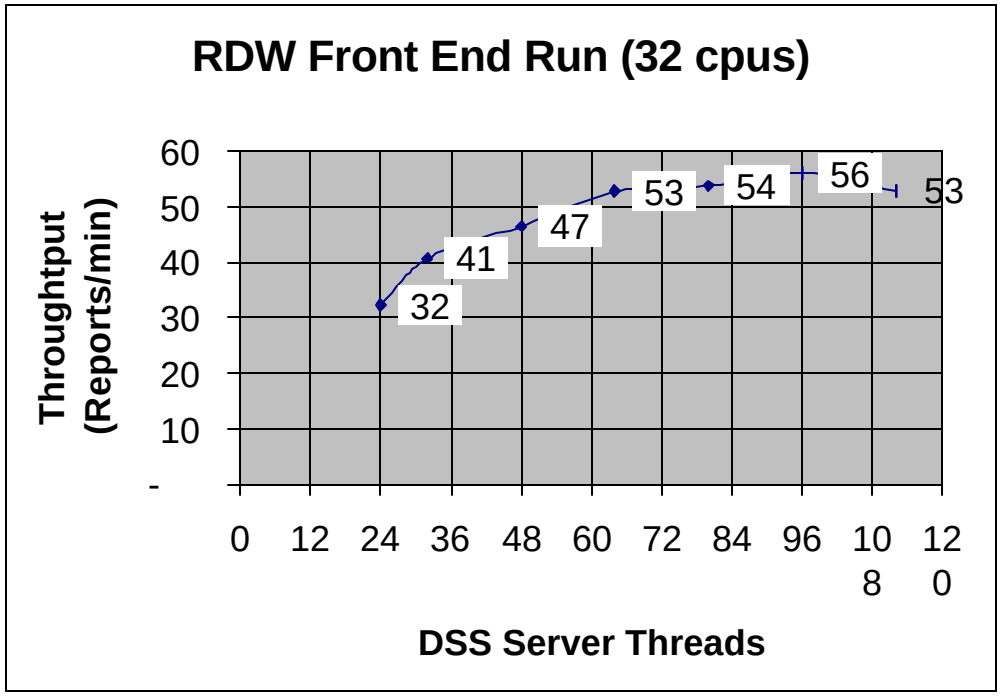

As can be seen from the graph above, the throughput rate (measured in reports executed per minute) begins to level off at about 64 Server threads, achieving a rate of approx. 53 reports/min.

The table below presents some more details about the front-end results:

Measure	600 GB	1 TB
Total Reports	2,784	2,935
Execution Time	0 hr 56 min	1 hr 3 min
Throughput	50 reports/min	46 reports/min
DSS Threads	48	48
Avg Execution/report	44 sec	50 sec
Avg Rows returned	177	115 rows
Avg Queries	3.9	4.2

As can be seen from the numbers above, the throughput rate was very similar between the 600 GB and 1 TB environments. It is worth noting that the reports per minute excluded reports that had errors or did not complete successfully.

The average execution time of the report is shown below. Note that this does not include queue time, which can vary based on number of threads and arrival rate of reports from the end users. In our case it was necessary to create a large queue of reports to ensure that the maximum number of reports was continuously executing against the database.

Response	Reports (600 GB)	Reports (1 TB)
1 – 10 sec	66%	69%
11 – 30 sec	12%	13%
31 sec – 2 min	12%	11%
2 - 7 min	7%	5%
8 – 15 min	3%	< 1%
16 – 30 min	<none>	2%
Total	100%	100%

The CPU utilization (32-cpu configuration and 48 threads) for the front-end testing is shown below:

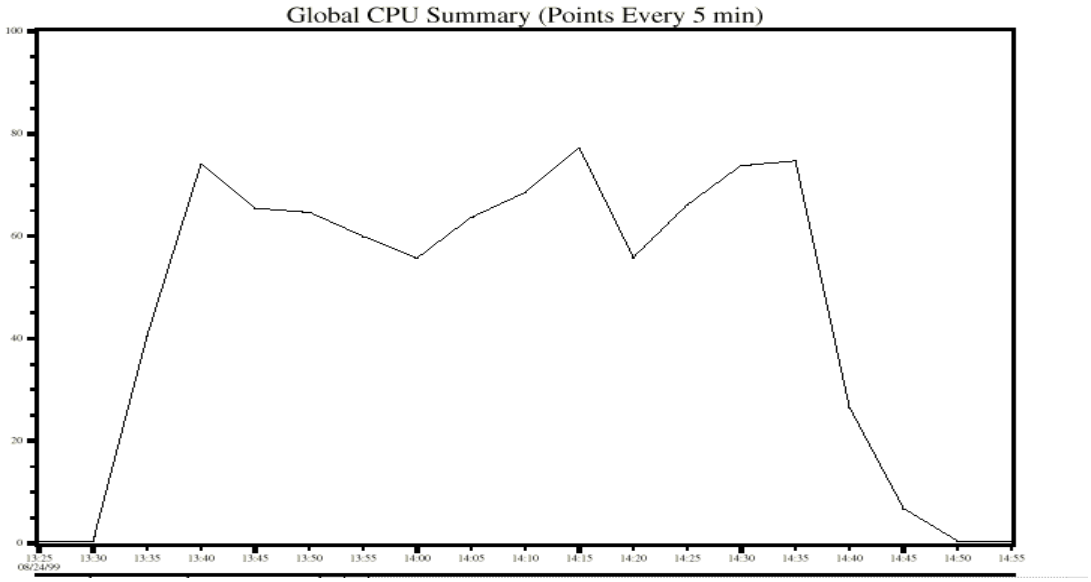

The Disk utilization (32-cpu configuration and 48 threads) for the front-end testing is shown below:

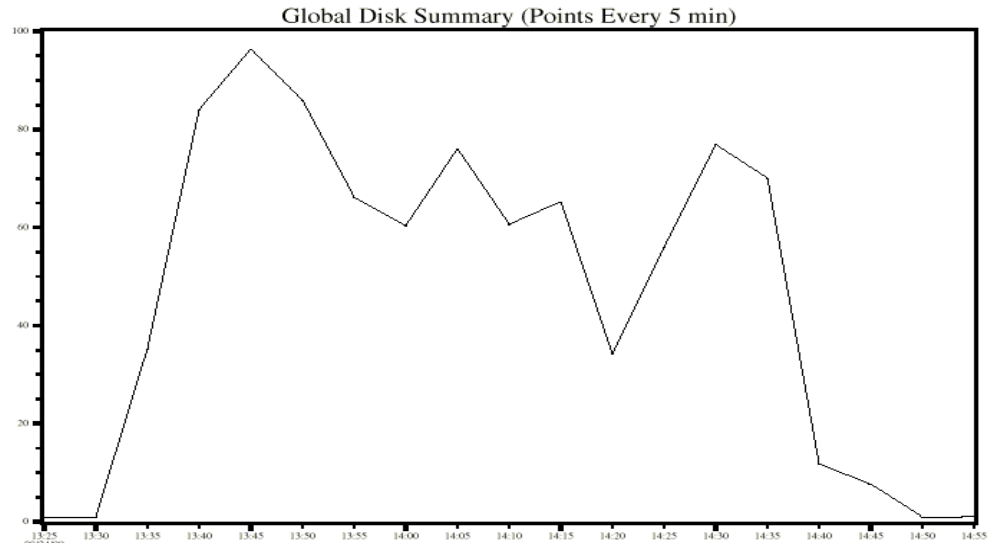

As can be seen from the graphs above, front-end execution of reports requires both high CPU and Disk utilization.

Database Details

Oracle 8i Features

Oracle 8.1.5.0 was used to conduct the RDW Performance Benchmark. At the time of the benchmark, only the 32-bit version was generally available for use. Since Retek's RDW application was not certified on the Oracle 8i platform it was not possible to use all the new features from Oracle. However, improvements made to the Oracle RDBMS kernel like optimization enhancements and sort efficiency had a positive effect on the benchmark results.

The following features & configuration options of Oracle 8i were used:

Sort Improvements

Sort performance in Oracle 8i is more rapid compared to previous releases. A new initialization parameter '`sort_multiblock_read_count`' provided more efficient I/O to temporary areas of the database for larger sorts. Additionally, sort area for a session could be dynamically increased by using '`alter session set sort_area_size = <value>` '.

Long Running Operations

In any typical datawarehouse, there can be several long running SQL operations both for batch loads and front-end queries. Oracle 8i introduced a new dynamic view `v$sessions_longops` that allow the monitoring and estimation of completion times for long running operations. The use of this feature allowed the team to monitor and estimate completion times for database operations, batch programs and user queries.

Initialization Parameters

The following initialization parameters were used for the benchmark. The parameters were unchanged in batch and front-end tests, as would be in a realistic RDW implementation. The initialization parameters also represent changes to obsolete parameters as of 8.0.5 and new parameters introduced in 8.1.5. Note there may be differences per specific environments as these are the parameters specifically used with Oracle 8.1.5 for the HPUX operating system.

Parameter	Value
O7_DICTIONARY_ACCESSIBILITY	TRUE
_db_block_hash_buckets	60000
_trace_files_public	TRUE
always_anti_join	HASH
always_semi_join	Standard
aq_tm_processes	0
audit_file_dest	$ORACLE_BASE/admin/$ORACLE_SID/adump
audit_trail	DB
background_core_dump	Partial
background_dump_dest	$ORACLE_BASE/admin/$ORACLE_SID/bdump
backup_tape_io_slaves	FALSE
bitmap_merge_area_size	10240000

Parameter	Value
blank_trimming	FALSE
buffer_pool_keep	
buffer_pool_recycle	
commit_point_strength	100
compatible	8.1.5
control_file_record_keep_time	7
control_files	/u01/oradata/rdw1/control1.ctl, /u01/oradata/rdw1/control2.ctl
core_dump_dest	$ORACLE_BASE/admin/$ORACLE_SID/cdump
cpu_count	32
create_bitmap_area_size	32768000
cursor_space_for_time	TRUE
db_block_buffers	60000
db_block_checking	FALSE
db_block_checksum	FALSE
db_block_lru_latches	64
db_block_max_dirty_target	60000
db_block_size	16384
db_domain	
db_file_direct_io_count	64
db_file_multiblock_read_count	64
db_file_name_convert	
db_files	2048
db_name	rdw1
db_writer_processes	1
dblink_encrypt_login	FALSE
dbwr_io_slaves	24
disk_asynch_io	TRUE
distributed_transactions	222
dml_locks	3564
enqueue_resources	5632
ent_domain_name	
event	
fast_start_io_target	60000
fast_start_parallel_rollback	low
fixed_date	
gc_defer_time	10
gc_files_to_locks	
gc_releasable_locks	0
gc_rollback_locks	
global_names	FALSE
hash_area_size	65536000
hash_join_enabled	TRUE
hash_multiblock_io_count	64
hi_shared_memory_address	0
hs_autoregister	TRUE
ifile	/u00/oracle/product/8.1.5/dbs/initdef.ora
instance_groups	
instance_name	
instance_number	0
java_max_sessionspace_size	0
java_pool_size	1M
java_soft_sessionspace_limit	0
job_queue_interval	60
job_queue_processes	0
large_pool_size	0
license_max_sessions	0
license_max_users	0
license_sessions_warning	0
lm_locks	12000
lm_procs	127
lm_ress	6000
local_listener	
lock_name_space	

Parameter	Value
lock_sga	FALSE
log_archive_dest	
log_archive_dest_1	
log_archive_dest_2	
log_archive_dest_3	
log_archive_dest_4	
log_archive_dest_5	
log_archive_dest_state_1	enable
log_archive_dest_state_2	enable
log_archive_dest_state_3	enable
log_archive_dest_state_4	enable
log_archive_dest_state_5	enable
log_archive_duplex_dest	
log_archive_format	%t%s.log
log_archive_max_processes	1
log_archive_min_succeed_dest	1
log_archive_start	TRUE
log_buffer	10240000
log_checkpoint_interval	99999999
log_checkpoint_timeout	7200
log_checkpoints_to_alert	FALSE
log_file_name_convert	
max_commit_propagation_delay	700
max_dump_file_size	51200
max_enabled_roles	128
max_rollback_segments	30
mts_dispatchers	
mts_listener_address	
mts_max_dispatchers	5
mts_max_servers	20
mts_multiple_listeners	FALSE
mts_servers	0
mts_service	rdw1
nls_calendar	
nls_comp	
nls_currency	
nls_date_format	
nls_date_language	
nls_dual_currency	
nls_iso_currency	
nls_language	AMERICAN
nls_numeric_characters	
nls_sort	
nls_territory	AMERICA
nls_time_format	
nls_time_tz_format	
nls_timestamp_format	
nls_timestamp_tz_format	
object_cache_max_size_percent	10
object_cache_optimal_size	102400
open_cursors	500
open_links	4
open_links_per_instance	4
optimizer_features_enable	8.1.5
optimizer_index_caching	0
optimizer_index_cost_adj	100
optimizer_max_permutations	80000
optimizer_mode	CHOOSE
optimizer_percent_parallel	0
optimizer_search_limit	5
oracle_trace_collection_name	
oracle_trace_collection_path	?/otrace/admin/cdf
oracle_trace_collection_size	5242880
oracle_trace_enable	FALSE

Parameter	Value
oracle_trace_facility_name	oracled
oracle_trace_facility_path	?/otrace/admin/fdf
os_authent_prefix	
os_roles	FALSE
parallel_adaptive_multi_user	FALSE
parallel_automatic_tuning	FALSE
parallel_broadcast_enabled	FALSE
parallel_execution_message_size	16384
parallel_instance_group	
parallel_max_servers	64
parallel_min_percent	0
parallel_min_servers	32
parallel_server	FALSE
parallel_server_instances	1
parallel_threads_per_cpu	2
partition_view_enabled	FALSE
plsql_load_without_compile	FALSE
plsql_v2_compatibility	TRUE
pre_page_sga	FALSE
processes	400
query_rewrite_enabled	FALSE
query_rewrite_integrity	enforced
rdbms_server_dn	
read_only_open_delayed	FALSE
recovery_parallelism	32
remote_dependencies_mode	TIMESTAMP
remote_login_passwordfile	NONE
remote_os_authent	FALSE
remote_os_roles	FALSE
replication_dependency_tracking	TRUE
resource_limit	FALSE
resource_manager_plan	
rollback_segments	rbs1,rbs2,rbs3.rbs4,rbs5,rbs6,rbs7,rbs8,rbs9,rbs10
row_locking	always
serial_reuse	DISABLE
serializable	FALSE
service_names	rdw1
session_cached_cursors	100
session_max_open_files	10
sessions	810
shadow_core_dump	partial
shared_memory_address	0
shared_pool_reserved_size	25M
shared_pool_size	256M
sort_area_retained_size	16000000
sort_area_size	32000000
sort_multiblock_read_count	64
sql92_security	FALSE
sql_trace	FALSE
standby_archive_dest	?/dbs/arch
star_transformation_enabled	FALSE
tape_asynch_io	TRUE
text_enable	FALSE
thread	0
timed_os_statistics	0
timed_statistics	FALSE
transaction_auditing	TRUE
transactions	891
transactions_per_rollback_segment	7
use_indirect_data_buffers	FALSE
user_dump_dest	$ORACLE_BASE/admin/$ORACLE_SID/udump
utl_file_dir	*

From an RDW implementation perspective at customer sites, most of the init.ora parameters for Oracle 8.1.5.0 can be the same as shown above; however, there is a small but critical subset of parameters that should be made site-specific based on the configuration and usage requirements.

The site-specific parameters are show below:

Parameter	Value
_db_block_hash_buckets	Same as db_block_buffers
bitmap_merge_area_size	8 – 16 Mb
Control_files	<site specific>
cpu_count	# cpus
create_bitmap_area_size	8 – 32 Mb
db_block_buffers	10,000+ (160+ Mb)
db_block_lru_latches	2 * cpus
db_file_multiblock_read_count	16 - 64
db_name	<instance specific>
dbwr_io_slaves	Start with # disk-arrays
disk_asynch_io	True, for async_io
hash_area_size	Start with 16 Mb
hash_multiblock_io_count	Same as db_file_multiblock_read_count
parallel_max_servers	2 * cpus
parallel_min_servers	# cpus
processes	Start w/ 100
rollback_segments	Start w/ rbs1,rbs2,rbs3.rbs4
sessions	2.1 * processes
shared_pool_reserved_size	10% of shared_pool_size
shared_pool_size	Start w/ 128 Mb
sort_area_retained_size	Start w/ 2 Mb
sort_area_size	Start w/ 16 Mb
sort_multiblock_read_count	Same as db_file_multiblock_read_count

Also, please note that other parameters may need to be modified in production environments (example, for archiving).

Tablespace Details

Tablespace organization is important in the physical database design both from a system management and performance perspective. Standard RDW implementation for tablespace assignment was used with the following enhancements:

- Partitioned fact tables (data) had their own tablespace
- Partitioned fact indexes had their own tablespace
- `AT_DIM_DATA, AT_FACT_DATA, DM_DIM_DATA, DM_FACT_DATA` tablespaces held both table/index due the smaller size and striped disks

The tablespace configuration and sizes for the entire database are shown in the following table. Note that the raw partition size was 6 GB, the size of an Oracle datafile and the minimum size of a tablespace.

Tablespace	Size(MB)	ID	FileName
AT_DIM_DATA	6,010	14	/u01/oradata/rdw1/at_dim_data1.dat
**********	----------		
sum	6,010		
AT_FACT_DATA	6,010	15	/u01/oradata/rdw1/at_fact_data1.dat
**********	----------		
sum	6,010		
CACHE_DATA	6,010	100	/u01/oradata/rdw1/cache_data1.dat
**********	----------		
sum	6,010		
DM_DIM_DATA	6,010	16	/u01/oradata/rdw1/dm_dim_data1.dat
**********	----------		
sum	6,010		
DM_FACT_DATA	6,010	17	/u01/oradata/rdw1/dm_fact_data1.dat
**********	----------		
sum	6,010		
RBS	6,010	2	/u01/oradata/rdw1/rbs1.dat
	6,010	3	/u01/oradata/rdw1/rbs2.dat
	6,010	4	/u01/oradata/rdw1/rbs3.dat
	6,010	5	/u01/oradata/rdw1/rbs4.dat
**********	----------		
sum	24,040		
SYSTEM	6,010	1	/u01/oradata/rdw1/system1.dat
**********	----------		
sum	6,010		
TEMP	6,010	6	/u01/oradata/rdw1/temp1.dat
	6,010	7	/u01/oradata/rdw1/temp2.dat
	6,010	8	/u01/oradata/rdw1/temp3.dat
	6,010	9	/u01/oradata/rdw1/temp4.dat
	6,010	10	/u01/oradata/rdw1/temp5.dat
	6,010	11	/u01/oradata/rdw1/temp6.dat
	6,010	12	/u01/oradata/rdw1/temp7.dat
	6,010	13	/u01/oradata/rdw1/temp8.dat
	6,010	111	/u01/oradata/rdw1/temp9.dat
	6,010	112	/u01/oradata/rdw1/temp10.dat
	6,010	113	/u01/oradata/rdw1/temp11.dat
	6,010	114	/u01/oradata/rdw1/temp12.dat
	6,010	115	/u01/oradata/rdw1/temp13.dat
	6,010	116	/u01/oradata/rdw1/temp14.dat
	6,010	117	/u01/oradata/rdw1/temp15.dat
	6,010	118	/u01/oradata/rdw1/temp16.dat
**********	----------		
sum	96,160		
TS_COST_SKU_LD_DM_DATA	6,010	18	/u01/oradata/rdw1/ts_cost_sku_ld_dm_data1.dat
	6,010	119	/u01/oradata/rdw1/ts_cost_sku_ld_dm_data2.dat
**********	----------		
sum	12,020		
TS_COST_SKU_LD_DM_INDX	6,010	19	/u01/oradata/rdw1/ts_cost_sku_ld_dm_indx1.dat
**********	----------		
sum	6,010		
TS_INV_MOVE_BLD_DM_DATA	6,010	20	/u01/oradata/rdw1/ts_inv_move_bld_dm_data1.dat
**********	----------		
sum	6,010		
TS_INV_MOVE_BLD_DM_INDX	6,010	21	/u01/oradata/rdw1/ts_inv_move_bld_dm_indx1.dat
**********	----------		
sum	6,010		
TS_INV_MOVE_BLW_DM_DATA	6,010	22	/u01/oradata/rdw1/ts_inv_move_blw_dm_data1.dat
	6,010	120	/u01/oradata/rdw1/ts_inv_move_blw_dm_data2.dat
**********	----------		
sum	12,020		
TS_INV_MOVE_BLW_DM_INDX	6,010	23	/u01/oradata/rdw1/ts_inv_move_blw_dm_indx1.dat
**********	----------		
sum	6,010		
TS_INV_MOVE_SLD_DM_DATA	6,010	24	/u01/oradata/rdw1/ts_inv_move_sld_dm_data1.dat

```
****************************** ----------
sum                               6,010

TS_INV_MOVE_SLD_DM_INDX           6,010       25 /u01/oradata/rdw1/ts_inv_move_sld_dm_indx1.dat
****************************** ----------
sum                               6,010

TS_INV_MOVE_SLW_DM_DATA           6,010       26 /u01/oradata/rdw1/ts_inv_move_slw_dm_data1.dat
****************************** ----------
sum                               6,010

TS_INV_MOVE_SLW_DM_INDX           6,010       27 /u01/oradata/rdw1/ts_inv_move_slw_dm_indx1.dat
****************************** ----------
sum                               6,010

TS_INV_SBC_LD_DM_DATA             6,010       36 /u01/oradata/rdw1/ts_inv_sbc_ld_dm_data1.dat
                                  6,010       37 /u01/oradata/rdw1/ts_inv_sbc_ld_dm_data2.dat
                                  6,010       38 /u01/oradata/rdw1/ts_inv_sbc_ld_dm_data3.dat
                                  6,010       39 /u01/oradata/rdw1/ts_inv_sbc_ld_dm_data4.dat
                                  6,010      126 /u01/oradata/rdw1/ts_inv_sbc_ld_dm_data5.dat
                                  6,010      127 /u01/oradata/rdw1/ts_inv_sbc_ld_dm_data6.dat
****************************** ----------
sum                              36,060


TS_INV_SBC_LD_DM_INDX             6,010       29 /u01/oradata/rdw1/ts_inv_sbc_ld_dm_indx1.dat
                                  6,010       30 /u01/oradata/rdw1/ts_inv_sbc_ld_dm_indx2.dat
****************************** ----------
sum                              12,020

TS_INV_SBC_LW_DM_DATA             6,010       31 /u01/oradata/rdw1/ts_inv_sbc_lw_dm_data1.dat
                                  6,010       32 /u01/oradata/rdw1/ts_inv_sbc_lw_dm_data2.dat
                                  6,010      101 /u01/oradata/rdw1/ts_inv_sbc_lw_dm_data3.dat
                                  6,010      133 /u01/oradata/rdw1/ts_inv_sbc_lw_dm_data4.dat
                                  6,010      134 /u01/oradata/rdw1/ts_inv_sbc_lw_dm_data5.dat
                                  6,010      135 /u01/oradata/rdw1/ts_inv_sbc_lw_dm_data6.dat
                                  6,010      136 /u01/oradata/rdw1/ts_inv_sbc_lw_dm_data7.dat
****************************** ----------
sum                              42,070

TS_INV_SBC_LW_DM_INDX             6,010       28 /u01/oradata/rdw1/ts_inv_sbc_lw_dm_indx1.dat
                                  6,010      121 /u01/oradata/rdw1/ts_inv_sbc_lw_dm_indx2.dat
****************************** ----------
sum                              12,020

TS_INV_SKU_LD_CUR_DM_DATA         6,010      106 /u01/oradata/rdw1/ts_inv_sku_ld_cur_dm_data1.dat
                                  6,010      107 /u01/oradata/rdw1/ts_inv_sku_ld_cur_dm_data2.dat
                                  6,010      108 /u01/oradata/rdw1/ts_inv_sku_ld_cur_dm_data3.dat
                                  6,010      109 /u01/oradata/rdw1/ts_inv_sku_ld_cur_dm_data4.dat
                                  6,010      110 /u01/oradata/rdw1/ts_inv_sku_ld_cur_dm_data5.dat
****************************** ----------
sum                              30,050

TS_INV_SKU_LD_CUR_DM_INDX         6,010      105 /u01/oradata/rdw1/ts_inv_sku_ld_cur_dm_indx1.dat
****************************** ----------
sum                               6,010

TS_INV_SKU_LD_DM_DATA             6,010       73 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data1.dat
                                  6,010       74 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data2.dat
                                  6,010       75 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data3.dat
                                  6,010       76 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data4.dat
                                  6,010       77 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data5.dat
                                  6,010       78 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data6.dat
                                  6,010       79 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data7.dat
                                  6,010       80 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data8.dat
                                  6,010       81 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data9.dat
                                  6,010       82 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data10.dat
                                  6,010      144 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data11.dat
                                  6,010      145 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data12.dat
                                  6,010      146 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data13.dat
                                  6,010      147 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data14.dat
                                  6,010      148 /u01/oradata/rdw1/ts_inv_sku_ld_dm_data15.dat
****************************** ----------
sum                              90,150


TS_INV_SKU_LD_DM_INDX             6,010       42 /u01/oradata/rdw1/ts_inv_sku_ld_dm_indx1.dat
                                  6,010       43 /u01/oradata/rdw1/ts_inv_sku_ld_dm_indx2.dat
                                  6,010       44 /u01/oradata/rdw1/ts_inv_sku_ld_dm_indx3.dat
                                  6,010       45 /u01/oradata/rdw1/ts_inv_sku_ld_dm_indx4.dat
                                  6,010      128 /u01/oradata/rdw1/ts_inv_sku_ld_dm_indx5.dat
```

```
                                 6,010      129 /u01/oradata/rdw1/ts_inv_sku_ld_dm_indx6.dat
****************************** ----------
sum                             36,060

TS_INV_SKU_LW_DM_DATA            6,010       83 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data1.dat
                                 6,010       84 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data2.dat
                                 6,010       85 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data3.dat
                                 6,010       86 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data4.dat
                                 6,010       87 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data5.dat
                                 6,010       88 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data6.dat
                                 6,010       89 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data7.dat
                                 6,010       90 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data8.dat
                                 6,010       91 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data9.dat
                                 6,010       92 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data10.dat
                                 6,010       93 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data11.dat
                                 6,010      102 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data12.dat
                                 6,010      166 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data13.dat
                                 6,010      167 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data14.dat
                                 6,010      168 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data15.dat
                                 6,010      169 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data16.dat
                                 6,010      170 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data17.dat
                                 6,010      171 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data18.dat
                                 6,010      172 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data19.dat
                                 6,010      173 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data20.dat
                                 6,010      174 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data21.dat
                                 6,010      175 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data22.dat
                                 6,010      176 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data23.dat
                                 6,010      177 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data24.dat
                                 6,010      178 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data25.dat
                                 6,010      179 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data26.dat
                                 6,010      180 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data27.dat
                                 6,010      181 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data28.dat
                                 6,010      182 /u01/oradata/rdw1/ts_inv_sku_lw_dm_data29.dat
****************************** ----------
sum                            174,290

TS_INV_SKU_LW_DM_INDX            6,010       46 /u01/oradata/rdw1/ts_inv_sku_lw_dm_indx1.dat
                                 6,010       47 /u01/oradata/rdw1/ts_inv_sku_lw_dm_indx2.dat
                                 6,010       48 /u01/oradata/rdw1/ts_inv_sku_lw_dm_indx3.dat
                                 6,010       49 /u01/oradata/rdw1/ts_inv_sku_lw_dm_indx4.dat
                                 6,010      149 /u01/oradata/rdw1/ts_inv_sku_lw_dm_indx5.dat
                                 6,010      150 /u01/oradata/rdw1/ts_inv_sku_lw_dm_indx6.dat
                                 6,010      151 /u01/oradata/rdw1/ts_inv_sku_lw_dm_indx7.dat
                                 6,010      152 /u01/oradata/rdw1/ts_inv_sku_lw_dm_indx8.dat


TS_INV_SKU_LW_DM_INDX            6,010      153 /u01/oradata/rdw1/ts_inv_sku_lw_dm_indx9.dat
                                 6,010      154 /u01/oradata/rdw1/ts_inv_sku_lw_dm_indx10.dat
****************************** ----------
sum                             60,100

TS_PRICING_SKU_LD_DM_DATA        6,010       33 /u01/oradata/rdw1/ts_pricing_sku_ld_dm_data1.dat
                                 6,010      122 /u01/oradata/rdw1/ts_pricing_sku_ld_dm_data2.dat
****************************** ----------
sum                             12,020

TS_PRICING_SKU_LD_DM_INDX        6,010       34 /u01/oradata/rdw1/ts_pricing_sku_ld_dm_indx1.dat
****************************** ----------
sum                              6,010

TS_SLS_CLASS_RW_DM_DATA          6,010       35 /u01/oradata/rdw1/ts_sls_class_rw_dm_data1.dat
                                 6,010      123 /u01/oradata/rdw1/ts_sls_class_rw_dm_data2.dat
****************************** ----------
sum                             12,020

TS_SLS_CLASS_RW_DM_INDX          6,010       40 /u01/oradata/rdw1/ts_sls_class_rw_dm_indx1.dat
****************************** ----------
sum                              6,010

TS_SLS_SBC_LD_DM_DATA            6,010       41 /u01/oradata/rdw1/ts_sls_sbc_ld_dm_data1.dat
                                 6,010      124 /u01/oradata/rdw1/ts_sls_sbc_ld_dm_data2.dat
****************************** ----------
sum                             12,020

TS_SLS_SBC_LD_DM_INDX            6,010       50 /u01/oradata/rdw1/ts_sls_sbc_ld_dm_indx1.dat
****************************** ----------
sum                              6,010

TS_SLS_SBC_LD_LY_DM_DATA         6,010       51 /u01/oradata/rdw1/ts_sls_sbc_ld_ly_dm_data1.dat
****************************** ----------
sum                              6,010
```

```
TS_SLS_SBC_LD_LY_DM_INDX            6,010       52 /u01/oradata/rdw1/ts_sls_sbc_ld_ly_dm_indx1.dat
**************************** ----------
sum                                 6,010

TS_SLS_SBC_LW_DM_DATA               6,010       53 /u01/oradata/rdw1/ts_sls_sbc_lw_dm_data1.dat
                                    6,010      103 /u01/oradata/rdw1/ts_sls_sbc_lw_dm_data2.dat
                                    6,010      130 /u01/oradata/rdw1/ts_sls_sbc_lw_dm_data3.dat
                                    6,010      131 /u01/oradata/rdw1/ts_sls_sbc_lw_dm_data4.dat
                                    6,010      132 /u01/oradata/rdw1/ts_sls_sbc_lw_dm_data5.dat
**************************** ----------
sum                                30,050

TS_SLS_SBC_LW_DM_INDX               6,010       54 /u01/oradata/rdw1/ts_sls_sbc_lw_dm_indx1.dat
**************************** ----------
sum                                 6,010

TS_SLS_SBC_LW_LY_DM_DATA            6,010       55 /u01/oradata/rdw1/ts_sls_sbc_lw_ly_dm_data1.dat
**************************** ----------
sum                                 6,010

TS_SLS_SBC_LW_LY_DM_INDX            6,010       56 /u01/oradata/rdw1/ts_sls_sbc_lw_ly_dm_indx1.dat
**************************** ----------
sum                                 6,010

TS_SLS_SKU_LD_DM_DATA               6,010       66 /u01/oradata/rdw1/ts_sls_sku_ld_dm_data1.dat
                                    6,010       67 /u01/oradata/rdw1/ts_sls_sku_ld_dm_data2.dat
                                    6,010       68 /u01/oradata/rdw1/ts_sls_sku_ld_dm_data3.dat
                                    6,010       69 /u01/oradata/rdw1/ts_sls_sku_ld_dm_data4.dat
                                    6,010      140 /u01/oradata/rdw1/ts_sls_sku_ld_dm_data5.dat
                                    6,010      141 /u01/oradata/rdw1/ts_sls_sku_ld_dm_data6.dat
                                    6,010      142 /u01/oradata/rdw1/ts_sls_sku_ld_dm_data7.dat
                                    6,010      143 /u01/oradata/rdw1/ts_sls_sku_ld_dm_data8.dat
**************************** ----------
sum                                48,080

TS_SLS_SKU_LD_DM_INDX               6,010       58 /u01/oradata/rdw1/ts_sls_sku_ld_dm_indx1.dat
                                    6,010       59 /u01/oradata/rdw1/ts_sls_sku_ld_dm_indx2.dat
                                    6,010      125 /u01/oradata/rdw1/ts_sls_sku_ld_dm_indx3.dat
**************************** ----------
sum                                18,030

TS_SLS_SKU_LD_LY_DM_DATA            6,010       60 /u01/oradata/rdw1/ts_sls_sku_ld_ly_dm_data1.dat
                                    6,010       61 /u01/oradata/rdw1/ts_sls_sku_ld_ly_dm_data2.dat
**************************** ----------
sum                                12,020

TS_SLS_SKU_LD_LY_DM_INDX            6,010       57 /u01/oradata/rdw1/ts_sls_sku_ld_ly_dm_indx1.dat
**************************** ----------
sum                                 6,010

TS_SLS_SKU_LW_DM_DATA               6,010       94 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data1.dat
                                    6,010       95 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data2.dat
                                    6,010       96 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data3.dat
                                    6,010       97 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data4.dat
                                    6,010       98 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data5.dat
                                    6,010       99 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data6.dat
                                    6,010      155 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data7.dat
                                    6,010      156 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data8.dat
                                    6,010      157 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data9.dat
                                    6,010      158 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data10.dat
                                    6,010      159 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data11.dat
                                    6,010      160 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data12.dat
                                    6,010      161 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data13.dat
                                    6,010      162 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data14.dat


TS_SLS_SKU_LW_DM_DATA               6,010      163 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data15.dat
                                    6,010      164 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data16.dat
                                    6,010      165 /u01/oradata/rdw1/ts_sls_sku_lw_dm_data17.dat
**************************** ----------
sum                               102,170

TS_SLS_SKU_LW_DM_INDX               6,010       62 /u01/oradata/rdw1/ts_sls_sku_lw_dm_indx1.dat
                                    6,010       63 /u01/oradata/rdw1/ts_sls_sku_lw_dm_indx2.dat
                                    6,010      137 /u01/oradata/rdw1/ts_sls_sku_lw_dm_indx3.dat
                                    6,010      138 /u01/oradata/rdw1/ts_sls_sku_lw_dm_indx4.dat
                                    6,010      139 /u01/oradata/rdw1/ts_sls_sku_lw_dm_indx5.dat
                                    6,010      183 /u01/oradata/rdw1/ts_sls_sku_lw_dm_indx6.dat
                                    6,010      184 /u01/oradata/rdw1/ts_sls_sku_lw_dm_indx7.dat
**************************** ----------
```

```
sum                                42,070

TS_SLS_SKU_LW_LY_DM_DATA            6,010       70 /u01/oradata/rdw1/ts_sls_sku_lw_ly_dm_data1.dat
                                    6,010       71 /u01/oradata/rdw1/ts_sls_sku_lw_ly_dm_data2.dat
                                    6,010       72 /u01/oradata/rdw1/ts_sls_sku_lw_ly_dm_data3.dat
****************************** ----------
sum                                18,030

TS_SLS_SKU_LW_LY_DM_INDX            6,010       64 /u01/oradata/rdw1/ts_sls_sku_lw_ly_dm_indx1.dat
****************************** ----------
sum                                 6,010

USER_DATA                           6,010       65 /u01/oradata/rdw1/user_data1.dat
                                    6,010      104 /u01/oradata/rdw1/user_data2.dat
****************************** ----------
sum                                12,020

                                ----------
sum                             1,105,840
```

Redo Log Configuration

The database was in 'noarchivelog' mode and thus it was not necessary to have more than 2 Redo Logs. The database was configured to have 2 Redo Log files, each 6 GB, on raw partitions and over the striped disk array. Since the drive array was in Raid 0/1 format, there was no need to have multiplexed members for each redo group.

Temp Table Configuration

The temporary tablespace `'temp'` was configured as (16) 6 GB datafiles on raw partitions over striped arrays for a total of 96 GB of temp space. The space was necessary to support large concurrent sorts and index rebuilds.

Rollback Configuration

The rollback tablespace `'rbs'` was configured as (4) 6 GB datafiles on raw partitions over striped arrays for a total of 24 GB of rollback space. A total of (10) rollback segments were used with the following configuration:

- Initial = 20m
- Next = 20m
- Minextents = 10
- Optimal = 200m

The transactions per rollback segment was set to 7.

Partitioning Details

Oracle offers a rich set of partitioning methods that allows the database administrator to choose the approach that will offer the best performance and database management based on the application profile.

The fact tables of the RDW functional areas of scope were partitioned using the range partitioning technique. The related indexes and keys of the partitioned tables were also partitioned using local prefixed indexes. The partitioning value was the time key, for example day_idnt. There are several advantages to partitioning by time:

- Rolling change windows of historical fact data are maintained by dropping/adding partitions.
- The daily batch programs insert/update data only in the current partition(s), making performance a factor of transaction volume and not of the overall table or database size.
- The front-end queries rely on the Oracle optimizer performing partition elimination hence improving the report response times.

The RDW application relies on Oracle range partitions to achieve scalability and manageability of large databases with high transaction volumes. A brief test was conducted using Oracle 8i composite partitions (range/hash), but no performance gains were achieved. This new feature will require more investigation at a later point.

The following table provides information on partitioned fact tables that were used for the benchmark:

OWNER	TABLE NAME	PARTITION_NAME	ROWS	SIZE(bytes)
RDW30_ADM_DM	COST_SKU_LD_DM	P_COST_SKU_LD_DM_199903	3,574,375	40
		P_COST_SKU_LD_DM_199904	3,574,375	41
		P_COST_SKU_LD_DM_199905	3,574,375	41
		P_COST_SKU_LD_DM_199906	3,574,375	40
. . .				
. . .				
		P_COST_SKU_LD_DM_199927	3,574,375	40
		P_COST_SKU_LD_DM_199928	3,574,375	40
		P_COST_SKU_LD_DM_199929	1,531,875	41
		P_COST_SKU_LD_DM_MAX	0	0
	**********	**********	**********	
	sum		94,465,625	
	INV_MOVE_BLD_DM	P_INV_MOVE_BLD_DM_199921	2,800,791	43
		P_INV_MOVE_BLD_DM_199922	2,800,791	43
. . .				
. . .				
		P_INV_MOVE_BLD_DM_199929	1,200,339	43
		P_INV_MOVE_BLD_DM_MAX	0	0
	**********	**********	**********	
	sum		23,606,667	
	INV_MOVE_BLW_DM	P_INV_MOVE_BLW_DM_199807	6,351,000	42
		P_INV_MOVE_BLW_DM_199808	7,938,750	42
		P_INV_MOVE_BLW_DM_199809	6,351,000	42
		P_INV_MOVE_BLW_DM_199810	6,351,000	41
RDW30_ADM_DM	INV_MOVE_BLW_DM	P_INV_MOVE_BLW_DM_199811	7,938,750	42
		P_INV_MOVE_BLW_DM_199812	6,351,000	42
. . .				
. . .				
		P_INV_MOVE_BLW_DM_199907	3,175,500	42
		P_INV_MOVE_BLW_DM_MAX	0	0
	**********	**********	**********	
	sum		85,738,500	
	INV_MOVE_SLD_DM	P_INV_MOVE_SLD_DM_199921	714,875	44
		P_INV_MOVE_SLD_DM_199922	714,875	44
		P_INV_MOVE_SLD_DM_199923	714,875	44
. . .				
. . .				
		P_INV_MOVE_SLD_DM_199929	306,375	44
		P_INV_MOVE_SLD_DM_MAX	0	0

```
            ******************************                                -------------
            sum                                                            6,025,375

            INV_MOVE_SLW_DM               P_INV_MOVE_SLW_DM_199807          1,634,000   43
                                          P_INV_MOVE_SLW_DM_199808          2,042,500   43
                                          P_INV_MOVE_SLW_DM_199809          1,634,000   43
. . .
. . .
                                          P_INV_MOVE_SLW_DM_199907          1,021,250   43
                                          P_INV_MOVE_SLW_DM_MAX                     0    0
            ******************************                                -------------
            sum                                                           22,263,250

            INV_SBC_LD_DM                 P_INV_SBC_LD_DM_199921           27,784,972   80
                                          P_INV_SBC_LD_DM_199922           27,786,507   80
                                          P_INV_SBC_LD_DM_199923           27,784,273   80
                                          P_INV_SBC_LD_DM_199924           27,785,801   80
                                          P_INV_SBC_LD_DM_199925           27,786,117   81

RDW30_ADM_DM INV_SBC_LD_DM                P_INV_SBC_LD_DM_199926           27,786,439   81
                                          P_INV_SBC_LD_DM_199927           27,785,562   80
                                          P_INV_SBC_LD_DM_199928           27,786,438   80
                                          P_INV_SBC_LD_DM_199929           11,908,125   81
                                          P_INV_SBC_LD_DM_MAX                       0    0
            ******************************                                -------------
            sum                                                          234,194,234

            INV_SBC_LW_DM                 P_INV_SBC_LW_DM_199801           15,877,500   78
                                          P_INV_SBC_LW_DM_199802           19,847,369   79
                                          P_INV_SBC_LW_DM_199803           15,877,500   79
                                          P_INV_SBC_LW_DM_199804           15,877,500   81
. . .
 . .
                                          P_INV_SBC_LW_DM_199905           19,846,833   80
                                          P_INV_SBC_LW_DM_199906           15,877,500   81
                                          P_INV_SBC_LW_DM_199907            7,938,750   79
                                          P_INV_SBC_LW_DM_MAX                       0    0
            ******************************                                -------------
            sum                                                          317,550,881

            INV_SKU_LD_DM                 P_INV_SKU_LD_DM_199801          204,260,575  136
                                          P_INV_SKU_LD_DM_199921           42,893,674  132
                                          P_INV_SKU_LD_DM_199922           42,894,972  132
. . .
. . .
                                          P_INV_SKU_LD_DM_199929           15,319,269  133
                                          P_INV_SKU_LD_DM_MAX                       0    0
            ******************************                                -------------
            sum                                                          534,141,566

            INV_SKU_LW_DM                 P_INV_SKU_LW_DM_199801          204,262,869  133
                                          P_INV_SKU_LW_DM_199802           71,490,183  127
                                          P_INV_SKU_LW_DM_199803           57,197,242  130
. . .
 . . .
                                          P_INV_SKU_LW_DM_199905           51,061,661  128
                                          P_INV_SKU_LW_DM_199906           40,852,555  129
                                          P_INV_SKU_LW_DM_199907           20,426,099  135
                                          P_INV_SKU_LW_DM_MAX                       0    0
            ******************************                                -------------
            sum                                                        1,176,552,764

            PRICING_SKU_LD_DM             P_PRICING_SKU_LD_DM_199903        3,574,375   51
                                          P_PRICING_SKU_LD_DM_199904        3,574,375   49
                                          P_PRICING_SKU_LD_DM_199905        3,574,375   48
. . .
. . .
                                          P_PRICING_SKU_LD_DM_199926        3,574,375   51
                                          P_PRICING_SKU_LD_DM_199927        3,574,375   50
                                          P_PRICING_SKU_LD_DM_199928        3,574,375   50
                                          P_PRICING_SKU_LD_DM_199929        1,531,875   49
                                          P_PRICING_SKU_LD_DM_MAX                   0    0

        ******************************                                -------------
            sum                                                           94,465,625

RDW30_ADM_DM SLS_CLASS_RW_DM              P_SLS_CLASS_RW_DM_199803            598,962   84
                                          P_SLS_CLASS_RW_DM_199804            184,296   85
                                          P_SLS_CLASS_RW_DM_199805            230,370   84
. . .
 . . .
```

```
                                P_SLS_CLASS_RW_DM_199905              230,370   84
                                P_SLS_CLASS_RW_DM_199906              184,296   84
                                P_SLS_CLASS_RW_DM_199907               92,148   83
                                P_SLS_CLASS_RW_DM_MAX                       0    0
                ****************************                      -------------
                sum                                                 3,685,920

                SLS_SBC_LD_DM                 P_SLS_SBC_LD_DM_199921        5,557,125   98
                                              P_SLS_SBC_LD_DM_199922        5,557,125   97
                                              P_SLS_SBC_LD_DM_199923        5,557,125   98
                                              P_SLS_SBC_LD_DM_199924        5,557,125   98
                                              P_SLS_SBC_LD_DM_199925        5,557,125   97
                                              P_SLS_SBC_LD_DM_199926        5,557,125   97
                                              P_SLS_SBC_LD_DM_199927        5,557,125   97
                                              P_SLS_SBC_LD_DM_199928        5,557,125   97
                                              P_SLS_SBC_LD_DM_199929                0    0
                                              P_SLS_SBC_LD_DM_MAX                   0    0
                ****************************                      -------------
                sum                                                44,457,000

                SLS_SBC_LD_LY_DM              P_SLS_SBC_LD_LY_DM_199828     5,557,125   98
                                              P_SLS_SBC_LD_LY_DM_199829     2,381,625   98
                                              P_SLS_SBC_LD_LY_DM_MAX                0    0
                ****************************                      -------------
                sum                                                 7,938,750

                SLS_SBC_LW_DM                 P_SLS_SBC_LW_DM_199801        6,351,000   96
                                              P_SLS_SBC_LW_DM_199802        7,938,750   97
                                              P_SLS_SBC_LW_DM_199803        6,351,000   97
                                              P_SLS_SBC_LW_DM_199804        6,351,000   97
                                              P_SLS_SBC_LW_DM_199805        7,938,750   97

RDW30_ADM_DM    SLS_SBC_LW_DM                 P_SLS_SBC_LW_DM_199806        7,938,750   97
                                              P_SLS_SBC_LW_DM_199807        4,763,250   96
                                              P_SLS_SBC_LW_DM_199808        7,938,750   96

. . .
. . .
                                              P_SLS_SBC_LW_DM_199906        6,351,000   97
                                              P_SLS_SBC_LW_DM_199907                0    0
                                              P_SLS_SBC_LW_DM_MAX                   0    0
                ****************************                      -------------
                sum                                               123,844,500

                SLS_SBC_LW_LY_DM              P_SLS_SBC_LW_LY_DM_199805             0    0
                                              P_SLS_SBC_LW_LY_DM_199806     6,351,000   98
                                              P_SLS_SBC_LW_LY_DM_199807     3,175,500   98
                                              P_SLS_SBC_LW_LY_DM_MAX                0    0
                ****************************                      -------------
                sum                                                 9,526,500

                SLS_SKU_LD_DM                 P_SLS_SKU_LD_DM_199921       42,880,409  111
                                              P_SLS_SKU_LD_DM_199922       42,877,071  110
                                              P_SLS_SKU_LD_DM_199923       42,879,796  111
                                              P_SLS_SKU_LD_DM_199924       42,899,420  110
                                              P_SLS_SKU_LD_DM_199925       35,746,687  110
                                              P_SLS_SKU_LD_DM_199926       35,746,222  110
                                              P_SLS_SKU_LD_DM_199927       35,742,296  110
                                              P_SLS_SKU_LD_DM_199928       35,741,929  111
                                              P_SLS_SKU_LD_DM_199929       15,317,861  111
                                              P_SLS_SKU_LD_DM_MAX                   0    0
                ****************************                      -------------
                sum                                               329,831,691

                SLS_SKU_LD_LY_DM              P_SLS_SKU_LD_LY_DM_199828     35,736,461  102
                                              P_SLS_SKU_LD_LY_DM_199829     15,319,409  102
                                              P_SLS_SKU_LD_LY_DM_MAX                0    0
                ****************************                      -------------
                sum                                                51,055,870

                SLS_SKU_LW_DM                 P_SLS_SKU_LW_DM_199801       49,007,164  102
                                              P_SLS_SKU_LW_DM_199802       61,260,399  101
                                              P_SLS_SKU_LW_DM_199803       49,021,825  102
                                              P_SLS_SKU_LW_DM_199804       49,040,098  101
. . .
. . .
                                              P_SLS_SKU_LW_DM_199904       40,847,658  102
                                              P_SLS_SKU_LW_DM_199905       51,037,184  102
                                              P_SLS_SKU_LW_DM_199906       40,846,619  101
                                              P_SLS_SKU_LW_DM_199907       20,428,013  102
                                              P_SLS_SKU_LW_DM_MAX                   0    0
                ****************************                      -------------
```

	sum		923,200,845	
	SLS_SKU_LW_LY_DM	P_SLS_SKU_LW_LY_DM_199805	0	0
		P_SLS_SKU_LW_LY_DM_199806	40,839,609	102
		P_SLS_SKU_LW_LY_DM_199807	20,426,081	101
		P_SLS_SKU_LW_LY_DM_MAX	0	0
	******************************		----------	
	sum		61,265,690	
**********			----------	
sum			4,143,811,253	

Conclusions

The Retek Data Warehouse has proven that the entire solution is scalable and can achieve high performance levels in the following areas:

- RDW application scalability
- HP V-Class hardware scalability
- HP FC60 storage array scalability
- Oracle 8i scalability

Retek will continue to make improvements in the RDW application as a result of the lessons learned from the Performance Benchmark. Additionally, Retek will continue to invest time and resources to performing benchmarks in more RDW functional areas to ensure that all parts of RDW can scale.

Contributions

The Retek Data Warehouse was a success due to contributions from all the team members:

Retek:

- Jay Bhow
- Mike Kramer
- Josh Lee
- Mike Miller
- Roger Ruckert
- Pam Rupar

HP:

- Steve Folkman
- Mike McCormack
- John Morris

In addition many thanks to HP management for providing the Performance Center and hosting Retek at Cupertino. Finally, thanks to Retek RDW management: Nick Whitney and Amy Gearing for supporting and sponsoring the benchmark.

Appendix

Sample SQL

Low Complexity SQL

```
create table   R0626883RT0 NOLOGGING  as
select  max(a2.SKU_DESC||', '||a2.SKU_IDNT) SKU_DESC,
       a2.SKU_KEY SKU_KEY,
       (SUM(a1.F_PRMTN_SLS_AMT)) SALESPROMOVALUE,
       (SUM(a1.F_RGLR_SLS_AMT)) SALESREGULARVALUE,
       (SUM(NVL(a1.F_CLRC_SLS_AMT,0)+NVL(a1.F_PRMTN_SLS_AMT,0)+NVL(a1.F_RGLR_SLS_AMT,0)))
SALESVALUE,
       (SUM(a1.F_RGLR_SLS_AMT))/(SUM(NVL(a1.F_CLRC_SLS_AMT,0)+NVL(a1.F_PRMTN_SLS_AMT,0)+N
VL(a1.F_RGLR_SLS_AMT,0))) SALESVALUEREGULART,
       (SUM(a1.F_PRMTN_SLS_AMT))/(SUM(NVL(a1.F_CLRC_SLS_AMT,0)+NVL(a1.F_PRMTN_SLS_AMT,0)+
NVL(a1.F_RGLR_SLS_AMT,0))) SALESVALUEPROMOTOT,
       (SUM(a1.F_CLRC_SLS_AMT)) SALESCLRNCEVALUE,
       (SUM(a1.F_CLRC_SLS_AMT))/(SUM(NVL(a1.F_CLRC_SLS_AMT,0)+NVL(a1.F_PRMTN_SLS_AMT,0)+N
VL(a1.F_RGLR_SLS_AMT,0))) SALESVALUECLRNCETO
from   SLS_SKU_LD_DM a1,
       PROD_SKU_DM a2
where  a1.SKU_KEY = a2.SKU_KEY
  and  ((a1.DAY_IDNT = 1999190))
  and    ((a1.LOC_KEY  IN ( 3,  4,  56,  62) ))
  and    (((a2.SKU_KEY < 5000)))
group by       a2.SKU_KEY
```

Medium Complexity SQL

```
create table T0626905NB2 NOLOGGING  as
select                                a1.SBCLASS_KEY,
SUM(NVL(a1.F_I_SOH_CLRC_RTL_AMT,0)+NVL(a1.F_I_SOH_PRMTN_RTL_AMT,0)+NVL(a1.F_I_SOH_RGLR_RT
L_AMT,0)) EOHVALUE,
                     SUM(a1.F_I_SOH_RGLR_RTL_AMT) EOHREGULARVALUE,
                      SUM(a1.F_I_SOH_PRMTN_RTL_AMT) EOHPROMOVALUE,
                      SUM(a1.F_I_SOH_CLRC_RTL_AMT) EOHCLRNCEVALUE,
                                    a1.DAY_IDNT
from                               V_INV_SBC_LD_DM a1
where         ((a1.DAY_IDNT  IN ( 1999190,  1999191,  1999192,  1999193) ))
  and   ((a1.LOC_KEY  IN ( 10,  11,  12,  13,  14,  15,  16,  17,  18,  19,  2,  20,  21,
3,  4,  5,  6,  7,  8,  9) ))
  and                             (((a1.SBCLASS_KEY < 50)))
group by                            a1.SBCLASS_KEY,
                                    a1.DAY_IDNT
create table T0626905NC3 NOLOGGING  as
select                          T0626905NB2.SBCLASS_KEY,
                          max(T0626905NB2.DAY_IDNT) DAY_IDNT
from                                T0626905NB2
group by                         T0626905NB2.SBCLASS_KEY
ANALYZE TABLE T0626905NB2 ESTIMATE STATISTICS
ANALYZE TABLE T0626905NC3 ESTIMATE STATISTICS
create table T0626905md1 NOLOGGING  as
select                          T0626905NB2.SBCLASS_KEY,
                             T0626905NB2.EOHVALUE EOHVALUE,
                     T0626905NB2.EOHREGULARVALUE EOHREGULARVALUE,
                       T0626905NB2.EOHPROMOVALUE EOHPROMOVALUE,
                      T0626905NB2.EOHCLRNCEVALUE EOHCLRNCEVALUE
from                                T0626905NB2 ,
                                    T0626905NC3
```

```
where               T0626905NB2.SBCLASS_KEY = T0626905NC3.SBCLASS_KEY
  and                 T0626905NB2.DAY_IDNT = T0626905NC3.DAY_IDNT
drop table T0626905NB2

drop table T0626905NC3

ANALYZE TABLE T0626905md1 ESTIMATE STATISTICS

create table   R0626905RT0 NOLOGGING  as
select                        a2.SBCLASS_KEY SBCLASS_KEY,
                  a2.SBCLASS_DESC||', '||a2.SBCLASS_IDNT SBCLASS_DESC,
                            T0626905md1.EOHVALUE EOHVALUE,
                        T0626905md1.EOHREGULARVALUE EOHREGULARVALUE,
                          T0626905md1.EOHPROMOVALUE EOHPROMOVALUE,
                         T0626905md1.EOHCLRNCEVALUE EOHCLRNCEVALUE,
         (T0626905md1.EOHREGULARVALUE/T0626905md1.EOHVALUE) EOHVALUEREGULARTOT,
          (T0626905md1.EOHPROMOVALUE/T0626905md1.EOHVALUE) EOHVALUEPROMOTOTOT,
          (T0626905md1.EOHCLRNCEVALUE/T0626905md1.EOHVALUE) EOHVALUECLRNCETOTO
from                                  T0626905md1 ,
                                     PROD_SBC_DM a2
where                   T0626905md1.SBCLASS_KEY = a2.SBCLASS_KEY
drop table T0626905md1
```

Complex SQL

```
create table T0626919SP2 NOLOGGING  as
select  a1.SBCLASS_KEY,
        SUM(NVL(a1.F_CLRC_SLS_AMT_LY,0)+NVL(a1.F_RGLR_SLS_AMT_LY,0)+NVL(a1.F_PRMTN_SLS_AMT
_LY,0)) SALESLYVALUE
from    SLS_SBC_LW_LY_DM a1,
        ORG_LOC_DM a2
where   a2.LOC_KEY = a1.LOC_KEY
  and   ((a1.WK_IDNT  IN ( 199924,  199925,  199926) ))
  and     ((a2.REGN_KEY = 1))
  and     (((a1.SBCLASS_KEY < 150)))
group by        a1.SBCLASS_KEY
create table T0626919SP3 NOLOGGING  as
select  a1.SBCLASS_KEY,
        SUM(NVL(a1.F_CLRC_SLS_AMT,0)+NVL(a1.F_PRMTN_SLS_AMT,0)+NVL(a1.F_RGLR_SLS_AMT,0))
SALESVALUE
from    SLS_SBC_LW_DM a1,
        ORG_LOC_DM a2
where   a2.LOC_KEY = a1.LOC_KEY
  and   ((a1.WK_IDNT  IN ( 199924,  199925,  199926) ))
  and     ((a2.REGN_KEY = 1))
  and     (((a1.SBCLASS_KEY < 150)))
group by        a1.SBCLASS_KEY

create table T0626919OJ4 NOLOGGING  as
select  T0626919SP2.SBCLASS_KEY
from    T0626919SP2
union
select  T0626919SP3.SBCLASS_KEY
from    T0626919SP3


ANALYZE TABLE T0626919OJ4 ESTIMATE STATISTICS

ANALYZE TABLE T0626919OJ4 ESTIMATE STATISTICS

ANALYZE TABLE T0626919SP3 ESTIMATE STATISTICS

ANALYZE TABLE T0626919SP2 ESTIMATE STATISTICS

create table T0626919md1 NOLOGGING  as
select  T0626919OJ4.SBCLASS_KEY,
        T0626919SP2.SALESLYVALUE SALESLYVALUE,
        T0626919SP3.SALESVALUE SALESVALUE
```

```
from    T0626919OJ4 ,
        T0626919SP3 ,
        T0626919SP2
where   T0626919OJ4.SBCLASS_KEY = T0626919SP3.SBCLASS_KEY (+)
  and   T0626919OJ4.SBCLASS_KEY = T0626919SP2.SBCLASS_KEY (+)
drop table T0626919OJ4
drop table T0626919SP3
drop table T0626919SP2
create table T0626919md5 NOLOGGING  as
select  a1.SBCLASS_KEY,
        SUM(NVL(a1.F_CLRC_SLS_AMT,0)+NVL(a1.F_PRMTN_SLS_AMT,0)+NVL(a1.F_RGLR_SLS_AMT,0))
SALESYTDWEEKVALUE
from    SLS_SBC_LW_DM a1,
ORG_LOC_DM a2
where   a2.LOC_KEY = a1.LOC_KEY
  and   ((a1.WK_IDNT < 199929)
  and     (a1.WK_IDNT >= 199901))
  and     (((a2.REGN_KEY = 1))
  and       (((a1.SBCLASS_KEY < 150))))
group by        a1.SBCLASS_KEY

create table T0626919md6 NOLOGGING  as
select  a1.SBCLASS_KEY,
        SUM(NVL(a1.F_CLRC_SLS_AMT,0)+NVL(a1.F_PRMTN_SLS_AMT,0)+NVL(a1.F_RGLR_SLS_AMT,0))
SALESYTDWEEKLYVALU
from    SLS_SBC_LW_DM a1,
        ORG_LOC_DM a2
where   a2.LOC_KEY = a1.LOC_KEY
  and   ((a1.WK_IDNT < 199829)
  and     (a1.WK_IDNT >= 199801))
  and     (((a2.REGN_KEY = 1))
  and       (((a1.SBCLASS_KEY < 150))))
group by        a1.SBCLASS_KEY
create table T0626919md7 NOLOGGING  as
select  a1.SBCLASS_KEY,
        SUM(NVL(a1.F_CLRC_SLS_AMT,0)+NVL(a1.F_PRMTN_SLS_AMT,0)+NVL(a1.F_RGLR_SLS_AMT,0))
SALESSTDWEEKVALUE
from    SLS_SBC_LW_DM a1,
        ORG_LOC_DM a2
where   a2.LOC_KEY = a1.LOC_KEY
  and   ((a1.WK_IDNT < 199929)
  and     (a1.WK_IDNT >= 199927))
  and     (((a2.REGN_KEY = 1))
  and       (((a1.SBCLASS_KEY < 150))))
group by        a1.SBCLASS_KEY
create table T0626919md8 NOLOGGING  as
select  a1.SBCLASS_KEY,
        SUM(NVL(a1.F_CLRC_SLS_AMT,0)+NVL(a1.F_PRMTN_SLS_AMT,0)+NVL(a1.F_RGLR_SLS_AMT,0))
SALESSTDWEEKLYVALU
from    SLS_SBC_LW_DM a1,
        ORG_LOC_DM a2
where   a2.LOC_KEY = a1.LOC_KEY
  and   ((a1.WK_IDNT < 199829)
  and     (a1.WK_IDNT >= 199827))
  and     (((a2.REGN_KEY = 1))
  and       (((a1.SBCLASS_KEY < 150))))
group by        a1.SBCLASS_KEY
create table T0626919OJ9 NOLOGGING  as
select  T0626919md1.SBCLASS_KEY
from    T0626919md1
union
select  T0626919md5.SBCLASS_KEY
from    T0626919md5
union
select  T0626919md6.SBCLASS_KEY
from    T0626919md6
union
select  T0626919md7.SBCLASS_KEY
from    T0626919md7
union
select  T0626919md8.SBCLASS_KEY
```

```
from    T0626919md8


ANALYZE TABLE T0626919OJ9 ESTIMATE STATISTICS

ANALYZE TABLE T0626919OJ9 ESTIMATE STATISTICS

ANALYZE TABLE T0626919md5 ESTIMATE STATISTICS

ANALYZE TABLE T0626919md6 ESTIMATE STATISTICS

ANALYZE TABLE T0626919md7 ESTIMATE STATISTICS

ANALYZE TABLE T0626919md8 ESTIMATE STATISTICS

ANALYZE TABLE T0626919md1 ESTIMATE STATISTICS

create table   R0626919RT0 NOLOGGING  as
select  a7.SBCLASS_KEY SBCLASS_KEY,
        a7.SBCLASS_DESC||', '||a7.SBCLASS_IDNT SBCLASS_DESC,
        T0626919md1.SALESLYVALUE SALESLYVALUE,
        T0626919md5.SALESYTDWEEKVALUE SALESYTDWEEKVALUE,
        T0626919md6.SALESYTDWEEKLYVALU SALESYTDWEEKLYVALU,
        T0626919md7.SALESSTDWEEKVALUE SALESSTDWEEKVALUE,
        T0626919md8.SALESSTDWEEKLYVALU SALESSTDWEEKLYVALU,
        T0626919md1.SALESVALUE SALESVALUE,
        T0626919md1.SALESVALUE/T0626919md1.SALESLYVALUE SALESVALUETYVSLY1,
        (T0626919md7.SALESSTDWEEKVALUE/T0626919md8.SALESSTDWEEKLYVALU) SALESSTDWEEKVALUET,
        (T0626919md5.SALESYTDWEEKVALUE/T0626919md6.SALESYTDWEEKLYVALU) SALESYTDWEEKVALUET
from    T0626919OJ9 ,
        T0626919md5 ,
        T0626919md6 ,
        T0626919md7 ,
        T0626919md8 ,
        T0626919md1 ,
        PROD_SBC_DM a7
where   T0626919OJ9.SBCLASS_KEY = T0626919md5.SBCLASS_KEY (+)
  and   T0626919OJ9.SBCLASS_KEY = T0626919md6.SBCLASS_KEY (+)
  and   T0626919OJ9.SBCLASS_KEY = T0626919md7.SBCLASS_KEY (+)
  and   T0626919OJ9.SBCLASS_KEY = T0626919md8.SBCLASS_KEY (+)
  and   T0626919OJ9.SBCLASS_KEY = T0626919md1.SBCLASS_KEY (+)
  and   T0626919OJ9.SBCLASS_KEY = a7.SBCLASS_KEY
drop table T0626919OJ9
drop table T0626919md5
drop table T0626919md6
drop table T0626919md7
drop table T0626919md8
drop table T0626919md1
```