AMDOCS/HP
HP 9000 Superdome server
Benchmark 2001

Version 1.0

Aliza Abulafia

Table of Content

1. Content

This document presents a short summary of Amdocs-HP benchmark 2001. The benchmark took place between Jan 22nd 2001 to mar 1 at HP laboratories in Cupertino, California, USA.

1. Objectives

- Examine official performance of typical ensemble modules on the HP 9000 Superdome server.
- Examine Solutions for fast growing customers.
  - OPS new versions
  - HP 9000 Superdome server scalability
- Examine new Technologies
  - HP SureStore Disk Array HP SURESTORE DISK ARRAY XP512
  - Oracle 64Bit Version
  - XACCTUsage mediation software
  - HP Internet Usage Manager (IUM) mediation software
  - Shareplex
  - Veritas
  - FSO - file sharing option
  - Veritas QIO Database accelerator – raw capabilities in FS

2. Environment

1.1. Application

The tested application modules were

- MAF
- IP-MAF
- Bill Preparation – flat files version.

Those modules represent Ensemble's different activity patterns. This allows measuring of the essential factors having an impact on the whole system performance. MAF exercises both CPU and I/O while Bill Preparation mostly is a heavy database consumer.

All application's source was taken from living production sites.

1.2.Data

Inflation processes generated the Data for MAF and Bill. There were two different inflation processes for bill and MAF. In both cases the generated data is duplication of 'seed data'. The 'seed data' is a small complete set of data taken from real life.

1.2.1 Bill Preparation

- 20 Bill Cycles
- BAN/CTNs ratio 1:1.25
- Data Volumes (# of BANs)
  - Base 24,000,000
  - Large 48,000,000

1.2.2 MAF

- Data Volume (# of CDRs)
  - Base 70,000,000
  - Medium 140,000,000
  - Large 350,000,000

- Data Nature:
  Format: GSM1205
  Records type: '01' (MOC), '95' (header) and '96' (trailer).
  Number of records per file: 80000.
  Errors: 7%.
  Drops: 2 records per file (header+trailer).

1.2.3 IP-MAF

- Data Volume (# of CDRs)
  - Base 349,920,000

1.3. Hardware

1.3.1 Servers

- **High-End**: 2 x HP 9000 Superdome server Servers
  64 CPUs of 550 MHz, 128GB memory.
  48 to 192 hot swap PCI I/O slots.
  Utilization:
  Only 96 PCI I/O slots were installed and utilized.

- **Mid Range**: 1 x HP 9000 V2500 SERVER server.
  32 CPUs 440 MHZ, 32 GB memory.
  28 2X PCI Slots.

- **Mid-range**: 3 x HP 9000 N4000 SERVER servers.
  8 CPUs of 550 MHz, 32GB memory.
  12 PCI I/O slots.

1.3.2 Storage

- **HP SURESTORE DISK ARRAY XP512**
  HP SURESTORE DISK ARRAY XP512 was used as the main storage at the benchmark.
  - Support RAID 0+1 and RAID 5.
  - Host interface fiber channel 0-32
  - Transfer rate Fiber channel 100 MB/S
  - Total Array capacity 72GB to 37 TB,

  Utilization
  - RAID 0+1
  - Total of 8.8 TB for all servers. 92 spindles of 18GB and 100 spindles of 72GB were installed and utilized.
- **FC-60**
  FC-60 served as additional help storage.
  - Raid 0+1
  - 1-2 Fiber channel interfaces, each 100 MB/s peak performance.

- **FC-10**
  FC-10 served as additional help storage.
  - Raid 0+1
  - 10-slot fiber channel
  - Bandwidth of 100 MB/s

3. Test types

The Benchmark will include the following test types

1. **Base**
   Applications: Bill Prep, MAF, IP-MAF
   Single Server: HP 9000 Superdome server, HP 9000 N4000 SERVER
   Volume: "Base" (24,000,000 BANs Database for Bill-Prep, processing cycle of 1,200,000 BANs. 70,000,000 CDRs for MAF and 350,000,000 SDRs for IP-MAF)
   Goal: define best Tuning Configuration for each application and server.

2. **CPU**
   Application: Bill Prep, MAF
   Single Server: HP 9000 Superdome server
   Volume: "Base"
   Goal: Examine the scalability of the applications under different number of CPUs. The test runs with the same configuration as in the 'Base test', but each time with different number of CPUs and different parallel processing level (optimal for the number of CPUs).

3. **Volume**
   Application: Bill Prep, MAF
   Single Server: HP 9000 Superdome server
   Volume: Bill Prep : 24,000,000 and 48,000,000 BANs.
   MAF 70,000,000 , 140,000,000 and 350,000,000 CDRs.
   Goal: Examine modules' performance stability under increasing volumes of data.

4. **OPS**
   Application: Bill Prep
   Clustered Servers: 2 x HP 9000 Superdome server configured as Cluster, using 100BT network with Oracle 64 bit version and Hyper fabric network for Oracle 32BIT.
   Goal: Examine scalability, availability and manageability aspects of Oracle's OPS version.

5. **File Sharing**
   Veritas's FSO enables read only access to remote files on a local network.
   Application: MAF Collection module.
   Servers: HP 9000 Superdome server, HP 9000 N4000 SERVER.
   Goal: Examine efficiency of File sharing options.
   Network: 100BT.

   Collection module ran with 3 options:
   1. Reading CDRs files from local Disk.
   2. Reading CDRs files from remote disk via NFS.
   3. Reading CDRs files from Remote Disk via FSO.

6. **Shareplex**

SharePlex is Quest fast replication product for Oracle. It provides log-based replication between Oracle instances, enabling to replicate complete or partial database.

Application: Bill Prep
Servers: HP 9000 V2500 SERVER, HP 9000 N4000 SERVER
Goal: Test a variety of SharePlex replication functionality's under high rate and high volume examine its limitations.

Test description:

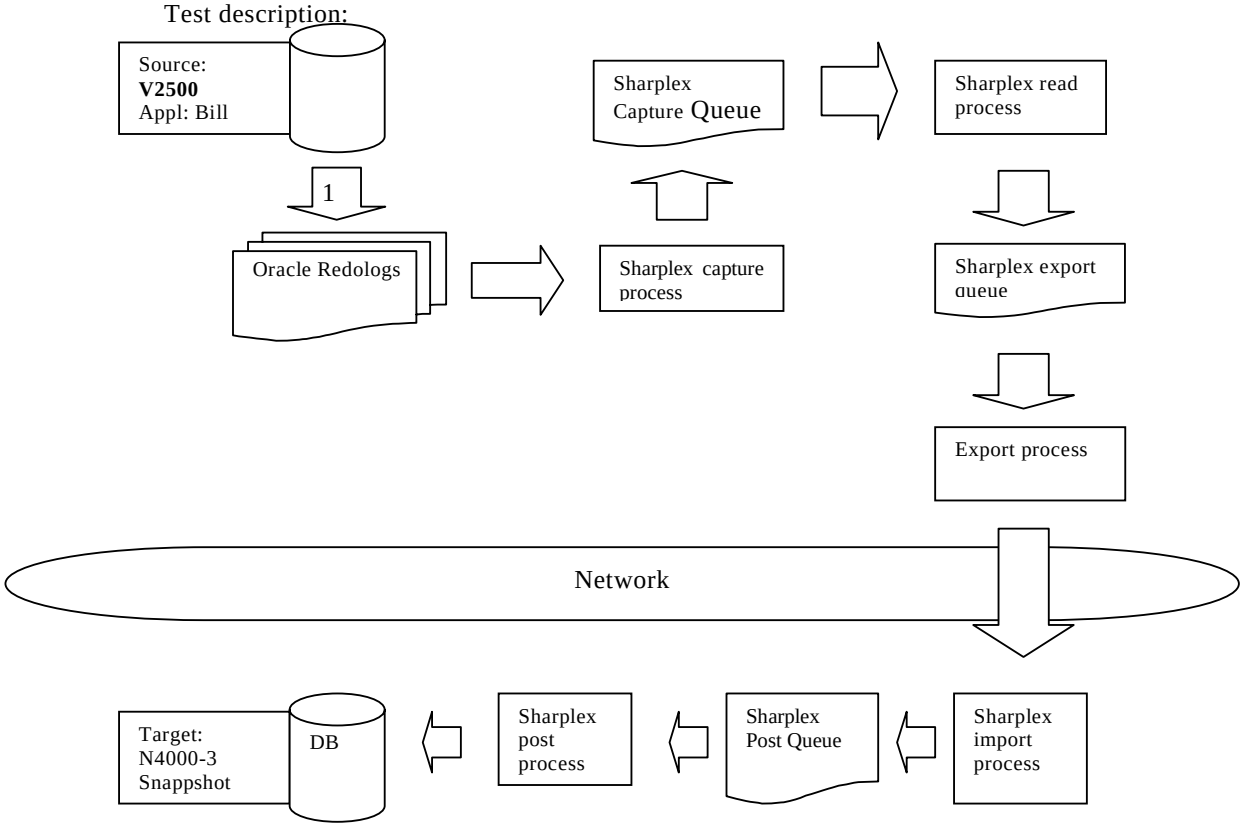

4. Technologies overview

SharePlex® (by Quest) – SharePlex is Quest's replication product for Oracle. It provides log-based replication between Oracle instances, enabling complete or partial replication of databases. SharePlex can be used for DRP purposes or for reporting database.

FSO (By Veritas) – enables read only access to remote files on a local network. FSO enables read only access without the overhead of TCP/IP connections between the machines. It can replace NFS or the need to copy files between servers in read only situations.

HP Internet Usage Manager (IUM) – mediation software, a key component of HP's "always-on" Internet infrastructure, can now extract customer usage information from Ericsson, Nokia and Motorola GPRS network equipment, TANTAU Software's Wireless Internet Platform, CMG's WAP Service Broker™ Product Line and Phone.com's WAP-based gateway, UP.Link server. This information enables content billing and marketing applications such as churn management and profit analysis.

XACCTUSAGE (part of Xacct Network to Business) - an engine which captures granular usage data from all network and service elements, in real time, from all layers of the network (routers, switches, and probes, and from service elements such as e-mail, Web, video, and application servers); dynamically synthesizes this data based on business policies; and transforms the data into actionable business intelligence.

Oracle Parallel Server (OPS) – a parallel server is designed to take advantage of HP cluster architecture by running multiple instances that "share" a single physical database. In appropriate applications, a parallel server enables access to a single database by users on multiple machines with increased performance in terms of speedup and improved scale-up to process larger workloads.

Oracle 64 Bit Version – running Oracle server 64 bit version, enables increasing Oracle data cache in memory, which reduces I/O and thus improves performance.

HP SURESTORE DISK ARRAY XP512 – HP's XP platform provides high availability, heterogeneity, and scalability (up to 512 disk, 37 TB and 32 host connections). RAID 0/1 and RAID 5 support optimizes the storage performance. 18 GB and 73 GB 10,000 RPM fiber channel drives provide scalability and reliability to match varying capacity requirements. The Crossbar architecture enables high performance and scalability. The 24x7x52 protective support identifies and resolves problems before they occur.

5. Results

1.4. Bill Prep

1.4.1 Base and CPU test

The following were executed on a single HP 9000 Superdome server, Using Oracle 8.0.5 client 32bit version and Oracle 8.1.7 64bit version, all with the same configuration.

Test No	Test Type	Server		No Of BANs	No Of BANs In The processed Cycle	Parallel Group	CPUs	Result
1373	Base	Single	**SuperDome**	24,000,000	1,200,000	128	64	**5,300,613**
1367	Cpu	Single	**SuperDome**	24,000,000	1,200,000	112	56	**4,837,626**
1368	Cpu	Single	**SuperDome**	24,000,000	1,200,000	96	48	**4,528,302**
1369	CPU	Single	**SuperDome**	24,000,000	1,200,000	64	32	**3,423,138**
1370	CPU	Single	**SuperDome**	24,000,000	1,200,000	48	24	**2,728,996**
1371	CPU	Single	**SuperDome**	24,000,000	1,200,000	32	16	**1,916,593**
1372	CPU	Single	**SuperDome**	24,000,000	1,200,000	16	8	**980,926**

The following was executed on a single HP 9000 Superdome server, Using Oracle 8.0.5 client 32bit version and Oracle 8.1.7 32bit version, all with the same configuration

Test No	Test Type	Server		No Of BANs	No Of BANs In The processed Cycle	Parallel Group	CPUs	Result
638	Base	Single	**N4000**	4,000,000	200,000	16	8	**875,912**

1.4.2 Volume Test

The following were executed on a single HP 9000 Superdome server, Using Oracle 8.0.5 client 32bit version and Oracle 8.1.7 32bit version, all with the same configuration.

Test No	Test Type	Server		No Of BANs	No Of BANs In The processed Cycle	Parallel Group	CPUs	Result
1011	Base	Single	SuperDome	12,000,000	600,000	120	64	**3,029,453**
1024	Volume	Single	SuperDome	24,000,000	1,200,000	120	64	**3,178,808**
1031	Volume	Single	SuperDome	48,000,000	2,400,000	120	64	**3,570,248**

1.4.3 OPS Test

The following were executed on 2 HP 9000 Superdome servers, configured as cluster, Using Oracle 8.0.5 client 32bit version and Oracle 8.1.7 32bit version, using hyper fabric network, all with the same configuration.
(Oracle OPS 64bit version has some problems working with hyperfabric connection)

Test No	Test Type	Server		No Of BANs	Parallel Group	CPUs	Result	First Side	Second Side
1206	Base	Cluster	SuperDome	12,000,000	120	64	**3,816,254**		
1204	Ops	Cluster	SuperDome	12,000,000	120	64	**6,419,510**	**3,209,510**	**3,210,000**

1.5.MAF

1.5.1 Base

The following were executed on a single HP 9000 Superdome server.

Test No	Test Type	Server	No Of CDRs	No Of Files	Records / File	Collections	Main Drivers	CPUs	CDRs / Hour
67	Base	Single **SuperDome**	140,040,000	2334	60,000	56	96	64	**369,065,886**

The following were executed on a single HP 9000 N4000 SERVER.

Test No	Test Type	Server	No Of CDRs	No Of Files	Records / File	Collections	Main Drivers	CPUs	CDRs / Hour
23	**Base**	Single **N4000**	**70,020,000**	**1167**	**60,000**	**8**	10	8	**48,132,901**

1.5.2 CPU

The following were executed on single HP 9000 Superdome server, all with the same configuration.

Test No	Test Type	Server	No Of CDRs	No Of Files	Records / File	Collections	Main Drivers	CPUs	CDRs / Hour
76	Base	Single **SuperDome**	140,040,000	2334	60,000	64	72	64	**339,034,297**
78	CPU		140,040,000	2334	60,000	56	71	56	**304,618,731**
81	CPU		140,040,000	2334	60,000	48	53	48	**267,450,398**
83	CPU		140,040,000	2334	60,000	32	36	32	**179,028,409**
86	CPU		140,040,000	2334	60,000	24	28	24	**133,230,444**
88	CPU		140,040,000	2334	60,000	16	20	16	**89,833,215**
95	CPU		140,040,000	2334	60,000	8	10	8	**44,165,046**

1.5.3 Volume

The following were executed on single HP 9000 Superdome server, all with the same configuration.

Test No	Test Type	Server	No Of CDRs	No Of Files	Records / File	Collections	Main Drivers	CPUs	CDRs / Hour
76	Volume	Single **SuperDome**	140,040,000	2334	60,000	64	72	64	**339,034,297**
89	Volume		350,040,000	5,834	60,000	64	72	64	**354,770,270**
92	Volume		70,020,000	1167	60,000	64	72	64	**355,531,735**

1.6. IP-MAF

IP-MAF was examined in batch and Online flows:

1. **Batch** : all files were ready, and IP-MAF processed all of them
2. **Online flow**: files were generated by one of the mediation devices (HP INTERNET USAGE MANAGER (IUM) or XACCT) on NT server, transferred to the HP 9000 Superdome server. IP-MAF processes on the HP 9000 Superdome server pooled the files, and processed them. Online flow total throughput is limited to mediation devices CDR's producing rate on the NT server. This rate was much lower than IP-MAF processing rate abilities on the HP 9000 Superdome server and HP 9000 N4000 SERVER.

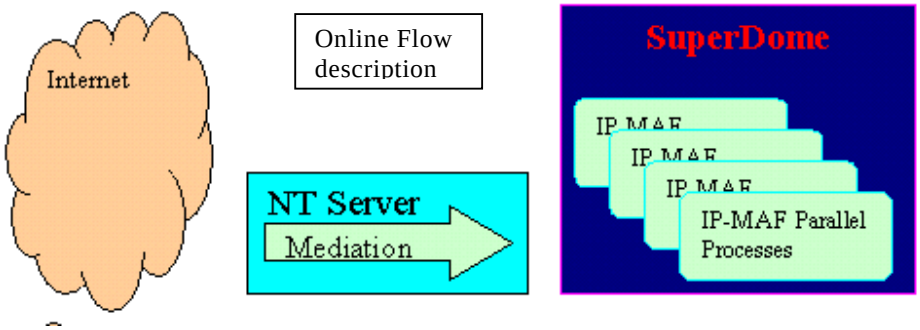

The following table presents IP-MAF batch processing rate on HP 9000 Superdome server and HP 9000 N4000 SERVER.

Test No	Test Type	Server	No Of CDRs	No Of Files	Records / File	Collections	Main Drivers	CPUs	CDRs / Hour
89	Base	Single SuperDome	349,920,000	972	360000	50	90	64	**2,131,492,386**
126	Base	Single N4000	70,020,000	1,167	60,000	8	8	8	**394,478,873**

1.7. File Sharing

The following test executed Collection module on a single HP 9000 Superdome server. 100BT network was used on the connection between HP 9000 Superdome server and HP 9000 N4000 SERVER.

In each test CDRs were read from a different resource as follow:

1. Local Disk.
2. Remote disk via NFS.
3. Reading CDRs files from Remote Disk via FSO.

FSO proved to be beneficial on a high loaded network environment. When the Network was not loaded (lower parallel level or same test having Collection run on HP 9000 N4000 SERVER) NFS performance were quite similar to those of the local.

File Sharing requires the writer to execute flash operation, and the reader to execute refreshes. Refreshing entire file system took 7 sec. Refreshing 1 directory took 1.2 sec.

Test No	Test Type	CDRs read from	Local Server	Remote Server	No Of CDRs	No Of Files	Records / File	Collections	CPUs	Total Time (sec)
93	File Sharing	Local Disk	SuperDome	N4000	140,040,000	389	360000	55	64	**58**
96	File Sharing	Remote via FSO	SuperDome	N4000	140,040,000	389	360000	50	64	**62**
98	File Sharing	Remote via NFS	SuperDome	N4000	140,040,000	389	360000	50	64	**1006**

1.8. Shareplex

Hardware:
Source: V-Class server (V182 - 32 CPU's and 32 GB of RAM)
Target: HP N-Class server (N3 - 4 CPU's and 32 GB of RAM)

Few feature were checked:

- Basic replication of some Ensemble tables.
  - Using single posting process, as in current production version.
  - Using multi-posting processes, as in next version.
- DB copy – copping database to a target machine, for replication.
- Data equator – compare in repair table replication when needed.

The new Shareplex version, with the multi threaded posting, is much more efficient.

Item	Descripotion
DB size	Relevant only for DB Copy. Inclued used and unused space in the database.
Used space	Used space withing the database
MPS	Messages per seconds. Amount of Inserts/Update/Delete that are applied at the target database.
Latency (catch up)	The amount of time that took to SharePlex to complete the posting process after bill prep completion.
Post Run Time	The total time that took to to SharePlex to complete the posting process

Test type	DB size	Used space	Inflation size	MPS	Bill-Prep run time	Post Run Time	Latency (catch up)	Time	Remarks
Tables replication(v→n)	-	25GB	1920000	450-500		31	19 min	-	Single Threaded posting
Tables replication(v→n)	-	25GB	1920000	800-1000		12	6 min	-	Multi Threaded Posting
Tables replication(n→v)	-	25GB	1920000	550		24	~	-	Single Threaded posting
Tables replication(n→v)	-	25GB	1920000	1000-1200		12	3 min	-	Multi Threaded Posting
Dbcopy(v→n)	64GB	25GB	-	-	-	-	-	101min	Compression=0% 687mb/min 12 threads

Dbcopy(v→n)	64GB	25GB	-	-	-	-	-	85min	Compression=50% 809mb/min 12 threads
Dbcopy(v→n)	64GB	25GB	-	-	-	-	-	106min	Compression=100% 661mb/min 12 threads
Data Equator(compare & repair deleted records)	-	-	881280	-	-	-	-	1h and 12min	Delete 1000 records on target (n)
Data Equator(compare & repair changed records)	-	-	881280	-	-	-	-	15min	Change 50 rows on target (n)
Data Equator(compare only)	-	-	920000	-	-	-	-	12min	Compare records.