

Case Study: Implementing the Oracle Grid Computing on Dell Hardware for Multiple ERP Applications



Kai Yu, System Engineer Senior Consultant, Dell Inc
Dan Brint, Supervising Analyst - Administrative Services
SUNY ITEC

Agenda

- ❖ Introduction to SUNY POC Project
- ❖ Oracle Enterprise Grid Model
- ❖ Grid Design and Implementation
- ❖ Database Grid Scalability
- ❖ Applications Test Methodology
- ❖ Performance Test Results
- ❖ QA



Introduction to SUNY POC Project

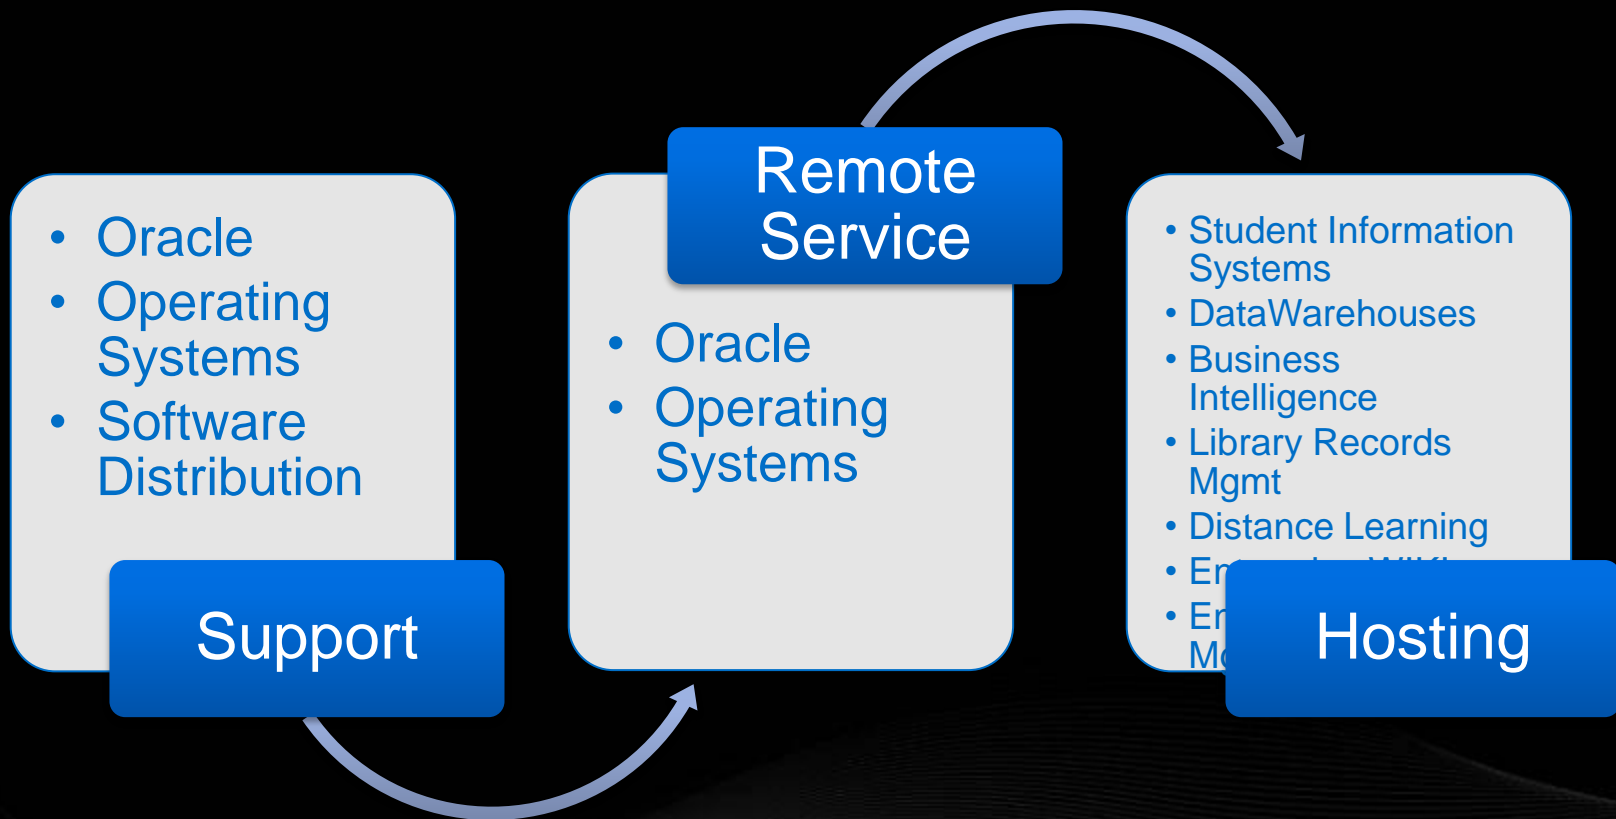
ITEC within SUNY

- ITEC is part of the AST – Alliance for Strategic Technologies
- Work with sister organizations:
 - SUNY Learning Network, SUNYNET, SUNY Connect
 - Center for Professional Development, SICAS
- Applications:
 - SunGuard: Banner – Student Info System, Luminis (Portal), ODS, EDW (data warehousing)
 - ExLibris: Aleph500 library automation
 - Atlassian: Confluence Enterprise Wiki
 - AtTask: @task Enterprise Project Management



Introduction to SUNY POC Project

History of ITEC



Introduction to SUNY POC Project

Challenges

Problems

A lot of systems

Too much variety

Difficult to Manage

Low Resource Utilization

Numerous Single Points of Failure

Desires

Fewer Systems

Less Variety

Ease of Management

Improved Resource Utilization

High Availability

Introduction to SUNY POC Project

Grid Computing to the Rescue

- Promises to address all our desires
- But – how to demonstrate it will work
- Without substantial investment?
- How about a Proof-of-Concept?

Desires

Fewer Systems

Less Variety

Ease of Management

Improved Resource Utilization

High Availability



Introduction to SUNY POC Project

Dell, SUNY ITEC, Oracle and SunGuard Joint

POC project:

- Consolidate 10 Banner ERP Applications
- In single Grid made of Dell hardware
- Capacity: 10 campuses, 170,000 students
- Performance Requirements:

11,000 students simultaneous actions

Response time: < 1 second at peak time

Throughputs: 70,000 courses registration
per hour during peak time



Oracle Enterprise Grid Model

- **Traditional Corporate Computing model:**
 - Consists of island-like systems
 - Little or no resource sharing
 - Hard to dynamically adapt changing workload
- **Enterprise Grid Computing Model**
 - Consolidate databases, applications, servers and storage connected by high speed network onto a common Grid platform.
 - Provide various server and storage resources as services to applications using databases
 - Integrate all the resources to allow provisioning on demand: dynamically provisioning to meet the workload needs

Oracle Enterprise Grid Model

- MegaGrid: a joint project by Oracle, Dell, EMC and Intel.
- Candidate for Grid model:
 - Multiple services by multiple tiered applications
 - Large number of resources: servers, network, storages
- A case study for SUNY ITEC Grid design for multiple ERP systems on Dell hardware
- Oracle 10g Features for Grid Computing
 - Clustering technology
 - Database services
 - Automatic Storage Management
 - Oracle Enterprise Manager Grid Control
 - Load balancing



Grid Design and Implementation

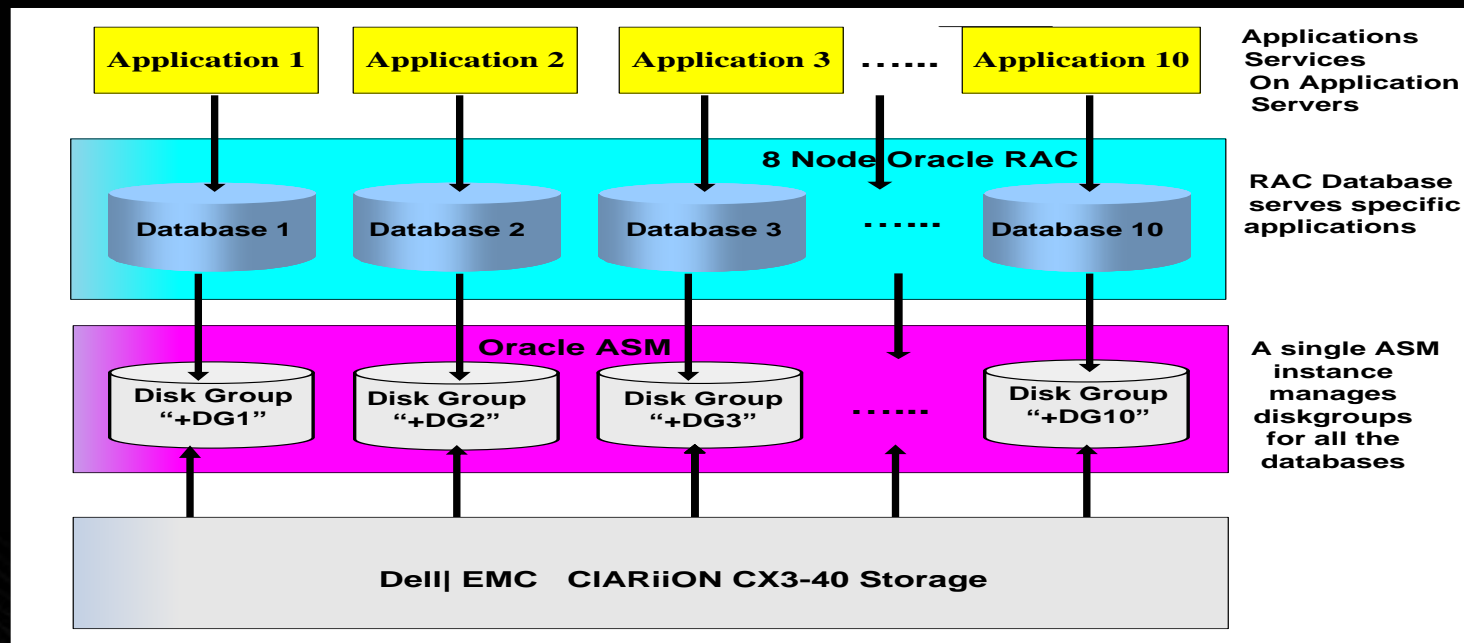
▪ Multi-Tier Proof of Concept (POC) Environment

– Server Grid:

- . Application servers offer application services using VM
- . Database servers offer database services for applications

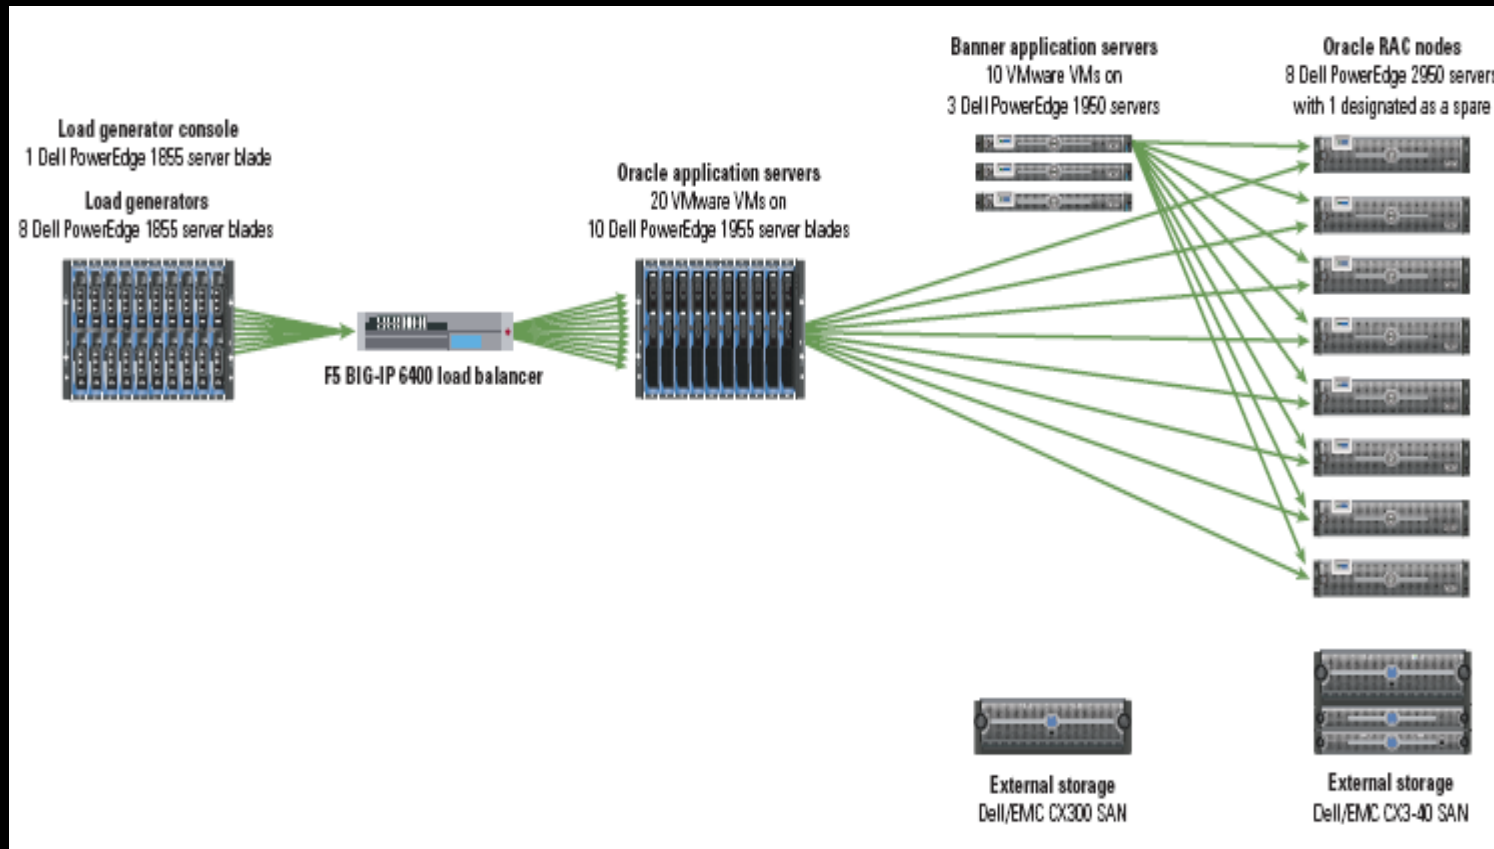
– Storage Grid:

- . ASM provides storage services for all the databases
- . ASM virtualizes the storage services using ASM diskgroups



Grid Design and Implementation

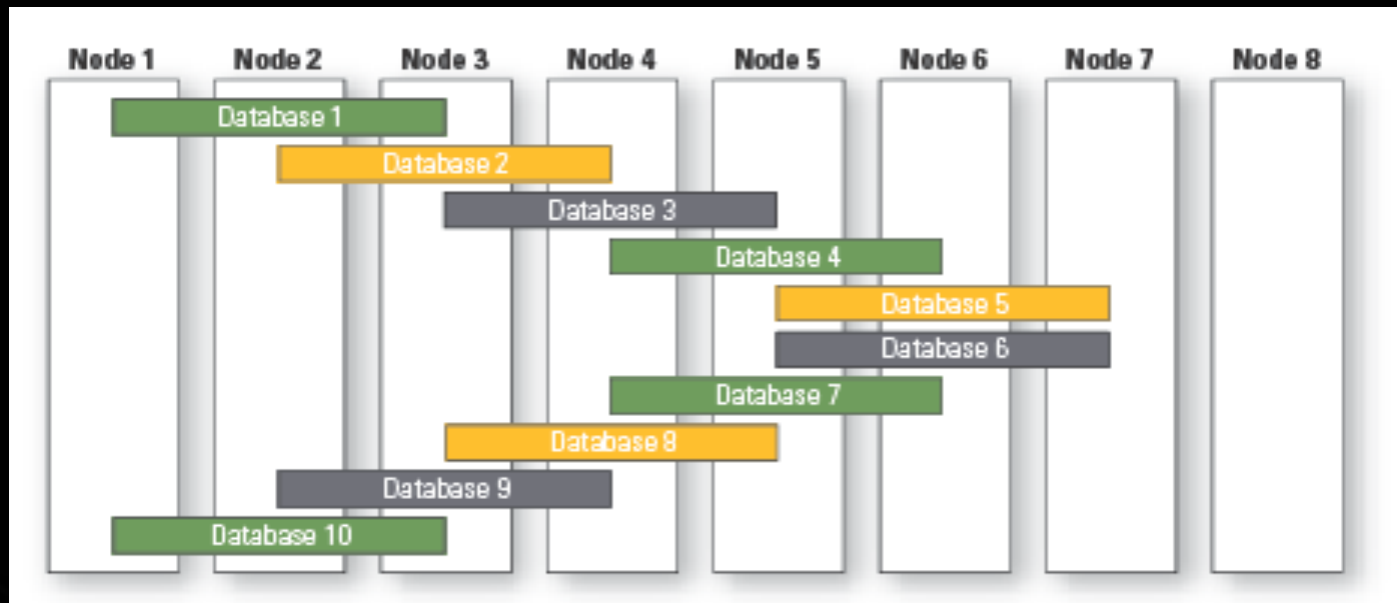
Multi-tier Hardware Configuration of the Grid



Grid Design and Implementation

▪ Database Grid Architecture Design :

- Consolidate 10 databases on a single 8 node RAC to provide 10 database services for 10 ERP applications
- Initial 3 instances for each database service
- Dynamic database instance reallocation
- Allow provisioning of additional nodes on demand
- Enterprise Manager Grid control for Grid Management



Grid Design and Implementation

▪ Database services

- Create 3 instances per each database service

```
$srvctl add service -d DB2 -s db2_srv -r db21, db22, db23
```

```
$srvctl status service -d db2
```

```
Service db2_s is running on instance(s) db23, db22, db21
```

- Connect to Database using services in tnsnames.ora

```
DB2_OLTP =
```

```
(DESCRIPTION =
```

```
(ADDRESS = (PROTOCOL = TCP)(HOST = bnode1-vip)(PORT = 1521))
```

```
(ADDRESS = (PROTOCOL = TCP)(HOST = bnode2-vip)(PORT = 1521))
```

```
(ADDRESS = (PROTOCOL = TCP)(HOST = bnode3-vip)(PORT = 1521))
```

```
(LOAD_BALANCE = yes)
```

```
(CONNECT_DATA =
```

```
(SERVER = DEDICATED)
```

```
(SERVICE_NAME = db2_srv.us.dell.com)
```

```
(FAILOVER_MODE =
```

```
(TYPE = SELECT)
```

```
(METHOD = BASIC)
```

```
(RETRIES = 180)
```

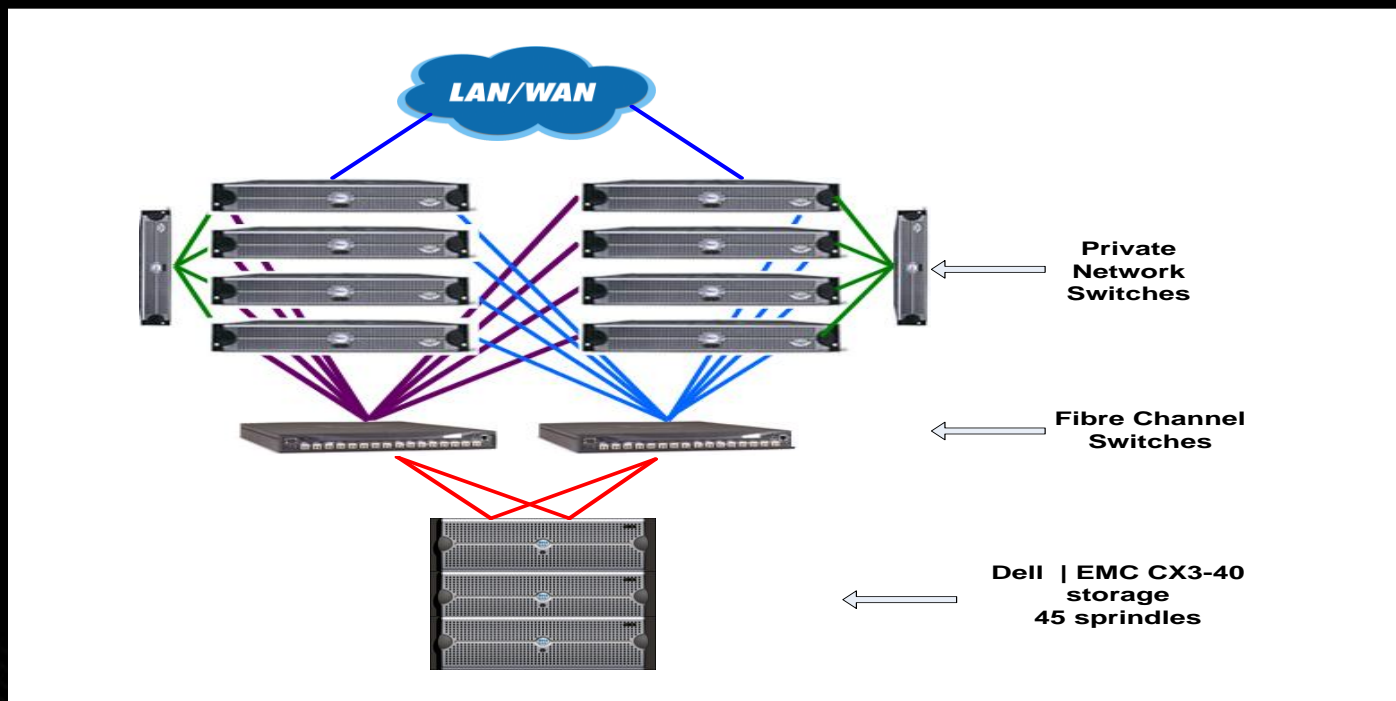
```
(DELAY = 5))))
```

- Dynamic database instance reallocation




Grid Design and Implementation

- **Database Grid built on 10g RAC**
 - Based on Dell | Oracle Tested and Validated Configuration:
 - Eight database servers
 - Two private interconnect network switches
 - Fibre Channel storage connections with dual HBAs
 - Dell EMC CX3-40 SAN with 45 spindles



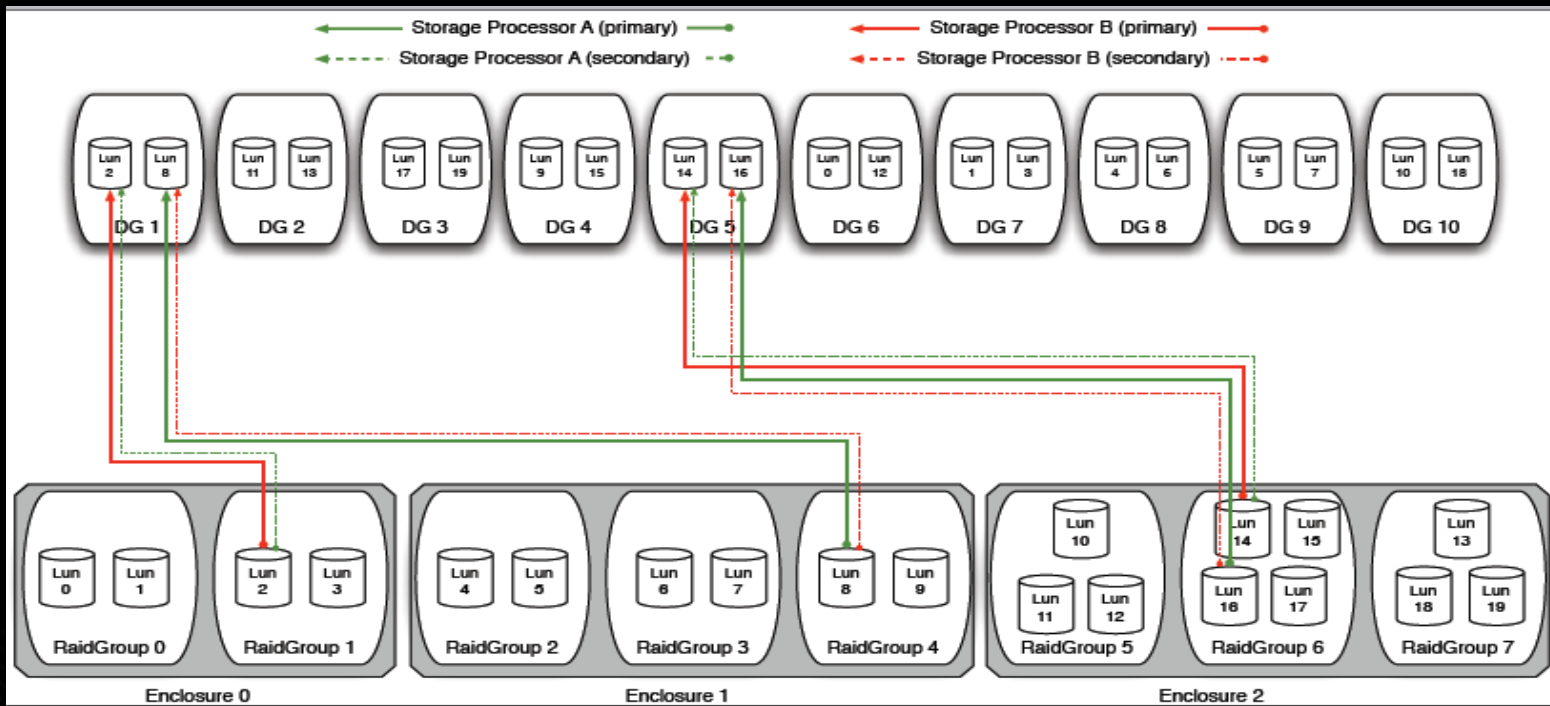
Grid Design and Implementation

- **Storage Grid Implementation for Storage Services**
 - Redundant IO Paths between PE 2950 and CX3-40 :
 - Two HBAs per server
 - Two Fiber Channel Switches
 - Two CX3-40 storage processors SPA and SPB



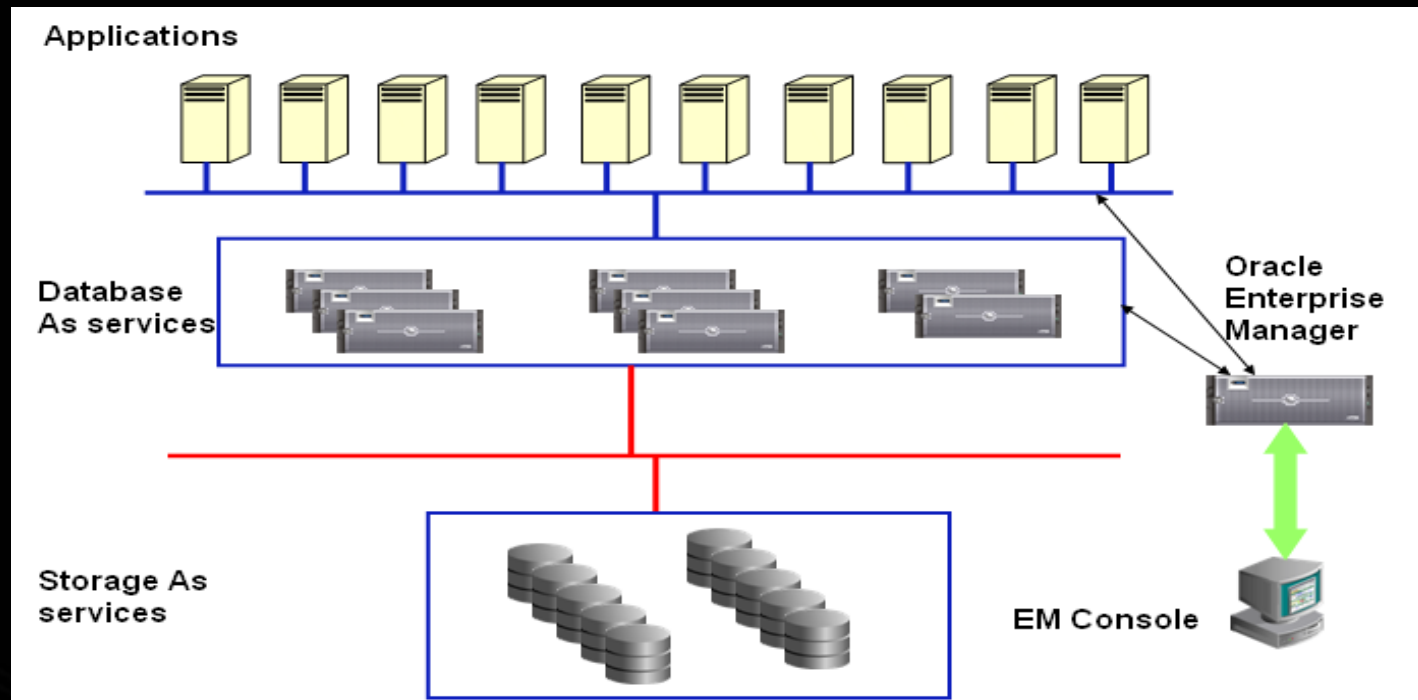
Grid Design and Implementation

- **Storage Grid Implementation for storage services**
 - A single ASM instance/server for all 10 databases
 - Ten ASM diskgroups created for 10 database
 - 7 Raid Groups: 4 disks per Raid Group in Raid 10
 - Each diskgroup: two LUNs from different SPs and and different Raid Groups



Grid Design and Implementation

- **Enterprise Manager Grid Control for Grid Management**
 - Oracle Enterprise Manager Grid Control(10.2.0.3)
 - Manage both Applications services and Database services
 - Dynamic Database Instance Management
 - Allow provisioning of additional nodes on demand
 - Performance Monitoring



Grid Design and Implementation

– Databases on the Grid

ORACLE Enterprise Manager 10g
Grid Control

Setup Preferences Help Logout
Home Targets Deployments Alerts Compliance Jobs Reports

Hosts | Databases | Application Servers | Web Applications | Services | Systems | Groups | All Targets

Databases

Page Refreshed Sep 26, 2007 7:39:18 PM CDT

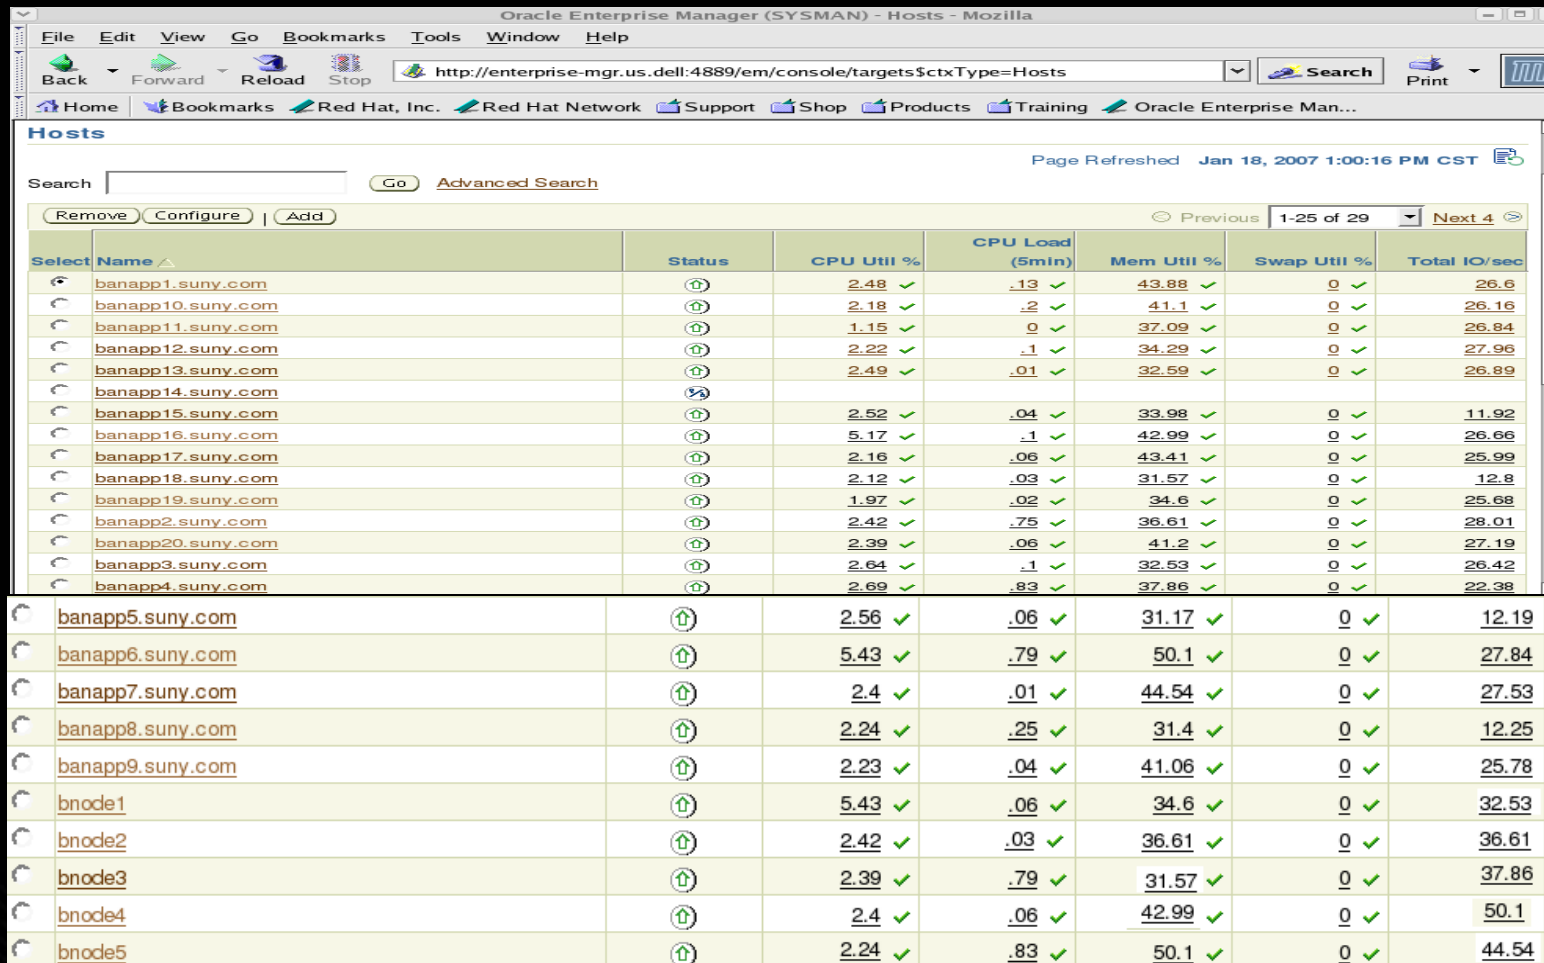
Search [Advanced Search](#)

Select	Name	Type	Status	Alerts	Policy Violations	Compliance Score (%)	Version	Sessions: CPU	Sessions: I/O	Sessions: Other	Instance CPU (%)
<input type="checkbox"/>	db1.us.dell.com	Cluster Database		0 4	21 29 7	89	10.2.0.3.0	n/a	n/a	n/a	n/a
<input type="checkbox"/>	db10.us.dell.com	Cluster Database		0 3	21 28 7	89		n/a	n/a	n/a	n/a
<input type="checkbox"/>	db2.us.dell.com	Cluster Database		0 3	20 28 6	89		n/a	n/a	n/a	n/a
<input type="checkbox"/>	db3.us.dell.com	Cluster Database		0 3	20 28 6	89	10.2.0.3.0	n/a	n/a	n/a	n/a
<input type="checkbox"/>	db4.us.dell.com	Cluster Database		0 3	20 29 6	88		n/a	n/a	n/a	n/a
<input type="checkbox"/>	db5.us.dell.com	Cluster Database		0 2	16 19 4	88		n/a	n/a	n/a	n/a
<input type="checkbox"/>	db6.us.dell.com	Cluster Database		0 1	8 10 2	91	10.2.0.3.0	n/a	n/a	n/a	n/a
<input type="checkbox"/>	db9.us.dell.com	Cluster Database		0 0	12 10 2	85	10.2.0.3.0	n/a	n/a	n/a	n/a
<input type="checkbox"/>	emrep.us.dell	Database Instance		0 2	10 11 4	91	10.1.0.4.0	.06	0		1.5



Grid Design and Implementation

- All the servers monitored by Oracle Enterprise Manager
 - Monitor the performance and workload of the entire Grid



The screenshot displays the Oracle Enterprise Manager (SYSMAN) interface for monitoring hosts. The browser address bar shows the URL: [http://enterprise-mgr.us.dell:4889/em/console/targets\\$cctxType=Hosts](http://enterprise-mgr.us.dell:4889/em/console/targets$cctxType=Hosts). The page title is "Hosts" and it indicates "Page Refreshed Jan 18, 2007 1:00:16 PM CST".

The main content area features a search bar and a table of host performance metrics. The table columns are: Select, Name, Status, CPU Util %, CPU Load (5min), Mem Util %, Swap Util %, and Total IO/sec. The table lists 25 hosts, including application servers (banapp) and nodes (bnode).

Select	Name	Status	CPU Util %	CPU Load (5min)	Mem Util %	Swap Util %	Total IO/sec
<input type="checkbox"/>	banapp1.suny.com		2.48 ✓	.13 ✓	43.88 ✓	0 ✓	26.6
<input type="checkbox"/>	banapp10.suny.com		2.18 ✓	.2 ✓	41.1 ✓	0 ✓	26.16
<input type="checkbox"/>	banapp11.suny.com		1.15 ✓	0 ✓	37.09 ✓	0 ✓	26.84
<input type="checkbox"/>	banapp12.suny.com		2.22 ✓	.1 ✓	34.29 ✓	0 ✓	27.96
<input type="checkbox"/>	banapp13.suny.com		2.49 ✓	.01 ✓	32.59 ✓	0 ✓	26.89
<input type="checkbox"/>	banapp14.suny.com						
<input type="checkbox"/>	banapp15.suny.com		2.52 ✓	.04 ✓	33.98 ✓	0 ✓	11.92
<input type="checkbox"/>	banapp16.suny.com		5.17 ✓	.1 ✓	42.99 ✓	0 ✓	26.66
<input type="checkbox"/>	banapp17.suny.com		2.16 ✓	.06 ✓	43.41 ✓	0 ✓	25.99
<input type="checkbox"/>	banapp18.suny.com		2.12 ✓	.03 ✓	31.57 ✓	0 ✓	12.8
<input type="checkbox"/>	banapp19.suny.com		1.97 ✓	.02 ✓	34.6 ✓	0 ✓	25.68
<input type="checkbox"/>	banapp20.suny.com		2.42 ✓	.75 ✓	36.61 ✓	0 ✓	28.01
<input type="checkbox"/>	banapp2.suny.com		2.39 ✓	.06 ✓	41.2 ✓	0 ✓	27.19
<input type="checkbox"/>	banapp3.suny.com		2.64 ✓	.1 ✓	32.53 ✓	0 ✓	26.42
<input type="checkbox"/>	banapp4.suny.com		2.69 ✓	.83 ✓	37.86 ✓	0 ✓	22.38
<input type="checkbox"/>	banapp5.suny.com		2.56 ✓	.06 ✓	31.17 ✓	0 ✓	12.19
<input type="checkbox"/>	banapp6.suny.com		5.43 ✓	.79 ✓	50.1 ✓	0 ✓	27.84
<input type="checkbox"/>	banapp7.suny.com		2.4 ✓	.01 ✓	44.54 ✓	0 ✓	27.53
<input type="checkbox"/>	banapp8.suny.com		2.24 ✓	.25 ✓	31.4 ✓	0 ✓	12.25
<input type="checkbox"/>	banapp9.suny.com		2.23 ✓	.04 ✓	41.06 ✓	0 ✓	25.78
<input type="checkbox"/>	bnode1		5.43 ✓	.06 ✓	34.6 ✓	0 ✓	32.53
<input type="checkbox"/>	bnode2		2.42 ✓	.03 ✓	36.61 ✓	0 ✓	36.61
<input type="checkbox"/>	bnode3		2.39 ✓	.79 ✓	31.57 ✓	0 ✓	37.86
<input type="checkbox"/>	bnode4		2.4 ✓	.06 ✓	42.99 ✓	0 ✓	50.1
<input type="checkbox"/>	bnode5		2.24 ✓	.83 ✓	50.1 ✓	0 ✓	44.54



Database Grid Scalability

- Use Grid Control to test:
 - Dynamic Database Instance Management
 - Dynamic Scale out Grid
- Dynamic Database Instance Management
 - Add instance to a database service:
For example, add the four instance db44 to db4:

ORACLE Enterprise Manager 10g
Grid Control

Home Targets Deployment

Hosts | Databases | Application Servers | Web Applications | Services | Systems | Groups | All Targets | F5 Big

Cluster Credentials Host Review

Add Instance: Host

Name of the Database Instance to be added

The following list of hosts have database software installed and are currently configured for this cluster. Select a host to which you want to add a database instance. This host should have access to the shared storage used by this database.

Select	Host	Existing Database Instances
<input type="radio"/>	bnode1	
<input type="radio"/>	bnode2	
<input checked="" type="radio"/>	bnode3	
<input type="radio"/>	bnode4	db41
<input type="radio"/>	bnode5	db42
<input type="radio"/>	bnode6	db43
<input type="radio"/>	bnode7	

ORACLE Enterprise Manager 10g
Grid Control

Home Targets Deployment

Hosts | Databases | Application Servers | Web Applications | Services | Systems | Groups | All Targets | F5 Big

Cluster: crs > Cluster Database: db4.us.dell.com >

Add Instance: Cluster Credentials

This wizard guides you through the steps required to add a database instance to the cluster database. The wizard will add the instance to the cluster database. At each step, checks are performed to ensure all prerequisites are satisfied for database instance creation.

Cluster Credentials

Enter the cluster credentials for the install owner of the Oracle Home from which the cluster database is being added.

* Username

* Password

ASM Credentials

Enter the credentials of the SYSDBA user running the asm instance.

* Username

* Password

ASM Instance **+ASM6_bnode6**

* Connect As

Save as Preferred Credential



Database Grid Scalability

- Add instance to a database

ORACLE Enterprise Manager 10g
Grid Control

Home Targets Deployments Alerts Compliance

Job Activity | Job Library

Job Run: ADDINSTANCE_DB4.US.DELL.COM_000027 >

Execution: db4.us.dell.com

Page Refreshed Aug 29, 2008 3:48:46 PM CDT [Delete Run]

Summary

The Stop and Suspend operations will wait for the current step to complete. A suspended job can be resumed later, at the next step.

Status	Running	Type	Add Instance
Scheduled	Aug 29, 2008 3:47:47 PM (UTC-05:00)	Owner	SYSMAN
Started	Aug 29, 2008 3:47:47 PM (UTC-05:00)	Description	AddInstance Job:
Ended		Oracle Home	/opt/oracle/product/10.2.0/db_1
Elapsed Time			
Notification	No		

Targets:

Status:

[Expand All](#) | [Collapse All](#)

Name	Targets	Status	Started	Ended
Execution: db4.us.dell.com	db4.us.dell.com	Running	Aug 29, 2008 3:47:47 PM (UTC-05:00)	
Step: add_instance	db4.us.dell.com_db43	Running	Aug 29, 2008 3:47:52 PM (UTC-05:00)	

```
oracle@bnode3:~  
[oracle@bnode3 ~]$ srvctl status database -d db4  
Instance db41 is running on node bnode4  
Instance db42 is running on node bnode5  
Instance db43 is running on node bnode6  
[oracle@bnode3 ~]$
```

```
oracle@bnode3:~  
[oracle@bnode3 ~]$ srvctl status database -d db4  
Instance db41 is running on node bnode4  
Instance db42 is running on node bnode5  
Instance db43 is running on node bnode6  
Instance db44 is running on node bnode3
```



Database Grid Scalability

- Drop an instance to a database
For example, drop the four instance db44 from db4:

ORACLE Enterprise Manager 10g
Grid Control

Home | Targets | Deployments | Alerts | Compliance

Delete Instance: Database Instance

The following list of database instances are currently part of this cluster database. Select the database instance and the files related to this database instance will be deleted.

Select	Database Instance	Host
<input type="radio"/>	db41	bnode4
<input type="radio"/>	db42	bnode5
<input type="radio"/>	db43	bnode6
<input checked="" type="radio"/>	db44	bnode3

ORACLE Enterprise Manager 10g
Grid Control

Home | Targets | Deployments | Alerts | Compliance

Job Run: DELETEINSTANCE_DB4.US.DELL.COM_000028

Page Refreshed Aug 29, 2008 4:11:38 PM CDT [Delete Run](#)

Scheduled **Aug 29, 2008 4:06:34 PM (UTC-05:00)**
Targets **db4.us.dell.com**

Type **Delete Instance**
Owner **SYSMAN**
Description **DeleteInstance Job:**

Executions

Targets
Status **All**

[Show All Details](#) | [Hide All Details](#)

Select	Details	Targets	Status	Started	Ended
<input checked="" type="radio"/>	Show	db4.us.dell.com	Succeeded	Aug 29, 2008 4:06:34 PM (UTC-05:00)	Aug 29, 2008 4:07:33 PM (UTC-05:00)

```
oracle@bnode3:~  
[oracle@bnode3 ~]$ srvctl status database -d db4  
Instance db41 is running on node bnode4  
Instance db42 is running on node bnode5  
Instance db43 is running on node bnode6  
[oracle@bnode3 ~]$
```



Database Grid Scalability

- **Dynamic scale out the Grid to a new node**
 - Prepare a new node (OS, network, access to the shared storage, EM agent install)
 - Scale out the RAC to a new node using EM Provisioning Pack
 - Use “One Click Extend Cluster Database” procedure
 - Predefined deployment procedure in EM provisioning

ORACLE Enterprise Manager 10g
Grid Control

Home Targets **Deployments** Alerts Compliance Jobs Reports

General | Provisioning

Deployment Procedure Manager

Procedures Procedure Completion Status Recycle Bin

Deployment procedures are best practices provided by Oracle for various Provisioning and Patching tasks. Procedures created by Oracle cannot be edited, but can be extended using 'Create Like', so that you can customize the procedure to fit your environment. For more details click Help.

Search Text Fields [Advanced Search](#)

View Run Edit Create Like Revert Delete | Upload Previous 1-25 of 33 Next 8

Select	Procedure	Type	Description	Last Modified By	Version	Last Updated
<input type="radio"/>	Oracle Clusterware / RAC Provisioning For Windows	RAC Provisioning	This procedure assists in installing/cloning and configuring a cluster database (a Real Application Cluster - RAC database) on a selection of hosts as specified by the Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide. i	Oracle	3.46	Sep 24, 2007 2:47:31 AM CDT
<input type="radio"/>	Oracle Clusterware / RAC Provisioning For UNIX	RAC Provisioning	This procedure assists in installing/cloning and configuring a cluster database (a Real Application Cluster - RAC database) on a selection of hosts as specified by the Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide. i	Oracle	3.46	Sep 24, 2007 2:47:30 AM CDT
<input checked="" type="radio"/>	One Click Extend Cluster Database	RAC Provisioning	This procedure will extend an existing cluster database to a set of new nodes. Oracle Clusterware and Oracle Database will be extended and configured by the procedure. i	Oracle	3.46	Sep 24, 2007 2:47:29 AM CDT
<input type="radio"/>	Delete/Scale down	RAC Provisioning	This procedure deletes nodes from Oracle Real	Oracle	3.46	Sep 24, 2007 2:47:28 AM CDT



Database Grid Scalability

- Select the database to be extended
- Select the new server and fill the server information
- Submit the RAC extend Job

Extend Real Application Clusters

Select Real Application Clusters (RAC)

Select the Oracle Real Application Clusters (RAC) you wish to extend. The Clusterware and Automatic Storage Management (ASM) will also be extended if these do not already exist.

Search

[Expand All](#) | [Collapse All](#)

Select Name	Member Nodes	Oracle Home	Platform	Product
Available Cluster Databases				
Previous 1 - 5 of 10				
<input checked="" type="radio"/> db5.us.dell.com (2)	bnode6, bnode5	/opt/oracle/product/10.2.0/db_1	Red Hat Enterprise Linux AS release 4 (Nahant Update 5)	Oracle Database 10.2.0.3.0
<input type="radio"/> db6.us.dell.com (2)	bnode6, bnode5	/opt/oracle/product/10.2.0/db_1	Red Hat Enterprise Linux AS release 4 (Nahant Update 5)	Oracle Database 10.2.0.3.0
<input type="radio"/> db7.us.dell.com (3)	bnode4, bnode5, bnode6	/opt/oracle/product/10.2.0/db_1	Red Hat Enterprise Linux AS release 4 (Nahant Update 5)	Oracle Database 10.2.0.3.0
<input type="radio"/> db8.us.dell.com (3)	bnode4, bnode5, bnode3	/opt/oracle/product/10.2.0/db_1	Red Hat Enterprise Linux AS release 4 (Nahant Update 5)	Oracle Database 10.2.0.3.0
<input type="radio"/> db9.us.dell.com (4)	bnode2, bnode3, bnode4, bnode1	/opt/oracle/product/10.2.0/db_1	Red Hat Enterprise Linux AS release 4 (Nahant Update 5)	Oracle Database 10.2.0.3.0
Next				

[Reference host options - \(bnode6\)](#)

Select New Nodes

Select the destination hosts and enter the respective Virtual Node Names.

[Hide Options](#)

Host	Private Node Name	Private IP (Optional)	Virtual Node Name	Virtual IP (Optional)	Working Directory	Remove
bnode7	<input type="text" value="bnode7-priv"/>	<input type="text" value="10.1.17.94"/>	<input type="text" value="bnode7-vip"/>	<input type="text" value="155.1.18.90"/>	<input type="text" value="/tmp"/>	<input type="button" value="Remove"/>



Database Grid Scalability

- Specify the reference host and the node to add
- Specify the credentials and submit the job

Reference host options - (bnode6)

Reference host

Select the reference host for extend.

Working directory

A working directory on the source host, is required to stage files for cloning. Supply the name of an existing directory on the host.

Home Location	Files To Exclude
/crs/oracle/oracle/product/10.2.0/crs	log,crs/init,racg/dump,srvm/log,cdata,EMStagedPatches
/opt/oracle/product/10.2.0/db_1	log,EMStagedPatches,oratab,*.trc,*.dbf,cdump
/opt/oracle/product/10.2.0/db_1	log,EMStagedPatches,oratab,*.trc,*.dbf,cdump

Oracle Home Shared Storage Options

Select New Nodes

Select the destination hosts and enter the respective Virtual Node Names.

Host	Private Node Name	Private IP (Optional)	Virtual Node Name	Virtual IP (Optional)	Working Directory
bnode7	bnode7-priv	10.1.17.94	bnode7-vip	155.1.18.90	/tmp

TIP Select hosts that are managed by agents of version 10.2.0.3.0



Database Grid Scalability

- Check the job status

The screenshot displays the Oracle Enterprise Manager 10g Grid Control interface. The top navigation bar includes 'Home', 'Targets', 'Deployments', 'Alerts', 'Compliance', 'Jobs', and 'Reports'. The 'Deployments' tab is active, and the 'Provisioning' sub-tab is selected. The main content area shows the 'Procedure Completion Status' for a job. The status is 'Succeeded'. The job details include: Run: crs_db5.us.dell.com_2006-12-03_12-01-24PM, Procedure: One Click Extend Cluster Database, Procedure Version: 3.46, Error Handling Mode: Stop On Error, Status: Succeeded, Owner: SYSMAN, Created On: Dec 3, 2006 12:01:38 PM CST, Scheduled: Dec 3, 2006 12:01:43 PM CST, Start Date: Dec 3, 2006 12:01:43 PM CST, Last Updated: Dec 4, 2006 12:42:14 PM CST, Completed Date: Dec 4, 2006 12:42:14 PM CST, and Elapsed Time: 88831 Seconds. Below the general information, the 'Status Detail' section shows a table of steps, all of which are 'Succeeded'.

General Information

Run	crs_db5.us.dell.com_2006-12-03_12-01-24PM	Created On	Dec 3, 2006 12:01:38 PM CST
Procedure	One Click Extend Cluster Database	Scheduled	Dec 3, 2006 12:01:43 PM CST
Procedure Version	3.46	Start Date	Dec 3, 2006 12:01:43 PM CST
Error Handling Mode	Stop On Error	Last Updated	Dec 4, 2006 12:42:14 PM CST
Status	Succeeded	Completed Date	Dec 4, 2006 12:42:14 PM CST
Owner	SYSMAN	Elapsed Time	88831 Seconds

Status Detail

Name	Status	Type	Description
One Click Extend Cluster Database	Succeeded		This procedure will extend an existing cluster database to a set of new nodes. Oracle Clusterware and Oracle Database will be extended and configured by the procedure.
Initialize Deployment Procedure	Succeeded	Computational	Initializes the current Deployment Procedure execution. Derived variables are set with computations. Do not disable or delete this step.
Create directory	Succeeded	Parallel	Creates first-level directories under / (requires root privileges).
Configure agents for all nodes	Succeeded	Parallel	Configures agents for all nodes for the job.



Database Grid Scalability

Before adding bnode7

Select	Name	Member Nodes
▼	Available Cluster Databases	
▲	Previous 1 - 5 of 11	
○	▶ db4.us.dell.com (3)	bnode5, bnode6, bnode4
⊙	▼ db5.us.dell.com (2)	bnode6, bnode5
	crs (6)	bnode1, bnode4, bnode5, bnode6, bnode2, bnode3
	+ASM6_bnode6 (2)	bnode6, bnode5
○	▶ db6.us.dell.com (2)	bnode6, bnode5
○	▶ db7.us.dell.com (3)	bnode5, bnode4, bnode6
○	▶ db8.us.dell.com (3)	bnode4, bnode5, bnode3

After adding bnode7

Select	Name	Member Nodes
▼	Available Cluster Databases	
▲	Previous 1 - 5 of 11	
○	▶ db4.us.dell.com (3)	bnode5, bnode6, bnode4
○	▼ db5.us.dell.com (3)	bnode7, bnode6, bnode5
	crs (7)	bnode1, bnode2, bnode3, bnode4, bnode5, bnode6, bnode7
	+ASM7_bnode7 (3)	bnode7, bnode6, bnode5
⊙	▶ db6.us.dell.com (2)	bnode6, bnode5
○	▶ db7.us.dell.com (3)	bnode5, bnode4, bnode6
○	▶ db8.us.dell.com (3)	bnode4, bnode5, bnode3

bnode7 added to db5

```

root@bnode7:~
[oracle@bnode7 ~]$ srvctl status database -d db5
Instance db51 is running on node bnode5
Instance db52 is running on node bnode6
Instance db53 is running on node bnode7
[oracle@bnode7 ~]$ srvctl status database -d db6
Instance db61 is running on node bnode6
Instance db64 is running on node bnode5
Instance db63 is running on node bnode7
[oracle@bnode7 ~]$ srvctl status nodeapps -n bnode7
VIP is running on node: bnode7
GSD is running on node: bnode7
Listener is running on node: bnode7
ONS daemon is running on node: bnode7
[oracle@bnode7 ~]$ ps -ef | grep pmon
oracle      3890      1289    0 23:30 pts/2        00:00:00 grep pmon
oracle      5645          1    0 14:23 ?                00:00:00 asm_pmon_+ASM7
oracle     10776          1    0 14:24 ?                00:00:00 ora_pmon_db53
oracle     12346          1    0 14:24 ?                00:00:00 ora_pmon_db63
    
```



Database Grid Scalability

▪ Lessons Learned

- CRS inventory was missing on the existing RAC nodes
- Required crs inventory recreated before adding new node
- Steps to recreate crs inventory from Oracle engineers
 - Shutdown clusterware and the databases
 - Backup the databases
 - Backup OCR and Votingdisk
 - Cleanup Oracle clusterware on all the nodes
 - Reinstall Oracle clusterware on all the nodes to have new crs inventory
 - Restore the OCR from its backup
 - Start all the database services using new clusterware
- Grid control issue with date
 - The tests required resetting system date
 - Inconsistent date between Grid control and targeted servers caused issue with Grid control



Applications Test Methodology

POC Test Basis

- Based on work done in 2006 with Texas Tech
- We narrowed the breadth of tests
- Increased the user load from 1 campus to 10
- Focused on peak user load: student registration
- Use LoadRunner workload generators to simulate simultaneous user actions



Applications Test Methodology

Users

Function	# Virtual Users <u>per</u> Large DB	# Virtual Users <u>per</u> Small DB
Student Registration (A)	320	80
Student Registration (B)	320	80
View Class List	200	50
Add/Drop Classes	160	40
View Grades	800	200
Total	1,800	450



Applications Testing methodology


Testing Points (users)

Function	1+1	3+3	5+5
Student Registration (A)	400	1200	2000
Student Registration (B)	400	1200	2000
View Class List	250	750	1250
Add/Drop Classes	200	600	1000
View Grades	1000	3000	5000
Total	2,250	6,750	11,250



Applications Testing methodology

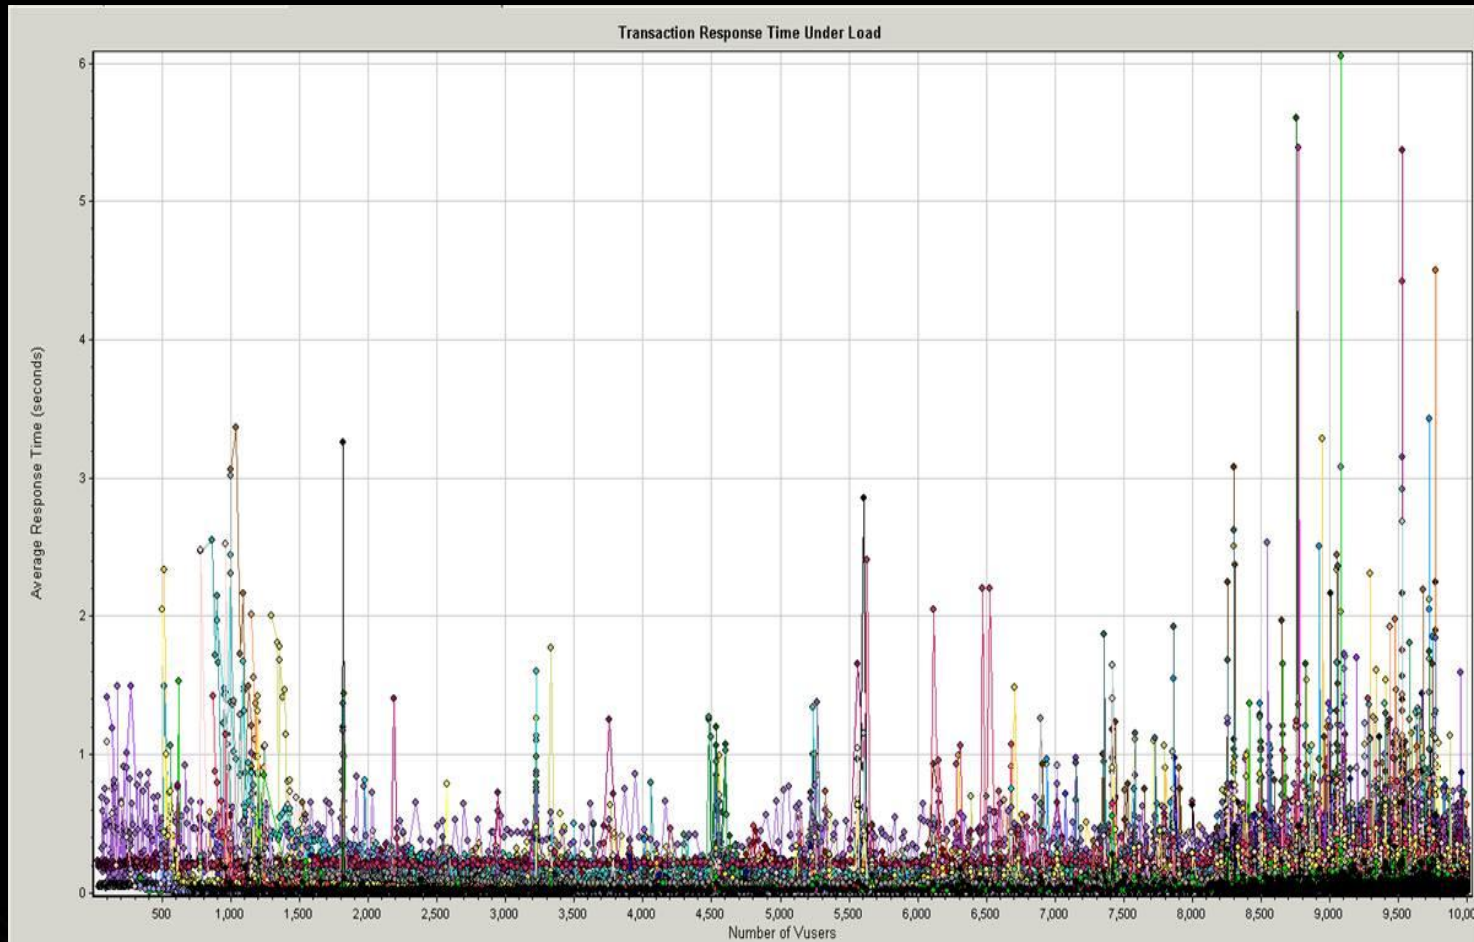
5+5



- 5 large + 5 small
- 11,250 users
- Bell-curve distribution of load (2,4,6,4,2)


Performance Test Results

Response time vs # of users



Performance Test Results

Traditional Registration Response Time



Performance Test Results

Traditional Registration Total Response Time

Workload	Total Response Time for Traditional Student Registration Process	Number of Database Servers	Average CPU Utilization
1 Large Campus + 1 Small Campus	0.7 sec	3	25%
3 Large Campuses + 3 Small Campuses	1.1 sec	5	31%
5 Large Campuses + 5 Small Campuses	0.8 sec	6	30%

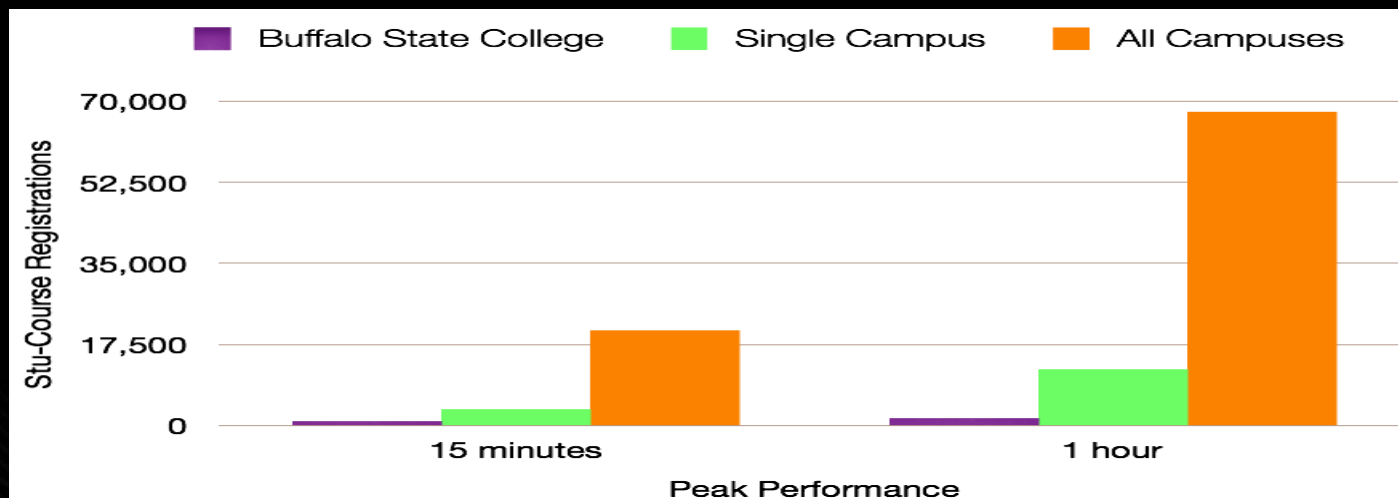


Performance Test Results

Application Throughput:

Student Registration Real World Comparison

	Dell-SUNY POC	Buffalo State (November 14, 2007)	Multiplication Factor
Maximum Student Course Registrations in a single 15 minute period	20,267	1,172	18x
Maximum Student Course Registrations over a one hour period	67,807	1,820	37x




Performance Test Results

Application Throughput:

Student Registration Real World Comparison

	Dell-SUNY POC	Buffalo State (November 14, 2007)	Multiplication Factor
Maximum Student Course Registrations in a single 15 minute period	20,267	1,172	18x
Maximum Student Course Registrations over a one hour period	67,807	1,820	37x



Conclusions

The Grid on Dell hardware can be scaled out to handle the needs of the multiple large campuses with the capacity of handling:

- Users loads from 11,000 simultaneous users actions with sub-second response times
- 70,000 courses registered in a hour, 37 times of a SUNY school of 11,000 students
- More than 6 sample schools of 175,000 students total
- The database instance on the Grid can be dynamically added, dropped and relocated on demand
- Grid infrastructure itself can be dynamically scaled out on demand



Acknowledgements

We would like to acknowledge the contributions by the following people to the POC project

- ITEC: Daniel Brint, Luke Clutter, Jeff Caughel, Paul Hebert, Erik Snyder, Mike Radomski, Mike Notarius, Ron Brown
- Sicas: Sue Smith, Pete Andrusyszyn, Charlie Young
- Dell: Aaron Burns, Orlando Gallegos, Ujjwal Rajbhandari, Roger Lopez, Jack Davis, Ericka Vilabobos, Dave Jeffe
- Oracle: Rajat Nigam, Tom Kopec, Austin Laird, John MacDougal, Marc Kelberman
- SunGrard: Pat Masterson, Rob McQue, Jennifer O'Brian
- F5: Keith Keeling



References

1. Scaling SunGuard Higher Education Banner Software on Dell Hardware, Dave Jaffe, Kai Yu, Dan Brint, Dell Power Solutions,, August 2008.
<http://www.dell.com/downloads/global/power/ps3q08-20080283-SUNY.pdf>
2. Scaling SunGuard Higher Education Banner Software on Dell Servers and Storage, The Dell-SUNY Proof of Concept, Dave Jaffe, Kai Yu, Dan Brint, Dell White Paper <http://www.dell.com/Downloads/Global/Solutions/Dell-SUNY-Banner-POC.pdf>
3. Project MEGAGRID: Practical Guidance for Deploying Large Clusters of GRID, An Oracle, Dell, EMC, Intel Joint White Paper, December 2004,
http://www.oracle.com/technology/products/database/clustering/pdf/project_megagrid_practical-guidance-for-deploying-large-clusters.pdf
4. Automating Oracle RAC Deployment and Implementing MegaGird Using Oracle EM Provisioning Pack, Kai Yu, Oracle OpenWorld 2007 Presentation. Metalink DOC ID: 388577.1
5. Using Deployment Procedures to Scale-up and Scale-down Oracle RAC, An Oracle White Paper, October 2007,
<http://www.oracle.com/technology/products/database/clustering/pdf/emgcdeploymentproc.pdf>
6. Dell | Oracle Tested and Validated Configurations.
http://www.dell.com/content/topics/global.aspx/alliances/en/oracle_builds?c=us&cs=555&l=en&s=biz&~tab=3



Q/A

THANK YOU

THANK YOU

