Reducing Risk of Surprises in Changing Oracle RAC Environment

Dr. Boris Zibitsker

BEZ Systems

www.bez.com

boris@bez.com

312-849-1583

Outline

- Major issues affecting Oracle RAC scalability for large EDW with mixed workloads
  - Shared Storage
  - Variable DOP affects consistency of RT and throughput
  - Memory management
  - Workload management
- Role of workload characterization and performance prediction in organizing Proactive Performance Management
- How to set up realistic expectations
- Systems approach includes analysis of all interdependent components of the multi-tier distributed environment, including web servers, application servers and DBMS servers
- Case studies illustrate justification of strategic, tactical and operational decisions during different phases of EDW life cycle
- Summary

Introduction

- Dr. Boris Zibitsker, Chairman, CTO BEZ Systems
- Boris founded BEZ Systems, a software development company focused on Predictive Performance Management, in 1983
- Boris and his colleagues developed performance prediction software supporting Oracle, DB2 UDB ESE, Teradata, Java EE applications in multi-tier distributed environment:
  - BEZPlus — 1992
  - BEZProphet — 2005
  - BEZVision — 2007
- Boris was Adjunct Associate Professor at DePaul University (1983 – 1990)
- He has worked on modeling and performance prediction for relational DBMS since 1983 and has taught seminars around the world for Relational Institute, which was organized by the pioneers of relational technology, Ted Codd and Chris Date
- Boris has consulted over 200 Fortune 500 companies on Proactive Performance Management

Causes of Performance Surprises

Change

- Workload Growth
- Volume of Data Increase
- New Application Implementation
- Change of Application Functionality
- Performance Tuning
- Server Consolidation
- Hardware Upgrade
- Software Version Upgrade

Surprise

- Unexpected Change of Response Time
- Unexpected Reduction of Throughput

Potential Bottlenecks Affecting RAC Performance

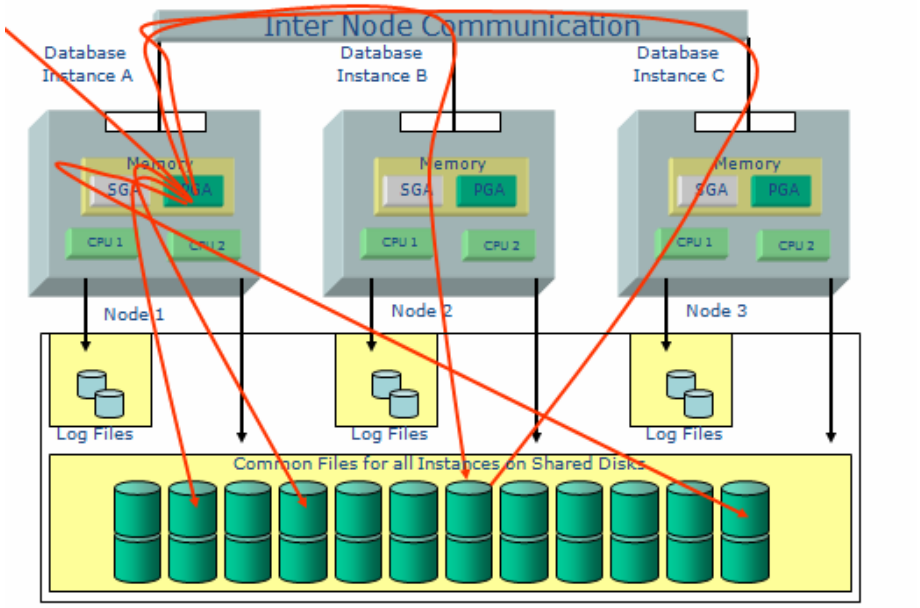

SQL request can be parallelized, but level of parallelism is not constant and depends on several factors

- **Shared Storage**
  - **Disk cannot handle more than 100-150 I/Os per second**
- **DOP**
  - **Variable DOP affects consistency of RT and Throughput**
- **Contention for Interconnect**
- **Memory Limitations**
- **Workload Management / Resource Management Limitations**

We Will Review How to Reduce Uncertainty and Risk of Performance Surprises in Constantly Changing Environment

- Know your workloads, and their performance, resource utilization and data usage profiles
- Know your hardware and software configuration, and factors affecting concurrency, parallelism and potential delays
- In addition to Oracle reactive self-healing features, implement proactive performance management approach
- Know trend and anticipated changes
- Apply analytical methods to predict the impact of anticipated changes, including growth, new applications, upgrades, consolidation, etc.
- Set realistic SLO, SLA and expectations prior to significant changes
- Implement necessary measures before it is too late
- Manage expectations

Perform Workload Characterization to Build Performance, Data & Resource Utilization Profiles for Each Workload

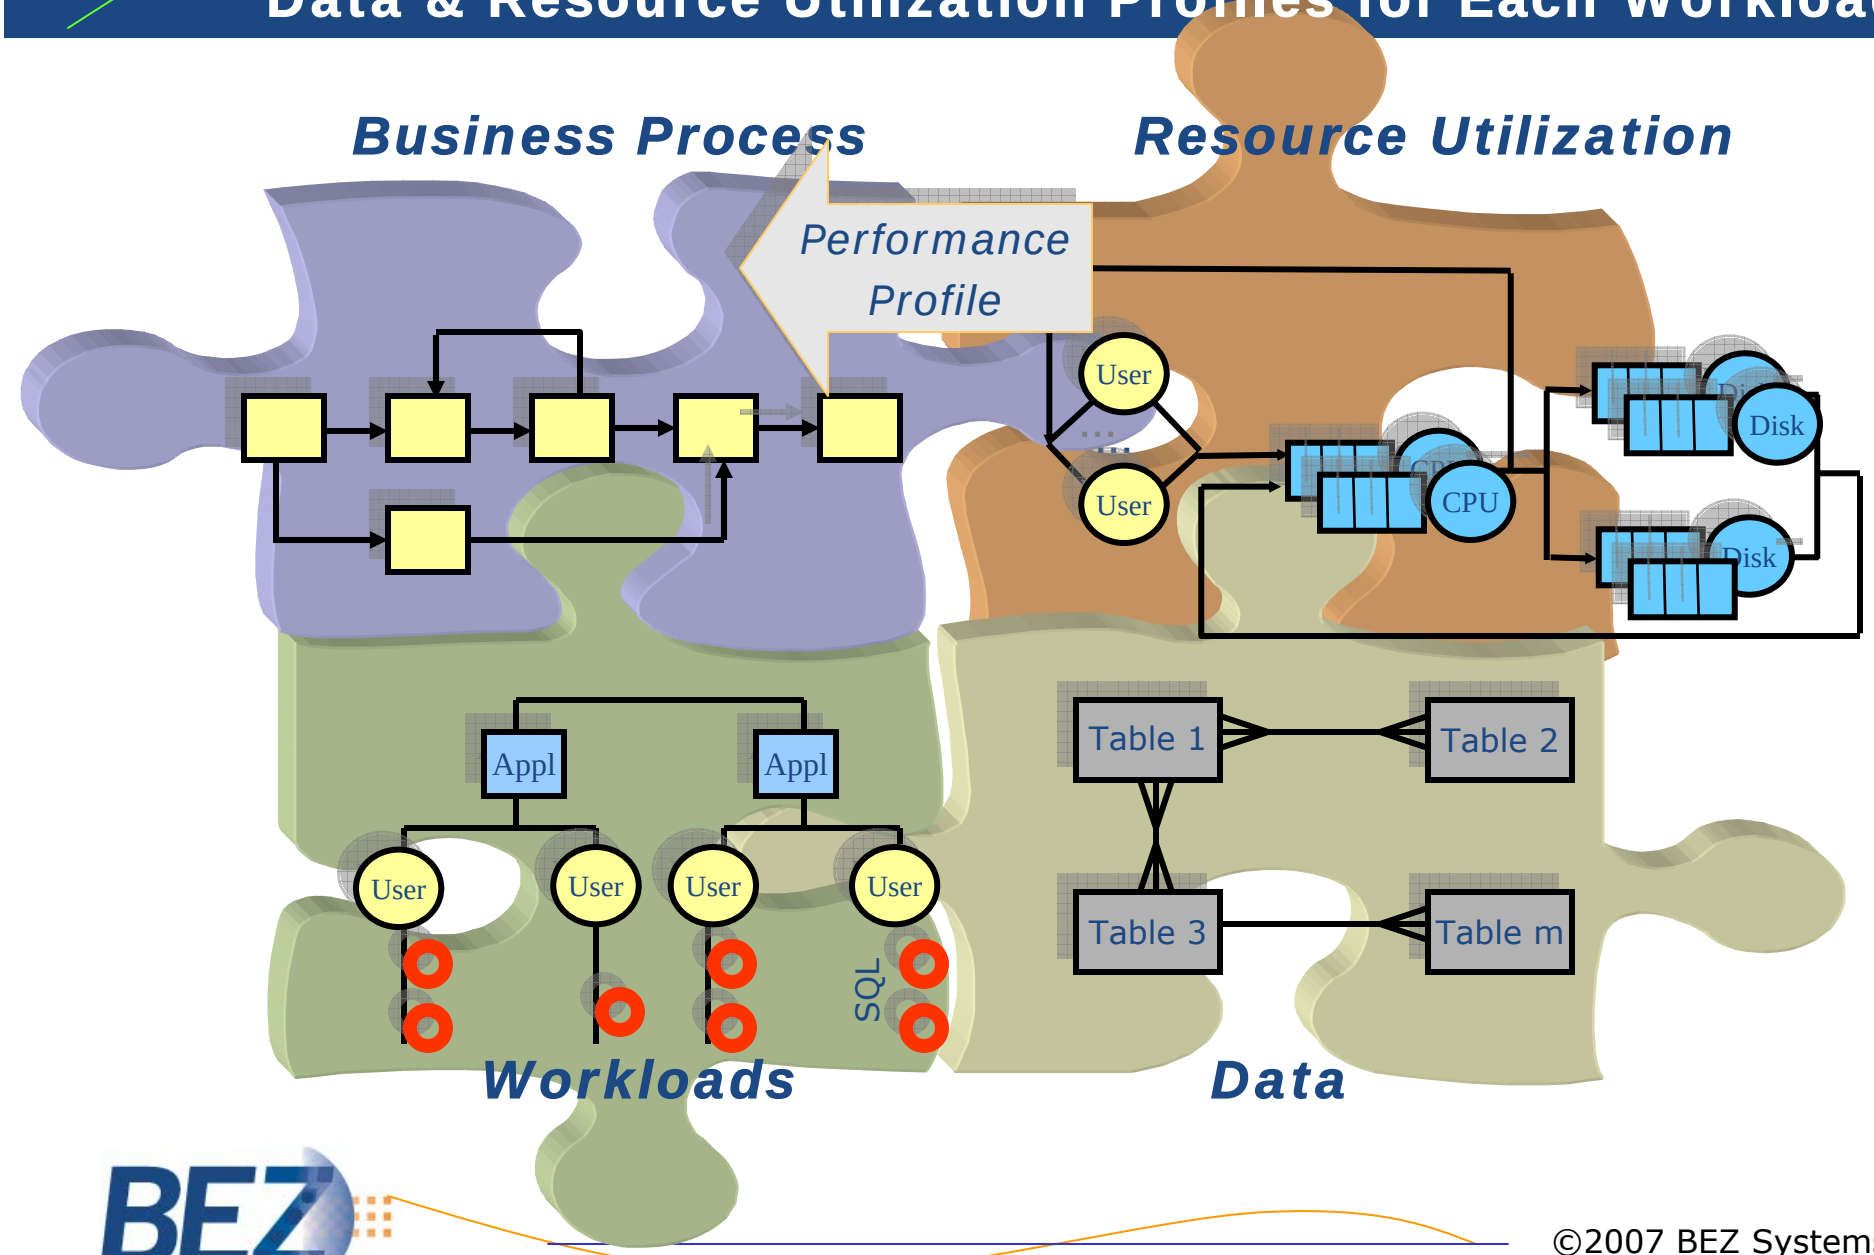

Use Modeling to Predict Impact of Change, Evaluate Different Options, Justify Decisions and Set Expectations

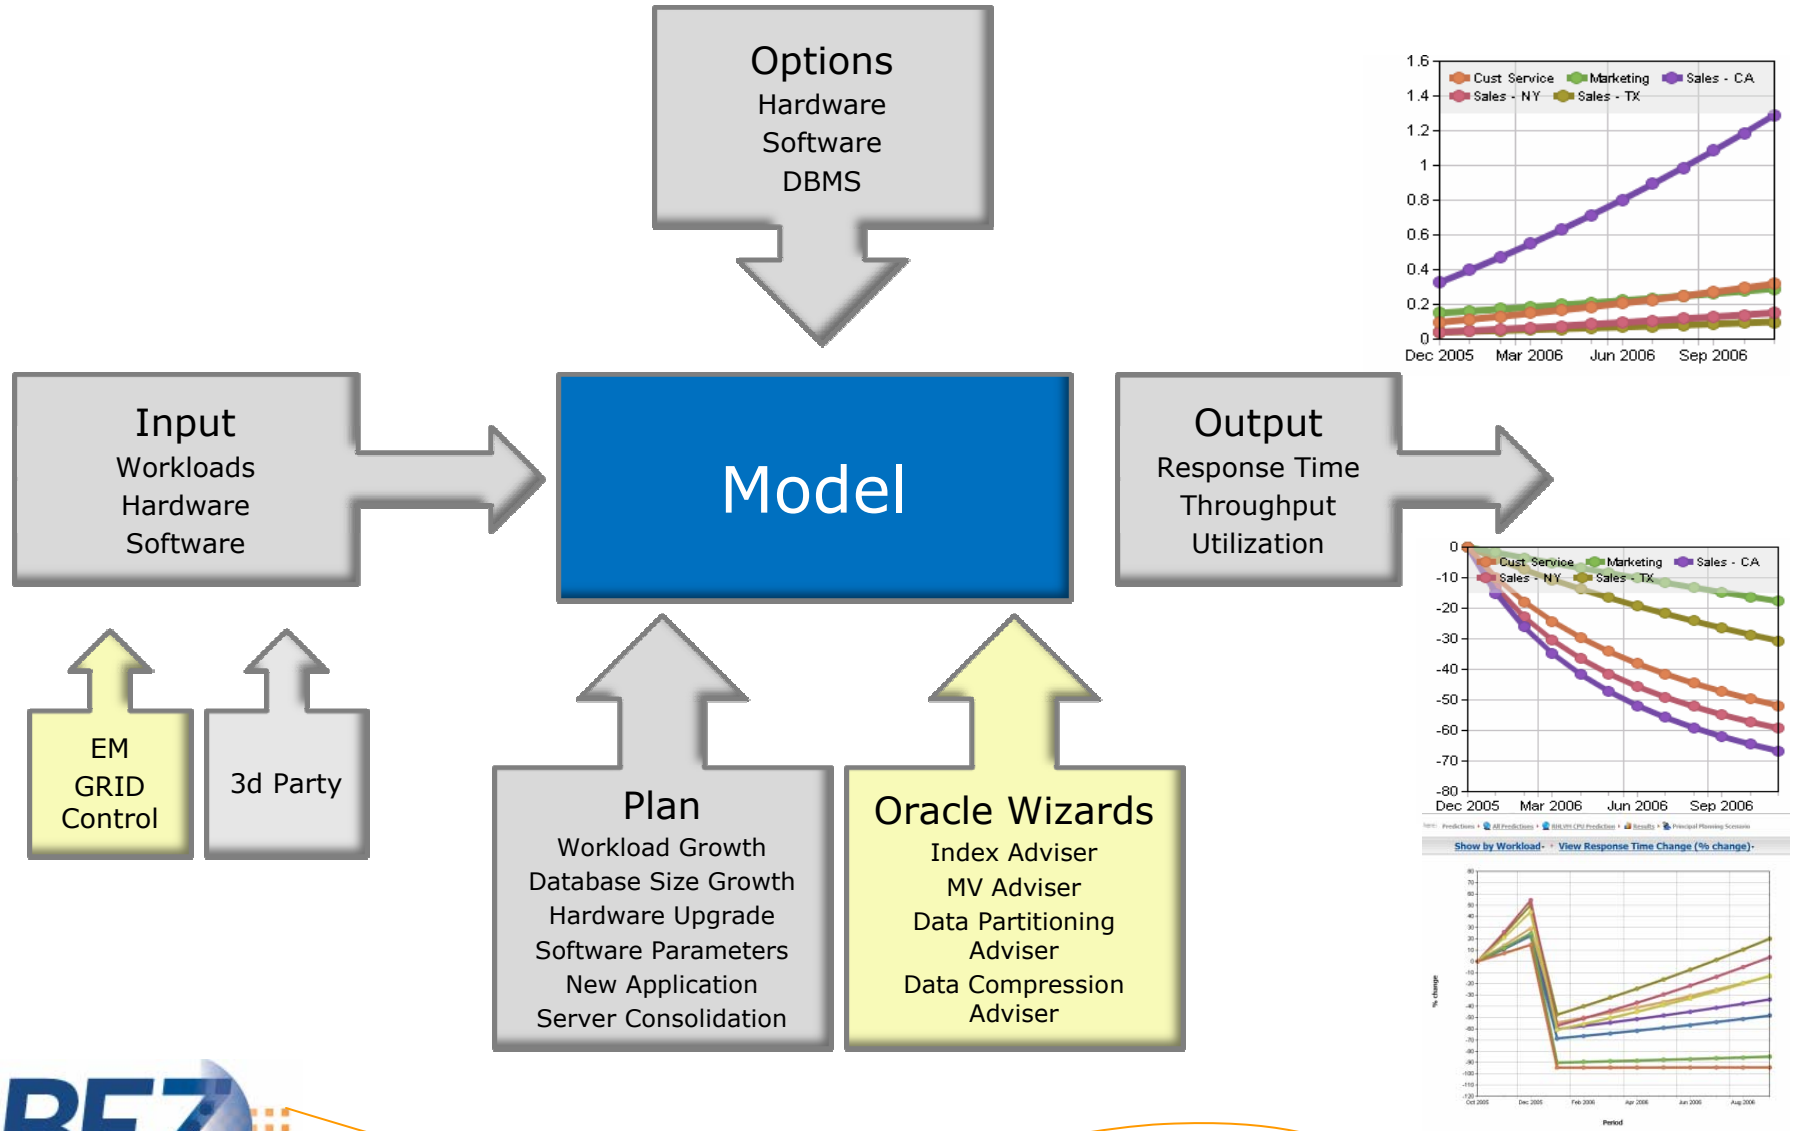

Several Vendors Offer Analytical Queueing Network Models

- Queueing Network Model represents RAC as a network of interconnected servers and queues where each device of the system is a server and a queue
- Software limitations affecting concurrency limitations, latching delays, and parallelism are represented as queues as well
- Input of the model includes result of the workload characterization, hardware configuration, and software parameters
- Plan describes expected growth, proposed changes, including new application implementation, hardware upgrade, database design changes, proposed change of the workload priority, etc.
- Interface with DBMS wizard allows us to process critical SQL and to obtain the wizards' DBMS, application, memory reallocation and other recommendations, which is also used as input to model
- Several vendors, including BEZ, BMC, Metron, Hyperformix, Team Quest and others offer their performance prediction tools

Simplified Example Illustrates Prediction Impact of Expected Workload Growth and Hardware Upgrade

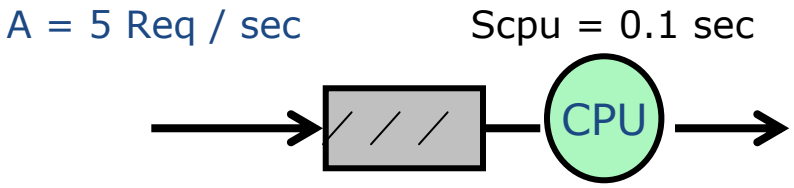

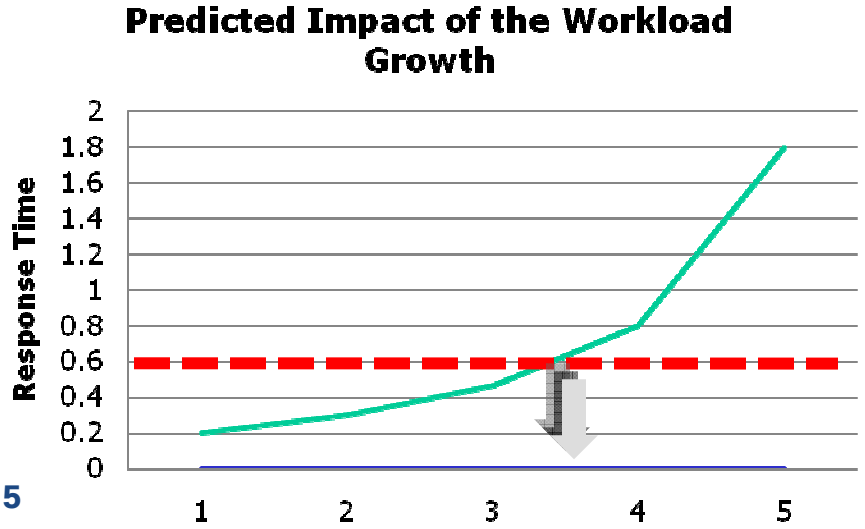

Utilization Ucpu = 5 Req/sec * 0.1sec = 0.5

Response Time Rcpu = 0.1 sec / (1 - 0.5) = 0.2 sec

Based on expected workload growth of 20% per month, predict when the system will not be able to meet SLO (0.6 sec).

What will be the impact of doubling CPU speed?

How long system will satisfy SLO?

	This Month	Next Month	In 2 Months	In 3 Months	In 4 Months
Arrival Rate (Req/sec)	5	6	7	8	9
Service Time (sec)	0.1	0.12	0.14	0.16	0.18
Utilization (%)	0.5	0.6	0.7	0.8	0.9
Response Time (sec)	0.2	0.3	0.46	0.8	1.8

What if We Double CPU Speed?

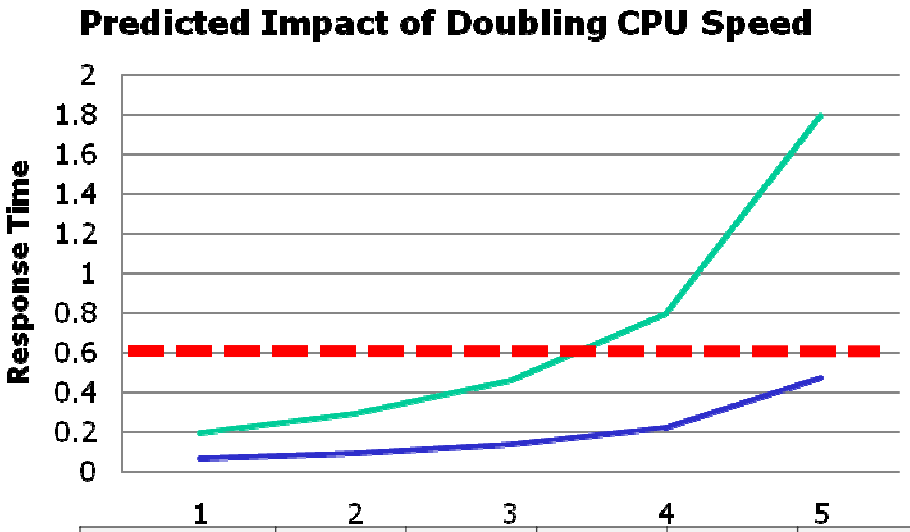

	This Month	Next Month	In 2 Months	In 3 Months	In 4 Months
Arrival Rate (Req/sec)	5	6	7	8	9
Service Time (sec)	0.1	0.12	0.14	0.16	0.18
Utilization (%)	0.5	0.6	0.7	0.8	0.9
Response Time (sec)	0.2	0.3	0.46	0.8	1.8
Doubling CPU Speed	0.06	0.09	0.13	0.22	0.47

Based on expected workload growth of 20% per month, predict when the system will not be able to meet SLO (0.6 sec).

What will be the impact of doubling CPU speed?

How long system will satisfy SLO?

Modeling and Optimization is a Foundation of Closed Loop Proactive Performance Management Supporting DW, Application and Information Life Cycle

Measure → Characterize → Plan → Model & Optimize → Advise → Manage

DW Life Cycle

DW architecture, hardware, software and DBMS platforms
Workload and DB size growth
DW database and application tuning
Data partitioning
Data compression
DW capacity planning
DR planning
DW and DM consolidation

Application Life Cycle

Design options
New application implementation
DB performance tuning
Application modification
Application consolidation

Information Life Cycle

Loading growing volume of data
Data loading strategies
DB backup and restore time
Data replication
Data consolidation
Enterprise data management strategies and alternatives
Information integration

Models Should Take into Consideration Interdependences between Servers and Concurrency Limitations to Support Decisions in a Multi-tier Distributed Environment Supporting Mixed Workload

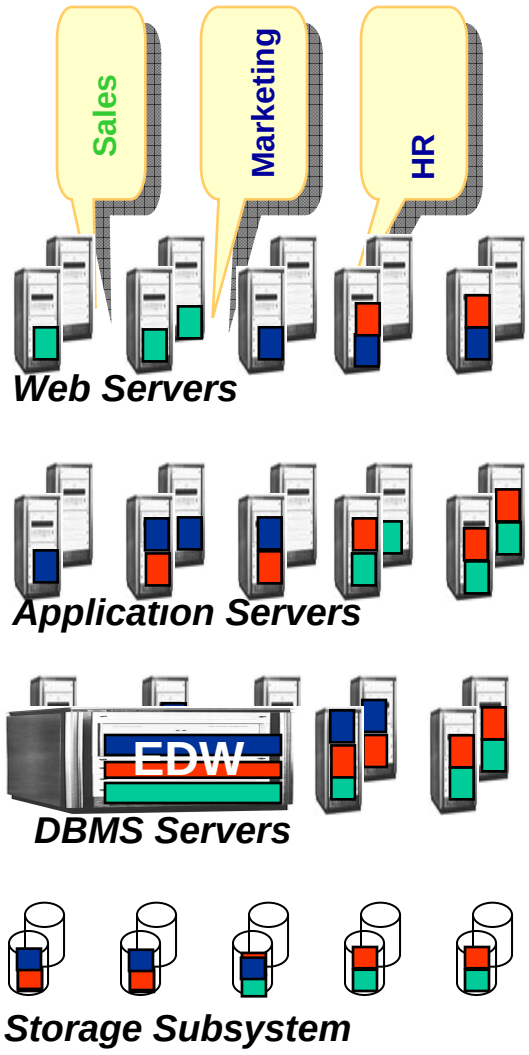

Strategic Decisions

- Setting realistic SLO and SLA for each workload
- New application implementation impact
- Justification of architecture, hardware, software, and database design to support workloads SLA
- Mixed workload management
- Minimum configuration for Disaster Recovery
- Server consolidation impact on each workload
- Support of distributed DBMS servers

Tactical Decisions

- Optimization of software parameters controlling concurrency, degree of parallelism and workload priorities
- Number of JVMs and application servers
- Number of JVM threads and connection pool size to support growing workload
- DBMS tuning

Operational Decisions

- Moving workloads between shifts and servers
- Change workload priority
- Change in number of concurrent ETL utilities for batch and online window
- Actual to expanded

Automation of Workload Characterization, Modeling and Optimization Enables Organization of Continuous Proactive Performance Management

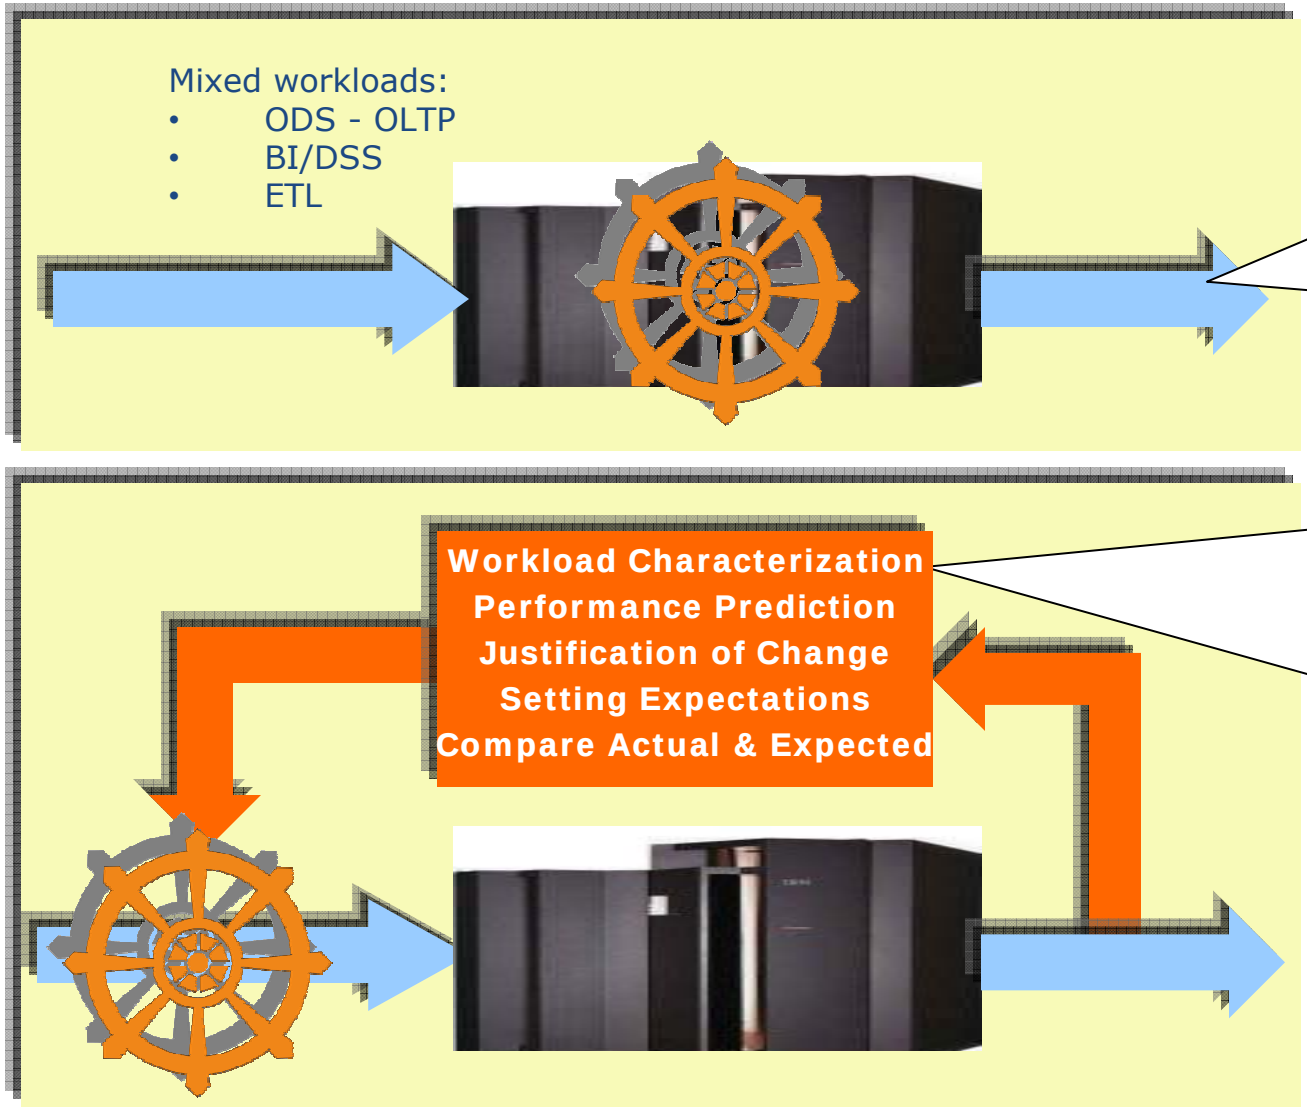

How can you manage if you cannot compare options, do not know what to expect and cannot verify results?

Automation of Workload Characterization, Modeling and Optimization enable evaluation of many options, justify solutions and organize a continuous, closed loop proactive performance management

1) Modeling Results Show How Contention for Shared Storage Affects RAC Scalability and Allow Setting Realistic SLO and SLA

Typical Actions to Reduce Contention for Shared Storage

- Disk cannot handle more than 100-150 I/Os per second and with a shared disk environment, this is a potential bottleneck
  - Graph illustrates that when I/O service time=5ms, increase of the arrival rate from 10 to 150 I/O per sec increases the contention for disk and I/O response time 4 times
- Memory management
- Data partitioning
- Data compression
- Data striping
- Increase number of disks
- Use more small disks to place critical data and solid state devices for indexes
- Consider ASM instead of Veritas
  - Acceptable for small environments
  - Large installations continue to use Veritas

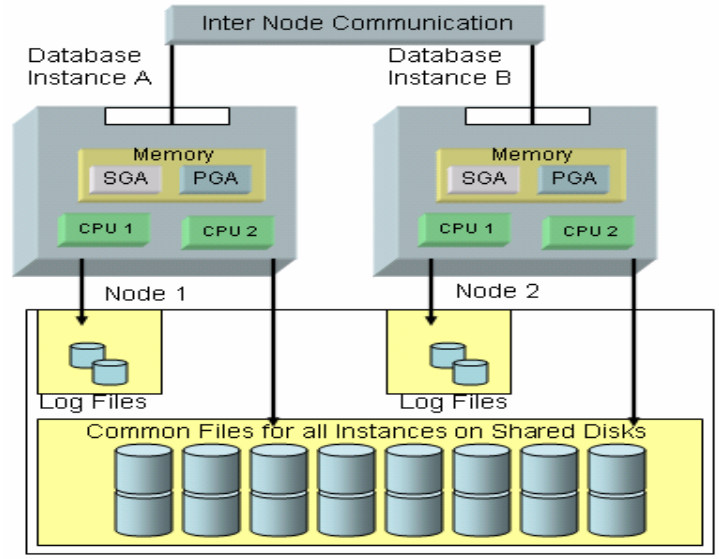

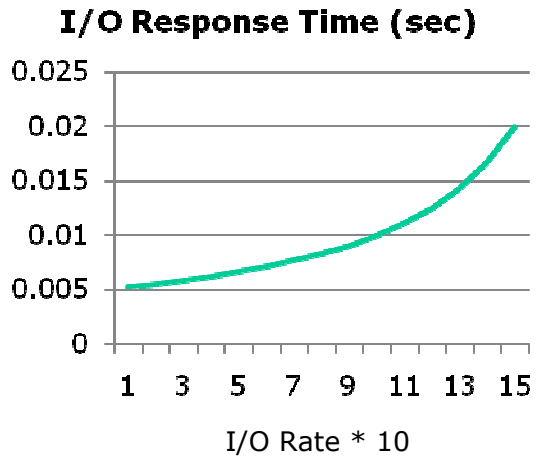

Prediction Results Show that Increase in # of RAC Nodes Reduces CPU utilization, Improves Response Time and Throughput, but Increases Contention for Disk

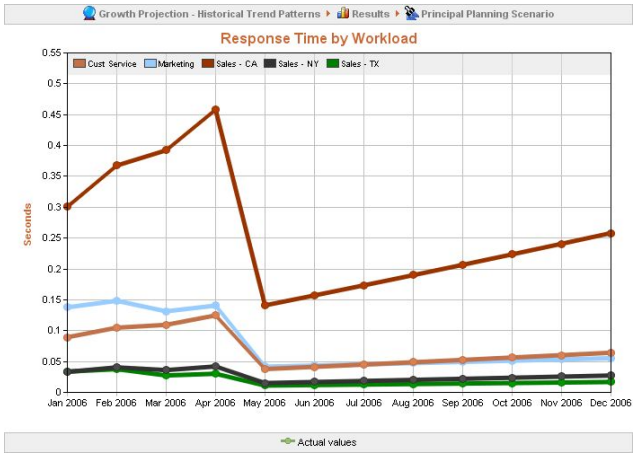

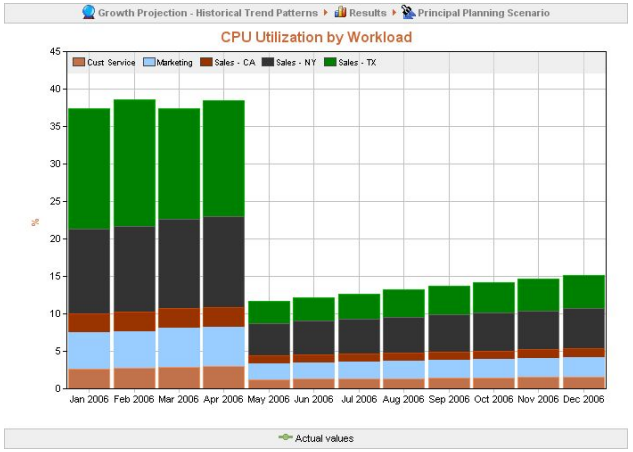

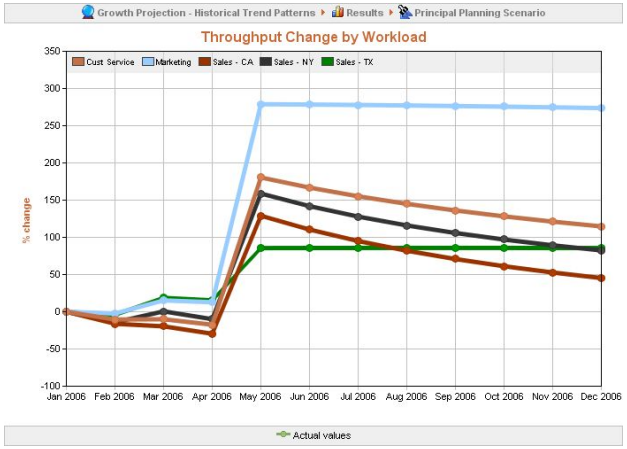

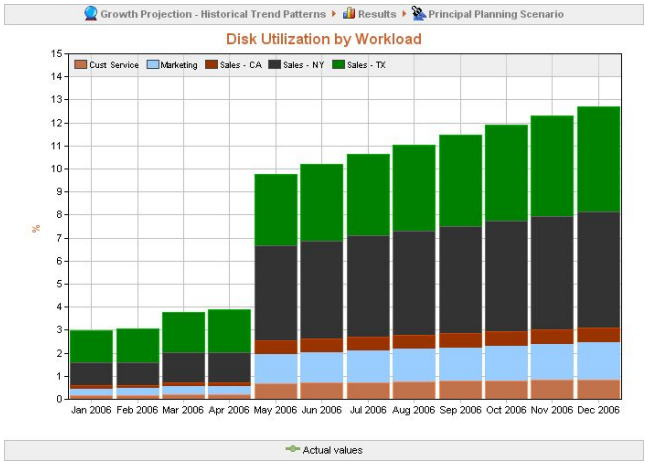

Modeling Results Help Customers to Set Up Realistic SLO and Negotiate SLA for Major Workload

Predicted RT for one of the Workloads

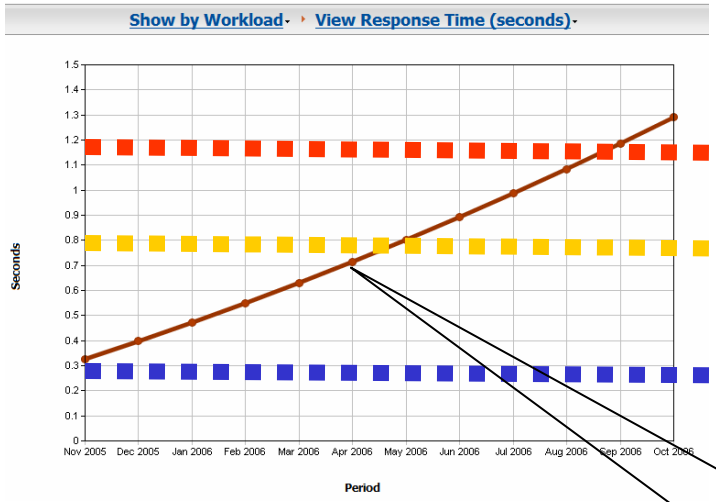

SLA

SLO

Hardware Configuration & TCO

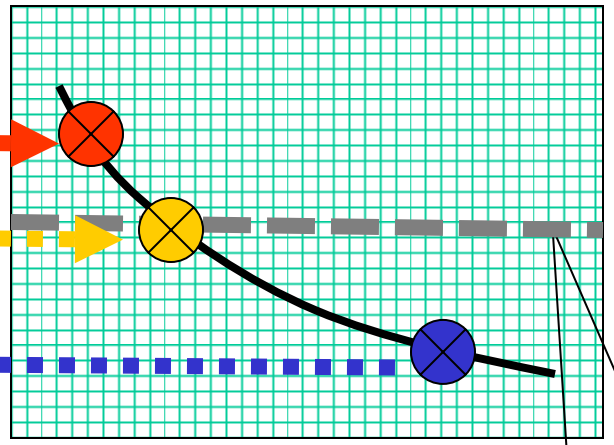

Expected workload & DB growth

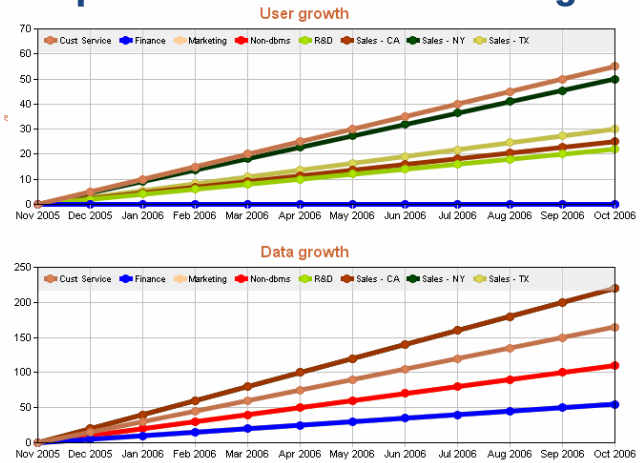

Prediction results allow customers to negotiate SLA between business and IT

For expected workload and database size growth, IT guarantees delivery of a certain level of responsiveness and throughput

Users and IT select SLO level that will provide acceptable performance with acceptable Total Cost of Ownership (TCO)

2) Modeling Results Show Impact of Parallel Processing and Allow Comparison of the Performance of Oracle RAC, Teradata and DB2 UDB ESE during Selection of the Platform for EDW

Difference in Parallel Processing by Oracle and Teradata

- Difference in parallel processing of Oracle RAC, DB2 UDB and Teradata affects consistency of service
- Oracle supports dynamic parallelism
  - For example, if we have 16 disks, but only 4 query servers available to support specific SQL, then each query server will read data from 4 disks
  - If tomorrow 8 query servers will be available to support the same SQL, then each query server will read data only from 2 disks and it will take less time to read data
- Number of virtual chunks of data processed by SQL in RAC environment varies
- In the case of Teradata, all AMPs read data from all disks/DSUs in parallel regardless of what else is going on in the system
- Level of concurrency is controlled by the number of connections, and TASM
- Teradata Optimizer can schedule join of several tables and reading data from other tables in parallel
- Oracle processes joins and reads from other tables sequentially

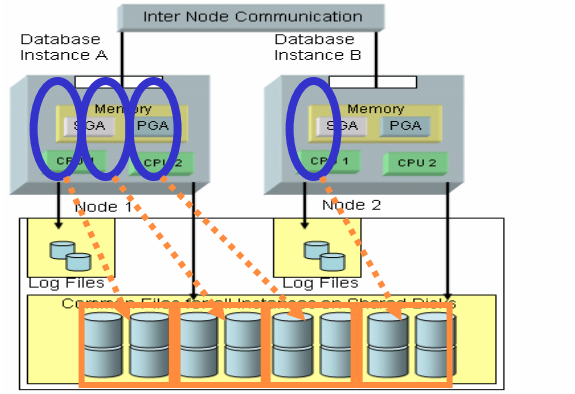

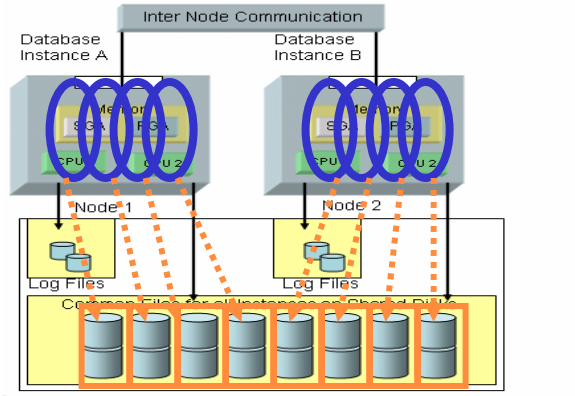

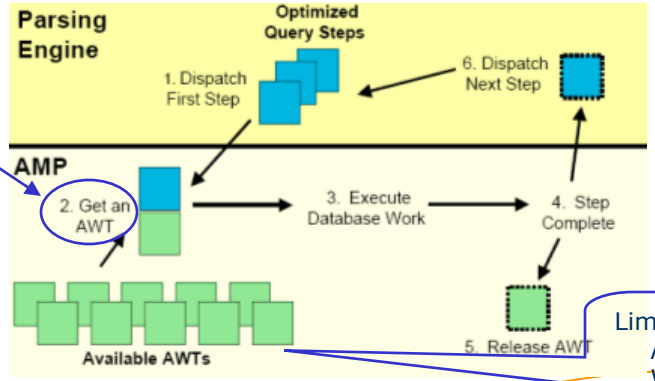

Limited # of Available AMP Worker Tasks

Prediction Results Comparing Response Time and CPU Utilization on Oracle SMP and RAC vs. Teradata During Online Window

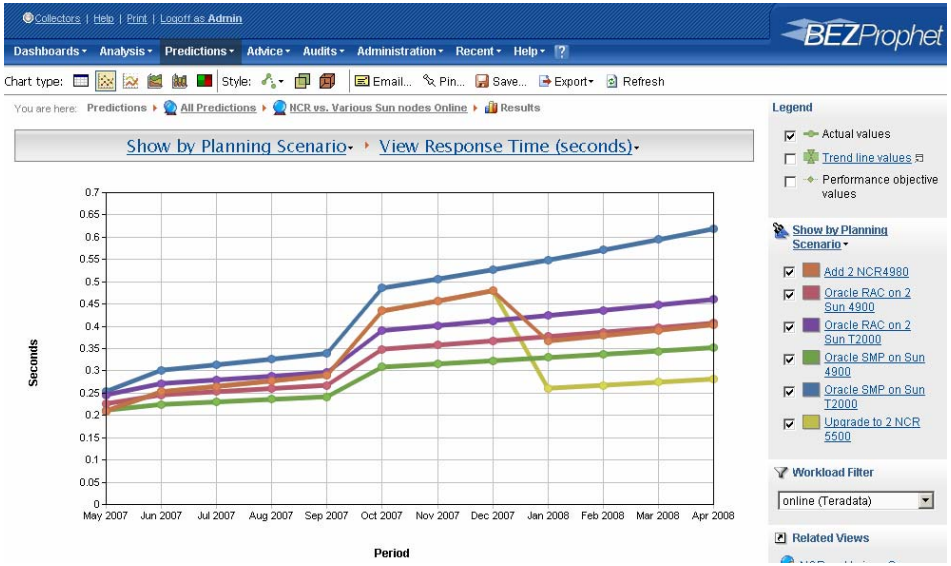

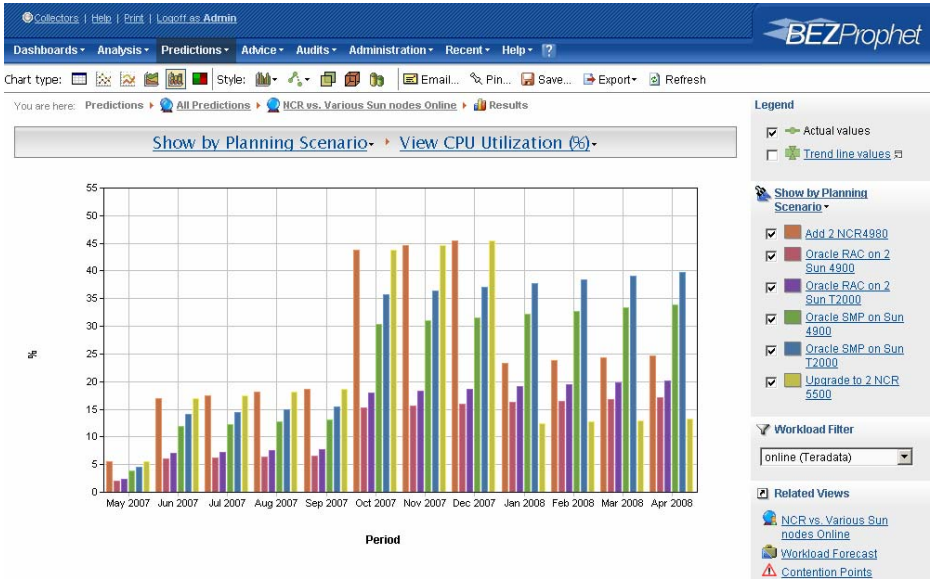

Requirement to Consider the Impact of Application Servers on EDW Performance

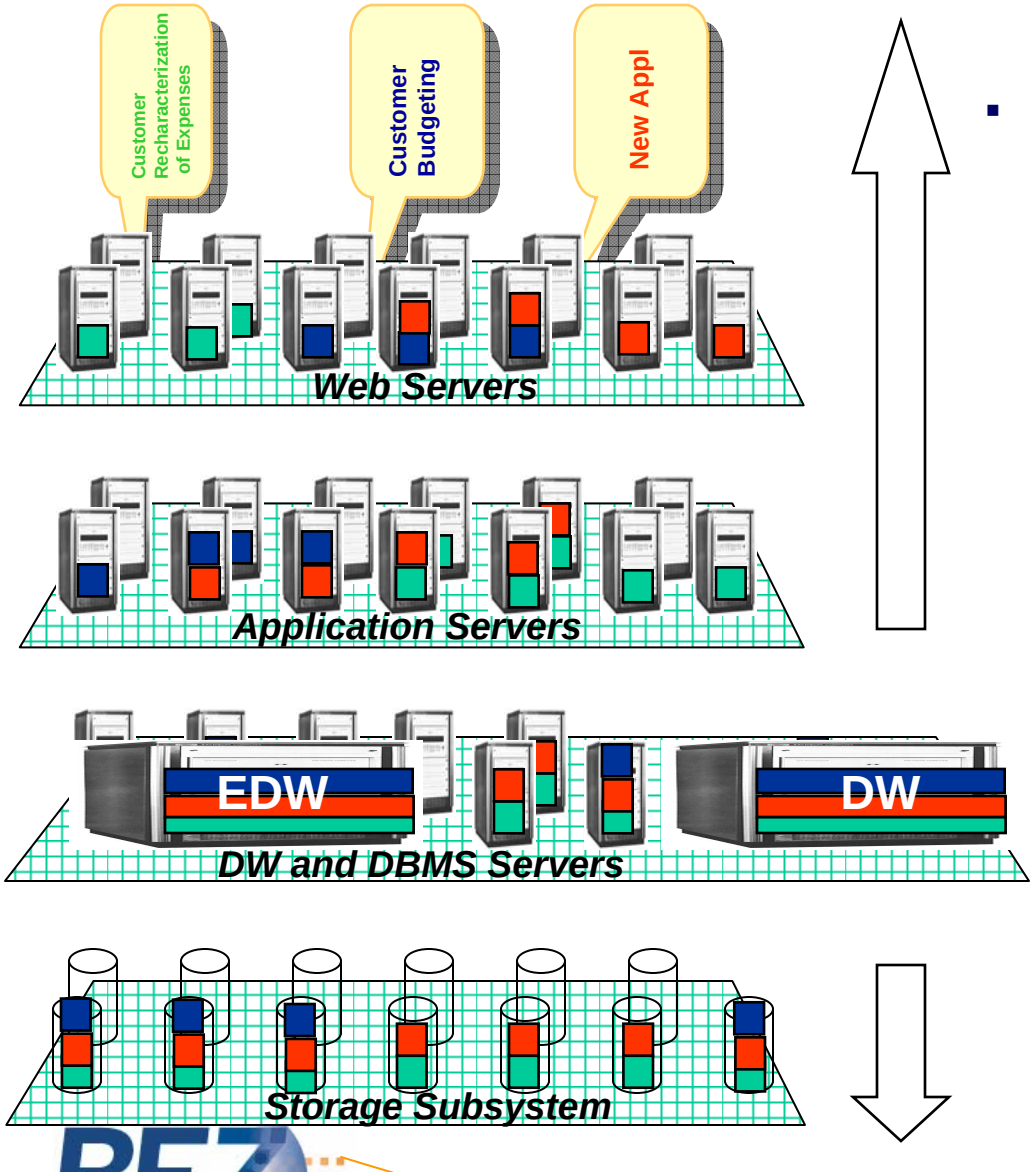

- Limitations of connection pool sizes and number of JVM threads in growing and changing environment can significantly affect the concurrency level and data warehouse performance

3) Modeling Results Show How Expected Changes Will Affect Contention for the Interconnect for Different Architectures

Example shows the evaluation of contention for the interconnect on Egenera vs. IBM BCU1 Power5 and BCU Power6 during selection of a platform for Oracle RAC EDW

BEZ
a business lens on IT

Predicting Interconnect Scalability on Different Platforms

- Increasing contention for the Interconnect caused by workload and DB size growth and increase in number of nodes can affect different workloads differently on different platforms
- Example below shows prediction results evaluating the difference between Egenera and IBM pSeries in supporting EDW RAC workload, taking into consideration the impact of the workload and volume of data growth on interconnect utilization and each workload's performance
- Performance prediction results compare Egenera and IBM pSeries Scaling Up and Scaling Out capabilities for EDW RAC workload

- How will workload growth increase the contention for internode communication?
- How will volume of data growth affect the traffic through the interconnect, and how will it affect the performance of each workload?
- How will adding more nodes affect the message traffic and performance of each workload?
- Predict each workload's interconnect response time, including service time, queueing time during transferring messages, and data and latency time

Egenera's Interconnect Is Used to Transfer Messages and Data Between pBlades and Connects Them with LUNs of Storage Subsystem and Application Servers

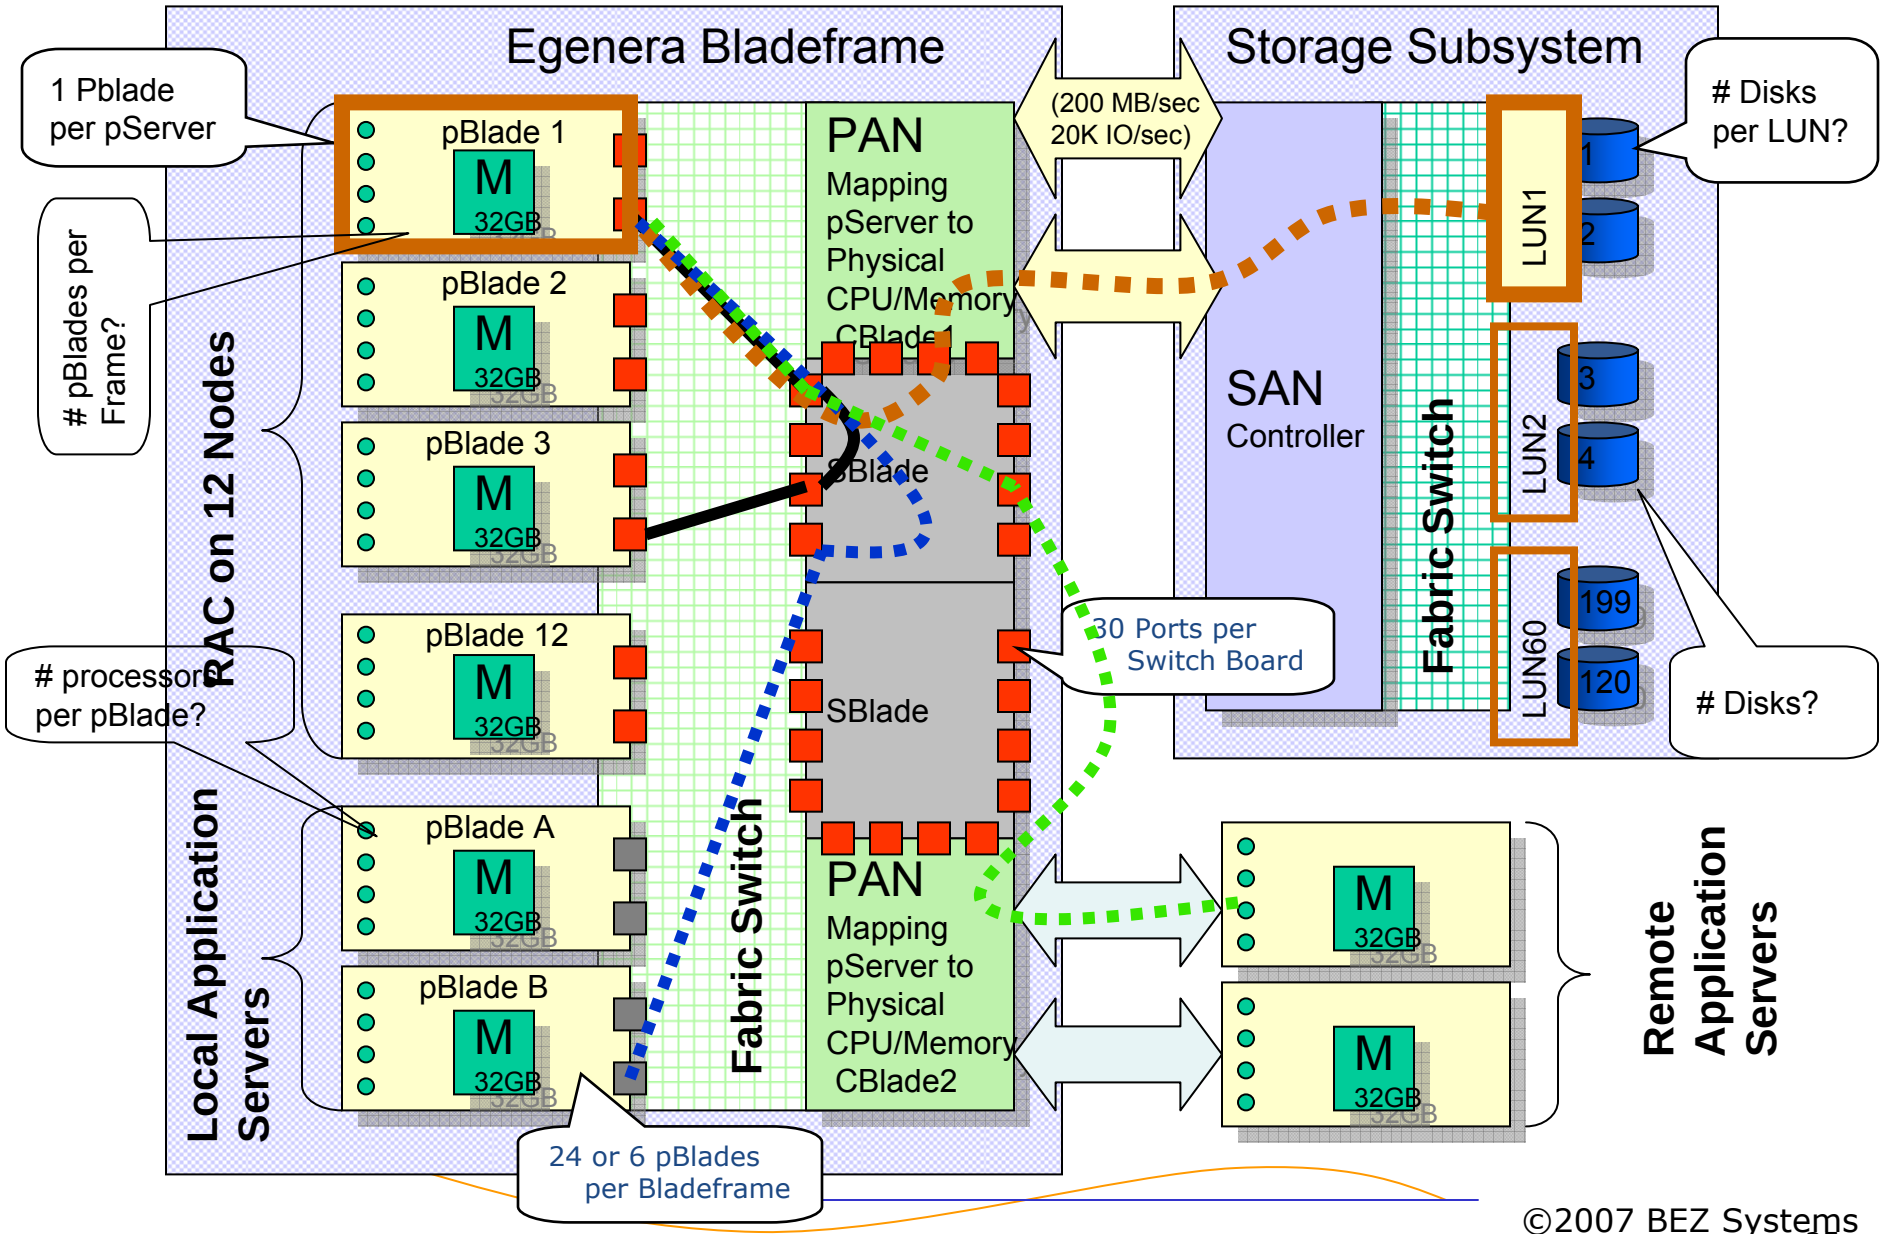

Predicted Impact of Workload Growth on Egenera Interconnect Utilization and Response Time

Interconnect Utilization by Workload

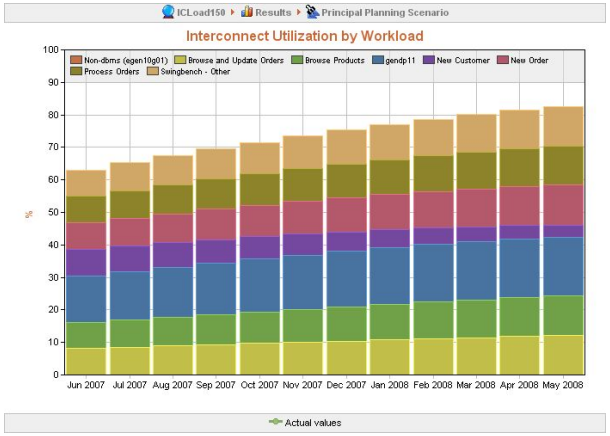

Interconnect utilization is growing from 63% to about 82%. This graph shows interconnect utilization by workload.

Workload Response Time Components (CPU/Disk/Interconnect)

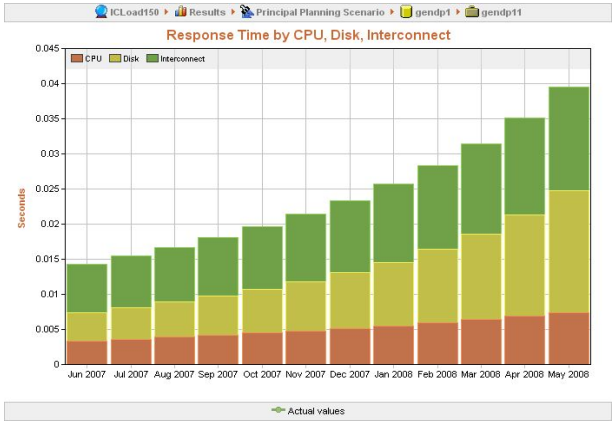

Interconnect response time is a significant component of the workload response time.

Service Time (CPU/Disk/Interconnect)

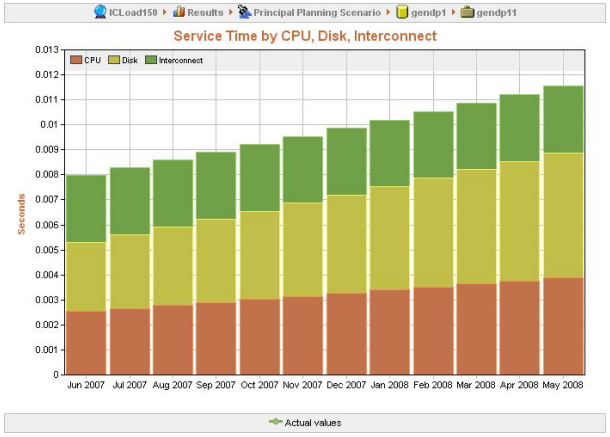

This graph shows the CPU, Disk and Interconnect service time.

Data volume growth will affect CPU and disk service time.

Interconnect service time will not change significantly.

Queueing Time (CPU/Disk/Interconnect)

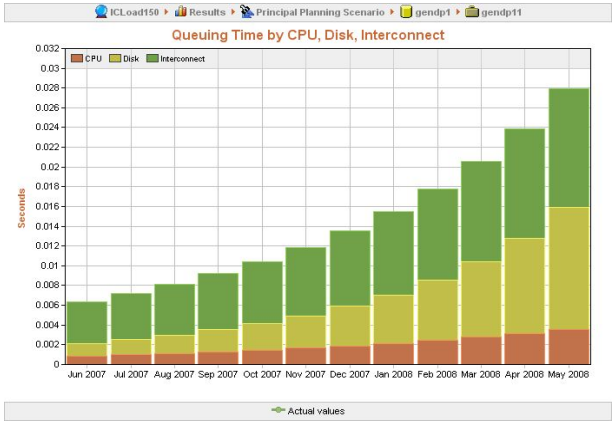

Contention for the interconnect will significantly increase queueing time.

Performance Prediction of the New Cache Fusion Protocols Impact on Interconnect and Each of the Workloads

- New cache fusion workload aware protocols of Oracle 11g will dynamically alter the inter-node RAC messaging depending on workload being processed
- When information about new read-optimized protocol, reduction in number of inter-node messages for read operations, and messaging behavior for update and table scan operations are available, the Prediction Engine will be able to predict how the performance of each workload will be affected

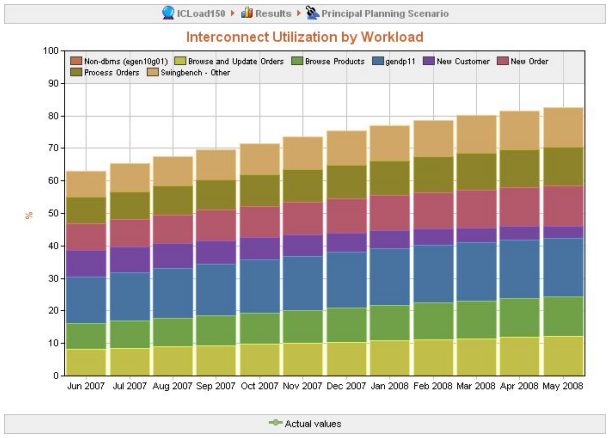

Interconnect utilization is growing from 63% to about 82%. This graph shows interconnect utilization by workload.

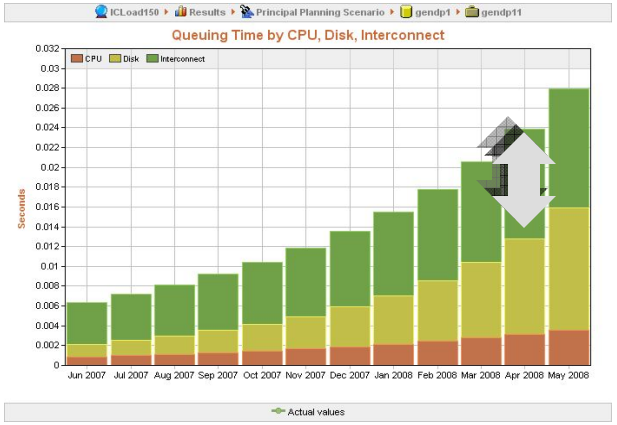

Modeling can be used to predict how new cache fusion protocols will reduce contention for the interconnect and how it will affect the queueing time.

EDW on Oracle RAC Based on Power5 and Power6 Nodes

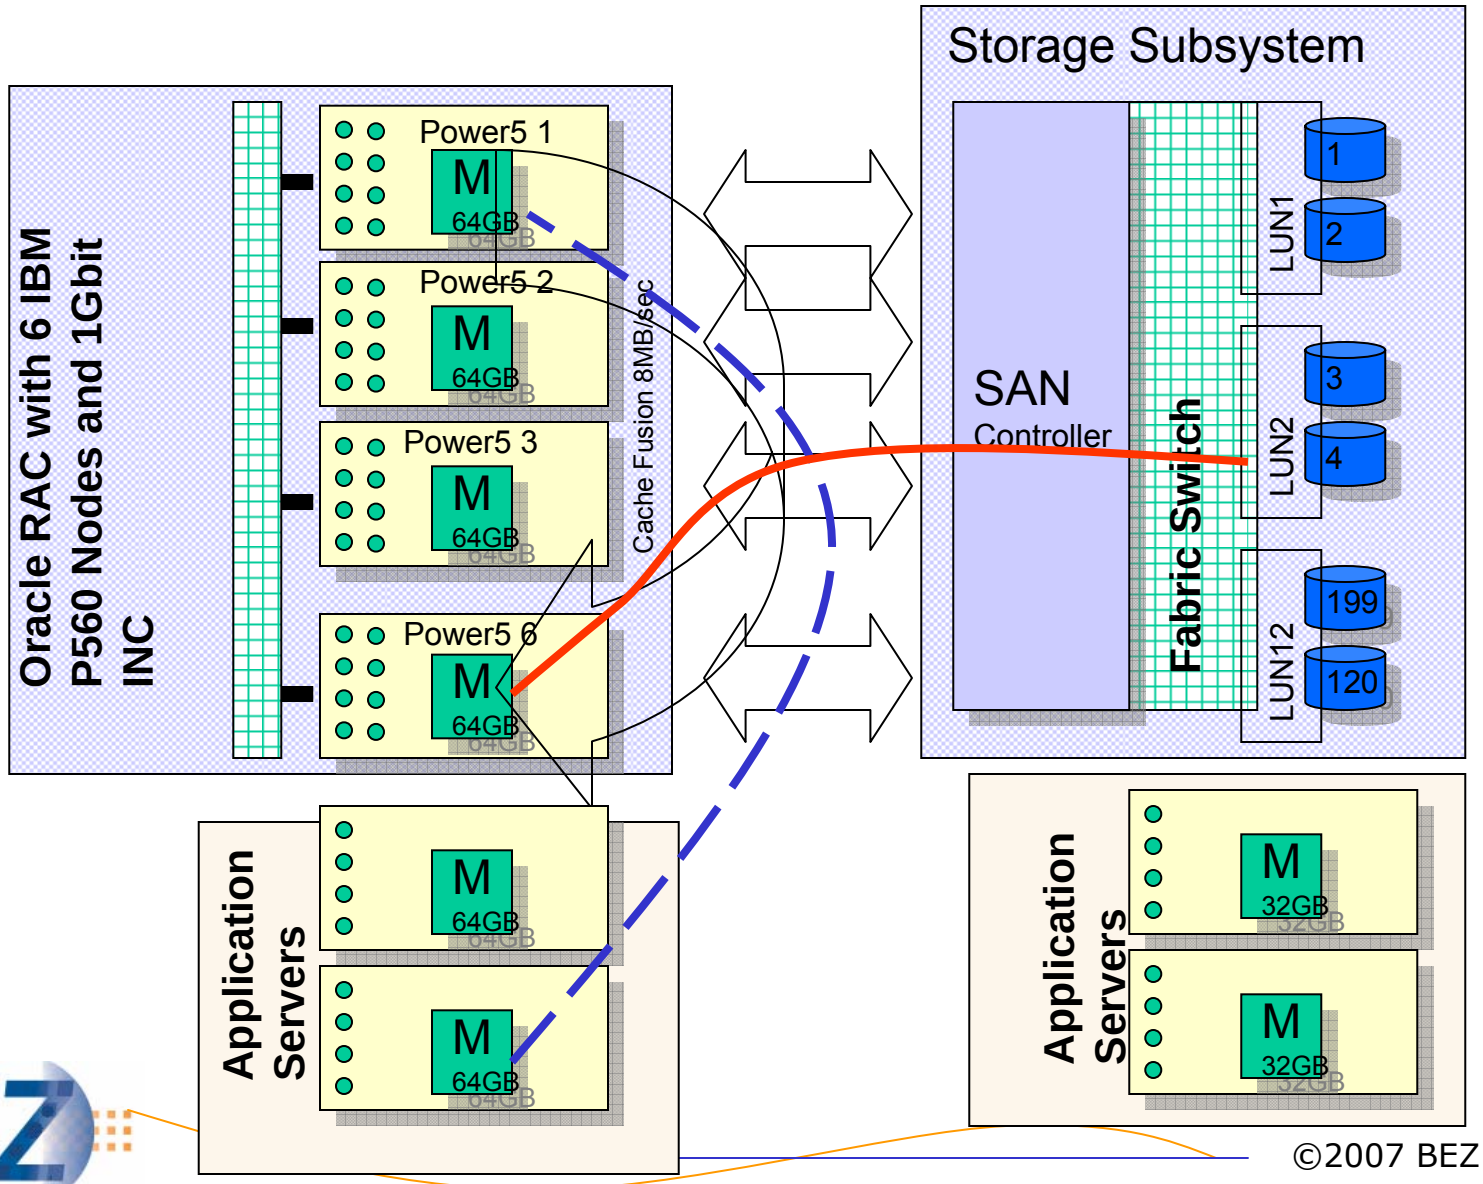

Scaling Out
2, 4 and 6 IBM Power5 and Power6 Nodes

Response Time	CPU Utilization	Disk Utilization	
	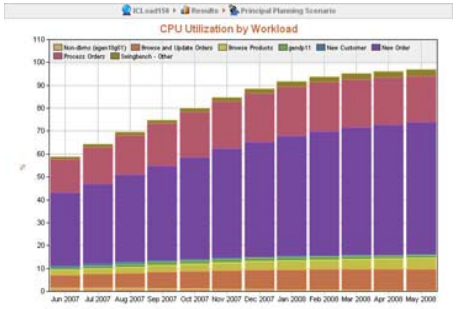	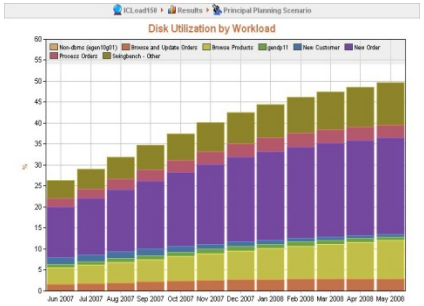	2 Nodes 8 Power5 CPU/node
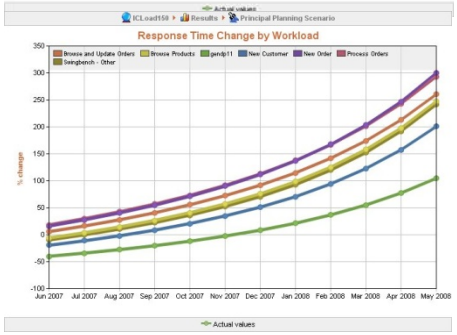	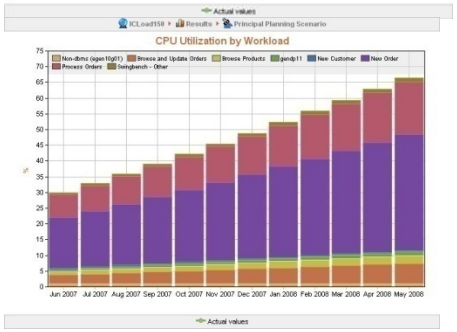	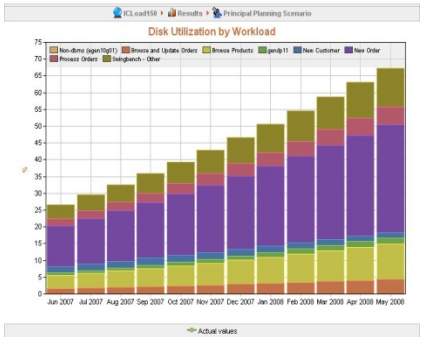	4 Nodes 8 Power5 CPU/node
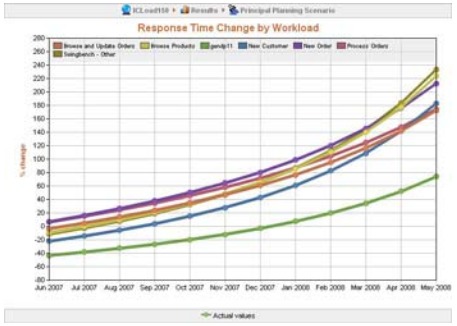	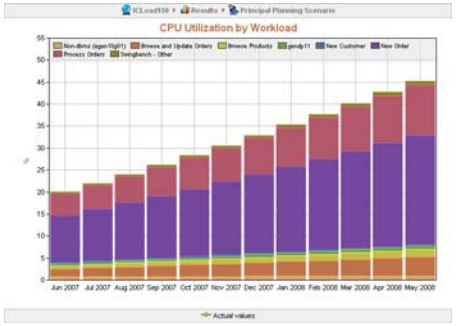	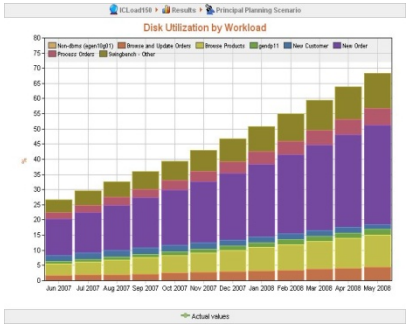	6 Nodes 8 Power5 CPU/Node

Scaling Up
2 IBM 550 Nodes with 2, 4 and 8 Power6 CPUs

Response Time	CPU Utilization	Interconnect	
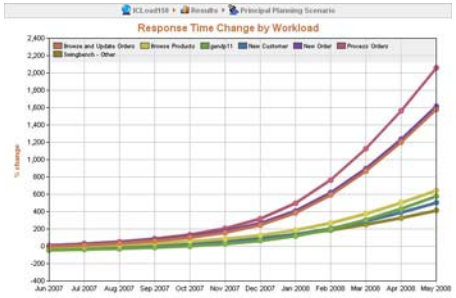	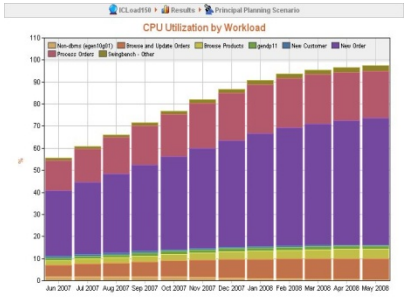	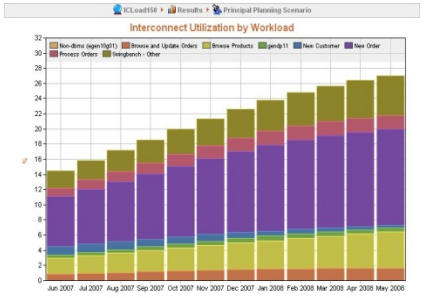	2 Nodes 2 Power 6 CPU
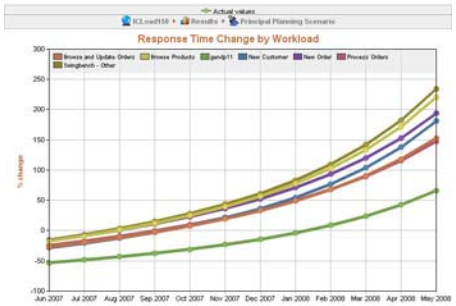	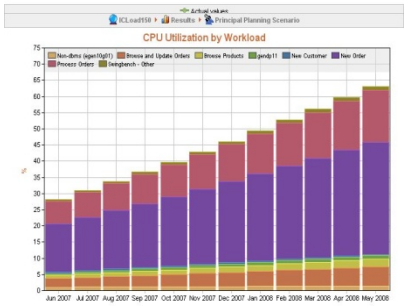	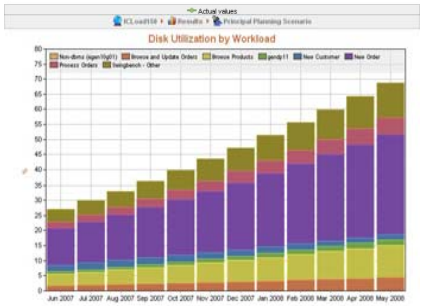	2 Nodes 4 Power 6 CPU
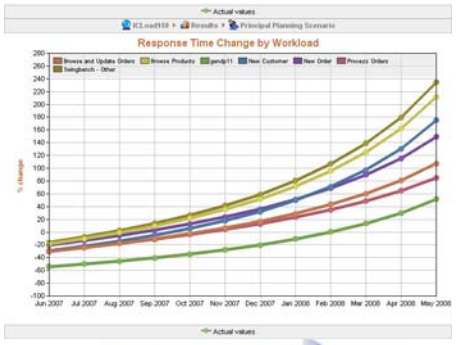	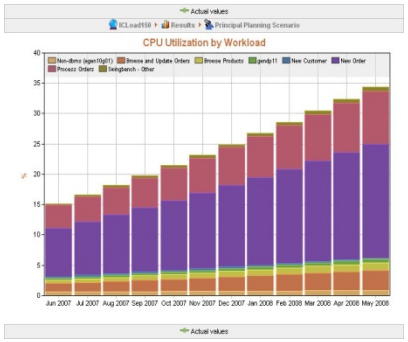	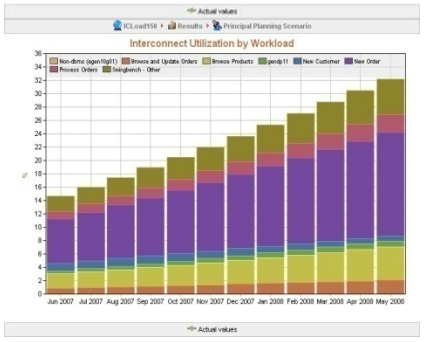	2 Nodes 8 Power 6 CPU

4) Modeling Results Show Impact of Memory Management Decisions in Mixed Workload Environment

Impact of Memory Management Decisions

- PGA size should be sufficient to support complex queries
- OLTP workload needs large SGA, but BI/DSS workload SGA should have just enough size to handle indexes, because complex requests bypass SGA and over-allocation of memory for SGA can negatively impact other workloads
- Performance prediction based on Oracle Memory Advisor recommendations can be used to evaluate different alternatives

Performance Prediction Results Based on Oracle Memory Advisor Recommendations Reflect the impact of the Workload Growth and Memory Pool Size Change

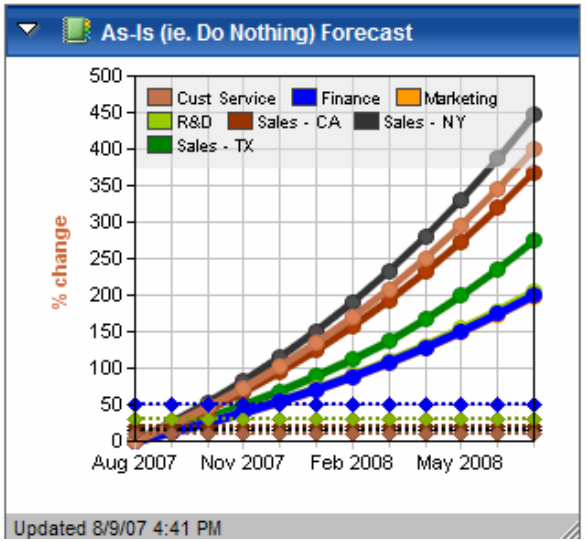

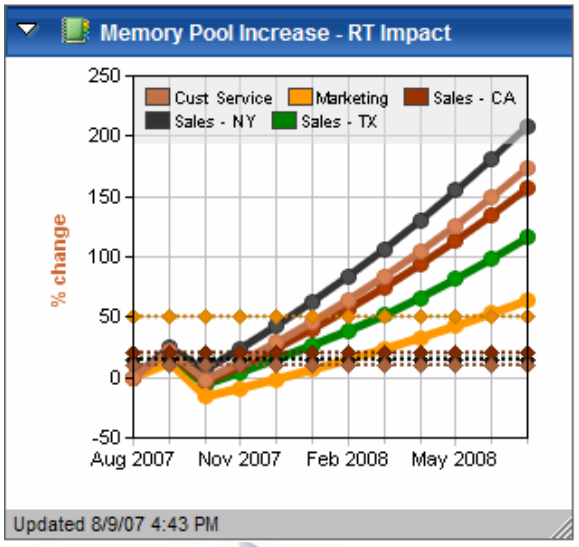

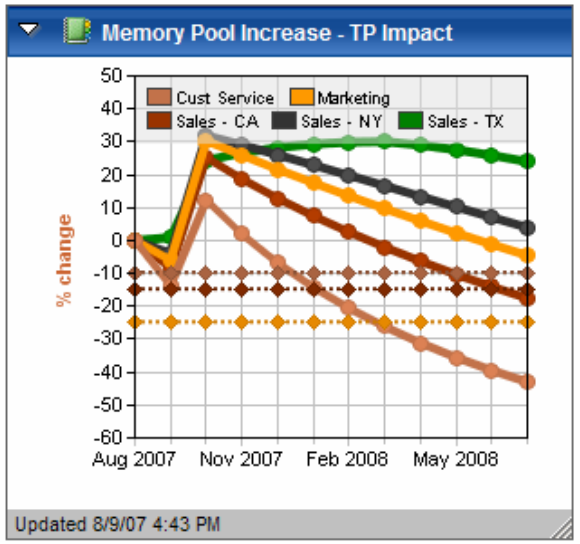

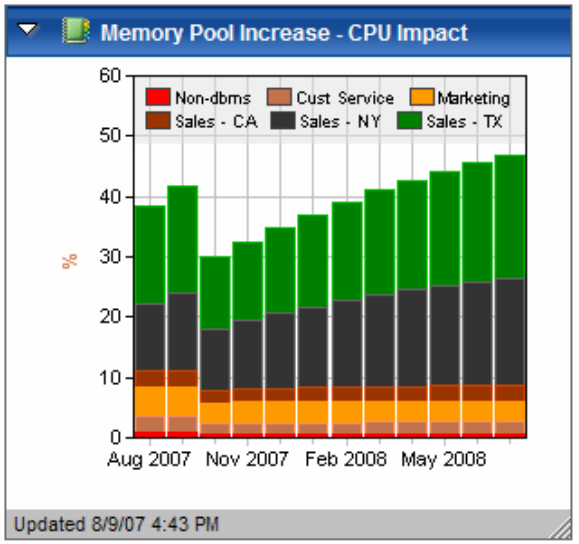

5) Modeling Results Show the Impact of the Workload Management, New Application Implementation and Resource Manager Changes on Mixed Workload Performance

Predicting How Changing the Workload's Priority Will Affect Performance

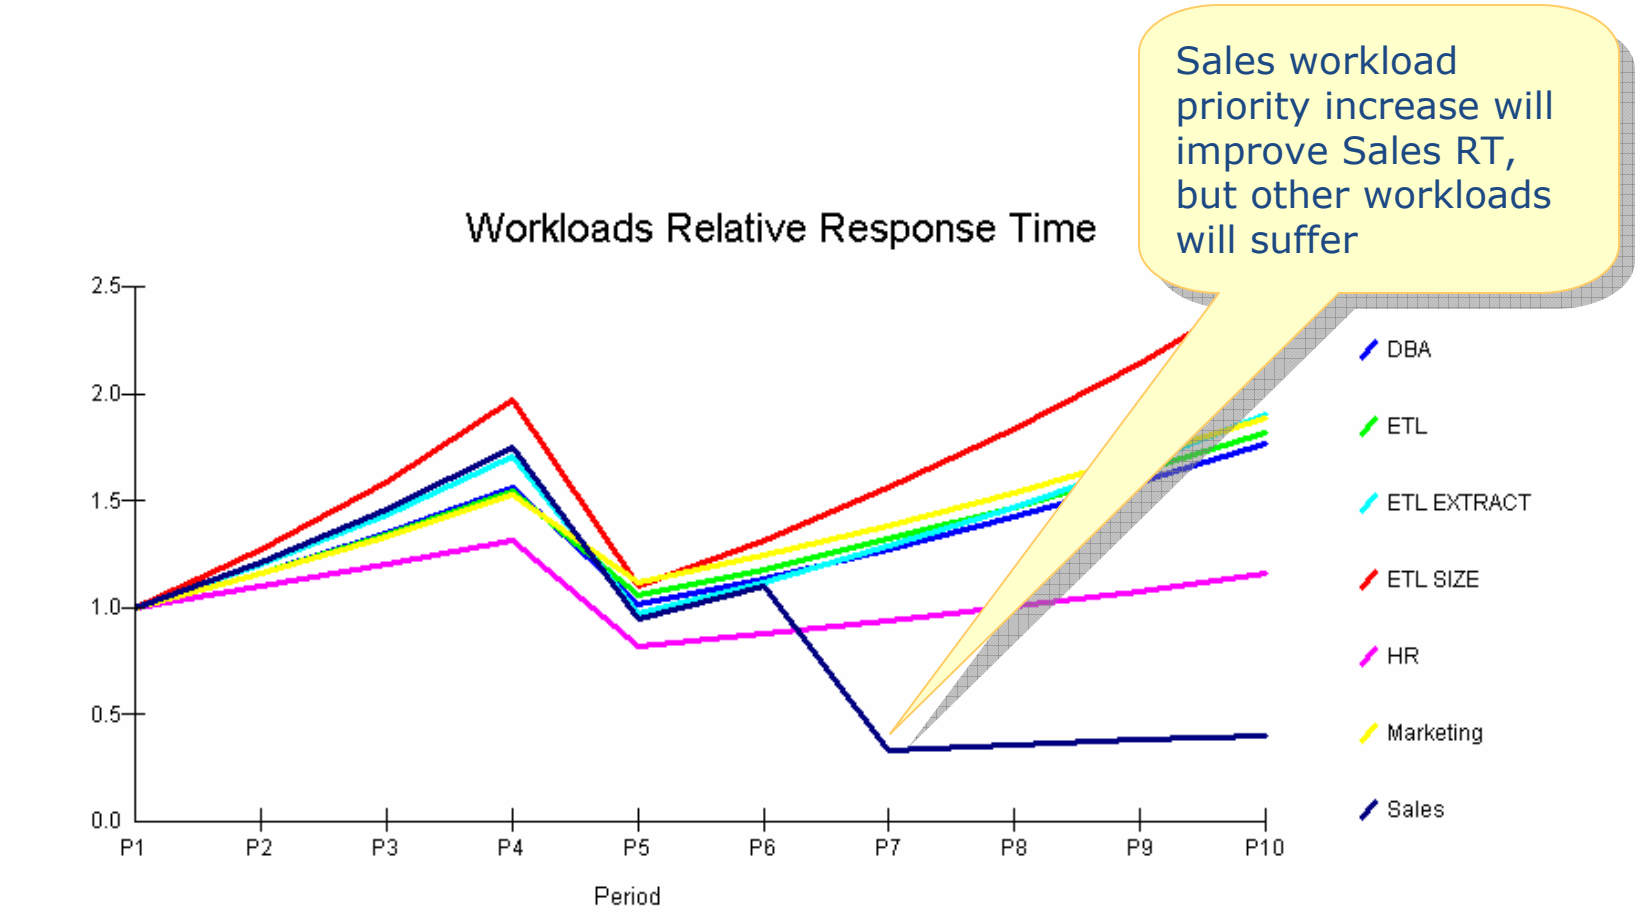

Predicting Impact of increasing # of ETL Utilities Loading Data in Parallel Starting Next Month (p2) and Hardware Upgrade (p5)

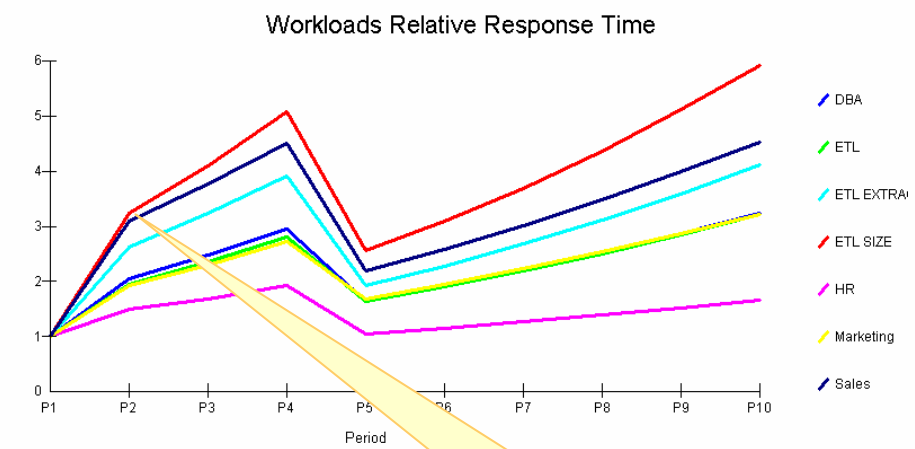

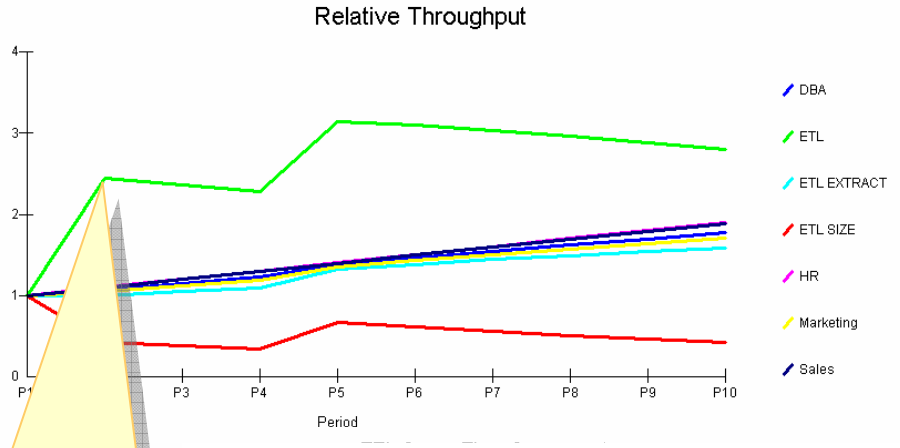

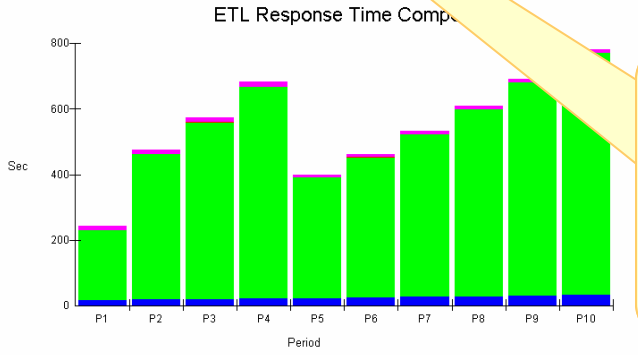

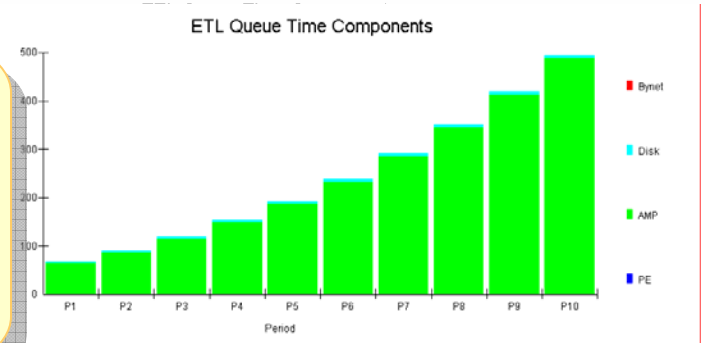

Increase in # of loads will allow a significant reduction of load time, but there will be a very significant elongation of the RT for HR, Marketing and Sales workloads

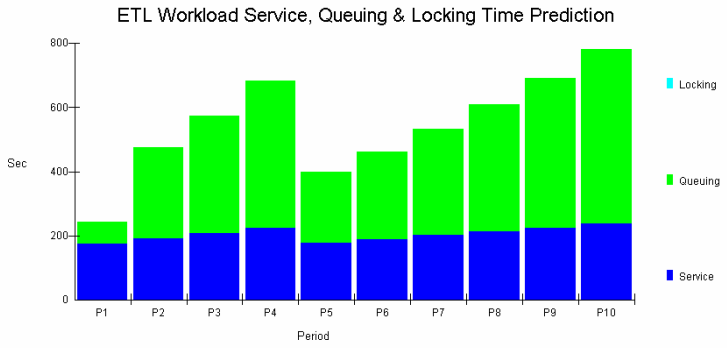

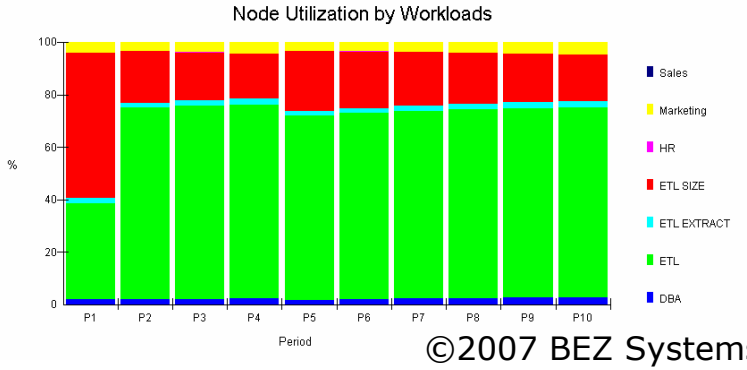

Predicted Impact of Adding a New Application

Set up realistic expectations and reduce risk of surprises

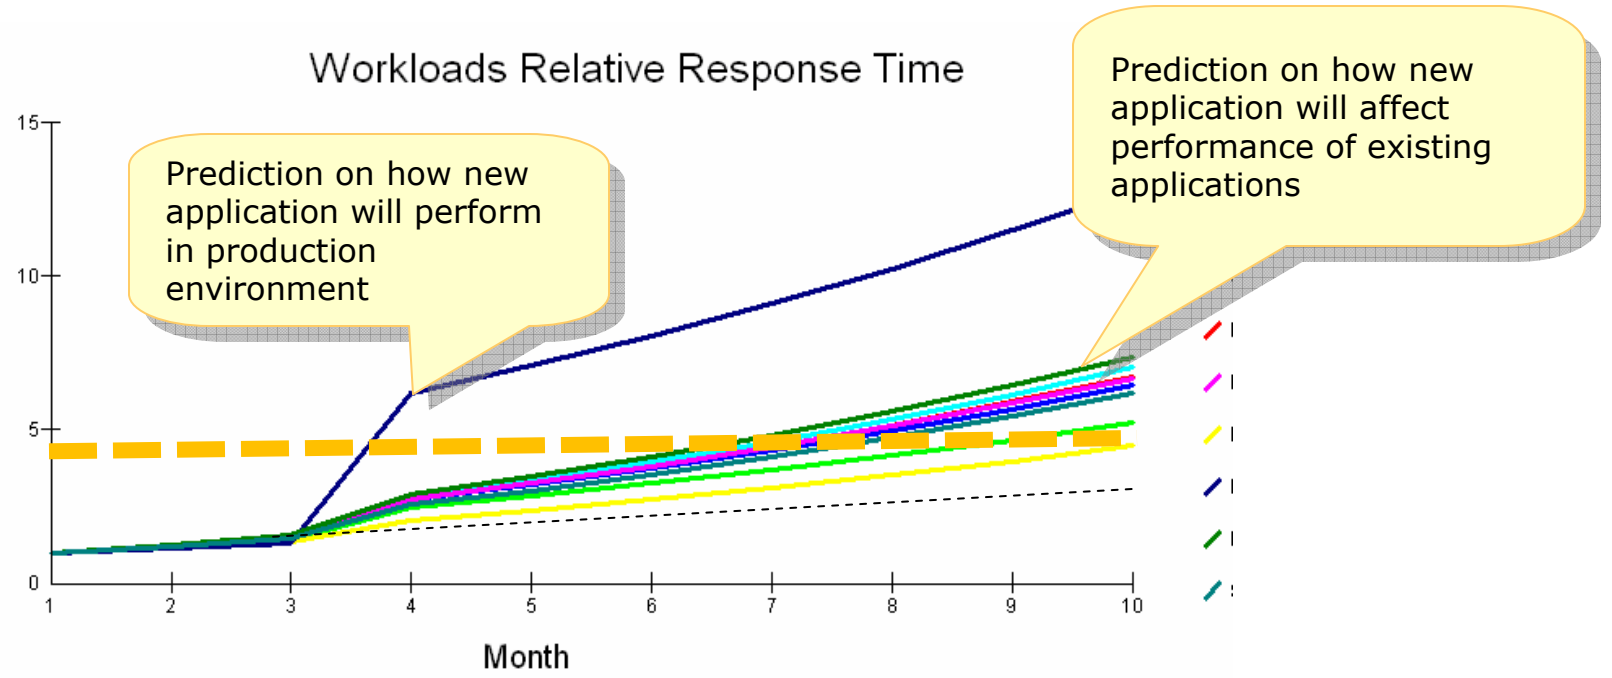

In a future Database Replay

Predicting New HR Application Implementation Impact

Workload Addition (HR) Scenario

This scenario models the addition of three new HR workloads to be merged with the existing workload mix to determine the performance and utilization impacts for this consolidation actiivty

Updated 8/9/07 5:00 PM

RT Impact Summary with HR added

Workload Name	Difference
Customer Service	+298%
Finance	+300%
Marketing	+303%
R&D	+315%
Sales-CA	+300%
Sales-NY	+343%
Sales-TX	+345%

Updated 8/9/07 5:00 PM

TP Impact Summary with HR added

Workload Name	Difference
Customer Service	0%
Finance	-6%
Marketing	0%
R&D	-300%
Sales-CA	0%
Sales-NY	-300%
Sales-TX	-400%

RT Projection: Existing WKLDs

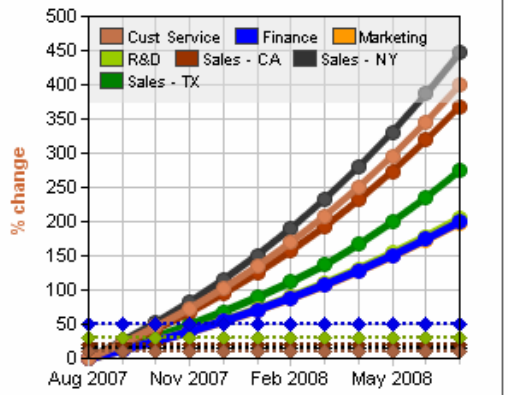

Updated 8/9/07 5:00 PM

RT Projection: Including HR WKLDs

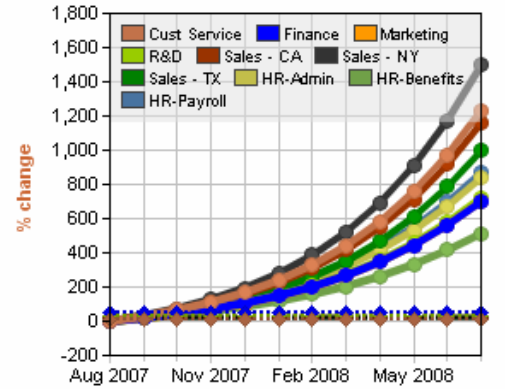

Updated 8/9/07 5:00 PM

TP Projection: Existing WKLDs

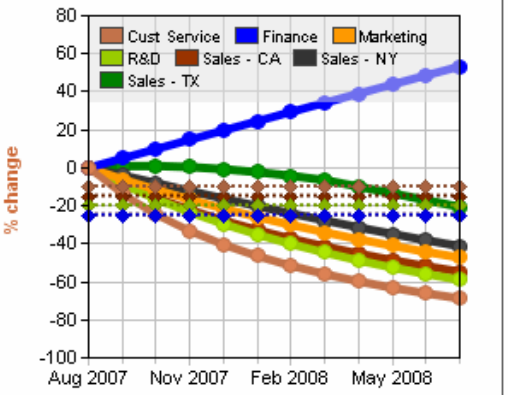

TP Projection: Including HR WKLDs

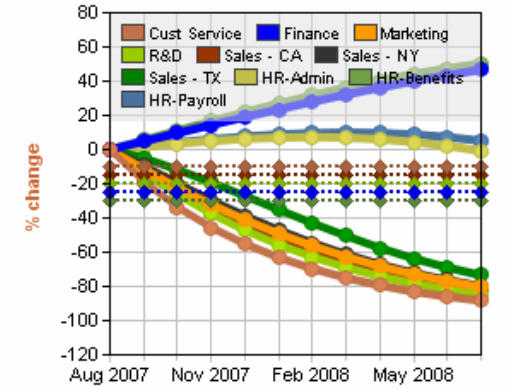

Predicting Workload Removal Impact

Workload Removal Scenario Summary

This scenario models the removal of one of the workloads (R&D) from consideration to determine the performance and utilization impacts to the remaining workloads

Updated 8/9/07 4:59 PM

RT Impact Summary

Workload Name	Difference
Customer Service	-11.43%
Finance	-11.50%
Marketing	-11.44%
Sales-CA	-11.47%
Sales-NY	-12.99%
Sales-TX	-12.54%

Updated 8/9/07 4:59 PM

TP Impact Summary

Workload Name	Difference
Customer Service	0%
Finance	-1.02%
Marketing	0%
Sales-CA	0%
Sales-NY	+25%
Sales-TX	+11%

RT Projection: All Workloads

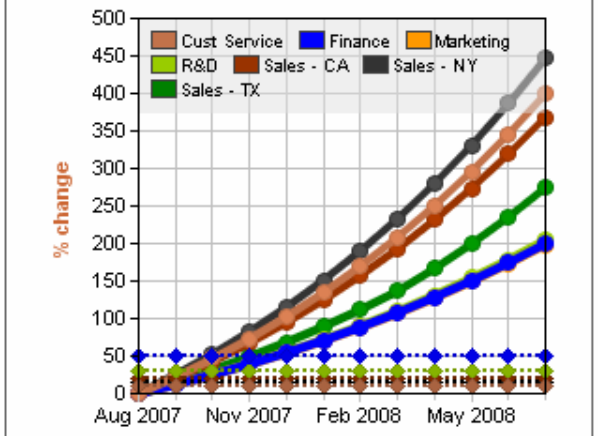

Updated 8/9/07 4:59 PM

RT Projection: R&D Removed

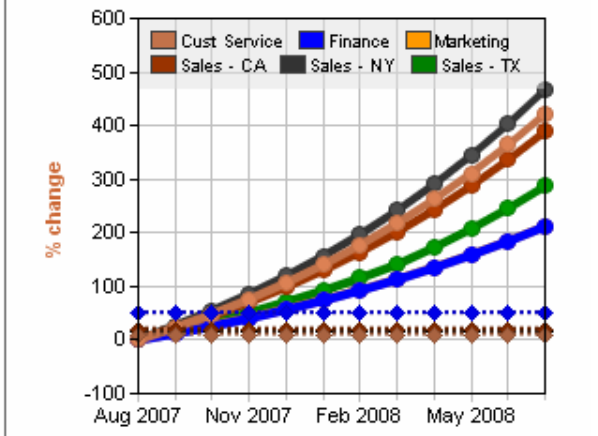

Updated 8/9/07 4:59 PM

TP Projection: All Workloads

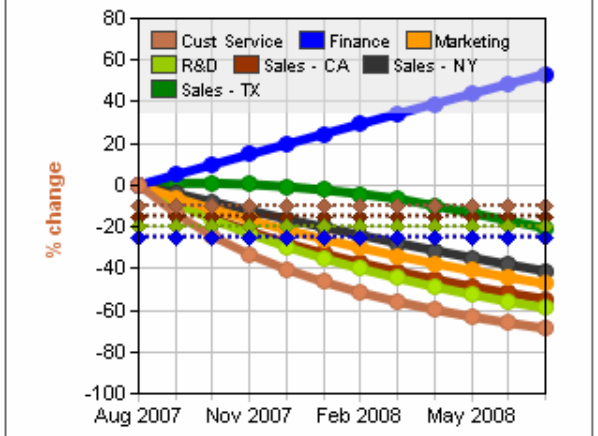

TP Projection: R&D Removed

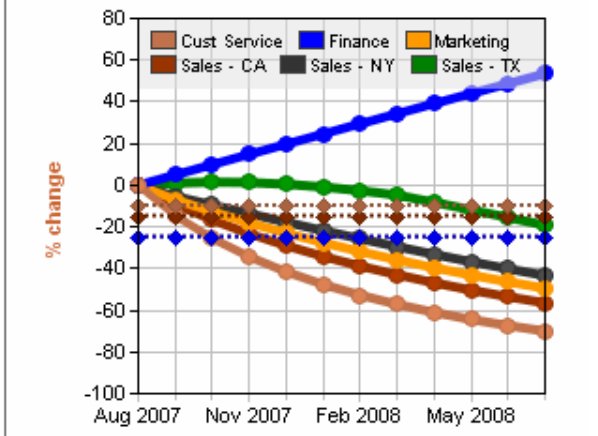

Example Showing How Modeling Results Were Used to Justify EDW and Server Consolidation by Financial Organization

Predicting How Data Mart Consolidation in EDW Will Affect ETL and Information Access Time

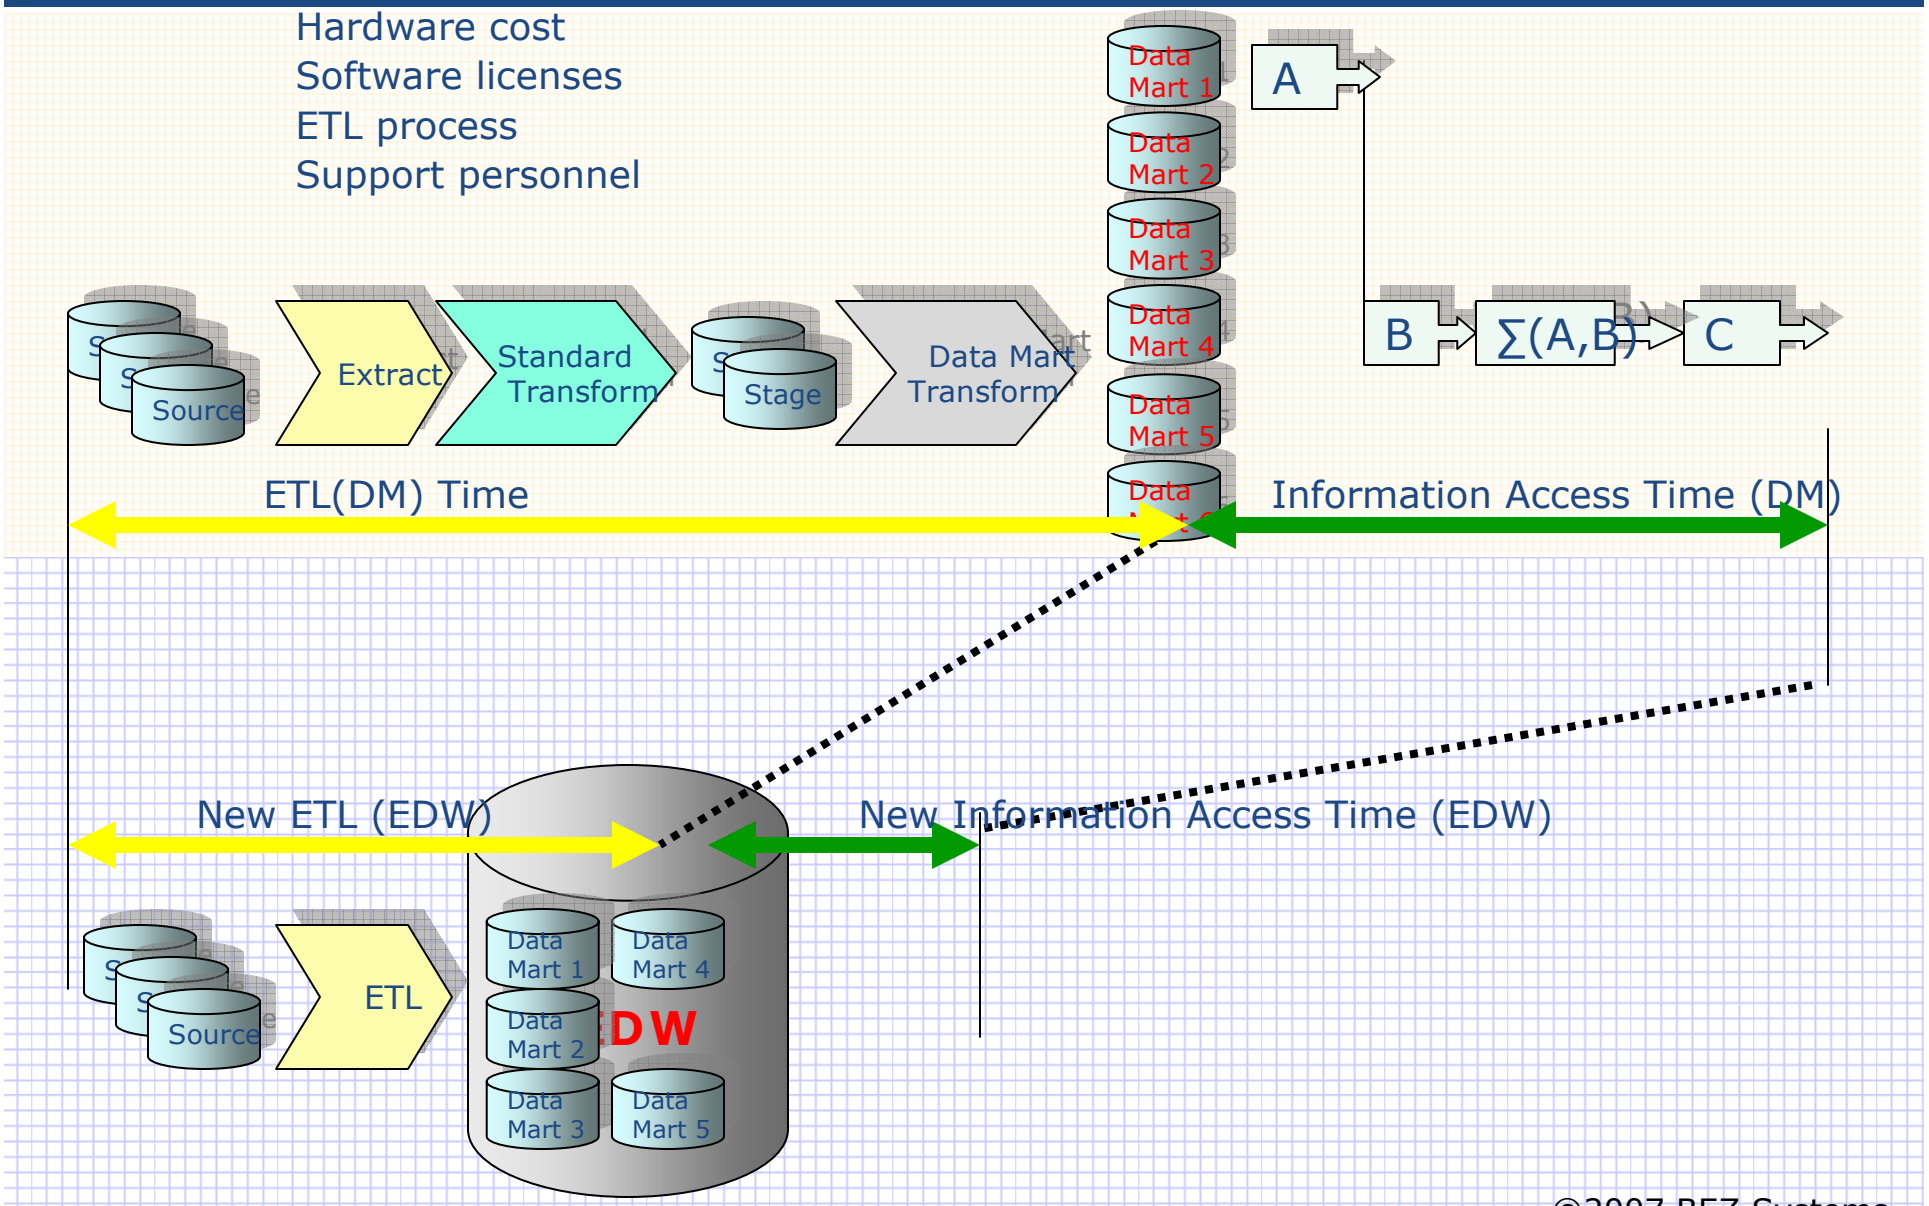

Evaluating Impact of Scaling Out to Different RAC Configurations

Predicting Impact of Different Vendor Platforms and Configurations on RAC Performance

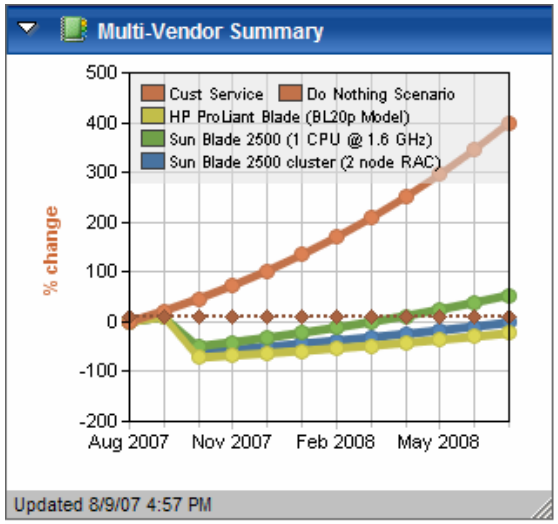

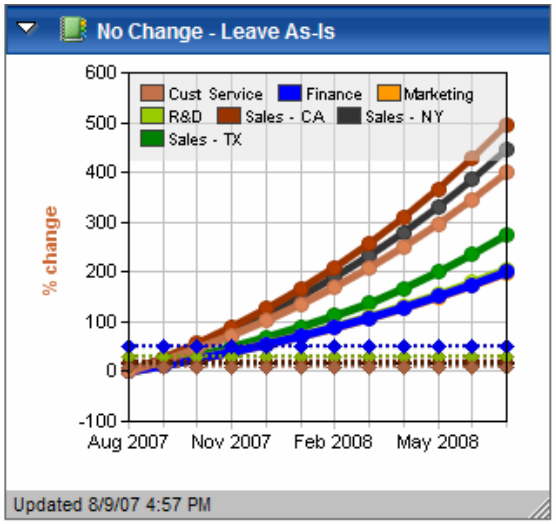

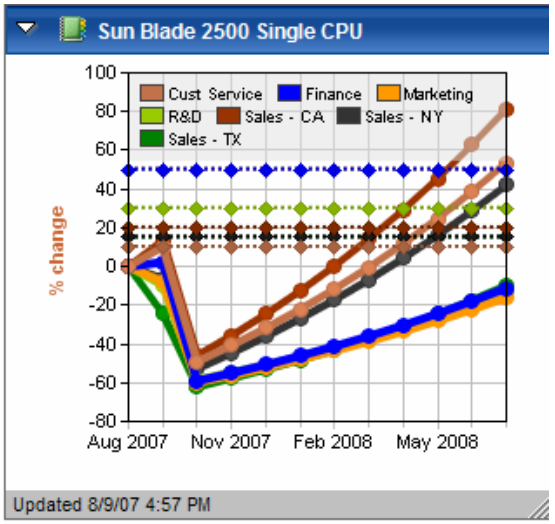

Insert tile here

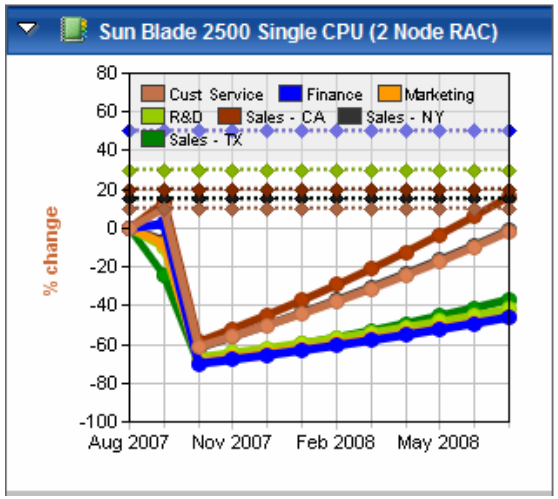

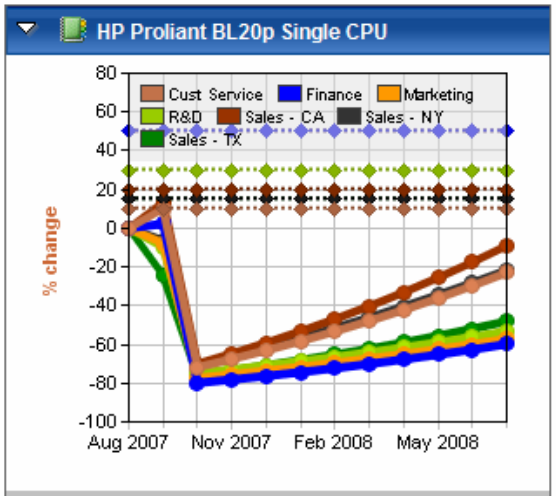

Predict the Impact of Expected Growth and Proposed Changes on Performance for Different Platforms

- Optimization of ETL process
  - ETL process is active 24 by 7
  - Requirement is to justify how many load utilities can run concurrently without affecting ability to satisfy SLO for critical workloads
- Optimization of the concurrency by controlling application server parameters
- Data compression
- Table partitioning
- New application implementation impact, modification / enhancement of application functionality
- How to proactively correct index, materialized views strategy
- Evaluate options and justify storage subsystem configuration

- Predict the impact of workload growth
- Predict the impact of database / table size growth
- Predict the impact of eliminating or moving the application / workload to another platform
- Model the impact of server consolidation

Data Loading Concerns

- How long will it take to load growing volume of data?
- How many loads should we run in parallel at night, and during the day time?
- How do we change priorities for ETL and other workloads to satisfy SLO?

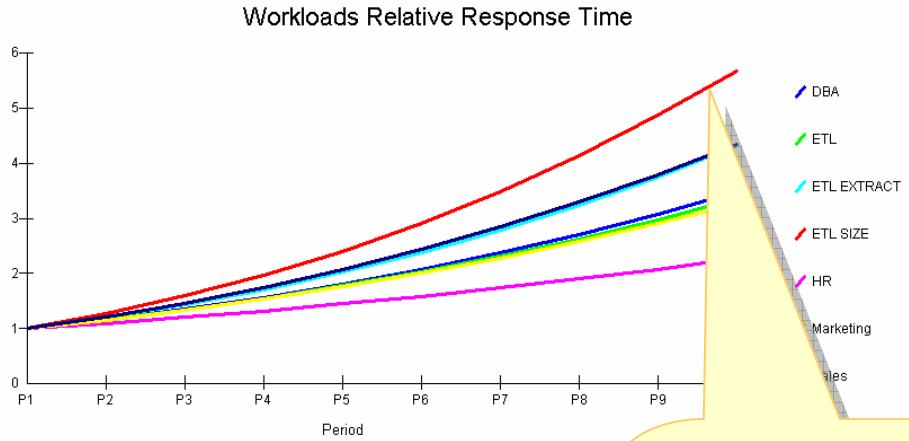

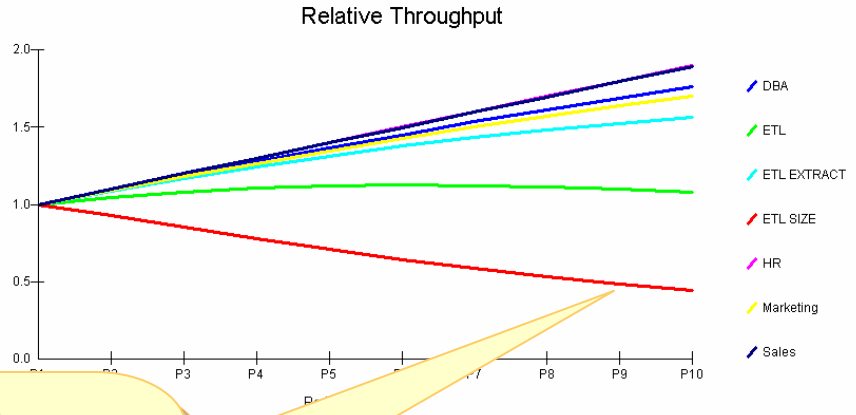

It will take 6 times longer to load growing volume of data in 10 months.
RT for HR application will increase almost 2 times & throughput for ETL will be reduced almost 2 times

What is the Minimum Hardware Upgrade Required to Load Growing Volume of Data On Time?

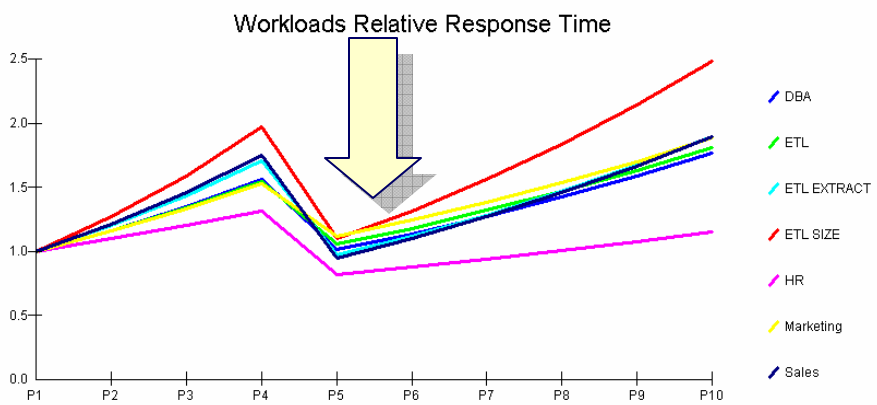

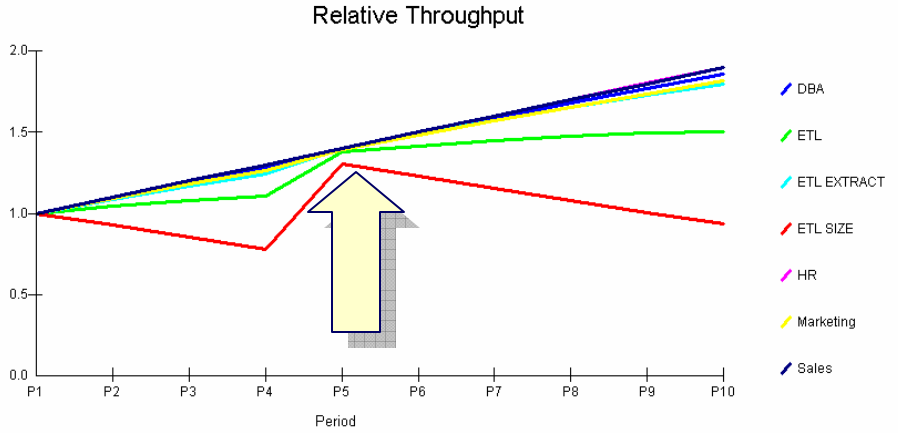

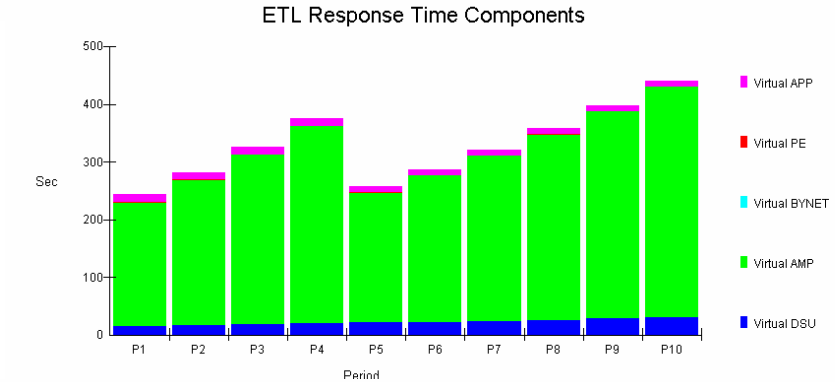

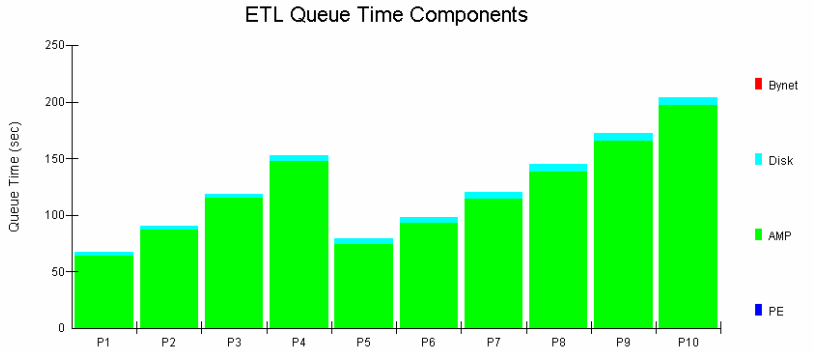

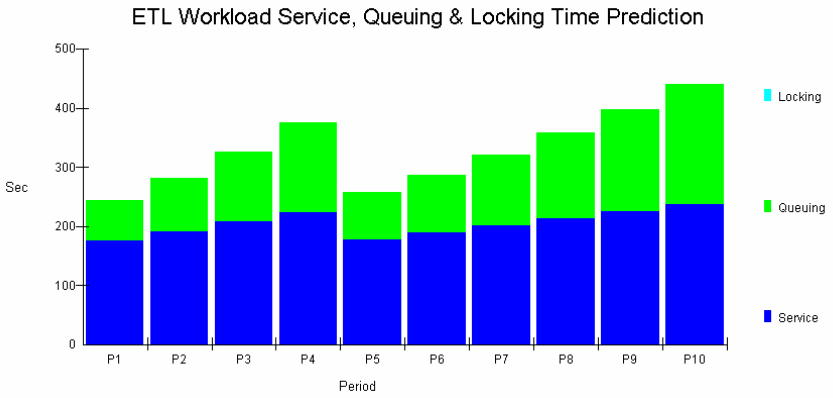

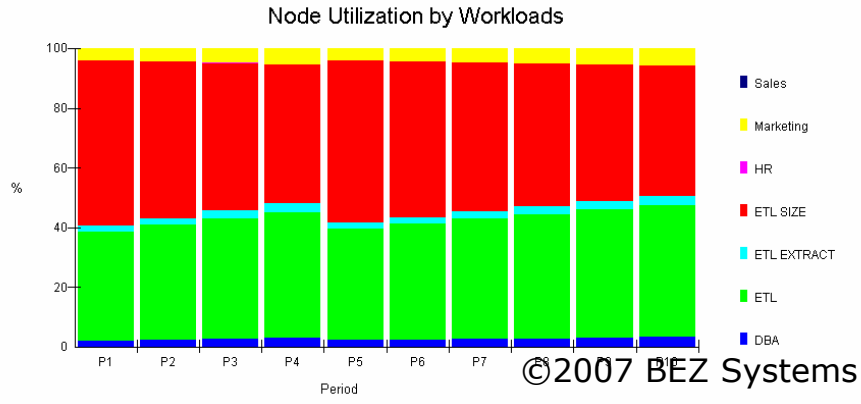

How to Identify Change of Data Usage Profile and Justify Proactive DB Changes

Each Workload Supporting Business Process Has a Unique Data Usage Profile

- Data size
- # Concurrent users
- # Concurrent SQL
- Data complexity
- Index use
- Data cardinality
- Locality of reference
- Data skewedness
- Complexity of SQL
  - %S/I/U/D
  - %SUM
  - #Joins
  - #Rows accessed & #rows retrieved

Tactical Performance Tuning Options

- Indexes
- Materialized views
- Data compression
- Data partitioning
- Parallel processing

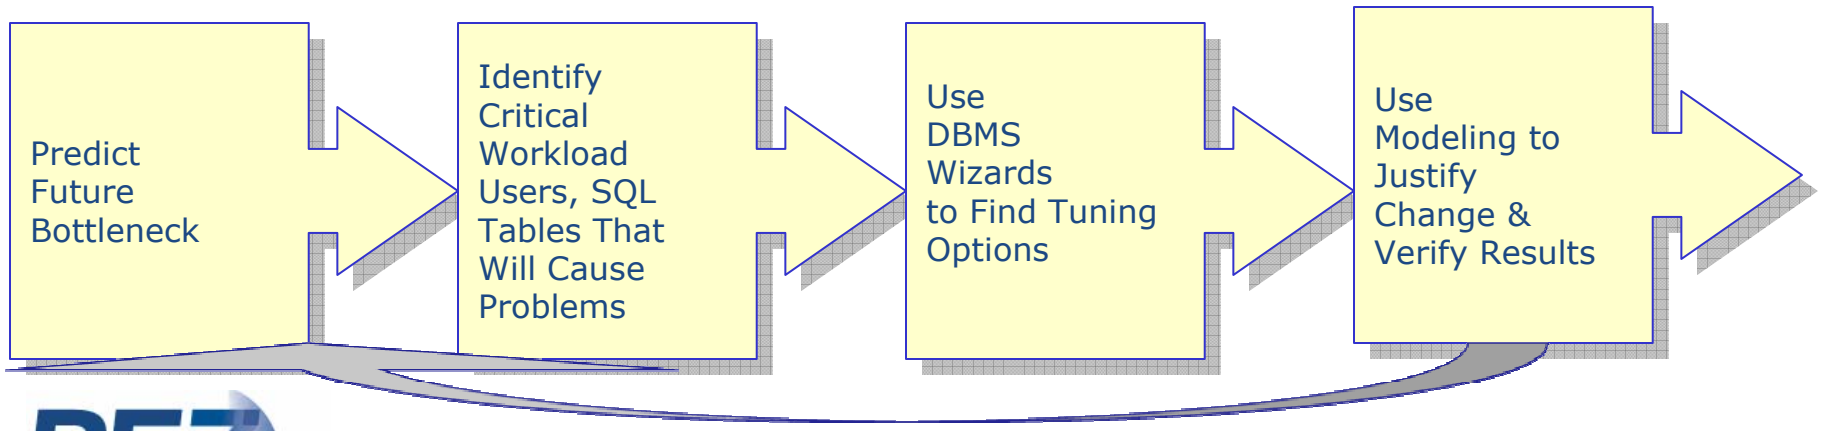

Performance Prediction Identifies SQL That Will Cause a Problem in the Future

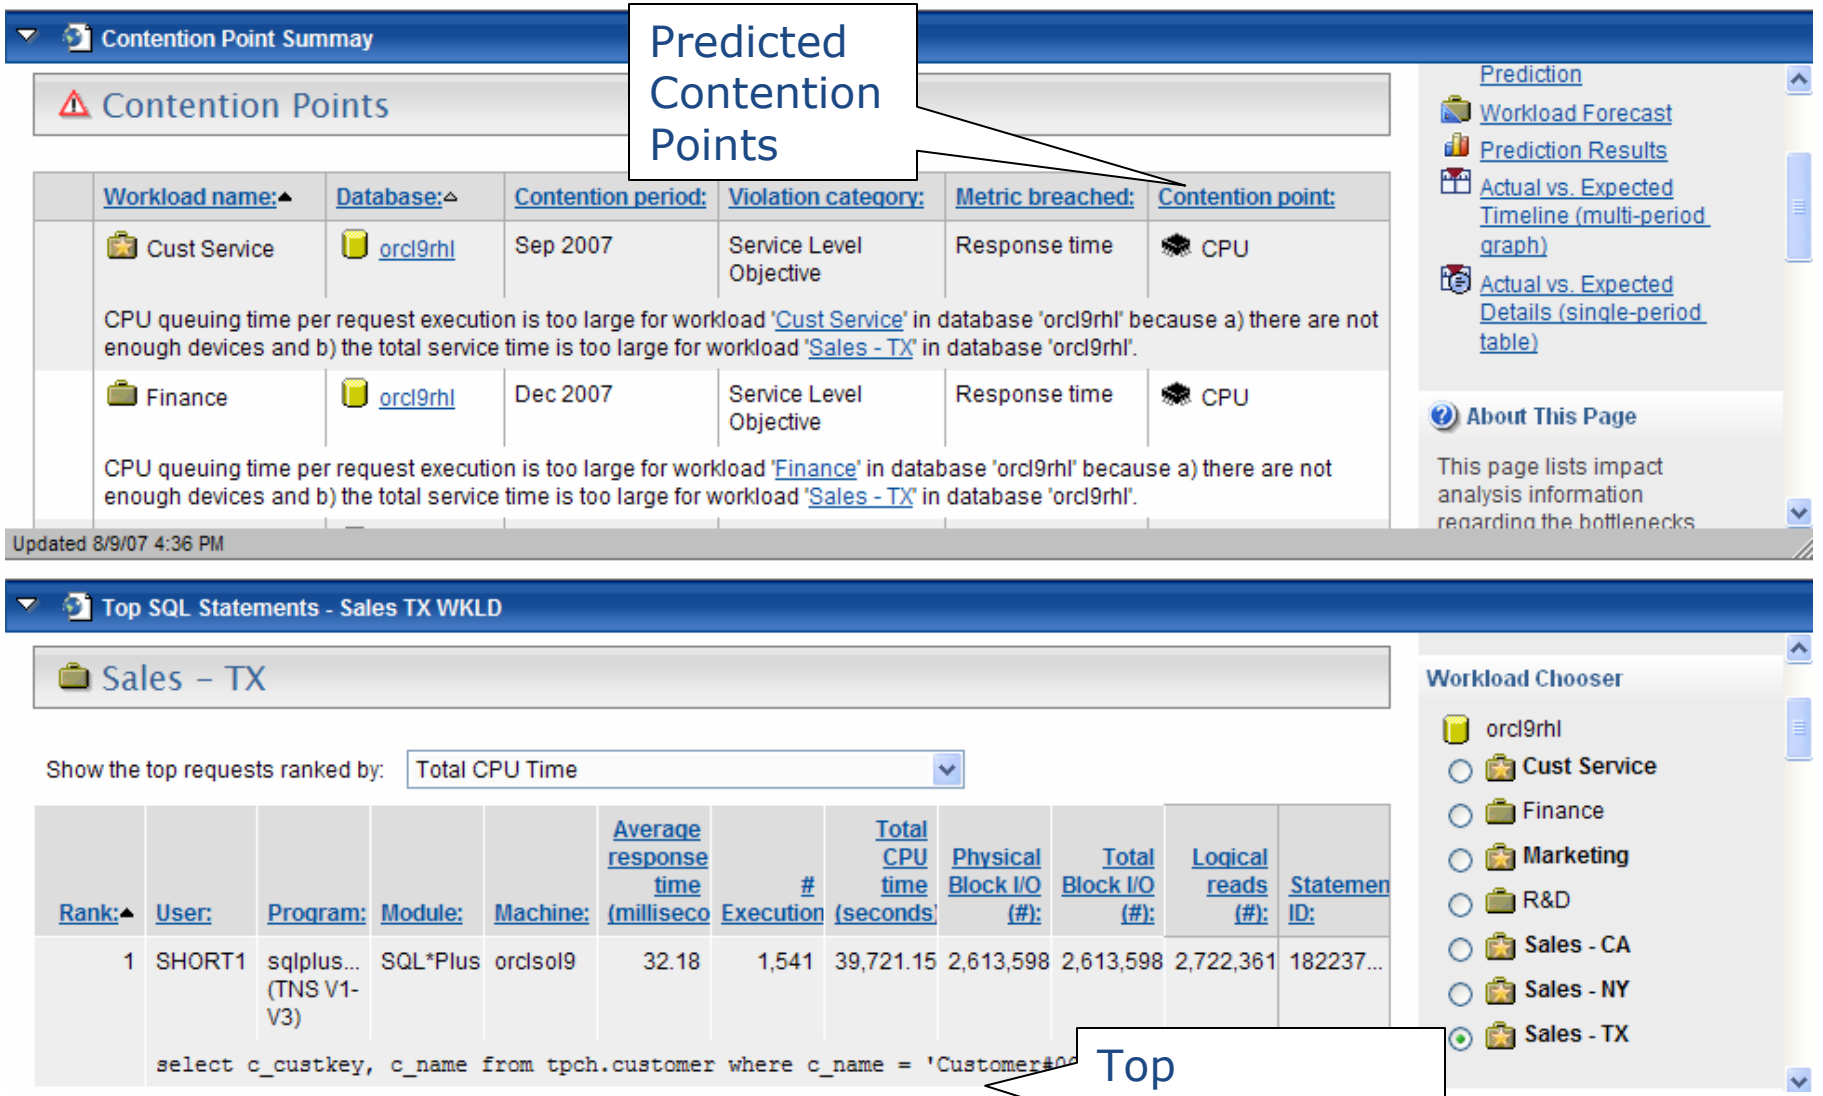

Processes SQL through SQL Access Advisor and Predicts the Impact of Recommended Changes

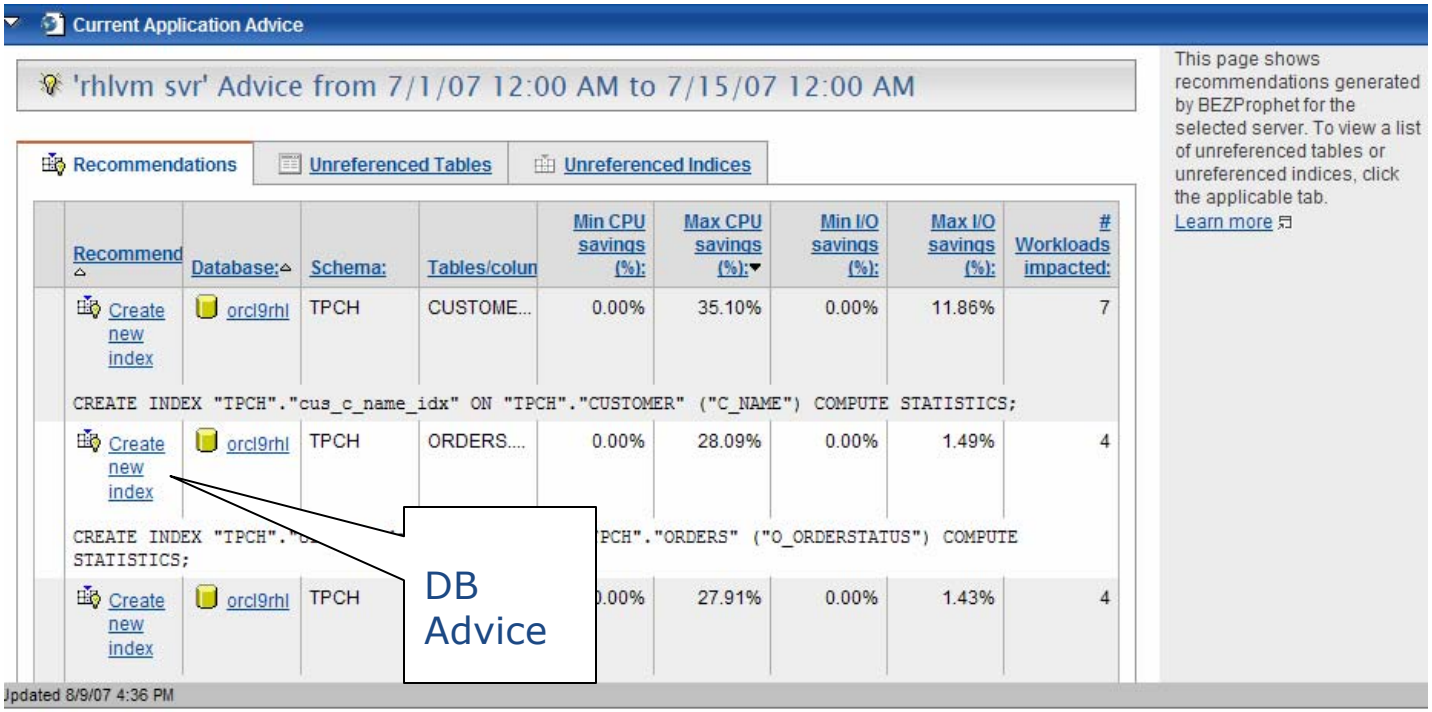

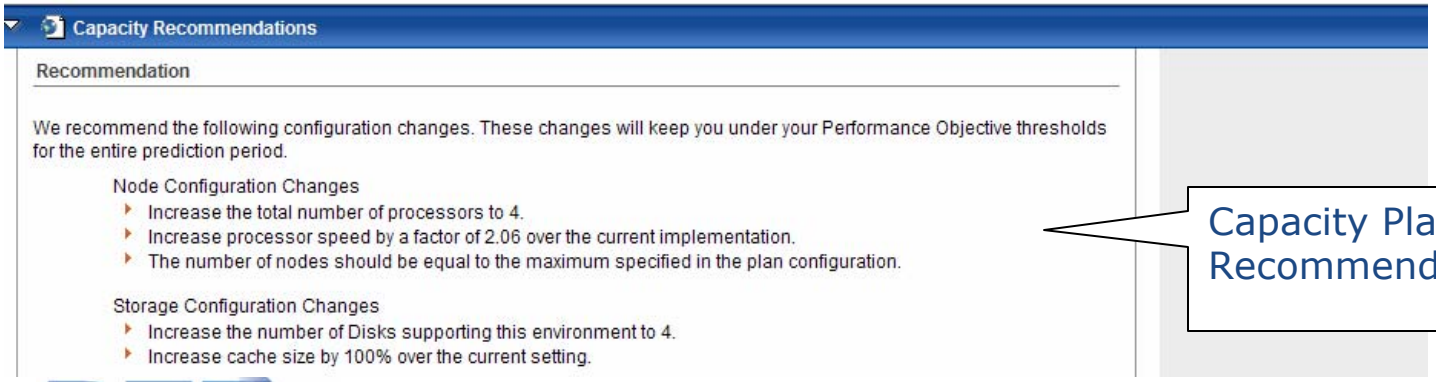

Performance Prediction Impact of Recommended Index

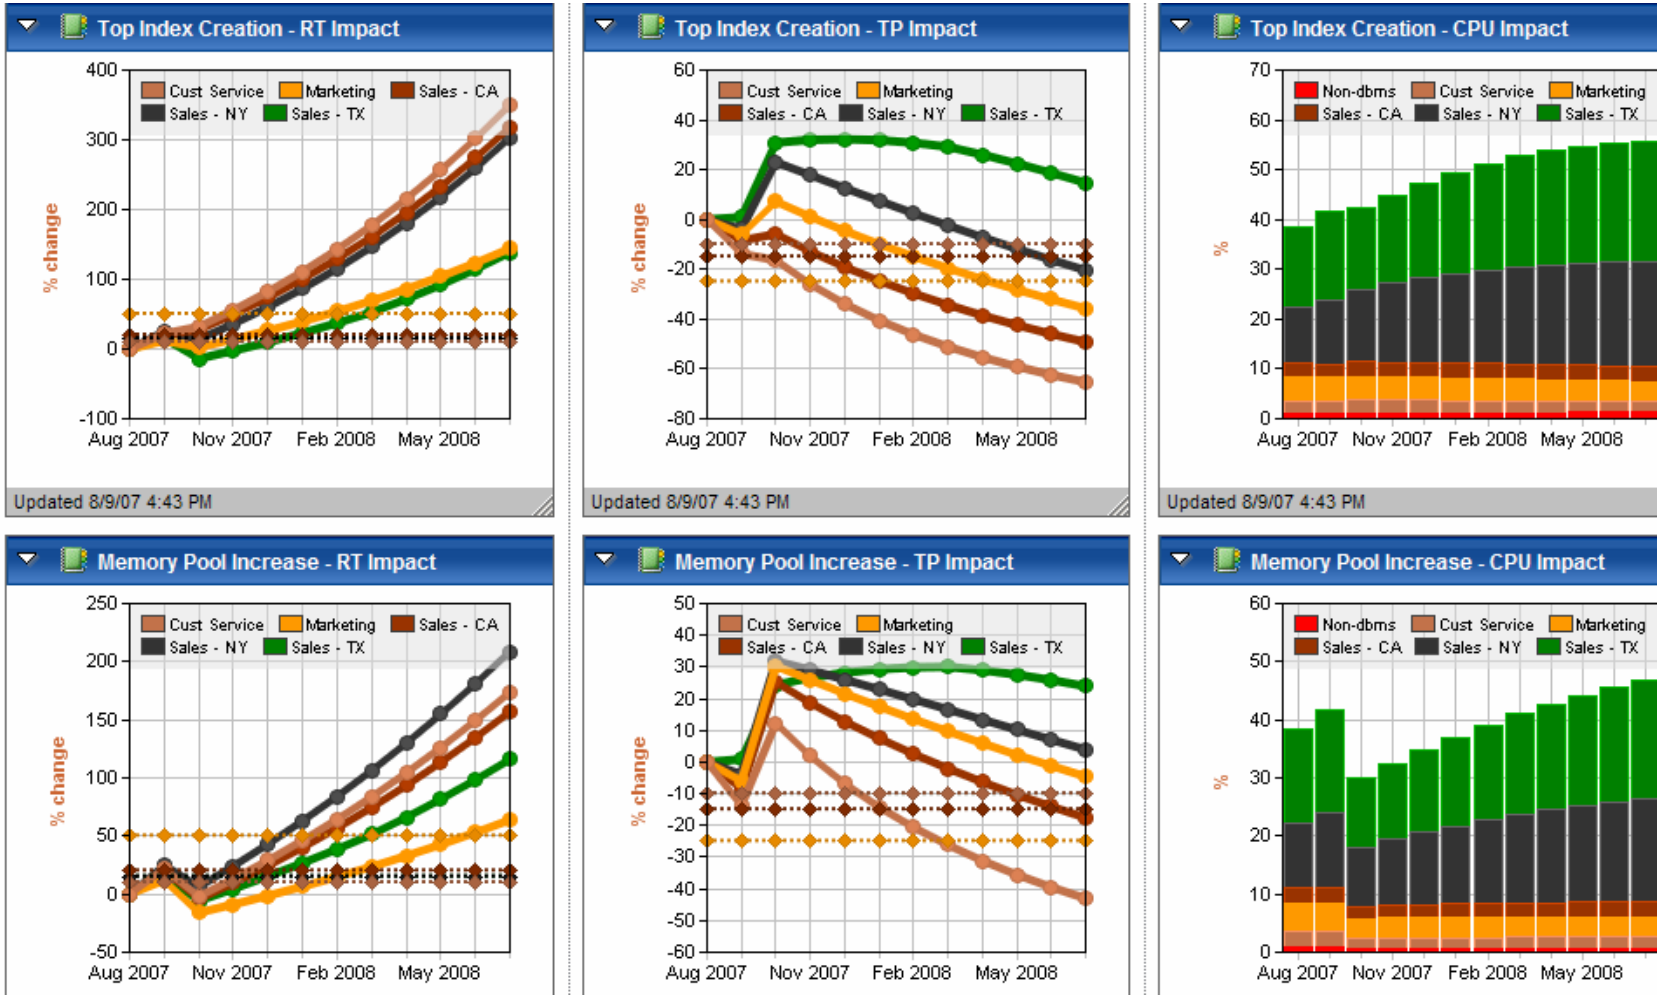

Predicted Impact of Adding New Index and Proposed Hardware Upgrade

Index implementation in 2 months will have a positive impact on 2 workloads.

Hardware upgrade in 4 months will significantly improve RT and throughput.

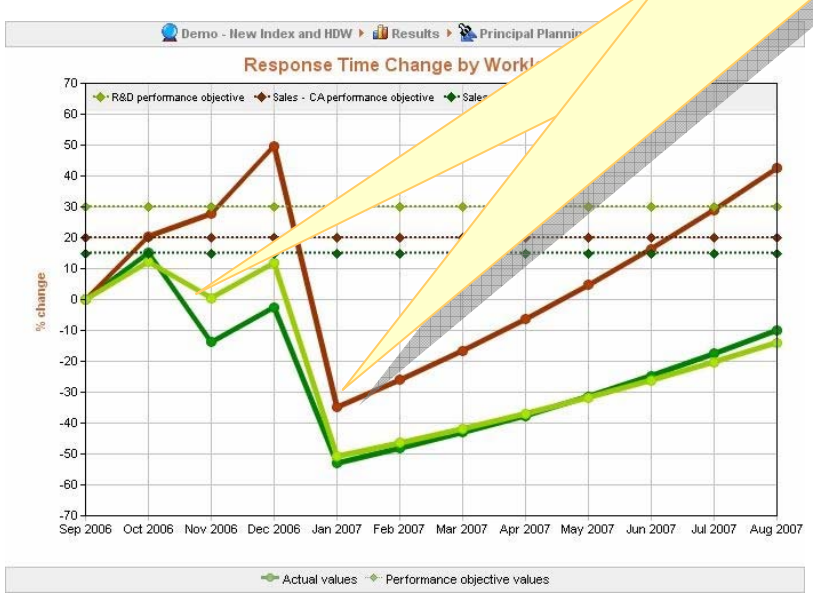

Data Compression Impact

- Oracle 11g expands DW data compression to OLTP tables and allows DML statements (INSERT, UPDATE, and DELETE) on compressed tables
  - Reduction in storage capacity demand
  - Increase in application performance
- Reduce physical IO and enhance cache efficiency
- As a result of data compression, users observe up to 50-70% savings in disk space for large tables
- Assumption that reduction in I/O rate to table is proportional to reduction in table space usage as result of data compression
- Data compression causes increase in CPU overhead due to data decompression for read operations and data compression for inserts and updates, and reduction in CPU overhead caused by reduction in number of I/O operations
- One of the important benefits of data compression is reduction of the volume of data that should be logged to support insert/update operations

Predicted Data Compression Impact on Different Workloads

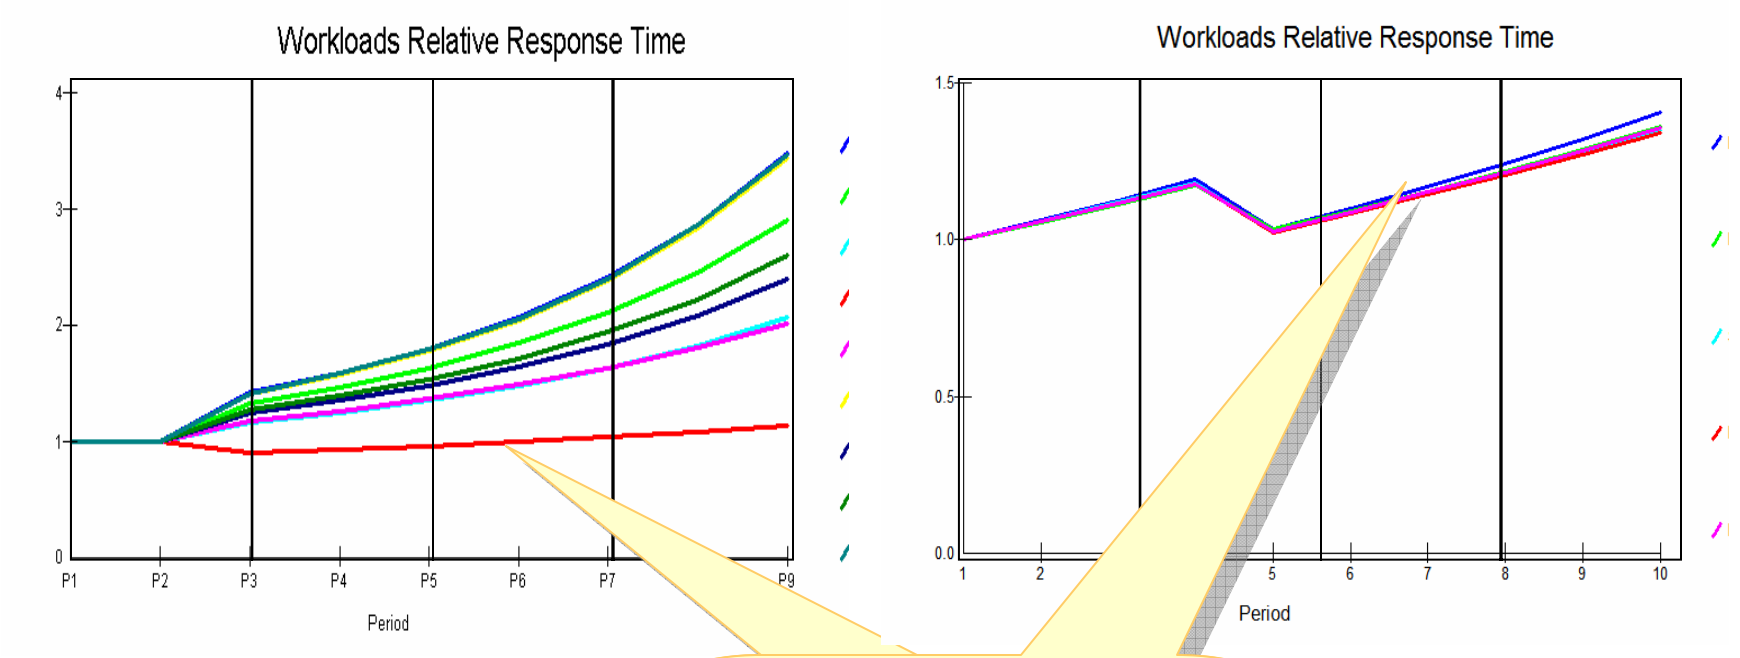

Data compression will have different impact on different workloads. DW workloads with primarily SELECT type of requests will benefit more.

Predicted Impact of Data Partitioning

- In Oracle 11g Interval Partitions, an extension of Range Partitions
  - DBA specifies the interval in which Oracle will automatically create partitions
  - You can now partition by date, one partition per month for example, with automatic partition creation
- It is significant improvement of manageability, but partitioning also can have a positive impact on performance
- Performance prediction results can be used to estimate potential improvements, set expectations, which can be compared with actual results

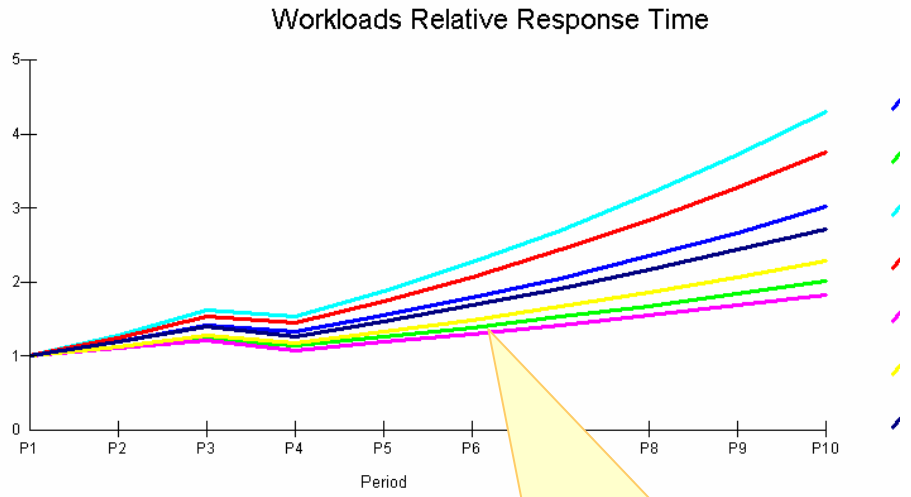

Data partitioning will have a positive impact on performance for all workloads

Example Showing How Modeling Results Can Be Used to Predict the Impact of Implementation of New Oracle 11g Features on Performance of Each of the Workloads

Predicting Impact of Oracle Server Result Cache

- Server Result Cache is a new component of the SGA that caches 'results' of queries and query fragments.
- These cached results are then used during future executions of similar queries or query fragments to bypass regular query processing and return the results faster.
- The cached results are completely shareable between sessions and SQL statements – as long as they share common execution plans, either partially or fully – and persist beyond the life of the initiating cursor.
- Assuming that Server Result Cache Advisor in 11g will have estimation of the Server Cache Hit Ratio, Prediction Engine will be able to predict impact of Server Result Cache on performance of each workload

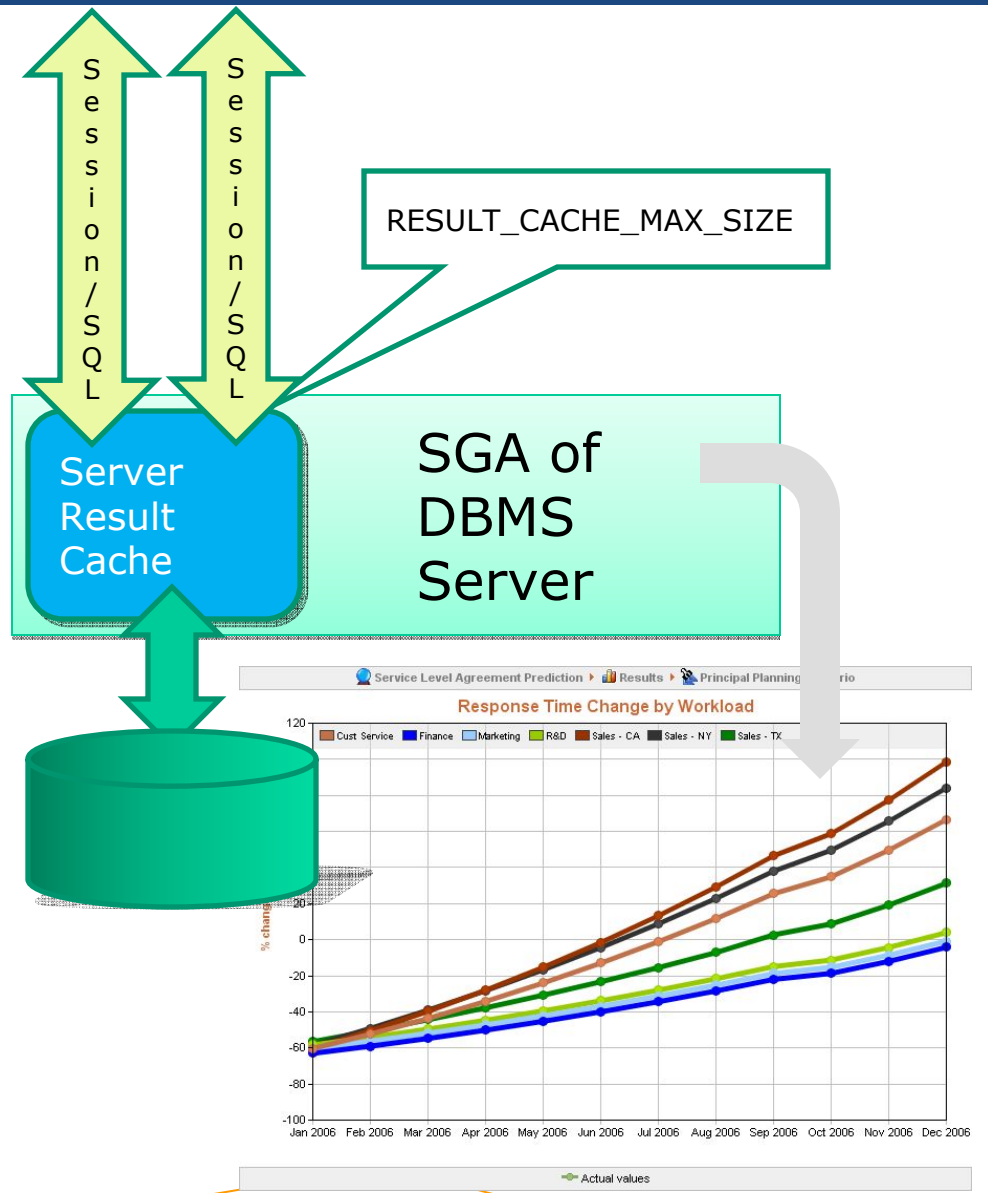

Client Cache Technology

- Client Cache technology reduces round-trips between the client and database server – reducing server CPU utilization
- The Client Cache contents can be shared across multiple sessions and/or threads
- Utilizing the Client Cache requires fewer round-trips between the client and database server, thus reducing CPU utilization on the server as a result of executing fewer SQL calls
- The Client Cache is optimal for queries of small lookup tables that are generally read-only or read-mostly
- Assuming that Client Cache Advisor in 11g will have estimation of the reduction in number of round-trips, modeling will be able to predict impact of Client Cache on performance of each workload

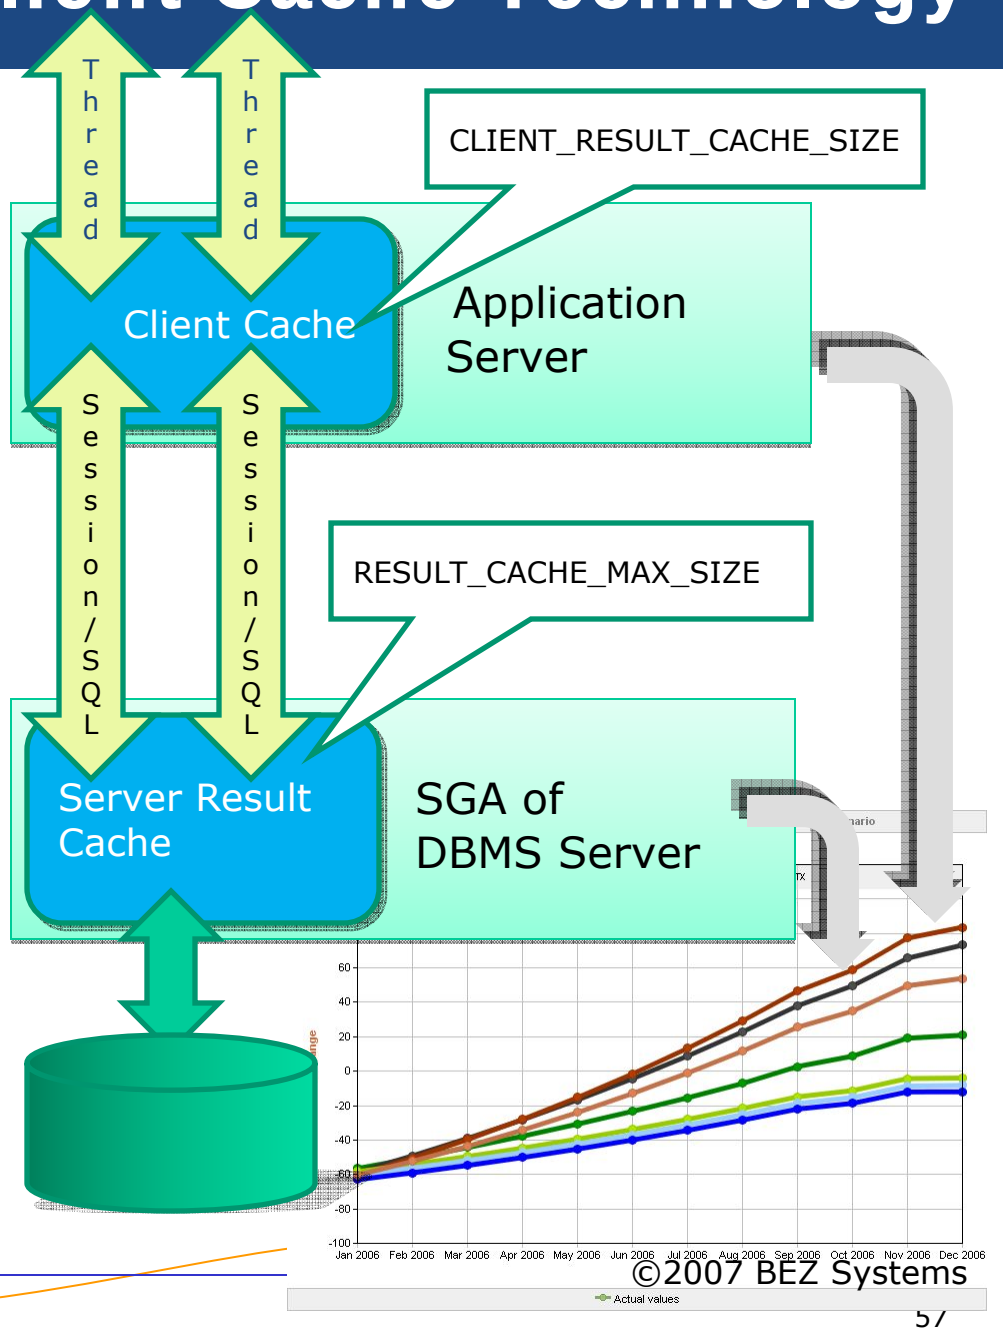

Comparison of the Actual Results with Expected and Organization of a Continuous Proactive Performance Management Process

BEZ
a business lens on IT

Comparison of Actual Results vs. Expected Enables Continuous Proactive Service Level Management

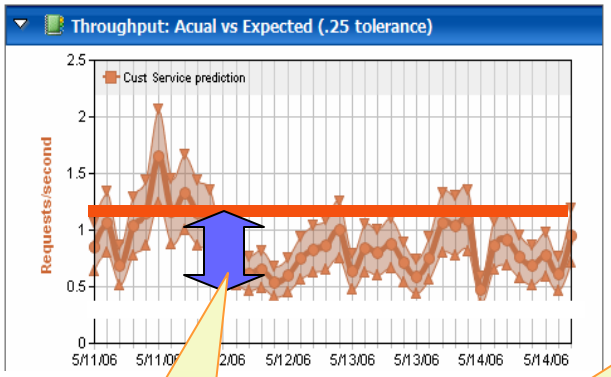

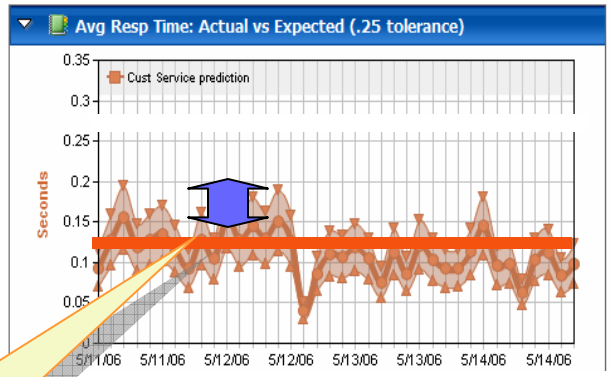

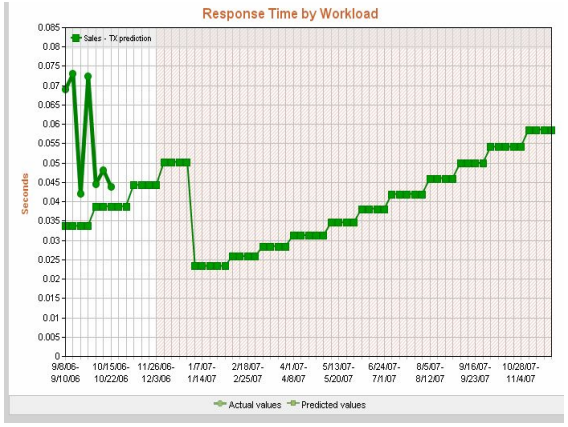

- Find difference between predicted results or expectations (red line) and actual measurement data
- Track how often the actual results do not meet expectation (SLA)
- When number of exceptions exceeds the threshold, generate alert
- Explain difference and develop new corrective recommendations

Comparing New Measurement Data to Expectations

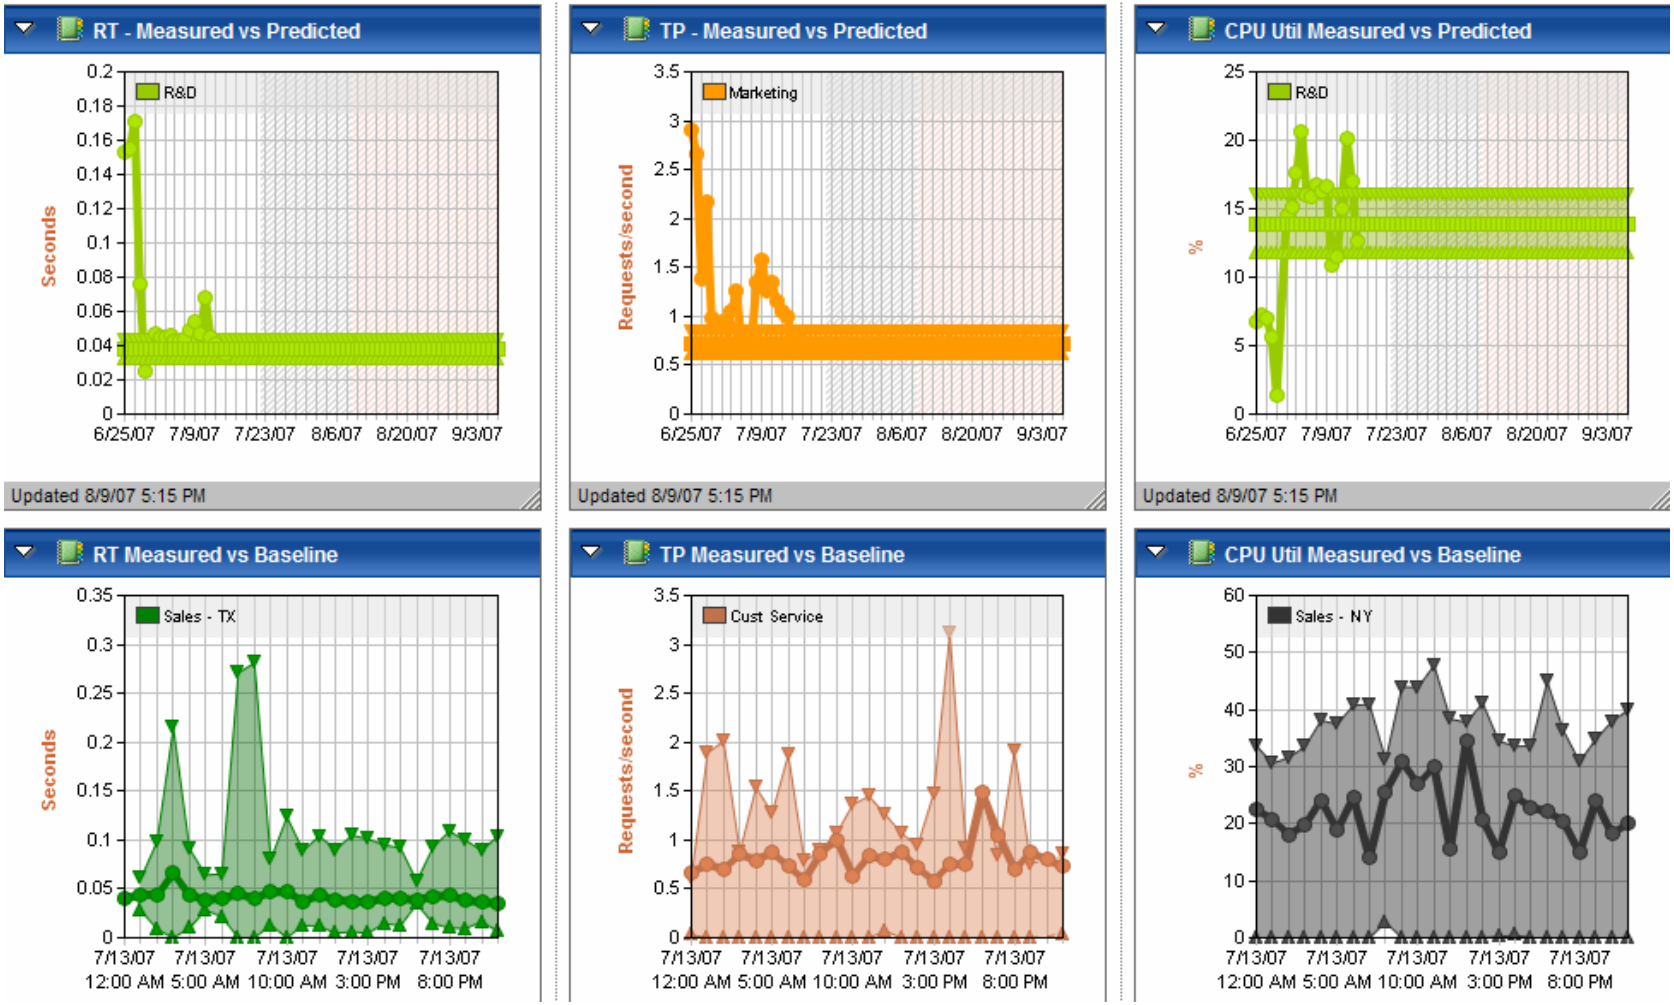

Conclusion

Conclusion

- We reviewed different factors affecting RAC performance in EDW environment with mixed workloads, including Contention for Shared Storage, DOP, Contention for Interconnect, Memory management, Workload management
- We discussed how modeling results can be used to evaluate and justify strategic, tactical and operational decisions, and minimize risk of performance surprises, including:
  - Platform selection for DW and EDW
  - New applications implementation
  - Hardware and storage upgrades
  - Disaster recovery
  - Data compression
  - Database tuning
  - Data partitioning
  - Server consolidation
- Presented methodology enables organization of the continuous proactive performance management process
- This process is based on automation of workload characterization, modeling, optimization and setting realistic expectations, and automation of comparing the actual results with expected and applying necessary changes before it is too late.

Thank You!

Dr. Boris Zibitsker
boris@bez.com
www.bez.com

BEZ
a business lens on IT

Thank You!

Questions?

Dr. Boris Zibitsker
boris@bez.com
www.bez.com
312-849-1583

BEZ
a business lens on IT