

Oracle

Tool

Tuning Frontier

2002.4.19



Agenda

...?

Tuning

Tuning

Tool

Tool

Tuning Tool

Performance Insight

Performance Insight

...?

- 1

가 ,

- 2



가

Transaction 가



- 3

가


- 4




- 1



- 2



Reasons for Poor Performance



Source: ORACLE Performance Tuning

1

Tuning


- 1 Tuning

- , ,
- 2) Response Time.

- 1 Tuning

Business , 가 ,

- 3 Tuning



,

- 4 Tuning

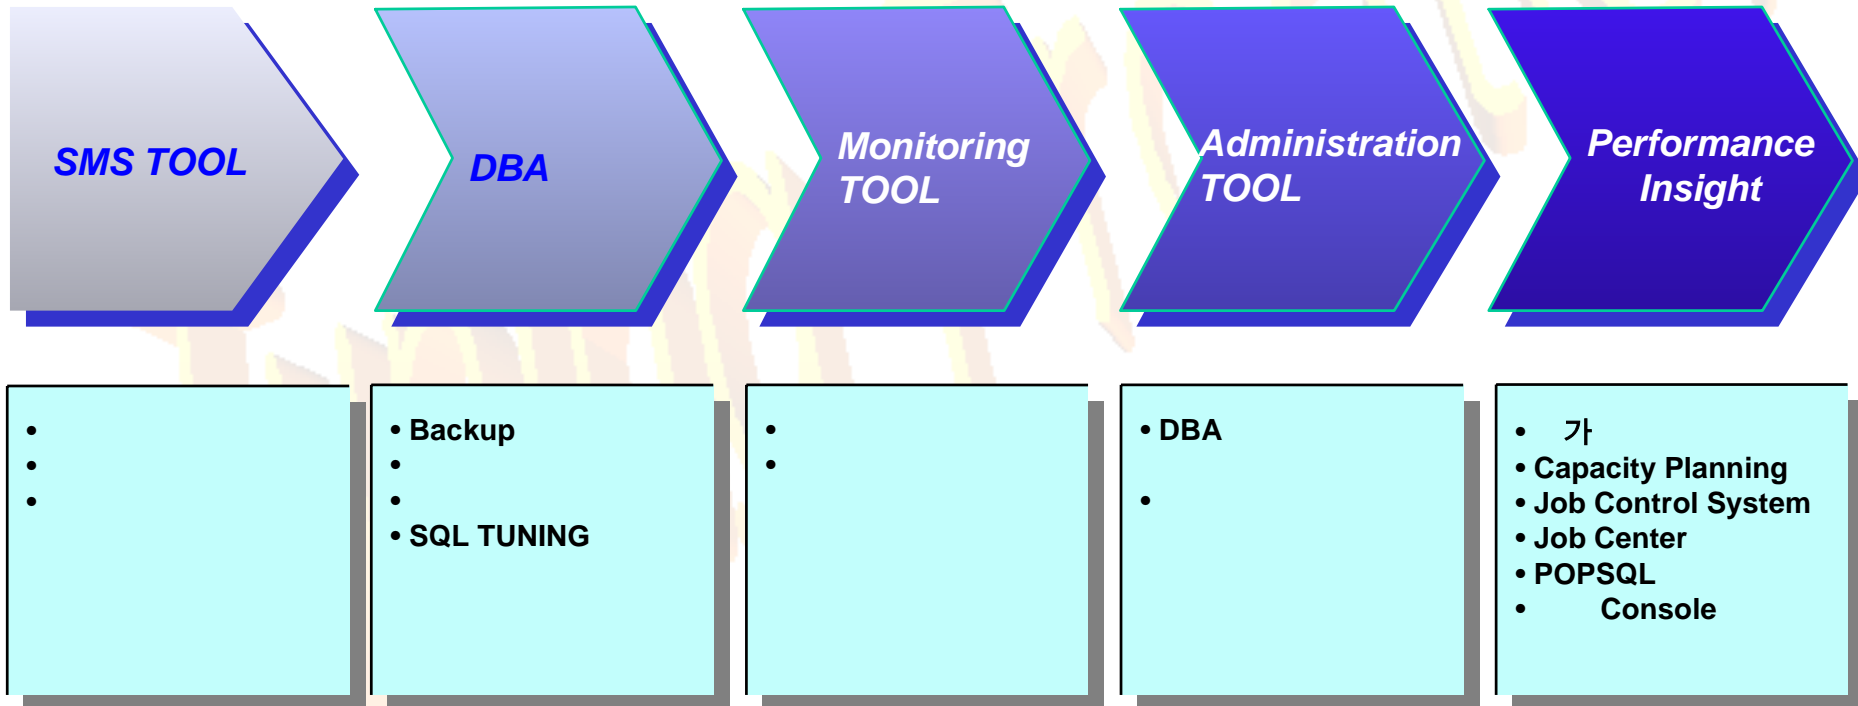
Tuning , Tuning
Tuning , Tuning

- 4 Tuning

: , , ,
: Tuning , ,
, 가

Tuning Tool

- 1 Tuning



- 2 Tuning Tool

Tuning

Tool

* 가

*
* Response
*
*

- 3 Tuning Tool



Tuning Tool

- 1 Tuning Tool


SQL	Current Data History (, , ,) 가	
	CPU	
Current	Historical	DISK I/O Buffer Wait Table/Index Full Scan Sort Merge Join Hash Join Parallel Option
	Manual/Automatic	

P/G(Service)	Batch
(Explain Plan)	SQL
SQL Tuning	Object Table

- 2 Tuning Tool


Tuning Tool , 가
Tool 가

Methodology : Performance




Implementation Support Feed-back	Performance 가
SQL Tuner Support Reporting Supplied Utilities	가 가 Performance
Customization Open Architecture	
Itself Easy Operation Work Load	

Knowledge - based Tuning Report
Integrated System Monitoring
Tuning Hint Message




H/W

Automated Tuning Report
Concentrated System Monitor
Customization
Capacity Planning Tool



Quality

Reliable Tuning Report
Capacity Planning Tool




User Response Time
가

- 1 Intro. Performance Insight
- 2. Performance Insight For Oracle
- 3. Performance Insight
- 4. Performance Insight
- 5. Performance Insight Architecture

- 1 Intro. Performance Insight

Performance System

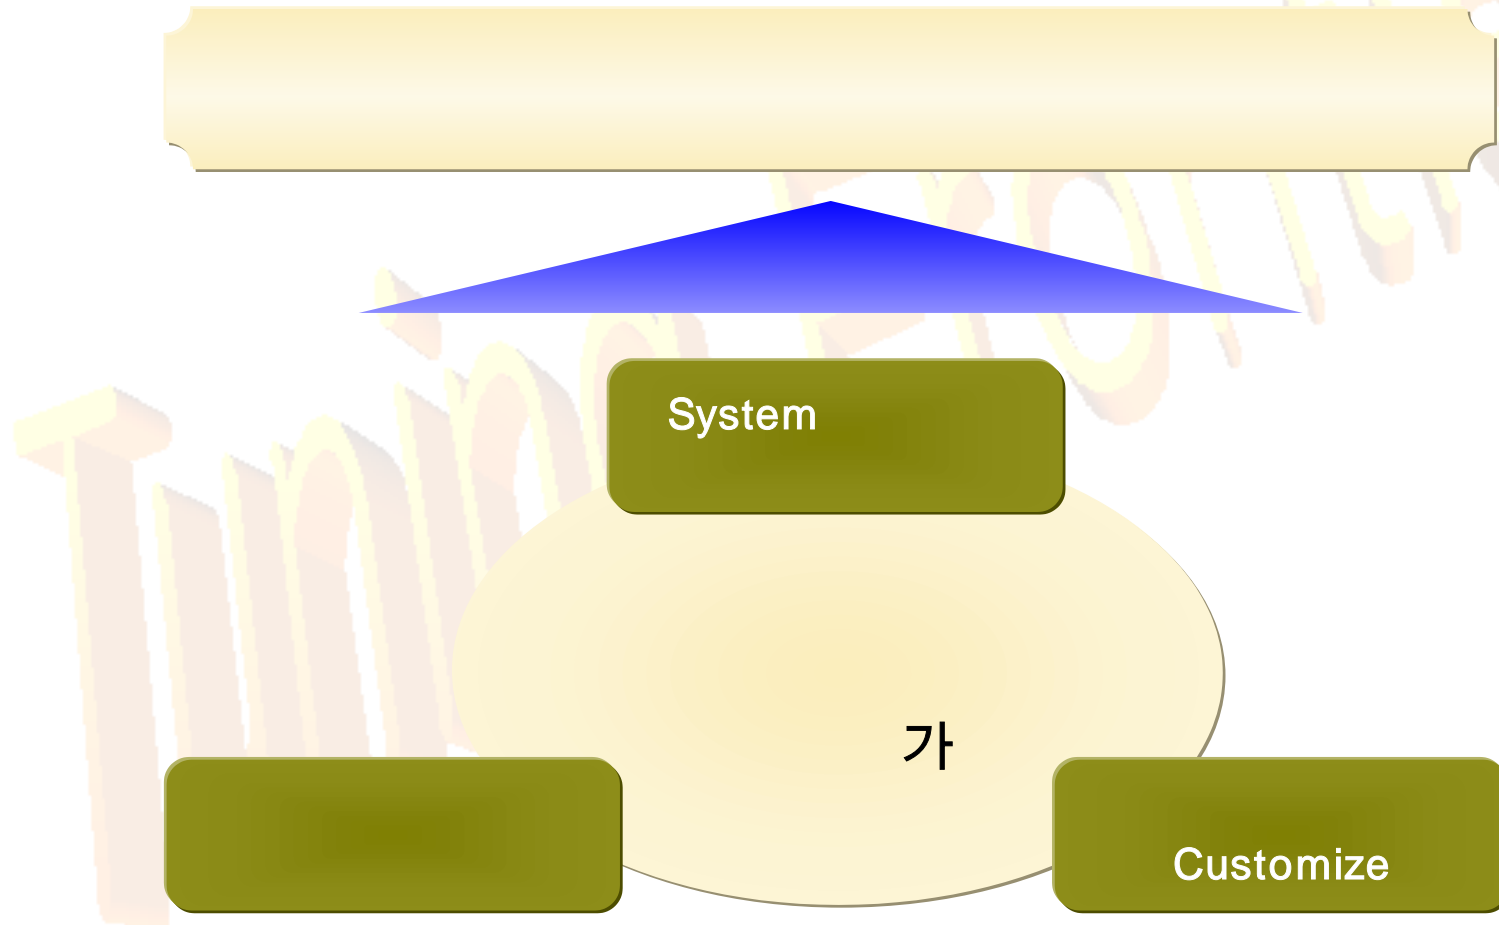


Oracle Performance

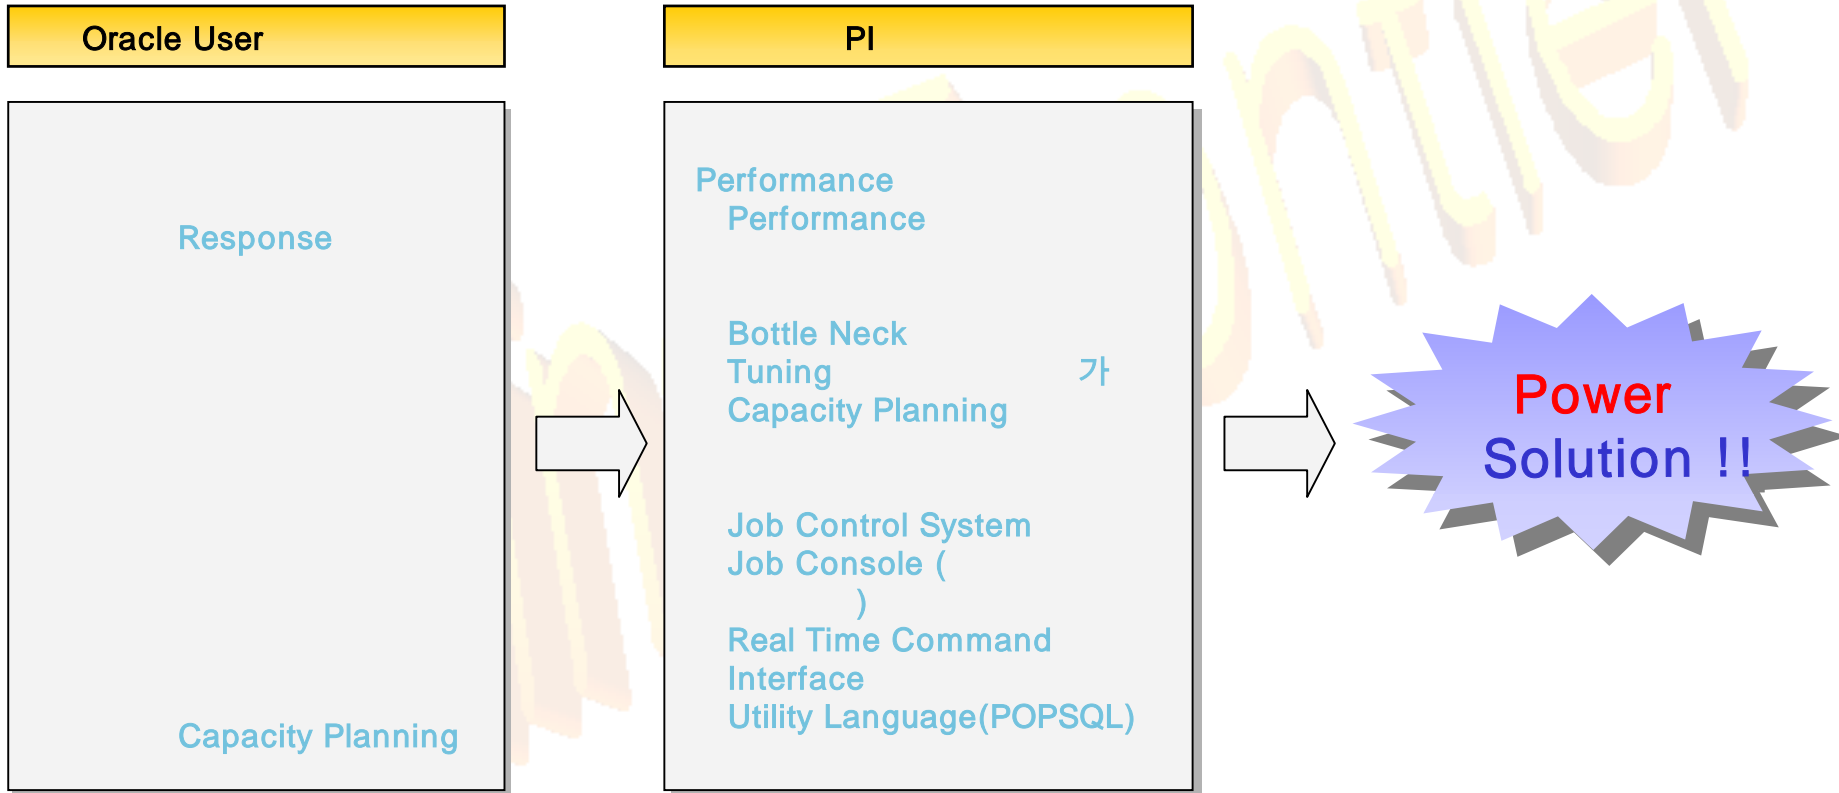
- 1 Intro. Performance Insight



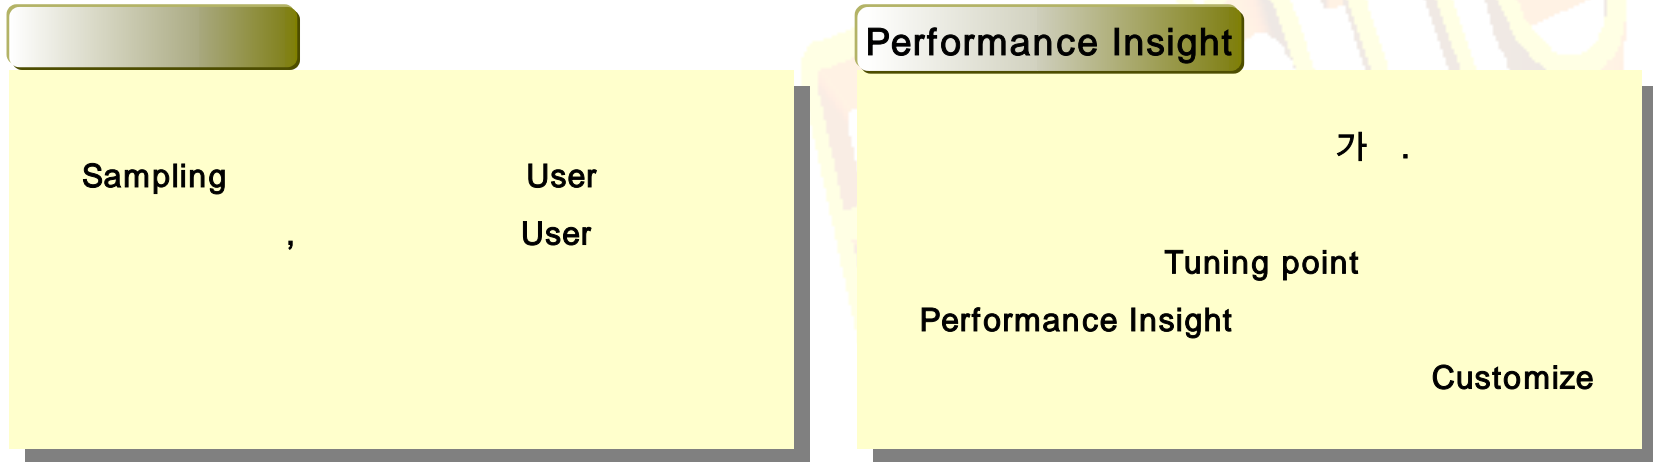
- 2. Performance Insight For Oracle



- 2. Performance Insight For Oracle



- 3. Performance Insight



- 4. Performance Insight

가

- Tuning
- Report

Capacity Planning

- Tuning (, ,Text)
-

Job Control system

- Application

Job Center

- Schedule (/ /)
-

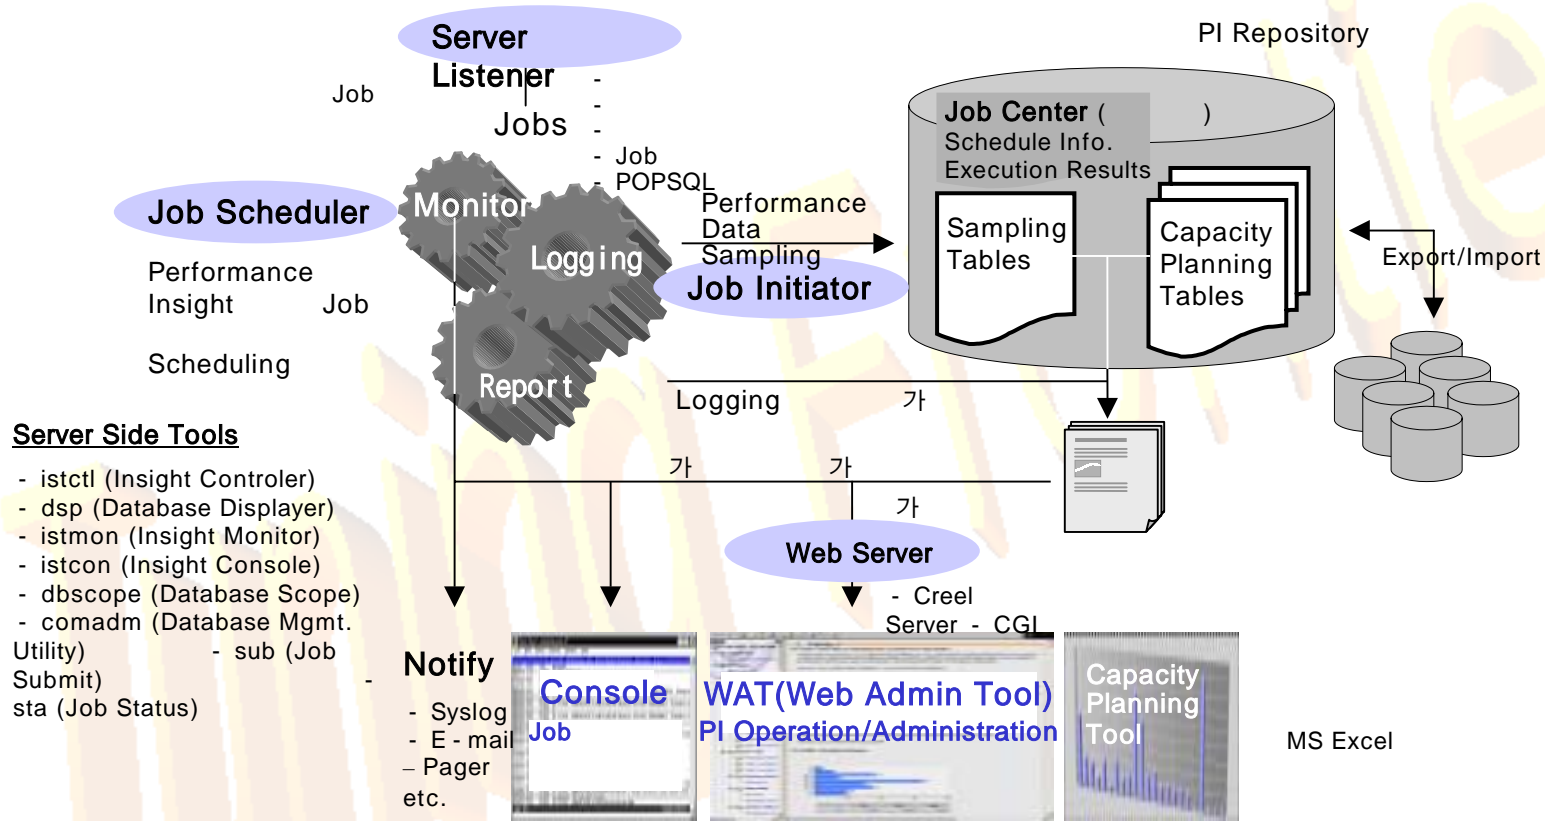
POPSQL

- PI module
- Default Source code
- (customize)

Console

- Network Host
- Message
- Console

- 5. Performance Insight Architecture



- ❑ Oracle Database Server Module : Agent & Manager (60M Byte Disk)
- ❑ Client Module : Excel Bridge (400K byte Disk)
- ❑ Interface : Console,VT(Text Mode), Web Browser(Internet Explore etc., GUI)

. Performance Insight

- 1. 가
- 2. Capacity Planning
- 3. Job Control
- 4. Job Center
- 5. POPSQL
- 6. Console
- 7.

- 1. Evaluation Report(1)

가

“Tuning 가 ”

SQL

Oracle

Bottle neck

Disk
I/O

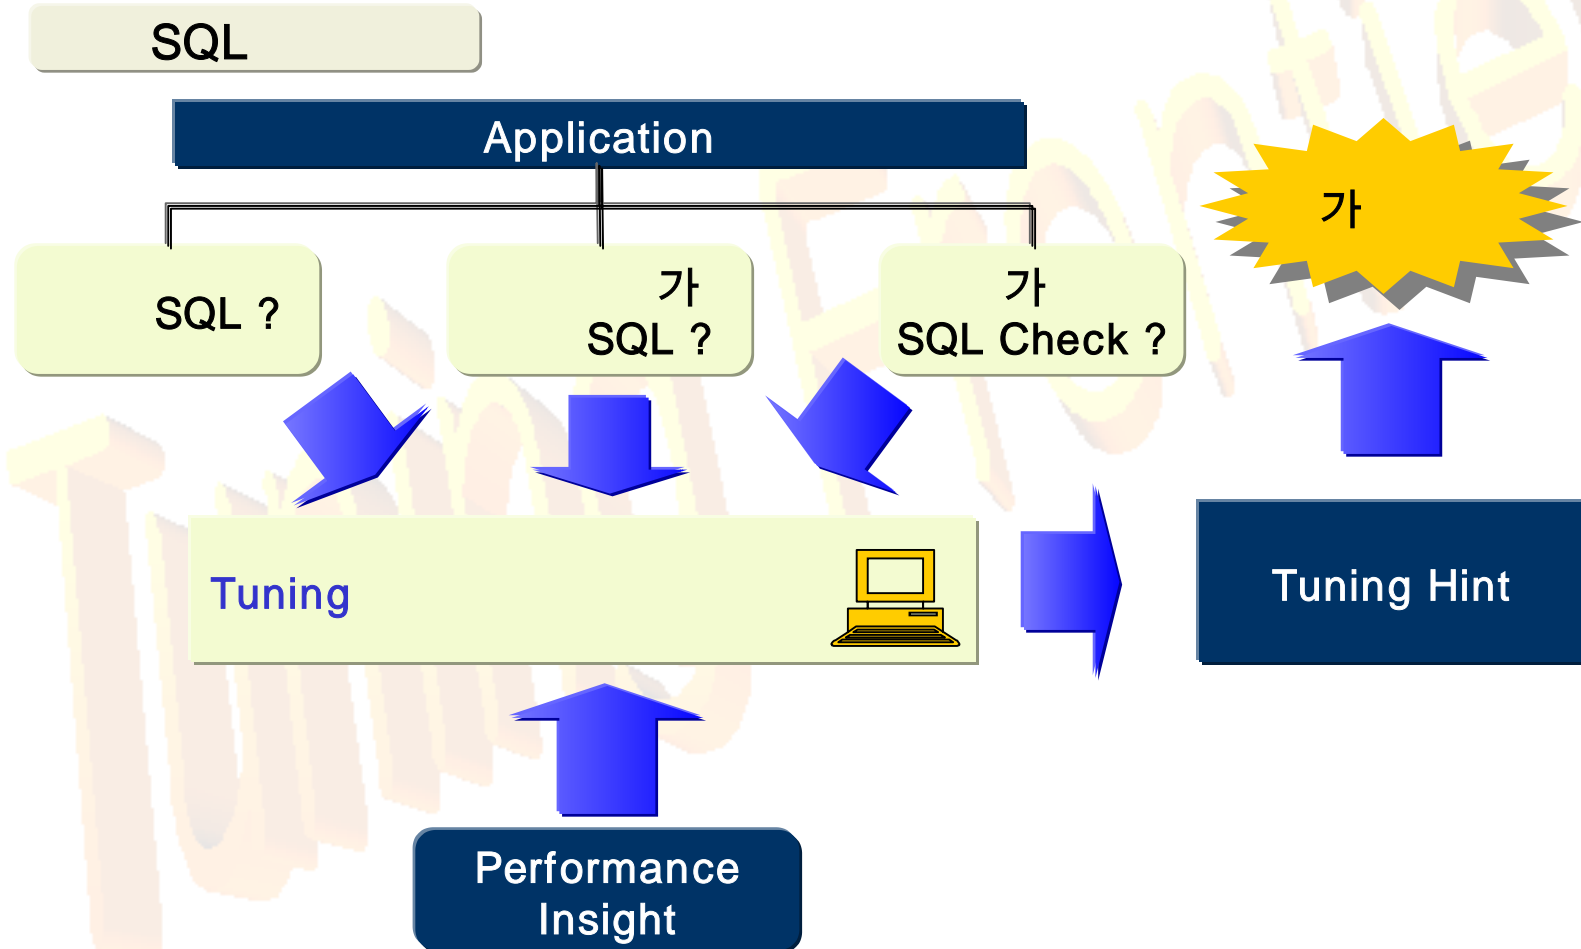
Network
I/O

Database

Parameter

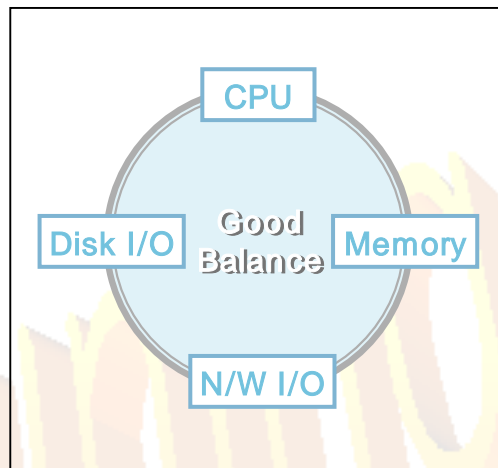
Dead Lock

- 1. Evaluation Report(2)






- 1. Evaluation Report(3)

Bottle Neck



OR



- 1. Evaluation Report(4)

System Analysis Report;

- , Oracle configuration, system resources (CPU, Memory, Disk I/O, Network I/O), SQL and Database design

-Performance (tuning hint)

가,

Ver4

XML

가

- 1. Evaluation Report(5)

The screenshot shows the 'Report Management' section of the Performance Insight Web Admin Tools. It displays a table of hints with columns for Hint File, Priority, Sampled, and Status Summary. The hints are sorted by priority, with Priority 1 (red) at the top and Priority 5 (blue) at the bottom. The status summaries provide details on physical reading I/O increases and error messages.

HINT FILE	PRIORITY	SAMPLED	STATUS SUMMARY
hintora734.010620.4	Priority 1	30 JUN 2001 09:15 (SATURDAY) - 30 JUN 2001 17:56 (SATURDAY)	IST-00015: There is an indication that the physical reading I/O increases by 12.23 per a day
hintora734.010620.316	Priority 1	28 JUN 2001 00:05 (THURSDAY) - 28 JUN 2001 18:32 (THURSDAY)	
hintora734.010620.315	Priority 3	27 JUN 2001 02:57 (WEDNESDAY) - 27 JUN 2001 22:55 (WEDNESDAY)	
hintora734.010620.313	Priority 1	27 JUN 2001 01:53 (WEDNESDAY) - 27 JUN 2001 02:37 (WEDNESDAY)	IST-00015: There is an indication that the physical reading I/O increases by 3.18 per a day
hintora734.010620.312	Priority 1	26 JUN 2001 23:35 (TUESDAY) - 27 JUN 2001 06:14 (WEDNESDAY)	
hintora734.010620.311	Priority 1	26 JUN 2001 19:30 (TUESDAY) - 26 JUN 2001 21:52 (TUESDAY)	IST-11000: An error message is included in ALERT LOG (app/product/7.3.4/rdtest/log/alert_ora734.log)
hintora734.010620.310	Priority 1	26 JUN 2001 16:25 (TUESDAY) - 26 JUN 2001 18:29 (TUESDAY)	
hintora734.010620.309	Priority 2	26 JUN 2001 00:05 (TUESDAY) - 26 JUN 2001 14:15 (TUESDAY)	
hintora734.010620.306	Priority 2	24 JUN 2001 00:05 (SUNDAY) - 24 JUN 2001 22:56 (SUNDAY)	
hintora734.010620.305	Priority 1	23 JUN 2001 00:05 (SATURDAY) - 23 JUN 2001 22:56 (SATURDAY)	IST-00015: There is an indication that the physical reading I/O increases by 14.96 per a day
hintora734.010620.304	Priority 1	22 JUN 2001 14:38 (FRIDAY) - 22 JUN 2001 22:59 (FRIDAY)	
hintora734.010620.303	Priority 1	22 JUN 2001 10:26 (FRIDAY) - 22 JUN 2001 13:44 (FRIDAY)	
hintora734.010620.302	Priority 1	22 JUN 2001 00:05 (FRIDAY) - 22 JUN 2001 18:17 (FRIDAY)	
hintora734.010620.207	Priority 3	21 JUN 2001 00:05 (THURSDAY) - 21 JUN 2001 13:41 (THURSDAY)	
hintora734.010620.1	Priority 1	21 JUN 2001 14:57 (THURSDAY) - 21 JUN 2001 22:59 (THURSDAY)	
hintora734.010620.206	Priority 1	20 JUN 2001 16:05 (WEDNESDAY) - 20 JUN 2001 23:00 (WEDNESDAY)	IST-00015: There is an indication that the physical reading I/O increases by 6.86 per a day
hintora734.010620.205	Priority 1	20 JUN 2001 01:11 (WEDNESDAY) - 20 JUN 2001 15:40 (WEDNESDAY)	IST-00015: There is an indication that the physical reading I/O increases by 9.65 per a day

- 1. Evaluation Report(6)

Performance Insight[®] Web Admin Tools

Report Management Job Ma

Report Management

Select all Reset Selection

Priority 1 Priority 2 Pri

EVENT FILE	PRIORIT
/u01/oradata/ora734/10620.4	
/u01/oradata/ora734/10620.316	
/u01/oradata/ora734/10621.315	
/u01/oradata/ora734/10621.313	
/u01/oradata/ora734/10621.312	
/u01/oradata/ora734/10626.311	
/u01/oradata/ora734/10626.310	
/u01/oradata/ora734/10626.309	
/u01/oradata/ora734/10624.306	
/u01/oradata/ora734/10621.305	
/u01/oradata/ora734/10621.304	
/u01/oradata/ora734/10621.303	
/u01/oradata/ora734/10621.302	
/u01/oradata/ora734/10621.207	
/u01/oradata/ora734/10621.2	
/u01/oradata/ora734/10620.206	
/u01/oradata/ora734/10620.205	

SAMPLED 30 JUN 2001 09:15 (SATURDAY) - 30 JUN 2001 17:56 (SATURDAY)

ora734 (C)
Oracle7 Server Release 7.3.4.0.0 - Production

Summary

- IST-05015 : There is an inclination that the physical reading I/O increases by 12.23 per a day

Priority-2

IST-01015: More than one checkpoint has occurred every sixty minutes

A REDO log is too small, or LOG_CHECKPOINT_INTERVAL or LOG_CHECKPOINT_TIMEOUT(INIT.ORA) is too small. At a checkpoint, all the contents of database buffer would be written to a disk, which might keep a user waiting for processing. Performance can be enhanced by decreasing the number of checkpoint executions.

Keyword: LOG_CHECKPOINT_TIMEOUT, LOG_CHECKPOINT_INTERVAL

IST-01003: Checkpoint

A checkpoint occurs when one of the REDO logs is full, which results in the writing of the information in the database (on database buffer) to data file (on the disk) by DBWR. The timing of this checkpoint occurrence can be controlled by LOG_CHECKPOINT_INTERVAL or LOG_CHECKPOINT_TIMEOUT of initial setup parameters. The reason why these parameters are set up is to control recovery time and to "shorten recovery time by occurring as many checkpoints as possible." However, from the viewpoint of tuning, you should "make as few checkpoints as possible," because they slow down database performance.

Speed of a disk drive is very important for checkpoint processing. Low disk speed results in slow recovery from REDO logs and affects the write to REDO logs by LGMR. Low disk speed causes many checkpoints. Delay caused by too many checkpoints makes the DBWR's delay write function (an Oracle function to enhance performance) meaningless.

Keyword: LOG_CHECKPOINT_TIMEOUT, LOG_CHECKPOINT_INTERVAL

IST-01020: LOG_CHECKPOINTS_TO_ALERT

In order to study the interval of CHECKPOINT, set up LOG_CHECKPOINTS_TO_ALERT in TRUE.

Keyword: LOG_CHECKPOINTS_TO_ALERT

IST-01008: Checkpoint Histogram

- 1. Evaluation Report(7)

XML

- (tuning hint)
 -

The screenshot displays an evaluation report with the following sections:

- Priority-2 (IST-05002):** Rate of long_table_full_scans. The search which does not use the index (full scan) affects performance. Consider preparing an appropriate index. $(\text{long_table_full_scans} / (\text{long_table_full_scans} + \text{short_table_full_scans})) * 100 = 6.62\%$. Below this is a histogram titled "Number of Full Scans Histogram(long tables)" with a table showing counts for different ranges.
- Priority-3 (IST-11006):** Archive area write busy. A REDO log, even after the check point, does not complete writing to archive area. It is possible that disk speed in archive area or the speed of a disk where a REDO log is saved is very slow. If this happens often and configuration of disk cannot be changed, consider to add another REDO log file.
- Priority-3 (IST-03002):** Processing volume of network doesn't correlate with database traffic. Therefore an analysis should be carried out excluding network load.
- Priority-3 (IST-06107):** Today's consumption in tablespace has exceeded free space.

Annotations in the image include:

- A blue box labeled "XML" pointing to the left sidebar.
- A blue box containing "- (tuning hint)" and "-" pointing to the top right of the Priority-2 section.
- A blue box pointing to the histogram table in the Priority-2 section.
- A blue box pointing to the bottom left of the report.
- A blue box pointing to the bottom left of the report.

Bin	Count
0-70	192
70-140	221
140-210	247

- 1. Evaluation Report(8)

Priority-2

IST-01015: More than one checkpoint has occurred every sixty minutes
 A REDO log is too small, or LOG_CHECKPOINT_INTERVAL or LOG_CHECKPOINT_TIMEOUT(INIT.ORA) is too small. At a checkpoint, all the contents of database buffer would be written to a disk, which might keep a user waiting for processing. Performance can be enhanced by decreasing the number of checkpoint executions.
 Keyword: LOG_CHECKPOINT_TIMEOUT , LOG_CHECKPOINT_INTERVAL

IST-01003: Checkpoint
 A checkpoint occurs when one of the REDO logs is full, which results in the writing of the information in the database (on database buffer) to data file (on the disk) by DBWR. The timing of this checkpoint occurrence can be controlled by LOG_CHECKPOINT_INTERVAL or LOG_CHECKPOINT_TIMEOUT of initial setup parameters. The reason why these parameters are set up is to control recovery time and to "shorten recovery time by occurring as many checkpoints as possible." However, from the viewpoint of tuning, you should "make as few checkpoints as possible," because they slow down database performance. Speed of a disk drive is very important for checkpoint processing. Low disk speed results in slow recovery from REDO logs and affects the write to REDO logs by LGWR. Low disk speed causes many checkpoints. Delay caused by too many checkpoints makes the DBWR's delay write function (an Oracle function to enhance performance) meaningless.
 Keyword: LOG_CHECKPOINT_TIMEOUT , LOG_CHECKPOINT_INTERVAL

IST-01020: LOG_CHECKPOINTS_TO_ALERT
 In order to study the interval of CHECKPOINT, set up LOG_CHECKPOINTS_TO_ALERT in TRUE.
 Keyword: LOG_CHECKPOINTS_TO_ALERT

IST-01008: Checkpoint Histogram

Checkpoint Histogram

Hour	Count
09	3
10	8
11	27
12	10
13	48
14	22
15	26
16	15
17	13

	Min	Max	Avg
	3 (09)	48 (13)	19.11

- 1. Evaluation Report(9)

The screenshot displays two windows from the Performance Insight tool. The top window, titled 'Priority-1', shows a report for 'Tablespace Free Status' with a table of free blocks and bytes. The bottom window, titled 'WAT DB Scope', shows a 'Tablespace Storage Allocation' report with a table of allocated, used, and free bytes for various tablespaces. A red box highlights a button in the top window, and a red arrow points from it to the 'Tablespace Storage Allocation' report in the bottom window.

가 REPORT TOOL

TABLESPACE_NAME	FREE_BLOCKS	FREE_BYTES	EXTENDED_KB
USERS	101,559	200,118	222
XPRT	14,667	29,334	5,294

TABLESPACE NAME	BYTES ALLOCATED(KB)	BYTES USED(KB)	BYTES FREE(KB)	PERCENT USED(%)	PERCENT FREE(%)
RBS	30,720	27,896	2,824	91	9
SYSTEM	343,040	72,796	270,244	21	79
TEMP	20,480	20,480	0	100	0
TOOLS	10,240	2	10,238	0	100
USERS	204,800	1,712	203,088	1	99
XPRT	67,440	34,236	33,204	51	49

- 1. Evaluation Report(10)

Priority-1
 IST-06103: Tablespace Free Status
 An empty situation with today's final sample in each tablespace is as follows:


TABLESPACE	FREE	USED	TOTAL
RBS	10,239	20,478	3,082
SYSTEM	5,119	10,238	0
TEMP	101,559	200,118	222
TOOLS			14,667
USERS			

가 REPORT TOOL

Capacity Planning Tool

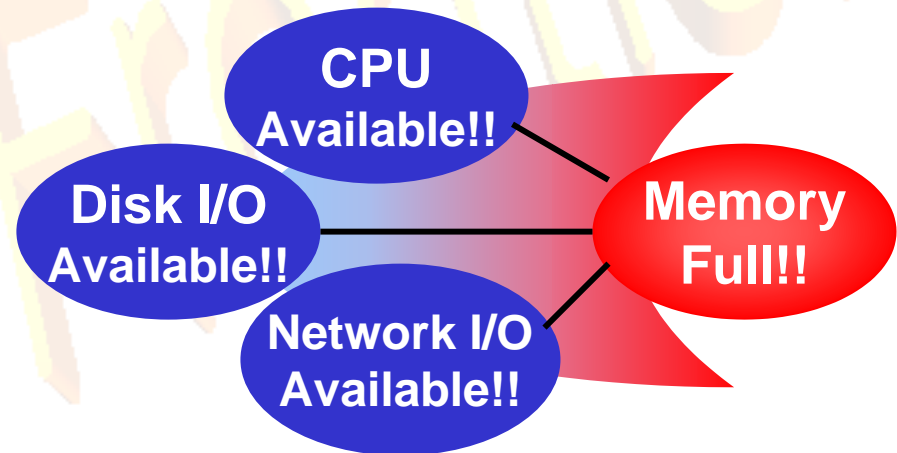
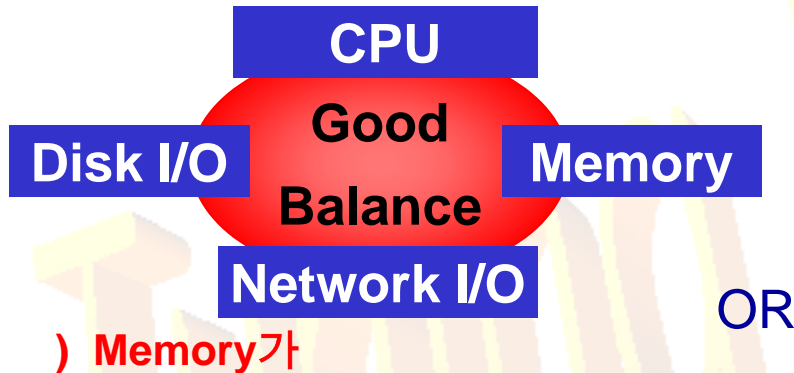
CP GROUP

- TOTAL CPU
- TOTAL SPOOL
- TOTAL MEMORY
- TOTAL NETWORK PACKET
- TOTAL INPUT NETWORK PACKET
- TOTAL OUTPUT NETWORK PACKET
- TOTAL NETWORK COLLISION
- MAX CONNECTIONS GENERATED - DATABASE
- MAX CONNECTIONS GRANTED - DATABASE
- TOTAL SQL EXECUTIONS
- TOTAL CHECKPOINT COUNT
- AVERAGE CPU USAGE
- DIFFERENCE ANALYSIS
- RESPONSE ANALYSIS
- TABLESPACE CONSUMPTION
- OTHER METRICS



- 1. Evaluation Report(11)

Bottleneck analysis (4)



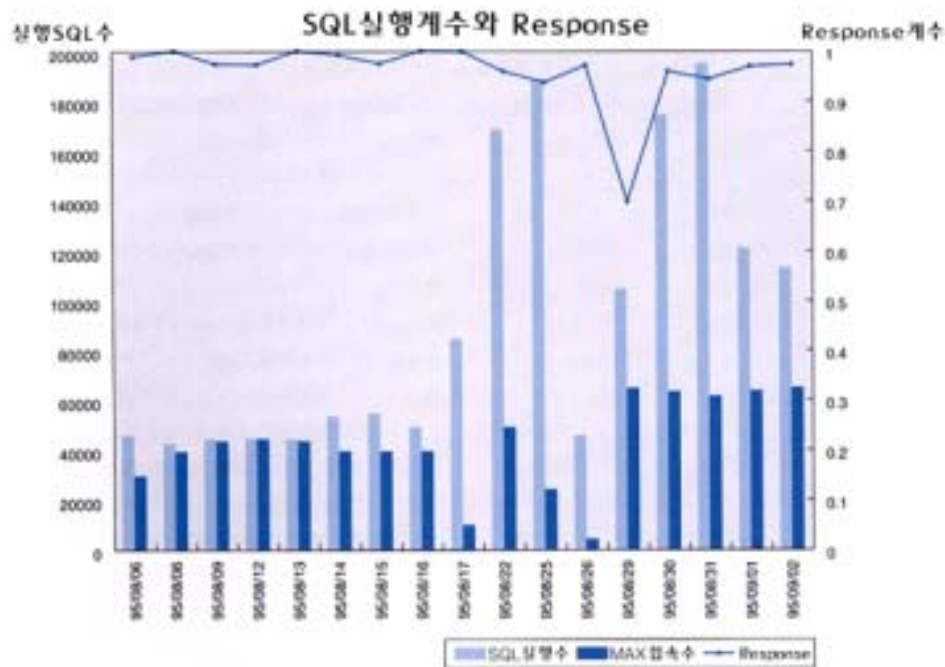
IST-02002: System overhead by paging increases
Paging influences response more than normal condition.

IST-02005: Use multi-thread server connections
It's possible to reduce paging by using multi-thread server(shared server) which is suitable for many users.

IST-02007: Reduce DB_BLOCK_BUFFERS
The hit rate of database buffer in peak time is maintained more than 90 %. Taking away paging overhead by reducing DB_BLOCK_BUFFERS(INIT.ORA) takes priority.

- 2. Capacity Planning (1)

Capacity Planning ?

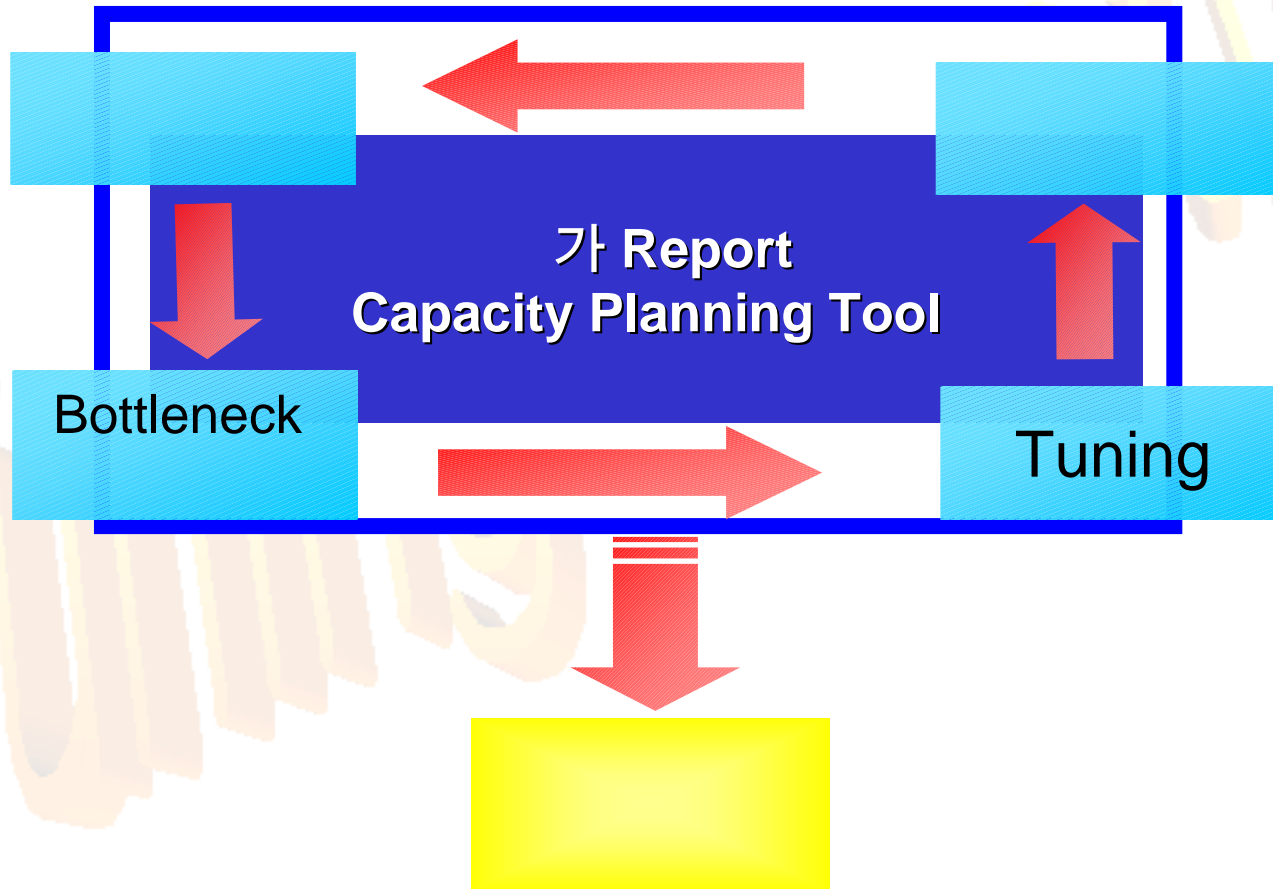


- 가 가
- 가

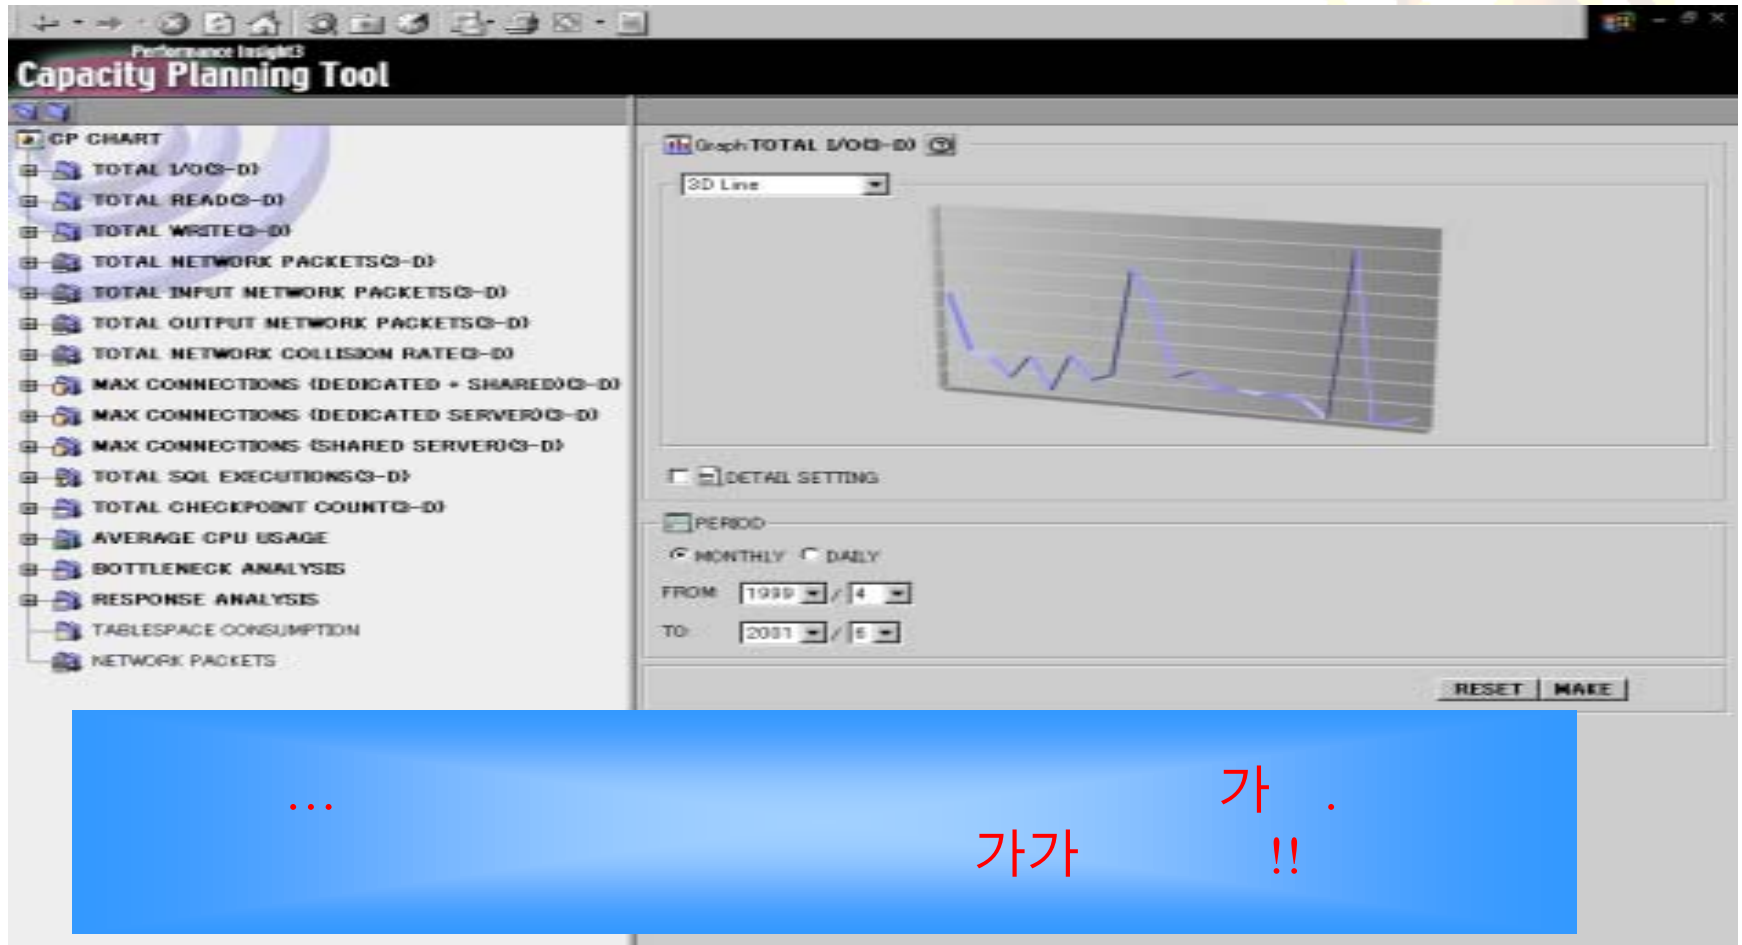


- 2. Capacity Planning (2)


- “ ” 가
- ,



- 2. Capacity Planning (3)




- 2. Capacity Planning (4)



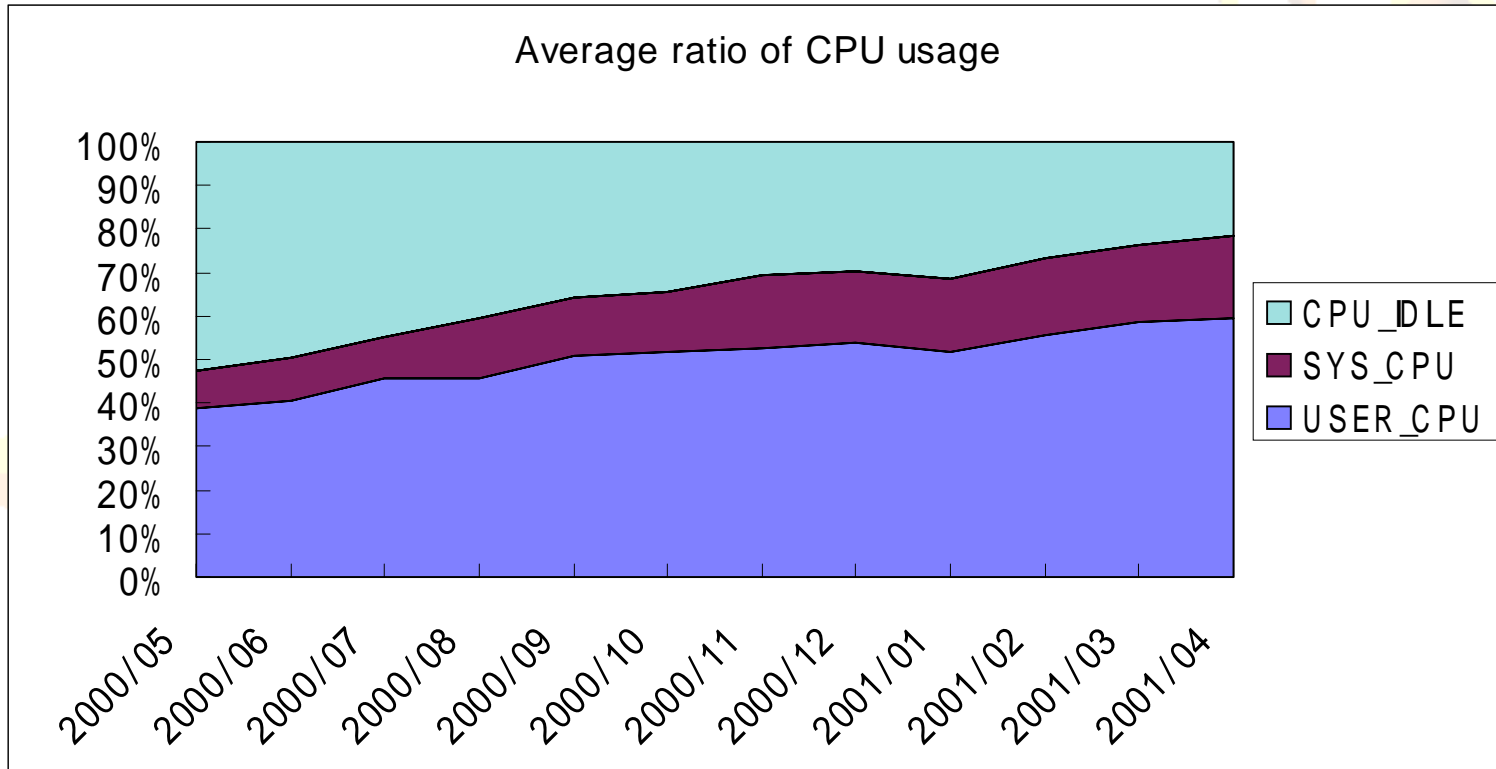
(가)

- 2. Capacity Planning (5)



가 !!

- 2. Capacity Planning (6)



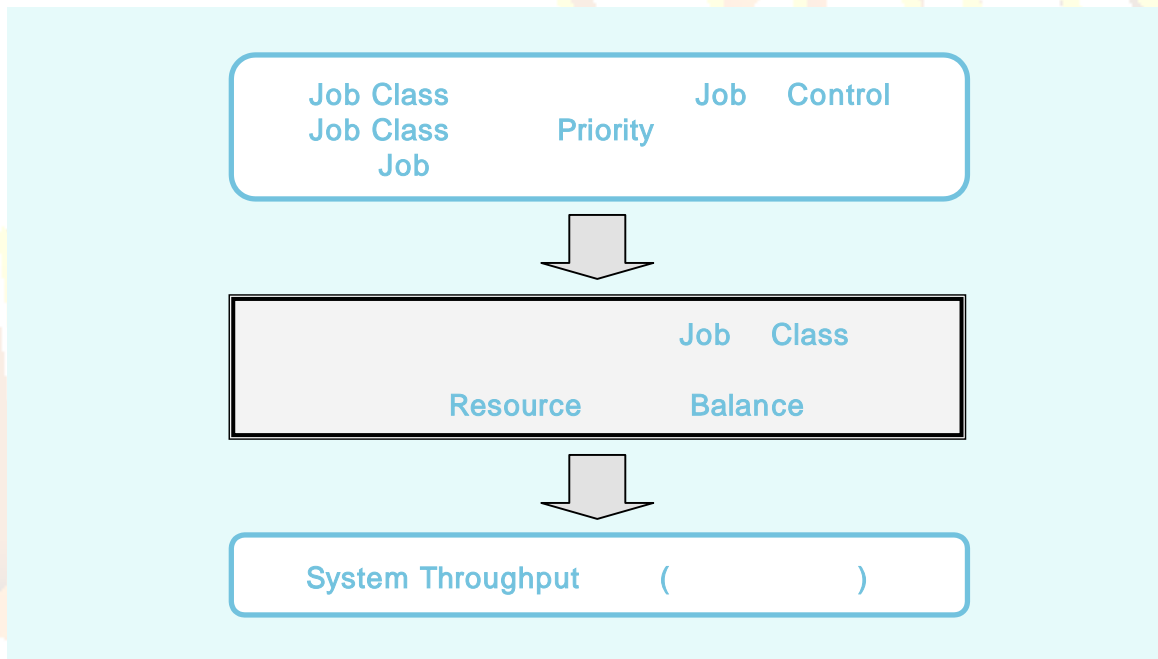
User CPU

System CPU

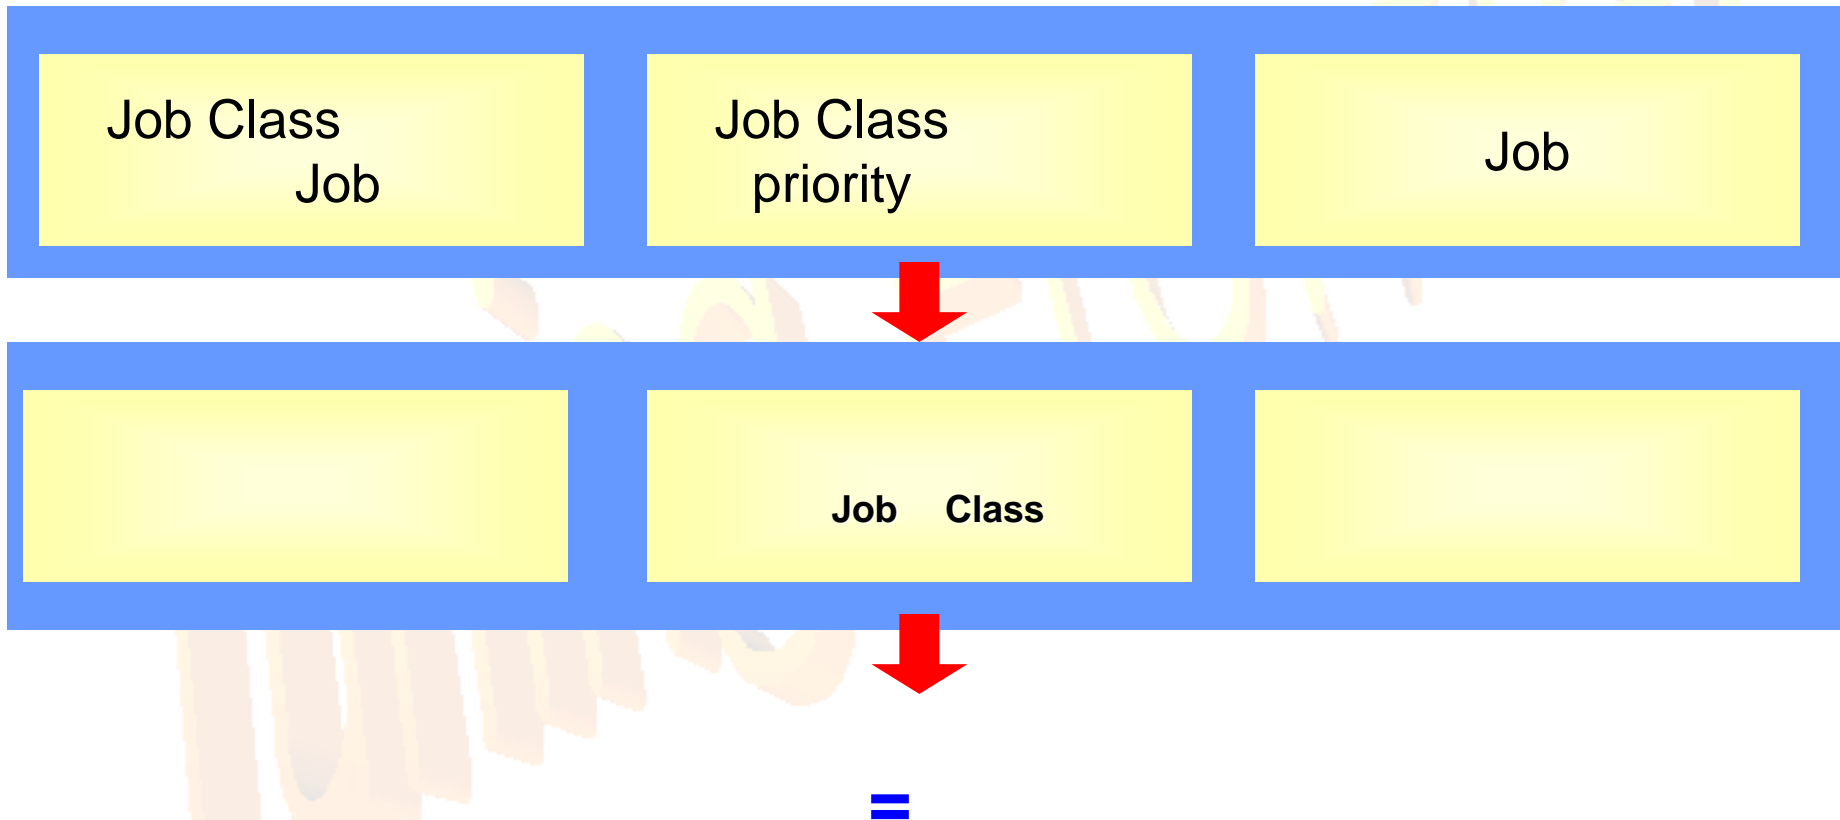
I/O

, IDLE

- 3. Job Control (1)

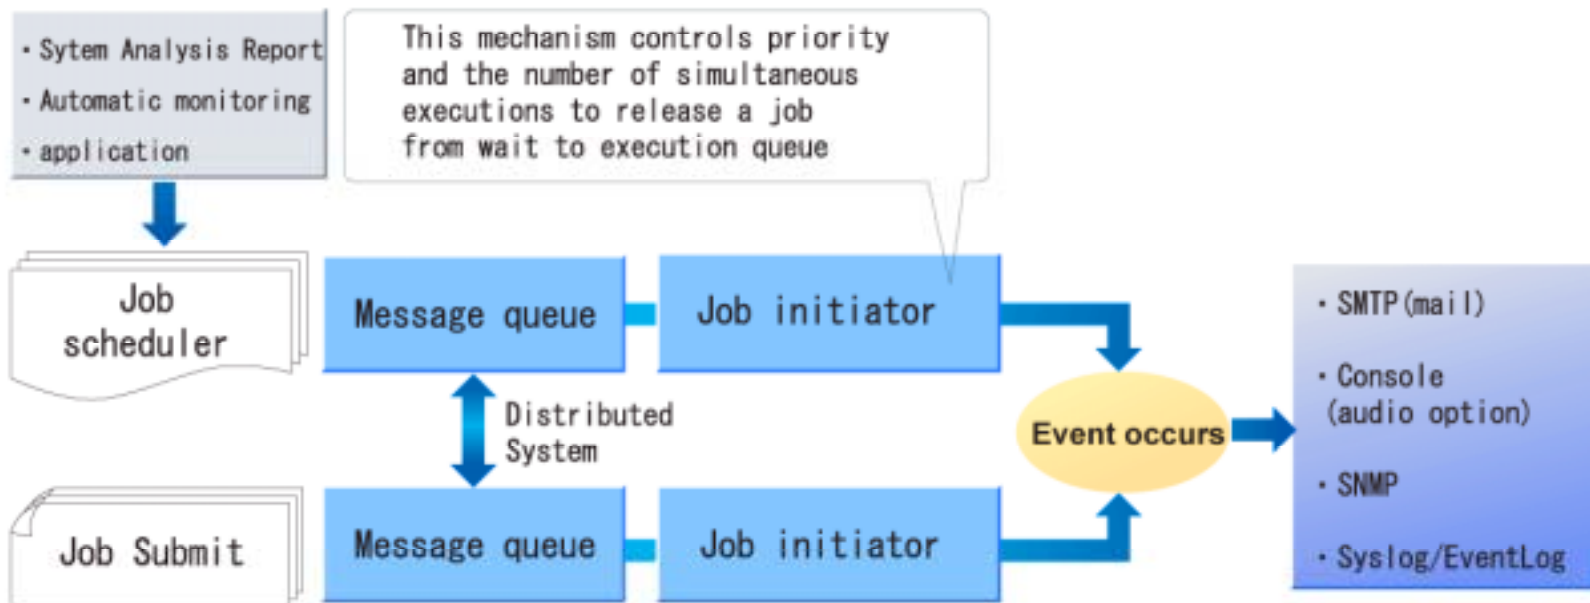


- 3. Job Control (2)



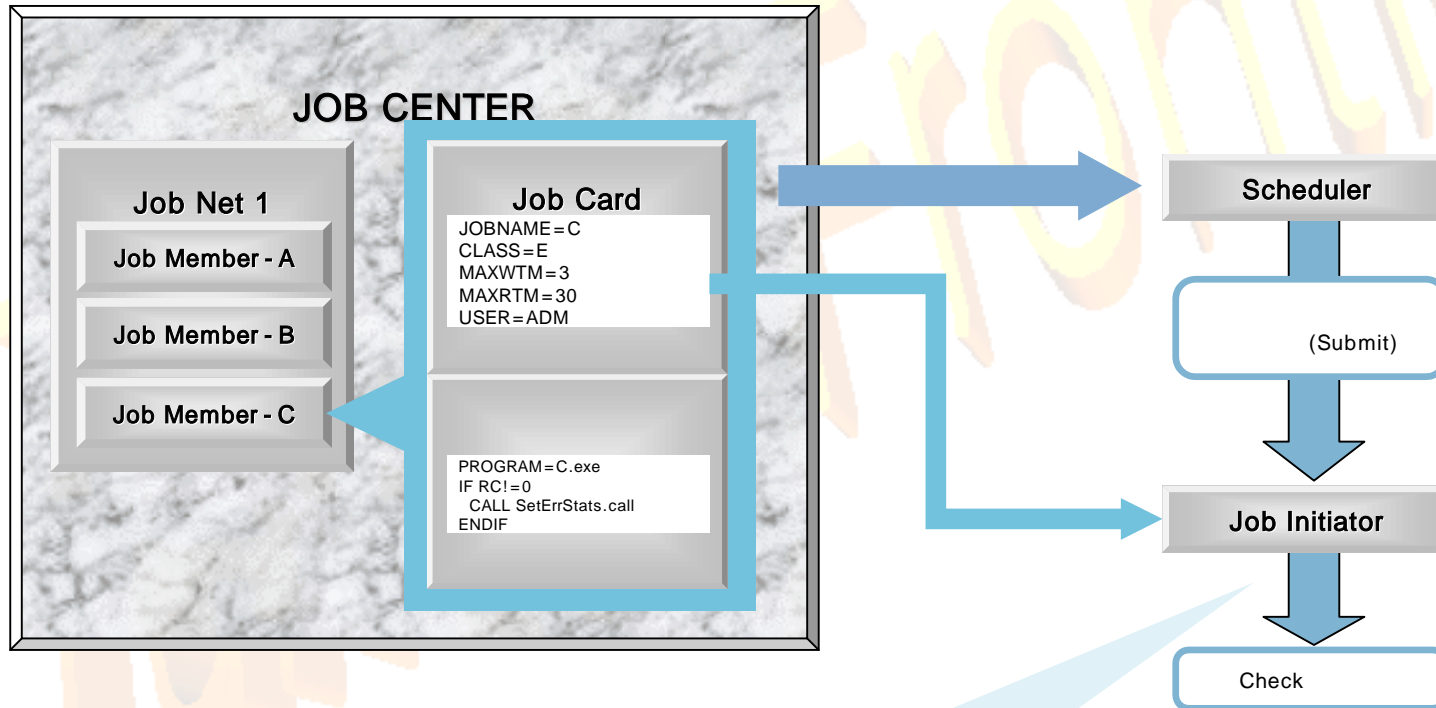
- 3. Job Control (3)

Message queuing system



- 4. Job Center (1)

Job Center



E Class Priority ?
 E Class Job 가 ?

- 4. Job Center (2)

Job Control

Schedule

Job
(Daily, weekly, monthly)
Job



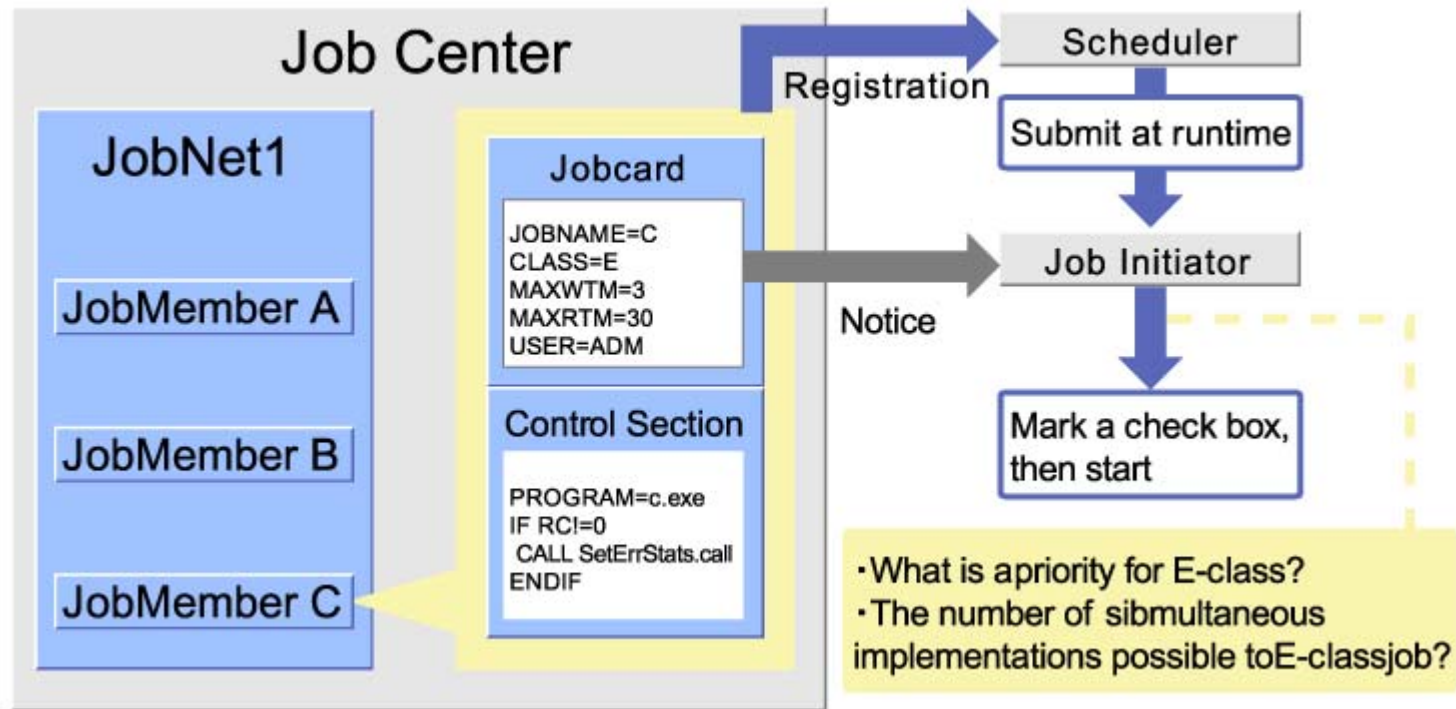
Calendar

-
- Calendar
(: ,)

가
가

가

- 4. Job Center (3)



- 5. POPSQL(SQEEL) (1)

POPSQL



SQeeL is new utility language of Performance Insight, provided as open source.

Database SQL Interface
Sequential File SQL
Interface
Report
 Console Monitor
 Interface
Network File , ,
Job File
Database Load, Update
Load, Unload
Pocket Bell Call
 Application
XA Interface

- 5. POPSQL(SQEEL) (2)

SQeel

- GUI
 - HTML
 - GUI
 - Demand Compile
 - 가
- Web
 SQeel
 , CGI
 , SQeel
 Compile
 "eeIVM"



GUI (Web)

가

,

- 5. POPSQL(SQEEL) (3)

Example of SQeeL

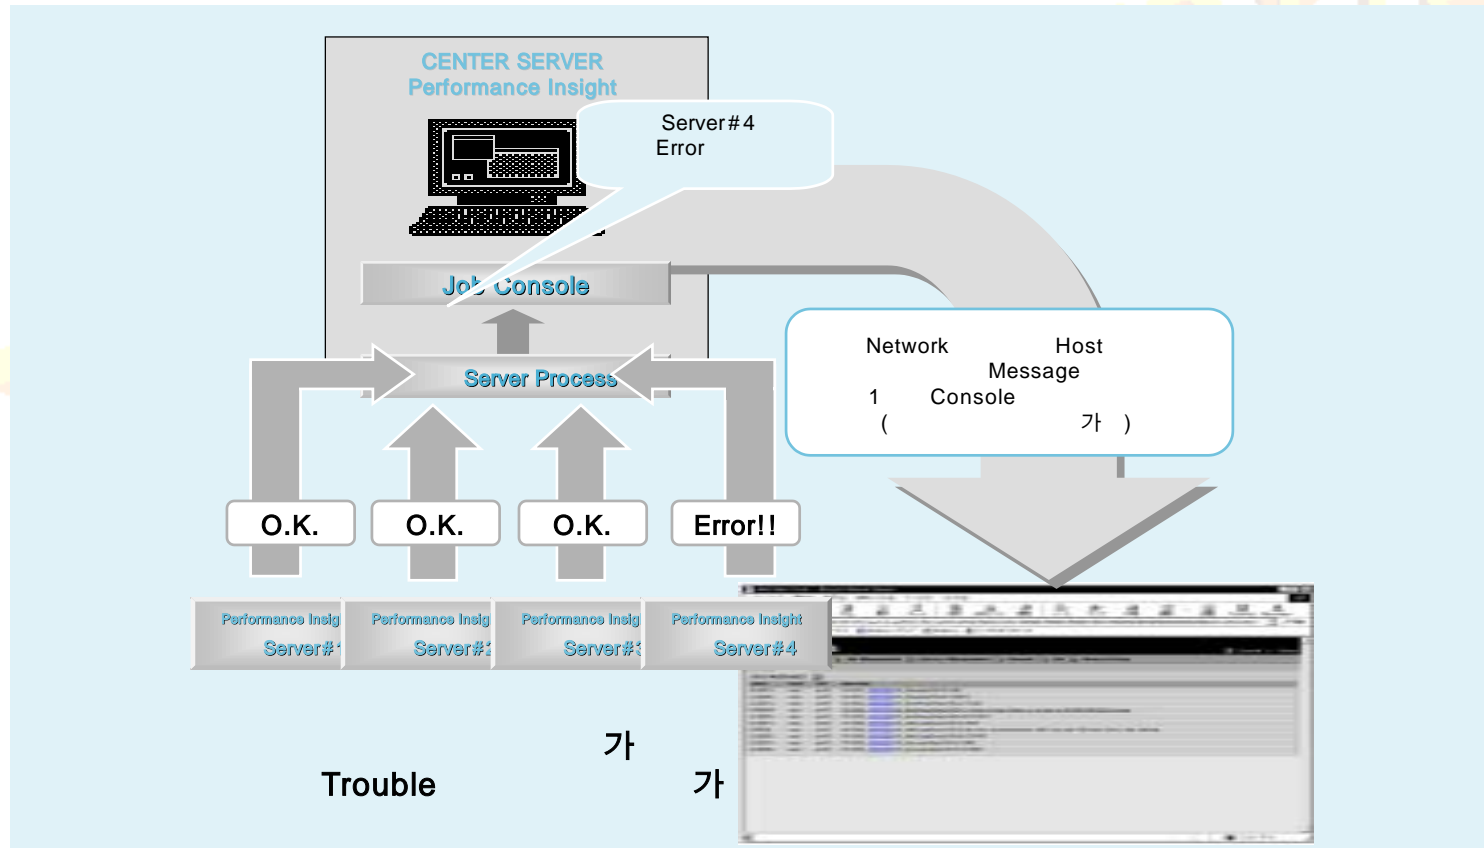


```
<eel" arguments resultVal ">
<eel" global i, j, col_name ">

<eel" col_name=keys(resultVal[0]) ">
<HTML>
<HEAD>
<TITLE>SQeeL demo(result)</TITLE>
</HEAD>
<BODY>
<eel" if count(resultVal) < 0 ">
    ERROR = ${resultVal}
<eel" endif ">
<TABLE border="1">
<eel" if count(resultVal) > 0 ">
    <TR>
<eel" loop i=0; i<count(col_name); i++ ">
        <TH>
                                ${col_name[i]}
        </TH>
<eel" endloop ">
    </TR>
<eel" loop j=0; j<count(resultVal); j++ ">
    <TR>
<eel" loop i=0; i<count(col_name); i++ ">
        <TD>
                                ${resultVal[j][col_name[i]]}
    </TD>
    </TR>
</eel>
</TABLE>
</BODY>
</HTML>
```

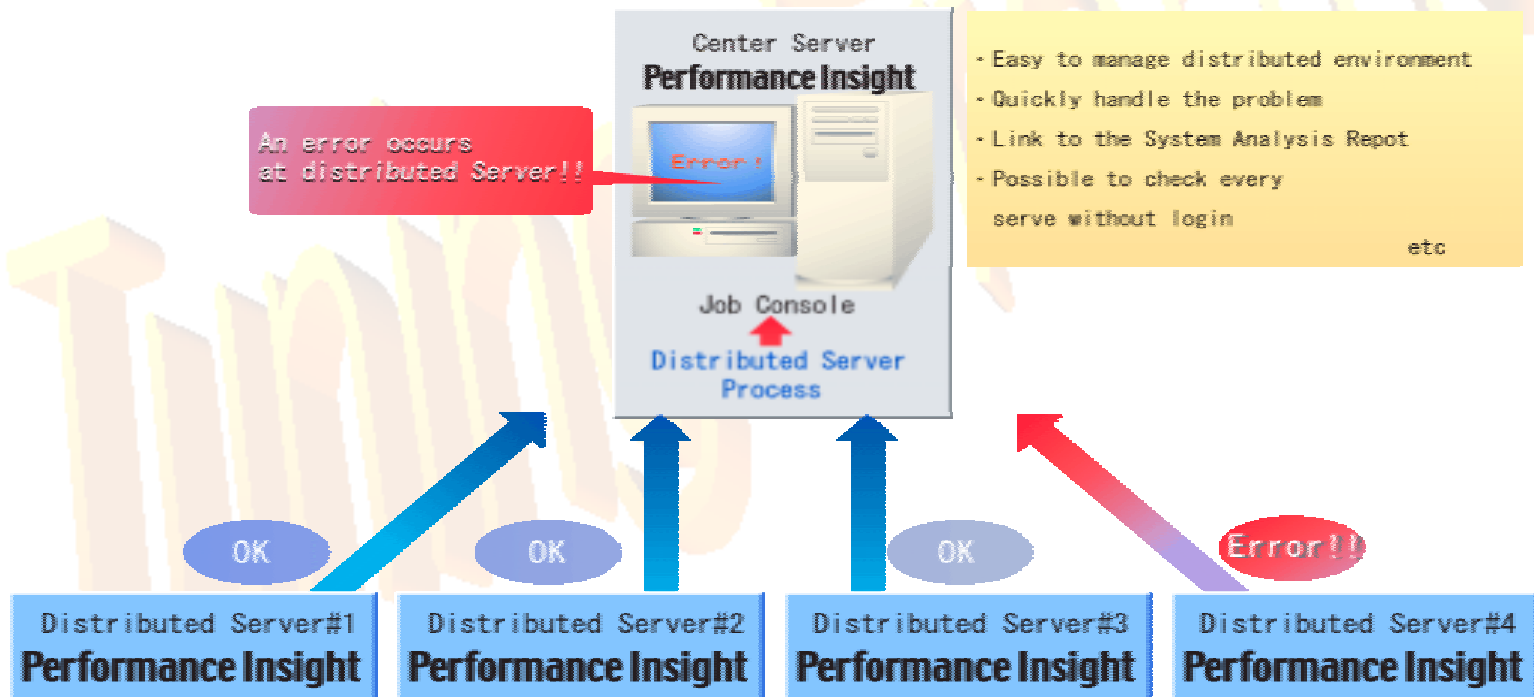
- 6. Console (1)

Console



- 6. Console (2)

Console



- 7. DB Scope (1)

SQL SCOPE

TABLE SCOPE

INDEX SCOPE

TABLESPACE SCOPE

- 7. DB Scope (2) – SQL Scope

SQL 3가 (가)

1. SQL
2. SQL
(가 HEAVY SQL SQL)
- 3.



--- SQL

--- SQL

--- SQL

- 7. DB Scope (3) – SQL Scope

The screenshot displays the Performance Insight WAT DB Scope interface. On the left, a tree view shows the hierarchy: DB Scope > Table > SQL. The main area is divided into two panes. The top pane, titled 'Execution Plan', shows the SQL statement: `select * from master2 where to_char(empno,999999) = '000000'`. A red circle highlights the `where` clause, with an arrow pointing to a light blue box containing the text 'WHERE'. The bottom pane, titled 'Execution Plan Tree', shows the execution plan details: `(0) SELECT STATEMENT`, `(0) SELECT STATEMENT ()`, `Opt_Mode=CBOOSE`, and `(1) TABLE ACCESS (FULL) MASTER2`. A red circle highlights the `(1) TABLE ACCESS (FULL) MASTER2` entry, with an arrow pointing to a light blue box containing the text 'Full Scan'. At the bottom right, a blue cloud contains the text 'Before SQL tuning'. The interface includes a 'Save to History' button and a 'GO' button at the top right.

- 7. DB Scope (4) – SQL Scope

The screenshot shows the Performance Insight WAT DB Scope interface. The left sidebar displays a tree view of the database structure, including Tablespaces, Tables, Schemas, Indexes, and SQL statements. The main window shows the 'Execute and Measure' results for a specific SQL statement. The SQL statement is: `select * from master2 where to_char(empno,999999) = '00000'`. The measurement results are as follows:

Metric	Value
EXECUTION TIME (secs)	0.360000
USER CPU TIME (secs)	0.190000
SYSTEM CPU TIME (secs)	0.030000
OPENED CURSORS CUMULATIVE	0
USER CALLS	4
RECURSIVE CALLS	0
LOGICAL READS	2866
STORED PROCEDURE SPACE	0
CPU USAGE	0
UGA MEMORY (bytes)	0
PGA MEMORY (bytes)	1456
DB BLOCKS GETS	2
CONSISTENT GETS(block)	2864
PHYSICAL READS(block)	2865
PHYSICAL WRITES(block)	0
REDO SIZE(bytes)	0
TABLE SCANS(about table)	0

A blue thought bubble contains the text "Before SQL tuning". A red circle highlights the "PHYSICAL READS(block)" row in the measurement results table, with an arrow pointing to a callout box that says "Physical read = 2865 blocks".

- 7. DB Scope (5) – SQL Scope

Performance Insight[®]
WAT DB Scope

GO

DB Scope

- Tablespace
- Table
- Tablespace
- Schema
- Index
- Tablespace
- Schema
- SQL
- From Running
- From Logfile
- From Text file
- From History
 - SQL_001
 - SQL_002
 - SQL_003
 - SQL_004
 - SQL_005
 - SQL_006
 - SQL_007

Summary | Display Execution Plan | Execute and Measure | Reference Object List | SQL Statement Edit | History

Reference Object List

where clause

```
to_char(empno, '999999') = '33333'
```

Reflected Objects Tree

Table Name(Owner)	Column Name	Index Name	Used?	NULL?	Data Type
MASTER2(SCOTT) (TABLE)	empno			NO	NUMBER
		MASTER2_IDX(I)			(INDEX)
	ename				VARCHAR2(10)
	job				CHAR(9)
	mgr				NUMBER(4)
	hiredate				DATE
	sal				NUMBER(7,2)
	comm				NUMBER(7,2)
	deptno				NUMBER(2)

Save to History

Tablespace | Table | Index | SQL

Before SQL tuning

Index not used

- 7. DB Scope (6) – SQL Scope

The screenshot displays the Performance Insight WAT DB Scope interface. The left sidebar shows a tree view with 'SQL' selected. The main area shows the 'Execution Plan' tab with the following SQL statement:

```
select * from master2 where empno = 20000
```

The Execution Plan Tree below shows the following details:

```
(1)SELECT STATEMENT  
(2)SELECT STATEMENT O  
Opt_Mode=CHOICE  
(1)--- TABLE ACCESS (BY ROWID) MASTER2  
(2)--- INDEX (RANGE SCAN) MASTER_IDX
```

Annotations include a red arrow pointing to the 'Edit SQL' button, a red arrow pointing to the 'Index scan' text in the execution plan, and a blue cloud containing the text 'After SQL tuning'.

- 7. DB Scope (7) – SQL Scope

The screenshot shows the Performance Insight WAT DB Scope interface. The left sidebar displays a tree view with 'SQL' selected. The main window shows the 'Execute and Measure' tab with the following SQL statement:

```
select * from master2 where empno = 28322
```

The 'Measurement result' table is displayed below:

EXECUTION TIME (secs)	0.250000
USER CPU TIME (secs)	0.170000
SYSTEM CPU TIME (secs)	0.050000
OPENED CURSORS CUMULATIVE	0
USER CALLS	4
RECURSIVE CALLS	0
LOGICAL READS	3
STORED PROCEDURE SPACE	0
CPU USAGE	0
UGA MEMORY (bytes)	0
PGA MEMORY (bytes)	0
DB BLOCKS GETS	0
CONSISTENT GETS(block)	3
PHYSICAL READS(block)	3
PHYSICAL WRITES(block)	0
REDO SIZE(byte)	0
TABLE SCANS(rowset tables)	0

A blue thought bubble contains the text "After SQL tuning". A red oval highlights the 'PHYSICAL READS(block)' row in the table, with an arrow pointing to a text box that reads "Physical read = 4 blocks (2865 blocks before tuning)".

- 7. DB Scope (8) – SQL Scope

Performance Insight[®]
WAT DB Scope

GO LOOON

DB Scope

- Tablespace
- Table
- Tablespace
- Schema
- Index
- Tablespace
- Schema
- SQL
- From Running
- From Logfile
- From Text file
- From History
- SQL_001
- SQL_002
- SQL_003
- SQL_004
- SQL_005
- SQL_006
- SQL_007

Summary | Display Execution Plan | Execute and Measure | Reference Object List | SQL Statement Edit | History

Reference Object List

where clause

```
empno = 30000
```

Reflected Objects Tree

Table Name(Owner)	Column Name	Index Name	Used?	NULL?	Data Type
MASTER2(SCOTT) (TABLE)	empno		NO		NUMBER
	ename	* MASTER2_ID00(1)	Used		(INDEX)
	job				VARCHAR2(10)
	mgr				VARCHAR2(9)
	hiredate				NUMBER(4)
	sal				DATE
	comm				NUMBER(7,2)
	deptno				NUMBER(7,2)
					NUMBER(2)

Save to History

Tablespace | Table | Index | SQL

After SQL tuning

Index in use

- 7. DB Scope (9) – TABLE Scope

1. STORAGE (PCTUSED, PCTFREE .)
2. EXTENT , BLOCK , BYTE
3. BLOCK / BLOCK
4. Chained row / Migrated row
Migrated row가 가 ,
가

- 7. DB Scope (10) – TABLE Scope

The screenshot displays the 'WAT DB Scope' interface for the SCOTT schema. The 'Row Migration' tab is active, showing a table with columns: SCHEMA, TABLE NAME, ROW MIGRATION, ROW MIGRATION SIZE, TABLESPACE NAME, TBS FOR TEMP WORK, and TEMP TABLE NAME. The 'CHAIN_TEST2' table is highlighted, showing 8 rows migrated to the USERS tablespace. A red circle highlights the 'ROW MIGRATION' column for this table, with a red arrow pointing to a callout box labeled 'migrated row'. A dropdown menu is open for the 'TBS FOR TEMP WORK' column, showing options: FBS (FREE SPACE 2.8M bytes), FBS (FREE SPACE 2.8M bytes), TOOLS (FREE SPACE 10.0M bytes), USERS (FREE SPACE 195.3M bytes), and SPPT (FREE SPACE 26.6M bytes). The 'TEMP TABLE NAME' column contains 'temp_table0'.

SCHEMA	TABLE NAME	ROW MIGRATION	ROW MIGRATION SIZE	TABLESPACE NAME	TBS FOR TEMP WORK	TEMP TABLE NAME
SCOTT	A	0	0	USERS		
SCOTT	BONUS	0	0	USERS		
SCOTT	CHAIN_TEST	0	0	USERS		
SCOTT	CHAIN_TEST2	8	3.9K	USERS	FBS (FREE SPACE 2.8M bytes)	temp_table0
SCOTT	CHAIN_TEST3	0	0	USERS		

migrated row

- 7. DB Scope (11) – INDEX Scope

- INDEX INDEX STORAGE
- 2. EXTENT , BLOCK , BYTE
- 3. INDEX INDEX 가
- 4. (B - TREE)
B - TREE

- 7. DB Scope (12) – INDEX Scope

The screenshot displays the Performance Insight WAT DB Scope interface. The left sidebar shows a tree view of the database structure, including Tablespaces, Schemas, and various tables and indexes. The main window is titled 'Index Storage Allocation' and shows a table of index details for the SCOTT schema.

SCHEMA	INDEX NAME	TABLESPACE NAME (INDEX)	SIZE	ORACLE BLOCKS	EXTENTS	TABLE NAME	TABLESPACE NAME (TABLE)
SCOTT	A_IDX	USERS	10KB	5	1	A	USERS
SCOTT	MASTER2_IDX	USERS	25MB	1265	4	MASTER2	USERS
SCOTT	OK_TEST	XPRT	20KB	10	1	TEST	USERS

- 7. DB Scope (14) – TABLESPACE Scope

Performance Insight[®] WAT DB Scope

Function Su | Tablespace | Tablespace | Datafile Info | Datafile Info | Datafile Info | Fragmentation Information | History

Fragmentation Information

SAVE TO HISTORY

TABLESPACE NAME	FRAGMENTS	EXTENTS_COALESCED	PERCENT EXTENTS_COALESCED	DB_BLOCK_SIZE(KB)	BLOCKS FREE	BYTES FREE(KB)	LARGE F CHUNK B
REG	3	3	0.0	2,048	1,452	2,904	
SYSTEM	5	5	0.0	2,048	125,092	270,184	
TOOLS	1	1	0.0	2,048	5,119	10,238	
USERS	322	16	95.0	2,048	84,059	183,198	
XPRT	19	19	0.0	2,048	13,602	27,204	

Fragment

Coalesce fragment

Coalesce


- 7. DB Scope (15) – ISTMON

5

(1) CPU/Memory	(4) Oracle DATAFILE I/O
(2) Oracle system I/O	(5) SQL
(3) Top 5 가	

15

(가)



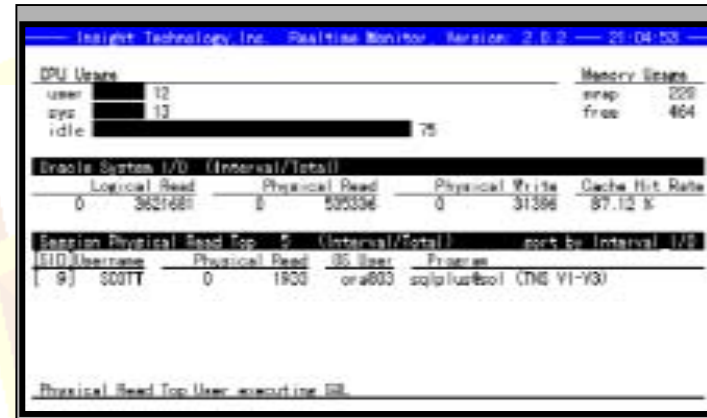
- 7. DB Scope (16) – ISTMON

<mode 1>



- (1) CPU/Memory Usage
- (2) Oracle system I/O
- (3) Top five sessions with heavy physical reads
- (4) Oracle datafile I/O

<mode 2>



- (1) CPU/Memory usage
- (2) Oracle system I/O
- (3) Top five sessions with heavy physical read
- (5) SQL statement with the heaviest load

- 7. DB Scope (17) – ISTMON

<mode 3>

Insight Technology, Inc. Realtime Monitor, Version 2.0.2, 31-05-08

Session Physical Read Top 10 (Interval/Total) sort by Interval 100				
SID	Username	Physical Read	OS User	Program
[9]	SCOTT	196	scott	sqlplus@scott (TNS V1-V3)

Physical Read Top User executing SQL

```
select * from t1@scott;
```

<mode 4>

Insight Technology, Inc. Realtime Monitor, Version 2.0.2, 31-05-08

Oracle Datafile I/O (Interval/Total) sort by File ID					
ID	Datafile Name	Physical Read	Physical Write	Total	Total
1	system01.dbf	0	487418	0	488536
2	tempo003.dbf	0	2954	0	2954
3	ibf0001.ora	0	1990	0	1990
4	spit.dbf	0	244	0	244
5	homsa.dbf	0	91	0	91
6	users01.dbf	0	1485	0	1485
7	users02.dbf	0	17175	0	17175
11	test.dbf	0	89	0	89
12	users02.dbf	56	20928	0	20984
13	curitall_data.db	0	89	0	89
14	ibf01	0	89	0	89
15	test_bsp_t.ora	0	89	0	89
16	test_bsp_l.ora	0	89	0	89
17	testbsp.ora	0	89	0	89
18	curitall_index.d	0	89	0	89
19	test_index.dbf	0	89	0	89
20	users03.dbf	0	144	0	144


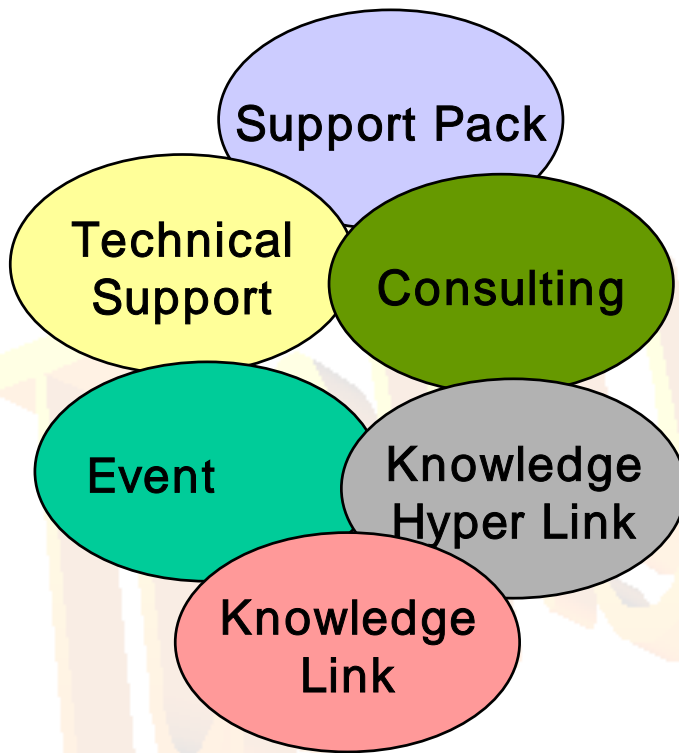
F4:change mode, s/p:toggle FID/Sort, r: L/R:refresh, q:quit

(3) Top five sessions with heavy physical read

(5) SQL statement with the heaviest load

(4) Oracle datafile I/O

- 8. Etc Functions



Q&A

Tuning Frontier

Oracle database server environment

OS	OS Ver		Oracle7.3.2	Oracle8.0.x	Oracle8.1.6
AIX(IBM)	4.2 - 4.3	32.64bit			
HP - UX(HP)	10.20				
	11.00	32.64bit			
Solaris(Sun)	2.5.1				
	2.6				
	2.7	32bit			
		64bit			
Windows(Microsoft)	NT/2000				
RedHat Linux	6.2				
Miracle Linux	1.1				
Turbo Linux Server	6.1				
Tru64UNIX (COMPAQ)	4.0d	64bit			
	5.0	64bit			