



Sensor & Task functions in IDS 11

David Jay
IBM Corporation


Session D12
Tuesday, April 29, 2008 • 4:40 – 5:40 p.m.

2008 IIUG Inform*i*x Conference




Topics Covered

- **What are Sensors and Tasks in IDS 11?**
- **How they are instantiated, stored and used**
- **Common Usage**
 - **Productivity/Performance**
- **How to build a sensor**
- **Q&A**
- **Resources**



What are Sensors and Tasks ?

These are **functions** which are part of the SQL Administration API, managed within the sysadmin database.



It is also known as the Admin API, or Scheduler API, since there is a scheduler process that manages the whole thing.


The Admin API

Aadmin functions return A number.

(an integer return status)

Task functions return Text.

(a string which describes the return status)



Both types of functions perform the command and also write a result to the command_history table.

If the function succeeds but the write to the command_history table fails, it will write a message to the online.log

(It is 2 separate transactions)

What do Sensors and Tasks do?

- **Help the dba to administer one *or more* IDS 11 instances using a system based on SQL commands**
- **Minimize complexity and syntax errors**
- **These task functions can**
 - Help manage resources
 - Allow on-the-fly configuration changes
 - Handle scheduling and running of routines
 - Perform system validation
 - Help to avoid complex shell scripts and command lines




Tasks can automate scheduling or can be set to trigger on a condition.

Tasks allow you to store results for executed commands in order to track trends.


What is a Task?

Tasks are functions used to implement an administrative job or routine. When the task is run, at a minimum, it will write a completion status message to the `command_history` table.




What is a Sensor?

A Sensor is a **task** which not only implements a routine and records a command `_history` reference, it can also *collect and save* information.



3 Main parts of a Task:

- **A definition**, with command parameters which are **stored** as a row **in the ph_task table**
- **A command function** which you **run**
- **A returned integer and lvarchar**, providing a **result** for the command



If the task is also a sensor...

The task definition in `ph_task` table will include:


- **A SQL object**, such as a stored procedure or function, to **collect** the information
- **A table definition** to create the table that will **hold** results for what you collect
- **Interval information** to determine **how and when** the sensor needs **to run, and how often to delete** data that has been stored for a while



Creating a Task or Sensor

We need to recognize –

- the elements that form the **command syntax**
- the defining elements in the **PH_TASK table**




Examples of the command syntax:

```
task ('command')
```

```
task ('command', 'argument' )
```

```
task ('command', size)
```

```
task ('command', 'argument', size)
```

**The Command syntax**

```
task ('command')
```

```
task ('command', 'argument' )
```


```
task ('command', size)
```

```
task ('command', 'argument', size)
```

- size is needed if the command specifies an offset, storage object, or buffer.
- size is a number followed by the unit abbreviation, with no space in between number and unit. (unit options include KB,B,MBGB,TB,PB). Default is KB when not specified.

Command examples:

```
execute function task('check extents',1);  
execute function task('print error',199);  
execute function  
    task("alter plog","physdbs","49 MB");  
execute function  
    task('check data', 2097214);
```



Check extents,1 == checks extents in the first dbspace.

“Check extents” == with no parameters – checks all the dbspaces.

Check data can only be done with a partnum.

The task and the admin functions are created through a routine called db_install.sql, which is visible under \$INFORMIXDIR/etc/sysadmin/.

If you look at the function definitions in db_install.sql, you will discover that task functions can handle up to 10 parameters.


Usage Notes:

- **These commands can be executed to remote instances, allowing a DBA to administer multiple instances through one command window.**
- **The SQL API has more than 80 command variations already.**



The sysadmin database:

- **PH_GROUP**: scheduler group names
- **PH_RUN**: how and when it happened
- **PH_ALERT**: errors, warnings, informational
- **PH_THRESHOLD**: information about thresholds for tasks
- **Results**: historical data about command execution
- **command_history** table –list of all commands run from the API, and the results of the run



Collectively, the sysadmin database contains the tables used by the Scheduler:

Using task or admin commands the administrator can execute functions and procedures or schedule them to run at predefined times.

There are 7 'PH' Tables altogether that contain and organize the Scheduler task information.

The results tables have different names. The built-ins that come with the engine usually have names that start with 'mon_' as in 'monitor'.

The sysadmin tables are created in the rootdbs, and can be moved to a different dbspace if necessary.


The scripts that create these tables are found in \$INFORMIXDIR/etc/sysadmin.

The definitions can be found in db_install.sql

PH_Task table:

This table has 28 columns to describe the parameter and descriptive information needed for each sensor or task:

- **tk_description** -- Description of task
- **tk_result_table** -- Result table name
- **tk_create** -- The create table statement
- **tk_execute** -- The SQL object to execute
- **tk_delete** -- Interval for deleting data
- **tk_frequency** -- How often the task runs



All task and sensor functions you create will be stored in `ph_task`. The column names in green are the fields we can change. On the next couple slides there are a few other columns that the server updates for us.

`tk_description` `lvarchar`

`tk_result_table` `varchar`

`tk_create` `lvarchar`

`tk_execute` `lvarchar`

`tk_delete` `interval day(2) to second`

`tk_frequency` `interval day(2) to second (max 99 days)`

The `tk_result_table` column is used only by sensors and the content matches the table created in `tk_create`.

When the `tk_delete` interval is achieved, data is deleted from `tk_result_table`.


`tk_create` - The `CREATE TABLE` statement which will be executed to create the result table, if needed.

Note: The `tk_create` column is used by sensors. and as necessary, it is executed to create the table needed to hold the data for the sensor we create.

PH_Task table:

- **tk_id** - Sequential task ID
- **tk_name** - Unique Task name
- **tk_start_time** - Starting time of the task
- **tk_stop_time** - Stop time for a task
- **tk_type** - Type of Task


TASK
SENSOR
STARTUP SENSOR



Tk_id	serial	
TK_name	char(36)	
Tk_start_time		datetime hour to second – when it is not a startup sensor
Tk_stop_time		datetime hour to second – ignored if is a startup sensor
tk_type	char(18)	
	Re: tk_type:	Sensor = collects info
		Startup Task – a task only executed when the server starts up
		Startup Sensor – a sensor which is only executed at server startup.
		undocumented -- STARTUP MONITOR -- present on 11.10, removed after.

PH_Task table:

- **tk_next_execution** - next time the task should be executed
- **tk_group** - Group name reference
- **tk_enabled** - scheduled or not scheduled ('t' = default)
- **tk_priority** - Job priority (1-5):
0=Default
5=execute first
1=execute last




tk_next_execution datetime year to second
 tk_group varchar(128) (ties to ph_group(group_name))
 There are 11 predefined groups:
 MISC,DISK,NETWORK,MEMORY,CPU,
 TABLES,INDEXES,SERVER,USER,
 BACKUP, and PERFORMANCE
 (default is MISC).

tk_enabled boolean
 tk_priority integer

ph_task table:

System controlled columns:


- **tk_sequence** - data collection number
- **tk_owner** - owner's thread ID
- **tk_attributes** - system flags for tracking status
- **tk_total_executions** - # of times to execute task
- **tk_total_time** - total time executing this task



Tk_sequence integer
Tk_owner integer
Tk_attributes integer
Tk_exec_num integer
Tk_exec_time integer

Common Usage

- Tailored scripts for
 - Update statistics
 - Monitoring disk capacity
 - Monitoring user activities
 - Maintaining the balance of resource usage
 - (e.g., memory, data distribution, locks)



Before the Admin API came along, the choices were limited to options like the ones in this slide.

Some Guidelines if you want to use the API tools:

If you would like a graphical interface, consider using OAT – the Open Admin Tool.

(this requires some knowledge of php and web interfacing components)

If you need a complete Toolbox package for discovery, analysis and response remediation, consider using Sentinel or Server Studio.

If you need only one tool at a time, consider building the tool with the API

```
LOGBUFF      32      # Logical log buffer size (Kbytes)
#CLEANERS    1      # Number of buffer cleaner processes
CLEANERS     10     # per tech support, 11/16/07 BC
#SHMVIRTSIZE 8192
#SHMVIRTSIZE 32000  # per tech support, 11/16/07 BC
#SHMVIRTSIZE 148000 # 11/19/07 BC
#SHMVIRTSIZE 100000 # 11/22/07 BC
#SHMVIRTSIZE 64000  # 11/22/07 BC
SHMVIRTSIZE  48000  # 11/19/07 BC
#SHMVIRTSIZE 64000  # 11/19/07 BC
SHMADD       8192
SHMADD       16000  #
#SHMADD      148000 # 11/19/07 BC
#SHMADD      100000 # 11/22/07 BC
#SHMADD      64000  # 11/22/07 BC
SHMADD       48000  # 11/19/07 BC
#SHMADD      64000  # 11/19/07 BC
EXTSHMADD    8192  # Size of new extension shared memory segments
```



An example of what the built-in sensors will do for us.


Over time, our imaginary dba, “Big Chuck”, has been busy trying to fine tune his ONCONFIG file...

It is not pretty. At least he documents his changes.

A built-in sensor: mon_config

- ph_task comes with 13 sensors and startup sensors. One of these is **mon_config**.


```
[ select * from sysadmin:ph_task ]
```



(Note: The 11.5 release has 16 sensors)

```
select * from sysadmin:ph_task where
tk_name='mon_config';
```

```
tk_id          2
tk_name        mon_config
tk_description  Collect information about database server's
                configuration file (onconfig). Only modified
                parameters are collected.
tk_type        SENSOR
tk_result_table mon_config
tk_create
create table mon_config (ID integer, config_id integer, config_value
    lvarchar(1024));
create view mon_onconfig as select ID ID, cf_name name, config_value
    value from mon_config, sysmaster:sysconfig where
    mon_config.config_id = sysmaster:sysconfig.cf_id;
```



Below is sample output from mon_onconfig.

In this output (which is not dramatically readable by itself), we can surmise that SHMVIRTSIZE has changed twice since the engine was initialized. In a later example, we can tie in the ph_run output to find out exactly when the changes happened.

```
select * from mon_onconfig where name="SHMVIRTSIZE"
```


```
id  1
name SHMVIRTSIZE
value 0x0x2000
```

```
id  2
name SHMVIRTSIZE
value 0x0x7d00
```

```
id  78
name SHMVIRTSIZE
value 0x0x7d00
```

```
select * from sysadmin:ph_task where
tk_name='mon_config'; (continued)
```

```
tk_dbs      sysadmin
tk_execute  onconfig_save_diffs
...
tk_frequency 1 00:00:00
tk_next_execution 2008-03-13 05:00:00
...
tk_group    SERVER
tk_enable    t
tk_priority  0
```



(The ellipsis on the slide above indicates there are a few fields I am not displaying).
Where is the function onconfig_save_diffs defined? ?

It is in \$INFORMIXDIR/etc/sysadmin/sch_tasks.sql
(first created when the engine was initialized):

```
CREATE FUNCTION onconfig_save_diffs(task_id INTEGER, ID INTEGER)
RETURNING INTEGER
DEFINE value  LVARCHAR(1024);
DEFINE conf_value LVARCHAR(1024);
DEFINE conf_id INTEGER;
LET value = NULL;
FOREACH select cf_id, trim(cf_effective)
INTO conf_id, conf_value
FROM sysmaster:syscfgtab
FOREACH select FIRST 1 config_value
INTO value
FROM sysadmin:mon_config
WHERE mon_config.config_id = conf_id
ORDER BY id DESC
END FOREACH
IF conf_value == value THEN
CONTINUE FOREACH;
END IF
INSERT INTO mon_config VALUES( ID, conf_id, conf_value );
END FOREACH
return 0;
END FUNCTION;
```


A Sensor example:

```
onstat -g seg
```

```
IBM Informix Dynamic Server Version 11.10.UC2  -- On-Line -- Up 28 days 07:20:12  
-- 36480 Kbytes
```

```
Segment Summary:
```


id	key	addr	size	ovhd	class	blkused	blkfree	
229378		1382107137	44000000	12189696	292924	R	2973	3
262147		1382107138	44ba0000	8388608	50184	V	2048	0
294916		1382107139	453a0000	8388608	50184	V	2033	15
327685		1382107140	45ba0000	8388608	50184	V	233	1815
Total:	-	-	37355520	-	-	7287	1833	



Big Chuck notices a bunch of virtual segments, and decides he better get around to tuning his memory; he would like to use the new scheduler API.


```
select * from ph_task where
tk_group="MEMORY"
```

```
tk_id          8
tk_name        mon_memory_system
tk_description  Server memory consumption
tk_type        SENSOR
tk_sequence    917
tk_result_table mon_memory_system
tk_create
create table mon_memory_system
(ID integer, class smallint, size int8, used int8, free int8 )
```



What does he have in the tool box already?

The possible groups are:

MISC,DISK,NETWORK,MEMORY,CPU,TABLES,INDEXES,SERVER,USER,BACKUP,PERFORMANCE

```
select * from ph_task where tk_group="MEMORY"
```

Some things to notice

our results are going into a table named mon_memory_system

- the table has no dates


- if we want dates, we have to do a sql join with the PH_RUN table, using the tk_id

- the does have 'class', 'size', 'used' and 'free' - similar to info from sysmaster:syssegments

- (class = 2 is the virtual class we are looking for)

```
select * from ph_task where  
tk_group="MEMORY" (continued)
```

```
tk_dbs      sysadmin  
tk_execute  insert into mon_memory_system  
select $DATA_SEQ_ID, seg_class, seg_size,  
seg_blkused, seg_blkfree FROM  
sysmaster:sysseglst  
tk_delete   15 00:00:00  
..  
tk_frequency 0 02:00:00  
tk_next_execution 2008-04-28 11:49:28
```



Notes:

-the execute statement for collecting data is grabbing from a new table named sysmaster sysseglst

(on old SMI queries we would have used sysmaster:syssegments)

-It collects 15 days of data before it recycles, and collects every 2 hours.

```
select date(a.run_time), b.size, b.used, b.free from  
ph_run a, mon_memory_system b where class=2 and  
b.id=a.run_id order by 1
```

(expression)	size	used	free
01/09/2008	8388608	2043	5
01/09/2008	8388608	533	1515
01/09/2008	8388608	2048	0
01/09/2008	8388608	2044	4
01/09/2008	8388608	533	1515
01/09/2008	8388608	2048	0
01/10/2008	8388608	2044	4
01/10/2008	8388608	532	1516
01/10/2008	8388608	2048	0
01/10/2008	8388608	2048	0
01/10/2008	8388608	2045	3
01/10/2008	8388608	531	1517
01/10/2008	8388608	2044	4
01/10/2008	8388608	2048	0
01/10/2008	8388608	532	1516

(For space considerations on this slide, I displayed date rather than datetime value.)


From the datetime references, Chuck was able to determine that he was averaging

3 virtual segments a day: he was able to correct this by setting SHMVIRTSIZE from 8192 to 25000.

Building your own tools –

Preliminaries:

- Only a dba can run task() and admin().
- By default, only informix user can connect to sysadmin database.
- All task and admin functions are executed against the sysadmin database.
- If you are creating tables, make sure you account for disk space requirements as rows are collected.



Preliminaries:


- Only a dba can run the task() and admin() functions.
- by default, informix is the only user that can connect to the sysadmin database.
- All task and admin functions are executed against the sysadmin database.

If you need to create a table, use this formula to estimate the amount of disk space needed:

(# of rows collected) * (size of the row collected) * (the # of records per day) *
(the retention period).

The Plan

1. Describe it.
2. Define the data table for storage.
3. Create the SQL to capture the data.
4. Determine the frequency.
5. Plug the entries in `ph_task`.
6. Verify that it works.



To set up a task, you need to plan the task first. You need to have:

- A description of the task you want to monitor
- The table where you want to store data
- The SQL command, stored procedure, or function to capture data
- Information on when and how often you want the task to run, and how often to clear the data.


1. Describe it

- **Monitor disk usage per database and dbspace.**

(DISK or TABLE group ?)




www.iiug.org



Do we have the tool already?


```
select count(*) from ph_task where  
                                tk_group="DISK"  
  
count = 0
```

```
select count(*) from ph_task where  
                                tk_group="TABLES"  
  
count =3
```



**...ph_task where
tk_group="TABLES"**


**mon_command_history
mon_table_names
mon_table_profile**



A closer look at ph_task to see what's in the tables group reveals that none of these tables appear to relate to what we need.

mon_table_profile contains:

**id integer, partnum , nextns , serialval, nptotal,
npused , npdata , lockid , nrows , ucount ,
ocount , pf_rqlock , pf_lockwait , pf_isread,
pf_iswrite, pf_isrwrite, pf_isdelete, pf_bfcread,
pf_bfcwrite, pf_seqscans , pf_dskreads ,
pf_dskwrites**




A closer look at `mon_table_profile` shows it tracks pages and i/o usage, but nothing we can use.

So, we will want to create our own table for tracking.

2. Define the data table for storage


Track:

- The dbspace by name
- The database by name
- Total space used per database
- Number of extents per database



2. Define the data table for storage


<u>dbspace</u>	<u>database</u>	<u>tot_space</u>	<u>no_of_exts</u>
db2	stores7	1263	133
rootdbs	most	1100	103
rootdbs	sysadmin	2232	212
rootdbs	sysmaster	1388	137
rootdbs	sysuser	1072	102
rootdbs	sysutils	1136	114



(A sample of what I want to monitor)


2. Define the data table for storage

```
create table mon_dbs_use (  
ID integer,  
dbspace char(10),  
database char(10),  
tot_space integer,  
no_exts integer)
```



3. Create the SQL to capture the data


```
select distinct
  substr(DBINFO("DBSPACE",partnum),1
,10) DBSpace, dbname[1,10] Database,
  sum(pe_size) tot_space, count(*)
  no_of_exts
from sysmaster:sysptnext,
  sysmaster:systabnames
where pe_partnum = partnum
and tabname != "TBLSpace"
group by 1,2
```



This is the query which I know will capture the data visible in the previous slide.

4. Determine the frequency

- **STARTUP SENSOR**
- **Delete every 90 days**




If this were a regular sensor, I would specify 'SENSOR', and a frequency interval at which to run.

5. Insert the entries in ph_task

Insert into ph_task

```
(tk_name, tk_description, tk_type,  
tk_result_table, tk_create, tk_dbs,  
tk_execute, tk_delete, tk_frequency,  
tk_group, tk_enable)
```

VALUES ...



5. Insert the entries in ph_task (continued)

```
VALUES
('mon_dbs_use',
'Monitor disk usage per database and dbspace',
'STARTUP SENSOR','mon_dbs_use',
'create table mon_dbs_use (ID integer, dbspace
char(10),database char(10), tot_space integer,no_exts
integer)','sysadmin','insert into mon_dbs_use select
$DATA_SEQ_ID, distinct substr(DBINFO(DBSPACE,par
tnum),1,10) DBSpace, dbsname[1,10]
Database,sum(pe_size) tot_space, count(*) no_of_exts from
sysmaster:sysptnext, sysmaster:systabnames where
pe_partnum = partnum and tablename != TBLSpace group by
1,2','90 00:00:00', '0 00:01:00','DISK','t')
```




1. When inserting your sql into the tk_execute column, if you have too much hardship getting the SQL expression (in blue) to work because of misplaced quotes or other errors, create the sql as a standalone function, then insert the name of the function to be called in the tk_execute column.
2. Always include a \$DATA_SEQ_ID in your tk_execute process, and an integer ID column in your result table, so you can keep track of the run_task sequence.

6. Verification

1. Verify our entry in ph_task

```
select * from ph_task where tk_name="mon_dbs_use"
```



When our task is inserted, it is also executed in order to create the result table.

In our first verification step, we are verifying insert syntax...


We should get back the row that we just put in:

```
select * from ph_task where tk_name="mon_dbs_use";
```

6. Verification

1. Verify the entry in Ph_task

```
select * from ph_task where tk_name="mon_dbs_use"
tk_id          14
tk_name        mon_dbs_use
tk_description Monitor disk usage per database and dbspace
tk_type        STARTUP SENSOR
tk_sequence    5
tk_result_table mon_dbs_use
tk_create      create table mon_dbs_use (ID integer,
              dbspace char(10),database char(10), tot_space
              integer,no_exts integer)
tk_dbs         sysadmin          ...
```



Note the tk_id assigned for our serial column: 14. (We will use this for the next verification step.)


6. Verification

2. Verify the existence of the result table:

```
select * from mon_dbs_use;
```

3. If our result table is empty, check ph_run for errors:

```
select * from ph_run where run_task_id=14
```



2. Verify the existence of the result table:


```
select * from mon_dbs_use;
```

If the table does not exist or it has no rows, we have a syntax error in the ph_task table.

6. Verification

```
select * from ph_run where run_task_id=14;
```


run_id	4760
run_task_id	14
run_task_seq	1
run_retcode	-201
run_time	2008-04-17 10:32:49
run_duration	0.021375137579
run_ztime	1205400729
run_btime	1205400729
run_mtime	1205767969



Selecting from the `ph_run` table for our task id, we find that there is a -201 (syntax error) in our executed statement (`tk_execute` has a problem).

(fix for the values clause used on step 5)

```
update ph_task set tk_execute= "insert into
mondbs_use select distinct $DATA_SEQ_ID,
substr(DBINFO('DBSPACE',partnum),1,10)
DBSpace, dbsname[1,10] Database,
sum(pe_size) tot_space, count(*) no_of_exts
from sysmaster:sysptnext,
sysmaster:systabnames where pe_partnum =
partnum and tabname != 'TBLSpace' group by
2,3" where tk_name="mon_dbs_use"
```




Be careful with the insertion clauses. Use double quotes to bracket the whole phrase,
single quotes for the parts inside.

6. Verification

```
select * from ph_run where run_task_id=14;
```

run_id	4775
run_task_id	14
run_task_seq	4
run_retcode	0
run_time	2008-04-17 11:32:11
run_duration	0.143057564702
run_ztime	1205774878
run_btime	1205774878
run_mtime	1205774951



After an update for our startup sensor, our tk_execute column is now fixed – but the task is not going to run automatically, since it is not newly instantiated.


To verify it at this point, we have to bounce the engine.

After the bounce, a check of the ph_run table will now show a 0 return code.

6. Verification


```
select * from mon_dbs_use;
```

id	dbspace	database	tot_space	no_exts
1	db2	stores7	1263	133
1	rootdbs	most	1100	103
1	rootdbs	sysadmin	1496	147
1	rootdbs	sysmaster	1388	137
1	rootdbs	sysuser	1072	102
1	rootdbs	sysutils	1136	114




And our mon_dbs_use table has data..

Q&A ?



Resources

- **IBM Informix Dynamic Server Administrator's Reference**
- **IBM Informix Guide to SQL: Syntax**
- **"Informix Dynamic Server 11:Advanced Functionality for Modern Business" a redbook located at** <http://www.redbooks.ibm.com/abstracts/sg247465.html?Open&pdfbookmark>
- **Certification Tutorial, part 2.**
- **Focused summary on monitoring at**
<http://www.serverstudio.com/>



Session D12
Sensor & Task functions in IDS 11

David Jay
IBM Corporation
djay@us.ibm.com

