WHAT'S NEW IN IBM INFORMIX V.14.10.XC7

v.1a

Agenda

- `onconfig` change to mirroring / mirror path
- Java updates
- GSKit upgrades
- Removal of the `ifxdeployassist`
- ER / CDC update
- `oncheck -w #`
- `onstat -I`
- `onstat -g rah` enhancements
- IHQ 1.6.0

`onconfig` change to mirroring / mirror path

onconfig change to mirroring

- This is an undocumented change
- The MIRROR parameter is now enabled by default and removed from onconfig
- If you use a FC6 (or earlier) file where MIRROR is disabled, the value is ignored
- If you use a FC6 (or earlier) file where MIRRORPATH is populated with default values, that value is used to create a mirror of the rootdbs regardless of any MIRROR value in the file!!
  - This only affects newly created instances

onconfig change to mirroring

- For example, an FC6 file used to create a new instance from scratch in FC7

```
ROOTNAME rootdbs
ROOTPATH rootspace
ROOTOFFSET 0
ROOTSIZE 100000
MIRROR 0
MIRRORPATH $INFORMIXDIR/tmp/demo_on.root_mirror
MIRROROFFSET 0
```

```
Inst_1_7: oninit -v
ISI API version mismatch: 2.2 required, 2.1 found
attn: Ignoring unknown or deprecated config parameter (MIRROR)
```

```
Dbspaces
address           number    flags      fchunk    nchunks   pgsize    flags     owner     name
45733028          1         0x2        1         1         2048      M  BA     informix  rootdbs
45734bf0          2         0x1        2         1         2048      N  BA     informix  data_space_1
467f8028          3         0x1        3         1         2048      N  BA     informix  log_space
467f8268          4         0x2001     4         1         2048      N TBA     informix  work_space
467f84a8          5         0x8001     5         1         2048      N SBA     informix  slob_space
467f86e8          6         0xa001     6         1         2048      N UBA     informix  slob_temp
 6 active, 2047 maximum

Chunks
address           chunk/dbs     offset     size       free       bpages     flags pathname
45733268          1      1      0          50000      37812                 PO-B-D rootspace
45734028          1      1      0          50000      0                     MD-B-- /opt/IBM/informix/14_10/tmp/demo_on.root_mirror
467f9028          2      2      0          102400     91223                 PO-B-D dataspace
46801028          3      3      0          102400     32347                 PO-B-D logspace
46802028          4      4      0          102400     102347                PO-B-- tmpspace
46803028          5      5      0          25600      23689      23802      POSB-D sbspace
                                Metadata   1745       1298       1745
46804028          6      6      0          102400     95432      95432      POSB-- sbtemp
                                Metadata   6915       5146       6915
 6 active, 32766 maximum
```

onconfig change to mirroring

- To prevent this mirror from being created when you initialize an instance, you must disable / erase the MIRRORPATH parameter

```
ROOTNAME rootdbs
ROOTPATH rootspace
ROOTOFFSET 0
ROOTSIZE 100000
# MIRROR 0
# MIRRORPATH $INFORMIXDIR/tmp/demo_on.root_mirror
MIRROROFFSET 0
```

```
Server Version 14.10.FC7DE -- On-Line -- Up 00:01:26 -- 234176 Kbytes

address           number   flags      fchunk   nchunks  pgsize   flags    owner    name
4596e028          1        0x20001    1        1        2048     N  BA    informix rootdbs
46fac028          2        0x20001    2        1        2048     N  BA    informix data_space_1
47520c20          3        0x20001    3        1        2048     N  BA    informix log_space
47594c20          4        0x2001     4        1        2048     N TBA    informix work_space
4752a660          5        0x28001    5        1        2048     N SBA    informix slob_space
47c9bd28          6        0xa001     6        1        2048     N UBA    informix slob_temp
 6 active, 2047 maximum

Chunks
address           chunk/dbs     offset     size       free       bpages     flags pathname
4596e268          1      1      0          50000      1046                  PO-B-- rootspace
46fac268          2      2      0          102400     102347                PO-B-D dataspace
47b89028          3      3      0          102400     102347                PO-B-D logspace
47ba1028          4      4      0          102400     102347                PO-B-- tmpspace
47c33028          5      5      0          25600      23802      23802      POSB-D sbspace
                                 Metadata 1745       1298       1745
47547028          6      6      0          102400     -1         -1         POSB-- sbtemp
 6 active, 32766 maximum
```

onconfig change to mirroring

- Why does this matter?
- You may have an instance try to create the mirror you were not expecting resulting is increased file system usage, potentially slower performance if `$INFORMIXDIR/tmp` (the default location) is not on the fastest disk and more
- If the instance is preexisting and you're performing an in-place upgrade to FC7, it will NOT create the mirror chunk
  - However you will get an error message on instance start that the root `MIRROR` chunk can not be located

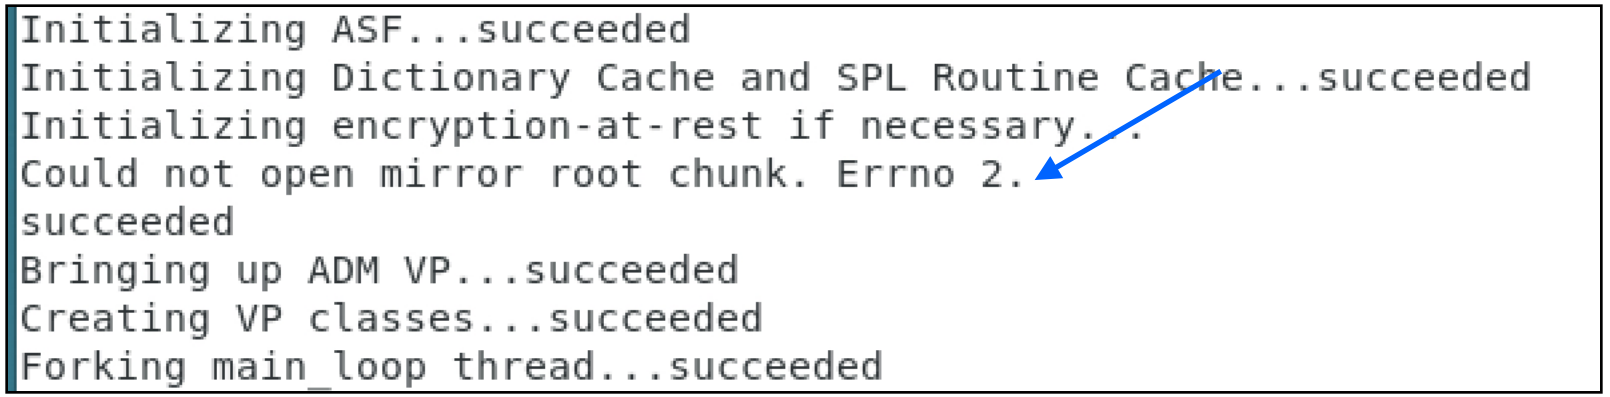

Questions

Java updates

Java updates

- So, ... what happened here??

- Well, according to development just general maintenance and being current with the most recent release
  - There isn’t anything visible to the end user

Questions

GSKit updates

GSKit upgrades

- So, ... what happened here??

- Well, according to development not a whole lot
  - There was some code cleanup and some internal changes but nothing that is visible to the end user

Questions

`IFXdeployassist` removal

ifxdeployassist

- The ifxdeployassist utility was introduced in Informix v.11.70
  - It was used to extract some or all of a source instance to create a snapshot to roll out as needed
  - It had a CLI and a graphical interface

  - It could capture
    - Some or all of the Informix binary
    - Some or all of an instance's data stored in cooked files
    - Both

- The utility was intended for use by partners to roll out their application and an Informix environment (binary and data) in one operation with the ifxdeploy utility

ifxdeployassist

- For example

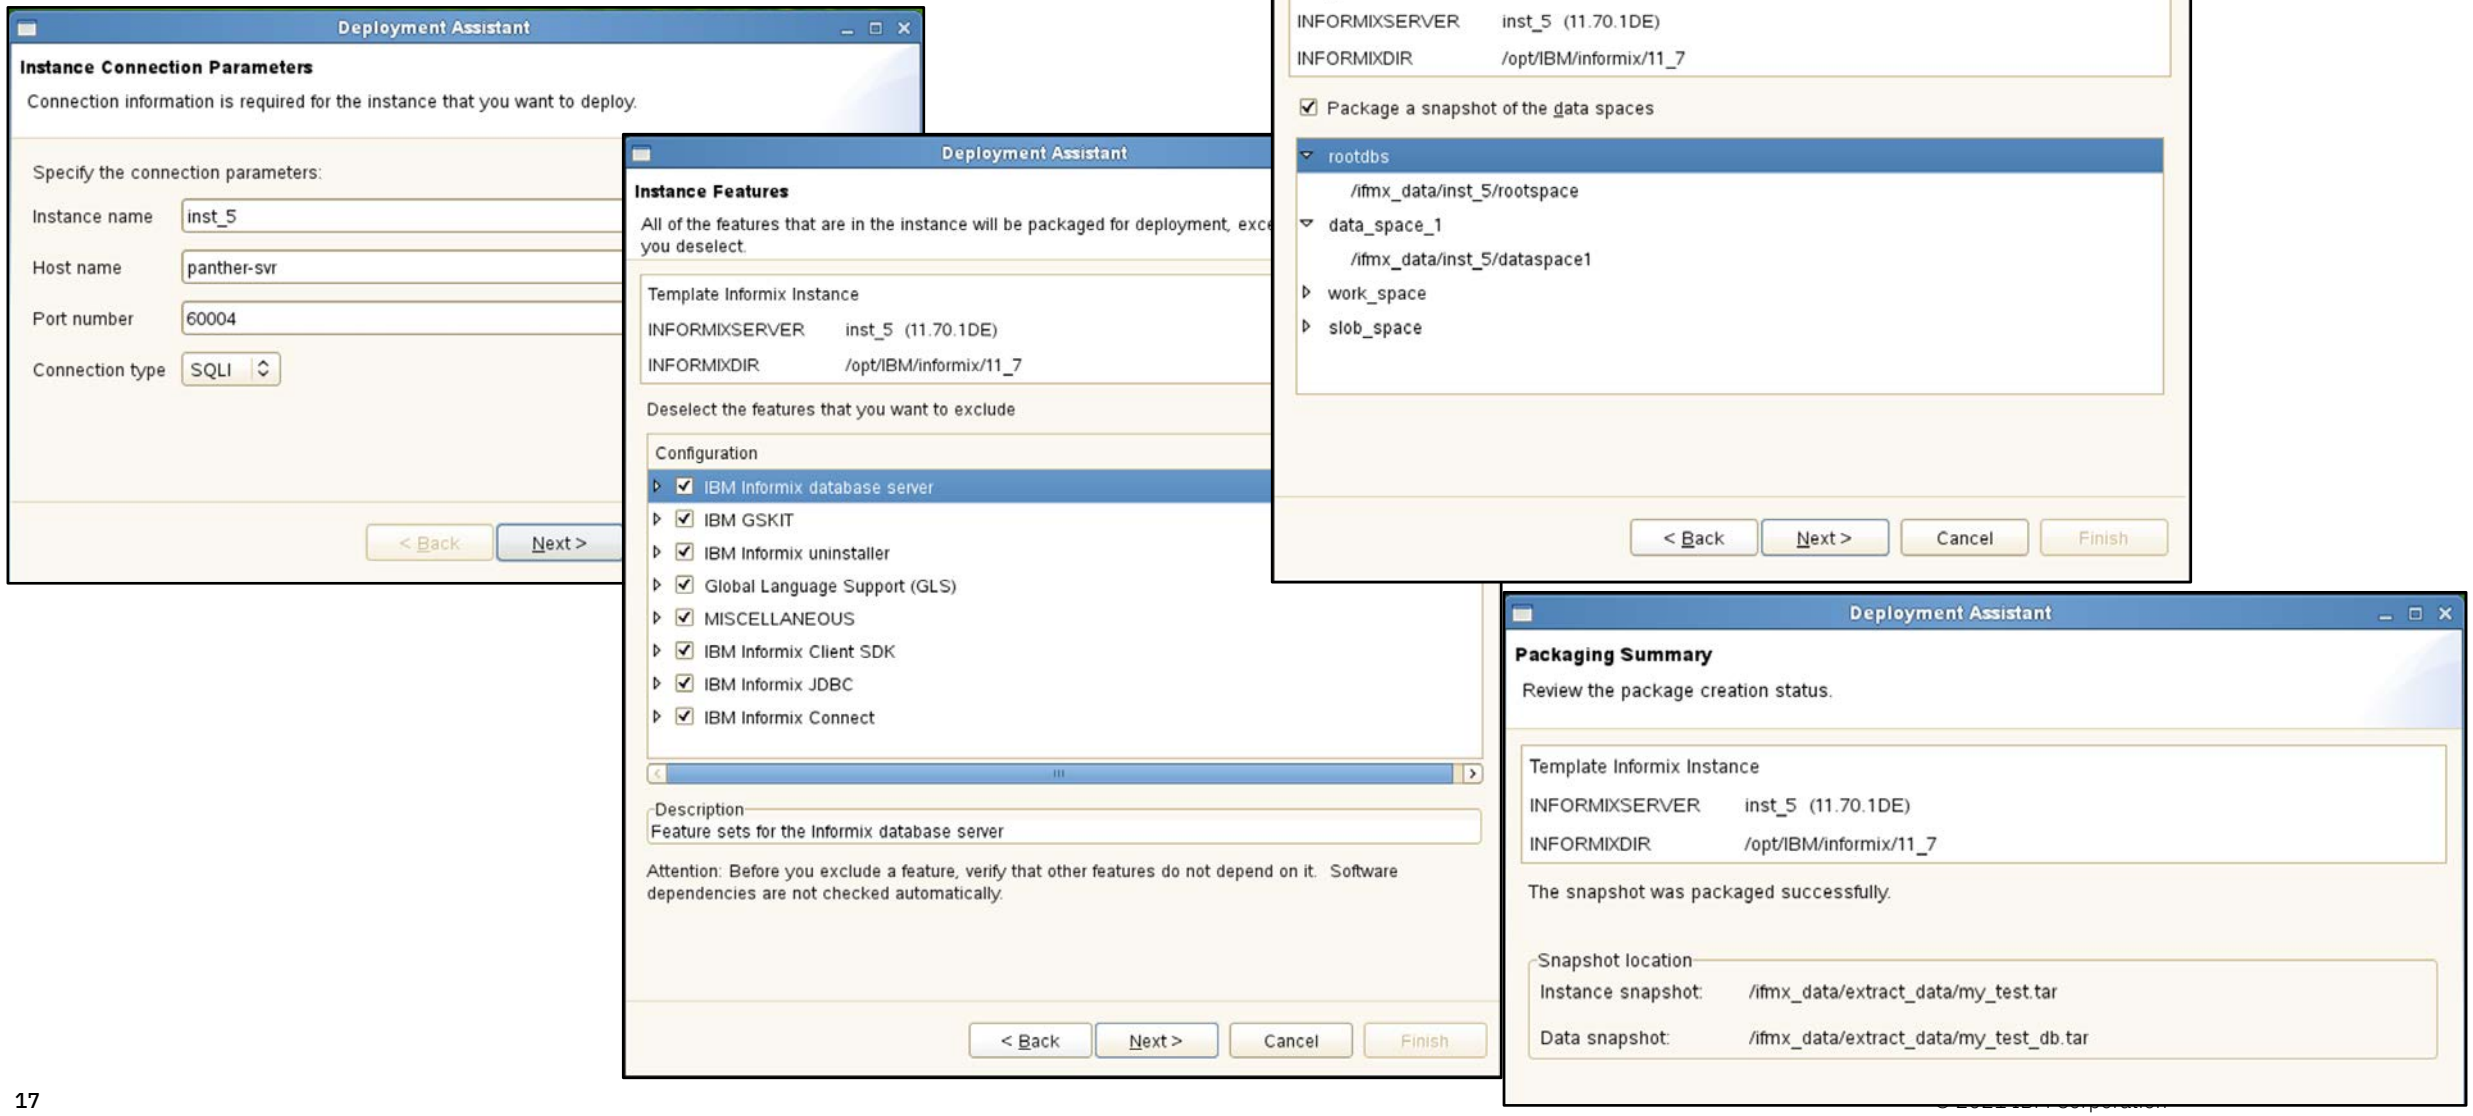

`ifxdeployassist`

- In real life however, it wasn't used much

- With the improved binary installer along with `ifxclone` and `cdr migrate server`, there are more and better options to populate a target

- As a result, the utility has been withdrawn
  - You can cry now if you want 😢😭 Haha!!

Questions

ER / CDC update

ER / CDC update

- Some background
  - Historically, when performing an in-place table alter and ER / CDC is active, the operation caused a flood of `page-out-of-sync` errors
  - In a *slow alter*, it was not a problem because the pages are immediately converted to the new table schema
    - With an IPA, the schema on page is updated when *data on the page changes*
      - This causes a schema version mismatch error to occur flooding the logs with messages
- With FC7, there are two significant changes to performing a table alter
  - First, CDC capture on the table must be stopped with the `cdc_endcapture()` command
  - Second, full row logging must be disabled with the `cdc_set_fullrowlogging()` command
  - The alter will be processed then you can reset full row logging and restart capture

ER / CDC update

- Syntax

```
cdc_set_fullrowlogging (table_name, [0, 1])
```

where

`0` (zero) - disable

`1` (one) - enable

Questions

`oncheck -w` (seconds)

oncheck wait

- A new -w (*seconds*) flag can be appended to any oncheck command

- If the command is blocked by a lock, it will wait the # of seconds specified
  - If the lock clears, the command continues
  - Otherwise, the command fails with the appropriate message

- Obviously, you should use a reasonable amount of time to wait
  - However, if you use the -w flag with NO time value, the command will wait indefinitely
    - Which might or might not be what you want it to do

Questions

onstat -I (capital “eye”)

onstat -I

- A new flag has been added to `onstat`
  - No!! Say it isn't true?!?!?!?!?!? 🤣😉😍
- The "I" (capital "eye") flag displays general index information such as
  - The number of times B-Tree indexes have been searched
  - The number of inserts or deletes into / from indexes
  - The number of reverse scans
  - The number of node splits
  - Index level merges
  - And more

```
Inst_1_7: onstat -I

IBM Informix Dynamic Server Version 14.10.FC7DE -- On-Line (Prim) --
2021-12-02 15:13:02

Index statistics
search      additem     delitem     retry       rev_scans   rs_retries
72179       22123       750         0           504         0

splits      page        slot        root        copyback
117         110         1           48          0

compress    merges      shuffles    root
0           0           0           0
```

Questions

`onstat -g rah` enhancements

Read-ahead reporting enhancements

- This is very much a work still in progress
- While some nice enhancements were added in FC7, it is reasonable to expect one or two additional changes will occur in the next fix pack

- Four new pseudo-tables have been added to sysmaster
  - `sys_ra_act_queues` — contains information about read-ahead queues with pending requests
  - `sys_ra_act_daemons` — contains information about read-ahead daemons currently processing requests
  - `sys_ra_ptn_stats` — contains read-ahead statistics broken down by partition
    - This information is returned by the `onstat -g rah partitions` command
  - `sys_ra_thrd_stats` — contains read-ahead statistics broken down by the requesting thread
    - The information is returned by the `onstat -g rah threads` command

Read-ahead reporting enhancements

- At this time, there is no session or SQL statement level R/A information available

```
{ Read Ahead Queues }
    create table informix.sys_ra_act_queues
        (
        ra_req  int8,
        type    char(10),
        req_id  integer,
        state   char(10),
        owner   int8,
        partnum integer,
        abort   smallint
        );
    revoke all on informix.sys_ra_act_queues
```

```
{ Read Ahead Threads }
    create table informix.sys_ra_act_daemons
        (
        rstcb           int8,
        ra_req          int8,
        req_id          integer,
        type            char(10),
        state           char(10),
        partnum         integer,
        pgs_read        integer,
        threshold       integer,
        pgs_ahead       integer
        );
    revoke all on informix.sys_ra_act_daemons
```

```
{ Partition Read Ahead Stats }
    create table informix.sys_ra_ptn_stats
        (
        partnum         integer,
        bfcread         integer,
        dskreads        integer,
        hit_ratio       smallint,
        data_reqs       integer,
        data_npages     integer,
        data_nios       integer,
        data_eff        smallint,
        leaf_reqs       integer,
        leaf_npages     integer,
        leaf_nios       integer,
        idx_eff         smallint,
        idxdata_reqs    integer,
        idxdata_npages  integer,
        idxdata_nios    integer,
        idxdata_eff     smallint,
        logrec_npages   integer,
        logrec_nios     integer,
        logrec_eff      smallint,
        lc_nreqs        integer,
        lc_nused        integer,
        lc_resched      integer,
        lc_fail         integer,
        lc_eff          smallint,
        ptn_reqs        integer,
        ptn_npages      integer,
        ptn_nios        integer,
        ptn_eff         smallint
        );
```

```
{ Thread Read Ahead Stats }
    create table informix.sys_ra_thrd_stats
        (
        tid             integer,
        bfcread         integer,
        dskreads        integer,
        hit_ratio       smallint,
        data_reqs       integer,
        data_npages     integer,
        data_nios       integer,
        data_eff        smallint,
        leaf_reqs       integer,
        leaf_npages     integer,
        leaf_nios       integer,
        idx_eff         smallint,
        idxdata_reqs    integer,
        idxdata_npages  integer,
        idxdata_nios    integer,
        idxdata_eff     smallint,
        logrec_npages   integer,
        logrec_nios     integer,
        logrec_eff      smallint,
        lc_nreqs        integer,
        lc_nused        integer,
        lc_resched      integer,
        lc_fail         integer,
        lc_eff          smallint,
        ptn_reqs        integer,
        ptn_npages      integer,
        ptn_niosq       integer,
        ptn_eff         smallint
        );
```

Read-ahead reporting enhancements

- This is a idle test system

```
Inst_1_7: onstat -g rah

IBM Informix Dynamic Server Version 14.10.FC7DE -- On-Lin
2021-12-02 15:20:31

Read Ahead

# Qs                            1
# RA Daemon Threads             1
# Requests                      361
cur. queued                     0
max. queued                     2
max. indiv. Q len               2
# Continued                     0
# Memory Failures               0
Way Behind                      0
Daemon Frees                    0
Last Thread Add                 12/02/2021.12:27:37

Active Read Ahead Queues


Active Read Ahead Daemons


Recovery ReadAhead Statistics
```

```
Partition ReadAhead Statistics

         Buffer           hit   Data                    Index                   I/D                     LogRec/PageList Last_Committed                          Partn_Pages
partnum  bfcread dskreads ratio #reqs pagecnt nios eff #reqs pagecnt nios eff #reqs pagecnt nios eff npages nios eff #reqs #used #resched #fail eff #reqs pagecnt nios eff
0x100004 17001   42       99    0     0       0    0   8     8       0    0   1     23      20   86  0      0    0   0     0     0       0     0   0     0       0    0
0x100005 23465   113      99    0     0       0    0   19    5       1    20  25    32      5    15  0      0    0   0     0     0       0     0   0     0       0    0
0x100006 1056    54       94    1     49      37   75  1     0       0    0   5     5       1    20  0      0    0   0     0     0       0     0   0     0       0    0
0x100007 519     31       94    0     0       0    0   8     7       6    85  5     6       0    0   0      0    0   0     0     0       0     0   0     0       0    0
0x100008 161     5        96    0     0       0    0   0     0       0    0   2     2       0    0   0      0    0   0     0     0       0     0   0     0       0    0
0x100009 1447    20       98    0     0       0    0   1     1       0    0   4     48      12   25  0      0    0   0     0     0       0     0   0     0       0    0
0x10000a 400     3        99    1     1       1    100 0     0       0    0   0     0       0    0   0      0    0   0     0     0       0     0   0     0       0    0
0x100013 667     135      79    0     0       0    0   1     19      18   94  1     110     109  99  0      0    0   0     0     0       0     0   0     0       0    0
```

```
Thread ReadAhead Statistics

    Buffer           hit   Data                    Index                   I/D                     LogRec/PageList Last_Committed                          Partn_Pages
tid bfcread dskreads ratio #reqs pagecnt nios eff #reqs pagecnt nios eff #reqs pagecnt nios eff npages nios eff #reqs #used #resched #fail eff #reqs pagecnt nios eff
0   2624    146      94    1     0       0    0   0     0       0    0   0     0       0    0   0      0    0   0     0     0       0     0   0     0       0    0
56  188     146      22    0     0       0    0   0     0       0    0   2     132     129  97  0      0    0   0     0     0       0     0   0     0       0    0
59  53556   871      98    1     38      13   34  12    60      0    0   13    14      1    7   0      0    0   0     0     0       0     0   0     0       0    0
57  11890   458      96    4     20      18   90  23    86      38   44  25    300     153  51  0      0    0   0     0     0       0     0   0     0       0    0
58  59858   762      98    0     0       0    0   17    26      1    3   14    19      4    21  0      0    0   0     0     0       0     0   0     0       0    0
```

Read-ahead reporting enhancements

- With this new information, the instance can calculate and print the effectiveness of R/A operations as well as break it down by R/A type (DATA, LEAF, DATA/INDEX, PARTITION, LASTCOMMIT, etc)

- This information helps understand the effectiveness of R/A based on the percentage of queried pages that came from the R/A cache

- There is currently a discussion going on about how to make this information even more valuable
  - For example, how many R/A cached pages are actually being used?
    - The use case here is that system memory may be used to cache too much data (result - low usage percentage)
    - In resource constrained environments, this can result in buffer churn trying to keep the R/A queues full but they are *not* used so it's wasted I/O
      - In other words, is too much R/A a bad thing???
  - Look for more in the next fix pack

Questions

IHQ 1.6.0

IHQ 1.6.0

- Please note — a new fix pack for v.14.10 and v.12.10 has been released that includes upgraded IHQ binaries that eliminate the log4js vulnerability!
  - Please use the w1 fix packs for both interims to get the new binaries

IHQ 1.6.0

- There are numerous changes and updates in this release
  - Additional documentation clarity
  - View query plan execution
  - Additional H/A cluster functionality
  - Search within system reports and expanded search throughout IHQ
  - Storage Manager enhancements
  - General stuff

Documentation clarity

- Two areas of IHQ functionality have additional documentation describing their use
  - Logging
  - Starting IHQ with Java

- For logging, IHQ uses the `log4j2` library with the `INFO` level of messages being sent to the server and agent logs
  - You can customize this in the `server.log4j.xml` / `agent.log4j.xml` files to display the level of information you need
  - There are options for displaying log messages to the console as well as a rolling log file with various trigger and rollover policies

Documentation clarity

- Starting with IHQ 1.2.0, a shell script that makes starting, stopping and managing server and agent services was introduced
  - With 1.4.0, there have been numerous changes to the functionality
  - With Informix v.14.10.FC6 and IHQ 1.5.0, there were additional changes to the shell script functionality
    - Korn shell for AIX
    - “Help” command functionality
    - Parameter consistency (all lower case)
    - Process naming removed due to AIX
    - `.jar` and `.properties` file renaming
    - The `list` parameter to display active IHQ server / agent processes
    - The stop process was significantly enhanced for accuracy and reliability
    - Connection property testing
    - Significant overhaul and enhancement to the Schema Manager
    - ISAM and incident reporting improvements

Documentation clarity

- While the shell scripts are my preferred method to interact with IHQ, you can use the raw Java commands
  - There is additional information about how the logging mechanism discussed earlier and the Java `start` commands interact
    - This affects what gets displayed when the IHQ server is started

IHQ 1.6.0 - query plan

- Within the **Schema Manager**, you can execute an SQL statement through the **SQL Editor**
  - With this release, you now have access to the query plan (in theory)
    - However it requires that SQL tracing is enabled

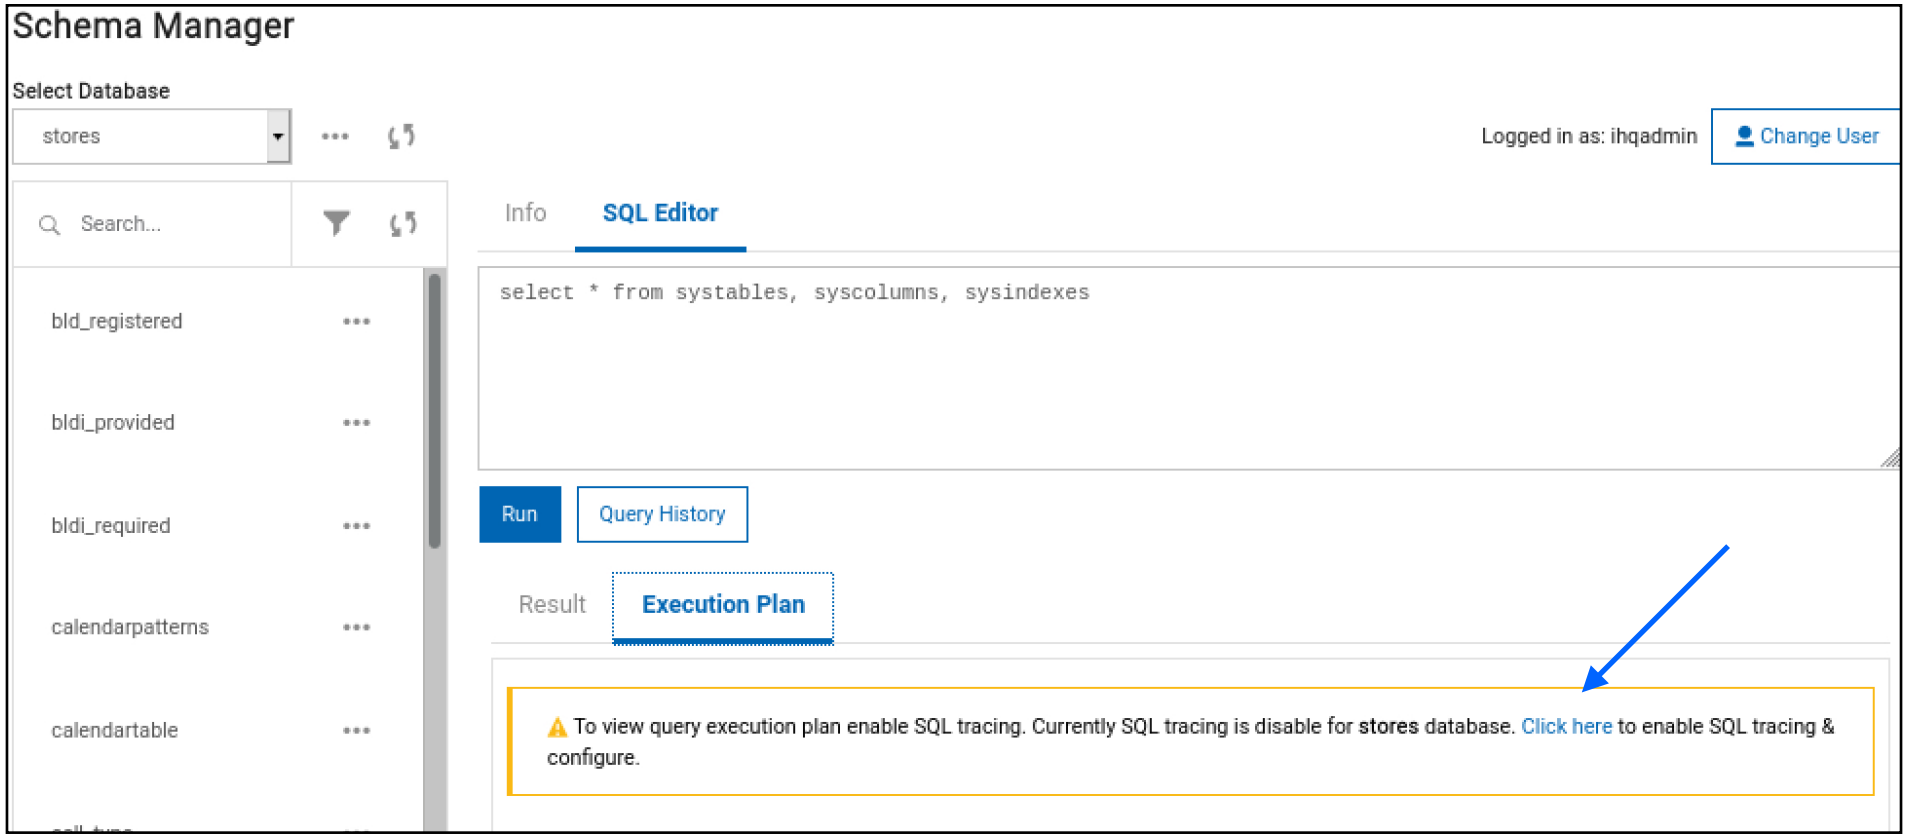

IHQ 1.6.0 - query plan

- Selecting the **Edit** button allows you to modify the tracing parameters to your needs
  - Then click the **Save** button

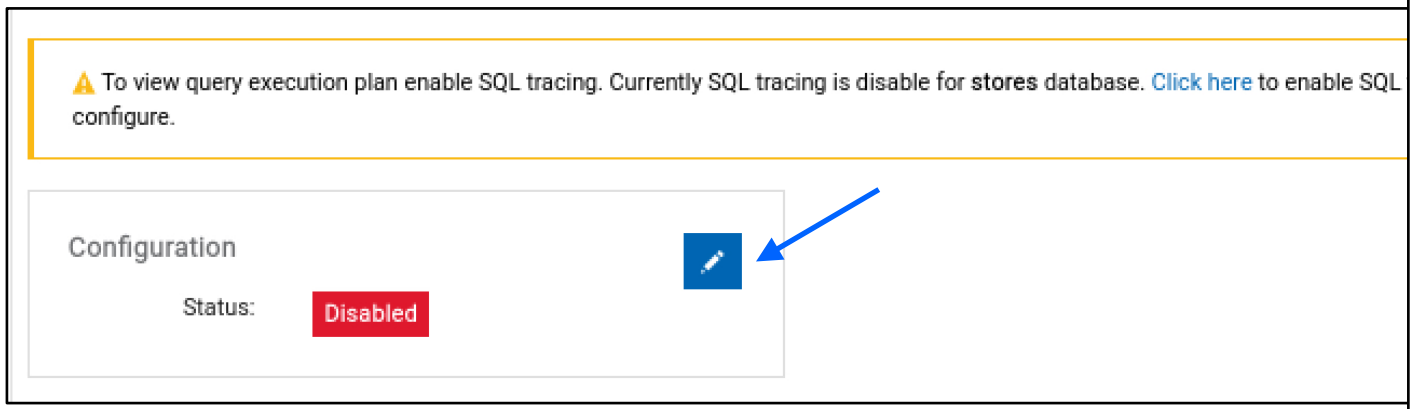

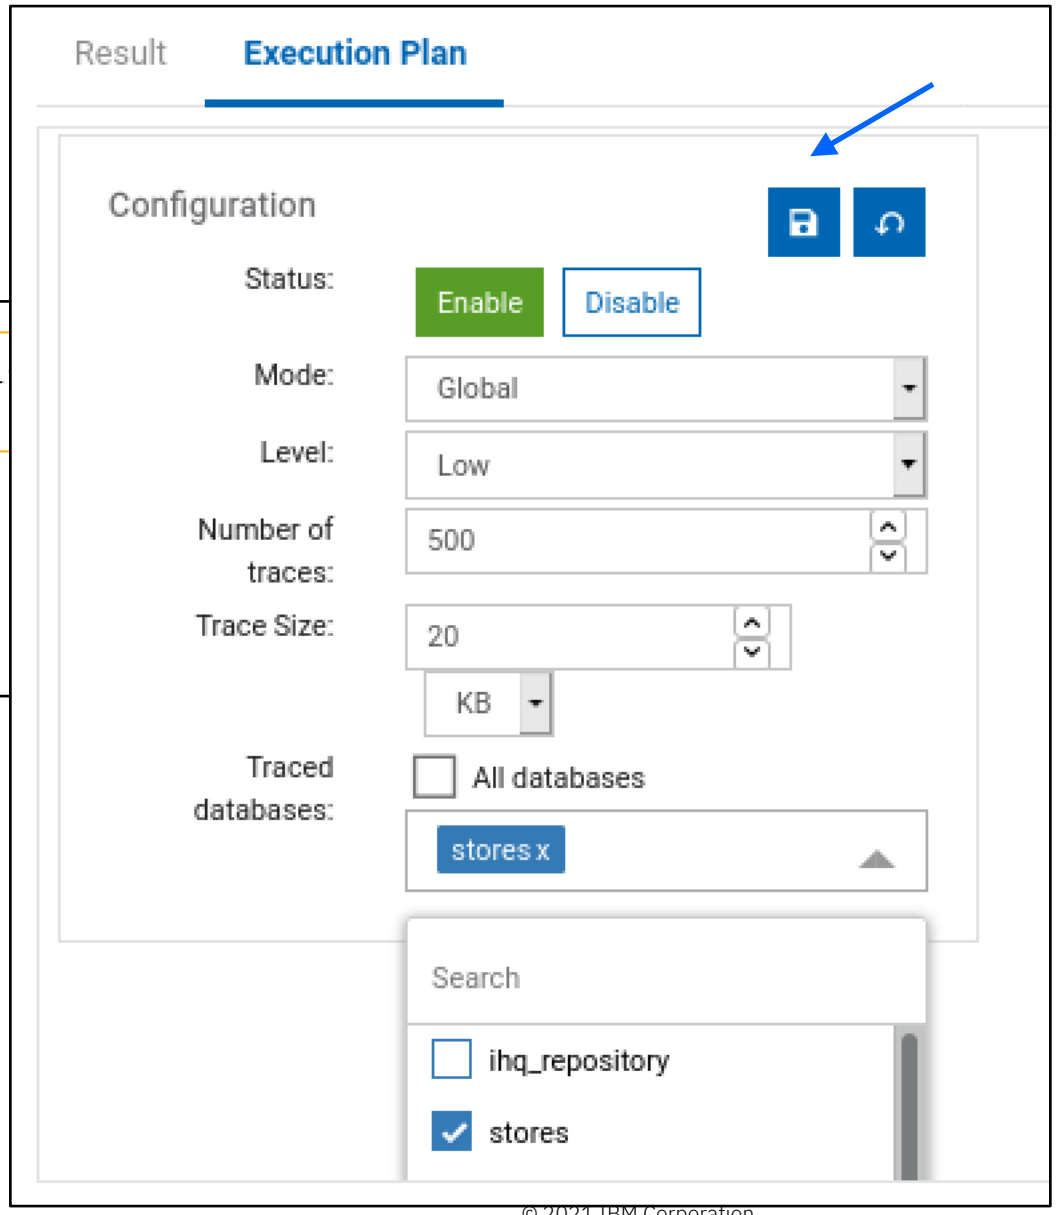

Note: Changing the SQL tracing parameters through this interface is NOT permanent. If you reboot IHQ or the instance, tracing reverts back to instance default.

IHQ 1.6.0 - query plan

- Run your query through the IHQ **SQL Editor**
  - When completed, click the **Execution Plan** option

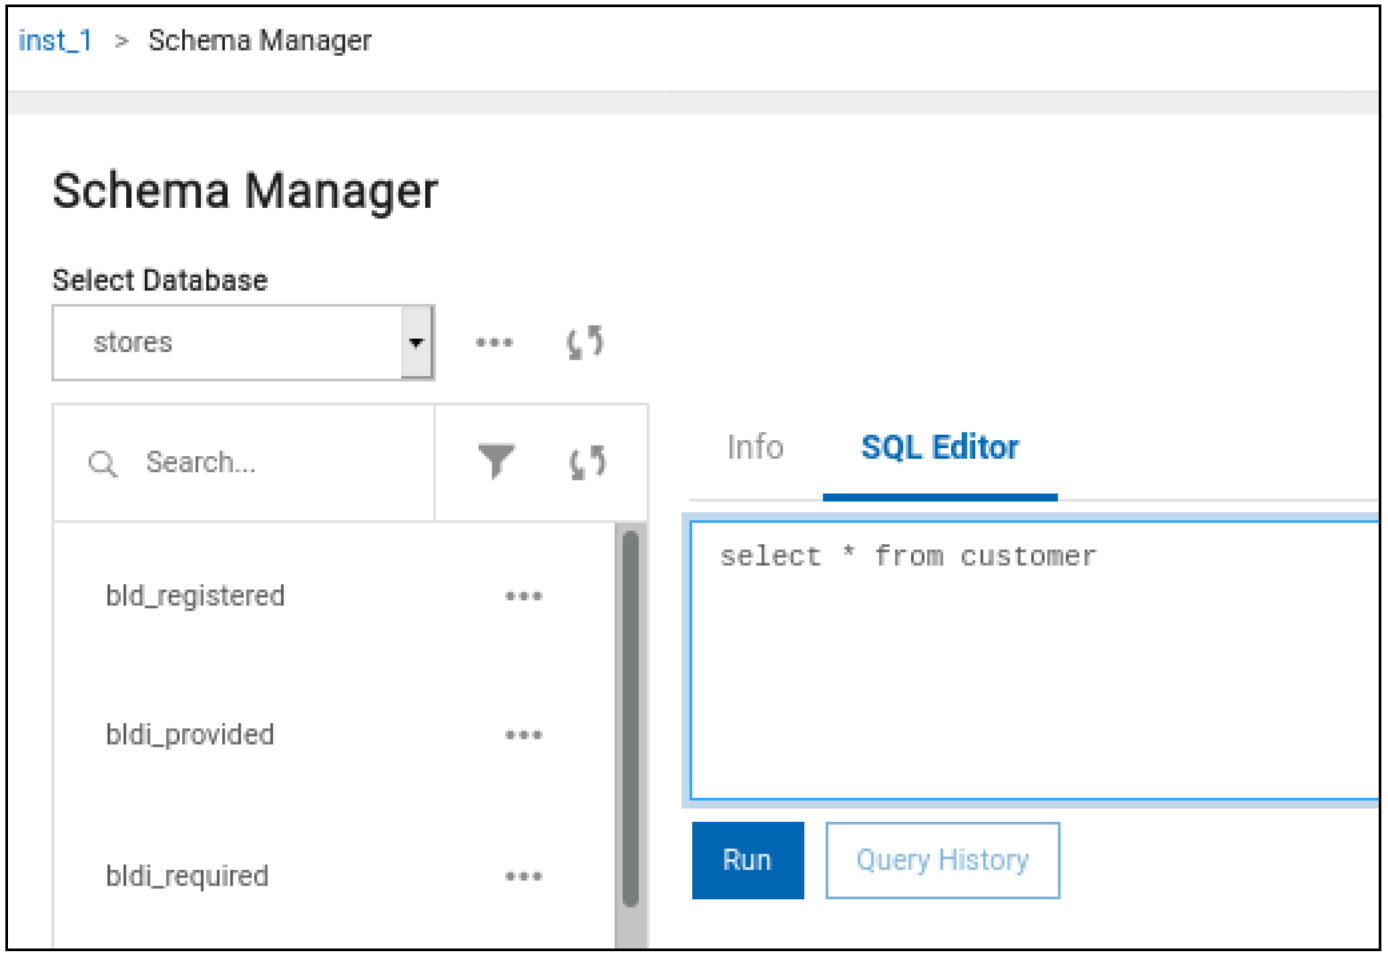

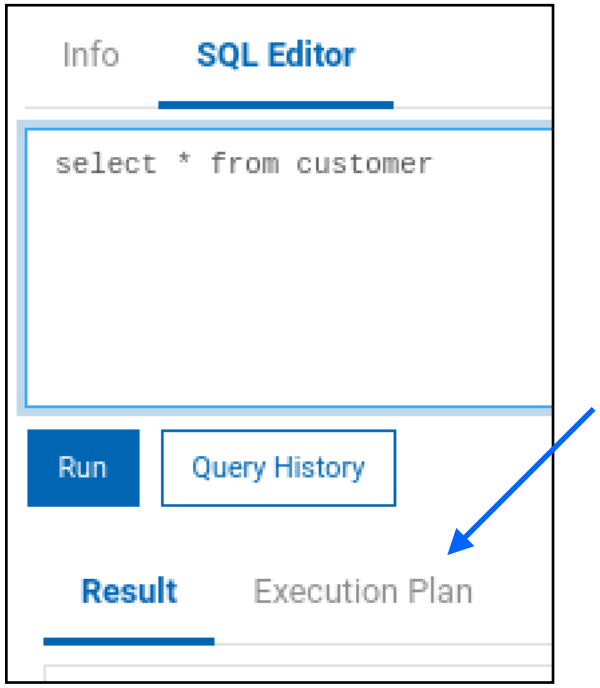

IHQ 1.6.0 - query plan

- In theory . . . . . the query plan should be displayed
  - I was never able to get it to work regardless of how simple or complex the query
- When I queried development:

> "Sometimes, based on sever configuration, [the] server may take longer to gather sql statistics and generate execution plan.
> Or sometime when IHQ UI makes call to Java API (which internally gets information from Informix server), it may not respond promptly.
>
> What we have implemented in IHQ, if it takes more than 10 seconds to fetch execution plan, it will show the pop up what you are getting. If it is not displaying query plan by "Try Again" button, then by re-executing query manually (clicking on run button), it should display execution plan."

  - As I said, it never worked for me

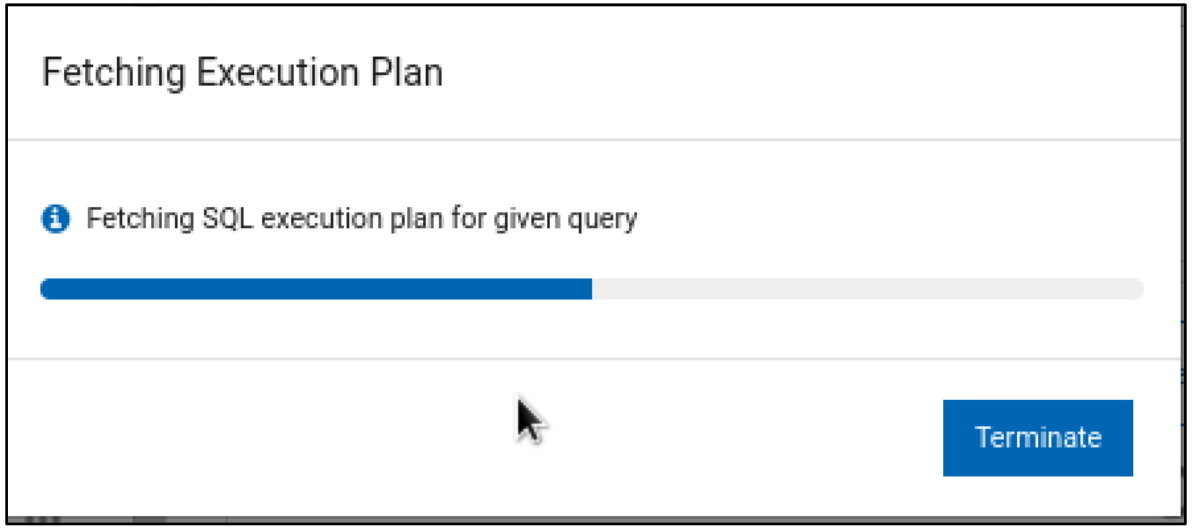

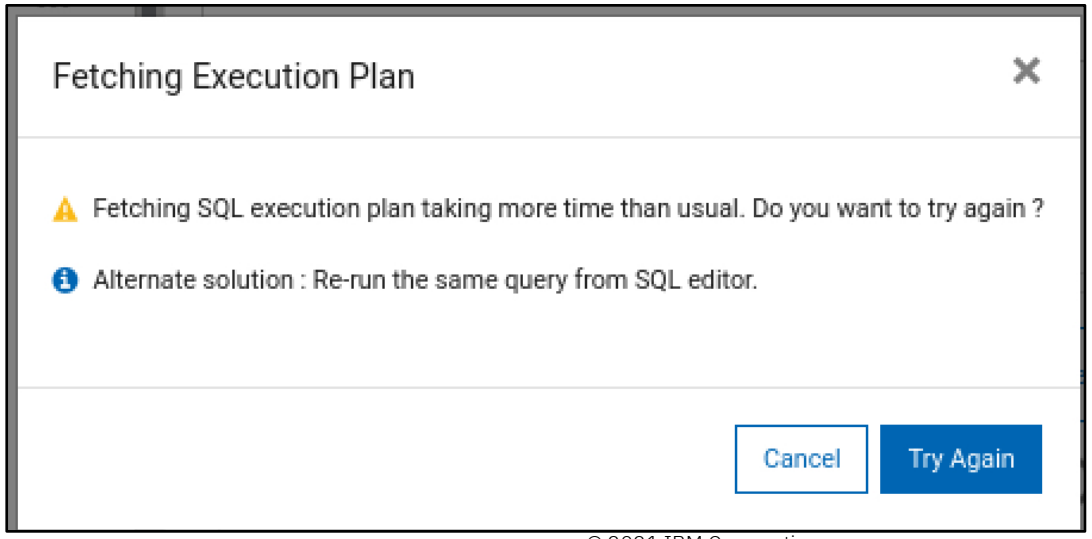

IHQ 1.6.0 - HA cluster

- In earlier versions, IHQ didn't provide a lot of functionality with respect to the HA cluster
  - For example, only basic information was available

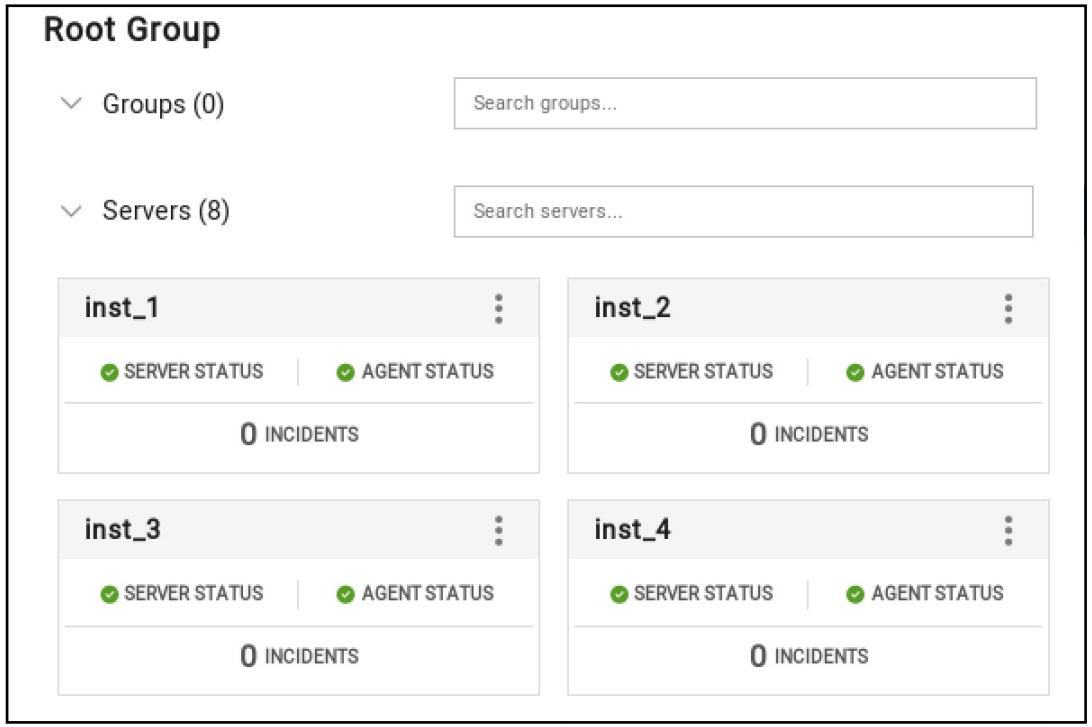

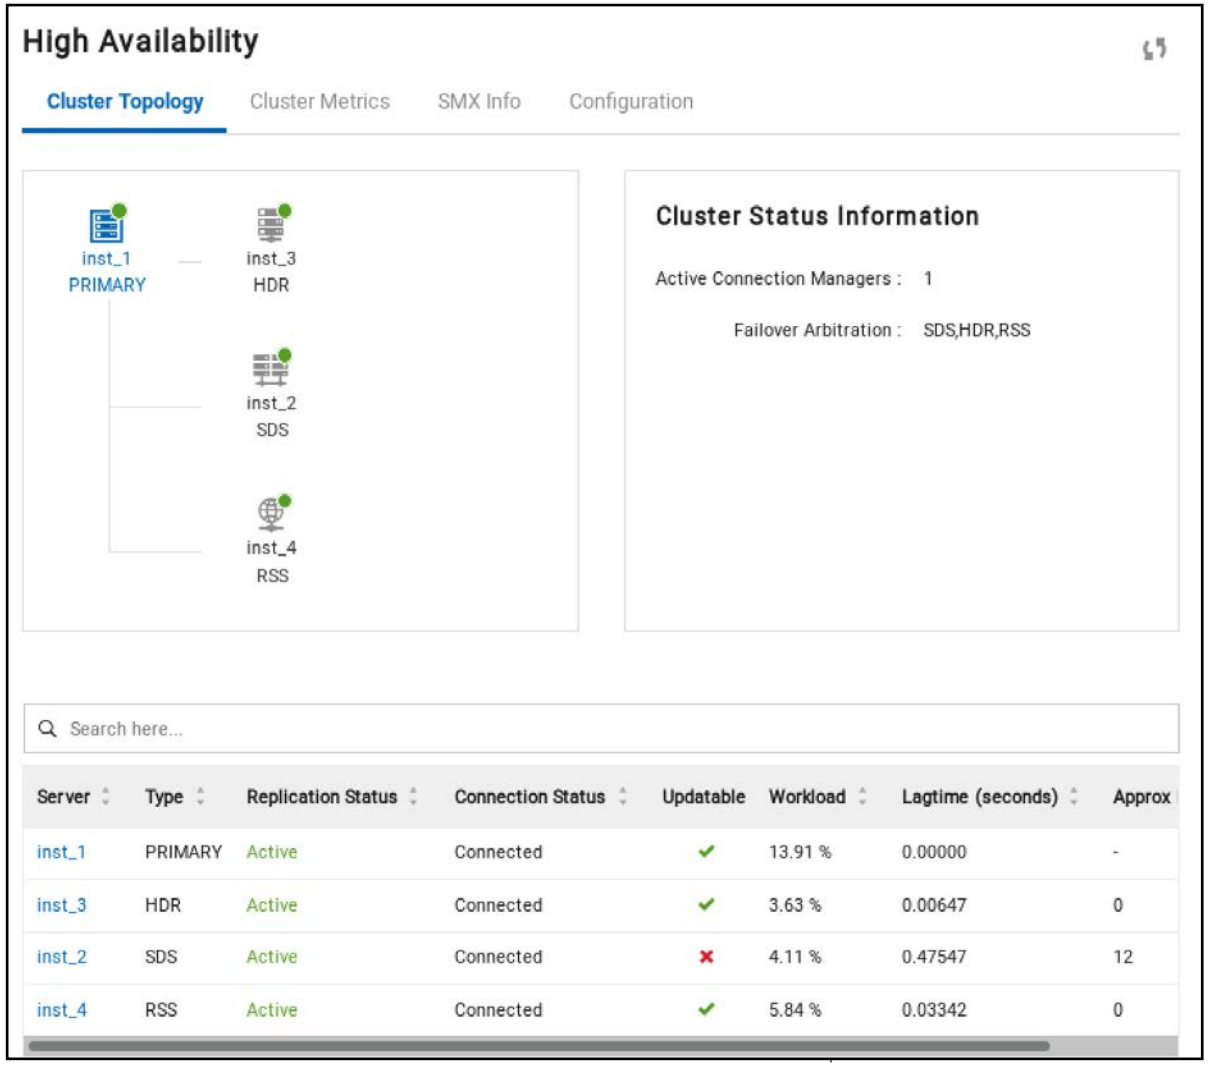

IHQ 1.6.0 - HA cluster

- With this release, it's more intelligent
  - At the highest level, it displays the cluster membership and state

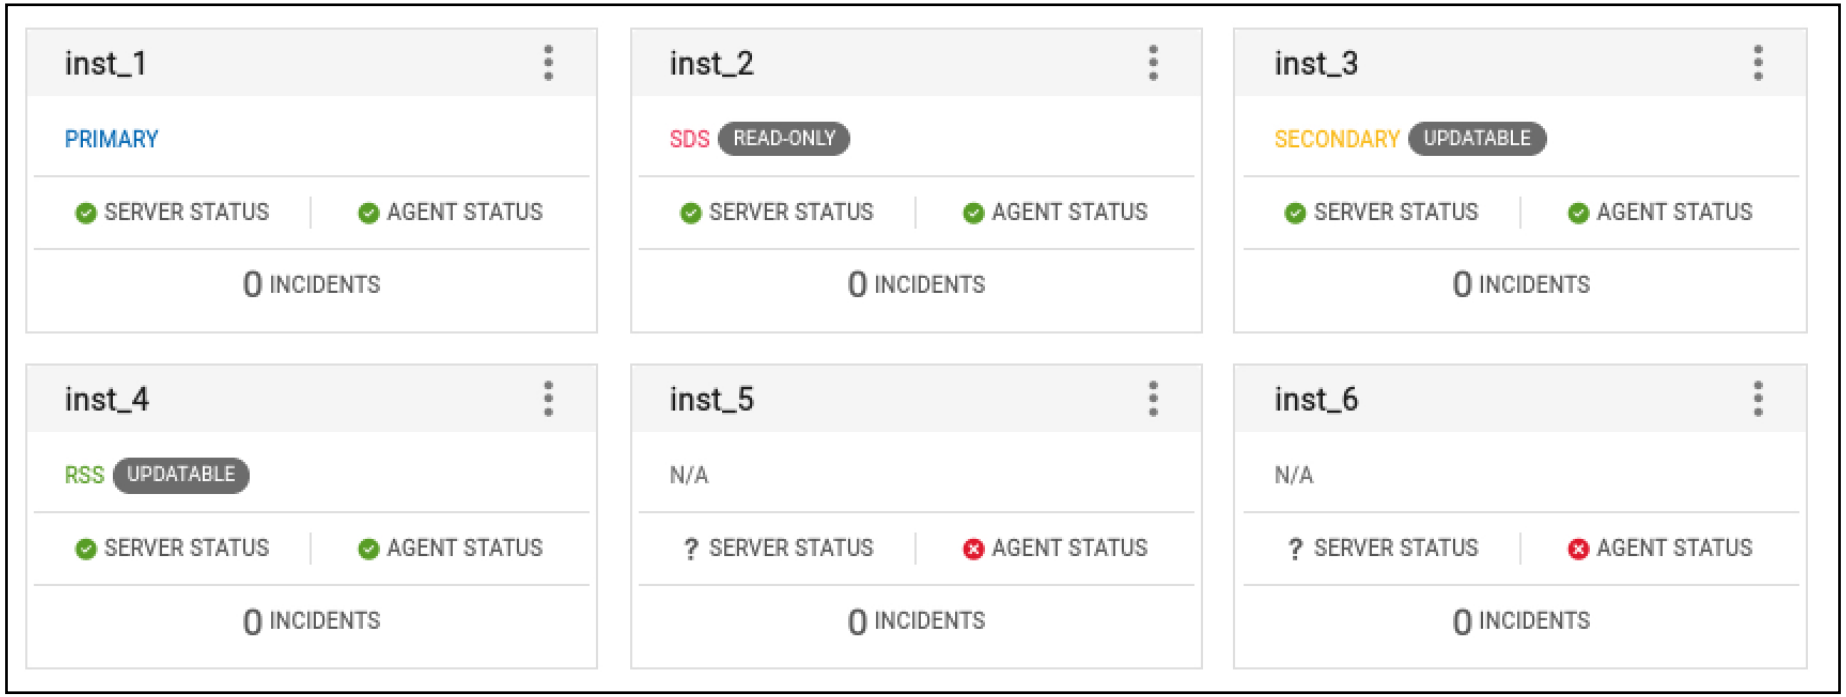

IHQ 1.6.0 - HA cluster

- The **Cluster Topology** page is similar but selecting a secondary displays additional statistical information such as replication latency

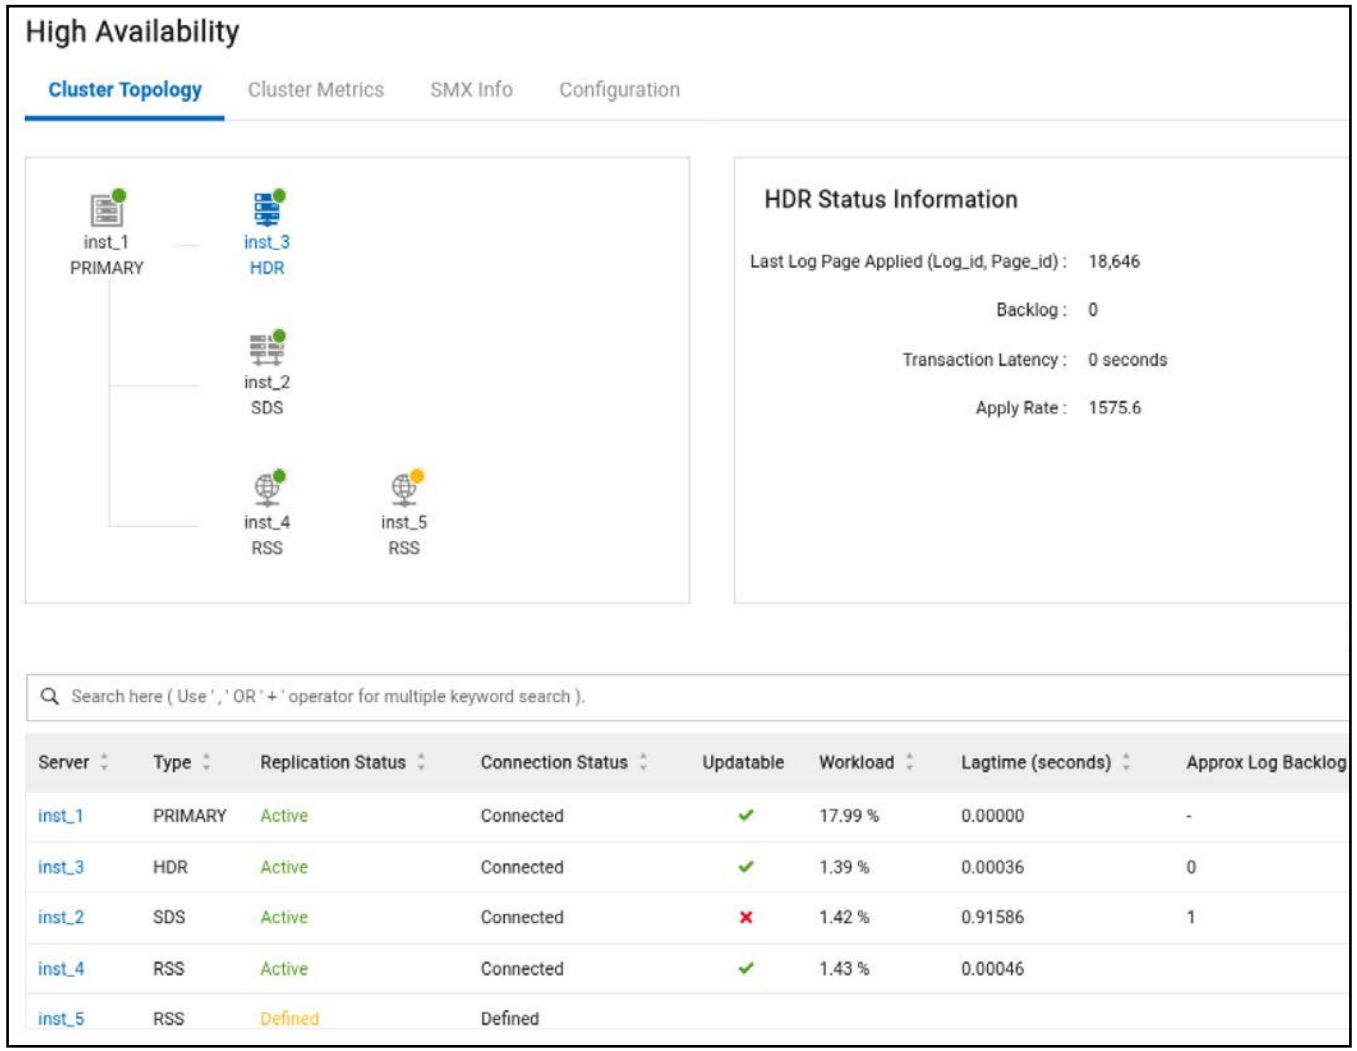

IHQ 1.6.0 - HA cluster

- Cluster metrics and SMX information is available from their respective tabs

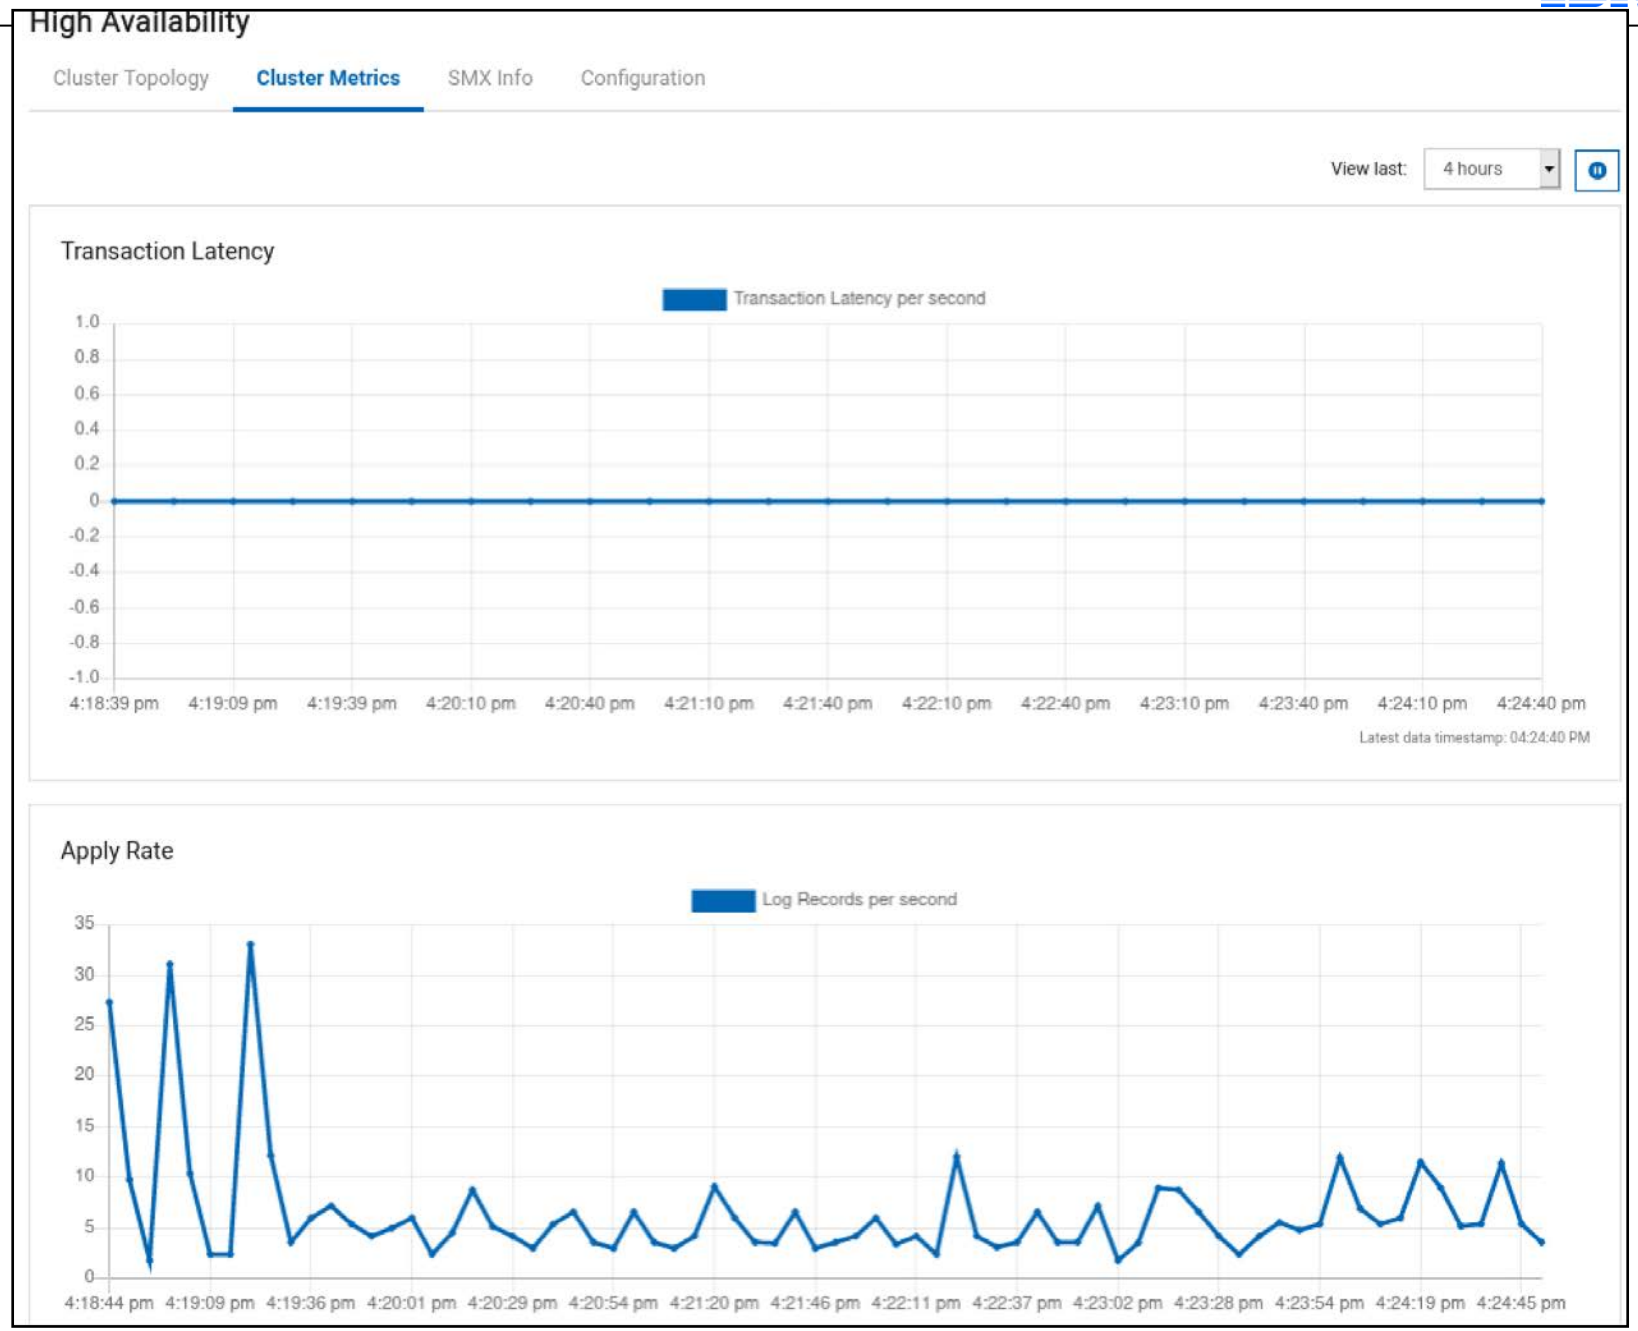

IHQ 1.6.0 - HA cluster

- Cluster metrics and SMX information is available from their respective tabs

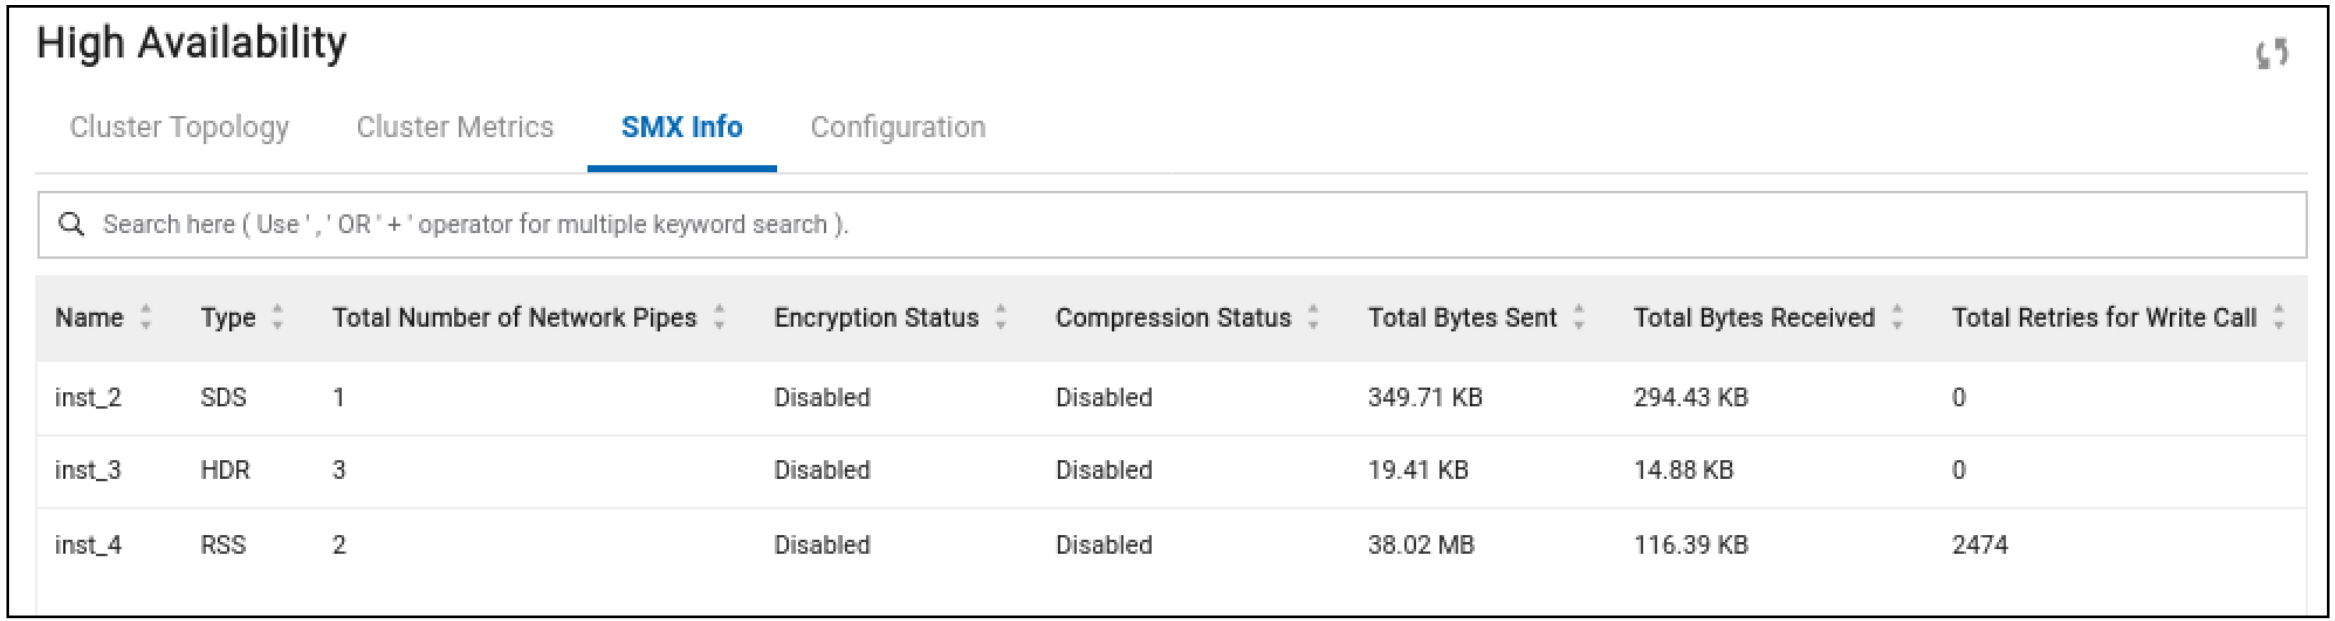

Name	Type	Total Number of Network Pipes	Encryption Status	Compression Status	Total Bytes Sent	Total Bytes Received	Total Retries for Write Call
inst_2	SDS	1	Disabled	Disabled	349.71 KB	294.43 KB	0
inst_3	HDR	3	Disabled	Disabled	19.41 KB	14.88 KB	0
inst_4	RSS	2	Disabled	Disabled	38.02 MB	116.39 KB	2474

IHQ 1.6.0 - HA cluste

- You can review, and in some cases modify, cluster parameters

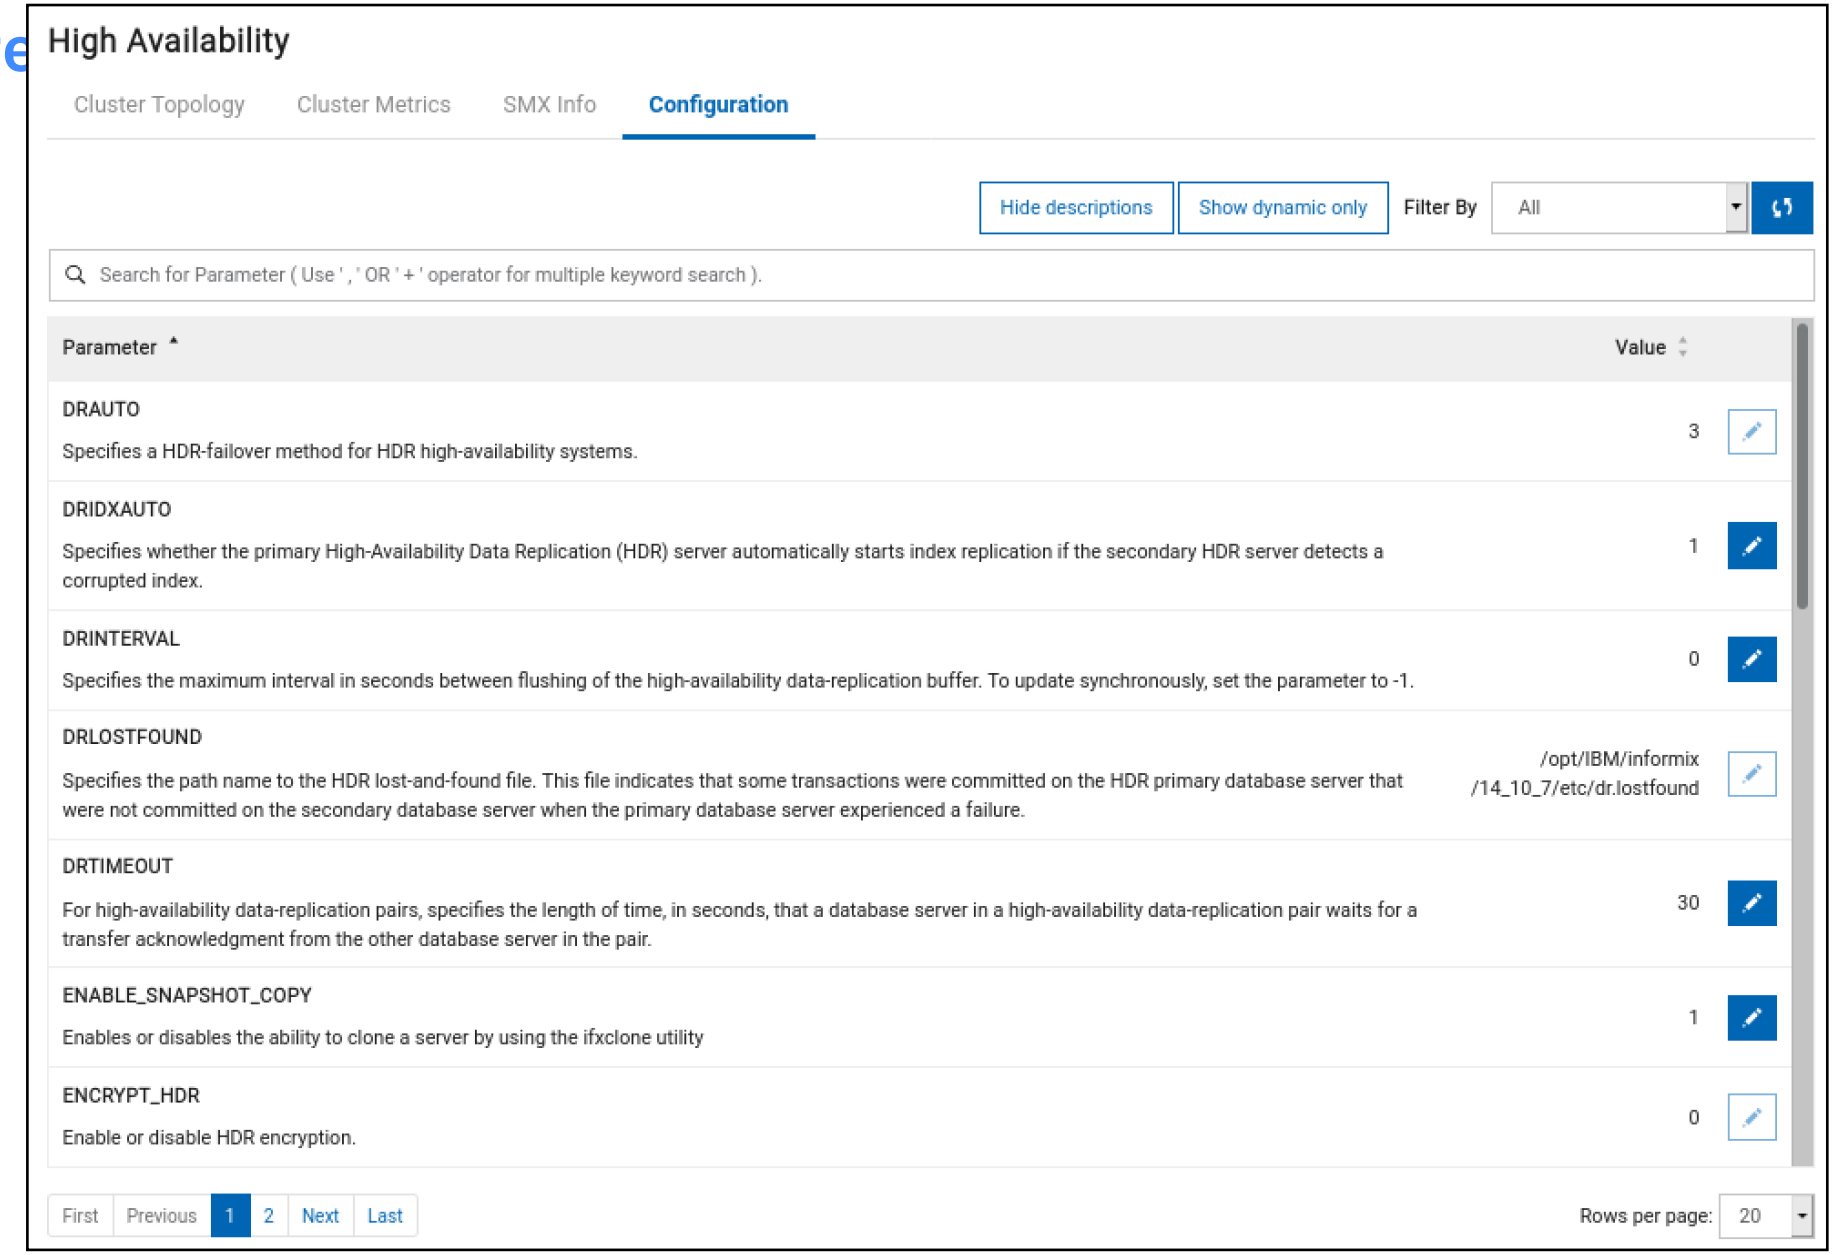

IHQ 1.6.0 - HA cluster

- The tool is “state” aware and displays that information when a node is selected

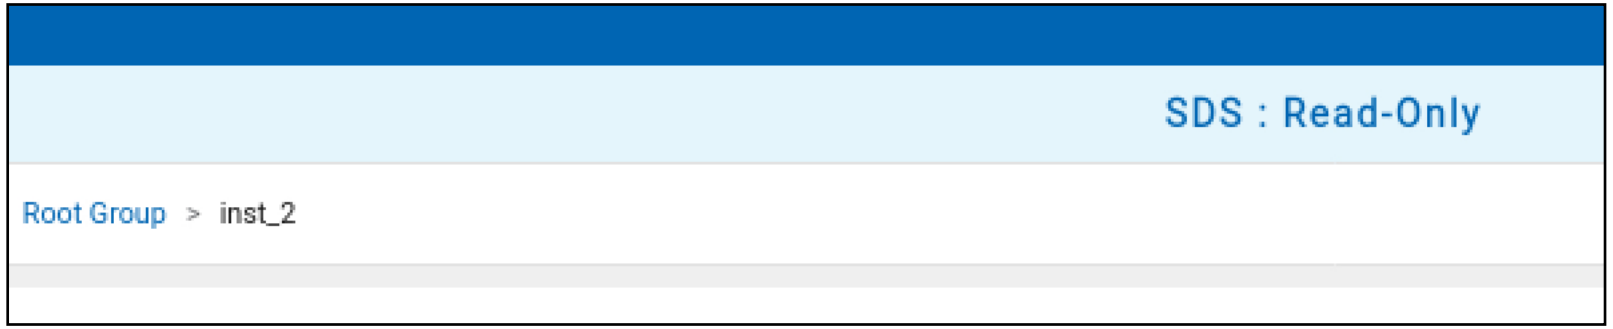

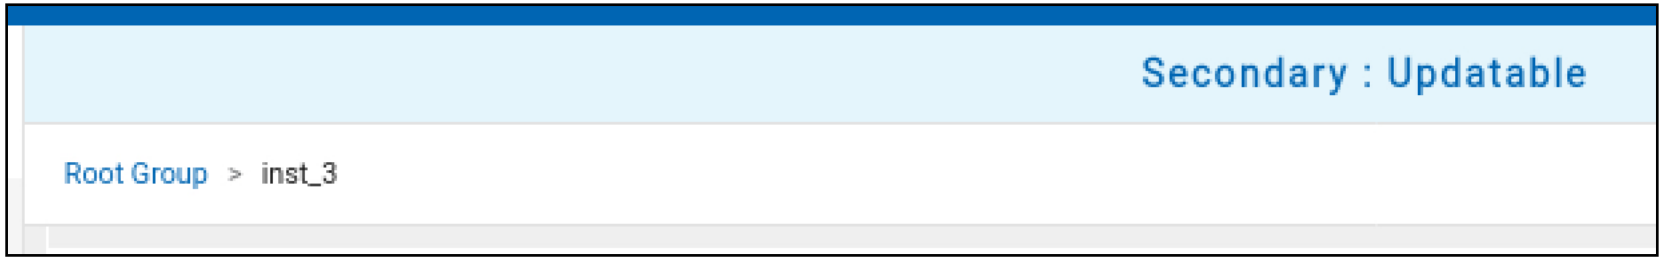

IHQ 1.6.0 - HA cluster

- The tool is “state” aware and will prevent secondaries from executing certain operations if not allowed
  - For example updatable operations on read-only configured nodes even if the IHQ userID has admin privileges

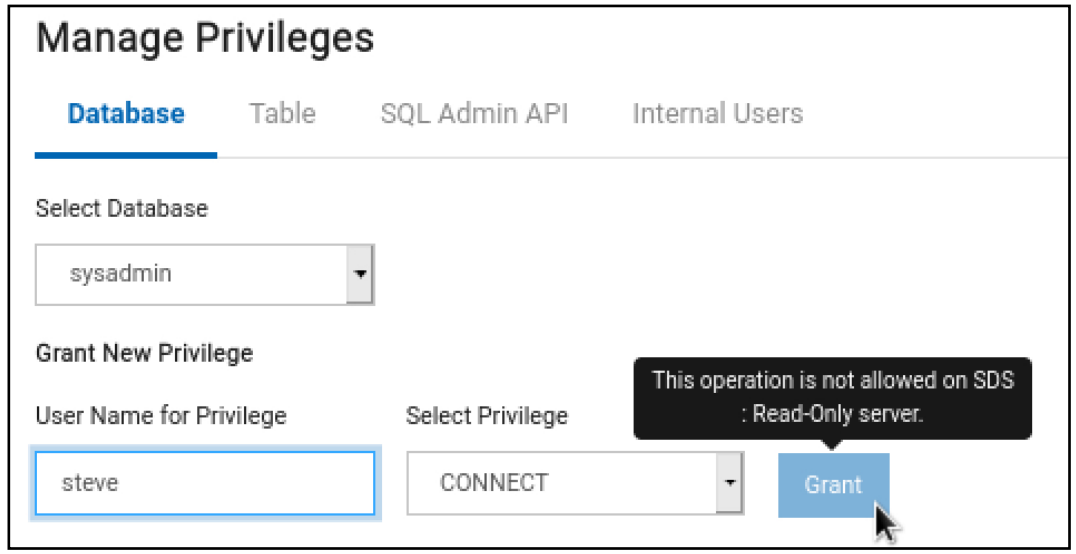

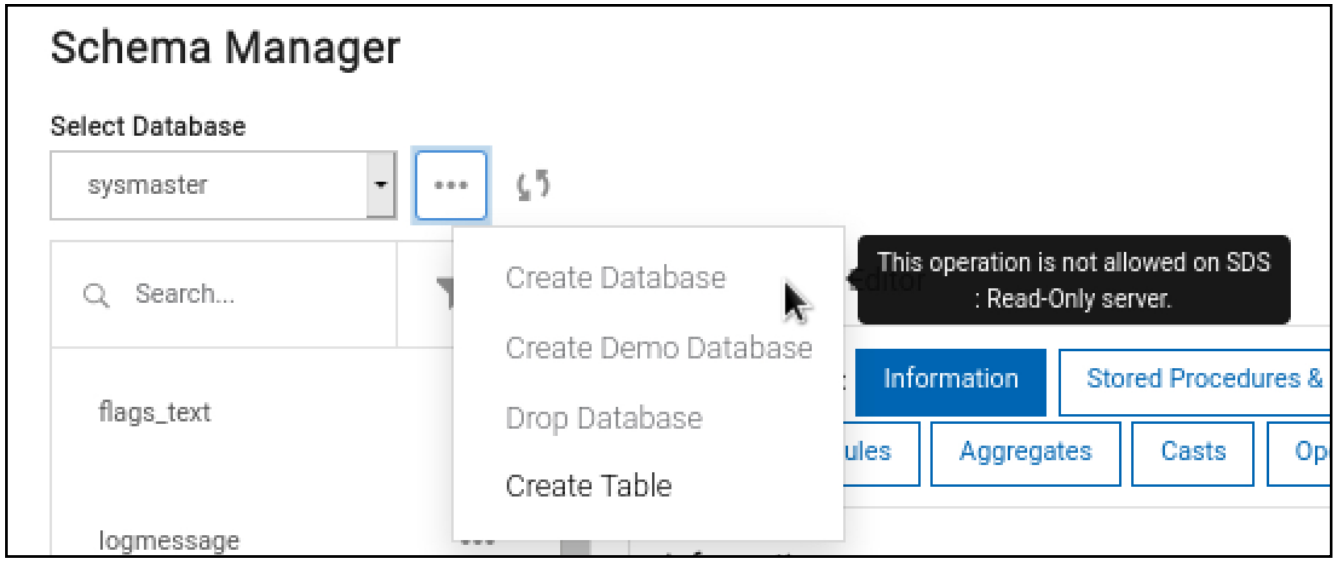

Search within System Repo

- You can now search for keywords in the name or description of reports listed in **System Reports**
  - These are the standard, built-in reports
  - You can now create your own, custom reports

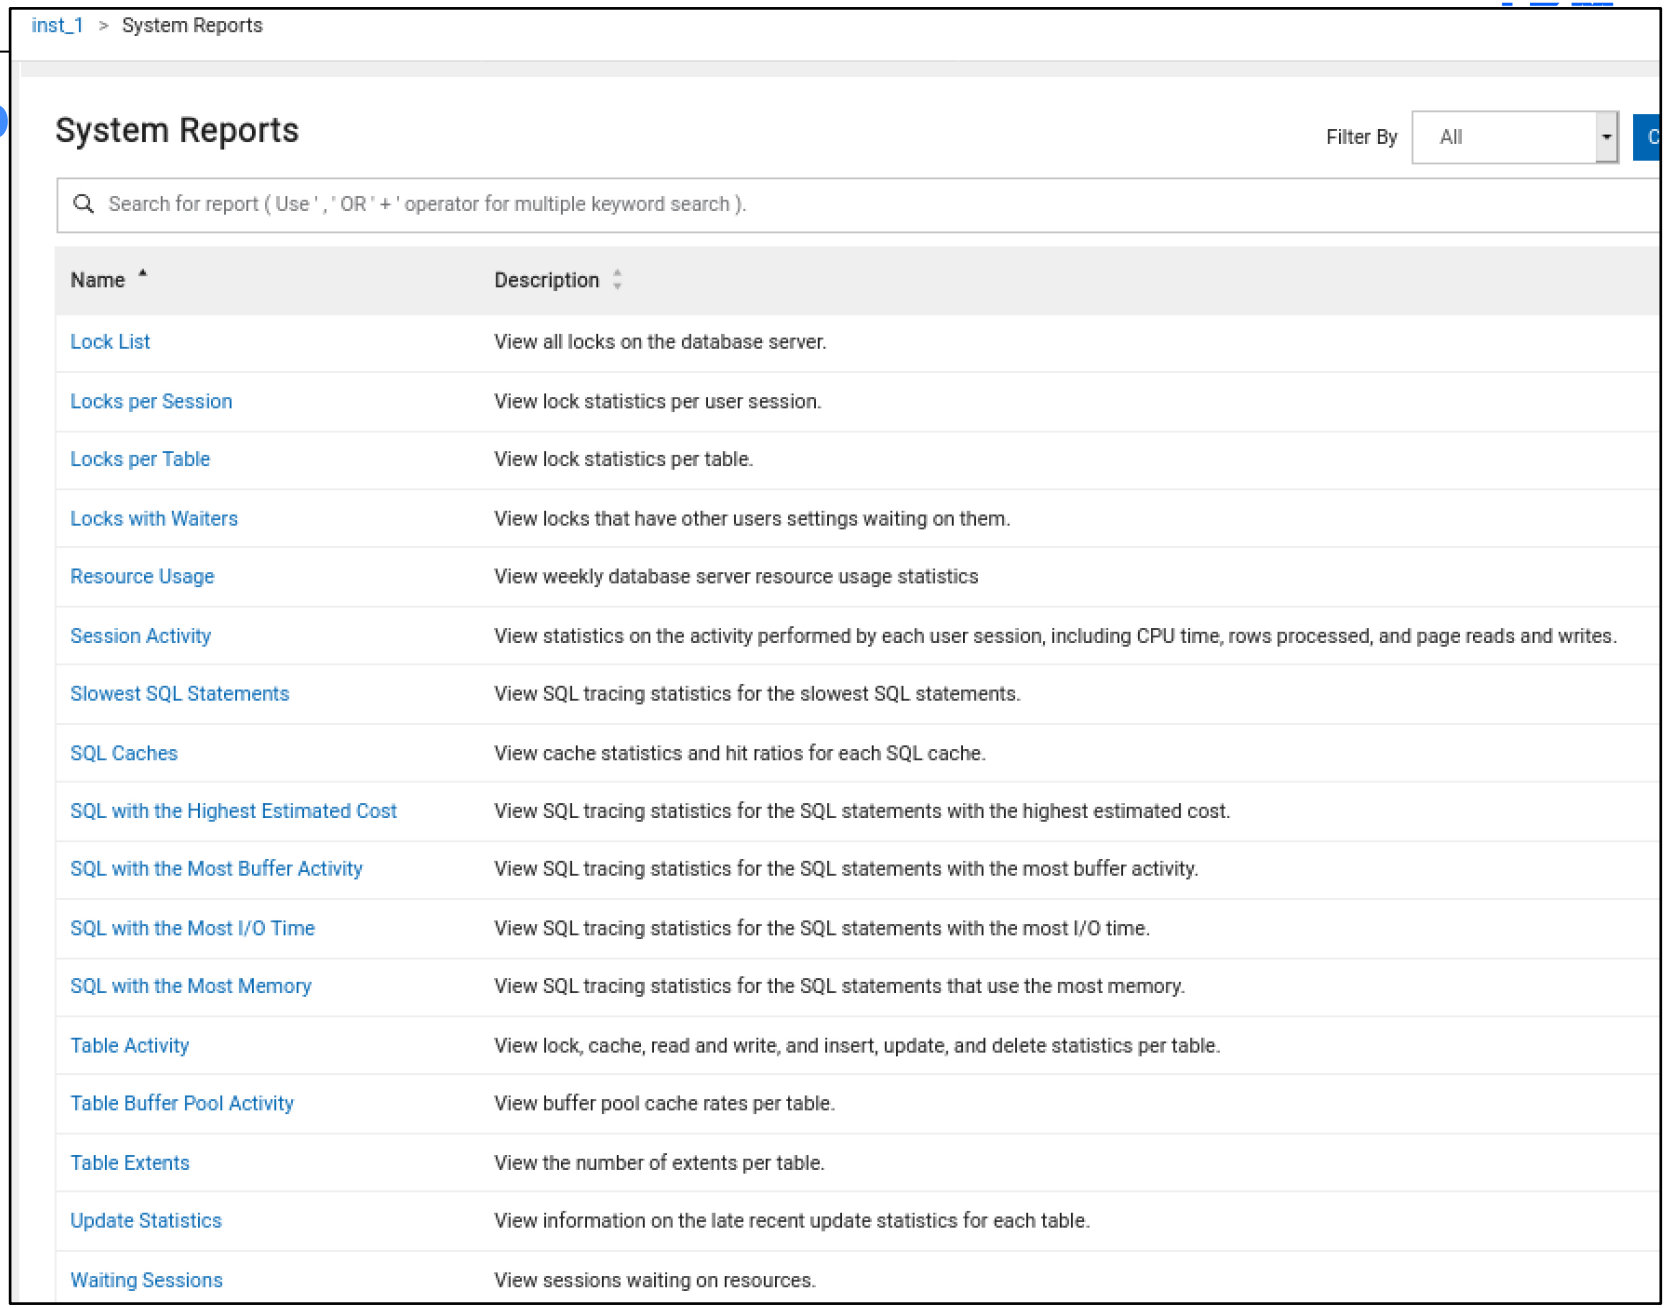

Search within System Reports

- If interested in a particular result type, you can filter by it
  - For example

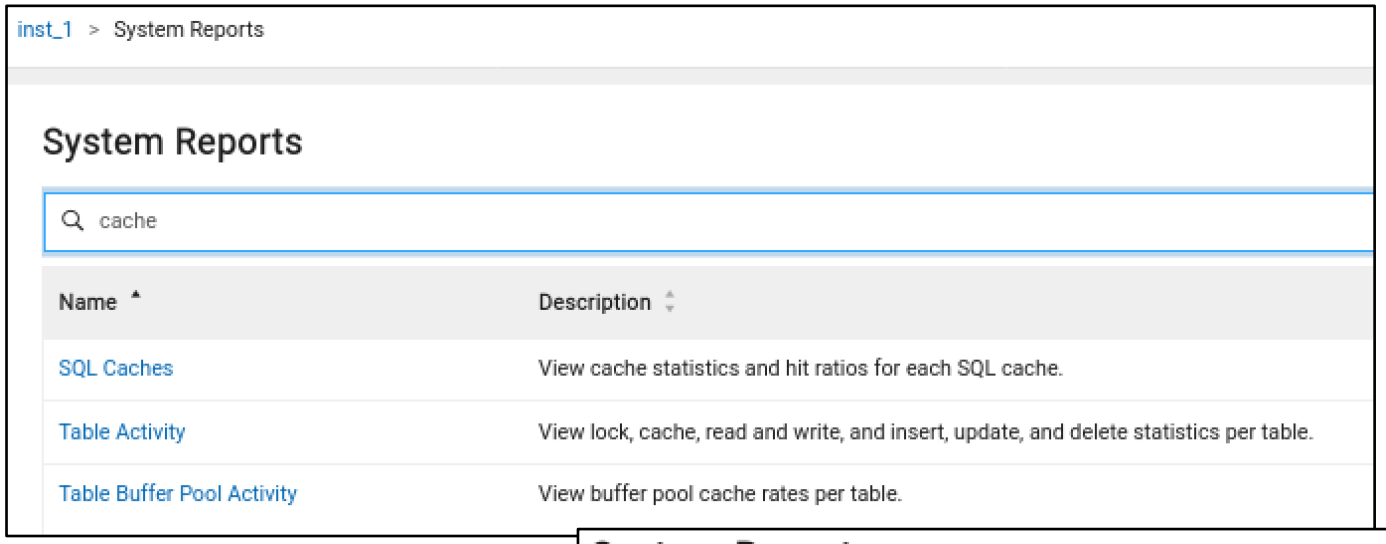

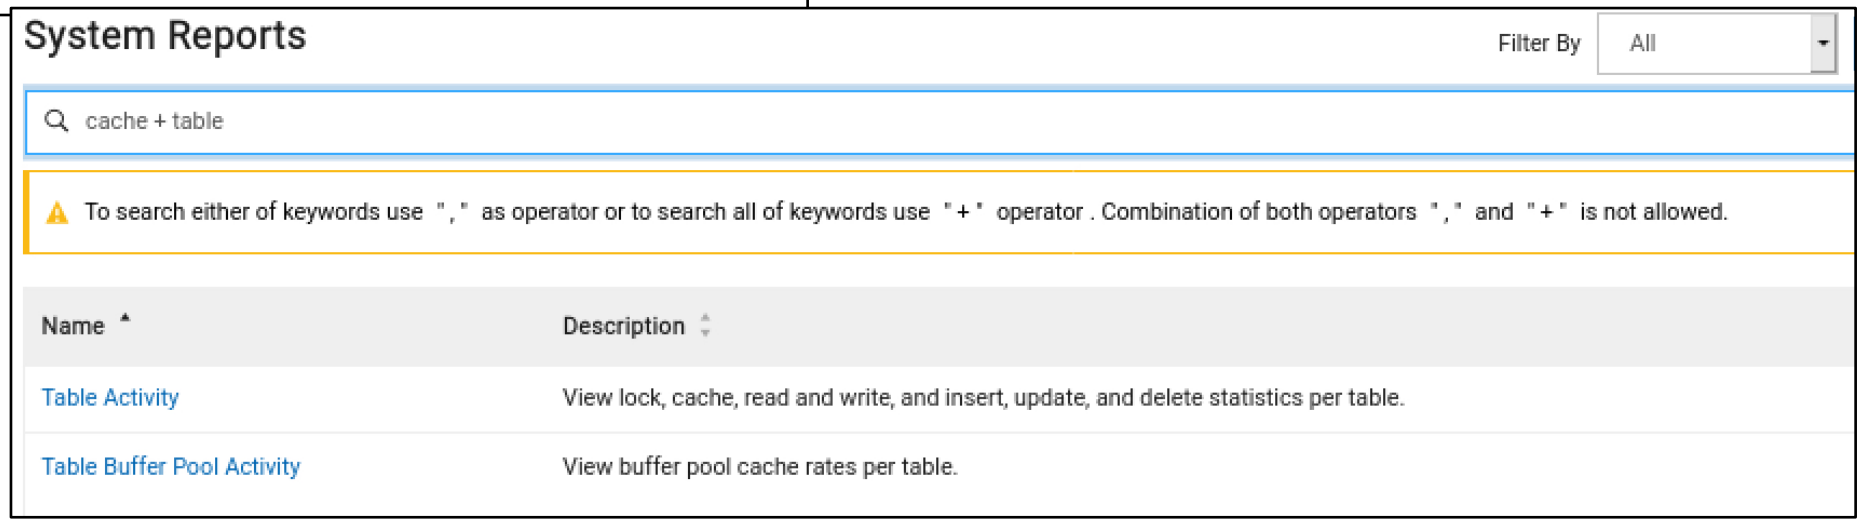

Note: you can not search or filter by the report *results*

For example, trying to find where "dirty" buffers or reads are returned

Expanded IHQ search

- In fact, many IHQ windows now have search capabilities to find keywords within names or descriptions
  - In some cases on some pages, you can filter by the data values, in other cases you can not
    - Experiment! 😁

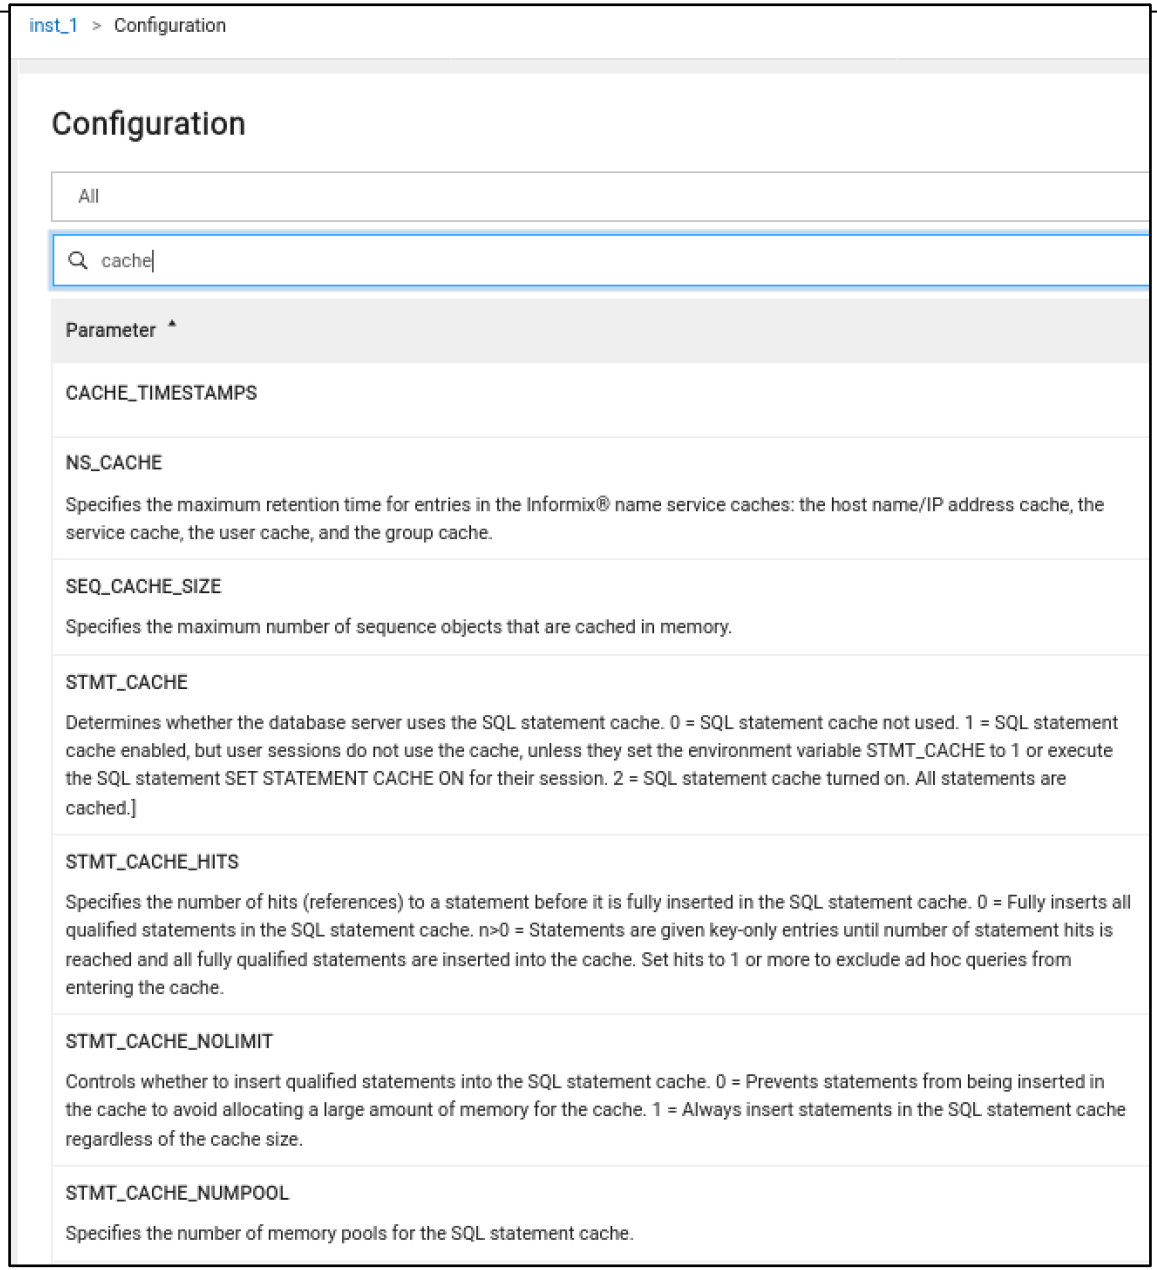

Storage Manager enhancements

- Within the **Storage Manager**, you can now display all tables and indexes in one view even if created outside of the “default” dbspace
  - For example:

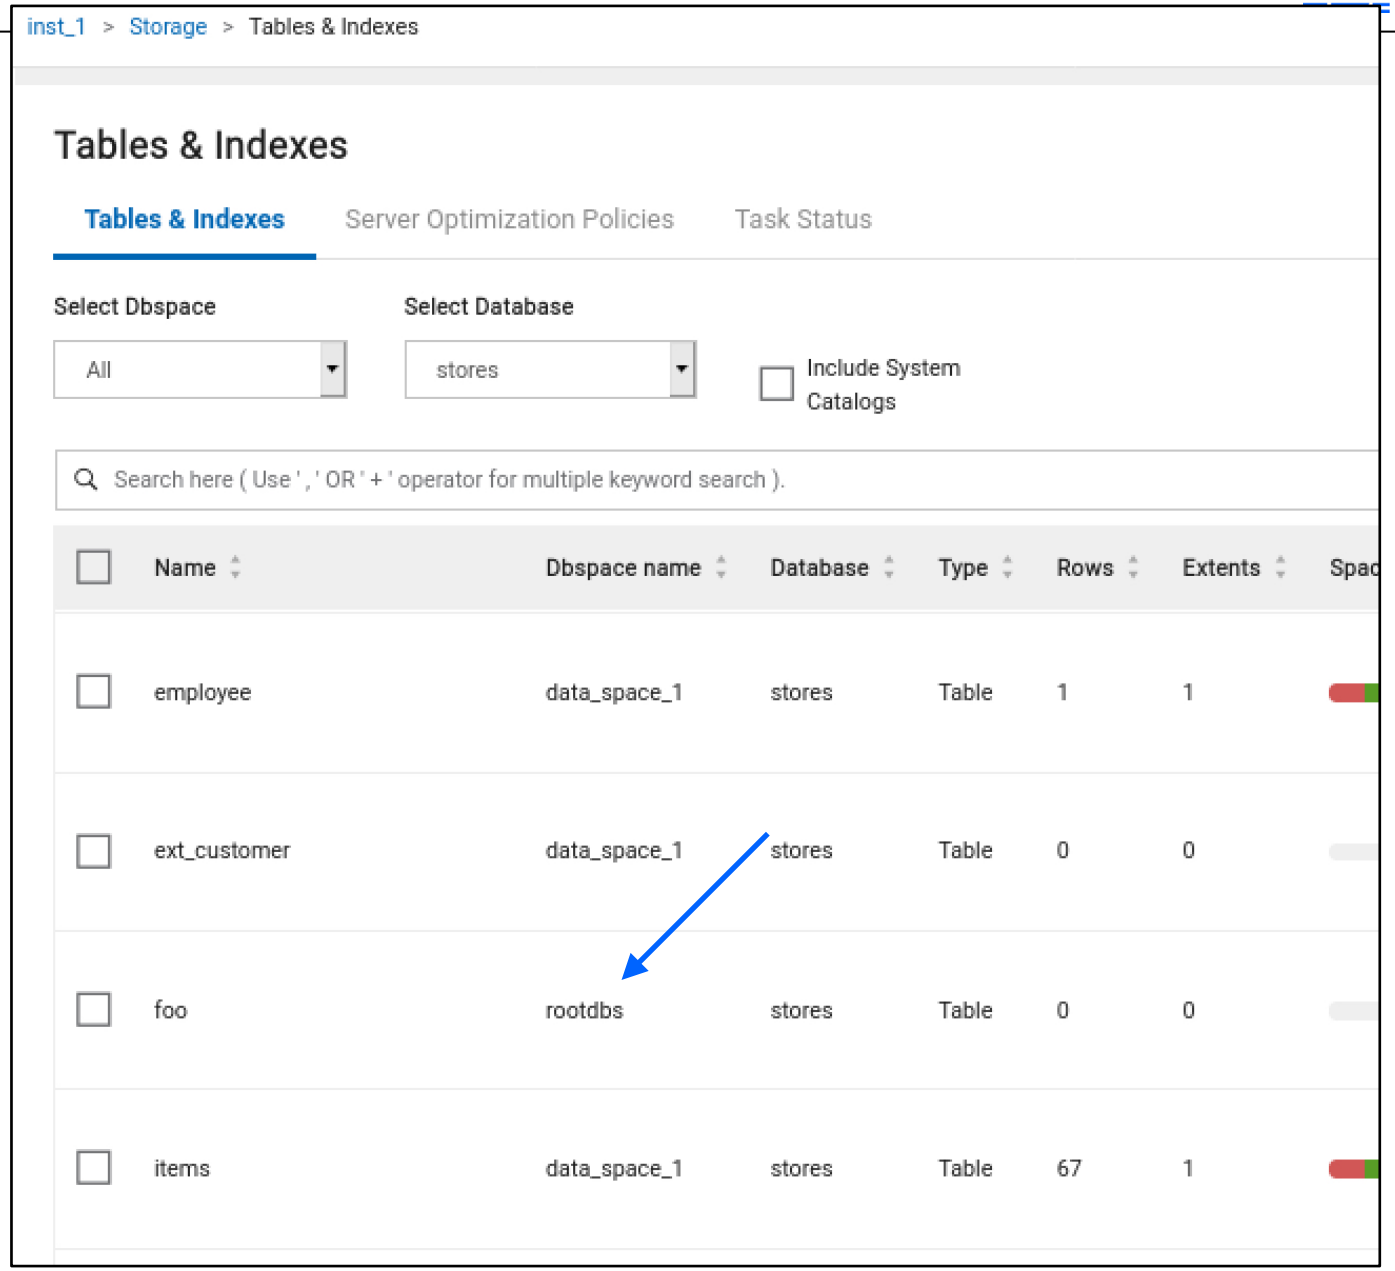

Storage Manager enhancements

- The **Storage Manager** page does allow filtering by values
  - You can use this to isolate fragmented tables or indexes to see where their partitions are
  - Filtering example:

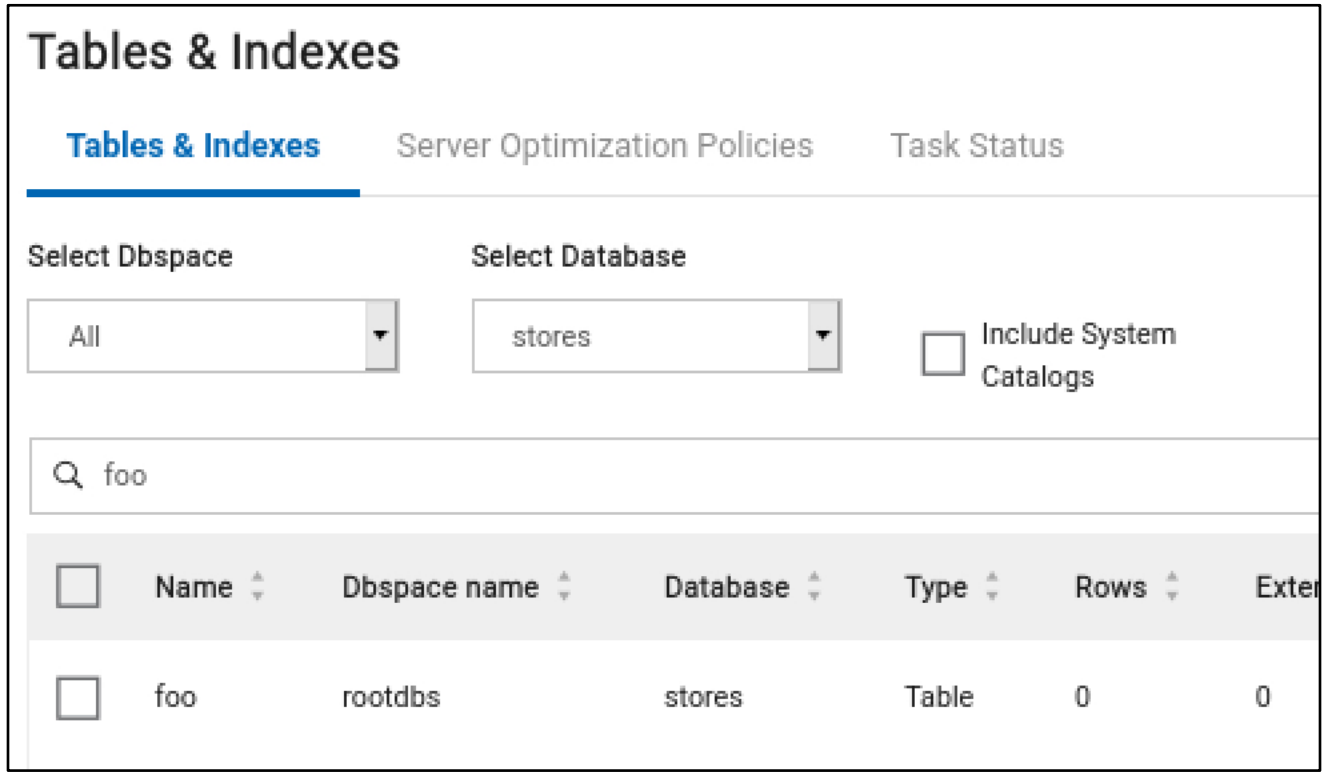

General IHQ stuff

- All open source libraries have been upgraded to the latest available
- White Source scan vulnerabilities have been addressed
- Some defects have been fixed

Questions