

## Consolidated AIX LPAR Performance Reporting Using Topas 25-Jan-2007

AIX5.3 TL5 (August 2007) introduced a new [topas](#) recording and reporting feature that allows a consolidated "system" view of both LPAR and system performance.

### LPAR That Will Be Recording & Reporting:

- Must be running AIX5.3 TL5 (or newer)

**Monitored LPARs:** On the same system/CEC can be running run a mix of:

- AIX5.3 TL5 (or newer)
- AIX5.3 Gold thru TL4 w/Performance Toolbox Agent (PAIDE) v3.1.0.14 (or newer)
- AIX5.2 w/Performance Toolbox Agent (PAIDE) v3.1.0.14 (or newer)

### Start Recording:

```
# /usr/lpp/perfagent/config_topas.sh add
```

The performance data is logged to: [/etc/perf/topas\\_cec.YMMMDD](#)

### Stop Recording:

```
# /usr/lpp/perfagent/config_topas.sh delete
```

Rename the [/etc/perf/topas\\_cec.YMMMDD](#) logfile so that a restart does not corrupt it.

### Summary Report:

```
# topasout -R summary /etc/perf/topas_cec.060816
```

```
#Report: CEC Summary --- hostname: localhost version:1.1
Start:08/16/06 17:24:43 Stop:08/16/06 17:37:43 Int: 5 Min Range: 13 Min
Partition Mon: 2 UnM: 0 Shr: 2 Ded: 0 Cap: 0 UnC: 2
-CEC----- -Processors----- -Memory (GB)-----
Time ShrB DedB Mon UnM Avl UnA Shr Ded PSz APP Mon UnM Avl UnA InU
17:29 0.01 0.00 0.3 0.0 0.0 0 0.3 0 2.0 2.0 3.0 0.0 0.0 0.0 1.4
17:34 0.01 0.00 0.3 0.0 0.0 0 0.3 0 2.0 2.0 3.0 0.0 0.0 0.0 0.9
17:37 0.01 0.00 0.3 0.0 0.0 0 0.3 0 2.0 2.0 3.0 0.0 0.0 0.0 0.9
```

[Check for Updates](#)

Copyright 2007, All Rights Reserved.  
Douglas R. Ranz, [doug\\_ranz@hotmail.com](mailto:doug_ranz@hotmail.com)

**Consolidated AIX LPAR Performance Reporting Using Topas**  
25-Jan-2007

**Detailed Report:**

# [topasout -R detailed /etc/perf/topas\\_cec.060816](#)

```
Time: 17:29:42 -----
Partition Info  Memory (GB)      Processors
Monitored      : 2  Monitored      : 3.0  Monitored      : 0.3  Shr Physcl Busy: 0.01
UnMonitored    : 0  UnMonitored    : 0.0  UnMonitored    : 0.0  Ded Physcl Busy: 0.00
Shared         : 2  Available      : 0.0  Available      : 0.0
Dedicated      : 0  UnAllocated    : 0.0  Unallocated    : 0.0  Hypervisor
Capped         : 0  Consumed       : 1.4  Shared         : 0.3  Virt Cntxt Swtch: 808
UnCapped       : 2                                Dedicated      : 0.0  Phantom Intrpt  : 0
                                                    Pool Size     : 2.0
                                                    Avail Pool    : 2.0

Host           OS  M  Mem  InU  Lp  Us  Sy  Wa  Id  PhysB  Ent  %EntC  Vcsw  PhI
-----shared-----
gilmore        A53 U  2.0  0.9  4   1   1   0  97  0.01  0.2  4.46  482   0
localhost      A53 U  1.0  0.4  2   0   2   0  97  0.00  0.1  4.48  325   0
-----dedicated-----
```

**Known Bugs:**

@ AIX5.3 TL5:

- When you start recording; a [nohup.out](#) file will be written in the current working directory and it will fill with the message: [Error in recording = 9](#)
- Multiple starts and stops of the recording subsystem can create logfile problem. When you stop recording; rename the [/etc/perf/topas\\_cec.YYMMDD](#) logfile before restarting recording.

[Check for Updates](#)

Copyright 2007, All Rights Reserved.  
Douglas R. Ranz, [doug\\_ranz@hotmail.com](mailto:doug_ranz@hotmail.com)

**Consolidated AIX LPAR Performance Reporting Using Topas**  
25-Jan-2007

**README.perfagent.tools**

```
* IBM_PROLOG_BEGIN_TAG
* This is an automatically generated prolog.
*
* perf53H src/perf/perfagent/usr/lpp/perfagent/README.perfagent.tools 1.7.1.4
*
*
* OBJECT CODE ONLY SOURCE MATERIALS
*
* (C) COPYRIGHT International Business Machines Corp. 1993,2006
* All Rights Reserved
*
* The source code for this program is not published or otherwise
* divested of its trade secrets, irrespective of what has been
* deposited with the U.S. Copyright Office.
*
* IBM_PROLOG_END_TAG
*
* COMPONENT_NAME: (CMDPERFT) AIX Performance Analysis Tools
*
* FUNCTIONS: README
*
* ORIGINS: 30
*
* (C) COPYRIGHT International Business Machines Corp. 2002
*
* All Rights Reserved.
* Licensed Materials - Property of IBM
*
* US Government Users Restricted Rights - Use, duplication or
* disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
*
```

=====

The information contained in this file consists of a description of a set of performance tools which may be used to evaluate and tune AIX system performance. For more detailed information on each of the tools, type "man 'toolname'".

In AIX Version 5, most of the performance tools were moved to the bos.perf.tools fileset.

=====

## Consolidated AIX LPAR Performance Reporting Using Topas

25-Jan-2007

### DESCRIPTION OF CONTENTS

#### Topas Cross-Partition View and Recording (-C | -R)

This panel displays metrics similar to the `lparstat` command for all the AIX partitions it can identify as belonging to the same hardware platform. Dedicated and shared partitions are displayed in separate sections with appropriate metrics. The top section represents aggregated data from the partition set to show overall partition, memory, and processor activity.

Remote enablement for this panel to collect from other partitions requires that latest updates to the `perfagent.tools` and `bos.perf.tools` used to support this function. For earlier versions of AIX, the `topas` command also collects remote data from partitions that have the Performance Aide product (`perfagent.server`) installed. `topas -C` may not be able to locate partitions residing on other subnets. To circumvent this, create a `$HOME/Rsi.hosts` file containing the fully qualified hostnames for each partition (including domains), one host per line.

Besides the display option, the `topas -R` option will run `topas` in the background (with no display) and record the Cross-Partition data. Recordings are made to the `/etc/perf/` directory, and are of the form `topas_cec.YYMMDD`. Recordings cover single-day periods (24 hours) and are retained for 8 days before being automatically deleted. This allows a weeks worth of data to be retained on the system at all times. The `topasout` command is used to covert these recordings into various text-based reports.

Customers using the `topas -R` function on an LPAR-capable system should identify a single partition for collecting and recording this data. The `/usr/lpp/perfagent/config_topas.sh` script can be used to automatically update the system `/etc/inittab` file to execute `topas -R` on startup.

The following metrics display in the initial Cross-Partition panel. Additional metrics with full descriptive labels can be displayed by using the key toggles identified in the Additional Cross-Partition Panel Subcommands section in the AIX commands manual pages:

#### Partition totals

|     |                                |
|-----|--------------------------------|
| Shr | Number of shared partitions    |
| Ded | Number of dedicated partitions |

#### Memory (in GB)

|       |                                       |
|-------|---------------------------------------|
| Mon   | Monitored partitions total memory     |
| Avl   | Memory available to partition set     |
| InUse | Memory in use on monitored partitions |

[Check for Updates](#)

Copyright 2007, All Rights Reserved.  
Douglas R. Ranz, [doug\\_ranz@hotmail.com](mailto:doug_ranz@hotmail.com)

**Consolidated AIX LPAR Performance Reporting Using Topas**  
25-Jan-2007

Processor

|           |  |
|-----------|--|
| Shr       | Number of shared processors  |
| Ded       | Number of dedicated processors   |
| PSz       | Active physical CPUs in the shared processor pool<br>being used by this LPAR |
| APP       | Available physical processors in the shared pool                             |
| Shr_PhysB | Shared Physical Busy   |
| Ded_PhysB | Dedicated Physical Busy  |

Individual partition data

|      |                        |
|------|------------------------|
| Host | Hostname               |
| OS   | Operating system level |
| M    | Mode                   |

For shared partitions

- C - SMT enabled & capped
- c - SMT disabled & capped
- U - SMT enabled & uncapped
- u - SMT disabled & uncapped

For dedicated partitions

- S - SMT enabled
- ' ' (blank) - SMT disabled

|       |  |
|-------|--|
| Mem   | Total memory in GB.  |
| InU   | Memory in use in GB.   |
| Lp    | Number of logical processors                                 |
| Us    | Percentage of CPU used by programs executing in user mode.   |
| Sy    | Percentage of CPU used by programs executing in kernel mode. |
| Wa    | Percentage of time spent waiting for IO.                     |
| Id    | Percentage of time the CPU(s) is idle.                       |
| PhysB | Physical Processors Busy                                     |
| Ent   | Entitlement granted (shared-only)                            |
| %Entc | Percent Entitlement consumed (shared-only)                   |
| Vcsw  | Virtual context switches average per second (shared-only)    |
| PhI   | Phantom interrupts average per second (shared-only)          |

Note: At least one partition to be monitored must have Pool Utilization Authority (PUA) configured for pool information metrics to be collected.

For cross-partition monitoring/recording, some global data is not available from any partition. The -o options allow users to specify those fields in the command line. Optionally, users can configure a system to allow topas to query the HMC directly for this information. This requires the following steps:

1. Install OpenSSH at the partition.
2. Enable remote command support on the HMC for user hscroot to allow ssh connections to be opened from the partition.

[Check for Updates](#)

Copyright 2007, All Rights Reserved.  
Douglas R. Ranz, [doug\\_ranz@hotmail.com](mailto:doug_ranz@hotmail.com)

## Consolidated AIX LPAR Performance Reporting Using Topas

25-Jan-2007

3. Configure ssh on the HMC to not require a password for the HMC user hscroot when queried from the selected partition. This requires the `.ssh/authorized_keys2` on the HMC for user login hscroot.
4. Try "ssh -l hscroot <hmc address> date" from the partition to check whether the date is displayed without entering password.
5. Utilize the topas -o options described in the usage table to specify the managed system and HMC names when executing topas.

Note: This functionality is currently available only for HMC version 5 and above, and should only be enabled after careful consideration of any security implications.

For customers who have configured the topas hosting partition to query the noHMC (see above), these options are used to specify the managed system and a HMC hostname:

- o managedsys = [managed system name under which this partition is configured available via: `lsrsrc -dx IBM.ManagementServer Hostname`]
- o hmc = [HMC name under which this partition is configured]

The HMC name as a parameter is optional and can be used when the resource class `IBM.ManagementServer` is not configured in RMC. It is preferred that the IP address of the HMC be provided instead of the host name. The managed system name needs to be specified as a parameter for topas for the query.

If the HMC platform cannot be configured for automatic queries, then the global data fields not available to the local partition may be set via these options:

- o availmem = [total amount of memory allocated to all partitions, in GB]
- o unavailmem = [total amount of memory unallocated from the HMC, in GB]
- o availcpu = [total number of physical processors allocated for all LPARs]
- o unavailcpu = [total number of physical processors unallocated from the HMC]
- o poolsize = [defined Pool Size required if HMC Processor Utilization Authority restricts access]
- o partitions = [number of partitions defined on the HMC]
- o reconfig = [number of seconds between checking for HMC configuration changes. Allowed values are 30, 60, 90, 120, 180, 240, 300 seconds. The default is 300 seconds]

Topasout formatted reports

[Check for Updates](#)

Copyright 2007, All Rights Reserved.  
Douglas R. Ranz, [doug\\_ranz@hotmail.com](mailto:doug_ranz@hotmail.com)

## Consolidated AIX LPAR Performance Reporting Using Topas

25-Jan-2007

Post-process topas CEC or xmwlm local recordings to generate a text report. The -s option should be used if a standard spreadsheet type is desired, no -R argument is required. Additional options are:

For local recordings generated by xmwlm -L:

type = detailed [prints a formatted output report similar to the topas main display]

type = summary [prints a single-line formatted output report of the system]

For recordings generated by topas -R:

type = detailed [prints a formatted output report similar to the topas CEC monitoring display]

type = summary [prints the CEC processor and memory summary report]

-i MM [split the recording reports into equal size time periods. This must be a multiple of the recording interval, which by default is 5 minutes. Allowed values are 5, 10, 15, 30, 60]

-b HHMM [begin time in hours (HH) and minutes (MM). This must be a range between 0000 and 2400]

-e HHMM [end time in hours (HH) and minutes (MM). This must be a time greater than the begin time]

If -b and -e options are not used, the daily recording will be processed from beginning to end.

[Check for Updates](#)

Copyright 2007, All Rights Reserved.  
Douglas R. Ranz, [doug\\_ranz@hotmail.com](mailto:doug_ranz@hotmail.com)