

IBM iDoctor for IBM i

How to collect CPU, memory from multiple LPARs

IBM iDoctor for IBM i Development Team

20 January 2022

Licensed Materials - Property of IBM

© Copyright International Business Machines Corporation 2022. All rights reserved.

Question:

I have the performance data of 5 IBM i LPAR's, and I want to consolidate into graphs showing CPU and memory utilization for each. Is that possible?

Answer:

Yes, in Collection Services only.

When using multiple partitions, it can be important to understand the overall utilization of processing capability, across all partitions, regardless of whether the partition is running IBM® i, AIX®, or Linux®. IBM i provides a way to collect and display this data.

Collecting the data

To collect CPU utilization for the physical system, you must meet the following configuration requirements:

- POWER6® hardware with IBM i 6.1 and a firmware level of xx340_061 or later.
- Enabled the collection of performance data for the partition you want it collected on. You only need to collect this data on one partition, and this partition must be an IBM i partition. The CPU utilization information that is collected will reflect work done in partitions that are running AIX and Linux as well as IBM i, but AIX and Linux do not support collecting this data.

Enabling the collection of this performance data requires the setting of a configuration parameter on the HMC or Integrated Virtualization Manager (IVM). On the HMC, there is an "Allow performance information collection" checkbox on the processor configuration tab. Select this checkbox on the IBM i partition that you want to collect this data. If you are using IVM, you use the `chgsyscfg` command, specifying the `allow_perf_collection` (permission for the partition to retrieve shared processor pool utilization) parameter. Valid values for the parameter are 0, do not allow authority (the default) and 1, allow authority.

Once the performance data collection support is enabled, Collection Services will collect this additional information. At each collection interval, Collection Services will collect partition configuration and utilization information from the hypervisor. The data is stored in the Collection Services database file QAPMLPARH. You also have the ability to get physical processor utilization in the Collection Services database file QAPMSYSPRC.

You only need to configure and collect this data on one partition. Sometimes the QAPMLPARH data is not fully available until after collection services restarts so if you wish to view the data as soon as possible then you will need to run: `ENDPFCOL FRCCOLEND(*YES)` then `STRPFCOL CYCCOL(*YES)`, wait about 15 minutes and verify data exists in the QAPMLPARH file.

For more information on capturing data:

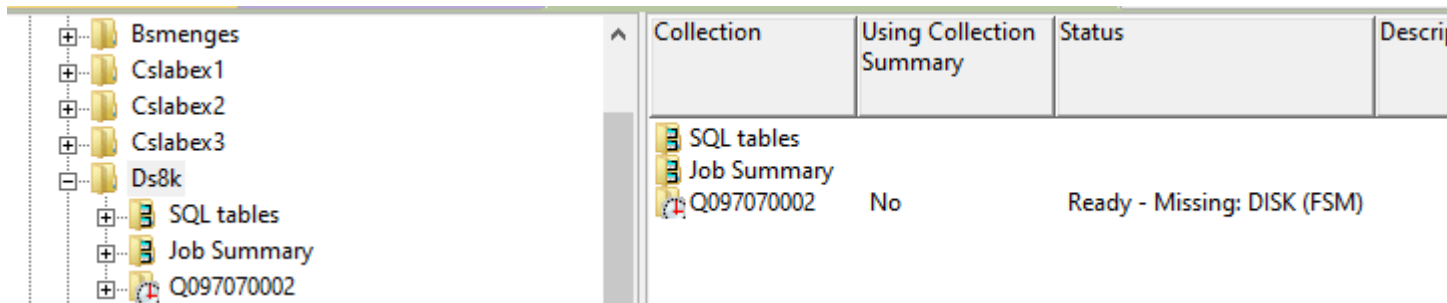
<https://www.ibm.com/docs/en/i/7.4?topic=services-collecting-displaying-cpu-utilization-all-partitions>

Section 14.10 of the Performance FAQ document: <https://www.ibm.com/downloads/cas/QWXA9XKN>

Displaying the data

Once the system has been configured to capture QAPMLPARH and QAPMSYSPRC data then iDoctor should be able to see it and display it.

The data library should list collections with a status of something like this (assuming the build is 1508 or higher):

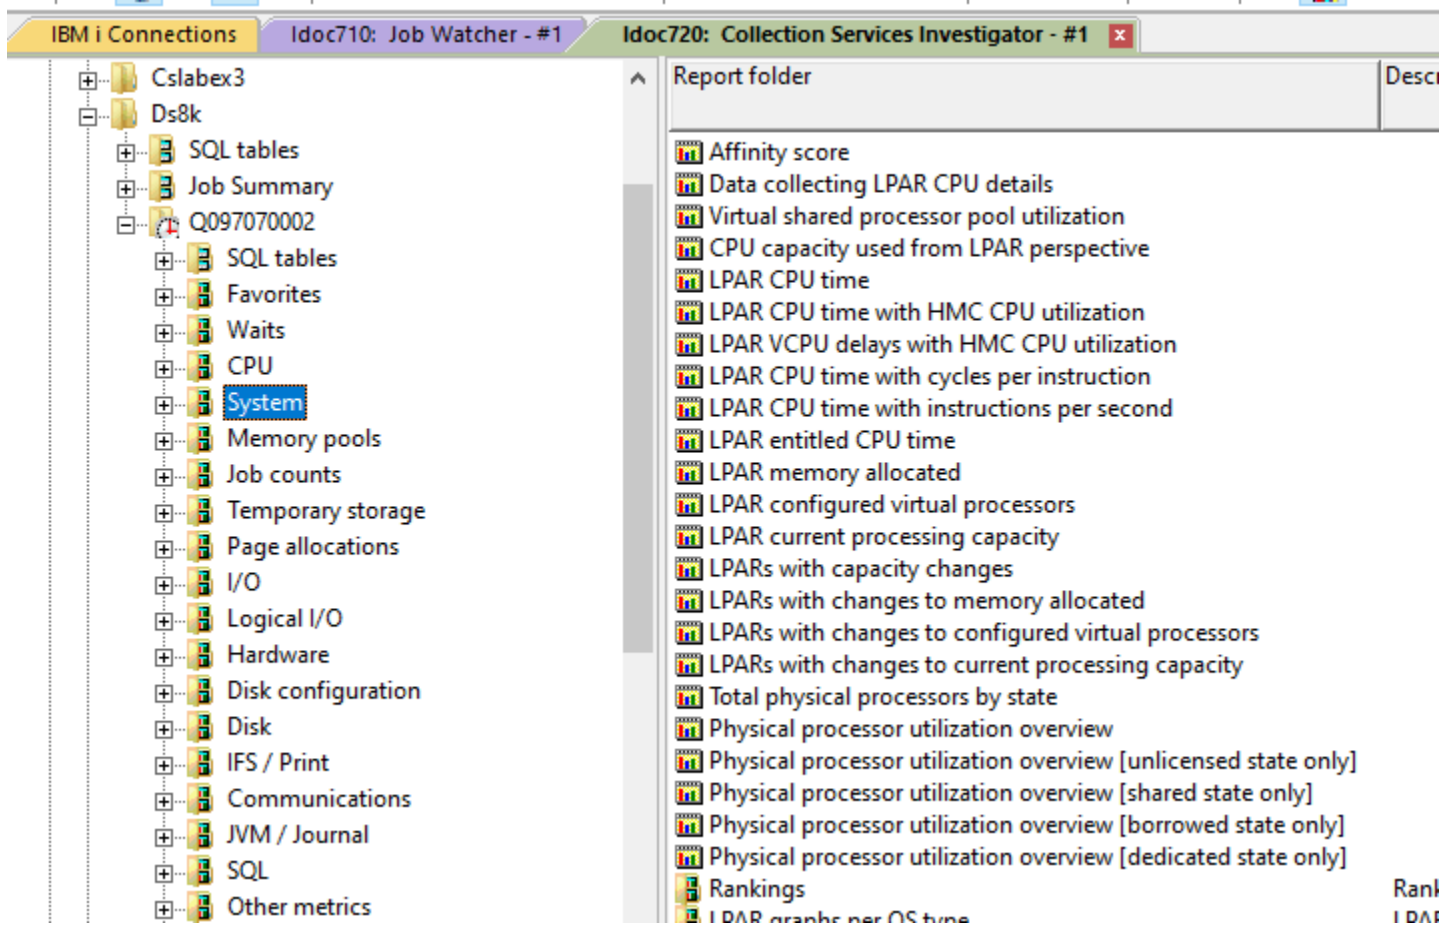


Collection	Using Collection Summary	Status	Descri
SQL tables			
Job Summary			
Q097070002	No	Ready - Missing: DISK (FSM)	

A collection in library DS8K with QAPMLPARH data in Collection Services Investigator

If the Status column says LPARH is missing, then it's not going to work, and something is wrong. If necessary, contact IBM I Performance Support for help.

Look for the graphs under the System folder once the data is available.



System folder and graphs in Collection Services Investigator