

This is a short description of the newly introduced additional metrics for Ganglia running with Linux on POWER5 systems.

Version: May 31st, 2006

For questions/tips etc. please contact Michael Perzl (michael@perzl.org).

The following additional metrics are defined for Linux:

- capped
- cpu_entitlement
- cpu_in_lpar
- cpu_in_machine
- cpu_in_pool
- cpu_pool_idle
- cpu_used
- disk_read
- disk_write
- kernel64bit
- lpar
- lpar_name
- lpar_num
- oslevel
- serial_num
- smt
- splpar
- weight

Despite the fact that most of these metrics are mostly only useful for Linux running in a POWER5 Shared Processor LPAR, some “reasonable” values must be returned if not running in that scenario. The metrics are now discussed in more detail.

Metric: **capped**

Return type:

- GANGLIA_VALUE_STRING

Notion:

- This metric either returns “yes” if the system is a POWER5 Shared Processor LPAR which is running in capped mode or “no” otherwise.

Remarks:

- A distinction has to be made if we are running on a POWER5 system (capped-capable) or not.
 - The `/proc/ppc64/lparcfg` **capped** entry is used to obtain that information.
-

Metric: **entitlement**

Return type:

- GANGLIA_VALUE_FLOAT

Notion:

- This function returns the Capacity Entitlement of the system in units of physical CPUs.

Remarks:

- The `/proc/ppc64/lparcfg` **partition_entitled_capacity** entry is used to obtain that information.
- If the `partition_entitled_capacity` entry is not found the number of CPUs in the system/LPAR is returned.

Metric: **cpu_in_lpar**

Return type:

- GANGLIA_VALUE_INT

Notion:

- This metric returns the number of CPUs the OS sees in the system.
- In a POWER5 Shared Processor LPAR this returns the number of virtual CPUs.
- When SMT is enabled this number is doubled.

Remarks:

- The `/proc/ppc64/lparcfg` **partition_active_processors** entry is used to obtain that information.
 - If the `partition_active_processors` entry is not found the number of CPUs in the system/LPAR is returned.
-

Metric: **cpu_in_machine**

Return type:

- GANGLIA_VALUE_INT

Notion:

- This metric returns the number of physical CPUs in the whole system.

Remarks:

- The `/proc/ppc64/lparcfg` **system_potential_processors** entry is used to obtain that information.
- If the `system_potential_processors` entry is not found the number of CPUs in the system/LPAR is returned.

Metric: **cpu_in_pool**

Return type:

- GANGLIA_VALUE_INT

Notion:

- This metric returns the number of physical CPUs in the Shared Processor Pool.

Remarks:

- The `/proc/ppc64/lparcfg` **pool_num_procs** entry is used to obtain that information.
 - If the `pool_num_procs` entry is not found the number of CPUs in the system/LPAR is returned.
-

Metric: **cpu_pool_idle**

Return type:

- GANGLIA_VALUE_FLOAT

Notion:

- This metric returns in fractional numbers of physical CPUs how much the Shared Processor Pool is idle.
- For example, if 7 physical CPUs are in the Shared Processor Pool, a value of 4.69 might be returned meaning that only an amount of $(7 - 4.69) = 2.31$ physical CPUs were used since the last time this metric was measured.

Remarks:

- The `/proc/ppc64/lparcfg` **pool_idle_time** entry is used to obtain that information.
- For good numerical results the time stamps are measured in μ -seconds.
- As the Shared Processor Pool idle time is returned in CPU clock cycles this value must still be divided by the CPU clock rate (obtained from the **timebase** entry from `/proc/cpuinfo`).
- If the `pool_idle_time` entry is not found a value of 0.0 is returned assuming there is no Shared Processor Pool available.

Metric: **cpu_used**

Return type:

- GANGLIA_VALUE_FLOAT

Notion:

- This metric returns in fractional numbers of physical CPUs how much compute resources this shared processor has used since the last time this metric was measured.
- For example, if the LPAR is running in uncapped mode and has a Capacity Entitlement of 0.2 physical CPUs and a value of 0.5 is measured then this LPAR has used $2.5 \times$ its entitled capacity since the last time this metric was measured (i.e., basically using 250% of its entitled CPU resources for this measured time interval).

Remarks:

- The `/proc/ppc64/lparcfg` **purr** entry is used to obtain that information. PURR is the acronym for **P**rocessor **U**tutilization of **R**esources **R**egister.
 - For good numerical results the time stamps are measured in μ -seconds.
 - As the real used CPU resources value is returned in CPU clock cycles this value must still be divided by the CPU clock rate (obtained from the **timebase** entry from `/proc/cpuinfo`).
 - If the purr entry is not found or some other errors occur the total number of CPUs found in the system/LPAR is returned.
-

Metric: **disk_read**

Return type:

- GANGLIA_VALUE_FLOAT

Notion:

- This metric returns in units of kB the total read I/O of the system.

Remarks:

- The `/proc/diskstats` file is used to obtain that information.
- For good numerical results the time stamps are measured in μ -seconds.

Metric: **disk_write**

Return type:

- GANGLIA_VALUE_FLOAT

Notion:

- This metric returns in units of kB the total write I/O of the system.

Remarks:

- The `/proc/diskstats` file is used to obtain that information.
 - For good numerical results the time stamps are measured in μ -seconds.
-

Metric: **kernel64bit**

Return type:

- GANGLIA_VALUE_STRING

Notion:

- This metric either returns “yes” if the running AIX kernel is a 64-bit kernel or “no” otherwise.

Remarks:

- Some distinction has to be made for all possible Linux distributions (SUSE, Red Hat, Debian, etc.).
-

Metric: **lpar**

Return type:

- GANGLIA_VALUE_STRING

Notion:

- This metric either returns “yes” if the system is a LPAR or “no” otherwise.

Remarks:

- Some efforts have to be made to make a decision on this metric as some defines available in kernel space (during kernel compilation) which would make that a very easy decision seem not to be available in user space.

Metric: `lpar_name`

Return type:

- GANGLIA_VALUE_STRING

Notion:

- This metric returns the name of the LPAR as defined on the Hardware Management Console (HMC) or some reasonable message otherwise.

Remarks:

- - The Open Firmware `/proc/device-tree/ibm,partition-name` entry is used to obtain that information.
 - If the `ibm,partition-name` entry is not found in the Open Firmware device-tree it is assumed that the system is not a LPAR. If it can't be read from the Open Firmware device-tree an appropriate error message is returned.
-

Metric: `lpar_num`

Return type:

- GANGLIA_VALUE_INT

Notion:

- This metric returns the partition ID of the LPAR as defined on the Hardware Management Console (HMC) or some reasonable message otherwise.

Remarks:

- The `/proc/ppc64/lparcfg partition_id` entry is used to obtain that information.
- If the `partition_id` entry is not found a value of 0 is returned.

Metric: **oslevel**

Return type:

- GANGLIA_VALUE_STRING

Notion:

- This metric returns a more detailed version string similar as provided by the AIX command 'oslevel'.

Remarks:

- This metric is retrieved only once and then “cached” for subsequential calls.
 - A distinction has to be made for the different Linux distributions.
 - Also the maximum length of a GANGLIA_VALUE_STRING (32 characters as defined by MAX_G_STRING_SIZE) has to be taken into account. This leads to some short forms as for instance:
 - SUSE Linux Enterprise Server 9 Service Pack 3 would become “SLES 9 SP 3”
 - Red Hat Enterprise Linux AS release 4 Update 2 would become “RHEL AS 4 Update 2”
 - etc.
 - If nothing reasonable can be found out an appropriate error message is returned.
-

Metric: **serial_num**

Return type:

- GANGLIA_VALUE_STRING

Notion:

- This metric returns the serial number of the system.

Remarks:

- The /proc/ppc64/lparcfg **serial_number** entry is used to obtain that information.
- If the serial_number entry is not found an appropriate error message is returned.

Metric: **smt**

Return type:

- GANGLIA_VALUE_STRING

Notion:

- This metric either returns “yes” if SMT is enabled or “no” otherwise.

Remarks:

- The `/proc/ppc64/lparcfg` **partition_active_processors** entry is used to obtain that information.
 - If the `partition_active_processors` entry is not found it is assumed that the system is not SMT-capable.
-

Metric: **splpar**

Return type:

- GANGLIA_VALUE_STRING

Notion:

- This metric either returns “yes” if the system is running in a shared processor LPAR or “no” otherwise.

Remarks:

- The `/proc/ppc64/lparcfg` **shared_processor_mode** entry is used to obtain that information.
- If the `shared_processor_mode` entry is not found it is assumed that this system is not shared processor LPAR capable.

Metric: **weight**

Return type:

- GANGLIA_VALUE_INT

Notion:

- This metric returns the weight of the LPAR running in uncapped mode.

Remarks:

- The `/proc/ppc64/lparcfg` **capacity_weight** entry is used to obtain that information.
- If the `capacity_weight` entry is not found or an error occurs a value of -1 is returned.