



Sandia
National
Laboratories

Broader and simpler LDMS data sampling

Lightning update

Benjamin Allan, 8/19/2021

SAND2021-xxxx C



Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Overview

1. Broader

- *Process metrics sampler* without code instrumentation or SLURM plug-in
- *AMD GPUs* (ROCM 4.x rdc sampler)
- *tx2mon* (ARM SOC)
- *loadavg*
- *syspapi*

2. Simpler to deploy: *dynamic & auto-adaptive*

- Lustre: client, mdt, ost
- TCP network data
- LNET stats
- AMD GPUs

Broader: per-process metrics: Who & Why

- Sampler(s) developed by Boston U, OGC, SNL
- Watch login nodes in detail
 - detect node misuse
- Watch complex jobs in detail
 - e.g.
 - tracking many-step automated builds
 - track through deeply scripted workflows, such as machine learning
- See what resources system services are consuming
 - e.g.
 - slurm launch steps

Broader: per-process metrics: What

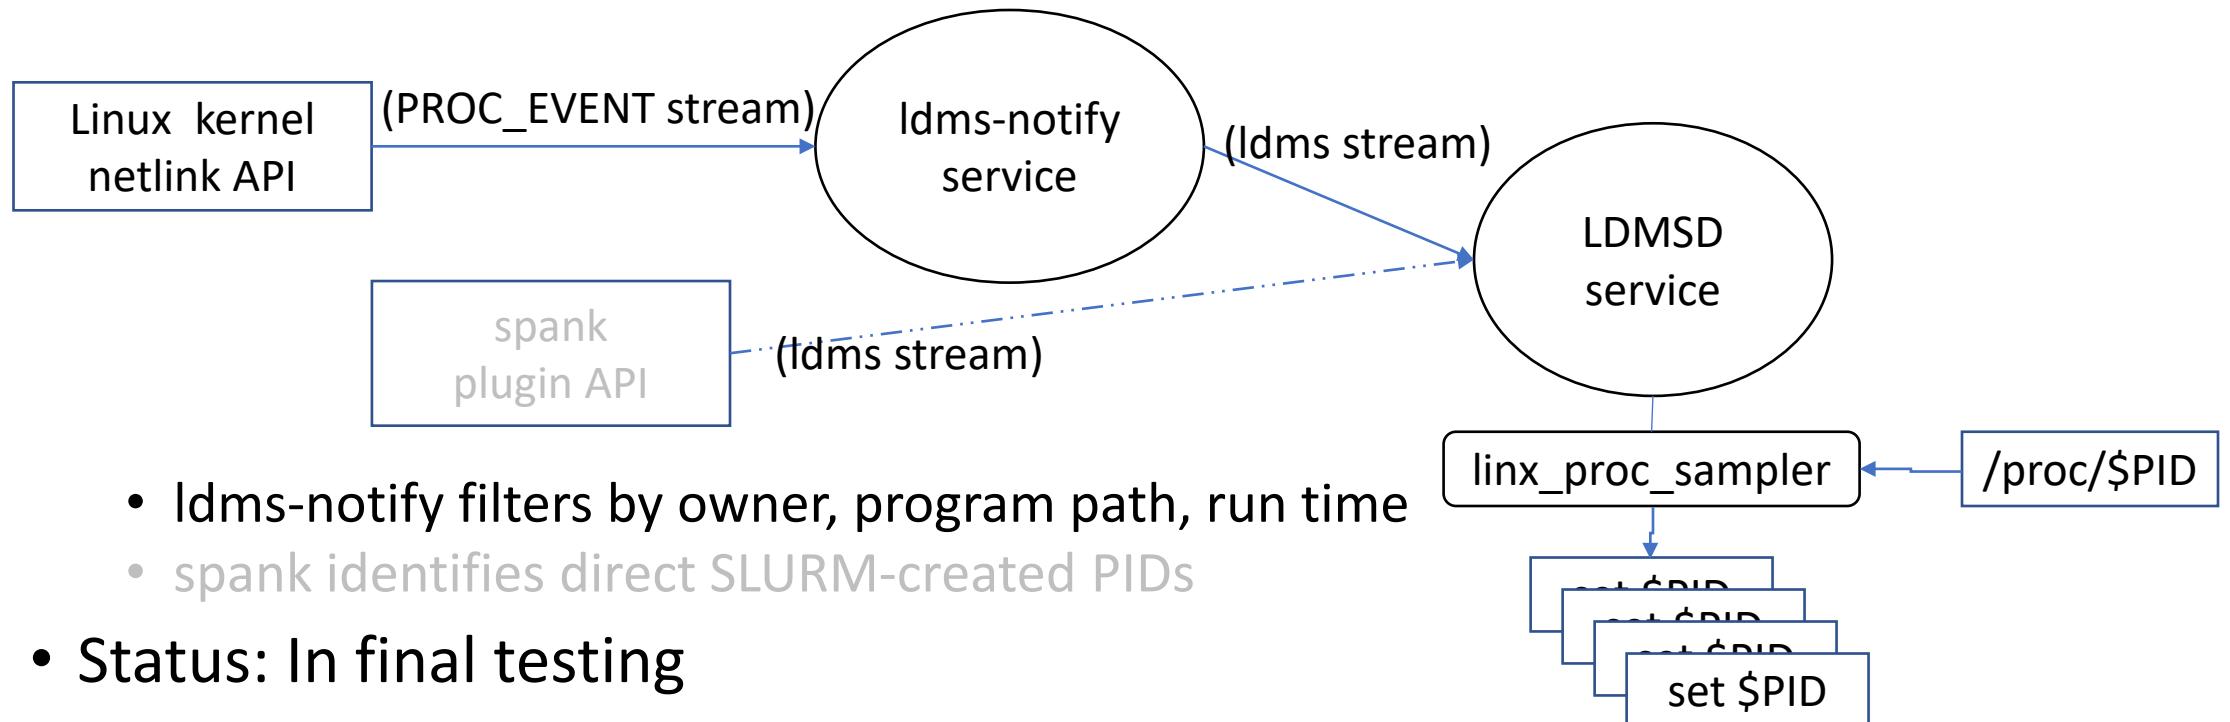
- Command line text
- Executable full path
- 130+ from /proc/\$PID/
 - RM job id, task ids if available
 - memory usages
 - CPU
 - I/O
 - file handle count, types
 - parent PID and PID, thread count
 - etc

JSON process sampler config:

```
{ "stream" : "slurm",
  "metrics" : [
    "stat_pid", "stat_state", "stat_rss",
    "stat_utime", "stat_stime",
    "stat_cutime", "stat_cstime",
    "stat_num_threads", "n_open_files",
    "io_read_b", "io_write_b",
    "status_vmdata",
    "status_rssfile", "status_vmswap",
    "status_hugetlbpages",
    "status_voluntary_ctxt_switches",
    "status_nonvoluntary_ctxt_switches",
    "syscall", "stat_comm", "cmdline"
  ]
}
```

Broader: per-process metrics: How

- Get process IDs from a Linux kernel channel (via Idms-notify daemon)
- Collect a metric set periodically per process
- How it works:



Broader: rdc_sampler for AMD GPU

- Contributed by AMD, SNL
- ROCM 4.x RDC API
 - Default RDC_FI_metrics:
 - GPU_CLOCK, MEM_CLOCK,
 - MEMORY_TEMP, GPU_TEMP,
 - POWER_USAGE,
 - PCIE_TX, PCIE_RX,
 - GPU_UTIL,
 - MEMORY_USAGE, MEMORY_TOTAL
 - Evolving with ROCM releases and hardware releases
 - Collaboration with SNL Advance Architecture testbeds team

Broader: tx2mon (ARM SOC)

- Contributed by HPE, SNL
- HPE ARM system-on-chip metrics
 - temperature, power, voltage, frequency (per core or bus)
 - CPU throttle events
- One metric set per socket

Broader: loadavg

- Contributed by SNL
- Turns `/proc/loadavg` into a metric set
 - These numbers are not trivially derivable from existing metric sets
- Data can be used to distinguish CPU saturation (`/proc/stat`) from the load numbers users see in ***top*** and similar utilities

Broader: syspapi

- Per-core hardware counters (machine specific registers)
 - PAPI used only for configuration file names of the counters
 - Kernel perf API used for data collection
- JSON configuration to select counters:

```
{  "schema": "syspapi-E5-2695.5",  "events": [    "INST_RETIRED:ANY_P",    "CPU_CLK_THREAD_UNHALTED:THREAD_P",    "CPU_CLK_THREAD_UNHALTED:REF_XCLK",    "MEM_LOAD_UOPS_RETIRED:L2_MISS",    "MEM_LOAD_UOPS_RETIRED:L3_MISS"  ]}
```

- Known conflict with HPC toolkit
 - SLURM ‘feature’ option has been used to pause syspapi at user’s choice.
 - LDMS stream message to pause syspapi is available for testing.

Simpler to deploy

- New & updated samplers automatically manage sampling multiple similar devices and their dynamic disappearance or appearance.
- Administrators no longer need to invent site-specific schema names correlated to lists of device identifiers.
- Very large schemas (high column count CSV) encapsulating multiple devices are no longer needed.

Simpler to deploy: sampler specifics

- Main-line now includes LLNL Lustre plugins: client, mdt, ost sampling
 - Makes 1 set per mount point or device.
 - *lustre2_client* sampler is no longer recommended.
- TCP network data
 - New *procnet* sampler handles ports dynamically & ignores inactive ports.
 - Makes 1 set per port, skips set updates with no data changes.
 - Old *procnetdev* sampler works as before.
- LNET traffic
 - Implementation in *Inet_stats* now handles the case when all the mount points are temporarily missing.
 - Merge to main-line is pending.
- RDC sampler (AMD GPU) is refactored to make 1 set per GPU.

Broader & Simpler

- Samplers are improving
- Let us know what other samplers could work better for you
- Let us know what other samplers are needed