SANDIA REPORT
SAND2019-4915

Printed April 2019

Sandia National Laboratories

Two Weeks In The Life of Skybridge

Benjamin A. Allan

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185
Livermore, California 94550

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology & Engineering Solutions of Sandia, LLC.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: http://www.osti.gov/scitech

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Road
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: https://classic.ntis.gov/help/order-methods

ABSTRACT

This report documents a new large public data set for researchers studying the behavior of large commodity high-performance computing systems. Such large data sets are typically confined within institutions and access to them is limited to institutional partners. We provide it to promote HPC research more widely.

The data set provides a two week time series of performance data collected once per minute using the Lightweight Distributed Metric Service from the system Skybridge at Sandia National Laboratories and the corresponding job-level accounting information. General system log information is not provided.

ACKNOWLEDGMENT

Thanks to Sophia Lefantzi for tabulating in 2016 most of the metric explanations used in the appendix.

CONTENTS

LIST OF TABLES

PREFACE

The work load of the Skybridge cluster during the two weeks was not specially selected for this report. Laboratory computing demand within a calendar year is highly variable and Skybridge is one of more than a dozen HPC systems sited at Sandia, so any inference about the total annual work load of Skybridge or Sandia HPC based on this data will be invalid.

The intended initial use of this data set is open research on algorithms for automatically discovering behaviors of jobs or of the system as a whole. The data has not been "cleaned up" or modified in any way other than anonymization. Dealing with messy data at scale is a critical challenge in HPC systems analysis. The sample frequency (once per minute) is insufficient to resolve high frequency behaviors.

The data set documented in this report has been anonymized. No application, user, or host identifiers contained in the data set or the examples within this report are the actual identifiers found in the Skybridge environment. The node count and general scalable-unit characteristics of Skybridge are well known from its appearances on the TOP500 list of supercomputers, so there is no attempt to anonymize the cluster as a whole.

SUMMARY

This data set is available by contacting the author (baallan@sandia.gov) or visiting https://ovis.ca.sandia.gov/index.php/Downloads_and_documentation#OVIS_Research_Data and selecting the data set skybridge-2019-1.

The aggregate data size is about 43 gigabytes compressed.

The time series data are collected from the LDMS [1] version 3.4.11 plugins sysclassib, jobid, lustre2_client, meminfo, procnetdev, procnfs, procstat, and vmstat.

The per-job job data included are collected from the job accounting command 'sacct' of slurm. Approximately 7000 jobs and 27000 job steps are included.

For automated processing, metadata files are provided to link column numbers, metric names, and data types.

1. THE DATA SET

1.1. SKYBRIDGE

This section provides data necessary when interpreting the performance metrics. The period of the data collected on Skybridge is from March 25, 2018 through April 10, 2018. The skybridge cluster has 1848 compute nodes, 36 Lustre router nodes, 12 administrative nodes, and 12 login nodes. All nodes host dual Xeon E5-2670 8-core processors with hyperthreading disabled. In the data sets provided, these nodes are anonymized as in Table 1-1. The data sets also contain a numeric LDMS-specific identifier component_id which maps to host names and eliminates the need for string comparison when sorting data.

Table 1-1. This provides the role of the hosts.

names	role	component_id
cb[1-1848]	compute	200001-201848
cbg[1-36]	Lustre router	230001-230036
cba[1-12]	administrative	210001-210012
cb-sl[1-12]	login	220001-220012

The slurm resource manager has two partitions, batch (multi-day) and short (4 hour time limit). Slurm version 17 was used. The short partition has a small number of nodes.

The Infiniband network (IB) is a three level tree as presented in [2]. IB data is collected from the node interface cards; however, switch port data are not available for the collection period. The IB hardware counters are reset by a supervisory process approximately every 10 minutes, but a reset command may be missed or never sent in some circumstances.

The Lustre routers connect the compute nodes IB network to the Lustre storage network which is shared with other clusters. In all nodes, the device qib0.1 is the interface to the cluster interconnect, and in the Lustre router nodes, the additional device mlx4_0.1 is the interface to the Lustre storage fabric.

A 1GB/sec Ethernet network connects all hosts for supervisory purposes and for access to NFS version 3 and GPFS services. No GPFS performance data is included in this set.

1.2. SCHEMAS AND METADATA

This section provides general notes on the data and specific notes on the individual data schemas represented in the set. The "meta/" subdirectory of the data set provides the files described in Table 1-2. Each meta* file contains space-separated columns giving the data file column number, the metric name, and the metric type identifier.

Table 1-2. Sources to interpret columnar information

File	Content
coltypes.txt	type identifiers and their definitions
metadata.jobid	data collected on node from files written by slurm prolog and epilog scripts
metadata.Lustre_Client	client statistics for two mount points
metadata.meminfo	content of /proc/meminfo
metadata.procnetdev	content of /proc/net/dev including IP-over-IB traffic
metadata.procnfs	content of /proc/net/rpc/nfs
metadata.procstat_16	content of /proc/stat, less interrupt statistics
metadata.slurmdata	sacct output, with timestamps in seconds since UNIX epoch
metadata.gw_sysclassib	Lustre router IB counters
metadata.sysclassib	IB counters from all nodes with only qib0.1
metadata.vmstat	content of /proc/vmstat

1.2.1. A note on rates

The Lustre and IB data sets contain rate metrics (named *.rate#DEVICE) computed by the sampler plugins at time of primary metric collection.

The derived rate metrics in the data set exist for the convenience of on-line analysis and reporting tools; post-hoc analyses should compute derivatives of counters from the raw counter data. In some cases, integer overflow may not be taken into account, e.g. yielding electronically impossible bandwidth measures. Down-stream filters suppressed these values. In other cases, where a sample collection (or a series of same) is missed by the data aggregator daemon, the rate reported when collection resumes will only apply to the previous single interval, not the entire data gap.

1.3. ANONYMIZATION AND SLURM

1.3.1. Special slurm considerations

To facilitate correlating relationships among similar jobs, anonymized UID, job name, and work directory paths are provided in the slurm data set. Automated detection of related jobs is key to

discovering anomalies in a family of jobs. The JobID and job step id strings are not anonymized, being numeric and expired. The slurm JobId is included (once slurm prolog has run) in all data sets as the 'job_id' metric; for node types where jobs do not run, job_id is always 0. Inclusion of job_id in all metric sets eliminates complicated join operations with the slurm data file and provides a check on the job End time recorded in the slurm data. Where the slurm accounting data and time series recorded job_id are inconsistent, the time series value derived from the execution of the slurm epilog is more reliable for analytical purposes.

1.3.2. Correlating anonymized data

Real UIDs are mapped to large integers. Job names are mapped to "n" appended with an integer I. If the same string name is used for multiple jobs, the same "nI" string will be used for all these jobs. Work directory path elements are mapped individually to "p" appended with an integer (as for job names). This allows relative directory structures to be constructed and a series of related jobs running from the same sub-tree to be discovered.

1.4. FILES

The "anonymized/" subdirectory contains the checksum, header and data files, as described in Table 1-3. The values for $schema are the names of the files documented in Table 1-2, stripped of the "metadata." prefix. The values for $date are the dates as described in subsection 1.4.1.

Table 1-3. Sources to validate and interpret columnar information

File(s)	Content
sha1sums.gzipped sha1sums.uncompressed	sha1sum output for each data file gzipped sha1sum output for each data file uncompressed
slurmdata.psv.gz slurmdata.HEADER.psv.gz	|-separated sacct data from slurm. |-separated column headings from sacct
$schema.HEADER.$date $schema.$date	Comma separated metric names for file $schema.$date. Comma separated LDMS metric data.

The SHA1 sums of the sha1sums.* files are provided in Table 1-4 as an independent means to verify the archives when duplicated.

Table 1-4. Summary checksums

sha1sum	File
5ed723cfeb04b4cae3d36171ffbf9e07d0fb14cd	sha1sums.gzipped
95b725a6b49a4cd527db2df7a3372d49f439c66a	sha1sums.uncompressed

1.4.1. Notes on time and aggregation

The numeric suffix on uncompressed metric files is the start time of aggregation for that file's data. Generally it will be near midnight Albuquerque time. Any date stamps contained in the data that are earlier than the date indicated in the file name by more than two minutes should be discarded; these are indicative of old data in a first level aggregator that has not been replaced because the producing node is down. The data files in this set were collected with a second level aggregator connected to first level aggregators that run on the administrative nodes of skybridge.

If you have LDMS installed, it provides the 'lsdate' command which provides augmented directory listings for timestamp suffixed file names. Absent LDMS, a useful Linux utility invocation is "TZ=America/Denver date --date=@EPOCH" where EPOCH is the integer time from the data file names. It will return the Albuquerque time if the American time zones are known to your computer. Omitting the TZ specification will return your local time when the data files were originated.

1.5. CONCLUSION

The data described in this report is provided on a best-effort basis. Any researchers publishing results derived from or developed with use of this data set are requested to provide attribution by citing this report. Please contact the author with any corrections to errors or omissions in appendix tables.

REFERENCES

[1] A. Agelastos, B. Allan, J. Brandt, P. Cassella, J. Enos, J. Fullop, A. Gentile, S. Monk, N. Naksinehaboon, J. Ogden, M. Rajan, M. Showerman, J. Stevenson, N. Taerat, and T. Tucker. Lightweight Distributed Metric Service: A Scalable Infrastructure for Continuous Monitoring of Large Scale Computing Systems and Applications. In *Proc. IEEE/ACM International Conference for High Performance Storage, Networking, and Analysis (SC14)*, 2014.

[2] Benjamin A. Allan, Michael Aguilar, and Serge Polevitzky. Comprehensive, synchronous, high frequency measurement of Infiniband networks in production HPC systems. Technical report SAND2018-3506 C, Sandia National Laboratories, Albuquerque, New Mexico 87185, 2018.

APPENDICES

A. METRIC DEFINITIONS

Analysis and reporting of metric-derived information requires the information on meaning, bounds and units shown in Tables A-2 to A-9. The tables are formatted for use with a large display device or the zoom feature of any PDF reader. The metric descriptions are derived from Linux kernel 3.x and Lustre documentation or from inspection of Linux source code.

The columns of these tables are defined in Table A-1. For any table missing one of these columns, the column value can be assumed to be 0 or False or undefined as appropriate.

Table A-1. reference information column definitions

column name	explanation
Sampler	name of the LDMS plugin.
Name	name of the metric; those with CORE or DEVICE suffix match many columns with unique identifiers substituted.
description	best available description we have found for the monitoring context.
Mode	the behavior of the metric: capacity (e.g. RAM), label (e.g. status string), or integral (e.g. network transmitted bytes or CPU jiffies)
Units	units of the raw metric as collected from the system. No data is scaled during collection, other than derived '.rate' values.
Dimension	underlying character of the units; units with identical dimensions are interconvertible.
Use Rate Only or URO	analysis and user presentation policy Boolean setting. Metrics marked URO would only be plotted in raw form as a quality control check.
Delta/Job Useful or DJU	analysis and user presentation policy Boolean setting. This is True if typical users should be presented with the time integral of the metric derivative.
Local rate UB or LRUB	upper bound on the node-local metric time derivative. This field is 0 or empty if no useful bound is known. The value derives from hardware specifications and is specific to the cluster. For 64 bit counters, using UINT64_MAX as an upper bound for analyses has not proven useful.
UB Local or LUB	upper bound on the node-local metric value. This field is 0 or empty if no useful bound is known. The value derives from the bit width of the counter sampled, not the width of the LDMS data type in which the metric is stored.
max by dt	theoretical limit on the metric time derivative value. Noise in calculations may cause this limit to be slightly exceeded in practice. The limit is expressed in seconds/second for CPU usages. A 0 or empty value means no useful limit is known.
itype	for integral mode metrics only, this describes the behavior of a counter which reaches its maximum limit: roll (loops back to 0 automatically), infinite or inf (not expected to reach the maximum limit in the life of the machine), saturates (sticks at the maximum limit until reset by an external process).
notes:	anything noteworthy.
kind	ldms data storage type or kernel storage type if smaller.

Table A-2. common reference information

Sampler	Name	description	Mode	Units	Dimension	URO	DJU	LRUB	kind
General	Time	wall clock time	integral	second	time	TRUE	TRUE	1	%d.%06d
General	component_id	data source numeric id	label						uint64
General	uid	UNIX numeric user id	label						uint64
General	job_id	resource manager job id	label						uint64
jobid	anonymized_host_ProducerName	hostname of node	label						char_array
jobid	anonymized_int_uid	fake user id number	label						uint64
jobid	anonymized_name_username	fake user name	label						char_array

Table A-3. meminfo reference information

Sampler	Name	description	Mode	Units	Dimension	Use Rate Only	Delta/Job Useful	kind
meminfo	Active(anon)	Active memory not mapped to file or device.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	Active(file)	Active memory mapped to file or device.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	Active	The total amount of buffer or page cache memory, in kilobytes, that is in active use. This is memory that has been recently used and is usually not reclaimed for other purposes.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	AnonHugePages		capacity	kB	data	FALSE	FALSE	uint64
meminfo	AnonPages	Active(Anon)+Active(File)	capacity	kB	data	FALSE	FALSE	uint64
meminfo	Bounce		capacity	kB	data	FALSE	FALSE	uint64
meminfo	Buffers	The amount of physical RAM, in kilobytes, used for file buffers.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	Cached	The amount of physical RAM, in kilobytes, used as cache memory.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	CmaFree	Free CMA (Contiguous Memory Allocator) pages.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	CmaTotal	Total CMA (Contiguous Memory Allocator) pages.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	CommitLimit		capacity	kB	data	FALSE	FALSE	uint64
meminfo	Committed_AS	The total amount of memory, in kilobytes, estimated to complete the workload. This value represents the worst case scenario value, and also includes swap memory.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	DirectMap1G	kB mapped in 1G chunks	capacity	kB	data	FALSE	FALSE	uint64
meminfo	DirectMap2M	kB mapped in 2M chunks	capacity	kB	data	FALSE	FALSE	uint64
meminfo	DirectMap4k	kB mapped in 4k chunks	capacity	kB	data	FALSE	FALSE	uint64
meminfo	Dirty	The total amount of memory waiting to be written back to the disk.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	HardwareCorrupted		capacity	kB	data	FALSE	FALSE	uint64
meminfo	HugePages_Free	The total number of hugepages available for the system.	capacity	pages	data	FALSE	FALSE	uint64
meminfo	Hugepagesize	The size for each hugepages unit in kilobytes	capacity	kB	data	FALSE	FALSE	uint64
meminfo	HugePages_Rsvd		capacity	pages	data	FALSE	FALSE	uint64
meminfo	HugePages_Surp		capacity	pages	data	FALSE	FALSE	uint64
meminfo	HugePages_Total	The total number of hugepages for the system. The number is derived by dividing Hugepagesize by the megabytes set aside for hugepages specified in /proc/sys/vm/hugetlb_pool.	capacity	pages	data	FALSE	FALSE	uint64
meminfo	Inactive(anon)	Inactive memory not mapped to file or device and available for reclaim.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	Inactive(file)	File buffer/page cache memory available for reclaim.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	Inactive	The total amount of buffer or page cache memory, in kilobytes, that are free and available. This is memory that has not been recently used and can be reclaimed for other purposes.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	KernelStack		capacity	kB	data	FALSE	FALSE	uint64
meminfo	Mapped	The total amount of memory, in kilobytes, which have been used to map devices, files, or libraries using the mmap command.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	MemAvailable	An estimate of how much memory is available for starting new applications, without swapping.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	MemFree	The amount of physical RAM, in kilobytes, left unused by the system.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	MemTotal	Total amount of physical RAM.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	Mlocked	Locked memory	capacity	kB	data	FALSE	FALSE	uint64
meminfo	NFS_Unstable		capacity	kB	data	FALSE	FALSE	uint64
meminfo	PageTables		capacity	kB	data	FALSE	FALSE	uint64
meminfo	Shmem		capacity	kB	data	FALSE	FALSE	uint64
meminfo	Slab	The total amount of memory, in kilobytes, used by the kernel to cache data structures for its own use.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	SReclaimable		capacity	kB	data	FALSE	FALSE	uint64
meminfo	SUnreclaim		capacity	kB	data	FALSE	FALSE	uint64
meminfo	SwapCached	The amount of swap, in kilobytes, used as cache memory.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	SwapFree	The total amount of swap free.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	SwapTotal	The total amount of swap available, in kilobytes.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	Unevictable	Pages from shm_lock, vm lock, RAMFS, etc that cannot be evicted (reclaimed).	capacity	kB	data	FALSE	FALSE	uint64
meminfo	VmallocChunk	The largest contiguous block of memory, in kilobytes, of available virtual address space.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	VmallocTotal	The total amount of memory, in kilobytes, of total allocated virtual address space.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	VmallocUsed	The total amount of memory, in kilobytes, of used virtual address space.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	Writeback	The total amount of memory actively being written back to the disk.	capacity	kB	data	FALSE	FALSE	uint64
meminfo	WritebackTmp		capacity	kB	data	FALSE	FALSE	uint64

Table A-4. Lustre reference information

Sampler	Name	description	Mode	Units	Dimension	URO	DJU	LRUB	LUB	itype	notes:	kind
Lustre_Client	client.lstats.alloc_inode#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.brw_read#llite.DEVICE	bytes read in bandwidth test mode llite.DEVICE	integral	pages	data	T	T	4294967296	4294967296	rolls	units need validation	uint64
Lustre_Client	client.lstats.brw_write#llite.DEVICE	bytes written in bandwidth test mode on llite.DEVICE	integral	pages	data	T	T	4294967296	4294967296	rolls	units need validation	uint64
Lustre_Client	client.lstats.close#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.create#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.dirty_pages_hits#llite.DEVICE	dirty_pages_hits for llite.DEVICE	integral	events	event	T	T		0	rolls		uint64
Lustre_Client	client.lstats.dirty_pages_misses#llite.DEVICE	dirty_pages_misses for llite.DEVICE	integral	events	event	T	T		0	rolls		uint64
Lustre_Client	client.lstats.flock#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.fsync#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.getattr#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.getxattr#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.inode_permission#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.ioctl#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.link#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.listxattr#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.mkdir#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.mknod#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.mmap#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.open#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.osc_read#llite.DEVICE		integral	bytes	data	T	T	4294967296	4294967296	rolls		uint64
Lustre_Client	client.lstats.osc_write#llite.DEVICE		integral	bytes	data	T	T	4294967296	4294967296	rolls		uint64
Lustre_Client	client.lstats.read_bytes#llite.DEVICE	bytes read (which may be from cache or network) on llite.DEVICE	integral	bytes	data	T	T	4294967296	4294967296	rolls		uint64
Lustre_Client	client.lstats.readdir#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.removexattr#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.rename#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.rmdir#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.seek#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.setattr#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.setxattr#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.statfs#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.status#llite.DEVICE	current status of llite.DEVICE	label	none	none	F	F	0	0	not		uint64
Lustre_Client	client.lstats.symlink#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.truncate#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.unlink#llite.DEVICE	calls for llite.DEVICE	integral	calls	call	T	T	0	0	inf		uint64
Lustre_Client	client.lstats.write_bytes#llite.DEVICE	bytes written (which may be to network or local cache) on llite.DEVICE	integral	bytes	data	T	T	4294967296	4294967296	rolls		uint64

Table A-5. procnetdev reference information

Sampler	Name	description	Mode	Units	Dimension	URO	DJU	LRUB	itype	notes:	kind
procnetdev	rx_bytes#DEVICE	Bytes received on DEVICE	integral	bytes	data	TRUE	TRUE	varies	rolls	eth devices	uint64
procnetdev	rx_compressed#DEVICE	Compressed packets received by 0	integral	compressed_packets	data	TRUE	TRUE		rolls	1Gbit/s	uint64
procnetdev	rx_drop#DEVICE	Packets dropped on DEVICE	integral	packets	data	TRUE	TRUE		rolls		uint64
procnetdev	rx_errs#DEVICE	Receive errors on DEVICE	integral	errors	event	TRUE	TRUE		rolls	IB devices	uint64
procnetdev	rx_fifo#DEVICE	Received fifo buffer errors on DEVICE	integral	errors	event	TRUE	TRUE		rolls	4GByte/s	uint64
procnetdev	rx_frame#DEVICE	The frame error count on DEVICE	integral	errors	event	TRUE	TRUE		rolls		uint64
procnetdev	rx_multicast#DEVICE	multicast frames received by 0	integral	multicast_packets	data	TRUE	TRUE		rolls		uint64
procnetdev	rx_packets#DEVICE	Packets received on DEVICE	integral	packets	data	TRUE	TRUE		rolls		uint64
procnetdev	tx_bytes#DEVICE	bytes transmitted on DEVICE	integral	byte	data	TRUE	TRUE	varies	rolls		uint64
procnetdev	tx_carrier#DEVICE	Carrier losses detected in transmit on DEVICE	integral	errors	event	TRUE	TRUE		rolls		uint64
procnetdev	tx_colls#DEVICE	Transmit collision detected on DEVICE	integral	errors	event	TRUE	TRUE		rolls		uint64
procnetdev	tx_compressed#DEVICE	Compressed packets transmitted on DEVICE	integral	compressed_packets	data	TRUE	TRUE		rolls		uint64
procnetdev	tx_drop#DEVICE	Transmit packets dropped by 0	integral	packets	data	TRUE	TRUE		rolls		uint64
procnetdev	tx_errs#DEVICE	Total transmit errors detected on DEVICE	integral	errors	event	TRUE	TRUE		rolls		uint64
procnetdev	tx_fifo#DEVICE	Transmit FIFO buffer errors on DEVICE	integral	errors	event	TRUE	TRUE		rolls		uint64
procnetdev	tx_packets#DEVICE	Packets transmitted on DEVICE	integral	packets	data	TRUE	TRUE		rolls		uint64

Table A-6. procnfs reference information

Sampler	Name	description	Mode	Units	Dimension	URO	DJU	itype	notes:	kind
procnfs	commit	nfs client-side calls to commit	integral	calls	call	T	T	infinite		uint64
procnfs	pathconf	nfs client-side calls to pathconf	integral	calls	call	T	T	infinite		uint64
procnfs	fsinfo	nfs client-side calls to fsinfo	integral	calls	call	T	T	infinite		uint64
procnfs	fsstat	nfs client-side calls to fsstat	integral	calls	call	T	T	infinite		uint64
procnfs	readdirplus	nfs client-side calls to readdirplus	integral	calls	call	T	T	infinite		uint64
procnfs	readdir	nfs client-side calls to readdir	integral	calls	call	T	T	infinite		uint64
procnfs	link	nfs client-side calls to link	integral	calls	call	T	T	infinite		uint64
procnfs	rename	nfs client-side calls to rename	integral	calls	call	T	T	infinite		uint64
procnfs	rmdir	nfs client-side calls to rmdir	integral	calls	call	T	T	infinite		uint64
procnfs	remove	nfs client-side calls to remove	integral	calls	call	T	T	infinite		uint64
procnfs	mknod	nfs client-side calls to mknod	integral	calls	call	T	T	infinite		uint64
procnfs	symlink	nfs client-side calls to symlink	integral	calls	call	T	T	infinite		uint64
procnfs	mkdir	nfs client-side calls to mkdir	integral	calls	call	T	T	infinite		uint64
procnfs	create	nfs client-side calls to create	integral	calls	call	T	T	infinite		uint64
procnfs	write	nfs client-side calls to write	integral	calls	call	T	T	infinite		uint64
procnfs	read	nfs client-side calls to read	integral	calls	call	T	T	infinite		uint64
procnfs	readlink	nfs client-side calls to readlink	integral	calls	call	T	T	infinite		uint64
procnfs	access	nfs client-side calls to access	integral	calls	call	T	T	infinite		uint64
procnfs	lookup	nfs client-side calls to lookup	integral	calls	call	T	T	infinite		uint64
procnfs	setattr	nfs client-side calls to setattr	integral	calls	call	T	T	infinite		uint64
procnfs	getattr	nfs client-side calls to getattr	integral	calls	call	T	T	infinite		uint64
procnfs	retransmitts	nfs client-side calls to retransmitts	integral	event	event	T	T	infinite	here for naming bug in sampler code	uint64
procnfs	retransmits	nfs client-side calls to retransmits	integral	event	event	T	T	infinite		uint64
procnfs	numcalls	total calls to all client functions	integral	calls	call	T	T	rolls		uint64

Table A-7. procstat reference information

Sampler	Name	description	Mode	Units	Dimension	URO	DJU	LRUB	LUB	max by dt	itype	kind
procstat	guest_nice#CORE	time in jiffies since boot core CORE spent in guest_nice#	integral	jiffies	time	TRUE	TRUE	100		1	infinite	uint64
procstat	guest#CORE	time in jiffies since boot core CORE spent in guest#	integral	jiffies	time	TRUE	TRUE	100		1	infinite	uint64
procstat	steal#CORE	time in jiffies since boot core CORE spent in steal#	integral	jiffies	time	TRUE	TRUE	100		1	infinite	uint64
procstat	softirq#CORE	time in jiffies since boot core CORE spent in softirq#	integral	jiffies	time	TRUE	TRUE	100		1	infinite	uint64
procstat	irq#CORE	time in jiffies since boot core CORE spent in irq#	integral	jiffies	time	TRUE	TRUE	100		1	infinite	uint64
procstat	iowait#CORE	time in jiffies since boot core CORE spent in iowait#	integral	jiffies	time	TRUE	TRUE	100		1	infinite	uint64
procstat	idle#CORE	time in jiffies since boot core CORE spent in idle#	integral	jiffies	time	TRUE	TRUE	100		1	infinite	uint64
procstat	sys#CORE	time in jiffies since boot core CORE spent in sys#	integral	jiffies	time	TRUE	TRUE	100		1	infinite	uint64
procstat	nice#CORE	time in jiffies since boot core CORE spent in nice#	integral	jiffies	time	TRUE	TRUE	100		1	infinite	uint64
procstat	user#CORE	time in jiffies since boot core CORE spent in user#	integral	jiffies	time	TRUE	TRUE	100		1	infinite	uint64
procstat	guest_nice	sum of jiffies across all cores spent in guest_nice	integral	jiffies	time	TRUE	TRUE	1600		16	infinite	uint64
procstat	guest	sum of jiffies across all cores spent in guest	integral	jiffies	time	TRUE	TRUE	1600		16	infinite	uint64
procstat	steal	sum of jiffies across all cores spent in steal	integral	jiffies	time	TRUE	TRUE	1600		16	infinite	uint64
procstat	softirq	sum of jiffies across all cores spent in softirq	integral	jiffies	time	TRUE	TRUE	1600		16	infinite	uint64
procstat	irq	sum of jiffies across all cores spent in irq	integral	jiffies	time	TRUE	TRUE	1600		16	infinite	uint64
procstat	iowait	sum of jiffies across all cores spent in iowait	integral	jiffies	time	TRUE	TRUE	1600		16	infinite	uint64
procstat	idle	sum of jiffies across all cores spent in idle	integral	jiffies	time	TRUE	TRUE	1600		16	infinite	uint64
procstat	sys	sum of jiffies across all cores spent in sys	integral	jiffies	time	TRUE	TRUE	1600		16	infinite	uint64
procstat	user	sum of jiffies across all cores spent in user	integral	jiffies	time	TRUE	TRUE	1600		16	infinite	uint64
procstat	cpu_enabled#CORE	0 if core CORE is downed	capacity	cores	available	FALSE	FALSE		1	0	1	uint64
procstat	cpu_enabled	0 if core is downed	capacity	cores	available	FALSE	FALSE		1	0	1	uint64
procstat	procs_blocked	processes waiting	capacity	processes	process_count	FALSE			4294967296	0	infinite	uint32
procstat	procs_running	processes active	capacity	processes	process_count	FALSE			4294967296	0	infinite	uint32
procstat	processes	processes started since boot	integral	processes	process_count	FALSE			4294967296	0	rolls	uint32
procstat	context_switches	cpu context switches since boot (or rollover of counter)	integral	event	event	TRUE	TRUE		4294967296	0	rolls	uint32
procstat	hwintr_count	number of hardware interrupts since boot	integral	event	event	TRUE	TRUE		4294967296	0	rolls	uint32
procstat	softirq_count	number of softirq since boot	integral	event	event	TRUE	TRUE		4294967296	0	rolls	uint32

Table A-8. sysclassib and gw_sysclassib reference information

Sampler	Name	description	Mode	Units	Dimension	URO	DJU	LRUB	LUB	itype	kind
sysclassib	ib.port_multicast_rcv_packets#DEVICE	multicast packets received on DEVICE	integral	packets	data	T	T		4294967296	saturates	uint64
sysclassib	ib.port_multicast_xmit_packets#DEVICE	multicast packets sent via DEVICE	integral	packets	data	T	T		4294967296	saturates	uint64
sysclassib	ib.port_unicast_rcv_packets#DEVICE	unicast packets received on DEVICE	integral	packets	data	T	T		4294967296	saturates	uint64
sysclassib	ib.port_unicast_xmit_packets#DEVICE	unicast packets sent via DEVICE	integral	packets	data	T	T		4294967296	saturates	uint64
sysclassib	ib.port_xmit_wait#DEVICE	The number of ticks during which the port selected by PortSelect had data to transmit but no data was sent during the entire tick either because of insufficient credits or because of lack of arbitration.	integral	ticks	event	T	T		4294967296	saturates	uint64
sysclassib	ib.port_rcv_packets#DEVICE	packets received on DEVICE	integral	packets	data	T	T		4294967296	saturates	uint64
sysclassib	ib.port_xmit_packets#DEVICE	packets transmitted onDEVICE	integral	packets	data	T	T		4294967296	saturates	uint64
sysclassib	ib.port_rcv_data#DEVICE	4-byte words received on DEVICE	integral	word	data	T	T	4294967296	4294967296	saturates	uint64
sysclassib	ib.port_xmit_data#DEVICE	4-byte words transmitted on DEVICE	integral	word	data	T	T	4294967296	4294967296	saturates	uint64
sysclassib	ib.VL15_dropped#DEVICE	Number of incoming VL15 packets dropped due to resource limitations (e.g., lack of buffers) in the port.	integral	errors	event	F	T		65536	saturates	uint64
sysclassib	ib.excessive_buffer_overrun_errors#DEVICE	The number of times that OverrunErrors consecutive flow control update periods occurred on DEVICE, each having at least one overrun error.	integral	errors	event	F	T		16	saturates	uint64
sysclassib	ib.local_link_integrity_errors#DEVICE	The number of times that the count of local physical errors on DEVICE exceeded the threshold specified by LocalPhyErrors	integral	errors	event	F	T		16	saturates	uint64
sysclassib	ib.COUNTER_SELECT2_F#DEVICE	junk	label	input	none	F	F		0		uint64
sysclassib	ib.port_rcv_constraint_errors#DEVICE	Total number of packets received on the switch physical port that are discarded for the following reasons: (a) FilterRawInbound is true and packet is raw (b) PartitionEnforcementInbound is true and packet fails partition key check or IP version check.	integral	errors	event	F	T		256	saturates	uint64
sysclassib	ib.port_xmit_constraint_errors#DEVICE	Total number of packets not transmitted from the switch physical port for the following reasons: (a) FilterRawOutbound is true and packet is raw, or (b) PartitionEnforcementOutbound is true and packet fails partition key check or IP version check.	integral	errors	event	F	T		256	saturates	uint64
sysclassib	ib.port_xmit_discards#DEVICE	Total number of DEVICE outbound packets discarded by the port because the port is down or congested. Reasons for this include: (a) Output port is not in the active state, or (b) Packet length exceeded NeighborMTU, or (c) Switch Lifetime Limit exceeded, or (d) Switch HOQ Lifetime Limit exceeded. This may also include packets discarded while in VLStalled State.	integral	packets	event	F	T		65536	saturates	uint64
sysclassib	ib.port_rcv_switch_relay_errors#DEVICE	Total number of packets received on DEVICE that were discarded because they could not be forwarded by the switch relay. Reasons for this include: (a) DLID mapping (see the description of PortDLIDMappingErrors in Table 250 PortRcvErrorDetails on page 1045), or (b) VL mapping, or (c) Looping (output port = input port).	integral	errors	event	F	T		65536	saturates	uint64
sysclassib	ib.port_rcv_remote_physical_errors#DEVICE	Total number of packets marked with the EBP delimiter received on DEVICE.	integral	errors	event	F	T		65536	saturates	uint64
sysclassib	ib.port_rcv_errors#DEVICE	Total number of packets containing an error that were received on DEVICE. These errors include: (a) Local physical errors (ICRC, VCRC, LPCRC, and all physical errors that cause entry into the BAD PACKET or BAD PACKET DISCARD states of the packet receiver state machine), or (b) Malformed data packet errors (LVer, length, VL), or (c) Malformed link packet errors (operand, length, VL), or (d) Packets discarded due to buffer overrun.	integral	errors	event	F	T		65536	saturates	uint64
sysclassib	ib.link_downed#DEVICE	The total number of times the DEVICE Port Training state machine has failed the link error recovery process and downed the link.	integral	fails	event	F	T		256	saturates	uint64
sysclassib	ib.link_error_recovery#DEVICE	The total number of times the DEVICE Port Training state machine has successfully completed the link error recovery process.	integral	errors	event	T	T		256	saturates	uint64
sysclassib	ib.symbol_error#DEVICE	The total number of minor link errors detected on one or more physical lanes. This includes 8B/10B coding violations and is typically an indication of a bit error on the line.	integral	errors	event	T	T		65536	saturates	uint64

Table A-9. vmstat reference information

Sampler	Name	description	Mode	Units	Dimension	URO	DJU	itype	kind
vmstat	thp_split	number of hugepage splits	integral	event	event	T	T	rolls	uint64
vmstat	thp_collapse_alloc_failed	Is incremented if khugepaged found a range of pages that should be collapsed into one huge page but failed the allocation.	integral	event	event	T	T	rolls	uint64
vmstat	thp_collapse_alloc	Is incremented by khugepaged when it has found a range of pages to collapse into one huge page and has successfully allocated a new huge page to store the data.	integral	event	event	T	T	rolls	uint64
vmstat	thp_fault_fallback	Is incremented if a page fault fails to allocate a huge page and instead falls back to using small pages.	integral	event	event	T	T	rolls	uint64
vmstat	thp_fault_alloc	Is incremented every time a huge page is successfully allocated to handle a page fault. This applies to both the first time a page is faulted and for copy-on-write faults.	integral	event	event	T	T	rolls	uint64
vmstat	unevictable_pgs_mlockfreed		integral	event	event	T	T	rolls	uint64
vmstat	unevictable_pgs_stranded		integral	event	event	T	T	rolls	uint64
vmstat	unevictable_pgs_cleared		integral	event	event	T	T	rolls	uint64
vmstat	unevictable_pgs_munlocked		integral	event	event	T	T	rolls	uint64
vmstat	unevictable_pgs_mlocked		integral	event	event	T	T	rolls	uint64
vmstat	unevictable_pgs_rescued		integral	event	event	T	T	rolls	uint64
vmstat	unevictable_pgs_scanned		integral	event	event	T	T	rolls	uint64
vmstat	unevictable_pgs_culled		integral	event	event	T	T	rolls	uint64
vmstat	htlb_buddy_alloc_fail		integral	event	event	T	T	rolls	uint64
vmstat	htlb_buddy_alloc_success		integral	event	event	T	T	rolls	uint64
vmstat	compact_success	is incremented if the system compacted memory and freed a huge page for use.	integral	event	event	T	T	rolls	uint64
vmstat	compact_fail	is incremented if the system tries to compact memory but failed	integral	event	event	T	T	rolls	uint64
vmstat	compact_stall	is incremented every time a process stalls to run memory compaction so that a huge page is free for use.	integral	event	event	T	T	rolls	uint64
vmstat	compact_pagemigrate_failed	is incremented when the underlying mechanism for moving a page failed.	integral	event	event	T	T	rolls	uint64
vmstat	compact_pages_moved	is incremented each time a page is moved. If this value is increasing rapidly, it implies that the system is copying a lot of data to satisfy the huge page allocation. It is possible that the cost of copying exceeds any savings from reduced TLB misses.	integral	event	event	T	T	rolls	uint64
vmstat	compact_blocks_moved	is incremented each time memory compaction examines a huge page aligned range of pages.	integral	event	event	T	T	rolls	uint64
vmstat	pgrotated		integral	event	event	T	T	rolls	uint64
vmstat	allocstall	Number of direct reclaim calls since boot	integral	event	event	T	T	rolls	uint64
vmstat	pageoutrun		integral	event	event	T	T	rolls	uint64
vmstat	kswapd_skip_congestion_wait		integral	event	event	T	T	rolls	uint64
vmstat	kswapd_high_wmark_hit_quickly		integral	event	event	T	T	rolls	uint64
vmstat	kswapd_low_wmark_hit_quickly		integral	event	event	T	T	rolls	uint64
vmstat	kswapd_inodesteal		integral	event	event	T	T	rolls	uint64
vmstat	kswapd_steal	Number of pages reclaimed by kswapd since boot	integral	event	event	T	T	rolls	uint64
vmstat	slabs_scanned	Number of scanned slap objects since boot	integral	event	event	T	T	rolls	uint64
vmstat	pginodesteal	Number of normal pages reclaimed via inode frees since boot	integral	event	event	T	T	rolls	uint64
vmstat	zone_reclaim_failed		integral	event	event	T	T	rolls	uint64
vmstat	pgscan_direct_movable		integral	event	event	T	T	rolls	uint64
vmstat	pgscan_direct_normal	Number of normal pages reclaimed since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgscan_direct_dma32	Number of dma32 page reclaimed since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgscan_direct_dma	Number of dma page reclaimed since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgscan_kswapd_movable		integral	event	event	T	T	rolls	uint64
vmstat	pgscan_kswapd_normal	Number of normal pages scanned by kswapd since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgscan_kswapd_dma32	Number of dma32 page scanned by kswapd since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgscan_kswapd_dma	Number of dma page scanned by kswapd since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgsteal_movable		integral	event	event	T	T	rolls	uint64
vmstat	pgsteal_normal	Number of normal page steals since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgsteal_dma32	Number of dma32 page steals since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgsteal_dma	Number of dma page steals since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgrefill_movable		integral	event	event	T	T	rolls	uint64
vmstat	pgrefill_normal	Number of normal page refills since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgrefill_dma32	Number of dma32 page refills since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgrefill_dma	Number of dma page refills since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgmajfault	Number of major page faults since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgfault	Number of minor page faults since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgdeactivate	Number of page deactivate since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgactivate	Number of page activates since boot	integral	event	event	T	T	rolls	uint64
vmstat	pgfree	Number of page frees since boot	integral	event	event	T	T	rolls	uint64

Table A-10. vmstat reference information (continued)

Sampler	Name	description	Mode	Units	Dimension	URO	DJU	itype	kind
vmstat	pgalloc_movable	movable zone pages in use	integral	pages	data	T	T	rolls	uint64
vmstat	pgalloc_normal	normal zone pages in use	integral	pages	data	T	T	rolls	uint64
vmstat	pgalloc_dma32	dma32 zone pages in use	integral	pages	data	T	T	rolls	uint64
vmstat	pgalloc_dma	dma zone pages in use	integral	pages	data	T	T	rolls	uint64
vmstat	pswpout	swap pages out count	integral	event	event	T	T	rolls	uint64
vmstat	pswpin	swap pages in count	integral	event	event	T	T	rolls	uint64
vmstat	pgpgout	disk pages out since boot- possibly includes all io	integral	event	event	T	T	rolls	uint64
vmstat	pgpgin	disk pages in since boot - possibly includes all io	integral	event	event	T	T	rolls	uint64
vmstat	nr_anon_transparent_hugepages		integral	event	event	T	T	rolls	uint64
vmstat	numa_other	pages allocated in RAM attached to this cpu while code using the pages was running elsewhere.	integral	pages	data	T	T	rolls	uint64
vmstat	numa_local	pages allocated in RAM attached to this cpu while code using the pages runs on the same cpu.	integral	pages	data	T	T	rolls	uint64
vmstat	numa_interleave	page allocated under an interleave strategy	integral	pages	data	T	T	rolls	uint64
vmstat	numa_foreign	small is better- large indicates we borrow RAM some hops away from this core	integral	pages	data	T	T	rolls	uint64
vmstat	numa_miss	small rate is better- large indicate other cores borrow locally attached RAM	integral	pages	data	T	T	rolls	uint64
vmstat	numa_hit	large rate is ok- this core allocated nearest ram HW	integral	pages	data	T	T	rolls	uint64
vmstat	nr_shmem		capacity	pages	data	F	F	rolls	uint64
vmstat	nr_isolated_file		capacity	pages	data	F	F	rolls	uint64
vmstat	nr_isolated_anon		capacity	pages	data	F	F	rolls	uint64
vmstat	nr_writeback_temp	small is better	capacity	pages	data	F	F	rolls	uint64
vmstat	nr_vmscan_write		capacity	pages	data	F	F	rolls	uint64
vmstat	nr_bounce		capacity	pages	data	F	F	rolls	uint64
vmstat	nr_unstable	unstable page count	capacity	pages	data	F	F	rolls	uint64
vmstat	nr_kernel_stack	small is better	capacity	pages	data	F	F	rolls	uint64
vmstat	nr_page_table_pages	small is better	capacity	pages	data	F	F	rolls	uint64
vmstat	nr_slab_unreclaimable	small is better	capacity	pages	data	F	F	rolls	uint64
vmstat	nr_slab_reclaimable	small is better	capacity	pages	data	F	F	rolls	uint64
vmstat	nr_writeback	count of pages scheduled out but not done. persistent value should be 0.	capacity	pages	data	F	F	rolls	uint64
vmstat	nr_dirty	count of pages waiting to be scheduled to output device. persistent value should be 0	capacity	pages	data	F	F	rolls	uint64
vmstat	nr_file_pages		capacity	pages	data	F	F	rolls	uint64
vmstat	nr_mapped	Number of pages mapped by files	capacity	pages	data	F	F	rolls	uint64
vmstat	nr_anon_pages		capacity	pages	data	F	F	rolls	uint64
vmstat	nr_mlock		capacity	pages	data	F	F	rolls	uint64
vmstat	nr_unevictable		capacity	pages	data	F	F	rolls	uint64
vmstat	nr_active_file		capacity	pages	data	F	F	rolls	uint64
vmstat	nr_inactive_file		capacity	pages	data	F	F	rolls	uint64
vmstat	nr_active_anon		capacity	pages	data	F	F	rolls	uint64
vmstat	nr_inactive_anon		capacity	pages	data	F	F	rolls	uint64
vmstat	nr_free_pages		capacity	pages	data	F	F	rolls	uint64

DISTRIBUTION

Email—Internal (encrypt for OUO)

Name	Org.	Sandia Email Address
Ann Gentile	9328	gentile@sandia.gov
Technical Library	01177	libref@sandia.gov

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.