NAME

sinfo − view information about Slurm nodes and partitions.

SYNOPSIS

sinfo [*OPTIONS*...]

DESCRIPTION

sinfo is used to view partition and node information for a system running Slurm.

OPTIONS

−a, **−−all**
Display information about all partitions. This causes information to be displayed about partitions that are configured as hidden and partitions that are unavailable to user's group.

−d, **−−dead**
If set only report state information for non−responding (dead) nodes.

−e, **−−exact**
If set, do not group node information on multiple nodes unless their configurations to be reported are identical. Otherwise cpu count, memory size, and disk space for nodes will be listed with the minimum value followed by a "+" for nodes with the same partition and state (e.g., "250+").

−−federation
Show all partitions from the federation if a member of one.

−h, **−−noheader**
Do not print a header on the output.

−−help Print a message describing all **sinfo** options.

−−hide Do not display information about hidden partitions. By default, partitions that are configured as hidden or are not available to the user's group will not be displayed (i.e. this is the default behavior).

−i <seconds>, **−−iterate=<seconds>**
Print the state on a periodic basis. Sleep for the indicated number of seconds between reports. By default, prints a time stamp with the header.

−−local
Show only jobs local to this cluster. Ignore other clusters in this federation (if any). Overrides −−federation.

−l, **−−long**
Print more detailed information. This is ignored if the **−−format** option is specified.

−M, **−−clusters**=*<string>*
Clusters to issue commands to. Multiple cluster names may be comma separated. A value of of '*all*' will query to run on all clusters. Note that the SlurmDBD must be up for this option to work properly. This option implicitly sets the **−−local** option.

−n <nodes>, **−−nodes=<nodes>**
Print information only about the specified node(s). Multiple nodes may be comma separated or expressed using a node range expression. For example "linux[00−07]" would indicate eight

nodes, "linux00" through "linux07." Performance of the command can be measurably improved for systems with large numbers of nodes when a single node name is specified.

−−noconvert
Don't convert units from their original type (e.g. 2048M won't be converted to 2G).

−N, **−−Node**
Print information in a node−oriented format with one line per node and partition. That is, if a node belongs to more than one partition, then one line for each node−partition pair will be shown. If **−−partition** is also specified, then only one line per node in this partition is shown. The default is to print information in a partition−oriented format. This is ignored if the **−−format** option is specified.

−o <output_format>, **−−format=<output_format>**
Specify the information to be displayed using an **sinfo** format string. Format strings transparently used by **sinfo** when running with various options are

default	"%#P %.5a %.10l %.6D %.6t %N"
−−summarize	"%#P %.5a %.10l %.16F %N"
−−long	"%#P %.5a %.10l %.10s %.4r %.8h %.10g %.6D %.11T %N"
−−Node	"%#N %.6D %#P %6t"
−−long −−Node	"%#N %.6D %#P %.11T %.4c %.8z %.6m %.8d %.6w %.8f %20E"
−−list−reasons	"%20E %9u %19H %N"
−−long −−list−reasons	"%20E %12U %19H %6t %N"

In the above format strings, the use of "#" represents the maximum length of any partition name or node list to be printed. A pass is made over the records to be printed to establish the size in order to align the sinfo output, then a second pass is made over the records to print them. Note that the literal character "#" itself is not a valid field length specification, but is only used to document this behaviour.

The field specifications available include:

%all Print all fields available for this data type with a vertical bar separating each field.

%a State/availability of a partition

%A Number of nodes by state in the format "allocated/idle". Do not use this with a node state option ("%t" or "%T") or the different node states will be placed on separate lines.

%b Features currently active on the nodes, also see **%f**

%B The max number of CPUs per node available to jobs in the partition.

%c Number of CPUs per node

%C Number of CPUs by state in the format "allocated/idle/other/total". Do not use this with a node state option ("%t" or "%T") or the different node states will be placed on separate lines.

%d Size of temporary disk space per node in megabytes

%D Number of nodes

%e Free memory of a node

%E The reason a node is unavailable (down, drained, or draining states).

%f Features available the nodes, also see **%b**

%F Number of nodes by state in the format "allocated/idle/other/total". Note the use of this format option with a node state format option ("%t" or "%T") will result in the

different node states being be reported on separate lines.

%g Groups which may use the nodes

%G Generic resources (gres) associated with the nodes

%h Jobs may oversubscribe compute resources (i.e. CPUs), "yes", "no", "exclusive" or "force"

%H Print the timestamp of the reason a node is unavailable.

%I Partition job priority weighting factor.

%l Maximum time for any job in the format "days−hours:minutes:seconds"

%L Default time for any job in the format "days−hours:minutes:seconds"

%m Size of memory per node in megabytes

%M PreemptionMode

%n List of node hostnames

%N List of node names

%o List of node communication addresses

%O CPU load of a node

%p Partition scheduling tier priority.

%P Partition name followed by "*" for the default partition, also see **%R**

%r Only user root may initiate jobs, "yes" or "no"

%R Partition name, also see **%P**

%s Maximum job size in nodes

%S Allowed allocating nodes

%t State of nodes, compact form

%T State of nodes, extended form

%u Print the user name of who set the reason a node is unavailable.

%U Print the user name and uid of who set the reason a node is unavailable.

%v Print the version of the running slurmd daemon.

%V Print the cluster name if running in a federation

%w Scheduling weight of the nodes

%X Number of sockets per node

%Y Number of cores per socket

%Z Number of threads per core

%z Extended processor information: number of sockets, cores, threads (S:C:T) per node

%.<*>
right justification of the field

%<Number><*>
size of field

−O <output_format>, **−−Format=<output_format>**
Specify the information to be displayed. Also see the **−o <output_format>**, **−−format=<output_format>** option described below (which supports greater flexibility in formatting, but does not support access to all fields because we ran out of letters). Requests a comma separated list of job information to be displayed.

The format of each field is "type[:[.]size]"

size is the minimum field size. If no size is specified, 20 characters will be allocated to print the information.

. indicates the output should be right justified and size must be specified. By default, output is left justified.

Valid *type* specifications include:

all Print all fields available in the −o format for this data type with a vertical bar separating each field.

allocmem
Prints the amount of allocated memory on a node.

allocnodes
Allowed allocating nodes.

available
State/availability of a partition.

cluster
Print the cluster name if running in a federation

cpus Number of CPUs per node.

cpusload
CPU load of a node.

freemem
Free memory of a node.

cpusstate
Number of CPUs by state in the format "allocated/idle/other/total". Do not use this with a node state option ("%t" or "%T") or the different node states will be placed on separate lines.

cores Number of cores per socket.

defaulttime
Default time for any job in the format "days−hours:minutes:seconds".

disk Size of temporary disk space per node in megabytes.

features
Features available on the nodes. Also see **features_act**.

features_act
Features currently active on the nodes. Also see **features**.

groups
Groups which may use the nodes.

gres Generic resources (gres) associated with the nodes.

gresused
Generic resources (gres) currently in use on the nodes.

maxcpuspernode
The max number of CPUs per node available to jobs in the partition.

memory
Size of memory per node in megabytes.

nodes Number of nodes.

nodeaddr
List of node communication addresses.

nodeai
Number of nodes by state in the format "allocated/idle". Do not use this with a node state option ("%t" or "%T") or the different node states will be placed on separate lines.

nodeaiot
Number of nodes by state in the format "allocated/idle/other/total". Do not use this with a node state option ("%t" or "%T") or the different node states will be placed on separate lines.

nodehost
List of node hostnames.

nodelist
List of node names.

oversubscribe
Jobs may oversubscribe compute resources (i.e. CPUs), "yes", "no", "exclusive" or "force".

partition
Partition name followed by "*" for the default partition, also see **%R**.

partitionname
Partition name, also see **%P**.

port Node TCP port.

preemptmode
PreemptionMode.

priorityjobfactor
Partition factor used by priority/multifactor plugin in calculating job priority.

prioritytier or **priority**
Partition scheduling tier priority.

reason
The reason a node is unavailable (down, drained, or draining states).

root Only user root may initiate jobs, "yes" or "no".

size Maximum job size in nodes.

statecompact
State of nodes, compact form.

statelong
State of nodes, extended form.

sockets
Number of sockets per node.

socketcorethread
Extended processor information: number of sockets, cores, threads (S:C:T) per node.

time Maximum time for any job in the format "days−hours:minutes:seconds".

timestamp
Print the timestamp of the reason a node is unavailable.

threads
Number of threads per core.

user Print the user name of who set the reason a node is unavailable.

userlong
Print the user name and uid of who set the reason a node is unavailable.

version
Print the version of the running slurmd daemon.

weight
Scheduling weight of the nodes.

−p <partition>, **−−partition=<partition>**
Print information only about the specified partition(s). Multiple partitions are separated by commas.

−r, **−−responding**
If set only report state information for responding nodes.

−R, **−−list−reasons**
List reasons nodes are in the down, drained, fail or failing state. When nodes are in these states Slurm supports optional inclusion of a "reason" string by an administrator. This option will display the first 20 characters of the reason field and list of nodes with that reason for all nodes that are, by default, down, drained, draining or failing. This option may be used with other node filtering options (e.g. **−r**, **−d**, **−t**, **−n**), however, combinations of these options that result in a list of nodes that are not down or drained or failing will not produce any output. When used with **−l** the output additionally includes the current node state.

−s, **−−summarize**
List only a partition state summary with no node state details. This is ignored if the **−−format** option is specified.

−S <sort_list>, **−−sort=<sort_list>**
Specification of the order in which records should be reported. This uses the same field specification as the <output_format>. Multiple sorts may be performed by listing multiple sort fields separated by commas. The field specifications may be preceded by "+" or "−" for ascending (default) and descending order respectively. The partition field specification, "P", may be preceded by a "#" to report partitions in the same order that they appear in Slurm's configuration file, **slurm.conf**. For example, a sort value of "+P,−m" requests that records be printed in order of increasing partition name and within a partition by decreasing memory size. The default value of sort is "#P,−t" (partitions ordered as configured then decreasing node state). If the **−−Node option is selected, the default sort value is "N" (increasing node name).**

−t <states> , **−−states=<states>**
List nodes only having the given state(s). Multiple states may be comma separated and the comparison is case insensitive. Possible values include (case insensitive): ALLOC, ALLOCATED, COMP, COMPLETING, DOWN, DRAIN (for node in DRAINING or DRAINED states), DRAINED, DRAINING, FAIL, FUTURE, FUTR, IDLE, MAINT, MIX, MIXED, NO_RESPOND, NPC, PERFCTRS, POWER_DOWN, POWERING_DOWN, POWER_UP, RESV, RESERVED, UNK, and UNKNOWN. By default nodes in the specified state are reported whether they are responding or not. The **−−dead** and **−−responding** options may be used to filtering nodes by the responding flag.

−T, **−−reservation**
Only display information about Slurm reservations.

−−usage
Print a brief message listing the **sinfo** options.

−v, **−−verbose**
Provide detailed event logging through program execution.

−V, **−−version**
Print version information and exit.

OUTPUT FIELD DESCRIPTIONS

AVAIL Partition state: **up** or **down**.

CPUS Count of CPUs (processors) on each node.

S:C:T Count of sockets (S), cores (C), and threads (T) on these nodes.

SOCKETS
Count of sockets on these nodes.

CORES
Count of cores on these nodes.

THREADS
Count of threads on these nodes.

GROUPS
Resource allocations in this partition are restricted to the named groups. **all** indicates that all groups may use this partition.

JOB_SIZE
Minimum and maximum node count that can be allocated to any user job. A single number indicates the minimum and maximum node count are the same. **infinite** is used to identify partitions without a maximum node count.

TIMELIMIT
Maximum time limit for any user job in days−hours:minutes:seconds. **infinite** is used to identify partitions without a job time limit.

MEMORY
Size of real memory in megabytes on these nodes.

NODELIST
Names of nodes associated with this configuration/partition.

NODES
Count of nodes with this particular configuration.

NODES(A/I)
Count of nodes with this particular configuration by node state in the form "available/idle".

NODES(A/I/O/T)
Count of nodes with this particular configuration by node state in the form "available/idle/other/total".

PARTITION
Name of a partition. Note that the suffix "*" identifies the default partition.

PORT Local TCP port used by slurmd on the node.

ROOT Is the ability to allocate resources in this partition restricted to user root, **yes** or **no**.

OVERSUBSCRIBE
Will jobs allocated resources in this partition oversubscribe those compute resources (i.e. CPUs). **no** indicates resources are never oversubscribed. **exclusive** indicates whole nodes are dedicated to jobs (equivalent to srun −−exclusive option, may be used even with select/cons_res managing individual processors). **force** indicates resources are always available to be oversubscribed. **yes** indicates resource may be oversubscribed or not per job's resource allocation.

STATE State of the nodes. Possible states include: allocated, completing, down, drained, draining, fail, failing, future, idle, maint, mixed, perfctrs, power_down, power_up, reserved, and unknown plus Their abbreviated forms: alloc, comp, down, drain, drng, fail, failg, futr, idle, maint, mix, npc, pow_dn, pow_up, resv, and unk respectively. Note that the suffix "*" identifies nodes that are presently not responding.

TMP_DISK
Size of temporary disk space in megabytes on these nodes.

NODE STATE CODES

Node state codes are shortened as required for the field size. These node states may be followed by a special character to identify state flags associated with the node. The following node sufficies and states are used:

***** The node is presently not responding and will not be allocated any new work. If the node remains non−responsive, it will be placed in the **DOWN** state (except in the case of **COMPLETING**, **DRAINED**, **DRAINING**, **FAIL**, **FAILING** nodes).

~ The node is presently in a power saving mode (typically running at reduced frequency).

The node is presently being powered up or configured.

% The node is presently being powered down.

$ The node is currently in a reservation with a flag value of "maintenance".

@ The node is pending reboot.

ALLOCATED
The node has been allocated to one or more jobs.

ALLOCATED+
The node is allocated to one or more active jobs plus one or more jobs are in the process of COMPLETING.

COMPLETING
All jobs associated with this node are in the process of COMPLETING. This node state will be removed when all of the job's processes have terminated and the Slurm epilog program (if any) has terminated. See the **Epilog** parameter description in the **slurm.conf** man page for more information.

DOWN The node is unavailable for use. Slurm can automatically place nodes in this state if some failure occurs. System administrators may also explicitly place nodes in this state. If a node resumes normal operation, Slurm can automatically return it to service. See the **ReturnToService** and **SlurmdTimeout** parameter descriptions in the **slurm.conf**(5) man page for more information.

DRAINED The node is unavailable for use per system administrator request. See the **update node** command in the **scontrol**(1) man page or the **slurm.conf**(5) man page for more information.

DRAINING The node is currently executing a job, but will not be allocated to additional jobs. The node state will be changed to state **DRAINED** when the last job on it completes. Nodes enter this state per system administrator request. See the **update node** command in the **scontrol**(1) man page or the **slurm.conf**(5) man page for more information.

FAIL The node is expected to fail soon and is unavailable for use per system administrator request. See the **update node** command in the **scontrol**(1) man page or the **slurm.conf**(5) man page for more information.

FAILING The node is currently executing a job, but is expected to fail soon and is unavailable for use per system administrator request. See the **update node** command in the **scontrol**(1) man page or the **slurm.conf**(5) man page for more information.

FUTURE The node is currently not fully configured, but expected to be available at some point in the indefinite future for use.

IDLE The node is not allocated to any jobs and is available for use.

MAINT The node is currently in a reservation with a flag value of "maintainence".

REBOOT The node is currently scheduled to be rebooted.

MIXED The node has some of its CPUs ALLOCATED while others are IDLE.

PERFCTRS (NPC)
Network Performance Counters associated with this node are in use, rendering this node as not usable for any other jobs

POWER_DOWN
The node is currently powered down and not capable of running any jobs.

POWERING_DOWN
The node is currently powering down and not capable of running any jobs.

POWER_UP
The node is currently in the process of being powered up.

RESERVED The node is in an advanced reservation and not generally available.

UNKNOWN The Slurm controller has just started and the node's state has not yet been determined.

ENVIRONMENT VARIABLES

Some **sinfo** options may be set via environment variables. These environment variables, along with their corresponding options, are listed below. (Note: Commandline options will always override these settings.)

SINFO_ALL **−a, −−all**

SINFO_FEDERATION
Same as **−−federation**

SINFO_FORMAT **−o <output_format>, −−format=<output_format>**

SINFO_LOCAL Same as **−−local**

SINFO_PARTITION **−p <partition>, −−partition=<partition>**

SINFO_SORT **−S <sort>, −−sort=<sort>**

SLURM_CLUSTERS Same as **−−clusters**

SLURM_CONF The location of the Slurm configuration file.

SLURM_TIME_FORMAT
Specify the format used to report time stamps. A value of *standard*, the default value, generates output in the form "year−month−dateThour:minute:second". A value of *relative* returns only "hour:minute:second" if the current day. For other dates in the current year it prints the "hour:minute" preceded by "Tomorr" (tomorrow), "Ystday" (yesterday), the name of the day for the coming week (e.g. "Mon", "Tue", etc.), otherwise the date (e.g. "25 Apr"). For other years it returns a date month and year without a time (e.g. "6 Jun 2012"). All of the time stamps use a 24 hour format.

A valid strftime() format can also be specified. For example, a value of "%a %T" will report the day of the week and a time stamp (e.g. "Mon 12:34:56").

EXAMPLES

Report basic node and partition configurations:

```
> sinfo
PARTITION AVAIL TIMELIMIT NODES STATE  NODELIST
batch    up    infinite    2 alloc  adev[8-9]
batch    up    infinite    6 idle   adev[10-15]
debug*   up       30:00    8 idle   adev[0-7]
```

Report partition summary information:

```
> sinfo -s
PARTITION AVAIL TIMELIMIT NODES(A/I/O/T) NODELIST
batch    up    infinite 2/6/0/8        adev[8-15]
debug*   up       30:00 0/8/0/8        adev[0-7]
```

Report more complete information about the partition debug:

```
> sinfo --long --partition=debug
PARTITION AVAIL TIMELIMIT JOB_SIZE ROOT OVERSUBS GROUPS NODES STATE NODELIST
debug*    up       30:00        8 no   no       all       8 idle  dev[0-7]
```

Report only those nodes that are in state DRAINED:

```
> sinfo --states=drained
PARTITION AVAIL NODES TIMELIMIT STATE  NODELIST
debug*    up        2     30:00 drain  adev[6-7]
```

Report node-oriented information with details and exact matches:

```
> sinfo -Nel
NODELIST    NODES PARTITION STATE  CPUS MEMORY TMP_DISK WEIGHT FEATURES REASON
adev[0-1]       2 debug*    idle      2   3448    38536     16 (null)   (null)
adev[2,4-7]     5 debug*    idle      2   3384    38536     16 (null)   (null)
adev3           1 debug*    idle      2   3394    38536     16 (null)   (null)
adev[8-9]       2 batch     allocated 2    246    82306     16 (null)   (null)
adev[10-15]     6 batch     idle      2    246    82306     16 (null)   (null)
```

Report only down, drained and draining nodes and their reason field:

```
> sinfo -R
REASON                    NODELIST
Memory errors              dev[0,5]
Not Responding             dev8
```

COPYING

SEE ALSO

scontrol(1), **smap**(1), **squeue**(1), **slurm_load_ctl_conf** (3), **slurm_load_jobs** (3), **slurm_load_node** (3), **slurm_load_partitions** (3), **slurm_reconfigure** (3), **slurm_shutdown** (3), **slurm_update_job** (3), **slurm_update_node** (3), **slurm_update_partition** (3), **slurm.conf**(5)