

Solaris Performance Metrics

The Uptime Infrastructure Monitor Solaris agent uses a number of utilities resident on a Solaris server to collect performance metrics including `sar`, `awk`, `grep` and `ps`. Note that the `sadc` utility (which collects system activity information and saves it as a binary file) is used by `sar` and must be readable by the Uptime Infrastructure Monitor agent user account on the system.

The Uptime Infrastructure Monitor Solaris agent collects the following performance metrics from the systems on which it is installed:

- CPU
- Multi-CPU
- Memory
- Disk
- Network
- Process
- Workload
- Veritas Volume Manager
- User

Whenever the `sar` command uses the `-f` option to specify a file, that file is generated using the `sadc 1 1` command. The `sadc` command polls the system counters at a one-second interval and writes the information that it receives to a file. The `sar` command then reads this file.

Each set of performance metrics is averaged between the interval at which the Uptime Infrastructure Monitor monitoring station polls the agent (e.g. every 10 minutes).

CPU

The Uptime Infrastructure Monitor agent uses the `sar -u -f` command to collect CPU metrics from a Solaris system. The statistics returned by the agent are averaged for all CPUs on the system and the `sar` command compares the system counters during a one-second interval. In multi-CPU systems, the CPU statistics output by the agent are an average of all the CPUs on the server.

| Metric | Explanation |
|---------------------|--|
| %USR | The percentage of time that the processor spends in user mode (a processing mode for applications and subsystems). |
| %SYS | The percentage of time that the kernel spends processing system calls. |
| %WIO | The amount of waiting time that a runnable process for a device takes to perform an I/O operation. |
| %Total | The total amount of User %, System %, and Wait I/O % |
| Run Queue Length | The percentage of time that one or more services or processes are waiting to be served by the CPU. |
| Run Queue Occupancy | The percentage of time that one or more services or processes are waiting to be served by the CPU. |

Multi-CPU

The Uptime Infrastructure Monitor agent outputs statistics for the entire Solaris system, per CPU. The `sar` command (`mpstat 1 2`) averages the statistics for each CPU and compares the system counters during a one-second interval

| Me tric | Explanation |
|--------------|---|
| Us er % | The percentage of CPU user processes that are in use. |
| Sy ste m % | The percentage of CPU kernel processes that are in use. |
| Wa it I /O % | The percentage of time that a process which can be run must wait for a device to perform an I/O operation. |
| SM TX | The number of read or write locks that a thread was not able to acquire on the first attempt, as reported by the <code>mpstat</code> command. |
| XC AL | The number of interprocess cross-calls. In a multi-processor environment, one processor sends cross-calls to another processor to get that processor to do work. Cross-calls can also be used to ensure consistency in virtual memory. Heavy file system activity (such as NFS) can result in a high number of cross-calls. |

| | |
|--------------------|--|
| Int err upts | The number of CPU interrupts. |
| Tot al % | The total amount of User %, System %, and Wait I/O%. |

Memory

The Uptime Infrastructure Monitor agent uses the `vmstat 1 2` command to average the metrics that are collected for the entire system. The agent also compares system counters during a one-second interval using the `sar` command and the following options:

- `-b -f` (cache metrics)
- `-g -f` and `-p -f` (paging activity)
- `-q -f` (the average queue length while it is occupied, and the percentage of time that the queue is occupied)
- `-c -f` (system calls)

The free swap metric is collected from the "available" value in the `swap -s` command.

| Metric | Explanation |
|---------------------|--|
| Free Memory | The amount of physical memory available to the operating system, system library files, and applications. |
| Cache Hit Rate | How often the system accesses the CPU cache. |
| Page-outs/s | The rate at which pages were written to disk. |
| Page-ins/s | The rate at which pages were read from or written to the disk. |
| Page Free/s | The number of pages that are freed from memory each second. |
| Attaches/s | The number of pages that get attached to memory each second. |
| Page-out Requests/s | The number of requests to perform a write operation that occur each second. |
| Page-in reqs/s | The number of requests to perform a read operation that occur each second. |
| PageScans/s | The number of pages that are scanned each second. |
| PageFaults/s | The number of page faults that occur each second. |
| Software Locks/s | The number of software locks that are issued each second. |
| Virtual Faults/s | The number of virtual memory faults that occur each second. |
| Free Swap | The amount of available free swap space, as a percentage of total available free swap space. |

Disk

The Uptime Infrastructure Monitor agent uses the following commands to collect disk statistics:

- `df -lk` to gather file system capacity statistics, for file system.
- `sar -d -f` to output disk statistics (e.g. %busy, Read/Write/s) per disk, and compare those statistics between polling intervals.

| Metric | Explanation |
|----------------------|---|
| Disk (Spindle) Name | The names of each disk on the system. |
| Usage (% Busy) | The percentage of time during which the disk drive is handling read or write requests. |
| Throughput (Blk/s) | The number of read and write operations on the disk that occur each second. |
| Read/Writes/s | The average number of bytes that have been transferred to or from the disk during write or read operations. |
| Average Queue Length | The number of threads that are waiting for processor time. |
| Average Service Time | The average amount of time, in milliseconds, that is required for a request to be carried out. |
| Average Wait Time | The average time, in milliseconds, that a transaction is waiting in a queue. The wait time is directly proportional to the length of the queue. |

Network

The Uptime Infrastructure Monitor agent uses the `netstat -s` command to collect network metrics from a Solaris server. Except for TCP retransmits, the agent averages all statistics per interface. Other statistics (e.g. kbps, errors and collisions) are collected per interface by the `kstat` command.

| Metric | Explanation |
|-----------------|--|
| In Kbps | The rate, in kilobytes per seconds, at which data is received over a specific network adapter. |
| Out Kbps | The rate, in kilobytes per seconds, at which data is sent over a specific network adapter. |
| In Errors | The number of inbound packets that contained errors, which preventing those packets from being delivered to a higher-layer protocol. |
| Out Errors | The number of outbound packets that could not be transmitted because of errors. |
| Collisions | The number of signals from two separate nodes on the network that have collided. |
| TCP Retransmits | The number of packets that have been re-sent over a network interface. |

Process

The Uptime Infrastructure Monitor agent gathers process information directly from the `/proc` filesystem using the `procfs` command.

| Metric | Explanation |
|---------------------------------|--|
| PID | The unique identifier of a specific process. |
| PPID | The identifier of the process that the process that is currently running. |
| UID | A value that identifies the current user. |
| GID | A value that identifies a group of users. |
| Memory Consumed | The amount of memory that is being used by a process. |
| RSS | The amount of physical memory that is being used by a process. |
| CPU % Utilization by Process | The percentage of CPU time that is being used by individual processes. |
| Memory % Utilization by Process | The amount of physical memory that is being used by individual processes. |
| Process Start Time | The time at which the process started. |
| Process Run Time | The time at which the process started. |
| Number of Processes Running | The total number of processes that are currently running on the system. |
| Number of Blocked Processes | The total number of processes that are blocking resources. |
| Number of Waiting Processes | The total number of processes that are waiting to be executed by the CPU. |
| Execs per Second | The total number of system calls that are executed each second. |
| Process Creation Rate | The total number of processes that are being spawned over a specified time period. |

Workload

Workload statistics are sorted within Uptime Infrastructure Monitor's core but are the same 20 processes that were gathered from the Process method ([see above](#)). The workload processes gathered by the agent include user/group/process name and their individual statistics. The Uptime Infrastructure Monitor core then sorts based on the selected graph (e.g. user, group or process name).

| Metric | Explanation |
|---------------------|--|
| Workload by Process | The demand that network and local services are putting on a system, based on the processes that are running. |
| Workload by User | The demand that network and local services are putting on the system, based on the IDs of the users who are logged into a system. |
| Workload by Group | The demand that network and local services are putting on the system, based on the IDs of the user groups that are logged into a system. |

| | |
|----------------------------|---|
| Workload Top 10 by Process | The 10 processes that are consuming the most CPU resources. |
| Workload Top 10 by User | The 10 processes the are consuming the most CPU resources, based on user ID. |
| Workload Top 10 by Group | The 10 processes the are consuming the most CPU resources, based on group ID. |

Veritas Volume Manager

The Uptime Infrastructure Monitor agent uses the `vxdg list` command to collect statistics from disk volumes that are managed by the Veritas Volume Manager. These statistics are gathered for each volume, first by retrieving the contents of the disk groups using the `vxdisk list` command, and then by collecting statistics using the `vxstat -g <diskgroup>` command.

| Metric | Explanation |
|----------------------|---|
| DG/Volume/Subdisk | The name of the disk, volume, or subdisk. |
| I/O Operations | The number of times, per second, that data is written to and read from the volume being managed by Veritas Volume Manager. |
| Block Throughput | The amount of disk traffic, in blocks of 512 bytes, that is flowing to and from the volume being managed by Veritas Volume Manager. |
| Average Service Time | The average amount of time, in milliseconds, that is required for a request to be carried out. |

User

The Uptime Infrastructure Monitor agent uses the following commands to collect user statistics from a Solaris system:

- `ps -eo`
- `last | head 10` (login history for the last 10 users on the system)
- `who` (lists who is currently logged into the system)

| Metric | Explanation |
|---------------|---|
| Login History | The number of times or frequency at which a user has logged into a system during any 30 minute time interval. |
| Sessions | The number of sessions or number of distinct users who are logged into a system during any 30 minute time interval. |