

Difference Points for Commands

Revision 2.4

Revision History

| Rev. | Date | Updates / Remarks |
|------|---------------|---|
| 1.3 | 28-Feb-2018 | the first version |
| 1.4 | 24-May-2018 | Changelog: <ul style="list-style-type: none"> - 'sar and sadf' commands will show VEOS RESTART instead of LINUX RESTART log message. - Added difference point for newly ported "ipcs & ipcrm" commands. - "ve_sysstat" service will restart only for the node who's VEOS is restarted. - Added 'blocked processes' related difference point in 'vmstat' and 'sar' command. - |
| 1.5 | 20-June-2018 | Changelog: <ul style="list-style-type: none"> - Updated difference point for psacct-ve service to handle it node wise. |
| 1.6 | 08-Feb-2019 | This revision covers VEOS v2.0.1 or later. Changelog: <ul style="list-style-type: none"> - Changed the format of top page. |
| 1.7 | 15-April-2019 | This revision covers VEOS v2.1 or later. Changelog: <ul style="list-style-type: none"> - Updated difference points in multiple commands to use default VE node 0. - Updated difference point to allow the execution of only VE binary using taskset, prlimit, time and strace command. - Updates in dump-acct command related to PPID value in process accounting file. - Updates in dump-acct and sa command related to elapsed time value. |
| 1.8 | July-2019 | This revision covers VEOS v2.1.3 or later. Changelog: <ul style="list-style-type: none"> - Added difference point in 'strace' command to mention behavior when execve() system call is invoked from traced process. - Removed difference point of dump-acct command related to PPID value in process accounting file. - Changed the format of Revision History. |
| 1.9 | May-2020 | This revision covers VEOS v2.5 or later. Changelog: <ul style="list-style-type: none"> - Added difference point in 'strace' command to mention non supported options for RHEL8.1. |

| | | |
|-----|-----------|--|
| | | <ul style="list-style-type: none"> - Added VE sysstat service related difference point for RHEL8.1. |
| 2.0 | July-2020 | <p>This revision covers VEOS v2.6.2 or later.</p> <p>Changelog:</p> <ul style="list-style-type: none"> - Added scenarios in 'dump-acct' and 'lastcomm' commands where controlling terminal (tty) value will be 'null'. - Added difference point in 'dump-acct' and 'lastcomm' command which shows additional vector information. - Added 'convert-acct' tool details in 'dump-acct', 'lastcomm' and 'sa' command which is used to read the accounting file which have different file format (version 3/version 14) records. |
| 2.1 | Aug-2020 | <ul style="list-style-type: none"> - Fix errors of 3.Enhanced process accounting |
| 2.2 | Sep-2020 | <p>This revision covers VEOS v2.7 or later.</p> <ul style="list-style-type: none"> - Added difference point in 'ps', 'top', 'w' and 'pidstat' command related to abnormal termination. - Updated the description of field 'NUMA' showed by 'dump-acct' and 'lastcomm' |
| 2.3 | Oct-2020 | <p>This revision covers VEOS v2.7.2 or later.</p> <ul style="list-style-type: none"> - Added the description of difference point in 'ps' about RSS |
| 2.4 | Sep-2021 | <p>This revision covers VEOS v2.9.1 or later.</p> <ul style="list-style-type: none"> - Added difference point in 'strace' command to mention behavior of tracee when tracer (strace command) killed with SIGKILL - Updates in 'uname' and 'arch' commands related to show information corresponding to first online node |

1. Introduction

This document aims at listing down all the differences between ported commands for VE and as-is VH commands.

2. List of difference points in commands

Following are the difference points between ported commands for VE and as-is VH commands:

| Package Name | Command Name | Difference Point | Reason |
|-------------------|--------------|---|---|
| coreutils-arch-ve | uname | In case of VE, the environment variable VE_NODE_NUMBER can be given: <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to first online VE node. | There are multiple nodes in VE architecture. |
| coreutils-arch-ve | arch | In case of VE, the environment variable VE_NODE_NUMBER can be given: <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to first online VE node. | There are multiple nodes in VE architecture. |
| coreutils-ve | nproc | In case of VE, the environment variable VE_NODE_NUMBER can be given: <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. | There are multiple nodes in VE architecture. |
| time-ve | time | 1. In case of VE, the environment variable VE_NODE_NUMBER can be given: <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Runs the specified program on given node. | 1. There are multiple nodes in VE architecture. |

| | | | |
|------------|---------|---|--|
| | | <ul style="list-style-type: none"> - VE_NODE_NUMBER is not set: Runs the specified program on VE node 0. <p>2. VE 'time' command is restricted to execute only VE binary and VH binary cannot be executed directly using this command.</p> <p>3. The following values will not be applicable for VE, so the values will be zero:</p> <ul style="list-style-type: none"> - Total number of CPU-seconds that the process spent in kernel mode - Number of major page faults that occurred while the process was running - Number of minor page faults - Number of times the process was swapped out of main memory | <p>2. VH process execution using VE commands are not allowed.</p> <p>3. VE architecture do not support the given fields.</p> |
| sysstat-ve | pidstat | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. - VE_NODE_NUMBER is not set and command executed with interval: Command shows the information corresponding to default VE node 0. <p>2. The command <code>"/opt/nec/ve/bin/pidstat <interval>"</code> will consider interval value if it is able to fetch the required information in given interval time period. Else it will ignore the interval value.</p> <p>3. The command <code>"/opt/nec/ve/bin/pidstat <interval>"</code></p> | <p>1. There are multiple nodes in VE architecture.</p> <p>2. Ported command retrieves the information from VEOS via IPC, which takes more time as compared to x86_64</p> <p>3. In case of VE, the value of "user" will be updated as per</p> |

| | | | |
|--|--|---|--|
| | | <p>can sometimes display more than or less than 100% value in “%usr” field while processes running on all the VE cores</p> <p>4. The command “/opt/nec/ve/bin/pidstat –p SELF” will not display statistics.</p> <p>5. The following values will not be applicable for VE, so the values will be zero:</p> <ul style="list-style-type: none"> - %system : Percentage of CPU used by the task while executing at the system level, kernel - %guest : Percentage of CPU spent by the task in virtual machine (running a virtual processor). - minflt/s : Number of minor faults the task has made per second - minflt-nr : Minor faults made by the task and all its children, and collected during the interval of time. - majflt/s : Number of major faults the task has made per second - majflt-nr : Major faults made by the task and all its children, and collected during the interval of time. - system-ms : Total number of milliseconds spent by the task and all its children while | <p>the timer interval only (default value is 100 milli-seconds). Sometimes a scenario can arrive where, when the command request VEOS to fetch the latest user time but the user time returned can be the value updated on last scheduler timer expiry and vice-versa. Hence, the values retrieved from VEOS can cause some percentage difference for the command.</p> <p>4. The SELF keyword indicates that statistics are to be reported for the “pidstat” process itself and pidstat is a VH process not VE process.</p> <p>5. VE architecture do not support the given fields.</p> |
|--|--|---|--|

| | | | |
|------------|--------|---|--|
| | | <p>executing at the system level (kernel)</p> <ul style="list-style-type: none"> - guest-ms : Total number of milliseconds spent by the task and all its children in virtual machine (running a virtual processor) - StkRef: Memory in kilobytes used as stack, referenced by the task. <p>6. <code>"/opt/nec/ve/bin/pidstat"</code> command can terminate abnormally when some VE processes terminates.</p> | <p>6. If any VE process's checking at VEOS is succeeded but process get terminated while trying to fetch statistics from VEOS. Then command will terminate with failure error.</p> |
| sysstat-ve | mpstat | <p>1. In case of VE, the environment variable <code>VE_NODE_NUMBER</code> can be given:</p> <ul style="list-style-type: none"> - <code>VE_NODE_NUMBER</code> is set: Command shows the information corresponding to given node. - <code>VE_NODE_NUMBER</code> is not set: Command shows the information corresponding to all online VE nodes. - <code>VE_NODE_NUMBER</code> is not set and command executed with interval: Command shows the information corresponding to default VE node 0. <p>2. <code>"/opt/nec/ve/bin/mpstat -l"</code> will show error message "Interrupts are not applicable for VE".</p> <p>3. <code>"/opt/nec/ve/bin/mpstat -A"</code> will display only CPUs statistics not interrupts statistics.</p> <p>4. The command <code>"/opt/nec/ve/bin/mpstat <interval>"</code> can sometimes display more than or less than 100% value in <code>"%usr"</code> field while processes running on all the VE cores.</p> | <p>1. There are multiple nodes in VE architecture.</p> <p>2. There are no interrupts on VE.</p> <p>3. There are no interrupts on VE.</p> <p>4. In case of VE, the value of "user" will be updated as per the timer interval only (default value is 100 milli-seconds). Sometimes a scenario can arrive where, when the</p> |

| | | | |
|--|--|--|---|
| | | <p>5. The command <code>"/opt/nec/ve/bin/mpstat -P {cpu [,...] ON ALL }</code> will display information of single VE node (either for the given node or first online node).</p> <p>6. The following values will not be applicable for VE, so the values will be zero:</p> <ul style="list-style-type: none"> - %nice : Percentage of CPU utilization while executing at the user level with nice priority - %sys : Percentage of CPU used by the task while executing at system level, kernel - %iowait : Percentage of time that the CPU or CPUs were idle during which the system had an outstanding disk I/O request - %steal : Percentage of time spend by a CPU (which is virtualized), for resources from the physical CPU - %irq : Percentage of time spent by the CPU or CPUs to service interrupt - %soft : Percentage of time spent by the CPU or CPUs to service softirqs - %guest : Percentage of CPU spent by the task in virtual machine (running a virtual processor | <p>command request VEOS to fetch the latest user time but the user time returned can be the value updated on last scheduler timer expiry and vice-versa. Hence, the values retrieved from VEOS can cause some percentage difference for the command.</p> <p>5. As per the design of this command, CPU information can be retrieved only for one node and different VE nodes can have different number of CPUs.</p> <p>6. VE architecture do not support the given fields.</p> |
|--|--|--|---|

| | | | |
|------------|--------|--|---|
| | | <ul style="list-style-type: none"> - %gnice : Percentage of time spent by the CPU or CPUs to run a niced guest. | |
| sysstat-ve | iostat | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. - VE_NODE_NUMBER is not set and command executed with interval: Command shows the information corresponding to default VE node 0. <p>2. The command "/opt/nec/ve/bin/iostat can sometimes display more than or less than 100% value in "%user" field while processes running on all the VE cores</p> <p>3. The following values will not be applicable for VE, so the values will be zero:</p> <ul style="list-style-type: none"> - %nice : Percentage of CPU utilization while executing at the user level with nice priority - %sys : Percentage of CPU used by the task while executing at system level, kernel - %iowait : Percentage of time that the CPU or CPUs were | <p>1. There are multiple nodes in VE architecture.</p> <p>2. In case of VE, the value of "user" will be updated as per the timer interval only (default value is 100 milli-seconds). Sometimes a scenario can arrive where, when the command request VEOS to fetch the latest user time but the user time returned can be the value updated on last scheduler timer expiry and vice-versa. Hence, the values retrieved from VEOS can cause some percentage difference for the command.</p> <p>3. VE architecture do not support the given fields.</p> |

| | | | |
|------------|-----|---|---|
| | | <p>idle during which the system had an outstanding disk I/O request</p> <ul style="list-style-type: none"> - %steal : Percentage of time spend by a CPU (which is virtualized), for resources from the physical CPU | |
| sysstat-ve | sar | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. - VE_NODE_NUMBER is not set and command executed with interval: Command shows the information corresponding to default VE node 0. <p>2. Command <code>"/opt/nec/ve/bin/sar -d"</code> shows error message "Block devices data is not applicable for VE"</p> <p>3. <code>"/opt/nec/ve/bin/sar -n"</code> shows error message "Network statistics is not applicable for VE".</p> <p>4. <code>"/opt/nec/ve/bin/sar -I"</code> shows error message "Interrupts are not applicable for VE"</p> <p>5. <code>"/opt/nec/ve/bin/sar -A"</code> will not display network statistics, interrupts statistics and block devices.</p> <p>6. The command <code>"/opt/nec/ve/bin/sar <interval>"</code> can sometimes display more than or less than 100% value in "%user" field while processes running on all the VE cores.</p> | <p>1. There are multiple nodes in VE architecture.</p> <p>2. There is no device data for VE.</p> <p>3. There are no network statistics for VE.</p> <p>4. There are no interrupts on VE.</p> <p>5. There are no network, interrupts and block device statistics for VE.</p> <p>6. In case of VE, the value of "user" will be updated as per the timer interval only (default value is 100 milli-seconds). Sometimes a scenario can arrive where, when the command request VEOS to fetch the latest user time but</p> |

| | | | |
|--|--|---|---|
| | | <p>7. When filename is not given, VE specific 'sar' command uses the standard system activity daily data file, "/var/opt/nec/ve/log/sa/sa<dd>_<node_number>", where the dd parameter indicates the current day.</p> <p>8. VE "sar" command will display "VEOS RESTART" instead of "LINUX RESTART" at restart of VEOS.</p> <p>9. Count of 'blocked processes for i/o' ("blocked" field) in "sar -q" command is unused for VE.</p> <p>10. The following values will not be applicable for VE, so the values will be zero:</p> <ul style="list-style-type: none"> - %nice : Percentage of CPU utilization while executing at the user level with nice priority - %system & %sys : Percentage of CPU used by the task while executing at system level, kernel - %iowait : Percentage of time that the CPU or CPUs were idle during which the system had an outstanding disk I/O request - %steal : Percentage of time spend by a CPU (which is virtualized), for resources from the physical CPU) - %irq : Percentage of time spent by the CPU or CPUs to service interrupts | <p>the user time returned can be the value updated on last scheduler timer expiry and vice-versa. Hence, the values retrieved from VEOS can cause some percentage difference for the command.</p> <p>7. There are multiple nodes in VE architecture. So, we maintained separate system activity daily data files for each node.</p> <p>8. "ve_sysstat" service is restarted at VEOS restart. So, VEOS RESTART message looks more appropriate.</p> <p>9. VE architecture do not maintain i/o specific blocked processes.</p> <p>10. VE architecture do not support the given fields.</p> |
|--|--|---|---|

| | | | |
|--|--|--|--|
| | | <ul style="list-style-type: none">- %soft : Percentage of time spent by the CPU or CPUs to service softirqs- %guest : Percentage of CPU spent by the task in virtual machine (running a virtual processor)- %gnice : Percentage of time spent by the CPU or CPUs to run a niced guest- pswpin/s : Total number of swap pages the system brought in per second- pswpout/s : Total number of swap pages the system brought out per second- fault/s : Number of page faults (major + minor) made by the system per second- majflt/s : Number of major faults the system has made per second- pgfree/s : Number of pages placed on the free list by the system per second- pgscank/s : Number of pages scanned by the kswapd daemon per second- pgscand/s : Number of pages scanned directly per second- pgsteal/s : Number of pages the system has reclaimed from cache (pagecache and swapcache) per second to satisfy its memory demands- %vmeff : Calculated as $\text{pgsteal} / \text{pgscan}$- Kbhugfree : Amount of hugepages memory in kilobytes that is not yet allocated.- %hugused : Percentage of total hugepages memory that has been allocated- bufpg/s : Number of additional memory pages used as buffers by the system per second | |
|--|--|--|--|

| | | | |
|------------|------|---|--|
| | | <ul style="list-style-type: none"> - campg/s : Number of additional memory pages cached by the system per second - Kbuffers : Memory used as buffers by the kernel in kilobytes - Kbcached : Memory used to cache data by the kernel in kilobytes - Kbcommit: Memory in kilobytes needed for current workload - %commit : Percentage of memory needed for current workload in relation to the total amount of memory (RAM+swap) - Kbactive : Active memory in kilobytes - Kbinact : Inactive memory in kilobytes - Kbdirty : Memory in kilobytes waiting to get written back to the disk - Kbswpcad : Amount of free swap space in kilobytes. - kbswpused : Amount of used swap space in kilobytes - %swpcad : Percentage of used swap space. - Kbswpcad : Amount of cached swap memory in kilobytes - %swpcad : Percentage of cached swap memory in relation to the amount of used swap space - Dentunusd : Number of unused cache entries in the directory cache. | |
| sysstat-ve | sadc | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command collects the information corresponding to given node. | <p>1. There are multiple nodes in VE architecture.</p> |

| | | | |
|--|--|--|---|
| | | <ul style="list-style-type: none"> - VE_NODE_NUMBER is not set and given interval is "1": Command collects the information corresponding to all online VE nodes. - VE_NODE_NUMBER is not set and command executed with interval (greater than "1"): Command shows the information corresponding to default VE node 0. <p>2. "/opt/nec/ve/lib64/sa/sadc -S" does not support options INT, DISK, SNMP, IPV6, XDISK. It will show the following error messages:</p> <ul style="list-style-type: none"> a) "/opt/nec/ve/lib64/sa/sadc -S INT": Interrupts are not applicable for VE b) "/opt/nec/ve/lib64/sa/sadc -S DISK": Block devices data is not applicable for VE c) "/opt/nec/ve/lib64/sa/sadc -S SNMP": SNMP statistics are not applicable for VE d) "/opt/nec/ve/lib64/sa/sadc -S IPV6": IPV6 statistics are not applicable for VE e) "/opt/nec/ve/lib64/sa/sadc -S XDISK": Partition and disk statistics is not applicable for VE <p>3. "/opt/nec/ve/lib64/sa/sadc" command collects the information in file "sa<dd>_<node_number>" at path "/var/opt/nec/ve/log/sa". If outfile (a file to collect information) is set to "-" then sadc uses the standard system activity daily data file which is</p> | <ul style="list-style-type: none"> - The command "sadc" is internally called by "sa1" which is designed to be started automatically by the cron command and collect system activity daily data at 1 sec of interval. So, it calls sadc with interval "1" and we need to collect system activity data for all online nodes. <p>2. Only power management specific data can be collected in VE, so, "sadc -S" option supports only POWER, ALL and XALL options. The other options will display error message.</p> <p>3. There are multiple nodes in VE architecture. Hence we need to create files according to VE nodes to collect system activity data.</p> |
|--|--|--|---|

| | | | |
|------------|------|---|---|
| | | <p>“var/opt/nec/ve/log/sa/sa<dd>_<node_number>” for VE. But the standard (x86_64) ‘sadc’ command collects system activity information in file "sa<dd>" at path "/var/log/sa", (dd parameter indicates the current day).</p> | |
| sysstat-ve | sadf | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. - VE_NODE_NUMBER is not set and command executed with interval: Command shows the information corresponding to default VE node 0. <p>2. “/opt/nec/ve/bin/sadf” can sometimes display more than or less than 100% value in “%user” field while processes running on all the VE cores.</p> <p>3. VE specific “/opt/nec/ve/bin/sadf” command uses file “/var/opt/nec/ve/log/sa/sa<dd>_<node_number>” to display system activity data. The same file will be used, if</p> | <p>1. There are multiple nodes in VE architecture.</p> <p>2. sadf command reads the data collected by sadc. In case of VE, the value of “user” will be updated as per the timer interval only (default value is 100 milli-seconds). Sometimes a scenario can arrive where, when the command request VEOS to fetch the latest user time but the user time returned can be the value updated on last scheduler timer expiry and vice-versa. Hence, the values retrieved from VEOS can cause some percentage difference for the command.</p> <p>3. There are multiple nodes in VE architecture. So, we maintained separate system activity daily data files for each node.</p> |

| | | | |
|------------|-----|--|---|
| | | <p>outfile (file to extract data and write to standard output) is omitted. But the standard (x86_64) 'sadf' command uses file "var/log/sa/sa<dd>", (dd parameter indicates the current day).</p> <p>4. VE "sadf" command will display "VEOS-RESTART" instead of "LINUX-RESTART" at restart of VEOS.</p> | <p>4. "ve_sysstat" service is restarted at VEOS restart. So, VEOS-RESTART message looks more appropriate.</p> |
| sysstat-ve | sa1 | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command collects the information corresponding to given node. - VE_NODE_NUMBER is not set and given interval is "1": Command collects the information corresponding to all online VE nodes. - VE_NODE_NUMBER is not set and command executed with interval (greater than "1"): Command shows the information corresponding to default VE node 0. <p>2. VE specific "/opt/nec/ve/lib64/sa/sa1" command which is started automatically by the cron command will collect the system activity daily data in file "/var/opt/nec/ve/log/sa/sa<dd>_<node_number>" for all online VE nodes. But the standard (x86_64) 'sa1' command collects system activity information in file "/var/log/sa</p> | <p>1. There are multiple nodes in VE architecture.</p> <ul style="list-style-type: none"> - The command "sa1", which internally calls "sadc", is designed to be started automatically by the cron job and collect system activity daily data at 1 sec of interval. So, it calls sadc with interval "1" and we need to collect system activity data for all online nodes. <p>2. There are multiple nodes in VE architecture. So, we maintained separate system activity daily data files for each node.</p> |

| | | | |
|------------|------------------|---|---|
| | | sa<dd>", (dd parameter indicates the current day). | |
| sysstat-ve | sa2 | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command writes a daily report corresponding to given node. - VE_NODE_NUMBER is not set: Command writes a daily report corresponding to all online VE nodes. <p>2. VE specific "/opt/nec/ve/lib64/sa/sa2" command which get started automatically by the cron command will collect the system activity daily data in file "/var/opt/nec/ve/log/sa/sar<dd>_<node_number>" for all online VE nodes. But the standard (x86_64) 'sa2' command collects system activity information in file "/var/log/sa/sar<dd>", (dd parameter indicates the current day).</p> | <p>1. There are multiple nodes in VE architecture.</p> <p>2. There are multiple nodes in VE architecture. So, we maintained separate system activity daily data files for each node.</p> |
| sysstat-ve | sysstat services | <p>1. The following services related files will be different for ported sysstat package:</p> <ul style="list-style-type: none"> - Sysstat service for collecting system activity information of VEs are defined in "/usr/lib/systemd/system/ve_sysstat@.service". The standard sysstat package defines this information in "/usr/lib/systemd/system/sysstat.service". Hence VE specific service will be started by command "systemctl start ve_sysstat@N.service". (Where N indicates VE node number) - In case of VE, whenever any/all VEOS is restarted then ve_sysstat@ service of | <p>1. There are multiple nodes in VE architecture. VE specific service will collect the system activity information for specified/all VE nodes. Hence, we need to maintain separate VE specific services and configuration files.</p> |

| | | | |
|---------------|---------|--|---|
| | | <p>corresponding node/s will be restarted.</p> <ul style="list-style-type: none"> - Sysstat service for collecting system activity information automatically by cron command of VEs are defined in “/etc/cron.d/ve_sysstat”. The standard sysstat package defines this information in “/etc/cron.d/sysstat”. - The configuration file which includes definition of multiple macros will be defined in “/etc/sysconfig/ve_sysstat”. The standard sysstat package defines this information in “/etc/sysconfig/sysstat”. <p>2. In RHEL8 environment, VE ported sysstat package enables and starts the ve_sysstat service at package installation.</p> | <p>2. In RHEL8, VH specific sysstat service is only enabled but not started at sysstat package installation. And user needs to start it whenever required to collect system statistics.</p> |
| util-linux-ve | taskset | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Runs the specified process or search the given PID on given node. - VE_NODE_NUMBER is not set: Runs the specified program on VE node 0 or search the given PID on all online nodes. <p>2. VE ‘taskset’ command is restricted to execute only VE binary and VH binary cannot be executed directly using this command.</p> | <p>1. There are multiple nodes in VE architecture.</p> <p>2. VH process execution using VE commands are not allowed.</p> |
| util-linux-ve | lscpu | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. | <p>1. There are multiple nodes in VE architecture.</p> |

| | | | |
|---------------|---------|---|--|
| | | <ul style="list-style-type: none"> - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. <p>2. Command "/opt/nec/ve/bin/lscpu -c -e" will show error message "No offline CPU exists for VE".</p> <p>3. Command "/opt/nec/ve/bin/lscpu" with '-s or --sysroot' option will show error message "-s or --sysroot option is not supported for VE".</p> | <p>2. In case of VE, we can't make a CPU offline.</p> <p>3. In case of VE, user can not specify the directory to gather CPU data.</p> |
| util-linux-ve | prlimit | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Runs the specified process or search the given PID on given node. - VE_NODE_NUMBER is not set: Runs the specified program on VE node 0 or searches the given PID on all online nodes. <p>2. In case of VE, PID is mandatory to run "/opt/nec/ve/bin/prlimit" and "/opt/nec/ve/bin/prlimit [RESOURCE OPTIONS] to show the resource limit otherwise it will give error. But in case of x86_64, it shows the resource limits of the current process</p> <p>3. In case of VE, command "/opt/nec/ve/bin/prlimit> -p <pid>" shows the resource limit for NICE and RTPRIO as blank (-).</p> <p>4. Get/set the limits for "nice" and "rtprio" are not supported for VE. So the following commands are not supported either with or without PID and shows the error "Resource not supported":</p> <ul style="list-style-type: none"> a) /opt/nec/ve/bin/prlimit -e=<limits> | <p>1. There are multiple nodes in VE architecture.</p> <p>2. In case of Linux, it shows the resource limits of the current process, i.e. running instance of 'prlimit' command which is a VH process and not a VE process". So in case of VE, prlimit cannot be run without specifying PID</p> <p>3. Priority scheduling is not supported in VE, hence, getpriority()/setpriority() system calls are not supported.</p> <p>4. Priority scheduling is not supported in VE, hence, getpriority()/setpriority() system calls are not supported.</p> |

| | | | |
|---------------|---------|--|--|
| | | <p>b) /opt/nec/ve/bin/prlimit -- nice=<limits> -p <pid></p> <p>c) /opt/nec/ve/bin/prlimit - r=<limits></p> <p>d) /opt/nec/ve/bin/prlimit -- rtprio=<limits> -p <pid></p> <p>5. Get/set the limits for “memlock” will successfully set the given limit but will not affect the VE process memory.</p> <p>6. We cannot run a VE process using command “/opt/nec/ve/bin/prlimit <VE_process>” with given stack limit specified by the environment variable VE_STACK_LIMIT.</p> <p>7. VE ‘prlimit’ command is restricted to execute only VE binary and VH binary cannot be executed directly using this command.</p> | <p>5. There is no swapped memory in VE, so the complete memory is locked.</p> <p>6. In this case, VEOS will get the resource limits given by prlimit command and the command cannot parse the binary and do not have stack information to calculate stack limit.</p> <p>7. VH process execution using VE commands are not allowed.</p> |
| util-linux-ve | lslocks | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. | <p>1. There are multiple nodes in VE architecture.</p> |
| util-linux-ve | ipcs | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. <p>2. VE command “ipcs” and “ipcs -a/--all” option only displays shared memory information.</p> | <p>1. There are multiple nodes in VE architecture.</p> <p>2. There are no VE specific message queues and semaphores.</p> |

| | | | |
|--|--|--|--|
| | | <p>3. The following options are not supported for VE:</p> <ul style="list-style-type: none"> • -q/--queue • -s/--semaphores <p>4. The shared memory limit (say 4096) displayed using command <code>"/opt/nec/ve/bin/ipcs -l"</code> includes the shared memories in VH and all VE nodes and there are no impact of <code>VE_NODE_NUMBER</code> to display shared memory limit.</p> <p>5. When VH shared memory limit is exhausted and VE <code>"ipcs"</code> command is executed then it will return error <code>"No space left on device"</code>.</p> <p>6. The number of shared memory segment which a user can create through VE process will be less than <code>"max number of segments"</code>.</p> <p>7. The following value will not be applicable for VE, so the values will be zero: pages swapped: Total number of swapped shared memory pages</p> | <p>3. There are no VE specific message queues and semaphores.</p> <p>4. VH and VE both shares the same shared memory limit.</p> <p>5. VH command reads the information from <code>"/proc"</code> file system but VEOS creates a shared memory on VH to write all VE specific shared memory information and the command <code>"ipcs"</code> reads this information to display the output. So when shared memory limit is exhausted and VEOS failed to create a shared memory then command will return error. To overcome from this error, user need to remove some shared memories using VH <code>'ipcrm'</code> command.</p> <p>6. <code>"ve_exec"</code> itself consumes a shared memory and uses the memory until its termination. So let's say, a VE process create 4095 shared memory than it will totally create 4096 segments on VH. One of them will be created by <code>"ve_exec"</code>.</p> <p>7. VE architecture do not support the given fields.</p> |
|--|--|--|--|

| | | | |
|----------------------|--------------|---|--|
| <p>util-linux-ve</p> | <p>ipcrm</p> | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. <p>2. The following options are not supported for VE:</p> <ul style="list-style-type: none"> • -Q/--queue-key <msgkey> • -q/--queue-id <msgid> • -S/--semaphore-key <semkey> • -s/--semaphore-id <semid> • --all=[msg sem] <p>3. VE command “ipcrm” and “ipcrm -a” will remove only shared memory.</p> <p>4. When VH shared memory limit is exhausted and VE “ipcrm -a” or “ipcrm --all=shm” command is executed then it will return error “No space left on device”.</p> | <p>1. There are multiple nodes in VE architecture.</p> <p>2. There are no VE specific message queues and semaphores.</p> <p>3. There are no VE specific message queues and semaphores.</p> <p>4. VEOS creates a shared memory on VH to write information about removed shared memories when command is executed with ‘-a’ or ‘--all=shm’ option. So when shared memory limit is exhausted and VEOS failed to create a shared memory then command will return error. You can perform following steps to overcome from this error:</p> <ol style="list-style-type: none"> i. Display all shared memory by VH’s ‘ipcs’ command ii. Delete some shared memories by VH’s ‘ipcrm’ command. iii. Now, execute /opt/nec/ve/bin/ipcrm -a |
|----------------------|--------------|---|--|

| | | | |
|-----------|----|--|---|
| psacct-ve | sa | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. <p>2. If VE_NODE_NUMBER is not set then by default VE node 0 will be used to execute following commands: a) /opt/nec/ve/sbin/sa <filename> b) /opt/nec/ve/sbin/sa --other-acct-file <filename></p> <p>3. If VE_NODE_NUMBER is not set then by default VE node 0 will be used to execute following commands: a) /opt/nec/ve/sbin/sa --other-usracct-file <filename> -s b) /opt/nec/ve/sbin/sa --other-savacct-file <filename> -s</p> <p>4. "/opt/nec/ve/sbin/sa" command with - -ahz option will not have any effect on the values printed in STDOUT.</p> <p>5. VE specific ported 'sa' command collects the information in file usracct_<node_number> and savacct_<node_number>" at path "/var/opt/nec/ve/account". But in case of x86_64, 'sa' command collects the process accounting</p> | <p>1. There are multiple nodes in VE architecture.</p> <p>2. These commands will use 'usracct and savacct' files (if these files exist). To pick the correct 'usracct and savacct' file in multiple VE node environment, we need to specify node number. Otherwise it will use any random VE specific file. Hence, will display wrong information.</p> <p>3. These commands will use pacct file to generate usracct/savacct file. To pick the correct 'pacct' file in multiple node environment, we need to specify node number. Otherwise it will use any random VE specific 'pacct' file. Hence, will display wrong information.</p> <p>4. AHZ value is used in calculation of time related options. In case of VE, this value is not used, as the time is received in seconds/microseconds from VEOS.</p> <p>5. There are multiple nodes in VE architecture. Hence we need to create accounting files according to VE nodes.</p> |
|-----------|----|--|---|

| | | | |
|-----------|--------|---|---|
| | | <p>information in file "usracct and savacct" at path "/var/log/sa".</p> <p>6. When a VH process is executed from VE process using 've_exec', VE accounting file captures the elapsed time of VE process only. VH process execution time is not included in elapsed time.</p> <p>7. The following values will not be applicable for VE, so the values will be zero:</p> <ul style="list-style-type: none"> - min & min/c : Number of minor page faults - maj & maj/c : Number of major page faults - swp & swp/c: Number of swap pages - system time of process - Disk I/O operations (io) <p>8. The '/opt/nec/ve/sbin/sa' command shows an error message and terminate when it found a record whose 'ac_version' is 3.</p> <p>9. The tool '/opt/nec/ve/sbin/convert-acct' is provided to convert records from 'ac_version 3' to 'ac_version 14'. Like "/opt/nec/ve/sbin/convert-acct <filename_to_convert_accounting>", user sets the file wanted to convert as an argument.</p> | <p>6. As per VEOS design, VE process is exited as soon as VH process is executed using execve () system call.</p> <p>7. VE architecture do not support the given fields.</p> <p>8. Accounting file can contain version 3 type records when user updates veos from v2.5 or earlier to v2.6 or later.</p> <p>9. Same accounting file can have two different type of records, version 3 format and version 14 format. Hence, this 'convert-acct' tool will be used to convert record into same v14 format.</p> |
| psacct-ve | accton | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command enables/disable accounting corresponding to given node. - VE_NODE_NUMBER is not set: Command enables/disable accounting corresponding to all online VE nodes. <p>2. "/opt/nec/ve/sbin/accton on" will not show any error if</p> | <p>1. There are multiple nodes in VE architecture.</p> <p>2. In case of VH, the acct file is created at the time of package</p> |

| | | | |
|-----------|-----------|---|---|
| | | <p>“/var/opt/nec/ve/account/pacct_<N>” file does not exist. But in case of x86_64, if “/var/account/pacct” file is not present, <accton on> will show error.</p> <p>3. VE specific ported accton command collects the information in file "pacct_<node_number> at path "/var/opt/nec/ve/account". But in case of x86_64, accton command collects the process accounting information in file "pacct" at path "/var/log/sa".</p> <p>4. In case of VE, if the required file doesn't exist while running command “/opt/nec/ve/sbin/accton <filename/on>” then it will display error “No such file and directory”. But in case of x86_64, command displays the error “permission denied”.</p> <p>5. If VE_NODE_NUMBER is not set and “/opt/nec/ve/sbin/accton <filename>” command is executed then by default accounting will be enabled only for VE node 0.</p> | <p>installation, but in case of VE the file is created when ‘accton on’ command is executed, because the acct files are created per online node and at the time of installation of package, the number of nodes which are online are not known</p> <p>3. There are multiple nodes in VE architecture. Hence we need to create files according to VE nodes to enable accounting.</p> <p>4. The design of VE and VH command is different. The VE specific command will check the file existence before checking permission and VH command will check for permission before checking the file existence.</p> |
| psacct-ve | dump-acct | <p>1. When a VH process is executed from VE process using ‘ve_exec’, VE accounting file captures the elapsed time of VE process only. VH process execution time is not included in elapsed time.</p> <p>2. On RHEL8 environment, ‘dump-acct’ command shows the NULL in ‘tty’ field in following scenarios:</p> <ul style="list-style-type: none"> - If VE process is executed with “nohup” command. - If VE process is executed in background using system () system call. | <p>1. As per VEOS design, VE process is exited as soon as VH process is executed from it using execve () system call.</p> <p>2. When no terminal is associated with a VE process then it will show ‘tty’ as NULL.</p> |

| | | | |
|-----------|----------|--|--|
| | | <ul style="list-style-type: none"> - If VE process is executed in background through any bash script. - If VE process ('conftest' binary) is executed from './configure' command. <p>3. Records in process account file contains vector information when it's 'ac_version' is 14. Like <code>"/opt/nec/ve/sbin/dump-acct --ve-info <filename_to_read_accounting>"</code>, the command 'dump-acct' with the option '--ve-info' shows vector information.</p> <p>4. The <code>"/opt/nec/ve/sbin/dump-acct"</code> command shows an error message and terminate when it found a record whose 'ac_version' is 3.</p> <p>5. The tool <code>"/opt/nec/ve/sbin/convert-acct"</code> is provided to convert records from 'ac_version 3' to 'ac_version 14'. Like <code>"/opt/nec/ve/sbin/convert-acct <filename_to_convert_accounting>"</code>, user sets the file wanted to convert as an argument.</p> | <p>3. Additional vector information is introduced in accounting file with version 14.</p> <p>4. Accounting file can contain version 3 type records when user updates veos from v2.5 or earlier to v2.6 or later.</p> <p>5. Same accounting file can have two different type of records, version 3 format and version 14 format. Hence, this 'convert-acct' tool will be used to convert record into same v14 format.</p> |
| psacct-ve | lastcomm | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command will show the accounting information for given node only. - VE_NODE_NUMBER is not set: Command will show the accounting information for all online nodes. <p>2. The 'lastcomm' command shows the NULL in 'tty' field in following scenarios:</p> <ul style="list-style-type: none"> - If VE process is executed with "nohup" command. - If VE process is executed in background using system () system call. | <p>1. There are multiple nodes in VE architecture.</p> <p>2. When no terminal is associated with a VE process then it will show 'tty' as NULL.</p> |

| | | | |
|-----------|-----------------|---|--|
| | | <ul style="list-style-type: none"> - If VE process is executed in background through any bash script. - If VE process ('conftest' binary) is executed from './configure' command. <p>3. Records in process account file contains vector information when it's 'ac_version' is 14. Like "/opt/nec/ve/bin/lastcomm --ve-info <filename_to_read_accounting>", the command 'lastcomm' with the option '--ve-info' shows vector information.</p> <p>4. The '/opt/nec/ve/bin/lastcomm' command shows an error message and terminate when it found a record whose 'ac_version' is 3.</p> <p>5. The tool '/opt/nec/ve/sbin/convert-acct' is provided to convert records from 'ac_version 3' to 'ac_version 14'. Like "/opt/nec/ve/sbin/convert-acct <filename_to_convert_accounting>", user sets the file wanted to convert as an argument.</p> | <p>3. Additional vector information is introduced in accounting file with version 14.</p> <p>4. Accounting file can contain version 3 type records when user updates veos from v2.5 or earlier to v2.6 or later.</p> <p>5. Same accounting file can have two different type of records, version 3 format and version 14 format. Hence, this 'convert-acct' tool will be used to convert record into same v14 format.</p> |
| psacct-ve | Psacct Services | <p>The following points related to services will be different for ported psacct-ve package:</p> <ul style="list-style-type: none"> - Psacct service for collecting process accounting information for VE is defined in "/usr/lib/systemd/system/psacct-ve@.service". The standard psacct package defines this information in "/usr/lib/systemd/system/psacct.service". - VE 'psacct-ve' service gets started by command "systemctl start psacct-ve@\$N.service" and stopped by command "systemctl stop psacct-ve@\$N.service" | <p>There are multiple nodes in VE architecture. So, separate services are required to handle process accounting for specified/all VE nodes.</p> <ul style="list-style-type: none"> - There are multiple nodes in VE architecture. So, VE psacct-ve@\$N.service used to start and stop VE process's |

| | | | |
|-----------|--------|--|---|
| | | <p>(Where \$N specifies VE node number).</p> <ul style="list-style-type: none"> - VE psacct-ve service enable/disable accounting for all VE nodes. It is not possible to enable/disable psacct-ve service for specified VE node. So, if psacct-ve service is enabled by command "systemctl enable psacct-ve@\$N.service" then it will enable psacct service for all VE nodes (Where \$N specifies VE node number). - To disable psacct-ve service, we need to disable all services which was enabled earlier. Let's say, psacct-ve service gets enabled by command: "systemctl enable psacct-ve@\$N.service" To disable this service, we need to use following command: "systemctl disable psacct-ve@\$N.service" (Where \$N specifies VE node number). - The logrotate file which is used by psacct-ve package will be defined in "/etc/logrotate.d/psacct-ve". The standard psacct package defines this information in "/etc/logrotate.d/psacct". | <p>accounting for specified/all VE nodes.</p> <ul style="list-style-type: none"> - When service gets enabled then it creates link with psacct-ve@.service at path "/etc/systemd/system/multi-user.target.wants/" which will enable accounting for all VE nodes. - When VE psacct-ve service needs to be disabled then all the links exist at path "/etc/systemd/system/multi-user.target.wants/" with "psacct-ve@.service" should be removed. |
| strace-ve | strace | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Runs the specified process or search the given PID on given node. - VE_NODE_NUMBER is not set: Runs the specified program on | <p>1. There are multiple nodes in VE architecture.</p> |

| | | | |
|--|--|---|--|
| | | <p>VE node 0 or searches the given PID on all online nodes.</p> <p>2. System calls which are executed at the time of program loading, are not captured in ported "strace" command.</p> <p>3. '-D' option is not supported.</p> <p>4. To trace multiple VE PIDs, all the given PIDs must be running on same VE node, i.e in command "/opt/nec/ve/bin/strace -p pid1,pid2" should be running on same node</p> <p>5. Command "/opt/nec/ve/bin/strace -S" will show stime as '0' for all the sytem calls.</p> <p>6. /opt/nec/ve/bin/strace -p pid: System call tracing, which was executed right before the command</p> | <p>2. VE strace command cannot attach a process which is not running on VE. So, first we have to "execv" to run VE program with --traceme flag and then will continue to trace system call. In this case system call executed at loading time are missed.</p> <p>3. With -D option, tracer process runs as a detached grandchild, not as parent of the tracee. In x86_64, the process is first attached and then loaded with execve (In parent). But In VE, the process is loaded using execve (in parent) with --traceme flag instead to attaching it. VE Ptrace gets its ppid and considers it as its tracer. But with -D option, tracee process's tracer is its detached grandchild not its parent process and VE ptrace will get its parent as 0. So, ported strace command cannot trace VE process using its detached grandchild.</p> <p>4. Strace internally uses 'ptrace' system call for tracing, and in case of VE, this a limitation of ptrace() system call.</p> <p>5. There is no system time in case of VE.</p> <p>6. When tracing is enabled on any running process, the system call running at that time get interrupted and then</p> |
|--|--|---|--|

| | | |
|--|---|--|
| | <p>"/opt/nec/ve/bin/strace -p pid" fired, will get skipped.</p> <p>7. The command "strace" cannot trace VE specific "ve_grow" system call.</p> <p>8. VE 'strace' command is restricted to execute only VE binary and VH binary cannot be executed directly using this command.</p> <p>9. If VE strace command is used to trace a VE process which invokes execve() system call or exec family of library functions then VE strace command will detach itself from the traced process.</p> <p>In case of Linux strace command, traced process will not be detached and Linux strace command will continue to trace the process.</p> <p>10. VE 'strace' command options "-e inject", "-e fault" and "-e kvm" are not supported.</p> <p>11. If VE 'strace' command is killed with SIGKILL then tracee process exits immediately.</p> <p>In case of x86, when strace command is killed with SIGKILL then tracee</p> | <p>restarted (move some instruction back) using PTRACE_SYSCALL and PRACE_CONT calls of ptrace. This is handled by kernel. So, it can trace that system call. But in case of VE, ptrace with PTRACE_SYSCALL is handled by libveptrace instead of kernel. So, we cannot handle such scenario and system call tracing is skipped.</p> <p>7. "ve_grow" system calls will not write its arguments on registers. So, 'strace' command cannot read its arguments from registers to show tracing.</p> <p>8. VH process execution using VE commands are not allowed.</p> <p>9. No support of PTRACE_O_TRACEEXEC flag in VE ptrace.</p> <p>10. These options are new in RHEL8 specific strace package. And RHEL7 do not have these options.</p> <p>11. "PTRACE_O_EXITKILL" flag is provided in VE 'strace' command so that tracee process can exit gracefully at tracer process's (VE strace command) termination.</p> |
|--|---|--|

| | | | |
|--------------|------|--|--|
| | | process executes normally and later exits. | |
| procps-ng-ve | pmap | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command will search the given PID on given node. - VE_NODE_NUMBER is not set: Command will search the given PID on all online nodes. <p>2. The following values will not be applicable for VE, so the values will be zero:</p> <ul style="list-style-type: none"> - Shared_Clean: Shared Pages not modified since they were mapped - Shared_Dirty: Shared Pages modified since they were mapped - Private_Clean: Private Pages not modified since they were mapped - Private_Dirty: Private Pages modified since they were mapped - Referenced: Amount of memory currently marked as referenced or accessed - Swap: Swap memory - Locked: Locked Pages which cannot be swapped out | <p>1. There are multiple nodes in VE architecture.</p> <p>2. VE architecture do not support the given fields.</p> |
| procps-ng-ve | w | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. <p>2. “/opt/nec/ve/bin/w” command can terminate abnormally when some VE processes terminates.</p> | <p>There are multiple nodes in VE architecture.</p> <p>2. If any VE process’s checking at VEOS is succeeded but process get terminated while trying to fetch statistics from</p> |

| | | | |
|--------------|--------|--|--|
| | | | VEOS. Then command will terminate with failure error. |
| procps-ng-ve | tload | <ul style="list-style-type: none"> - In case of VE, the environment variable VE_NODE_NUMBER can be given: - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to default VE node 0. | <p>There are multiple nodes in VE architecture.</p> <ul style="list-style-type: none"> - 'tload' command runs continuously without exiting. So the command cannot display information for all nodes. |
| procps-ng-ve | vmstat | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. - VE_NODE_NUMBER is not set and command executed with some delay: Command shows the information corresponding to default VE node 0. <p>2. "/opt/nec/ve/bin/vmstat -m" shows error message "slabinfo not supported in this configuration".</p> <p>3. Count of 'blocked processes for i/o' ("b" field) is unused in vmstat command.</p> <p>4. The following values will not be applicable for VE, so the values will be zero:</p> <ul style="list-style-type: none"> - swpd : Swap memory used - buff: Memory used as buffers - cache : Memory used as cache - si : Memory swapped in from disk - so : Memory swapped to disk | <p>1. There are multiple nodes in VE architecture.</p> <p>2. There is no slabinfo for VE.</p> <p>3. VE architecture do not maintain i/o specific blocked processes.</p> <p>4. VE architecture do not support the given fields.</p> |

| | | | |
|--|--|--|--|
| | | <ul style="list-style-type: none"> - bi : Blocks received from a block device - bo : Blocks sent to a block device - in : Number of interrupts per second - active memory : Memory that has been used recently - inactive memory : Memory which has been less recently used - swap cache : In-memory cache for files read from the disk - total swap : Total swap space size - used swap : Total used swap memory - free swap : Available swap memory size - sy : Time spent running kernel code. (system time) - st : Time stolen from a virtual machine. - ni (nice user cpu ticks) : Time spent by all CPU's to execute niced processes in user mode - wa (IO-wait cpu ticks): Time spent by all CPU's waiting for I/O to complete - IRQ cpu ticks : Time spent by all CPU's in servicing interrupts - softirq cpu ticks : Time spent by all CPU's in servicing softirqs - stolen cpu ticks : Time spent by all CPU's during involuntary wait - pages paged in - pages paged out - pages swapped in - pages swapped out - interrupts : counts of interrupts serviced since boot time, for each of the possible system interrupts | |
|--|--|--|--|

| | | | |
|---------------------|-------------|---|---|
| <p>procps-ng-ve</p> | <p>free</p> | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. - VE_NODE_NUMBER is not set and command executed with -c or -s option: Command shows the information corresponding to default VE node 0. <p>2. The following values will not be applicable for VE, so the values will be zero:</p> <ul style="list-style-type: none"> - Mem: (buffers) :: Memory used by buffers - Mem: (cache) :: In-memory cache for files read from the disk - Low: (total) :: Total low memory - Low: (used) :: Used low memory - Low: (free) :: Free low memory - High: (total) :: Total high memory - High: (used) :: Used high memory - High: (free) :: Free high memory - '-/+ buffers/cache (total) :: Total memory for buffer and cache - '-/+ buffers/cache (used) :: Total memory used for buffer and cache - Swap: (total) :: Total swap space size - Swap: (used) :: Used swap space size | <p>1. There are multiple nodes in VE architecture.</p> <p>2. VE architecture do not support the given fields.</p> |
|---------------------|-------------|---|---|

| | | | |
|--------------|--------|---|--|
| | | <ul style="list-style-type: none"> - Swap: (free) :: Memory which has been evicted from RAM, and is temporarily on the disk | |
| procps-ng-ve | uptime | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. <p>2. In case of VE, <code>"/opt/nec/ve/bin/uptime -p"</code> command output sometimes display "up"</p> | <p>1. There are multiple nodes in VE architecture.</p> <p>2. Whenever VE "uptime -p" command is executed just after VEOS started then the command output displays "up" without showing any minutes because VEOS is started just 0 minutes before.</p> <p>However, In case of X86_64, when system is restarted and reaches at terminal to execute the same command then it consists of some value in minutes. So 'uptime -p' command displays "up <value> minutes".</p> |
| procps-ng-ve | ps | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command shows the information corresponding to all online VE nodes. <p>2. In case of VE, <code>"/opt/nec/ve/bin/ps"</code> command displays blank (-) for 'priority' and 'nice' value with all the applicable options.</p> | <p>1. There are multiple nodes in VE architecture.</p> <p>2. Priority scheduling is not supported in VE, hence, <code>getpriority()/setpriority()</code> system calls are not supported.</p> |

| | | |
|--|---|--|
| | <p>3. <code>"/opt/nec/ve/bin/ps"</code> command displays blank (-) for all the namespaces (IPC, MNT, NET, PID, USER, UTS) values.</p> <p>4. <code>"/opt/nec/ve/bin/ps"</code> command will not display the current instruction pointer (EIP) and stack pointer (ESP) values of VE process.</p> <p>5. <code>"/opt/nec/ve/bin/ps s"</code> command will not display PENDING signal for VE</p> <p>6. The following values will not be applicable for VE, so the values will be zero:</p> <ul style="list-style-type: none"> - majflt : major page faults that have occurred with this process - minflt : minor page faults that have occurred with this process - nwchan : Address of the kernel function where the process is sleeping - wchan : Name of the kernel function in which the process is sleeping - size : Swap space that would be required if the process were to dirty all writable pages and then be swapped out. <p>7. <code>"/opt/nec/ve/bin/ps"</code> command executed with some option (like, -L, -H, -m, -f, -T or --sort) can terminate abnormally when some VE processes terminates. Also if command is executed with '-p' option for any specified VE PID and any running VE process terminated</p> | <p>3. Namespaces are not supported for VE.</p> <p>4. VEOS doesn't fetch instruction and stack pointer values from the running VE core at the time of command request for it. But VEOS provides the last updated values of it.</p> <p>5. VEOS cannot distinguish between shared pending signals and signals pending for a particular TID.</p> <p>6. VE architecture do not support the given fields.</p> <p>7. If any VE process's checking at VEOS is succeeded but process get terminated while trying to fetch statistics from VEOS. Then command will terminate with failure error.</p> |
|--|---|--|

| | | | |
|--------------|-----|--|--|
| | | <p>then command can terminate abnormally.</p> <p>8. In case of VE, RSS(Resident set size) is the sum of USS(Unique set size) and PSS(Proportional set size). USS of process is the total size of non shared memory of a process. PSS of process is the total size of shared memory of a process. Each PSS region is divided by the number of attaching processes.</p> | |
| procps-ng-ve | top | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command shows the information corresponding to given node. - VE_NODE_NUMBER is not set: Command show the information corresponding to default VE node 0. <p>2. “/opt/nec/ve/bin/top” command displays blank (-) for ‘priority’ and ‘nice’ values.</p> <p>3. “/opt/nec/ve/bin/top” command displays blank (-) for all the namespaces (IPC, MNT, NET, PID, USER, UTS) values.</p> <p>4. “/opt/nec/ve/bin/top” command can sometimes display incorrect percentages values in “us” and “id” of %CPU<core_id> field when pressing 1.</p> | <p>1. There are multiple nodes in VE architecture.</p> <ul style="list-style-type: none"> - ‘top’ command runs continuously without exiting. So the command cannot display information for all nodes. <p>2. Priority scheduling is not supported in VE, hence, getpriority()/setpriority() system calls are not supported.</p> <p>3. Namespaces are not supported for VE.</p> <p>4. In case of VE, the value of “user” will be updated as per the timer interval only (default value is 100 milli-seconds). Sometimes a scenario can arrive where, when the command request VEOS to fetch the latest user time but the user time returned can be the value updated on last scheduler timer expiry and vice-versa. Hence, the values</p> |

| | | | |
|-----------|---------|---|--|
| | | <p>5. The following values will not be applicable for VE, so the values will be zero:</p> <ul style="list-style-type: none"> - Percentage of the CPU for system processes - Percentage of the CPU processes waiting for I/O operations - Percentage of the CPU serving hardware interrupts - Percentage of the CPU serving software interrupts - Time stolen from a virtual machine. - Memory used by buffers - Total swap memory - Swap memory in use currently - Free swap memory - Cached memory by system - wchan : Name or the address of the kernel function in which the task is currently sleeping. - nDRT : Dirty pages count - nMaj : Major page fault count - nMin : Minor page fault count - vMj : Major page fault count delta - vMn : Minor page fault count delta <p>6. “/opt/nec/ve/bin/top” command can terminate abnormally when some VE processes terminates. Also if command is executed with ‘-p’ option for any specified VE PID and any running VE process terminated then also the command will terminate abnormally.</p> | <p>retrieved from VEOS can cause some percentage difference for the command.</p> <p>5. VE architecture do not support the given fields.</p> <p>6. If any VE process’s checking at VEOS is succeeded but process get terminated while trying to fetch statistics from VEOS. Then command will terminate with failure error.</p> |
| psmisc-ve | prtstat | <p>1. In case of VE, the environment variable VE_NODE_NUMBER can be given:</p> <ul style="list-style-type: none"> - VE_NODE_NUMBER is set: Command will search the given PID on given node. | <p>1. There are multiple nodes in VE architecture.</p> |

| | | | |
|-------------|----------|--|--|
| | | <ul style="list-style-type: none"> - VE_NODE_NUMBER is not set: Command will search the given PID on all online nodes. <p>2. “/opt/nec/ve/bin/prtstat” command displays blank (-) for ‘priority’, ‘rt_priority’ and ‘nice’ values.</p> <p>3. “/opt/nec/ve/bin/prtstat” command will not display the current instruction pointer (EIP) and stack pointer (ESP) values of VE process.</p> <p>4. The following values will not be applicable for VE, so the values will be zero :</p> <ul style="list-style-type: none"> - minflt, majflt : This Process minor & major faults - cminflt, cmajflt : Child processes minor & major faults - stime : Process's system time - guest_time : Process's guest time - delayaccr_blkio_ticks : Process's blkio - cstime : Child processes system time - cguest_time : Child processes guest time - wchan : Address where process went to sleep - nswap : Size of swap space of the process - cnsnap : Size of swap space of children of the process | <p>2. Priority scheduling is not supported in VE, hence, getpriority()/setpriority() system calls are not supported.</p> <p>3. VEOS doesn't fetch instruction and stack pointer values from the running VE core at the time of command request for it. But VEOS provides the last updated values of it.</p> <p>4. VE architecture do not support the given fields.</p> |
| autmake-ve | automake | N/A | |
| autoconf-ve | autoconf | N/A | |
| libtool-ve | libtool | In case of VE, command “/opt/nec/ve/bin/libtool --mode=finish <libname> <path>” will | In case of VE, programs are compiled using musl-libc and it does not provide "ldconfig". So ported libtool command will |

| | | | |
|--|--|---|---|
| | | not complete the installation of libtool libraries. | not be able to create the necessary links and cache for dynamic linker to the generated shared libraries found at specified path. |
|--|--|---|---|

3. Enhanced process accounting

Start of the process accounting service

If you use the process accounting, start the psacct-ve service with the following command.

```
# for i in `seq 0 7`; do if [ -e /dev/veslot$i ]; then systemctl enable psacct-ve@$i; systemctl restart psacct-ve@$i; fi done
```

When the psacct-ve service is active, the process accounting information of VE processes is recorded when they terminate. You can read this information with the lastcomm command or dump-acct command specifying a process accounting file corresponding to a VE number.

```
# /opt/nec/ve/bin/lastcomm -f /var/opt/nec/ve/account/pacct_N (Where $N specifies VE node number)
```

For example, the following command shows the accounting information in the process accounting file of the VE node#0.

```
# /opt/nec/ve/bin/lastcomm -f /var/opt/nec/ve/account/pacct_0
```

Stop of the process accounting service

To stop using the process accounting, stop the psacct-ve service with the following command.

```
# for i in `seq 0 7`; do systemctl stop psacct-ve@$i; systemctl disable psacct-ve@$i $i; done
```

Additional VE specific fields

On command lastcomm and dump-acct, the command with the option "--ve-info" shows additional VE specific fields in the process accounting information.

The fields are described as follows:

Command lastcomm

| COMMAND | Command name |
|----------------|--|
| FLAG | Accounting flags: S -- command executed by the super-user F -- command executed after a fork but without a following exec D -- command terminated with the generation of a core file X -- command was terminated with the signal SIGTERM |
| OWNER | The name of the user who ran the process |
| TTY | Terminal on which the process was executed |
| EXECUTION TIME | Time the process executed |
| START TIME | Time the process started |
| SID | Session ID |

| | |
|-----------|---|
| TIMESLICE | Timeslice [μ s] |
| NTHREADS | Max number of threads |
| NUMA | NUMA node number |
| TOTAL MEM | VE's total memory usage in clicks [kb * tick] |
| MAX MEM | VE's max memory usage [kb] |
| SYSCALL | The number of system calls |
| TRANSDATA | Data transfer amount between VE-VH [kb] |
| EX | Execution count |
| VX | Vector execution count |
| FPEC | Floating point data element count |
| VE | Vector element count |
| L1LMC | L1 instruction cache miss count |
| VECC | Vector execution in microseconds |
| L1MMC | L1 cache miss in microseconds |
| L2MMC | L2 cache miss in microseconds |
| VE2 | Vector element count 2 |
| VA REC | Vector arithmetic execution in microseconds |
| L1LMCC | L1 instruction cache miss in microseconds |
| VLDEC | Vector load execution in microseconds |
| L1OMCC | L1 operand cache miss in microseconds |
| PCCC | Port conflict in microseconds |
| LTRC | Load instruction traffic count |
| VLPC | Vector load packet count |
| STRC | Store instruction traffic count |
| VLEC | Vector load element count |
| VLCME | Vector load cache miss element count |
| VLCME2 | Vector load cache miss element count 2 |
| FMAEC | Fused multiply add element count |
| PTCC | Power throttling in microseconds |
| TTCC | Thermal throttling in microseconds |

Command dump-acct

| | |
|-------------|--|
| COMMAND | Command name |
| VERSION | Acct version |
| UTIME | User time |
| ETIME | Elapsed time in clock ticks [tick] |
| UID | User ID |
| GID | Group ID |
| PID | Process ID |
| PPID | Parent process ID |
| FLAG | Accounting flags: S -- command executed by the super-user F -- command executed after a fork but without a following exec D -- command terminated with the generation of a core file X -- command was terminated with the signal SIGTERM This field is only on RHEL8 environment. |
| EXIT STATUS | Process exit status This field is only on RHEL8 environment. |
| TTY | Terminal name |

| | |
|------------|---|
| | This field is only on RHEL8 environment. |
| START TIME | Process creation time |
| SID | Session ID |
| TIMESLICE | Timeslice [μ s] |
| NTHREADS | Max number of threads |
| NUMA | NUMA node number |
| TOTAL MEM | VE's total memory usage in clicks [kb * tick] |
| MAX MEM | VE's max memory usage [kb] |
| SYSCALL | The number of system calls |
| TRANSDATA | Data transfer amount between VE-VH [kb] |
| EX | Execution count |
| VX | Vector execution count |
| FPEC | Floating point data element count |
| VE | Vector element count |
| L1LMC | L1 instruction cache miss count |
| VECC | Vector execution in microseconds |
| L1MMC | L1 cache miss in microseconds |
| L2MMC | L2 cache miss in microseconds |
| VE2 | Vector element count 2 |
| VA REC | Vector arithmetic execution in microseconds |
| L1LMCC | L1 instruction cache miss in microseconds |
| VLDEC | Vector load execution in microseconds |
| L1OMCC | L1 operand cache miss in microseconds |
| PCCC | Port conflict in microseconds |
| LTRC | Load instruction traffic count |
| VLPC | Vector load packet count |
| STRC | Store instruction traffic count |
| VLEC | Vector load element count |
| VLCME | Vector load cache miss element count |
| VLCME2 | Vector load cache miss element count 2 |
| FMAEC | Fused multiply add element count |
| PTCC | Power throttling in microseconds |
| TTCC | Thermal throttling in microseconds |

'convert-acct' tool

The tool 'convert-acct' is provided to convert records from 'ac_version 3' to 'ac_version 14'. If you update VEOS from v2.5 or later to v2.6 or later, two different formats of records may be recorded in the same accounting file: version 3 format and version 14 format. Hence, this 'convert-acct' tool will be used to convert record into same v14 format. The convert-acct tool is more useful in conjunction with the lastcomm command than alone. For example: /opt/nec/ve/sbin/convert-acct <file> | /opt/nec/ve/sbin/lastcomm -f -