Collect
Version 2.0 for Tru64 UNIX
Release Notes

The following release notes for Collect for Tru64™ UNIX® (formerly DIGITAL® UNIX) describe enhancements, new features, and restrictions for this release. These notes are applicable to Version 5.0 and earlier releases of Tru64 UNIX.

Contents

Disclaimer

This is the **external field test (EFT)** release of Collect for Tru64 UNIX 5.0A. This version has been tested extensively and is a release candidate.

Preface

Collect is a data collector for operating system and process statistics, designed for high reliability and very low system overhead. Unless you set Collect to gather all possible data at frequent intervals, it does not have a significant performance impact.

For more information, see the `collect`(8) reference page.

Year 2000 Ready

Collect is Year 2000 Ready.

"Year 2000 Ready" products are defined by Compaq as products capable of accurately processing, providing, and/or receiving date data from, into, and between the twentieth and the twenty-first centuries, and the years 1999 and 2000, including leap year calculations, when used in accordance with the associated product documentation and provided that all hardware, firmware, and software used in combination with such products properly exchange accurate date data with the products.

For additional information visit the DIGITAL Brand area on Compaq's Year 2000 Ready web site located at `http://ww1.digital.com/year2000/warranty.asp.`

To ensure that this product is Year 2000 Ready, the following testing process/methods were used:

> This product has been certified as fully Year 2000 Ready. The system time was advanced to December 31, 1999 and the application operated properly as the system time progressed beyond January 1, 2000. The log files showed continuous logging activity through this time period, with accurate time stamping.

Operation Notes

The following operation notes may help you with your use of Collect:

- Some Collect operations use kernel data that is only accessible to `root`. Because good system administration practice does not encourage lengthy operations as `root`, by default Collect installs with permissions set as 04750, allowing group (typically system) members to run Collect with owner `setuid` permissions. If this is inappropriate in your environment, you may reset permissions to fit your needs.
- Collect can be configured to automatically start when a system is rebooted. This is particularly useful for continuous monitoring. To do this, use the `rcmgr` command (found in `/usr/sbin`) with the `set` operation to affect the following values:

  `/usr/sbin/rcmgr set COLLECT_AUTORUN 1`

  A value of 1 sets Collect to automatically start on reboot. A value of 0 (default) causes Collect to not start on reboot.

  `/usr/sbin/rcmgr set COLLECT_ARGS ""`

  A null value causes Collect to start with the default values of `-i60,120-f /var/adm/collect.dated/collect -H d0:L5,1w`. You may select other values.

  `/usr/sbin/rcmgr set COLLECT_COMPRESSION 1`

  A value of 1 sets compression on. A value of 0 sets compression off.
- `RAID` disks are not supported in this release.
- The binary file format has changed to Version 12, but older versions are read transparently.
- To use the file command to identify binary Collect files, as `root` append the contents of `magic.rules` in the top-level directory to your `/etc/magic` files.

Recent Changes

The following is the complete list of changes made in the last year:

- In the current Tru64 UNIX release, certain metrics are no longer available to Collect. For performance reasons these metrics are not available outside the kernel. To avoid breaking scripts that depend on these values, Collect now returns them as 0 (zero).

- Increased the size of the datatype used to store kilobits/second for LSM volumes — removed 64 megabits/second barrier.
- Changed LSM average read and write service times to floats.
- Implemented proper disk average active and pending queue service times and queue lengths.
- Changed disk average queue length and average service times to floats — queue lengths are more accurate and service times of less than 1 millisecond can be observed.
- Fixed bugs with `-T` flag in playback mode.
- Changed interval from integer to float value — fractions of seconds and subsecond intervals can be specified.
- Fixed bug that caused incorrect display of RSS and VSZ in abbreviated mode.

Historic Changes

The following changes may be of interest to early adopters of Collect:

August — October, 1998 (Version 1.09 — 1.10)

- Bug in `-T` option fixed (produced `can't read from kernel` error message).
- Fixed bug with `-PP <pidlist>` on playback not working correctly.
- Added call to `plock()` to lock program in memory (can't be swapped out anymore). Collect is now compiled statically so that all the necessary pages get locked.
- Added call to `renice()` to give Collect highest priority by default. There is a flag (`-on`) to disable.
- Added ability to play back and convert multiple data files.
- Added ability to read and write compressed data files. Compression is enabled by default for writing. Compressed input files are recognized automatically.
- Added switch (`-R`) to set a duration for Collect to run, after which it will automatically stop.
- Collect can now read from `stdin` using the `-p` option and write to `stdout` using the `-f` option.

collgui

- Added companion ability to select multiple data files.

June — August, 1998 (Version 1.08d — 1.09)

- Support for Tru64 UNIX.
- Bugs fixed that caused unusual disk configurations to produce a core dump (when SCSI floppy drive present, for example).
- Added PPID (parent PID) to data collected for each process (for determining creation hierarchy of processes in collected data). Also added the ability to select on PPID during playback.
- Added fork/vfork statistics to CPU subsystem.
- Changed LSM output such that average service time is displayedseparately for read and write operations.
- LSM not tested under Tru64 UNIX Version 5.0.
- Fixed bug when selecting processes during playback by username.
- Added new disk statistics (only available in Tru64 UNIX Version 5.0):

- AVW - Average wait in milliseconds.
- WTQ - Number of requests in wait queue.
- %WT - Percent of requests that must spend time in wait queue.

- Changed the system and user time statistics in the process subsystem to display by default a normalized delta — that is, the accumulated CPU time since the last sample normalized to 1 second. Therefore, the units are CPU-seconds per second.
- The cumulative time can still be displayed with a switch (see `collect -h`).

collgui

- Added `-VGA` switch for VGA resolution screens (640x480).
- Added `-size <size>` to allow font size adjustment.
- Hard-coded text foreground to black and text background to beige — should eliminate problems with white-on-white text under CDE.

Miscellaneous

- Added entries to `/etc/magic` so that file `collect.output` works.

June — July, 1997 (Version1.08b - 1.08c)

- Fixed bug in NET subsystem of BW being shown in increments of 3%.
- Added TTY subsystem.
- Added CD-ROMs to disk subsystem.
- Improved the *Add* list boxes — multiple objects can now be selected.
- Sped up filtering when processes are being selected. Now Collect is only asked to extract data for the PIDs requested, so cfilt doesn't have to wade through so much data.
- Division can now be used in cfilt with impunity. Division-by-zero errors are trapped, and a zero is returned as result.
- An optional arithmetic expression can now be evaluated after normalization has been done in cfilt.
- On systems for which there is no pre-build binary, a custom version will be built and installed.
- *Incompatible change*: The `normalize` directive is now a hash (#) sign instead of a percent (%) sign.

May, 1997 (Version 1.07 - 1.08)

- Added a Tape subsystem for statistics about SCSI tape drives.
- Added message queue subsystem.
- Fixed an oversight by which Collect wasn't able to seek to the last record when the records were > 32K. Records can be up to 256K before this causes problems.
- Raid disks can now also be specified using `-DreX,reY`,...
- Total flag (`-T`) now causes totals for disks and tapes to be displayed.
- Added `-PP<pid-list>`, `-PC<command-list>`, `-PU<uid-list>` options to allow collection for processes having specified parent-PIDS or belonging to specified process-group (`-PP`), having specified string in command (`-PC`), or owned by specified UID (`-PU`). Fixed bug when using the `-n` switch in playback mode.
- Added support for regular expressions for the `-L` and `-D` flags

- Added ability to convert Version 1.07 data files to Version 1.08, and simultaneously extract records for a particular interval and select subsystems.
- Changed absolute memory values from pages to megabytes.
- Disk and tape counting (for data structure allocation) and discovery has been made consistent. A side-effect of this is that a SCSI floppy drive will no longer cause Collect to dump core. (It was not being counted, but was being discovered, so Collect tried to write data into non-existent data structures.)

cfilt

- Fixed bug in cfilt that caused problems with uppercase letters in names (such as LSM volumes).

collgui

- Added new selection mechanism for lists of objects longer than can be displayed in a popup menu.
- Added a new selection mechanism for processes.
- Added a file browser for opening files.
- Added image file output (jpeg, ppm, pbm).

December, 1996 (Version 1.06 - 1.07)

- The binary format has changed. Version 1.07 will not be able to read pre-1.07 data files. If this affects you, rename the 1.06 executable to `collect1.06` (or something) and move it out of `/usr/opt/COL106/bin` if you used the `setld` kit.
- Fixed bug in collgui when two or more objects are selected and normalization is enabled. Normalization wasn't being done. This example didn’t work properly:

```
lsm:name=vol01,vol02:rkb/s+wkb/s%
```

- Fixed bug in Collect that affected the selection of disks with non-null LUNs (mostly on HSZs), such as rza33, rzb33, etc.
- New reference pages for cfilt and collgui.
- Collect now flushes `stdout` when writing to it, and cfilt will immediately process data if no normalization is being done, so it is now possible to do the following:

```
collect -i1 -F -sp | cfilt proc+:user=smith:rss
```

  to sum RSS for user `smith` for all processes owned by `smith`.

- Fixed bug in load average that caused all load averages over 2.55 to be misrepresented.
- Superficially tested RAID (SWXCR) support.
- In collgui, changed the `Memory Used` expression in the memory subsystem to use the cfilt expression `Active+Inactive+Wired` because `Active` already contains the UBC pages.

October, 1996 (Version 1.05 - 1.06)

- Added to process info input blocks, output blocks, major faults and minor faults. This caused the line to be longer than 80 columns, requiring the next fix.
- Added the `-F` option. Normally, not all fields are printed, and some things, such as RSS and VSZ are converted to compact format using `K`, `M`, and `G` tags. The `-F` option causes all information to print, in expanded form.
- Added to memory info processes swapped per second, UBC hits/sec, UBC pages pushed (written)/sec, and UBC pages allocated/sec.

- Improved the `-C` (chop out a timeslice) flag. The format is somewhat different: Either the start or end time can be in the format [+]Year:Month:Day:Hour:Minute:Second. If there is a plus-sign at the beginning, then the time specified is relative to the beginning of the collection period. If there is no plus sign, then the time is absolute. Any of Year, Month, Day, Hour, or Minute can be left off, in which case the values from the beginning of the collection period are taken. Therefore, if you collected from 10:23:34 to 10:44:15 on October 16th, you could use -C24:00,25:00 to extract samples from 10:24:00 to 10:25:00. You could also use: -C10:24:00,10:25:00, or even C16:10:24:00,16:10:25:00, etc. If no start time (-C,10:25:00) or no end time (-C10:24:00,) is given, then the collection start and end times are used, respectively.
- Added (`-l`) option to seek to and print the last valid record. This aids collgui, reducing the time to open a large binary data file.
- Adjusted collgui to use the new time format of Collect.
- New reference page for collgui.
- New reference page for cfilt.
- Added call to task_set_notify_port (suggested by Eric Schott) so that kernel doesn't store non-deliverable mach messages and eat memory.

September, 1996 (Version 1.04 - 1.05)

- Improved cfilt to understand almost any arbitrary expression, such as `(100-rkb/s+wkb/s)/100"` or `"log(idle)`.
- Added a graphical front-end for collect, cfilt, and gnuplot called collgui.
- Integrated functionality of `cavg` into cfilt with the `-a num` switch.
- Added the `-p` flag to cfilt to select only those samples that contain process data.
- Saved 8 bytes/process record (but these will soon be gobbled up by new process-record elements).

September, 1996 (Version 1.03 - 1.04)

- Tested and fixed LSM support under DIGITAL UNIX Version 3.x and Version 4.0.
- Added `-L` option for specifying LSM volumes to Collect or playback.
- Cleaned up `on_disk.h.`
- Added the %busy calculation for disks. This is based on:

$$\frac{TimeSpentDoingIOs}{RealTimeInterval} \times 100$$

  Generally, a disk will be 100% busy during a transfer, so it's just a question of for how much of the measurement interval was there an outstanding request. This has nothing to say about number of transfers, kilobits or kilobits/second transferred, etc.

- Added `-DLSM` to make file to explicitly build collector with LSM support.
- Added reference page for Collect.
- Added -v flag to print program and data file version.

August 26, 1996 (Version 1.02 - 1.03)

- Changed `-u` (use) to `-s` (select). Sorting is now `-S` (capital).

- There is a new output field for disks, `%BSY` (percent busy). This only works under DIGITAL UNIX Version 4.0x. It is simply the service time (the time it took for the disk to complete all requests) of the disk over the interval divided by the interval.
- Added the magic number to the main header as well, so when a file is given for playback that was not written by Collect, a reasonable error message is printed.

August 8, 1996 (Version 1.01 - 1.02)

- Fixed sorting bug: `-s -nX` now works properly.
- `-D<disklist>`, `-P<proclist>`, `-s -nX` options now work in playback mode.
- New option, `-C[starttime,endtime]`, added to extract samples from a particular time-range.
- Collector now waits one interval before writing first sample. This solves the problem of the first sample having bogus values due to false delta values.
- Fixed bug with non-printable characters in header after `CPU-TYPE`.
- Added Perl script Scripts/filter that is suitable for extracting values to be read into Microsoft® Excel or gnuplot.

July 14, 1996 (Version 1.00 - 1.01)

- A list of disks for which statistics are to be gathered can now be given to the collector in the form `-DrzXX,rzXY,...` (only works in collect mode, not playback).
- When a non-privileged user runs Collect and Collect is setuid-root, if a binary-output file is opened, the open is performed as the non-privileged user. This should make the program completely safe for setuid-root. If the binary-output file already exists, the user is prompted for confirmation before overwriting.

Kit location

Collect will be included with the next release of Tru64 UNIX, and it will also be included in the next patch kit for Tru64 UNIX, expected in early 2000. Currently, Collect is freely available from the following FTP locations:

- Within Compaq

  `ftp://ftp.zso.dec.com/pub/collect` (primary site, United States)

  `ftp://nasaxp.rto.dec.com:collect/` (secondary site, Germany)

- Outside Compaq

  `ftp://ftp.digital.com/pub/DEC/collect` (US)

  `ftp://ftp.digital.de:/pub/DEC/collect/` (Germany)

These are the Collect kits at these sites:

`COLLECT200SETLD.tar` (`setld` kit)

`collect-2.00.tar.Z` (compressed version)

These Perl kits are also available if you want to use the GUI or filter scripts:

`PERL5003SETLD.tar` (DIGITAL UNIX Version 3.0+)

`PERL5004SETLD.tar` (DIGITAL UNIX Version 4.0+, and all versions of Tru64 UNIX)

Perl

The above setld kit for Perl includes the Tk extension. If you want to use the graphical user interface for Collect, collgui, you must get this kit and install it using the standard software installation utility setld.

If you already have Perl, you may not have the Tk extension for Perl (not the same as Tcl/Tk). You can still use cfilt. Edit /usr/bin/cfilt and set the path of perl to where you have Perl installed. If you also have Tk for Perl, you need to edit /usr/bin/collgui and set the path in the same way.

Questions and Comments

If you have any questions or comments, please send email to the Collect team at this address:

collect_support@zso.dec.com