

Installing T4 Version 4.0

29 June 2005

WELCOME TO T4 VERSION 4.0

Thank you for downloading the latest T4 kit from the public web site located at:

<http://h71000.www7.hp.com/openvms/products/t4/index.html>

To get maximum benefit from the T4 data you will be able to collect with this kit, you will want to download the TLViz and CSVPNG kits which are available from the same web site. TLViz and CSVPNG are specifically designed to save you time when you are analyzing the CSV (Comma Separated Value) files that the T4 collector creates.

If you have questions about T4, TLViz, CSVPNG or other Friends of T4 , or if you would like to be on the mailing list, please send mail to T4@hp.com.

This document contains the following sections to help you get started with T4

INSTALLING THE KIT

- THE T4 V4.0 KIT
- PRIVILEGES REQUIRED FOR T4 INSTALLATION
- RESET FILE ATTRIBUTES
- DEASSIGN T4\$SYS LOGICAL NAME
- INSTALL THE LATEST T4 KIT
- MAKING THE T4 LOGICAL NAMES PERMANENT
- CREATING A T4\$DATA STORAGE AREA
- SET UP A LOCAL BATCH QUEUE TO RUN T4

RUNNING T4

- REQUIRED QUOTAS TO RUN THE T4 COLLECTION
- PRIVILEGES NEEDED TO RUN T4 COLLECTION

DISK SPACE

- ESTIMATING DISK SPACE REQUIREMENTS
- RUNNING A T4 V4.0 TRIAL SESSION
- LAUNCHING T4 IN LONG TERM HISTORY MODE
- TYPES OF FILE CREATED IN T4\$DATA

TERMS & CONDITIONS FOR USE OF T4 V4.0

INSTALLING THE KIT

THE T4 V4.0 KIT

Your zipped up T4 kit includes the following files:

- 1) **HP-VMS-T4-V0400--1.PCSI\$COMPRESSED;1** - The T4 V4.0 PCSI Kit
- 2) **How-to-install-T4.PDF** The PDF version of this document
- 3) **How-to-install-T4.TXT** The text file version of this document.

PRIVILEGES REQUIRED FOR T4 INSTALLATION

You will need the SYSNAM privilege to establish the logical names required to successfully install T4 and make it ready to run.

RESET FILE ATTRIBUTES

Once you have moved this PCSI kit to your OpenVMS system, you may need to RESET the FILE ATTRIBUTES. These can sometimes be altered in transit depending on how the transfer was accomplished. If necessary, the appropriate reset can be accomplished with the following command:

```
$ set file/att=(rfm:fix,lrl:512) HP-VMS-T4-V0400--1.pcsi$compressed
```

DEASSIGN T4\$SYS LOGICAL NAME

If you are already running T4 V3.3 or T4 V3.4, be sure you DEASSIGN your current T4\$SYS system wide logical name before installing the new kit.

```
$ DEASSIGN T4$SYS /SYS
```

INSTALL THE LATEST T4 KIT

You can now install the latest T4 v4.0 kit using:

```
$ PRODUCT INSTALL T4
```

This installs the kit to: SYS\$SYSDEVICE:[VMS\$COMMON.T4\$SYS]

To install to a different directory use the /DESTINATION option

```
$ PRODUCT INSTALL T4/DESTINATION=disk$somewhere:[dir.tree]
```

If you install the T4 kit to a disk drive that is common to an OpenVMS cluster, then all the nodes on that cluster will be able to launch T4 sessions from that common location.

Once you have installed the T4 V4.0 kit, you will find many more files of additional reference material.

MAKING THE T4 LOGICAL NAMES PERMANENT

This is a very important step that helps make sure you get full use out of T4 on your systems. By making key T4 related logical names permanent, you enable interrupted T4 sessions to restart following a system reboot and make sure that you don't miss any vital data.

Define the T4\$SYS logical and make it permanent by adding the definition to the system startup SYLOGICALS.COM command procedure. For example, try the following command.

```
$ DEFINE /SYSTEM /EXEC T4$SYS SY$$SYSDEVICE:[VMS$COMMON.T4$SYS]
```

CREATING A T4\$DATA STORAGE AREA

When you run T4 collections you will need a convenient and suitably sized disk location to store the T4 generated performance data. Set up your data disk in advance and point to it with the T4\$DATA logical name. It's equally important to make this logical name permanent by adding the definition to the system startup code SYLOGICALS.COM command procedure which allows T4 sessions to resume following system reboots.

```
$ DEFINE /SYSTEM /EXEC T4$DATA Your_Data_Disk:[000000.T4$DATA]
```

You can establish a separate T4\$DATA area for each OpenVMS system or you can set up a single T4\$DATA disk that is available to all nodes on your OpenVMS cluster.

WARNING: Do not use your system disk for T4\$DATA.

SET UP A LOCAL BATCH QUEUE TO RUN T4

For each OpenVMS node that you will want to monitor with T4, you will need to use a **local batch queue** to run the T4 jobs on that node. This can be an existing

queue (assuming it has available job slots) or you can set up a special T4 batch queue for that node as shown below.

First, start the queue manager if not already started

```
$ start/queue/manager/new
```

Create and start up a new batch queue

```
$ init/que/start/batch T4$batch /job=10
```

```
$ sho queue/batch/full
```

Batch queue T4\$BATCH, idle, on NODEX::

```
/BASE_PRIORITY=4 /JOB_LIMIT=10 /OWNER=[SYSTEM]  
/PROTECTION=(S:M,O:D,G:R,W:S)
```

You will need one batch queue and one T4 collection session for each OpenVMS node you are interested in. As noted above, all the collected and processed data from these sessions will be saved to the T4\$DATA directory.

RUNNING T4

REQUIRED QUOTAS TO RUN T4 COLLECTION

The User Account (e.g. T4_USER) that you plan to use to run T4 will require the following QUOTAS.

PRCLM must be at least 20.

TQELM must be at least 100.

PGFLQUOTA must be at least 500000.

```
$ mc authorize
```

```
UAF> mod T4_USER /prclm=20/tqelm=100/pgflquota=500000
```

```
UAF> exit
```

PRIVILEGES NEEDED TO RUN T4 COLLECTION

For the user who will be launching T4 history creation sessions, the ALTPRI privilege is needed so that the OpenVMS Monitor Utility can run at the recommended process priority of 15.

DISK SPACE

ESTIMATING DISK SPACE REQUIREMENTS

For your most important and performance sensitive systems, we suggest that you launch T4 in long term history mode by answering yes when asked to the question:

Re-Submit data collection job daily [N] : Y

This way it will build up a detailed performance history for you day by day.

The actual size you will need for your T4 history area depends on several key factors:

- * the number of nodes under measurement,
- * the number of hours of measurement each day,
- * the number of devices to be measured
- * the number of processes on each system,
- * the sampling rate (default = 60 seconds)
- * the OpenVMS version

To gauge the amount of disk space that you will require, we suggest you start with at least 500,000 blocks. Run a trial session (see quick instructions below) of 1 to 2 hours and determine how much disk space is needed for that run.

Then adjust the size of the T4\$DATA storage area as necessary to meet your needs.

WARNING: Do not use your system disk for T4\$DATA

T4 includes some rudimentary capabilities for assisting you in the management of your T4\$DATA performance history area. Since the data that T4 helps you collect may prove invaluable to you in the future, you will want to think through and apply your standard local site policies for backing up, archiving, and preserving this potentially priceless historical system information.

RUNNING A T4 V4.0 TRIAL SESSION

Once you have installed T4, set up your T4 user account with the appropriate privileges and process limits, established a T4 Batch queue on each node you will be monitoring, and created a T4\$DATA area, you are now ready to launch your first T4 V4.0 collection session.

We suggest you run a trial session to help calculate the disk space that you will require to run T4 in its recommended long term history mode (see the OpenVMS Technical Journal Article at <http://h71000.www7.hp.com/openvms/journal/v3/t4.pdf> for a full explanation of the benefits of long term performance histories

To help estimate the amount of disk space you will require, run the following command and respond to the questions you are asked following the guidelines shown below in answering the questions poses to you.

\$ @T4\$SYS:T4\$CONFIG.COM

Start Time pick a time that is 5 minutes in the future as this will give you enough time to work your way through all the questions. That way the T4 collection session will be launched prior to the time you specify and you can make sure that you get a full hour of data..

End Time pick a time that will give you a total run time of one hour.

Batch queue remember to use your T4\$BATCH queue that is local to the node you are measuring

Network Interface Device Enter '?' to get a list of all available network interface devices. Then, highlight that list, and paste it in as the answer to the question. This will cause T4 to monitor each and every one of your Network adapters.

Sampling Interval use the default 60 seconds

Data Directory use T4\$DATA

Automatically manage T4 data storage answer YES to this question and then use the default values. This will give you a preview of the rudimentary file storage management capabilities that are part of T4 V4.0.

Re-submit data collection job daily answer NO to this question as you are using this run to help you determine storage needs before launching T4 in long term history collection mode.

Email Address if you can send email from the VMS node being measured, fill this in so it sends email notifying you that the data has all been collected and processed.

Once the run is complete, check the sizes of the files created in the T4\$DATA directory to discover the approximate storage costs per hour for this OpenVMS node. See the sections below describing the types of files created during a T4 session and the suggested retention periods of each type of data.

LAUNCHING T4 IN LONG TERM HISTORY MODE

Having established a reasonable size for your T4\$DATA area by following the steps outlined above, you are now ready to launch T4 in long term history mode. Long term history mode for T4 is explicitly designed to help you maximize the benefits that T4 can provide. At the same time it helps to minimize how much time you will have to spend and keep that down to an absolute bare minimum.

Here are some recommendations for parameter values to use in response to the questions triggered by running the following command

`$ @T4$SYS:T4$CONFIG.COM`

Start Time use the default which will be tomorrow at one minute after midnight.

End Time use the default which will be tomorrow at 23:59. These settings for Start Time and End Time mean that you will have round the clock performance data for the systems that are most vital to you.

Batch queue remember to use your T4\$BATCH queue that you created and that is local to the node you are measuring.

Network Interface Device type in the names of the devices that are a regular part of your production environment. You can Enter '?' to get a list of all available network interface devices. If you want to measure all of them, you can highlight that list, and paste it in as the answer to this question.

Sampling Interval use the default 60 seconds. This has proved to be an excellent default compromise value for long term history creation.

Data Directory use T4\$DATA

Automatically manage T4 data storage answer YES to this question.

Number of days to retain raw data we recommend setting this as long as possible with a minimum value of 45 days. We suggest that you consider archiving this data and retaining it for at least 18 months if you are unable to keep it all on-line.

Number of days to retain intermediate data use the default of 3 days.

Number of days to retain reduced files we strongly recommend you use as large a value as possible with a minimum of 500 days. These reduced files are the T4 style CSV files that contain your complete performance history, minute by minute, day by day, for the most important performance metrics on your system. They are the prime reason for running T4 in long term history mode. We strongly recommend that old data of this type be archived for at least 3 years.

Re-submit data collection job daily answer YES to this question. When you do this, each time T4 runs, the first thing it will do is submit a new batch job for the following day. If you have established T4\$DATA and T4\$SYS as system wide logical names in your SYLOGICALS.COM file, then this single launch operation will create a full long term history for this node.

Email Address using this or not using this feature is your call.

Your T4 collections will start tomorrow at one minute after midnight and continue collecting and saving vital performance data for you until the batch jobs are deleted.

Welcome to the world of T4. We look forward to hearing of your experiences in using these handy tools.

TYPES OF FILE CREATED IN T4\$DATA

T4 produces a composite CSV file daily for each node being monitored. The names of these files are of the form:

T4_<nodename>_<Collection_Date>_<start_time>_<end_time>_COMP.CSV

For example, a one hour run on node PRFE40 might look like:

T4_PRFE40_28JUN2005_1400_1500_COMP.CSV

These *COMP.CSV files are normally the first thing that we look at.

The output of a T4 collection session includes

CSV files
ZIP files
LOG files
DAT files (for OpenVMS MONITOR and for T4FCMON)

The key CSV files are:

- *COMP.CSV
- *DISK.CSV
- *SCS.CSV
- *T4FCMON.CSV

The remaining CSV files are what we refer to as intermediate files and they can be deleted after a few days.

We recommend that you retain all the CSV files or their zip files equivalent for at least eighteen months. We recommend retaining the MONITOR.DAT files for at least 45 days.

Even beyond eighteen months, if you find you must remove the older CSV files, we suggest you send these to a permanent archival storage location that you can access as needed. You'll never know when you will have a need to search back 2 or 3 years to compare what is happening right now to how things used to be. You will only be able to do this if you carefully guard all of these invaluable performance timeline files.

TERMS & CONDITIONS FOR USING T4 V4.0

T4 V4.0 is provided to you, subject to the following terms and conditions. Any comments or questions regarding T4 V4.0 or its terms and conditions can be directed to T4@hp.com.

- (a) T4 V4.0 is supplied 'as is.' without warranties, either expressed or implied.
- (c) T4 V4.0 is not a Commercial 'Off-The-Shelf' software product.
- (d) T4 V4.0 is neither freeware nor shareware, and cannot be freely distributed other than by Hewlett-Packard Company, subsidiaries, successors and assignees.
- (e) T4 V4.0 remains the sole intellectual property of Hewlett-Packard Company, subsidiaries, successors and assignees.
- (f) T4 V4.0 may not be redistributed or supplied to any third party, either for commercial gain or otherwise without the prior written consent of Hewlett-Packard Company.
- (g) T4 V4.0 is deemed to be supported *only* by consulting services purchased directly from Hewlett-Packard Company, subsidiaries, successors and assignees, specific for this purpose and by prior agreement.
- (h) The recipient must not request any kind of support from any other Hewlett-Packard Company entity, such as any Customer Support Center or Engineering group. This support restriction also applies to any issues of coexistence or interoperability with any other software or hardware, including (but not limited to) supported Hewlett-Packard Company products.

WHERE TO GET MORE INFORMATION

If you have questions, suggestions, comments, or if you want to be on the T4 & Friends mailing list then please send mail to T4@hp.com.