



Taking T4 to the Next Level.....	1
BIO.....	2
T4 origins – our time is our most precious resource	2
Necessity is the mother of invention.....	3
Time is of the essence	3
The T4 name	3
T4 - the early years.....	3
T4 successes and improvements.....	5
T4-compatible collectors and extractors.....	6
The latest T4 news.....	7
T4 future trajectory – building on the foundation	7
T4 resources.....	8

BIO

Steve has made hands-on practical use of T4 from its earliest beginnings. He actively employs T4, TLViz & CSVPNG as essential ingredients of his performance work for TrendsThatMatter. Steve's [audio segments on T4 technical tips](#) are a regular feature of the [VMS Podcast Network](#) (VPN) program and Steve wrote the [first in-depth article on T4](#) that appeared in [Volume 3 of the OpenVMS Technical Journal](#). The [easy reader version](#) of the article (just created recently) is a great way to begin learning about T4 and to get a sense for the visual nature and the power of this approach.

Steve was a member of the OpenVMS Engineering Performance Group for 11 years and he has more than 25 years of enterprise system performance management experience. Steve has been a regular speaker about T4 at every [OpenVMS Boot Camp](#) and at the popular OpenVMS Technical Updates in Europe and in the United States. He has delivered seminars and hands-on workshops on practical performance management at more than two dozen locations in the U.S, Canada, Europe, and the Pacific Rim. Steve also participated on the Transaction Processing Council in its formative years while it was developing and defining the TPC-C benchmark. You can learn more about Steve at [TrendsThatMatter.com](#).

Overview

It is hard to believe that the T4 approach has now been around for almost seven years, and that year by year it has achieved wider acceptance while incrementally improving its underlying capabilities. This is easily seen in the long list of enhancements added to the T4 collector, the large number of productivity changes for TLViz and CSVPNG, and the step-by-step increase in the number of T4 compatible collectors, now totaling more than 20. We have seen time and again that the T4 approach encourages collaboration while doing performance work.

This article briefly looks back at T4's origins and where it has been, and then projects the likely trajectory it will take in the years to come as it builds even further improvements onto an already solid foundation. This article also provides you with references to all the latest T4-related resources and to where you can turn for help and advice so that your OpenVMS performance work can achieve the maximum benefit from this approach.

T4 origins – our time is our most precious resource

T4 began its life seven years ago during a large-scale benchmark just before the release of the AlphaServer GS160 system. Tom Cafarella (OpenVMS Engineering) and I were assisting an important customer compare the GS160 performance to that of the largest AS8400 system available.

The benchmark was especially complex because of its use of captured log data from a live production environment to drive the load. Each test showed highly variable behavior during its brief, 30-minute run – a far cry from steady-state industry-standard benchmarks, which prove much easier to analyze.

We were using the OpenVMS Monitor utility (MONITOR) to watch OpenVMS performance while the customer was using their own software logging program to measure and report on application response time and throughput.

The customer wanted to test many different variations of CPU count, memory size, disk configuration, tuning parameters, and so on. However, the results we were getting were hard to understand and even harder to explain. We were finding that it took 1 to 2 days to investigate and analyze even a single head-to-head comparison between the GS160 and the AS8400 systems. The customer hoped to test at least a dozen different options per day in order to complete the work in the allotted time. Needless to say, Tom and I were under a lot of pressure.

Necessity is the mother of invention

The T4 approach was born into this cauldron of activity. Tom Cafarella wrote a DCL script that started with a `Monitor.DAT` file and ended up producing a comma-separated value (CSV) text file. The text file contained what Tom and I considered the 30 most important OpenVMS performance factors. These included CPU idle, kernel mode time, interrupt mode time, mpsynch time, buffered I/O, paging I/O, and direct I/O. Tom's code replaced a manual process we had been using that took us hours or even days to finish and that was not nearly as complete. We could now get a timeline view of how all 30 of these important factors varied during the complex course of the benchmark. And we could now complete this first level analysis within a few minutes after the end of the test. This development changed everything.

The CSV file that we created later became known as T4-style format. The file contained a table that had a few header rows, including one where each performance factor was listed by name, followed by a set of rows, with exactly one row for each sample time period. Each column of the table represented a single performance factor. Because it was in CSV format, and because all the data for a single factor was contained in a single column, we could open this file with Microsoft® Excel software and then, with about six mouse clicks, plot the trend lines for that factor.

Time is of the essence

This time savings achieved with Tom's DCL script made a huge difference in our ability to analyze these complex benchmark test runs in a reasonable amount of time. Whereas it had been taking days per run, we were now able to discover important lessons about a run within an hour or two of the end of the test. We could then use that information for feedback into the ongoing process, thereby enabling more focused selection for the next round of benchmark tests.

The T4 name

In honor of Tom's yeoman and timely efforts to create the DCL conversion script, the original name "T4" was meant to represent Tom's Terrific Timeline Tool. Since then, as T4 evolved and a new compiled version of the extractor was constructed, the T4 name came to mean Total Timeline Tracking Tool. Given how much time this approach has saved us over the years, we might think of renaming it yet again to something like Time-saving Trend Tracking Tool. Whatever name we give it, because of its significant ability to save our time, the T4 approach has completely transformed the way we carry out performance work on OpenVMS systems.

T4 - the early years

Within a very short time after its first use, the OpenVMS Engineering organization began to apply T4 to real-life production performance problems. Just like the original benchmark, T4 was perfect in these situations because production systems rarely show steady-state behavior. To deal with the ongoing variability of each important factor, it is essential to first grab a time-series view of how that factor changes over time, and then to be able to examine it visually. The T4 approach evolved rapidly in response to a spike in performance issues related to the introduction of the GS160 system.

The extraction of MONITOR data was rewritten by Ian Megarity (OpenVMS Engineering) as compiled code, and many new performance factors were included, such as detailed data about each CPU and about each disk.

New collectors that generated output in the T4-style CSV format were then added that captured important performance factors not available with MONITOR. These included a TCP/IP collector called TPCMON (developed by Matt Muggerridge) and a dedicated lock manager collector (developed by Ian Megarity). The collectors were bundled into a new DCL script (that we called T4)

so that MONITOR and all the associated collectors started and stopped at the same time, and all used the same sampling rate.

Another new tool called Append Record (APRC) was created by Ian Megarity to consolidate the CSV files generated by the different collectors into a single, larger CSV file (known as the *COMP.CSV file) containing all the factors from different collectors.

Saving more time with TLViz. Another magic piece of puzzle was completed when Ian Megarity created a new timeline visualization tool called TLViz (TimLine Vlsualizer). TLViz made it incredibly easy to visualize the trend data of each individual factor. And TLViz made it just as easy to view any interesting combination of factors. Our experience with this tool proved that we could analyze data at least 5 times, and often 10 times, faster than when we read the same CSV file into Microsoft Excel. Kevin Jenkins (from OpenVMS Engineering) suggested the idea of opening *before and after* CSV files and having TLViz overlay the data. Ian Megarity implemented version 1 of that idea literally overnight, changing forever and for better the way OpenVMS performance engineers carried out their analysis and comparison of complex production systems.

Because of its visual nature, TLViz also proved to be amazingly good for doing collaborative work with a group of analysts, as well as being a most effective tool for demonstrating the results of analysis to both technical and non-technical audiences. If you care about performance but you have not yet tried TLViz, I guarantee that you are in for a treat when you first experience its features and its built-in time-saving capability.

To give you a flavor of TLViz and its powerful Before and After feature, Figure 1 shows a visual example from the more than two dozen charts contained in the [simplified version](#) of the original T4 article in the OpenVMS Technical Journal. While these kinds of graphs in a static format are useful, we have found that direct use of TLViz can be even more powerful. We have also discovered that interactive use of TLViz to present results turns out to be much more effective than creating static reports.

The Power of Before & After Upgrading a GS160 to a GS1280

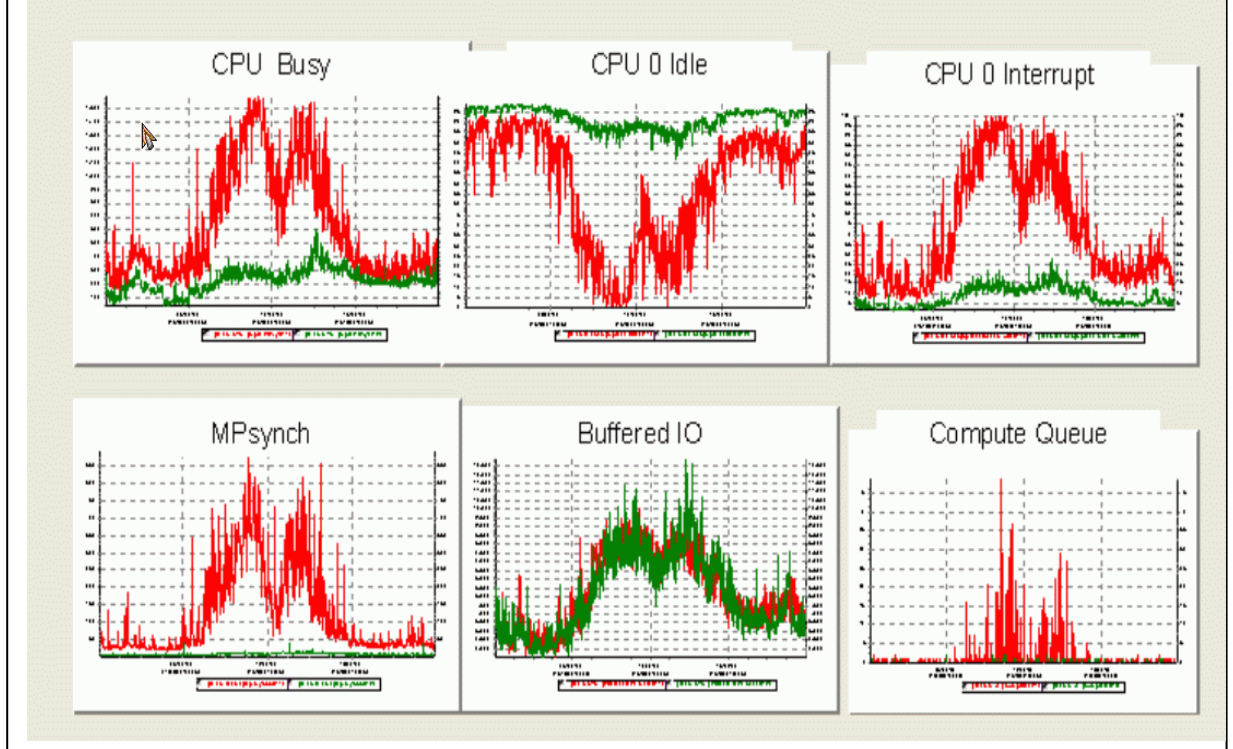


Figure 1 – Type and extent of changes as reported by TLViz

Automating repetitive work with CSVPNG. Building on the T4 and TLViz foundation, Pat Moran (HP Support, Ireland) invented yet another new tool, called CSVPNG. Among other things, this tool can automatically convert T4-style CSV files to PNG graphics format. This has proved to be the perfect complement to TLViz in that it allows the analyst or system manager to program repetitive performance-analysis steps using a DCL script that includes one or more CSVPNG command-line commands. Because CSVPNG runs on OpenVMS Alpha and HP Integrity servers, you can add your customized CSVPNG string to the tail end of each T4 collection session. This is exactly the kind of job that fits well with the skill set of trained OpenVMS system managers.

T4 successes and improvements

As the years rolled by, the T4 approach continued to evolve, often at a rapid pace. More and more OpenVMS customers pressed it into service on their most important and largest-scale systems. It proved helpful right out of the box and even more so when small custom enhancements were added. The following sections briefly highlight some of the successes and improvements that have been added.

T4 collector improvements. The first T4 involved manually running MONITOR and then manually running the T4 DCL extractor code to produce a CSV file that could then be read into Microsoft Excel. Very early on, Ian Megarity combined these two steps and added an automatic way to run the T4 collector once and instruct it to turn on what might be called History Mode. In this mode, the collector continued measuring and creating T4-style CSV files, once per day, 365 days a year. With History Mode turned on, we were assured that we would always have the baseline data we needed for comparison whenever anything changed or whenever a new problem arose.

The T4 DCL script continued to add new collectors and to integrate them as part of a standard T4 collection. These now include a network adapter collector for each network adapter, a basic XFC collector, and a Fibre Channel Monitor collector named FCMON. FCMON provides detailed response time and kilobytes per second KB/s rates for each disk from the OpenVMS server point of view.

Optional collectors for spinlock data and for TDC data are also now available with the prerelease version of T4 version 4.2.

TLViz improvements. Thanks to the continued efforts of Ian Megarity, TLViz experienced a steady progress of new features that made it easier to use and that allowed the analyst to dig deeper and deeper in reasonably short periods of time. Some of the most useful features added include moving averages, calculation of the ratio of two factors, scatter plots and correlations between a pair of factors, area charts and 100% area charts, and bar charts for cases with relatively few data points.

TLViz added some very powerful capabilities for quickly navigating CSV data sets that contain a large number of different performance factors. This meant we were able to apply its time-saving visualization to more and more complex situations, including cases where there were literally thousands of separate factors to consider.

Recently, Pat McConnell (OpenVMS Engineering) has added some terrific new features in version 2.0-1. There are now more “radio” buttons on the bottom left of the screen that simplify and speed up some common and important operations that arise frequently during analysis situations.

CSVPNG improvements. In the capable hands of Pat Moran, CSVPNG has continued to evolve, improve, and add new functional capability. CSVPNG now includes more than 30 separate command-line directives that work individually or together to manipulate, automate, visualize, and report on T4-style CSV files. CSVPNG is particularly powerful now for trimming down a huge data file by selecting exactly the performance factors you are most concerned with and further selecting the precise time periods that interest you. CSVPNG contains multiple options for combining T4-style CSV files in different and useful ways. These include assembling a week or month of data into a single file and combining data from multiple nodes in a cluster.

CSVPNG can also input and apply its magic to an increasing number of files of slightly different formats, including tab-separated files from the Linux COLLECT-L utility, and outputs from the Microsoft Windows® PerfMon utility.

T4-compatible collectors and extractors

A growing number of T4-compatible collectors and extractors continue to appear. Recently, we have seen the introduction of the VEVAMON collector that runs on OpenVMS and collects detailed storage performance factors from EVA5000 and EVA8000 systems. We also have heard that a new approach for extracting HP’s XP storage data and transforming it into T4-style format is also in the works.

Other T4-style collectors have been built for Rdb, Oracle®, additional OpenVMS memory factors, the Collect utility for Tru64 UNIX®, and NFS.

With its latest releases for the EVA, HP Storage now includes instructions for how to convert EvaPerf data into T4-style, TLViz-compatible format.

© Copyright 2007 Hewlett-Packard Development Company, L.P

These developments are part of a wave of growing appreciation for the idea that the T4 approach is not just for OpenVMS. The conclusion: When data is organized in T4-style files, time-saving, enhanced productivity, improved depth-of-analysis benefits, strengthened collaboration, and greater ease of reporting are immediately available.

The latest T4 news

New versions of TLViz and CSVPNG are now available. For more information, see the links at the end of this article.

A number of developments are underway to build additional analysis tools that will accept T4-style data as input and provide additional analysis benefits not currently available. We will be reporting about these on the [TrendsThatMatter blog](#) as soon as they become available.

T4 future trajectory – building on the foundation

The T4 approach has proved that it works and that it saves huge amounts of time for the analyst. It has proved that it benefits OpenVMS performance work. It has also proved that any trend data can be assembled into a T4-style format and thereby tap into T4's time-saving benefits.

A solid foundation is now in place. A substantial percentage of mission-critical OpenVMS systems now in production operation have been collecting T4 data and building performance histories that go back one or more years.

New collectors, when they became available, have proved to be easily integrated into the picture because the basic T4 approach is inherently extendable. This is a huge area of opportunity because we have found that while "out of the box" benefits are good, the benefits of local customization are even better.

Even after 7 years of T4 development and the existence of more than 20 collectors, some vital metrics still are not regularly collected in a form that allows them to be viewed easily as trend data or integrated with existing performance data. For example, there is no easy way today to observe CPU chip-level behavior such as Data Cache Miss Rates as they change over time. We are certain that once such a chip-level trend data collector becomes available, it will open the door to identifying opportunities for radical performance improvements. For more information about this topic, see [Extending the T4 Approach to New Territory](#).

In addition, the users of TLViz and CSVPNG continue to come up with good ideas for how to make these two important tools even more useful and powerful and even better at saving our precious time.

For those who want to extend the life of their current AlphaServer systems, the T4 approach can provide you with the data and analysis tools that help you push your systems as far as they are capable of going.

For those moving toward using the new powerful, Dual-Core Intel® Itanium® 2-based Integrity server systems, the T4 tool set can help you get ready for that transition, make sure that you know how much spare capacity you have and how fast you are using it, assist you in sizing your new system, and validate that the new system delivers the performance level you expected.

For those who are thinking of consolidating many current OpenVMS systems into a new, larger system and of taking advantage of the growing virtualization capabilities, your T4-based performance history on your existing systems will prove to be essential input data. When dealing with consolidation, CSVPNG's ability to combine data from separate CSV files will surely prove to be an invaluable time saver that can help you achieve the best possible results.

T4 resources

The following links show you where to find the most important T4-related information and where you can turn for help and advice along the way. If you are interested in [The T4 Universe](#), these resources will help you both get started and keep up with the latest developments.

If you have T4 issues or questions: please send email to [and](#) and to T4@TrendsThatMatter.com.

Want an advance copy of T4 V4.2 with the option for Spinlock Tracing and TDC collection? Send your request to [and](#)

Struggling with a particularly complex performance situation where the T4 data just doesn't add up and you are looking for a quick second opinion? You can send email to [T4 Mini-Evaluation](#).

Would you like to know more about the TLViz Formatter that you can use with EVAPerf data? You can find details in the [HP StorageWorks Command View EVA user guide](#).

Looking for the current version of the T4 collector for OpenVMS and other T4 approach tools? You can download these from the [HP T4 Home Page](#). Also see the [HP T4 FAQ](#).

The latest T4 & Friends news can be found by checking on the [TrendsThatMatter blog](#).

You will find some useful downloads of new versions of T4-style tools at the [TrendsThatMatter download page](#).

New to T4? The [TimeLine Collaboration](#) article from the [OpenVMS Technical Journal, Volume 3](#) is a good source of detailed information about the T4 approach to performance and about all the many people who have played a role in its development.

A streamlined version of the TimeLine Collaboration article ([Quick Intro to T4](#)) is a great place to start learning about T4. This work is inherently visual and is most understandable when you look at the actual performance trend data with tools like TLViz and CSVPNG. This quick introductory version includes many easy-to-understand examples.

Interested in using TLViz or CSVPNG to examine the rich range of T4-style data available from the many different collectors? Please send mail to: [T4-style data samples](#).

Rather learn more about T4 by audio? You can check out the full set of T4 segments from the [VMS Podcast Network](#) (VPN) at: [T4 Audio Segments](#).