



An Oracle White Paper
September 2011

Tuning the SPARC CPU to Optimize Workload Performance on SPARC T4

Tuning the SPARC CPU to Optimize Workload Performance on SPARC T4 Systems

Starting with the [Oracle VM Server for SPARC 2.1](#) patch release (147507-01), you can use dynamic CPU threading controls to optimize workload performance on [SPARC T4 systems](#). These threading controls enable you to specify the number of hardware threads to be activated per core. Existing applications can take advantage of the SPARC dynamic threading performance benefits without having to be rewritten or recompiled.

This paper describes how to use the Oracle VM Server for SPARC 2.1 CPU threading controls to optimize CPU performance on SPARC T4 platforms. CPU performance can be optimized for CPU-bound workloads by tuning CPU cores to maximize the number of instructions per cycle (IPC). Or, CPU performance can be optimized for maximum throughput by tuning CPU cores to use a maximum number of CPU threads. By default, the CPU is tuned for maximum throughput.

This paper covers the following topics:

- “CPU Threading Modes and Workloads”
- “Setting the CPU Threading Mode”
- “Threading Control Limitations”

CPU Threading Modes and Workloads

On SPARC T4 systems, you can optimize CPU performance by specifying the CPU threading mode. The threading mode can be set dynamically and independently for each domain on the system. A reboot is not required to change the threading mode, and the set mode is maintained across domain reboots or platform power cycles. Note that you *must* perform an initial reboot to properly configure the system.

By selecting the appropriate CPU threading mode, you can improve the performance of applications and workloads that are running on a domain. You can select a threading mode that either maximizes throughput or maximizes the number of instructions per cycle, as follows:

- **Maximizing for throughput (`max-throughput`).** Workloads that benefit most from high throughput run a lot of software and perform a lot of I/O operations. When you optimize for maximum throughput, you enable CPU cores to concurrently run a maximum number of hardware threads. This mode is best for running heavily threaded workloads, such as those performed by web servers, database servers, and file servers. This mode is used by default and is also used on older SPARC T-series platforms, such as SPARC T3 platforms.
- **Maximizing for IPC (`max-ipc`).** Workloads that benefit most from high IPC are CPU intensive, such as systems that run intensive arithmetic computations. When you optimize for maximum IPC, you enable a CPU thread to execute more instructions per CPU cycle. This optimization is achieved by reducing the number of CPU threads that are concurrently active on the same CPU core.

Selecting the CPU Threading Mode

Select the CPU threading mode of a domain by using the `ldm add-domain` or `ldm set-domain` command to set the `threading` property.

```
ldm add-domain [mac-addr=num] [hostid=num]
  [failure-policy=ignore|panic|reset|stop]
  [extended-mapin-space=on] [master=master-ldom1,...,master-ldom4]
  [threading=max-throughput|max-ipc] ldom
```

```
ldm set-domain [mac-addr=num] [hostid=num]
  [failure-policy=ignore|panic|reset|stop]
  [extended-mapin-space=[on|off]]
```

```
[master=[master-ldom1, ..., master-ldom4]]
[threading=max-throughput|max-ipc] ldom
```

The `threading` property is used to dynamically change the threading mode by specifying one of the following values:

- **max-throughput.** Use this value to select the threading mode that maximizes throughput. This mode activates all threads that are assigned to the domain. This mode is used by default and is also selected if you do not specify any mode (`threading=`).
- **max-ipc.** Use this value to select the threading mode that maximizes the number of instructions per cycle (IPC). When you use this mode on the SPARC T4 platform, only one thread is active for each CPU core that is assigned to the domain. Selecting this mode requires that the domain is configured with the whole-core constraint.

Use the `ldm add-vcpu -c` or `ldm set-vcpu -c` command to configure the whole-core constraint. See the `ldm (1M)` man page.

Note that changing the threading mode dynamically activates or deactivates CPU threads. So, the number of virtual CPUs that are available in the domain also dynamically changes.

The `max-ipc` threading mode leverages the whole-core constraint, so you must abide by the whole-core constraint requirements and restrictions to do the following:

- Change the number of cores that are allocated to a domain.
- Enable or disable the whole-core constraint.

Thus, to dynamically change the threading mode of a running domain to `max-ipc` mode, you must configure the domain with the whole-core constraint.

For information about the restrictions, see “Threading Control Limitations” on page 7. For more information about the `add-domain` and `set-domain` subcommands, see the `ldm(1M)` man page.

Viewing the `threading` Property Value

You can use the following commands to view the `threading` property value:

- The `ldm list -o resmgt ldg1` command shows the constraints. The following example output shows that the `threading` property is set to `max-ipc`:

```
# ldm list -o resmgt ldg1
NAME
```

```
ldg1
CONSTRAINT
whole-core
max-cores=3
threading=max-ipc
```

- The `ldm list -o cpu ldg1` command shows the deactivated virtual CPUs by specifying a value of 0 in the UTIL column. The bold text in the following `max-ipc` example shows that only one thread is activated per CPU:

```
# ldm list -o cpu ldg1
NAME
ldg1
VCPU
VID PID CID UTIL STRAND
0 8 1 0.3% 100%
1 9 1 0 100%
2 10 1 0 100%
3 11 1 0 100%
4 12 1 0 100%
5 13 1 0 100%
6 14 1 0 100%
7 15 1 0 100%
8 24 2 0.4% 100%
...
```

- The `ldm list -l` command includes all the information about the specified domain. The bold text in the following example shows that the `threading` property is set to `max-ipc`:

```
# ldm list -l ldg1
...
VID PID CID UTIL STRAND
0 8 1 0.6% 100%
1 9 1 0 100%
2 10 1 0 100%
3 11 1 0 100%
4 12 1 0 100%
5 13 1 0 100%
```

```
6    14    1    0    100%  
...  
CONSTRAINT  
whole-core  
max-cores=3  
threading=max-ipc  
...
```

Threading Control Limitations

The threading controls feature has the following restrictions:

- The whole-core constraint restrictions apply. See “Allocating CPUs” in [Oracle VM Server for SPARC 2.1 Administration Guide](#).
- The `threading` property value does not persist across a domain migration.
- The `threading` property cannot be set to `max-ipc` while Power Management (PM) is enabled. When PM runs, all domains must have the `threading` property set to `max-throughput`.

Conclusion

The Oracle VM Server for SPARC 2.1 software enables you to use dynamic CPU threading controls to optimize workload performance on SPARC T4 systems. You can optimize CPU performance for CPU-bound workloads or for high throughput workloads. By default, the CPU is tuned for maximum throughput.

For more information about Oracle’s virtualization, visit www.oracle.com/virtualization.



Tuning the SPARC CPU to Optimize Workload
Performance on SPARC T4 Systems
September 2011
Author: Cathleen Reiher
Contributing Authors: Alexandre Chartre, Rory
Foster, Menno Lageman, Jorge Osario

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0611

Hardware and Software. Engineered to Work Together