

FairCom White Paper

# FairCom CPU Accounting, Definitions, and Procedures

**Audience**

**Developers**

**Subject**

**FairCom's high-performance NAV and SQL  
database technology.**



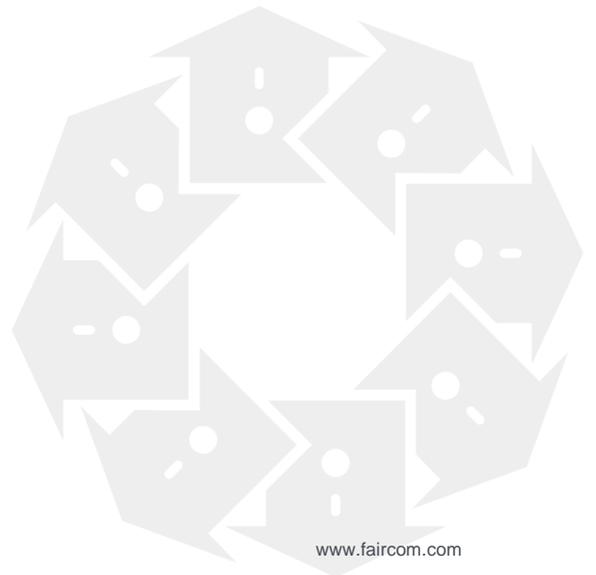
© Copyright 2021, FairCom Corporation. All rights reserved. For full information, see the FairCom Copyright Notice (page xi).



**FairCom**<sup>®</sup>

# Contents

<b>1.</b>	<b>CPU Counting .....</b>	<b>1</b>
1.1	CPU and Threads .....	1
	Virtualization.....	2
1.2	Performance .....	2
<b>2.</b>	<b>Counting CPUs &amp; Threads .....</b>	<b>4</b>
2.1	AIX.....	5
2.2	Linux.....	8
2.3	Solaris.....	9
2.4	Mac OS X .....	9
2.5	Windows .....	10
<b>3.</b>	<b>Document History .....</b>	<b>10</b>
<b>4.</b>	<b>Index .....</b>	<b>13</b>



# 1. CPU Counting

This white paper discusses the way FairCom DB handles actual physical processors, multi-core processors, multithreaded processors, and virtual machines.

Modern processor design allows multiple processor cores to be packaged in one integrated circuit chip. Each of those cores is essentially a complete CPU. Simultaneous multithreading (SMT) is a technique that allows a single processor core to execute multiple threads (Intel's hyper-threading is an example).

FairCom builds multi-threading into all FairCom DB Servers so they can take advantage of any available processors or threads. The architecture and performance (page 2) benefits are discussed later in this paper.

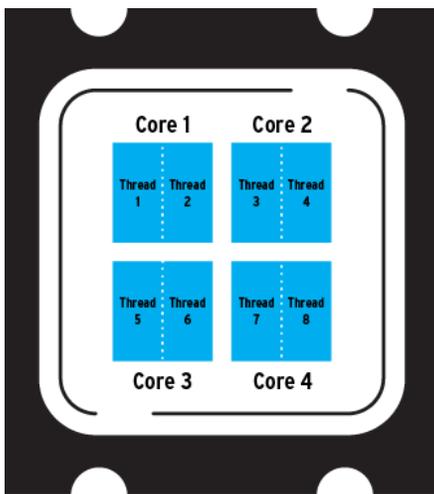
## Licensing

The section titled Counting CPUs & Threads (page 4) explains of how processors are counted for CPU-based pricing. Virtualization (page 2), running more than one instance of an operating system on a physical server, is discussed from a licensing perspective.

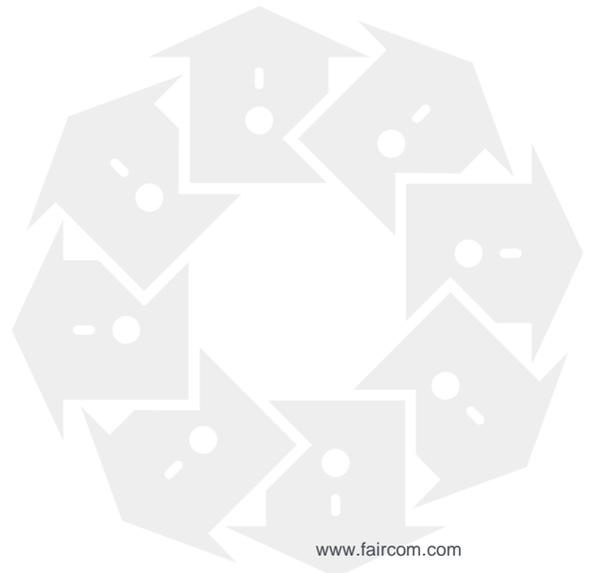
## 1.1 CPUs and Threads

With years of experience developing multi-threaded applications, FairCom has built multi-threading into all FairCom DB Servers since the late 1980s. This allows FairCom DB Server to take advantage of any processors or threads available to it.

Simultaneous multithreading (SMT) is a technique some processor manufacturers use so that a single processor can execute multiple threads. Intel's hyper-threading technology is an implementation of SMT that supports two threads on each processor core.



A single CPU chip (1 "socket") with 4 cores and 8 threads





In the limiting case, SMT would be as effective as multiple processors or CPUs; in actual practice the performance gains depend on the application. In particular, applications that are I/O bound, such as databases that must wait for disks, can benefit from SMT because one thread can execute while another is waiting for I/O.

Modern processor design allows multiple processor cores to be packaged in one integrated circuit chip. Each of those cores is essentially a complete CPU. For example, a single chip may contain four processor cores, so it is equivalent to four separate CPUs. In addition, each core may use SMT, causing it to appear to be two or more processors. The result is a single chip could contain four cores each of which appears as two processors, for a total of eight processors reported by the operating system. A large server could use more than one of these chips, resulting in a large number of reported processors.

The amount of performance improvement provided by these technologies depends on the specifics of your application.

## Virtualization

From a licensing perspective, virtualization is an issue related to CPU counting.

Virtualization is broadly defined as running more than one instance of an operating system on a physical server machine. This gives one physical server the ability to run multiple applications each inside of its own virtualized environment.

Each virtual environment that requires a FairCom DB Server must have its own server license. The FairCom DB Servers in each virtual environment can have unique user, connection, and CPU counts if desired.

## Reporting

On a machine running a VMware server or IBM's VIOS manager, the operating system of the virtual machine reports the number of processors assigned to the virtual machine, not the total number of processors on the physical machine. In turn, the FairCom DB server reports only the number of processors indicated by the operating system.

## 1.2 Performance

FairCom DB Server takes advantage of any processors available to it. The performance depends on the saturation point with your specific application. Specifically, performance depends on whether your application is bound by CPU, memory, or disk I/O.

We have seen situations where simultaneous multithreading (SMT) with two threads on a single CPU core provides close to the same performance as two CPU cores, and we have seen situations where it does not. The only way to definitively understand your performance is to run your application in both configurations and monitor the system throughput.

The FairCom DB statistics monitoring program, **ctstat**, provides a good way to monitor your application's throughput. You will need to run your application for a period of time in both manners to see which yields the best throughput. Documentation for **ctstat** can be found in the *FairCom DB Server Administrator's Guide* <http://docs.faircom.com/doc/ctserver/52856.htm>.



To use this utility, use the `-text` option as it dumps everything FairCom DB monitors except for the function timings. Note you can enable function timings using the `ctstat -wrktime on` support, but this will impact performance as the function timings are a costly exercise. The rest of the monitoring does not impact performance because the Server is already keeping these metrics.

The output from `ctstat -text` will go to a text file called `snapshot.fcs`. With SMT enabled, run this utility 5 or 6 times every 10 minutes over the course of an hour (or longer as you feel appropriate). Move the `snapshot.fcs` file from your first run to a new location. Then repeat the same exercise with the application running the same tests with SMT disabled.

Now that you have two `snapshot.fcs` files, you can view them with an editor and compare selected metrics. The `ctsnpr.exe` utility allows you to read a `snapshot.fcs` and dump it to a comma-delimited file you can open in Excel. The usage for `ctsnpr` is:

```
ctsnpr.exe snapshot.fcs > snapshot1.out
```

Be sure to update snapshot file parsing utility, `ctsnpr`, to a version that supports the correct `SNAPSHOT.FCS` file format.

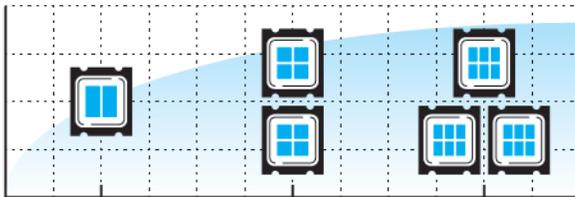
Repeat this for both `snapshot.fcs` files and open them in Excel to compare the results.

## 2. Counting CPUs & Threads

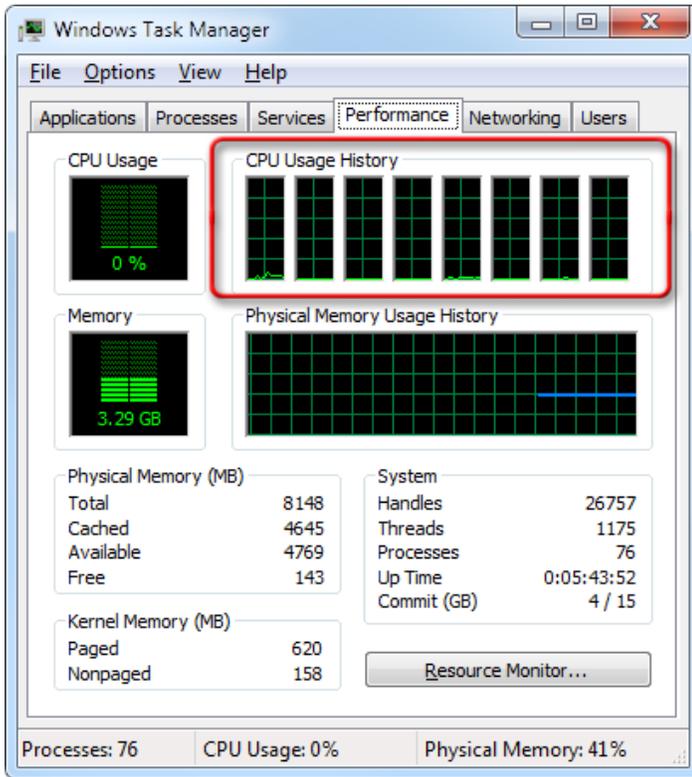
A question frequently arises about the definition of a CPU for purposes of counting.

When determining the number of CPUs in a host machine, **FairCom counts each CPU reported by the operating system** (be that a *physical CPU*, a *CPU core*, an *SMT thread* (page 1), or a *virtual machine* (page 2)).

Unlike many database vendors, we allow the end-user to bind the FairCom DB Server to a given number of CPUs in the machine. For example, a FairCom DB Server supporting 4 CPUs can be run on a machine with 8 CPUs by binding the FairCom DB Server to 4 of those CPUs. Further, the pricing for each additional CPU is an incremental price rather than a 100% increase as with most other database providers. This helps FairCom pricing to better reflect the benefits of multiple cores and threads.



In some cases, the system calls provided by the operating system do not distinguish between actual processors and SMT threads. This causes the operating system to report the number of CPUs as the number of processor cores times the number of SMT threads running on each core. For example, a system may have four processor cores each of which has two SMT threads causing a total of eight CPUs to be reported. For example, the Windows Task Manager shows eight CPUs on systems that have a single, quad-core, hyper-threaded chip:



## 2.1 AIX

An AIX system with two CPUs using the Workload Partitioning feature (WPAR), which supports two SMT threads per processor, will report four CPUs. The AIX system call that detects the number of CPUs counts the virtual processors created by the threads, which causes it to report four CPUs in this example. If the FairCom DB Server was licensed for two CPUs, the following message would be displayed:

```
*****
LICENSE NOTICE:
FairCom DB is licensed for 2 CPU's, but 4 CPU's have
been detected in the host machine. Either upgrade
the FairCom DB license to support a greater number
of CPU's or bind FairCom DB to specific CPU's.
*****
```

One of the following options can be used to correct this situation:

1. Change the binding to a single CPU.
2. Disable the SMT support.
3. Purchase a larger FairCom DB Server license that supports 4 cores.

AIX does not support the CPU\_AFFINITY keyword.



## Binding to Specific CPUs

### **execrset**

On AIX, the FairCom DB Server can be launched via the **execrset** command to operate with a given resource set.

#### **Usage:**

```
execrset [-P][-F] rsetname [-e] command [args]
```

```
execrset [-P][-F] -c cpuNum [cpuNum] [-m memNum [memNum]] -e command [args]
```

For example, **execrset -c 0 -e ctreesql** will launch the server and bind it to processor 0.

To use the **execrset** command, the user must have root permissions or have the CAP\_NUMA\_ATTACH and CAP\_PROPAGATE capabilities assigned. To assign these capabilities to a user, the **chuser** command can be used:

```
chuser capabilities=CAP_NUMA_ATTACHED,CAP_PROPAGATE user
```

To check if the capabilities are assigned to a user, you can use the **lsuser** command:

```
lsuser -a capabilities user
```

### **bindprocessor**

If the FairCom DB Server is already running, the **bindprocessor** command can be used. This command may also be used to query the list of available processors.

### **taskset**

A process can be bound to a given set of CPUs by using the **taskset** command (see the *Linux section* (page 8)). This utility is installed as part of the *util-linux* package.



## An AIX Program

The simple program listed below compiles on AIX 5.3 and later showing one way to programmatically determine the CPU count, whether or not the partition supports SMT, and whether or not it is on. This program could easily be incorporated into the server startup logic.

```
/* sample program cpustats.c */
#include <stdio.h>
#include <libperfstat.h>

int main(int argc, char* argv[])
{
    perfstat_cpu_total_t cpustats;
    perfstat_partition_total_t partstats;

    if (!perfstat_cpu_total(NULL, &cpustats, sizeof(perfstat_cpu_total_t), 1)) {
        perror("perfstat_cpu_total");
        exit(-1);
    }

    if (!perfstat_partition_total(NULL, &partstats, sizeof(perfstat_partition_total_t), 1)) {
        perror("perfstat_partition_total");
        exit(-1);
    }

    printf("\nCurrent number of active CPUs: %d", cpustats.ncpus);

    if (partstats.type.b.smt_capable)
        printf("\nOS supports SMT mode");
    else
        printf("\nOS DOES NOT support SMT mode");

    if (partstats.type.b.smt_enabled)
        printf("\nSMT mode is ON\n\n");
    else
        printf("\nSMT mode is OFF\n\n");

    return(0);
}
```

To compile the program, use the command line:

```
- cc cpustats.c -o cpustats -lperfsta
```

## FairCom DB Server physical CPU counting method on AIX

**Note:** This is a Compatibility Change.



The physical CPU counting method used by FairCom DB Server on AIX has been changed to account for LPAR CPU assignment. The physical CPU count returns the total number of CPUs on the system. However, on AIX, an LPAR can be defined that has access to a subset of the CPUs. To accommodate this, we now base our count on the "maximum virtual CPUs" assigned to the LPAR.

## 2.2 Linux

On Linux, a process can be bound to a given set of CPUs by using the **taskset** command shown below. The CPU affinity is represented as a bitmask with the lowest order bit corresponding to the first logical CPU and the highest order bit corresponding to the last logical CPU. For example, 0x00000001 is processor #0 (first processor), 0x00000003 is processors #0 and #1, and 0x00000004 is processor #2 (third processor). You can use the **-p** option to set the affinity for an already running process, or you can start a process with affinity by specifying the mask and the command to follow as shown in the first example below.

### Usage:

```
taskset [options] [mask | cpu-list] [pid | cmd [args...]]
```

Set or get the affinity of a process:

- **-p, --pid** - Operate on existing given pid
- **-c, --cpu-list** - Display and specify CPUs in list format
- **-h, --help** - Display this help
- **-v, --version** - Output version information

The default behavior is to run a new command:

```
taskset 03 sshd -b 1024
```

You can retrieve the mask of an existing task:

```
taskset -p 700
```

Or set it:

```
taskset -p 03 700
```

List format uses a comma-separated list of CPUs instead of a mask:

```
taskset -pc 0,3,7-11 700
```

Ranges in list format can take a stride argument, for example:

```
0-31:2
```

is equivalent to mask:

```
0x55555555
```



## 2.3 Solaris

On Solaris systems, `CPU_AFFINITY` accepts a single numeric value, which is interpreted as a processor set number. For example, `CPU_AFFINITY 2` configures FairCom DB to run on processor set 2.

**Note:** To use `CPU_AFFINITY` under Solaris you need to run the FairCom DB Server with root permission. This because Solaris requires any process using a processor set to have such permission.

To create a processor set on Solaris, use the Solaris command `psrset`. For example, to create a set comprising processors 4 through 7, use:

```
psrset -c 4-7
```

where:

- `-c` - Create processor set.
- `4-7` - The processor numbers included in the set.

The ID of the newly created processor set is returned:

```
created processor set ps_id
```

To bind a process to this processor set, use:

```
psrset -b ps_id pid
```

where:

- `-b` - Bind.
- `ps_id` - The ID returned by the command when the processor set was created.
- `pid` - The ID of the process to be bound to the processor set.

If the process does not have permission to assign itself to the specified processor set (or if an invalid processor set is specified), FairCom DB logs the following message to `CTSTATUS.FCS`, where `<configuration_file_name>` is the name of the FairCom DB configuration file, `<line_number>` is the line number on which the `CPU_AFFINITY` option was specified, and `<error_code>` is the system error code returned by the OS.

```
Configuration error: <configuration_file_name>, line <line_number>: Failed to set CPU affinity: system error code <error_code>.
```

## 2.4 Mac OS X

On Mac OS X, there is no check by FairCom DB for the processor count. There is no system command to set the CPU affinity for this platform.



## 2.5 Windows

On Windows systems, CPU\_AFFINITY server keyword can be used to set the processor affinity mask for the FairCom DB process. The option accepts a comma-delimited list of processor numbers. For example: CPU\_AFFINITY 0,1,2,3,8,9,10,11 indicates that FairCom DB is to be run on the eight specified CPUs.

If FairCom DB successfully sets the CPU affinity to the specified CPUs, the following message is logged to CTSTATUS.FCS, where <cpulist> is the list of CPUs specified for the CPU\_AFFINITY option:

```
Successfully set CPU affinity to: <cpulist>
```

The following error situations can occur when using the CPU\_AFFINITY option:

If the list of CPUs specifies a CPU number that is out of range on the system, a message is logged to CTSTATUS.FCS:

```
Configuration error: <config_file_name>, line <line_number>: The CPU_AFFINITY option specifies an invalid CPU number for this system.
```

## 3. Document History

04/26/2019	Updated AIX instructions.
02/12/2015	Added AIX program.
02/11/2015	Removed an unnecessary reference in the Linux section.
11/11/2014	Created separate sections for <i>Solaris</i> (page 9) and <i>Windows</i> (page 10) and added information specific to each.
06/20/2014	Document published.

Last published Thursday, April 15, 2021.

# Copyright Notice

Copyright © 1992-2021 FairCom USA Corporation. All rights reserved.

No part of this publication may be stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise without the prior written permission of FairCom USA Corporation. Printed in the United States of America.

Information in this document is subject to change without notice.

## Trademarks

FairCom DB, FairCom EDGE, c-treeRTG, c-treeACE, c-treeAMS, c-treeEDGE, c-tree Plus, c-tree, r-tree, FairCom, and FairCom's circular disc logo are trademarks of FairCom USA, registered in the United States and other countries.

The following are third-party trademarks: Btrieve is a registered trademark of Actian Corporation. Amazon Web Services, the "Powered by AWS" logo, and AWS are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries. AMD and AMD Opteron are trademarks of Advanced Micro Devices, Inc. Macintosh, Mac, Mac OS, and Xcode are trademarks of Apple Inc., registered in the U.S. and other countries. Embarcadero, the Embarcadero Technologies logos and all other Embarcadero Technologies product or service names are trademarks, service marks, and/or registered trademarks of Embarcadero Technologies, Inc. and are protected by the laws of the United States and other countries. HP and HP-UX are registered trademarks of the Hewlett-Packard Company. AIX, IBM, POWER6, POWER7, POWER8, POWER9, POWER10 and pSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. Intel, Intel Core, Itanium, Pentium and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. ACUCOBOL-GT, Micro Focus, RM/COBOL, and Visual COBOL are trademarks or registered trademarks of Micro Focus (IP) Limited or its subsidiaries in the United Kingdom, United States and other countries. Microsoft, the .NET logo, the Windows logo, Access, Excel, SQL Server, Visual Basic, Visual C++, Visual C#, Visual Studio, Windows, Windows Server, and Windows Vista are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Oracle and Java are registered trademarks of Oracle and/or its affiliates. QNX and Neutrino are registered trademarks of QNX Software Systems Ltd. in certain jurisdictions. CentOS, Red Hat, and the Shadow Man logo are registered trademarks of Red Hat, Inc. in the United States and other countries, used with permission. SAP® Business Objects, SAP® Crystal Reports and SAP® BusinessObjects™ Web Intelligence® as well as their respective logos are trademarks or registered trademarks of SAP. SUSE" and the SUSE logo are trademarks of SUSE LLC or its subsidiaries or affiliates. UNIX and UNIXWARE are registered trademarks of The Open Group in the United States and other countries. Linux is a trademark of Linus Torvalds in the United States, other countries, or both. Python and PyCon are trademarks or registered trademarks of the Python Software Foundation. isCOBOL and Veryant are trademarks or registered trademarks of Veryant in the United States and other countries. OpenServer is a trademark or registered trademark of Xinuos, Inc. in the U.S.A. and other countries. Unicode and the Unicode Logo are registered trademarks of Unicode, Inc. in the United States and other countries.

All other trademarks, trade names, company names, product names, and registered trademarks are the property of their respective holders.

Portions Copyright © 1991-2016 Unicode, Inc. All rights reserved.

Portions Copyright © 1998-2016 The OpenSSL Project. All rights reserved. This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>).

Portions Copyright © 1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved. This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Portions © 1987-2020 Dharma Systems, Inc. All rights reserved.

This software or web site utilizes or contains material that is © 1994-2007 DUNDAS DATA VISUALIZATION, INC. and its licensors, all rights reserved.

Portions Copyright © 1995-2013 Jean-loup Gailly and Mark Adler.

Portions Copyright © 2009-2012 Eric Haszlkiewicz.

Portions Copyright © 2004, 2005 Metaparadigm Pte Ltd.

Portions Copyright © 2008-2020, Hazelcast, Inc. All Rights Reserved.

Portions Copyright © 2013, 2014 EclipseSource.

Portions Copyright © 1999-2003 The OpenLDAP Foundation.

## Open Source Components

Like most software development companies, FairCom uses third-party components to provide some functionality within our technology. Often those third-party components are selected because they are a standard in the industry, they offer specific functionality that is easier to license than to develop and maintain in the long run, or they provide a proven and inexpensive solution to a particular business need. Examples of third-party software FairCom uses are the OpenSSL toolkit that provides Transport Layer Security (TLS) for secure communications and the ICU Unicode libraries to provide wide character support (think international characters and emojis).

Some of these third-party components are the subject to commercial licenses and others are subject to open source licenses. For open source solutions that we incorporate into our technology, we include the package name and associated license in a notice.txt file found in the same directory as the server.

The notice.txt file should always stay in the same directory as the server. This is particularly important in instances where your company has redistribution rights, such as an ISV who duplicates server binaries and (re)distributes those to an eventual end-user at a third-party company. Ensuring that the notice.txt file "travels with" the server binary is important to maintain third-party and FairCom license compliance.

4/15/2021

# 4. Index

**A**  
 AIX .....5

**B**  
 bind (Solaris).....9  
 Binding to Specific CPUs (AIX) .....5  
 bindprocessor .....5

**C**  
 CAP\_NUMA\_ATTACH and CAP\_PROPAGATE .....5  
 chuser .....5  
 Copyright Notice ..... xi  
 Counting CPUs & Threads .....4  
 CPU .....1, 2, 4  
 CPU Counting.....1  
 CPU\_AFFINITY .....9, 10  
 CPUs and Threads .....1  
 cpustats.c (AIX) .....5  
 ctsnpr.exe .....2  
 ctstat .....2  
   -text option .....2  
   -wrktme on option .....2

**D**  
 Document History .....10

**E**  
 execrset .....5

**H**  
 hyper-threading.....1

**I**  
 I/O .....1  
 IBM VIOS manager.....2  
 Intel hyper-threading.....1

**L**  
 Linux .....8  
 lsuser .....5

**M**  
 Mac OS X .....9

**O**  
 operating system .....4

**P**  
 Performance .....2  
 performance, measuring.....2  
 processor set .....9  
 psrset (Solaris).....9

**S**  
 server license.....2

simultaneous multithreading .....1  
 SMT ..... 1, 2, 4  
 snapshot.fcs .....2  
 Solaris .....9

**T**  
 taskset..... 5, 8  
 threads ..... 1, 4

**V**  
 VIOS manager .....2  
 Virtualization .....2  
 VMware .....2

**W**  
 Windows ..... 10  
 Workload Partitioning feature (WPAR) .....5

