
# IBM Tivoli® Monitoring

## *Resource Model Reference*

*Version 5.1*

SH19-4570-00


# IBM Tivoli® Monitoring

## *Resource Model Reference*

*Version 5.1*

SH19-4570-00

**Tivoli Monitoring Resource Model Reference, Version 5.1**

### Copyright Notice



### Trademarks



### Notices

References in this publication to Tivoli Systems or IBM products, programs, or services do not imply that they will be available in all countries in which Tivoli Systems or IBM operates. Any reference to these products, programs, or services is not intended to imply that only Tivoli Systems or IBM products, programs, or services can be used. Subject to valid intellectual property or other legally protectable right of Tivoli Systems or IBM, any functionally equivalent product, program, or service can be used instead of the referenced product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by Tivoli Systems or IBM, are the responsibility of the user. Tivoli Systems or IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, New York 10504-1785, U.S.A.

### ISO 9001 Certification

This product was developed using an ISO 9001 certified quality system.

Certification has been awarded by Bureau Veritas Quality International (BVQI) (Certification No. BVQI - 92086 / A).

BVQI is a world leader in quality certification and is currently recognized by more than 20 accreditation bodies.

# Contents

# Preface

IBM Tivoli Monitoring enables you to monitor availability and performance status of *resources* on your systems to identify bottlenecks and potential resource problems. This reference provides key information about the resource models included in this product.

IBM Tivoli Monitoring was formerly known as Tivoli Distributed Monitoring (Advanced Edition).

## Who Should Read This Reference

The target audiences for this reference are system administrators and system installers.

To make effective use of the product you require knowledge and practical experience of the following:

- Importing, installing, and managing the Tivoli Management Framework® and the Tivoli management environment
- System administration on the systems where you plan to install Tivoli Monitoring components, and which you plan to monitor using the product

You should also be familiar with the Tivoli Enterprise Console® product.

## What is New in this Release

IBM Tivoli Monitoring Version 5.1 is a development of, and upgrades, Tivoli Distributed Monitoring (Advanced Edition), Version 4.1 and provides the following additional functionality:

**Mdist2 Support**
The product now uses Mdist2 to distribute profiles to the endpoints.

**Web-based Health Console**
The former Health Console has become a web-based Health Console .

**Tools and guidelines for migrating from Tivoli Distributed Monitoring (Classic Edition)**
A set of guidelines is available for the migration, and a compatibility mode is now provided to enable Tivoli Monitoring 5.1 users to use Tivoli Distributed Monitoring (Classic Edition) monitors inside a Tivoli Monitoring 5.1 resource model.

**New response actions**
The product is now able to act on an event by sending an e-mail to a specified address, sending a notice to a notice group, running a program.

**Tools and guidelines for migrating from Tivoli Web Component Manager**
A tool and a set of resource models are now available to enable Tivoli Web Component Manager users to migrate to Tivoli Monitoring.

**Serviceability Features**
Several new serviceability features have been added to the product.

# What This Reference Contains

This reference contains the following sections:

**Part I. Resource Models for Windows®**

- "Chapter 1, "Event Log Resource Model""

  Describes the Event Log resource model. It includes the indications, thresholds, logging information, built-in actions, and default settings for the resource model.

- "Chapter 2, "Logical Disk Resource Model""

  Describes the Logical Disk resource model. It includes the indications, thresholds, logging information, and default settings for the resource model.

- "Chapter 3, "Memory Resource Model""

  Describes the Memory resource model. It includes the indications, thresholds, logging information, and default settings for the resource model.

- "Chapter 4, "Network Interface Card Resource Model""

  Describes the Network Interface Card resource model, which runs on Windows NT®. It includes the indications, thresholds, logging information, built-in actions, and default settings for the resource model.

- "Chapter 5, "Parametric Event Log Resource Model""

  Describes the Parametric Event Log resource model. It includes the indications, thresholds, logging information, parameters, and default settings for the resource model.

- "Chapter 6, "Parametric Services Resource Model""

  Describes the Parametric Services resource model. It includes the indications, thresholds, logging information, parameters, and default settings for the resource model.

- "Chapter 7, "Parametric TCP/IP Ports Resource Model""

  Describes the Parametric TCP/IP resource model. It includes the indications, thresholds, logging information, parameters, and default settings for the resource model.

- "Chapter 8, "Physical Disk Resource Model""

  Describes the Physical Disk resource model. It includes the indications, thresholds, logging information, and default settings for the resource model.

- "Chapter 9, "Printer Resource Model""

  Describes the Printer resource model. It includes the indications, thresholds, logging information, and default settings for the resource model.

- "Chapter 10, "Process Resource Model""

  Describes the Process resource model. It includes the indications, thresholds, logging information, and default settings for the resource model.

- "Chapter 11, "Processor Resource Model""

  Describes the Processor resource model. It includes the indications, thresholds, logging information, and default settings for the resource model.

- "Chapter 12, "Server Performance Prediction Resource Model""

  Describes the Server Performance Prediction resource model. It includes the parameters, logging information, and default settings for the resource model.

- "Chapter 13, "Services Resource Model""

Describes the Services resource model. It includes the indications, thresholds, logging information, built-in actions, and default settings for the resource model.

- "Chapter 14, "TCP/IP Resource Model""

  Describes the TCP/IP resource model. It includes the indications, thresholds, logging information, and default settings for the resource model.

**Part II. Resource Models for UNIX®/Linux**

- "Chapter 15, "CPU Resource Model""

  Describes the CPU resource model. It includes the indications, thresholds, logging information, and default settings for the resource model.

- "Chapter 16, "File Resource Model""

  Describes the File resource model. It includes the indications, thresholds, logging information, parameters, and default settings for the resource model.

- "Chapter 17, "File System Resource Model""

  Describes the File System resource model. It includes the indications, thresholds, logging information, and default settings for the resource model.

- "Chapter 18, "Memory Resource Model""

  Describes the Memory resource model. It includes the indications, thresholds, logging information, and default settings for the resource model.

- "Chapter 19, "Network Interface Resource Model""

  Describes the Network Interface resource model. It includes the indications, thresholds, logging information, and default settings for the resource model.

- "Chapter 20, "Network RPC-NFS Resource Model""

  Describes the Network RPC-NFS resource model, which is available only for Solaris platforms. It includes the indications, thresholds, logging information, and default settings for the resource model.

- Chapter 21, "Chapter 21, "Process Resource Model""

  Describes the Process resource model. It includes the indications, thresholds, parameters, logging information, and default settings for the resource model.

- "Chapter 22, "Security Resource Model""

  Describes the Security resource model. It includes the indications, thresholds, parameters, logging information, and default settings for the resource model.

- "Chapter 23, "Server Performance Prediction Resource Model""

  Describes the Server Performance Prediction resource model. It includes the parameters, logging information, and default settings for the resource model.

**Part III. Appendixes**

- "Appendix A, "Return Codes for Built-in Actions (Windows)""

  Lists return codes for the Windows resource models that have built-in actions.

- "Appendix B, "Correlated Events (Windows)""

  Describes correlated events and the indications from which each one is generated.

# Publications

This section lists publications in the *IBM Tivoli Monitoring* library and any other related documents. It also describes how to access Tivoli publications online, how to order Tivoli publications, and how to make comments on Tivoli publications.

## IBM Tivoli Monitoring Library

The following documents are available in the *IBM Tivoli Monitoring* library:

- *IBM Tivoli Monitoring: Resource Model Reference*, SH19-4570

  Provides information about using and customizing the resource models that can be used with IBM Tivoli Monitoring.

- *IBM Tivoli Monitoring: User's Guide*, SH19-4569

  Describes how to install, customize, and use IBM Tivoli Monitoring to manage system and application resources.

- *IBM Tivoli Monitoring: Workbench User's Guide*, SH19-4571

  Describes how to use the IBM Tivoli Monitoring Workbench to create new resource models or to modify existing ones.

- *IBM Tivoli Monitoring: Rapid Deployment Roadmap*, GI11-0398

  Describes how to use the Rapid Deployment tool to make a fresh installation of IBM Tivoli Monitoring.

- *IBM Tivoli Monitoring: Release Notes*, GI10-5793

  Provides the most current information about IBM Tivoli Monitoring.

Versions of these documents in PDF and HTML formats can be found on the IBM Tivoli Monitoring and IBM Tivoli Monitoring Workbench product CDs. They are stored in the `Books` directory, and can be accessed by selecting the file `Books/infocenter.html` with your Web browser. This displays an HTML page from which all of the documents can be accessed in either format.

## Prerequisite and Related Publications

Before beginning the installation, read the following prerequisite and related documentation for more information about the management options that IBM Tivoli Monitoring provides:

- *Tivoli Management Framework: User's Guide*

  Provide prerequisite information about setting up the Tivoli Management Framework and Tivoli environment.

- *Tivoli Management Framework: Planning and Installation*

  Provide prerequisite information about setting up the Tivoli Management Framework and Tivoli environment.

- *Tivoli Management Framework: Reference Manual*

  Provide prerequisite information about setting up the Tivoli Management Framework and Tivoli environment.

- *Tivoli Software Installation Service: User's Guide*

  Provides task-oriented information on how to import, manage, and install Tivoli Management Environment® software on selected machines and managed nodes within your Tivoli management region.

- *Tivoli Software Installation Service: Release Notes*

Provides important information about using and installing the Tivoli Software Installation Service (SIS).

- *Tivoli Enterprise Console: Rule Builder's Guide*

  Provides information about using the Tivoli Enterprise Console rule editor and graphical rule builder to modify existing rules and create new rules to match your specific event management needs.

- *Tivoli Enterprise Console: Release Notes*, Version 3.7 or later

  Provides the most current information about Tivoli Enterprise Console.

- *Tivoli Business Systems Management: Distributed User's Guide*, Version 1.5 (if you intend to use Tivoli Business Systems Manager to monitor IBM Tivoli Monitoring events)

  Describes how to use the Tivoli Business Systems Manager product.

- *Tivoli Business Systems Management: Console User's Guide*, Version 1.5 (if you intend to use Tivoli Business Systems Manager to monitor IBM Tivoli Monitoring events)

  Describes how to use the Tivoli Business Systems Manager console.

## Accessing Publications Online

You can access many Tivoli publications online at the Tivoli Customer Support Web site: http://www.tivoli.com/support/documents/

These publications are available in PDF or HTML format, or both. Translated documents are also available for some products.

## Ordering Publications

You can order many Tivoli publications online at the following Web site: http://www.elink.ibmlink.ibm.com/public/applications/publications/cgibin/pbi.cgi

You can also order by telephone by calling one of these numbers:

- In the United States: 800-879-2755
- In Canada: 800-426-4968
- In other countries, for a list of telephone numbers, see the following Web site: http://www.tivoli.com/inside/store/lit_order.html

## Providing Feedback about Publications

We are very interested in hearing about your experience with Tivoli products and documentation, and we welcome your suggestions for improvements. If you have comments or suggestions about our products and documentation, contact us in one of the following ways:

- Send an e-mail to pubs@tivoli.com.
- Complete our customer feedback survey at the following Web site: http://www.tivoli.com/support/survey/

# Contacting Customer Support

If you have a problem with any Tivoli product, you can contact Tivoli Customer Support. See the *Tivoli Customer Support Handbook* at the following Web site: http://www.tivoli.com/support/handbook

The handbook provides information about how to contact Tivoli Customer Support, depending on the severity of your problem, and the following information:

- Registration and eligibility
- Telephone numbers and e-mail addresses, depending on the country you are in
- What information you should gather before contacting support

# Conventions Used in This Guide

This book uses several conventions for special terms and actions, operating system-dependent commands and paths, and margin graphics.

## Typeface Conventions

The following typeface conventions are used in this book:

**Bold** Lowercase and mixed-case commands, command options, and flags that appear within text appear like **this,** in **bold** type.

Graphical user interface elements (except for titles of windows and dialogs) and names of keys also appear like **this,** in **bold** type.

*Italic* Variables, values you must provide, new terms, and words and phrases that are emphasized appear like *this,* in *italic* type.

**`Monospace`** Commands, command options, and flags that appear on a separate line, code examples, output, and message text appear like `this`, in `monospace` type.

Names of files and directories, text strings you must type, when they appear within text and HTML and XML tags also appear like `this`, in `monospace` type.

## Operating System-dependent Variables and Paths

This book uses the UNIX convention for specifying environment variables and for directory notation.

When using the Windows command line, replace **$***variable* with ***%variable%*** for environment variables and replace each forward slash (**/**) with a backslash (\) in directory paths.

**Note:** If you are using the bash shell on a Windows system, you can use the UNIX conventions.

# Part I. Resource Models for Windows

# 1

# Event Log Resource Model

This chapter describes the Event Log resource model for Windows systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | TMW_EventLog |
| Category | Windows |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | Yes |
| Clearing events | No |
| Default cycle time | 600 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to Windows NT and Window 2000 endpoints.

The Event Log resource model examines the Windows System Event Log for events that normally require immediate attention, or can be rectified with a simple registry edit. The Event Log resource model highlights the following problems:

- Client connectivity problems

  Busy networks can cause clients to time out. This indication can be trapped. When it is, the registry value that determines the time out period can be enlarged. This solution must only be tried a few times. After which, the problem is more likely to be caused by a problem with the network card or drivers.

- Detection of malfunctioning devices

  Events 9, 11, and 15 are the most common events that occur when a device or driver is misbehaving. Trapping these errors helps to correct the devices as soon as possible.

- Server connectivity problems

  Bad IRPStackSize and OEM installs of Windows are common situations that prevent the server service from operating correctly. If the server service is not functioning properly, clients cannot connect to the affected machine.

# Indications and Events

The following table lists the events that can be generated by the Event Log resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| TMW_EventID9 | Event ID 9 | Warning | 8 |
| TMW_EventID11 | Event ID 11 | Warning | 8 |
| TMWEvent_ID15 | Event ID 15 | Warning | 9 |
| TMW_EventID2011 | Event ID 2011 | Warning | 9 |
| TMW_EventID2511 | Event ID 2511 | Warning | 9 |
| TMW_EventID3013 | Event ID 3013 | Warning | 10 |
| TMW_EventID7023 | Event ID 7023 | Warning | 10 |

## Event ID 9

This indication is sent when a device or device driver times out. See “Indication Properties” on page 11 for information on the properties for this event.

No particular threshold is exceeded to generate this indication. If an event in the Event Log is within the first set of events (determined by the threshold Maximum Number Of Logs) and has the event ID of 9, this indication is sent.

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Event ID 11

This indication is sent when an error is detected on a device or device controller. See “Indication Properties” on page 11 for information on the properties for this event.

No particular threshold is exceeded to generate this indication. If an event in the Event Log is within the first set of events (determined by Maximum Number Of Logs) and has the event ID of 11, this indication is sent.

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Event ID 15

This indication is sent when a drive (typically removable media, for example, CD-ROM) is not ready for access or is not functioning properly. See “Indication Properties” on page 11 for information on the properties for this event.

No particular threshold is exceeded to generate this indication. If an event in the Event Log is within the first set of events (determined by Maximum Number of Logs) and has the event ID of 15, this indication is sent.

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Event ID 2011

This indications is sent when the IRPStackSize is too small. Event 7023 will typically occur with this event. See “Indication Properties” on page 11 for information on the properties for this event.

No particular threshold is exceeded to generate this indication. If an event in the Event Log is within the first set of events (determined by Maximum Number of Logs) and has the event ID of 2011, this indication is sent.

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Event ID 2511

This indication is sent when a shared directory no longer exists. See “Indication Properties” on page 11 for information on the properties for this event.

No particular threshold is exceeded to generate this indication. If an event in the Event Log is within the first set of events (determined by Maximum Number of Logs) and has the event ID of 2511, this indication is sent.

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Event ID 3013

This indication is sent when the redirector is timing out. See "Indication Properties" on page 11 for information on the properties for this event.

No particular threshold is exceeded to generate this indication. If an event in the Event Log is within the first set of events (determined by Maximum Number of Logs) with the event ID of 3013, this indication is sent.

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Event ID 7023

This indication is sent when there is not enough server storage space. When it occurs by itself, the problem usually has to do with OEM installs of either Windows or Service Packs. Microsoft®'s Knowledge Base article Q151427 details this event. The solution is usually to reinstall of the latest service pack. See "Indication Properties" on page 11 for information on the properties for this event.

No particular threshold is exceeded to generate this indication. If an event in the Event Log is within the first set of events (determined by Maximum Number of Logs) with the event ID of 7023, this indication is sent.

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Indication Properties

All the indications in the Event Log resource model have the following properties.

**EvtID** Identifies the event by number (key attribute)

**RepeatCount**
Number of Event ID events received in a cycle

**Logfile**
The name of the log file in which the event was generated (for example, system, security, or application)

**SourceName**
The source of the event (key attribute)

**TimeGenerated**
The time that the event was generated

**TimeWritten**
The time that the event was written to the Event Log

**ComputerName**
The name of the computer where the event occurred (key attribute)

**Message**
The message text of the event

# Thresholds

The following table lists the thresholds that can be set for the Event Log resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| Maximum Number of Logs (MostNumberOfLogs) | This threshold determines the maximum number of newest events in the event log that are examined. | 100 |

# Built-in Actions

This resource model has the following built-in actions:

| Indication | Built-in Actions |
|---|---|
| Event ID 2011 | 1. Notifies the system administrator that the IRPStackSize is too small.<br>2. Edits the registry key Computer Name\HKEY_LOCAL_MACHINE\SYSTEMS\CurrentControlSet\ Services\LanmanServer\Parameters\IRPStackSize and<br>3. Sets the value to 0xCh (Decimal value 12).<br>The system must be rebooted for this change to take effect. |

| Indication | Built-in Actions |
|---|---|
| Event ID 2511 | 1. Notifies the system administrator that a shared directory no longer exists.<br>2. Edits the registry key Computer Name\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\LanmanServer\Shares and the system must be rebooted for this change to take effect.<br>3. Deletes the key that represents the sharename. |
| Event ID 3013 | 1. Notifies the system administrator that the redirector is timing out.<br>2. Edits the registry key Computer Name\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters\SessTimeOut to a higher value. When the value is not present, the default value is 0x2Dh (decimal value 45) seconds. Corrective actions consist of raising this value by an increment of 0xAh (decimal value 10).<br>The system must be rebooted for this change to take effect. After three attempts, user intervention is required. The system administrator should examine the network subsystem and possibly reinstall or repair network drivers or devices. |

# 2

# Logical Disk Resource Model



This chapter describes the Logical Disk resource model for Windows systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | TMW_LogicalDisk |
| Category | Windows |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 120 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to Windows endpoints.

The Logical Disk resource model highlights the following problems:

- Bytes Transferred Per Second

  This indicates the raw amount of data that is transferred through the disk. The maximum bytes per second that a disk subsystem can handle depends on the type of hard drive, the type of controller, the drive configuration, and, in the case of most IDE-based hard drives, the speed of the CPU.

- Disk Space

  Low disk space can often make it impossible to reboot a Windows NT machine. The threshold value is based on the percentage of free space on the disk. Administrators monitoring systems with disks of less than one gigabyte should raise this value. A good guideline is to keep enough free disk space to allow the pagefile to reach its maximum size limit.

- Percentage Usage

  The percentage of used disk often needs to be altered when monitoring file servers for bottlenecks. File servers, print servers, and mail servers use a high percent of disk space.

  However, workstations typically use a high percentage of disk space in short bursts. Performance in CPU speed and general system response decline.

# Prerequisites

The Logical Disk resource model requires disk counters to be enabled on the Windows endpoints.

To enable disk counters, perform the following steps on each endpoint:

1. On Windows NT systems, issue the **diskperf -y** command.
   On Windows 2000 systems, issue the **diskperf -yv** command.
2. Reboot the system.

# Indications and Events

The following table lists the events that can be generated by the Logical Disk resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| TMW_HighLogicalDiskReadBytesSec | High Read Bytes per Second | Minor | 15 |
| TMW_HighLogicalDiskWriteBytesSec | High Write Bytes per Second | Minor | 16 |
| TMW_HighLogicalDiskXferRate | High Transfer Rate | Minor | 17 |
| TMW_HighLogicalPercentDiskTime | High Percentage Disk Time | Warning | 18 |
| TMW_LogicalPossibleFrag | Logical Disk Possible Fragmentation | Minor | 19 |
| TMW_LowLogicalDiskSpace | Low Disk Space | Warning | 20 |
| TMW_SlowLogicalDrive | Slow Logical Drive | Warning | 21 |

## High Read Bytes per Second

This indication is sent when a logical disk is reading a lot of data.

The indication has the following attributes:

**PercentDiskRead**
The percentage of the logical drive being used to read data

**DiskReadBytesSec**
The number of bytes read per second on the logical disk

**DiskReadSec**
The number of transactions that are read per second on the logical disk

**PhysicalDisk**
The physical disk on which the logical disk resides (Key attribute)

**LogicalDisk**
The logical disk that is being analyzed (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Bytes per Second (HighBytesSec) | The maximum amount of bytes transferred per second (being written or read) exceeds this threshold. | 1572864 |
| High Percent Usage (HighPercentUsage) | The time that the selected disk drive spends for read or write requests exceeds this threshold. | 90 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 2 |

## High Write Bytes per Second

This indication is sent when a logical disk is writing a lot of data.

The indication has the following attributes:

**DiskWriteBytesSec**
The number of bytes written per second on the logical disk

**DiskWriteSec**
The number of transactions that are written per second on the logical disk

**PercentDiskWrite**
The percent of the logical drive used during write operations.

**PhysicalDisk**
The physical disk on which the logical disk resides (Key attribute)

**LogicalDisk**
The logical disk that is being analyzed (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Bytes per Second (HighBytesSec) | The amount of bytes transferred per second (being written or read) exceeds this threshold. | 1572864 |
| High Percent Usage (HighPercentUsage) | The percent of time that the selected disk drive spends for read or write requests exceeds this threshold. | 90 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 2 |

## High Transfer Rate

This indication is sent when a logical disk is used heavily to transfer data, both for reading and writing.

The indication has the following attributes:

**DiskXfersSec**
The rate of bytes read or written per second on the logical disk

**DiskReadsSec**
The rate of bytes read per second on the logical disk

**DiskWritesSec**
The rate of bytes written per second on the logical disk

**PercentDiskReadTime**
The percentage of the logical drive used for read operations

**PercentDiskWriteTime**
The percentage of the logical drive used for write operations

**PhysicalDisk**
The physical disk on which the logical disk resides (Key attribute)

**LogicalDisk**
The logical disk that is being analyzed (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Bytes per Second (HighBytesSec) | The disk read or write bytes per second exceed this threshold. | 1572864 |
| High Percent Usage (HighPercentUsage) | The percentage of time that the logical drive spends for read or write requests exceeds this threshold. | 90 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 2 |

## High Percentage Disk Time

This indication is sent when a logical disk is being used a high percent of the time.

The indication has the following attributes:

**PercentDiskTime**
The percentage of time that the logical drive is being used

**PercentWriteTime**
The percentage of time that the logical drive is used for write operations

**PercentReadTime**
The percentage of time that the logical drive is being used for read operations

**PhysicalDisk**
The physical disk on which the logical disk resides (Key attribute)

**LogicalDisk**
The logical disk that is being analyzed (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Queue Length (HighQLength) | The maximum number of outstanding requests, including those in progress, on the disk exceeds this threshold. | 3 |
| High Percent Usage (HighPercentUsage) | The time spent by the selected disk to work exceeds this threshold. | 90 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 1 |

## Logical Disk Possible Fragmentation

This indication is sent when a high percentage of logical disk is being used, but the queue length is not long, and bytes per second reading and writing are not high.

The indication has the following attributes:

**PercentDiskTime**
The percentage of time that the logical drive is being used

**DiskBytesSec**
The read and write transfer rate per second

**PhysicalDisk**
The physical disk on which the logical disk resides (Key attribute)

**LogicalDisk**
The logical disk that is being analyzed (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Bytes per Second (HighBytesSec) | The disk bytes per seconds value does *not* exceed this threshold. | 1572864 |
| High Percent Usage (HighPercentUsage) | The time spent by the selected disk to work exceeds this threshold. | 90 |
| High Queue Length (HighQLength) | This threshold must *not* be exceeded. | 3 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 1 |

## Low Disk Space

This indication is sent when a logical drive has very little free space.

The indication has the following attributes:

**PercentFreeSpace**
The percentage of free space on the logical drive

**FreeMB**
The actual size of free space on the logical drive in megabytes

**PhysicalDisk**
The physical disk on which the logical disk resides (Key attribute)

**LogicalDisk**
The logical disk that is being analyzed (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Low Disk Space (LowDiskSpace) | The free disk space drops below this threshold. | 5 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 5 |
| Holes | 0 |

## Slow Logical Drive

This indication is sent when the logical disk is too slow to keep up with the work it must perform.

The indication has the following attributes:

**CurrentDiskQLength**
The number of jobs waiting to be processed by the logical disk

**PercentDiskTime**
The percentage of time that the logical disk is being used

**AvgQLength**
The average queue length

**AvgReadQLength**
The average queue length for read operations

**AvgWriteQLength**
The average queue length for write operations

**DiskReadBytesSec**
The number of bytes read per second on the logical disk

**DiskWriteBytesSec**
The number of bytes read per second on the logical disk

**PhysicalDisk**
The physical disk on which the logical disk resides (Key attribute)

**LogicalDisk**
The logical disk that is being analyzed (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Bytes per Second (HighBytesSec) | The disk bytes per seconds value exceeds this threshold. | 1572864 |
| High Queue Length (HighQLength) | The number of outstanding requests, including those in progress, on the disk exceeds this threshold. | 3 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 1 |

# Thresholds

The following table lists the thresholds that can be set for the Logical Disk resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| High Bytes per Second (HighBytesSec) | This threshold determines the maximum amount of bytes transferred per second (being written or read). The default is based on the average logical disk. | 1572864 |
| High Queue Length (HighQLength) | This threshold determines the maximum number of outstanding requests, including those in progress, on the disk. This is an exact length, not an average over the time interval. | 3 |
| High Percent Usage (HighPercentUsage) | This threshold determines what percent of time that the selected disk drive spends for read or write requests is too high. | 90 |
| Low Disk Space (LowDiskSpace) | This threshold determines what percentage of free disk space is too low. | 5 |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource  | Context  | Properties  |
|---|---|---|
| LogicalDisk | Bytes Transferred | **LogicalDisk**<br>The logical disk to be measured<br>**DiskBytesSec**<br>The amount of bytes per second being written or read |
| | Queue Length | **LogicalDisk**<br>The logical disk to be measured<br>**AvgQLength**<br>The average queue length for jobs |
| | Percent Space | **LogicalDisk**<br>The logical disk to be measured<br>**PercentFreeSpace**<br>The percentage of free space on the logical disk |
| | Percent Disk Usage | **LogicalDisk**<br>The logical disk to be measured<br>**PercentDiskTime**<br>The percentage of time that the logical drive is being used |

# 3

# Memory Resource Model

This chapter describes the Memory resource model for Windows systems. The Memory resource model detects bottlenecks specifically associated with RAM. Because nearly every device within the system uses RAM, most bottlenecks and problems will appear here.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | TMW_MemoryModel |
| Category | Windows |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 60 seconds |



## Overview

**Resource model distribution**

This resource model should be distributed to Windows endpoints.

The Memory resource model highlights the following problems:

- Cache

  When cache is used, four different areas are examined:

  - Copy Reads
  - Data Maps
  - MDL Reads
  - Pin Reads

  The Hit Percentage Rate is examined for each area. This gives an idea of how much data was supplied by the cache and how much was supplied from another location. If cache is used, ideally it should provide at least 70% of the data. A lower percentage indicates data is being lost in the cache or that the cache is not large enough.

- Committed bytes

  Committed bytes indicate the amount of virtual memory, in bytes, that have been committed. The commit limit is the current limit of physical space (either RAM or in the pagefile) that is available for the committed bytes.

  If the committed bytes become larger than the commit limit, the pagefile must be enlarged. Having a pagefile that is constantly in flux causes some performance problems because the system spends time extending the pagefile or shrinking it. This process also involves the slowest part of the system and is prone to becoming a bottleneck.

  The committed bytes has an upper limit. This is dictated by a setting in the system control applet. If the pagefile reaches its maximum limit, it is almost guaranteed that available memory will be extremely low. In this situation a major performance degradation can be expected.

  If committed bytes continues to grow and there is not enough available physical space to hold the committed memory, a core dump will probably occur.

  The initial size of the pagefile is best determined by the amount of RAM in the system. A good guideline is 150% of the amount of RAM.

- Low available memory

  Available memory is checked constantly to ensure it does not drop below 10 MB. When available memory is low, performance declines slightly, but the system still works correctly. However, the administrator needs to know that memory is low and may need to monitor it.

  As available memory drops, Windows NT systems try to keep the available memory above 4 MB. At this point memory is too low and the operating system spends more time keeping memory available than processing requests. Additionally, high paging begins to occur, and the pagefile reaches its maximum size.

- Memory leaks

  These three different areas of memory need to be analyzed for a possible leak:

  - Private Bytes
  - System Code
  - System Drivers

  The top five processes are examined in the Private Bytes. The total pools of System Code and System Drivers are also examined for growth. This generally indicates a memory leak in the affected area of memory.

- Paging and page faulting

  Page Faults occur when Windows NT moves a piece of data to another portion of RAM or to the pagefile. Paging moves data to the hard drive, so it is incorporated in the counters for Page Faults.

  Page faults are considered high when they reach approximately 350 page faults per second. Paging is considered high at approximately 60 pages per second. Performance declines because the system moves data to and from the hard drive, typically the slowest component of the system. The hard drive and processor become busier and overall performance of the machine decreases.

# Indications and Events

The following table lists the events that can be generated by the Memory resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| TMW_HighPaging | High Paging | Warning | 26 |
| TMW_LowAvail | Low Available Memory | Warning | 27 |
| TMW_LowAvailCausingHardPaging | Low Available Memory Causing Hard Paging | Warning | 28 |
| TMW_LowAvailCausingSoftPagePagefileResize | Low Available Memory Causing Soft Paging and Pagefile Resizing | Warning | 30 |
| TMW_LowAvailCausingSoftPaging | Low Available Memory Is Causing Excessive Soft Paging | Warning | 30 |
| TMW_LowAvailCausingManyProblems | Low Available Memory Causing Many Problems | Critical | 31 |
| TMW_LowAvailHighCache | Low Available Memory with High Cache | Warning | 32 |
| TMW_LowAvailHighWS | Low Available Memory with High Working Set | Warning | 33 |
| TMW_LowAvailWithSmallPageFile | Low Available Memory With A Small Pagefile | Warning | 35 |
| TMW_LowCopyReadHits | Low Copy Read Hits | Minor | 36 |
| TMW_LowDataMapHits | Low Data Map Hits | Minor | 37 |
| TMW_LowMDLReadHits | Low MDL Read Hits | Minor | 38 |
| TMW_LowPinReadHits | Low Pin Read Hits | Minor | 39 |
| TMW_PageFileResizing | Pagefile Is Resizing | Warning | 40 |
| TMW_MemoryLeakInPB | Memory Leak In Private Bytes | Critical | 41 |
| TMW_MemoryLeakInSC | Memory Leak in System Code | Critical | 42 |
| TMW_MemoryLeakInSD | Memory Leak in System Drivers | Critical | 43 |

## High Paging

This indication is sent when there is a high amount of paging to the pagefile.

The indication has the following attributes:

**PagesSec**
The number of pages swapped per second

**TotalAvail**
The total of available memory in bytes

**TotalWorkingSet**
The total working set size in bytes

**TotalCache**
The total cache size in bytes

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| ExcessivePaging | The number of pages per second exceeds this threshold. | 60 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 15 |
| Holes | 1 |

## Low Available Memory

This indication is sent when the available memory is low, but it is not clear in which area of memory the problem is located.

The indication has the following attributes:

**TotalAvail**
The total of available memory in bytes

**TotalWorkingSet**
The total working set size in bytes

**TotalCache**
The total cache size in bytes

**PercentAvail**
The percentage of available memory in comparison to the size of the working set and the cache

**PercentWS**
The percentage of working set memory in comparison to the size of available memory and the cache

**PercentCache**
The percentage of cache memory in comparison to the size of the working set and the available memory

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Minimum Available Bytes (MinimumAvail) | Available memory is lower than this threshold. | 10485760 |
| N/A | This indication is sent when the cache must be larger than the total amount of Working Set and available memory. | N/A |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 2 |

## Low Available Memory Causing Hard Paging

This indication is sent when low available memory causes too much paging to the pagefile.

The indication has the following attributes:

**TotalAvail**
The total of available memory in bytes

**PageFaultsSec**
The current rate of page faults per second

**PagesSec**
The current rate of pages per second

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Minimum Available Bytes (MinimumAvail) | Available memory is lower than this threshold. | 10485760 |
| Minimum Committed Bytes (MinimumCommitted) | This threshold is *not* crossed. | 5242880 |
| ExcessivePageFaults | Paging to disk and to other areas of memory is higher than this threshold. | 350 |
| ExcessivePaging | The number of pages per second exceeds this threshold. | 60 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 1 |

## Low Available Memory Causing Soft Paging and Pagefile Resizing

This indication is sent when excessive soft paging and pagefile resizing is being caused by low available memory.

The indication has the following attributes:

**TotalAvail**
The total of available memory in bytes

**PageFaultsSec**
The current rate of page faults per second

**CommittedBytes**
The current total of committed bytes

**CommittedLimit**
The upper limit of the committed bytes

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Minimum Available Bytes (MinimumAvail) | Available memory is lower than this threshold. | 10485760 |
| ExcessivePageFaults | Paging to disk and to other areas of memory is higher than this threshold. | 350 |
| Excessive Paging (ExcessivePaging) | This threshold must *not* be exceeded. | 60 |
| Minimum Committed Bytes (MinimumCommitted) | The difference between the current committed bytes and the peak current bytes is smaller than this threshold. | 5242880 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Low Available Memory Is Causing Excessive Soft Paging

This indication is sent when the available memory is low and is causing a high amount of page faults

The indication has the following attributes:

**PageFaultsSec**
The current value for page faults per second

**TotalAvail**
The total of available memory in bytes

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Minimum Available Bytes (MinimumAvail) | Available memory is lower than this threshold. | 10485760 |
| ExcessivePageFaults | Page faults per second are higher than this threshold. | 350 |
| Excessive Paging (ExcessivePaging) | This threshold must *not* be exceeded. | 60 |
| Minimum Committed Bytes (MinimumCommitted) | This threshold must *not* be exceeded. | 5242880 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 1 |

## Low Available Memory Causing Many Problems

This indication is sent when multiple problems are being caused by lack of available memory.

The indication has the following attributes:

**TotalAvail**
The total of available memory in bytes

**PageFaultsSec**
The current rate of page faults per second

**PagesSec**
The current rate of pages per second

**CommittedBytes**
The current total of committed bytes

**CommittedLimit**
The upper limit of the committed bytes

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Minimum Available Bytes (MinimumAvail) | Available memory is lower than this threshold. | 10485760 |
| ExcessivePageFaults | Paging to disk and to other areas of memory is higher than this threshold. | 350 |
| ExcessivePaging | The number of pages per second exceeds this threshold. | 60 |
| Minimum Committed Bytes (MinimumCommitted) | The difference between the current committed bytes and the peak current bytes is smaller than this threshold. | 5242880 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Low Available Memory with High Cache

This indication is sent when available memory is low due to a large amount of memory being taken up by cache.

The indication has the following attributes:

**TotalAvail**
The total of available memory in bytes

**TotalWorkingSet**
The total working set size in bytes

**TotalCache**
The total cache size in bytes

**PercentAvail**
The percentage of available memory in comparison to the size of the working set and the cache

**PercentWS**
The percentage of working set memory in comparison to the size of available memory and the cache

**PercentCache**
The percentage of cache memory in comparison to the size of the working set and the available memory

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Minimum Available Bytes (MinimumAvail) | Available memory is lower than this threshold. | 10485760 |
| N/A | This indication is sent when the cache must be larger than the total amount of Working Set and available memory. | N/A |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

## Low Available Memory with a High Working Set

This indication is sent when available memory is low and the working set is high.

The indication has the following attributes:

**TotalAvail**
The total of available memory in bytes

**TotalWorkingSet**
The total working set size in bytes

**TotalCache**
The total cache size in bytes

**PercentAvail**
The percentage of available memory in comparison to the size of the working set and the cache

**PercentWS**
The percentage of working set memory in comparison to the size of available memory and the cache

**PercentCache**
The percentage of cache memory in comparison to the size of the working set and the available memory

**HighProcessWorkingSet**
The working set of the process with the highest working set

**HighProcessPrivateBytes**
The private bytes of the process with the highest working set

**HighProcessPoolNPBytes**
The pool non-paged bytes of the process with the highest working set

**HighProcessPoolPPBytes**
The pool paged bytes of the process with the highest working set

**HighProcessVirtualBytes**
The virtual bytes of the process with the highest working set

**HighProcessIDProcess**
The numeric ID of the process with the highest working set

**TotalProcesses**
The total number of processes

**HighProcess**
The string name of the process with the highest working set

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Minimum Available Bytes (MinimumAvail) | Available memory is lower than this threshold. | 10485760 |
| N/A | This indication is sent when the working set must be larger than the total amount of cache and available memory. | N/A |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 2 |

## Low Available Memory with a Small Pagefile

This indication is sent when available memory is low and the pagefile is being resized.

The indication has the following attributes:

**CommittedBytes**
The current total of committed bytes

**CommittedLimit**
The upper limit of the committed bytes

**TotalAvail**
The total of available memory in bytes

**TotalWorkingSet**
The total working set size in bytes

**TotalCache**
The total cache size in bytes

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Minimum Committed | The difference between the committed bytes and committed limit is lower than this threshold. | 5242880 |
| Minimum Available Bytes (MinimumAvail) | Available memory is lower than this threshold. | 10485760 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Low Copy Read Hits

This indication is sent when a copy read hit from cache does not retrieve a high percentage of data.

The indication has the following attributes:

**AvrgCopyReadHitsPercent**
The percentage average of copy read hits

**TotalAvail**
The total of available memory in bytes

**TotalWorkingSet**
The total working set size in bytes

**TotalCache**
The total cache size in bytes

**PercentAvail**
The percentage of available memory in comparison to the size of the working set and the cache

**PercentWS**
The percentage of working set memory in comparison to the size of available memory and the cache

**PercentCache**
The percentage of cache memory in comparison to the size of the working set and the available memory

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Low Cache Hits Percent (LowCacheHitsPercent) | The number of copy read hits is lower than the threshold. | 70 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 1 |

## Low Data Map Hits

This indication is sent when a data map hit from cache does not retrieve a high percentage of data.

The indication has the following attributes:

**AvrgDataMapHitsPercent**
The percentage average of data map hits

**TotalAvail**
The total of available memory in bytes

**TotalWorkingSet**
The total working set size in bytes

**TotalCache**
The total cache size in bytes

**PercentAvail**
The percentage of available memory in comparison to the size of the working set and the cache

**PercentWS**
The percentage of working set memory in comparison to the size of available memory and the cache

**PercentCache**
The percentage of cache memory in comparison to the size of the working set and the available memory

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Low Cache Hits Percent (LowCacheHitsPercent) | The number of copy read hits is lower than the threshold. | 70 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 1 |

## Low MDL Read Hits

This indication is sent when the MDL read hit from cache does not retrieve a high percentage of data. Continuous occurrences of this indication show that there is not enough cache. Either adding more RAM or adjusting the registry will help this situation.

The indication has the following attributes:

**AvrgMDLReadHitsPercent**
The percentage average of MDL read hits

**TotalAvail**
The total of available memory in bytes

**TotalWorkingSet**
The total working set size in bytes

**TotalCache**
The total cache size in bytes

**PercentAvail**
The percentage of available memory in comparison to the size of the working set and the cache

**PercentWS**
The percentage of working set memory in comparison to the size of available memory and the cache

**PercentCache**
The percentage of cache memory in comparison to the size of the working set and the available memory

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Low Cache Hits Percent (LowCacheHitsPercent) | The number of MDL read hits is lower than this threshold | 70 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 1 |

## Low Pin Read Hits

This indication is sent when a pin read hit from cache does not retrieve a high percentage of data. The pin read hits are too low. Continuous occurrences of this indication show that there is not enough cache. Either adding more RAM or adjusting the registry will help this situation.

The indication has the following attributes:

**AvrgPinReadHitsPercent**
The percentage average of pin read hits

**TotalAvail**
The total of available memory in bytes

**TotalWorkingSet**
The total working set size in bytes

**TotalCache**
The total cache size in bytes

**PercentAvail**
The percentage of available memory in comparison to the size of the working set and the cache

**PercentWS**
The percentage of working set memory in comparison to the size of available memory and the cache

**PercentCache**
The percentage of cache memory in comparison to the size of the working set and the available memory

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Low Cache Hits Percent (LowCacheHitsPercent) | The number of pin read hits is lower than this threshold. | 70 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 1 |

## Pagefile Is Resizing

The pagefile is resizing indication. This indication is sent when the pagefile is being resized.

The indication has the following attributes:

**CommittedBytes**
The current total of committed bytes

**CommittedLimit**
The upper limit of the committed bytes

**TotalAvail**
The total of available memory in bytes

**TotalWorkingSet**
The total working set size in bytes

**TotalCache**
The total cache size in bytes

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Minimum Committed Bytes (MinimumCommitted) | The difference between the current committed bytes and the peak current bytes is smaller than this threshold. | 5242880 |
| Minimum Available Bytes (MinimumAvail) | This threshold must *not* be crossed. | 10485760 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Memory Leak In Private Bytes

This indication is sent when there is a memory leak in the private memory, which is the memory pages that cannot be shared with other processes.

The indication has the following attributes:

**IDProcess**
The numeric ID of the process with the memory leak

**CurrentWorkingSet**
The current working set of the process

**CurrentPrivateBytes**
The current private bytes of the process

**CurrentPoolNonPagedBytes**
The current pool of non-paged bytes of the process

**CurrentPoolPagedBytes**
The current pool of paged bytes of the process

**CurrentVirtualBytes**
The current virtual bytes set of the process

**Process**
The name of the process

No specific threshold causes this indication to be sent. The decision tree examines the top five processes in this memory area and compares them with the top five processes from the last cycle. It tries to find identical processes by matching the process IDs and then comparing the amount of memory it is using for this particular pool of memory. If the memory size has grown since the last cycle, this indication is sent.

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 40 |
| Holes | 1 |

## Memory Leak In System Code

This indication is sent when there is a memory leak in the system code. The system code is the area of pageable memory in ntoskrnl.exe, hal.dll, and the boot drivers and file systems loaded by ntldr/osloader.

The indication has the following attributes:

**OrigSize**
The original size of the system code memory

**SysDriverBytes**
The current size of the system code memory in bytes

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| N/A | The decision tree examines the system code pool of memory and finds that the memory size has grown since the previous cycle. | N/A |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 40 |
| Holes | 10 |

## Memory Leak In System Drivers

This indication is sent when there is a memory leak in the system drivers.

The indication has the following attributes:

**OrigSize**
The original size of the system driver memory

**SysCode**
The current size of the system driver memory

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| N/A | The decision tree examines the system drivers pool of memory and finds that the memory size has grown since the previous cycle. | N/A |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 40 |
| Holes | 10 |

# Thresholds

The following table lists the thresholds that can be set for the Memory resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| Excessive Page Faults (ExcessivePageFaults) | This threshold indicates how much paging to disk and to other areas of memory is too high. | 350 |
| Excessive Paging (ExcessivePaging) | This threshold indicates when pages per second are too high. | 60 |
| Low Cache Hits Percent (LowCacheHitsPercent) | This threshold indicates what percentage is too low when reading data from the cache. | 70 |
| Minimum Available Bytes (MinimumAvail) | This threshold indicates at what point available memory is too low (in bytes). | 10485760 |
| Minimum Committed Bytes (MinimumCommitted) | This threshold indicates the lowest limit for the difference between total committed bytes and committed limit values. | 5242880 |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource  | Context  | Properties  |
|---|---|---|
| Process | Process Memory Usage | **Process** The process being monitored<br>**WorkingSet** The working set<br>**PrivateBytes** The number of private bytes used by the process<br>**VirtualBytes** The number of virtual bytes used by the process<br>**ID** The process ID |
| Memory | Memory Usage | **Memory** The memory instance<br>**TotalAvail** The total available memory<br>**TotalCache** The total cache memory<br>**CommittedBytes** The bytes committed to this memory |
| | Paging | **PagesSec** The number of pages per second<br>**PageFaultsSec** The rate of page faults per second |

# 4

# Network Interface Card Resource Model

The Network Interface Card (NIC) resource model detects bottlenecks specifically associated with all network interface cards installed on Windows NT systems.

**Note:** The Network Interface Card (NIC) resource model is not supported on Windows 2000.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | TMW_NetworkIntCard |
| Category | System |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | Yes |
| Clearing events | Yes |
| Default cycle time | 150 seconds |

## Prerequisites

The Network Interface Card resource model requires Windows NT. SNMP and the Network Monitoring Agent must be installed in the Network properties.

## Overview

**Resource model distribution**

This resource model should be distributed to *Windows NT* endpoints.

The Network Interface Card resource model highlights the following problems:

- Broadcast frames

  If too much of the network is broadcast frames for an extended period of time, the network administrator should segment the network further or redesign the network to use fewer broadcast frames.

- Network interface card

  The core of the network subsystem is the network interface card (NIC). If a network card is not fast enough to handle the number of requests from other nodes on the

network or from internal counterparts, all aspects of the network subsystem are affected. Tests are made to determine if a network card is too slow, or if there is too much data trying to be sent through the network.

- Percentage usage

  To identify the network component being overused, tests are performed on the percentage of the network segment being used, the server service, and the workstation service (redirector). This reveals if the system is part of a busier network, is servicing too many requests, or is requesting too much.

  The system administrator must determine whether to upgrade the server or workstation service, or whether to distribute the workload differently.

- Server and workstation services

  The server and workstation services are vital to the network operations of Windows NT and Windows 2000. Performance statistics for each service, such as bytes per second in and out, sessions with errors, or high requests to the services are examined.

## Indications and Events

The following table lists the events that can be generated by the Network Interface Card resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| TMW_AdjustWorkItems | Adjust Work Items | Warning | 47 |
| TMW_HighBroadcastFrames | High Broadcast Frames | Warning | 48 |
| TMW_HighCurrentCommands | High Current Commands | Warning | 49 |
| TMW_HighErroredRatio | High Errored Ratio | Critical | 50 |
| TMW_NICOverworked | Network Interface Card Overworked | Warning | 51 |
| TMW_NICTooSlow | Network Interface Card Too Slow | Warning | 52 |
| TMW_RedirectorAffectingServer | Redirector Affecting Server | Warning | 53 |
| TMW_RedirectorOverloaded | Redirector Overloaded | Warning | 55 |
| TMW_RedirectorOverloadedAffectingSegment | Redirector Overloaded Affecting Segment | Warning | 56 |
| TMW_SegmentAffectingRedirector | Segment Affecting Redirector | Warning | 58 |
| TMW_SegmentAffectingServer | Segment Affecting Server | Warning | 60 |
| TMW_ServerAffectingRedirector | Server Affecting Redirector | Warning | 61 |
| TMW_ServerOverloaded | Server Overloaded | Warning | 62 |
| TMW_ServerOverloadedAffectingSegment | Server Overloaded Affecting Segment | Warning | 63 |

## Adjust Work Items

The work item shortage is too high and the Windows NT registry needs to be adjusted. This indication occurs when there is only a high work item shortage. This can be remedied by adjusting the max work items in the registry.

The indication has the following attributes:

**WorkItemShortages**

The number of times that work items cannot be allocated to service incoming requests to the server

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Percent Bytes Per Second (HighPercentBytesSec) | To test for a high amount of bytes per second, a dynamic threshold value needs to be created based on the current bandwidth of the network interface card. Because there will be an internal dynamic calculation, the threshold value is a percentile that indicates what percentage of the bandwidth is considered too high. For example: 830 KBPS is considered high for 10 MB ethernet networks. This value is about 8 percent. | 9 |
| High Percent Utilization (HighPercentUtilization) | This threshold determines at what point the network utilization of the internal segment is too high. | 40 |
| High Work Item Shortages (HighWorkItemShortages) | This threshold must not be exceeded. | 10 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

## High Broadcast Frames

This indication is sent when broadcast frames are too high on the internal segment.

The indication has the following attributes:

**PercentNetworkUtil**
The percentage value of network utilization on the segment

**PercentBroadcastFrames**
The percentage value of broadcast frames on the network segment

**NetworkSegment**
The identity of the network segment being analyzed (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Percent Broadcast (HighPercentBroadcast) | This threshold measures the percentage of broadcast frames on the internal network segment. | 90 |
| High Percent Utilization (HighPercentUtilization) | This threshold determines at what point the network utilization of the internal segment is too high. | 40 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

## High Current Commands

This indication is sent when current commands are too high on the redirector. Because the percentage utilization of bytes per second is normal, this is most likely due to a problem with the workstation service or a software component tied into the service.

The indication has the following attributes:

**CurrentCommands**
The current commands that are queued for service on the redirector

**RedirectorBPS**
The bytes per second being transferred through the redirector

**CurrentBandWidth**
The bandwidth of the network interface card

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Current Commands Modifier (HighCurrCmdsMod) | This threshold counts the requests to the redirector that are currently queued for service. If this number is much larger than the number of network adapter cards installed in the computer, the networks, servers, or both being accessed have serious bottlenecks. This value is used in conjunction with the number of network cards installed on the system. The final threshold value is calculated by multiplying the number of network cards in the system by this threshold. If the current commands rise above this level, there is a bottleneck occurring. When the value of current commands is greater than the product of this threshold times the total number of installed network cards, this indication is generated. | 10 |
| High Percent Bytes Per Second (HighPercentBytesSec) | This threshold must *not* be exceeded. | 9 |
| High Percent Utilization (HighPercentUtilization) | This threshold determines at what point the network utilization of the internal segment is too high. | 40 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

## High Errored Ratio

This indication is sent when the ratio between the number of sessions errored out, and the sum of sessions timed out and logged off is high. This indicates how frequently network problems are causing dropped sessions on the server.

The indication has the following attributes:

**SessionsErroredRatio**
The ratio of sessions errored out to the total sessions on the server

**SessionsErroredOut**
The number of sessions errored out on the server

**SessionsForcedOff**
The number of sessions forced off the server bandwidth of the network interface card

**SessionsLoggedOff**
The number of sessions logged off the server

**ServerTotalSessions**
The total number of sessions on the server

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Errored Out Ratio (HighErroredOutRatio) | This threshold indicates the ratio of sessions being errored out in comparison to sessions either logging out or being timed out. This number is a percentage generated by dividing the number of sessions errored out by the sum of the sessions forced off (timed out) and logged off. | 50 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

## Network Interface Card Overworked

This indication is sent when the network interface card is overworked. Bytes per second through the network card is too high and the outbound queue length is getting too large.

The indication has the following attributes:

**OutputQueueLength**
The number of requests queued up to be sent out the network interface card

**NICBPS**
The bytes per second transferred through the network interface card

**CurrentBandWidth**
The bandwidth of the network interface card

**NetworkInterface**
The identity of network interface card being examined (Key attribute)

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Output Queue Length Modifier (HighOutputQueueLengthMod) | The network card has an output queue that is used to queue network requests. This threshold is a modifier that can be adjusted by the administrator to adjust the threshold for the output queue length for the network interface. This threshold is calculated using the total number of network interface cards in the system plus this value. | 3 |
| High Percent Bytes Per Second (HighPercentBytesSec) | To test for a high amount of bytes per second, a dynamic threshold value needs to be created based on the current bandwidth of the network interface card. Because there will be an internal dynamic calculation, the threshold value is a percentile that indicates what percentage of the bandwidth is considered too high. For example: 830 KBPS is considered high for 10 MB ethernet networks. This value is about 8 percent. | 9 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 1 |

## Network Interface Card Too Slow

This indication is sent when the outbound queue length is getting too large but the bytes per second traveling through the network interface card are normal.

The indication has the following attributes:

**OutputQueueLength**
The number of requests queued up to be sent out the network interface card

**NICBPS**
The bytes per second transferred through the network interface card

**CurrentBandWidth**
The bandwidth of the network interface card

**NetworkInterface**
The identity of network interface card being examined (Key attribute)

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Output Queue Length Modifier (HighOutputQueueLengthMod) | The network card has an output queue that is used to queue network requests. This threshold is a modifier that can be adjusted by the administrator to adjust the threshold for the output queue length for the network interface. This threshold is calculated using the total number of network interface cards in the system plus this value. | 3 |
| High Percent Bytes Per Second (HighPercentBytesSec) | To test for a high amount of bytes per second, a dynamic threshold value needs to be created based on the current bandwidth of the network interface card. Because there will be an internal dynamic calculation, the threshold value is a percentile that indicates what percentage of the bandwidth is considered too high. For example: 830 KBPS is considered high for 10 MB ethernet networks. This value is about 8 percent. | 9 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 1 |

## Redirector Affecting Server

This indication is sent when there is a high work item shortage, but bytes per second on the server are not exceeding the threshold, and bytes per second on the redirector are exceeding the threshold. This affects the percentage utilization of the network segment.

The indication has the following attributes:

**WorkItemShortages**
The number of times that work items cannot be allocated to service incoming requests to the server

**RedirectorBPS**
The number of bytes per second transferring through the redirector

**CurrentBandWidth**
The bandwidth of the network interface card

**PercentNetworkUtil**
The percentage of network utilization on the segment

**NetworkSegment**
The identity of the network segment being analyzed (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Percent Bytes Per Second (HighPercentBytesSec) | To test for a high amount of bytes per second, a dynamic threshold value needs to be created based on the current bandwidth of the network interface card. Because there will be an internal dynamic calculation, the threshold value is a percentile that indicates what percentage of the bandwidth is considered too high. For example: 830 KBPS is considered high for 10 MB ethernet networks. This value is about 8 percent. The bytes per second through the redirector are greater than the number of bytes per second. This number is reached by multiplying the current bandwidth with this threshold, and dividing the result by 100. | 9 |
| High Percent Utilization (HighPercentUtilization) | This threshold determines at what point the network utilization of the internal segment is too high. | 40 |
| High Work Item Shortages (HighWorkItemShortages) | This occurs when no work item is available, or no work item can be allocated to service the incoming request. A work item is used when the server service receives a request. HighWorkItemShortages indicates that the initial work items or maximum work items parameters may need tuning. | 10 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

## Redirector Overloaded

This indication is sent when current commands are high and bytes per second on the redirector are high, but the network segment percentage utilization is normal. The redirector is unable to keep up with the number of requests being made to it.

The indication has the following attributes:

**CurrentCommands**
The current commands that are queued for service on the redirector

**RedirectorBPS**
The number of bytes per second transferring through the redirector

**CurrentBandWidth**
The bandwidth of the network interface card

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Current Commands Modifier (HighCurrCmdsMod) | This threshold counts the requests to the redirector that are currently queued for service. If the number of current commands is greater than the product of this threshold times the total number of installed network cards, this indication is sent. | 10 |
| High Percent Broadcast (HighPercentBroadcast) | This threshold must *not* be exceeded. | 90 |
| High Percent Bytes Per Second (HighPercentBytesSec) | The bytes per second through the redirector are greater than this threshold, which is calculated by multiplying the current bandwidth with the HighPercentBytesSec threshold, and dividing the result by 100. | 9 |
| High Percent Utilization (HighPercentUtilization) | This threshold determines at what point the network utilization of the internal segment is too high. | 40 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

## Redirector Overloaded Affecting Segment

This indication is sent when the redirector, or workstation service, is being overloaded and it is generating a bottleneck on the network segment. The current commands and bytes per second on the redirector are exceeded. This usually causes the percentage utilization of the network segment to be exceeded.

The indication has the following attributes:

**CurrentCommands**
The current commands that are queued for service on the redirector

**RedirectorBPS**
The number of bytes per second transferring through the redirector

**CurrentBandWidth**
The bandwidth of the network interface card

**PercentNetworkUtil**
The percentage of network utilization on the segment

**NetworkSegment**
The network segment being analyzed (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Current Commands Modifier (HighCurrCmdsMod) | This threshold counts the requests to the redirector that are currently queued for service. If this number is much larger than the number of network adapter cards installed in the computer, the networks, servers, or both being accessed have a serious bottleneck. This value is used in conjunction with the number of network cards installed on the system. The final threshold value is calculated by multiplying the number of network cards in the system by this threshold. | 10 |
| High Percent Broadcast (HighPercentBroadcast) | This threshold must *not* be exceeded. | 90 |
| High Percent Bytes Per Second (HighPercentBytesSec) | To test for a high amount of bytes per second, a dynamic threshold value needs to be created based on the current bandwidth of the network interface card. When the bytes per second through the redirector are greater than the number of high bytes per second, this indication is sent. This value is calculated by multiplying the current bandwidth with the High Percentage Bytes per Second threshold, and dividing the result by 100. | 9 |
| High Percent Utilization (HighPercentUtilization) | This threshold determines at what point the network utilization of the internal segment is too high. | 40 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

## Segment Affecting Redirector

This indication is sent when the current commands are high, but bytes per second on both the redirector and server services are not high. It is then assumed that the current commands are high due to a high percentage utilization of the network segment.

The indication has the following attributes:

**CurrentBandWidth**
The bandwidth of the network interface card

**CurrentCommands**
The current commands that are queued for service on the redirector

**PercentNetworkUtil**
The percentage of network utilization on the segment

**NetworkSegment**
The network segment being analyzed (Key attribute)

**ServerBPS**
The bytes per second being transferred through the server

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Current Commands Modifier (HighCurrCmdsMod) | This threshold counts the requests to the redirector that are currently queued for service. If this number is much larger than the number of network adapter cards installed in the computer, the networks, servers, or both being accessed have a serious bottleneck. This value is used in conjunction with the number of network cards installed on the system. The final threshold value is calculated by multiplying the number of network cards in the system by this threshold. | 10 |
| High Percent Broadcast (HighPercentBroadcast) | This threshold must *not* be exceeded. | 90 |
| High Percent Bytes Per Second (HighPercentBytesSec) | To test for a high amount of bytes per second, a dynamic threshold value needs to be created based on the current bandwidth of the network interface card. Because there will be an internal dynamic calculation, the threshold value is a percentile that indicates what percentage of the bandwidth is considered too high. For example: 830 KBPS is considered high for 10 MB ethernet networks. This value is about 8 percent. | 9 |
| High Percent Utilization (HighPercentUtilization) | This threshold determines at what point the network utilization of the internal segment is too high. | 40 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

## Segment Affecting Server

This indication is sent when there is a high work item shortage, but bytes per second on the redirector and the server are not high. It is assumed that the high work item shortage is due to a percentage utilization on the network segment.

The indication has the following attributes:

**WorkItemShortages**
The number of times that work items cannot be allocated to service incoming requests to the server

**PercentNetworkUtil**
The percentage of network utilization on the segment

**NetworkSegment**
The network segment being analyzed (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Percent Broadcast (HighPercentBroadcast) | This threshold must *not* be exceeded. | 90 |
| High Percent Bytes Per Second (HighPercentBytesSec) | To test for a high amount of bytes per second, a dynamic threshold value needs to be created based on the current bandwidth of the network interface card. Because there will be an internal dynamic calculation, the threshold value is a percentile that indicates what percentage of the bandwidth is considered too high. For example: 830 KBPS is considered high for 10 MB ethernet networks. This value is about 8 percent. | 9 |
| High Percent Utilization (HighPercentUtilization) | This threshold determines at what point the network utilization of the internal segment is too high. | 40 |
| High Work Item Shortages (HighWorkItemShortages) | This occurs when no work item is available, or no work item can be allocated to service the incoming request. A work item is used when the server service receives a request. HighWorkItemShortages indicates that the initial work items or maximum work items parameters may need tuning. | 10 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

## Server Affecting Redirector

This indication is sent when current commands are high, server bytes per second are high, but bytes per second on the redirector are not high. This causes the network segment percentage utilization to rise. The redirector for the segment is overloaded due to the high amount of information coming through the server.

The indication has the following attributes:

**CurrentCommands**
The current commands that are queued for service on the redirector

**ServerBPS**
The bytes per second being transferred through the server

**CurrentBandWidth**
The bandwidth of the network interface card

**PercentNetworkUtil**
The percentage of network utilization on the segment

**NetworkSegment**
The network segment being analyzed (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Current Commands Modifier (HighCurrCmdsMod) | This threshold counts the requests to the redirector that are currently queued for service. If this number is much larger than the number of network adapter cards installed in the computer, the networks, servers, or both being accessed have a serious bottleneck. This value is used in conjunction with the number of network cards installed on the system. The final threshold value is calculated by multiplying the number of network cards in the system by this threshold. | 10 |
| High Percent Broadcast (HighPercentBroadcast) | This threshold must *not* be exceeded. | 90 |
| High Percent Bytes Per Second (HighPercentBytesSec) | The bytes per second through the server are greater than the number of high bytes per second. This value is calculated by multiplying the current bandwidth with this threshold, and dividing the result by 100. | 9 |
| High Percent Utilization (HighPercentUtilization) | This threshold determines at what point the network utilization of the internal segment is too high. | 40 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

## Server Overloaded

This indication is sent when there is a high work item shortage and bytes per second are high on the server service. Because the percentage utilization of the network segment is normal, the work demand on the server service is probably too high.

The indication has the following attributes:

**WorkItemShortages**
The number of times that work items cannot be allocated to service incoming requests to the server (Key attribute)

**ServerBPS**
The bytes per second being transferred through the server

**CurrentBandWidth**
The bandwidth of the network interface card

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Percent Bytes Per Second (HighPercentBytesSec) | The bytes per second through the server are greater than the number of high bytes per second. This value is calculated by multiplying the current bandwidth with this threshold, and dividing the result by 100. | 9 |
| High Percent Utilization (HighPercentUtilization) | This threshold determines at what point the network utilization of the internal segment is too high. | 40 |
| High Work Item Shortages (HighWorkItemShortages) | This occurs when no work item is available, or no work item can be allocated to service the incoming request. A work item is used when the server service receives a request. HighWorkItemShortages indicates that the initial work items or maximum work items parameters may need tuning. | 10 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

## Server Overloaded Affecting Segment

This indication is sent when there is a high work item shortage and a high bytes per second through the server service. The server for a segment is overloaded due to the high amount of data being requested from the server.

The indication has the following attributes:

**WorkItemShortages**
The number of times that work items cannot be allocated to service incoming requests to the server (Key attribute)

**ServerBPS**
The bytes per second being transferred through the server

**CurrentBandWidth**
The bandwidth of the network interface card

**PercentNetworkUtil**
The percentage value of network utilization on the segment

**NetworkSegment**
The identity of the network segment being analyzed (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Percent Broadcast (HighPercentBroadcast) | This threshold must *not* be exceeded. | 90 |
| High Percent Bytes Per Second (HighPercentBytesSec) | The bytes per second through the server are greater than the number of high bytes per second. This value is calculated by multiplying the current bandwidth with this threshold, and dividing the result by 100. | 9 |
| High Percent Utilization (HighPercentUtilization) | This threshold determines at what point the network utilization of the internal segment is too high. | 40 |
| High Work Item Shortages (HighWorkItemShortages) | This occurs when no work item is available, or no work item can be allocated to service the incoming request. A work item is used when the server service receives a request. HighWorkItemShortages indicates that the initial work items or maximum work items parameters may need tuning. | 10 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

# Thresholds

The following table lists the thresholds that can be set for the Network Interface Card resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| High Current Commands Modifier (HighCurrCmdsMod) | This threshold counts the requests to the redirector that are currently queued for service. If this number is much larger than the number of network adapter cards installed in the computer, the networks, servers, or both being accessed have a serious bottleneck. This value is used in conjunction with the number of network cards installed on the system. The final threshold value is calculated by multiplying the number of network cards in the system by this threshold. | 10 |
| High Errored Out Ratio (HighErroredOutRatio) | This threshold indicates the ratio of sessions being errored out in comparison to sessions either logging out or being timed out. This number is a percentage generated by dividing the number of sessions errored out by the sum of the sessions forced off (timed out) and logged off. | 50 |
| High Output Queue Length Modifier (HighOutputQueueLengthMod) | The network card has an output queue that is used to queue network requests. This threshold is a modifier that can be adjusted by the administrator to adjust the threshold for the output queue length for the network interface. This threshold is calculated using the total number of network interface cards in the system plus this value. | 3 |
| High Percent Broadcast (HighPercentBroadcast) | This threshold measures the percentage of broadcast frames on the internal network segment. | 90 |
| High Percent Bytes Per Second (HighPercentBytesSec) | To test for a high amount of bytes per second, a dynamic threshold value needs to be created based on the current bandwidth of the network interface card. Because there will be an internal dynamic calculation, the threshold value is a percentile that indicates what percentage of the bandwidth is considered too high. For example: 830 KBPS is considered high for 10 MB ethernet networks. This value is about 8 percent. | 9 |
| High Percent Utilization (HighPercentUtilization) | This threshold determines at what point the network utilization of the internal segment is too high. | 40 |
| High Work Item Shortages (HighWorkItemShortages) | This occurs when no work item is available, or no work item can be allocated to service the incoming request. A work item is used when the server service receives a request. HighWorkItemShortages indicates that the initial work items or maximum work items parameters may need tuning. | 10 |

# Built-in Actions

This resource model has the following built-in actions:

- Adjust the InitWorkItems registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\LanManServer\Parameters\InitWorkItems to a value of 512.
- Adjust the MaxWorkItems registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\LanManServer\Parameters\MaxWorkItems. This value can only go to 64 on Windows NT Workstations and 65 535 on Windows NT Servers.

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource  | Context  | Properties  |
|---|---|---|
| Redirector | Bytes Rate | **BytesTotalSec** The rate at which the redirector is processing data bytes.<br>**Redirector** The redirector being monitored |
| | Current Commands | **Redirector** The redirector being monitored<br>**CurrentCommands** The number of requests that are currently queued for the redirector. |
| Network Segment | Percent Broadcast | **Segment** The segment being monitored<br>**PercentNetworkUtil** The percent of the network being used<br>**PercentBroadcastFrames** The percent of the network that is broadcast frames |
| Server | Bytes Rate | **BytesTotalSec** The total number of bytes passing through the network per second<br>**Server** The server being monitored |
| | Server Activity | **ServerTotalSession** The number of sessions on the server<br>**SessionsErroredOut** The number of sessions that ended in error<br>**SessionsForcedOff** The number of sessions forced offline<br>**SessionsLoggedOff** The number of sessions logged off<br>**WorkItemShortages** Identifies work item shortages<br>**server** Identifies server being monitored |

| Resource | Context | Properties |
|---|---|---|
| Network Interface Card | Network Traffic | **BytesTotalSec**<br>The total number of bytes passing through the network per second<br>**CurrentBandwidth**<br>The network bandwidth<br>**NetworkInterfaceCard**<br>The network interface card being monitored |
| | Output Queue Length | **Output Queue Length**<br>The length of the queue for output<br>**Network Interface Card**<br>The card being monitored |

# 5

# Parametric Event Log Resource Model

This chapter describes the Parametric Event Log resource model for Windows systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | TMW_ParamEventLog |
| Category | Windows |
| Thresholds | No |
| Parameters | Yes |
| Built-in actions | No |
| Clearing events | No |
| Default cycle time | 60 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to Windows endpoints.

This resource model examines the Windows NT or Windows 2000 Event Log for events that are specified by the user through the parameter configuration.

## Indications and Events

The following table lists the event that can be generated by the Parametric Event Log resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| TMW_NTEventLogOccurred | Windows Event Logged | Warning | 68 |

## Windows Event Logged

This indication is sent when the user-specified event occurs and is logged in the Windows event log.

The indication has the following attributes:

**EventID**
Identifies the event

**RepeatCount**
The number of Event ID events received during a cycle

**Logfile**
The name of the log file in which the event was generated (for example, system, security, or application).

**SourceName**
The source of the event (Key attribute)

**TimeGenerated**
The time that the event was generated

**TimeWritten**
The time that the event was written to the Event Log

**ComputerName**
The name of the computer where the event occurred (Key attribute)

**Message**
The message text of the event

**Severity**
The severity of the event

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

# Thresholds

No particular threshold is exceeded to generate this indication. If an event in the Event Log is within the first set of events (determined by MostNumberOfLogs) and has the Event ID of 9, this indication is sent.

# Parameters

The following table lists the parameters that can be set for this resource model.

| Parameter | Description | Default |
|---|---|---|
| Event IDs (Eids) | Allows you to specify the number of the required event IDs to be monitored. This must be the value of the event ID from the Windows event log. | Numeric list |
| Event Severity (Severity) | Allows you to specify the severity of the events to be monitored. You can choose any of the following severity levels:<br>■ Information: Set to Harmless on Tivoli Enterprise Console server. Default: False<br>■ Warning. Default: False<br>■ Error: Set to Minor on Tivoli Enterprise Console server. Default: False<br>■ Audit success: Set to Harmless on Tivoli Enterprise Console server. Default: False<br>■ Audit failure: Set to Minor on Tivoli Enterprise Console server. Default: False | Boolean list |
| Log File Type (LogType) | The type of log file that you are interested in. You can choose any of the following types:<br>**Security**<br>Records security events. Default: False<br>**System**<br>Records events logged by the Windows NT system components. Default: False<br>**Application**<br>Records events logged by applications. Default: False | Boolean list |
| Windows 2000 Logs (Win2kLogs) | The type of Windows 2000 file that you are interested in. You can choose any of the following types:<br>■ Directory Services. Default: False<br>■ File Replication. Default: False<br>■ DNS Server. Default: False | Boolean list |
| Source (Source) | The source of the event log. This is the software that has logged the event. The software can be either an application or a system component, such as a driver. The name of the source must be as shown in the Windows NT or Windows 2000 event log. | String list |

| Parameter | Description | Default |
|---|---|---|
| Computers (Computers) | The computer which originated the event. You must specify the exact name of the computer, as shown in the Windows NT or Windows 2000 event log. | String list |
| Filter Type (FilterType) | The type of filter applied to the parameters specified with this dialog. You can select an AND or OR filter. The AND filter will trigger an event only when all of the conditions specified by the other parameters are met. The OR filter will trigger an event when any one of the conditions specified by the other parameters is met. The default value is AND. | Choice list |

# 6

# Parametric Services Resource Model



This chapter describes the Parametric Services resource model for Windows systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | TMW_ParamServices |
| Category | Windows |
| Thresholds | No |
| Parameters | Yes |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 300 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to Windows endpoints.

The resource model allows you to monitor the services you require.

## Indications and Events

The following table lists the event that can be generated by the Parametric Services resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| TMW_ParamServicesFailingService | Services Failing Service | Critical | 72 |
| TMW_ParamServicesStoppedService | Services Stopped Service | Critical | 73 |

## Services Failing Service

This indication is sent when one of the listed services is not stable. This can cause problems for the local machine and for connected machines over the network. Services that are not stable should be stopped to prevent them from causing any bottlenecks or damage.

The indication has the following attributes:

**Name**
The name of the service being examined (Key attribute)

**StartMode**
The start mode of the service

**State** The current state of the service

**ServiceStatus**
The current status of the service

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Services Stopped Service

This indication is sent when one of the listed services is stopped. If a key service is stopped, it must be restarted to ensure Windows is working properly.

The indication has the following attributes:

**Name** The name of the service being examined (Key attribute)

**StartMode**
The start mode of the service

**ServiceStatus**
The current status of the service

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 2 |
| Holes | 0 |

# Parameters

The single parameter that can be set for this resource model is Services (services). It is a string list. You must specify the exact name of the service, as defined in the Registry under the following key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services
```

The following lists the default services that can be monitored with the Parametric Services resource model:

- lcfd
- LanmanWorkstation
- LanmanServer
- NtLmSsp
- Netlogon
- EventLog
- Browser

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource | Context | Properties |
|---|---|---|
| Service | Services Status | **State** The state of the service<br>**Status** The status of the service<br>**Service** The service being monitored |

# 7

# Parametric TCP/IP Ports Resource Model

This chapter describes the Parametric TCP/IP Ports resource model for Windows systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | TMW_ParamPorts |
| Category | Windows |
| Thresholds | No |
| Parameters | Yes |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 300 seconds |



## Overview

**Resource model distribution**
This resource model should be distributed to Windows endpoints.

This resource model checks the TCP and UDP port numbers you specify and generates events whenever these ports are in a specified state or states. You can customize this resource model to monitor the state of TCP/IP ports used by an application.

The resource model includes the information provider file *NetworkPortInfoProvider.dll*.

## Prerequisites

The parametric TCP/IP ports resource model requires TCP/IP and SNMP to be installed.

## Indications and Events

The following table lists the event that can be generated by the Parametric TCP/IP Ports resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| TMW_ParamPortStatus | State of the Defined Port | Warning | 76 |

## State of the Defined Port

This indication is sent when both of the conditions set by the parameters Port Number and Possible State of a Port are met.

The indication has the following attributes:

**LocalPort**
The port number (Key attribute)

**LocalPortName**
The name of the port being monitored

**State** Indicates the state of the port being monitored

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 2 |
| Holes | 0 |

# Parameters

The following table lists the parameters that can be set for this resource model. Together, the parameters function as a filter.

| Parameter | Description | Default |
|---|---|---|
| Port Numbers (ParamPortList) | Allows you to specify the port numbers to be monitored. | Numeric list |
| Possible State of a Port (ParamPortStatusList) | Allows you to specify the state of the defined port or ports to be monitored. Ports can be monitored for one or more of the following states:<br>■ Unknown. Default: False<br>■ LISTENING. Default: False<br>■ ESTABLISHED. Default: False<br>■ TIME_WAIT. Default: False<br>■ Fin_wait. Default: False<br>■ FIN_WAIT2. Default: False<br>■ CLOSE_WAIT. Default: False<br>■ CLOSING. Default: False<br>■ CLOSED. Default: True | Boolean list |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource | Context | Properties |
|---|---|---|
| Port | Port State | **State** The state of the port<br>**Port** The port |

# 8

# Physical Disk Resource Model

This chapter describes the Physical Disk resource model for Windows systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | TMW_PhysicalDiskModel |
| Category | Windows |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 120 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to Windows endpoints.

The Physical Disk resource model detects bottlenecks specifically associated with all physical disks that are configured on the system. This model highlights the following problems:

- Bytes transferred per second

  Bytes transferred per second indicates the raw amount of data that is transferred through the disk. The maximum bytes per second a disk subsystem can handle depends on the type of hard drive, the type of controller, the drive configuration, and, in the case of most IDE-based hard drives, the speed of the CPU. Bytes transferred per second can be a good indicator of whether or not too much data is being pushed through the disk.

- Percentage usage

  Another indicator that a disk is being over worked or is causing a bottleneck is the percentage usage of the disk. This particular threshold often needs to be altered when monitoring file servers for bottlenecks.

  File servers, along with print and mail servers, can expect the percentage usage to be consistently high and are generally expected to show high use. However, in the case of workstations, these machines should typically see high percentage usage of disks in short bursts.

Workstations experiencing high disk usage will often see declining performance in CPU speed and general system response. In these situations, a high percentage usage of the disk is generating a bottleneck.

## Prerequisites

This resource model requires disk counters to be enabled on the Windows endpoints.

To enable disks counters, perform the following steps on each endpoint:

1. On Windows NT systems, issue the **diskperf -y** command.
   On Windows 2000 systems, issue the **diskperf -yv** command.
2. Reboot the system.

## Indications and Events

The following table lists the events that can be generated by the Physical Disk resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| TMW_HighPhysicalPercentDiskTime | High Percent Disk Time | Warning | 81 |
| TMW_HighPhysicalDiskReadBytesSec | High Read Bytes per Second | Minor | 82 |
| TMW_HighPhysicalDiskXferRate | High Transfer Rate | Minor | 83 |
| TMW_HighPhysicalDiskWriteBytesSec | High Write Bytes per Second | Minor | 84 |
| TMW_PhysicalPossibleFrag | Physical Disk Possible Fragmentation | Minor | 85 |
| TMW_SlowPhysicalDrive | Slow Physical Drive | Warning | 86 |

## High Percent Disk Time

This indication is sent when a physical disk is in heavy use, and the drive cannot process the requests fast enough. This lengthens the queue.

The indication has the following attributes:

**PercentDiskTime**
The percentage of time that the physical drive is being used

**PercentWriteTime**
The percentage of time that the physical drive is used for write operations

**PercentReadTime**
The percentage of time that the physical drive is being used for read operations

**PhysicalDisk**
The physical disk on which the physical disk resides (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Bytes per Second (HighBytesSec) | This threshold must *not* be exceeded. | 1 572 864 |
| High Queue Length (HighQLength) | The maximum number of outstanding requests, including those in progress, on the disk exceeds this threshold. | 3 |
| High Percent Usage (HighPercentUsage) | The time spent by the selected disk to work exceeds this threshold. | 90 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 1 |

## High Read Bytes per Second

This indication is sent when a physical drive is in heavy use reading data.

The indication has the following attributes:

**PercentDiskRead**
The percentage of the physical drive being used to read data

**DiskReadBytesSec**
The number of bytes read per second on the physical disk

**DiskReadSec**
The number of transactions that are read per second on the physical disk

**PhysicalDisk**
The physical disk being monitored (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Bytes per Second (HighBytesSec) | The amount of bytes transferred per second (being written or read) exceeds this threshold. | 1 572 864 |
| High Percent Usage (HighPercentUsage) | The percent of time that the selected disk drive spends for read or write requests exceeds this threshold. | 90 |
| High Queue Length (HighQLength) | This threshold must *not* be exceeded. | 3 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 2 |

## High Transfer Rate

This indication is sent when a physical disk is used heavily to transfer data, both for reading and writing.

The indication has the following attributes:

**DiskXfersSec**
The rate of bytes read or written per second on the physical disk

**DiskReadsSec**
The rate of bytes read per second on the physical disk

**DiskWritesSec**
The rate of bytes written per second on the physical disk

**PercentDiskReadTime**
The percentage of the physical drive used for read operations

**PercentDiskWriteTime**
The percentage of the physical drive used for write operations

**PhysicalDisk**
The physical disk being monitored (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Bytes per Second (HighBytesSec) | The disk read or write bytes per second exceed this threshold. | 1572864 |
| High Percent Usage (HighPercentUsage) | The percentage of time that the physical drive spends for read or write requests exceeds this threshold. | 90 |
| High Queue Length (HighQLength) | This threshold must *not* be exceeded. | 3 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 2 |

## High Write Bytes per Second

This indication is sent when a physical disk is writing a lot of data.

The indication has the following attributes:

**DiskWriteBytesSec**
The number of bytes read per second on the physical disk

**DiskWriteSec**
The number of bytes read per second on the physical disk

**PercentDiskWrite**
The percent of the physical drive used during write operations

**PhysicalDisk**
The physical disk being monitored (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Bytes per Second (HighBytesSec) | The amount of bytes written per second exceeds this threshold. | 1572864 |
| High Percent Usage (HighPercentUsage) | The percent of time that the selected disk drive spends for read or write requests exceeds this threshold. | 90 |
| High Queue Length (HighQLength) | This threshold must *not* be exceeded. | 3 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 2 |

## Physical Disk Possible Fragmentation

This indication is sent when a high percentage of physical disk is being used, but the queue length is not long, and bytes per second reading and writing are not high.

The indication has the following attributes:

**PercentDiskTime**
The percentage of time that the physical drive is being used

**DiskBytesSec**
The read and write transfer rate per second

**PhysicalDisk**
The physical disk being monitored (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Bytes per Second (HighBytesSec) | The disk bytes per seconds value does *not* exceed this threshold. | 1572864 |
| High Percent Usage (HighPercentUsage) | The time spent by the selected disk to work exceeds this threshold. | 90 |
| High Queue Length (HighQLength) | This threshold must *not* be exceeded. | 3 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 1 |

## Slow Physical Drive

This indication is sent when the physical disk is too slow to keep up with the work it must perform.

The indication has the following attributes:

**CurrentDiskQLength**
The number of jobs waiting to be processed by the physical disk

**PercentDiskTime**
The percentage of time that the physical disk is being used

**AvgQLength**
The average queue length

**AvgReadQLength**
The average queue length for read operations

**AvgWriteQLength**
The average queue length for write operations

**DiskReadBytesSec**
The number of bytes read per second on the physical disk

**DiskWriteBytesSec**
The number of bytes read per second on the physical disk

**PhysicalDisk**
The physical disk being monitored (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Bytes per Second (HighBytesSec) | The disk bytes per seconds value exceeds this threshold. | 1572864 |
| High Queue Length (HighQLength) | The number of outstanding requests, including those in progress, on the disk exceeds this threshold. | 3 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 1 |

# Thresholds

The following table lists the thresholds that can be set for the Physical Disk resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| High Bytes per Second (HighBytesSec) | This threshold determines how many bytes per second (either reading per second or writing per second) is too high. Different disk sub-systems have different maximums, so the default threshold value is targeted towards the average physical disk. | 1 572 864 |
| High Percent Usage (HighPercentUsage) | This threshold determines what percent of time that the selected disk drive spends for read or write requests is too high. | 90 |
| High Queue Length (HighQLength) | This threshold determines the maximum number of outstanding requests, including those in progress, on the disk. This is an exact length, not an average over the time interval. | 3 |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource  | Context  | Properties  |
|---|---|---|
| PhysicalDisk | Bytes Transferred | **PhysicalDisk**<br>The physical disk to be measured<br>**DiskBytesSec**<br>The amount of bytes per second being written or read |
| | Percent Disk Usage | **PhysicalDisk**<br>The physical disk to be measured<br>**PercentDiskTime**<br>The percentage of time that the Physical drive is being used |
| | Queue Length | **PhysicalDisk**<br>The physical disk to be measured<br>**AvgQLength**<br>The average queue length for jobs |

# 9

# Printer Resource Model

This chapter describes the Printer resource model for Windows 2000 systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | TMW_PrintModel |
| Category | Windows |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 1200 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to Windows 2000 endpoints only.

The Printer resource model determines if there are any bottlenecks in the print engine of Windows 2000. These bottleneck checks include determining if a print queue is generating an abnormally high amount of errors, if a print device is being overworked, and general functionality of the print server.

The Printer resource model highlights the following problems:

- Printer Errors

  If the number of errors is too high then this should be addressed as soon as possible. Various errors include not ready errors (the print device is off-line or has a pending message), out of paper errors (if a print device is consistently running out of paper, then the administrator should look into upgrading the paper trays), and other job related errors.

- Percentage processor time

  As with all processes, the print server utilizes the CPU. Monitoring the percentage processor time of the print server will help catch a potential bottleneck.

# Indications and Events

The following table lists the events that can be generated by the Printer resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| TMW_HighCurrentPercentTime | High Current Percentage Time | Warning | 91 |
| TMW_HighJobErrors | High Job Errors | Warning | 92 |
| TMW_HighJobErrorsPerDay | High Job Errors Per Day | Harmless | 93 |
| TMW_HighNotReadyErrors | High Not Ready Errors | Warning | 94 |
| TMW_HighNotReadyErrorsPerDay | High Not Ready Errors Per Day | Harmless | 95 |
| TMW_HighOutOfPaperErrors | High Out Of Paper Errors | Minor | 94 |
| TMW_HighOutOfPaperErrorsPerDay | High Out Of Paper Errors Per Day | Harmless | 97 |

## High Current Percent Time

This indication is sent when a job uses a lot of CPU on the print server.

The indication has the following attributes:

**CurrentPercentTimeJob**
The current percentage processor time being used by the print job

**CurrentPercentTime**
The total current percentage usage of the CPU

**JobObject**
The identity of the print job being examined (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percent Processor (PercentProcessor) | This threshold indicates when a job has a percentage usage of the processor that is too high. | 80 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 2 |
| Holes | 0 |

## High Job Errors

This indication is sent when there is a high amount of job errors per cycle. This is typical of corrupted or bad drivers installed on either the workstation sending the print job, or the server in which the printer is installed.

The indication has the following attributes:

**JobErrors**
The total number of job errors

**PrintQueue**
The print queue, or printer in Microsoft terms, being examined (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Job Errors (JobErrors) | This threshold indicates the number of job errors per cycle that can occur before determining that a printer or print device is experiencing excessive problems. | 0 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## High Job Errors Per Day

This indication is sent when there is a high amount of job related errors per day. This is typical of corrupted or bad drivers installed on either the workstation sending the print job, or the server in which the printer is installed.

The indication has the following attributes:

**JobErrorsPerDay**
The total number of job errors per day

**TotalJobErrors**
The total number of job errors since the last reboot

**SystemUpTimeDays**
The time, in days, that the system has been running

**PrintQueue**
The print queue, or printer in Microsoft terms, being examined (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Job Errors Per Day (JobErrorsPerDay) | This threshold indicates the number of job errors per day that can occur before determining that a printer or print device is experiencing excessive problems. | 10 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

## High Not Ready Errors

This indication is sent when the printing device is either offline or experiencing a paper jam.

The indication has the following attributes:

**NotReadyErrors**
The total number of not ready errors

**PrintQueue**
The print queue, or printer in Microsoft terms, being examined (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Not Ready Errors (NotReadyErrors) | This threshold indicates how many not ready errors can occur per cycle before determining that a printer or print device is experiencing excessive problems. | 8 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## High Not Ready Errors Per Day

This indication is sent when the printing device is either off-line or experiencing a paper jam.

The indication has the following attributes:

**NotReadyErrorsPerDay**
The total number of not ready errors, calculated per day

**TotalNotReadyErrors**
The total number of not ready errors since the last reboot

**SystemUpTimeDays**
The time that the system has been running, in days

**PrintQueue**
The print queue, or printer in Microsoft terms, being examined (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Not Ready Errors Per Day (NotReadyErrorsPerDay) | This threshold indicates how many not ready errors can occur per day before determining that a printer or print device is experiencing excessive problems. | 8 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 2 |
| Holes | 0 |

## High Out Of Paper Errors

This indication is sent when the paper supply is insufficient for the job demand.

The indication has the following attributes:

**OOPErrors**
The total number of out-of-paper errors

**PrintQueue**
The print queue, or printer in Microsoft terms, being examined (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Out Of Paper Errors (OutOfPaperErrors) | This is the maximum number of times per cycle, that a print device can report an out-of-paper error before determining that the print device is experiencing too many out-of-paper errors in a cycle. | 6 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## High Out Of Paper Errors Per Day

This indication is sent when the paper supply is insufficient for the daily job demand.

The indication has the following attributes:

**OOPErrorsPerDay**
The total number of out-of-paper errors, calculated per day

**TotalOOPErrors**
The total number of out of paper errors since the last reboot

**SystemUpTimeDays**
The time that the system has been running, in days

**PrintQueue**
The print queue, or printer in Microsoft terms, being examined (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Out Of Paper Errors Per Day (OutOfPaperErrorsPerDay) | This threshold indicates the maximum number of times per day, that a print device can report an out of paper error before determining that the print device is experiencing too many out of paper errors in a day. | 6 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

# Thresholds

The following table lists the thresholds that can be set for the Printer resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| Job Errors (JobErrors) | This threshold indicates the number of job errors per cycle that can occur before determining that a printer or print device is experiencing excessive problems. | 0 |
| Job Errors Per Day (JobErrorsPerDay) | This threshold indicates the number of job errors per day that can occur before determining that a printer or print device is experiencing excessive problems. | 10 |
| Maximum Print Jobs (MostPrintJobs) | This threshold indicates the maximum number of print jobs that will be examined when looking for bottlenecks such as job errors per day, or a job's percentage processor usage. | 60 |
| Maximum Print Queues (MostPrintQueues) | This threshold indicates the maximum number of print queues (or in Microsoft terms, printers), that will be examined when looking for bottlenecks such as not ready errors and out of paper errors. | 30 |
| Not Ready Errors (NotReadyErrors) | This threshold indicates how many not ready errors can occur per cycle before determining that a printer or print device is experiencing excessive problems. | 0 |
| Not Ready Errors Per Day (NotReadyErrorsPerDay) | This threshold indicates how many not ready errors can occur per day before determining that a printer or print device is experiencing excessive problems. | 8 |
| Out Of Paper Errors (OutOfPaperErrors) | This is the maximum number of times per cycle, that a print device can report an out-of-paper error before determining that the print device is experiencing too many out-of-paper errors in a cycle. | 0 |
| Out Of Paper Errors Per Day (OutOfPaperErrorsPerDay) | This threshold indicates the maximum number of times per day, that a print device can report an out of paper error before determining that the print device is experiencing too many out of paper errors in a day. | 6 |
| Percent Processor (PercentProcessor) | This threshold indicates when a job has a percentage usage of the processor that is too high. | 80 |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource  | Context  | Properties  |
|---|---|---|
| PrintQueue | Errors | **PrintQueue**<br>The queue<br>**JobErrors**<br>The number of jobs in error<br>**NotReadyErrors**<br>The number of not ready errors<br>**OutOfPaperErrors**<br>The out-of-paper errors |

# 10

# Process Resource Model



This chapter describes the Process resource model for Windows systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | TMW_Process |
| Category | Windows |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 60 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to Windows endpoints.

The Process resource model looks for bottlenecks relating to the processes running on the system. The Process resource model highlights the following problems:

- Handle leaks

  The resource model performs a check to make sure there are no processes with a handle leak. If handle leaks go unchecked, other processes may not be able to run later.

- High CPU usage

  Processes that have high CPU usage threaten the overall performance of the system because the CPU will not have much chance to perform requests made by other processes and devices.

# Indications and Events

The following table lists the events that can be generated by the Process resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| TMW_ProcessHandleLeak | Process Handle Leak | Warning | 102 |
| TMW_ProcessHighCPU | Process High CPU | Critical | 103 |

## Process Handle Leak

This indication is sent when a process leaks handles.

The indication has the following attributes:

**CurrentHanProcessID**
The process being monitored (Key attribute)

**CurrentHandleCount**
The current number of handles allocated to the process

**Process**
The process being monitored (Key attribute)

The top *n* processes (where *n* is the value of the Maximum Processes threshold) that have a number of handles in excess of the Maximum Handles threshold, are compared with the top *n* processes from the previous cycle. Process IDs are matched and for each process present in both cycles the number of handles in use is compared. If the handle count has increased since the last cycle for a process, an indication is sent. The following thresholds are used:

| Threshold | Description | Default |
|---|---|---|
| Maximum Processes (MaxProcesses) | This threshold indicates the maximum number of processes that will be compared. These processes will be ordered from the process with the highest handle count to the lowest. | 5 |
| Maximum Handles (MaxHandles) | This threshold indicates the number of handles that a process can use, above which the process will be considered as a candidate for the Process Handle Leak indication. | 300 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 40 |
| Holes | 1 |

## Process High CPU

This indication is sent when a process is using too much CPU time.

The indication has the following attributes:

**IDProcess**
The process being monitored (Key attribute)

**PrcPercentUserTime**
The percentage usage of the CPU that is being used by the process

**PrcPercentPrivilegedTime**
The percentage privileged time of the CPU that is being used by the process

**PrcPriorityBase**
The base priority of the process

**Process**
The process being monitored (Key attribute)

An indication is sent for each of the top *n* processes (where *n* is the value of the Maximum Processes threshold) that have a CPU usage in excess of the High CPU Usage threshold. The full details of the thresholds are as follows:

| Threshold | Description | Default |
|---|---|---|
| N/A | The process ID is not equal to zero. | N/A |
| High CPU Usage (HighCPUUse) | This threshold indicates when a process is using too much of the processor in terms of percentage usage. | 60 |
| Maximum Processes (MaxProcesses) | This threshold indicates the maximum number of processes that will be taken into consideration for this indication. These processes will be ordered from the process with the highest CPU usage to the lowest. | 5 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 20 |
| Holes | 3 |

# Thresholds

The following table lists the thresholds that can be set for the Process resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| High CPU Usage (HighCPUUse) | This threshold indicates when a process is using too much of the processors in terms of percentage usage. | 60 |
| Maximum Processes (MaxProcesses) | For the Process Handle Leak indication, this threshold indicates the maximum number of processes that will be compared. These processes will be ordered from the process with the highest handle count to the lowest. For the Process High CPU indication, this threshold indicates the maximum number of processes that will be taken into consideration for this indication. These processes will be ordered from the process with the highest CPU usage to the lowest. | 5 |
| Maximum Handles (MaxHandles) | This threshold indicates the number of handles that a process can use, above which the process will be considered as a candidate for the Process Handle Leak indication. | 300 |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource  | Context  | Properties  |
|---|---|---|
| Process | Handle Usage | **Process** The name of the process<br>**HandleCount** The number of handles allocated to the process<br>**PercentProcessorTime** The percent of processor time used by the process<br>**ID** The process |
| | CPU Usage | **Process** The name of the process<br>**PercentUserTime** The percentage usage of the CPU that is being used by the process<br>**PercentPrivilegedTime** The percentage privileged time of the CPU that is being used by the process<br>**ID** The process |

# 11

# Processor Resource Model

This chapter describes the Processor resource model for Windows systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | TMW_Processor |
| Category | Windows |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 60 seconds |



## Overview

**Resource model distribution**

This resource model should be distributed to Windows endpoints.

The Processor resource model detects bottlenecks happening within or from the Central Processing Unit (CPU). The Processor resource model highlights the following problems:

- High CPU usage

  Because every process and every device on the system has to interact with the CPU, either directly or through a controller, it is imperative that the CPU is not being extensively used by a single process or device for an extended period of time. There may be exceptions to this rule for certain application-based servers.

  The system administrator needs ultimately to decide how much utilization one process makes of the CPU for these situations. Tests to determine how busy a CPU are include checking the percentage usage, the percentage interrupts, and the queue length.

- Multiple CPU problems

  With systems running with more than one processor, it is not always obvious that a CPU may not be working. A check is made to ensure that all CPUs are being utilized at approximately the same rate.

  If there are problems with a CPU, the CPU activity will either be abnormally high or totally inactive, in either case this check will detect the CPU that is faulty.

# Indications and Events

The following table lists the events that can be generated by the Processor resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| TMW_BusyHardware | Busy Hardware | Warning | 107 |
| TMW_CPUCantKeepUpWithHW | CPU Cannot Keep up with Hardware | Critical | 108 |
| TMW_HWKeepingCPUBusy | Hardware Keeping CPU Busy | Warning | 109 |
| TMW_HighPercentUsageDelta | High Percent Usage Delta | Minor | 110 |
| TMW_HighProcesses | High Processes | Critical | 111 |
| TMW_ProcessorBusy | Processor Busy | Critical | 112 |

## Busy Hardware

This indication is sent when one or more devices installed on the system are tying up the CPU.

The indication has the following attributes:

**PercentProcessorTime**
The total current percentage usage of the CPU

**PercentInterruptTime**
The current percentage usage of the CPU as it handles interrupt requests

**InterruptsSec**
The number of interrupts per second that are passed to the CPU

**Processor**
The processor being monitored (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High CPU Usage (HighCPUUsage) | This threshold indicates when a high percent of the CPU is being used. CPUs are expected to reach 100% marks quite often, but generally not expected to sustain that rate for an extended amount of time. | 80 |
| High CPU Usage Interrupt (HighCPUUsageInterrupt) | This threshold indicates when a high percent of the CPU is being used for interrupt requests, generally made by hardware devices. | 20 |
| High Interrupts per Second (HighInterruptsSec) | This threshold indicates the maximum rate at which the devices installed on the system are accessing the CPU (interrupts per second) | 600 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 15 |
| Holes | 3 |

## CPU Cannot Keep up with Hardware

This indication is sent when the CPU can not keep up with hardware requests.

The indication has the following attributes:

**PercentProcessorTime**
The total current percentage usage of the CPU

**PercentInterruptTime**
The current percentage usage of the CPU as it handles interrupt requests

**InterruptsSec**
The number of interrupts per second that are passed to the CPU

**Processor**
The processor being monitored (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High CPU Usage (HighCPUUsage) | This threshold indicates when a high percent of the CPU is being used. CPUs are expected to reach 100% marks quite often, but generally not expected to sustain that rate for an extended amount of time. | 80 |
| High CPU Usage Interrupt (HighCPUUsageInterrupt) | This threshold indicates when a high percent of the CPU is being used for interrupt requests, generally made by hardware devices. | 20 |
| High Interrupts per Second (HighInterruptsSec) | This threshold must *not* be exceeded. | 600 |
| Total CPUs Modifier (TotalCPUsMod) | This threshold is a modifier that is added to the total number of CPUs in the system. The result is the threshold that determines when the queue for the systems processors is too long. | 10 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 15 |
| Holes | 3 |

## Hardware Keeping CPU Busy

This indication is sent when the queue length is normal but CPU usage is high and a high percent of the CPU is being used for interrupt requests. This keeps the CPU busy.

The indication has the following attributes:

**PercentProcessorTime**
The total current percentage usage of the CPU

**PercentInterruptTime**
The current percentage usage of the CPU as it handles interrupt requests

**InterruptsSec**
The number of interrupts per second that are passed to the CPU

**Processor**
The processor being monitored (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High CPU Usage (HighCPUUsage) | This threshold indicates when a high percent of the CPU is being used. CPUs are expected to reach 100% marks quite often, but generally not expected to sustain that rate for an extended amount of time. | 80 |
| High CPU Usage Interrupt (HighCPUUsageInterrupt) | This threshold indicates when a high percent of the CPU is being used for interrupt requests, generally made by hardware devices. | 20 |
| High Interrupts per Second (HighInterruptsSec) | This threshold must *not* be exceeded. | 600 |
| N/A | The queue length must be normal. | N/A |
| Total CPUs Modifier (TotalCPUsMod) | This threshold must *not* be exceeded. | 10 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 15 |
| Holes | 3 |

## High Percent Usage Delta

This indication is sent when there is a high percentage of usage between multiple processors installed on a system.

The indication has the following attributes:

**PercentUsageFirst**
Identifies what percent of the most used processor is being used

**PercentUsageLast**
Identifies what percent of the least used processor is being used

**ProcessPercentProcessorTime**
The percentage processor time of the process with the highest percentage processor time

**IDProcess**
The process with the highest percentage processor time

**Processor**
The processor being monitored (Key attribute)

**LowProcessor**
The least used processor

**Process**
The name of the process with the highest percentage process time

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Percent Usage Delta (HighPercentUsageDelta) | This threshold indicates the difference in percentage usage between the highest used CPU and the lowest used CPU. | 5 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | No |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 3 |

## High Processes

This indication is sent when one or more processes have a CPU usage above the High CPU Usage Process threshold.

The indication has the following attributes:

**HighProcesses**
The total number of high usage processes

**IDProcess**
The process with the highest CPU usage

**PercentProcessorTime**
The total percentage usage of the CPU

**Process**
The name of the process with the highest CPU usage that is active on the CPU

**Processor**
The processor being monitored (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High CPU Usage (HighCPUUsage) | This threshold indicates when a high percent of the CPU is being used. CPUs are expected to reach 100% marks quite often, but generally not expected to sustain that rate for an extended amount of time. | 80 |
| High CPU Usage Interrupt (HighCPUUsageInterrupt) | This threshold indicates when a high percent of the CPU is being used for interrupt requests, generally made by hardware devices. | 20 |
| High CPU Usage Process (HighCPUUsageProcess) | This threshold indicates the maximum percent of CPU usage for a single process allowed. | 40 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 10 |
| Holes | 2 |

## Processor Busy

This indication is sent when CPU usage is high, but there do not seem to be any particular process or hardware requests.

The indication has the following attributes:

**PercentProcessorTime**
The total current percentage usage of the CPU

**PercentInterruptTime**
The current percentage usage of the CPU as it handles interrupt requests

**PercentUserTime**
The percentage user time of the CPU

**PercentPrivilegedTime**
The percentage privileged time of the CPU

**Processor**
The processor being monitored (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| High CPU Usage (HighCPUUsage) | This threshold indicates when a high percent of the CPU is being used. CPUs are expected to reach 100% marks quite often, but generally not expected to sustain that rate for an extended amount of time. | 80 |
| High CPU Usage Process (HighCPUUsageProcess) | This threshold must *not* be exceeded. | 40 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 20 |
| Holes | 1 |

# Thresholds

The following table lists the thresholds that can be set for the Processor resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| High CPU Usage (HighCPUUsage) | This threshold indicates when a high percent of the CPU is being used. CPUs are expected to reach 100% marks quite often, but generally not expected to sustain that rate for an extended amount of time. | 80 |
| High CPU Usage Interrupt (HighCPUUsageInterrupt) | This threshold indicates when a high percent of the CPU is being used for interrupt requests, generally made by hardware devices. | 20 |
| High CPU Usage Process (HighCPUUsageProcess) | This threshold indicates the maximum percent of CPU usage for a single process allowed. | 40 |
| High CPU Usage User Privilege (HighCPUUsageUserPriv) | This threshold indicates the maximum percent of CPU usage by all the processes running. This value is obtained by adding the percent user time and the percent privileged time | 70 |
| High Interrupts per Second (HighInterruptsSec) | This threshold indicates the maximum rate at which the devices installed on the system are accessing the CPU (interrupts per second) | 600 |
| High Percent Usage Delta (HighPercentUsageDelta) | This threshold indicates the difference in percentage usage between the highest used CPU and the lowest used CPU. | 5 |
| Total CPUs Modifier (TotalCPUsMod) | This threshold is a modifier that is added to the total number of CPUs in the system. The result is the threshold that determines when the queue for the systems processors is too long.The processor queue length is the instantaneous length of the processor queue in units of threads. All processors use a single queue in which threads wait for processor cycles. This length does not include the threads that are currently running. | 10 |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource | Context | Properties |
|---|---|---|
| Processor | Load Balance | **Delta** The percentage difference of usage between the highest used processor and the least used<br>**HighestPercentUsage** Identifies what percent of the most used processor is being used<br>**LowestPercentUsage** Identifies what percent of the least used processor is being used |
| | Interrupt Time | **Processor** The processor being monitored<br>**InterruptsSec** The number of interrupts per second that are passed to the CPU |
| | CPU Usage | **Processor** The processor being monitored<br>**PercentInterruptTime** The current percentage usage of the CPU as it handles interrupt requests<br>**PercentProcessorTime** The total current percentage usage of the CPU<br>**PercentPrivilegedTime** The percentage privileged time of the CPU<br>**PercentUserTime** The percentage user time of the CPU |
| System | Processor Queue Length | **Processor** The processor being monitored<br>**ProcessorQueueLength** The queue length of the processor |

# 12

# Server Performance Prediction Resource Model

This chapter describes the Server Performance Prediction resource model for Windows systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | TMW_Spp |
| Category | Windows |
| Thresholds | No |
| Parameters | Yes |
| Built-in actions | No |
| Clearing events | No |
| Default cycle time | 30 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to Windows endpoints.

The Server Performance Prediction resource model gathers data from Windows systems for the Tivoli Decision Support Guide for Server Performance Prediction (Advanced Edition). This resource model does not have thresholds, indications or events.

# Parameters

The following table lists the parameters that can be set for the Server Performance Prediction resource model.

| Parameter | Description | Default |
|---|---|---|
| Physical Disk Configuration (PhysicalDiskConf) | The instance for which you want the disk I/O and disk transfer rates calculated.<br><br>Values: 0-3. | Choice list |
| Network Interface Configuration (NetworkInterfacesConf) | The network interface for which you want the network interface values calculated.<br><br>Values: 1-4. | Choice list |
| Available Disk Space Configuration (DiskSpaceAvailableConf) | The instance for which you want the available disk space calculated.<br><br>Values: System Drive, All, C:, D:, E:, A:, all other drives. | Choice list |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource  | Context  | Properties  |
|---|---|---|
| LogicalDisk | SPP Guide for NT | **LogicalDisk** The logical disk being monitored<br>**FreeMB** The free megabytes on the disk |
| Memory | SPP Guide for NT | **Memory** The memory being monitored<br>**PageInputSec** The number of page inputs per second<br>**PageOutputSec** The number of page outputs per second<br>**Avail** The available memory |
| NetworkInterface | SPP Guide for NT | **PacketsOutboundErrors** The number of errors sending outbound packets<br>**PacketsReceivedErrors** The number of errors for packets being received<br>**PacketsSentSec** The number of packets sent per second<br>**PacketsReceivedSec** The number of packets received per second<br>**NetworkInterface** The network interface card |
| Objects | SPP Guide for NT | **Objects** The objects being monitored<br>**NumberOfProcesses** The number of processes running |

| Resource  | Context  | Properties  |
|---|---|---|
| PhysicalDisk | SPP Guide for NT | **PhysicalDisk**<br>The physical disk being monitored<br>**DiskBytesSec**<br>The disk bytes per second being written or read<br>**DiskXfersSec**<br>The rate of bytes read or written per second on the physical disk |
| System | SPP Guide for NT | **System** The name of the system being monitored<br>**ProcessorQueueLength**<br>The length of the processor queue<br>**PrcTotPrivTime**<br>The percentage privileged time of the CPU that is being used<br>**PrcTotUserTime**<br>The percentage of total user time<br>**PrcTotCpuTime**<br>The percentage of total CPU time |

# 13

# Services Resource Model

This chapter describes the Services resource model for Windows systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | TMW_Services |
| Category | Windows |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | Yes |
| Clearing events | Yes |
| Default cycle time | 300 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to Windows endpoints.

The Services resource model checks all installed services to make sure they are functioning correctly and that key services are running. It highlights the following problems:

- Key services

  Some services are vital to the operation of Windows. The resource model checks that key services are not stopped to ensure proper functioning of the system.

- Unstable services

  Services, like any other installed software, can become corrupted or unstable. The resource model checks that all installed services are stable. Unstable services must be stopped to ensure they do not harm other functions of the system.

# Indications and Events

The following table lists the events that can be generated by the Services resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| TMW_ServicesFailingService | Services Failing Service | Critical | 120 |
| TMW_ServicesStoppedService | Services Stopped Service | Critical | 121 |

## Services Failing Service

This indication is sent when any service is not stable. Services that are not stable can cause problems for the local machine and for connected machines over the network and should be stopped.

The indication has the following attributes:

**Name**
The name of the service being examined (Key attribute)

**StartMode**
The start mode of the service

**State** The current state of the service

**ServiceStatus**
The current status of the service

Because the Service resource model determines primarily whether a service is running correctly, there are no real numeric thresholds to be measured or exceeded.

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Services Stopped Service

This indication is sent when one of the key services is stopped. If a key service is stopped, it must be restarted to ensure Windows is working properly.

The indication has the following attributes:

**Name** The name of the service being examined (Key attribute)

**StartMode**
The start mode of the service

**ServiceStatus**
The current status of the service

Because the Service resource model determines primarily whether a service is running correctly, there are no real numeric thresholds to be measured or exceeded.

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 2 |
| Holes | 0 |

# Thresholds

Because the Service resource model determines primarily whether a service is running correctly, there are no real numeric thresholds to be measured or exceeded. However, in this model, the threshold values act as numerical flags where 0 is false and any other number is true. With this, the user can determine which services they want to monitor.

For example, in the threshold section, LanmanServer (Server service) is set to 1 meaning the resource model checks if the Server service is running. System Administrators may decide that on certain systems the Server service does not need to run. On those systems they can change the profile to set the Server threshold value to 0 and the resource model does not check if the Server service is running.

The following table lists the thresholds that can be set for the Services resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| Browser (Browser) | This threshold determines the smooth running of the browser service. If the browser threshold is set to 0, the browser service is not checked to ensure it has started. Use the browser service to create the list of computers and networks in the network neighborhood. It is also needed so that the local machine is present on other network neighborhoods throughout the local network. | 1 |
| EventLog (EventLog) | This threshold determines whether the EventLog service is running correctly. If this threshold is set to 0, the EventLog service is not checked to ensure it has started. This service logs errors and information regarding the local machine. Important data can be found in the event log when an administrator is trying to correct a malfunctioning component of Windows. | 1 |
| LanmanServer (LanmanServer) | This threshold determines whether the LanmanServer service is running correctly. This threshold is set to 0, the LanmanServer service is not checked to ensure it has started. Better known as the Server service, this service is used to manage shares that are accessible to other workstations on the network. Without this service, no machine can connect to the local machine. | 1 |
| LanmanWorkstation (LanmanWorkstation) | This threshold determines whether the LanmanWorkstation service is running correctly. This threshold is set to 0, the LanmanWorkstation service can not be checked to ensure it has started. This service is the workstation service; it is the counterpart of the server service. If it is not running, the local machine can not connect to other machines on the network. | 1 |
| Tivoli endpoint (lcfd) | This threshold determines whether the Tivoli endpoint service, which provides the endpoint with connectivity to the Tivoli Management Environment, is running correctly. If the Tivoli endpoint threshold is set to 0, the Tivoli endpoint service will not be checked to ensure that it has started. | 1 |

| Threshold | Description | Default |
|---|---|---|
| Netlogon (Netlogon) | This threshold determines whether this service is running correctly. If the Netlogon threshold is set to 0, the Netlogon service is not checked to ensure it has started. Use the Netlogon service when logging onto the local workstation as well as processing logon attempts from remote machines. If the Net Logon service is not functioning properly, it is possible that nobody can log on to the system. | 1 |
| NtLmSsp (NtLmSsp) | This threshold determines whether this service is running correctly. If the NtLmSsp threshold is set to 0, the NtLmSsp service is not checked to ensure it has started. The NtLmSsp is the Windows LM Security Support Provider. It provides Windows security to remote procedure call (RPC) programs that use transports other than named pipes. This service does not have to be running at all times, but if it is, this does not cause any problems. | 1 |

## Built-in Actions

This resource model has the following built-in action:

**Restart Service**

If a key service is in a stopped or paused state, this action restarts the service. The system administrator is also notified.

# 14

# TCP/IP Resource Model



This chapter describes the TCP/IP resource model for Windows systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | TMW_TCPIP |
| Category | Windows |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 180 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to Windows endpoints.

The TCP/IP resource model looks for bottlenecks found within the TCP/IP protocol. It highlights the following problems:

- Fragmented data

  Fragmented datagrams are datagrams sent through the network that are incomplete. Fragments have to wait at their destination point to be reassembled.

  When datagrams have to be reassembled, it requires CPU time that could otherwise be put to use servicing current processes or interrupt request. Therefore, excessive fragmented datagrams need to be minimized whenever possible.

- Network congestion

  Congestion can be a big problem on any type of network. One way of diagnosing congestion is to look for retransmitted TCP segments and high ping rates.

## Prerequisites

The TCP/IP resource model requires TCP/IP and SNMP to be installed.

# Indications and Events

The following table lists the events that can be generated by the resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| TMW_HighFragRatio | High Fragment Ratio | Warning | 127 |
| TMW_HighPing | High Ping | Warning | 128 |
| TMW_SegmentsReXmit | Segments ReXmit | Warning | 129 |

## High Fragment Ratio

This indication is sent when the ratio between fragmented datagrams received to total datagrams received is greater than the High Fragment Ratio threshold. This ratio rises the more the local network becomes congested, particularly with IP.

This particular indication is more of an informative indication, as it is usually a result of some other bottleneck. Indications coming from the Network Interface Card resource model can help determine the cause.

The indication has the following attributes:

**FragsToDGRatio**
The ratio of fragmented datagrams to total datagrams

**DGReceivedSec**
The number of datagrams received per second

**DGSentSec**
The number of datagrams being sent per second

**DGSec**
The number of datagrams transmitted per second

**FragReceivedSec**
The number of fragmented datagrams received per second

**FragReassembledSec**
The number of fragmented datagrams reassembled per second

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Fragment Ratio (HighFragRatio) | This threshold indicates a ratio between fragmented datagrams and datagrams received whole. It is used to determine when this value becomes too high. | 80 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

## High Ping

This indication is sent when a high amount of datagrams is being received compared to the number of segments, which is common when a system is being pinged excessively.

The indication has the following attributes:

**DGReceivedSec**
The number of datagrams received per second

**SegmentsRcvdSec**
The number of segments received per second

**DGSentSec**
The number of datagrams being sent per second

**SegmentsSentSec**
The number of segments being sent per second

**DGSec**
The number of datagrams transmitted per second

**SegmentsSec**
The number of segments transmitted per second

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Low Segments (LowSegments) | This threshold indicates the minimum TCP segments per second. This threshold, used with the Moderate DG threshold, determines whether a system is receiving a high number of pings. | 10 |
| Moderate DG (ModerateDG) | This threshold indicates when the rate of datagrams per second are moderate. This should not be treated as a threshold indicating a high amount of datagrams, but as one that indicates an above normal threshold. | 60 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

## Segments ReXmit

This indication is sent when TCP segments have to be resent. This indicates congestion on the network with the TCP protocol. Indications from the Network Interface resource model can lend more insight.

The indication has the following attributes:

**SegmentsReXmitSec**
The number of segments retransmitted per second

**DGSec**
The number of datagrams transmitted per second

**SegmentsSentSec**
The number of segments being sent per second

**FragsToDGRatio**
Ratio of fragmented datagrams to total datagrams

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| High Segment Retransmitted (HighSegmenReXmit) | This threshold indicates the maximum TCP segments being retransmitted. | 1 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

# Thresholds

The following table lists the thresholds that can be set for the TCP/IP resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| High Fragment Ratio (HighFragRatio) | This threshold indicates a ratio between fragmented datagrams and datagrams received whole. It is used to determine when this value becomes too high. | 80 |
| High Segment Retransmitted (HighSegmenReXmit) | This threshold indicates the maximum TCP segments being retransmitted. | 1 |
| Low Segments (LowSegments) | This threshold indicates the minimum TCP segments per second. This threshold, used with the Moderate DG threshold, determines whether a system is receiving a high number of pings. | 10 |
| Moderate DG (ModerateDG) | This threshold indicates when the rate of datagrams per second are moderate. This should not be treated as a threshold indicating a high amount of datagrams, but as one that indicates an above normal threshold. | 60 |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource | Context | Properties |
|---|---|---|
| TCPIP | Segments Retransmitted | **TCP** The protocol<br>**SegmentsRetrasmitted** The number of segments retransmitted |
| | Segments Traffic | **TCP** Identifies the protocol<br>**SegmentSentSec** The number of segments sent per second<br>**SegmentsSec** The number of segments per second<br>**SegmentRcvdSec** The number of segments received per second |
| IP | Datagrams Traffic | **IP** Identifies the protocol<br>**DGSentSec** The number of datagrams sent per second<br>**DGSec** The number of datagrams per second<br>**DGReceivedSec** The number of datagrams received per second |
| | Fragments Received | **IP** Identifies the protocol<br>**FragmentsReceivedSec** The number of fragments received per second |

# Part II. Resource Models for UNIX/Linux

# 15

# CPU Resource Model

This chapter describes the CPU resource model for UNIX/Linux systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | DMXCpu |
| Category | UNIX/Linux |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 60 seconds |



## Overview

**Resource model distribution**

This resource model should be distributed to UNIX and Linux endpoints.

This resource model detects problems with the central processing unit (CPU) of a computer, for example, how long processes wait in the queue to be processed.

## Indications and Events

The following table lists the events that can be generated by the CPU resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| Low_IdleCPUUsage | High CPU Overload | Warning | 136 |
| High_SysCPUUsage | High CPU Usage by System | Warning | 137 |

## High CPU Overload

This indication is sent if a low percentage of CPU is idle.

The indication has the following attributes:

**percidlecpu**
The percent of CPU that is idle

**name** Identifies the single instance (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of CPU in Idle (IdleCPUTimeThr) | This threshold indicates the minimum percentage of CPU that should be idle to maintain satisfactory system performance. | 20 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 4 |
| Holes | 1 |

## High CPU Usage by System

This indication is sent when a high percentage of CPU is used for system requests.

The indication has the following attributes:

**percsyscpuusage**
The percent of CPU time spent on system requests

**name** Identifies the single instance (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of CPU Used by System (SysCPUTimeThr) | This threshold measures the percentage of CPU required for service system requests. | 80 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 4 |
| Holes | 1 |

# Thresholds

The following table lists the thresholds that can be set for the CPU resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| Percentage of CPU in Idle (IdleCPUTimeThr) | This threshold indicates the minimum percentage of CPU that should be idle to maintain satisfactory system performance. | 20 |
| Percentage of CPU Used by System (SysCPUTimeThr) | This threshold measures the percentage of CPU required for service system requests. | 80 |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource  | Context  | Properties  |
|---|---|---|
| CPU | Percent usage | **name** Identifies the single instance<br>**prcIdleTime** The percent of the time that the CPU is idle.<br>**prSysTime** The percent of the time that the CPU in system mode.<br>**prcUserTime** The percent of the time that the CPU is in user mode. |
| CPU | Average Loading | **name** Identifies the single instance<br>**loadAvg1** The number of processes running every minute<br>**loadAvg5** The number of processes running every 5 minutes<br>**loadAvg15** The number of processes running every 15 minutes |

# 16

# File Resource Model

This chapter describes the File resource model for UNIX/Linux systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | DMXFile |
| Category | UNIX/Linux |
| Thresholds | No |
| Parameters | Yes |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 120 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to UNIX and Linux endpoints.

The File resource model gives information about files in the system.

## Indications and Events

The following table lists the events that can be generated by the File resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| FileChanged | File Changed | Warning | 140 |
| FilesAttributeChange | File Attributes Changed | Warning | 141 |
| FileNotPresent | File not Present | Warning | 142 |

## File Changed

This indication is sent when the file checksum or the file modification time changes.

The indication has the following attributes:

**pathname**
The path name of the file being monitored (Key attribute)

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## File Attributes Changed

This indication is sent when the date and time of last updates to file are different from the previous cycle.

The indication has the following attributes:

**pathname**

The path name of the file being monitored (Key attribute)

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## File not Present

This indication is sent when the user tries to monitor a file that does not exist on the machine, for example, because it was deleted.

The indication has the following attributes:

**pathname**
The path name of the file being monitored (Key attribute)

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

# Parameters

The following table lists the parameters that can be set for the File resource model.

| Parameter | Description | Parameter Type |
|---|---|---|
| Files to Be Checked (FileList) | The files that you want to monitor. Enter the file details, separated by the pipe character, in the field to the left of the Add button in the following format:<br><br>path_filename | String list |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource  | Context  | Properties  |
|---|---|---|
| File | File Times | **pathname**<br>The file path<br>**modificationTime**<br>The time when the contents of file change<br>**changeTime**<br>The time when the file attributes change |
| | File Checksum | **fileName** The filename<br>**size** The filesize<br>**checksum**<br>The checksum that changes |

# 17

# File System Resource Model

This chapter describes the File System resource model for UNIX/Linux systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | DMXFileSystem |
| Category | UNIX/Linux |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 120 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to UNIX and Linux endpoints.

This resource model measures how efficiently the file systems are being used.

## Indications and Events

The following table lists the events that can be generated by the File System resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| LowKAvail | Low Space Available | Critical | 146 |
| FragmentedFileSystem | Fragmented File System | Minor | 147 |
| LowPercInodesAvail | Low Percentage of Available I-nodes | Warning | 148 |

## Low Space Available

This indication is sent when the available number of kilobytes on the specified file system goes below the threshold.

The indication has the following attributes:

**n_avail**
The number of kilobytes available on the specified disk drive.

**mountname**
The mount point on which the file system is mounted. (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Available Space (KB) (AvailableSpace) | This threshold indicates the minimum amount of available space (KB) that is required on the file system. | 7000 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 4 |
| Holes | 1 |

## Fragmented File System

This indication is sent when the number of used i-nodes is high but the used file system space is relatively low.

The indication has the following attributes:

**percinodeused**
The percentage of used i-nodes in the file system

**perckbused**
The percentage of space being used by the file system

**mountname**
The mount point on which the file system is mounted (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of File System Space Used (PrcUsedKspace) | This threshold indicates the minimum percentage of file system space that is being used. | 60 |
| Percentage of I-nodes Used (PrcUsedInodes) | This threshold indicates the maximum percentage of file system space that can be dedicated to i-nodes without causing high fragmentation of the file system. | 80 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 4 |
| Holes | 1 |

## Low Percentage of Available I-nodes

This indication is sent when the percentage of available i-nodes is below the threshold.

The indication has the following attributes:

**percavailinodes**
The percentage of available i-nodes

**mountname**
The mount point on which the file system is mounted (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of Available I-nodes (PrcAvailInodes) | This threshold indicates the minimum percentage of file system space available for i-nodes. | 20 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 4 |
| Holes | 1 |

# Thresholds

The following table lists the thresholds that can be set for the File System resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| Available Space (KB) (AvailableSpace) | This threshold indicates the minimum amount of available space (KB) that is required on the file system. | 7000 |
| Percentage of Available I-nodes (PrcAvailInodes) | This threshold indicates the minimum percentage of file system space available for i-nodes. | 20 |
| Percentage of File System Space Used (PrcUsedKspace) | This threshold indicates the minimum percentage of file system space that must be available. | 60 |
| Percentage of I-nodes Used (PrcUsedInodes) | This threshold indicates the maximum percentage of file system space that can be dedicated to i-nodes without causing high fragmentation of the file system. | 80 |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource  | Context  | Properties  |
|---|---|---|
| File System | File System Availability | **mountpoint**<br>The mount point on which the file system is mounted.<br>**percUsed**<br>The percent of file system space that is being used<br>**percInodesUsed**<br>The percent of i-nodes being used |

# 18

# Memory Resource Model



This chapter describes the Memory resource model for UNIX/Linux systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | DMXMemory |
| Category | UNIX-Linux |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 60 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to UNIX and Linux endpoints.

This resource model provides information about how the memory is used.

## Indications and Events

The following table lists the events that can be generated by the Memory resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| LowStorage | Low Storage Space | Critical | 152 |
| LowSwap | Low Swap Space | Critical | 153 |
| Thrashing | System Thrashing | Critical | 154 |

## Low Storage Space

This indication is sent when the percentage of available storage space in the system is lower than the threshold.

The indication has the following attributes:

**percavailstorage**
The percentage of available storage space

**totavailstorage**
The available storage space in kilobytes

**name** Identifies the single instance (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of Available Virtual Storage (AvailVirtualStorage) | This threshold measures the percentage of total virtual storage that is available. | 40 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 20 |
| Holes | 0 |

## Low Swap Space

This indication is sent when the percentage of swap space available on the system goes below the threshold.

The indication has the following attributes:

**percavailswap**
The percentage of available swap space

**totavailswap**
The available swap space in kilobytes

**name** Identifies the single instance (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of Available Swap Space (SwapSpacePrc) | This threshold indicates the minimum percentage of all swap space that must be available. | 30 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 4 |
| Holes | 0 |

## System Thrashing

This indication is sent when there is excessive or unusual system paging, for example, when the page-in or page-out rate exceeds the threshold.

The indication has the following attributes:

**pageins** The rate of paging in

**pageouts** The rate of paging out

**name** Identifies the single instance (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Memory Page-in Rate (PageInRate) | This threshold indicates the maximum occurrences of paging in per second, as averaged over the cycle. | 400 |
| Memory Page-out Rate (PageOutRate) | This threshold indicates the maximum occurrences of paging out per second, as averaged over the cycle. | 400 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 0 |

# Thresholds

The following table lists the thresholds that can be set for the Memory resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| Memory Page-in Rate (PageInRate) | This threshold indicates the maximum occurrences of paging in per second, as averaged over the cycle. | 400 |
| Memory Page-out Rate (PageOutRate) | This threshold indicates the maximum occurrences of paging out per second, as averaged over the cycle. | 400 |
| Percentage of Available Swap Space (SwapSpacePrc) | This threshold indicates the minimum percentage of all swap space that must be available. | 30 |
| Percentage of Available Virtual Storage (AvailVirtualStorage) | This threshold measures the percentage of total virtual storage that is available. | 40 |



# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource | Context | Properties |
|---|---|---|
| Memory | Memory Paging | **name** Identifies the single instance<br>**pageInsRate** The rate of paging in for the cycle<br>**pageOutsRate** The rate of paging out for the cycle |
| Memory | Memory Availability | **name** Identifies the single instance<br>**PrcAvailSwap** The percent of swap space that is available<br>**PrcAvailStorage** The percentage of storage space that is available |

# 19

# Network Interface Resource Model

This chapter describes the Network Interface resource model for UNIX/Linux systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | DMXNetworkInterface |
| Category | UNIX/Linux |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 150 seconds |



## Overview

**Resource model distribution**

This resource model should be distributed to UNIX and Linux endpoints.

The Network Interface resource model detects problems with the following installed network interfaces

- Loopback
- Ethernet
- Token-Ring

Events are generated when performance data, such as bytes per second in and out and sessions with errors or requests, become critical.

## Indications and Events

The following table lists the events that can be generated by the resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| HighInputErrPacks | High Input Packets in Error | Warning | 158 |
| HighOutErrorPacks | High Output Packets in Error | Warning | 159 |

| Event | Indication | Severity | Page |
|---|---|---|---|
| HighPacktsCollision | High Percentage Packet Collisions | Critical | 160 |
| InterfaceNotEnabled | Interface Not Enabled | Warning | 161 |
| InterfaceNotOperat | Interface Not Operational | Warning | 162 |
| IntNotSupported | Network Interface Card Not Supported | Warning | 163 |
| IntStatUnknown | Unknown Interface Status | Critical | 164 |

## High Input Packets in Error

This indication is sent when a high percentage of input packets are going into error.

The indication has the following attributes:

**totinpacks**
The total input packets

**percinerrpacks**
The percentage of input packets that are going into error

**interfacename**
The network interface (Key attribute)

**interfacetype**
The network interface type

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| Input Packets in Error (PercInPacketErrThr) | This threshold indicates the maximum percentage of input packets in error allowed. | 20 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

## High Output Packets in Error

This indication is sent when a high percentage of output packets are going into error.

The indication has the following attributes:

**totoutpacks**
The total output packets

**percouterrpacks**
The percentage of output packets that are going into error

**interfacename**
The network interface (Key attribute)

**interfacetype**
The network interface type

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| Output Packets in Error (PercOutPacketErrThr) | This threshold indicates the maximum percentage of output packets in error allowed. | 10 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

## High Percentage Packet Collisions

This indication is sent when a high percentage of output packets are colliding. This indicates that the packets being sent might be corrupted or that there might be a network overload.

The indication has the following attributes:

**totoutpacks**
The total output packets

**percPacktsCollision**
The percentage of packets that are colliding

**interfacename**
The interface name (Key attribute)

**interfacetype**
The interface card type

The indication is generated only if the system uses an Ethernet card and when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| Packet Collision Percentage (PercPacketCollisionThr) | This threshold indicates the maximum percentage of corrupted packets allowed. | 10 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 2 |
| Holes | 0 |

## Interface Not Enabled

This indication is sent when the interface status is not up

The indication has the following attributes:

**interfacename**
The name of the network interface (Key attribute)

**interfacetype**
The type of the interface card

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Interface Not Operational

This indication is sent when the interface has been enabled, but is not in a running state. The interface driver might not be correctly installed.

The indication has the following attributes:

**interfacename**
The name of the network interface (Key attribute)

**interfacetype**
The type of interface card

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Network Interface Card Not Supported

This indication is sent when the network interface card is not supported.

The indication has the following attributes:

**interfacename**
The name of the network interface (Key attribute)

**intercafetype**
The type of interface card

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Unknown Interface Status

This indication is sent when the status of the monitored interface can not be retrieved.

The indication has the following attributes:

**interfacename**
The name of the network interface (Key attribute)

**intercafetype**
The type of interface card

**interfacestatus**
The status of the interface card

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

# Thresholds

The following table lists the thresholds that can be set for the Network Interface resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| Input Packets in Error (PercInPacketErrThr) | This threshold indicates the maximum percentage of input packets in error allowed. | 20 |
| Output Packets in Error (PercOutPacketErrThr) | This threshold indicates the maximum percentage of output packets in error allowed. | 10 |
| Packet Collision Percentage (PercPacketCollisionThr) | This threshold indicates the maximum percentage of corrupted packets allowed. | 10 |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource  | Context  | Properties  |
|---|---|---|
| Network | Interface Card | **Interfacename** The name of the interface card<br>**InPacks** The total number of input packets<br>**InPacksErr** The percent of input packets in error<br>**OutPacks** The total number of output packets<br>**OutPackErr** The percent of output packets in error<br>**OutPackColl** The percent of output packets colliding |

# 20

# Network RPC-NFS Resource Model

This chapter describes the Network RPC-NFS resource model for Solaris systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | DMXNetworkRPCNFS |
| Category | Solaris |
| Thresholds | Yes |
| Parameters | No |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 150 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to Solaris endpoints.

The Network RPC-NFS resource model detects problems and monitors the performance of the RPC and NFS servers and clients.

## Indications and Events

The following table lists the events that can be generated by the resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| HighNFSBufferSize | High NFS Buffer Size | Warning | 168 |
| HighNFSSrvGetattr | High NFS Server Get-attribute Operations | Warning | 169 |
| HighNFSSrvRead | High NFS Server Read Operations | Warning | 169 |
| HighNFSSrvReadLink | High NFS Server Readlink Operations | Warning | 171 |
| HighNFSSrvWrites | High NFS Server Write Operations | Warning | 172 |
| HighPercDupReqs | High Duplicate RPC Server Calls | Warning | 173 |
| HighPercRetrans | High Retransmitted Calls | Critical | 174 |
| HighPercRPCBadCalls | High RPC Bad Calls | Warning | 175 |

| Event | Indication | Severity | Page |
|---|---|---|---|
| HighTimeoutsAnd_Badxids | High Timeouts and Badxids | Warning | 176 |
| NetworkBusy | High Network Traffic | Warning | 177 |
| NetworkSlow | Slow Network | Warning | 178 |

## High NFS Buffer Size

This indication is sent when the percentage of client RPC calls that were retransmitted or timed out is too high, and there are no duplicate acknowledgements.

The indication has the following attributes:

**totrpccalls**
The total client RPC calls

**percrpcretrans**
The percentage of client RPC calls that were retransmitted

**percrpctimeouts**
The percentage of client RPC calls that timed out

**name** Identifies the single instance (Key attribute)

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of Client RPC Calls in Time-out (RPCCLCallsToutPercThr) | This threshold indicates the maximum percentage of timed-out client RPC calls allowed to maintain acceptable network traffic level. | 5 |
| Percentage of Client RPC Retransmissions (RPCCLRetransPercThr) | This threshold indicates the maximum of client RPC retransmissions allowed. | 5 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

## High NFS Server Get-attribute Operations

This indication is sent when the percentage of NFS server calls to read the client attribute cache is too high.

The indication has the following attributes:

**totnfssrvcalls**
The total NFS server calls

**percnfssrvgetattr**
The percentage of NFS server requests to read the client attribute cache

**name** Identifies the single instance (Key attribute)

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of NFS Server Getattr Operations (NFSSrvGettatrThr) | This threshold indicates the maximum percentage of NFS server calls to read the client cache attribute without impacting NFS server performance. | 40 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

## High NFS Server Read Operations

This indication is sent when the percentage of NFS server calls for read operations is too high.

The indication has the following attributes:

**totnfssrvcalls**
The total NFS server calls

**percnfssrvread**
The percentage of NFS server calls for read operations

**name** Identifies the single instance (Key attribute)

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of NFS Server Read Operations (NFSSrvReadThr) | This threshold indicates the maximum percentage of NFS server read operations allowed. | 30 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

## High NFS Server Readlink Operations

This indication is sent when the percentage of NFS server calls for readlink operations compared to the total lookup calls on the NFS server is too high.

The indication has the following attributes:

**totnfssrvlookup**
The total lookup calls on the NFS server

**percnfssrvreadlink**
The percentage of NFS server calls for readlink operations

**name** Identifies the single instance (Key attribute)

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of NFS Server Readlink Operations (NFSSrvReadlinkThr) | This threshold indicates the maximum percentage of NFS server read operations using symbolic links on file systems exported by the server. | 10 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

## High NFS Server Write Operations

This indication is sent when the percentage of NFS server calls for write operations is too high.

The indication has the following attributes:

**totnfssrvcalls**
The total number of NFS server calls

**percnfssrvwrites**
The percentage of NFS server calls for write operations

**name** Identifies the single instance (Key attribute)

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of NFS Server Write Operations (NFSSrvWriteThr) | This threshold indicates the maximum percentage of NFS server write operations allowed without impacting NFS performance. | 15 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

## High Duplicate RPC Server Calls

This indication is sent when the percentage of RPC server calls that are duplicates is too high. This can be due to transmission problems.

The indication has the following attributes:

**totrpcsrvdupreqs**
The total number of RPC server calls that are duplicate requests

**percrpcsrvdupreqs**
The percentage of RPC server calls that are duplicate requests

**name** Identifies the single instance (Key attribute)

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of Server RPC Duplicate Requests (RPCSrvDupCallsThr) | This threshold indicates the maximum percentage of duplicate RPC calls allowed. | 60 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

## High Retransmitted Calls

This indication is sent when the percentage of client RPC calls being retransmitted is too high.

The indication has the following attributes:

**totrpccalls**
The total number of RPC client calls

**percrpcretrans**
The percentage of RPC client calls being retransmitted

**name** Identifies the single instance (Key attribute)

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of Client RPC Retransmissions (RPCCLRetransPercThr) | This threshold indicates the maximum of client RPC retransmissions allowed. | 5 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 5 |
| Holes | 0 |

## High RPC Bad Calls

This indication is sent when the percentage of RPC server or client calls being rejected by the RPC is too high.

The indication has the following attributes:

**totrpcsrvcalls**
The total number of RPC server calls

**percrpcsrvbadcalls**
The percentage of RPC server calls being rejected

**totrpcclcalls**
The total number of RPC client calls

**percrpcclbadcalls**
The percentage of RPC client calls being rejected

**name** Identifies the single instance (Key attribute)

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of Bad RPC Calls (RPCBadCallsPercThr) | This threshold indicates the maximum number of bad RPC calls allowed to maintain acceptable network performance. It refers to both client and server RPC bad calls. | 30 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 5 |
| Holes | 0 |

## High Timeouts and Badxids

This indication is sent when the percentages of timed-out RPC client calls and badxids (duplicate acknowledgements) are too high.

The indication has the following attributes:

**totrpccalls**
The total number of RPC client calls

**percrpctimeouts**
The percentage of timed-out RPC client calls

**percrpcbadxids**
The percentage of badxids

**name** Identifies the single instance (Key attribute)

The indication is generated when the following thresholds are exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of Client RPC Badxids (RPCCLBAdXidsPercThr) | This threshold indicates the maximum percentage of badxids (duplicate acknowledgements) allowed by RPC client maintaining acceptable network traffic level. | 5 |
| Percentage of Client RPC Calls in Time-out (RPCCLCallsToutPercThr) | This threshold indicates the maximum percentage of timed-out client RPC calls allowed to maintain acceptable network traffic level. | 5 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

## High Network Traffic

This indication is sent when the number of client RPC calls being retransmitted is close to the client RPC calls timing out.

The indication has the following attributes:

**numrpcretrans**
The number of RPC retransmissions

**numrpctimeouts**
The number of timed-out RPC client calls

**name** Identifies the single instance (Key attribute)

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

## Slow Network

This indication is sent when the number of timed-out client RPC calls approaches the number of client RPC duplicate acknowledgements (badxids) for a single NFS request.

The indication has the following attributes:

**numrpcbadcalls**
The number of timed-out client RPC calls

**numrpcdupacks**
The number of client RPC duplicate acknowledgements (badxids) for a single NFS request

**name** Identifies the single instance (Key attribute)

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

# Thresholds

The following table lists the thresholds that can be set for the Network resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| Percentage of Bad RPC Calls (RPCBadCallsPercThr) | This threshold indicates the maximum number of bad RPC calls allowed to maintain acceptable network performance. It refers to both client and server RPC bad calls. | 30 |
| Percentage of Client RPC Badxids (RPCCLBAdXidsPercThr) | This threshold indicates the maximum percentage of badxids (duplicate acknowledgements) allowed by RPC client maintaining acceptable network traffic level. | 5 |
| Percentage of Client RPC Calls in Time-out (RPCCLCallsToutPercThr) | This threshold indicates the maximum percentage of timed-out client RPC calls allowed to maintain acceptable network traffic level. | 5 |
| Percentage of Client RPC Retransmissions (RPCCLRetransPercThr) | This threshold indicates the maximum of client RPC retransmissions allowed. | 5 |
| Percentage of NFS Server Getattr Operations (NFSSrvGettatrThr) | This threshold indicates the maximum percentage of NFS server calls to read the client cache attribute without impacting NFS server performance. | 40 |
| Percentage of NFS Server Readlink Operations (NFSSrvReadlinkThr) | This threshold indicates the maximum percentage of NFS server read operations using symbolic links on file systems exported by the server. | 10 |
| Percentage of NFS Server Read Operations (NFSSrvReadThr) | This threshold indicates the maximum percentage of NFS server read operations allowed. | 30 |
| Percentage of NFS Server Write Operations (NFSSrvWriteThr) | This threshold indicates the maximum percentage of NFS server write operations allowed without impacting NFS performance. | 15 |
| Percentage of Server RPC Duplicate Requests (RPCSrvDupCallsThr) | This threshold indicates the maximum percentage of duplicate RPC calls allowed. | 60 |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource | Context | Properties |
|---|---|---|
| NFS | Server | **name** Identifies the single instance<br>**NFScalls** The NFS calls<br>**NFSbadcalls** The NFS calls that are timed out<br>**NFSgetattr** The NFS requests to read the client attribute cache<br>**NFSread** The NFS read operations<br>**NFSwrite** The NFS write operations<br>**NFSlookup** The lookup calls on the NFS server<br>**NFSreadlink** The NFS server calls for readlink operations |
| | Client | **name** Identifies the single instance<br>**NFScalls** The NFS calls<br>**NFSbadcalls** The NFS calls that are timed out |
| RPC | Client | **name** Identifies the single instance<br>**RPCcalls** The client RPC calls<br>**RPCbadcalls** The number of timed-out client RPC calls<br>**RPCretrans** The RPC client calls being retransmitted<br>**RPCbadxids** The badxids<br>**RPCtimeouts** The client RPC calls that timed out |
| | Server | **name** Identifies the single instance<br>**RPCCalls** The client RPC calls<br>**RPCBadcalls** The number of timed-out client RPC calls<br>**RPCDuprequests** The RPC server calls that are duplicate requests<br>**RPCDupchecks** The number of RPC server calls that are looked up in the duplicate request cache. |

# 21

# Process Resource Model

This chapter describes the Process resource model for UNIX systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | DMXProcess |
| Category | UNIX/Linux |
| Thresholds | Yes |
| Parameters | Yes |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 60 seconds |

## Overview

**Resource model distribution**

This resource model should be distributed to UNIX and Linux endpoints.

The Process resource model looks for bottlenecks in running processes. Problems highlight include:

- A process uses too much CPU time
- Too many zombie processes in the system
- A process is stopped or killed
- A process that was requested does not exist

# Indications and Events

The following table lists the events that can be generated by the Process resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| HighZombieProcesses | High Number of Zombie Processes | Warning | 182 |
| ProcessHighCPU | Process Consuming High CPU | Minor | 183 |
| ProcessKilledOrNotExisting | Process Killed or Nonexistent | Critical | 184 |
| ProcessStopped | Process Stopped | Critical | 185 |

## High Number of Zombie Processes

This indication is sent when the number of processes in a zombie state is too high. A process is a zombie when it has been terminated and its results have not been gathered by the parent process.

The indication has the following attributes:

**numZombie**
The number of zombie processes

**name** Identifies the single instance (Key attribute)

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| Maximum Number of Zombie Processes (HighZombieProcess) | This threshold indicates the maximum number of zombie processes allowed without impacting system performance. | 20 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 6 |
| Holes | 1 |

## Process Consuming High CPU

This indication is sent when a process uses too much CPU time.

The indication has the following attributes:

**IDProcess** The ID of the process being monitored (Key attribute)

**PrcProcessorTime**
The percentage of CPU time used for the process

**Process** The process being monitored (Key attribute)

**state** The status of the process

The indication is generated when the following threshold is exceeded:

| Threshold | Description | Default |
|---|---|---|
| Percentage of CPU Used (HighCPUUsed) | This threshold indicates the maximum percentage of CPU that can be used by a single process without impacting system performance. When a process uses a high percentage of CPU, it reduces available CPU time for other processes. | 60 |

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 20 |
| Holes | 3 |

## Process Killed or Nonexistent

This event is sent when a process is stopped using the command **sh ... stop**, or when the resource model detects that the process is no longer present in the system. Specify any processes in the Processes parameters.

The indication has the following attributes:

**name** The process being monitored (Key attribute)

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Process Stopped

This event is generated when the process has the state *stopped*. The process continues to be present in the system, and can be monitored using the **ps -ef** command. Specify any processes in the Processes parameters.

The indication has the following attributes:

**ProcessID**
The ID of the process being monitored (Key attribute)

**ParentProcessID**
The parent of the process being monitored

**name** The process being monitored (Key attribute)

**ProcessStatus**
The status of the process being monitored

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 3 |
| Holes | 0 |

# Thresholds

The following table lists the thresholds that can be set for the Process resource model. For each threshold it shows the name, a short description, and the default value:

| Threshold | Description | Default |
|---|---|---|
| Maximum Number of Zombie Processes (HighZombieProcess) | This threshold indicates the maximum number of zombie processes allowed without impacting system performance. | 20 |
| Percentage of CPU Used (HighCPUUsed) | This threshold indicates the maximum percentage of CPU that can be used by a single process without impacting system performance. When a process uses a high percentage of CPU, it reduces available CPU time for other processes. | 60 |

# Parameters

The following table lists the parameter that can be set for the Process resource model.

| Parameter | Description | Default |
|---|---|---|
| Processes (processes) | The single processes to be monitored for the Process Stopped and Process Killed or Nonexistent events. | String list |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource | Context | Properties |
|---|---|---|
| Process | Processor Usage | **Process** The name of the process<br>**PID** The ID of the process<br>**PercentProcessUsage** The percentage of CPU that the process is using |

# 22

# Security Resource Model



This chapter describes the Security resource model for UNIX/Linux systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | DMXSecurity |
| Category | UNIX/Linux |
| Thresholds | No |
| Parameters | Yes |
| Built-in actions | No |
| Clearing events | Yes |
| Default cycle time | 120 seconds |

## Overview

**Resource model distribution**
This resource model should be distributed to UNIX and Linux endpoints.

The Security resource model provides information about files and the users logged onto the system. It highlights the following items or changes that might indicate security breaches:

- Property changes, such as the owner, group, or attributes, for certain files
- The number of logons onto the system by the same user
- A suspect superuser
- An account that is not valid for root

## Indications and Events

The following table lists the events that can be generated by the Security resource model, the name of the indication from which each event is generated, the default severity of the event, and where you can find a detailed description of the indication:

| Event | Indication | Severity | Page |
|---|---|---|---|
| DuplicatedAccount | Duplicate Account | Minor | 188 |
| FileNotExisting | Nonexistent File | Warning | 189 |
| HighLoggingNumber | High Log-in Number for User | Minor | 189 |

| Event | Indication | Severity | Page |
|---|---|---|---|
| IllegalGroup | Illegal Group | Critical | 190 |
| IllegalOwner | Illegal Owner | Critical | 190 |
| NotRegularRootAccount | Account not Valid for Root | Harmless | 191 |
| PasswdNull | Null Password | Critical | 191 |
| SuspectSuperGroup | Suspect Supergroup | Critical | 192 |
| SuspectSuperUser | Suspect Superuser | Critical | 192 |
| WrongMode | Wrong File Mode | Critical | 193 |

## Duplicate Account

This indication is sent when two users or groups have the same user or group ID.

The indication has the following attribute:

**id** The ID that has been duplicated (Key attribute)

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 4 |
| Holes | 0 |

## Nonexistent File

This indication is sent when the user wants to monitor a file that does not exist on the machine, for example, because it was deleted.

The indication has the following attribute:

**pathname**
The pathname of the file specified (Key attribute)

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## High Log-in Number for User

This indication is sent when a user is logged in too many times.

The indication has the following attributes:

**numlogged** The number of times the user has logged in

**username** The ID of the user (Key attribute)

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 4 |
| Holes | 1 |

## Illegal Group

This indication is sent when a security-sensitive file does not have a superuser group.

The indication has the following attributes:

**fullname** The full file path (Key attribute)

**group** The file group being monitored

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Illegal Owner

This indication is sent when a security-sensitive file has an owner who is not a superuser.

The indication has the following attributes:

**owner** The owner of the file

**fullname** The full file path (Key attribute)

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Account not Valid for Root

This indication is sent when the account ID is not valid for root.

The indication has the following attributes:

**id** The ID of the root account that was created (Key attribute)

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Null Password

This indication is sent when a user and/or a group has a password that is set to null.

The indication has the following attributes:

**id** The ID of the user or group (Key attribute)

**name** Identifies the single instance

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Suspect Supergroup

This indication is sent when the supergroup is suspected of being an intruder.

The indication has the following attributes:

**groupid**
The ID of the suspicious group

**groupname**
The group name of the file that is suspicious (Key attribute)

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Suspect Superuser

This indication is sent when superuser is suspected of being an intruder.

The indication has the following attributes:

**id** The user ID that is suspicious

**userName** The user name (Key attribute)

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

## Wrong File Mode

This indication is sent when security-sensitive files have a file mode that is different from the default mode.

The indication has the following attributes:

**actualmode** The type of access permission that the file actually has on the system.

**rightmode** The type of access permission that the file should have in the parameters as defined from the IBM Tivoli Monitoring desktop.

**filename** The name of the file (Key attribute).

The following table shows the default settings for this indication:

| Setting | Default |
|---|---|
| Send indications to Tivoli Enterprise Console | Yes |
| Send indications to Tivoli Business Systems Manager | No |
| Occurrences | 1 |
| Holes | 0 |

# Parameters

The following table lists the parameters that can be set for the Security resource model.

| Parameter | Description | Parameter Type |
|---|---|---|
| Alternative Groups (AlternativeGroup) | Groups other than root that can own certain security files in the system | String list |
| Alternative Owners (AlternativeOwners) | Users that can own certain security files in the system | String list |
| Defined Users (Users) | Users to be monitored in the following format:<br><br>*username* \| *number_logins* | String list |
| Files to Be Monitored (FilesList) | The files that you want to monitor. Enter the file details, separated by the pipe character, in the field to the left of the Add button in the following format:<br><br>*filename* \| *string_mode* (as listed with the ls -la command) \| *owner* \| *group* | String list |
| Special Groups (Supergroups) | The groups that can become supergroups in the system | String list |
| Special Users (Superusers) | Users that can become superusers in the system | String list |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource | Context | Properties |
|---|---|---|
| User | Logging | **userName** The name of the user logging on<br>**id** The ID of the suspicious user<br>**numLogged** The number of times the user is logged in |
| File | File Usage | **fileName** The file name being monitored<br>**group** The file group being monitored<br>**owner** The file owner being monitored<br>**mode** The type of access permission to the file<br>**size** The file size being monitored |

# 23

# Server Performance Prediction Resource Model

This chapter describes the Server Performance Prediction resource model for UNIX/Linux systems.

The following table shows the key characteristics of this resource model:

| Resource Model at a Glance | |
|---|---|
| Internal name | DMXSpp |
| Category | UNIX/Linux |
| Thresholds | No |
| Parameters | Yes |
| Built-in actions | No |
| Clearing events | No |
| Default cycle time | 10 seconds |



## Overview

**Resource model distribution**

This resource model should be distributed to UNIX and Linux endpoints.

The Server Performance Prediction resource model gathers data from UNIX and Linux systems for the Tivoli Decision Support Guide for Server Performance Prediction (Advanced Edition). This resource model does not have thresholds, indications or events.

## Parameters

The following table lists the parameters that can be set for the Server Performance Prediction resource model.

| Parameter | Description | Default |
|---|---|---|
| Network Interface Configuration (NetworkInterfacesConf) | The network interface for which you want the network interface values calculated. Values: 1-4. | Choice list |
| Available File System Space Configuration (FileSystemSpaceAvailableConf) | The instance for which you want the available file system space calculated. Values: /, /usr, /tmp, /var, /home. | String list |

# Logging

The following table shows the resource, context and properties for which data can be logged:

| Resource | Context | Properties |
|---|---|---|
| Cpu | SPP Guide for UNIX | **idleTime** The idle cpu time as a percentage of total cpu time<br>**userTime** The cpu time spent on user applications as a percentage of total cpu time<br>**sysTime** The cpu time spent by the system as a percentage of total cpu time<br>**loadAvg1** The average number of processes running in the most recent calendar minute |
| FileSystem | SPP Guide for UNIX | **mountPoint** The directory on which the file system is mounted<br>**availKBytes** The free kilobytes on the disk<br>**prcInodeUsed** The percentage of Inodes used in the file system |
| Memory | SPP Guide for UNIX | **availSwapSpace** The amount of space available to be swapped<br>**pageInsRate** The number of page inputs per second<br>**pageOutsRate** The number of page outputs per second<br>**pctusedVirtualStorage** The percentage of total memory (including cache memory and swap space) being used |
| NetworkInterface | SPP Guide for UNIX | **deltaInPackets** The delta of packets received, compared with the previous cycle<br>**deltaInPacketsErr** The delta of errors for packets received, compared with the previous cycle<br>**deltaOutPackets** The delta of packets sent, compared with the previous cycle<br>**deltaOutPacketsErr** The delta of errors for packets sent, compared with the previous cycle<br>**deltaCollisions** The delta of packet collisions, compared with the previous cycle<br>**networkInterface** The network interface card |
| NetworkRPCNFS | SPP Guide for UNIX | **deltaNFSClientCalls** The number of calls from the NFS client, compared with the previous cycle. |

| Resource  | Context  | Properties  |
|---|---|---|
| Processes | SPP Guide for UNIX | **numberWaitProcesses**<br>The number of processes waiting to run<br>**numberOfProcesses**<br>The total number of processes |

# Part III. Appendixes

# A

# Return Codes for Built-in Actions (Windows)

This appendix describes the return codes for built-in actions for resource models running on Windows systems.

The following table lists the return codes by category:

| Return Code Category | See Page |
|---|---|
| Common to all built-in actions | 201 |
| Built-in actions of the Event Log resource model only | 202 |
| Built-in actions of the Network Interface Card resource model only | 203 |
| Built-in actions of the Service resource model only | 204 |

## All Built-in Actions

The return codes common to all built-in actions are as follows:

**RETURN CODE 99**
It was not possible to get the definition of the output parameters of the called method.

**RETURN CODE 98**
It was not possible to spawn an instance of the output parameters of the called method. This might be because there is not enough memory.

**RETURN CODE 97**
The called method is not implemented by this method provider.

**RETURN CODE 9**
The resource model does not implement this method.

# Built-in Actions of the Event Log Resource Model

The Event Log resource model has the following built-in actions:

- EnlargeIRPStackSize action
- DeleteRegistry action
- RaiseSessTimeOut action

## Enlarge IRP Stack Size Action

The return codes specific to the EnlargeIRPStackSize built-in action of the Event Log resource model are as follows:

**RETURN CODE 0**
The action has been successfully executed.

**RETURN CODE 1**
An error occurred while accessing the system registry.

**RETURN CODE 2**
The registry key: SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters\IRPStackSize was not updated because its value is already greater than 12 (its suggested value).

## Delete Registry Action

The return codes specific to the DeleteRegistry built-in action of the Event Log resource model are as follows:

**RETURN CODE 0**
The action has been successfully executed.

## Raise Session Time Out Action

The return codes specific to the RaiseSessTimeOut built-in action of the Event Log resource model are as follows:

**RETURN CODE 0**
The action has been successfully executed.

**RETURN CODE 1**
The registry key: SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters\SessTimeOut was not updated because its value is already greater than 70 (its maximum suggested value).

**RETURN CODE 2**
The registry key: SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters\SessTimeOut was not found in the registry, it has been created and its value has been set to 55.

# Built-in Actions of the Network Interface Card Resource Model

The Network Interface Card resource model has the following built-in actions:

- AdjustInitWorkItems action
- AjustMaxWorkItems action

## Adjust Initial Work Items Action

The return codes specific to the AdjustInitWorkItems built-in action of the Network Interface Card resource model are as follows:

**RETURN CODE 0**
The action has been successfully executed.

**RETURN CODE 1**
The registry key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet \Services\LanmanServer\Parameters\InitWorkItems was set to a value that is already greater than the suggested value.

## Adjust Maximum Work Items Action

The return codes specific to the AdjustMaxWorkItems built-in action of the Network Interface Card resource model are as follows:

**RETURN CODE 0**
The action has been successfully executed.

**RETURN CODE 1**
The registry key: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet \Services\LanmanServer\Parameters\MaxWorkItems was set to a value that is already greater than the suggested value.

# Built-in Actions of the Service Resource Model

The Service resource model has only one built-in action, the RestartService action.

## Restart Service Action

The return codes specific to the RestartService built-in action of the Service resource model are as follows:

**RETURN CODE 0**
The action has been successfully executed.

**RETURN CODE 1**
An error occurred while retrieving the input parameters.

**RETURN CODE 2**
Wrong type of Input Parameters.

**RETURN CODE 3**
Unable to open the service control manager.

**RETURN CODE 4**
Unable to open the service. This might be because the service does not exist on the target machine.

**RETURN CODE 5**
Unable to query the status of the service.

**RETURN CODE 6**
It was not possible to start the service. This might be because some service prerequisites have not been met.

**RETURN CODE 7**
It was not possible to start the service. This might be because the service is corrupted.

**RETURN CODE 8**
It was not possible to restart the service. This might be because some service prerequisites have not been met.

# B

# Correlated Events (Windows)

This appendix describes all correlated events and the indications from which each one is generated. Correlated events are only generated for Windows resource models.

The following table lists the correlated events:

# Busy Drive from High Paging

The *TMW_BusyDriveFromPaging* indication is generated from the following indications:

- *TMW_HighLogicalPercentDiskTime* indication from the Logical Disk resource model
- *TMW_HighPaging* indication from the Memory resource model

This event is generated when it is determined that the local logical disk is busy which is most likely being caused by excessive hard page faulting.

*TMW_BusyDriveFromPaging* has the following event properties:

| | |
|---|---|
| **Disk** | The logical disk being analyzed |
| **PercentDiskTime** | The percentage usage of the logical drive |
| **PagesSec** | The number of pages per second |
| **TotalAvail** | The total of available memory in bytes |
| **TotalWorkingSet** | The total working set size in bytes |
| **TotalCache** | The total cache size in bytes |
| **Severity** | The default value is warning |

# Busy Drive from Low Avail

The *TMW_BusyDriveFromLowAvail* indication is generated from the following indications:

- *TMW_LowAvailCausingManyProblems* indication from the Memory resource model
- *TMW_HighLogicalPercentDiskTime* indication from the Logical Disk resource model

This event is a relatively serious event because it indicates not only that the local disk is unusually busy, but also that memory is low, paging is high, and the page file is changing. This combination will ultimately cause a core dump if left unchecked.

*TMW_BusyDriveFromLowAvail* has the following event properties:

| | |
|---|---|
| **Disk** | The logical disk that is being analyzed |
| **PercentDiskTime** | The percentage usage of the logical drive |
| **TotalAvail** | The total of available memory in bytes |
| **PageFaultsSec** | The current value for page faults per second |
| **PagesSec** | The current value of pages per second |
| **CommittedBytes** | The current total of committed bytes |
| **CommittedLimit** | The upper limit of the committed bytes |
| **Severity** | The default value is critical |

# Congested TCP Network

The *TMW_CongestedTCPNetwork* indication is generated from the following indications:

- *TMW_NICOverworked* indication from the NIC resource model
- *TMW_SegmentsReXmit* indication from the TCP/IP resource model

This indicates that the TCP/IP network segment that the originating machine is on may be congested. *Congested TCP Network* is generated when an indication of an overworked NIC and a high amount of retransmitted segments are sent.

*TMW_CongestedTCPNetwork* has the following event properties:

| | |
|---|---|
| **NetworkInterface** | The identity of network interface card being examined |
| **SegmentsReXmitSec** | The number of segments retransmitted per second |
| **DGSec** | The number of datagrams transmitted per second |
| **SegmentsSec** | The number of segments being sent per second |
| **FragsToDGRatio** | Ratio of fragmented datagrams to total datagrams |
| **OuputQueueLength** | The length of the output queue |
| **NICBPS** | The bytes per second transferred through the network interface card |
| **CurrentBandWidth** | The bandwidth of the network interface card |
| **Severity** | The default value is warning |

# Critically Low Disk Space

The *TMW_CriticallyLowDiskSpace* indication is generated from the following indications:

- *TMW_LowAvailCausingSoftPagePagefileResize* indication from the Memory resource model
- *TMW_LowLogicalDiskSpace* indication from the Logical Disk resource model

This indicates that disk space is low, and it may soon come into demand. If the system needs additional space to resize a pagefile, and the logical disk where the pagefile resides is low on space, a core dump may occur.

*TMW_CriticallyLowDiskSpace* has the following event properties:

| | |
|---|---|
| **Disk** | The logical disk that is being analyzed |
| **PercentFreeSpace** | The percentage amount of free space on the logical drive |
| **FreeMB** | The actual size of free space on the logical drive in Megabytes |
| **CommittedBytes** | The current total of Committed Bytes |
| **CommittedLimit** | The upper limit of the Committed Bytes |
| **TotalAvail** | The total of available memory in bytes |
| **PageFaultsSec** | The current value for page faults per second |
| **Severity** | The default value is critical |

# Critical Memory Leak

The *TMW_CriticalMemoryLeakInWS* indication is generated from the following indications:

- *TMW_MemoryLeakInPB* indication from the Memory resource model
- *TMW_LowAvailHighWS* indication from the Memory resource model

A memory leak together with an indication of low available memory due to a high working set will generate this event. This indicates that not only is memory low, but it is also being consumed and will soon run out.

*TMW_CriticalMemoryLeakInWS* has the following event properties:

| | |
|---|---|
| **LeakyProcess** | The process with the memory leak |
| **IDLeakyProcess** | The numeric ID of the process with the memory leak |
| **CurrentWorkingSet** | The current working set of the process |
| **CurrentPrivateBytes** | The current private bytes of the process |
| **TotalAvail** | The total of available memory in bytes |
| **TotalWorkingSet** | The total working set size in bytes |
| **TotalCache** | The total cache size in bytes |
| **PercentAvail** | The percentage of available memory in comparison to the size of the working set and the cache |
| **PercentWS** | The percentage of working set memory in comparison to the size of available memory and the cache |
| **PercentCache** | The percentage of cache memory in comparison to the size of the working set and the available memory |
| **IDHighWSProcess** | The numeric ID of the process with the highest working set |
| **NumProcesses** | The total number of processes |
| **HighWSProcess** | The working set of the process with the highest working set |
| **Severity** | The default value is critical |

# Faulty Disk Subsystem

The *TMW_FaultyDiskSubsystem* indication is generated from the following indications:

- *TMW_HighPhysicalPercentDiskTime* indication from the Physical Disk resource model
- *TMW_BusyHardware* indication from the Processor resource model
- *TMW_HighLogicalPercentDiskTime* indication from the Logical Disk resource model

A faulty disk subsystem event is generated when a disk is very busy both in terms of just itself and how busy it is keeping the CPU. Sometimes, as in the case of major file servers, this can be expected. However, for most machines, an extremely busy drive over an extended period of time indicates a disk or controller is failing.

*TMW_FaultyDiskSubsystem* has the following event properties:

| | |
|---|---|
| **LogicalDisk** | The logical disk that is being analyzed |
| **PhysicalDisk** | The physical disk that is being analyzed |
| **LogicalPercentDiskTime** | The percentage usage of the logical drive |
| **PhysicalPercentDiskTime** | The percentage usage of the physical drive |
| **Processor** | The identity of the CPU |
| **PercentProcessorTime** | The current percentage use of the CPU |
| **PercentInterruptTime** | The current percentage usage of the CPU as it handles interrupt requests |
| **InterruptsSec** | The number of interrupts per second that are passed to the CPU |
| **Severity** | The default value is critical |

# High Disk Read Bytes per Second

The *TMW_HighDiskReadBytesSec* indication is generated from the following indications:

- *TMW_HighPhysicalDiskReadBytesSec* indication from the Physical Disk resource model
- *TMW_HighLogicalDiskReadBytesSec* indication from the Logical Disk resource model

This correlation is generated when a high read bytes per second is indicated in both the logical drive and the physical drive. This usually indicates that bytes per second read from the physical drive is high.

*TMW_HighDiskReadBytesSec* has the following event properties:

| | |
|---|---|
| **LogicalDisk** | The logical disk that is being analyzed |
| **PhysicalDisk** | The physical disk on which the logical disk resides |
| **LogicalDiskReadBytes** | The number of bytes read per second on the logical disk |
| **PhysicalDiskReadBytesSec** | The number of bytes read per second on the physical disk |
| **LogicalDiskReadSec** | The number of read transactions per second on the logical disk |
| **PhysicalDiskReadSec** | The number of read transactions per second on the physical disk |
| **LogicalPercentDiskRead** | The percentage usage of the logical drive during a read operation |
| **PhysicalPercentDiskRead** | The percentage usage of the physical drive during a read operation |
| **Severity** | The default value is warning |

# High Disk Write Bytes per Second

The *TMW_HighDiskWriteBytesSec* indication is generated from the following indications:

- *TMW_HighLogicalDiskWriteBytesSec* indication from the Logical Disk resource model
- *TMW_HighPhysicalDiskWriteBytesSec* indication from the Physical Disk resource model

This correlation is generated when a high write bytes per second is indicated in both the logical drive and physical drive. This usually indicates that bytes per second written to the physical drive is high.

*TMW_HighDiskWriteBytesSec* has the following event properties:

| | |
|---|---|
| **LogicalDisk** | The logical disk that is being analyzed |
| **PhysicalDisk** | The physical disk on which the logical disk resides |
| **LogicalDiskWriteBytesSec** | The number of bytes read per second on the logical disk |
| **PhysicalDiskWriteBytesSec** | The number of bytes read per second on the physical disk |
| **LogicalDiskWriteSec** | The number of read transaction per second on the logical disk |
| **PhysicalDiskWriteSec** | The number of read transaction per second on the physical disk |
| **LogicalPercentDiskWrite** | The percentage usage of the logical drive during a write operation |
| **PhysicalPercentDiskWrite** | The percentage usage of the physical drive during a write operation |
| **Severity** | The default value is warning |

# High Drive Transfer Rate

The *TMW_HighDriveXferRate* indication is generated from the following indications:

- *TMW_HighLogicalDiskXferRate* indication from the Logical Disk resource model
- *TMW_HighPhysicalDiskXferRate* indication from the Physical Disk resource model

This correlation is generated when a high read and written bytes per second is indicated in both the logical drive and physical drive. This usually indicates that bytes per second transferred through the physical drive is high.

*TMW_HighDriveXferRate* has the following event properties:

**LogicalDisk** The logical disk that is being analyzed.

**PhysicalDisk** The physical disk on which the logical disk resides.

**LogicalDiskXfersSec** The rate of bytes transmitted (read or write) per second on the logical disk.

**PhysicalDiskXfersSec** The rate of bytes transmitted (read or write) per second on the physical disk.

**LogicalPercentDiskReadTime**
The percentage usage of the logical drive during a read operation.

**PhysicalPercentDiskReadTime**
The percentage usage of the physical drive during a read operation.

**Severity** The default value is warning.

# High Percent Disk Time

The *TMW_HighPercentDiskTime* indication is generated from the following indications:

- *TMW_HighPhysicalPercentDiskTime* indication from the Physical Disk resource model
- *TMW_HighLogicalPercentDiskTime* indication from the Logical Disk resource model

This correlation is generated when a high percentage disk time is indicated in both the logical drive and physical drive. A busy physical disk is the most likely cause.

*TMW_HighPercentDiskTime* has the following event properties:

| | |
|---|---|
| **LogicalDisk** | The logical disk that is being analyzed. |
| **PhysicalDisk** | The physical disk on which the logical disk resides. |
| **LogicalPercentDiskTime** | The percentage usage of the logical drive. |
| **PhysicalPercentDiskTime** | The percentage usage of the physical drive. |
| **LogicalPercentReadTime** | The percentage usage of the logical drive during a read operation. |
| **PhysicalPercentReadTime** | The percentage usage of the physical drive during a read operation. |
| **LogicalPercentWriteTime** | The percentage usage of the logical drive during a write operation. |
| **PhysicalPercentWriteTime** | The percentage usage of the physical drive during a write operation. |
| **Severity** | The default value is warning. |

# Possible Disk Fragmentation

The *TMW_PossibleFrag* indication is generated from the following indications:

- *TMW_LogicalPossibleFrag* indication from the Logical Disk resource model
- *TMW_PhysicalPossibleFrag* indication from the Physical Disk resource model

This correlation is generated when a possible fragmentation is indicated in both the logical drive and physical drive. It signifies that there is possibly some fragmentation on the physical drive.

*TMW_PossibleFrag* has the following event properties:

| | |
|---|---|
| **LogicalDisk** | The logical disk that is being analyzed |
| **PhysicalDisk** | The physical disk that is being analyzed |
| **LogicalPercentDiskTime** | The percentage usage of the logical disk |
| **PhysicalPercentDiskTime** | The percentage usage of the physical disk |
| **LogicalDiskBytesSec** | The transfer rate (both read and write) per second of logical disk |
| **PhysicalDiskBytesSec** | The transfer rate (both read and write) per second of physical disk |
| **Severity** | The default value is minor |

# Process Hogging CPU

The *TMW_ProcessHoggingCPU* indication is generated from the following indications:

- *TMW_HighProcesses* indication from the Processor resource model
- *TMW_ProcessHighCPU* indication from the Process resource model

This indication is generated when CPU has a high percentage utilization and one or more processes use an unusually high percentage of CPU.

*TMW_ProcessHoggingCPU* has the following event properties:

| | |
|---|---|
| **Process** | The name of the process with the highest CPU usage that is active on the CPU. |
| **IDProcess** | The identity of the process with the highest CPU usage. |
| **Processor** | The CPU being examined. |
| **PercentProcessorTime** | The total percentage usage of the CPU. |
| **PrcPercentUserTime** | The percentage usage of the CPU that is being used by the process. |
| **PrcPercentPriviledgedTime** | The percentage privileged time of the CPU that is being used by the process. |
| **PrcPriorityBase** | The base priority of the process. |
| **Severity** | The default value is critical. |

# Slow Hard Drive

The *TMW_SlowHardDrive* indication is generated from the following indications:

- *TMW_SlowPhysicalDrive* indication from the Physical Disk resource model
- *TMW_SlowLogicalDrive* indication from the Logical Disk resource model

This correlation is generated when a slow drive is indicated in both the logical drive and the physical drive. In this situation, the physical hard drive is possibly too slow.

*TMW_SlowHardDrive* has the following event properties:

| | |
|---|---|
| **LogicalDisk** | The logical disk that is being analyzed |
| **PhysicalDisk** | The physical disk on which the logical disk resides |
| **CurrentLogicalQLength** | The current queue length for the logical disk |
| **CurrentPhysicalQLength** | The current queue length for the physical disk |
| **PercentLogicalDiskTime** | The percentage usage of the logical disk |
| **PercentPhysicalDiskTime** | The percentage usage of the physical disk |
| **Severity** | The default value is warning |

# Glossary

## A

**Adapter Configuration Facility.** In the Tivoli Enterprise Console, a graphical user interface that enables a Tivoli administrator to easily configure and customize event adapters.

## C

**cache.** A buffer storage that contains frequently accessed instructions and data; it is used to reduce access time.

**configuration file.** A file that specifies the characteristics of a system device or network.

## D

**data view.** In IBM Tivoli Monitoring, a way of displaying data in the Web Health Console.

**database browser.** In IBM Tivoli Monitoring, the Web Health Console component that provides hierarchical access to information about current resource problems. This information can come from any endpoint to which the Web Health Console has access, and it might pertain to any resource model installed on that endpoint. See also Web Health Console and resource model.

**default policy.** In a Tivoli environment, a set of resource property values that are assigned to a resource when the resource is created.

**defragmentation.** The process of running a software utility to rewrite fragmented data to contiguous sectors of a computer storage medium to improve access and retrieval time.

**Desktop Management Task Force (DMTF).** An alliance of computer vendors that was convened to define streamlined management of the diverse operating systems commonly found in an enterprise.

## E

**endpoint.** In a Tivoli environment, a Tivoli client that is the ultimate recipient for any type of Tivoli operation.

**event.** In a Tivoli environment, any significant change in the state of a system resource, network resource, or network application. An event can be generated for a problem, for the resolution of a problem, or for the successful completion of a task. Examples of events are the normal starting and stopping of a process, the abnormal termination of a process, and the malfunctioning of a server. See also *indication*.

**event class.** In the Tivoli Enterprise Console, a classification for an event that indicates the type of information that the event adapter will send to the event server.

## F

**fragmentation.** An operating system's process of writing different parts of a file to discontiguous sectors on a computer storage medium when contiguous space that is large enough to contain the entire file is not available. When data is thus fragmented, the time that it takes to access the data may increase because the operating system must search different tracks for information that should be in one location.

## G

**gateway.** In a Tivoli environment, software running on a managed node that provides all communication services between a group of endpoints and the rest of the Tivoli environment. This gateway includes the multiplexed distribution (MDist) function, enabling it to act as the fanout point for distributions to many endpoints.

## H

**Web Health Console.** In IBM Tivoli Monitoring, a component that displays real-time and historical data for any resource model at any endpoint. Using the graphical user interface, users can locate individual problems associated with one or more resources. The status is displayed as a value between 0 (representing an identified problem, that is, an event) and 100 (representing no recent indications). Users can select views of resource problems as tabular data, different types of charts, and so on.

## I

**IBM Tivoli Monitoring.** A Tivoli application that applies preconfigured, automated best practices to the automated monitoring of essential system resources. The application detects bottlenecks and other potential problems and provides for the automatic recovery from critical situations, which eliminates the need for system administrators to manually scan through extensive performance data. The application also integrates seamlessly with other Tivoli Availability solutions, including the Tivoli Business Systems Manager and the Tivoli Enterprise Console. Previously called Tivoli Distributed Monitoring (Advanced Edition).

**indication.** An entity triggered by the occurrence of a problem in an endpoint relating to one or more resources. Indications are consolidated into events within the endpoint being monitored. See also *event*.

**installation repository (IR).** In Tivoli Software Installation Service (SIS), the directory that contains reusable installation images and other data that is used by SIS.

## M

**managed node.** In a Tivoli environment, any managed resource on which the Tivoli Management Framework is installed.

**monitoring collection.** In IBM Tivoli Monitoring, a collection of predefined monitors. Several monitoring collections are packaged with IBM Tivoli Monitoring, but Tivoli administrators can also use custom-developed and third-party monitoring collections. See also *custom monitor*.

## O

**object.** In object-oriented design or programming, a concrete realization of a class that consists of data and the operations associated with that data.

## P

**policy region.** In a Tivoli environment, a group of managed resources that share one or more common policies. Tivoli administrators use policy regions to model the management and organizational structure of a network computing environment. The administrators can group similar resources, define access to and control the resources, and associate rules for governing the resources.

**policy subregion.** In a Tivoli environment, a policy region created or residing in another policy region. When a policy subregion is created, it initially uses the resource and policy properties of the parent policy region. The Tivoli administrator can later change or customize these properties to reflect the specific needs and differences of the subregion.

**profile.** In a Tivoli environment, a container for application-specific information about a particular type of resource. A Tivoli application specifies the template for its profiles; the template includes information about the resources that can be managed by that Tivoli application. A profile is created in the context of a profile manager; the profile manager links a profile to the Tivoli resource (for example, a managed node) that uses the information contained in the profile. A profile does not have any direct subscribers.

**profile manager.** In a Tivoli environment, a container for profiles that links the profiles to a set of resources, called subscribers. Tivoli administrators use profile managers to organize and distribute profiles. A profile manager is created in the context of a policy region and is a managed resource in a policy region.

## R

**reference model.** In the context of Tivoli software, the model configuration for a system or set of systems that is used to maintain consistent configurations in a distributed environment. In Tivoli Inventory, reference models are created in the configuration repository.

**resource.** In IBM Tivoli Monitoring context, a resource is anything that affects the operation of the system and includes physical and logical disks, CPUs, memory, printers, as well as the processes running, services, such as LanMan and the Windows event log, and TCP/IP.

**resource model.** In IBM Tivoli Monitoring, the logical modeling of one or more resources, along with the logic on which cyclical data collection, data analysis, and monitoring are based. Related events and actions are triggered, if required. For any resource model, users can specify individual thresholds and event aggregation rules. See also *event*.

**rule base.** In the Tivoli Enterprise Console, one or more rule sets and the event class definitions for which the rules are written. The Tivoli Enterprise Console uses the rule base in managing events. An organization can create many rule bases, with each rule base fulfilling a different set of needs for network computing management.

## S

**Software Installation Service (SIS).** A Tivoli product that provides an easy-to-use, efficient interface for installing Tivoli Enterprise™ software. SIS uses Tivoli's MDist technology and provides automated checking for prerequisite software, a reusable repository of installation images, and both graphical and command line interfaces for deploying Tivoli products to a large number of computers.

**subscriber.** In a Tivoli environment, a managed node, a profile manager, an endpoint, or another Tivoli client that is subscribed to a profile manager. Although profiles are distributed to a subscriber, the subscriber may or may not be the final destination of the profile distribution.

## T

**task library.** In a Tivoli environment, a container in which a Tivoli administrator can create and store tasks and jobs.

**threshold.** (1) In software products, a value that defines a limit for a monitored condition. (2) In IBM Tivoli Monitoring, a threshold is a named property with a user-defined value. Typically, the value specified for a threshold represents a significant level of a performance-related entity, which, if exceeded, a system administrator might want to know about.

**Tivoli Distributed Monitoring.** Previous name of Tivoli Distributed Monitoring (Classic Edition).

**Tivoli Distributed Monitoring (Advanced Edition).** Previous name of IBM Tivoli Monitoring.

**Tivoli Distributed Monitoring (Classic Edition).** A Tivoli application that provides distributed monitors for monitoring system resources. The application initiates necessary corrective actions and informs system administrators of potential problems. These monitors can be centrally configured and deployed to monitor individual machines. The application also

integrates seamlessly with other Tivoli Availability solutions, including the Tivoli Enterprise Console. Previously called Tivoli Distributed Monitoring.

**Tivoli Enterprise Console.** A Tivoli product that collects, processes, and automatically initiates corrective actions for system, application, network, and database events; it is the central control point for events from all sources. The Tivoli Enterprise Console provides a centralized, global view of the network computing environment; it uses distributed event monitors to collect information, a central event server to process information, and distributed event consoles to present information to system administrators.

**Tivoli environment.** The Tivoli applications, based upon the Tivoli Management Framework, that are installed at a specific customer location and that address network computing management issues across many platforms. In a Tivoli environment, a system administrator can distribute software, manage user configurations, change access privileges, automate operations, monitor resources, and schedule jobs.

**Tivoli management agent.** In the Tivoli environment, an agent that securely performs administrative operations.

**Tivoli Management Framework.** The base software that is required to run the applications in the Tivoli product suite. This software infrastructure enables the integration of systems management applications from Tivoli Systems Inc. and the Tivoli Partners. In a Tivoli environment, the Tivoli Management Framework is installed on every client and server; however, the TMR server is the only server that holds the full object database.

**Tivoli management region.** In a Tivoli environment, a Tivoli server and the set of clients that it serves. An organization can have more than one Tivoli management region. A Tivoli management region addresses the physical connectivity of resources whereas a policy region addresses the logical organization of resources.

# Index

## E



## F



## H

## I

## J



## K



## L

## M



## N

## O



## P

## R



## S

## T

## U



## V



## W

SH19-4570-00