



# WebSphere Studio Application Monitor Operator's Guide

3.1





WebSphere Studio Application Monitor



# WebSphere Studio Application Monitor Operator's Guide

*3.1*

**Note:**

Before using this information and the product it supports, read the information in "Notices" on page 51 .

**First Edition (November 2004)**

This edition applies to WebSphere Studio Application Monitor (product number 5697-J18) and to all subsequent releases and modifications until otherwise indicated in new editions.

You can order publications through your IBM representative or the IBM branch office serving your locality.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Licensed Materials - Property of IBM.

WebSphere Studio Application Monitor (program number 5697-J18).

© Copyright, IBM Corp. 2004 All Rights Reserved.

US Government User Restricted by GSA ADP Schedule Contract with IBM Corp.

IBM, the IBM logo, and WebSphere are trademarks of IBM Corp. in the U.S., other countries, or both.

© Copyright International Business Machines Corporation 2004. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>Preface</b> .....	<b>v</b>
About this book .....	v
Who should read this book .....	v
Where to find more information .....	v
Documentation Conventions .....	v

<b>Tables</b> .....	<b>vii</b>
---------------------	------------

---

## **Part 1. The Application Monitor** .....

---

**1**

### **Chapter 1. System Level Operations** .....

**3**

Starting DB2 UDB .....	3
Starting the Oracle Database .....	4
Starting the Web Server .....	4
Starting Virtual Frame Buffer .....	4
Linux .....	4
AIX .....	4
Starting the Application Monitor Visualization Engine .....	5
Starting the rest of the Application Monitor .....	5
Stopping the Application Monitor .....	5

### **Chapter 2. Component Level Operations** .....

**7**

Starting the Kernel, Publish Server, Global Publish Server, Archive Agent, Message Dispatcher and Polling Agent .....	7
Stopping the Kernel, Publish Server, Global Publish Server, Archive Agent, and Message Dispatcher .....	8
Checking the status of the Kernel, Publish Server, Global Publish Server, Archive Agent, and Message Dispatcher .....	9
Watchdog Operations .....	9
Starting the Visualization Engine .....	10
Performing Database Optimization .....	10
Data Maintenance .....	10
To perform data maintenance on the database: .....	10
Oracle Database Optimization .....	11
The Data Trimmer .....	11
Log Files .....	11

---

## **Part 2. The Common Services Address Space (z/OS and OS/390)** .....

---

**13**

### **Chapter 3. The Common Services Address Space** .....

**15**

### **Chapter 4. Commands and Parameters** .....

**17**

Starting CYN1 .....	17
Stopping CYN1 .....	18
Modifying CYN1 .....	18
The ACT CYN1 Command .....	18
The INACT CYN1 Command .....	19
The DUMP CYN1 Command .....	19
The SMFWRITE CYN1 Command .....	20
The TRACE CYN1 Command .....	20
The SPINLOG CYN1 Command .....	20
The ZAP CYN1 Command .....	20
The STAMP CYN1 Command .....	21
The VER CYN1 Command .....	21
The STATS CYN1 Command .....	22

<b>Part 3. The Data Collector</b> .....	<b>23</b>
<b>Chapter 5. Disabling the Data Collector</b> .....	<b>25</b>
PMI Module Settings .....	25
<b>Part 4. Appendixes</b> .....	<b>27</b>
<b>Appendix A. Starting Multiple Instances of Application Monitor Components</b> .....	<b>29</b>
<b>Appendix B. Sample CRON Jobs</b> .....	<b>31</b>
<b>Appendix C. Managing Server and Data Collector Tuning</b> .....	<b>35</b>
Properties Files .....	35
aa.properties .....	35
ve.properties .....	37
md.properties .....	37
kl.properties .....	38
datacollector.properties .....	39
ps.properties .....	42
sam1.properties .....	44
Miscellaneous .....	45
<b>Appendix D. Data Trimmer</b> .....	<b>47</b>
Data Trimmer Properties Files .....	47
<b>Notices</b> .....	<b>51</b>
Trademarks .....	53
<b>Index</b> .....	<b>55</b>
<b>Bibliography</b> .....	<b>57</b>

---

## Preface

---

### About this book

This book is the Operator's Guide for WebSphere® Studio Application Monitor (Application Monitor). It contains information about the operation of the Application Monitor and the common services address space.

---

### Who should read this book

Those who want to use the Application Monitor.

---

### Where to find more information

The following list shows the books in the Application Monitor library:

- *WebSphere Studio Application Monitor User's Guide* contains instructions and user information for the Application Monitor.
- *WebSphere Studio Application Monitor Operator's Guide* contains information about the operation of Application Monitor and the common services address space.
- *WebSphere Studio Application Monitor Installation and Customization Guide* contains instructions on installing user's exits and customizing the Application Monitor.
- *WebSphere Studio Application Monitor for CICS® Data Collector Product Guide* contains information about the installation, configuration and use of the Application Monitor CICS Data Collector.
- *WebSphere Studio Application Monitor for IMS Data Collector Product Guide* contains information about the installation, configuration and use of the Application Monitor IMS Data Collector.
- *WebSphere Studio Application Monitor Messages and Codes* contains information about messages and codes generated by the Application Monitor.
- *WebSphere Studio Application Monitor Program Directory* contains complete installation instructions for the Application Monitor Engine.
- *WebSphere Studio Application Monitor Program Directory for the CICS Data Collector* contains complete installation instructions for the Application Monitor CICS Data Collector Engine.
- *WebSphere Studio Application Monitor Program Directory for the IMS Data Collector* contains complete installation instructions for the Application Monitor IMS Data Collector Engine.
- *WebSphere Studio Application Monitor* has an online help system that describes all of the commands and dialogs available from its graphical user interface.

---

### Documentation Conventions

The following table describes the conventions used throughout this book.

*Table 1. Documentation Conventions*

<b>Convention</b>	<b>Item</b>
Ctrl+Tab	Indicates that you must press two or more keys simultaneously.
<installation directory>	Indicates that you should insert the directory appropriate for your environment.
monospace text	Indicates code samples.



---

## Tables

1. Documentation Conventions . . . . .	vi	3. Example of Zap CYN1 Command . . . . .	21
2. Example of Dump CYN1 Command Output . . . . .	20	4. Trademarks. . . . .	53



---

## Part 1. The Application Monitor



---

## Chapter 1. System Level Operations

The Application Monitor runs on either the RS6000/AIX®, Intel/Linux, Sparc/Solaris, or z/Linux platform. It consists of the following components:

- Kernel
- Publish Server
- Archive Agent
- Message Dispatcher
- Visualization Engine
- Database
- SAM/GPS
- Polling Agent

It is by default that the code base for the Application Monitor is located at /opt/cyaneaone. If you installed the Application Monitor elsewhere, please adjust the documentation accordingly.

The four sub-directories under the /opt/cyaneaone are bin, etc, lib, and logs.

1. Executable scripts are located under the /opt/cyaneaone/bin directory.
2. Configuration files are located under the /opt/cyaneaone/etc directory.
3. Log files generated by the components are located under /opt/cyaneaone/logs.

All commands are executed as the "cyanea" user.

The Application Monitor will start 2 instances of the Kernel, the Publish Server and Archive Agent and a single instance of the remaining components by default. (For more information on starting multiple instances of the Application Monitor Component, please see the Appendix A, "Managing Server and Data Collector Tuning," on page 35).

**To start the Application Monitor in the following sequence after a reboot:**

1. Start the database.
2. Start the Web Server.
3. Start Virtual Frame Buffer.
4. Start the Application Monitor Visualization Engine.
5. Start the rest of the Application Monitor components.

---

### Starting DB<sup>®</sup>2 UDB

If you are using DB2 for your managing server database perform the following steps.

To manually start DB2 UDB, type in:

```
su - db2inst1
db2start
```

---

## Starting the Oracle Database

To manually start the Oracle database, type in

```
su - oracle
```

1. Check if ORACLE\_SID is octigate, if not, do:  

```
export ORACLE_SID=octigate
```
2. Start Server Manager:  

```
svrmgrl
```
3. Connect to octigate database:  

```
connect internal
```
4. Start octigate database:  

```
startup
```
5. Exit from Server Manager Command line mode:  

```
quit
```
6. Check if Oracle database listener is running and if octigate is part of the service:  

```
lsnrctl stat
```
7. Start Oracle database listener if it has not yet been up running or octigate is not in the service:  

```
lsnrctl start
```

---

## Starting the Web Server

To manually start the Web Server, type in:

```
su - root
```

```
/opt/IBMHHTTP/bin/apachectl start
```

---

## Starting Virtual Frame Buffer

You will need Virtual Frame Buffer if you are starting WebSphere for the managing server from a telnet terminal window. Virtual Frame Buffer is required for the operation of Performance Analysis & Reporting (PAR). If it is not installed, the graphs in PAR will not display.

### Linux

To implement a virtual X11 environment on Linux:

To start Xvfb server, type in:

```
$ /usr/bin/nohup /usr/X11R6/bin/Xvfb :1 1>/dev/null 2>/dev/null &
```

### AIX

To start Xvfb server, type in:

```
$ X -vfb -force :1 &
```

OR

```
$ /bin/nohup X -vfb -force :1 1>/dev/null 2>/dev/null &
```

---

## Starting the Application Monitor Visualization Engine

To manually start the Application Monitor visualization engine, type in:

```
su - root
```

Start virtual frame buffer as described above.

```
DISPLAY = :1 /opt/WebSphere/AppServer/bin/startupserver.sh &
```

WebSphere will be started.

**Note:** If you followed the instructions in the Installation and Customization Guide, the Application Monitor Visualization Engine Application Server should also be started too.

---

## Starting the rest of the Application Monitor

To start the rest of the Application Monitor manually, type in:

```
su - cyanea
```

```
cd /opt/cyaneaone/bin
```

```
sh cyanea-start.sh
```

---

## Stopping the Application Monitor

To stop the Application Monitor manually, type in:

```
su - cyanea
```

```
sh cyanea-stop.sh
```

This will stop all the stand alone Java™ components. It will not stop the Cyanea application server or the database.





---

## Chapter 2. Component Level Operations

This section gives you guidelines of how to operate on specific components within the Application Monitor.

The Kernel, Publish Server, Archive Agent, Message Dispatcher, Polling Agent, and Global Publish Server are all controlled via a common script "cyaneactl". The syntax is:

```
su - cyanea
```

```
./cyaneactl.sh <property file name> <command>
```

**Note:** Property file name refers to the name of the property file associated with the instance of the component.

*Command* refers to one of "start", "stop", "ping" or "status".

---

### Starting the Kernel, Publish Server, Global Publish Server, Archive Agent, Message Dispatcher and Polling Agent

The syntax of starting the components:

```
su - cyanea
```

```
./cyaneactl.sh <property file name> start
```

By default, the following property files are available on the Application Monitor:

- K11 - first instance of the kernel
- K12 - second instance of the kernel
- ps1 - first instance of the publish server
- ps2 - second instance of the publish server
- sam - the one instance Global Publish Server
- aa1 - first instance of the archive agent
- aa2 - second instance of the archive agent
- md- message dispatcher
- pa - polling agent

For Example:

To start the first instance of the kernel, type in

```
su - cyanea
```

```
./cyaneactl.sh k11 start
```

To start the second instance, type in

```
./cyaneactl.sh k12 start
```

To start the first instance of the archive agent, type in (Check that the two instances of the archive agent have different Server Port Number settings.)

```
su - cyanea
```

```
./cyaneact1.sh aa1 start
```

To start the second instance of the archive agent, type in

```
su - cyanea
```

```
./cyaneact1.sh aa2 start
```

To check whether the component started successfully, you can type in

```
su - cyanea
```

```
./cyaneact1.sh <property file name> ping
```

---

## Stopping the Kernel, Publish Server, Global Publish Server, Archive Agent, and Message Dispatcher

The syntax of stopping the components:

```
su - cyanea
```

```
./cyaneact1.sh <property file name> stop
```

By default, the following property files are available on the Application Monitor:

- K11 - first instance of the kernel
- K12 - second instance of the kernel
- ps1 - first instance of the publish server
- ps2 - second instance of the publish server
- sam - the one instance Global Publish Server
- aa1 - first instance of the archive agent
- aa2 - second instance of the archive agent
- md- message dispatcher
- pa - polling agent

For example:

To stop the first instance of the Kernel, type in

```
su - cyanea
```

```
./cyaneact1.sh k11 stop
```

To stop the second instance, type in

```
su - cyanea
```

```
./cyaneact1.sh k12 stop
```

To stop the first instance of the archive agent, type in

```
su - cyanea
```

```
./cyaneact1.sh aa1 stop
```

To stop the second instance of the archive agent, type in

```
su - cyanea
```

```
./cyaneact1.sh aa2 stop
```

---

## Checking the status of the Kernel, Publish Server, Global Publish Server, Archive Agent, and Message Dispatcher

The syntax of checking the status of the components:

```
su - cyanea  
./cyaneact1.sh <property file name> status
```

By default, the following property files are available on the Application Monitor:

- K11 - first instance of the kernel
- K12 - second instance of the kernel
- ps1 - first instance of the publish server
- ps2 - second instance of the publish server
- sam - the one instance Global Publish Server
- aa1 - first instance of the archive agent
- aa2 - second instance of the archive agent
- md- message dispatcher

For example:

To check the status of the first instance of the Kernel, type in

```
su - cyanea  
./cyaneact1.sh k11 status
```

To check the status of the second instance, type in

```
su - cyanea  
./cyaneact1.sh k12 status
```

To check the status of the first instance of the archive agent, type in

```
su - cyanea  
./cyaneact1.sh aa1 status
```

To check the status of the second instance of the archive agent, type in

```
su - cyanea  
./cyaneact1.sh aa2 status
```

---

## Watchdog Operations

The Watchdog is a light weight component that manages a particular kernel. If the kernel does not respond, the watchdog process restarts the kernel. Each kernel has its own watchdog process. The Watchdog process must be on the machine where the kernel is running.

Watchdog properties file:

```
k1wd1.properties
```

Tune the following properties in the file:

```
kernelwd.renewal.failure=3
```

This property takes an integer value. It's the number of times the Watchdog pings the kernel before it restarts the kernel.

```
kernelwd.ping.interval=5000
```

This property takes an integer value. It's units are in milliseconds. It is the time between individual pings from the Watchdog to the Kernel.

To start the kernel with the Watchdog:

```
cyaneactl.sh kll start
```

To stop the kernel with the Watchdog:

```
cyaneactl.sh kll stop
```

To start the kernel without the Watchdog:

```
klctl.sh start kll.properties
```

To stop the kernel without the Watchdog:

```
klctl.sh stop kll.properties
```

---

## Starting the Visualization Engine

To start the Visualization Engine, type in:

```
su - cyanea
```

```
/opt/WebSphere/AppServer/bin/startupserver.sh &
```

---

## Performing Database Optimization

### Data Maintenance

The script run-stat-cmds.sh is provided as a command line utility so that users can simply schedule CRON jobs to periodically maintain the database. Execute run-stat-cmds.sh daily. If there is still a performance problem in the Performance Analysis and Reporting, then perform the following procedures.

The data maintenance utilities are as follows:

- REORGCHK
- REORG
- RUNSTATUS

### To perform data maintenance on the database:

1. Login to WSAM db2 server.
2. Db2 connect to octigate user db2inst1.
3. Db2 reorgchk update statistics on table <cyanea.tablename>. Go to step 4 and 5 for table reorganization if an asterisk is shown in the REORG column of the output. Otherwise, skip them.
4. Db2 reorg table <cyanea.tablename> index <cyanea.indexname>.
5. Db2 runstatus on table <cyanea.tablename>.

**Note:** If tables are being reorganized, then Archive Agent will not be able to insert any data into the tables.

## Oracle Database Optimization

To perform database maintenance on the Oracle database, perform the following steps:

1. Login to the database server machine:  
`su - oracle`
2. Connect to the octigate database:  
`sqlplus <cyanea username>/<password>@octigate`
3. Execute the command to collect schema statistics:  
`exec dbms_stats.gather_schema_stats ('cyanea')`
4. Exit the database:  
`quit`

---

### The Data Trimmer

The Data Trimmer is used to trim old data that does not need to be maintained in the database any longer. It is provided as a command line utility so that users can simply schedule CRON jobs to periodically trim appropriate data. Please be aware that running this command could take a very long time depending on the amount of data being trimmed.

The syntax is as follows:

```
/opt/cyaneaone/bin/dataTrimmer.sh dbuser dbpassword startdate starttime enddate  
endtime [maxrow]
```

starttime and endtime must be specified using military (24-hour) notation. The optional parameter maxRow is the number of rows is committed to database in a transaction. The default for maxRow is 1000.

For example, to delete data collected from January 1 2003 to March 31 2003, use the following invocation:

```
/opt/cyaneaone/bin/dataTrimmer.sh dbuser dbpassword 01/01/03 00:00:00 03/31/03  
11:59:59
```

---

### Log Files

The following parameters pertain to the file size and rotation of the log files. The log files are located in <cyanea\_home>/etc directory. The log files available are as follows:

- log-kl1.properties
- log-kl2.properties
- LogTest.properties
- log-aa.properties
- log-aa1.properties
- log-aa2.properties
- log-am.properties
- log-td.properties
- log-md.properties
- log-ps1.properties

- log-ps2.properties
- log-sam1.properties
- log-ps1.properties
- log-pa.properties

The parameters are as follows:

MaxFileSize=1000KB

MaxBackupIndex=10

---

## **Part 2. The Common Services Address Space (z/OS<sup>®</sup> and OS/390<sup>®</sup>)**





---

## Chapter 3. The Common Services Address Space

The Common Services Address Space CYN1 only exists in a z/OS or OS/390 environment where applications running inside WebSphere are being monitored. It serves as a repository for SMF records generated by the WebSphere Application Server and captured by the SMF address space. This repository takes copies of SMF records and provides the basis for real time analysis when used in conjunction with the Monitoring Console. It also provides other functions that are unique to the z/OS or OS/390 environment.



---

## Chapter 4. Commands and Parameters

---

### Starting CYN1

CYN1 is defined as an MVS™ subsystem; it also runs as a started task and the JCL resides in one of the system PROCLIB datasets. The PROCLIB member name and the subsystem name do not have to be the same; if they are, the sub parameter "SUB=JES2" must be added to the start command to inform the system to start the job under the JES2 procedure library, instead of the MSTR library, which is the default when the task being started is a subsystem.

The syntax of the start command is as follows:

```
START xxxx[.yy] [,SUB=JES2] [,PARM=zzzz] [SMFWRITE=ON | OFF]
```

Where:

**xxxx** The name of the PROCLIB member containing the JCL for CYN1. "CYN1" is also the name of the MVS subsystem. The PROCLIB member is usually called "CYN1", but it does not have to be. In fact, in order to avoid specifying "SUB=JES2", it is recommended to keep those names distinct and use "CYN1PROC" as the name for the Proclib member.

**yy** The optional identifier of the started task for future reference is the MODIFY or STOP commands. For short yy = "CC". If yy is not specified, the system uses xxxx as the identifier.

#### **SUB=JES2**

If the PROCLIB member and the subsystem share the same name, that is "CYN1", this specification is required so that the JCL to be selected is to be found in the JES2 procedure library, e.g., USER.PROCLIB, instead of the default MSTR.

#### **PARM=zzzz**

It is appended to the string "CYNP" to form the name of a parameter module DSECT, member of the STEPLIB library, in which the run parameters are specified. If omitted, the name specified in the PARM keyword, in the JCL EXEC statement will be selected. If there is no PARM keyword, this specification is required.

#### **SMFWRITE=ON | OFF**

Either ON or OFF must be used if the SMFWRITE directive is specified. If omitted, the name specified in the SMFWRITE keyword, in the JCL EXEC statement will be selected. If there is no SMFWRITE keyword, this specification is required. This specification sets the initial value of the SMFWRITE variable, which can be subsequently changed with the SMFWRITE command. IF SMFWRITE is ON, WebSphere Application Server generated SMF records (type 120, subtypes 3, 6, and 8) will be written to the SMF datasets. If SMFWRITE is OFF, writing of these records will be disabled.

Example:

```
START CYN1.CC,SUB=JES2
```

```
START CYN1PROC.DS
```

```
START CYN1PROC,PARM=BIG
```

```
START CYN1PROC,SMFWRITE=ON
```

```
START CYN1PROC.CC,SMFWRITE=OFF,PARMS=TEST
```

---

## Stopping CYN1

The syntax of the stop command is as follows:

```
P xxxx or P yy
```

Where:

**xxxx** The name of the PROCLIB member containing the JCL for CYN1. This form is required when no identifier name was supplied in the start command.

**yy** The optional identifier for the started task, in the start command. If yy is not specified, the system uses xxxx as the identifier.

**Note:** When stopping CYN1, outstanding requests will be completed. Subsequent requests, however, may produce unpredictable results. Therefore use the INACT CYN1 command (see below) prior to issuing command to insure that no additional requests will be initiated.

---

## Modifying CYN1

The syntax of the modify command is as follows:

```
F xxxx,verb[,param...[,param]...]
```

Or

```
F yy,verb[,param...[,param]...]
```

Where:

**xxxx** The name of the PROCLIB member containing the JCL for CYN1. This form is required when no identifier name was supplied in the start command.

**yy** The optional identifier for the started task, in the start command. If yy is not specified, the system uses xxxx as the identifier.

**verb** The command verb. Currently there are eight verbs: "ACT", "INACT", "DUMP", "TRACE", "SPINLOG", "ZAP", "STAMP", "VER", for the CYN1 commands described below.

**param** The CYN1 command parameters. They are command dependent and their presence is optional.

---

## The ACT CYN1 Command

The form of the ACT command is as follows:

```
F xxxx,ACT
```

Or

```
F yy,ACT
```

The ACT command is only useful to render CYN1 active again in the event the INACT command has been issued. When CYN1 is started, it is initiated in active mode, ready to process Clients' requests.

---

## The INACT CYN1 Command

The form of the INACT command is as follows:

```
F xxxx, INACT
```

Or

```
F yy, INACT
```

The INACT command puts CYN1 in a state where no further Client requests will be accepted, however, currently active requests will be completed. The use of the INACT command is recommended when "trouble" is detected, either at the Server, or the Clients' level, or when CYN1 is about to be brought down.

---

## The DUMP CYN1 Command

The DUMP command is used to display in hexadecimal and character format areas of storage in the CYN1 address space.

The form of the DUMP command is as follows:

```
F {xxxx | yy}, DUMP, AS, hexaddr[, length]
```

**AS** AS stands for address space.

### Hexaddr

It specifies an eight-digit hexadecimal starting address of the memory location to be displayed. Alternately, the starting address can be specified with keywords describing what control blocks to be displayed.

The current valid keywords are:

\$CYNCYNA	The CYN1 anchor area in CSA
\$CYNSETB	The server directory table
\$CYNSTITB	The server/container interval table
\$CYNSSPL	The server section entry pool
\$CYNBSPL	The bean section entry pool
\$CYNBMPL	The bean method section entry pool
\$CYNHSPL	The HttpSessionManager section entry pool
\$CYNWSPL	The WebApplication section entry pool
\$CYNVPL	The Servlet section entry pool
\$LOGBOOK	The in-core SYSLOG book
\$CMDBOOK	The in-core command book

**Note:** The corresponding DSECTs are part of the macro library.

### length

It specifies the length in bytes of the area to be displayed. The length is specified in decimal and has a maximum of 4 digits (9999 bytes). Default is 16 bytes (one dump line).

Example:

Table 2. Example of Dump CYN1 Command Output

F CC,DUMP,AS,00000000	It displays 16 bytes of the PSA.
F CC,DUMP,AS,\$CYNSETB,192	It displays entries 0 and 1 of the server directory table.
F CC,DUMP,AS,02A0C100,4096	It displays one memory page starting at location 02A0C100.

---

## The SMFWRITE CYN1 Command

The SMFWRITE command is used to turn on/off the writing of WebSphere Application Server generated SMF records (type 120, subtypes 3, 6, and 8) to the SMF datasets. It is also used to display the current value of SMFWRITE.

The form of the SMFWRITE command is as follows:

```
F {xxxx | yy},SMFWRITE {ON | OFF| SHOW}
```

**Note:** ON will turn the SMF writing on, while OFF will turn it off. When CYN1 comes up, the writing is governed by the SMFWRITE specification in the CYN1 START command, or is hard-coded in the JCL of the corresponding PROCLIB member JCL. SHOW will display the current value (ON or OFF) of SMFWRITE.

---

## The TRACE CYN1 Command

The TRACE command is used to turn on/off the SYSLOG trace, which currently consists of information messages regarding the reception of SMF records. The SYSLOG dataset is buffered in a circular in-core book, the size of which is specified in the PARM member.

The form of the TRACE command is as follows:

```
F {xxxx | yy},TRACE {ON | OFF}
```

**Note:** ON will turn the trace on, while OFF will turn it off. When CYN1 comes up, the trace is off.

---

## The SPINLOG CYN1 Command

The SPINLOG command is used to reset pointers to the in-core SYSLOG book, close the current SYSLOG dataset, and reopen a new one. The closed SYSLOG dataset can be subsequently retrieved by report writer programs. SYSLOG traces can be separated.

The form of the SPINLOG command is as follows:

```
F {xxxx | yy},SPINLOG
```

---

## The ZAP CYN1 Command

The ZAP command is used to modify areas of storage in the CYN1 address. The form of the ZAP command is as follows:

```
F {xxxx | yy},ZAP,{AS|SS},hexaddr,hexcont,hexrepl
```

## AS|SS

AS stands for address space dump. SS stands for address space authorized. Either one must be specified.

## hexaddr

It specifies an eight-digit hexadecimal address of the word to be zapped. Low core (4 kilobyte PSA) zapping is always disabled. System region, i.e. virtual addresses 00001001 through 00005000 can only be zapped if SS address space authorized option is specified.

## hexcont

It specifies the eight-digit hexadecimal current contents of the word to be zapped, for verification.

## Hexrepl

It specifies the eight-digit hexadecimal value to replace hexcont with at address hexaddr.

Example:

Table 3. Example of Zap CYN1 Command

Valid	Invalid
F CC,ZAP,SS,0000100F,E8000000,000000 00	F CC,ZAP,AS,00000000,040C0000,00000000 or F CC,ZAP,SS,00000000,040C0000,00000000 <b>Note:</b> It is invalid because the area to be zapped is in the PSA.
F CC,ZAP,AS,028AC300,001F007A,0000 0000  zeroes the word at virtual address 028AC300 location.	F CC,ZAP,AS,0000100F,E8000000,00000000 <b>Note:</b> It is invalid because the area to be zapped is in the system region.

---

## The STAMP CYN1 Command

The STAMP command is used to turn on/off time stamping, which consists in outputting to the SYSLOG dataset stamp records corresponding to critical events in processing clients' requests. The SYSLOG dataset is buffered in a circular in-core book, the size of which is specified in the PARM member.

The form of the STAMP command is as follows:

F {xxxx | yy},STAMP {ON | OFF}

**Note:** ON will turn the STAMP on, while OFF will turn it off. When CYN1 comes up, time stamping is off.

---

## The VER CYN1 Command

The VER command is used to display the version information in CYN1 main task entry module CYN1DRVR. The information consists of the name of the module followed by the "version.level.change" number followed by the date and time of the module assembly.

The form of the VER command is as follows:

---

## The STATS CYN1 Command

The STATS command is used to reset the statistical information that is gathered by CYN1 from the WebSphere generated SMF records. SMF records of type 120 are generated on a regular basis (recommended every 300 seconds) by every WebSphere application server instances, summarizing server activity (subtype 3), ejb container activity (subtype 6) and Web container activity (subtype 8) for that period. Normally CYN1 gathers this information and accumulates/averages these statistics across periods for each server instance, so that what the user sees on the System Resource screen is a running activity history for that server instance. The RESET option of the STATS command will cause the next generated statistics to override the current accumulated statistics for all server instances, thereby starting a new history interval. After the override, statistics accumulation resumes until the next time the RESET option is used. If the RESET option was entered by mistake, the CUMUL option will nullify it provided the CUMUL option is entered prior to WebSphere generation of the next SMF records.

The form of the STATS command is as follows:

```
F {xxxx | yy},STATS {RESET | CUMUL}
```

**Note:** RESET will reset the start of the statistics accumulation interval. CUMUL will undo RESET if entered before the next SMF record cycle. CUMUL has no effect if RESET was not entered.



---

## Part 3. The Data Collector



---

## Chapter 5. Disabling the Data Collector

For the Distributed Platform only (WebSphere 4.x, 5.x).

Follow these instructions to disable the data collector on your application server. There are two methods for disabling the data collector. The first method disables the data collector inside the application. After the data collector is disabled, it still sends a heartbeat to the managing server. The second method for disabling the data collector completely terminates all data collector functions.

### To disable the Data Collector (first method):

1. Open the GUI for the Application Monitor product.
2. From the top navigation, select **Administration > Server Management > Data Collector Configuration**.

The Data Collector Overview page opens.

3. Under the Configured Data Collectors, click **Disable** next to the data collector you want to disable.

The unconfigured data collector displays in the Unconfigured Data Collector table.

### To disable the Data Collector (second method):

1. Go to the WebSphere Advanced Administrative Console.
2. Click on the application server where the data collector is installed.
3. Click on the **JVM Settings** tab.
4. Click the **Advanced JVM Settings** button at the bottom of the screen.
5. Under Command Line Arguments replace **-Xruncyanea** with **-DXruncyanea**.
6. Click **OK** and click **Apply**.
7. Click the **Custom** tab.
8. Select Cyanea Custom Service.
9. Click **Edit**.
10. Uncheck Enabled.
11. Click **OK** and click **Apply**.
12. Restart your application server and the data collector will be disabled.

---

## PMI Module Settings

To view the data for the System Resources section, the Performance Monitoring Settings for each module in WebSphere 4 must be set to Maximum.

To view the PMI data for the WebSphere 5 data collector, the individual PMI module settings must be as follows:

JVM Runtime – High

ORB – High

ThreadPools – High

Transaction Manager – High

Enterprise Beans – Maximum  
Servlet Session Manager – High  
WebApplications – High  
J2C Connection Pool – High  
Db Connection Pool – High

---

## Part 4. Appendixes



---

## Appendix A. Starting Multiple Instances of Application Monitor Components

By default, the following property files are available on the Application Monitor:

- K11 - first instance of the Kernel
- K12 - second instance of the Kernel
- ps1 - first instance of the Publish Server
- ps2 - second instance of the Publish Server
- aa1 - first instance of the Archive Agent
- aa2 - second instance of the Archive Agent
- md - Message Dispatcher
- pa - Polling Agent
- sam - the one instance Global Publish Server

You may create additional instances of the Kernel, Publish Server or Archive Agent. Each instance must have its own property file. The simplest way to create another instance of one of these components is to make a copy of an existing property file and change the component ID and the port number to a unique value.

For example, to create an additional instance of the Publish Server, you would copy ps1.properties to ps3.properties and modify the lines in that file that define the component ID and port number to be unique values distinct from component ID and port number used by any other component instance.





---

## Appendix B. Sample CRON Jobs

You can find sample CRON jobs on the managing server in `/opt/cyaneaone/bin` for regular housekeeping tasks.

### HTTP Server Log Files

Rotate the HTTP server log files, created by the IHS or Apache HTTP server running on the same node with the VE. This is necessary on Linux, but is recommended on other supported managing server platforms as well.

The crontab and the script must be edited before they are used, and the `APACHE_HOME` variable must be defined correctly. The crontab must be run as the person who controls the HTTP server processes and has permissions to delete server log files (often the root user).

### DB2 Database Tuning

Improve the performance of our DB2 (Octigate) database. The crontab and the script must be edited before they are used especially the script's schema, database, and `db2inst_home` variable must be correctly set to the owner of the tables, the database name, and the DB2 instance home directory, respectively. The script must be run as the instance owner (e.g. `db2inst1`), NOT the schema owner (e.g. `cyanea`).







---

## Appendix C. Managing Server and Data Collector Tuning

The following instructions will assist you with tuning the application to best suit the needs of your environment.

---

### Properties Files

The following properties support the Application Monitor. Use the parameters listed with the specific property if you want to adjust or change it.

#### Acronyms

- AA - Archive Agent
- MD - Message Dispatcher
- GC - Garbage Collection
- GPS - Global Publish Server
- PAR - Performance Analysis & Reporting
- PMI - Performance Monitoring Infrastructure
- PS - Publish Server
- RFS - Request for Stub
- RMI - Remote Method Invocation
- SAM - synonym for the Global Publish Server
- URI - Uniform Resource Identifier
- VE - Visualization Engine

### aa.properties

The following are the properties listed in the aa.properties file. If you have more than one instance of the Archive Agent, you will have aa1.properties, aa2.properties, and so on:

#### **serverPortNumer**

The port where the AA listens for the socket connections from the other components.

#### **queues.pmi.size**

The default value is 1024 PMI records.

Size of the PMI queue. This is a circular queue and the AA wraps the queue (cleans the queue). If the queue reaches this set size, the PS pushes PMI data to this queue.

#### **queues.pmi.priority**

The default value is 5. Priority of the PMI dequeuing thread (aka PMI Reader Thread). Default values range from minimum 1 to maximum 10.

#### **queues.method.size**

The default value is 5120 request records.

Size of the method/request queue. This is a circular queue and the AA wraps the queue (cleans the queue). If the queue reaches this set size, the PS pushes request/method data to this queue.

#### **queues.method.priority**

The default value is 6 method availability records.

Priority of the method/request dequeuing thread (aka Request Reader Thread).

**queues.availability.size**

The default value is 3072 availability records. Size of the availability queue. This is a circular queue and the AA wraps the queue (cleans the queue). If the queue reaches this set size, the PS and Kernel push availability data to this queue.

**queues.availability.priority**

The default value is 7.

Priority of the availability dequeuing thread (aka Availability Reader Thread).

**queues.memory.size**

The default value is 1024 memory records.

Size of the memory queue. This is a circular queue and the AA wraps the queue (cleans the queue). If the queue reaches this set size, the PS pushes memory related data (GC data and Heap Dump data) to this queue.

**queues.memory.priority**

The default value is 5.

Priority of the memory dequeuing thread (aka Memory Reader Thread).

**TEXTSTRING\_SEQ\_BLOCK**

The default value is 10.

This value specifies how many keys the primary key generator should cache in memory for the table TEXTSTRING. A higher number increases efficiency by reducing the number of calls to the database to fetch the next block of keys. However this may cause large gaps in the primary key sequence, because when the archive agent is stopped, all unused keys in memory are discarded.

**PROBES\_SEQ\_BLOCK**

The default value is 10.

This value specifies how many keys the primary key generator should cache in memory for the table PROBES. A higher number increases efficiency by reducing the number of calls to the database to fetch the next block of keys. However, this may cause large gaps in the primary key sequence, because when the archive agent is stopped, all unused keys in memory are discarded.

**REQUEST\_SEQ\_BLOCK**

The default value is 20.

This value specifies how many keys the primary key generator should cache in memory for the table REQUEST. A higher number increases efficiency by reducing the number of calls to the database to fetch the next block of keys. However, this may cause large gaps in the primary key sequence, because when the archive agent is stopped, all unused keys in memory are discarded.

**REQUESTDATA\_READER\_SLEEPTIME**

The default value is 2000 milliseconds.

Time for which the Request reader thread sleeps if no data is available in the request/method queue for dequeuing (queue is empty). After this time elapses, the thread wakes up and re-checks the queue to see if data is available in the queue. The request data reader reads data from the request/method queue and processes the data and puts it in a temporary request buffer. A different thread (Request Data writer thread) reads data from this buffer and inserts it into the database.

**REQUESTDATA\_WRITER\_SLEEPTIME**

The default value is 1000 milliseconds.

Time interval at which the Request writer thread will be inserting data into the database. Request writer thread reads data from the request buffer and inserts it into the database.

**REQUESTDATA\_BUFFER\_THRESHOLD**

The default value is 10.

If the number of requests in the request buffer exceeds this size, the writer thread wakes up and inserts requests into the database irrespective of the REQUESTDATA\_WRITER\_SLEEPTIME property.

### **BATCH\_THRESHOLD**

The default threshold is 10000 inserts in a batch.

This specifies the maximum number of database inserts that can be executed in a batch. When this number is exceeded, the database insert buffer is flushed to avoid running out of memory space.

### **AVAILABILITY\_READER\_SLEEPTIME**

The default value is 2000 milliseconds.

Time that the Availability reader thread sleeps if no data is available in the availability queue for dequeuing (queue is empty). After this time elapses, the thread wakes up and re-checks the queue to see if data is available in the queue.

### **PMI\_READER\_SLEEPTIME**

The default value is 2000 milliseconds.

Time that the PMI reader thread sleeps if no data is available in the PMI queue for dequeuing (queue is empty). After this time elapses, the thread wakes up and re-checks the queue to see if data is available in the queue.

### **MEMORY\_READER\_SLEEPTIME**

The default value is 2000 milliseconds.

Time that the memory reader thread sleeps if no data is available in the memory queue for dequeuing (queue is empty). After this time elapses, the thread wakes up and re-checks the queue to see if data is available in the queue.

### **TRAP\_SEQ\_BLOCK**

The default value is 100. This value specifies how many keys the primary key generator should cache in memory for the tables TRAP\_HISTORY, TRAP\_THREAD\_DUMP, TRAP\_METHOD\_DETAIL, and TRAP\_STACK\_TRACE. A higher number increases efficiency by reducing the number of calls to the database to fetch the next block of keys. However, this may cause large gaps in the primary key sequence, because when the archive agent is stopped, all unused keys in memory are discarded.

## **ve.properties**

The following are the properties listed in the ve.properties file:

### **kernel.codebase**

The value of this property is filled in during installation time by the installer. It specifies where the Kernel codebase can be found.

### **kernel.rfs.address**

The value of this property is filled in during installation time by the installer. This is used by the VE to locate the Kernel components.

### **com.cyanea.ve.report.sql.MAX\_ROW\_COUNT**

The default value is 1000. Specifying too high a value may cause the VE to run out of memory when running the reports.

This specifies the maximum number of rows that should be retrieved from the database when displaying the detail reports (request, method, SQL) in the PAR component.

## **md.properties**

The following are the properties listed in the md.properties file:

### **component.listen.port**

Default port number is 9106. The port number that the MD listens to accepts incoming requests. Check that this port number is available in the machine running the Application Monitor. You may also change this value to any other port number available on the machine.

Example:

```
component.listen.port=9106
```

#### **component.id**

A unique identifier for a component in the Application Monitor. This ID is generated during the installation.

Example:

```
component.id=30ca500d-3798-d601-3ccc-00065b0ea116
```

#### **template.directory**

Default value is [installation directory]/etc. Unless you want to move the template directory to a different path, changing this value is not recommended. The directory that keeps the message template files of the Application Monitor. When sending an outgoing message, the MD searches for the target message template file based on this value.

Example:

```
template.directory=/opt/cyaneaone/etc
```

#### **smtp.host**

Default value is the IP address or hostname that is entered in the installation program. The IP address or hostname of the SMTP server. In order to send an outgoing email, an SMTP server is required by the MD.

Example:

```
smtp.host=192.168.3.12
```

## **kl.properties**

The following are the properties listed in the kl.properties file. If you have more than one instance of the Kernel, you will have kl1.properties, kl2.properties, and so on.

#### **cyanea.usecustomrmi**

Specifies whether or not to use custom RMI socket factory. The default value is true. The custom RMI socket factory observes a timeout limit, defined in the KernelRMIThreadPool.threadpool.threadTaskTimeout property.

Example:

```
cyanea.usecustomrmi=true
```

#### **KernelRMIThreadPool.threadpool.threadTaskTimeout**

If the custom RMI socket factory is used (which is specified in the cyanea.usecustomrmi property,) this property defines the timeout limit for RMI socket connections, in milliseconds. Default is 3 minutes.

Example:

```
KernelRMIThreadPool.threadpool.threadTaskTimeout=6000
```

#### **kernel.codebase.port**

Default port numbers are 9122 and 9123, when there are two instances of the Kernel up and running. Check that these port numbers are available in the machine running the Application Monitor. You may also change this value to any other port number available on the machine. The port number that the codebase server in the Kernel listens to accepts incoming requests.

Example:

```
kernel.codebase.port=8122
```

#### **kernel.codebase.libpath**

Default value is [installation directory]/lib. Unless you want to move the codebase directory to a different path, changing this value is not recommended.



The directory that keeps the Java binary files of the Application Monitor. After accepting a request, the codebase server in the Kernel searches for the target Java binary file based on this value.

Example:

```
kernel.codebase.libpath=/opt/cyaneaone/lib
```

#### **kernel.rmi.port**

Default port numbers are 9118 and 9119, when there are two instances of the Kernel up and running. Check that these port numbers are available in the machine running the Application Monitor. You may also change this value to any other port number available on the machine. The port number that the RMI-based Kernel listens to that accepts incoming requests.

Example:

```
kernel.rmi.port=8118
```

#### **kernel.rfs.port**

Default port numbers are 9120 and 9121. Check that these port numbers are available in the machine running the Application Monitor. You may also change this value to any other port number available on the machine. The port number that the RFS server in the Kernel listens to that accepts incoming requests. The RFS server returns the stub object of the Kernel to the components in the Application Monitor. With the stub object, the components establish a connection with the Kernel and then make use of the services that the Kernel provides.

Example:

```
kernel.rfs.port=8120
```

#### **contract.duration**

Default value is 15 seconds. Changing this value to below 10 seconds is not recommended. Valid values are 5 to 60 seconds. The time interval in seconds that each component in the Application Monitor renews the contract with the Kernel. Each component in the Application Monitor needs to maintain a heartbeat with the Kernel; this is done by renewing a time-limited contract with the Kernel. The Kernel determines the availability of services in the Application Monitor based on the heartbeat.

Example:

```
contract.duration=15
```

#### **avm.max.renewal.failure**

Default value is 1. Changing this value to higher than 3 is not recommended. Valid values are 1 to 5. The maximum number of times that a component fails to renew the contract before the Kernel determines that the component is not available.

Example:

```
avm.max.renewal.failure=1
```

#### **component.id**

The component ID is a unique identifier of a component. When there are two instances of the Kernel up and running, a unique ID is assigned to each instance.

## **datacollector.properties**

The following are the properties listed in the datacollector.properties file:

#### **kernel.codebase**

The value of this property is filled in during installation time by the installer. It specifies where the Kernel codebase can be found.

**kernel.rfs.address**

The value of this property is filled in during installation time by the installer. This is used by the VE to locate the Kernel components.

**probe.library.name**

Default value is cyanea. This property specifies the name of the native shared library which the Data Collector needs to run. If the value of the property is cyanea, the Data Collector searches for a shared library with the name libcyanea.so on UNIX<sup>®</sup> platforms and cyanea.dll on the Windows<sup>®</sup> platform. In normal cases, this property does not need to be specified or changed from the default. Only when the user needs to run a native shared library with a different name does this property need to change.

Example:

```
probe.library.name=cyanea
```

**internal.probe.event.packet.size**

Default value is 5000 or (5000 X 1024 bytes). Changing below the default of 5000 is not recommended. Valid values are 500 to 4000000 (or up to available process memory on the server).

This property specifies the size of the Data Collector's internal send buffer. The send buffer controls how much data the Data Collector can be sent to the PS at a given time. In normal situations, this property does not have to be changed, as the default send buffer size is more than adequate. However, if the user is seeing a problem with the amount of data the Data Collector sends to the PS, this property allows the user to configure the size of the send buffer.

**internal.memory.limit**

Default value is 60 (MB).

This property limits the amount of memory the Data Collector may use.

**internal.memory.accept.threshold**

Default value is 2 (MB).

This property specifies the minimum free memory after which the Data Collector starts accepting data once it reaches the upper limit, as specified by the property internal.memory.limit.

**internal.url.limit**

Default value is 1000.

This property controls the maximum URL length accepted by Data Collector.

**internal.sql.limit**

Default value is 1000.

This property controls the maximum SQL length accepted by Data Collector.

**internal.probe.event.queue.size.limit**

Default value is 5000.

This property controls the maximum size of the queue of events maintained by the Data Collector. When the queue is full, the Data Collector will drop events.

**com.cyanea.jcaprops**

If you are using WebSphere, ManagedConnectionFactory name published in the JCA driver that needs to be managed. If multiple JCA adapters are there, they must be separated by a colon ":" or a semicolon ";".

Example:

```
com.cyanea.jcaprops=com.ibm.ssys.helloworldra.HelloWorldManagedConnectionFactoryImpl
```

**proxy.host, proxy.port**

If there is a firewall between the managing server and the Data Collectors, by default two additional ports need to be opened in the firewall for each Data Collector on a host. However, only one port is needed for all Data Collectors if the communication is done through proxy servers. A proxy server should be run on each Data Collector host.

In order for the Data Collector to communicate with the managing server through the proxy, the properties `proxy.host` and `proxy.port` must be set in the `datacollector.properties`. Property `proxy.host` is the host name of the computer which runs the proxy server.

Example:

```
proxy.host=qaapp-aix-s01
```

**proxy.port**

Property `proxy.port` is a port opened in the firewall for the proxy server. As long as the proxy servers reside on different hosts, they can share the same port number. Example:

```
proxy.port=8080
```

**deploymentmgr.rmi.port**

It is not necessary to define the property `deploymentmgr.rmi.port` if you are running a standalone WebSphere. This property is needed for WebSphere 5 clusters or application servers controlled by a deployment manager.

Example:

```
deploymentmgr.rmi.port=<Deployment Manager RMI (bootstrap) port>
```

**deploymentmgr.rmi.host**

It is not necessary to define the property `deploymentmgr.rmi.host` if you are running a standalone WebSphere. This property is needed for WebSphere 5 clusters or application servers controlled by a deployment manager. Example:

```
deploymentmgr.rmi.host=<Deployment Manager host>
```

**Network Agent Socket Reset Time**

The default is no reset. Time interval after which the connection between the Data Collector and the Publish Server will be reset.

```
# networkagent.socket.resettime=-1
```

**IBM® JVMs**

The following properties pertain to platforms that use IBM JVMs with WebSphere AIX and Windows.

**internal.app.ext.dir**

This variable specifies the full path name of the installed applications directory for the application server.

The directory contains the application's set of classes (within Ears) installed on the application server.

When `CY_HIGH_PERF_TRACING` is set to YES, the data collector will only apply method tracing to the application's set of classes.

WAS 4.0.X:

Uses default set to `$WAS_HOME/installedApps`

WAS 5.0:

Uses default set to `$WAS_HOME/installedApps/$WAS_CELL`

**internal.bootstrap.file**

Specifies a Java property file which contains the two properties `classes_to_trace` and `classes_not_to_trace`.

**internal.tracing.mode**

Specifies either "new" or "classic". Default is "new". Setting internal.tracing.mode to "classic" provides better performance, but prohibits dynamic change to monitoring level L3.

**classes\_to\_trace=...**

This property specifies an additional list of classes which are to be method traced. These classes are in addition to those in the user's application set. The property only applies when CY\_HIGH\_PERF\_TRACING is set to YES.

**classes\_not\_to\_trace=...**

This property specifies a list of classes which are not to be method traced. The property only applies when CY\_HIGH\_PERF\_TRACING is set to YES.

A class list consists of class specifiers and uses either : or ; as separators between the class specifiers.

A class specifier consists of the package name plus its class name and allows use of \* as wildcard to describe matching packages or classes.

Example of a property setting:

```
trace.classlist=java/net/SocketInputStream:com/cyanea/kernel/common/HTTPClassLoader:java/net/*:java/util/Vector:java/io/*:com/ibm/websphere/pmi/server/LoadDataImpl:java/util/*
```

## ps.properties

The following are the properties listed in the ps.properties file. If you have more than one instance of the PS, you will have ps1.properties, ps2.properties, and so on.

**PORT\_PS**

The port where the PS listens for socket connections.

**TIMEOUT\_LIMIT**

Default value is 5 minutes. Value is in minutes. Time after which the request will be deleted from the PS's local cache if no activity occurred on this request (no request/method records received by PS) for TIME\_OUT interval.

**ve.host**

The hostname of the J2EE application server upon which the Cyanea web application has been deployed. This property is used to generate URLs in outgoing trap email messages. If omitted, this property defaults to the current hostname.

**ve.port**

The HTTP port of virtual host of the J2EE application server upon which the Cyanea web application has been deployed. This property is used to generate URLs in outgoing trap email messages. If omitted, this property defaults to 80.

**cyanea.aa.recheckTime**

Default value is 15000 milliseconds. Values are in milliseconds.

If PS loses the connection to AA, the frequency at which it will try to reconnect to AA.

**cyanea.ps.dataSource.useQueue**

Default value is true. Values can be true or false.

Use separate threads for reading data from the socket (data written by probe) and for processing that data.

**cyanea.ps.dataSource.maxQueueLength**

Default value is 1000000.

Max allowable queue size if the `cyanea.ps.dataSource.useQueue` property is set to true. This sets the max queue size between thread reading data from the socket and thread processing that data.

**cyanea.socket.timeout**

Default value is 3 hours. Values are in milliseconds.

PS will wait on the socket for this length of time interval when it is trying to read data from the sockets. If no data is received in this interval, PS will close the socket. If this property is set to 0 PS will wait infinitely. (This setting is advisable so that PS won't close sockets to the probe if no activity occurs on the probe for a long time.)

**cyanea.socket.keepAlive**

Default value is true. Valid values are true or false.

When the keepAlive option is set for a TCP socket and no data has been exchanged across the socket in either direction for 2 hours, TCP automatically sends a keepAlive probe to the peer.

**cyanea.socket.sendBuffer**

Default value is 65356.

This sets a suggested size for the underlying buffers used by the platform for outgoing network I/O. This is a suggestion to the Kernel from the application about the size of buffers to use for the data to be sent over the socket. Increasing buffer size can increase the performance of network I/O for high-volume connection, while decreasing it can help reduce the backlog of incoming data.

**cyanea.socket.recvBuffer**

Default value is 65356.

This sets a suggested size for the underlying buffers used by the platform for incoming network I/O. This is a suggestion to the Kernel from the application about the size of buffers to use for the data to be received over the socket.

**cyanea.socket.noDelay**

Default is true. Valid values are true or false.

True=Disable Nagle's algorithm for this connection. Written data to the network is not buffered pending acknowledgement of previously written data. False=Enable Nagle's algorithm.

**cyanea.ps.aa.maxQueueLength**

Default value is 10000.

Maximum queue length between the PS and the AA. The PS buffers events in a queue if the PS is trying to write data to the AA faster than the AA can read or if the AA is down.

**com.cyanea.ps.recent\_activity\_max\_requests**

Default value is 100.

This is the number of completed requests that the user wants to view at one time. Maximum number of recent activity data applies to each server. When the queue is full, the newest request data will replace the oldest data.

**DAILY\_STATS\_ARCHIVE\_INTERVAL**

Default value is 15 minutes.

Interval at which we will be collecting SMF daily statistics. We put the data collected at this interval in a working table in the database and push that data to the actual table in the database if we notice that the probe goes down or if it's a midnight poll.

**user.name**

The return email address to be used on all outgoing trap email. By default, it is the username of the UNIX user which starts the message dispatcher process.

**sam1.properties**

The following are the properties listed in the sam1.properties file, which is the properties file for the Global Publish Server.

**PORT\_SAM**

The port where the GPS listens for socket connections.

**SOCKET\_BACKLOG**

The default value is 5 connections.

The size of the backlog the GPS will allow while this dispatching connection attempts to Listeners.

**MAX\_SOCKET\_WRITES**

The default value is 100,000 sockets.

The number of times the PS can write to a socket before it closes the socket and opens a new one.

**WRITE\_TIMEOUT**

The default is 60,000 milliseconds. 0 is infinite.

For the PS to GPS socket connection

**SAM\_ID**

The component ID of the GPS, which is used by the Kernel to uniquely identify the GPS.

**CONN\_POOL\_SIZE**

The minimum default is 5. Do not recommend setting the pool size less than 5.

The number of database connection pools.

**INSERT\_BATCH\_SIZE**

The default value is 1000.

**BACKUP\_THREAD\_TIMER**

The default value is 120000 in milliseconds.

**CLEANUP\_THREAD\_TIMER**

The default value is 60000 in milliseconds.

**DEQUEUEING\_THREAD\_TIMER**

The default value is 2000 in milliseconds.

**SAMPLING\_THREAD\_TIMER**

The default value is 2000 in milliseconds.

**TOKEN\_RESOLUTION\_TIMEOUT**

The default value is 120 hours.

**ORPHAN\_RESOLUTION\_TIMEOUT**

The default value is 5 minutes.

**THREAD\_END\_TIMEOUT**

The default value is 5 minutes.

**RESOLVED\_TRANSACTIONS\_TIMEOUT**

The default value is 1 minute.

## Miscellaneous

To sustain large amounts of trace, change the following settings:

1. Increase the method count limit of the PS;

The PS will truncate method trace larger than a specified number. To increase the limit, change `methodCount=10000` to `methodCount=50000` in `/opt/cyaneaone/etc/ps1.properties` and `/opt/cyaneaone/etc/ps2.properties` files.

2. Increase the Java heap size of the AA;

The runtime memory of the AA needs to be increased as potentially bigger method trace from the PS. Set `HEAP_MIN_SIZE_ARCHIVE_AGENT=1024` and `HEAP_MAX_SIZE_ARCHIVE_AGENT=2048` in file `/opt/cyaneaone/bin/setenv.sh` to increase the Java heap size.

3. Reduce the batch size of the AA;

The AA will try to batch process a specified number of requests, the default is 10 requests. It will hit the limit of the DB2 driver with the large requests. Change `REQUESTDATA_BUFFER_THRESHOLD=10` to `REQUESTDATA_BUFFER_THRESHOLD=1` in `/opt/cyaneaone/etc/aa.properties`.

4. Increase the Java and JNI stack size of the AA; and

Even one big request may still hit the DB2 driver limit. To avoid it, add `-Xoss8M` and `-Xss8M` in the Java startup command. i.e. `{JAVA_HOME}/bin/java -classpath $CLASSPATH -Xoss8M -Xss8M \` in `/opt/cyaneaone/bin/aactl`.

5. Increase the query heap size of DB2.

Also increase the query heap size of the DB2 server to accommodate the large batch query from the AA. To do so, issue `db2 update dbm cfg using query_heap_sz 8000` using user `db2inst1`.





---

## Appendix D. Data Trimmer

The Data Trimmer is used to trim old data that does not need to be maintained in the database any longer. It is provided as a command line utility so that users can simply schedule CRON jobs to periodically trim appropriate data. For REQUEST and METHOD tables, this process has two parts: marking records to be deleted and deleting marked records. Be aware that running this command could take a very long time depending on the amount of data being trimmed.

This release of the Data Trimmer does not support more than one instance of the trimmer running concurrently.

You will need to use the cyanea user account, and you should make the properties file accessible to the cyanea user. You will need the following files to use the Data Trimmer:

- `$CYANEA_HOME/lib/datatrim.jar`
- `$CYANEA_HOME/etc/datatrim.xml`
- `$CYANEA_HOME/etc/markdatadeleting.xml`
- `$CYANEA_HOME/bin/datatrim.sh`

The Data Trimmer will produce the following file after it runs:

- `$CYANEA_HOME/logs/datatrim.log`

### To execute the Data Trimmer:

1. Start the Data Trimmer process on the Managing Server (If you have your Managing Server and Database Server on two different machines, start the Data Trimmer on the Managing Server machine.), the syntax is as follows:

```
$CYANEA_HOME/bin/datatrim.sh <dbname> <dbuser> <dbpassword>
```

Where the dbname is the WSAM database name, dbuser and dbpassword combination will allow the Data Trimmer to access the WSAM database.

2. The process log file is located under  
`$CYANEA_HOME/logs/datatrim.log`
3. The default JVM minimum heap size is 128 and the maximum is 256. You can increase the JVM heap size in `datatrim.sh` for optimum performance.
4. See the property files (`markdatadeleting.xml` and `datatrim.xml`) to setup the parameters.
5. In case of failure, rerun the `datatrim.sh` script.

---

## Data Trimmer Properties Files

### `markdatadeleting.xml`

The `markdatadeleting.xml` file is found in the `<CYANEA_HOME>/etc` directory.

To schedule deleting records in the REQUEST and METHOD tables, set the following properties:

### `commitcount`

Number of records to be committed to database in a transaction (commitcount). The default for commitcount is 500. This number should be set to a number lower than the maximum number of locks allowed to be held per transaction.

#### **useoracle**

Indicator of database type (useoracle). The default for useoracle is false. If an Oracle database is used, then the value should be true.

#### **daystokeep**

Number of days to keep the data (daystokeep). If you want to trim all data up to 7 days ago, then daystokeep=7. For example, if today is July 8, 2004, and you specify daystokeep=7, the Data Trimmer will do the following:

- Find out what the current date is (July 8, 2004).
- Keep today's data.
- Start counting 7 days back from 24 hours ago yesterday (July 7, 2004) – July 1st, 2004.
- Delete all data up to June 30th, 2004, inclusive.
- Data from July 1st onwards is kept.

#### **startdate**

Start date for deleting data (startdate). Use this to specify a time period for deleting data. To delete from and not include January 1 2003, enter startDate= 01/01/03 00:00:00. If "daystokeep" property is specified, then this property will not be used.

#### **enddate**

End date for deleting data (enddate). Use this to specify a time period for deleting data. To delete up to and not including March 31, 2003, enter endDate=03/31/03 00:00:00. If "daystokeep" property is specified, then this property will not be used.

#### **datatrim.xml**

The datatrim.xml file is found in the <CYANEA\_HOME>/etc directory.

To schedule deleting records in the database, set the following properties:

#### **Properties are applied to all tables**

##### **commitcount**

Number of records to commit to database in a transaction (commitcount). The default for commitcount is 500. This number should be set to a number lower than the maximum number of locks allowed to be held per transaction.

##### **useoracle**

Indicator of a database type (useoracle). The default for useoracle is false. If an Oracle database is used, then the value is true.

**Properties are applied to each individual table, such as PMISTATS, SERVERSTATS, and VOLUMESTAT.**

##### **daystokeep**

Number of days to keep the data (daystokeep). If you want to trim all data up to 7 days ago, then daystokeep=7. For example, if today is July 8, 2004, and you specify daystokeep=7, the Data Trimmer will do the following:

- Find out what the current date is (July 8, 2004)
- Keep today's data
- Start counting 7 days back from 24 hours ago yesterday (July 7, 2004) – July 1st, 2004
- Delete all data up to June 30th, 2004, inclusive
- Data from July 1st onwards is kept

**startdate**

Start date for deleting data (startdate). Use this to specify a time period for deleting data. To delete from and not include January 1 2003, enter startDate= 01/01/03 00:00:00. If "daystokeep" property is specified, then this property will not be used.

**enddate**

End date for deleting data (enddate). Use this to specify a time period for deleting data. To delete up to and not including March 31, 2003, enter endDate=03/31/03 00:00:00. If "daystokeep" property is specified, then this property will not be used.



---

## Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described 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:

IBM Director of Licensing

IBM Corporation

North Castle Drive

Armonk, NY 10504-1785

U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation

Licensing

2-31 Roppongi 3-chome, Minato-ku

Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION • AS IS • WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you. This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you. Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including

this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation

P.O. Box 12195, Dept. TL3B/B503/B313

3039 Cornwallis Rd.

Research Triangle Park, NC 27709-2195

U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only. This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

**COPYRIGHT LICENSE:** This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to Cyanea Systems Corp., for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM or Cyanea, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM or Cyanea Systems Corp. for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's and Cyanea's application programming interfaces.

---

## Trademarks

The following terms are trademarks of the IBM Corporation or its subsidiaries in the United States or other countries or both:

*Table 4. Trademarks*

IBM	MVS
WebSphere	WebSphere Studio
z/OS	DB2
S/390	OS/390
AIX	AIX-5L
Tivoli Enterprise Console	Tivoli
USS	TSO
RACF	TME

The following terms are trademarks of other companies:

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and JavaScript are registered trademarks and Sun Solaris and Solaris are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Pentium is a registered trademark of Intel Corporation in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

HP-UX is a trademark of Hewlett-Packard Corporation in the United States, other countries, or both.

RedHat is a trademark of RedHat Corporation in the United States, other countries, or both.

Linux is a trademark of Linux Torvalds in the United States, other countries, or both.

Other company, product and service names may be trademarks or service marks of others.





---

# Index

## A

- aa.properties 35
- AIX
  - Start
    - Virtual Frame Buffer 4
- Application Monitor 5
  - Start 5
  - Start Multiple Instances 29
  - Stop 5
- Archive Agent
  - Check Status 9
  - Start 7
  - Stop 8
- Archive Agent 3
- Availability Manager
  - Check Status 9
  - Start 7

## C

- Check Status
  - Archive Agent 9
  - Availability Manager 9
  - Email Engine 9
  - Kernel 9
  - Publish Server 9
- Commands and Parameters 17
- Common Services Address Space 15
- Component Level Operations 7
- CRON jobs
  - Daily 10
- CYN1
  - Start 17
  - Stop 18

## D

- Data Collector
  - Disable 25
- Data Trimmer 11
- Database 3
  - DB2 10
  - Maintenance 10
  - Optimization 10
  - Oracle 11
- datacollector.properties 39
- DB2
  - Start 3
- DB2 Database Tuning 31

## E

- Email Engine
  - Check Status 9
  - Start 7

## G

- Global Publish Server
  - Check Status 9
  - Stop 8

## H

- HTTP Server Log Files 31

## I

- INACT CYN1 Command 19

## K

- Kernel 3
  - Check Status 9
  - Start 7
  - Stop 8
- kl.properties 38

## L

- Linux
  - Start
    - Virtual Frame Buffer 4

## M

- Message Dispatcher 3
  - Check Status 9

## O

- Oracle Database
  - Start 4

## P

- Performing Database Optimization 10
- PMI Module Settings 25
- Polling Agent 3
- Properties Files
  - aa.properties 35
  - datacollector.properties 39
  - kl.properties 38
  - md.properties 37
  - ps.properties 42
- Properties Files 35
- ps.properties 42
- Publish Server 3
  - Check Status 9
  - Start 7
  - Stop 8

## S

- SAM/GPS 3
- SPINLOG CYN1 Command 20
- Start
  - Application Monitor 5
  - Application Monitor Visualization Engine 5
  - Archive Agent 7
  - Availability Manager 7
  - CYN1 17
  - DB2 UDB 3
  - Email Engine 7
  - Kernel 7
  - Publish Server 7

Virtual Frame Buffer  
AIX 4  
Linux 4  
Visualization Engine 10  
Web Server 4

**Stop**

Application Monitor 5  
Archive Agent 8  
Availability Manager 8  
CYN1 17  
Email Engine 8  
Kernel 8  
Publish Server 8  
System Level Operations 25

**T**

TRACE CYN1 Command 20

**V**

ve.properties 37  
Virtual Frame Buffer  
Start  
AIX 4  
Linux 4  
Visualization Engine 3  
Start 5, 10

**W**

Web Server  
Start 4

**Z**

ZAP CYN1 Command 20

---

## Bibliography

- *WebSphere Studio Application Monitor  
User's Guide (SC31-6865-00)*
- *WebSphere Studio Application Monitor  
Operator's Guide (SC31-6866-00)*
- *WebSphere Studio Application Monitor  
Installation and Configuration Guide (SC31-6867-00)*
- *WebSphere Studio Application Monitor  
CICS Data Collector Product Guide (SC31-6868-00)*
- *WebSphere Studio Application Monitor  
IMS Data Collector Product Guide (SC31-6869-00)*
- *WebSphere Studio Application Monitor  
Messages and Codes (SC31-6864-00)*
- *WebSphere Studio Application Monitor  
Program Directory for  
IBM WebSphere Studio Application Monitor (GI10-3349-00)*
- *WebSphere Studio Application Monitor  
Program Directory for the  
CICS Data Collector (GI10-3350-00)*
- *WebSphere Studio Application Monitor  
Program Directory for the  
IMS Data Collector (GI10-3351-00)*



---

# Readers' Comments—We'd Like to Hear from You

WebSphere Studio Application Monitor  
WebSphere Studio Application Monitor Operator's Guide  
3.1

Publication No. SC31-6866-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

Thank you for your response. May we contact you?  Yes  No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

_____ Name	_____ Address
_____ Company or Organization	_____
_____ Phone No.	_____

SC31-6866-00



Cut or Fold  
Along Line


IBM Corporation  
P.O. Box 12195, Dept. TL3B/B503/B313  
3039 Cornwallis Rd.  
Research Triangle Park, NC  
U.S.A.27709-2195



Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE  
NECESSARY  
IF MAILED IN THE  
UNITED STATES

# BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

SC31-6866-00



Fold and Tape

Please do not staple

Fold and Tape

Cut or Fold  
Along Line

**Overall, how satisfied are you with the information in this book?**

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

**How satisfied are you that the information in this book is:**

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

**Please tell us how we can improve this book:**

Thank you for your response. May we contact you?  Yes  No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

_____ Name	_____ Address
_____ Company or Organization	_____
_____ Phone No.	_____

**Overall, how satisfied are you with the information in this book?**

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

**How satisfied are you that the information in this book is:**

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

**Please tell us how we can improve this book:**

Thank you for your response. May we contact you?  Yes  No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

_____ Name	_____ Address
_____ Company or Organization	_____
_____ Phone No.	_____







Part Number: 5697-J18

SC31-6866-00



(1P) P/N: 5697-J18

