


An Industry Around the Tivoli Framework: Examples from the 10/Plus Association

Stefan Uelpenich, Gary Emmerton, Martin Keller, Andreas Rohrbach



International Technical Support Organization

<http://www.redbooks.ibm.com>



International Technical Support Organization

SG24-2122-00

**An Industry Around the Tivoli Framework:
Examples from the 10/Plus Association**

February 1998

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix B, "Special Notices" on page 297.

First Edition (February 1998)

This edition applies to Tivoli's TME 10 family of products and partner products from Tivoli's 10/Plus Association for use with UNIX and Windows NT operating systems.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 678
P.O. Box 12195
Research Triangle Park, NC 27709-2195

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

© **Copyright International Business Machines Corporation 1998. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	vii
Tables	xiii
Preface	xv
The Team That Wrote This Redbook	xv
Comments Welcome	xvii
Chapter 1. Introduction	1
1.1 What Is TME 10?	1
1.2 What Is the 10/Plus Association?	3
1.2.1 Partner Levels	4
1.2.2 Integration Products	7
1.3 Management Working Groups	8
1.4 What Are the Benefits of the 10/Plus Association?	9
1.5 How Do Tivoli/Plus and 10/Plus Relate?	9
1.6 Why Does Tivoli Want and Need Partners?	10
1.7 10/Plus Modules from IBM	11
1.8 The Partner Modules Described in This Book	11
Chapter 2. BEST/1 from BGS Systems Inc.	13
2.1 About BGS Systems Inc.	13
2.2 About BEST/1	14
2.3 About BEST/1 and TME 10	15
2.4 Overview and Objective	16
2.4.1 Overview of the BEST/1 Architecture	18
2.4.2 Hardware and Software Prerequisites	20
2.5 Installing BEST/1 in Our Environment	21
2.5.1 Installation Process	22
2.5.2 Installation Errors	26
2.5.3 Starting BEST/1 after Installation	26
2.6 Using the Functions of BEST/1	27
2.6.1 BEST/1 Collect	28
2.6.2 BEST/1-Monitor	35
2.6.3 Analyze	47
2.6.4 BEST/1-Predict	59
2.6.5 BEST/1 Manager	71
2.6.6 BEST/1-Visualizer	79
2.7 Setting Up BEST/1+ for TME 10 in Our Environment	81
2.7.1 Overview of the TME 10 Environment	81
2.7.2 Installing BEST/1+ for TME 10	83
2.7.3 Overview of BEST/1+ for TME 10 Functionality	90
2.8 How Can BEST/1+ for TME 10 Help Us in Managing BEST/1?	98
2.9 Examples Using BEST/1+ for TME 10 in Our Environment	98
2.9.1 Deploying BEST/1 Collect Using TME 10 Software Distribution	98
2.9.2 Monitoring BEST/1 Using TME 10 Distributed Monitoring	124
2.9.3 Event Management for BEST/1 Using TME 10 Enterprise Console	135
2.9.4 Installing BEST/1-Visualizer Using an AutoPack	144
2.10 How Can BEST/1 Help Us with Capacity Management for TME 10	149
2.10.1 Analyzing TME 10 Enterprise Console Performance	150

2.10.2	TEC Stress Test Implementation	150
2.10.3	Definition of a Transaction Class for Tivoli	151
2.10.4	Collecting the Baseline Data	152
2.10.5	Running the Stress Test	152
2.10.6	Interpreting the Test Results	154
2.10.7	"What if?" Modeling	155
2.11	Summary	157
Chapter 3. DCEmgmt from Santix Software GmbH		159
3.1	About Santix Software GmbH	159
3.2	About DCEmgmt	160
3.3	What is DCE?	161
3.4	What Does DCEmgmt?	162
3.5	DCEmgmt/Security Manager	162
3.6	DCEmgmt/Event Manager	163
3.7	DCEmgmt/Cell Manager	163
3.8	DCEmgmt/DFS Manager	163
3.9	DCEmgmt Examples	163
3.9.1	Integration of DCE Management into the TME 10 Desktop	163
3.9.2	DCEmgmt/Security Manager Examples	166
3.9.3	DCEmgmt/Event Manager	169
3.10	How Does DCEmgmt Complement TME 10?	172
Chapter 4. EcoTOOLS from Compuware Corporation		173
4.1	About Compuware Corporation	173
4.2	About EcoTOOLS	173
4.3	About EcoTOOLS and TME 10	174
4.4	Overview and Objective	174
4.5	Overview of EcoTOOLS Architecture	175
4.6	Hardware and Software Prerequisites	176
4.6.1	General Requirements	176
4.6.2	Installation Requirements and Supported Databases	177
4.6.3	Software Requirements	178
4.7	Installing EcoTOOLS	179
4.7.1	Installation Overview	179
4.7.2	Installation Details	179
4.7.3	Unloading EcoTOOLS	184
4.7.4	Installing EcoTOOLS	185
4.7.5	Set Up Sybase SQL Servers for EcoTOOLS Monitoring	188
4.7.6	EcoTOOLS Agent Installation on Remote UNIX Nodes	190
4.7.7	EcoTOOLS Agent Installation on Remote Windows NT Nodes	190
4.7.8	Starting the EcoTOOLS GUI	193
4.8	Functions of EcoTOOLS	194
4.9	Defining Domains	194
4.9.1	Computer Map Navigation	197
4.9.2	Scenarios	200
4.9.3	Inventory	212
4.9.4	EcoTOOLS Logs	213
4.10	Installing EcoTOOLS Plus Module for TME 10 in Our Environment	214
4.11	Overview of the EcoTOOLS Plus Module for TME 10 Functionality	221
4.11.1	Components to Set Up Agent Stations	221
4.11.2	Components to Set Up TME 10 Enterprise Console and TME 10 Distributed Monitoring	222
4.11.3	Components to Monitor EcoTOOLS Processes	224

4.11.4 Components to Maintain EcoTOOLS	225
4.12 Examples Using EcoTOOLS Plus Module for TME 10 in Our Environment	227
4.12.1 Deploying EcoTOOLS Using TME 10 Software Distribution	227
4.12.2 Monitoring EcoTOOLS Using TME 10 Distributed Monitoring	236
4.12.3 Event Management for EcoTOOLS Using TME 10 Enterprise Console	241
4.12.4 EcoTOOLS Plus Module for TME 10 Tasks to Manage EcoTOOLS	272
4.13 How Can EcoTOOLS Complement TME 10?	282
Chapter 5. TME 10 Solution Design Involving 10/Plus Partner Products	285
5.1 The Solution Design Challenge	285
5.2 The Tivoli Approach	287
5.3 Experiences Working with 10/Plus Applications	288
5.4 Simplified Design Steps	289
Appendix A. Tivoli 10/Plus Summary	291
A.1 10/Plus Partners	291
A.2 More Information about the 10/Plus Association	293
A.3 Press Release BGS Systems/BMC Software	294
Appendix B. Special Notices	297
Appendix C. Related Publications	299
C.1 International Technical Support Organization Publications	299
C.2 Redbooks on CD-ROMs	299
C.3 Other Publications	300
How to Get ITSO Redbooks	301
How IBM Employees Can Get ITSO Redbooks	301
How Customers Can Get ITSO Redbooks	302
IBM Redbook Order Form	303
Index	305
ITSO Redbook Evaluation	307

Figures

1.	TME 10 Architecture	2
2.	TME 10 Architecture Including Partners	3
3.	The 10/Plus Association Logo	4
4.	Management Working Groups	8
5.	BEST/1 Overview	15
6.	Tivoli - BEST/1 Environment	17
7.	Best/1 Architecture	19
8.	Installation Medium Selection	22
9.	BEST/1 for UNIX Installation Menu	23
10.	BEST/1 for UNIX Installation Options Menu	23
11.	BEST/1-Predict for UNIX Customization Menu	24
12.	Changing Customization Options	24
13.	BEST/1 Multiple Version Setup Menu	25
14.	BEST/1-Monitor for UNIX Customization Menu	25
15.	BEST/1 for UNIX Collect Customization Menu	26
16.	.bestrc Script	27
17.	BEST/1 Main Menu	27
18.	Collect - Edit Domain Attributes Window	29
19.	BEST/1 for UNIX - Collect Window	30
20.	Managed Node Properties in the Collect - Edit Node Window	31
21.	Data Collection Options - Schedule	32
22.	Data Collection Options - Frequencies	33
23.	Data Collection Options - Collect	34
24.	BEST/1-Monitor Window	36
25.	BEST/1-Monitor-Add Node Window	37
26.	BEST/1-Monitor-Add Alert Window	38
27.	BEST/1-Monitor: Condition Dialog Window	39
28.	BEST/1-Monitor-Add Filter	40
29.	BEST/1-Monitor: Mail Action	41
30.	Sample Alert Mail Message	42
31.	BEST/1-Monitor: Script Action Window	43
32.	Action - SNMP Trap Configuration	44
33.	Monitor Graph Icon	45
34.	Condition Parameters for Graphs	45
35.	Disk Parameter Breakdown	45
36.	Multiple Metric Graph Button	46
37.	BEST/1-Monitor: Multiple Metric Charting Window	46
38.	Example Disk I/O Graph	47
39.	BEST/1-Analyze - Main, after Adding All Nodes and Classes	48
40.	BEST/1 Analyze with Two Managed Nodes	49
41.	Analyze - Browse Collected Data Window	49
42.	Analysis Interval Configuration Window	50
43.	BEST/1 for UNIX - Analyze - User Classes - Main	51
44.	Analyze - Edit User Class ITSO.RAL	52
45.	BEST/1 for UNIX - Analyze - Transaction Classes - Main Window	53
46.	Analyze - Create Transaction Class Window	54
47.	BEST/1 for UNIX - Analyze - Workloads - Main Window	55
48.	New Workload Properties	56
49.	Analyze - Analyze Collected Data - Options Window	57
50.	Analyze Completed Window	57

51.	Sample Summary Report from Analyze	59
52.	BEST/1 for UNIX - Predict Window	61
53.	Predict - Model Evaluation Status Window	61
54.	BEST/1 for UNIX - Predict - //garye/itso.ral.md	62
55.	BEST/1-Predict Pie Chart Detail	62
56.	BEST/1 for UNIX - Predict - Nodes - Main Window	63
57.	Predict - Edit Node rs600024 Window, Showing Node Configuration	64
58.	Predict - Edit Node rs600024 Window, Showing Service Objectives	65
59.	Predict - Edit Wkl zzz@rs600024 Window	67
60.	Predict - Model Evaluation Window	68
61.	BEST/1 for UNIX - Predict - //garye/itso.ral.md	68
62.	BEST/1 for UNIX - Predict - Plans - Main Window	69
63.	Predict - Edit Plan Window	70
64.	Predict - Run Plan Window, after Completing Scenario Run	71
65.	BEST/1 for UNIX - Manager - Main Window	72
66.	Manager - Advanced Features Window - Collect	74
67.	Manager - Advanced Features Window - Analyze	75
68.	Manager - Advanced Features Window - Predict	76
69.	Manager - Advanced Features Window - Visualizer	77
70.	Manager - Advanced Features Window - Other	78
71.	BEST/1-Visualizer Example 1	80
72.	BEST/1-Visualizer Example 2	80
73.	The TME 10 Desktop	84
74.	TME 10 Install Product Window	85
75.	BEST/1+ for TME 10 Install Options	86
76.	The TivoliPlus Icon on the TME 10 Desktop	87
77.	Select Open...	88
78.	The TivoliPlus Window	88
79.	The BEST/1+ for TME 10 Window	90
80.	The Task Icon	91
81.	The File Package Icon	91
82.	The Profile Manager Icon	91
83.	The TME 10 Distributed Monitoring Profile Icon	92
84.	TME 10 Distributed Monitoring Indicator Collection Icon	92
85.	Collect Setup on Managed Nodes	93
86.	TME 10 Enterprise Console and TME 10 Distributed Monitoring Setup for BEST/1	94
87.	TME 10 Distributed Monitoring and BEST/1	95
88.	The Indicator Window	96
89.	TME 10 Components for Using BEST/1	97
90.	Profile Managers to Distribute BEST/1	100
91.	Profile Manager for HP-UX Machines	101
92.	Icon for the Setup BEST/1 Collect Task	102
93.	Properties of the Setup BEST/1 Collect Task	102
94.	Context Menu of Setup BEST/1 Collect Icon	103
95.	Parameters for the File Package	104
96.	Failure in File Package Generation	105
97.	Success File Package Generation	105
98.	The File Package Icon	106
99.	Open File Package Properties	106
100.	File Package Properties	107
101.	Excerpt of the File Package Definition @Collect_accessories	108
102.	Unix Options Window	109
103.	Open Subscribers for All File Packages	109

104. Subscribers Window	110
105. Open the Subscribers Window for a Certain File Package	110
106. Select the TME 10 Software Distribution Notice Group	111
107. The TME 10 Software Distribution Notice Groups Window	112
108. Log Distribution File after Successful Distribution	112
109. Notice Reporting a Failed Distribution	113
110. The Log Distribution File in Case of an Error	113
111. Select All Necessary Installation Scripts	114
112. File Package Properties Window	115
113. Excerpt of Modified collect_install_after.sh	116
114. Select Modify	117
115. Edit Job Window for BEST/1 Collect Task	118
116. Select Run Job...	118
117. Task Arguments	119
118. Output of Task BEST/1 Collect	120
119. Active BEST/1 Collect Processes	120
120. BEST/1 Collect Files	121
121. BEST/1 Collect Log File	121
122. BEST/1 Collect Log File bgscollect.log	122
123. Stop BEST/1 Collect Icon	122
124. Arguments for the Stop BEST/1 Collect Task	123
125. Output from Stop BEST/1 Collect Task	123
126. BEST/1 Collect Log File after Process Stopped	124
127. Select Properties	125
128. BEST/1+ for TME 10 Monitors	126
129. Properties of BEST/1 Daemons Available Monitor	128
130. Properties of BEST/1 File System Full Monitor	129
131. Schedule Definition	129
132. Select Subscribers	130
133. Subscribers of the Monitor Profile	130
134. Select Distribute	131
135. Distribute Profiles Window	131
136. Output from odstat Command	132
137. Output from the wlseng Command	133
138. Response Level Warning	135
139. Response Level Critical	135
140. Rule Bases before Setup BEST/1 EventServer	136
141. Event Groups before Running Setup BEST/1 EventServer Task	136
142. Event Sources before Setup BEST/1 EventServer	137
143. Event Server Setup Parameters	138
144. Setup BEST/1 EventServer Task Output	139
145. BAROC Files of Best1 Rule Base	140
146. Rule Bases after Running Setup Up BEST/1 EventServer Task	140
147. Event Groups after Setting Up BEST/1 EventServer	141
148. Event Sources after Setting Up BEST/1 EventServer	141
149. Events Arriving at TEC Server	142
150. Events of best1_PLUS Event Group	143
151. BEST/1 Test Event Pushed into the BEST/1 Monitor Logfile	144
152. AutoPack Control Center	145
153. Create Baseline Snapshot Window	146
154. Scan Status Window	146
155. Install Software Window	147
156. Setup Type	147
157. Build AutoPack Window	148

158. AutoPack Properties Window	148
159. AutoPack Window	149
160. Visualizer Installed on the Target System	149
161. Stress Test Script	151
162. Analyzing TEC Workload	152
163. wtdumpri Output	153
164. Analyzing TEC Performance	154
165. Analyzing Transaction Classes	154
166. Analyzing Processes	155
167. "What if?" Modeling	156
168. Selecting a CPU Type	157
169. DCEmgmt Sample Screen	160
170. Creating a DCEmgmt Cell Object	164
171. DCE Cell Object on the TME 10 Desktop	165
172. DCE Resources on the TME 10 Desktop	165
173. Example TME 10 User Administration Profile for DCE Groups	166
174. Edit Group Profile Window	167
175. DCE User Profile	168
176. User Properties Window	169
177. TME 10 Distributed Monitoring Profile for DCE	170
178. TEC Adapter for DCE	171
179. EcoTOOLS Components on Test Systems	176
180. Setting EcoTOOLS Installation Directory	180
181. Main Menu - EcoTOOLS Installation	181
182. EcoTOOLS Configuration Menu	181
183. System Type Selection Menu	183
184. Current Configuration Screen	184
185. Transfer Agent Software Confirmation	185
186. Install EcoTOOLS Menu	186
187. Choose Destination Location Window	191
188. Choose Destination Location Window	192
189. Services Window	192
190. Service Window for EcoTOOLS Agent Dispatcher	193
191. EcoTOOLS 3.4.1: MainWindow	195
192. EcoTOOLS : Define Window	196
193. EcoTOOLS: Computer Map - ITSO.RAL(Sybase) Window	198
194. Client (left) and Server (right) Icons from Computer Map	198
195. EcoTOOLS: Client-Server Map - ITSO.RAL(Sybase) Window	199
196. EcoTOOLS: Client Information - ITSO.RAL: rs600024: VICTOR Window	200
197. EcoGRAPH: Test_Scenario: rs600024	201
198. EcoMeter 3.4.1: root@rs600024	202
199. EcoTABLE 3.4.1:MainWindow	203
200. View Running Scenarios Window	204
201. EcoTOOLS: Edit Scenario Window	205
202. EcoTOOLS: Scenario Time & Properties Window	206
203. Agent Parameters Section of EcoTOOLS: Edit Scenario Window	207
204. Event Actions Section of EcoTOOLS: Edit Scenario Window	207
205. Data Output Section of EcoTOOLS: Edit Scenario Window	209
206. EcoTOOLS: Running Scenarios	211
207. EcoTOOLS: Databases Window	212
208. EcoTOOLS: Computer Inventory - ITSO.RAL: rs600026 Window	212
209. EcoTOOLS Log Window	213
210. The TME 10 Desktop	215
211. Install Product Window	216

212.	Product Install Window	217
213.	Product Install Window	218
214.	TivoliPlus Window	218
215.	EcoTOOLS Plus Module for TME 10 Window	220
216.	Components to Set Up Agent Stations	221
217.	Components to Set Up TME 10 Enterprise Console and TME 10 Distributed Monitoring	222
218.	Components to Monitor EcoTOOLS Processes	224
219.	The Sentry Indicator Window for EcoTOOLS	225
220.	Components to Maintain EcoTOOLS	226
221.	The Configure Agent Station Filepack Icon	228
222.	Edit Job Window	228
223.	Select Run Job from the Configure Agent Station Filepack Icon	229
224.	Agent File Package Generation for Different Machine Types	229
225.	File Package Output for AIX and HP-UX Machines	230
226.	File Packages for AIX and HP Machines	230
227.	EcoTOOLS File Package Properties	231
228.	Definition of the EcoTOOLS Agent Update for aix6000 File Package	233
229.	Subscribers for AIX File Package	234
230.	EcoTOOLS File Package Distribution	234
231.	TME 10 Software Distribution Notice Group	235
232.	Log File of EcoTOOLS File Package Distribution	235
233.	Sentry Profile Properties Window	236
234.	Edit Sentry Monitor Window	237
235.	Set Monitoring Schedule Window	238
236.	Sentry Profile Properties Window	238
237.	Sentry Alert Window	239
238.	TME 10 Enterprise Console Window	240
239.	Tivoli Plus Window	240
240.	Event Group Message Viewer Window	241
241.	Select Run Job from the Setup T/EC EventServer Icon	242
242.	Setup TEC EventServer Parameters for EcoTOOLS	243
243.	Setup EventServer Task Output	244
244.	Manual Restart of the Event Server	244
245.	BAROC Files of EcoTOOLS Rule Base	245
246.	Rule Files of EcoTOOLS Rule Base	245
247.	Rule Bases after Setup EcoTOOLS Event Server	245
248.	Event Groups after Setup EcoTOOLS Event Server	246
249.	Event Sources after Setup EcoTOOLS Event Server	246
250.	ecoplus.baroc File (Part 1 of 5)	248
251.	ecoplus.baroc File (Part 2 of 5)	249
252.	ecoplus.baroc File (Part 3 of 5)	250
253.	ecoplus.baroc File (Part 4 of 5)	251
254.	ecoplus.baroc File (Part 5 of 5)	252
255.	ECO_via_Sentry2.0.baroc File (Part 1 of 3)	253
256.	ECO_via_Sentry2.0.baroc File (Part 2 of 3)	254
257.	ECO_via_Sentry2.0.baroc File (Part 3 of 3)	255
258.	ECO_Monitors.baroc File	255
259.	Excerpt from the WTDUMPRL Command Output	256
260.	EcoTOOLS Test Events on the TME 10 Enterprise Console	256
261.	TEC Logfile Adapter and EcoTOOLS Plus Module for TME 10	257
262.	Configure EcoTOOLS Logfile Adapter Window	258
263.	Error Message When Configuring Logfile Adapter	259
264.	Install Product Window	260

265. Product Install Window	261
266. Another Error Message When Configuring Logfile Adapter	262
267. Excerpt from eco.conf File	263
268. Successful Configuration of Logfile Adapter	264
269. /etc/Tivoli/tecad/conf/tecad_logfile.conf File	265
270. EcoTOOLS Computer Map Window	266
271. EcoTOOLS Edit Scenario Window	267
272. EcoTOOLS Time and Properties Definition Window	268
273. EcoTOOLS Running Scenarios Window	269
274. New Events in /tmp/eventfile	270
275. EcoTOOLS Event/Error Log	270
276. TEC Event Display	271
277. Event Group Message Viewer	272
278. Executing EcoTOOLS Agent Help Task	273
279. Output of EcoTOOLS Agent Help Task	274
280. Backup EcoTOOLS Password File	275
281. Output of Backup Password Task	275
282. Executing Backup EcoTOOLS Database Task	276
283. Output of the Backup EcoTOOLS Database Task	277
284. Executing Restore EcoTOOLS Database Task	277
285. Output of the Restore EcoTOOLS Database Task	278
286. Executing Start/Stop ecolic Daemon Task	279
287. Output of the Start/Stop ecolic Daemon Task	279
288. Executing Start/Stop eamon Daemon Task	279
289. Output of the Start/Stop eamon Daemon Task	280
290. Executing Start/Stop ecoad Daemon Task	280
291. Output of the Start/Stop ecoad Daemon Task	281
292. Executing Run Scenario via ecocli Task	281
293. Executing List User Access Roles Task	282
294. Output of the List User Access Roles Task	282
295. EcoTOOLS Monitors	283
296. The Solution Design Challenge	286

Tables

1. TME 10 Environment	81
2. General EcoTOOLS Module Installation Requirements	177
3. EcoTOOLS Installation Prerequisites	177
4. Overview of Tivoli/Plus and 10/Plus Modules	291

Preface

This redbook positions the 10/Plus Association, the Tivoli partner initiative, and shows selected network computing management solutions that were developed by Tivoli partners.

This redbook will help you understand what the 10/Plus Association is, what products are involved and how you can use these products in conjunction with TME 10 to create an enterprise systems management solution.

This redbook is intended for professionals who need to understand or design a network computing management solution involving TME 10 and partner products.

We explain what the 10/Plus Association is, how it works and which benefits it creates, and show the following partner products as examples:

- BEST/1+ for TME 10 from BGS Systems Inc.
- EcoTOOLS Plus Module for TME 10 from Compuware Corporation
- DCEmgmt from Santix Software GmbH

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the Systems Management and Networking ITSO Center, Raleigh.

This project was designed and managed by:

Stefan Uelpenich, is an Advisory ITSO Representative, working as a project leader at the Systems Management and Networking ITSO Center, Raleigh. He applies his extensive field experience as an I/T architect and project leader to his work at the ITSO, where he writes extensively and consults worldwide on all areas of systems management. Before joining the ITSO, he worked in IBM Germany's Professional Services organization as an Advisory I/T Architect for Systems Management, consulting major IBM customers. In this role, Stefan architected the configuration management solution for one of Germany's largest client/server networks.

The other authors of this redbook are:

Gary Emmerton, is a Systems Architect for IBM Global Services Australia. He has worked extensively in systems management across multiple platforms including Windows NT, Novell NetWare and UNIX. His recent projects include helping to establish the NWS Hub for the Australia and Asia Pacific region and working as the Architect for the NWS Transition Project for the Department of Veterans' Affairs, the first Australian Federal Government department to enter into a full outsourcing contract.

Martin Keller, is President of GlobalVision GmbH, a consulting firm in Germany, providing professional services for system and network management. He has 10 years experience in network and systems management. Before founding GlobalVision he worked as a project leader and consultant for the German software companies Danet, Softlab and Santix. He has been involved in development and consulting projects using OSI network management software, SNMP development

kits and Tivoli products. Martin is currently involved as a consultant in the banking industry.

Andreas Rohrbach is an IT Consultant with SerCon GmbH, a German subsidiary of IBM. He has several years of experience in client/server environments, including systems management. His areas of expertise include UNIX and SAP R/3 basis implementation. Andreas has experience in the design, planning and installation of complex client/server environments and is currently involved in projects in the German banking industry.

Thanks to the following people for their invaluable contributions to this project:

Thomas J. Murphy, Jr.
Product Manager
BGS Systems, Inc.
Waltham, MA USA

David Wagner
Manager, UNIX Product Marketing
BGS Systems, Inc.
Waltham, MA USA

Jeanne Ikemoto
Director, Alliances
Compuware Corporation
Campbell, CA USA

Scott Fosdick
Alan Ghahramany
Rob O'Brian
Tish Peterson
Lisa White
Compuware Corporation

Michael Santifaller
President and CEO
Santix Software GmbH
Munich, Germany

Anne Cuddy
Strategic Account Manager
Tivoli Systems, an IBM Company, Raleigh

Heinz Johner
International Technical Support Organization, Austin Center

Rob Macgregor
IBM ISC Tivoli Group, Hursley, United Kingdom

Kathryn Casamento
Linda Robinson
Dave Shogren
Margaret Ticknor
Shawn Walsh

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in "ITSO Redbook Evaluation" on page 307 to the fax number shown on the form.
- Use the electronic evaluation form found on the Redbooks Web sites:
For Internet users <http://www.redbooks.ibm.com>
For IBM Intranet users <http://w3.itso.ibm.com>
- Send us a note at the following address:
redbook@vnet.ibm.com

Chapter 1. Introduction

In this chapter we introduce TME 10 and the Tivoli 10/Plus Association and explain the benefits created by this partner initiative.

1.1 What Is TME 10?

We only give a brief overview of Tivoli Management Environment (TME) 10 in this section. For a detailed explanation of TME 10 and its architecture, refer to the redbook *Understanding Tivoli's TME 3.0 and TME 10*, SG24-4948.

TME 10 is the resulting product of the merger between IBM's systems management division and Tivoli Systems, which formed Tivoli Systems, an IBM company. The product set is therefore the result of combining Tivoli's TME 3.0 with IBM's SystemView family of products.

TME 10 is based on the Tivoli Management Framework (TMF) that provides common services to TME 10 applications. On top of this framework reside a number of Tivoli core applications that utilize services provided by the TME 10 Framework. These applications themselves provide essential systems management functions that are quite generic and create the basis for management of any system, network, application or database.

These core applications are:

- TME 10 Software Distribution
- TME 10 Distributed Monitoring
- TME 10 User Administration
- TME 10 Enterprise Console
- TME 10 Inventory

Systems Management and therefore the core applications can be divided into management disciplines. The disciplines defined by Tivoli are as follows:

- Deployment
- Availability
- Security
- Operations

For example, TME 10 Software Distribution and TME 10 Inventory belong to the deployment area, whereas TME 10 Distributed Monitoring belongs to the availability area.

The following figure illustrates the TME 10 Framework architecture.

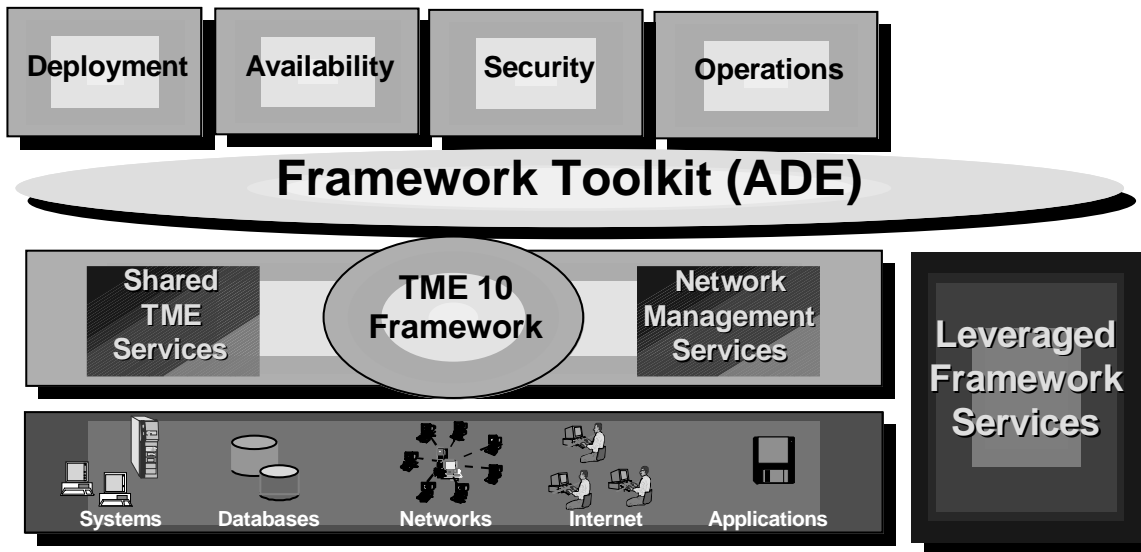


Figure 1. TME 10 Architecture

Management applications for specific components, such as applications or systems, can often be built around the TME 10 core applications.

For example, a management module for a specific application will usually require software distribution of the application's components, monitoring of the application's behavior and function, management of application events and maybe even management of application users.

Now, creating a management module for that application based on TME 10 and its core applications can take advantage of the proven strength and flexibility of the TME 10 Framework and the core applications while at the same time creating all the component-specific management parts within the module.

Not only does Tivoli provide the TME 10 Framework and the core applications but also a comprehensive and open management strategy for applications, including Application Management Specification (AMS) and corresponding toolkits to make applications manageable by TME 10.

While for any business application you can create a TME 10 management module using this approach, there are also systems management applications themselves that benefit when integrated into TME 10.

The reason for this is that Tivoli does not try to provide a management solution for every particular area involved in systems management, that is, any component in your I/T environment that needs to be managed.

Hence, there are specialized management applications available that deal with managing a specific component or focus on a specific management area. By linking these applications with TME 10 and the TME 10 Framework you can benefit from both the expertise of a vendor in managing a specific component or area and Tivoli's open and flexible framework, integrating all components into a scalable enterprise-wide management solution.

The following figure shows how partner applications complement TME 10 in creating a comprehensive systems management solution.

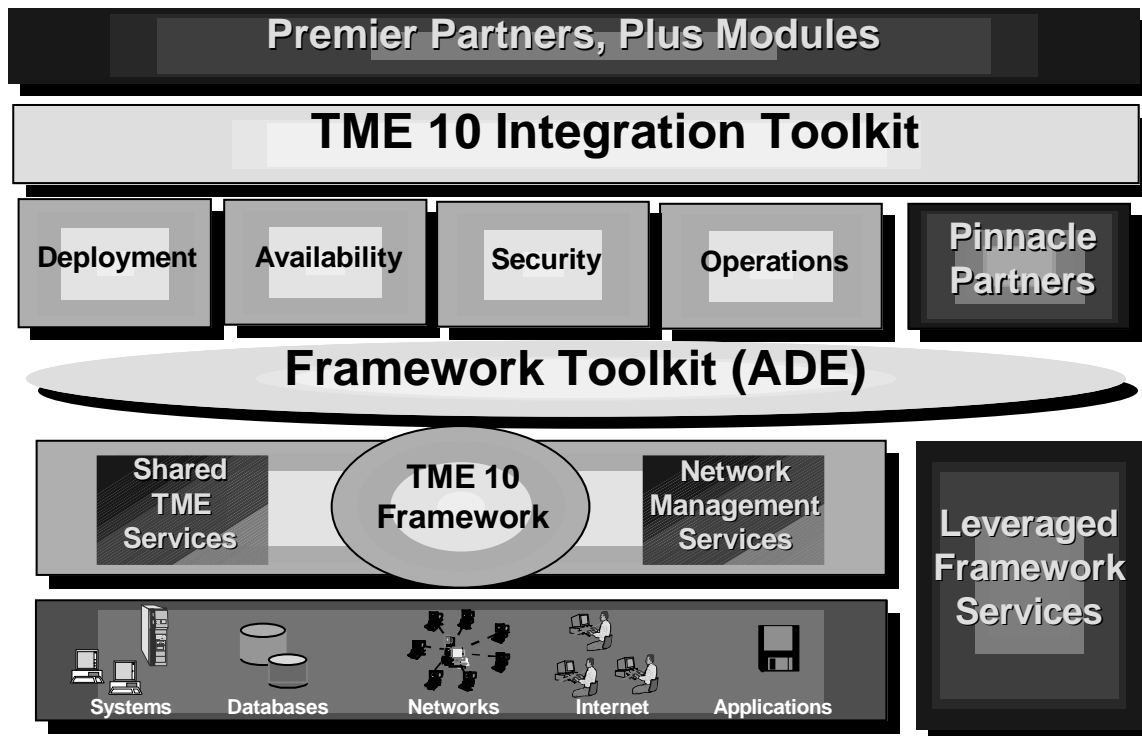


Figure 2. TME 10 Architecture Including Partners

Pinnacle partners create applications that are on the same integration level as the TME 10 core applications. They use core framework services and therefore exploit the full functionality offered by the TME 10 Framework.

Premier partners use the TME 10 Integration Toolkit to create management applications that use TME 10 application services, that is, they utilize the TME 10 core applications.

In the next sections we explain in more detail the structure of the 10/Plus Association, integration levels and integration tools.

1.2 What Is the 10/Plus Association?

The forming of the 10/Plus association was based on recognizing that no single vendor can provide a solution that covers *all* network computing management needs.

In the past, development of systems management applications often led to point products in specific areas of systems management that provided good solutions in their areas, but did not integrate well or even at all with other systems management applications.



Figure 3. The 10/Plus Association Logo

A few attempts have been made by vendors to create *integrated solutions*, but these are often closed solutions. They integrate a few products of one vendor quite nicely but are limited to these few products and are not really open to third parties.

On the other hand there have been efforts to provide integration through vendor-independent standardization groups, but they also mostly failed because there was no technology around somebody could build on.

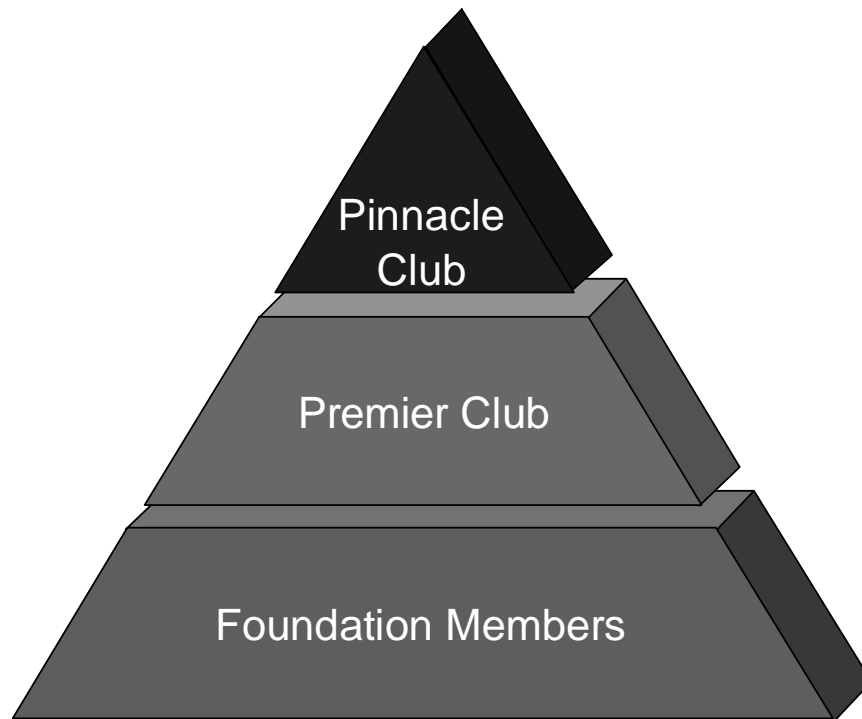
The 10/Plus Association is built around TME 10 and the TME 10 Framework that provide a solid and proven platform for integrating systems management applications. While TME 10 incorporates standards from the above-mentioned standardization groups (such as OMG's CORBA) at the same time TME 10 provides the first actual implementation of a truly open integration platform for systems management applications.

With this platform as a base TME 10 allows other vendors to link their products into TME 10 and the TME 10 Framework thus creating an integrated, best-of-breed solution for any environment.

1.2.1 Partner Levels

Tivoli's 10/Plus partners are organized in three levels:

- Foundation Members
- Premier Club
- Pinnacle Club



These levels differ in the degree of integration the members provide to TME 10. Having different integration levels recognizes the fact that certain management areas require only basic integration while others require tight integration.

Segmenting the degree of integration also allows each vendor to decide how deeply they want to integrate their products into TME 10.

If you think of different components that could be managed, you will immediately find differences in the degree of TME 10 integration they might require. For example, management modules for RDBM systems are likely to integrate at the Pinnacle level as they can make use of TME 10 Framework services in managing an RDBMS. On the other hand, a certain hardware management application might only need to be able to send events to TME 10 Enterprise Console and can therefore be integrated on the Foundation level.

As you will see in Appendix A, “Tivoli 10/Plus Summary” on page 291 the majority of partner modules today are integrating on the Premier level.

1.2.1.1 Foundation Members

Members at this level provide basic integration of their applications into TME 10, such as:

- Application launch from TME 10 desktop
- Event integration
- Integration with TME 10 menu system
- Integration with TME 10 help system
- MIB integration with TME 10 NetView

The following benefits are offered by Tivoli to members at this Partner level:

- Internal and external Tivoli Web-site listing
- Training
- Development consulting
- Access to porting centers
- Trade show participation
- Blind mailing opportunities to Tivoli customer base
- Listing in 10/Plus journal
- Ability to submit articles for publication in 10/Plus journal
- Tivoli customer event participation
- Sponsorship opportunities
- Program Manager
- Inclusion in Pointcast Channel

Pointcast is a service for casting information on the Web. Tivoli has its own channel on Pointcast for distributing information on network computing management.

- No-fee certification

The Foundation level is useful for partner products that only require a basic integration into TME 10 and has the advantage for the application vendor that the TME 10 integration module is easy to create.

1.2.1.2 Premier Club

In addition to the integration provided by Foundation members, Premier Club partners integrate with TME 10 application services such as software distribution, task management, inventory management and network management. Premier Club partners develop applications that use TME 10 application services and Tivoli core applications, while remaining focused on their respective areas of expertise.

Besides the benefits available to the Foundation level, Premier members enjoy the following additional benefits:

- Development support
- Certification support
- Tivoli Ready branding (for certified products).
- Inclusion in the SuperSystem permanent demonstration. The SuperSystem is located in Dallas and is used to demonstrate integrated systems management with TME 10 to customers.
- Road show participation
- Advertising

The Premier level is useful for almost any systems management application that requires the standard set of key functions provided by the Tivoli core applications, including software distribution, event management, user management, etc. Tools are provided by Tivoli to assist in creating a Premier level integrated module.

1.2.1.3 Pinnacle Club

Pinnacle Club partners develop products based on the TME 10 Framework. These products are highly integrated, leverage the entire breadth of TME 10, and utilize the characteristics of a Tivoli application. Framework integration vendors are Pinnacle Club members.

Besides the benefits available to Premier and Foundation members, Pinnacle partners enjoy the following additional benefits:

- Custom marketing programs
- Dedicated marketing budget
- Reselling/OEM opportunities

Note

Some of the benefits listed above and in the previous sections may require a fee to be paid to Tivoli. For detailed information contact your Tivoli/IBM representative.

1.2.2 Integration Products

To assist a 10/Plus partner to develop applications that are integrated with TME 10, there are several tools available from Tivoli. These tools differ in the level of integration into TME 10 they provide and therefore match the different levels of 10/Plus partners.

The following tools are available:

- TME 10 Integration Toolkit
- TME 10 Framework Toolkit
- Application Management Specification (AMS)
- AMS Developer's Toolkit (DevKit)

The TME 10 Integration Toolkit can be used to easily take advantage of Tivoli application services, such as software distribution, task management, inventory management and monitoring. Also, it provides a simplified approach in integrating application events with TME 10 Enterprise Console.

The TME 10 Integration Toolkit is available to Premier Club members of the 10/Plus Association.

If you want to learn more about the TME 10 Integration Toolkit, you can refer to the redbook *Developing Plus Modules for Tivoli with the TME 10 Integration Toolkit*, SG24-2007.

The TME 10 Framework Toolkit is only available to Pinnacle Club members that integrate with TME 10 on the highest level. This toolkit allows an application to exploit the full range of services offered by the TME 10 Framework.

Both toolkits provide tools to develop and test integration into TME 10. For example, provided are wizards, sample applications, APIs and documentation. Also

included are test criteria that are used for certification of an integration module by Tivoli.

AMS is a formal specification that can be used to describe the management needs of an application. As not only the content but also the format of the information in AMS is standardized, application developers have a standard way of communicating their management needs to the systems management tool.

Moreover, TME 10 core applications such as TME 10 Distributed Monitoring or TME 10 Software Distribution can import AMS-compliant Component Description Files (CDF) and provide instant management.

For example, to help create CDF files for an application the AMS Developer's Toolkit can be used. With this toolkit, the developer can specify the management needs of an application and the toolkit produces the corresponding CDF files. These can then be imported into TME 10.

For more information about this topic, you can refer to the redbook *Using the Applications Management Specification*, SG24-4809.

1.3 Management Working Groups

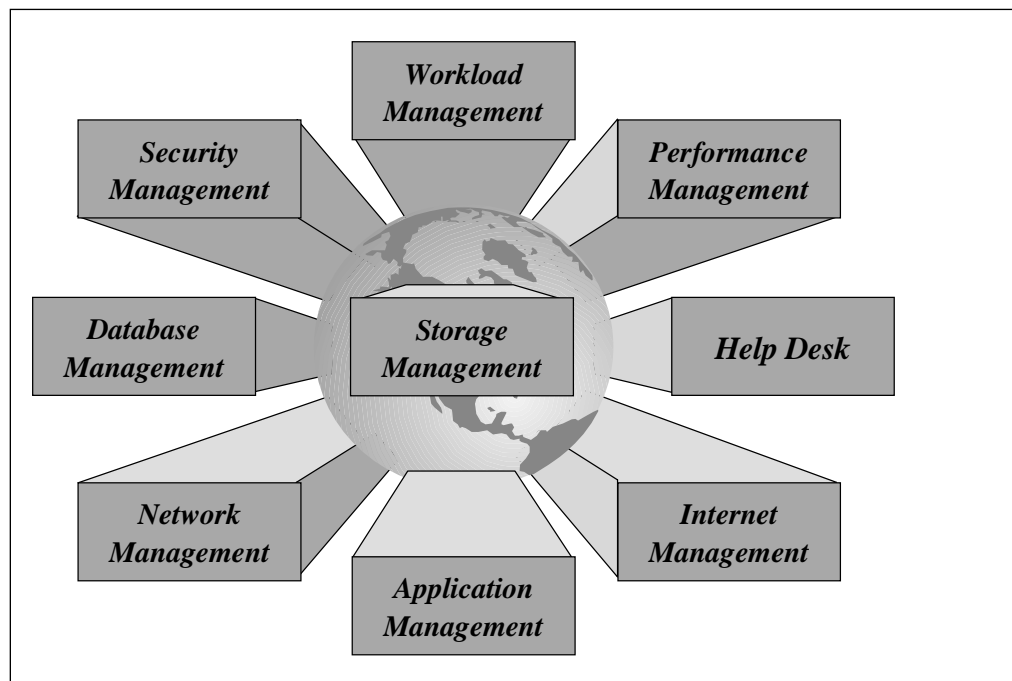


Figure 4. Management Working Groups

In the 10/Plus Association's management working groups Tivoli and its partners work on defining requirements, specifications and architectures in specific areas of systems management. The different working groups, divided by topic, are shown in Figure 4.

Results from the working group include, for example, APIs, specifications and toolkits that help integrating the workgroups dedicated area of management into a comprehensive management solution.

1.4 What Are the Benefits of the 10/Plus Association?

For the customer the benefit of TME 10 and the 10/Plus Association is that they can choose their systems management applications from a best-of-breed selection offered by Tivoli and its partners.

This selection of specific modules will likely be different among customers and reflects customer-specific needs. However, the modules they pick around the TME 10 core applications are based on the same principles, that is they are based on TME 10 and can be easily combined and integrated with other TME 10 applications.

TME 10 as the backbone provides the scalability to deploy the specific management applications at an enterprise level and holds all the components together.

The partner approach also means that development of management applications is truly driven by market demand and not by a specific vendor. As a result, vendors can specialize in their field of expertise and leverage their knowledge in a certain field and at the same time leverage TME 10 as a solid and flexible platform for the fast development of management applications.

For example, a vendor might be a hardware manufacturer that consequently has excellent knowledge about the management of its own hardware. Instead of creating a management application for hardware that stands alone, or not creating a management application at all (and thus not using unique expertise), the vendor can now easily create a management application based on TME 10 that is automatically integrated, scalable and has a lot bigger market.

Another example are systems management vendors that are highly specialized and experienced in one particular area of systems management, such as user help desk or performance management. The 10/Plus Association allows the vendor to integrate with TME 10, thus allowing a customer to integrate the vendor's application with other systems management applications and create a total solution. At the same time the vendor module enhances the value of TME 10 by increasing its scope, therefore increasing the choices for a customer and creating a systems management solution that is constantly being enhanced.

1.5 How Do Tivoli/Plus and 10/Plus Relate?

The 10/Plus Association is the result of combining the partner programs of Tivoli and the former NetView and SystemView partner initiatives.

As a result, the 10/Plus Association combines former Tivoli partners, former NetView/SystemView partners of IBM and a large number of new partners.

In the past (before Tivoli and IBM merged) there were already a number of integration modules, called Tivoli/Plus modules. For example, at that time there already was a module integrating IBM's ADSTAR Distributed Storage Manager (ADSM) into Tivoli's TME.

These modules already provided integration with TME that extended far beyond just iconic launch integration. The 10/Plus Association continues here and offers its

integration partners additional value through TME 10 Framework services and applications.

All 10/Plus modules provide a common look and feel, thus avoiding the need for a system administrator to learn numerous different procedures, techniques or user interfaces for operating different management tools.

To simplify integration into TME 10, the 10/Plus Association offers a number of toolkits that partners can use to write integration modules. These toolkits have been significantly enhanced and extended since the forming of the 10/Plus Association.

The NetView Association partners, mainly focused on network management, are now also included in the 10/Plus Association.

1.6 Why Does Tivoli Want and Need Partners?

You might wonder why Tivoli puts a high focus on partnerships instead of claiming to have a solution for all problems in systems management itself.

As mentioned before, any attempt of a single vendor to create a systems management solution that can manage anything has always and will always fail. The reason is simply that no single vendor has enough resources and knowledge to create products to manage every resource that needs to be managed in an I/T system.

While nobody would argue much about that fact, the actual reaction to that fact in the past was quite surprising in that a number of vendors (including IBM) created systems management applications that were excellent point products but did not integrate well. And, as these products were not designed to be open, attempts to "grow" them into integrated solutions also failed.

Tivoli recognized the challenge behind systems management years ago and built that knowledge into its product, TME. Instead of creating point-products, the focus was put on creating an open, standards-based architecture that would lay the groundwork for any application.

Also, Tivoli recognized that it could not create dedicated management applications for every kind of management but instead decided to put the focus on management applications that address the base requirements of systems management, such as software distribution, monitoring, event management, etc. These efforts led to the TME 10 core applications; these core applications constitute industry-leading solutions in their respective area.

Tivoli's focus is to provide all the core technology, including the TME 10 Framework and the core applications as well as selected management modules for certain tasks.

In areas of extended function, Tivoli partners provide additional solutions that typically are leaders in their respective area and, combined with TME 10, add to create the most comprehensive solution for almost any environment.

1.7 10/Plus Modules from IBM

Interestingly IBM is (and has been) a Tivoli partner itself and therefore Tivoli/IBM provided and continues to provide modules that integrate IBM products with TME 10.

For example, the Tivoli/Plus Module for ADSM integrates IBM's leading edge backup/restore solution ADSM with TME 10.

Other modules that are available from Tivoli/IBM include TME 10 Module for SAP R/3, TME 10 Module for Lotus Notes, TME 10 Module for MQSeries and DB2 Enterprise Control Center.

Besides, IBM and Tivoli are working on a number of new modules that link applications, databases and systems to TME 10 management.

More information about specific TME 10 modules can be found in the following redbooks:

- *Examples of Using the TME 10 Module for MQSeries*, SG24-2134
- *ADSM Operation and Management with TME 10*, SG24-2214
- *Managing a Notes Environment with TME 10 Module for Lotus Notes*, SG24-2104

1.8 The Partner Modules Described in This Book

In the following chapters we describe three partner applications and their integration into TME 10, namely:

- BEST/1 from BGS Systems Inc.
- DCEmgmt from Santix Software GmbH
- EcoTOOLS from Compuware Corporation

BEST/1 and EcoTOOLS are integrated with TME 10 on the Premier level, while DCEmgmt is a Pinnacle level application that was developed using TME 10 Framework technology.

In Chapter 2, "BEST/1 from BGS Systems Inc." on page 13 and Chapter 4, "EcoTOOLS from Compuware Corporation" on page 173 we describe in detail how to install BEST/1 and EcoTOOLS in our existing TME 10 environment, install and configure the Plus Modules to link BEST/1 and EcoTOOLS to TME 10, and then show practical examples of how the integrated products can be used.

We show unique features of BEST/1 and EcoTOOLS to show how these products enhance the breadth of TME 10 and show how the integration of these products adds benefit to the products *and* to TME 10.

In Chapter 3, "DCEmgmt from Santix Software GmbH" on page 159 we introduce DCEmgmt, describe its architecture and show examples of its use. In the case of DCEmgmt we do not describe the full setup and configuration, as this is described in detail in the redbook *Tivoli TME 10 Management for DCE and DFS*, SG24-2027. (Available at a later date.)

Feedback

If you have any comments or suggestions regarding this redbook, please send an e-mail to uelpenich@vnet.ibm.com.

Chapter 2. BEST/1 from BGS Systems Inc.

In this chapter we describe the BEST/1 product from BGS Systems Inc. and how it can be combined with TME 10. BEST/1 is a framework for performance management that is integrated into TME 10 using the BEST/1+ for TME 10 module from BGS Systems Inc.



We show how to install and set up BEST/1 and then explain how to integrate BEST/1 with TME 10 by installing and configuring BEST/1+ for TME 10. Once this is done, we show practical examples of how TME 10 can be used to manage BEST/1.

Note

BEST/1 is a complex client/server application that has a large number of features. Hence, in this chapter we can give only an overview of some key components in order to illustrate the functions of BEST/1. For a complete description of all features of BEST/1 refer to the appropriate product documentation or contact your BGS Systems Inc. representative.

2.1 About BGS Systems Inc.

BGS Systems Inc. is a software company located in Waltham, MA USA that develops performance management and capacity planning software for MVS, OS/390, VM/ESA, AS/400, OpenVMS, UNIX and Windows NT.

The key product is BEST/1 Performance Assurance that helps customers to measure, analyze and predict performance and capacity of distributed computer systems and applications.



Headquartered in Waltham, MA BGS Systems Inc. distributes its products worldwide through offices in 30 countries.

More information about BGS Systems Inc. can be obtained under www.bgs.com.

Important Note

By the time this redbook is written, BMC Software, Inc. has announced the execution of a definitive merger agreement with BGS Systems, Inc.

While this merger is not expected to change the name of the BEST/1 product, you should be aware that BGS Systems, Inc. might be run under the name of BMC Software, Inc. in the future.

You can find the complete press release regarding the announced merger in A.3, "Press Release BGS Systems/BMC Software" on page 294.

2.2 About BEST/1

The following figure gives a first overview of the BEST/1 architecture. Later in this chapter we discuss the components in detail.

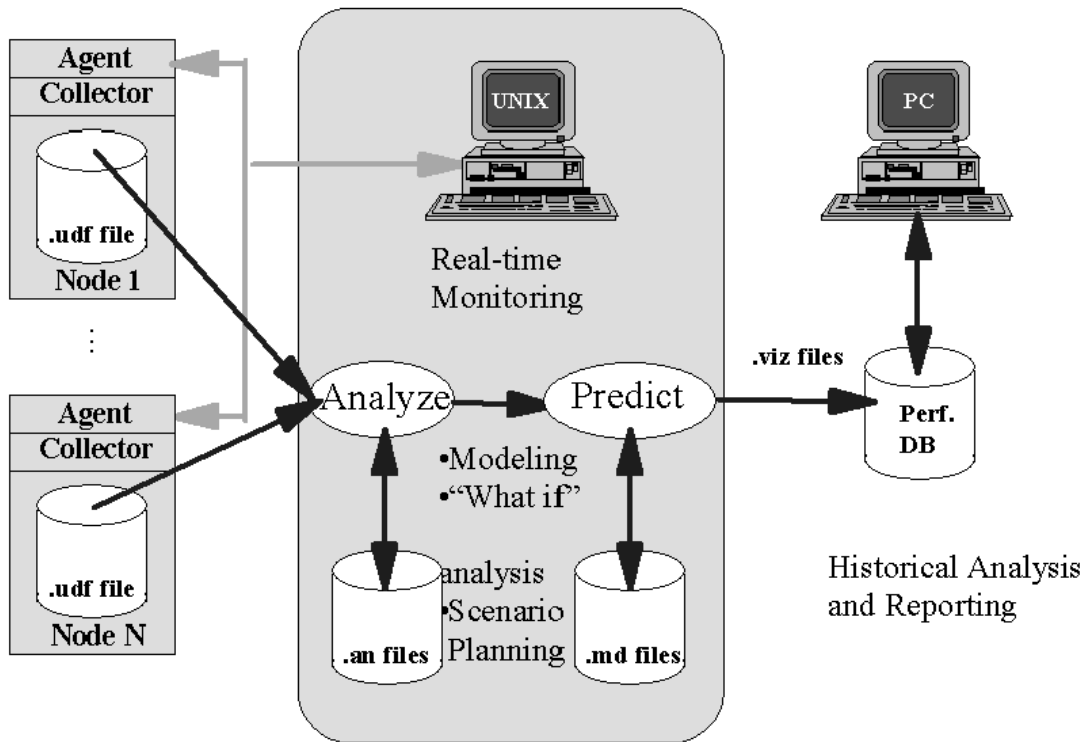


Figure 5. BEST/1 Overview

BEST/1 consists of a managing component (the BEST/1 Management Console) and an agent component (the BEST/1 Agent). The agent contains a collector component that collects performance data about various system components on the systems to be monitored, stores this information in files and then sends these files to the server.

The information is stored in a standard file format (.udf files). Once on the server, performance data can be displayed in graphs, actions can be triggered depending on thresholds and other actions can be performed on the performance data.

Further, the performance data can be analyzed in specific scenarios to see if service levels are met and also future configurations can be predicted. We show a lot of these features in this chapter.

Finally, performance data can be transformed into presentation quality graphical representations using a Windows NT-based PC (the BEST/1-Visualizer).

2.3 About BEST/1 and TME 10

BEST/1 is a solution for performance management that is available on a wide range of operating systems, including the most popular UNIX operating systems, such as AIX, HP-UX and Solaris.

The key strength of BEST/1 is its capability to analyze performance data and predict how changes in the system configuration would affect future performance.

The BEST/1+ for TME 10 module can be used to link the strength of BEST/1 in the performance management field to the strength of TME 10 in enterprise systems management.

BGS Systems Inc. is a Premier level member of the 10/Plus Association. BEST/1+ for TME 10 was developed using the Tivoli Integration Toolkit that can be used to rapidly integrate existing systems management applications into TME 10. BGS Systems Inc. is also a member of the TME 10 Management Working Group for Performance Management.

Similar to other 10/Plus modules, BEST/1+ for TME 10 provides integration with the following TME 10 functions:

- TME 10 Framework
- TME 10 Software Distribution
- TME 10 Distributed Monitoring
- TME 10 Enterprise Console
- TME 10 Tasks and Jobs
- Application launch from the TME 10 desktop

All functions are integrated in the familiar TME 10 desktop, so that all actions can be performed from within the familiar TME 10 environment.

2.4 Overview and Objective

In this chapter we first describe how to set up our environment containing the TME 10 and BEST/1 components. We then show how to use BEST/1 in combination with TME 10. We start with giving an overview of the already installed Tivoli components, and an overview of BEST/1.

An overview of our environment is shown in the following figure:

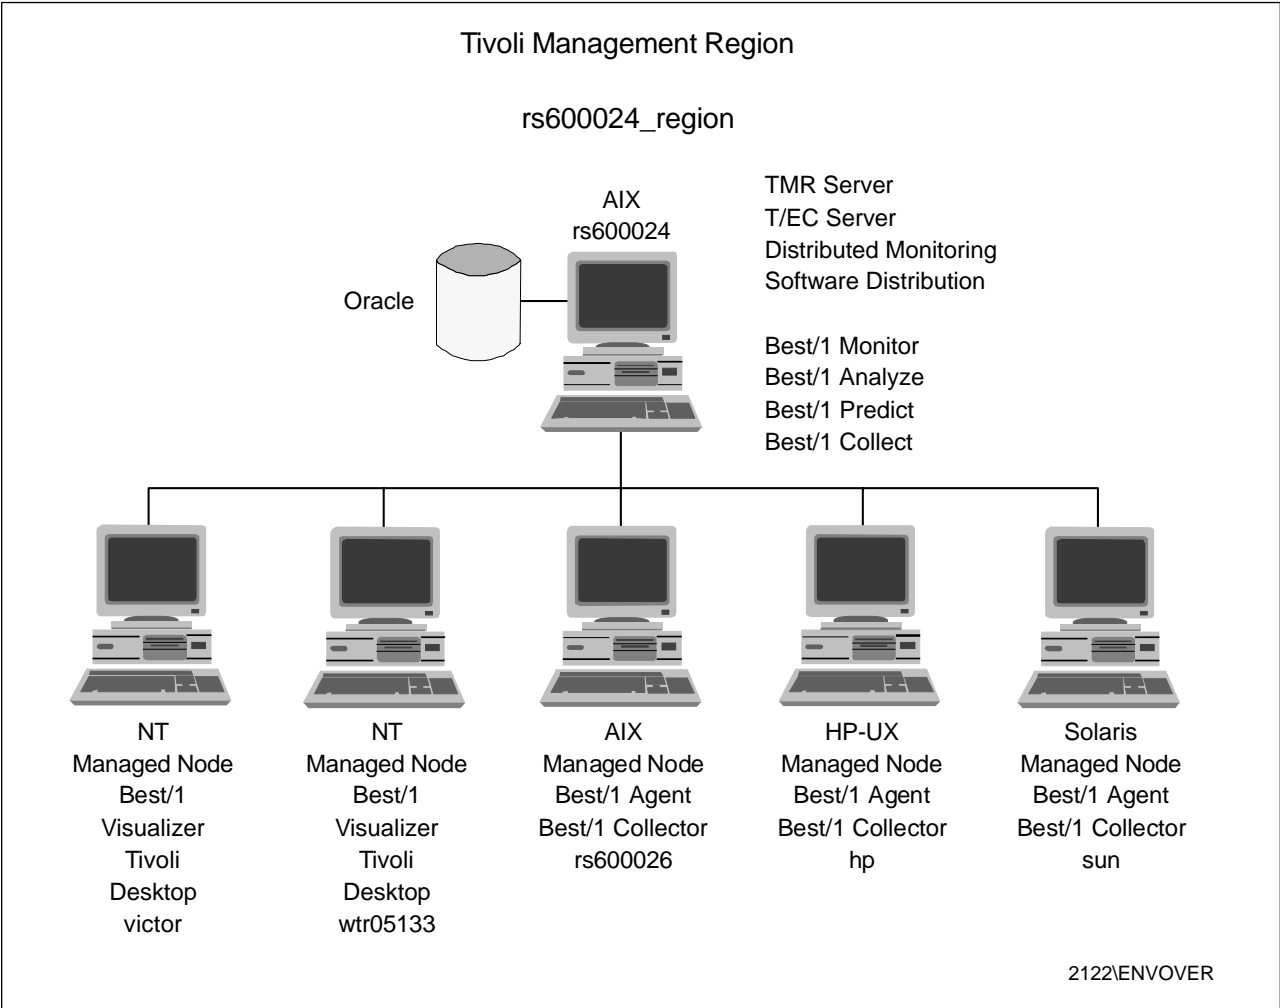


Figure 6. Tivoli - BEST/1 Environment

We use the AIX server rs600024 to act as the BEST/1 Management Console as well as the TMR server in our Tivoli environment. In addition, we use rs600026 running AIX, hp running HP-UX and sun running Solaris as managed nodes for BEST/1 as well as TME managed nodes in our Tivoli environment.

Also, we use two Windows NT PCs to install the BEST/1-Visualizer component for BEST/1 and to run the TME 10 Desktop for Windows.

Note

We do not show how to install the basic Tivoli environment, that is the TME 10 Framework, in this chapter. However, we do describe in detail how to install and use BEST/1+ for TME 10 in 2.7, "Setting Up BEST/1+ for TME 10 in Our Environment" on page 81.

The objective of this chapter is to show the operation of a solution for performance management that is well integrated in Tivoli.

To do so we perform the following steps:

1. Install the TME 10 environment on all systems involved. This is the basic Tivoli setup not described in this chapter.

2. Install the BEST/1 product and show its key features. This is described in 2.5, “Installing BEST/1 in Our Environment” on page 21.
3. Install the BEST/1+ for TME 10 product. This is described in 2.7, “Setting Up BEST/1+ for TME 10 in Our Environment” on page 81.
4. Use BEST/1+ for TME 10 to install BEST/1 client components, monitor BEST/1 resources and forward BEST/1 alerts to TME 10 Enterprise Console. This is described in 2.9, “Examples Using BEST/1+ for TME 10 in Our Environment” on page 98.

We first show how to use the basic functions of BEST/1, then show how to install BEST/1+ for TME 10 and use BEST/1+ for TME 10 to manage BEST/1. Also, we show in this chapter how to analyze the performance of a TME 10 environment including TME 10 Enterprise Console using BEST/1 in 2.10, “How Can BEST/1 Help Us with Capacity Management for TME 10” on page 149.

If you are already familiar with the BEST/1 product, you can skip the BEST/1 overview and directly go to 2.7, “Setting Up BEST/1+ for TME 10 in Our Environment” on page 81.

2.4.1 Overview of the BEST/1 Architecture

BEST/1 consists of a suite of client/server software components for the following management disciplines:

- Monitoring
- Performance management and analysis
- Capacity and performance planning

The following picture shows a more detailed overview of the BEST/1 components than Figure 5 on page 15. We explain the components in the following.

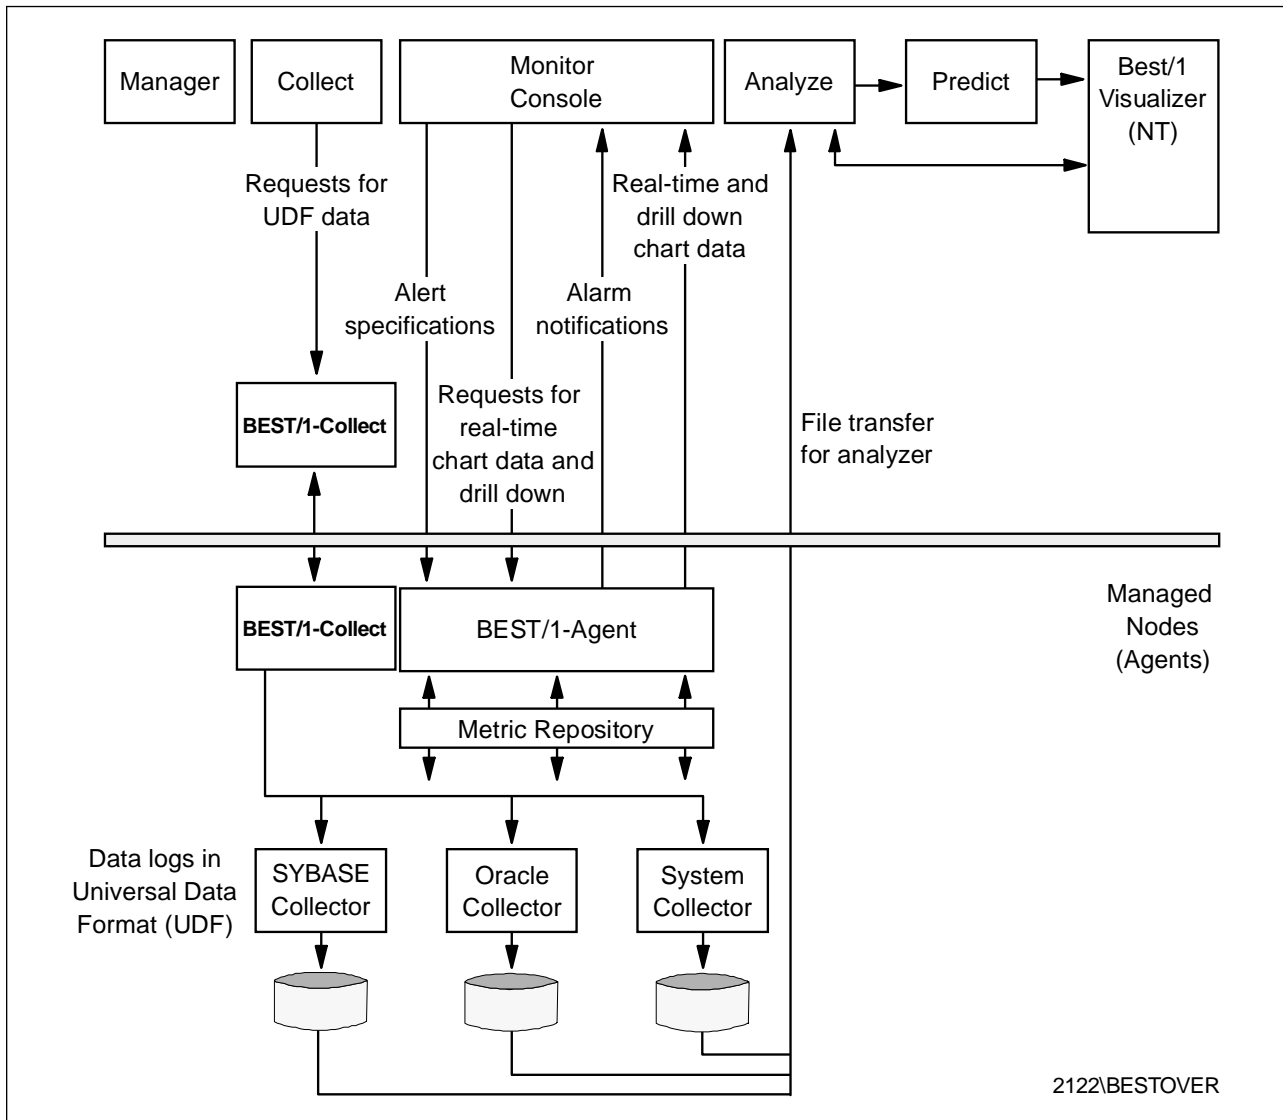


Figure 7. Best/1 Architecture

To oversee the actual performance of UNIX machines, databases and hard disks and other system components the BEST/1-Monitor console and the BEST/1 Agent components are used.

With the BEST/1-Monitor console alert specifications are defined and activated on the BEST/1 Agent. The Monitor Console has two other functions: one is to send requests for real-time chart data and the second is the drill-down function. With this you can drill down through the monitoring data to rapidly focus on problems and causes.

The BEST/1 Agents send alarm notifications back to the Monitor Console and responses to the chart and drill-down data requests.

The term managed node is also used for the components where BEST/1 Agents reside. Since managed node is also a commonly used term in the Tivoli environment, we use the term BEST/1 managed node when we refer to BEST/1 and we use managed node only when we refer to the Tivoli managed node.

For performance management BEST/1 Collectors for UNIX, Sybase and Oracle databases are used to capture performance data and store it in a file at the BEST/1 Managed Node. The data is stored in the so-called Universal Data Format (UDF). These files can be forwarded to BEST/1 Analyze for evaluation.

BEST/1 Analyze can figure out where system resources are consumed, which users and processes are responsible for the workload, and what work is being performed.

To plan future hardware and software procurement, the Predict component is used to build future scenarios. You can either assess the effects of increased load on the current hardware or you can simulate the impact of using different hardware, for example, a change of processors or hard disks. This hardware simulation is based on standard benchmark results available for the hardware. Updated benchmark results for new hardware can be downloaded from the BGS Web page (www.bgs.com). In practice, when simulating different hardware you can select from a wide range of available hardware components from the most important vendors. BEST/1 has the benchmark information for these devices stored in its database.

The modelling techniques provided by the BEST/1 Predict module are one of the key strengths of the BEST/1 product as they allow you to plan future configurations *before* actually making the decision to buy new hardware, such as CPUs or hard disks.

To manage all the BEST/1 components, the Manager is used to collect and transfer monitoring data. The daily operation tasks can be automated using the Manager.

The BEST/1-Visualizer shows analyzed data on Windows NT-based PCs. Visualizer generates a variety of performance graphs and can also prepare graphs for publishing on the Web.

2.4.2 Hardware and Software Prerequisites

In the following we give a brief overview of the hardware and software prerequisites required for BEST/1. For a more detailed description refer to the appropriate BEST/1 documentation.

2.4.2.1 Hardware

BEST/1 Performance Assurance for UNIX is running on the most popular UNIX single processor and multiprocessor machines, including:

- IBM RS/6000 running AIX
- Digital Alpha running Digital UNIX
- HP/9000 Series 700 and 800 running HP-UX
- NCR running MP RAS
- Sun Sparc Series running Solaris

BEST/1 requires at least 32 MB of RAM and 64 MB of paging space.

A desktop display is required with a resolution of at least 1024 * 768 pixels, color or grey scale.

The *BEST/1 Installation Guide* recommends 86 MB of disk space for the installation, so that only one tape read is required. Since we installed directly from the CD we actually required only 55 MB, measured on AIX Version 4.2. Additional disk space is required for collection of performance data. Depending on the number of clients and the length of measurement periods, an additional 25 to 100 MB are needed. The BEST/1+ for TME 10 module needs an additional 0.2 MB disk space.

2.4.2.2 Software

BEST/1 Performance Assurance for UNIX Version 5.01 supports the following operating system versions:

- AIX 3.2.5, 4.1.1, 4.1.3, 4.1.5, and 4.2
- HP-UX 9.04, 9.05, 9.07, 10.01, 10.10, and 10.20
- Solaris 2.4, 2.5, and 2.5.1
- SunOS 4.1.3_U1, and 4.1.4
- Digital UNIX 3.2C, 3.2D, and 4.0
- NCR MP-RAS 2.0, and 3.0 on Intel platforms

BEST/1 Performance Assurance Version 6.0, released in October 1997, supports Windows NT in addition to the above.

Motif and X11 Release 5.01 run-time libraries are required for the BEST/1 GUI.

Some shared libraries and font libraries are required. For details see *BEST/1 Performance Assurance for UNIX - Installation Guide Release 5.01*, 9708100.

BEST/1 Agent and BEST/1-Collector use some semaphores and shared memory segments. For Solaris, Digital UNIX, HP-UX, and NCR some configuration settings must be made. Since AIX dynamically manages these resources, there is nothing to do on AIX.

To make use of all BEST/1+ for TME 10 features, the following Tivoli TME 10 products must be installed:

- TME 10 Software Distribution (on BEST/1 Server and all BEST/1 Managed Nodes)
- TME 10 Distributed Monitoring (on BEST/1 Server and all BEST/1 Managed Nodes)
- TME 10 Enterprise Console (on central TEC server)

For information about the TME 10 product versions used in our environment see 2.7.1, "Overview of the TME 10 Environment" on page 81.

2.5 Installing BEST/1 in Our Environment

This section deals with the installation of the BGS Systems BEST/1 product in the environment used during the writing of this redbook. Where possible, actual screen captures have been used to aid in understanding the installation procedure.

Each screen indicates what selection to make from the menu, or where multiple selections are necessary, and it shows the correct status of the selections before you accept them and proceed.

2.5.1 Installation Process

Follow the installation process below to install BEST/1 for UNIX on your UNIX server:

1. Ensure that the BEST/1 installation media (tape or CD) is accessible from an X terminal window on the target machine for the installation. In our case the installation medium is a CD-ROM.
2. Create the installation directory on the target UNIX server (for example, /best1).
3. Change to the installation directory as the current directory by typing `cd /best1` and then press Enter.

Installation Note

Manually copy the OP_SYS file from the installation media to the installation directory before running the installation script (see next step) to reduce the installation time.

4. Run the `/<CD_Mount_point>/<UNIX_Version>/install.best1` script.

The messages shown in Figure 8 will appear requesting the location of the input device and file for the installation. Type in the correct location of the `best1.tar` file on the installation medium (include the file name at the end of the path as indicated in the example).

Figure 8 indicates that the AIX version of BEST/1 is being installed (indicated by the install source directory of `/cdrom/AIX/`). There are separate installation subdirectories on the BEST/1 tape/CD for each operating system supported (for example, Solaris, HP-UX, AIX).

The installation script extracts the OP_SYS file from the `best1.tar` file. If an incorrect path is specified or the OP_SYS file cannot be extracted, then the installation will fail at this point.

```
Please specify the input device/file from which to install BEST/1 for UNIX
or press "return" to accept the following default value:
/dev/rmt0
/cdrom/AIX/best1.tar

Confirm the input device/file from which to install BEST/1 for UNIX:
/cdrom/AIX/best1.tar ? (y or n)
y
```

Figure 8. Installation Medium Selection

Confirm the new path by typing `y` and pressing Enter.

The BGS Systems Copyright notice appears, and then after a short delay the menu shown in Figure 9 on page 23 appears.

2.5.1.1 Installation Menu

```
BEST/1 for UNIX Installation Menu

Please select one of the following choices:

1) Proceed with BEST/1 product family installation.
2) Create remote installation tar file from distribution tape.
3) Create remote collect tar file from existing installation.
4) Create remote collect tar file from distribution tape.
5) De-install existing BEST/1 installation.

x) Exit.

Selection: 1
```

Figure 9. BEST/1 for UNIX Installation Menu

Select option 1 and press Enter to continue the installation. (The other menu options are not discussed in the installation section of this document.) The BEST/1 for UNIX Installation Options dialog will appear (see Figure 10).

2.5.1.2 Installation Options

```
BEST/1 for UNIX Installation Options

To modify the selection of an installation option, type the option number
To proceed, enter the action letter.

Installation Options:                                Selection:      Size:
1) Install BEST/1-Predict/Visualizer for UNIX ?    yes            24 Meg
2) Install BEST/1-Monitor for UNIX ?                yes            25 Meg
3) Install BEST/1-Collect for Oracle ?              no             0 Meg
4) Install BEST/1-Collect for Sybase ?              no             0 Meg
5) Install Collect Examples ?                       no             0 Meg
6) Install calling scripts to a local directory ?   yes
7) Install man pages ?                              yes

                                         Size of install.collect file: 10 Meg
                                         Additional files:              19 Meg
                                         Total disk space required:     78 Meg
                                         Total disk space available in default filesystem: 96.4 Meg

a) Accept selections and proceed.
x) Save selections and exit.

Action:
```

Figure 10. BEST/1 for UNIX Installation Options Menu

Type the appropriate menu item number to toggle the selections between yes and no so that they are set as shown above in the Selection column (for basic installation). Then type a and press Enter to accept the selections you have made and proceed to the next menu.

2.5.1.3 Predict Customization

```
BEST/1-Predict for UNIX Customization Menu

1) Change installation directory from:
   /best1
2) Change installation owner from:
   adm
3) Change installation owner's group from:
   adm
4) Change input device/file from:
   /cdrom/AIX/best1.tar
5) Change script directory from:
   /usr/local/bin
6) Change man page directory from:
   /usr/man/man8
7) Change data repository from:
   /best1/local/collect/data

a) Accept selections and proceed.
x) Save selections and exit.

Action: 1
```

Figure 11. BEST/1-Predict for UNIX Customization Menu

Select the installation and source file paths using the menu shown in Figure 11. The screen shows the paths used for installation in the test environment. The source path indicated (in option 4) is from a locally mounted CD-ROM drive.

To change the details under each menu item select the corresponding menu option and press Enter. A message similar to the one in Figure 12 will appear requesting you to enter a new value or accept the default. Once you have entered the new value, press Enter.

```
Please specify where to install the BEST/1 directory
or press "return" to accept the following default value:
 /best1/install
 /best1

Confirm where to install the BEST/1 directory:
 /best1 ? (y or n)
 y
```

Figure 12. Changing Customization Options

Confirm your selection by typing `y` and then pressing Enter.

Note

It is important to include the correct *tar* file in the value in option 4 of the menu in Figure 11 and not just the path. Otherwise the installation will fail after the last menu has been accepted.

Once all values are correct, type `a` and press Enter.

2.5.1.4 Multiple Version Configuration

```
BEST/1 Multiple Version Setup Menu

Select the override options below if another version of BEST/1 is currently
running on this node, and you wish to install and operate this new version in
parallel with the old.

1) Change best1server installation directory from:
   /etc
2) Alter BEST/1-Monitor port configuration.
   Service:                Default Port:
   -----                -
   Agent Port              = 6767
   Manager Daemon Port    = 6768

a) Accept selections and proceed.
x) Save selections and exit.

Action: a
```

Figure 13. BEST/1 Multiple Version Setup Menu

The menu in the figure above is used to change the parameters that allow more than one version of BEST/1 to run on the same machine. To run multiple versions, change both menu options to specify a new installation directory and TCP/IP port numbers. Do not change these options if you are only running a single version of BEST/1 on the target machine.

Type a and then Enter to proceed once the desired parameters have been set.

Important Note - TCP/IP Ports

When changing TCP/IP port addresses take care not to use a port that is required by other applications. Otherwise, those applications (as well as BEST/1) will not function correctly.

2.5.1.5 BEST/1-Monitor for UNIX Customization Menu

```
BEST/1-Monitor for UNIX Customization Menu

1) Change log directory from:
   /best1/bgs/monitor/log
2) Change operational directory from:
   /best1/bgs/monitor/tmp

a) Accept selections and proceed.
x) Save selections and exit.

Action: a
```

Figure 14. BEST/1-Monitor for UNIX Customization Menu

This menu sets the location of the BEST/1-Monitor log file and operational directory. Change these paths to subdirectories of the BEST/1 installation directory if they are not already correct.

Press a and then Enter to proceed once the desired paths have been set.

2.5.1.6 Collect Customization

```
BEST/1 for UNIX Collect Customization Menu

1) Change automatic e-mail notification to BGS and adm
   upon first time data is collected on this node:
   yes
2) Change BGS e-mail address:
   custdata@bgs.com

a) Accept selections and proceed.
x) Save selections and exit.

Action: a
```

Figure 15. BEST/1 for UNIX Collect Customization Menu

This menu is not critical to the installation, but is used by BGS Systems for enhanced problem resolution via electronic mail. Change menu item 1 to no if you do not want to use this feature. An external e-mail gateway is required for this option to function.

The installation script now installs the application onto the target machine. A message appears when the installation is complete.

2.5.1.7 BEST/1 Control File

The licensing of BEST/1 is handled by a control file that is supplied by BGS Systems. Copy or create the control file `bgsproducts.cntl` in the `/$BEST1_INSTALL` directory before attempting to run BEST/1.

This needs to contain product license information that is sent to you by BGS Systems Inc.

2.5.2 Installation Errors

Check the following files after installation for any errors that occurred. These files are text files:

- `/$BEST1_HOME/Best1.Installation.Log.<date>`
- `/$BEST1_HOME/Filelist.<date>`

Where `<date>` is the date of installation.

2.5.3 Starting BEST/1 after Installation

The procedure below explains how to start the BEST/1 application after installation is complete.

We created a script named `.bestrc`, which sets up the BEST/1 environment variables.

This script is shown below:

```
export BEST1_HOME=/best1
export BEST1_COLLECT_HOME=/usr/adm/best1_5.01
if [[ -z $(echo $PATH | grep $BEST1_HOME/bgs/scripts) ]]
then
  export PATH=$PATH:$BEST1_HOME/bgs/scripts
fi
```

Figure 16. *.bestrc* Script

1. Run the command `./bestrc`. This command sets up the environment for running BEST/1.
2. Run the command `best1`. This command starts the BEST/1 application.

The window in Figure 17 appears on the UNIX console when the application runs. Each button represents a module of the BEST/1 application. Click on any button to start the associated BEST/1 function.




Figure 17. *BEST/1* Main Menu

We show in 2.7.3.4, “Components for Using BEST/1” on page 97 how to launch the BEST/1 application from the TME 10 desktop.

In the following sections we show how to use all of the modules contained in the BEST/1 application.

2.6 Using the Functions of BEST/1

Now that we have successfully installed BEST/1 in our environment we explain the functions of the various BEST/1 modules, which are:

- Manager
- Collect
- Monitor
- Analyze
- Predict
- Hardware

Each of the following sections describes the function of the corresponding module and the fundamentals behind using and configuring the module. Context-sensitive online help is also available in each module to assist in the configuration.

Note

The Manager module is the last module covered as knowledge of all the other BEST/1 modules is required to use the Manager functions.

2.6.1 BEST/1 Collect

The BEST/1 Collect module of BEST/1 is responsible for all performance data collection whether it is for use by monitor (real time) or by the analyze and predict modules. This section explains the configuration processes required to use the BEST/1 Collect module.

2.6.1.1 Collect Process

The order of tasks required to set up the BEST/1 Collect module is:

1. Create/open domain
2. Edit domain properties (for new domain)
3. Add BEST/1 managed nodes to domain (for new domain)
4. Configure data collection options by starting collection
5. Save domain

These steps are described in detail in the following sections.

2.6.1.2 Domains

After starting the BEST/1 Collect module, domains are set up to form logical groupings of nodes within the management environment. These domains do not follow other domain models (such as DNS) and are defined only for BEST/1.

Set up a domain within BEST/1 Collect by selecting **New Domain** from the **File** pull-down menu. Select the **Domain Name and Description** option from the **Edit** pull-down menu to display the window shown in Figure 18 on page 29.

Enter the domain name and description in the Collect - Edit Attributes window as shown in Figure 18 on page 29. The domain name can be any descriptive text that identifies the nodes within the domain. The name must conform to UNIX file naming conventions.

Click **OK** when complete.




Figure 18. Collect - Edit Domain Attributes Window

The example window in Figure 19 on page 30 shows two nodes, rs600024 and rs600026 in the current domain (itso.ral.ibm.com).



Figure 19. BEST/1 for UNIX - Collect Window

Note

The above figure shows two nodes, one of which had already been added as shown in the next section. If you start with a new BEST/1 installation, this window will initially only contain the BEST/1 Management Console.

2.6.1.3 Adding Nodes to the Domain

BEST/1 managed nodes are UNIX machines that have the BEST/1 agent installed. Managed nodes must be added to the collect domain before any data collection can take place.

Add managed nodes to the domain by selecting **Create** from the **Edit** pull-down menu. Enter the properties for the new node in the Collect - Edit Node window. An example of managed node properties is shown in Figure 20 on page 31.

Repeat this process until all managed nodes have been added to the domain.

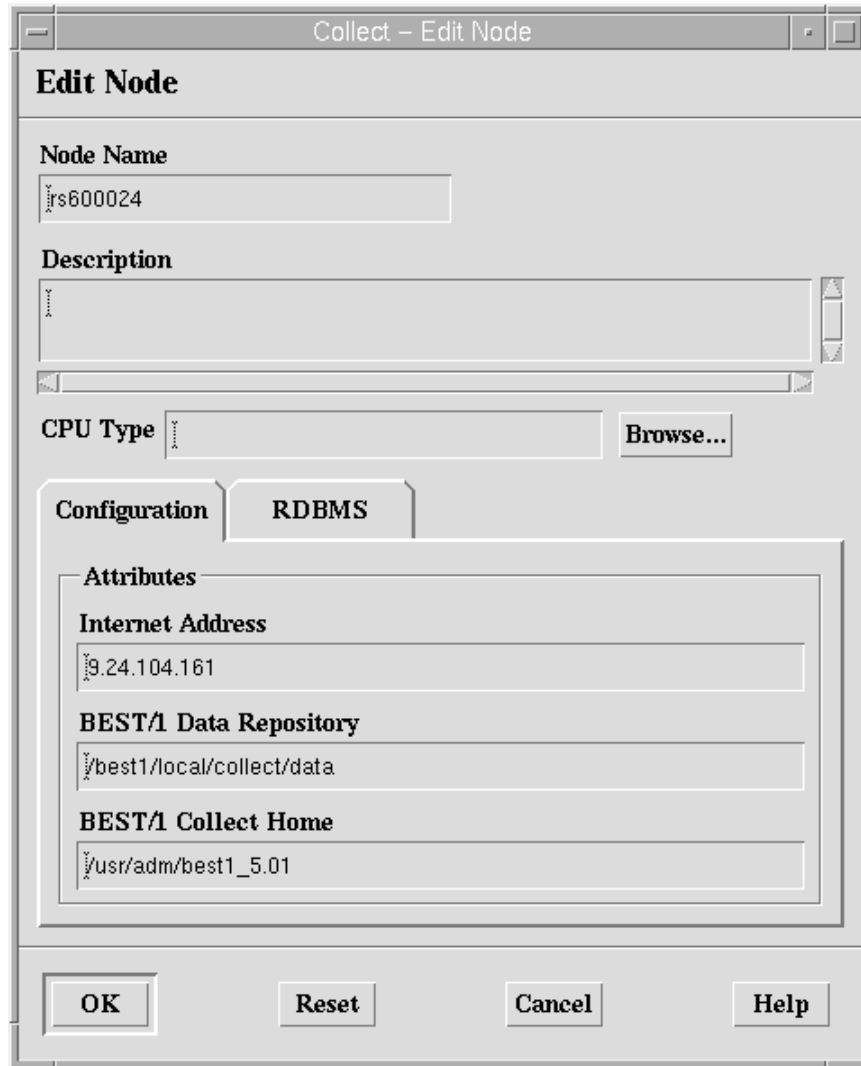


Figure 20. Managed Node Properties in the Collect - Edit Node Window

The input fields in the dialog are as follows:

- Node Name
Node name for use by BEST/1. This name does not have to be the same as the TCP/IP hostname of the device, but it is advisable to make it the same to avoid confusion when using BEST/1.
- Description
Free-form text to describe the node, purely for reference.
- CPU Type
Identifier for the type of CPU used by the UNIX server. It is not necessary to enter a value in this box as BEST/1 can automatically detect the CPU type during the running of data collection. You can use the **Browse...** button to select a CPU type.
- Internet Address
Enter the TCP/IP address of the target node here.
- BEST/1 Data Repository

This identifies the location where the performance data will be stored on the node as it is collected.

- BEST/1 Collect Home

The collect process can only run on machines that have BEST/1 installed. The value entered in this field identifies the location of BEST/1 on the node to be monitored by BEST/1 Collect. This reflects the environment variable \$BEST1_COLLECT_HOME on the target node and hence will generally not change from node to node. The value shown in Figure 20 on page 31 represents a link to /best1/admbest1_5.01, which in turn links to /etc/best1server and best1/local/collect/.

2.6.1.4 Starting Data Collection

Select the managed nodes that you want to monitor using the mouse and then select the **Start Collection** button in the middle of the BEST/1 for UNIX - Collect window. A window will appear for BEST/1 Collect configuration.

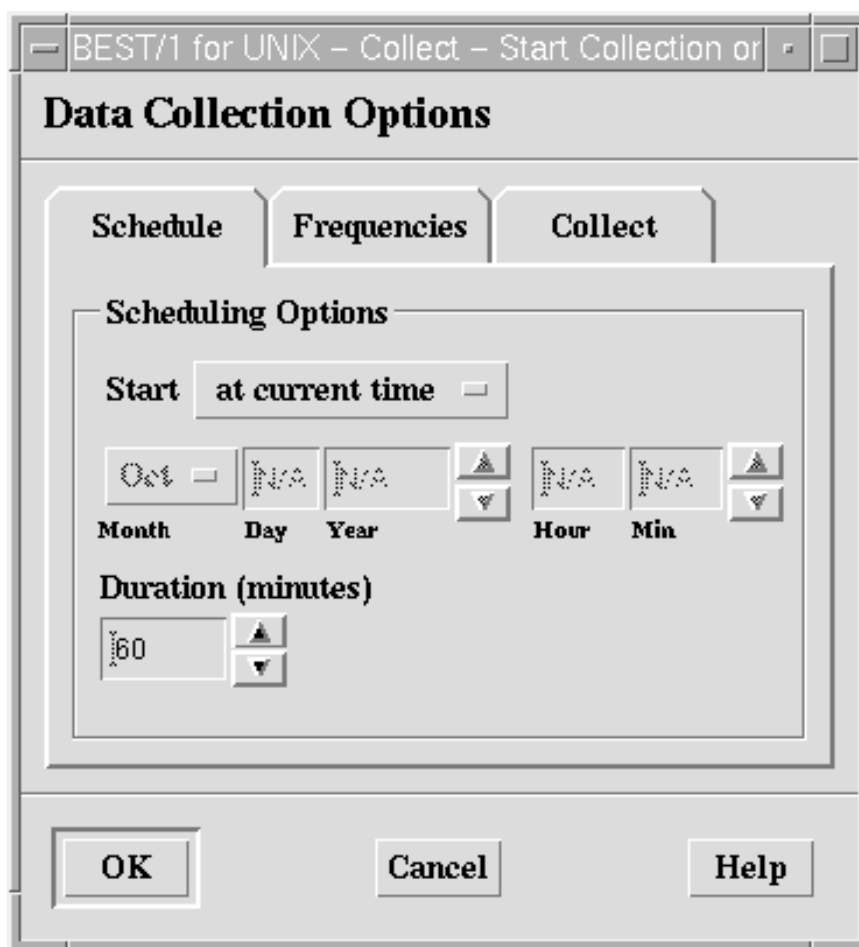


Figure 21. Data Collection Options - Schedule

Set up the schedule and duration of the collection process in the window shown in Figure 21. Click on the button beside Start to toggle between starting the collection at a defined date and time, or starting collection immediately after configuration is complete (at current time). Specify the date and time to start the collection below Start if you have chosen **Start - at specified time**.

The duration is the length of time that the collection will run for after starting, measured in minutes and can be a minimum of 1 minute. Click on the arrows beside duration to increase or decrease the duration in increments of 15 minutes.

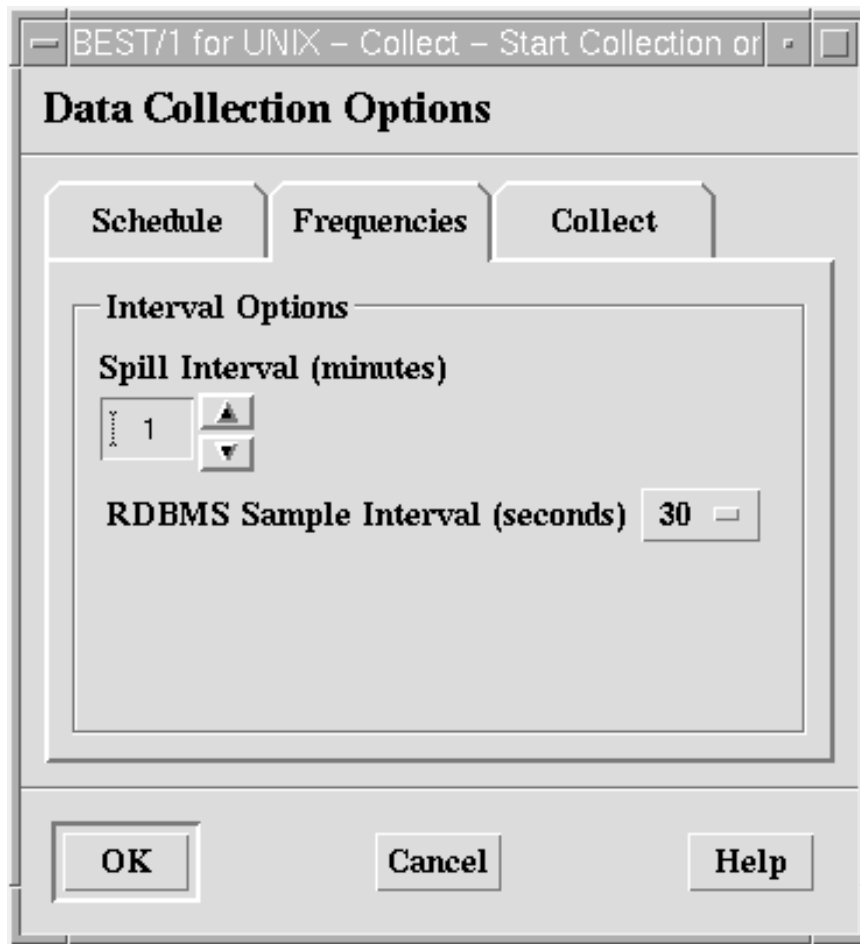


Figure 22. Data Collection Options - Frequencies

The Spill Interval (minutes) is the time in minutes between data samples on the hosts. Change the Spill interval (minutes) by clicking on the up and down arrows beside the value.

The RDBMS Sample Interval (seconds) is the interval between data samples for RDBMS data. Set the **RDBMS Sample Interval (seconds)** to the desired period by clicking on the button and choosing a value from the pull-down list.

Click on the **Collect** tab when settings are correct.



Figure 23. Data Collection Options - Collect

Select a name for the collection in the System Collect Name field, and enter the Electronic Mail Destination address of a user on the system (usually root) to notify by e-mail when collection is complete.

Click **OK** when all parameters are configured as desired. The display will return to the window shown in Figure 19 on page 30 showing that the collector is running on the selected managed nodes.

2.6.1.5 Stopping Data Collection

With reference to Figure 19 on page 30 select any node(s) from the list and click on **Stop Collection** to stop collecting data on those nodes.

2.6.1.6 Update and Status Window

Status information is displayed in the Messages window on the BEST/1 for UNIX - Collect screen in Figure 19 on page 30. Select a node from the list and then click on the **Update Status** button to update information regarding that node in the Messages window.

2.6.2 BEST/1-Monitor

The BEST/1-Monitor module of BEST/1 is a real-time performance analysis tool for UNIX. It can be used to monitor any UNIX server that is running the BEST/1 Collect processes and BEST/1 Agent.

2.6.2.1 Monitor Process

This section explains the processes involved in configuring Monitor for real-time system analysis.

1. Create a policy
2. Create domain(s)
3. Add nodes to the domain(s)
4. Create alerts
5. Create conditions
6. Activate the policy
7. View monitors

Refer to the following sections for further details on each of the above tasks.

2.6.2.2 Monitor Object Hierarchy

In the BEST/1-Monitor module, the general flow of hierarchy is:

Policy --> Domain --> Nodes --> Alerts --> Conditions and Actions

This indicates that domains contain nodes and policies contain domains. This structure is not rigid, in that domains are not necessarily required, and an alert may exist at the same level as a node in a domain or policy. These exceptions are explained in the following sections.

Reference Figure 24 on page 36 for an example of this hierarchy.

2.6.2.3 Creating Policies

The first step in using BEST/1-Monitor is to create a policy. The policy contains all information regarding the nodes being monitored and with what parameters they are monitored.

Select the **New** option from the **File** pull-down menu in BEST/1-Monitor to create a new policy. Save the policy file when all appropriate nodes, alerts, conditions, etc. have been added (see sections below) using the **Save** or **Save As** options from the **File** menu. Policy files are saved with a .plc extension.

2.6.2.4 Creating Domains

Domains are used in BEST/1 as containers to create logical groups of BEST/1 managed nodes. Multiple domains can be created in a single policy file, but domains are not compulsory; nodes may exist directly under the policy.

With reference to Figure 27 on page 39, select the **Policy** icon using the left mouse button, and then click on the **Domain** icon on the left side of the window. This action will display a dialog box requesting a name for the domain. Enter a descriptive name for the domain and then click **OK**. The domain icon appears to the right of the Policy icon in the window. Repeat this process to create additional domains.

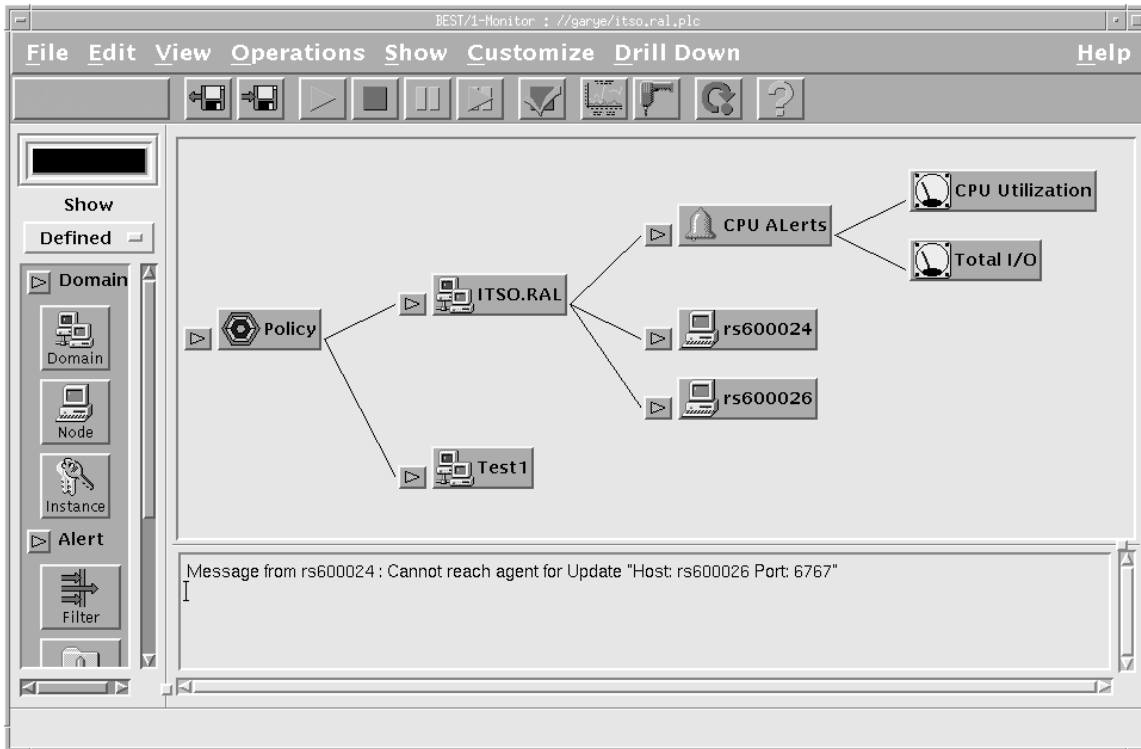


Figure 24. BEST/1-Monitor Window

2.6.2.5 Adding Nodes to a Domain

Nodes are added to the domain in the same way that a domain is added to the policy. Select the **Domain** in the window to which the node is to be added, then click on the **Node** icon on the left side of the BEST/1-Monitor window as shown in Figure 24. Select the nodes to add from the BEST/1-Monitor-Add Node window that appears and then click **OK**. The Discover option can also be used to add nodes.




Figure 25. BEST/1-Monitor-Add Node Window

2.6.2.6 Alerts and Conditions

Use alert objects to group together one or more conditions, so that when *all* of the conditions grouped under the alert object are true, an alert is generated.

Define an alert at the same level as nodes (for example, policy or domain level) for that alert to affect all of those nodes. If an alert is defined *below* a node, it will only affect that node.



Figure 26. BEST/1-Monitor-Add Alert Window

Use conditions to specify thresholds for values of a specific monitored parameter. The window in Figure 27 on page 39 shows what parameters are required to configure a condition.

Select options by clicking on the arrows beside the parameters, which will show a selection list with all available values. The only parameter that does not have a selection list is the Threshold Definition/Value parameter. Enter a value manually in this field. Click **OK** when all options are set as desired.

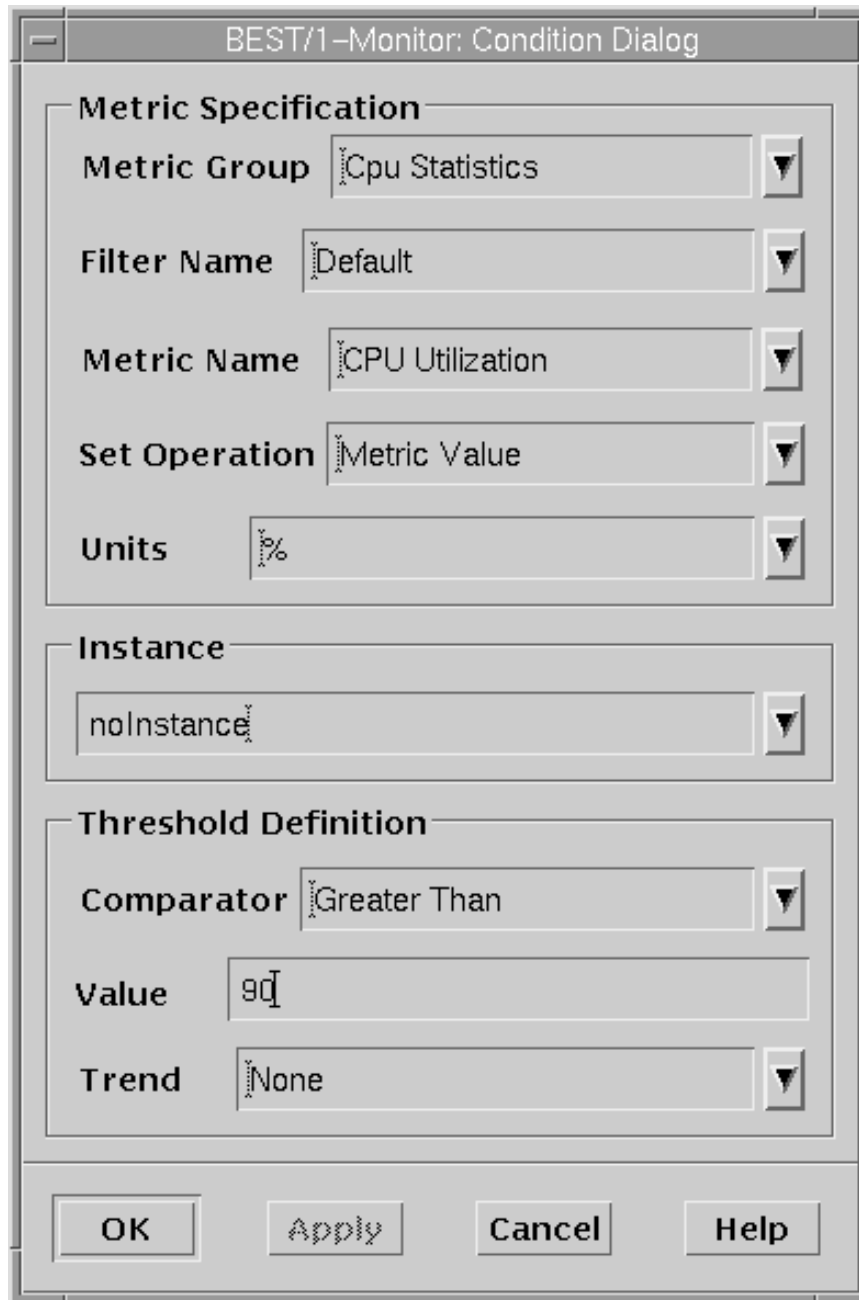


Figure 27. BEST/1-Monitor: Condition Dialog Window

Note

The Instance parameter is only necessary when monitoring a database.

2.6.2.7 Define Filters

Use filters to reduce the amount of data that is viewed either in a chart or table. This is useful, for example, where a chart is required of CPU utilization of processes running on a UNIX server. If the chart shows all processes then it is difficult to read and hence obtain any useful information. A filter can be used to limit the chart to show only the 10 processes that use the most CPU time.

Define filters at the same level as conditions and actions (that is, below alerts). With reference to Figure 24 on page 36, click on an alert icon in the policy under which the filter is to be placed, then click on the **Filters** button under the Alerts section on the left of the window to add a filter. Set the filter options in the Add Filter window. Add arguments by selecting **Metric Group** from the pull-down list and defining the filter characteristics in the lower part of the window. The items and arguments that are available in the Filters section of the window depend on what metric group has been selected. The Metric Group is the same as those used in conditions. The first part of the actual filter corresponds to the Metric Name also used in conditions.

Add multiple filters by clicking on the + symbol at the left side of the window.

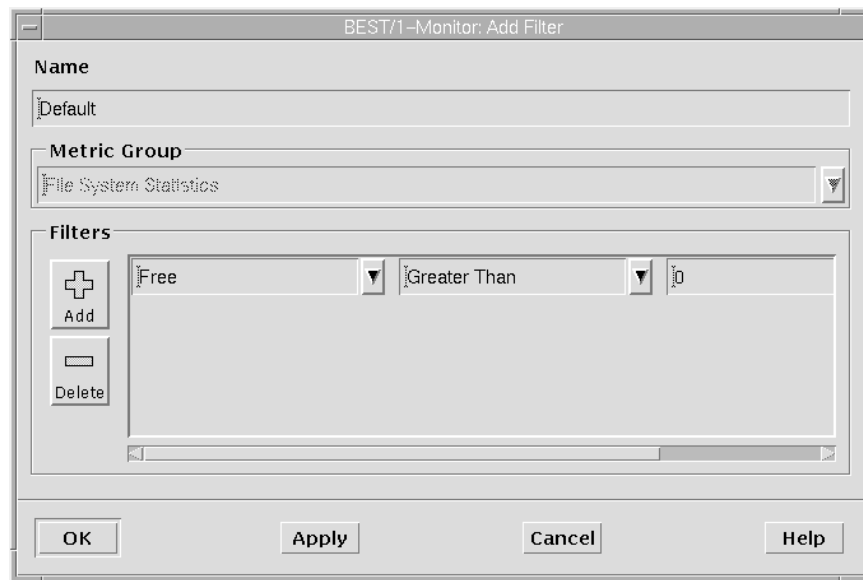


Figure 28. BEST/1-Monitor-Add Filter

2.6.2.8 Define Actions

Actions are initiated once an alert has been triggered.

An action can be:

- Execute a script (local or remote)
- Send an e-mail message
- Send an SNMP trap

Select the desired alert in the policy displayed in the BEST/1-Monitor window Figure 24 on page 36 and then click on an action (SNMP, script or e-mail) on the left-hand side of the monitor window.

Configure the action parameters for each action set for an alert.

E-Mail Configuration




Figure 29. BEST/1-Monitor: Mail Action

Enter the user name to whom you want to send the e-mail in the Send mail to field, and a subject. Monitor creates the body of the mail message automatically. An example message is shown in the following figure.

```

adm Tue Nov 4 17:25:29 1997
Date: Tue, 4 Nov 1997 17:25:28 -0500 (EST)
From: adm
To: root
Subject: !!! ALERT !!!

Tue Nov 4 17:25:27 1997

On node rs600024, the following alert was asserted

Alert Name:          CPU Alerts
Console:             9.24.104.161
Alert Type:          AUTO_CLEAR
Severity:             6
Check Interval:      10 second(s)
Number of Intervals: 1

Condition 1
Metric Group: Cpu Statistics
Filter:         Default
Instance:       noInstance
Metric:         CPU Utilization
Set Operation: Metric Value
Record IDs:     CPU ID = 0
Current Value:  22.756 %
Threshold Definition:
  Comparator:   Greater Than
  Value:        20
  Trend:        None

Condition 2
Metric Group: Disk Statistics
Filter:         Default
Instance:       noInstance
Metric:         Total I/O
Set Operation: Metric Value
Record IDs:     Disk Name = hdisk0
Current Value:  15.82 Pages/Sec
Threshold Definition:
  Comparator:   Greater Than
  Value:        10
  Trend:        None

Send Mail
User List: root
Subject:  !!! ALERT !!!
Execute Remote Command
Hostname: localhost
Username: garye
Command:  /usr/bin/a.sh
Send SNMP Trap
Trap Id:  14
Manager:  127.0.0.1

```

Figure 30. Sample Alert Mail Message

This example indicates the following condition:

- The top part of the message is the e-mail header and contains addressing information.
- Below the message header is the alert configuration information.
- Condition 1 indicates that the CPU Utilization statistic exceeded the set threshold of greater than 20.

- Condition 2 indicates that the Total I/O for hdisk0 exceeded the set threshold of greater than 10 pages/second.
- The last part of the message shows the Actions configuration, including the mail, script and SNMP configuration.

Script Execution




Figure 31. BEST/1-Monitor: Script Action Window

Enter the hostname of the device on which the script is located, the user name that is to be used to execute the script and the location of the script file. Click **OK** when finished.

SNMP Traps: Shown below is the configuration window for SNMP Actions with respect to alerts.



Figure 32. Action - SNMP Trap Configuration

Enter the ID of the SNMP trap to send in the Trap ID field, and the SNMP Manager that the trap will be sent to. Click **OK** when complete.

Note

SNMP traps can be used to send status and alert information from BEST/1 to an SNMP manager, such as TME 10 NetView.

2.6.2.9 Policy Activation

The alerts, conditions and actions only work once the policy is activated. To activate the selected policy, either press Ctrl+A or select **Activate** from the Operations pull-down menu in the BEST/1-Monitor window shown in Figure 24 on page 36.

2.6.2.10 Graphs and Charts

The output statistics of BEST/1-Monitor can be in the form of graphs and charts. This section describes how to produce graphs of Monitor statistics. Follow the steps below to produce an output graph:

1. In the main Monitor window click on the **Graph** icon.
2. Select the parameter(s) and instance to graph.
3. Click on the **Graph** icon.

These steps will vary slightly depending on the type of information and number of parameters that you want to display on the graph.

Select any icon (Domain, Node, Alert, Condition, etc.) in the monitor main window and using the mouse, select the icon shown below:




Figure 33. Monitor Graph Icon

Depending on the type of icon selected prior to clicking on the graph icon, a table is displayed showing parameters that can be displayed on the graph. If a condition icon was selected then only that parameter will appear in the table, but if a node, alert, domain, or other icon was selected prior to clicking on the graph icon, then parameters for all conditions that apply to that object and all objects below it will appear. The example below shows disk and CPU parameters as these are the only conditions set up in BEST/1-Monitor for our example.

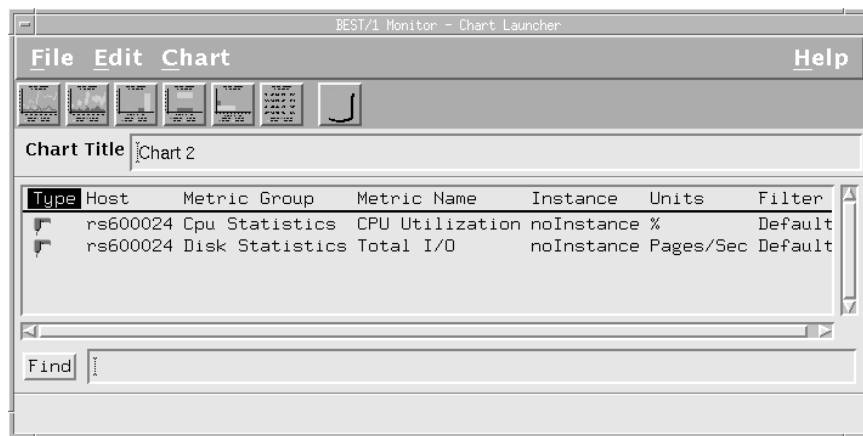


Figure 34. Condition Parameters for Graphs

Select one of the parameters (for example, either CPU or Disk) by clicking on the appropriate row of the table, and then click on the **Graph** icon below the File menu. This action will display another table with an additional breakdown of parameters. The example shown in Figure 35 below is produced by selecting the **Disk Statistics** row and clicking on the graph icon.

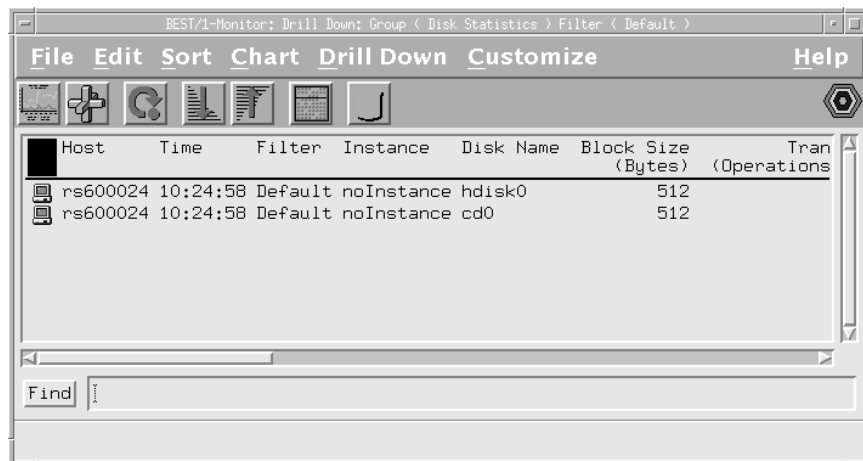


Figure 35. Disk Parameter Breakdown

The screen shows two *instances* of disks, a hard disk (hdisk0) and a CD-ROM (cd0). Select the row of the device that you want to graph, then do the following.

To Produce a Single Parameter Graph: Using the horizontal scroll bar, scroll the screen to find the single parameter to graph (in the top row of the table) and then click on that parameter. Click on the graph icon (Figure 33 on page 45) in the icon bar to display the graph.

To Produce a Multiple Metric Graph: Click on the icon in Figure 36 to display the window in Figure 37. Then select the parameters to graph by clicking the mouse on the box to the left of each parameter. Click on the graph icon of that window when all desired parameters are selected.




Figure 36. Multiple Metric Graph Button

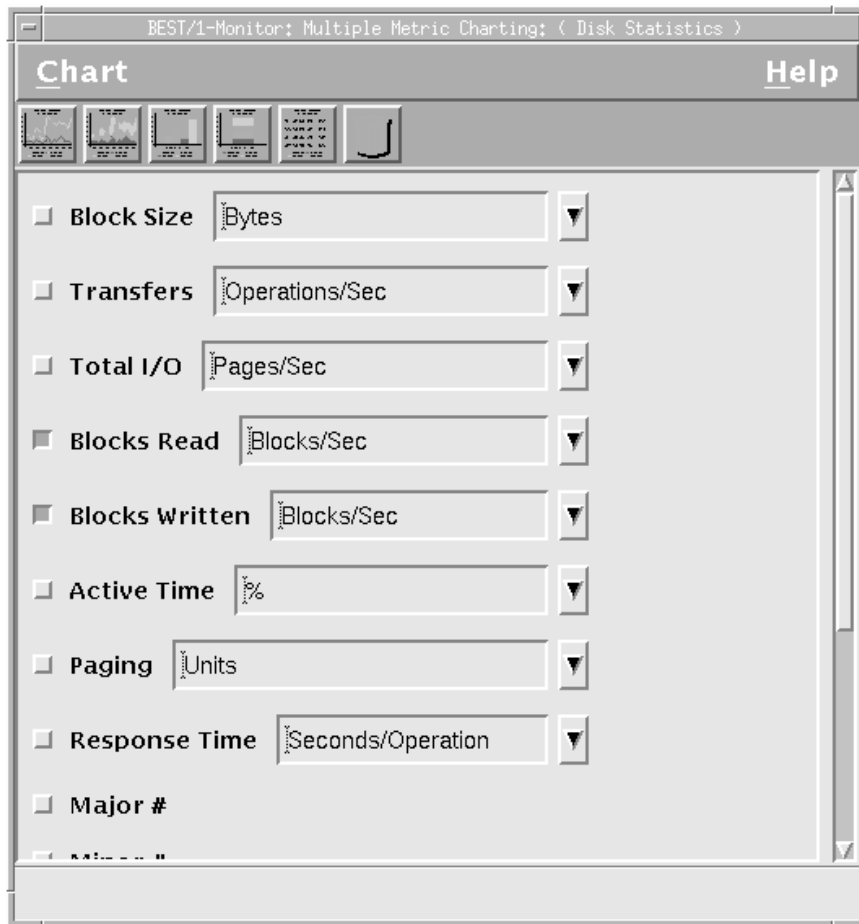


Figure 37. BEST/1-Monitor: Multiple Metric Charting Window

The selections shown in Figure 37 produce the graph shown in Figure 38 on page 47.

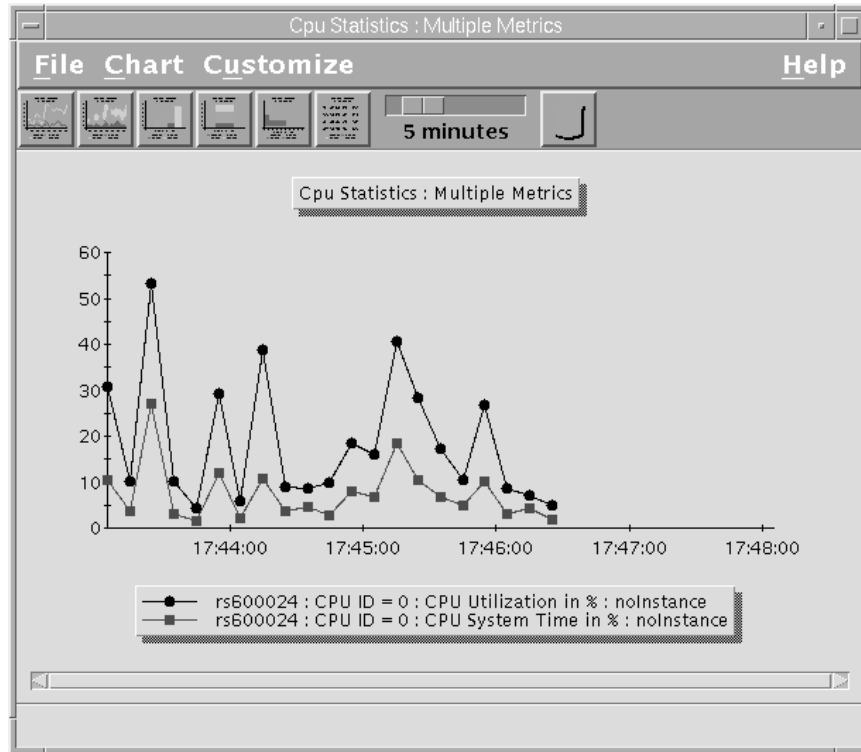


Figure 38. Example Disk I/O Graph

2.6.3 Analyze

The BEST/1 Analyze module is used to analyze and create reports from BEST/1 Collect data. This section describes the configuration required to use the BEST/1 Analyze module.

2.6.3.1 Analyze Process

The order of tasks required to set up the Analyze module is:

1. Create a new or open an existing command file.
2. Add nodes to analyze.
3. Set analysis interval.
4. Set user classes, transaction classes and workloads.
5. Run analysis on data.
6. Print reports.

The following sections describe each of these steps in detail.

2.6.3.2 Creating/Opening a Command File

1. Select **File** from the menu bar of the main Analyze window. Then select either the **Open Commands** or **New Commands** option from the pull-down menu.
 2. If you selected **Open Commands**, a new dialog box appears prompting for the location of *.an files. Select the location of the command file, and then the appropriate command file.
- Or,
3. When you select the **New Commands** option, the Analyze window will appear as shown below.

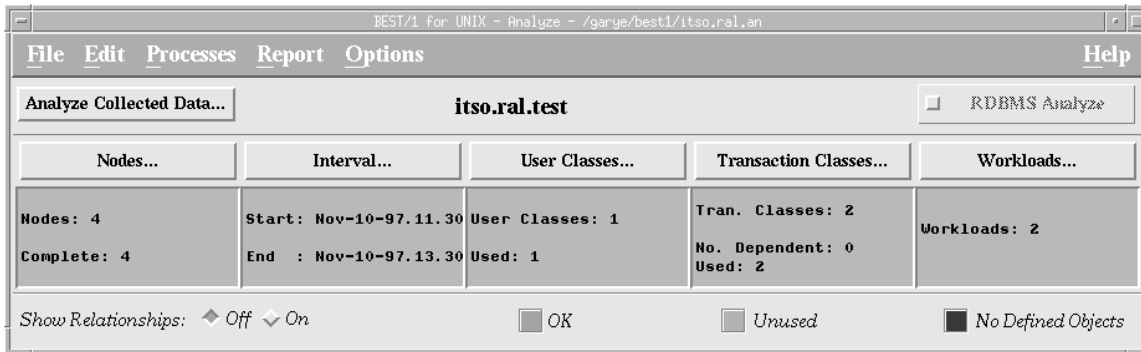


Figure 39. BEST/1-Analyze - Main, after Adding All Nodes and Classes

2.6.3.3 Adding Nodes for Analysis

Add nodes to the analysis model either by adding them individually in the same way as they are added to the collect domain, or select a domain file with the appropriate nodes already added. Click on the **Nodes...** button (refer to Figure 39) to add new nodes.

A new window is displayed as shown in Figure 40 on page 49, showing what nodes are included in the analysis. To add new nodes, from the File menu select **Load from Domain File** or select **Create** from the Edit menu to add nodes individually. Click on **Close** when complete.

If loading nodes from a domain file, when you select this option from the File menu, select either **Merge New Nodes and Update** or **Update Only**. The **Update Only** option will only update nodes currently in the active window but will not add any additional nodes that it finds in the domain file. The **Merge New Nodes and Update** option will update any existing nodes and add any new nodes that it finds in the domain file.

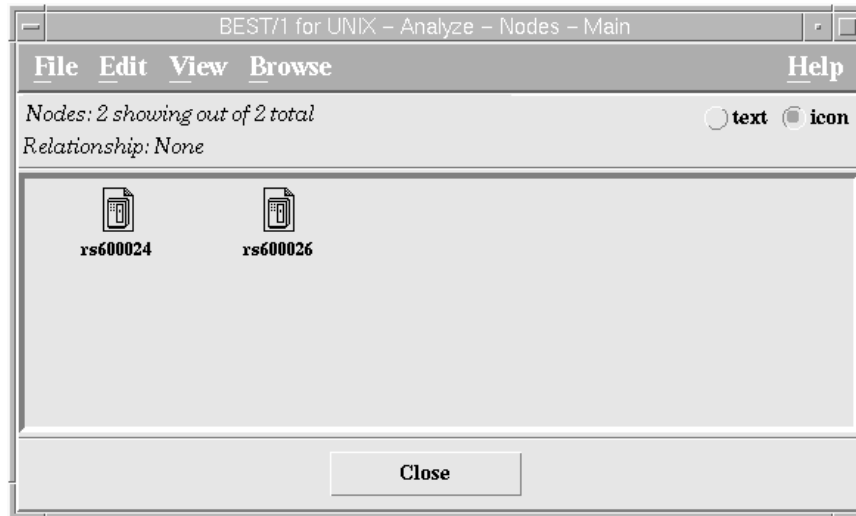


Figure 40. BEST/1 Analyze with Two Managed Nodes

2.6.3.4 Setting Intervals

Click on the **Interval...** button in the BEST/1-Analyze main window to access the Browse Collected Data window.

This configuration window defines the scope of collected data that you can browse and the scope of data to be included in any analysis. An example of this window is shown in Figure 41.

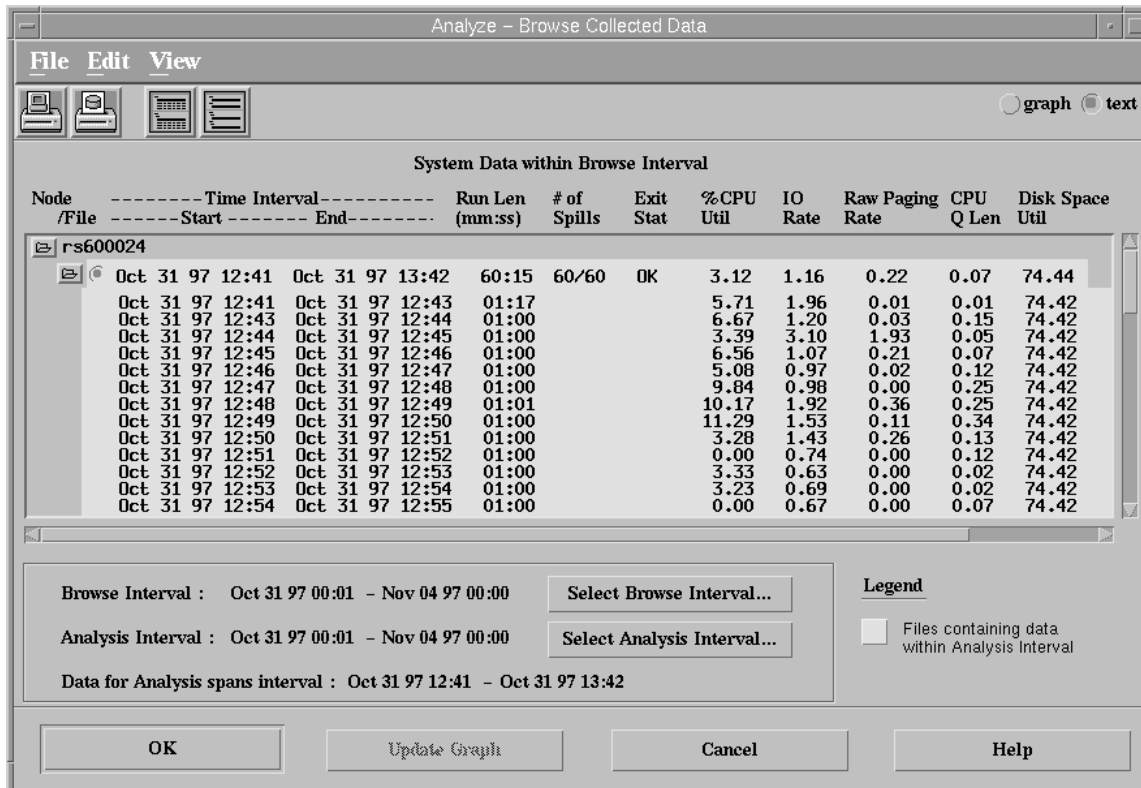


Figure 41. Analyze - Browse Collected Data Window

To set the scope of data to browse, click on the **Select Browse Interval...** button at the bottom of the window. Set the desired browse interval (similar to the window shown in Figure 42 on page 50) and then click **OK** to return to the Analyze - Browse Collected Data window.

Click on the **Select Analysis Interval...** button to define the scope of the data to be used for analysis. This action will display the window shown in Figure 42.

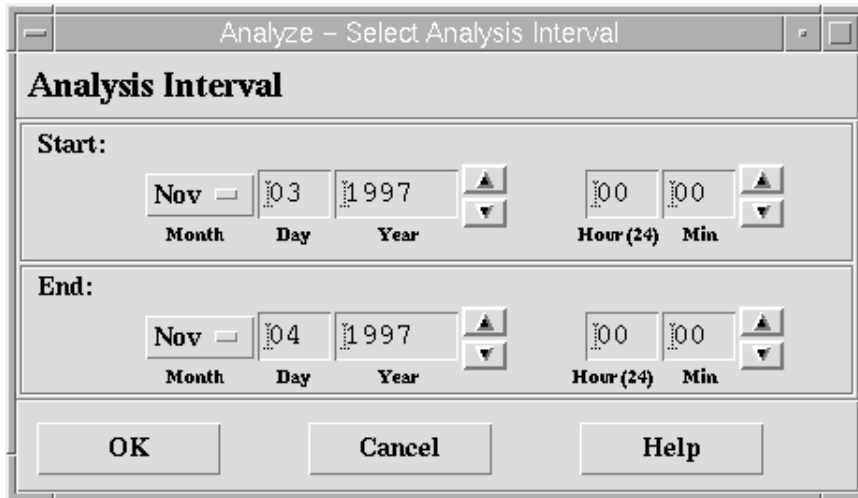


Figure 42. Analysis Interval Configuration Window

Set the start and end date/time for the data to be included in the analysis. Click **OK** when the desired interval is set.

To set the scope of data to browse, click on the **Select Browse Interval...** button from the Analyze - Browse Collected Data window (see Figure 41 on page 49). Click **OK** when you have set the desired interval.

2.6.3.5 User Classes

We used the default user classes. This section is a basic guide to setting up new user classes.

Create user classes to narrow the scope of the analysis carried out by the BEST/1 Analyze module. When creating a user class, define which patterns to include and/or omit by entering appropriate patterns in the fields provided (refer to Figure 44 on page 52). The *BEST/1 User's Guide* provides detailed information regarding acceptable patterns. The include and omit information can be set for one or more of the following areas:

- User name
- UID
- Group name
- GID

Select the **User Classes...** button from the BEST/1 Analyze main window. This displays the window shown in Figure 43 on page 51 where you select either an existing user class or create a new class by selecting **Create** from the Edit pull-down menu. We used the default user classes.

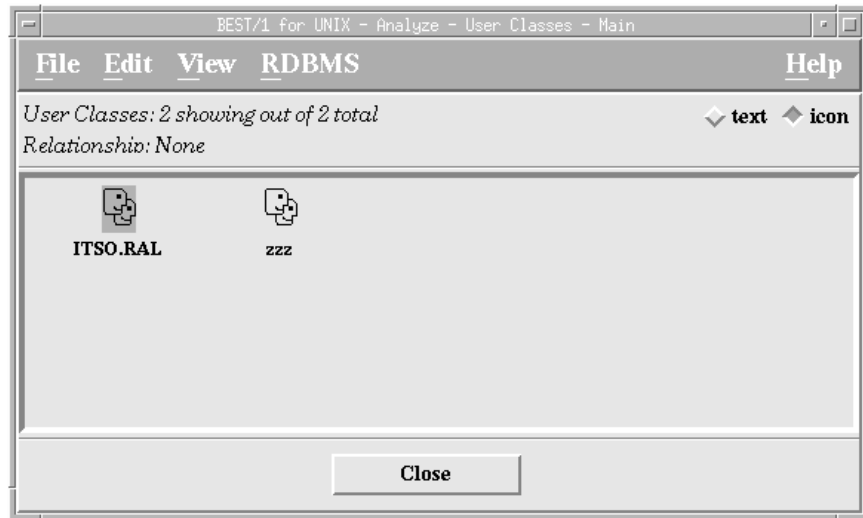


Figure 43. BEST/1 for UNIX - Analyze - User Classes - Main

Creating a New User Class: If you choose to create a new user class, the window in Figure 44 on page 52 will appear for you to fill in the relevant fields.

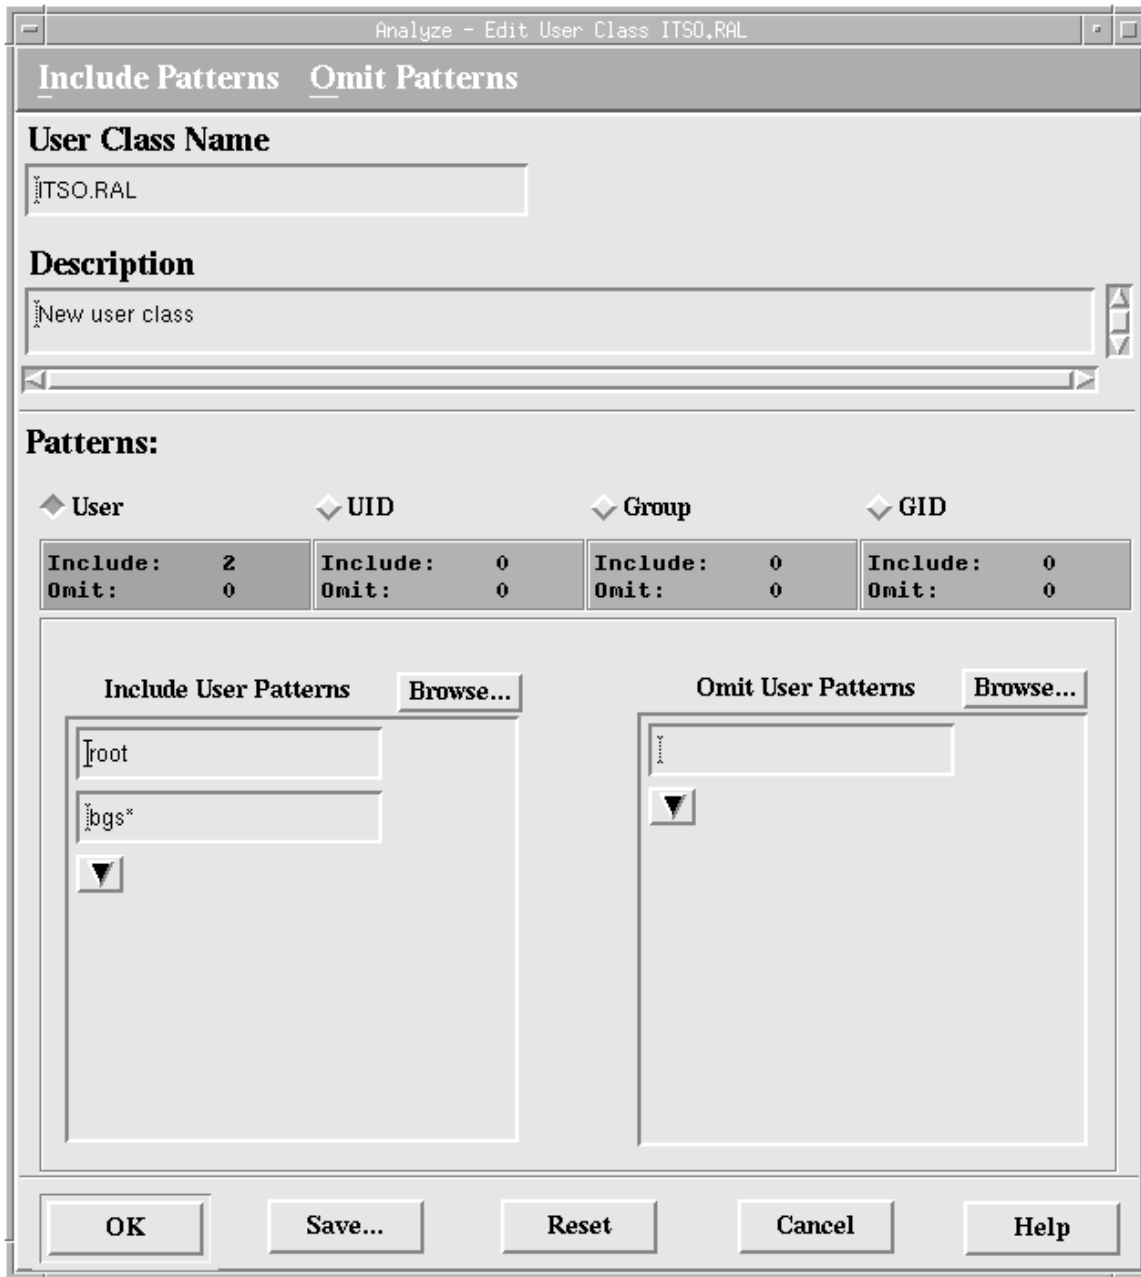


Figure 44. Analyze - Edit User Class ITSO,RAL

You can create *include* and/or *omit* ranges for each of the four patterns (User, UID, etc). Select the required pattern and then enter values in the Include User Patterns and/or Omit User Patterns boxes. Create multiple patterns by clicking on the arrow below the pattern value. Click on the **Save** button when finished.

The example patterns shown in Figure 44 indicates that the users listed below are included in any analysis that uses this user class:

- *root* user
- Any user name that begins with *bgs*, such as *bgsadmin*

Note

The bgs* pattern shown will actually include all users that begin with bg. BEST/1-Analyze will ignore the character immediately prior to the * character.

In the example, no users are excluded from the user class.

2.6.3.6 Transaction Classes

In the same way as user classes, transaction classes are used to reduce the scope of the analysis so that the analysis data is more specific to certain areas of interest. Transaction classes are used to group UNIX processes together to form a logical group that is treated as a single entity when monitoring performance.

Transaction classes are either dependent or independent. Independent transactions are initiated directly by users in workloads, and dependent transactions are initiated by independent transactions.

Select the **Transaction Classes...** button from the BEST/1-Analyze main window. This displays the window shown in Figure 45 where you select either an existing transaction class or create a new class by selecting **Create** from the Edit pull-down menu, at which point you can select either to create an **Independent...** transaction or a **Dependent...** transaction.

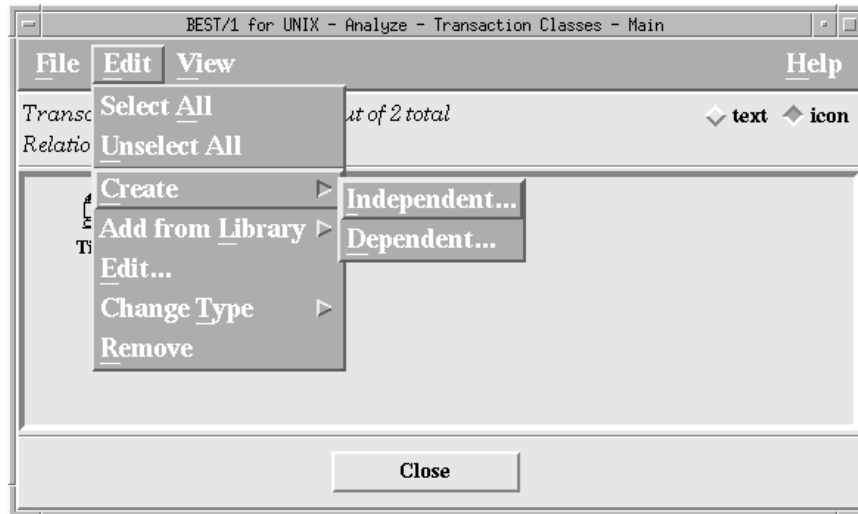


Figure 45. BEST/1 for UNIX - Analyze - Transaction Classes - Main Window

If you choose to create a new transaction class, the window shown below will appear for you to fill in the relevant fields.

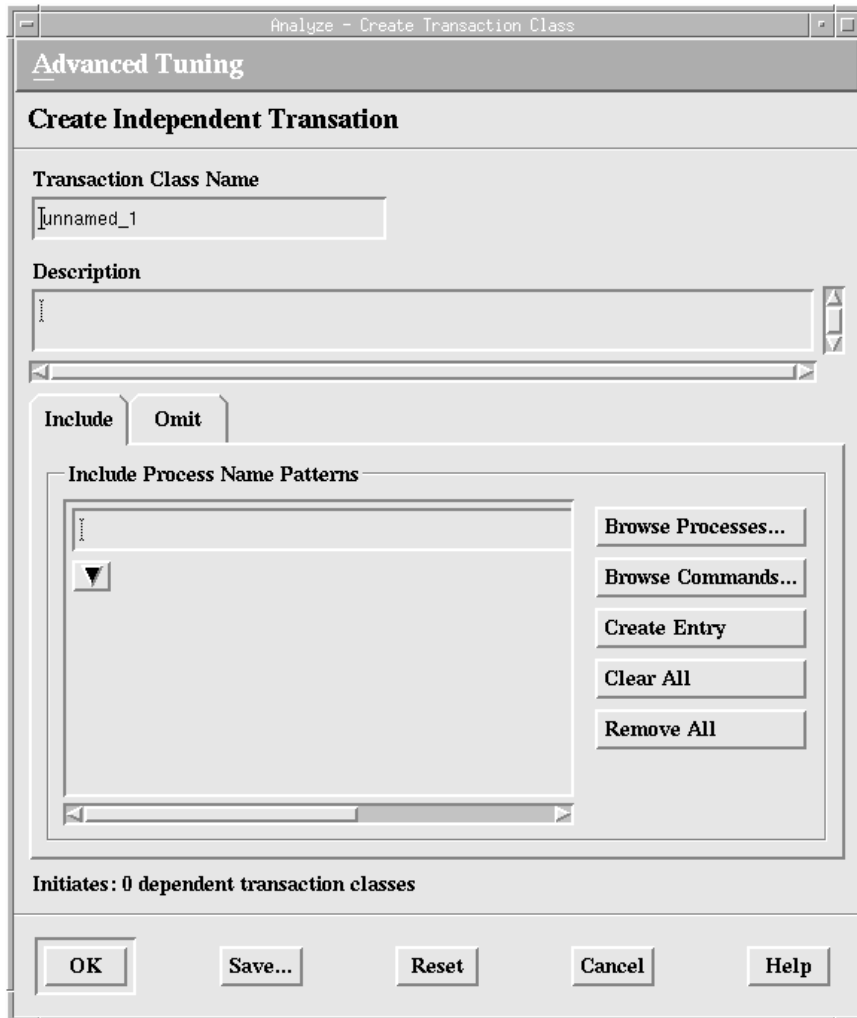


Figure 46. Analyze - Create Transaction Class Window

Enter the process name patterns in the Include and/or Omit areas of this configuration window. Click on the arrow in the pattern areas to specify multiple patterns for the class. Click on the **Browse Processes...** or **Browse Commands...** buttons to view processes and commands that are available in the current model. Use wildcard characters such as * in name patterns to specify multiple processes.

Normally, processes are specified in a transaction class using only the process name, but in some circumstances it may be necessary to specify the complete command, such as when two processes have the same name but you want them analyzed in separate transaction classes.

Advanced options are available for more detailed configuration of the transaction classes via the Advanced Tuning menu, in which you can configure options such as child processes, etc. Click on the **Save** button when you have finished editing the transaction class.

2.6.3.7 Workloads

Workloads are used by BEST/1 Analyze to group users, nodes and transactions together to form a logical unit of work that can be analyzed. An example of the practical use of a workload is where you want to analyze a group of users (for example, the sales department) using a specific application (for example, a sales database) on a server (node). In this example you would create a user class containing all the sales department users, a transaction class containing the appropriate database processes running on the server, and add the appropriate node to the model. Once these are defined, create a workload containing these classes and the node.

Creating a New Workload: In the BEST/1 for UNIX - Analyze - Main window click on the **Workloads...** button to open the Available Workloads window. The window contains a default workload name zzz. Any workloads that you create will appear in this window. Select either the desired workload and click on the **Close** button to return to the BEST/1 for UNIX - Analyze - Main window or select **Create** from the **Edit** pull-down menu to create a new workload.

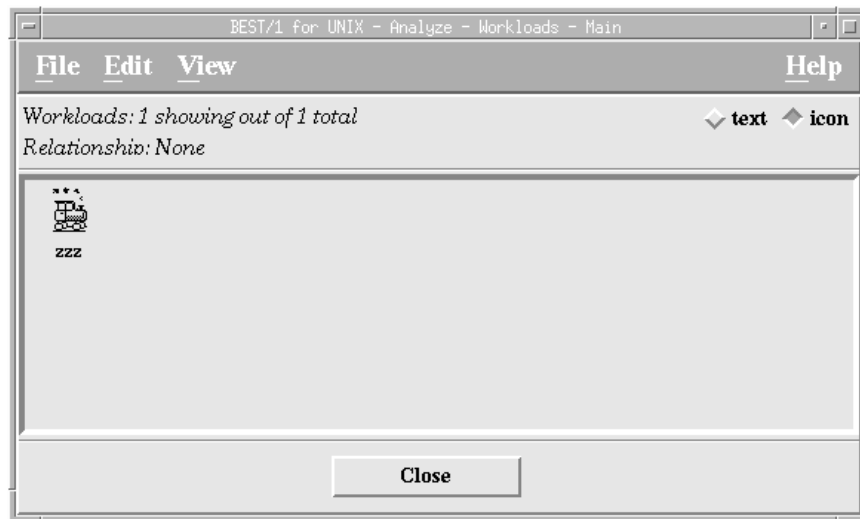


Figure 47. BEST/1 for UNIX - Analyze - Workloads - Main Window

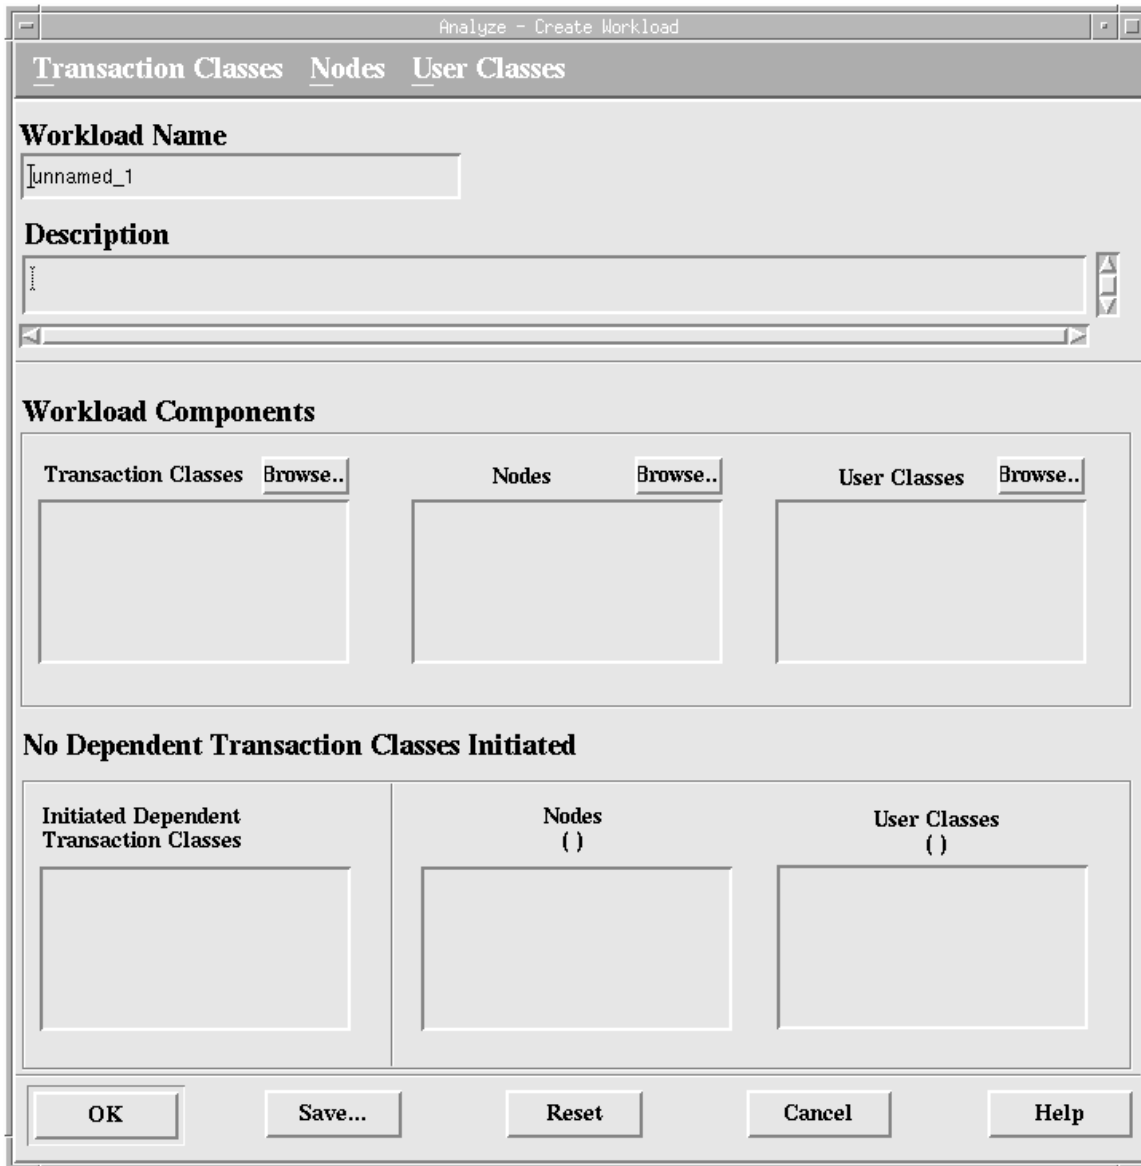


Figure 48. New Workload Properties

After selecting **Create...** from the Edit pull-down menu the window shown in Figure 48 appears. Enter the Workload Name, Description and the individual workload components (transaction classes, nodes and user classes) in the appropriate boxes. Click on the **Browse..** buttons for each workload component to select from a list of available objects in the current model.

Select one of the menu options; **Transaction Classes**, **Nodes** or **User Classes** to remove items from the lists.

2.6.3.8 Performing Analysis

Select the **Analyze Collected Data...** button from the BEST/1 for UNIX - Analyze - Main window to start the data analysis process (refer to Figure 39 on page 48).

Enter the Output Filename in the indicated box of the Analyze Collected Data - Options window, as shown in Figure 49 on page 57. Click **OK** when finished.

The analyze process automatically creates a report file with the same name and location as the model file being evaluated (*.md), but with a .anb extension. Change the output file name in the Analyze - Analyze Collected Data - Options window shown in Figure 49 on page 57. Do not add the extension to the file name as indicated in the Output Filename field.

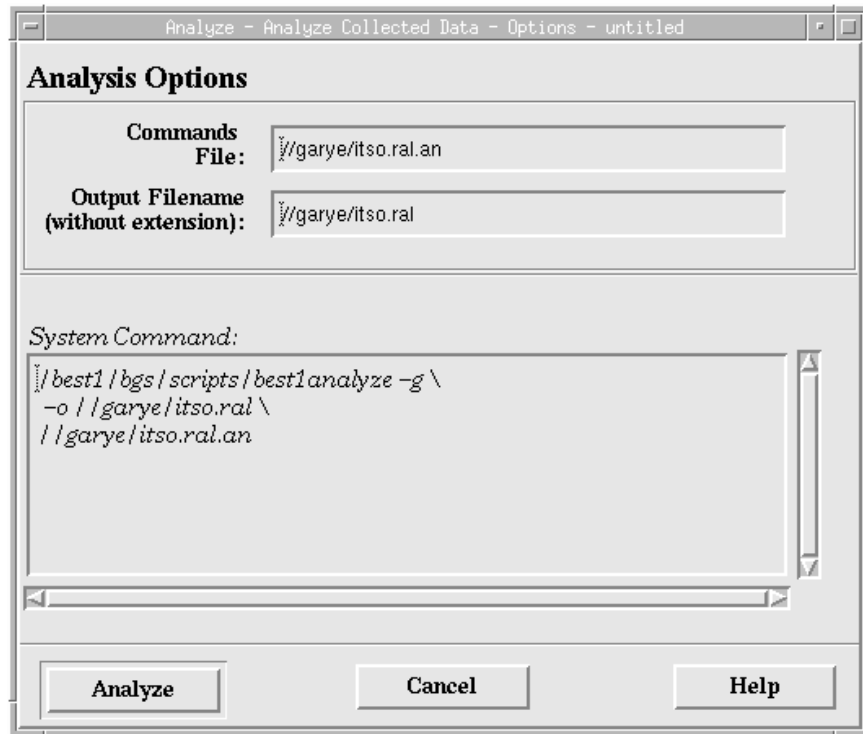


Figure 49. Analyze - Analyze Collected Data - Options Window

Click **OK** on the next dialog box, which simply indicates to you that BEST/1 will inform you once the data analysis is complete. The box shown in Figure 50 will appear once the analysis is complete.



Figure 50. Analyze Completed Window

2.6.3.9 Analyze Module Output

The output from the BEST/1 Analyze module is in one of two forms:

- Performance reports
- Models for use in BEST/1-Predict

Performance Reports: With reference to Figure 39 on page 48 select **Report** from the menu bar of the BEST/1 for UNIX - Analyze - Main window as shown in Figure 39 on page 48. Select either of the types below depending on the type of analysis performed and required reports:

- System
- or
- RDBMS

Selecting the System reports allows you to access the following reports:

- Nodes
- Disks
- Users
- Workloads
- Transactions
- Processes
- Networks

Select any of the above to produce an on-screen report. Shown below is a sample summary report.

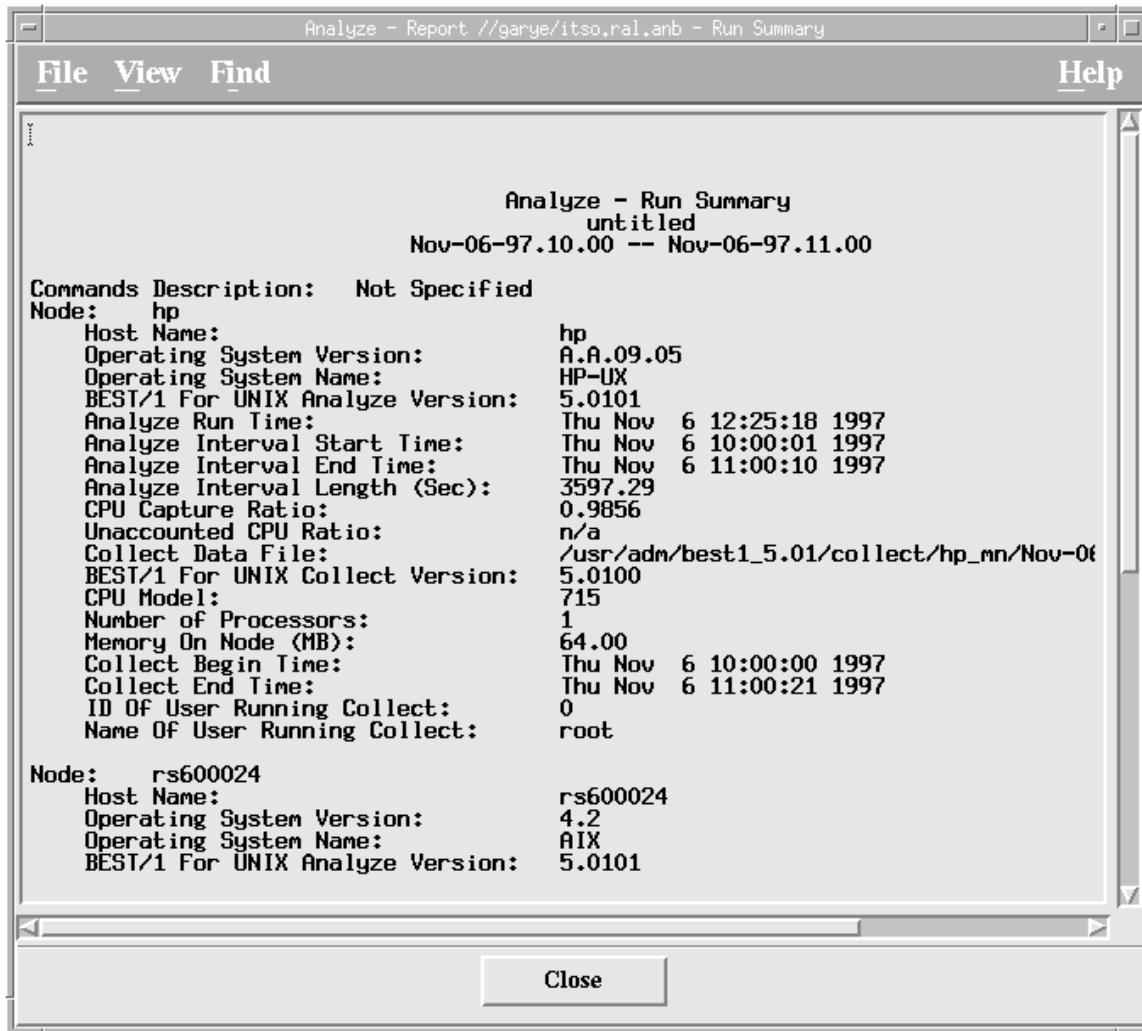


Figure 51. Sample Summary Report from Analyze

Output Model File for Predict: The Analyze module also creates an output file for use by the Predict module. This file has a .md file extension and is created in the same directory as the analyze command file (*.an) when the collect files are analyzed.

2.6.4 BEST/1-Predict

Use the Predict module to analyze the model created by the analyze module and predict performance and capacity issues depending on changes to the model, for example:

- Upgrading hardware in a node
- Increasing or decreasing the number of users on a certain node or application
- Changes in applications
- Moving users across servers

When using BEST/1-Predict to perform "What-if ?" scenarios, it is important that the data analyzed from the BEST/1 Collect module is as accurate as possible. The data must also be applicable to the type of analysis being performed.

An example could be as follows.

You want to analyze the effect of allowing an additional 20 users access to an application running on a server. To perform valid predictions you would collect data from this node during the peak utilization periods to assess whether the node can handle the additional load without degradation in performance. There would not be any value in performing a prediction on load data that was collected at midnight on a Saturday because this is not a correct representation of the performance of the server and application.

If there are instances in collected data where server performance and load vary, break these periods up and analyze them separately, otherwise results will not provide an accurate prediction for scenarios that you have created.

Another important consideration when evaluating a model is to take into account not just the systems that you are evaluating, but also consider systems that you are not analyzing but that will have an effect on the systems that you are evaluating. One example of this effect is when you consider network performance: You may only wish to evaluate a single UNIX server, but other servers and devices add a load to the network and that will affect the performance of the server that you are evaluating. Allowances for this effect are included in BEST/1 by the object ROW. ROW signifies "Rest Of the World". Each area of the model (network, nodes, transactions, etc.) have configuration objects for this information.

2.6.4.1 Predict Process

Follow the steps identified below to use the Predict module:

1. Open a model (*.md) file
2. Evaluate the model
3. Verify the model
4. Calibrate the model
5. Produce baseline
6. Modify service level objectives
7. Perform "What if...?" analysis
8. Create growth scenarios

2.6.4.2 Opening a Model

With reference to the BEST/1-Predict main window shown in Figure 52 on page 61, select **Open Model** from the File menu. Locate the .md file in the directory browse list and click **OK** when complete. The status of the Predict main window will be identical to that shown below.



Figure 52. BEST/1 for UNIX - Predict Window

2.6.4.3 Evaluating the Model

The window shown in Figure 52 indicates that the loaded model requires evaluation. Click on the **Evaluate Model** button to begin the evaluation. When the evaluation starts, a status window appears identifying any errors in the evaluation and indicating when the evaluation is completed. An example of this window is shown in Figure 40 on page 49. Click **Close** when you have finished viewing any messages that are displayed.

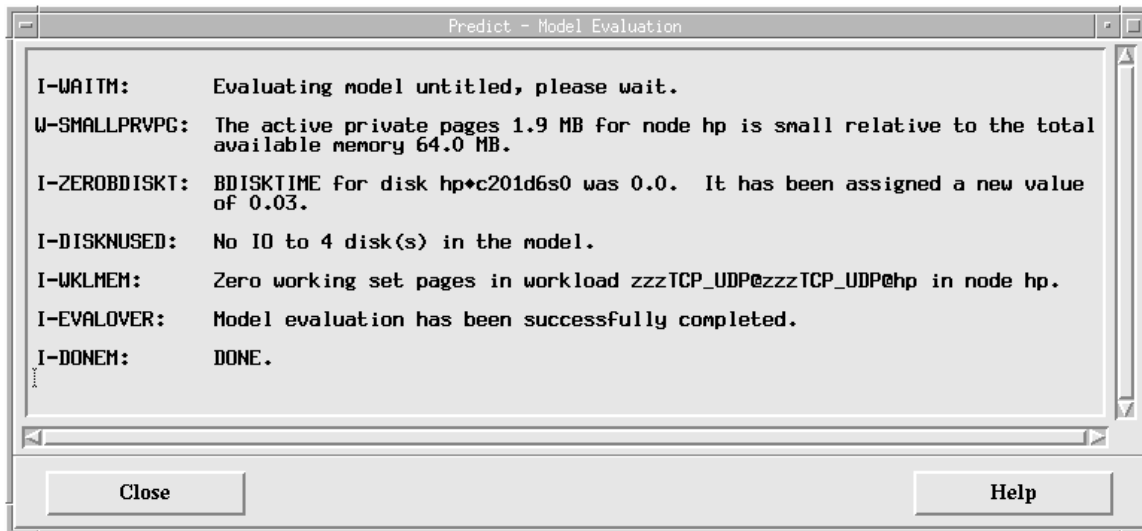


Figure 53. Predict - Model Evaluation Status Window

When the evaluation has completed successfully, a window similar to the one shown in Figure 54 on page 62 is displayed.



Figure 54. BEST/1 for UNIX - Predict - //garye/itso.ral.md

The window shown in Figure 54 represents the separate areas of the model: networks, nodes, transactions, workloads, transactions, disks and logical volumes (LVs) as pie charts. The pie charts indicate by color coding using green, yellow and red the status of that specific area of the model as meeting the objectives, exceeding guidelines or exceeding thresholds, respectively.




Figure 55. BEST/1-Predict Pie Chart Detail

An example of this is that if the Disks... pie chart was displayed as indicated above, then it would be interpreted in the following way:

- There are six disks in the model.
- Three disks meet objectives.
- Two disks exceed guidelines.
- One disk exceeds thresholds.

The pie charts are created from the performance figures in the model file (.md) compared to preset service level objectives. These service level objectives have default values set which may not be applicable to your requirements. The pie charts will change as necessary if you modify the service level objectives (see 2.6.4.5, "Modifying Service Level Objectives" on page 64).

2.6.4.4 Verifying the Model

Verify the accuracy of the model to ensure that it is a true representation of the systems that you are evaluating. The verification is done by clicking on each of the buttons above the individual pie charts as seen in Figure 54. Clicking on these buttons displays the objects in that category of the model, for example, clicking on the **Nodes...** button displays a list of all nodes in the model. Edit each object by clicking on the object and then selecting **Edit...** from the **Edit** menu of the window (see Figure 56 on page 63).

Re-evaluate the model if any changes have been made to any object.

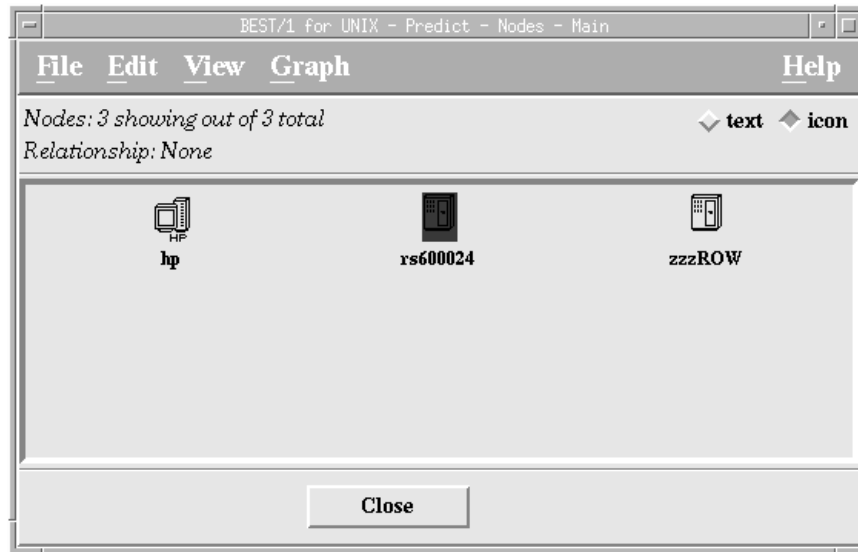


Figure 56. BEST/1 for UNIX - Predict - Nodes - Main Window

Modify any fields that show incorrect values (refer to Figure 57 on page 64 for an example showing node fields that can be changed). In the example of node configuration, the configurable parameters in Figure 57 on page 64 are Count and Memory. Refer to online help for more information regarding these parameters.

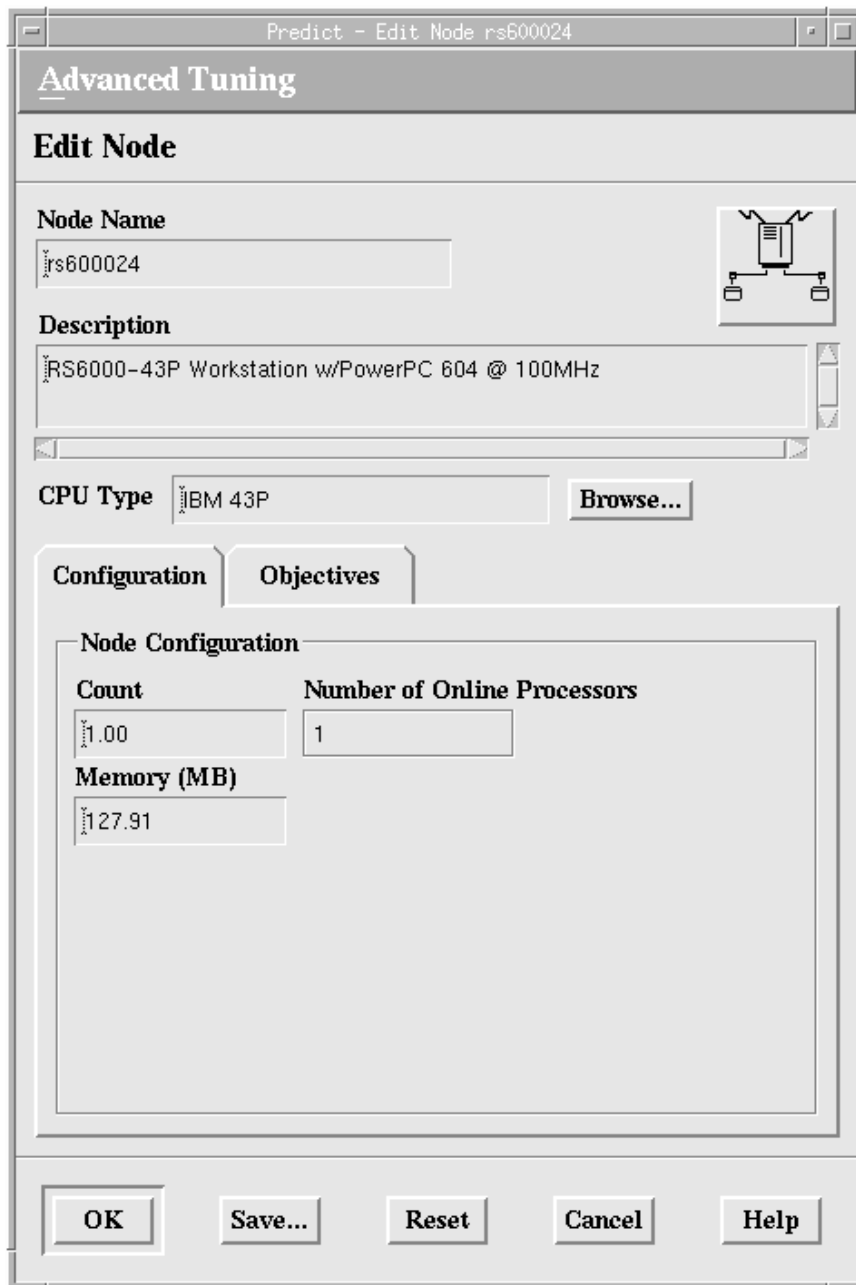


Figure 57. Predict - Edit Node rs600024 Window, Showing Node Configuration

2.6.4.5 Modifying Service Level Objectives

Service Level Objectives (SLOs) can be modified at any stage of the prediction process, but it is logical to set them prior to performing any changes to the model, other than those required during model verification.

To modify the service level objectives and thresholds click on any button above the pie charts in the window shown in Figure 54 on page 62 to show available objects in the model. The example window in Figure 57 shows the objectives for node rs600024. Change the Guideline and Threshold values as required and click **OK** when complete. Click **Close** on the BEST/1 for UNIX - Predict - Nodes - Main

window to return to the BEST/1 for UNIX - Predict window. Re-evaluate the model after making changes to assess the impact of those changes.

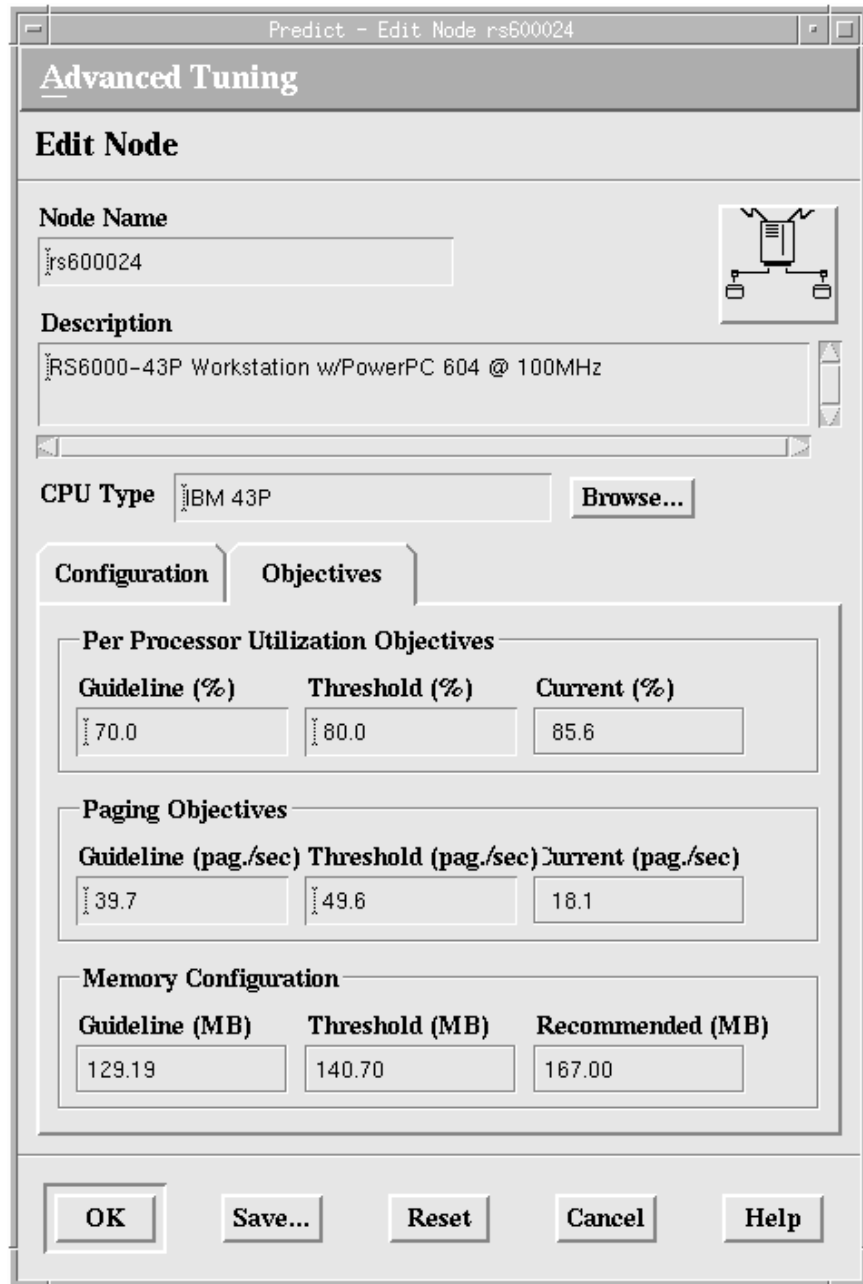


Figure 58. Predict - Edit Node rs600024 Window, Showing Service Objectives

2.6.4.6 Baselining the Model

Baselines are used so that any predictions made can be compared back to the original model to ascertain changes in performance.

Once you are happy that the model represents your systems accurately and has been evaluated, click on the baselined check box in the top right corner of the BEST/1 for UNIX - Predict window.

2.6.4.7 Performing "What if...?" Analysis

This section describes how to evaluate the impact of system changes on the performance of the systems being analyzed. After baselining of the system is complete variables can be changed and the impact of these changes assessed prior to implementing them.

Modify objects in the same way as described in 2.6.4.4, "Verifying the Model" on page 62, by clicking on the appropriate button in the BEST/1 for UNIX - Predict window and editing the properties of the objects, as described below.

The model used in the analysis and prediction so far meets all objectives, shown by the pie charts in Figure 52 on page 61, but being a test system with no "real" users it is not easy to simulate a high load on the system. For the purposes of this example, the workload zzz@rs600024 was modified from a single user and two terminals to 20 users and 20 terminals. This modification of the model is achieved by clicking on the **Workloads...** button in the BEST/1 Predict window, selecting the **zzz@rs600024** workload, and then selecting **Edit...** from the Edit menu. These actions display the window shown in Figure 59 on page 67.

Model Defaults

The BEST/1 model defaults for users, transactions and workloads are named zzz. The workload zzz@rs600024 mentioned in the above paragraph signifies the default workload, zzz, applied to the node rs600024.

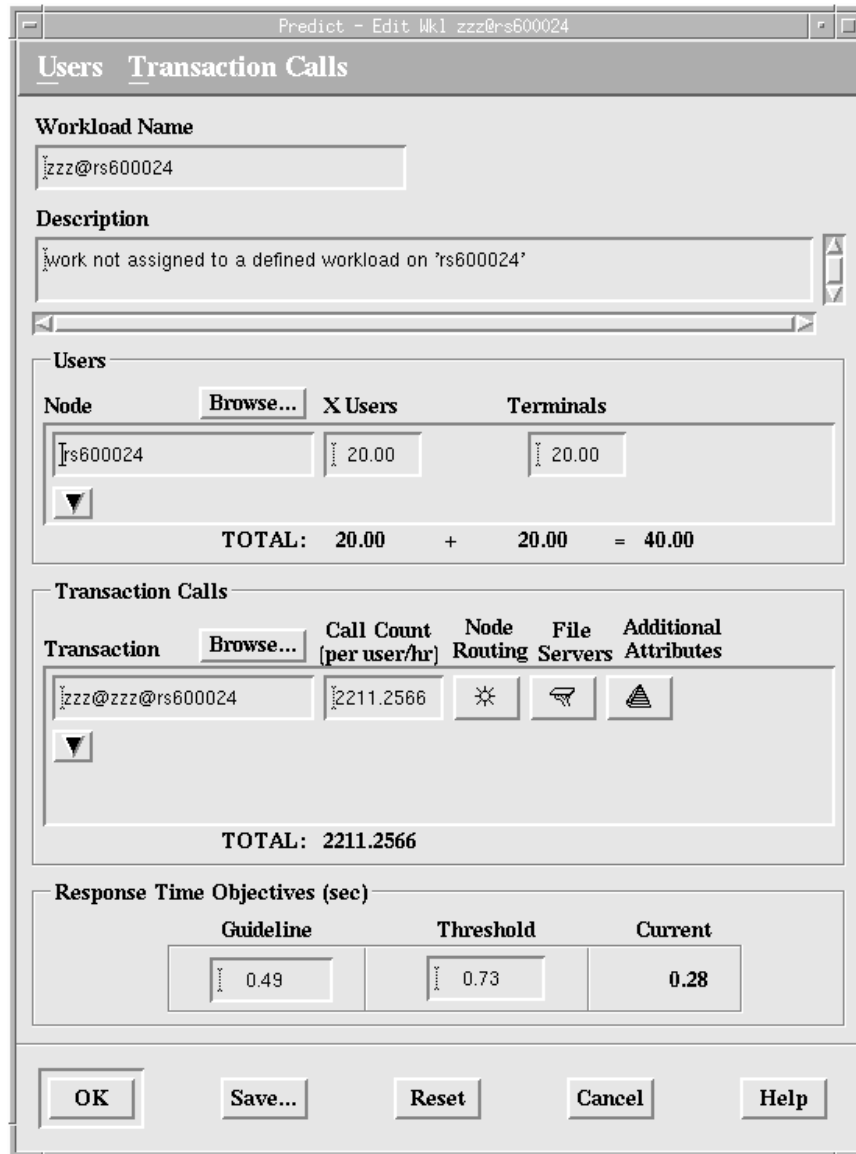


Figure 59. Predict - Edit Wkl zzz@rs600024 Window

Set the values in the X Users field and the Terminals field to 20. This sets the total users to 40, with 20 X Windows users and 20 terminal users. Click on the **OK** button when you have modified the values.

The evaluation of this modified model produced the output shown in Figure 60 on page 68. The output information indicates any areas that now need addressing.

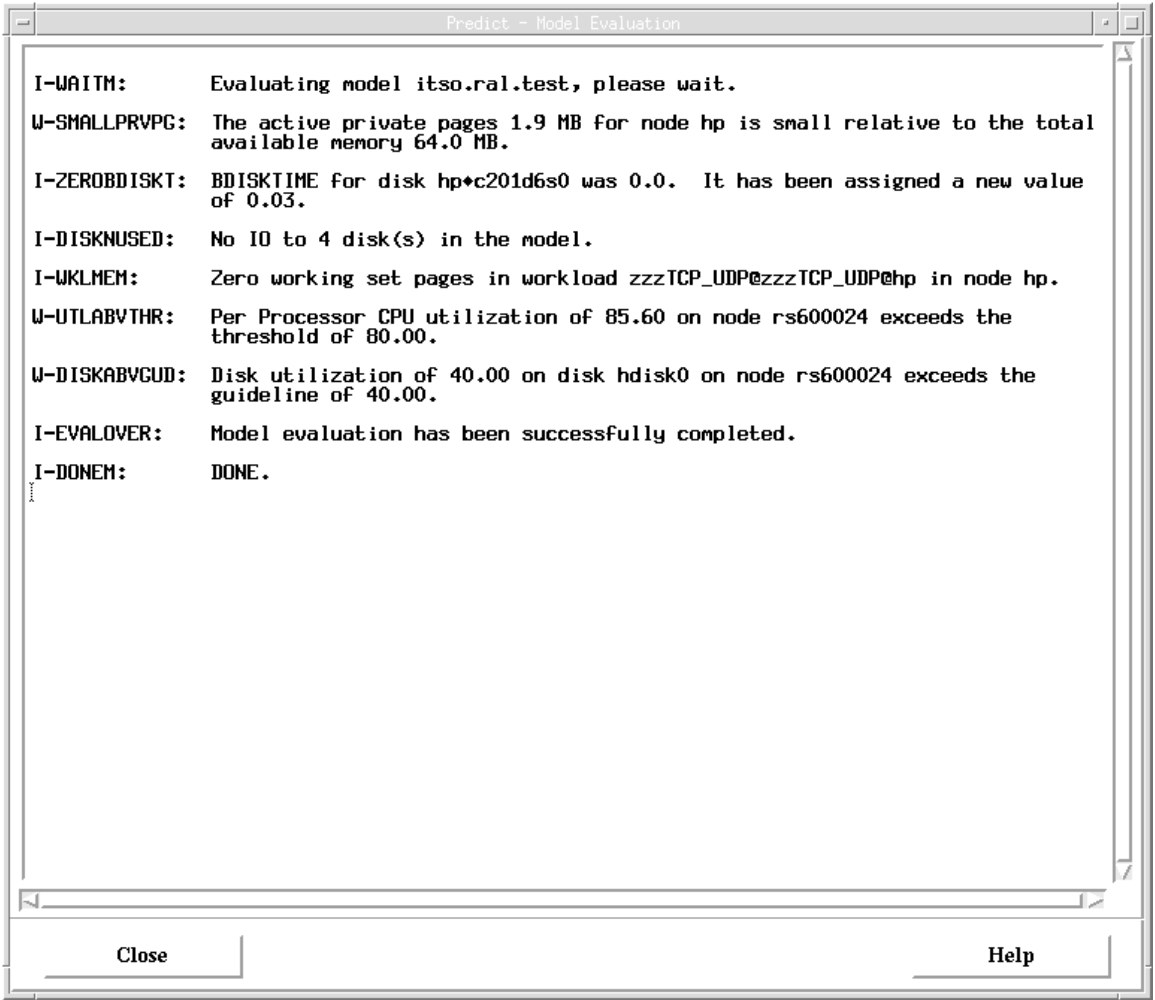


Figure 60. Predict - Model Evaluation Window

The new evaluation results in the pie charts showing that:

- One node exceeds thresholds.
- One disk exceeds guidelines.

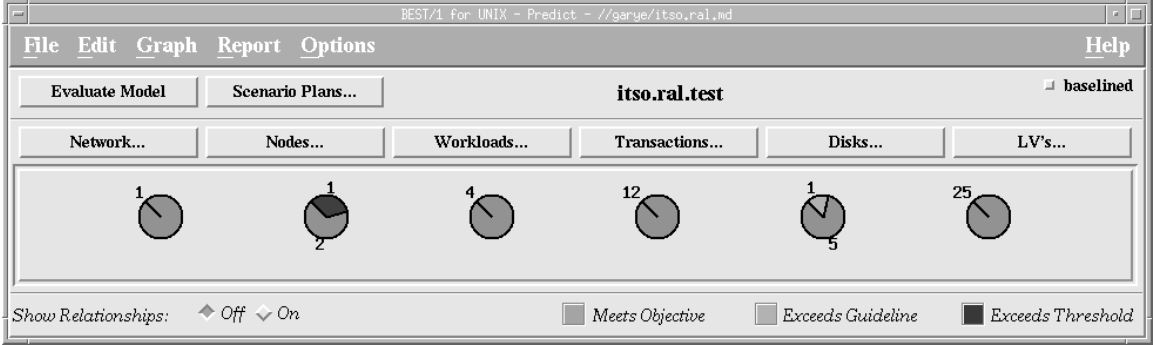


Figure 61. BEST/1 for UNIX - Predict - //garye/itso.ral.md

Further investigation reveals that the disk that is exceeding guidelines is a disk subsystem on the same host that is exceeding thresholds. This condition is not

surprising as these are devices directly affected by the changes (addition of users and terminals to rs600024) made to the model.

Once problem areas have been identified, modify parameters of the model and re-evaluate the model to ascertain the effect. For example, consider the following: What is the effect of increasing the memory of the affected node to 256 MB, or rather, how much memory needs to be added to bring response times to acceptable levels, or would upgrading the processor be a more effective solution?

2.6.4.8 Scenarios

BEST/1-Predict does capacity planning by allowing simulations of workload growth increase in the model. Click on the **Scenario Plans...** button of the BEST/1 for UNIX - Predict window to access the scenario plans configuration. The window shown in Figure 62 appears showing all available scenario plans. (BEST/1-Predict does not supply any default plans.)

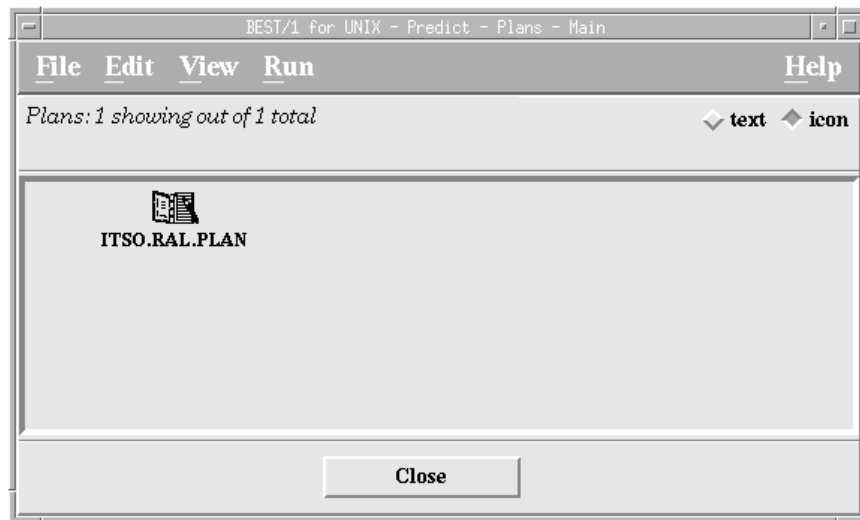


Figure 62. BEST/1 for UNIX - Predict - Plans - Main Window

Click on an existing plan and select **Edit...** from the Edit menu to amend, or select **Create...** from the Edit pull-down menu to create a new plan.

The window shown in Figure 63 on page 70 shows the Predict - Edit Plan window with the configuration for the plan ITSO.RAL.PLAN. This plan shows two configured growth periods: Q1 and Q2, with growths of 20% and 10% respectively. The periods Q1 and Q2 represent quarter 1 and quarter 2 of a year. These periods can be any period for which you want to predict system growth.

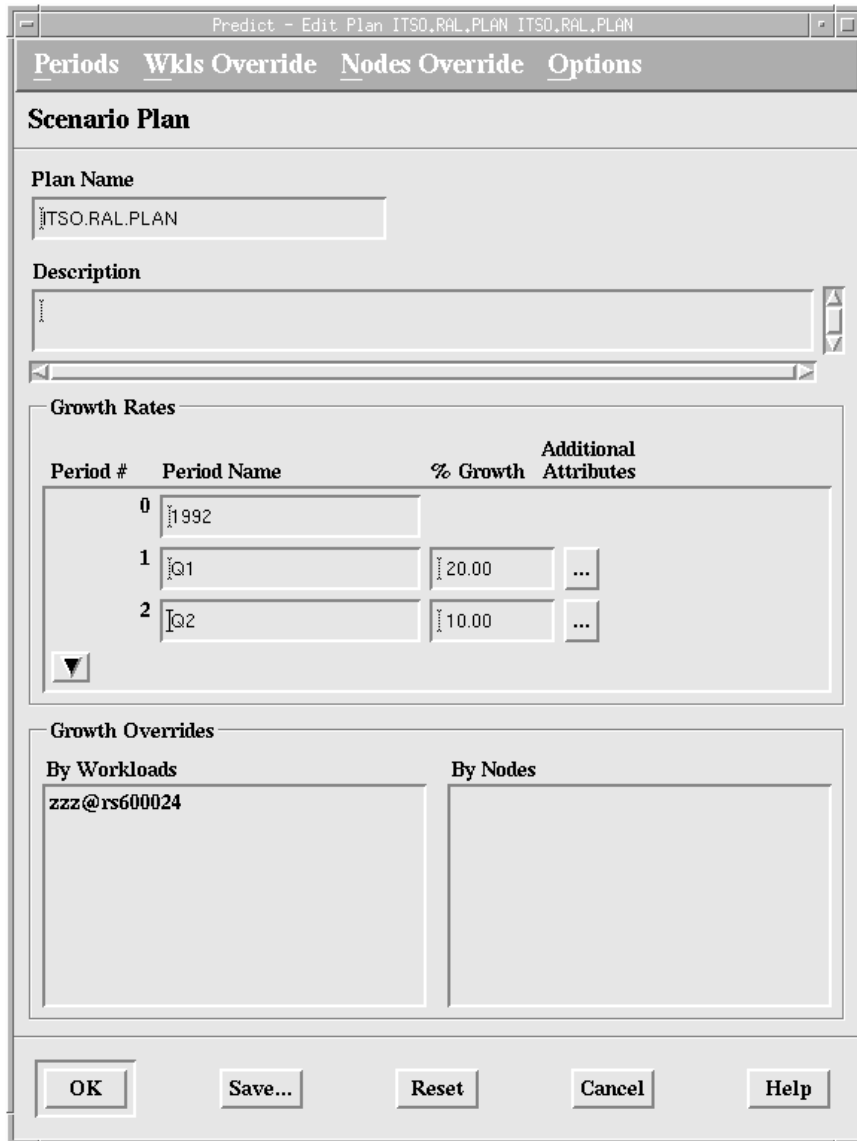


Figure 63. Predict - Edit Plan Window

Exemptions can be included in a scenario plan by adding objects to the Growth Overrides section of the configuration window shown in Figure 63. The exceptions can be in the form of exempt workloads or exempt nodes. Click on the **Add...** option of either the Wkls Override or Nodes Override menu to add overrides for workloads or nodes, respectively. When adding overrides, you have the option of adding overrides only for workloads or nodes in the current model. Click **OK** when configuration is complete.

Once the scenario plan is complete as described above, click on the plan that you want to run and then select the **Run** menu from the BEST/1 for UNIX - Predict - Plans - Main window as shown in Figure 62 on page 69. The window shown in Figure 64 on page 71 appears when you run a plan, except that the Plan Summary area of the window is blank until the plan has completed running.

Select the options required at the base of the window:

- Period of the plan to run

- Mode of execution
- Creation of Visualizer files

Click on the **Run** button at the base of the window when the desired options are selected. A message window will appear briefly as the plan runs. Click the **Close** button when the run is complete. The Plan Summary then contains the information breakdown of the specified growth for each period.

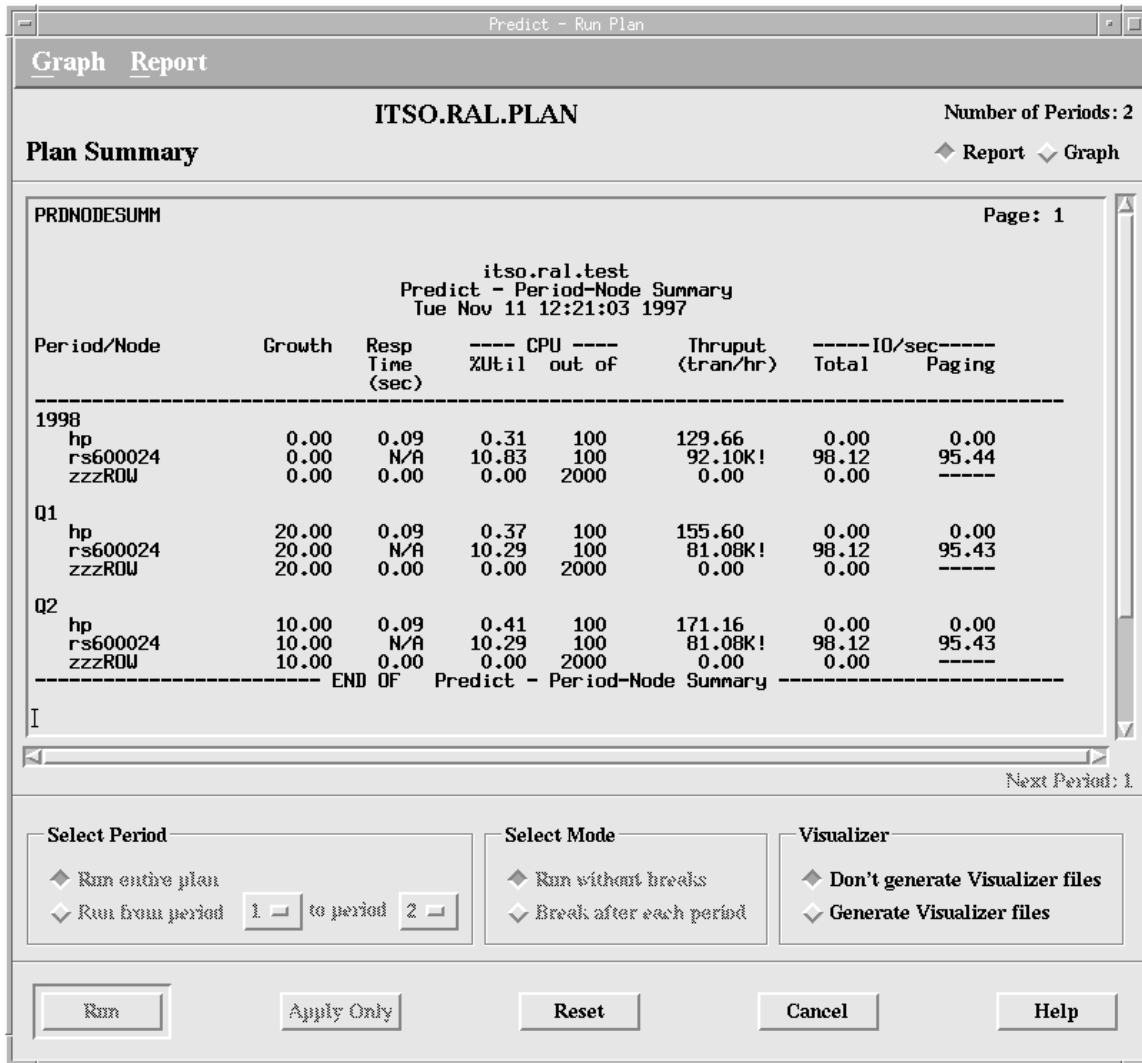


Figure 64. Predict - Run Plan Window, after Completing Scenario Run

2.6.5 BEST/1 Manager

This module is used to schedule periodic tasks for the modules of BEST/1.

The window shown in Figure 65 on page 72 appears when the BEST/1 Manager is started for the first time.



Figure 65. BEST/1 for UNIX - Manager - Main Window

2.6.5.1 Configuring Basic Options

The basic options for BEST/1 Manager are configured in the BEST/1 for UNIX - Manager - Main window.

These options include:

- Data Source and Destination
- Data Collection Time
- Workload Definition

Data Source and Destination: In this section of the Manager - Main window (Figure 65) configuration of the data source allows you to specify whether to Use old data, or Collect new data. Click on the appropriate button in the window to

select one of these options. If selecting to collect new data, you can collect system data, RDBMS data or both by clicking on the button to the right of **Collect new data**. If you select **Use old data**, then BEST/1 checks for the collected data in the location of the BEST/1 Collect repository.

Enter the name of the domain file that contains the nodes that are to be included as part of this process in the Domain Files box. Click on the **Browse...** button to locate the domain file.

Data Collection Time: Select the date and time of the data collection, if collecting new data, in the Data Collection Time section of the Manager - Main window. Select the start and end dates by selecting the Month, Day and Year. The Daily Begin and Daily End times specify the time at which the data collection will start and end on a daily basis. If the Start Date and End Date are the same date, then the collection will only happen once at the time specified, but if the start and end dates are different, then the collection will occur on every day in between the set dates (inclusive) at the times specified. The data collection time is still specified if you select to use old data, in which case the data collection time is used to extract the collected data from the UDF files in the collect repository of the BEST/1 Management Console.

Workload Definition: In this area of the window select whether to use the default workload supplied with BEST/1 or use another workload created previously in the BEST/1-Analyze module.

2.6.5.2 Configuring Advanced Options

BEST/1 Manager allows configuration of other BEST/1 modules, listed below, to give a complete analysis and data prediction:

- Collect
- Analyze
- Predict

Click on the **Advanced Options** item of the Options pull-down menu to access these additional configuration parameters. For more information on each of the options see the online help in each window by clicking on the **Help** button.

The following sections describe the available options for advanced configuration for each module. The options available are the ones that are normally configured in each specific module (rather than BEST/1 Manager), details of which are included in the section describing the functions of that module earlier in this document.

Collect Configuration

The screenshot shows a window titled "Manager - Advanced Features" with a sub-header "Advanced Features". Below the header are five tabs: "Collect", "Analyze", "Predict", "Visualizer", and "Other". The "Collect" tab is selected. The "Collect Options" section contains a text box for "System Collect Name" with the value "holInstance", a text box for "Electronic Mail Destination", and a checked checkbox for "Join RDBMS Data after Collection". The "Interval Options" section contains a "Spill Interval (minutes)" spinner set to 15 and an "RDBMS Sample Interval (seconds)" spinner set to 30. At the bottom are "OK", "Cancel", and "Help" buttons.

Figure 66. Manager - Advanced Features Window - Collect

In the Collect Options field of advanced configuration specify the name of the user to send an e-mail to upon completion of the collect process in the Electronic Mail Destination box, the Spill Interval (minutes) and RDBMS Sample Interval (seconds). Click on the **Analyze** tab when finished.

Analyze Configuration



Figure 67. Manager - Advanced Features Window - Analyze

The following options are available when configuring advanced options of the Analyze module for BEST/1 Manager.

- Run Analyze
- Generate Analyze system reports
- Process RDBMS data
- Generate Analyze RDBMS reports
- Override repositories
- CPU benchmark selection

With reference to Figure 67, click on the button beside the desired option to activate it. In the example shown, Run Analyze and Process RDBMS data are activated.

Predict Configuration



Figure 68. Manager - Advanced Features Window - Predict

The following options are available when configuring advanced options of the Predict module for BEST/1 Manager:

- Run Predict
- Generate Predict reports
- Summarize I/O Server for Visualizer
- Predict commands file name

Enter the name of a command file only if you want to create Predict reports or Visualizer input files for a specific "What if?" scenario. See the file `$BEST1_HOME/bgs/docs/predict_syntax.doc` for information regarding the syntax of commands for this file.

Visualizer Configuration



Figure 69. Manager - Advanced Features Window - Visualizer

The following options are available when configuring advanced options of the Visualizer module for BEST/1 Manager:

- Visualizer interval (minutes)
- Module to create Visualizer files in

The Visualizer interval (minutes) is the amount of collected data to be included in the Visualizer files. This number must be a multiple of the Spill interval (minutes), which you set in the advanced options of Collect.

The Visualizer files are, by default, created by the Predict module, but you can change this to the Analyze module by clicking on the **Predict** button beside the text "Create Visualizer files in".

Other Configuration: The other options available for advanced configuration are:

- Preprocess command file
- Shell command file
- Processing delay
- Use time stamp for output directory



Figure 70. Manager - Advanced Features Window - Other

The preprocess command file is a script file to execute before the collected data is analyzed.

A template for this file is provided with BEST/1:

```
$BEST1_HOME/bgs/scripts/preprocess.template
```

The Shell command file is a file that will be run at the end of the collect run each day.

The Processing delay time, if specified, is the delay in minutes before processing of the collected data occurs. This parameter is useful when new data is being collected and scripts are used to copy the data files from remote hosts to the local repository for analysis. The number specified here is added to the date set for the end of the collect run.

The directory time stamp option, if selected, creates a new subdirectory for each Manager run with the appropriate time stamp. If not selected, then data for each run is written directly into the output directory specified in the Manager - Main window.

Running Manager Configuration: Click on the **Run Manager Commands** button when you have configured all options and advanced options to execute the Manager tasks that you have defined.

2.6.6 BEST/1-Visualizer

BEST/1-Visualizer is a PC-based graphical data analysis tool that allows users to analyze, display, and report performance of UNIX applications on a central workstation.

The performance data is viewed through a series of historical and analytical reports and graphs, and stored in a database where it is automatically maintained.

BGS value-added expertise has been built in to BEST/1-Visualizer to enhance the management of performance.

Using BEST/1-Visualizer, supporting graphs can be used to justify expenditures to management. The need for additional resources can easily be communicated with a trend graph showing the life expectancy of the resource. By capturing and presenting the utilization within the current environment, the planner gains credibility when forecasting the future requirements. It shows thorough understanding of the current computing activity and the ability to understand computing requirements.

BEST/1-Visualizer's intelligence allows for the product to learn normal performance behavior and report on deviations. This functionality, known as MASF (Multivariate Adaptive Statistical Filter), enhances the performance management capabilities of the BEST/1-Visualizer user. BEST/1-Visualizer also has turnkey performance graphs that help identify resource utilization trends and graphically display the results of BEST/1-Predict modelling.

With BEST/1-Visualizer, users simply double-click on a performance problem area within a graph and a graph of finer detail is displayed exposing the origin of a performance problem. This drill-down capability is useful when correlating multiple performance problems with the same characteristics and identifying trends.

In addition to hardcopy and console graphics, BEST/1-Visualizer generates HTML code that can be automatically published on the corporate Web. Hourly, nightly, weekly or monthly, performance information can be disseminated to satisfy the requests of management and end users. Performance information, once difficult to communicate, is now readily available to managers and users at any browser-equipped desktop.

The following figures show example graphs created with BEST/1-Visualizer.




Figure 71. BEST/1-Visualizer Example 1

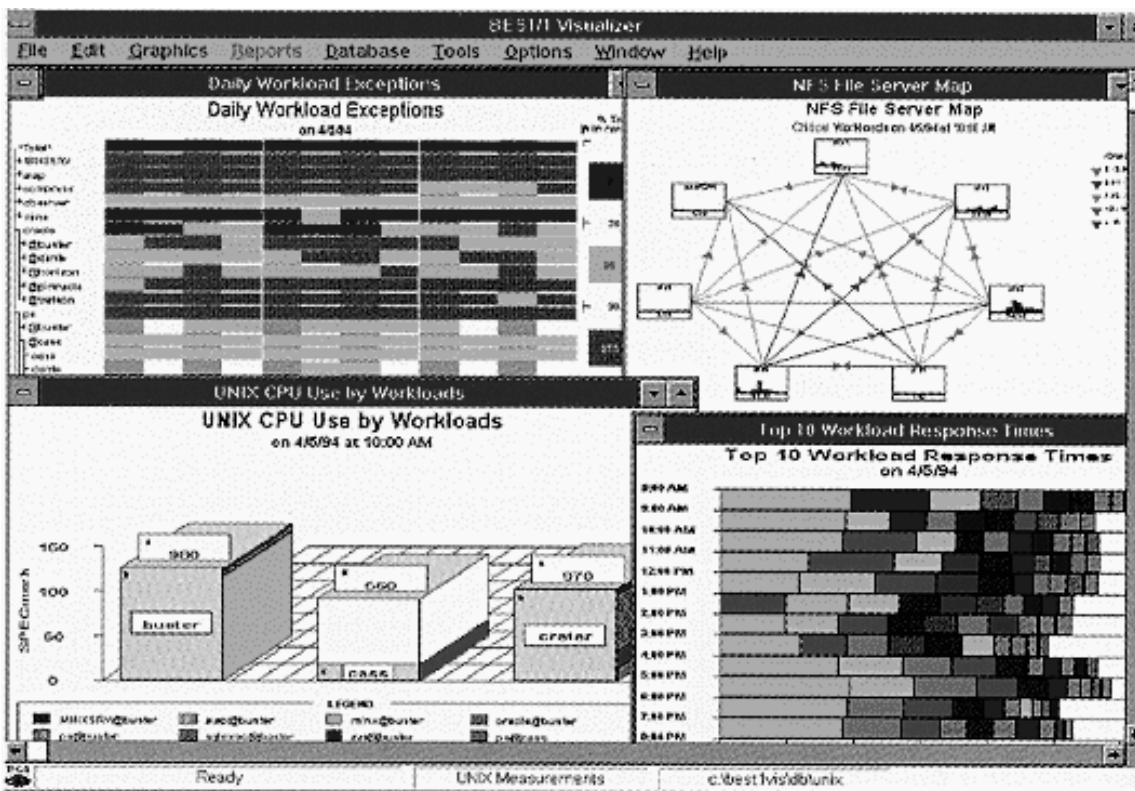


Figure 72. BEST/1-Visualizer Example 2

Now, that we have installed BEST/1 and have looked at some of the most important functions, we go ahead and install BEST/1+ for TME 10.

2.7 Setting Up BEST/1+ for TME 10 in Our Environment

In this section we describe how to install BEST/1+ for TME 10 in our existing TME 10 environment. We do not explain how to install and configure the TME 10 Framework and the Tivoli core applications but focus on the installation of BEST/1+ for TME 10.

Once BEST/1+ for TME 10 is installed, we show examples of how this module can be used to administer BEST/1 from TME 10.

2.7.1 Overview of the TME 10 Environment

At the ITSO test site we have set up a Tivoli environment with four UNIX machines and two Windows NT PCs.

All of them are in one Tivoli Management Region named rs600024-region. Figure 6 on page 17 gives an overview of our TME 10 environment.

The following table shows the installed TME servers, managed nodes, their operating system, function and installed software.

Hostname	Operating System	Function	Software
rs600024	AIX 4.2	TMR Server TEC Server BEST/1 Server	TME 10 Framework 3.1 Revision C with Patch 3.1.2 TME 10 Enterprise Console 3.1. Server Tivoli/Sentry Version 3.0 Tivoli/Sentry Patch 3.0.2 Tivoli/Sentry 3.0 Monitors Tivoli/Sentry 3.0.2 Monitors Patch TME 10 Software Distribution Release 3.1 BEST/1 Performance Assurance for UNIX Release 5.01 BEST/1+ for TME 10 Release 1.0

Table 1 (Page 2 of 2). TME 10 Environment

Hostname	Operating System	Function	Software
rs600026	AIX 4.2	Tivoli Managed Node BEST/1 Managed Node	TME 10 Framework 3.1 Managed Node Tivoli/Sentry Version 3.0 Tivoli/Sentry Patch 3.0.2 TME 10 Software Distribution Release 3.1 BEST/1 Agent 5.01 BEST/1 Collector 5.01 BEST/1+ for TME 10 Release 1.0
hp	HP-UX 9.05	Tivoli Managed Node BEST/1 Managed Node	TME 10 Framework 3.1 Managed Node Tivoli/Sentry Version 3.0 Tivoli/Sentry Patch 3.0.2 TME 10 Software Distribution Release 3.1 BEST/1 Agent 5.01 BEST/1 Collector 5.01 BEST/1+ for TME 10 Release 1.0
sun	Solaris 2.5.1	Tivoli Managed Node BEST/1 Managed Node	TME 10 Framework 3.1 Managed Node Tivoli/Sentry Version 3.0 Tivoli/Sentry Patch 3.0.2 TME 10 Software Distribution Release 3.1 BEST/1 Agent 5.01 BEST/1 Collector 5.01 BEST/1+ for TME 10 Release 1.0
hp	HP-UX 9.05	Tivoli Managed Node BEST/1 Managed Node	TME 10 Framework 3.1 Managed Node Tivoli/Sentry Version 3.0 Tivoli/Sentry Patch 3.0.2 TME 10 Software Distribution Release 3.1 BEST/1 Agent 5.01 BEST/1 Collector 5.01 BEST/1+ for TME 10 Release 1.0

Note

In the above table, the term Tivoli/Sentry is used for the Tivoli core application now called TME 10 Distributed Monitoring. When speaking of Tivoli/Sentry or TME 10 Distributed Monitoring we refer to the same application.

For our examples we used one TMR where the TMR server resides on node rs600024. On this machine we installed the following Tivoli core applications:

- TME 10 Enterprise Console to handle the BEST/1 events
- TME 10 Software Distribution to distribute BEST/1 software to the managed nodes
- TME 10 Distributed Monitoring to monitor the BEST/1 application

All other UNIX machines (rs600026, sun and hp) are connected as managed nodes to the TMR.

2.7.2 Installing BEST/1+ for TME 10

In the following we describe the installation of the BEST/1+ for TME 10 module. This has to be done on the TMR server.

You should always refer to *BEST/1+ for TME 10 User's Guide*, 97095-00 as at this point we describe only the parts of the installation that are specific to our environment.

2.7.2.1 Preparing the Installation

We recommend taking a backup of the Tivoli object database before you install BEST/1+ for TME 10. To do this, check the consistency of the Tivoli database first. Run the following command from the command line:

```
# wchkdb
```

If the command reports any discrepancies, try to repair the database by specifying the -u option. You may have to run the command multiple times until consistency is restored.

By default, the database backups will be created in the directory \$DBDIR/backups when you execute the following command on the command line of the TMR server:

```
# wbkupdb
```

It is a good idea to assign self-explaining names to your database backups by using the -d flag, so you can identify your backups more easily. Do as follows:

```
# wdbckup -d $DBDIR/backups/before_best1_installation
```

You will receive an output like this:

```
Starting the snapshot of the database files for rs600024...
.....
Starting the snapshot of the database files for hp...
.....
Starting the snapshot of the database files for rs600026...
...
Starting the snapshot of the database files for sun...
```

As you see we backed up the databases of all managed nodes. Although you can specify single nodes to be backed up, we think it is more secure to include all managed nodes.

2.7.2.2 Installation Process

We used the TME 10 Desktop as seen in Figure 73 to install the BEST/1+ for TME 10 module on our TMR server (rs600024). Remember that the module must always be installed on the TMR server.

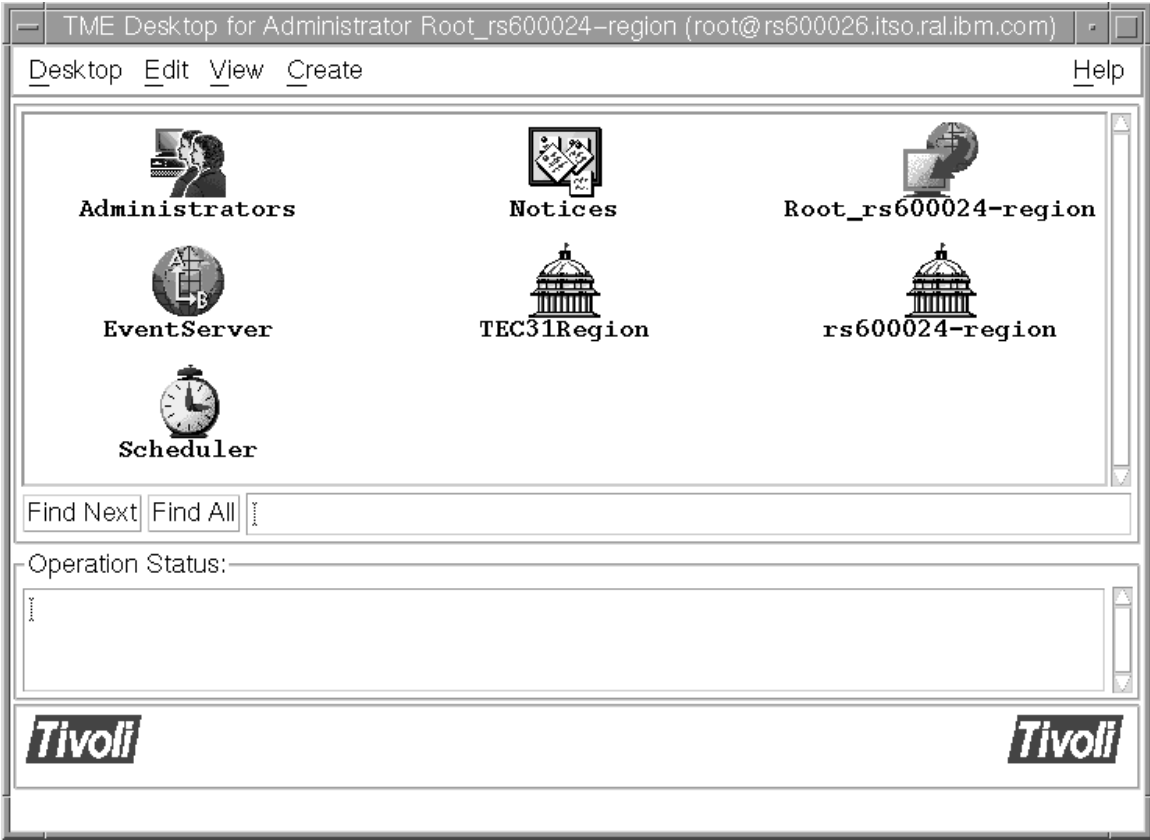


Figure 73. The TME 10 Desktop

From the **Desktop** menu select the **Install -> Install Product...** options to display the Install Product window.

Note

You must have the Tivoli install_product TMR role assigned in order to be able to install additional products in your TMR.

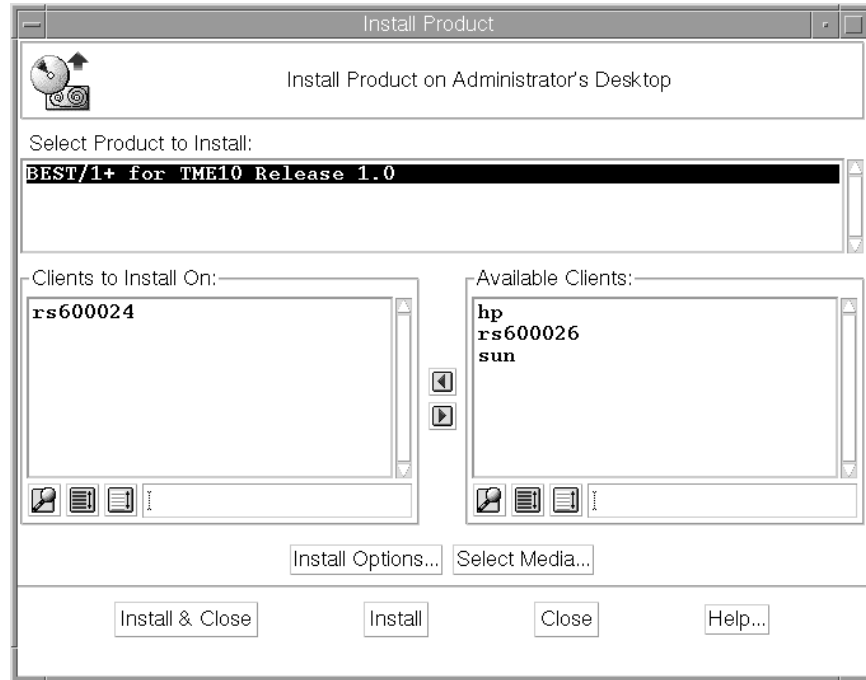


Figure 74. TME 10 Install Product Window

Here, you have to specify the following:

- The installation media (click the **Select Media...** button)
- The product to install (click **BEST/1+ for TME 10 Release 1.0**)
- The installation options (click the **Install Options...** button)

This is described in the *BEST/1+ for TME 10 User's Guide*, 97095-00 in full detail.

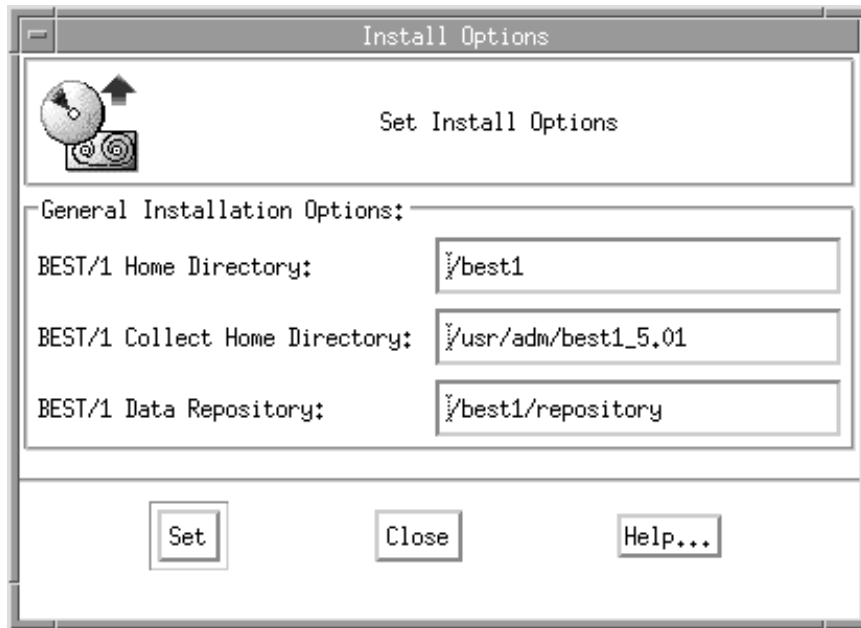


Figure 75. BEST/1+ for TME 10 Install Options

In the Install Product window click on **Install Options...** to get the Install Options window as displayed in Figure 75.

The installation options include the following:

- The BEST/1 Home Directory contains all executable and miscellaneous files of BEST/1. The environment variable BEST1_HOME points to this directory.
- The BEST/1 Collect Home Directory is the location of the BEST/1 Collect. It is referenced with the environment variable BEST1_COLLECT_HOME. The system assumes this directory to be always /usr/adm/best1_5.01, since it will actually be a link to the BEST/1 home directory. So, be sure that you have entered the correct value.
- The BEST/1 Data Repository is used by the BEST/1 Collect to store all collected performance data. Since the amount of data may be rather large (see 2.4.2.1, “Hardware” on page 20) you might enter the name of a file system. We used /best1/repository as the data repository.

Select **Set** to return to the Install Product window.

In our environment the resulting Install Product windows looks as displayed in Figure 74 on page 85. When you click the **Install & Close** or the **Install** button the installation of the module will be started. The system will bring up the Product Install window, where you can see the single steps of the installation. Wait for the Finished product installation message within this window indicating the completion of the installation.

Note

If any problem occurs, analyze the following text files for troubleshooting:

- /tmp/BEST1PLUS_ALIDB_after.output
- /tmp/BEST1PLUS_ALIDB_after.error

The first file contains a list of executed tasks that have been performed. The second file gives detailed information why one of those tasks failed.

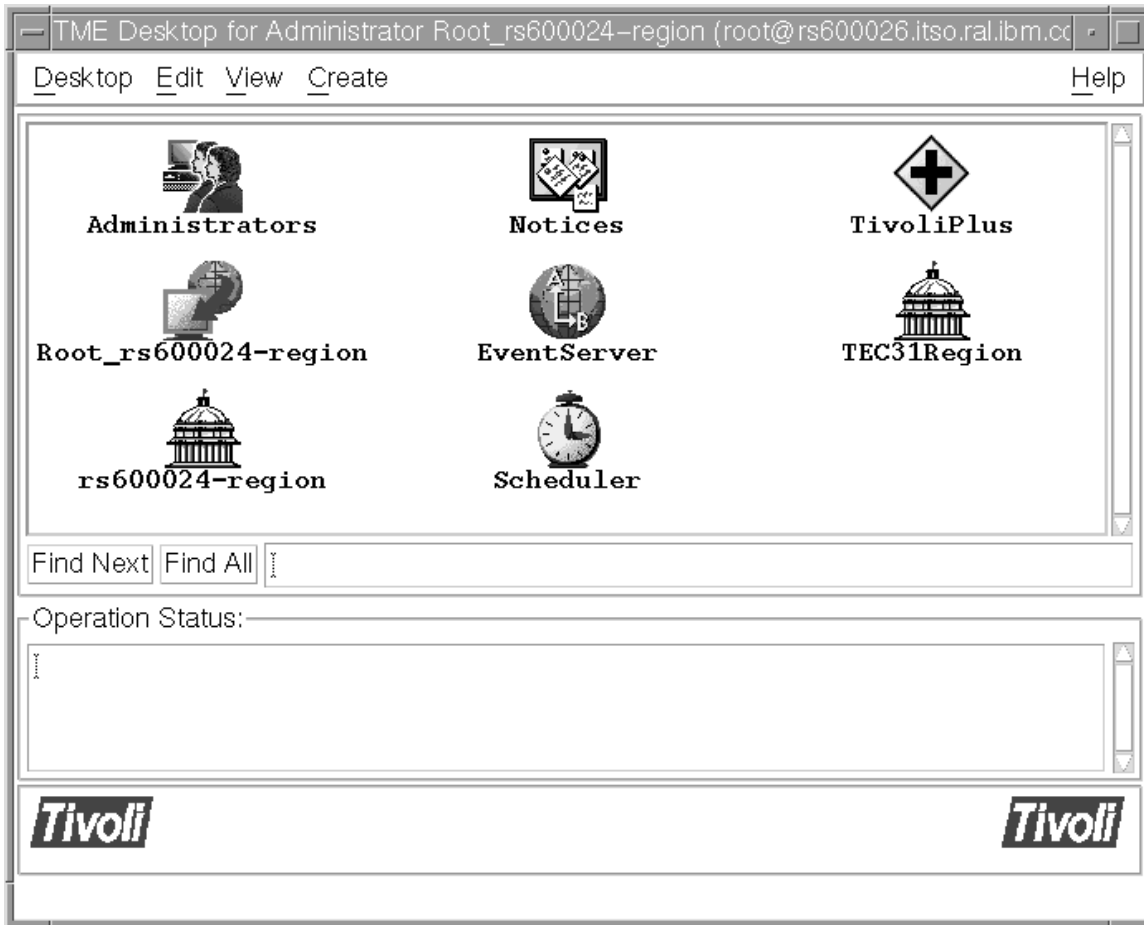


Figure 76. The TivoliPlus Icon on the TME 10 Desktop

After the successful installation you will see the TivoliPlus icon on the TME 10 Desktop as shown in Figure 76.



Figure 77. Select Open...

To access the TivoliPlus collection double-click on the **TivoliPlus** icon on the TME 10 Desktop or select **Open...** from the context menu of the icon (see Figure 77).

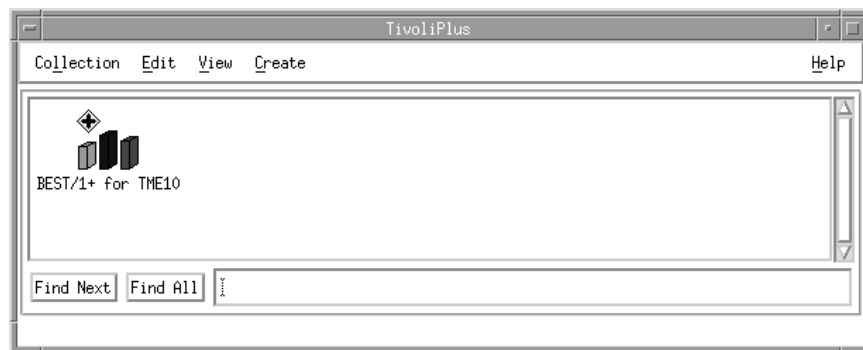


Figure 78. The TivoliPlus Window

You will see the BEST/1+ for TME 10 icon in the TivoliPlus window (Figure 78). This is the starting point for working with BEST/1+ for TME 10.

Before you continue we recommend that you back up the Tivoli Database again (see 2.7.2.1, "Preparing the Installation" on page 83).

2.7.2.3 De-installation of BEST/1+ for TME 10

During the installation some subdirectories will be created in the Tivoli installation directory. The location of these directories can be determined with the following command sequence:

```
# oserv=objcall 0.0.0 get_oserv # This lists object id of Tivoli's oserv
# objcall $oserv query install_dir
```

The following list of subdirectories assumes the Tivoli installation directory to be /usr/local/Tivoli/bin:

- /usr/local/Tivoli/bin/aix4-r1/.installed
This directory contains flag files for all installed Tivoli components. The entry for BEST/1+ for TME 10 is called BEST1PLUS_BIN.
- /usr/local/Tivoli/bin/aix4-r1/TME/PLUS/best1

This directory contains one executable file called `seminfo`. This program is part of the distributed monitoring package contained in the module and is used to evaluate the number of available semaphores and semaphore IDs.

- `/usr/local/Tivoli/bin/generic_unix/.installed`

Similar to the first one, this directory contains flag files. For BEST/1+ for TME 10 two entries are created: `BEST1PLUS_LINK` and `BEST1PLUS_GBIN`.

- `/usr/local/Tivoli/bin/generic_unix/TME/PLUS/best1`

This directory contains executable and script files that are used to manage BEST/1 from the TME 10 Desktop.

- `/usr/local/Tivoli/msg_cat/.installed`

This directory contains the flag file for the installed message catalogs.

- `/usr/local/Tivoli/msg_cat/C`

All installed Tivoli message catalog files are stored here. Two files belong to BEST/1+ for TME 10: `bgsCat.cat` and `bgsHelpCat.cat`.

Further, some entries are made in the object database. To evaluate the ID of your managed node, use the following command:

```
# mid='objcall 0.0.0 get_host_location'
```

Subsequently, you can list the stored location of the BEST/1 data repository and the Collect home directory using the following commands:

```
# wgetattr _BEST1_DataRepository $mid
# wgetattr _BEST1_CollectHome $mid
```

At install time, some TME 10 Distributed Monitoring are registered in the object database, which can be listed with the following commands:

```
# wlsmon "BEST/1 Remote Monitors"
# wlsmon "BEST/1 Central Monitors"
```

While working with the module, additional files will be created in these directories:

- `/usr/local/Tivoli/bin/aix4-r1/TME/TEC/adapters/bin`

This directory contains the source files of the TEC adapters (`tecad_best1.pl` and `tecad_collect.pl`) and a script file (`tecad_best1.sh`) to start and stop the TEC event adapter. Another file `postmsg.sh` is used to generate events.

- `/usr/local/Tivoli/bin/aix4-r1/TME/TEC/adapters/etc`

Here, a configuration file called `tecad_best1.conf` is stored. It defines the environment variables for TME 10 Enterprise Console adapter and for error logging.

- `/usr/local/Tivoli/itso_rulebases/best1`

TME 10 Enterprise Console uses rules to determine the appropriate actions to perform when a certain event from an event adapter arrives. We used this directory to store the rule base for BEST/1 events. You can define the location when you set up TME 10 Enterprise Console.

Whenever you want to uninstall the module, you have to perform the following steps:

1. Delete the icons from the TME 10 Desktop.

2. Execute the de-installation script `deinstall.sh` on the BEST/1+ for TME 10 CD-ROM. This will delete all files and directories created at installation time. It will also remove the BEST/1 entries in the object database.
3. Remove the TME 10 Enterprise Console adapters and the rule bases created for BEST/1+ for TME 10.

Finally, run the command `wchkdb -u` to make sure the Tivoli database contains no inconsistencies, and back up the database as explained in 2.7.2.1, "Preparing the Installation" on page 83.

2.7.3 Overview of BEST/1+ for TME 10 Functionality

The BEST/1+ for TME 10 module contains a task collection for all necessary steps during installation, usage and administration of the BEST/1 software on the server and on the managed nodes. Furthermore, it provides distributed monitoring profiles, event adapters for TME 10 Enterprise Console and an indicator collection to monitor the availability of BEST/1 daemons and for the translation of BEST/1 events into TEC events.

To access these tasks double-click on the **BEST/1+ for TME 10** icon in the TivoliPlus window (see Figure 78 on page 88).

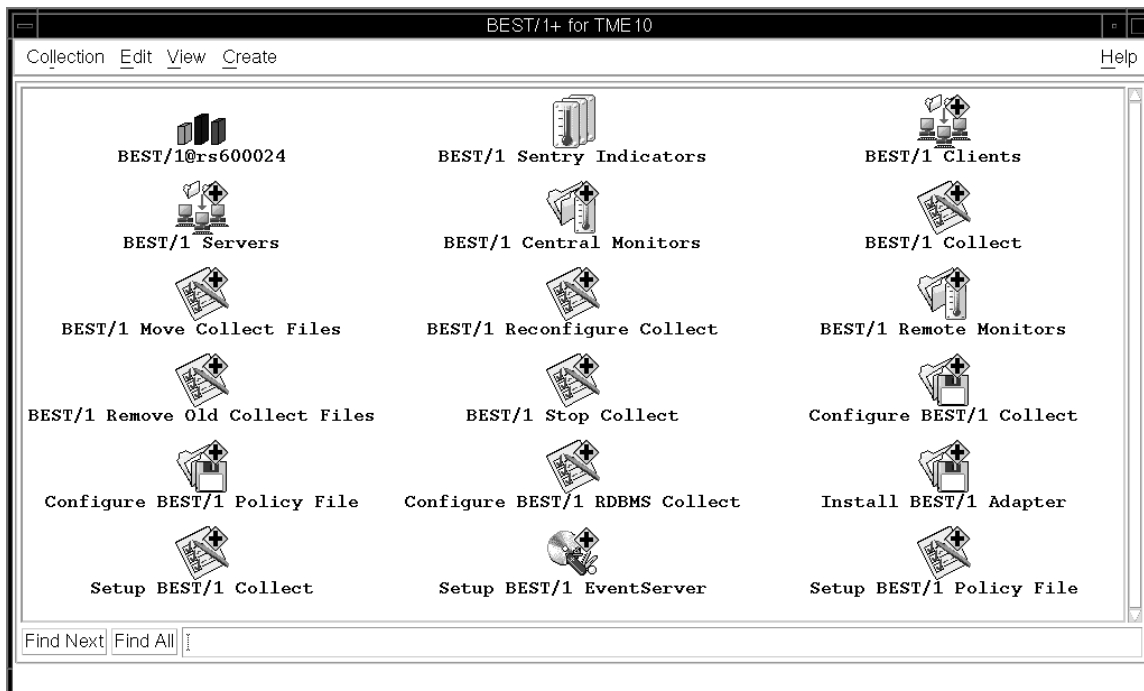


Figure 79. The BEST/1+ for TME 10 Window

This will open the BEST/1+ for TME 10 window containing the icons for all predefined components as shown in Figure 79.

We first give a short overview of all the available functions and then explain these functions in more detail in concrete examples.



Figure 80. The Task Icon

The task icon (Figure 80) stands for a Tivoli resource, representing a specific action or operation performed within the TME 10 environment. Adding specific task endpoints, format and destination of output, execution parameters and execution mode definitions can be specified with a task.

When a task is to be scheduled, it needs to be converted to a TME 10 job. A job is simply a task that is to be scheduled for execution instead of executing it immediately.

The following tasks exist in the BEST/1+ for TME 10 collection:

- Setup BEST/1 Collect
- Setup BEST/1 Policy File
- Configure BEST/1 RDBMS Collect
- BEST/1 Reconfigure Collect
- BEST/1 Stop Collect
- BEST/1 Move Collect Files
- BEST/1 Remove Old Collect Files
- BEST/1 Collect



Figure 81. The File Package Icon

File packages are used by TME 10 Software Distribution to distribute software from a TMR server to a number of TME endpoints, such as TME Managed Nodes or TME PC Managed Nodes.

File packages designate the directories, files, configuration programs and platform-specific options applicable for a specific distribution. The definition of a file package contains only references to the location of the data and does not contain the actual data.

These file packages may be generated within the BEST/1+ for TME 10 collection when executing certain predefined tasks:

- Configure BEST/1 Policy File
- Install BEST/1 Adapter
- Configure BEST/1 Collect



Figure 82. The Profile Manager Icon

A profile is a collection of information corresponding to a system resource. The file package described above is such a profile. Usually, this information is common to a certain number of managed resources, so it makes sense to manage this information at a central point and to distribute it to the managed resources when changes occur. A profile manager is a container for profiles, for example, file packages. Managed resources, for example, managed nodes or other profile managers, subscribe to a profile manager. During a profile distribution these subscribers receive copies of the profiles and their contents.

The collection contains the following profile managers:

- BEST/1 Clients
- BEST/1 Servers

The BEST/1 Clients profile manager is supposed to contain all systems that are clients in the BEST/1 environment, whereas BEST/1 Servers contains all BEST/1 servers.



Figure 83. The TME 10 Distributed Monitoring Profile Icon

A monitor is an entity that monitors a specific aspect of a resource on a managed node. A monitor is defined by response levels, associated response threshold and actions to be performed when a certain response level is reached. Usually, a certain number of (related) monitors are combined in one profile and then distributed to the managed resource altogether.

The following TME 10 Distributed Monitoring collections exist:

- BEST/1 Central Monitors
- BEST/1 Remote Monitors




Figure 84. TME 10 Distributed Monitoring Indicator Collection Icon

An indicator collection simplifies the monitoring of numerous resources by combining multiple indicator icons and status alerts in one window. Each indicator icon within the indicator collection refers to one certain TME 10 Distributed Monitoring profile and designates the current status of the resource. Since a TME 10 Distributed Monitoring profile combines several monitors and since there is only one indicator for each TME 10 Distributed Monitoring profile the indicator reports only the highest severity level.

The BEST/1+ for TME 10 module contains one indicator collection with two indicators, one for the central and one for the remote monitor.

The following sections give an overview of the components of the BEST/1+ for TME 10 module.

2.7.3.1 Components Needed to Set Up Managed Nodes

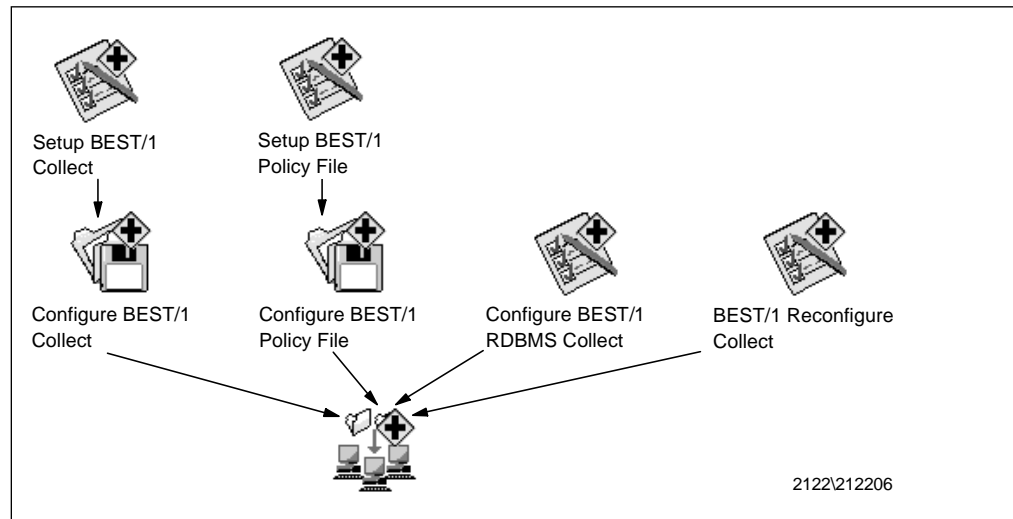


Figure 85. Collect Setup on Managed Nodes

The following components help you to generate TME 10 Software Distribution file packages and to distribute them to the BEST/1 managed nodes:

- Setup BEST/1 Collect: Task to prepare TME 10 Software Distribution file packages for distributing BEST/1 Collect to managed nodes
- Setup BEST/1 Policy File: Generates TME 10 Software Distribution file package containing policy files to be distributed to the managed nodes. Run the Configure BEST/1 Policy to actually distribute the file packages afterwards.
- Configure BEST/1 Collect: Collection of TME 10 Software Distribution file packages containing the BEST/1 Collect software to be distributed to the managed nodes.
- Configure BEST/1 Policy file: Collection of TME 10 Software Distribution file packages containing BEST/1 policy files to be distributed to the managed nodes. By default, it is subscribed to the BEST/1 clients profile manager.
- Configure BEST/1 RDBMS Collect: Begins BEST/1 collection of RDBMS data on remote systems.
- BEST/1 Reconfigure Collect: Task to reconfigure a managed node that has been installed without using BEST/1+ for TME 10 or on which you want to change the data repository or the collect home directory.
- BEST/1 Clients: Profile manager for BEST/1 managed nodes.

2.7.3.2 Components to Set Up TME 10 Enterprise Console and TME 10 Distributed Monitoring for BEST/1

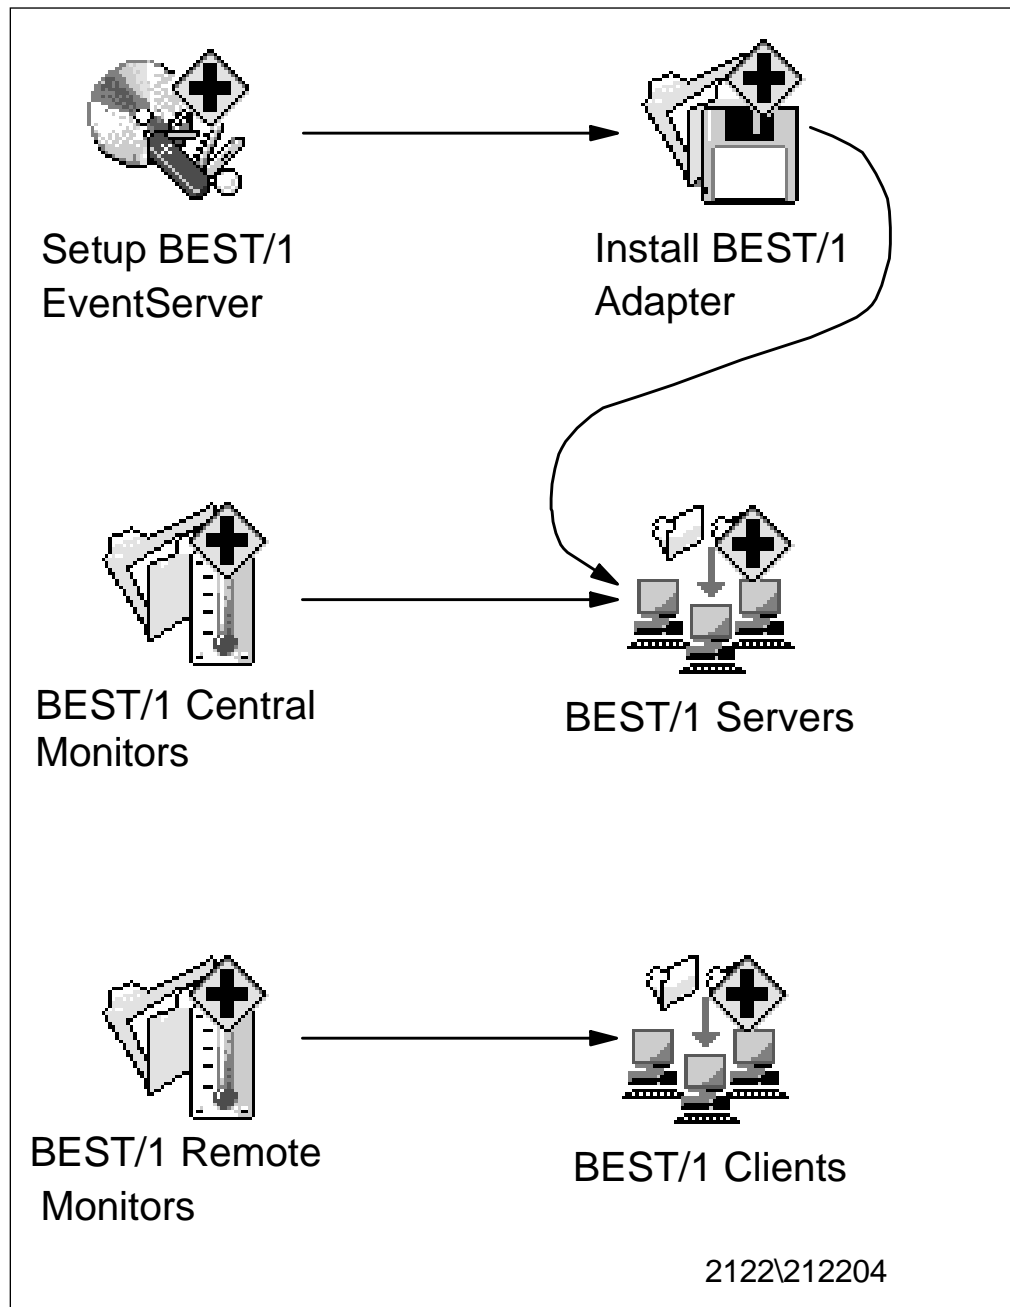


Figure 86. TME 10 Enterprise Console and TME 10 Distributed Monitoring Setup for BEST/1

These components distribute the TME 10 Distributed Monitoring profile to the BEST/1 managing and managed nodes and prepare the TME 10 Enterprise Console to receive BEST/1 events.

- Setup BEST/1 EventServer: Task to prepare the TME 10 Enterprise Console to receive events from the BEST/1 event adapter. Run this task before the Install BEST/1 Adapter task.

- Install BEST/1 Adapter: Task to distribute the BEST/1 event adapter to BEST/1 servers. The event adapter forwards BEST/1-Monitor alerts and errors occurring during collect file processing to the TME 10 Enterprise Console. Run the Setup BEST/1 EventServer task prior to this task.
- BEST/1 Central Monitors: TME 10 Distributed Monitoring profile to check the status of the BEST/1 servers. By default this profile subscribes to the BEST/1 Servers profile manager.
- BEST/1 Remote Monitors: Monitors to watch the BEST/1 managed nodes. By default they are subscribed to the BEST/1 Clients profile manager.
- BEST/1 Servers: Profile manager for BEST/1 managing nodes.
- BEST/1 Clients: Profile manager for BEST/1 managed nodes.

2.7.3.3 Components for Monitoring BEST/1

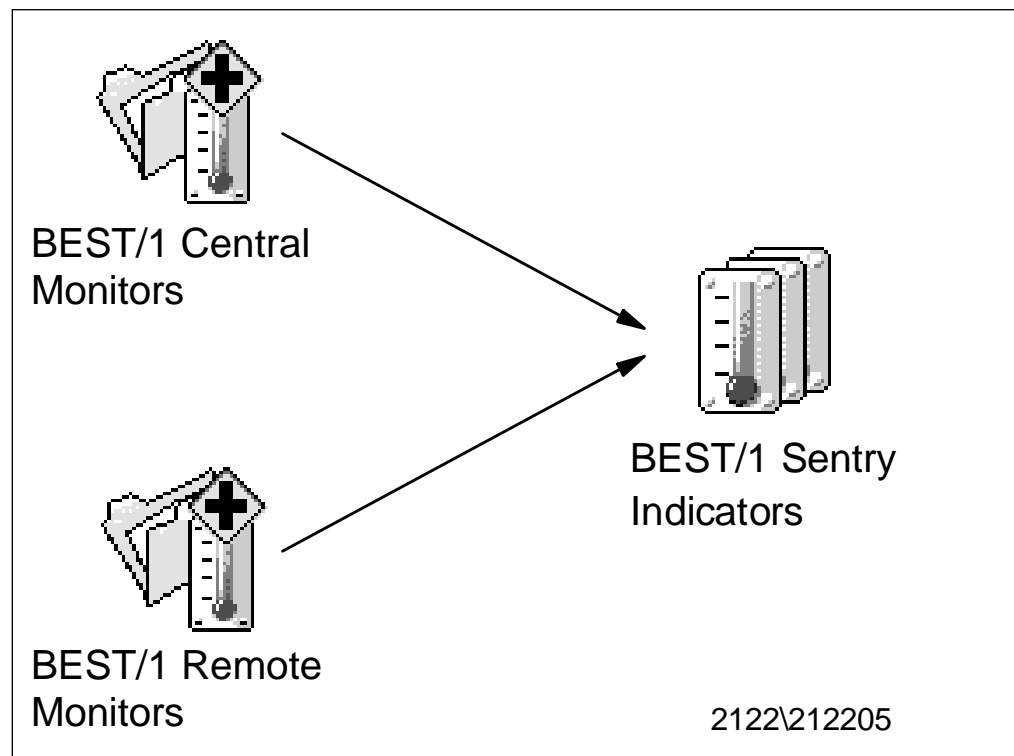


Figure 87. TME 10 Distributed Monitoring and BEST/1

The monitoring components allow you to check the status of the managing and managed nodes and the BEST/1 processes by means of Tivoli.

- BEST/1 Central Monitors: TME 10 Distributed Monitoring profile to check the status of the BEST/1 servers.
- BEST/1 Remote Monitors: TME 10 Distributed Monitoring profile to monitor the BEST/1 managed nodes. The profile includes a large number of monitors observing the BEST/1 processes, watching system resources such as shared memory, semaphores, ports and file systems, and monitoring the size of the collect file.
- BEST/1 Sentry Indicators: Collection of indicators for the central and the remote monitors. These indicators may be used instead of the TME 10 Enterprise

Console to view the status of the monitored resources. To access the indicator collection double-click on the icon.

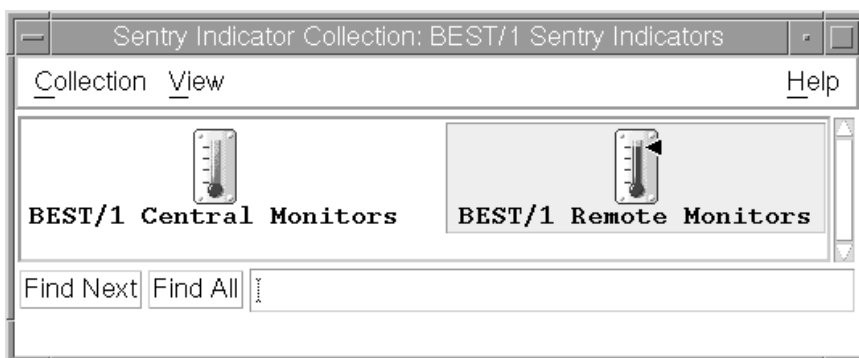


Figure 88. The Indicator Window

Figure 88 shows the Sentry Indicator window with the indicators for the central and the remote monitors. The thermometer rises as the monitored resource becomes more severe.

Four levels of severity exist:

- Normal
- Warning
- Severe
- Critical

Depending on the monitored resource the monitors define different actions to be performed, when a certain severity level is reached.

2.7.3.4 Components for Using BEST/1




Figure 89. TME 10 Components for Using BEST/1

Using BEST/1 The application components provide interfaces to access BEST/1 directly from the TME 10 Desktop.

- BEST/1: The icon to launch the BEST/1 application.
- BEST/1 Collect: Task to start the collection process on managed nodes.
- BEST/1 Move Collect Files: Task to move collected data from a remote managed node to the current machine.
- BEST/1 Remove Old Collect Files: Task to remove collect files older than a specified number of days.
- BEST/1 Stop Collect: Task to stop the collect daemon on a remote managed node.

In the following sections we explain the BEST/1+ for TME 10 components mentioned in concrete examples in our environment.

2.8 How Can BEST/1+ for TME 10 Help Us in Managing BEST/1?

So far, we have shown how to install the BEST/1 server and the BEST/1+ for TME 10 module in our environment.

In the section where we worked with the BEST/1 functions, we worked on the BEST/1 server and a BEST/1 client. However, we did not show, for example, how to install the BEST/1 agent on that client.

The first example, where BEST/1+ for TME 10 can help us in managing our environment by automatically installing BEST/1 on the client systems using TME 10, is shown in 2.9.1, "Deploying BEST/1 Collect Using TME 10 Software Distribution."

Although BEST/1 provides very sophisticated performance management and capacity planning functions, it is not a complete systems management solution. Therefore, it can benefit from being integrated into a systems management solution such as TME 10 in a number of ways.

Besides distributing the BEST/1 clients using TME 10 we also show how to implement synchronous monitors to control BEST/1 application availability with TME 10 Distributed Monitoring, process alerts from BEST/1 using TME 10 Enterprise Console and other features of BEST/1+ for TME 10 in the remainder of this chapter.

2.9 Examples Using BEST/1+ for TME 10 in Our Environment

In this part we show how we use the BEST/1+ for TME 10 module to manage BEST/1.

We show the following examples:

- Distribution of BEST/1 Collect to managed nodes
- Monitoring BEST/1 with TME 10 Distributed Monitoring
- Event Management using TME 10 Enterprise Console
- Installing BEST/1-Visualizer using AutoPack

2.9.1 Deploying BEST/1 Collect Using TME 10 Software Distribution

Usually, the first task in managing a BEST/1 with TME 10 is to distribute the BEST/1 agents to all systems that are to be managed by the BEST/1 application in regard to performance management.

This can be done by TME 10 Software Distribution in combination with BEST/1+ for TME 10.

In our environment, we deploy the BEST/1 agent to the machines rs600026, hp and sun.

As the distribution is done by means of Tivoli, the BEST/1 managed nodes must be attached as Tivoli managed nodes to your TMR. Be sure you have TME 10 Software Distribution installed on each of your destination machines.

The installation of BEST/1 Agent and BEST/1 Collect has to be done in three phases:

1. Extraction of BEST/1 architecture-specific Collect install script
2. Preparation of the profile managers
3. Preparation of the TME 10 Software Distribution file packages

See the following sections for a detailed description of each phase.

2.9.1.1 Extraction of BEST/1 Architecture-Specific Install Script

The BEST/1+ for TME 10 CD-ROM provides you with prepared installation scripts for several UNIX operating systems so that you can easily build up a heterogeneous performance assurance environment. The install files are located in different directories and are named as follows (we mounted the CD-ROM on /cdrom):

- /cdrom/AIX/install.collect-aix for IBM AIX
- /cdrom/HP/install.collect-hp for HP-UNIX
- /cdrom/NCR/install.collect-ncr
- /cdrom/OSF/install.collect-osf for DEC-UNIX
- /cdrom/Solaris/install.collect-solaris for SUN Solaris
- /cdrom/SunOS/install.collect-sunos for SUN OS

It is possible to install directly from the CD-ROM, but we recommend you copy all necessary files according to your environment to a directory on the TMR server so they are available for future use. As we had one managed node running HP-UX (hp) and one running Solaris (sun) we had to copy the files install.collect-hp and install.collect-solaris. Since the managed node rs600026 runs the same operating system (AIX) as our TMR server and the installation procedure of BEST/1 copies the corresponding file to the managing node, we did not need to copy the install.collect-aix file as we still have it on our system from installing the BEST/1 server. It is located in the directory /best1/install as this was our installation directory for BEST/1 (see Figure 75 on page 86).

We choose this directory as the destination for the copy commands:

```
# cp /cdrom/HP/install.collect-hp /best1/install
# cp /cdrom/Solaris/install.collect-solaris /best1/install
```

After we have finished this we are ready for the next step.

2.9.1.2 Preparation of the Profile Managers

The profile managers define the destinations for the distribution of the file packages. They are especially useful in large environments but are needed in any environment. If you have only a few machines to be integrated, you can omit this step and set up the list of subscribers directly after you have prepared the file package (see 2.9.1.3, "Preparation of the File Packages" on page 101).

In a homogenous environment you just need one profile manager. You can use the predefined one that comes with the BEST/1+ for TME 10 module. In heterogeneous system environments we recommend working with several profile managers. Define one profile manager for each UNIX flavor of your environment.

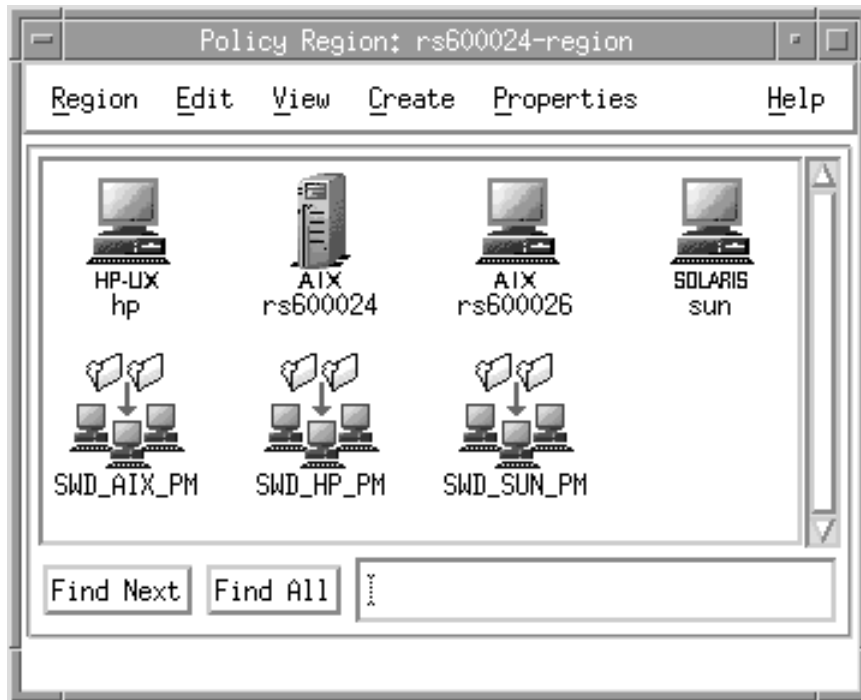


Figure 90. Profile Managers to Distribute BEST/1

Since we use three different UNIX flavors (AIX, HP-UX, Solaris) in our environment we use three profile managers. They all belong to our policy region rs600024_region as shown in Figure 90. The profile manager SWD_AIX_PM is used for AIX, SWD_HP_PM is used for HP-UX and SWD_SUN_PM is used for Solaris.

Alternatively, you can specify one policy region for each UNIX flavor latent in your environment. Each policy region must then own one separate profile manager. Since we had only a small number of machines it was sufficient to integrate all profile managers into one policy region.

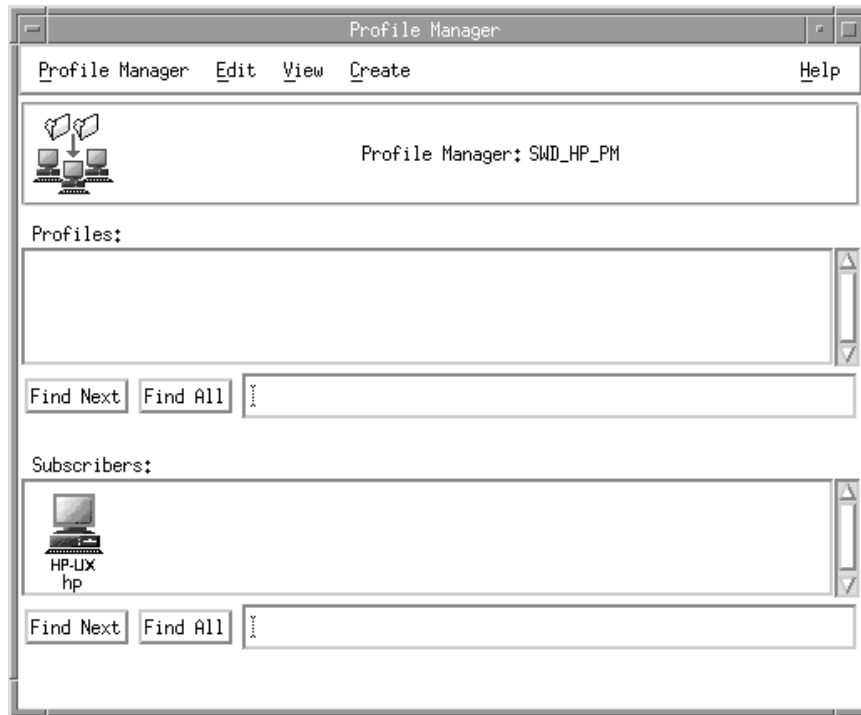


Figure 91. Profile Manager for HP-UX Machines

For each profile manager you have subscribe the list of machines that run the corresponding flavor of the UNIX operating system and that you want to become a BEST/1 managed node. As an example Figure 91 shows the profile manager for our machine running HP-UX. The profile managers for our other two machines are defined accordingly.

2.9.1.3 Preparation of the File Packages

In this step you generate the file packages that will be distributed to the BEST/1 managed nodes. You need to create one package for each operating system type in your environment. To prevent misunderstandings we define the following terms:

- *Target machine* to refer to the TMR server as this is actually the machine where the file package generation takes place
- *Remote machine* to refer to the managed node to which the software will be distributed




Figure 92. Icon for the Setup BEST/1 Collect Task

Before you execute the task to create the file package you can modify the task properties. Select **Modify job...** with the right mouse button from context menu of the task icon to open the Edit Job window as shown in Figure 93.

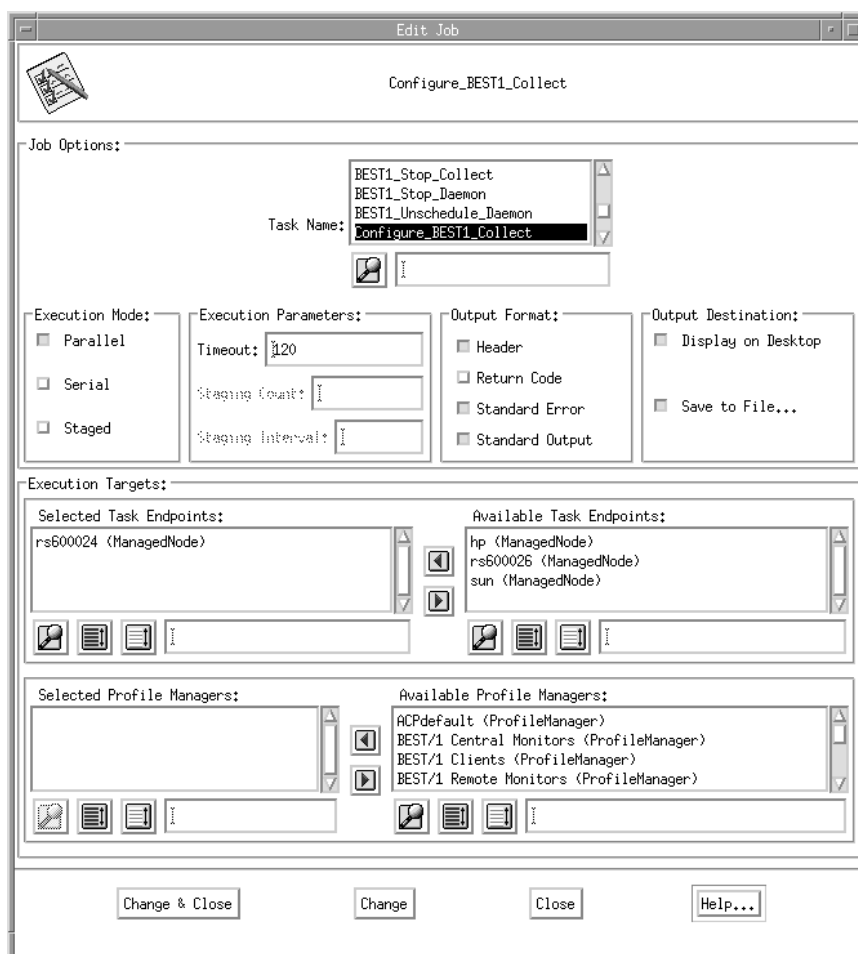


Figure 93. Properties of the Setup BEST/1 Collect Task

In the Edit Job window you can modify the following parameters:

- The Execution mode may be parallel, serial or staged. Staged means that you specify a stage interval and the tasks are executed in order according to this interval on the target machines. Since this task is always executed on one machine only, the selected value is of no concern.
- The Output Format may contain a header, the return code, the output on standard error and the output on standard output.
- The Output Destination may be the TME 10 Desktop or a file.
- The Execution Target has to be the TMR server.

The last parameter is very important. Make sure that you specify here the machine where the installation script resides. Usually this is the TMR server. Otherwise the installation scripts won't be found during task execution and the task cannot be accomplished properly.

The other parameters are not that important. You may select the values according to your preferences. For example, you can send the output to a file instead of displaying it on the desktop. This will be helpful for troubleshooting when problems arise.

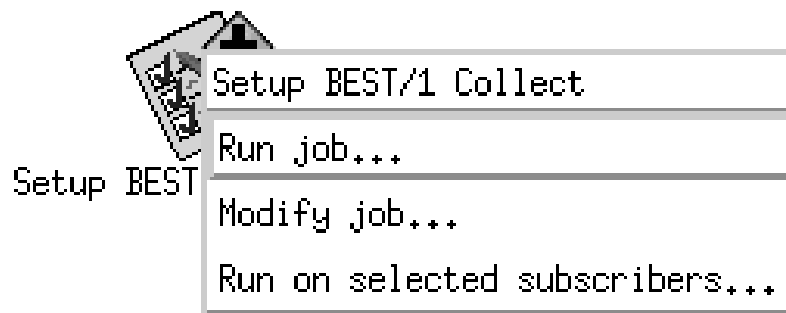


Figure 94. Context Menu of Setup BEST/1 Collect Icon

After checking the task parameters you are ready to execute the task. Therefore, select **Run job...** from the context menu of the **Setup BEST/1 Collect** icon as shown in Figure 94.

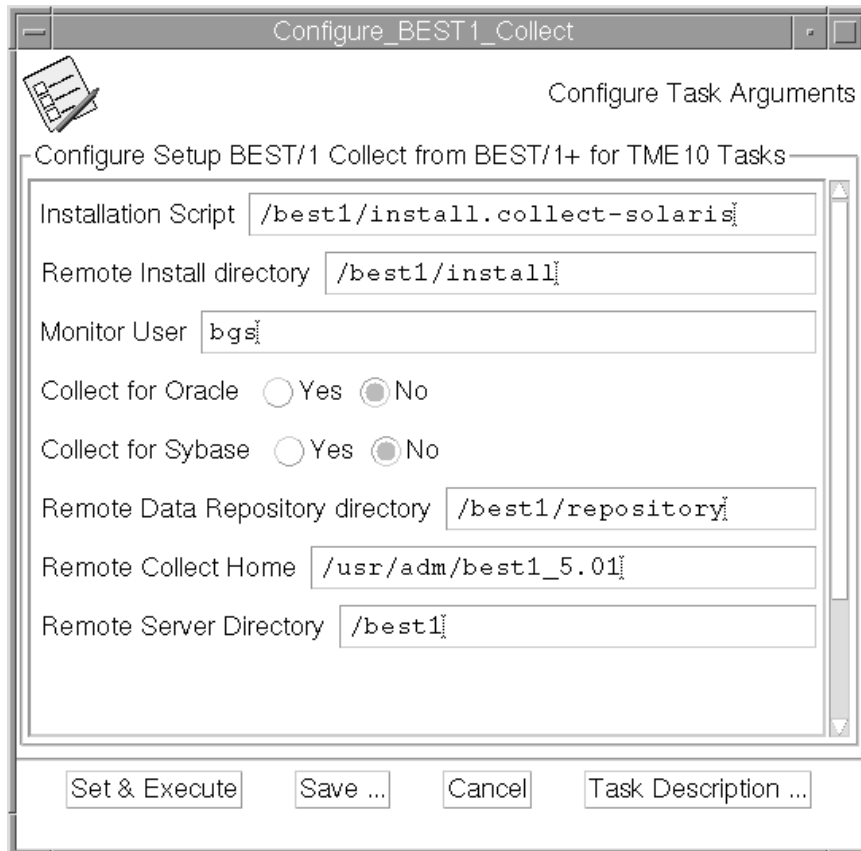


Figure 95. Parameters for the File Package

The following list describes the installation parameters to be specified in the `Configure_BEST/1_Collect` window as shown in Figure 95:

1. The Installation Script is the architecture-specific script that will be executed on the remote machines after distribution. Specify here the full path of the script on the target machine. For example, this is `/best/install/install.collect-hp` for the hp machine in our environment (see 2.9.1.1, "Extraction of BEST/1 Architecture-Specific Install Script" on page 99).
2. The Remote Install directory specifies where the BEST/1 Collect files will be installed during distribution. We used `/best1` as the installation target.
3. The Monitor User will be the owner of BEST/1-Monitor agent processes. The user is only needed when you want to use the BEST/1-Monitor on the remote machines. Be sure that the user exists on all remote machines before you initiate the distribution process. If you do not want the BEST/1-Monitor agent to be installed, enter No (without quotes).
4. Collect for Oracle must be set to Yes if you want to install BEST/1 Collect for Oracle. Since we had no Oracle server in our environment, we set this option to No.
5. Collect for Sybase must be set to Yes if you want to install BEST/1 Collect for Sybase. Since we had no Sybase server in our environment, we set this option to No.
6. The Remote Data Repository directory holds all results of the BEST/1 Data Collection on the remote machine

7. Remote Collect Home gives the path for the home directory of BEST/1 Collect. Actually, this will be a link to an equally named directory beneath the installation directory. You have to use the default value /usr/adm/best1_5.01, because the installation will not complete if you change this path. The 5.01 refers to the current version of BEST/1.
8. The Remote Server Directory locates the BEST/1 Server program. It is called best1server.

Finally, select the **Set & Execute** button in the Configure_BEST1_Collect window to start the file package generation. Assuming the output destination to be the desktop, this will create the following output:

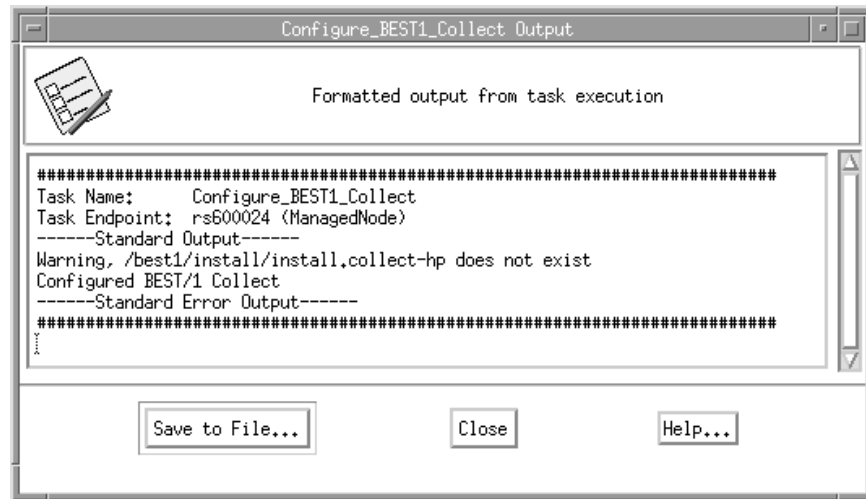


Figure 96. Failure in File Package Generation

Figure 96 shows an example of a failing file package generation. In this case the architecture-specific installation script was not found and the task execution has to be repeated.

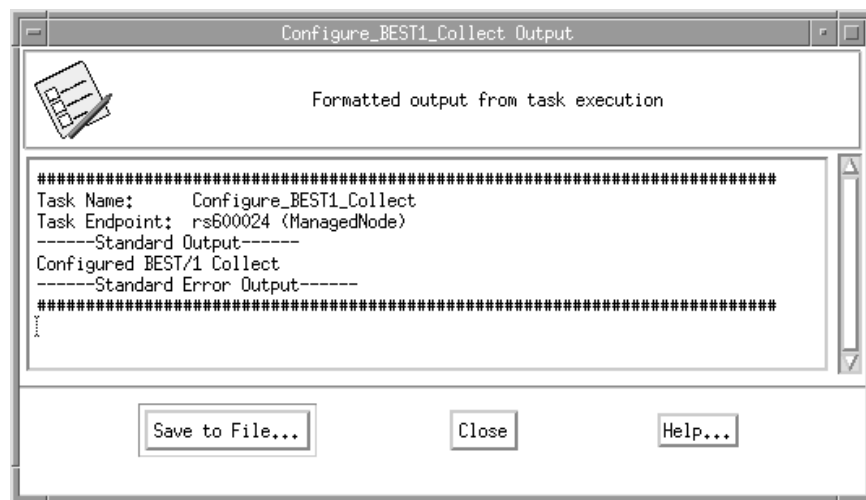


Figure 97. Success File Package Generation

When the file package generation terminates without errors the output will look as in Figure 97. At this point, the distribution of the file package to the managed nodes can be initiated.

Note

The task always assigns the same name to the file package and it overwrites the existing file package each time it is executed. This means that this step has to be repeated whenever a file package for a new machine type is needed. You have to rerun this task for one machine type if you created a file package for another machine type in the mean time.

2.9.1.4 Distribution of the File Package

After you have generated a TME 10 Software Distribution file package as described in the last section you are ready to distribute it. Before we are going to do this, we want to show what happens when you started the generation process by selecting the **Set & Execute** button in the Configure_BEST1_Collect window.



Figure 98. The File Package Icon

Selecting **Open...** from the context menu of the Configure BEST/1 Collect icon opens the Configure BEST/1 Collect window (Figure 99) containing one icon for the generated file package.

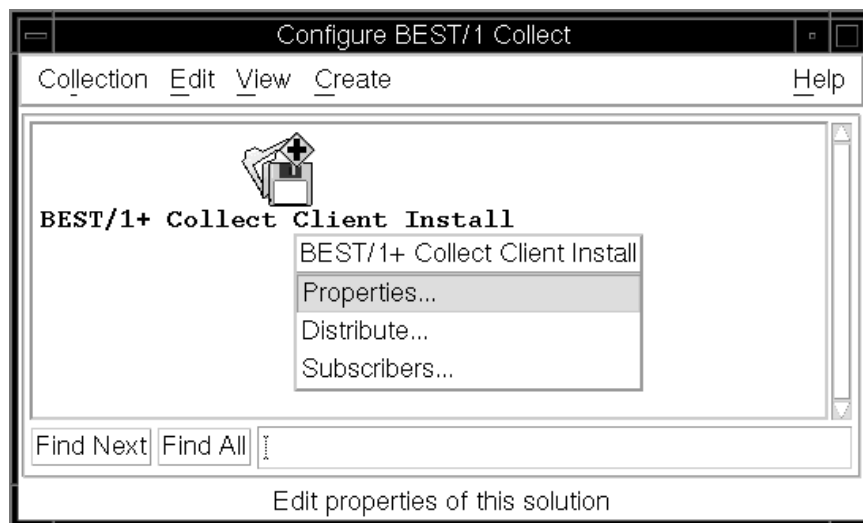


Figure 99. Open File Package Properties

To inspect the definition of the file package select **Properties...** from the context menu of the file package icon.

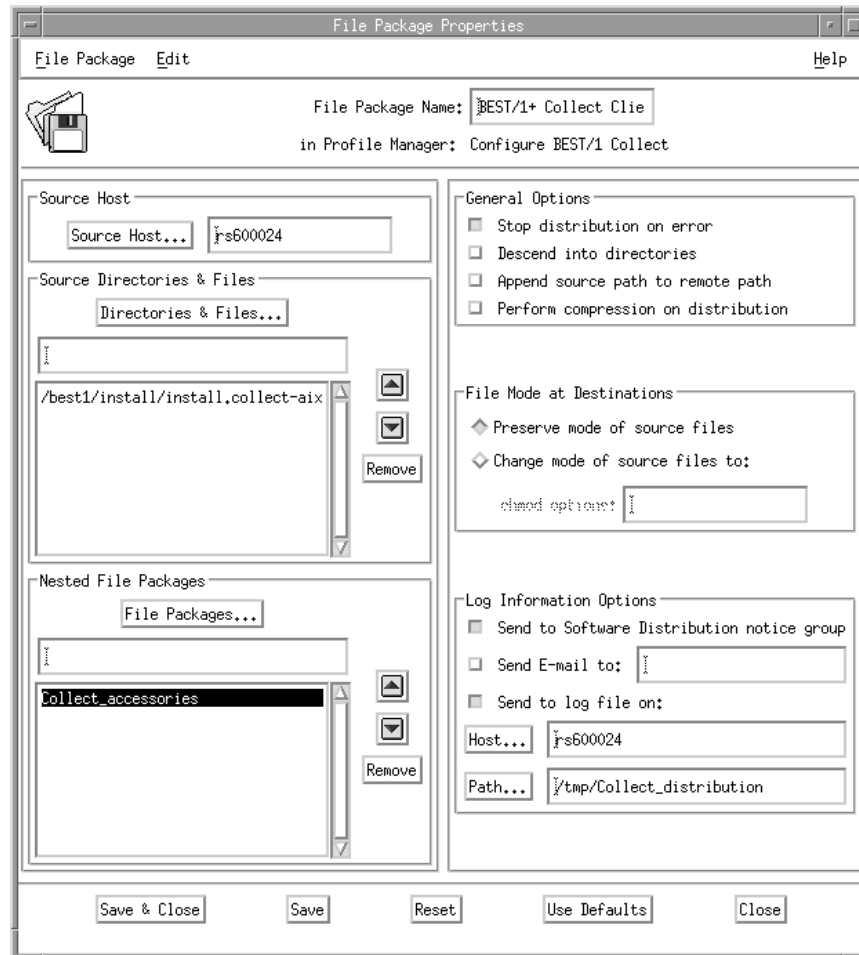


Figure 100. File Package Properties

The upcoming window (Figure 100) shows the source host of the file package and the name of the included source files. The source host should be the TMR server which is in our case the UNIX machine rs600024. The source file is the architecture-specific installation script (install.collect-hp).

Besides some general options and file node destination options that should not be changed, the window offers to modify the log information options. By default, a note is sent to the TME 10 Software Distribution notice group and a log file written on the source host. This information might be helpful and can be changed according to your needs.

As you can see the file package contains a so-called nested file package. This is nothing more than a normal file package except that no icon for this file package exists. In general, this offers the possibility to imbed file packages from other hosts or to make slight differences for certain destinations. The nested file package contains information common to all destinations. The imbedding file package holds destination-specific information.

The definition of a file package is stored in the Tivoli object database and can be exported with the following command sequence:

```
# wlookup -r FilePackage -a
# wexprtfp -c -f /tmp/fp.out @Collect_accessories
```

The first command lists all defined file packages and their resource IDs. The second exports the specified file package (Collect_accessories) to a file specified with the -f flag. If this flag is omitted, the output is printed to stdout.

```
##TFP-v2.02      Tivoli Filepack (version v2.02)
do_checksum=n
do_compress=n
modifiers=0
rm_empty_dirs=n
stop_on_error=y
descend_dirs=n
keep_paths=n
rm_extraneous=n
follow_links=y
create_dirs=y
skip_older_src=n
no_overwrite=n
backup_fmt=
list_path=
nested_first=n
src_re1path=
file_cksums=n
unix_platform_prefix=/tmp/Collect_access
log_host=rs600024
log_file=/tmp/Collect_distribution_access
append_log=n
....
prog_env="DATA_REPOSITORY='/best1/repository'
INSTALL_DIR='/best1 MONITOR_USER='no' ORACLE='no'
SYBASE='no' SERVER_DIR='/etc' INSTALL_COLLECT='install.collect-aix'
COLLECT_HOME='/usr/adm/best1_5.01'"
....
%
/usr/local/Tivoli/bin/generic_unix/TME/PLUS/best1/collect_install_after.sh
```

Figure 101. Excerpt of the File Package Definition @Collect_accessories

Figure 101 shows an extract of the file package definition of the nested file package. Apart from the general options at the beginning you see the program environment reflecting the parameter entered as shown in Figure 95 on page 104. At the end of the database export the list of included files is appended, in our case a shell script called collect_install_after.sh.

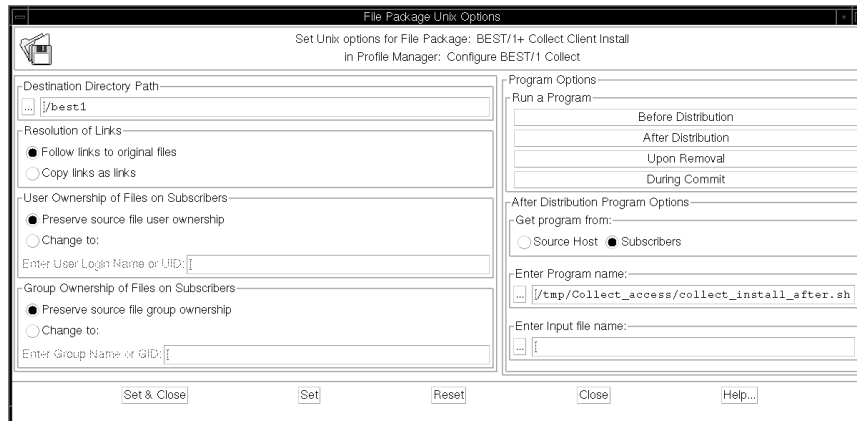


Figure 102. Unix Options Window

To open the Unix Options window (Figure 102) select **Platform-Specific Options -> Unix Options...** from the **Edit** menu of the File Package Properties window (Figure 100 on page 107). This window allows you to specify options for machines running UNIX operation system. The different phases of the distribution process, such as Before Distribution or After Distribution, are taken into account. For each a script or executable program can be specified.

Figure 102 shows the UNIX options for the BEST/1+ Collect Client Install file package. If you click on **After Distribution** within the Program Options section of this window, you will see that the collect_install_after.sh shell script contained in the nested file package is called as soon as everything is installed on the destination machine.



Figure 103. Open Subscribers for All File Packages

Now let us go back to the actual topic of this section, the distribution of the file package. This can be done in two ways. Choosing the first one you select the **Subscribers...** option from the Configure BEST/1 Collect icon's pop-up menu (Figure 103).

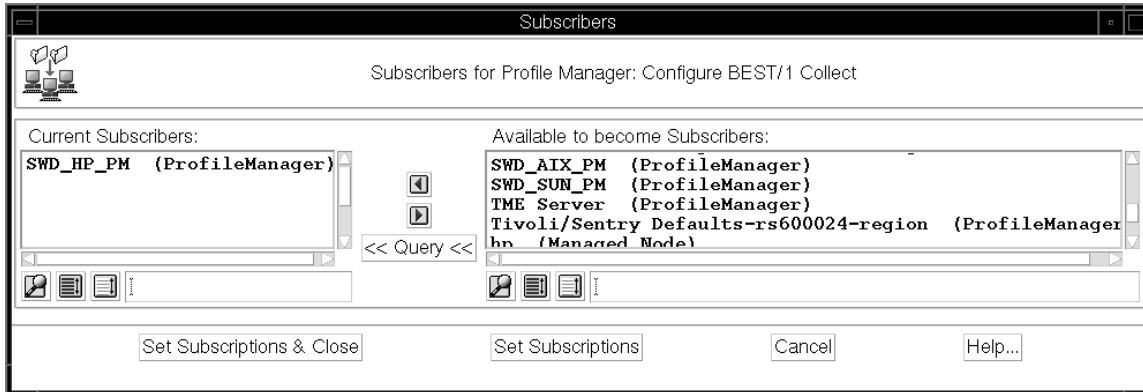


Figure 104. Subscribers Window

You will see the Subscribers window as shown above.

The next step is to specify the subscription list to which you want to distribute the file package. Remember that the machine types of the subscribers must correspond to the installation script included in your file package. We recommend to use the profile managers proposed in 2.9.1.2, "Preparation of the Profile Managers" on page 99.

After you have completed the subscription list, select the **Set Subscription & Close** button or the **Set Subscription** button to save your settings. Finally, select **Distribute** from the context menu of the BEST/1+ Collect Client Install icon (Figure 103 on page 109) to actually start the distribution.



Figure 105. Open the Subscribers Window for a Certain File Package

As mentioned above there are two ways to initiate the file package distribution. The second way differs in opening the Subscribers window by selecting **Subscribers...** from the pop-up menu of the icon representing the file package in the Configure BEST/1 window (Figure 105). You obtain a Subscribers window as shown in Figure 104. Here, you specify the subscribers list as described above and select the **Set Subscription & Close** button or the **Set Subscription** button to actually start the distribution. Finally, select **Distribute** from the context menu of the BEST/1+ Collect Client Install icon (Figure 105) to actually start the distribution.

The difference between both ways is that the first one distributes all file packages to their subscribers altogether. The second way allows you to specify the subscribers for a single file package and distribute them separately. This way, it is even possible to install the file package on a single machine to avoid the overhead of installing it on machines that already have it installed.

Assuming the standard log information options are set, the distribution process generates both, a note in the TME 10 Software Distribution notice group and a log file.

To view the note in the TME 10 Software Distribution notice group double-click on the **Notices** icon on the TME 10 Desktop.

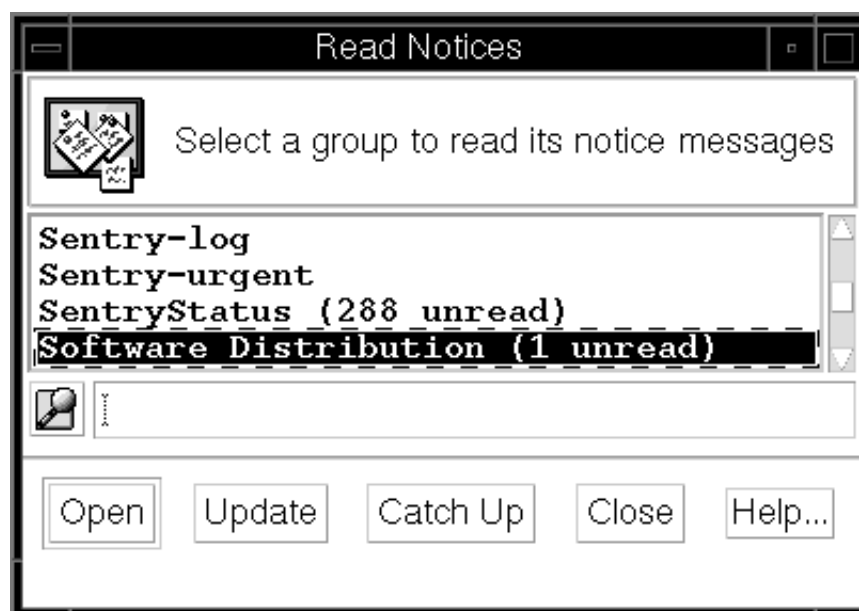


Figure 106. Select the TME 10 Software Distribution Notice Group

This will open the Notice Groups window (Figure 106). From this window select the **TME 10 Software Distribution** notice group by double-clicking on it.

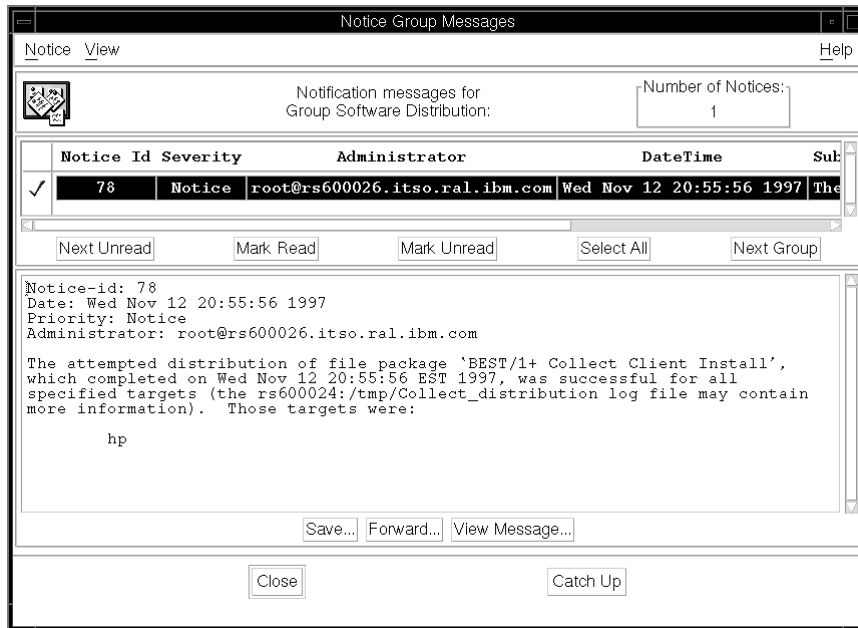


Figure 107. The TME 10 Software Distribution Notice Groups Window

The Notice Groups window in Figure 107 displays all unread messages. The contents of the message will be printed in the area beneath the message list as soon as you double-click on the message.

In our example the distribution to the hp machine was successful as the message indicates. All machines contained in the list of subscribers will be named. In the message, successes and failures will be listed separately.

```

File Package: "BEST/1+ Collect Client Install"
Operation:    install (m=5)
Finished:     Wed Nov 12 20:55:56 1997
-----
Source messages:
<none>
-----
hp:SUCCESS
starting script: /tmp/Collect_access/collect_install_after.sh
script stderr: mkdir: cannot create /best1/repository: File exists
best1 stream tcp nowait root /etc/best1server best1server
best1port      5070/tcp      best1
script complete: status=0
=====

```

Figure 108. Log Distribution File after Successful Distribution

The corresponding log file looks as depicted in Figure 108. In contrast to the TME 10 Software Distribution notice the log file contains an error message resulting from an existing data depository. Fortunately, this does not do any harm to the distribution process and the installation succeeds.

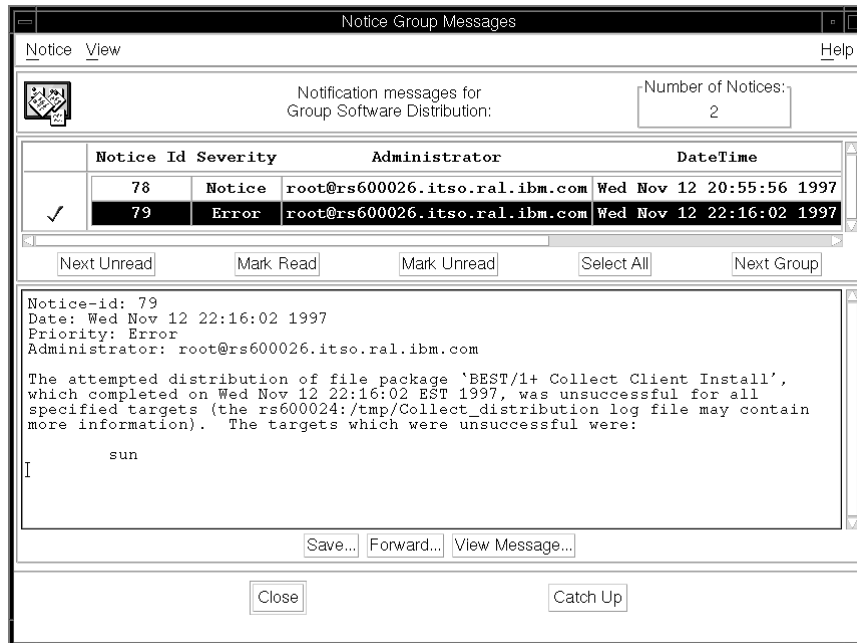


Figure 109. Notice Reporting a Failed Distribution

Figure 109 shows a TME 10 Software Distribution Notice received after a failed file package distribution. The problem in this situation is that we have chosen the wrong profile manager. We chose the profile manager for Sun machines and the system tried to install the file package prepared for HP machines resulting in the above error.

```
File Package: "BEST/1+ Collect Client Install"
operation:    install (m=5)
Finished:    Mon Nov 10 14:45:30 1997
-----
Source messages:
<none>
-----
hp:FAILED
An internal error occurred: IPC shutdown
Please refer to the TME 10 Framework Planning and Installation Guide,
"TME Maintenance and Troubleshooting" for details on diagnosing
internal errors or your Tivoli support provider.
=====
```

Figure 110. The Log Distribution File in Case of an Error

Figure 110 shows the log distribution file. We got the same message when there was not enough free space in the remote installation file system.

2.9.1.5 Manipulating the File Package

In the previous section we described the structure of the file package. We use this knowledge to simplify the distribution process a little bit. The idea is to extend the file package with all necessary installation scripts. Necessary means the file package must contain the installation scripts for all operating systems in your environment. The script `collect_install_after.sh` that is supplied with the nested file

package must then call the correct installation script regarding the actual operating system type.

This procedure has one disadvantage. All installation scripts will be copied to the remote machines. This results in longer distribution times and a higher need of free disk space since every installation script is about 6 MB in size.

Two steps have to be performed to realize this idea:

1. Extension of the file package
2. Manipulation of the script `collect_install_after.sh`

Let us begin with the description of the first step. In Figure 100 on page 107 you see the File Package Properties window right after the task Setup BEST/1 Collect has been executed as described in 2.9.1.3, "Preparation of the File Packages" on page 101. In the Source Directories & Files section you see the embedded installation script you have specified before.

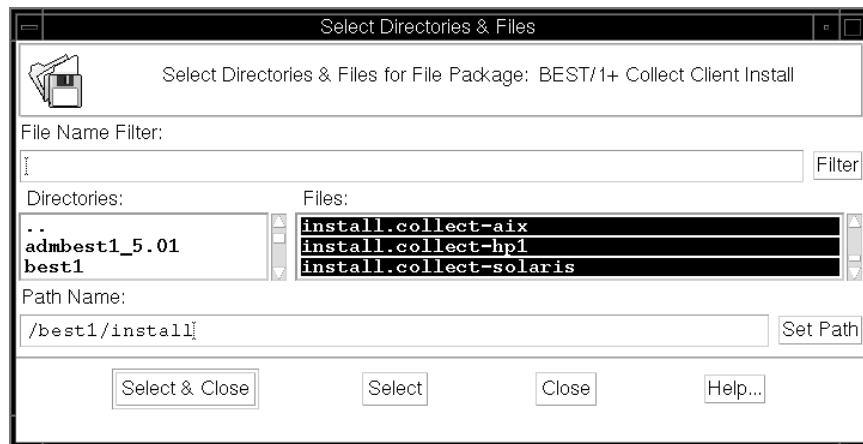


Figure 111. Select All Necessary Installation Scripts

Clicking on the **Directories & Files...** button displays the Select Directories & Files window on your desktop (Figure 111). Select all installation scripts you need in your environment and then click the **Select & Close** button.

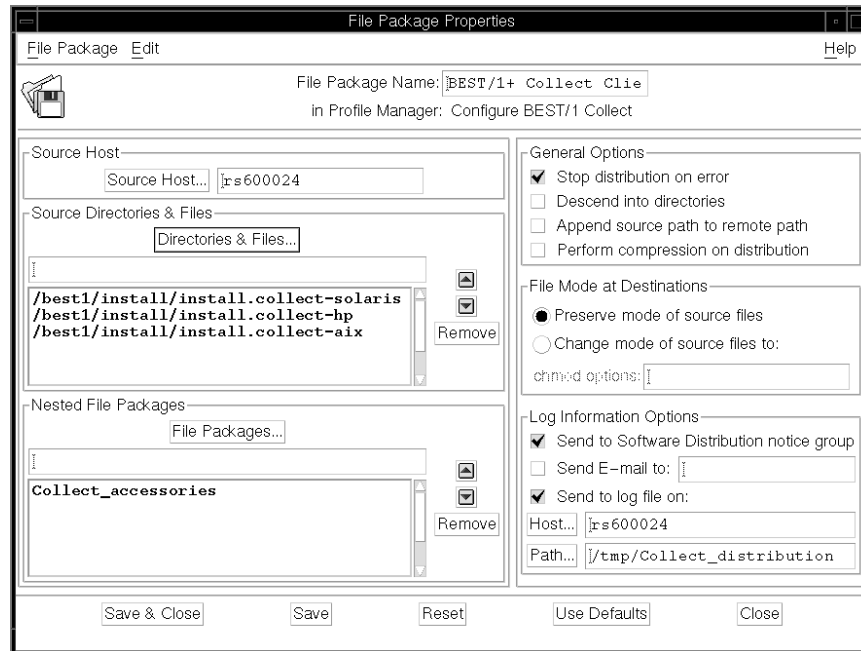


Figure 112. File Package Properties Window

The resulting File Package Properties window looks as displayed in Figure 112. For our environment we need the installation scripts for AIX, HP-UX and Solaris.

Finally, select **Save** from the File Package menu of the File Package Properties window to make your changes permanent. This completes the first step.

The second step describes the manipulation of the script called right after the distribution.

```

.....
26 if [ ! -d $DATA_REPOSITORY ]; then
27   mkdir $DATA_REPOSITORY
28 fi
29 osname=$ uname
30 echo "$osname" > /tmp/best1_inst.log
31 case $osname in
32   "AIX") osname="aix";;
33   "HP-UX") osname="hp";;
34   "SunOS")
35     version= `uname -a` | cut -f3 -d ' '
36     if [ $version -gt 3 ]; then
37       osname=solaris;
38     else
39       osname=sunos
40     fi;;
41   *)
42     echo "not a supported OS" >> /tmp/best1_inst.log
43     exit 1;;
44 esac
45 INSTALL_COLLECT="install.collect-$osname"
46 print "cd $INSTALL_DIR" >> /tmp/best1_inst.log
47 cd $INSTALL_DIR
48 print $INSTALL_COLLECT $DATA_REPOSITORY $MONITOR_USER $ORACLE $SYBAS
   $SERVER_DIR >> /tmp/best1_inst.log
49 ./$INSTALL_COLLECT $DATA_REPOSITORY $MONITOR_USER $ORACLE $SYBASE
   $SERVER_DIR << EOF
50 y
51 custdata@bgs.com
52 y
53 y
54 y
55 EOF
.....
"/usr/local/Tivoli/bin/generic_unix/TME/PLUS/best1/collect_install_after.

```

Figure 113. Excerpt of Modified collect_install_after.sh

In Figure 113 you see the changes that have to be made. Line 49 contains the call of the architecture-specific installation script. In the original version the name of this script used to be stored in the variable `INSTALL_COLLECT` which is set according to the settings in Figure 95 on page 104. This variable must be changed.

In lines 29 to 45 we derive the name of the install script from the operating system type by calling the `uname` command. The result of the `uname` command must be transformed to match the suffix of the corresponding installation script. For each operating system at least one operation is needed within the case statement. Finally, the name of installation script is the concatenation of the constant `install.collect-` and the operating system name.

Unfortunately, the `uname` command always returns "SunOS" on Sun machines though it should be Solaris for versions greater than 5.5.1. Therefore, we have to insert the lines 35 through 40. Here we use the `uname -a` command to determine the current version and change the variable `osname` to `solaris` if necessary.

By performing these two steps you don't have to run the Setup BEST/1 Collect again unless you want to change the configuration. During distribution it is no longer necessary to pay attention to different kinds of operating systems.

Note

We tested the modified script for the operating systems AIX, HP-UX and Solaris. If other operating systems are in use in your system environment, simply extend the above case statement.

2.9.1.6 Testing the BEST/1 Distribution

In this section we show how to check the file package distribution after it has been indicated successful. The easiest way to do this is to start and stop the collect daemon on the remote systems. With reference to Figure 79 on page 90 the BEST/1+ for TME 10 collection provides two tasks to do this:

- BEST/1 Collect
- Stop BEST/1 Collect

We will execute both tasks to see if everything works well.

Before you execute the tasks, you should make sure that the list of subscribers is set up correctly. This must be done for each task separately.



BEST/1 Collect




Figure 114. Select Modify ...

To access the subscribers list of the BEST/1 Collect, select **Modify Job...** from the context menu of the task's icon as depicted in Figure 114.

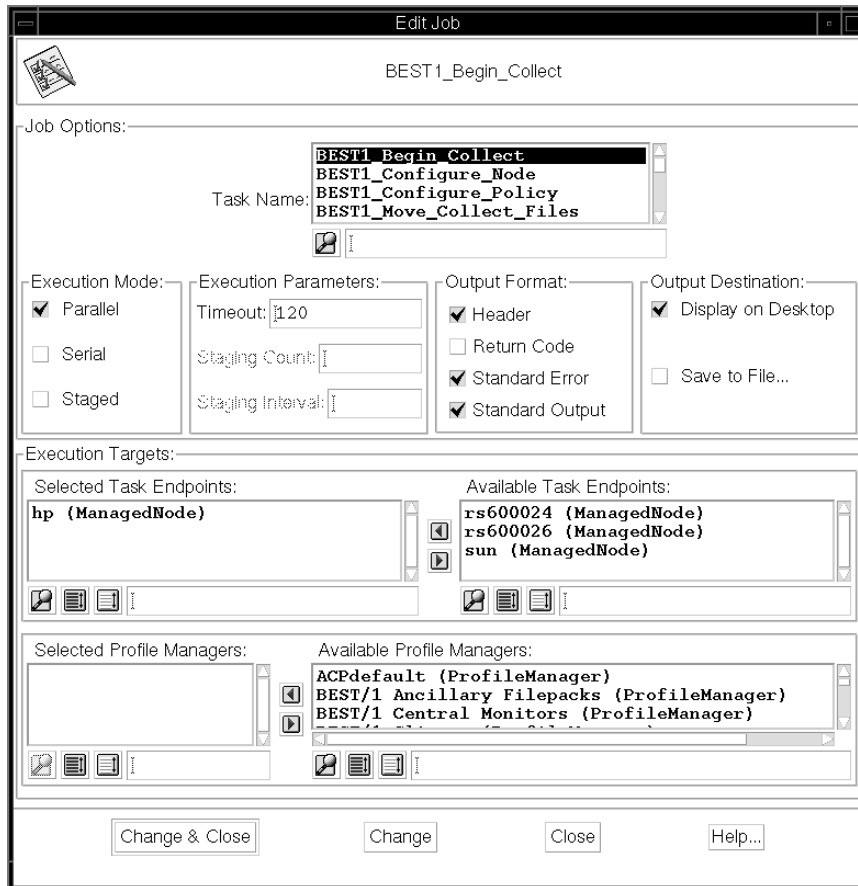


Figure 115. Edit Job Window for BEST/1 Collect Task

The Edit Job window for the BEST/1 Collect task will appear. The window is similar to the one displayed in Figure 93 on page 102. For our intention an output displayed on the TME 10 Desktop is sufficient. Hence, we change the Output Destination option and leave the other options unchanged. We choose the HP machine to be the test destination and specify it as Task Endpoint. Make sure that you have entered your target machine properly in Select Task Endpoints or activated the correct profile manager in the Select Profile Managers list. Then, select the **Select & Close** or the **Close** button.

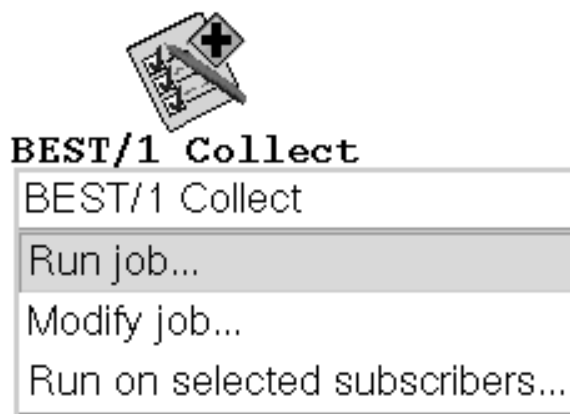


Figure 116. Select Run Job...

To initiate the task execution select **Run job...** from the pop-up menu of the BEST/1 Collect icon (Figure 116).

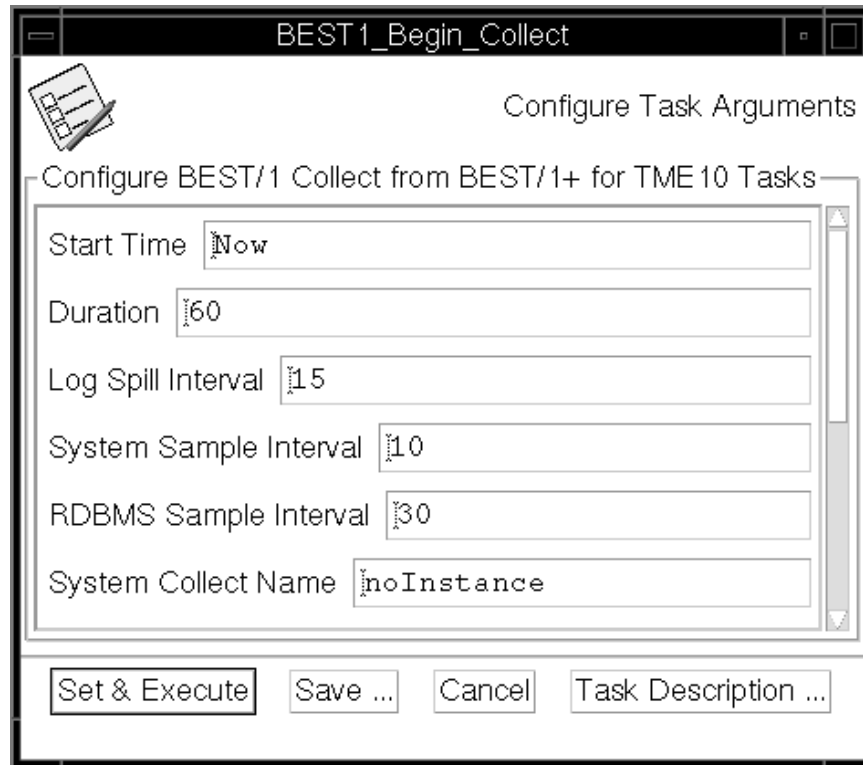


Figure 117. Task Arguments

Before the task is executed you have to specify some parameters in the BEST1_Begin_Collect window. The window comes up with default values that can be accepted unchanged for the test run. You can obtain a description of the parameters by clicking on the **Task Description...** button.

In order to make sure that the BEST/1 Collect process will generate some data during the test interval, you should set the start time to now and enter some small values for the Log Spill Interval (minutes) and the System Sample Interval (seconds). The duration should be long enough, so that the collect process can be stopped later with the Stop BEST/1 Collect task.

The Log Spill Interval represents the time period between two data samples written to disk. The System Sample Interval specifies the time period between two samples taken from the system kernel.

The last two parameters, RDBMS Sample reflecting the sample period for Database Systems and System Collect Name needed when multiple processes are running, are of no interest in our case.

Select the **Set & Execute** button to execute the task when all settings are correct.

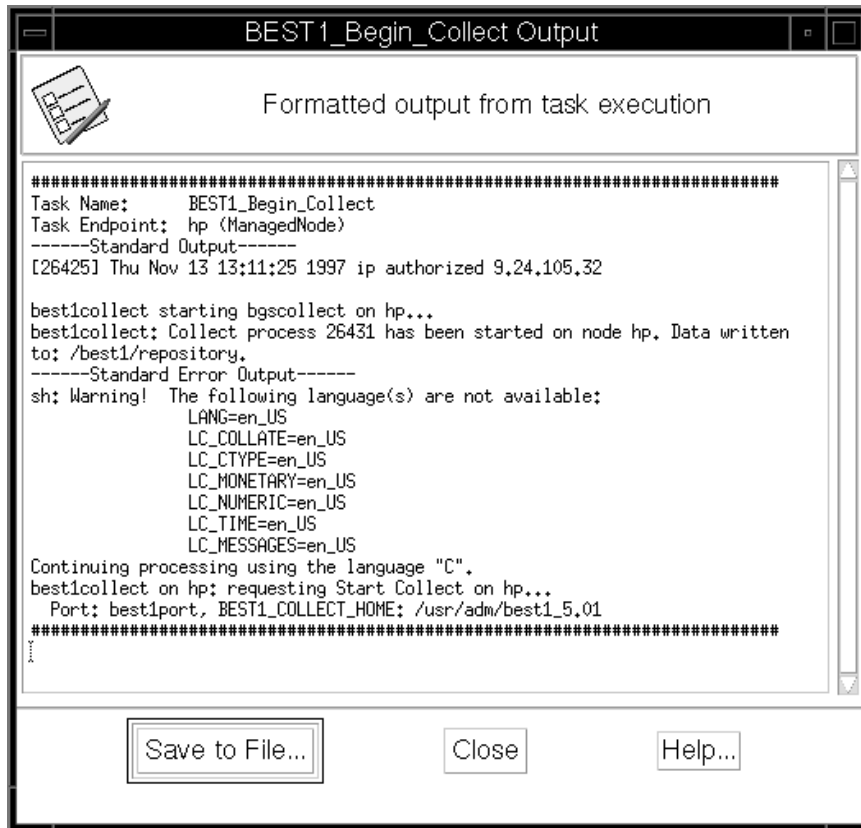


Figure 118. Output of Task BEST/1 Collect

The BEST/1_Begin_Collect_Output window reflects the result of the task execution. Figure 118 shows the output of a successful start of the BEST/1 Collect Process.

In general, the output gives the instance name and the process ID of the BEST/1 Collect process as well as the location of the data repository on the task endpoint.




Figure 119. Active BEST/1 Collect Processes

To verify the task output, you may use the ps -ef command together with the process ID obtained from Figure 118 right after the task has been executed. This lists the active BEST/1 Collect process as displayed in Figure 119. You can see the parameters specified earlier in the Task Arguments window.

The running BEST/1 Collect process generates a new directory in the data repository each time it is started. The name of the directory consists of the date and time the process was started.

```
# ls -ltr /best1/repository/Nov-13-97.13.11

-rwxrwxrwx  1 root    root      38670 Nov 13 13:19 collect.udf
-rwxrwxrwx  1 root    root       734 Nov 13 13:11 collect.udf.log
```

Figure 120. BEST/1 Collect Files

In the directory, the BEST/1 Collect process creates two files as shown in Figure 120. The first one contains the collect data and grows while the process is running. The second one is a log file written at collect start time.

```
# cat collect.udf.log
{ /* udf_header_block --- Thu Nov 13 13:11:26 1997
--- */
    global_special_part_type          1
    special_part_record_size          0
    standard_availability_map         fffffffe00 00000000
    special_availability_map          00000000 00000000
    collect_name                      .../bgs/bin/bgscollect
    collect_version                   5.0100
    collect_built_on                  May  6 1997
    operating_system_name             HP-UX
    operating_system_version          A.A.09.05
    hostname                          hp
    inet_address                      hp
    hostid                            152594720
    collect_uid                       0
    collect_euid                      0
    local_time_zone                   EST5EDT
    start_time                        Thu Nov 13 13:11:26 1997
    end_time                          Thu Nov 13 13:11:26 1997
    expected_end_time                 Thu Nov 13 14:11:26 1997
    spill_interval                    900
    sample_interval                   10
    termination_status                -1
    ppid                              26431
    interval_flag                     0
} /* udf_header_block */
```

Figure 121. BEST/1 Collect Log File

The log file in Figure 121 again gives information about the run-time environment of the process. You can see values such as the start time and the expected end time as well as the sample period.

```

Options set at command line:
-u 0
-g 0
-s 15
-d /best1/repository
-i 10
-e +60
-Z
-X
Values used:
Start time = Mon Nov 24 16:01:38 1997
BEST1_COLLECT_HOME = /usr/adm/best1_5.01
Data and log files directory = /best1/repository
End time = Mon Nov 24 17:01:38 1997
Process sampling interval = 10 seconds
Instance Name = noInstance
Number of data spills = 6
Spill interval = 900 seconds
-Z is set, not using shared memory
End of Values used information

```

Figure 122. BEST/1 Collect Log File bgscollect.log

Another log file called bgscollect is created in the directory /\$COLLECT_HOME/bgs/log where the variable COLLECT_HOME is set to /usr/adm/best1_5.01 in our environment. Figure 122 gives an example.

The next step is to see if the collect process can be stopped with the Stop BEST/1 Collect task.

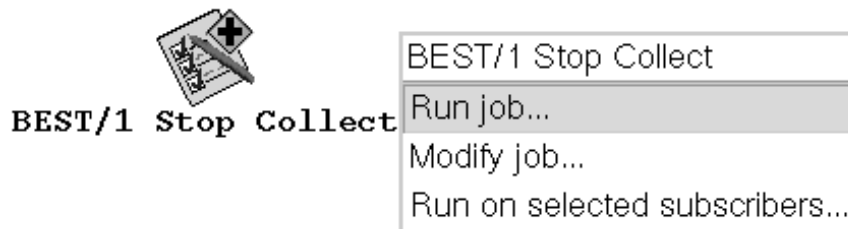


Figure 123. Stop BEST/1 Collect Icon

As shown in Figure 123 select **Run job...** from the context menu of the Stop BEST/1 Collect task icon.



Figure 124. Arguments for the Stop BEST/1 Collect Task

The system opens the BEST/1_Stop_Collect window (Figure 124). Enter the name of the instance you want to stop. Be sure to enter the same instance name as before when starting Collect.



Figure 125. Output from Stop BEST/1 Collect Task

The BEST/1_Stop_Collect Output window tells you the ID of the process that has been stopped. As before you can verify this using the ps command.

```

....
Terminated by signal 2
{ /* udf_header_block --- --- */
    global_special_part_type          1
    special_part_record_size          0
    standard_availability_map         fffffffe00 00000000
    special_availability_map          00000000 00000000
    collect_name                      ../bgs/bin/bgscollect
    collect_version                   5.0100
    collect_built_on                  May  6 1997
    operating_system_name             HP-UX
    operating_system_version          A.A.09.05
    hostname                          hp
    inet_address                      hp
    hostid                            152594720
    collect_uid                       0
    collect_euid                      0
    local_time_zone                   EST5EDT
    start_time                        Thu Nov 13 13:11:26 1997
    end_time                          Thu Nov 13 13:11:26 1997
    expected_end_time                 Thu Nov 13 14:11:26 1997
    spill_interval                    900
    sample_interval                   10
    termination_status                 -1
    ppid                             26431
    interval_flag                     0
} /* udf_header_block */

```

Figure 126. BEST/1 Collect Log File after Process Stopped

The system stops the Collect process with a signal 2, so it has time to extend the log file, as you can see in Figure 126.

As this point you can be sure that BEST/1+ for TME 10 is correctly installed on your managed node.

2.9.2 Monitoring BEST/1 Using TME 10 Distributed Monitoring

Monitoring is essential in a distributed systems environment. As soon as you have distributed the BEST/1 software to the managed nodes as described in the last section and started the BEST/1 Collect daemons you will want to know if all components integrate into the whole system.

You must be sure that the BEST/1 server is alive, that the processes are running on the remote sites and that the communication between server and client works fine. And finally, you will want to control the whole environment from a central point.

This section shows how all these requirements can be covered with the BEST/1+ for TME 10 module. We show how distributed monitoring of the BEST/1 environment has to be set up and how the module contents fit into the TME 10 Distributed Monitoring application.

As described in 2.7.3.3, “Components for Monitoring BEST/1” on page 95 the BEST/1+ for TME 10 module provides components to control the critical resources

of BEST/1. The first question to answer here is what are the critical resources of BEST/1.

The following list gives an overview:

- Hosts must be checked if they can be reached via the network.
- Collect processes and BEST/1 daemons must be up and running and they should not use too much CPU time.
- File systems used to store collected data must have sufficient capacity.
- Collect files must be growing to be sure that data is actually collected. On the other hand, they should not grow too fast.
- Semaphores and semaphore IDs are restricted.
- Shared memory is a restricted system resource.

For all these resources monitors are available in the BEST/1+ for TME 10 module. Two TME 10 Distributed Monitoring profiles are provided (see 2.7.3, “Overview of BEST/1+ for TME 10 Functionality” on page 90). The BEST/1 Central Monitor profile contains just one monitor to check the availability of the hosts. The BEST/1 Remote Monitor profile consists of 18 monitors covering all remaining resources.

As a representative example we show in this section how to activate the remote monitors and how to generate test events for the monitors watching the BEST/1 Collect process and the data repository.

We do not give a complete description of the monitors and their default actions for the certain levels of severity in this section, since this can be obtained from *BEST/1+ for TME 10 User's Guide*, 97095-00 in full detail.



Figure 127. Select Properties

Before the distribution of the TME 10 Distributed Monitoring profile, you can adjust the monitors to your needs in three ways:

1. Define the actions to be performed in response to severity level changes

2. Change the activity schedule
3. Enable or disable a monitor

Select **Properties...** from the context menu of the BEST/1 Remote Monitors icon (Figure 127 on page 125).

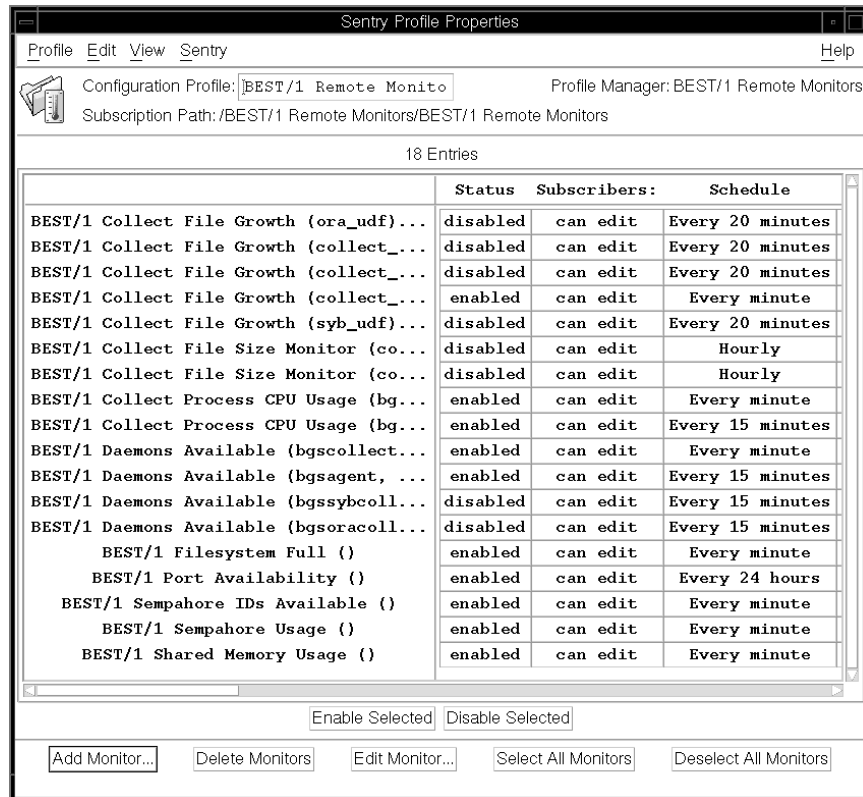


Figure 128. BEST/1+ for TME 10 Monitors

You obtain the Sentry Profile Properties window listing all monitors available in the profile. For each monitor its current properties are displayed. Use the lower scroll bar to see the remaining part of the window.

Because of the incomplete file type in the monitor description shown in Figure 128, we repeat the list of monitors again at this point:

1. BEST/1 Collect File Growth (ora_udf)
2. BEST/1 Collect File Growth (collect_syb)
3. BEST/1 Collect File Growth (collect_ora)
4. BEST/1 Collect File Growth (collect_udf)
5. BEST/1 Collect File Growth (syb_udf)
6. BEST/1 Collect File File Size Monitor (collect_syb)
7. BEST/1 Collect File File Size Monitor (collect_ora)
8. BEST/1 Collect Process CPU Usage (bgscoll)
9. BEST/1 Collect Process CPU Usage (bgsagent)
10. BEST/1 Daemons Available (bgscoll)

11. BEST/1 Daemons Available (bgsagent)
12. BEST/1 Daemons Available (bgssycollect)
13. BEST/1 Daemons Available (bgsoracollect)
14. BEST/1 File System Full
15. BEST/1 Port Availability
16. BEST/1 Semaphore IDs Availability
17. BEST/1 Semaphore Usage
18. BEST/1 Share Memory Usage

There are five monitors watching the file growth of the collect files, reflecting the different types of data. The UDF files contain data collected during system and database (Oracle or Sybase) monitoring. The collect_ora and collect_syb files are also used during database monitoring. Their content is merged into the udf files. Additionally, two monitors exist that take care of the growth rate measured in bytes per second of these files.

The profile also contains monitors predefined to watch the CPU usage of the different BEST/1 processes.

The last five monitors are responsible for the system resources such as file system, port, shared memory and semaphores.

Of course, it makes no sense to activate all monitors at the same time. If no collect daemons for database monitoring are running, there is no need to start the monitors controlling the size and growth rate of the database collect files. Even if you want to collect data from your database you only have to have the monitor running during collect time. That is the reason why only some of the monitors are enabled by default (see Figure 128 on page 126). To enable or disable a monitor select the monitor with a single mouse click and select the **Enable Selected** or the **Disable Selected** button respectively.

The monitors 14 - 18 should always be active. Depending on which application you are investigating you should activate the corresponding monitors. That means if you are investigating Oracle database servers the following monitors are needed:

- BEST/1 Collect File Growth (ora_udf)
- BEST/1 Collect File Growth (collect_ora)
- BEST/1 Collect File File Size Monitor (collect_ora)
- BEST/1 Daemons Available (bgsoracollect)

The monitors concerning a Sybase database server are the following:

- BEST/1 Collect File Growth (collect_syb)
- BEST/1 Collect File Growth (syb_udf)
- BEST/1 Collect File File Size Monitor (collect_syb)
- BEST/1 Daemons Available (bgssycollect)

If you want to change the attributes of a monitor, just double-click on its entry in the Sentry Profile Properties window and wait for the Edit Sentry Monitor window to appear on the desktop. The following two screen shots show the Edit Sentry window for monitors we want to use in our examples.

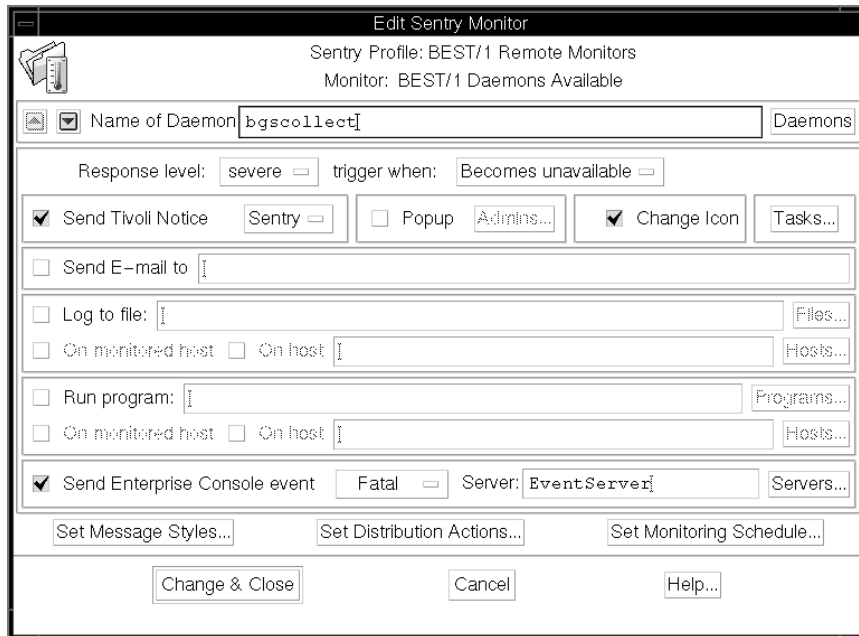


Figure 129. Properties of BEST/1 Daemons Available Monitor

Figure 129 shows the properties of BEST/1 Daemons Available monitor. The name of the daemon monitored is bgscollect. The figure shows the settings for the severity level "severe". This level is reached when the daemon becomes unavailable. Below the reactions for this response level are listed. You can choose between a Tivoli notice, a pop-up on an administrator's desktop, the icon to be changed, an e-mail, a log file, a program to be executed or a TEC event to be sent with a certain severity level.

By default, the BEST/1+ for TME 10 monitors send TEC events when triggered. This presumes that the TME 10 Enterprise Console is available in your environment. If not, disable TEC events in all monitors.

For our needs, pop-ups on our desktop work well and we activate them in both monitors. This way, we immediately get informed when an event occurs.



Figure 130. Properties of BEST/1 File System Full Monitor

The BEST/1 File System Full monitor is triggered as severe when 95% of the data repository is used. It is reset when the use of the file system decreases below 90% again.

For every monitor a schedule is defined as you can see in Figure 128 on page 126. For our test runs we want to change the monitoring period.

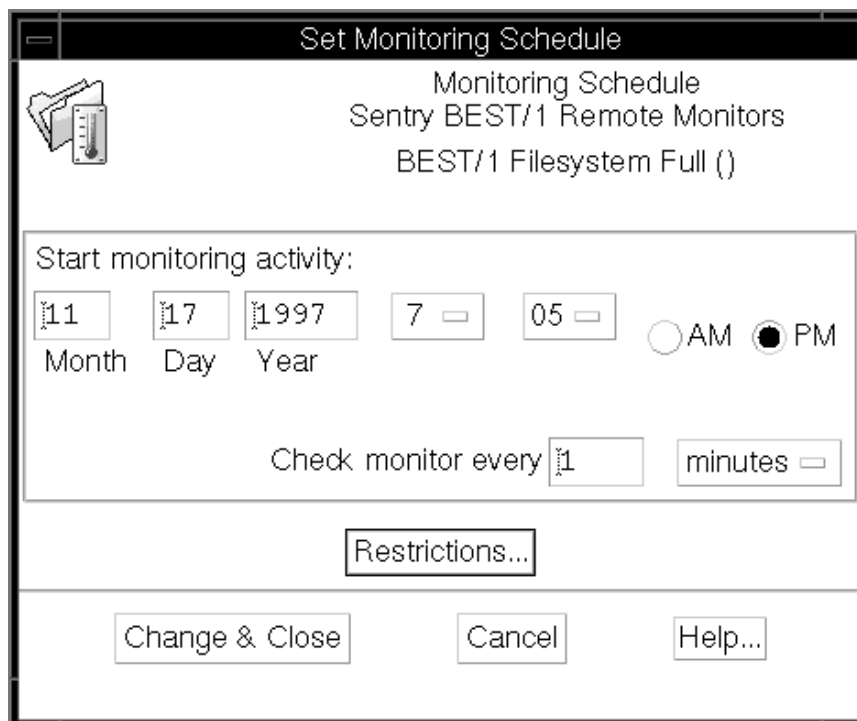


Figure 131. Schedule Definition

Click the **Set Monitoring Schedule...** button to receive the Set Monitoring Schedule window (Figure 131). You can select either a certain date or time period. We change the time period to 1 minute and leave the start date set to the current date as we want to be aware of system changes as soon as possible.

Select the **Change & Close** button when finished.

After you have made any changes to a monitor, you must save the profile by selecting **Save** from the Profile menu of the Sentry Profile Properties window (Figure 128 on page 126).

At this point you have all monitors adjusted as desired and you are ready to activate the monitors on the remote machines. The following steps illustrate the activation.



Figure 132. Select Subscribers ...

Select the **Select Subscribers...** option from the pop-up menu of the BEST/1 Remote Monitors icon.

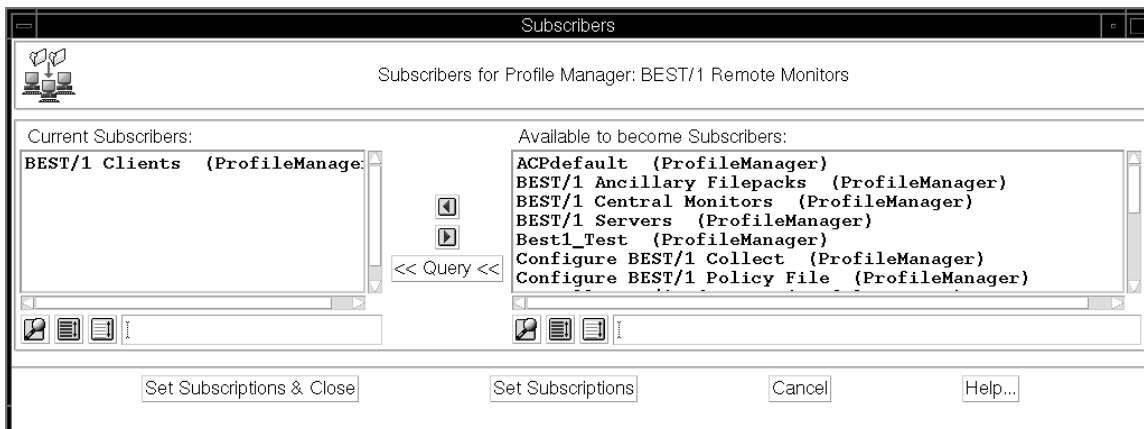


Figure 133. Subscribers of the Monitor Profile

In the Subscribers window (Figure 133) you have to select the machines you want to monitor remotely. We use the BEST/1 Clients profile manager to which all our client machines subscribe. This makes sense as the monitors are not dependent on machine types. You may use separate profile managers for groups of machines running different kinds of software.

Select **Set Subscriptions and Close** when all settings are made.



Figure 134. Select Distribute ...

To start the distribution select the **Distribute...** option from the pop-up menu of the BEST/1 Remote Monitors icon as illustrated in Figure 134.



Figure 135. Distribute Profiles Window

The Distribute Profiles window allows you to schedule the distribution. To immediately distribute the TME 10 Distributed Monitoring profiles click on the **Distribute Now** button. The distribution may take a little while depending on the number of subscribers.

The monitors should be active immediately after the termination of profile distributions. The heart of TME 10 Distributed Monitoring is a process called `sentry_engine`. It implements the entire TME 10 Distributed Monitoring functionality, such as:

- Activating the timers that are responsible for probe execution
- Executing monitoring probes implementing the actual monitors
- Triggering responses to certain events

Without the `sentry_engine` monitoring does not take place. The administrator has to check if the process is up and running and has to restart it if it is not running. To make sure that the `sentry_engine` is active on our managed nodes we issue the following TME 10 CLI command:

```
# odstat -d
daemons:
  pid  flags  count  uid  gid  oid  name
29818  MO      0      root  gid  1576275774.3.32#  /.../sentry_engine
....
```

Figure 136. Output from odstat Command

The output of the `odstat` command as shown in Figure 136 returns the process ID and the object ID when the `sentry_engine` is active. A similar result may be seen when using the `ps` command.

If the `sentry_engine` process is not running on one of your managed nodes you can simply restart the process issuing the TME 10 CLI command `wlseng`. This will start the `sentry_engine` process and additionally list all monitors currently configured on the managed node.

The `wlseng` command can also be used to check the TME 10 Distributed Monitoring configuration on the managed node.

The following figure shows the output of `wlseng` in our example.


```

Timer:
=====
<<BEST1_Monitors><bgsfreespace>
<BEST/1 Remote Monitors#BEST/1 Remote Monitors_1576275774.1.717
#Sentry::All#_29e1576275774.3.7>>
1:m:Y8,17:Y18,7:Y1,5:Y6,0:N0,23:N0,6:Custom hours:Custom days:879886500
Wakeup: 11/20/97 11:15:00 (56 seconds from now)

<<BEST1_Monitors><bgsDaemonStatus>
<BEST/1 Remote Monitors#BEST/1 Remote Monitors_1576275774.1.717
#Sentry::All#_20e1576275774.3.7>>
3:m:Y8,17:Y18,7:Y1,5:Y6,0:N0,23:N0,6:Custom hours:Custom days:879458700
Wakeup: 11/20/97 11:17:00 (2 minutes from now)

....
Monitoring probes:
=====
<<BEST1_Monitors><bgsfreespace>
<BEST/1 Remote Monitors#BEST/1 Remote Monitors_1576275774.1.717
#Sentry::All#_29e1576275774.3.7>>
("@_BEST/1 Remote Monitors","@G_BEST/1 Remote Monitors",
"BEST/1 Remote Monitors1576275774.3.7")
bgsfreespace() ->> 95 : "critical", ->> 90 : "warning",
-<< 90 : "reset", : "normal";

Last value: 9

<<BEST1_Monitors><bgsDaemonStatus>
<BEST/1 Remote Monitors#BEST/1 Remote Monitors_1576275774.1.717
#Sentry::All#_20e1576275774.3.7>>
("@_BEST/1 Remote Monitors","@G_BEST/1 Remote Monitors",
"BEST/1 Remote Monitors1576275774.3.7")
bgsDaemonStatus("bgscollect","") -> "down" : "severe",
-> "up" : "reset", : "normal";

Last value: down

Responses:
=====
<<BEST1_Monitors><bgsDaemonStatus>
<BEST/1 Remote Monitors#BEST/1 Remote Monitors_1576275774.1.717
#Sentry::All#_20e1576275774.3.7>>
severe
icon("@K_SentryBEST1"),
popup("0", " " 54x1576275774.1.179#TMF_Administrator:
:Configuration_GUI#Root_rs600024-region");

<<BEST1_Monitors><bgsfreespace>
<BEST/1 Remote Monitors#BEST/1 Remote Monitors_1576275774.1.717
#Sentry::All#_29e1576275774.3.7>>
warning
icon("@K_BEST/1 Remote Monitors"),
popup("0", " " 54x1576275774.1.179#TMF_Administrator:
:Configuration_GUI#Root_rs600024-region");

<<BEST1_Monitors><bgsfreespace>
<BEST/1 Remote Monitors#BEST/1 Remote Monitors_1576275774.1.717
#Sentry::All#_29e1576275774.3.7>>
critical
icon("@K_BEST/1 Remote Monitors"),
popup("0", " " 54x1576275774.1.179#TMF_Administrator:
:Configuration_GUI#Root_rs600024-region");

```

Figure 137. Output from the wlseng Command

The `wlseng` command consists of three parts reflecting the three sections of a monitor mentioned above. The Timer part contains as you may have expected timing details. There is one entry for each monitor of the distributed profile. We shortened this part so that only the activated monitors are left. For each entry the time of the next execution is given.

The Monitoring Probes part of the output has one entry for each activated monitor of the collection. In our case these are the monitors `bgsfreespace` and `bgsDaemonStatus`. This part details for every monitor the configured trigger values and actions for the response levels (normal, warning, severe, critical). For every probe the last value is given.

The last part is the Responses part, which lists all response levels triggered so far. The list is used to contain events of monitors that are not currently active. So, again we shortened this part in Figure 137 on page 133.

The monitor ID that is contained in each line of the `wlseng` output is grouped as follows:

```
<<collection><monitor><distribution>
```

For example, `BEST1_Monitors` is the name of the collection name and `bgsDaemonStatus` is the name of our monitor.

The `<distribution>` part of the monitor ID has the following structure:

- The name of the TME 10 Distributed Monitoring profile label, `BEST/1 Remote Monitors` in our case
- A separator `#`
- The profile label and the object ID of the profile
- Another separator `#Sentry::All#`
- A unique distribution number, for example, `__29e`
- The object ID of the endpoint

This information can be used when dealing with other TME 10 CLI commands such as `wlstime`, `wsprb` or `wlsra` which list the timer, probe or response information of a specified monitor, respectively.

Before the activated monitor can trigger events we have to start the `BEST/1 Collect` daemon. This process is watched by the `bgsDaemonStatus` monitor. It is the process of writing data to the repository which is controlled by the `bgsfreespace` monitor. Refer to 2.9.1.6, "Testing the BEST/1 Distribution" on page 117 for how to start this daemon on your managed node.

For testing reasons we adjusted the Log Spill Interval and the System Sample Interval to one minute and started several `Collect` daemons in parallel. This way we made sure the data repository would be filled up within a short period of time.

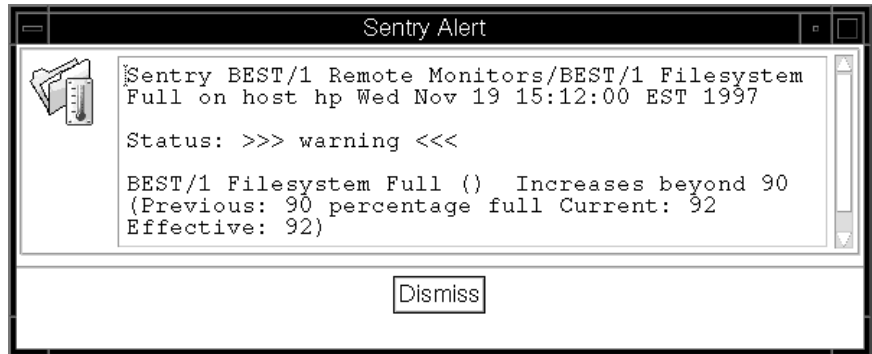


Figure 138. Response Level Warning

The bgsfreespace monitor checks the file system containing the BEST/1 data repository. It triggers the response level warning if the fill degree increases beyond a level of 90%. Figure 138 shows the indicator icon of the BEST/1 Sentry Indicators and corresponding pop-up window appearing on the desktop.

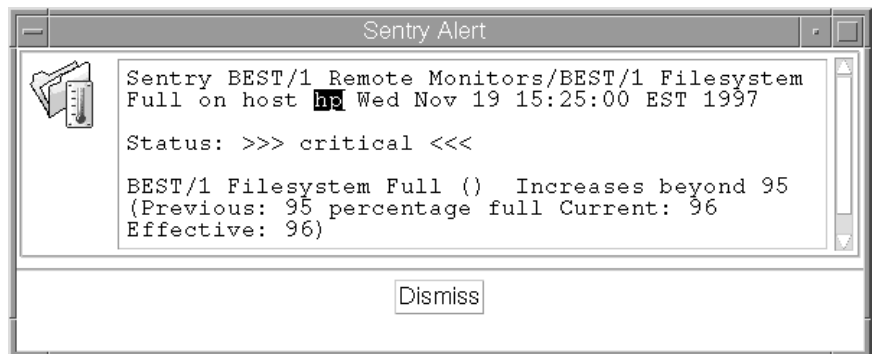


Figure 139. Response Level Critical

A bit later when the BEST/1 Collect daemon has generated some more data the fill level of the data repository will increase beyond the second threshold of 95%. This triggers the response level Critical. Figure 138 shows the indicator icon and the pop-up window.

2.9.3 Event Management for BEST/1 Using TME 10 Enterprise Console

In this section we show how to use TME 10 Enterprise Console to process alerts from BEST/1.

We do not explain the function or concepts of TME 10 Enterprise Console in detail. If you are not familiar with TME 10 Enterprise Console, you can refer to the redbook *TME 10 Cookbook for AIX - Systems Management and Networking Applications*, SG24-4867.

With BEST/1+ for TME 10 modules comes a task to set up the Tivoli Enterprise Console and a file package to install the BEST/1 event adapter. The BEST/1 event adapter takes alerts from the BEST/1 monitor and sends them to the TEC. The Distributed Monitoring profiles (see 2.9.2, "Monitoring BEST/1 Using TME 10 Distributed Monitoring" on page 124) are also prepared to send events to TEC. In the following we describe how we set up our environment to receive BEST/1 events in the TEC.

2.9.3.1 Setting Up the Event Server

To set up the TEC correctly to receive events, the task Setup BEST/1 EventServer is provided. Before this task can be started a decision must be made of where to store the Tivoli rule bases.

We recommend storing them in /var/spool/Tivoli, because the Tivoli database and the Tivoli database backups are there, and it is recommended you back up this directory to a tape from time to time, so the TEC rule bases can be easily backed up, too. For this we created the directory ITSORuleBases. The following three figures show the installed rule base, defined event sources and event groups before the setup task was started.

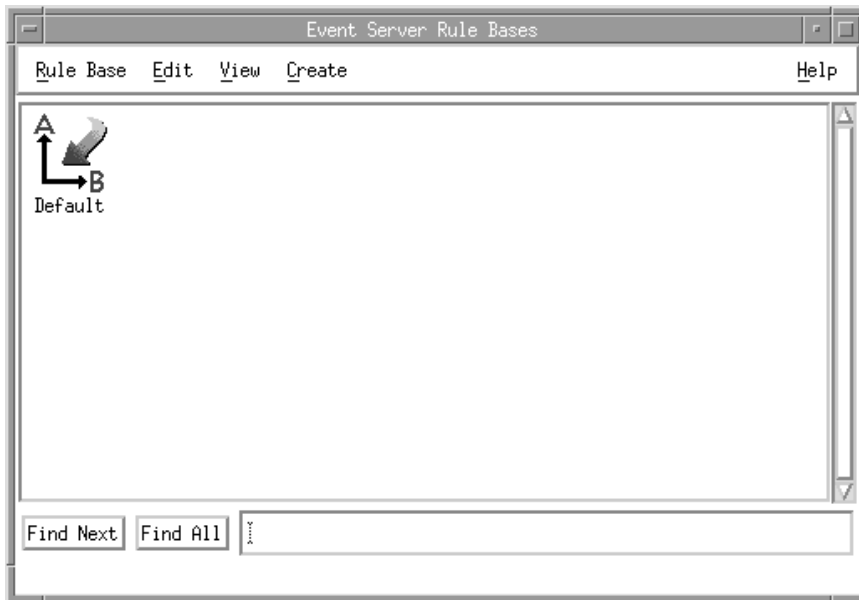


Figure 140. Rule Bases before Setup BEST/1 EventServer

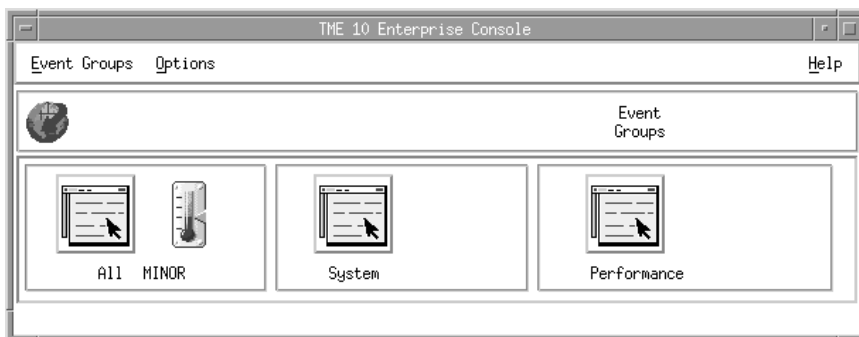


Figure 141. Event Groups before Running Setup BEST/1 EventServer Task

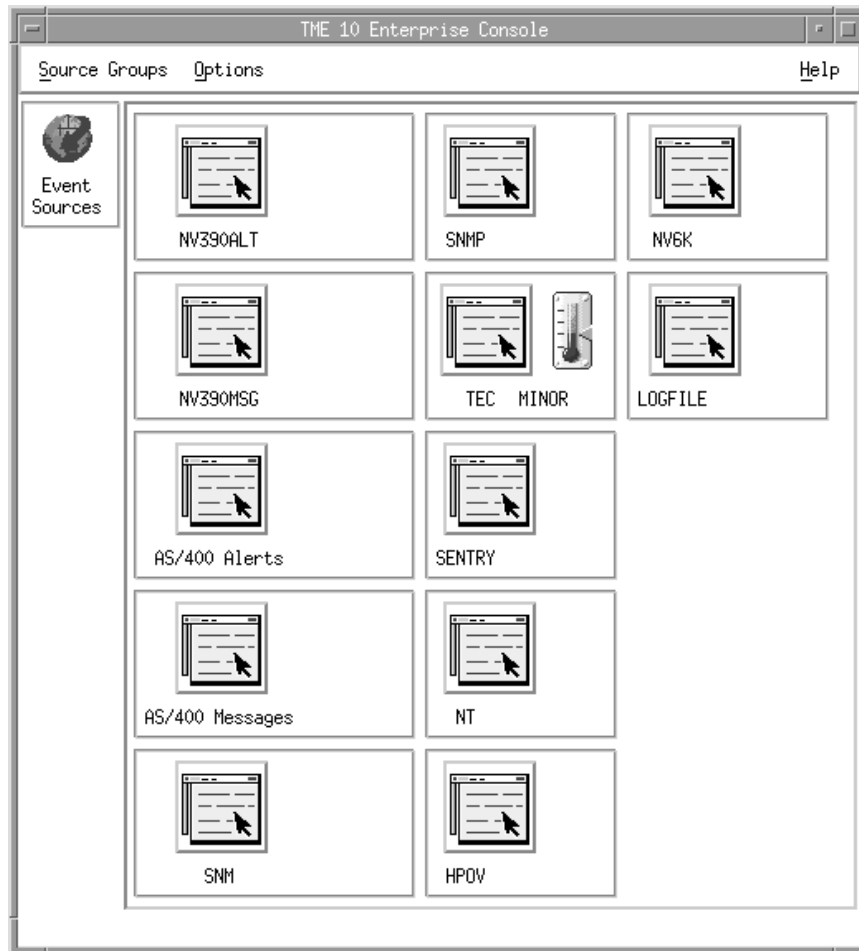


Figure 142. Event Sources before Setup BEST/1 EventServer

Now, we call the Setup BEST/1 Event Server task to prepare the TEC to receive events from BEST/1. There are two kinds of events, the BEST/1 alerts and the events that are triggered by the TME 10 Distributed Monitoring profiles, which also come with the BEST/1+ for TME 10 module.

The setup task asks for four parameters. The first is the name of the rule base. Here you can give the name of an existing rule base which is then extended, or the name of a new rule base. We created a new rule base named Best1_rb.

The second parameter is the rule base to clone. Usually you use the currently active and properly working rule base. This means all your defined event classes and rules are taken over to the new rule base and the BEST/1 event classes are added to this rule base. Since this is our first rule base, we clone the Tivoli Default rule base here.

The third parameter is the place where the rule bases are stored. As we stated above we want to have all our rule bases under /var/spool/Tivoli/ITSORuleBases. The BEST/1 rule base should be stored under Best1, so we enter /var/spool/Tivoli/ITSORuleBases/Best1. The Best1 directory should not exist, since it is created by the setup task.

The last parameter is optional. The setup task creates an event group named best1_PLUS. To see the events of a certain class at the event console of an

administrator we must have the event groups assigned for that administrator's event console. With the parameter here you can automatically assign the best1_PLUS event group to an administrator of your choice. To make the choice click on **Choose** and select the administrator you want to see the BEST/1 events. We are the root administrator, so we choose that.

The following figure shows the values we entered in our example.

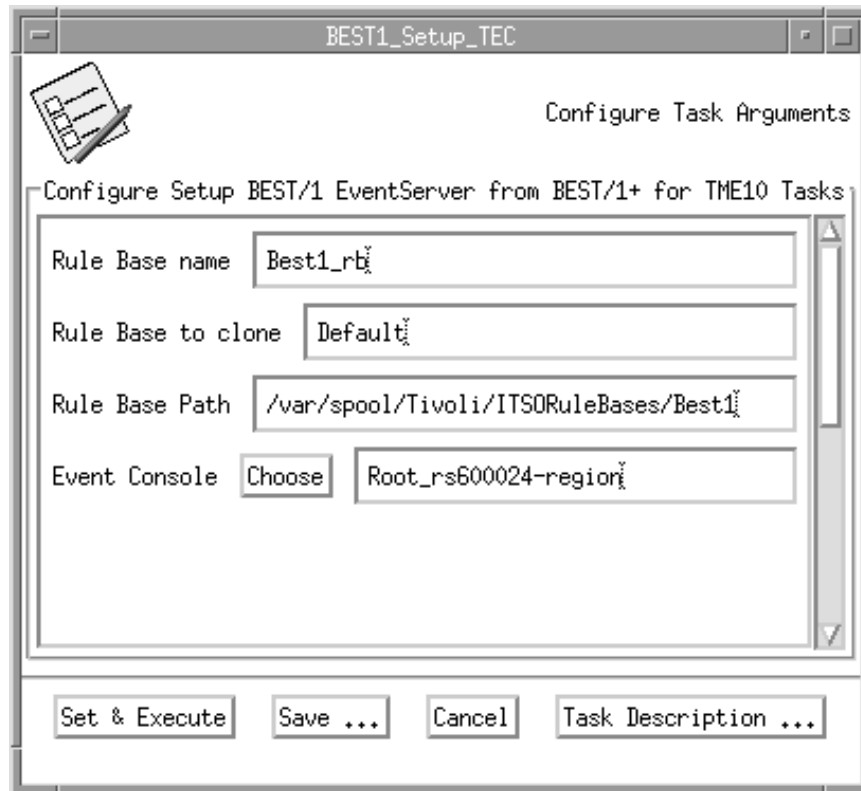


Figure 143. Event Server Setup Parameters

While the setup task is running a window appears showing the progress of the task. As you can see from the status messages everything goes well. The last steps are the shutdown and startup of the event server. In the example, startup went wrong; you don't see the red bold arrow anymore on the event server icon on your desktop. We executed the setup task several times and always got that error. We think this is due to the fast shutdown and restart. You have to wait a second to start up the event server again. You can start up the event server manually by clicking on the event server icon with the right mouse button and choosing **Start-up**. It should come up without any problems.



Figure 144. Setup BEST/1 EventServer Task Output

In addition to the task output you can check whether under the ITSORuleBases directory the right directory and file structure was created. To do this compare the structure with the directory structure of the Default rule base.

The next figure shows the output of `ls -al` in `/var/spool/Tivoli/ITSORuleBases/Best1`. Under this directory all the files concerning event classes of the Best1 rule base are created. The list contains two new baroc files. The baroc file `best1_sentry.baroc` contains the event classes to receive the TME 10 Distributed Monitoring events whereas the baroc file `tecad_best1.baroc` contains the event classes to receive the BEST/1 monitor alerts.

```

#pwd
/var/spool/Tivoli/ITSORuleBases/Best1/TEC_CLASSES
#ls -al
total 528
drwxr-xr-x  2 root sys      512 Nov 10 11:52 .
drwxr-sr-x  5 root sys      512 Nov 10 11:52 ..
-rwxr-xr-x  1 root root     224 Feb 08 1997 .event_sum_slots
-rw-r--r--  1 root nobody  211 Nov 10 11:52 .load_classes
-rwxr-xr-x  1 root root     105 Feb 08 1997 .task_slots
-rwxr-xr-x  1 root root    15957 Feb 08 1997 as400msg.baroc
-rwxr-xr-x  1 root bin     1581 May 19 15:42 best1_sentry.baroc
-rwxr-xr-x  1 root root     1901 Feb 08 1997 root.baroc
-rwxr-xr-x  1 root root      977 Feb 08 1997 tec.baroc
-rwxr-xr-x  1 root bin      512 May 19 15:42 tecad_best1.baroc
-rwxr-xr-x  1 root root     5386 Feb 08 1997 tecad_hpov.baroc
-rwxr-xr-x  1 root root    12742 Feb 08 1997 tecad_logfile.baroc
-rwxr-xr-x  1 root root     8035 Feb 08 1997 tecad_nt.baroc
-rwxr-xr-x  1 root root     5836 Feb 08 1997 tecad_nv6k.baroc
-rwxr-xr-x  1 root root     7944 Feb 08 1997 tecad_ov.baroc
-rwxr-xr-x  1 root root    14790 Feb 08 1997 tecad_snaevent.baroc
-rwxr-xr-x  1 root root     2826 Feb 08 1997 tecad_snm.baroc
-rwxr-xr-x  1 root root    146674 Feb 08 1997 tecad_snmp.baroc

```

Figure 145. BAROC Files of Best1 Rule Base

In the next three figures you see on the GUI that the BEST/1 rule base was created and loaded. The event source BEST1 was created and the event group best1_PLUS was also created.

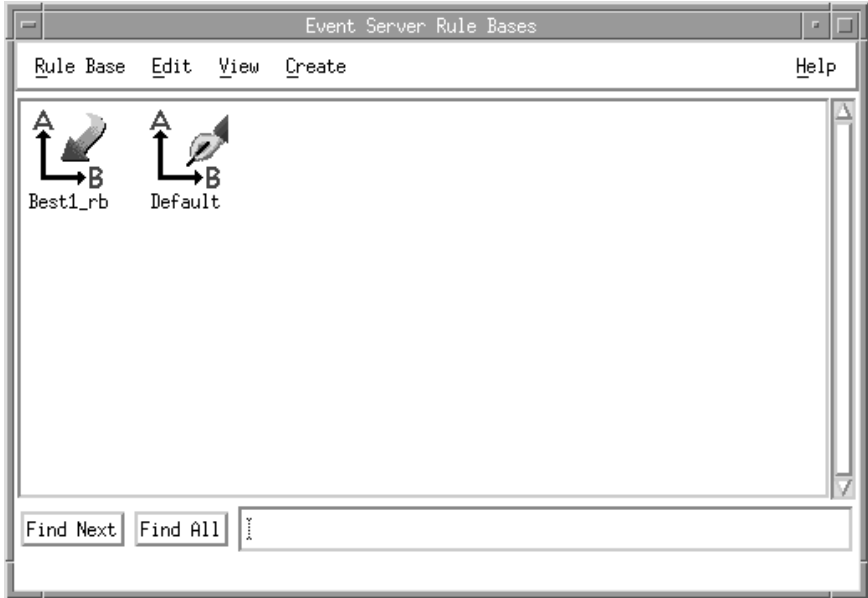


Figure 146. Rule Bases after Running Setup Up BEST/1 EventServer Task

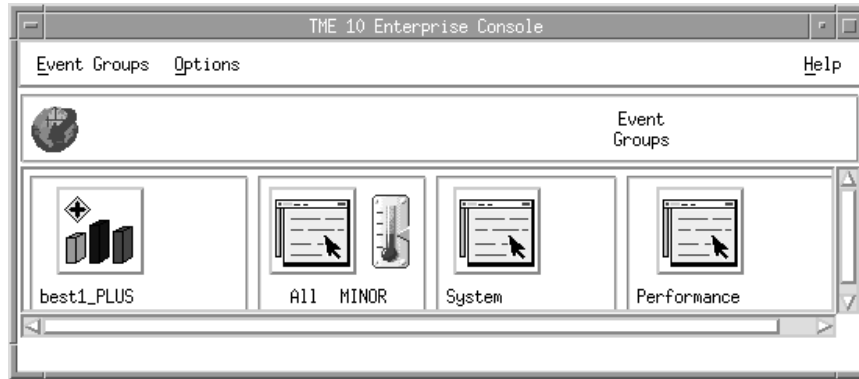


Figure 147. Event Groups after Setting Up BEST/1 EventServer

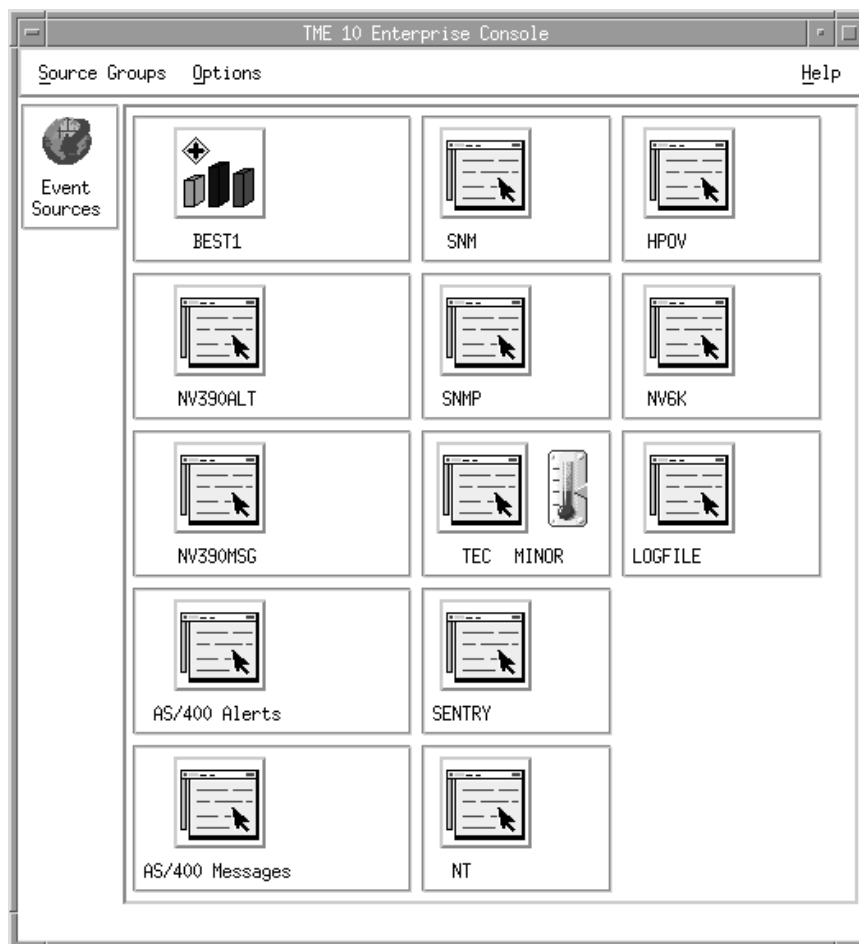


Figure 148. Event Sources after Setting Up BEST/1 EventServer

TME 10 Enterprise Console is now ready to receive and process events from BEST/1 and TME 10 Distributed Monitoring. To check the proper setup of the event server you can use the `wpostemsg` command. The following examples test two event classes:

```
wpostemsg -m "Global Vision GmbH, Germany" BEST1_EVENT BEST1
```

or

```
wpostmsg -m "Global Vision GmbH, Germany" BEST1_Repeat_Daemon BEST1
```

The following figure shows the result at the TEC server.

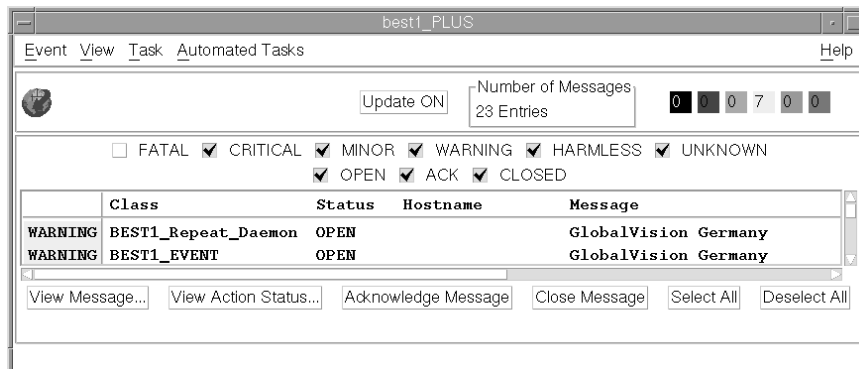


Figure 149. Events Arriving at TEC Server

Another way to check if events are received and processed correctly is the use of the `wtdumpri` command. No parameter is needed and all received events and executed processing actions performed on it are printed to standard output. For details refer to the TME 10 Enterprise Console documentation.

2.9.3.2 Install BEST/1 Adapter

To get real events from TME 10 Distributed Monitoring simply distribute the profiles (see 2.9.2, "Monitoring BEST/1 Using TME 10 Distributed Monitoring" on page 124). To get real events from BEST/1 we have to install the BEST/1 event adapter on the BEST/1 server, which is in our environment `rs600024`.

We have to distribute the file package Install BEST/1 Adapter. You can have a look inside this file package by clicking on the file package icon with the right mouse button and choosing **Properties**.

We have to distribute this file package to the BEST/1 server, so we have to set the subscribers. This is done by clicking on the file package icon with the right mouse button again and choosing **Subscribers**; select `rs600024` and then click on **Set Subscriptions & Close**. Now we can distribute the file package by selecting the file package icon with the right mouse button and then selecting **Distribute** from the pop-up menu.

Now we should be able to get the real events from the TME 10 Distributed Monitoring monitors and from the BEST/1 monitor. Prerequisite is of course that events are triggered. We fulfill this prerequisite by changing the trigger values for a particular monitor in the TME 10 Distributed Monitoring profile BEST/1 Remote Monitors and by setting up a monitoring policy in BEST/1 (see *BEST/1-Monitor for UNIX User's Guide*) and activating them.

Afterwards we have the following events in our event console for the event group `best1_PLUS`:

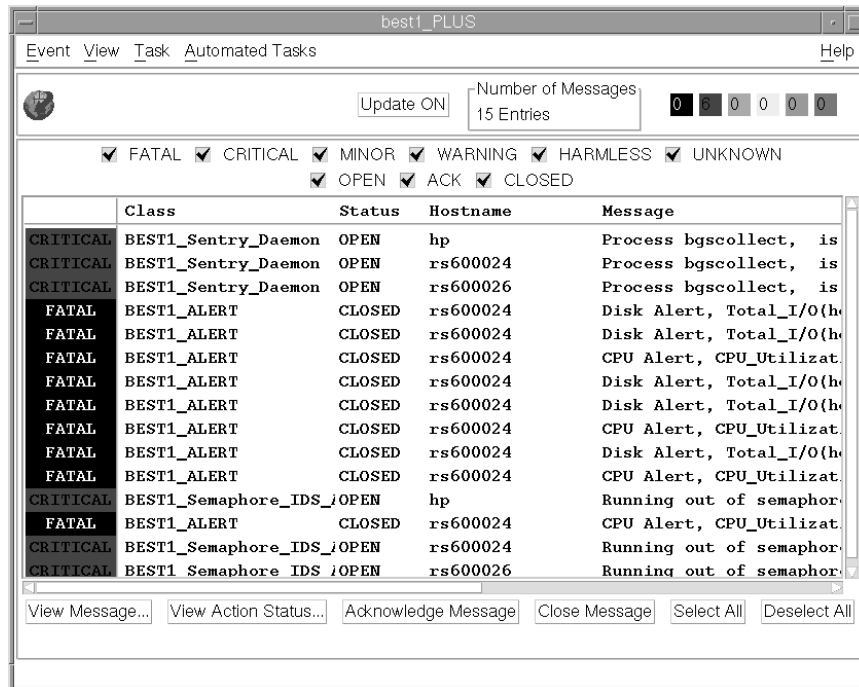


Figure 150. Events of best1_PLUS Event Group

As you can see we got several events from different hosts. The event of class BEST1_Sentry_Daemon comes from TME 10 Distributed Monitoring and monitors the status of the BEST/1 daemon bgscollect. Another TME 10 Distributed Monitoring event is the one of class BEST1_Semaphore_IDS showing the used semaphore IDs. We also see two different BEST/1 events of class BEST1_ALERT, one monitoring the CPU utilization the other the disk utilization. For a complete list of items to monitor see *BEST/1 Monitor for UNIX User's Guide*, 97083-00.

The BEST/1 event adapter is implemented as a perl script which listens to the BEST/1 monitor logfile. The BEST/1 event adapter can be started and stopped manually by calling the script `tecad_best1.sh [start|stop]`

in `/etc/Tivoli/tecad/bin`. To test the BEST/1 event adapter you can generate test log entries in the BEST/1 monitor logfile. In the file `/tmp/testlog1` we prepared a test log entry:

```
Thu_Nov_13_12:25:32_1997 rs600024 "Martin Keller GlobalVision" Sev=2 "
```

We pushed the test log entry to the end of the monitor logfile by typing:

```
cat /tmp/testlog >>/best1/bgs/monitor/log/b1monitor.log
```

The event appears in the event console:

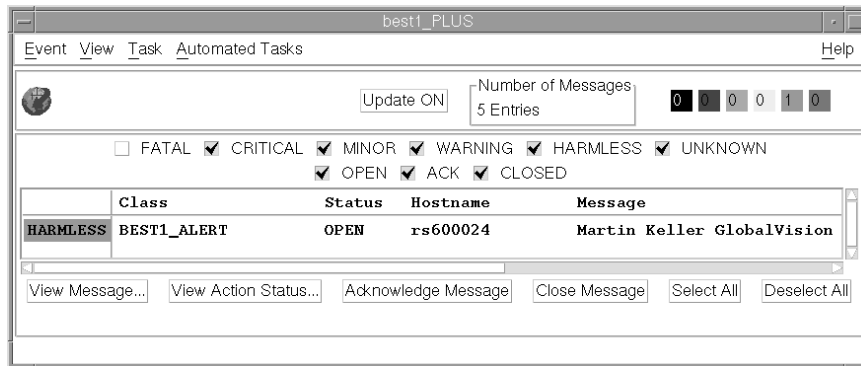


Figure 151. BEST/1 Test Event Pushed into the BEST/1 Monitor Logfile

2.9.4 Installing BEST/1-Visualizer Using an AutoPack

BEST/1+ for TME 10 does not provide a function to automatically install Visualizer using TME 10. Visualizer is a Windows NT application that is used to create graphical reports from data gathered by BEST/1.

In this section we use AutoPack which is a feature of TME 10 Software Distribution to create a file package that can be used to install Visualizer on any Windows NT system in the Tivoli TMR.

We do not explain the AutoPack feature in detail. If you need detailed information about AutoPack, refer to the redbook *The TME 10 Deployment Cookbook: Courier and Friends*, SG24-4976.

Visualizer is usually installed manually on a system using a standard Windows InstallShield installation.

With AutoPack, you can take a snapshot of the system before installing Visualizer, then install Visualizer manually and then take another snapshot of the system after the installation. By comparing the snapshots, AutoPack can then determine the files added to the system and the system parameters that were changed during the installation.

AutoPack contains an AutoPack file containing all changes that can then be included in a TME 10 Software Distribution profile and distributed to a number of systems, thus installing AutoPack.

In order to perform the following example, AutoPack Control Center, which can be found on the TME 10 Software Distribution CD-ROM, needs to be installed on the system where we install Visualizer.

After AutoPack Control Center has been started, you will see the following window:

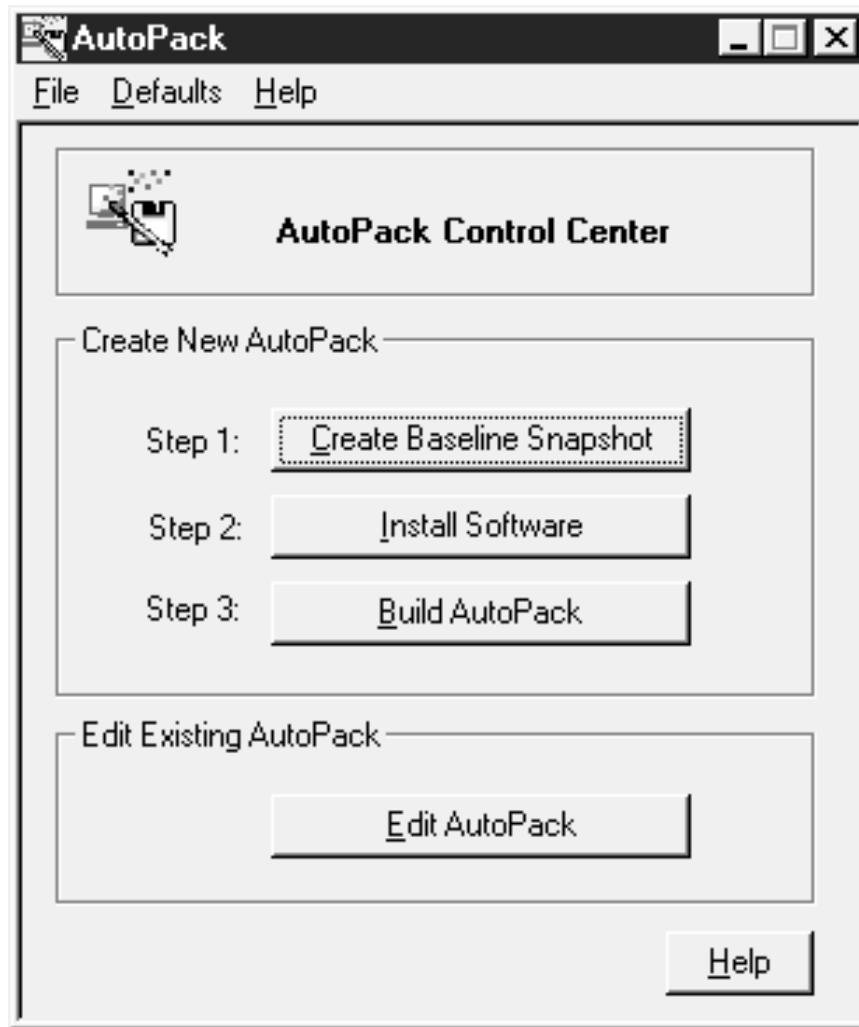


Figure 152. AutoPack Control Center

Click on **Create Baseline Snapshot** to perform the initial scan of the system.

The following window will appear:

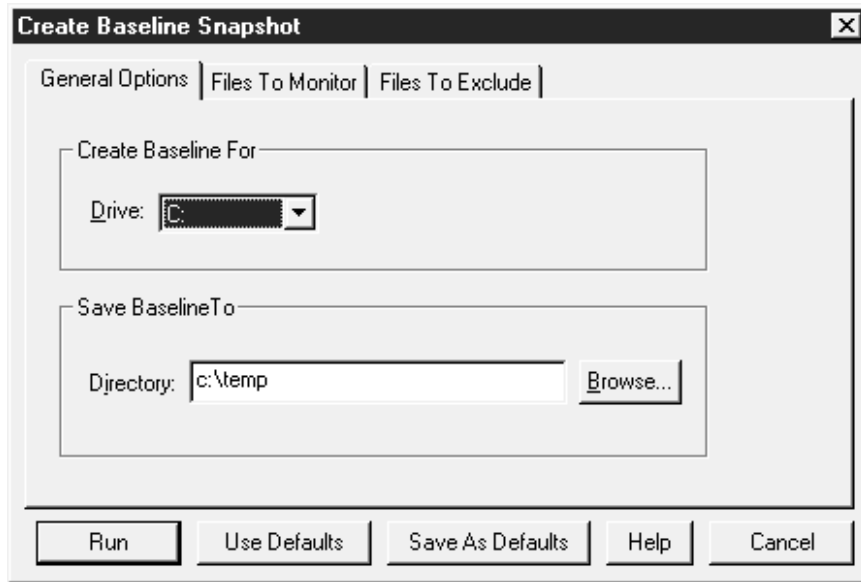


Figure 153. Create Baseline Snapshot Window

In this window you can specify a number of parameters influencing the scanning, for example, the files to be excluded from scanning.

When finished, click on the **Run** button.

The following window will appear:

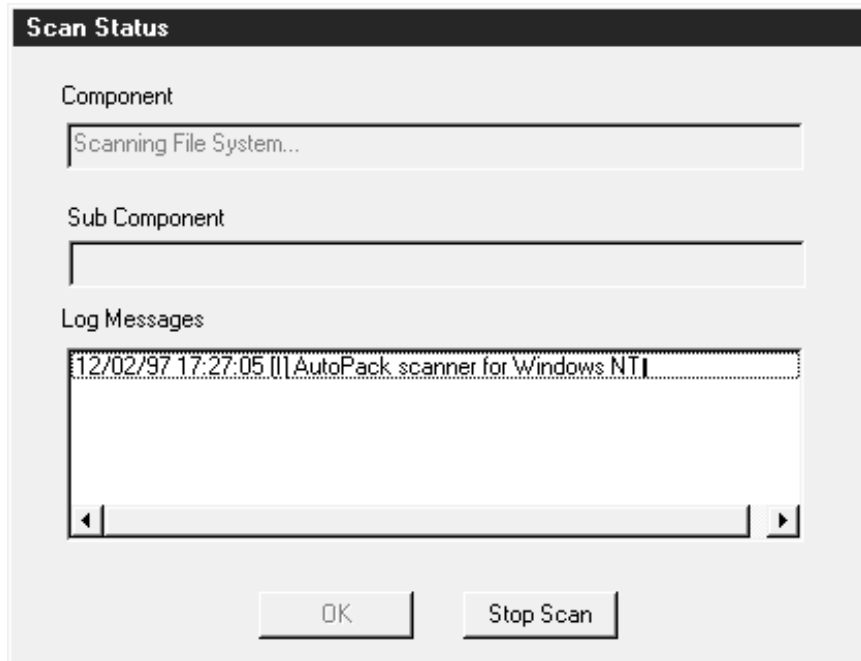


Figure 154. Scan Status Window

The above window displays information about the scanning in progress and informs you when the scanning process has completed. As soon as this is the case, click the **OK** button to get back to the AutoPack Control Center main window as displayed in Figure 152 on page 145.

This time, click the **Install Software** button. The following window will appear:

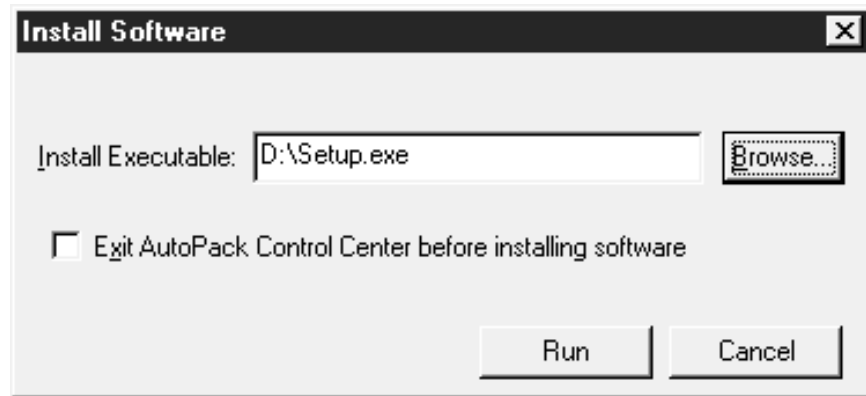


Figure 155. Install Software Window

Click on the **Browse...** button and select the installation program for Visualizer. In our case, we have the Visualizer CD in drive D:, so we select D:\Setup.exe as the install program.

When finished, we select the **Run** button. This will start the InstallShield installation for Visualizer.

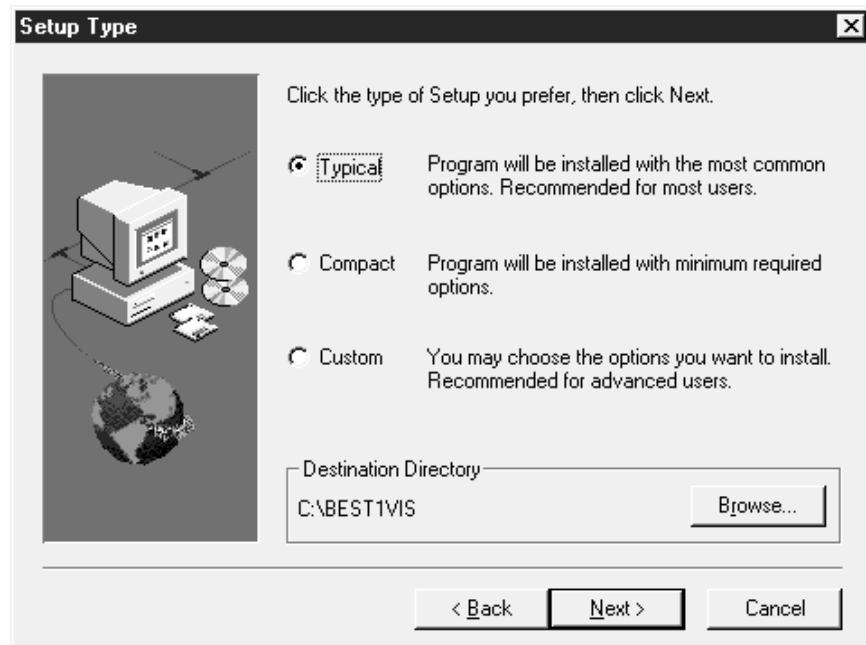


Figure 156. Setup Type

In the Visualizer Setup Type window we click on **Typical** to perform an installation with typical installation values. To actually start the installation we click the **Next>** button.

After a few more installation panels, the installation program will start copying Visualizer to the hard disk.

Once the installation has finished we go back to the AutoPack Control Center main window as shown in Figure 152 on page 145 and this time click on the **Build AutoPack** button.

The following window will appear, asking us to specify a name for the AutoPack file to create.

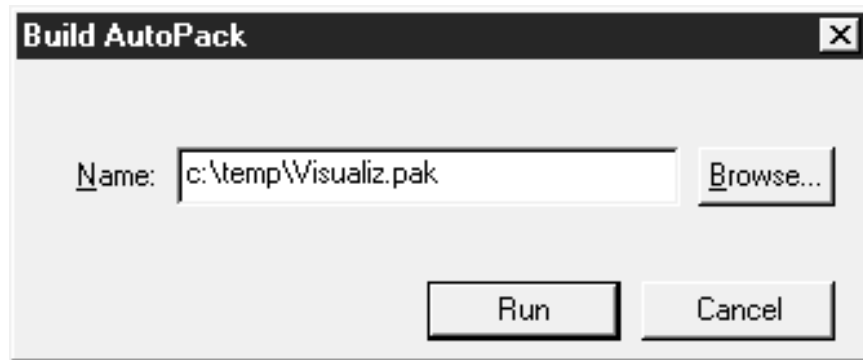


Figure 157. Build AutoPack Window

When we have finished entering a name of choice, we click the **Run** button, which will get us to the following window:

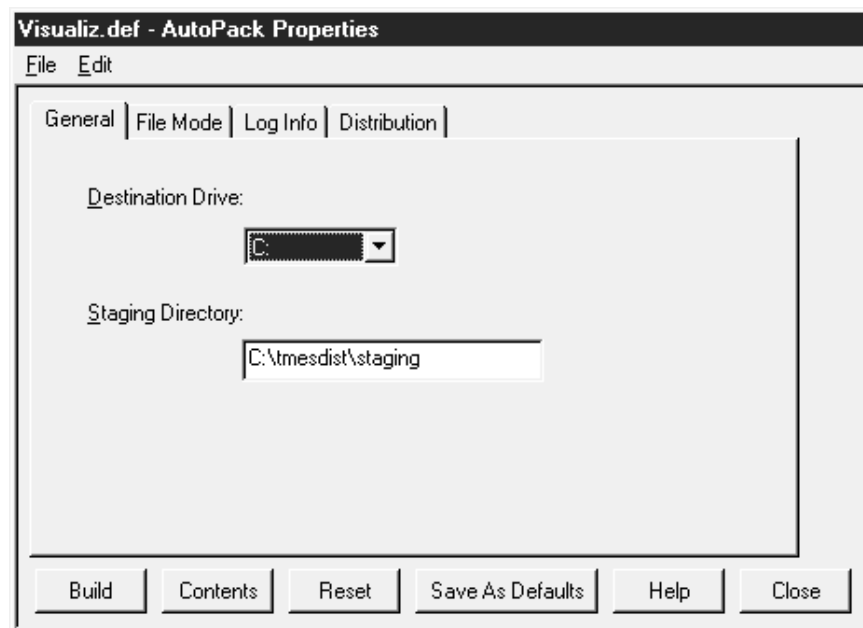


Figure 158. AutoPack Properties Window

In this window we can change the properties of the AutoPack, for example, set distribution options.

We leave the defaults and select the **Build** button. This will actually create the AutoPack file. When this process has finished, we get the following confirmation message:



Figure 159. AutoPack Window

We can now retrieve the AutoPack file we have just created into a TME 10 Software Distribution file on our TMR server and then distribute this profile using TME 10.

After distributing the profile to another Windows NT machine, Visualizer will be installed on this machine, including all system changes made by the Setup program, such as menus and folders.

The following figure shows the BEST_1 Visualizer menu on the target system.

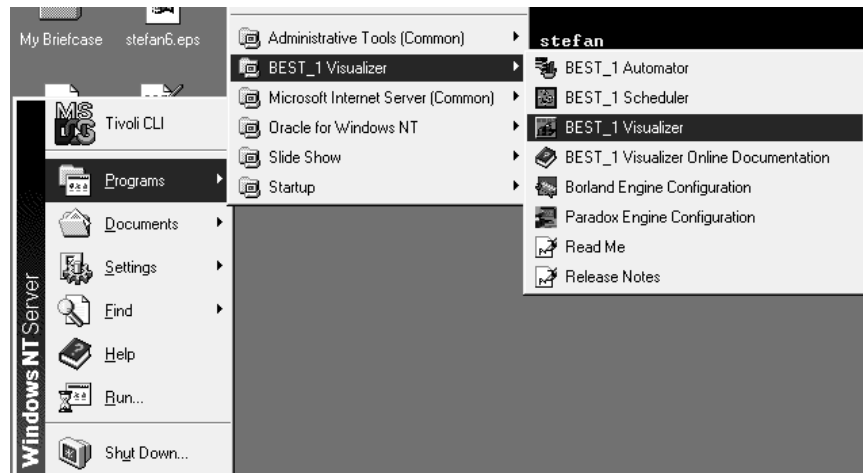


Figure 160. Visualizer Installed on the Target System

2.10 How Can BEST/1 Help Us with Capacity Management for TME 10

So far, we have shown the most important functions of the BEST/1 product and how BEST/1+ for TME 10 can help us manage BEST/1 from TME 10.

Now let's see how BEST/1 can help us improve and plan a TME 10 environment.

The most common challenge in designing and deploying a TME 10 environment is to size the components of the environment, such as servers, disk space, memory, network and so on.

In practice, everybody who has to design a TME 10 infrastructure faces questions such as:

- What processor model do I need to run a TME 10 Managed Node?
- How big does my TMR server need to be?
- How much disk space, memory and CPU power is needed for a TME 10 Enterprise Console server, assuming I have to process 10 events per second on average?
- Can the TME 10 Software Distribution server that I have sized handle 20 additional clients?

How can BEST/1 help us answer these and many other questions?

BEST/1's unique feature to predict future environments is very well suited for this task. With the BEST/1-Monitor we can first measure performance of a component in a test environment, then set our Service Level Objectives and simulate how different configurations would meet these SLOs.

In this section we show a simple example of how to measure and size a TME 10 Enterprise Console server. This example is just included to show you how to start and is not meant to be a full-blown sizing methodology for TME 10 Enterprise Console; however, it should give you a good understanding of the principle.

2.10.1 Analyzing TME 10 Enterprise Console Performance

Predicting the performance of TME 10 Enterprise Console and sizing the TEC servers is a very common challenge in a TME 10 environment. In this section we describe how we stress tested the TEC, measured the performance and found a suitable machine for a given load.

For this we had to:

- Develop a stress test
- Define a working model for BEST/1 Analyze
- Collect baseline data
- Run the stress test, while collecting the performance data again
- Analyze the collected data
- Perform "What if?" modeling

2.10.2 TEC Stress Test Implementation

To bring the TEC under load we wrote a script that sent events to the TEC using the wpostmsg command. We chose an event that also triggered the evaluation of a rule.

```

#####
#
# script to test T/EC performance
# sends events with wpostemsg command in an forever loop
# a running number is send as part of the message slot
#
# 11/14/97
# Martin Keller, GlobalVision Germany
#
#####

NUMBER=1

while [ $# -eq 0 ]
do
  wpostemsg -m "event number: $NUMBER" BEST1_Sentry_Fileystem_Full BEST1
  NUMBER= expr $NUMBER '+' 1
done

```

Figure 161. Stress Test Script

2.10.3 Definition of a Transaction Class for Tivoli

In the next step we defined a transaction class for Tivoli. A transaction class is used by the BEST/1 Analyze module to analyze performance for a given transaction. For details about defining transaction classes see 2.6.3.6, "Transaction Classes" on page 53.

This means we have to define which processes make up Tivoli. We added the following processes to our transaction class named Tivoli:

- oserv (Tivoli Object Request Broker process)
- dataserver (TEC Sybase database process)
- tec_rule (TEC Rule engine)
- tec_dispatch (TEC Dispatcher, bringing the events to the consoles)
- tec_reception (TEC Reception engine)
- tec_task (TEC process executing tasks associated with events)
- tec_server (TEC master process, controlling the others)
- tecad_best1.pl (event adapter for BEST/1)
- sentry_engine (process executing the distributed monitors)

Since we defined the appropriate advanced tuning option for this transaction class all, child processes also count in our performance analysis.

2.10.4 Collecting the Baseline Data

To have a reference to the high load we expect when calling our script, we started the BEST/1 collector during a period of normal load to collect performance data. We evaluated this data with BEST/1 Predict and got only green pie charts as you can see in the next figure:



Figure 162. Analyzing TEC Workload

2.10.5 Running the Stress Test

We then started our script and sent as many events to TEC as possible. As the end of the wtdumpri output shows 1432 events were sent.

```

1 2765 0 879541505(Nov 14 16:05:05 1997)
### EVENT ###
BEST1_Sentry_Filesystem_Full;
source=BEST1;
msg='event number: 1';
origin=9.24.104.161;
END

### END EVENT ###
PROCESSED

1 2766 0 879541505(Nov 14 16:05:05 1997)
### EVENT ###
BEST1_Sentry_Filesystem_Full;
source=BEST1;
msg='event number: 2';
origin=9.24.104.161;
END

### END EVENT ###
PROCESSED

1 2767 0 879541507(Nov 14 16:05:07 1997)
### EVENT ###
BEST1_Sentry_Filesystem_Full;

.
.
.
.
.
.
.

origin=9.24.104.161;
END

### END EVENT ###
WAITING

1 4216 0 879542384(Nov 14 16:19:44 1997)
### EVENT ###
BEST1_Sentry_Filesystem_Full;
source=BEST1;
msg='event number: 1432';
origin=9.24.104.161;
END

### END EVENT ###
WAITING

```

Figure 163. wtdumpri Output

Our BEST/1 Predict evaluation now has red parts in the pies:



Figure 164. Analyzing TEC Performance

2.10.6 Interpreting the Test Results

The pies can be interpreted as follows. The Nodes... pie indicates that there is one of two nodes where the defined service level threshold is reached. The Workloads... pie indicates that for one workload the threshold is reached. And the Transactions... pie shows that there is one transaction that reaches the threshold.

The next figure goes deeper into analyzing the transaction utilization. The list shows that the rs600024 CPU is 98.39% busy. The Tivoli transaction class is responsible for 78.86% of the load. All other transactions produce 19.5%. The X-Window processes are responsible for 0.02% of the load.

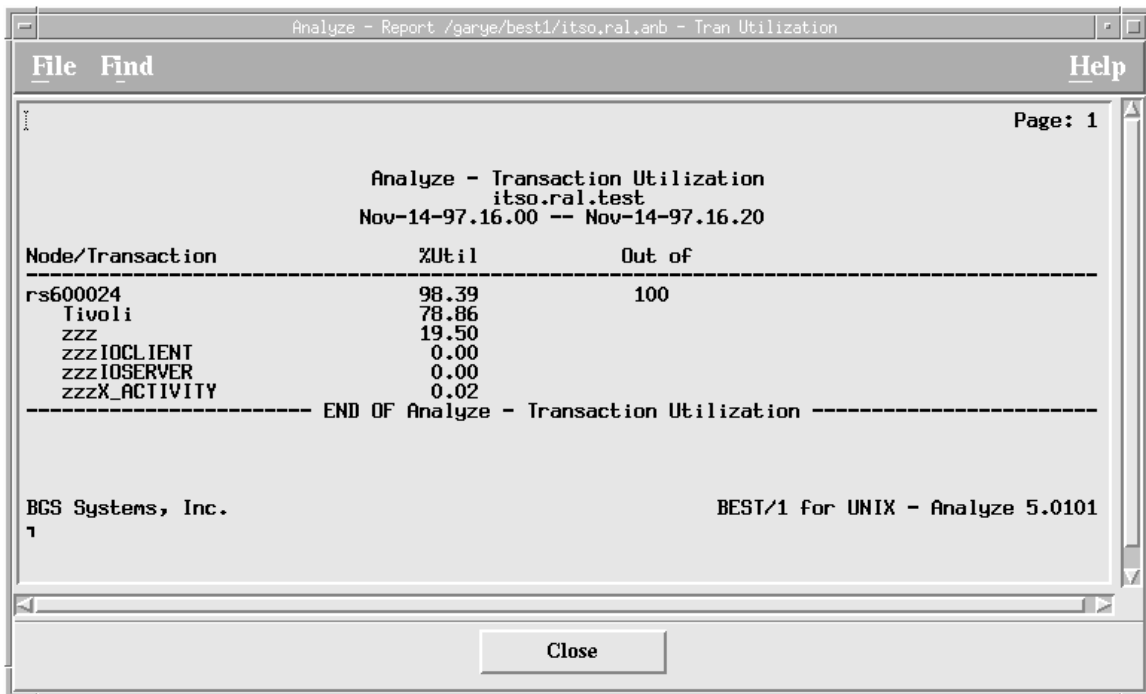


Figure 165. Analyzing Transaction Classes

We analyze the Tivoli transaction class more deeply and can see in Figure 166 on page 155 which processes caused the load. We see that the oserv (190.12 seconds of CPU time) and the dataserver (150.49 seconds of CPU time) processes caused the main load. This is because each event arrives through the oserv and is stored in the TEC database.

We also see that 94.59 seconds of CPU time are used by an oserv subprocess named RIM_Sybase_prog. TEC 3.1 uses RIM to support several databases. The RIM_Sybase_prog writes the events to the TEC Sybase RDBMS.

The tec_rule process consumes 60.83 seconds of CPU time, which is caused by the rule evaluation triggered by the event we have chosen for our stress test.

The tec_reception process (8.13 sec) had a low CPU time consumption, which shows that the reception and the network are not the bottleneck in our configuration.

The tec_task process (7.77 sec) had a low CPU time consumption too, since no task was triggered by the events we sent.

With these results a first conclusion can be made. The CPU load could be reduced by a faster CPU or more main memory, because a lot of CPU time is consumed by the oserv and internal TEC processes. But for sure a certain point will be reached where the network and the throughput of events might be the bottleneck.

Node/Process	PID	User Name	CPU Time (in secs)	Major Pg.Flts	Parent PID	Start Time
rs600024						
oserv	30488	root	198.12	268	7100	Nov 7 19:24:46 1997
dataserver	24996	root	150.49	23	1	Nov 7 15:56:50 1997
RIM_Sybase_prog	33326	nobody	94.59	1575	30488	Nov 7 19:24:54 1997
tec_rule	21558	root	60.83	2	47142	Nov 12 18:37:56 1997
tec_dispatch	36658	root	55.30	0	47142	Nov 12 18:37:56 1997
Collection_prog1	33572	nobody	8.68	0	30488	Nov 7 19:24:51 1997
ksh	34580	root	8.41	0	40992	Nov 14 16:05:04 1997
tec_reception	40244	root	8.13	0	47142	Nov 12 18:37:56 1997
tec_task	26424	root	7.77	10	47142	Nov 12 18:37:56 1997
TNR_prog1	27948	nobody	7.75	0	30488	Nov 7 19:24:55 1997
init	1	root	6.71	0	0	Nov 7 15:42:36 1997
cdsclerk	8576	root	5.58	0	8274	Nov 7 15:44:15 1997
dced	8186	root	3.98	0	1	Nov 7 15:43:50 1997

Figure 166. Analyzing Processes

2.10.7 "What if?" Modeling

The next picture shows a model based on our collected data. The actual machine used to collect data, was a 100 MHz RS/6000 43P workstation. For modeling we could choose from a couple of machines (see Figure 168 on page 157). To get this list you have to click the **Browse** button. After trying several machines, we chose a 200 MHz RS/6000 604E and got a processor utilization of 44.2%. This is beneath our defined threshold of 80.0% and beneath the guideline of 70.0%.

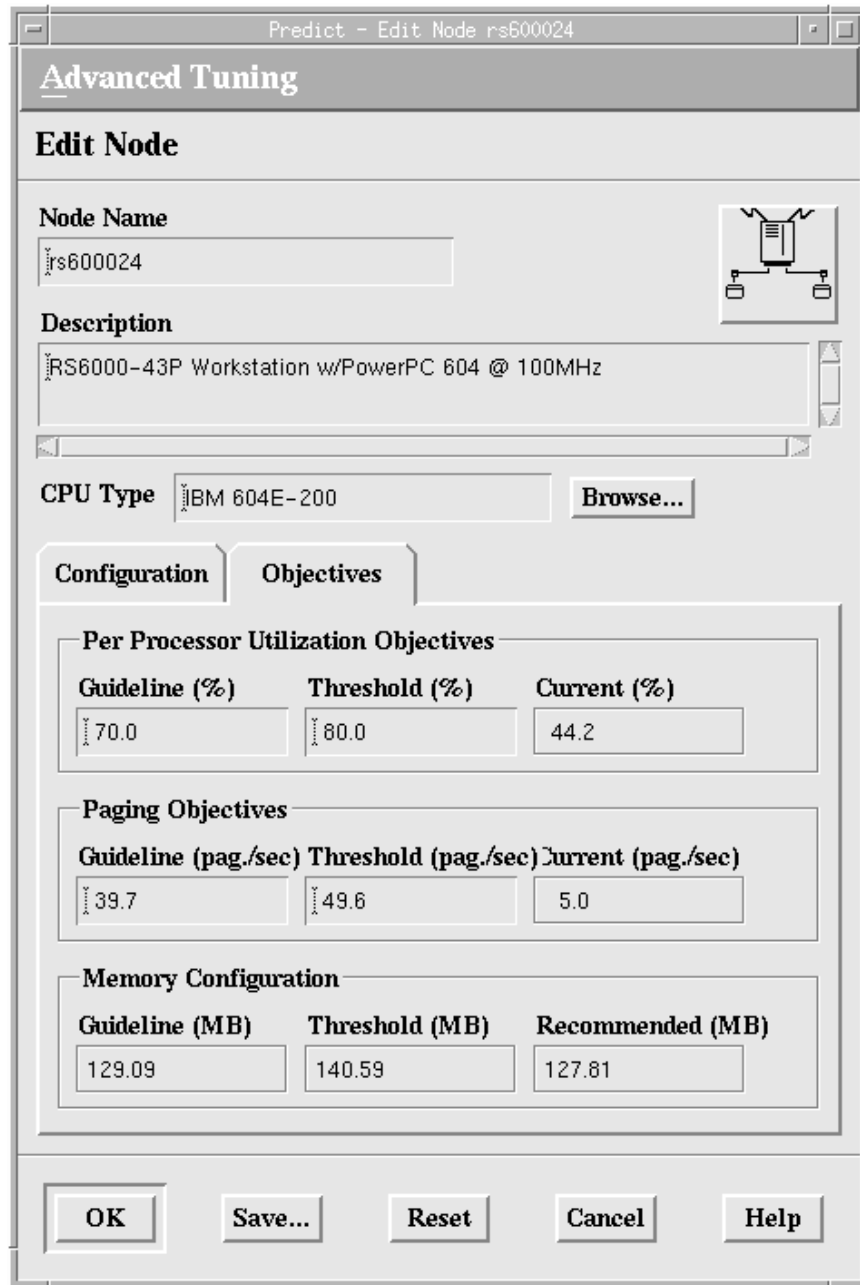


Figure 167. "What if?" Modeling

The following figure shows an excerpt of the processor models BEST/1 offers us to select from. Since we are using IBM RS/6000 systems we are interested particularly in how a larger RS/6000 processor will affect our environment. We could, however, also select from a wide range of processors from other vendors.

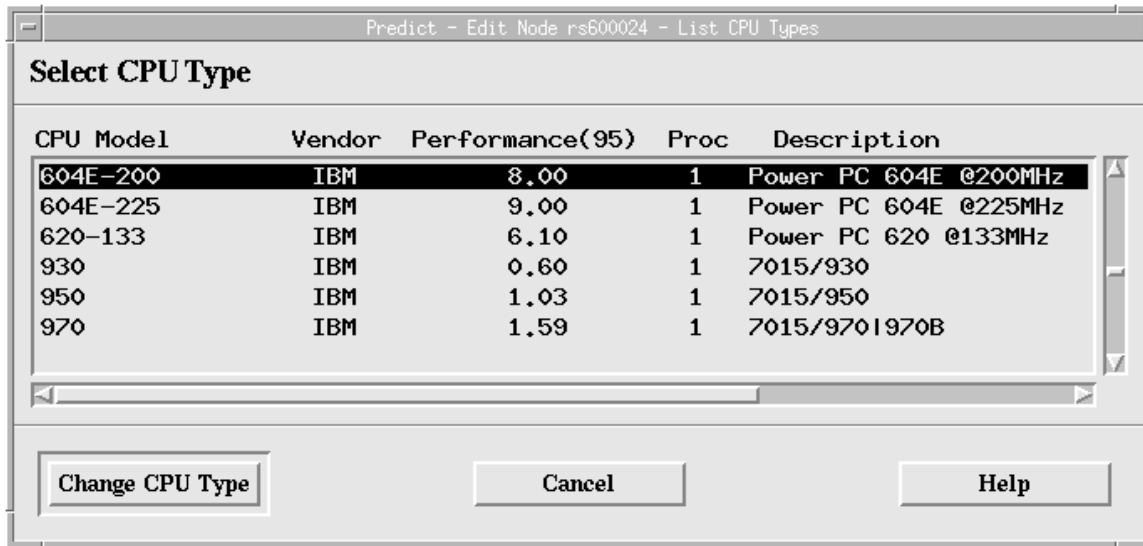


Figure 168. Selecting a CPU Type

Similar selections are also offered for hard disks and other common hardware devices.

2.11 Summary

In this chapter we have demonstrated how BEST/1, a premier solution for performance management and capacity planning can be used in combination with TME 10, the leading solution for enterprise systems management.

To achieve this, the BEST/1+ for TME 10 module from BGS Systems Inc. can be used to link BEST/1 to TME 10, using TME 10's core application functionality to manage and ensure availability of the BEST/1 application.

We have demonstrated how to use the key functions of the module, namely deployment of BEST/1 clients using TME 10 Software Distribution, monitoring BEST/1 using TME 10 Distributed Monitoring and forwarding alerts from BEST/1 to TME 10 Enterprise Console.

In a production environment, the BEST/1+ for TME 10 module allows you to rapidly deploy large numbers of BEST/1 clients automatically and keep control over a growing BEST/1 environment. All systems management functions can be performed from the TME 10 Desktop, while BEST/1 is operated using the BEST/1 user interface. This user interface can be launched from the TME 10 Desktop from any UNIX TME Managed Node in your TMR.

We have also demonstrated how BEST/1 can immediately add value to any TME 10 environment by showing an example of how to size TME 10 components using BEST/1.

Chapter 3. DCEmgmt from Santix Software GmbH

In this chapter we give an overview of the DCEmgmt product from Santix Software GmbH. DCEmgmt is a Pinnacle level application for management of Distributed Computing Environment (DCE) and is tightly integrated with TME 10.

We give a functional overview of the main DCEmgmt components and show figures to illustrate specific functions.



Detailed examples of how to use DCEmgmt to manage DCE can be found in the redbook *Tivoli TME 10 Management for DCE and DFS*, SG24-2027. (Available at a later date.)

Note

DCEmgmt is a complex client/server application that has a large number of features. Hence, in this chapter we can give only an overview of some key components in order to illustrate the functions of DCEmgmt. For a complete description of all features of DCEmgmt refer to the appropriate product documentation or contact your Santix Software GmbH representative.

3.1 About Santix Software GmbH

Santix Software GmbH is a software company in the area of distributed computing and systems management.

Besides software development, Santix provides services in the areas of consultancy, training and systems integration. Santix is a leading service provider for TME 10 professional services in Europe.

Headquartered in Munich, Germany, Santix is a Tivoli Certified Training Center and a Pinnacle Member of the 10/Plus Association.

More information about Santix can be obtained under www.santix.de.

3.2 About DCEmgmt

DCEmgmt is a suite of management applications that significantly simplify the management of Distributed Computing Environment (DCE). We describe DCE in 3.3, “What is DCE?” on page 161.



DCEmgmt is a Pinnacle level TME 10 partner application, that is it makes extensive use of core TME 10 Framework services and the TME 10 core applications. DCEmgmt was one of the first partner applications that were tightly integrated into Tivoli and is the most comprehensive product to manage enterprise-wide DCE installations. We look at DCEmgmt in more detail in 3.4, “What Does DCEmgmt?” on page 162.

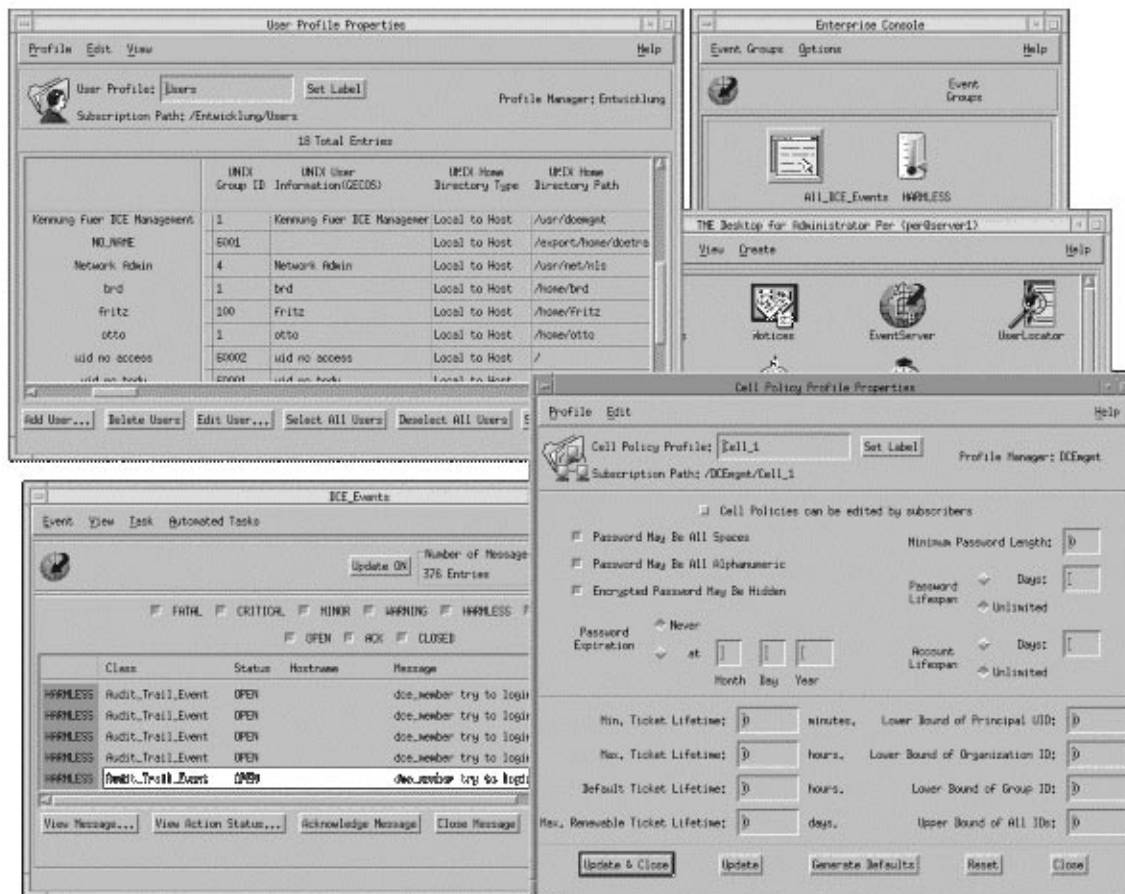


Figure 169. DCEmgmt Sample Screen

3.3 What is DCE?

DCE stands for Distributed Computing Environment. The specification for DCE was developed by the Open Software Foundation (OSF). DCE is a standard for client/server computing that offers a complete architecture model for client/server computing, including a wide range of services and APIs.

The DCE standard is implemented in products of almost all major operating system vendors. DCE implementations are available, for example, on UNIX, Windows, OS/390, VM, OS/400, DEC VMS and BS/2000.

On the one hand the services and APIs offered by DCE can be used to build client/server applications; on the other hand DCE offers a range of administration tools for client/server environments also. DCE is a middleware layer, that means it hides operating system and network specifics from the application and enables a server application to work across operating systems and networks.

The main components of DCE are:

- The DCE Cell
- DCE services:
 - DCE threads
 - DCE Remote Procedure Call (RPC)
 - DCE Cell Directory Service (CDS)
 - DCE Security Service
 - DCE Distributed Time Service (DTS)
- Distributed File System (DFS)

We briefly explain the components in the following.

The DCE cell is the means of organizing and grouping things in DCE. A DCE cell is a domain that contains systems, users and other DCE resources. A computer system can only belong to one DCE cell at a time.

Threads are used to divide a process into multiple concurrent execution units. Many operating systems support the concept of threads. DCE threads are implemented on an operating specific independent level, which means that DCE applications and services can use DCE threads and DCE will shield the operating specifics from these applications. When DCE is running on an operating system that does not support threads, DCE maps the DCE threads, for example, to an operating system process.

Remote Procedures Calls (RPCs) are used to create applications that invoke procedures that cannot only run locally but also on any system in the network. RPCs are one of key ingredients needed to create distributed applications. RPCs are available on a number of operating systems. What makes the DCE RPC special is that it is standardized and available on any system that implements DCE.

The DCE Cell Directory Service is used to identify resources within a DCE network, that is within one or interconnected DCE cells. The difference between network or operating system-specific naming services is again that DCE CDS is platform-independent.

The DCE Security Service provides a common means of securing data in the network and is used by the core DCE services, such as the DCE RPC and DFS.

The DCE Time Service provides a consistent mechanism for synchronizing time information in a network between different system types.

3.4 What Does DCEmgmt?

DCE has been traditionally difficult to manage. The reason for this is mainly that DCE comprises a wide range of functions and services, however, the tools and applications within DCE provide only a basic toolset to manage DCE components themselves.

The consequence of this is that DCE, although meant to ease the management of distributed systems, imposes a management challenge itself to the administrator.

DCE itself comes with a command line interface (CLI) that can be used to administer the DCE suite. To ease the administration, several vendors have added more user friendly interfaces, such as a SMIT interface for AIX; however, these interfaces are not standardized between platforms and they usually just provide a user interface on top of the function of the DCE command line interface.

This is where DCEmgmt comes into place. DCEmgmt provides a consistent interface on a wide range of platforms to manage DCE resources. Besides just providing a user friendly interface, DCEmgmt utilizes TME 10 services to make DCE management easier and more efficient.

Depending on the DCE component to manage, DCEmgmt ties into different TME 10 components, each of which we explain in further detail.

The components of DCEmgmt are:

- DCEmgmt/Security Manager
- DCEmgmt/Event Manager
- DCEmgmt/Cell Manager
- DCEmgmt/DFS Manager

3.5 DCEmgmt/Security Manager

DCEmgmt/Security Manager can be used to manage DCE users, DCE user groups and additional DCE components from the TME 10 desktop. To do so, DCEmgmt/Security Manager provides an extension to TME 10 User Administration that allows you to administer the specific attributes of DCE users and groups.

We show examples for this in 3.9.2, "DCEmgmt/Security Manager Examples" on page 166.

3.6 DCEmgmt/Event Manager

DCEmgmt/Event Manager provides integration between DCE and TME 10 Enterprise Console and TME 10 Distributed Monitoring, thus allowing you to monitor availability of DCE applications and services on any system in the network.

DCEmgmt/Event Manager contains a TEC event adapter that forwards DCE events to TME 10 Enterprise Console and a number of predefined TME 10 Distributed Monitoring monitors to monitor components of DCE.

This allows the administrator to control availability of DCE completely from the TME 10 desktop.

We show examples for this in 3.9.3, “DCEmgmt/Event Manager” on page 169.

3.7 DCEmgmt/Cell Manager

DCEmgmt/Cell Manager integrates DCE resources in the TME 10 desktop, thus allowing you to administer DCE completely from the familiar TME 10 environment. Unlike modules that are integrating on the Premier level and that basically use predefined integration icons to link to the TME 10 desktop, DCEmgmt/Cell Manager actually adds new icons, menus, etc. to the TME 10 desktop that allow you to control DCE as an integral part of the TME 10 environment.

3.8 DCEmgmt/DFS Manager

DCEmgmt/DFS Manager provides an elegant way to administer all components of Distributed File System (DFS) from the TME 10 desktop. This includes management of the various DFS parameters, such as aggregates, cache managers and file sets.

3.9 DCEmgmt Examples

We do not explain how to plan the installation of DCEmgmt or the installation of the DCEmgmt product in our environment, as this is documented in detail in the redbook *Tivoli TME 10 Management for DCE and DFS*, SG24-2027. (Available at a later date.) Refer to this redbook for complete coverage of supported platforms, installation prerequisites, etc.

We just show examples to give you a feeling of what DCEmgmt does.

3.9.1 Integration of DCE Management into the TME 10 Desktop

After installation of DCEmgmt you will have new resource types available in your TMR, one of which is the DCEmgmt cell object.

The starting point to integrating DCE resources is to create such a cell object on the TME 10 desktop as shown in the following figure.

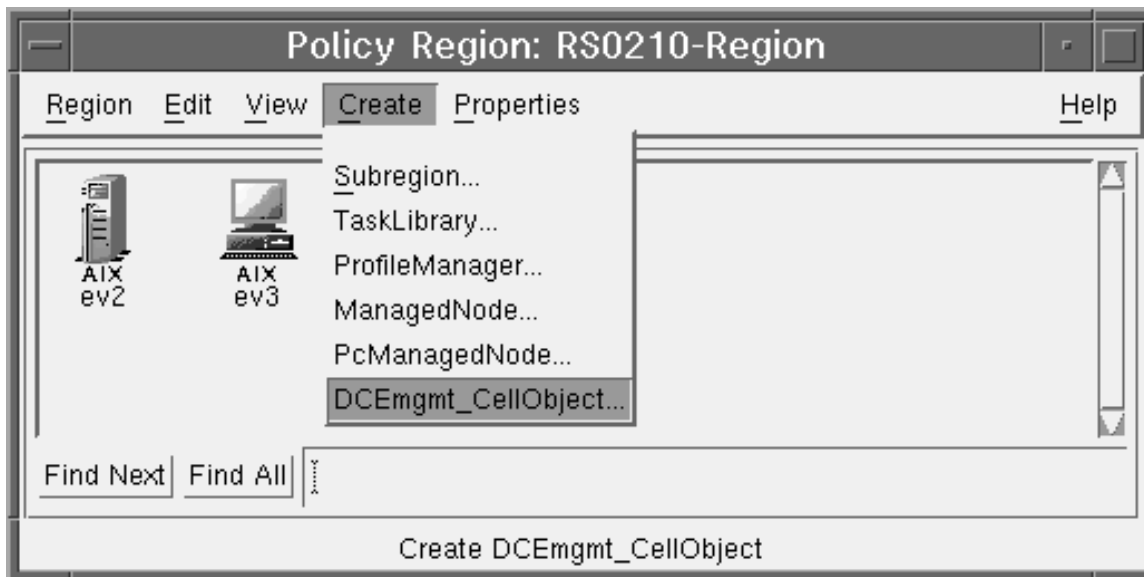


Figure 170. Creating a DCEmgmt Cell Object

In the Policy Region window, select **Create** from the menu bar and then select **DCEmgmt_CellObject...** from the pull-down menu.

Note

You should notice that in order for the DCEmgmt_CellObject option to show up on the pull-down menu, DCEmgmt_CellObject needs to be added as a managed resource in the TMR. Again, we do not explain the specific steps in detail here, but just show a few samples.

A window will appear asking you to provide information of the cell to manage, namely the name of the DCE cell, the cell administrator password and a label for the cell object.

Once this information is entered, the cell object is created.

The following figure shows an example where the cell object label is DCE_Cell.

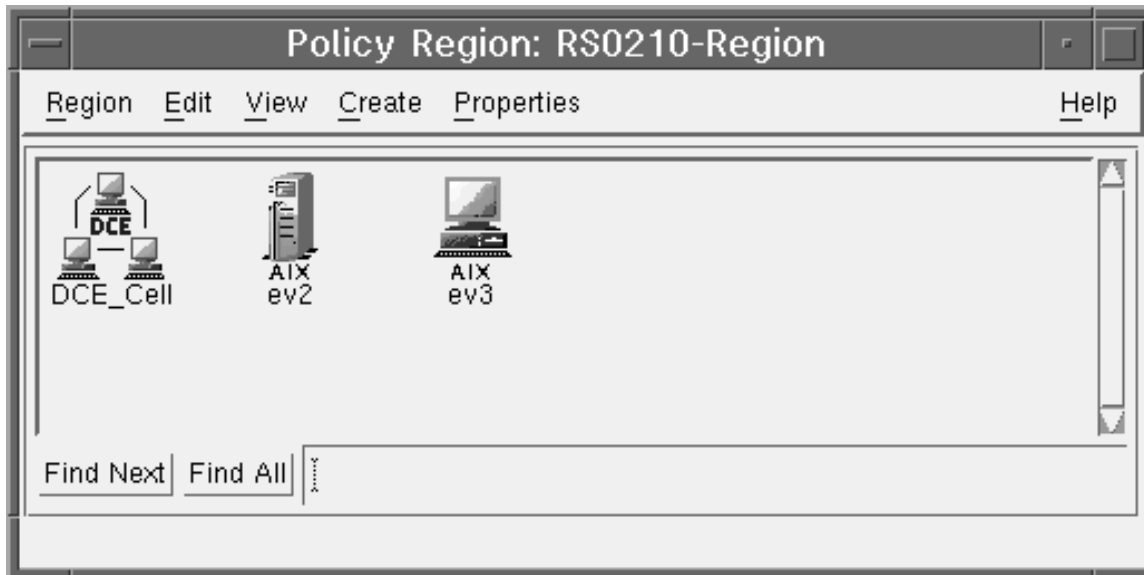


Figure 171. DCE Cell Object on the TME 10 Desktop

The cell object is the entry point to managing DCE from the TME 10 desktop.

Double-clicking on the **DCE_Cell** icon will open the Nodes and Services of Cell window as shown in the following figure.

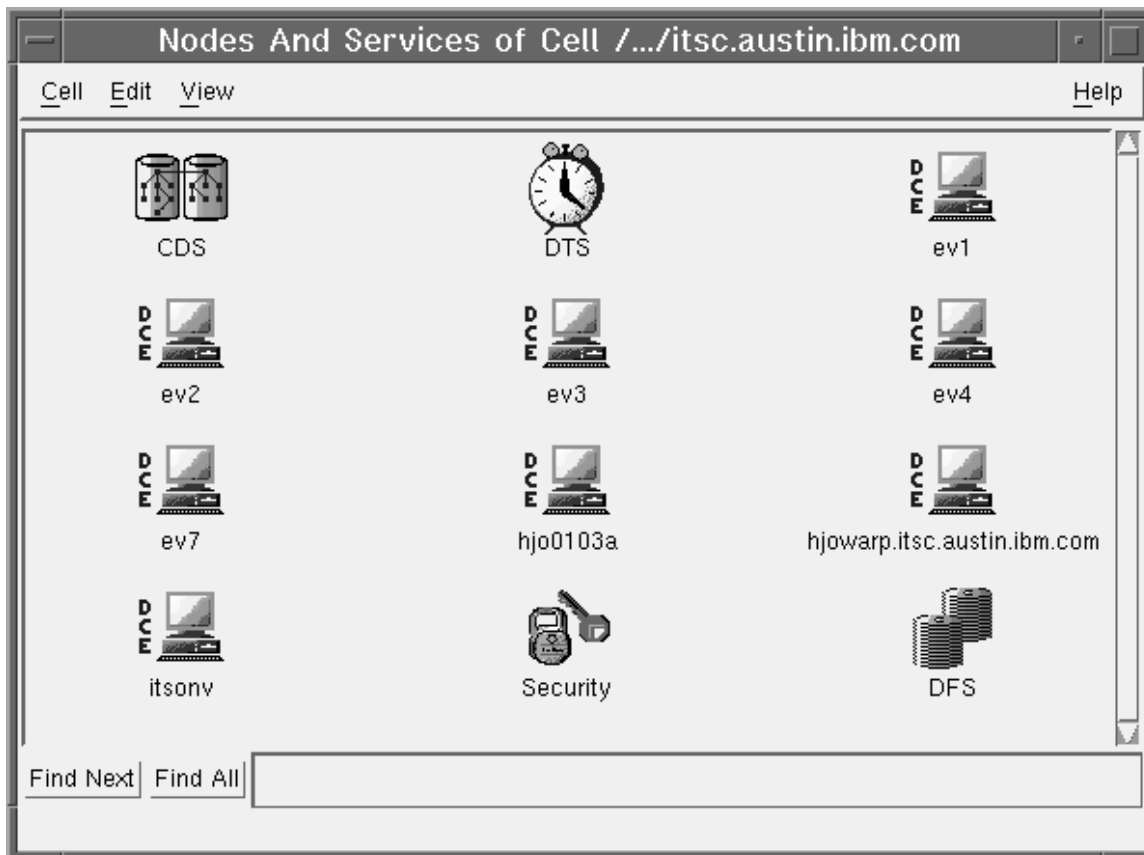


Figure 172. DCE Resources on the TME 10 Desktop

In the above figure you can see a number of DCE resources in their graphical representation on the TME 10 desktop, such as DCE nodes, the cell directory service, etc. DCEmgmt can auto-detect resources in the DCE cell and include them in this view.

The resources in the window can be used again as starting points of managing the different DCE components, such as the DCE nodes, the cell directory service, time service, etc.

3.9.2 DCEmgmt/Security Manager Examples

The DCEmgmt/Security Manager extends TME 10 User Administration in order to manage DCE users and DCE user groups.

As a result, DCE user groups can be managed by means of TME 10 User Administration (extended) profiles.

These profiles can be created, like any other profile, within a profile manager.

The following figure shows an example DCE group profile.

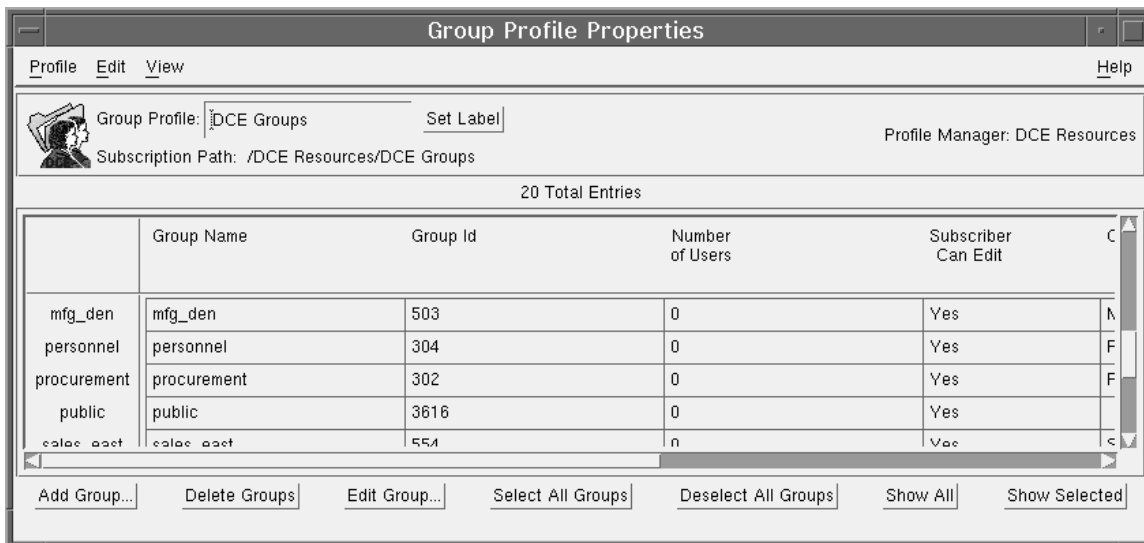


Figure 173. Example TME 10 User Administration Profile for DCE Groups

You can now manage DCE users by means of TME 10 User Administration, for example, populate profiles with existing groups, change the attributes of a group profile and then implement the change by just distributing the profile.

For example, double-clicking on any group in the above window will open the Edit Group Profile window as shown in the following figure.

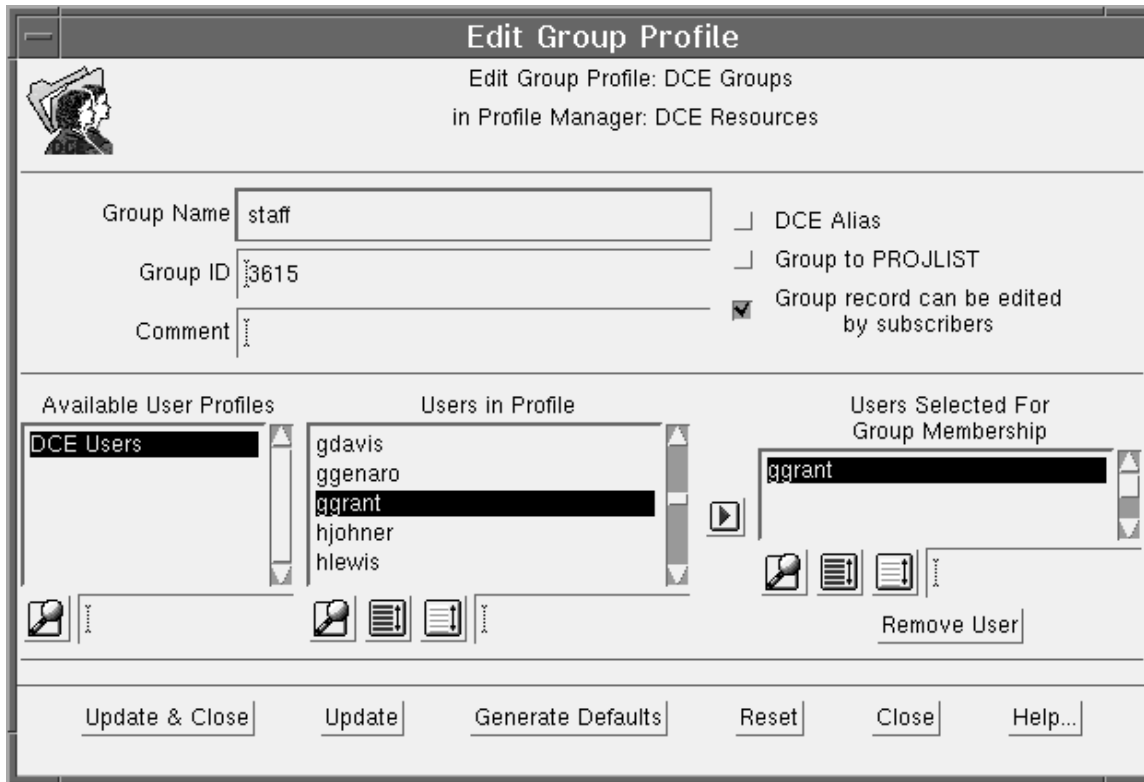


Figure 174. Edit Group Profile Window

You can see that the main attribute in this profile is the users actually being part of that DCE group or not. To add a user to the DCE group, you can just move the user name from the Users in Profile section to the Users Selected for Group Membership section. Saving the modified profile and then distributing the profile will then make the modifications necessary in the DCE registry.

An example DCE user profile is shown in the following figure.

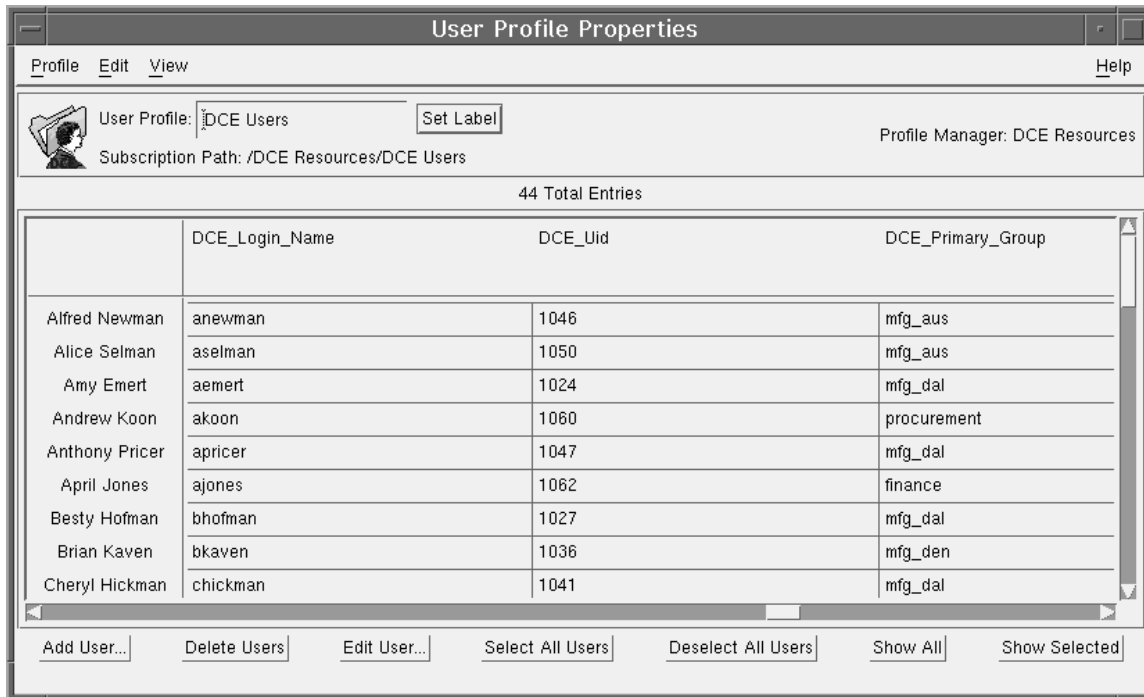
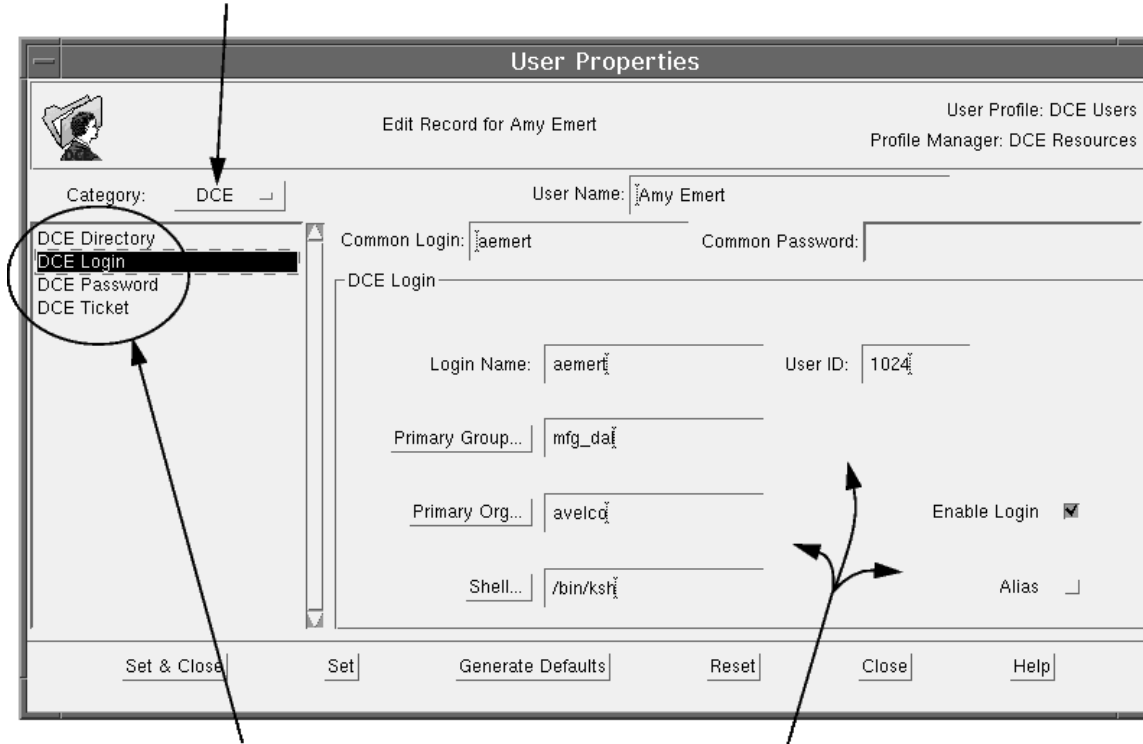


Figure 175. DCE User Profile

You can see the DCE users contained in the profile and their attributes, such as DCE_Login_Name and DCE_Uid.

To modify attributes for a user, you can just double-click on any user line in the profile which will open the User Properties window as shown in the following figure.

Select a **Category** (or Platform) first



Then, select a **Sub-Category** to view/edit the attributes shown on the right

Figure 176. User Properties Window

You can see the different attributes of the DCE user. Again, to modify the attributes, you can just modify and then distribute the TME 10 profile.

3.9.3 DCEmgmt/Event Manager

The DCEmgmt/Event Manager component provides integration between DCE and TME 10 Distributed Monitoring and TME 10 Enterprise Console.

For integration with TME 10 Distributed Monitoring, DCEmgmt/Event Manager provides custom monitors to ensure availability of the DCE system.

The following figure shows an example TME 10 Distributed Monitoring profile with DCE monitors.

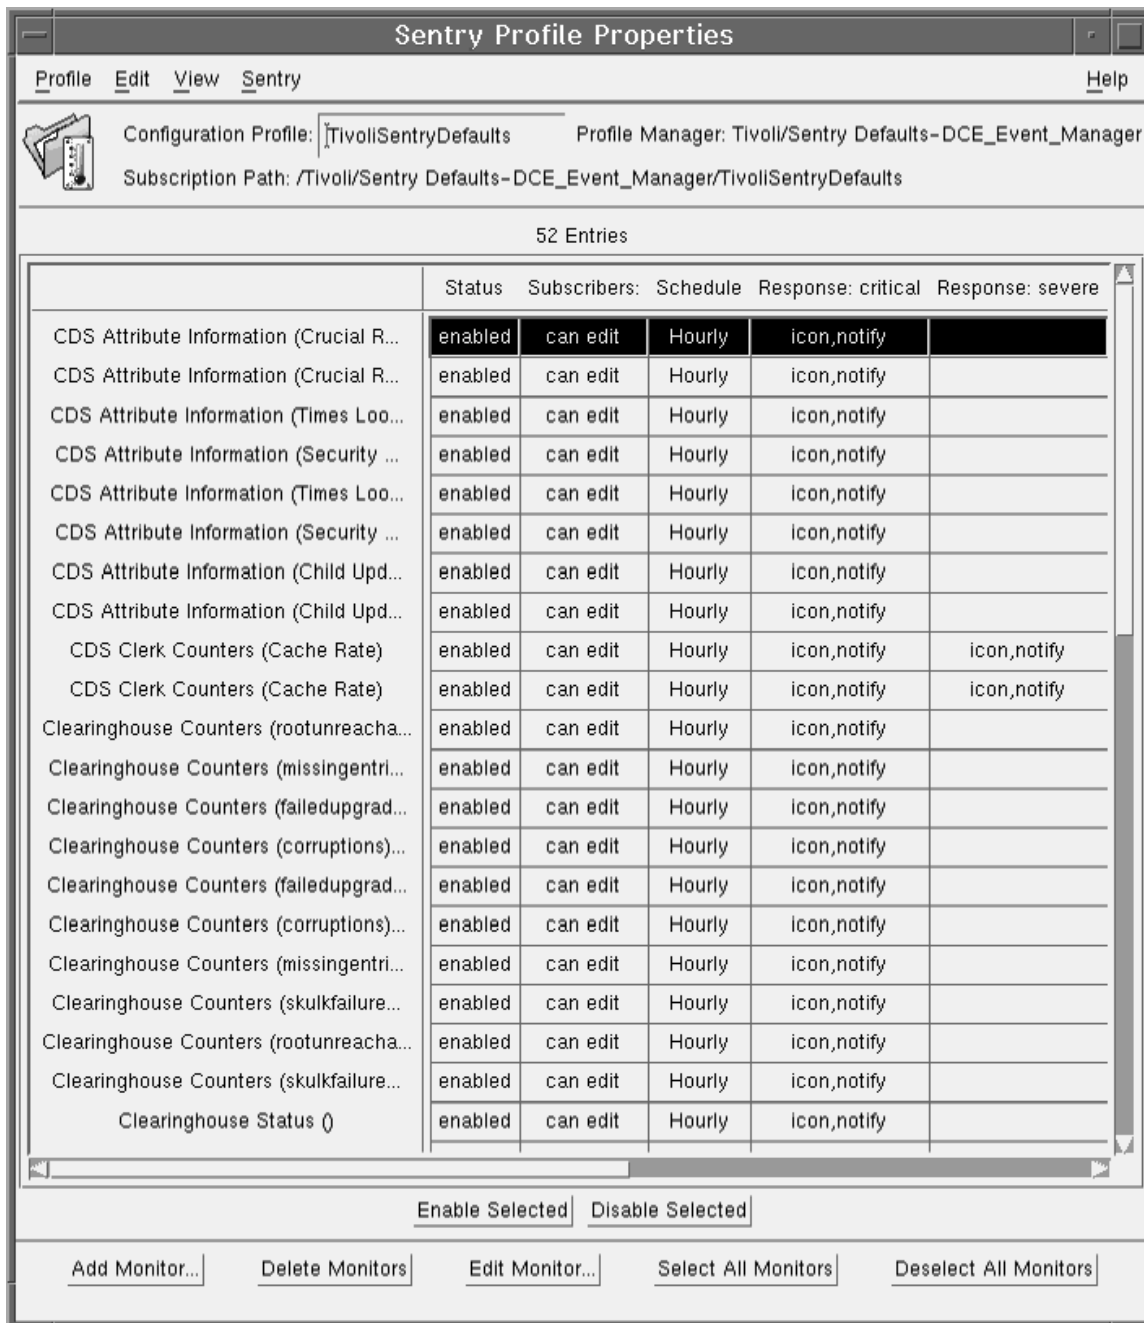


Figure 177. TME 10 Distributed Monitoring Profile for DCE

The handling of the TME 10 Distributed Monitoring profile, such as editing, distributing, etc. is exactly the same as for BEST/1+ for TME 10 and EcoTOOLS Plus Module for TME 10. Refer to Chapter 2, "BEST/1 from BGS Systems Inc." on page 13 and Chapter 4, "EcoTOOLS from Compuware Corporation" on page 173 for detailed explanations on how to handle TME 10 Distributed Monitoring profile.

The TME 10 Distributed Monitoring profiles supplied with DCEmgmt also provide integration with TEC, that is events from the TME 10 Distributed Monitoring profiles can be forwarded to TEC. A BAROC file is provided, so that events generated by the TME 10 Distributed Monitoring profiles for DCE can be classified upon arrival at the TEC.

The asynchronous monitoring of DCE by TME 10 is performed using a TEC adapter for DCE that is supplied by DCEmgmt

The following figure shows the architecture of the TEC adapter for DCE.

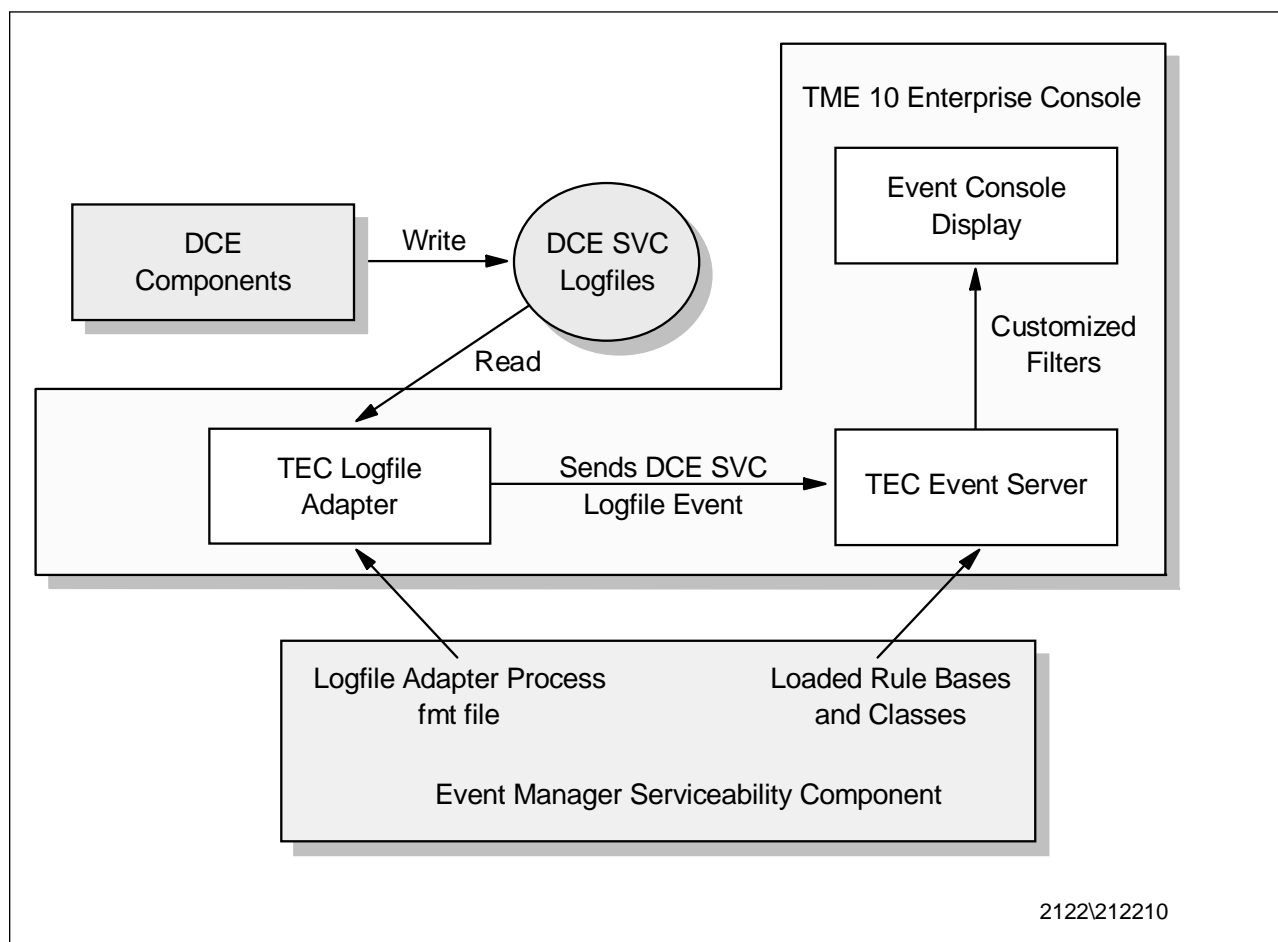


Figure 178. TEC Adapter for DCE

Like the TEC adapter for EcoTOOLS, described in Chapter 4, “EcoTOOLS from Compuware Corporation” on page 173 the TEC adapter for DCE uses the standard TEC Logfile Adapter. DCEmgmt provides the appropriate .fmt files to recognize and match DCE messages in the corresponding logfiles.

The TEC adapter monitors the logfiles that the DCE components that are to be monitored write messages to. If a message generates an event, this event is forwarded to the TEC server.

DCEmgmt/Event Manager provides BAROC files, so that the messages from the TEC adapter for DCE can be recognized and classified at the TEC server. At the TEC server the event is displayed at an event console according to the defined event filters and automatic actions or TEC rules can be defined to automate event processing.

Again, you can refer to Chapter 2, “BEST/1 from BGS Systems Inc.” on page 13 and Chapter 4, “EcoTOOLS from Compuware Corporation” on page 173 for detailed explanations on how to set up TEC, rule bases, filters, etc.

3.10 How Does DCEmgmt Complement TME 10?

DCEmgmt is somewhat different from the integration modules described in Chapter 2, "BEST/1 from BGS Systems Inc." on page 13 and Chapter 4, "EcoTOOLS from Compuware Corporation" on page 173.

While EcoTOOLS Plus Module for TME 10 and BEST/1+ for TME 10 link other systems management applications, namely EcoTOOLS and BEST/1 to TME 10, DCEmgmt includes management of an infrastructure or *middle-ware* component in TME 10 management.

As a result the integration is slightly different. While DCEmgmt also provides a TEC adapter, monitors for ensuring DCE availability using TME 10 Distributed Monitoring, etc. it further provides deeper integration into the TME 10 Framework. For example, DCEmgmt creates its own resources in the Tivoli object database and on the TME 10 desktop.

From a module development perspective that also means that for a Pinnacle module the TME 10 Framework Toolkit is used, rather than the TME 10 Integration Toolkit.

Chapter 4. EcoTOOLS from Compuware Corporation

In this chapter we describe EcoTOOLS from Compuware Corporation and discuss how it can be combined with TME 10. EcoTOOLS is a tool for comprehensive management of networked applications, focused on providing nonstop availability management that is integrated into TME 10 using the EcoTOOLS Plus Module for TME 10 from Compuware Corporation. We show how to install and set up EcoTOOLS and then explain how to integrate EcoTOOLS with TME 10 by installing and configuring the EcoTOOLS Plus Module for TME 10. Once this is done, we show practical examples of how TME 10 can be used to manage EcoTOOLS.

Note

EcoTOOLS is a complex client/server application that has a large number of features. Hence, in this chapter we can give only an overview of some key components in order to illustrate the functions of EcoTOOLS. For a complete description of all features of EcoTOOLS refer to the appropriate product documentation or contact your Compuware Corporation representative.

4.1 About Compuware Corporation

Compuware Corporation is a leading provider of software products that help clients to efficiently develop, implement and support the applications that run their businesses.



Founded in 1973, Compuware Corporation today is one of the largest independent software companies in the world. Compuware Corporation distributes its products worldwide through offices in many countries.

More information about Compuware Corporation can be obtained under www.compuware.com.

4.2 About EcoTOOLS

EcoTOOLS is a software product that helps managing mission critical applications in a highly distributed, heterogeneous environment.

The key strength of EcoTOOLS is what Compuware calls *application component management*. Application components can be the application software itself, the servers that run the application, RDBMS servers on which the application relies or the operating system and networks on which the application runs.

EcoTOOLS helps users to assure availability and performance of an application by managing all components that are relevant for a mission critical application.

EcoTOOLS collects information about the components, measures system performance and availability and can take automatic corrective actions.

For example, EcoTOOLS can manage most popular RDBMS systems, including Sybase, Oracle, DB2, Informix and Microsoft SQL.

For these and other applications and components, EcoTOOLS provides real-time monitors that can be used to help ensure availability of client/server applications.

Performance information can be displayed in many different ways, including graphs, EcoTOOLS gauges and tables.

4.3 About EcoTOOLS and TME 10

EcoTOOLS is focused on component management, more specifically the management of components that are relevant for client/server applications. TME 10 is a *systems management framework* that provides a solid platform on which you can integrate component management applications of various kinds to create a scalable and efficient enterprise management solution.

In order to combine the strengths of EcoTOOLS and TME 10, EcoTOOLS Plus Module for TME 10 can be used to integrate EcoTOOLS strength in managing applications with TME 10's strength in enterprise systems management.

Compuware Corporation is a Premier level member of the 10/Plus Association. EcoTOOLS was developed using the Tivoli Integration Toolkit that can be used to rapidly integrate existing systems management applications into TME 10.

Compuware Corporation is also a member of Tivoli's Performance Management Working Group.

Similar to other 10/Plus modules, EcoTOOLS Plus Module for TME 10 provides integration with the following TME 10 functions:

- TME 10 Framework
- TME 10 Software Distribution
- TME 10 Distributed Monitoring
- TME 10 Enterprise Console
- TME 10 Tasks and Jobs
- Application launch from the TME 10 desktop

4.4 Overview and Objective

The main objective of this chapter is to show the integration between the TME 10 Framework (and associated TME 10 products) and the EcoTOOLS product from Compuware. This section describes the environment used to test this integration and identifies the steps taken during the testing.

The tasks carried out during the testing include:

1. Installation of TME 10 Framework and applications on each test system
2. Installation of EcoTOOLS on the TMR server

3. Installation of EcoTOOLS agents on remote hosts
4. Performing basic EcoTOOLS tasks to illustrate the function of EcoTOOLS
5. Installation of the EcoTOOLS Plus Module for TME 10
6. Examples of using the EcoTOOLS Plus Module for TME 10

All of the above steps are described in detail in this chapter.

We start with giving an overview of the EcoTOOLS architecture and then perform each of the steps mentioned above.

4.5 Overview of EcoTOOLS Architecture

EcoTOOLS is a performance monitoring tool for UNIX and Windows NT systems including RDBMSs running on those supported platforms (refer to Table 3 on page 177).

The main focus of EcoTOOLS is on the availability side of performance management, that is measuring the performance of system components in order to insure availability and performance of these components.

Also, EcoTOOLS has sophisticated procedures to monitor specific client/server components, such as RDBMS servers.

The EcoTOOLS application consists of the following components:

- License station
- Management station
- Agent station

Each of the above components can have multiple instances; the obvious example of this is that the agent code runs on every machine that you want to monitor. The diagram below shows the basic configuration used during the writing of this redbook.

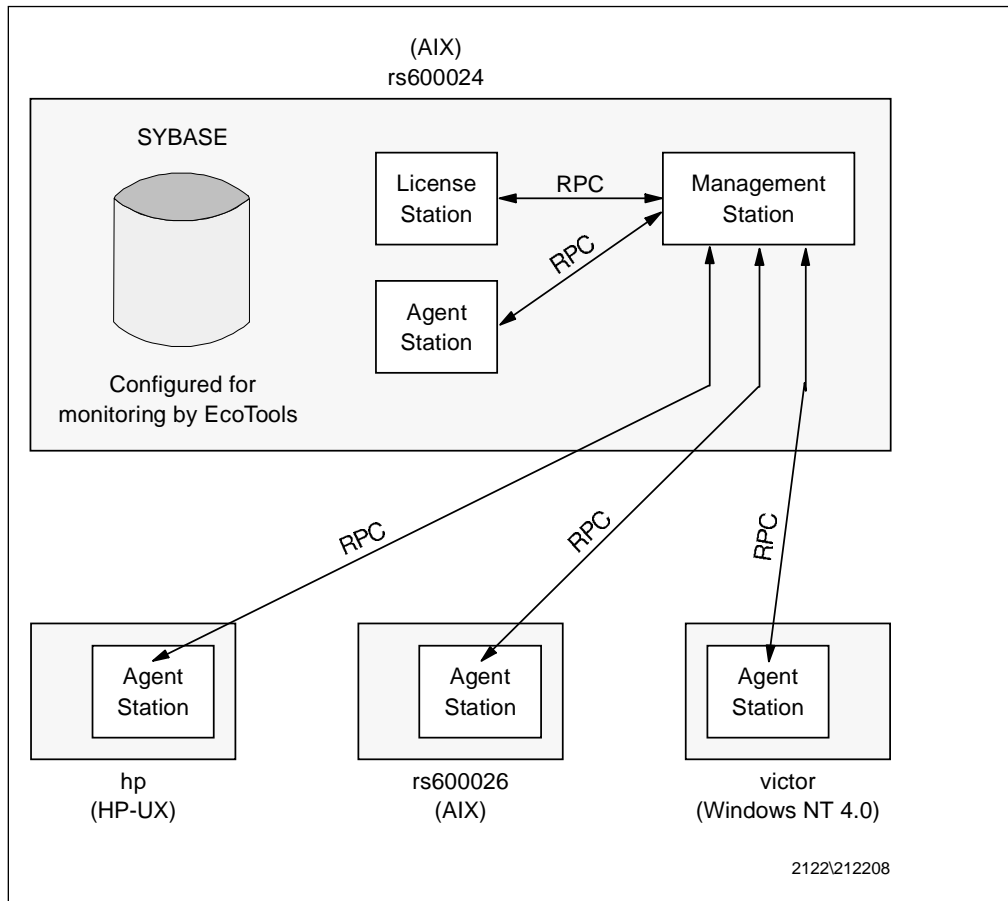


Figure 179. EcoTOOLS Components on Test Systems

The diagram in Figure 179 shows what components of EcoTOOLS are installed on which systems for the purposes of testing. The AIX server rs600024 has all three EcoTOOLS components installed, and all other systems have only the Agent Station code installed. The diagram also illustrates the communication channels between the components using remote procedure calls (RPC), with the Management Station acting as the hub of the EcoTOOLS environment.

The Management Station provides a graphical interface to the functions of EcoTOOLS, enabling a system administrator to run monitoring tasks on devices configured for EcoTOOLS. It also communicates with the license server to ensure compliance with EcoTOOLS licensing requirements.

4.6 Hardware and Software Prerequisites

This section lists the minimum requirements for installation of EcoTOOLS components, and what components can run on each platform or platform version.

4.6.1 General Requirements

In addition to the disk space requirements described in 4.6.2, "Installation Requirements and Supported Databases" on page 177 the following minimum hardware and communications requirements apply to nodes that you want to install the indicated EcoTOOLS modules onto. All figures are in MB unless otherwise specified.

Component	Data Storage	Swap Space	Memory	Other
Management Station	24	64	32	TCP/IP and RPC communications, X11R5 run-time environment
License Station	N/A	3	N/A	TCP/IP and RPC communications
UNIX Agent Station	24	64	32	TCP/IP and RPC communications, X11 run-time environment, 100 KB in /tmp
Windows NT Agent Station	N/A	50	32	Pentium class processor

Refer to the appropriate EcoTOOLS documentation for any additional prerequisites.

4.6.2 Installation Requirements and Supported Databases

Table 3 shows the databases supported by each platform and the disk space requirements for installation of EcoTOOLS. The disk space requirements do not include swap space, which is covered in Table 2. All figures shown are in MB unless otherwise specified.

Platform	Operating Systems	Supported Databases	Disk Space Requirements			
			Management Station	License Station	Agent Install and Package	Agent Package Only
IBM	AIX 4.1, 4.2, 4.3	Oracle, Sybase, Informix, DB2	29	1	19	19
DEC	Digital UNIX 3.2G, 4.0A-D	Oracle, Sybase, Informix	18	1	14	7
HP	HP-UX 10.10-10.20, 11.0	Oracle, Sybase, Informix, DB2, CA-Open Ingres, SAP R/3	25	1	37	18
AT&T (NCR)	GIS 3.0, 3.01	Oracle, Sybase, Informix	N/A	N/A	24	12
Sequent	DYNIX / ptx 4.2	Oracle, Sybase, Informix	N/A	N/A	24	24

Table 3 (Page 2 of 2). EcoTOOLS Installation Prerequisites

Platform	Operating Systems	Supported Databases	Disk Space Requirements			
			Management Station	License Station	Agent Install and Package	Agent Package Only
SGI	IRIX 6.2 (32 bit mode + 64 bit environment)	Oracle, Sybase, Informix	N/A	N/A	31	16
Sun SPARC	SunOS 4.1.3, 4.1.4.	Oracle, Sybase, Informix	N/A	N/A	15	11
Sun SPARC	Solaris 2.4, 2.5, 2.5.1, 2.6	Oracle, Sybase, Informix, CA-Open Ingres	21	1	16	8
Microsoft (Intel Only)	Windows NT 3.51, 4.0	Microsoft SQL Server 6.x, Informix	N/A	N/A	15	N/A
-	SCO UNIXWARE 2.1.1, 2.1.2	N/A	N/A	N/A	7	4

You must add together the agent package file sizes if you are distributing multiple platform packages. For example, you may want to transfer agent software from your Management Station to AIX, HP and Solaris machines, in which case the disk space requirements for storing the agent packages on the Management Station are: 19+18+8. This means that you must allow a total of 45 MB for all agent packages to be stored on the management station prior to transfer.

4.6.3 Software Requirements

In addition to the operating system requirements, there are no specific additional software requirements for EcoTOOLS.

In order to use EcoTOOLS Plus Module for TME 10 you have to integrate EcoTOOLS with TME 10 as shown in this chapter; however, you must have Release 3.6 of EcoTOOLS as the EcoTOOLS Plus Module for TME 10 is included with this release. EcoTOOLS Plus Module for TME 10 at the moment runs on the following operating systems:

- Sun Solaris 2.5.1
- IBM AIX 4.1
- HP-UX 10.20

4.7 Installing EcoTOOLS

This section details the initial installation of EcoTOOLS on an IBM RS/6000 running AIX. The installation media is a locally mounted CD-ROM on the host rs600024.

The list below identifies the major tasks involved in the installation and the appropriate order. More detail on these tasks is given further on in this section.

4.7.1 Installation Overview

The following sections describe each of these steps in detail. The installation of EcoTOOLS must be performed as the UNIX root user.

1. Mount the EcoTOOLS CD-ROM.
2. Create the installation directory (ecohome) on the host.
3. Change to the installation (ecohome) directory.
4. Run the installation script.
5. Pre-installation check (not mandatory).
6. Configure the EcoTOOLS installation.
7. Define the license key.
8. Unload the installation files.
9. Install EcoTOOLS on the management station.
10. Propagate software to agent stations.
11. Start the license server.

4.7.2 Installation Details

This section gives details of the menu selections made and system messages displayed during the installation of EcoTOOLS. The points below explain the parameters used for the test installation that is described in this section:

- The CD-ROM mount point is /cdrom.
- The CD-ROM is mounted on host rs600024.
- The EcoTOOLS installation target directory is /usr/eco.
- The target machines for the installation are: rs600024, rs600026 and hp.
- During installation, default menu options are shown in square brackets []. Press Enter to accept default options.

Note

The tag name and contract expiration data fields do not offer valid defaults.

4.7.2.1 Initiating the Installation

To initiate the installation process, follow the steps below:

1. Change to the installation directory using the command `cd /usr/eco` before running the installation script.

Note

We recommend you have a dedicated file system with a mount point that is identical to the EcoTOOLS installation directory, in our case /usr/eco.

2. Ensure that the CD-ROM is mounted and then type the following command to execute the installation script:

```
/cdrom/INSTALL/UNIX/install.cd
```

3. Follow the instructions below for each menu.

4.7.2.2 Setting the EcoTOOLS Home Directory

```
-----  
| Compuware EcoTOOLS CDRom Installation Utility |  
-----  
  
You must be in the EcoTOOLS Home Directory to install,  
apply a patch, or perform other maintenance functions.  
  
Your current directory is [/usr/eco]  
  
Is this the correct EcoTOOLS Home Directory ? (y/n) y
```

Figure 180. Setting EcoTOOLS Installation Directory

Type n and press Enter and then type the correct path to the directory that you will install EcoTOOLS into if the default is not correct; otherwise press Enter to accept the displayed default directory. The following messages are displayed once the directory settings are accepted.

```
Unloading EcoTOOLS installation module..
```

```
/usr/eco/log/install_ecotools.log
```

After a short delay, the Main Menu appears.

4.7.2.3 Configuring the Installation

The Main Menu initially only has two options. This section covers the second option on the menu, namely configuring the EcoTOOLS Installation.


```
-----  
|           Main Menu           |  
-----  
Please select one of the following options:  
  
  1) Pre Installation Check  
  2) Configure EcoTOOLS Installation  
  3) Unload EcoTOOLS from CDRom  
  4) Install EcoTOOLS  
  5) Maintenance  
  
  Q) Quit EcoTOOLS Installation.  
  
---> [2] 2
```

Figure 181. Main Menu - EcoTOOLS Installation

Main Menu Note

Options 3, 4 and 5 do not appear on the Main Menu until you have completed Configure EcoTOOLS Installation.

The Main Menu allows you to continue with the installation or carry out a pre-installation check. This check tests the systems on which you want to install EcoTOOLS and assesses whether the base criteria for an installation to complete successfully are satisfied. Use this option if you are unsure of the status of the target hosts.

Select option 2 if you do not want to carry out a pre-installation check and want to continue with the installation.

4.7.2.4 Configuration Menu

The installation options are configured prior to starting the installation, and are defined using the Configuration Menu shown in Figure 182. The installation of EcoTOOLS cannot proceed until all necessary configuration has been completed.

```
-----  
|           Configuration Menu           |  
-----  
Please select one of the following options:  
  
  1) Define Complete System Configuration  
  2) Setup Default Configuration  
  3) Define Partial System Configuration  
  4) Display Current Configuration  
  
  Q) Quit. (Return to previous menu)  
  
---> [1] 1
```

Figure 182. EcoTOOLS Configuration Menu

Select option 1 from the Configuration Menu to define the system configuration of your target systems. Select option 4 to display the current configuration. (An

example output is shown in Figure 184 on page 184.) Press Enter when you have made the menu selection.

Upon selecting option 1 to define the complete system configuration, you are prompted with several options screens:

- CD-ROM Path

Where prompted, enter the path to the EcoTOOLS installation CD-ROM and press Enter. Press Enter without entering a path to accept the default.

- Management Station Configuration

Enter the hostname of the station that will be used to manage EcoTOOLS. During our test installation, the host rs600024 is defined as the Management Station. Press Enter without entering a hostname to accept the default.

- License Station Configuration

Enter the hostname of the station that will be the EcoTOOLS license server. During the test installation, the host rs600024 is defined as the License Station. Press Enter without entering a hostname to accept the default.

- Agent Station Configuration

Agent stations are servers and clients that run the EcoTOOLS agent so that they can act as targets for monitoring. Enter the names of *all* systems that you want to install EcoTOOLS onto at this point. Separate each hostname by a space or comma. In the test installation case, the hosts rs600024 and rs600026 are entered at this point.

Installation Note

During later stages of the installation, the installation script will attempt to copy installation files to these hosts. If any of the listed hosts cannot be contacted, then installation errors will occur. Additional hosts can be installed at any time so it is not necessary to install all hosts during the initial installation of EcoTOOLS.

4.7.2.5 System Type Selection

Define the operating system type for each hostname entered in the Agent Station Configuration option of the installation. With reference to the menu shown in Figure 183 on page 183 select the appropriate operating system type for that host and press Enter.

```
For system 'rs600024' please answer the following questions:

System Types:

  1. Sun Solaris
  2. Sun SunOS
  3. HP 9000 HP-UX
  4. AIX
  5. Windows NT
  6. DEC UNIX
  7. NCR UNIX SVR4
  8. SCO UNIXWARE
  9. Sequent DYNIX/ptx 4.x
 10. SGI IRIX

What type of computer is rs600024? [4]
```

Figure 183. System Type Selection Menu

After selecting the operating system for the host, you are prompted for the following information:

- Installation path of the EcoTOOLS agent
Enter the path on the host where the agent should be installed. Press Enter when complete, or do not type a path and press Enter to accept the default.
- UNIX login name for the EcoTOOLS daemon
Enter an existing UNIX login name that the EcoTOOLS daemon can run as. The daemon cannot use root. We use bin in our environment. Press Enter when finished.

When the above configuration (OS type, path and user) has been configured for all target agent stations you are returned to the Configuration Menu.

4.7.2.6 Current Configuration Information

Check the current configuration from the Configuration Menu (Figure 182 on page 181) by selecting option 4 once you have completed entering the installation configuration information. An example screen is shown in Figure 184 on page 184.

```

The current configuration of EcoTOOLS as of the last update
on ----- Wed Nov 12 11:48:00 EST 1997 -----

CDROM Drive Subsystem (Configured)
  CDROM Drive Name: /cdrom

Management Station (Configured)
  Management System: rs600024
  Management system directory: /usr/eco

License Station (Configured)
  License server system: rs600024
  License system directory: /usr/eco
  License system hostid: 0x0x91868a1

Agent Station (Configured)
Defined agent systems: (2 Systems)
Host      Type      User ID  Pathname
rs600024  aix6000  bin     /usr/eco
rs600026  aix6000  bin     /usr/eco

End of configuration list

```

Figure 184. Current Configuration Screen

Review the configuration details and repeat the configuration process if necessary.

Select q and return to the main menu when configuration is complete.

4.7.3 Unloading EcoTOOLS

The installation files must be unloaded from the CD-ROM before the actual installation can take place. The Unload EcoTOOLS option appears on the Main Menu when installation configuration is complete. Select this option from the Main Menu to begin unloading the files. The messages shown below appear as the unloading commences, along with a copyright notice.

```

Unloading software for: EcoTOOLS Management Station version 3.6.0.0.0
EcoTOOLS License Station Version 3.6.0.0.0 EcoTOOLS Agent Station version
3.6.0.0.0

```

When prompted type y and press Enter to continue and you will see the following information:

```

Unloading EcoTOOLS Management software
version 3.6.0.0.0 into /usr/eco.
A list of extracted files can be found in /usr/eco/etc/installation.list.

```

A progress status is shown to indicate how much data has been unloaded from the installation CD-ROM, as shown below:

```

7752 Kbytes loaded. 27 percent.

```

Further installation status messages are displayed, until all files are copied, at which time the completion message appears:

```

Successfully loaded requested agent software from CDRom.

```

4.7.3.1 Copying Files to Remote Hosts

Once the software has been copied to the installation target, the copy to any other Agent Stations commences. Type `y` and press Enter when the message in Figure 185 is displayed to install to the indicated host, or type `a` and press Enter to commence copying to all hosts entered as Agent Stations.

```
Ready to attempt transfer of agent software to system 'rs600026'  
Do you want to attempt this transfer now?  
Enter (y)es (n)o (q)uit or (a)ll to send to all agents: [y]
```

Figure 185. Transfer Agent Software Confirmation

The message shown below appears upon completion of the transfer:

```
Successfully copied agent station software to system 'rs600026'
```

When all code transfers are complete, the Main Menu appears.

Note

In this example, the installation of EcoTOOLS on the agent stations is performed together with the installation of the EcoTOOLS management station. You should notice that you can also install the agent stations later using EcoTOOLS Plus Module for TME 10 and TME 10 Software Distribution to perform this task as described in 4.12.1, “Deploying EcoTOOLS Using TME 10 Software Distribution” on page 227.

4.7.4 Installing EcoTOOLS

Up to this point, only the appropriate code has been copied to each host defined in the installation configuration. The last part of the installation process is to actually install EcoTOOLS.

Select option 4 to Install EcoTOOLS (Figure 181 on page 181) from the Main Menu when it re-appears after the installation configuration and unloading of EcoTOOLS.

The Install EcoTOOLS Menu appears showing all available installation options.

```
-----  
|           Install EcoTOOLS Menu           |  
-----  
Please select one of the following options:  
  
  1) Install EcoTOOLS on System  
  2) Setup ORACLE Databases for EcoTOOLS Monitoring  
  3) Setup Sybase SQL Servers for EcoTOOLS Monitoring  
  4) Setup Informix Servers for EcoTOOLS Monitoring  
  5) Setup Tuxedo for EcoTOOLS Monitoring  
  6) Setup DB2 for EcoTOOLS Monitoring  
  7) Setup OpenIngres for EcoTOOLS Monitoring  
  
  Q) Quit. (Return to previous menu)  
  
---> [1]
```

Figure 186. Install EcoTOOLS Menu

Select option 1 from the Install EcoTOOLS Menu to install on the current host. The installation script displays information regarding the operations that take place as part of the installation process.

This installation performs the following operations:

1. Creates /etc/eco/eco.conf, /etc/eco/msg.cat, /etc/eco/msg.cat.m and /etc/eco/msg.idx.
2. Extracts and verifies the Agent Station software.
3. Adds or replaces entries at the end of /etc/inetd.conf. The new entries are identified by a comment referencing EcoTOOLS.
4. Adds or replaces entries at the end of /etc/rpc. The new entries are identified by a comment referencing EcoTOOLS.
5. Adds or replaces entries at the end of the system startup file /etc/rc.tcpip so that EcoTOOLS daemons will start at each system boot. The new entries are identified by a comment referencing EcoTOOLS.

The installation process then extracts the Agent software.

For information regarding the configuration of Sybase SQL servers for EcoTOOLS monitoring, refer to 4.7.5, "Set Up Sybase SQL Servers for EcoTOOLS Monitoring" on page 188.

4.7.4.1 EcoTOOLS License Configuration

The installation program prompts for license information after the agent software has been extracted. Some of the responses that you enter in this stage of the installation are provided by Compuware when you purchase EcoTOOLS. The responses shown are for information only.

The message below appears requesting the destination of error messages. The default is the current console. Press Enter to accept this default or type a new destination and press Enter.

Normally `ecoad` uses `/dev/console` to place the standard error message. You may choose a text file to replace `/dev/console`. If you do, then the standard error message will be shown in this text file. Please choose the destination of `stderr` [`/dev/console`]?

The configuration prompts shown below correspond to nominating a license server for EcoTOOLS. Enter the information for each prompt and press Enter. See the *EcoTOOLS User Guide* for commands to obtain the Host ID of your target server.

Would you like to set/change primary license server (y|n)?[y]

Primary License Server Name: rs600024

Primary License Server Host ID(hexdecimal): 6023864c

Would you like to set/change secondary license server (y|n)?[n]

Would you like to set/change license entry (y|n)?[y]

Note

The following commands can be used to obtain the Server Host ID, depending on the platform:

- AIX: `/usr/bin/hostid`
- Solaris: `hostid`
- HP-UX: `/bin/uname -i`

The license tag is a descriptive name for the EcoTOOLS license. Type a name and press Enter to continue.

Please provide an alphanumeric name for your license as a tag for this entry. This Tag could be 8 characters or less in length, and must begin with a letter. Tag name? rs600024

Enter the licensing information provided by Compuware in the following options when requested. Example values are shown.

1. EcoTOOLS License Server Password?
2. Site License?[n] y
3. Demo License?[n] y
4. Number of Licenses? 1
5. Number of Database Users? 20
6. Number of Database Servers? 5
7. Number of Client Machines? 1
8. Number of Non-Database Servers? 5
9. Number of Workstations? 5
10. Serial Number(X-X-X)? 10-10-10000
11. License Expiration Date (e.g. 12-31-1996)? 12-05-1997

12. Support Contract Expiration Date (e.g. 12-31-1996)? 12-30-1999
13. EcoCHARGEBACK Option (y|n)?[n] n

Note

The installation password used in our environment is not shown. Refer to your EcoTOOLS package to obtain the password for your environment.

Check all of the information once entered. When you press Enter after entering the last option a prompt will appear asking whether you would like to change any of the options. Type either *y* or *n* depending on your requirements and then press Enter.

When prompted type *y* and press Enter to start the license server, which will take approximately 3 minutes to start.

Note

There should be a 30 second delay before any dots start printing on the screen. If dots start to print right away, the license server has failed.

Further installation status messages appear, and also some warnings may appear, such as:

WARNING: There are some `>> exit <<` lines in your `/etc/rc.tcpip` file. Please make sure that all entries for `eamon` in `/etc/rc.tcpip` appear BEFORE the `exit` entry

Check any listed files to ensure that the entries added by the EcoTOOLS installation come before any `exit` commands. Other installation status messages appear as the installation program attempts to start the agent, monitor and discovery daemons.

The EcoTOOLS database conversion utility prompt appears when all daemons are started, which you can select to run if you are upgrading from a previous version of EcoTOOLS. The installation program returns to the Main Menu once installation is complete. Type *q* and press Enter to exit the installation program.

4.7.5 Set Up Sybase SQL Servers for EcoTOOLS Monitoring

This section describes how EcoTOOLS is configured to monitor Sybase SQL servers and databases. This installation makes modifications to the EcoTOOLS Manager and the target servers and databases. The process described is specific to Sybase databases, however, it has similarities to the process used to configure monitoring of other database types, requiring just different configuration parameters to be entered. (Refer to the appropriate EcoTOOLS documentation for information on the parameters for other database types.)

This installation can be done as part of the initial installation of EcoTOOLS or installed later, if more convenient. In the example that is described in this section the installation was not done as part of the original EcoTOOLS installation. The installation process is similar regardless of when it is done.

4.7.5.1 Starting the Installation

The installation process can be initiated from any node that has had the EcoTOOLS code unloaded onto it (see 4.7.3, “Unloading EcoTOOLS” on page 184).

If this installation is being done as part of the main installation process, proceed directly to step 5.

1. Open a telnet or X Windows session on the target node and change to the `/usr/eco` directory.
2. Switch to the UNIX root user.
3. Run the command:

```
./install_ecotools
```
4. When prompted, confirm the location of the EcoTOOLS installation directory (`/usr/eco`).
5. Select option **4) Install EcoTOOLS** from the Main Menu when it appears.
6. Select option **3) Setup Sybase SQL Servers for EcoTOOLS Monitoring**.

4.7.5.2 Installing Sybase Monitoring

For EcoTOOLS to configure the database(s) and respective server(s) for monitoring, the following information is required. The installation process requests input of these values.

1. Path of the Sybase Interfaces file
2. Name of Remote Server
This is the database server name and *not* the host name of the UNIX server on which the database resides. For example, the default name for the database server is SYBASE.
3. Sybase system administrator (sa) user name and password
4. Sybase user ID
The installation process creates this user ID in all databases on the database server. EcoTOOLS uses this ID for monitoring purposes. The default ID is ECO_IRM, but you may change this to anything you like.
5. Sybase user ID password
Provide a password for the user ID defined in the previous point.
6. Sybase environment variable
This variable points to the Sybase home directory, for example, `/usr/local/sybase`.
7. Database server hostname
This value is the UNIX hostname of the database server. In our environment this is set to rs600024.
8. Sybase server process owner user ID
The user name on the database server that owns the Sybase processes.
9. SQL monitor server name
Used only if SQL monitor server is configured on the database server.

10. Size of EcoToolsDb

Enter the desired size (in megabytes) of the EcoTOOLS database which will be created on the Sybase server and is used to store temporary data used by the EcoTOOLS agents. The size must be at least 2 MB, which is the default.

11. Device to create EcoToolsDb

The Sybase device name where the installation process creates the EcoToolsDb tables.

4.7.5.3 Post Installation Tasks

After installing, run the discovery process for the host in EcoTOOLS so that the database and associated clients are inserted into the EcoTOOLS maps.

4.7.6 EcoTOOLS Agent Installation on Remote UNIX Nodes

The EcoTOOLS agent code is only copied to the remote nodes during the initial installation process. Follow the instructions in this section to install the agent on nodes that have had the code distributed to them already.

1. Change to the EcoTOOLS home directory, for example:

```
cd /usr/eco
```

2. Execute the following command:

```
./install_ecotools
```

3. Follow the installation instructions from 4.7.3.1, "Copying Files to Remote Hosts" on page 185.

Note

You should notice that the installation of remote UNIX hosts can also be performed using EcoTOOLS Plus Module for TME 10 and TME 10 Software Distribution as described in 4.12.1, "Deploying EcoTOOLS Using TME 10 Software Distribution" on page 227.

4.7.7 EcoTOOLS Agent Installation on Remote Windows NT Nodes

The installation of the EcoTOOLS agent code on a Windows NT server differs from the UNIX agent installation. This section describes the installation process in detail.

Log on locally to the NT server as administrator or equivalent and insert the EcoTOOLS CD-ROM into the CD drive. Open a Windows Explorer window or command prompt and run the program:

```
<path to CD-ROM>:\INSTALL\ntx86\ECONTX86.EXE
```

Click on the **Setup** button when the WinZip Self-Extractor (ECONTX86.EXE) window appears. This starts the Windows InstallShield installation process. The first window that requires user input is the Choose Destination Location window shown in Figure 187 on page 191.

4.7.7.1 Installation Location



Figure 187. Choose Destination Location Window

Click on the **Browse** button to select the installation directory of EcoTOOLS, and then click on the **Next>** button. To accept the default (as displayed) and continue the installation click only on the **Next>** button.

4.7.7.2 File Extraction

The installation program now copies the EcoTOOLS files to the target Windows NT machine. The Select Options window appears when copying is complete.

4.7.7.3 Select Options

Choose what parts of EcoTOOLS to install from the Select Options window (see Figure 188 on page 192). The components that can be selected are:

- Install EcoTOOLS Services
Installs only the standard EcoTOOLS agent for Windows NT.
- Setup MS SQL Server for EcoTOOLS Monitoring
This option configures the EcoTOOLS agent and Microsoft SQL Server for monitoring with EcoTOOLS.
- Setup Informix for EcoTOOLS Monitoring
Same as the previous option, but for Informix databases rather than SQL.



Figure 188. Choose Destination Location Window

Click on the box beside the options that you want to install so that a check mark appears in the box (as shown for the Install EcoTOOLS Services option only in Figure 188). Click on the **Next>** button when finished.

4.7.7.4 Configuring Services

Installation status messages appear while the installation completes. The EcoTOOLS Windows NT services require configuration when the installation is complete.

To configure the services, start the **Control Panel** and double-click the **Services** icon. The window shown in Figure 189 appears listing all services.

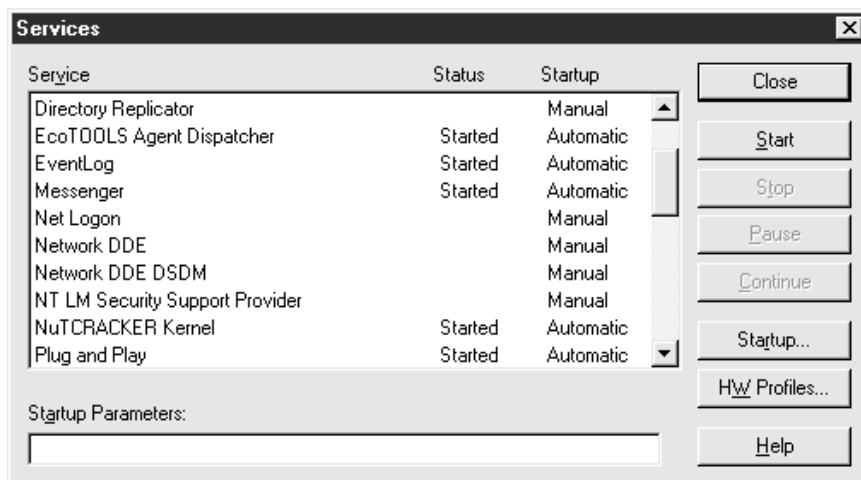


Figure 189. Services Window

Locate the EcoTOOLS Agent Dispatcher service (shown as second in list in the example). Select it using the left mouse button and then click on the **Startup** button. This action displays the window shown in Figure 190 on page 193. Also

shown in Figure 190 on page 193 is the other service created by the EcoTOOLS installation, NuTCRACKER Kernel.



Figure 190. Service Window for EcoTOOLS Agent Dispatcher

If not already set, click on the **This Account** option and enter the login name Administrator (or equivalent) as shown in Figure 190. Enter the correct password for the selected account in the Password and Confirm Password boxes. Click on the **OK** button when finished to return to the Services window.

4.7.7.5 Starting Services

You can start the EcoTOOLS Agent Dispatcher by clicking on the **Start** button in the Services window, but a reboot is recommended, at which point the agent will start automatically.

4.7.8 Starting the EcoTOOLS GUI

EcoTOOLS has a graphical user interface used to manage the functions of EcoTOOLS. Follow the instructions below to run the program once you have installed EcoTOOLS:

1. Ensure that the \$DISPLAY environment variable is set correctly.
2. Execute the command:

```
<ecohome>/bin/eco &
```

Where <ecohome> is the location of the EcoTOOLS installation, /usr/eco in our example. Refer to the appropriate EcoTOOLS documentation for information on options that can be used to run the EcoTOOLS GUI.

3. When prompted, enter a valid user name and password.

4.8 Functions of EcoTOOLS

This section of the document explains the basic functions of EcoTOOLS, which can be summarized as:

- Define domains

EcoTOOLS domains are logical groups of clients and servers.

- Adding nodes
- Create/edit scenarios

A scenario is a collection of monitoring agents and actions associated with them.

- Run scenarios
- Inventory

EcoTOOLS provides a hardware and OS inventory system.

- EcoPager

The first part of the functions of EcoTOOLS is the initial configuration required to set up EcoTOOLS to function in your environment.

When running EcoTOOLS you are prompted to enter a user name and a password. This must be a valid user name and password that exists on the license server. See the appropriate EcoTOOLS documentation for information on how to bypass the authentication process.

4.9 Defining Domains

After starting EcoTOOLS for the first time the EcoTOOLS main window appears as shown in Figure 191 on page 195. This window shows all available EcoTOOLS domains, but initially will not contain any icons as no domains will have been defined. In the example window shown a domain called ITSO.RAL has already been defined. This section explains how that domain was created.

Installation Note

There is an option during EcoTOOLS installation to create a startup domain. If this option is used during installation, then the main window will contain a domain icon when you first start EcoTOOLS.



Figure 191. EcoTOOLS 3.4.1: MainWindow

Click on the **Define** option of the Domain pull-down menu to create a new domain. This action will open a window the same as that shown in Figure 192 on page 196.




Figure 192. EcoTOOLS : Define Window

The EcoTOOLS: Define window is used to define the members of the EcoTOOLS domain. A domain in terms of EcoTOOLS is a group of nodes (including clients and servers) to be included in the EcoTOOLS discovery process. The discovery process checks the availability of the nodes in the map, and the client/server relationships between nodes. Discovery does not produce errors if the nodes included in discovery do not have the EcoTOOLS agents installed; the EcoTOOLS functions will just not function correctly. Select one of the three sources of domain member lists by clicking on one of the following buttons to the right of From:

- Map
List hosts currently defined in an EcoTOOLS map.
- NIS
List hosts from Network Information Service (NIS).
- Hosts file
List entries in the /etc/hosts file of the local host.

A list of available hosts appears in the Scope: section when you click on one of the From: buttons. You can also add hosts manually by typing a host name in the Add: box and clicking on the **Add** button to the right-hand side of the window.

Type a name for the domain in the Domain: box and select a class by clicking on one of the buttons to the right of Class:. The Class selection indicates what type of client/server relationships exist within the domain. EcoTOOLS currently supports the following systems (classes):

- Oracle

- Sybase
- Informix
- DB2
- OpenIngres
- Microsoft SQL Server

Each domain can only have one class associated with it.

Compuware recommends that Auto Discovery is configured to enable EcoTOOLS to check the availability of the nodes configured in the domain. Enter values in the Auto Discovery Interval section of the EcoTOOLS: Define window to configure auto discovery. If the values are all zero, auto discovery will not run.

Click **Apply** to run discovery on the nodes you have defined in the scope. This action creates the domain icon in the main window. Click **Cancel** to close the EcoTOOLS: Define window.

4.9.1 Computer Map Navigation

Double-click the left mouse button on any domain icon to display the EcoTOOLS: Computer Map of that domain. Most EcoTOOLS functions are accessed from the computer map window, either by the pull-down menus at the top of the window, or by selecting an object in the window and then clicking on the right mouse button. The menu that appears contains the following options that are explained in the subsequent sections. The options in the menu apply only to the currently selected node.

- Acknowledge alerts
- Show problems
- View scenarios running
- Stop scenarios running
- Run a scenario
- Edit scenarios (can apply to other nodes)
- Delete scenarios
- View hardware inventory
- Run discovery
- Process monitor
- Run commands

The pull-down menu is shown in Figure 193 on page 198. The currently selected node name (to which the menu functions apply) is shown at the top of the menu.

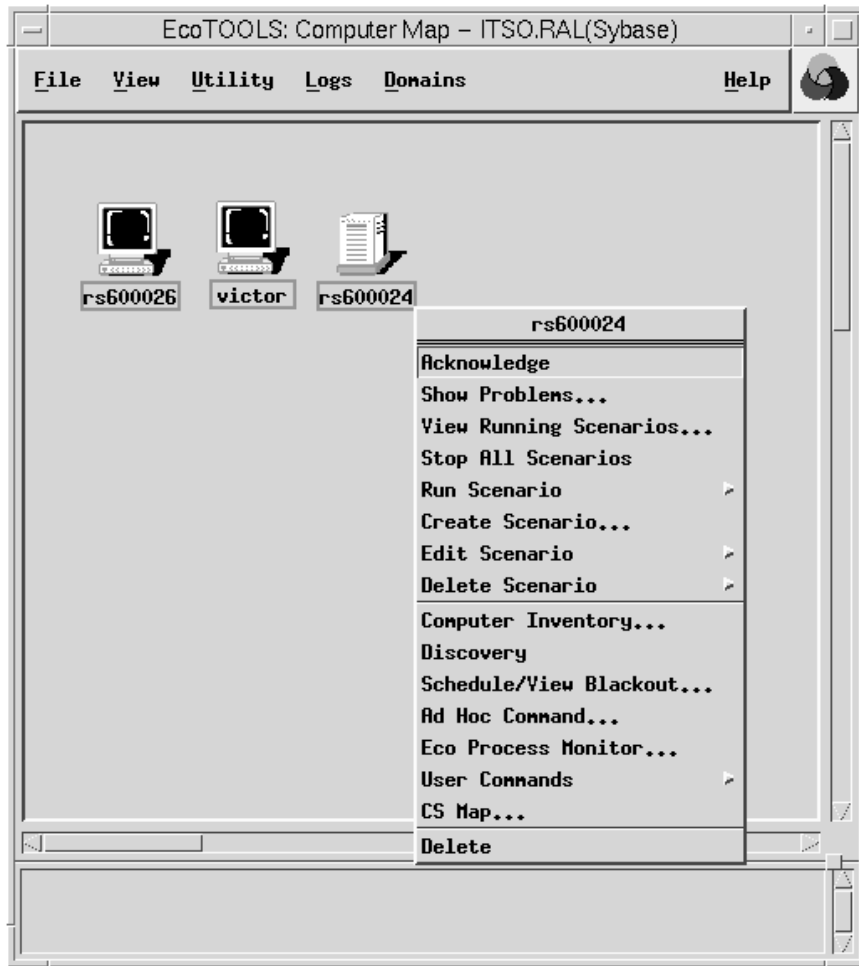


Figure 193. EcoTOOLS: Computer Map - ITSO.RAL(Sybase) Window




Figure 194. Client (left) and Server (right) Icons from Computer Map

The window in Figure 193 shows two client machines (rs600026 and victor) and a single server (rs600024). Also shown is the menu that appears when you click on a client or server with the right mouse button. Computers appear as clients by default. A computer will appear as a server only when it is running an RDBMS and EcoTOOLS is configured to monitor that database, at which time the discovery

process in EcoTOOLS will discover the client/server relationships between computers.

Double-click on the server icon to display the EcoTOOLS: Client- Server Map window for that server. An example of this window is shown in Figure 195, which shows the Sybase database named SYBASE and two clients of that database: rs600024 and VICTOR.

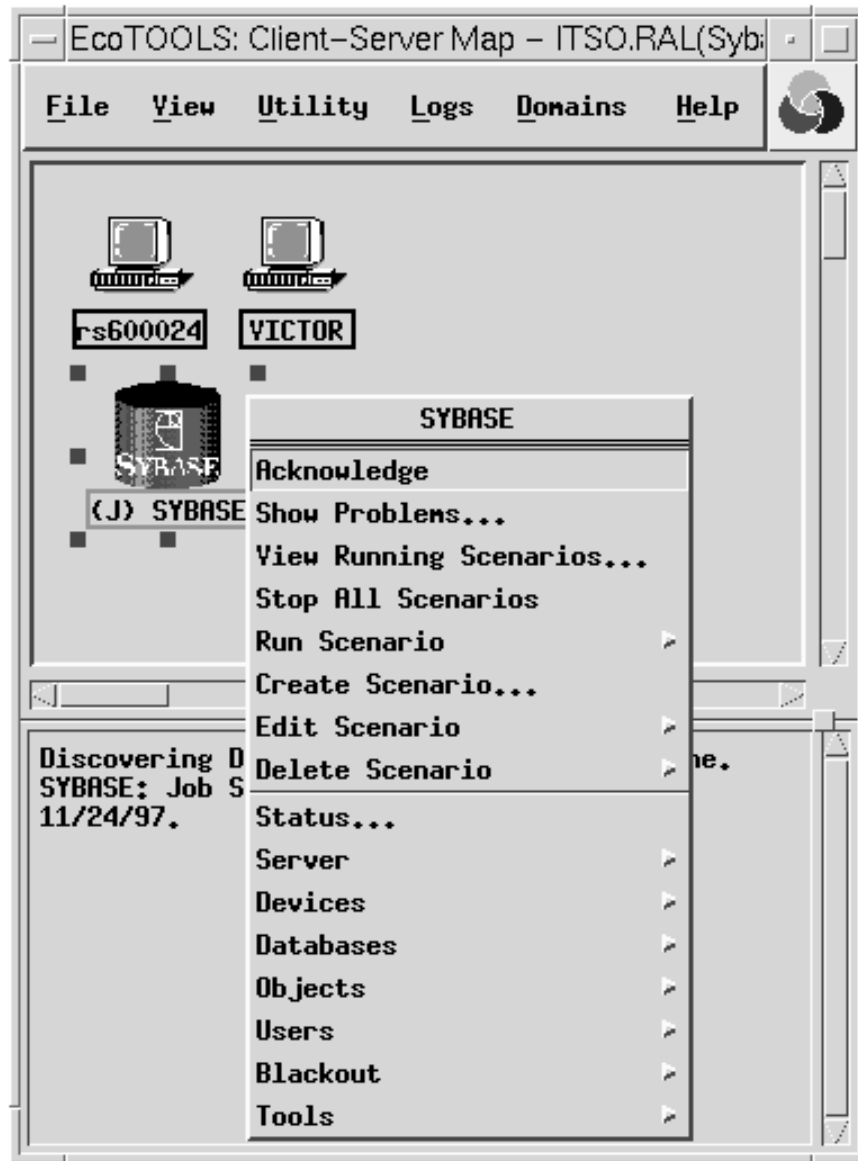


Figure 195. EcoTOOLS: Client-Server Map - ITSO.RAL(Sybase) Window

The window also shows the menu available from the right button of the mouse when selecting the Sybase object. The top part of the menu is similar to that which appears when clicking the right mouse button on a computer (client or server) in the EcoTOOLS: Computer Map window as shown in Figure 193 on page 198, which gives access to the various scenario functions. The lower half of the menu shown in Figure 195 (starting from Status...) gives access to functions specific to the Sybase SQL database.

Clicking the right mouse button on a client icon in the EcoTOOLS: Client-Server Map window produces a different menu for checking client information, such as the client application and user name of the database user.

4.9.1.1 Client Information

The following information is only available from the EcoTOOLS: Client-Server window (see Figure 198 on page 202). You cannot run scenarios on client machines from this window.

Click the right mouse button to display the Client Information menu. This menu has only one option, which allows you to display information regarding the type of database client software used on the node you selected. Choose the **Client Information...** option (the only option available) from the menu, which displays the window shown below.

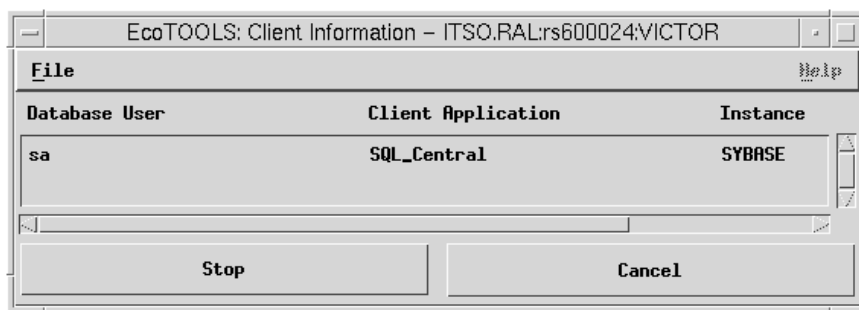


Figure 196. EcoTOOLS: Client Information - ITSO.RAL: rs600024: VICTOR Window

The window shows the client information for the node VICTOR.

The parameters in the window indicate that:

1. The Database User that the client is logged on as is sa.
2. The Client Application is SQL_Central.
3. The Instance of the database is SYBASE.

4.9.2 Scenarios

Scenarios are groups of EcoTOOLS monitoring agents that run on a node defined in an EcoTOOLS domain. When installed, EcoTOOLS has predefined scenarios to enable monitoring of system performance that are accessed by clicking the right mouse button on a node defined in a map window as shown in Figure 193 on page 198. The functions to create and edit scenarios are also accessed via this same menu.

4.9.2.1 Running Scenarios

Click on the node on which you want to run the scenario, then using the right mouse button click on the **Run Scenario** option from the menu that appears. This displays a further menu that contains a list of available scenarios (some in further sub-menus). Click on one of the listed scenarios to run it.

When a scenario is running, various windows are displayed depending on the configuration of the scenario. These new windows represent the data output configuration of the different agents in the scenario. The data output is configured for one or more of the following outputs:

- Graphs
- EcoMeter gauges
- Tables

Also displayed in the same window as any EcoMeters are event status monitors. Event status monitors are displayed as colored indicators, showing green if no events have occurred, and red if events have occurred. Examples of these windows are shown in the following figures.

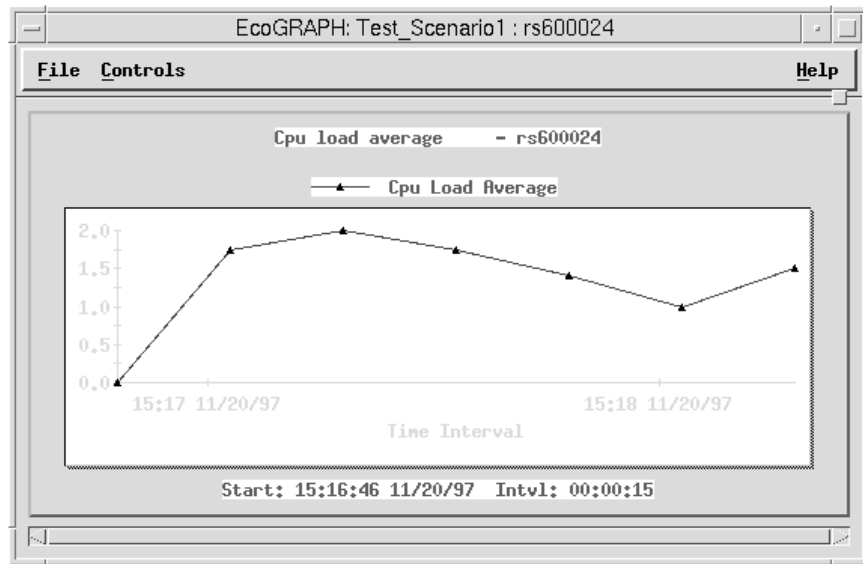


Figure 197. EcoGRAPH: Test_Scenario: rs600024

The graph in Figure 197 shows the Cpu Load Average for node rs600024, with an interval of 15 seconds between data points which is shown at the bottom of the window. The start time and date of the data collection is also shown at the bottom of the window (15:16:46 11/20/97).

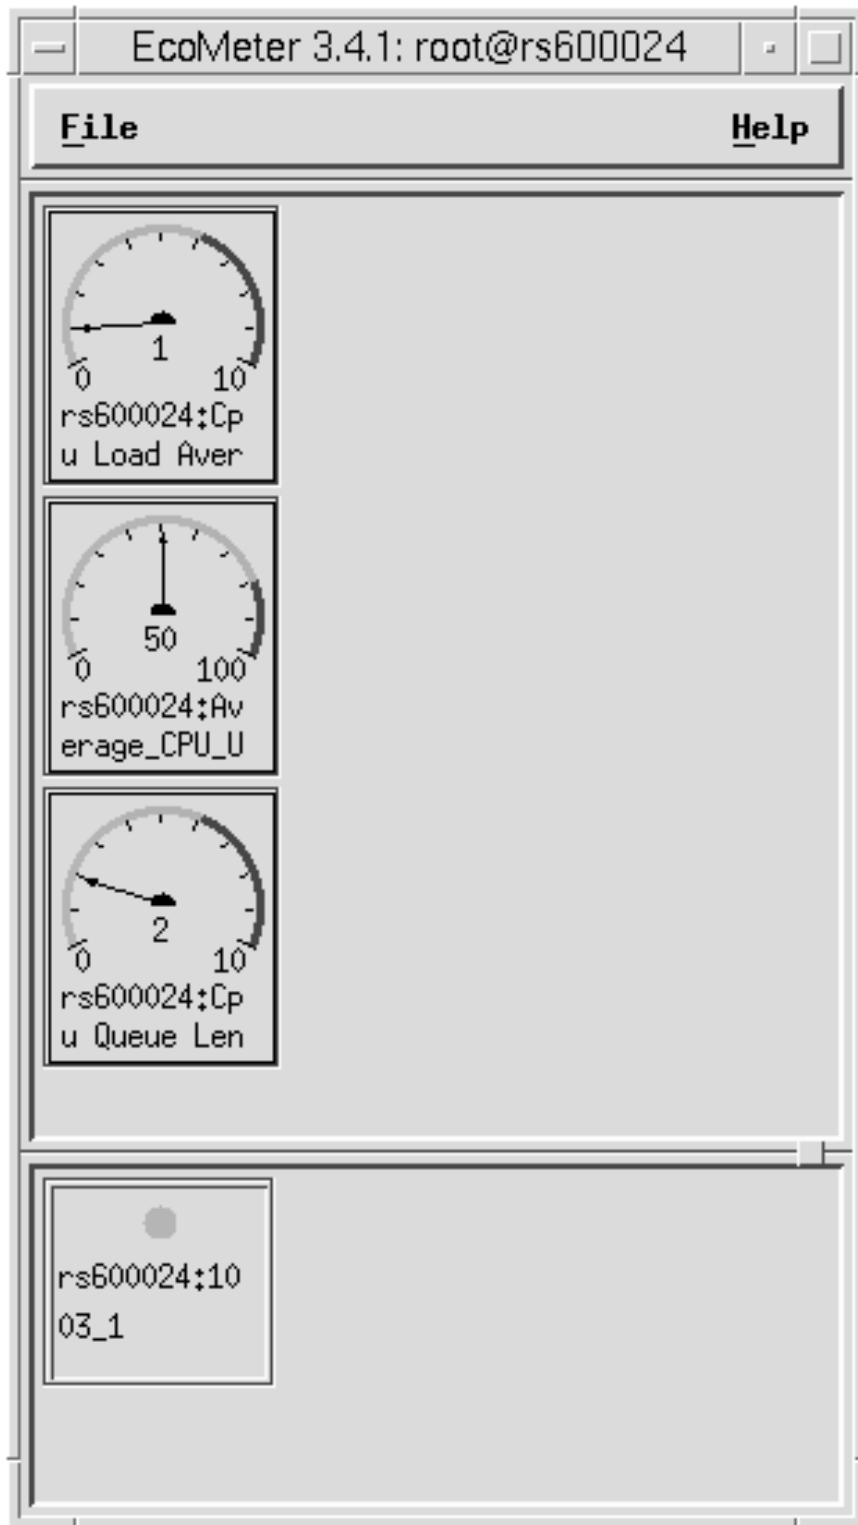


Figure 198. EcoMeter 3.4.1: root@rs600024

The window shown in Figure 198 shows three meters in the top part of the window, and one event status monitor in the lower part of the window. The three meters represent (top to bottom) CPU load average, average CPU utilization and CPU queue length, which are three separate EcoTOOLS agents. The event status monitor rs600024:1003_1 represents the agent CPU queue length. The number

1003_1 in the label for the event monitor refers to the agent number. The agent number is listed next to the agent name in the EcoTOOLS: Edit Scenario window shown in Figure 201 on page 205.

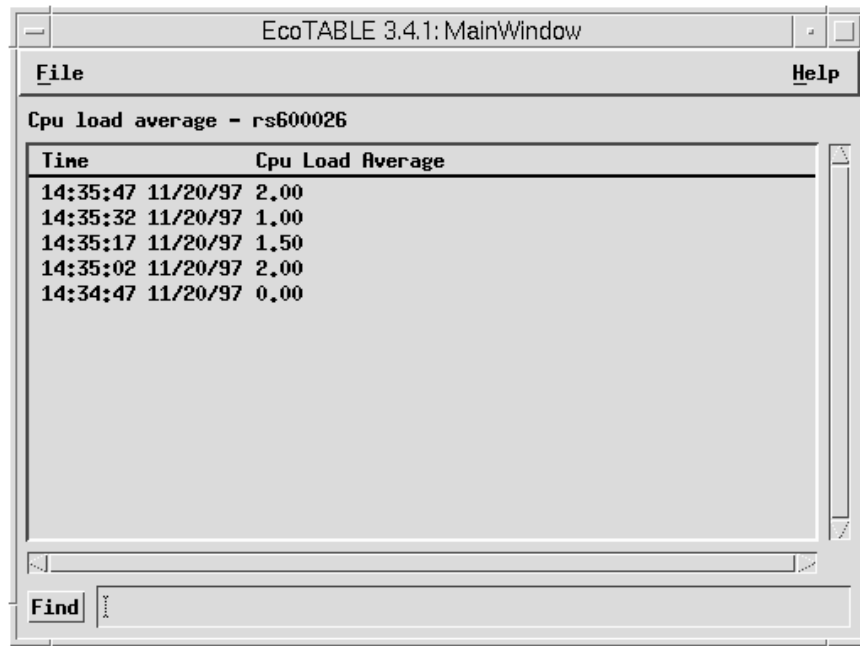


Figure 199. EcoTABLE 3.4.1:MainWindow

The example output table shown in Figure 199 shows the CPU load average for node rs600026. New values appear at the top of the table when the agent sends data. Note the times in the leftmost column of the table indicate an interval of 15 seconds.

4.9.2.2 Stopping Scenarios

There are two ways of stopping scenarios running on a node. You can either stop scenarios individually, or select the **Stop all scenarios** option from the menu that appears when you click the right mouse button on a node in the map.

To stop scenarios or agents in scenarios individually, select **View Running Scenarios...** from the node menu to see all scenarios running on the selected node. This action displays a window similar to the one shown below.

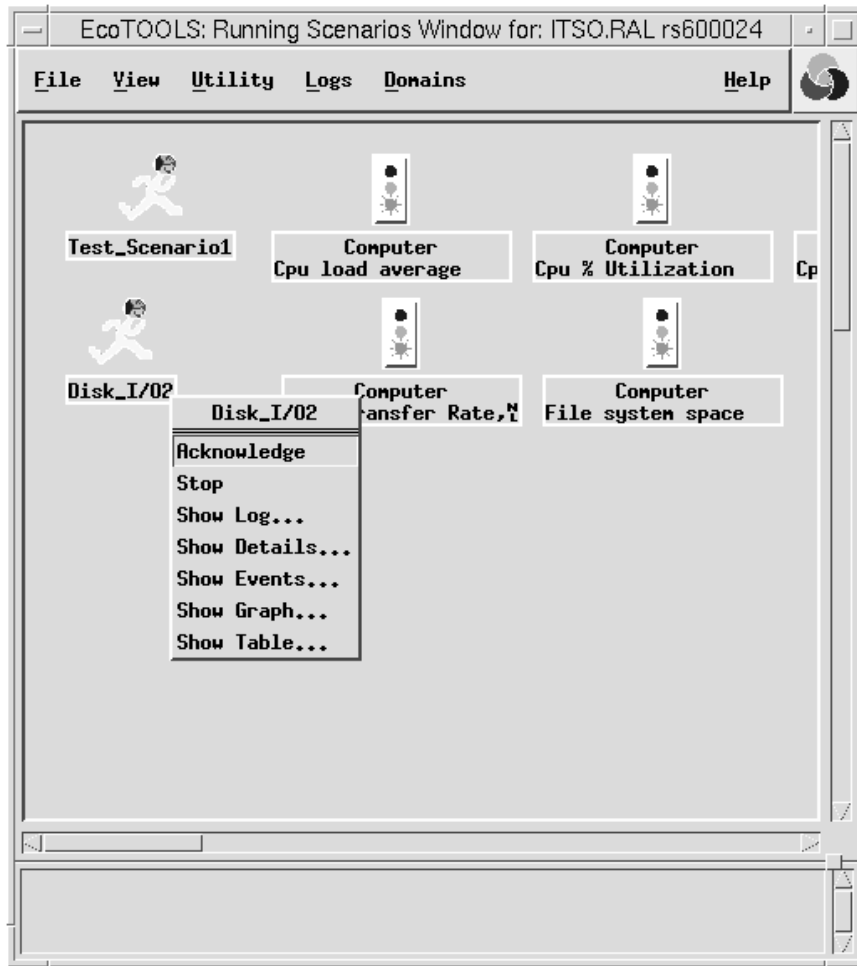


Figure 200. View Running Scenarios Window

A scenario is shown in the EcoTOOLS: Running Scenarios window as a running figure icon (Test_Scenario1 and Disk_I/O2 in Figure 200), and the individual agents running in a scenario are shown to the right of their respective scenario as a traffic light icon.

Click on either the scenario or individual agent icons using the right mouse button, and select **Stop** from the menu that appears to stop the selected scenario or agent. Click on the **Yes** button of the window that appears to confirm your choice.

4.9.2.3 Creating/Editing Scenarios

In addition to the default scenarios supplied with EcoTOOLS, customized scenarios can be created to suit the requirements of any situation. The stages involved in creating a new scenario are as follows:

1. Select monitoring agents to include.
2. Configure actions and outputs.
3. Save and add to scenario menu.

Select the option **Create new scenario...** from the right mouse button menu. This opens the EcoTOOLS: Edit Scenario window as shown below.

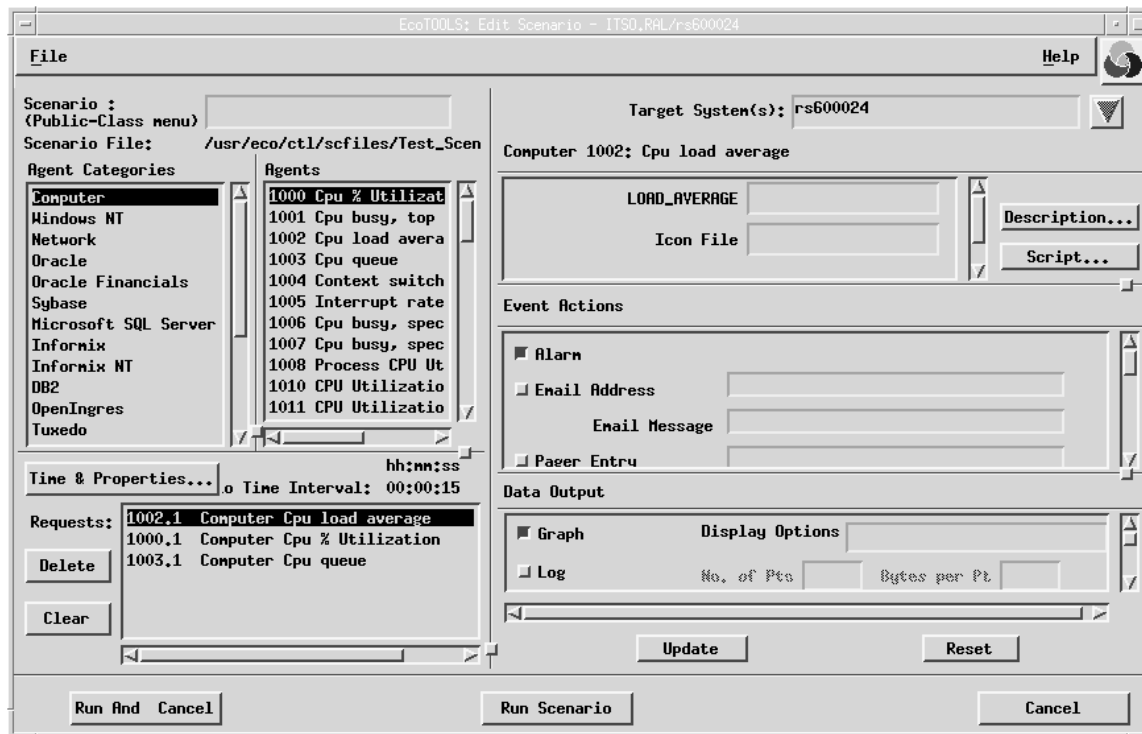


Figure 201. EcoTOOLS: Edit Scenario Window

There are four configuration areas for each agent configured in a scenario:

- Select Agents
- Edit Time and Properties
- Set Event Actions
- Select Data Output
- Save

Selecting Agents: Select an item in the Agent Categories box at the top left of the Edit Scenario window. When a category is selected a list of available agents appears in the Agents box. Select the agent that you want to add to the scenario and click on the **Add** button near the bottom of the window. The selected agent is added to the Requests box at the bottom left of the window.

Setting Time and Properties: Select an agent from the Requests box of the Edit Scenario window, then click on the **Time & Properties...** button. The window shown in Figure 202 on page 206 appears.

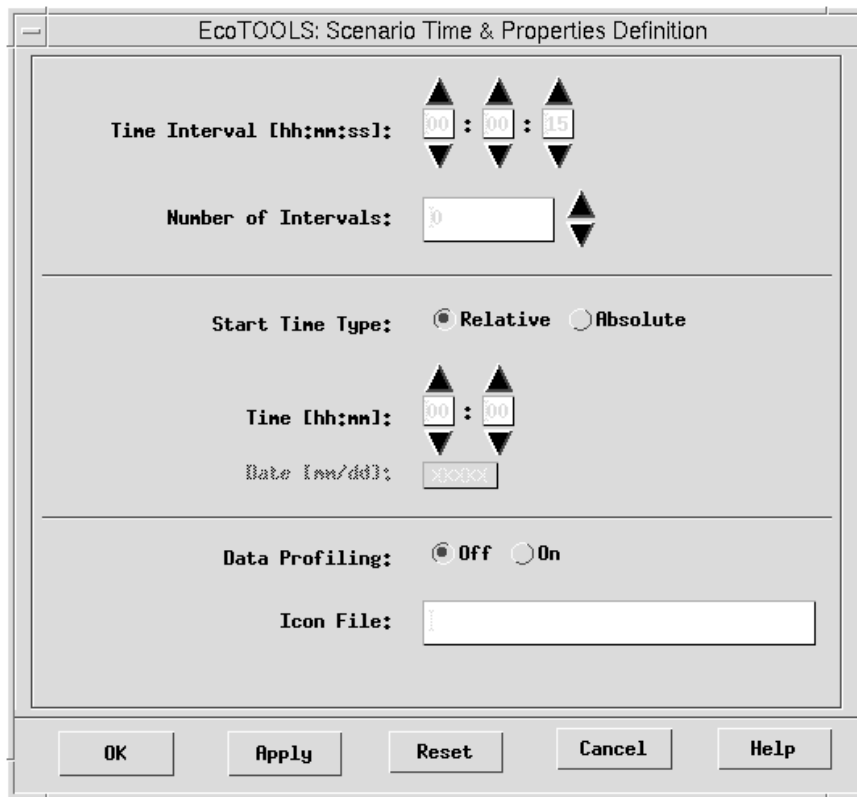
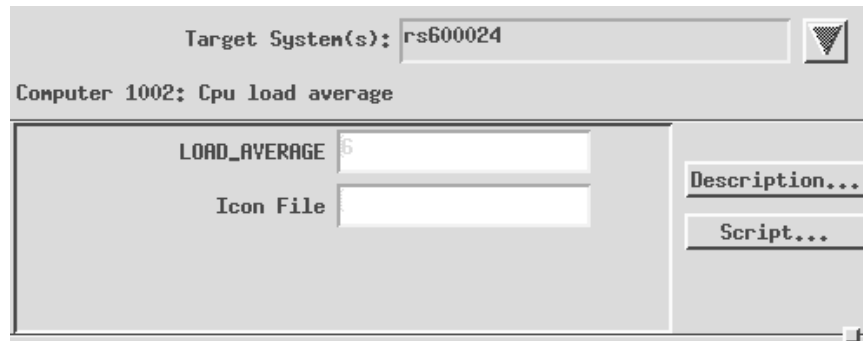


Figure 202. EcoTOOLS: Scenario Time & Properties Window

Configure the interval between times that the agent reports values to the management console in the top part of the EcoTOOLS Time & Properties Definition window by clicking on the arrows to adjust the hours, minutes and seconds intervals. Set the number of times the agent will report values by changing the Number of Intervals parameter. If the number of intervals is set to zero, the agent will run continuously until stopped manually.

The agent can be configured to start either relative to the current time or at a specific time (absolute). Click on either the Relative or the Absolute button beside the text Start Time Type and then enter the number of hours and minutes to elapse before starting the agent for a relative start time; or enter the actual time in hours and minutes at which the agent is to start for absolute configuration. Also set the Date (mm/dd) if using an absolute start time type.

Agent Parameters



Target System(s): rs600024

Computer 1002: Cpu load average

LOAD_AVERAGE 5

Icon File

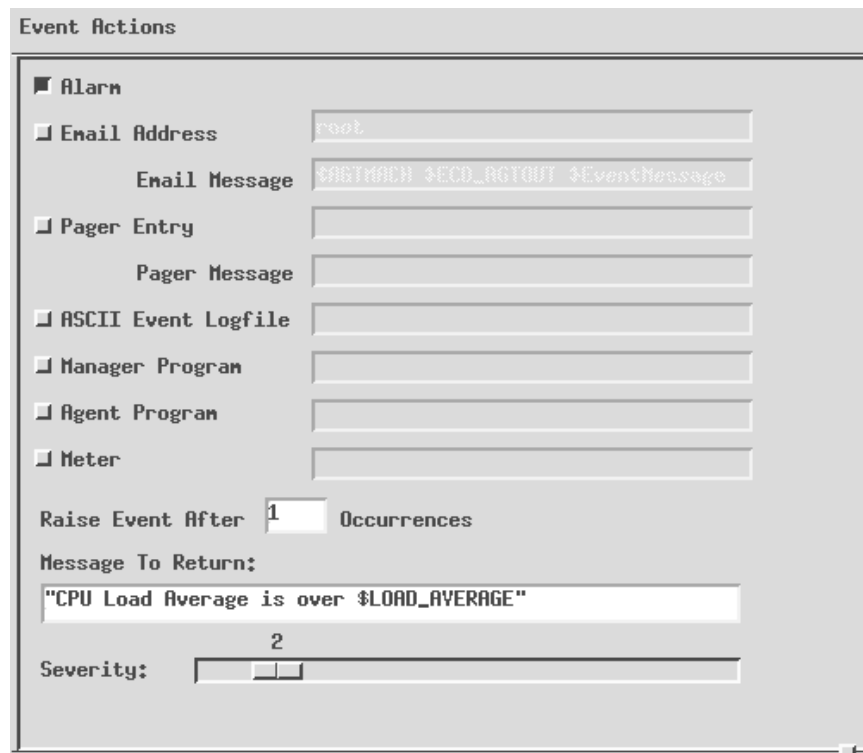
Description...

Script...

Figure 203. Agent Parameters Section of EcoTOOLS: Edit Scenario Window

Select an agent from the Requests box. The parameters for this agent are then displayed in the right hand portion of the Edit Scenario window. Set the parameter(s) for the agent to the desired level. The **Description...** button displays a window giving information about the parameters for this agent. The **Script...** button displays the agent script.

Event Actions: The Event Actions panel as shown (expanded) in Figure 204 defines what occurs when an agent has an event.



Event Actions

Alarm

Email Address root

Email Message \$AGENTNAME \$ECD_AGTOUT \$EventMessage

Pager Entry

Pager Message

ASCII Event Logfile

Manager Program

Agent Program

Meter

Raise Event After 1 Occurrences

Message To Return:

"CPU Load Average is over \$LOAD_AVERAGE"

Severity: 2

Figure 204. Event Actions Section of EcoTOOLS: Edit Scenario Window

An event is what occurs when an agent threshold is met or exceeded.

Activate event actions by clicking on the button to the left of each listed action. Multiple actions can be selected. The available actions are:

- Alarm

The map icon for the host will flash and a beep will sound. An event is also sent to the Event Console window.

- E-mail

Send an e-mail to the address entered.

- Pager

Send a pager message to the pager tag entered in the Pager Entry field. This refers to a tag in the ecopager.ctl file.

- ASCII Event Logfile

The event will cause an entry into an ASCII text file defined in the ASCII Event Logfile field.

- Manager Program

Execute a script or program on the management station when the event occurs. For example, this can be used to trigger a corrective action on the server when a critical condition on an agent has occurred.

- Agent Script

Execute a script or program on the agent station when the event occurs.

- Meter

Display the status of the event in the EcoMeter window. The status appears in the lower portion of the EcoMeter window as shown in Figure 198 on page 202. A green status indicates that no events have occurred, a red status indicates that events have occurred.

- Raise Event After X Occurrences

An event action will not occur until the event has occurred the number of times specified here. This is especially useful when an unusual condition occurs only once and this does not indicate a critical condition but is only an exception. If in that case you have set the number of occurrences to 2, no event will be triggered.

- Message To Return

The text specified here appears in the Event Console window when the event occurs.

- Severity

Set the severity by dragging the slider to the left or right. Severity one is the highest severity and severity ten is the lowest.

Data Output: The data output section of the EcoTOOLS: Edit Scenario window defines how the collected information/parameters are displayed on the Management Station.

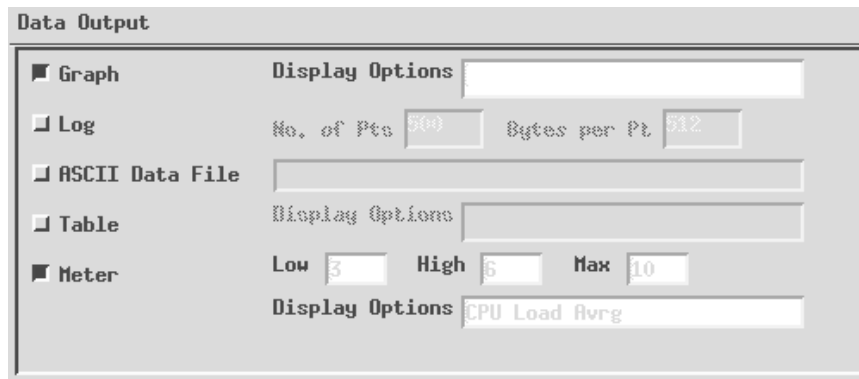


Figure 205. Data Output Section of EcoTOOLS: Edit Scenario Window

The following options are available:

- Graph

Plot a graph of the information collected. A separate graph is displayed for each agent configured with this option.
- Log

A circular log is kept of the data collected. Set the size of the log by entering values for the No. of Pts and Bytes per pt. These values signify the number of data points to collect before discarding the earliest data point, and the maximum size of each data point collected. (This will depend upon the agent.)
- ASCII Data File

Collected data is written to a user-defined text file if this option is selected. Type the text file name and location in the space provided.
- Table

This selection starts EcoTable when the first data point is received from the agent.
- Meter

This section displays an EcoMeter for the agent. The options Low, High and Max set the scale of the meter. From zero to low is shown as green on the scale; low to high is shown as yellow; and high to max is displayed as red. The Display Options option can be used to specify alternative display consoles. The example shown below causes the EcoMeter to be displayed on host rs600026.

```
-display rs600026:0.0
```

Saving and Updating Scenarios: Save the scenario that you created once you have finished configuration. If the scenario is a newly created one, as opposed to being a previously created scenario that you are editing, then you have the following options:

- Save as Private Scenario

The scenario that you have created will be accessible from the right click menu only by the user that you are currently authenticated to EcoTOOLS with.
- Save as Public Scenario

The saved scenario will appear on the right click menu for all EcoTOOLS users. A public scenario is preceded with an asterisk (*) in the scenario sub-menus.

When you select either of these options from the File menu of the EcoTOOLS: Edit Scenario window, you have the option to save the scenario for a particular object (the one you clicked on with the right mouse button to access the Create Scenario... menu option), or for all objects.

Choose the **Update Current Menu Item** option from the **File** menu if you are editing an existing scenario. The changes you have made to the scenario will become active the next time that you run the scenario.

4.9.2.4 Running Database Scenarios

Database scenarios are handled in a similar way as other scenarios. They have to be created first before they can be activated on the database server.

To create a database scenario open the Client-Server Map window as shown in Figure 195 on page 199. Then, use the right mouse button to select **Create Scenario...** from the context menu of a database server icon. In the Edit Scenario Window select the database type you want to monitor from the Agent Categories area. Select the desired monitors from the Agents area and click the **New Request** button to include them into the scenario. Finally, select each monitor with the left mouse button from the Requests box to enter the necessary monitoring parameters (Figure 204 on page 207) and to specify the event actions (see Figure 205 on page 209). Select the **Update** button when you are finished.

The scenario must be saved and integrated into the menu structure of EcoTOOLS. Therefore, enter a name in the Scenario field. Select **Save As Public Scenario** and the desired option from the File menu to save the changes.

As an example we created a scenario to monitor the CPU usage of a Sybase database. We want the monitor to report the five most active processes and send an alarm when the CPU utilization exceeds 90%.

To start the scenario you can either use the Run & Cancel or the Run Scenario button from the Edit Scenario window or select **Run Scenario** and the option representing your scenario from the context menu of the database server icon. Depending on the scenario configuration the system displays various windows to display the measured data.

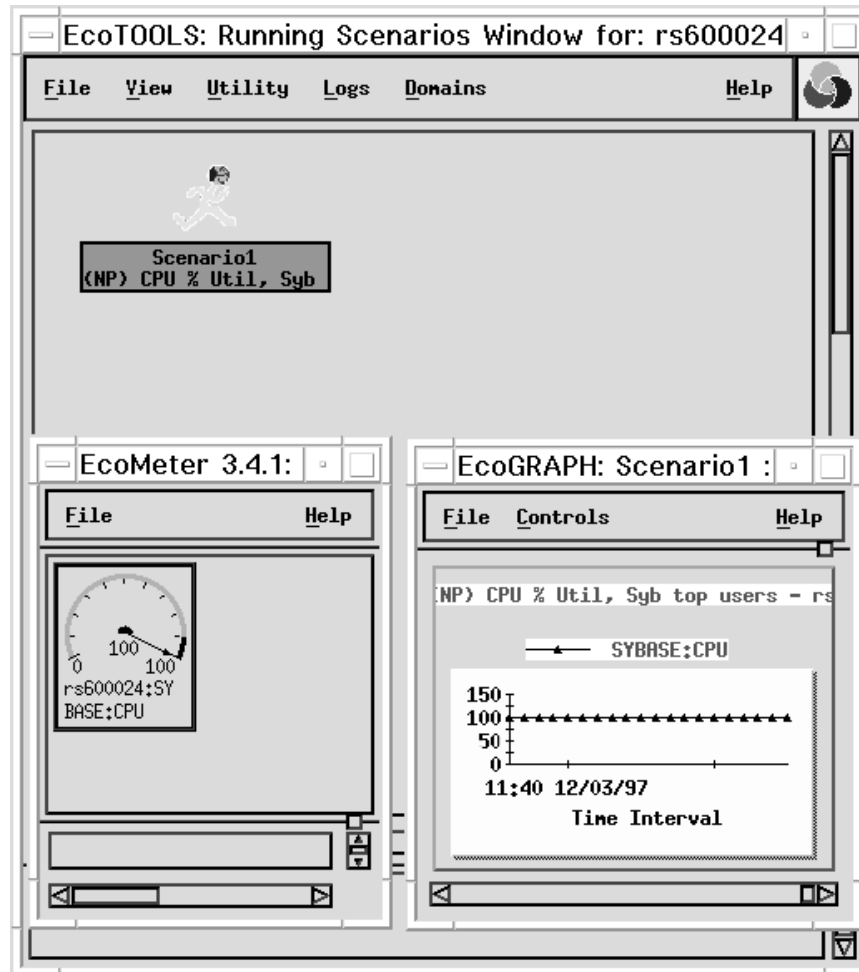


Figure 206. EcoTOOLS: Running Scenarios

For our example we selected an EcoMeter gauge and an EcoGraph as the output option. Both are displayed in Figure 206. As the CPU utilization is above the threshold of 90% the icon is blinking red to indicate a problem. The output windows are updated according to the interval specified in the Time & Properties window (Figure 203 on page 207), which is 30 seconds in our example.

Additionally, EcoTOOLS provides functions to manage databases. They may be accessed via the context menu of the database icon in the Client-Server Map window. They comprise basic functions to administer a database such as reporting the state of the SQL server or starting and stopping and stopping the database.

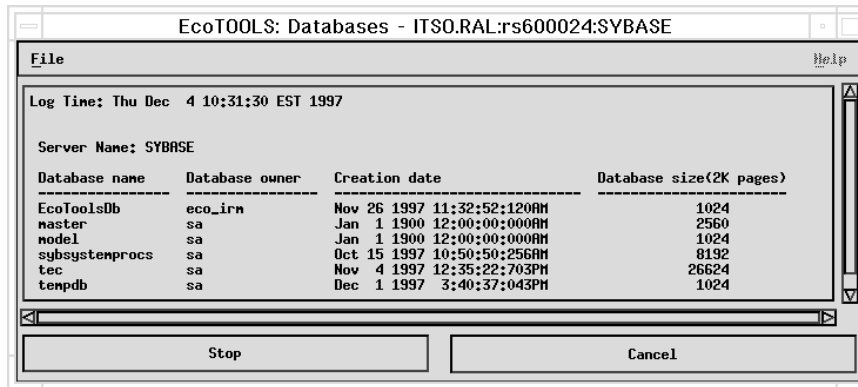


Figure 207. EcoTOOLS: Databases Window

Most functions display information about the database in question. Figure 207 gives an example and displays all databases managed by the server named SYBASE on rs600024 together with the sizes of databases and their dates of creation. To access this function select the **Server** and then the **Databases...** option from the context menu of a Sybase icon.

Other examples may be free space information and log space availability. In general, the menu items change to reflect the different types of databases.

4.9.3 Inventory

The EcoTOOLS Management Station GUI provides a means to view the hardware and operating system inventory of nodes running the EcoTOOLS agent code. Access this function via the right mouse button menu after selecting a node from one of the EcoTOOLS computer map screens. Select the **Computer Inventory...** option. The inventory information displayed is a snapshot of the host at the time of inventory selection. An example inventory display window is shown below in Figure 208.

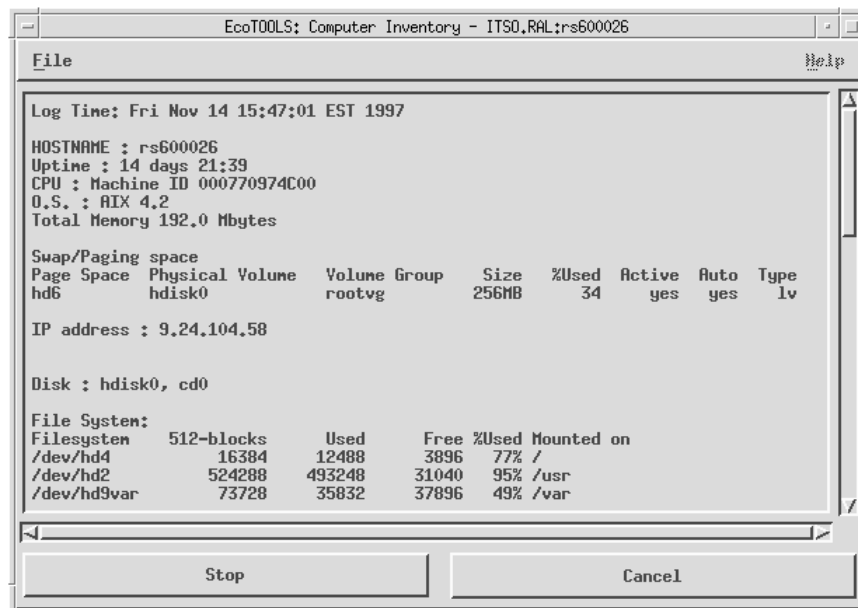


Figure 208. EcoTOOLS: Computer Inventory - ITS0.RAL: rs600026 Window

Save the inventory information displayed by clicking on the **Save** option from the **File** menu. The information is saved in an ASCII format.

Click on the vertical scroll bar to view inventory information further down the window. Click on **Cancel** when you are finished viewing the inventory.

4.9.4 EcoTOOLS Logs

The EcoTOOLS Management Station interface utilizes three logs to provide information to the console operator:

- EcoTOOLS log
- Event log
- Agent Error log

These logs are accessed via the **Logs** menu from the EcoTOOLS: MainWindow (see Figure 191 on page 195) or any of the Computer Map or Client-Server Map sub-windows.

4.9.4.1 The EcoTOOLS Log

This log file displays all of the scenarios that have been or are running and details about them, including the number of (if any) events or errors that have occurred. The window in Figure 209 shows an example of the EcoTOOLS Log window.

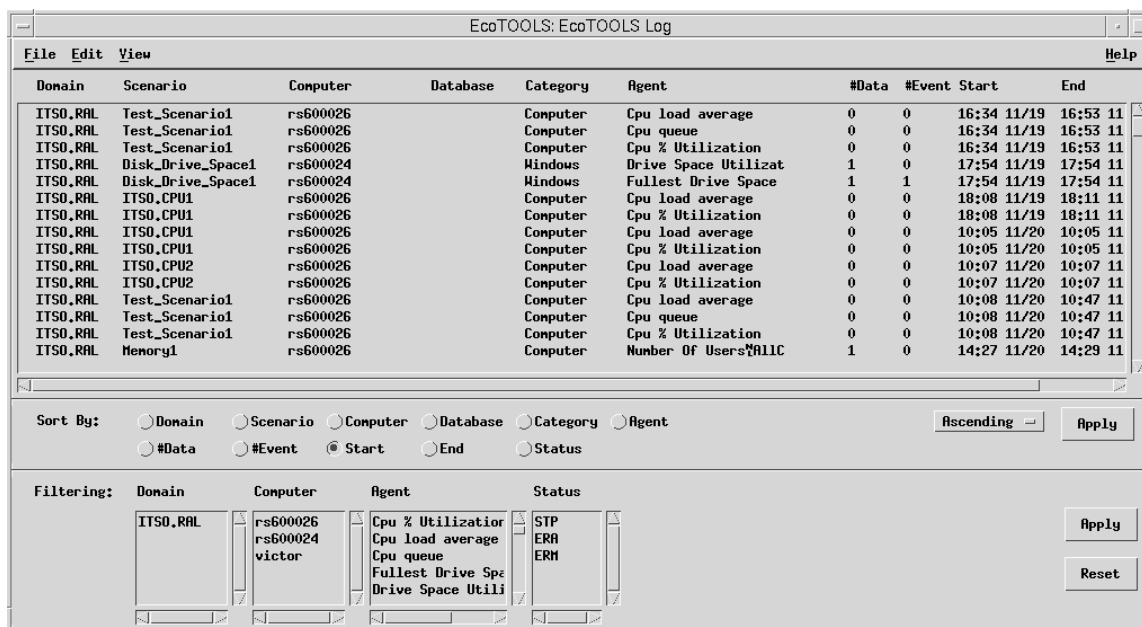


Figure 209. EcoTOOLS Log Window

The items in the log can be ordered or filtered by selecting options in the lower part of the window. Select a button to the right of Sort By: and then click **Apply** to group log entries by the selected field (for example, group all entries for each Computer together). Select items from the Filtering: section at the bottom of the window and then click **Apply** to display only entries that match your selections.

4.9.4.2 The Event Log

This log lists all events that have occurred as part of a scenario that has been run or is currently running. The event log file is also used to forward events to TME 10 Enterprise Console. Refer to 4.12.3.3, “Configuring the EcoTOOLS Event Adapter” on page 257 for more information.

4.9.4.3 The Agent Error Log

This window appears with a message when any errors concerning EcoTOOLS agents occur, such as the correct daemons are not running on the target host. This window is also accessible from the **Show Problems** item on the right mouse click menu.

4.10 Installing EcoTOOLS Plus Module for TME 10 in Our Environment

Now that we have installed EcoTOOLS in our environment and have described the basic functions of the product, we install EcoTOOLS Plus Module for TME 10 in our environment in order to integrate EcoTOOLS with TME 10.

EcoTOOLS Plus Module for TME 10, like any other 10/Plus module, is installed using the TME 10 Desktop.

However, before you can install EcoTOOLS Plus Module for TME 10 you need to install the EcoTOOLS application, as the code for EcoTOOLS Plus Module for TME 10 will be extracted to the EcoTOOLS installation directory when installing EcoTOOLS.

Note

At the time of writing of this redbook, EcoTOOLS Plus Module for TME 10 support TME 10 Version 3.1. EcoTOOLS 4.0 which should be available when you read this book, will support TME 10 Version 3.2.

In our environment, the code for EcoTOOLS Plus Module for TME 10 is in the following directory after the installation of EcoTOOLS:

```
/usr/eco/src/EcoCONNECT/Plus
```

In order to install EcoTOOLS Plus Module for TME 10 we start the TME 10 Desktop. The following window will appear:

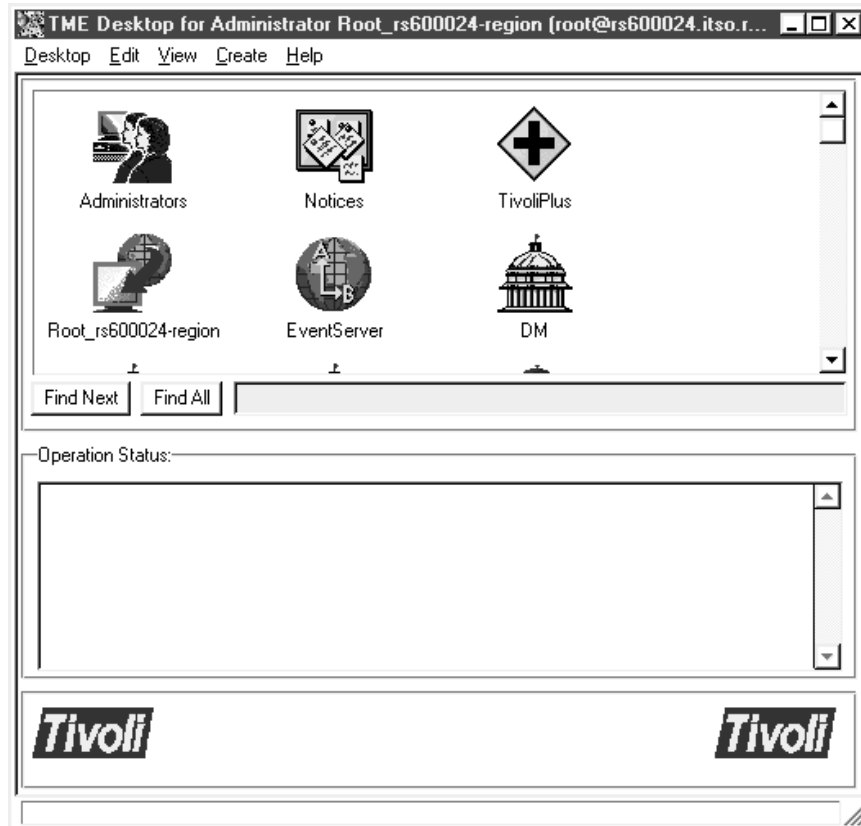


Figure 210. The TME 10 Desktop

You will notice that there is already a TivoliPlus icon on the desktop. This is because we use the same TMR server on which we had already installed BEST/1+ for TME 10 as described in Chapter 2, “BEST/1 from BGS Systems Inc.” on page 13.

Note

Before you install EcoTOOLS Plus Module for TME 10 you should make a backup copy of the Tivoli object database using the `wbkupdb` command. This will allow you to back up to the previous version in case there are any problems during the installation of EcoTOOLS Plus Module for TME 10.

In order to start the installation, select **Desktop** from the menu bar and then **Install** from the pull-down menu. From the sub-menu select **Install Product....**

The Install Product window as shown in Figure 211 on page 216 will appear. In order to change to the EcoTOOLS Plus Module for TME 10 installation directory, select the **Set Media...** button.

This will start the File Browser window. In this window select the source directory for EcoTOOLS Plus Module for TME 10. In our example this is:

```
/usr/eco/src/EcoCONNECT/Plus
```

Once you are finished, select the **Set Media & Close** button. This will get you back to the Install Product window as shown below.

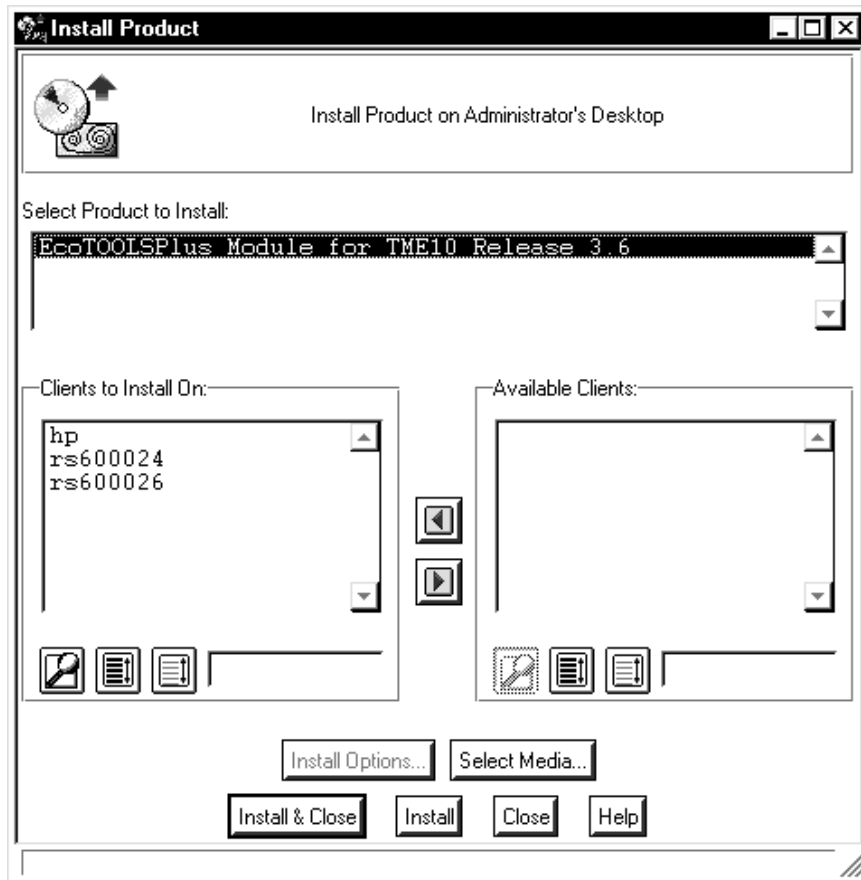


Figure 211. Install Product Window

Click on **EcoTOOLSPlus Module for TME 10 Release 3.6**. This will move all your managed nodes from the Available Clients section of the window to the Clients to Install On section.

If EcoTOOLS Plus Module for TME 10 does not show up in the Select Product to Install section, check that the code for EcoTOOLS Plus Module for TME 10 exists in the directory you specified as the installation source.

Note

You should notice that Release 3.6 of EcoTOOLS Plus Module for TME 10 refers to the version of EcoTOOLS the 10/Plus module is included with and not to the version of TME 10.

Now you can select the **Install & Close** or **Install** button in order to start the installation of EcoTOOLS Plus Module for TME 10.

The following window will appear:

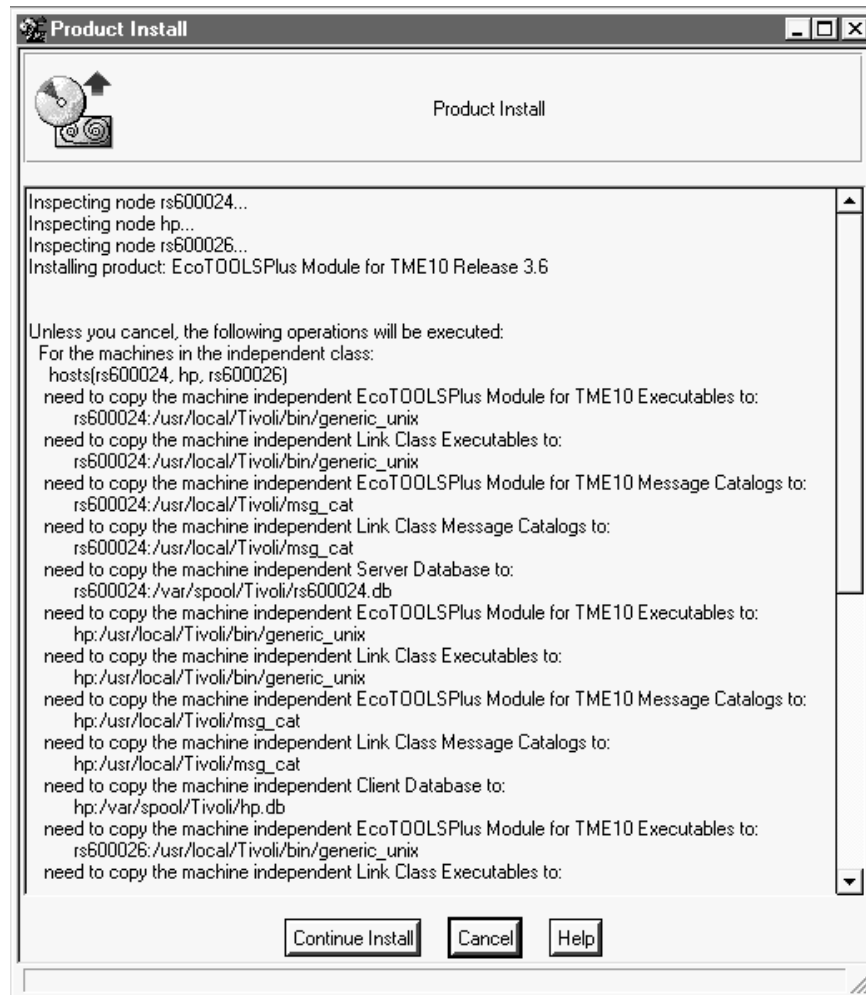


Figure 212. Product Install Window

Review the messages and then click on **Continue Install** to start the actual installation. In our environment, EcoTOOLS Plus Module for TME 10 will be installed on rs600024, rs600026 and hp.

Once the installation has finished (this can take a couple of minutes) you should see messages in the Product Install window similar to the ones shown at the bottom of the following figure.

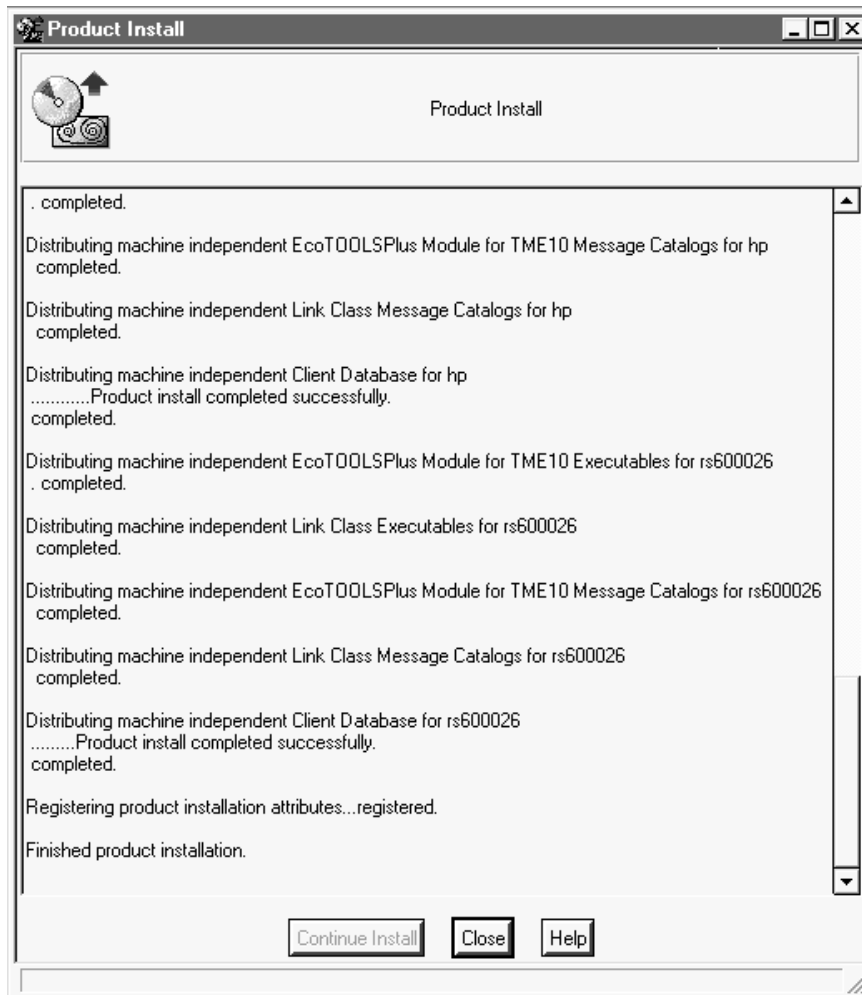


Figure 213. Product Install Window

Select the **Close** button to return to the TME 10 Desktop main window.

Double-click on the **TivoliPlus** icon in order to open the TivoliPlus window as shown in the following figure:

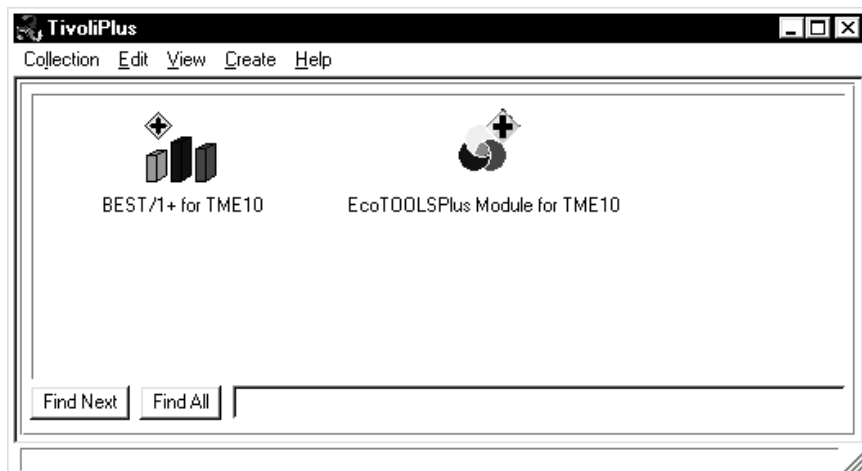


Figure 214. TivoliPlus Window

You will notice that the EcoTOOLS Plus Module for TME 10 icon has been added to the window, indicating that the installation of EcoTOOLS Plus Module for TME 10 was successful.

Note

You will also notice that there is an additional icon for BEST/1+ for TME 10. This is because we use the same TMR server for EcoTOOLS Plus Module for TME 10 that was already used in BEST/1.

In order to see the components of EcoTOOLS Plus Module for TME 10, double-click on the **EcoTOOLS Plus Module for TME 10** icon.

This will open the EcoTOOLS Plus Module for TME 10 window as shown below:

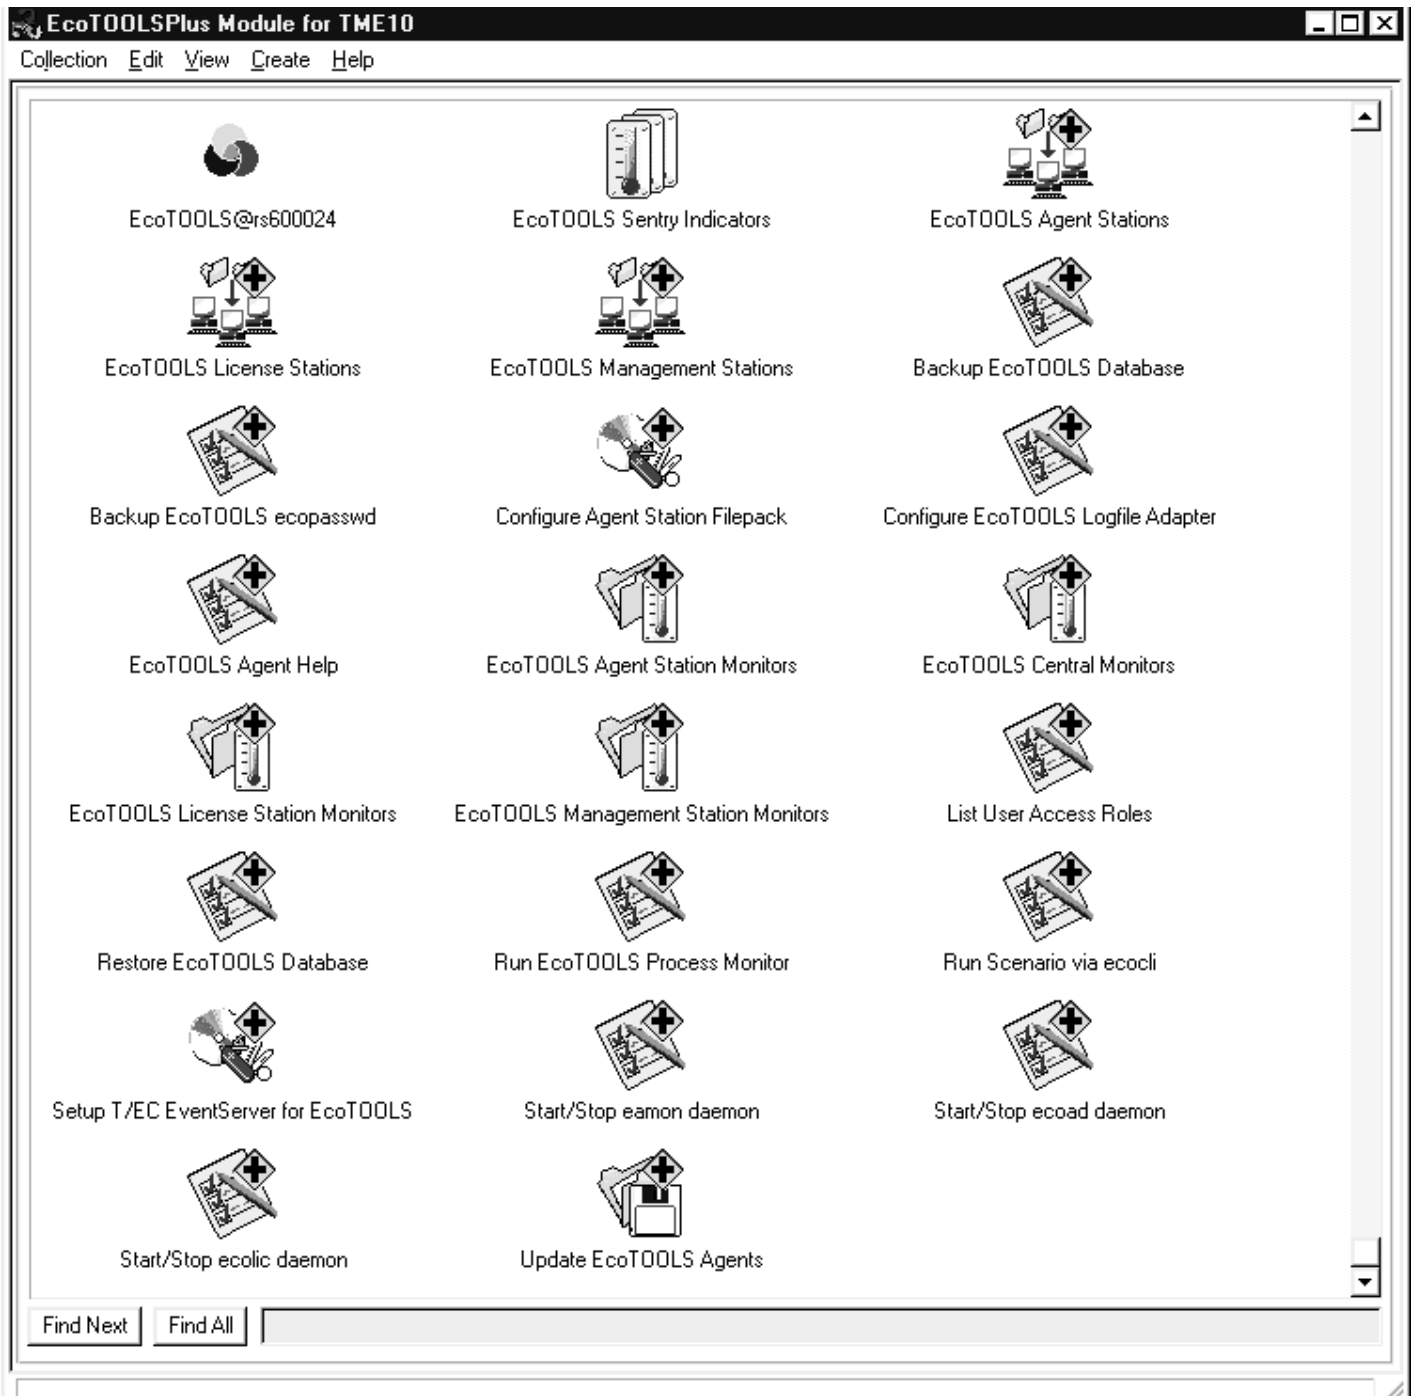


Figure 215. EcoTOOLS Plus Module for TME 10 Window

The icons shown represent the different options available in EcoTOOLS Plus Module for TME 10. In the following sections we look at all the functions in detail.

4.11 Overview of the EcoTOOLS Plus Module for TME 10 Functionality

The EcoTOOLS Plus Module for TME 10 provides a collection that supports the system administrator during installation and maintenance of EcoTOOLS on the management and agent stations. Refer to Figure 215 on page 220 for an overview of the components available in EcoTOOLS Plus Module for TME 10.

The collection contains profile managers that can be used during the distribution process of the EcoTOOLS client software. The profile managers can be used to distribute the profiles contained in the EcoTOOLS Plus Module for TME 10.

Profiles are supplied for TME 10 Software Distribution and TME 10 Distributed Monitoring.

The TME 10 Software Distribution profiles can be used to automatically deploy the agent component of EcoTOOLS.

The TME 10 Distributed Monitoring profiles allow you to monitor the various EcoTOOLS processes, such as the license server and agent processes. Additionally, you can configure event adapters for the TME 10 Enterprise Console to observe EcoTOOLS from a central point.

In the following sections we give a brief overview of the different components of the EcoTOOLS Plus Module for TME 10. In the subsequent sections we then explain the components in more detail in example scenarios.

4.11.1 Components to Set Up Agent Stations




Figure 216. Components to Set Up Agent Stations

Use the components shown above to generate TME 10 Software Distribution file packages and to distribute them to the EcoTOOLS agents. With the file package distribution the installation of the EcoTOOLS software is performed automatically.

- **Configure Agent Station Filepack:** This task generates TME 10 Software Distribution file packages containing all necessary software for the installation on the managed nodes. Regarding different platform types in your environment you may have to perform this task several times, once for each platform.
- **Update EcoTOOLS Agents:** This collection receives all file packages generated with the previous task. Each file package contains the machine-dependent software for the different destination platform types. References to shell scripts within the file packages perform the installation of the agent software. The shell scripts are part of the EcoTOOLS Plus Module for TME 10 and are copied to the agents during the installation of the EcoTOOLS Plus Module for TME 10.

- EcoTOOLS Agent Stations: Profile manager for EcoTOOLS agents. You can add all managed nodes as subscribers to this profile manager that you want to become EcoTOOLS agents.

In 4.12.1, “Deploying EcoTOOLS Using TME 10 Software Distribution” on page 227 you can find more detailed information about the distribution and installation process of the EcoTOOLS agent software.

4.11.2 Components to Set Up TME 10 Enterprise Console and TME 10 Distributed Monitoring

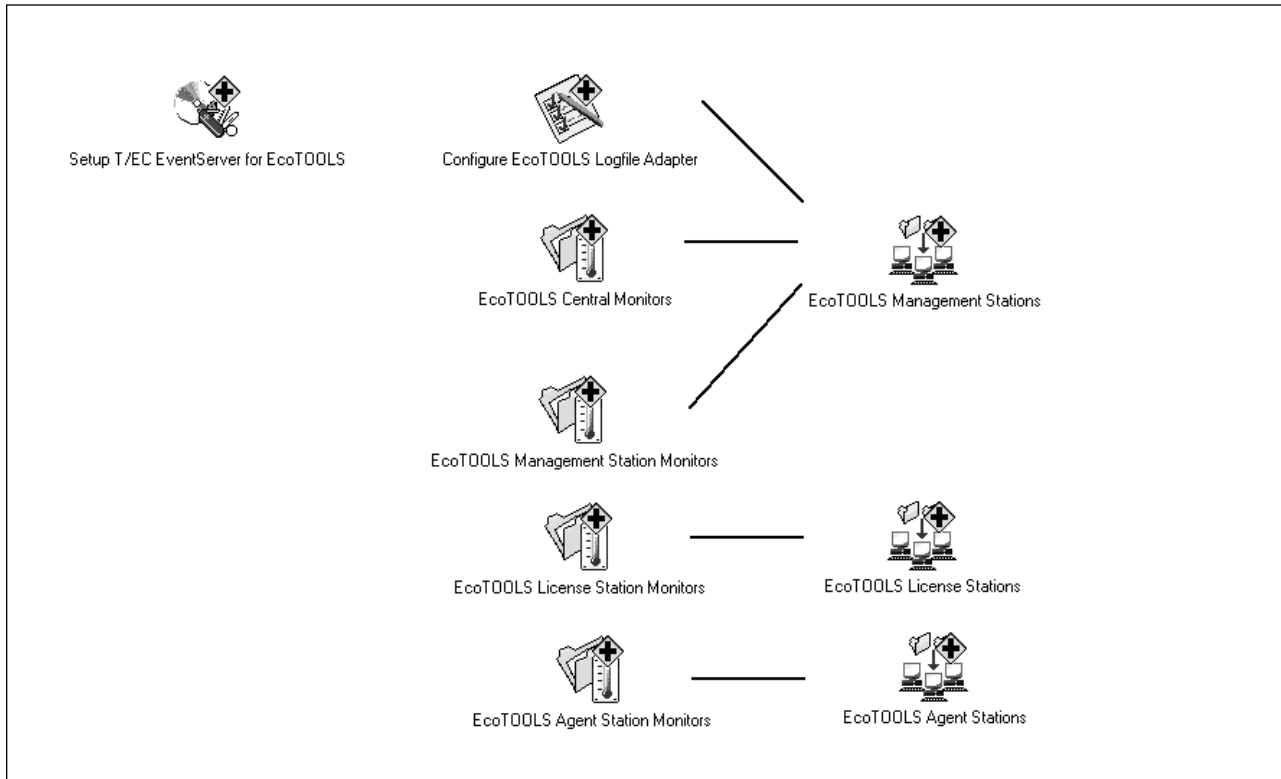


Figure 217. Components to Set Up TME 10 Enterprise Console and TME 10 Distributed Monitoring

These components allow you to configure the TME 10 Enterprise Console and the EcoTOOLS event adapter. Additionally, you can modify TME 10 Distributed Monitoring profiles and activate them on the EcoTOOLS machines.

- Setup TEC EventServer for EcoTOOLS: Task to prepare the TME 10 Enterprise Console to receive events forwarded from the EcoTOOLS event adapter. Run this task before the Configure EcoTOOLS Logfile Adapter.
- Configure EcoTOOLS Logfile Adapter: Task to set up the event adapter on the EcoTOOLS agents. The event adapter transforms EcoTOOLS events into TEC events and forwards them to the TME 10 Enterprise Console. Run the Setup TEC EventServer for EcoTOOLS task prior to this task.
- EcoTOOLS Central Monitors: Monitor collection to observe the status of the EcoTOOLS management station, license station and agent stations or any other resource by means of a ping request.
- EcoTOOLS Management Station Monitors: The monitors of this collection monitor several resources:

- The agent monitoring process called eamon, which is responsible for the execution of scenario requests. The monitor checks every minute whether the daemon is running.
- Available disk space in the EcoTOOLS database directory which usually is /usr/eco/db. The disk space is checked every 15 minutes by default and should be considered critical when less than 1 MB.
- Swap space available on the EcoTOOLS management station is checked every 15 minutes by default. The swap space should be at least 10 MB for EcoTOOLS.
- Size of the ASCII event log file, which should be less than 750 KB. This monitor runs every 15 minutes by default.

The subscriber list of this profile defaults to the EcoTOOLS management stations.

- EcoTOOLS License Station Monitors: This collection contains monitors to check the status of the ecolic daemon and the number of available licenses. The ecolic daemon is the license server that controls the number of available EcoTOOLS licenses. The daemon status is checked every minute and the number of licenses is controlled every 15 minutes. By default, the profiles of this collection are distributed to the subscribers of the EcoTOOLS License Station's profile manager.
- EcoTOOLS Agent Station Monitors: This collection contains two monitors to control the EcoTOOLS dispatcher daemon called ecoad and the availability of swap space on the agent stations where the dispatcher is running. The status of the daemon is controlled every minute, the swap space every 15 minutes. By default, the EcoTOOLS Agent Station's profile manager subscribes to this profile.
- EcoTOOLS Management Stations: Profile manager for the EcoTOOLS management server which is the TMR server, rs600024 in our environment.
- EcoTOOLS License Stations: Profile manager for the EcoTOOLS license server which must be a Tivoli managed node.
- EcoTOOLS Agent stations: Profile manager for EcoTOOLS agent stations which must be Tivoli managed nodes.

Refer to 4.12.3, "Event Management for EcoTOOLS Using TME 10 Enterprise Console" on page 241 for more detailed information about the configuration of the TME 10 Enterprise Console. The configuration of TME 10 Distributed Monitoring is described in 4.12.2, "Monitoring EcoTOOLS Using TME 10 Distributed Monitoring" on page 236.

4.11.3 Components to Monitor EcoTOOLS Processes

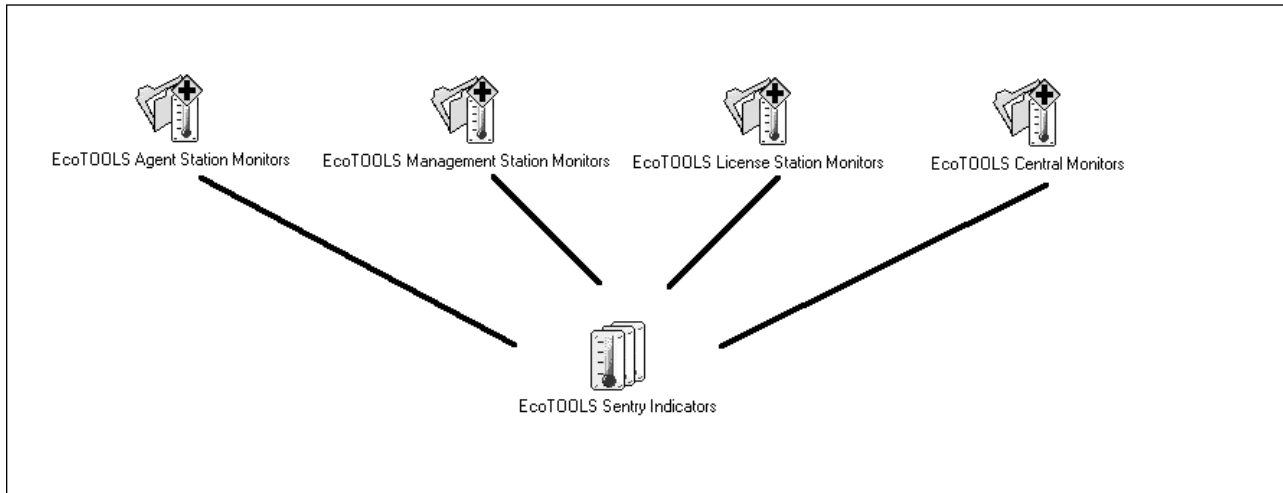


Figure 218. Components to Monitor EcoTOOLS Processes

The monitoring components provide mechanisms to integrate the observation of the EcoTOOLS machines and daemons into Tivoli. The following icons belong to this group:

- EcoTOOLS Management Station Monitors
- EcoTOOLS License Station Monitors
- EcoTOOLS Agent Station Monitors
- EcoTOOLS Central Monitors
- EcoTOOLS Sentry Indicators

The above listed icons represent TME 10 Distributed Monitoring profiles.

Note

For the EcoTOOLS Plus Module for TME 10 the term *Sentry* is still used when referring to indicators. However, you should notice that the official term for the TME 10 monitoring application is now TME 10 Distributed Monitoring.

The EcoTOOLS Sentry Indicators can be used as a central unit to quickly examine the status of the monitored resources. In order to access the indicators simply double-click on the **EcoTOOLS Sentry Indicators** icon.

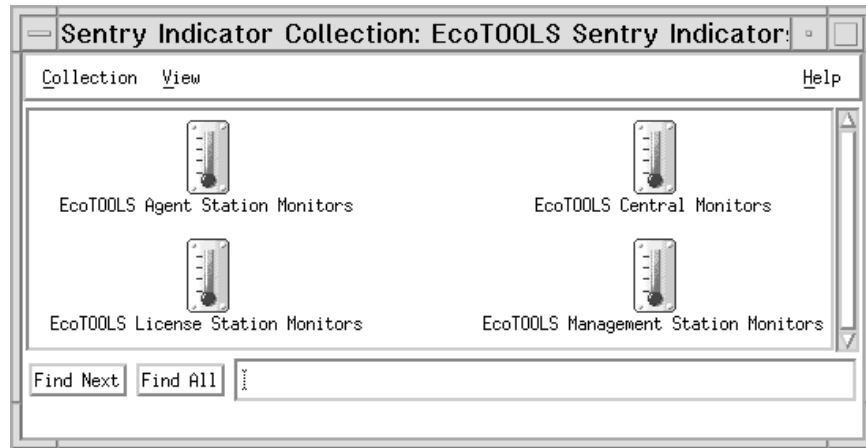


Figure 219. The Sentry Indicator Window for EcoTOOLS

Figure 219 shows the Sentry Indicator window containing four indicators, one for each TME 10 Distributed Monitoring profile of the EcoTOOLS Plus Module for TME 10. As soon as one of the active monitors triggers an event the corresponding thermometer rises and the operator can easily identify the origin of the problem.

The monitoring part of the EcoTOOLS components is described in more detail in 4.12.2, “Monitoring EcoTOOLS Using TME 10 Distributed Monitoring” on page 236.

4.11.4 Components to Maintain EcoTOOLS

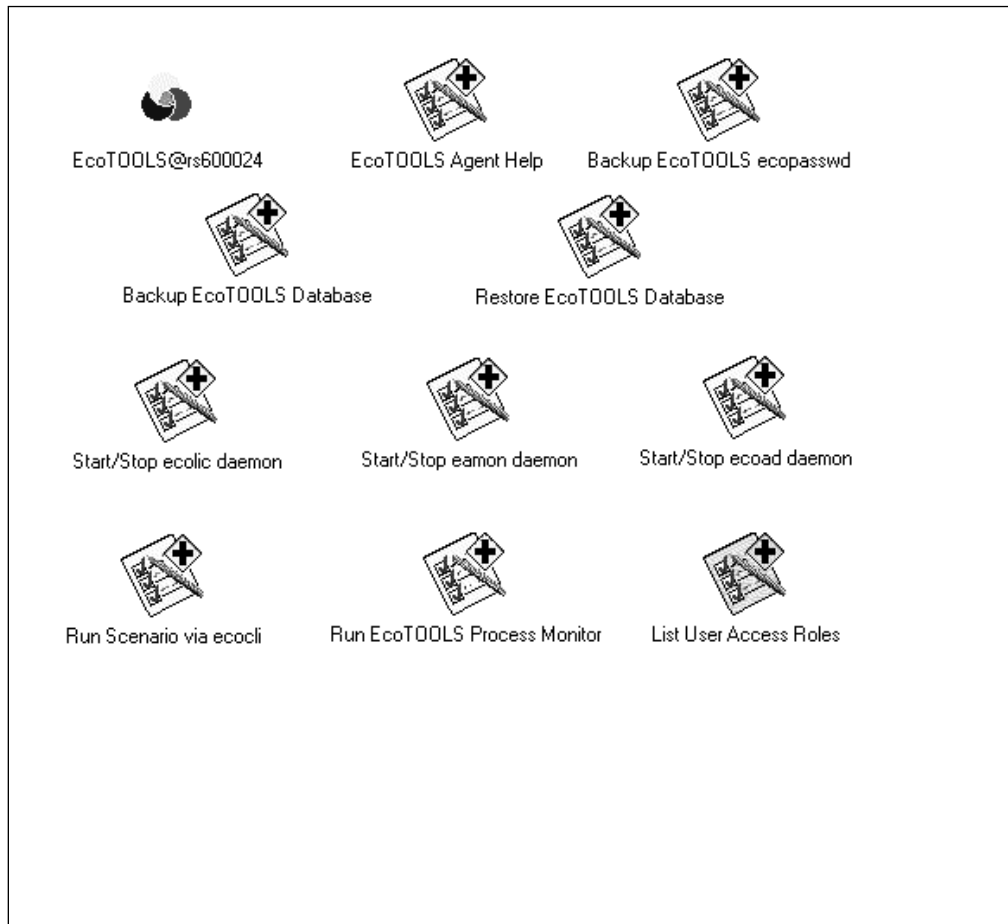


Figure 220. Components to Maintain EcoTOOLS

The following components exist to maintain EcoTOOLS directly from the TME 10 Desktop:

- EcoTOOLS: Use this icon to start the process monitor or the Management Console from the TME 10 Desktop.
- EcoTOOLS Agent Help: Get a detailed description of certain EcoTOOLS agents. This task is executed on the management station by default.
- Backup EcoTOOLS ecopasswd: Use this task to back up the EcoTOOLS password file ecopasswd which is located in the directory `/usr/eco/etc`.
- Backup EcoTOOLS Database: Task to back up the EcoTOOLS database.
- Restore EcoTOOLS Database: Task to restore the EcoTOOLS database.
- Start/Stop ecolic daemon: Task to maintain the EcoTOOLS license server. This daemon controls the number of running EcoTOOLS desktops. By default, this task is defined as an endpoint in the EcoTOOLS License Station's profile manager.
- Start/Stop eamon daemon: Task to maintain the EcoTOOLS agent monitoring process (which is responsible for the EcoTOOLS scenario requests) running properly. By default, this task is defined as an endpoint in the EcoTOOLS Management Station's profile manager.
- Start/Stop ecoad daemon: Task to control the EcoTOOLS dispatcher process, which is responsible for sending and receiving scenario requests. By default,

this task is defined as an endpoint in the EcoTOOLS Agent Station's profile manager.

- Run Scenario via ecocli: Execute a predefined EcoTOOLS scenario on the EcoTOOLS agents. By default, this task is defined as an endpoint in the EcoTOOLS Agent Station profile manager.
- Run EcoTOOLS Process Monitor: Task to start the GUI of the EcoTOOLS process monitor, which displays real-time information about UNIX processes. By default, this task is defined as an endpoint in the EcoTOOLS Agent Station profile manager.
- List User Access Roles: Display user information EcoTOOLS uses to access monitoring resources, such as databases.

Refer to 4.12.4, “EcoTOOLS Plus Module for TME 10 Tasks to Manage EcoTOOLS” on page 272 for a more detailed description of these tasks.

4.12 Examples Using EcoTOOLS Plus Module for TME 10 in Our Environment

Now that we have installed EcoTOOLS and EcoTOOLS Plus Module for TME 10 in our environment and have described the main components of EcoTOOLS Plus Module for TME 10, we show a number of examples in which we use these components.

We show the following examples:

- Distribution of the EcoTOOLS client software using TME 10 Software Distribution
- Monitoring EcoTOOLS with TME 10 Distributed Monitoring
- Event Management using TME 10 Enterprise Console
- Operational TME 10 tasks for handling EcoTOOLS

4.12.1 Deploying EcoTOOLS Using TME 10 Software Distribution

The distribution of the EcoTOOLS client software to the agent machines is the first of our examples using EcoTOOLS Plus Module for TME 10.

This task consists of the following steps:

- File package generation for all necessary machine types
- Distribution of the file packages
- Test of the installation

These steps are covered in the following sections.

4.12.1.1 File Package Generation for EcoTOOLS

In this section we explain the file package generation process. This step must be repeated for each operating system type available in the system environment. During the process references to several software package files and scripts are included into the package. We describe the structure of the file package at the end of this section.



Configure Agent Station Filepack




Figure 221. The Configure Agent Station Filepack Icon

Before we start the file package generation process we want to make sure that the correct task endpoint is specified. Select **Modify Job...** from the context menu of the the Configure Agent Station Filepack icon as shown in Figure 221.

The following window will appear:

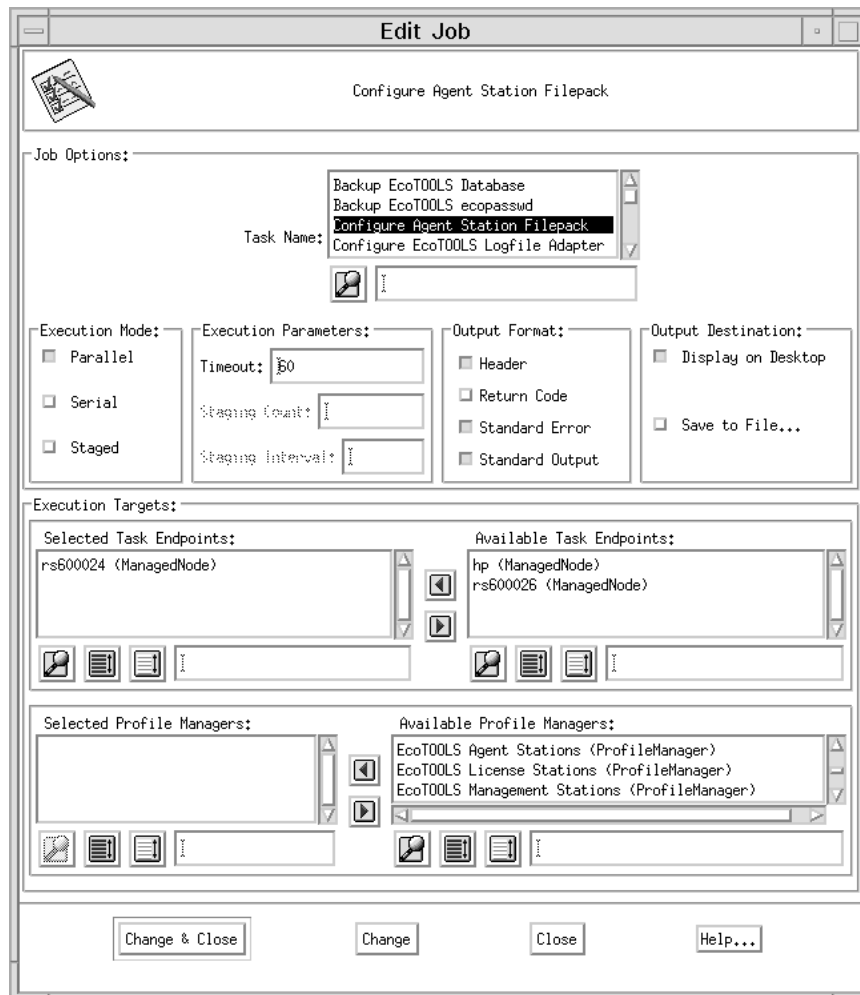


Figure 222. Edit Job Window

In the Edit Job window (Figure 222) enter the correct task endpoint in the Execution Targets section. The task endpoint must be the EcoTOOLS Management station, rs600024 in our example.



Configure Agent Station Filepack




Figure 223. Select Run Job from the Configure Agent Station Filepack Icon

The file package generation can now be initiated. Select **Run Job...** from the Configure Agent Station Filepack icon as shown above.

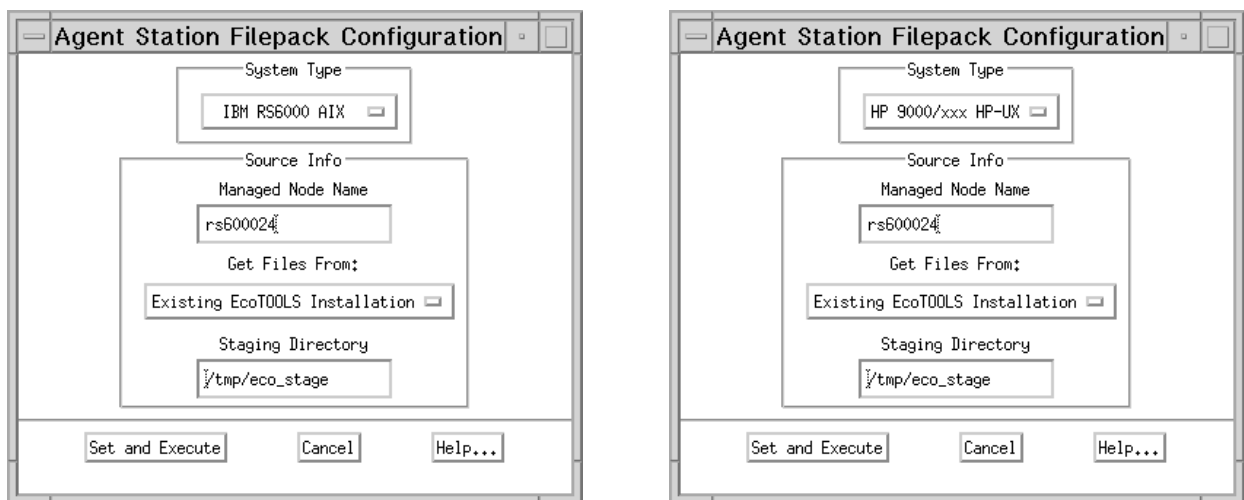


Figure 224. Agent File Package Generation for Different Machine Types

Figure 224 shows the Agent File Package Generation window for different machine types: IBM RS6000 running AIX and HP 9000 running HP-UX. The following parameters must be specified in this dialog:

- The machine type is selected via a pull-menu. Choose the correct type in the System Type field.
- The name of the managed node where the EcoTOOLS source binaries are located.
- The source location of the EcoTOOLS binaries is selected via the pull-down menu in the Get Files From field. Select either **Existing EcoTOOLS Installation** to take an existing EcoTOOLS directory or CD ROM to copy the files from a mounted CD-ROM file system.
- The name of the staging directory that holds the EcoTOOLS image during the installation. The default value is /tmp/eco_stage. Be sure to have enough space in this directory.

Click the **Set and Execute** button when finished entering the values to start the file package generation.

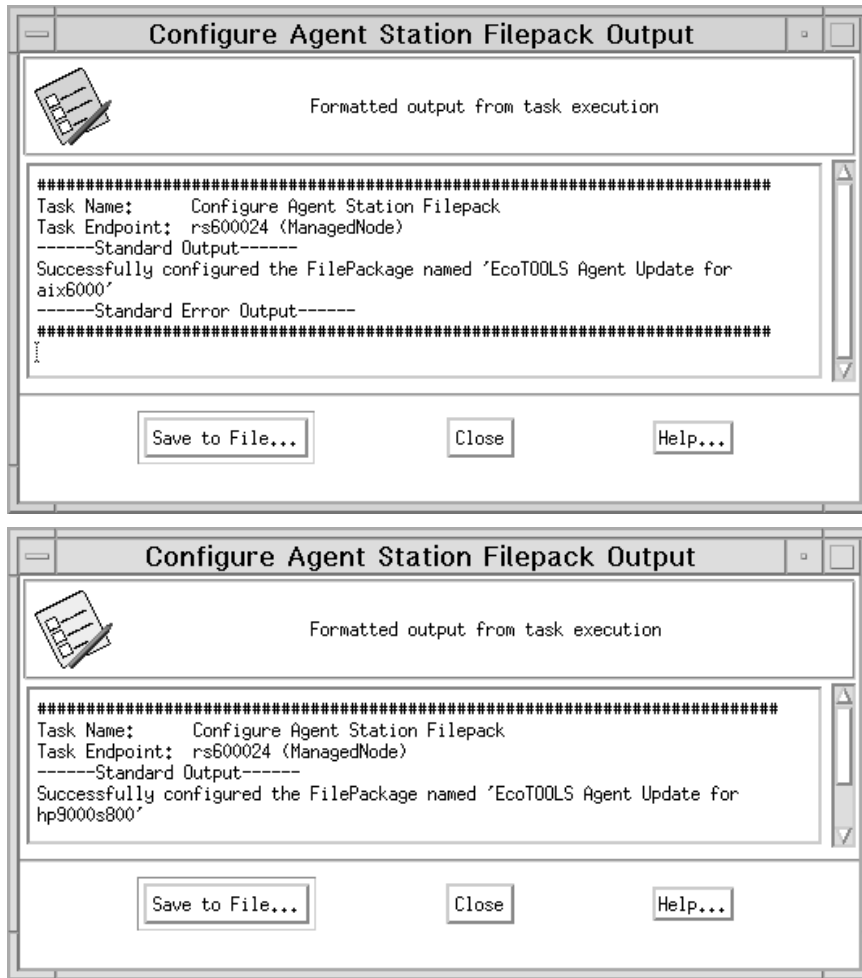


Figure 225. File Package Output for AIX and HP-UX Machines

Under the assumption that you have chosen to receive the task output on the desktop in the Edit Job window (Figure 222 on page 228) the system opens the Configure Agent Station Filepack Output window. In Figure 225 you see the output for our two machine types aix6000 (rs600024) and hp9000s800 (hp).

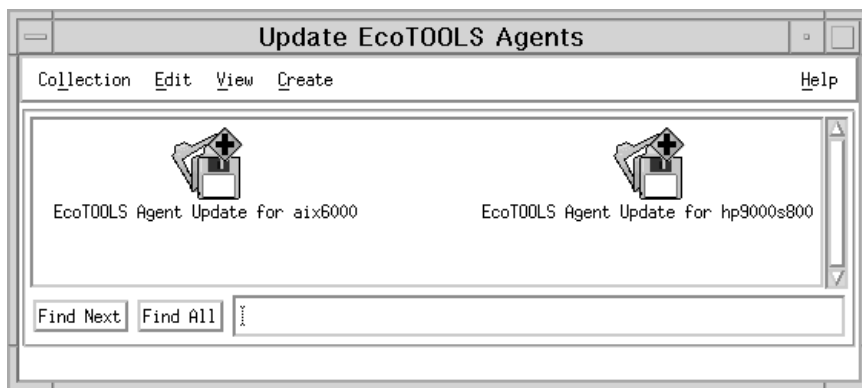


Figure 226. File Packages for AIX and HP Machines

As soon as the tasks complete the new file packages are displayed in the Update EcoTOOLS Agents window shown in Figure 226.

To examine the file package structure select **Properties...** from the pop-up menu of one of the EcoTOOLS Agent Update icons.

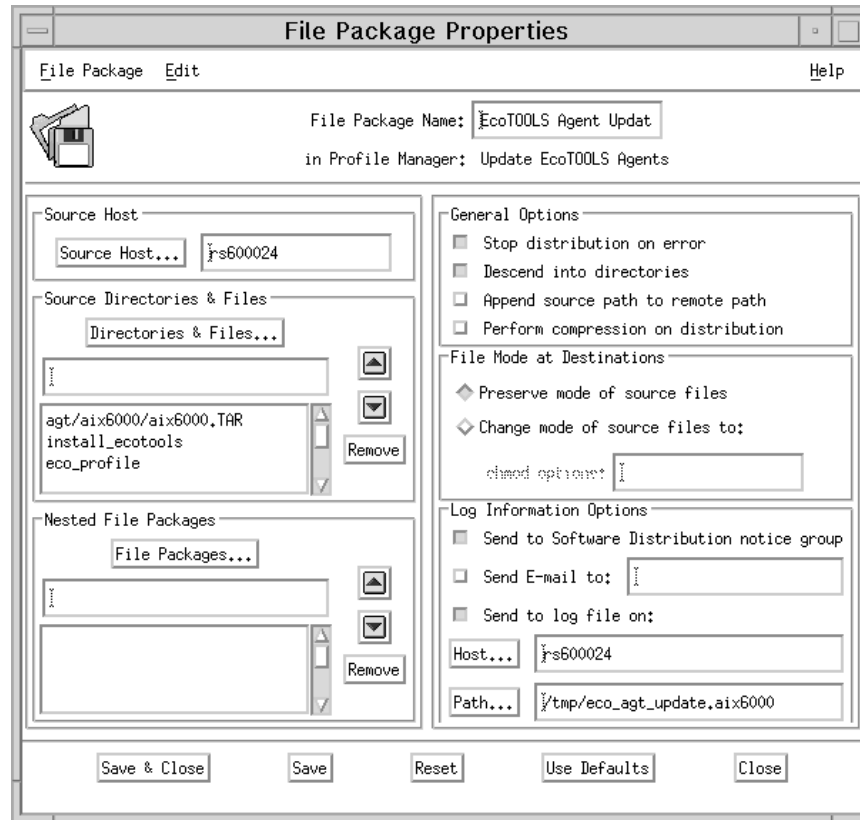


Figure 227. EcoTOOLS File Package Properties

Figure 227 shows the File Package Properties window for the AIX file package. In the Source Host section you see the managed node containing the software binaries, rs600024 in our case.

The location on this host and their names are specified in Source Directories & Files section of the window.

The following files should be included:

- /usr/eco/agt/aix6000/aix6000.TAR: The EcoTOOLS binaries
- /usr/eco/agt/aix6000/eco_profile: The EcoTOOLS profile
- /usr/eco/agt/install_ecotools: The installation script

These files are transferred to the agent machines during the distribution process. Be sure that the Source Host option is set in Get Program From section of the Unix Options window. You can get to this window by selecting **Edit** from the menu bar in the File Package Properties window and selecting **Platform Specific Options -> Unix Options** from the pull-down menu.

If you make any changes click **Save & Close** or select **Save** from the File Package menu of the File Package Properties window to make your changes permanent. The file package contains no nested file packages. The log file is written to the file /tmp/eco_agt_update.aix6000 as specified in the Log Information Options section.

To see the UNIX-specific options select **Platform-Specific Options -> Unix Options...** from the **Edit** pull-down menu of the File Package Properties window. This will open the File Package Unix Options window. Select **Before Distribution** or **After Distribution** in the Run a Program section and you will see that the following scripts will be called prior to and after the distribution process respectively:

- /usr/local/Tivoli/bin/generic_unix/TME/PLUS/ECO/eco_agt_fp_before.sh
- /usr/local/Tivoli/bin/generic_unix/TME/PLUS/ECO/eco_agt_fp_after.sh

To examine the definitions in the Tivoli object database use the following TME 10 CLI commands.

```
root@rs600024[/# wlookup -r FilePackage -a
```

In our environment, this produces the following output:

```
EcoT00LS Agent Update for aix6000 1576275774.1.946#FilePackage::FpoCore#  
EcoT00LS Agent Update for hp9000s800 1576275774.1.948#FilePackage::Fpore#
```

To see a more detailed definition of one of the above listed file packages, you can use the following command:

```
root@rs600024[/# wexpftfp -c -f /tmp/fp.out \  
"@EcoT00LSAgent Updadata for aix6000"
```

In our environment, the output file fp.out is shown in the following figure:

```

#*TFP-v2.01      Tivoli Filepack (version v2.01)
....
list_path=
nested_first=n
src_relpath=/tmp/eco_stage
file_cksums=n
unix_platform_prefix=/tmp/eco_agt_link
....
log_host=rs600024
log_file=/tmp/eco_agt_update.aix6000
...
src_before_prog_path=/usr/local/Tivoli/bin/generic_unix/TME/PLUS/ECO/eco
_agt_fp_src_before.sh
.
unix_before_prog_path=/usr/local/Tivoli/bin/generic_unix/TME/PLUS/ECO/ec
o_agt_fp_before.sh
unix_before_prog_from_src=n
unix_before_input_path=
unix_before_input_from_src=n
unix_before_skip_non_zero=y
unix_before_as_uid=
unix_after_prog_path=/usr/local/Tivoli/bin/generic_unix/TME/PLUS/ECO/eco
_agt_fp_after.sh
unix_after_prog_from_src=n
unix_after_input_path=
unix_after_input_from_src=n
unix_after_as_uid=
unix_removal_prog_path=
unix_removal_prog_from_src=n
unix_removal_input_path=
unix_removal_input_from_src=n
unix_removal_as_uid=
unix_commit_prog_path=
unix_commit_prog_from_src=n
unix_commit_input_path=
unix_commit_input_from_src=n
unix_commit_as_uid=
...
prog_env="TARGET='aix6000' SRC_MEDIA='client' SRC_STAGING='/tmp/eco_stag
e CHECK_MEMSIZE='2' CHECK_DISKSIZE='12000' DEST_PATH='/tmp/eco_agt_link'"
...
%
agt/aix6000/aix6000.TAR
install_ecotools
eco_profile
%

```

Figure 228. Definition of the EcoTOOLS Agent Update for aix6000 File Package

Figure 228 shows the definition of the file package as defined in the Tivoli database. At the top the scripts called before and after the software distribution is defined. Additionally, a script called `eco_agt_fp_src_before.sh` is defined. In the program environment you can see the variables that are passed to the installation scripts. At the end you can see the included files that have been mentioned before in the File Package Properties window.

4.12.1.2 File Package Distribution for EcoTOOLS

As soon as you have finished the file package generation for all machine types in your environment you are ready to distribute these file packages to the agent stations and install the software on the agents. In order to fulfill this process you need TME 10 Software Distribution installed on your managed nodes.




Figure 229. Subscribers for AIX File Package

Before you begin with the software distribution you must set the subscriber list of each file package. Open the Subscribers window by selecting **Subscribers...** from the file package icon in the Update EcoTOOLS Agents window and select the appropriate subscribers which may be profile managers or single hosts. We prepared separate profile managers for every machine type. Each profile subscribes to the corresponding file package.

When you have set the subscriber list you can distribute the file package.

To do so, click on the file package icon with the right mouse button and then select **Distribute...** from the pull-down menu as shown in the following figure.

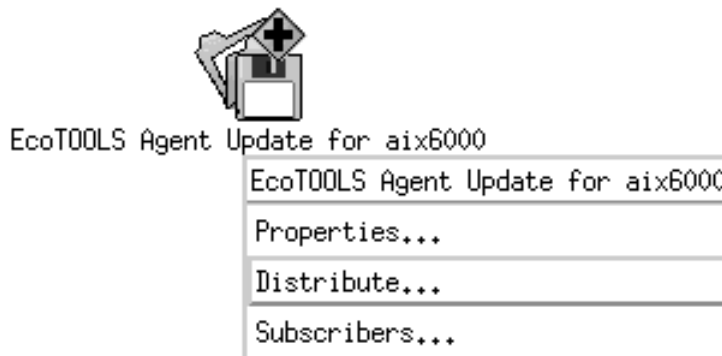


Figure 230. EcoTOOLS File Package Distribution

Once the distribution is finished, open the Tivoli Notice Board and you should see a notice similar to the following in the TME 10 Software Distribution Notice Group.



Figure 231. TME 10 Software Distribution Notice Group

You can also choose to write the distribution result to a log file.

```

File Package: "EcoTOOLS Agent Update for aix6000"
Operation:    install (m=5)
Finished:    Fri Dec 12 17:49:39 1997
-----
Source messages:
starting script:
/usr/local/Tivoli/bin/generic_unix/TME/PLUS/ECO/eco_agt_fp_src_before.sh
script complete: status=0
-----
rs600026:SUCCESS
temp script: unix_before: /tmp/unix_cKdXib
starting script: /tmp/unix_cKdXib
script complete: status=0
starting script:
/usr/local/Tivoli/bin/generic_unix/TME/PLUS/ECO/eco_agt_fp_after.sh
script complete: status=0
=====

```

Figure 232. Log File of EcoTOOLS File Package Distribution

Figure 232 shows the contents of the log file written during the software distribution. This log file was specified in the File Package Properties window in Figure 227 on page 231.

4.12.2 Monitoring EcoTOOLS Using TME 10 Distributed Monitoring

EcoTOOLS Plus Module for TME 10 provides a set of custom monitors for TME 10 Distributed Monitoring to monitor the different components of the EcoTOOLS installation, which are the Management Stations, Agent Stations and License Stations.

The primary function of the EcoTOOLS monitors is to insure availability of the EcoTOOLS application. Hence, the monitors provided can be used to monitor components that are relevant for the health of the EcoTOOLS application, such as daemon status, file system utilization, etc.

In this section we show an example of how to use the custom TME 10 Distributed Monitoring monitors for EcoTOOLS

In order to do so, we select the **EcoTOOLS Management Station Monitors** icon in the EcoTOOLS Plus Module for TME 10 window with the right mouse button and select **Properties...** from the pull-down menu.

The following window appears.

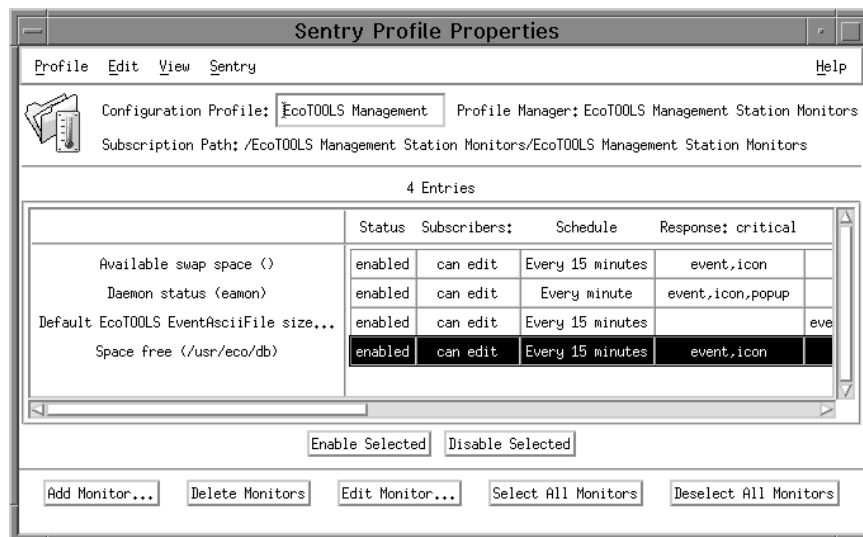


Figure 233. Sentry Profile Properties Window

This window displays the custom monitors for EcoTOOLS Management Stations. As you can see, there are monitors to check the available swap space, the EcoTOOLS daemon status, etc.

The monitoring profiles EcoTOOLS Agent Station Monitors, EcoTOOLS License Station Monitors and EcoTOOLS Central Monitors contain specific monitors for other EcoTOOLS components.

In this example we want to examine the Space free (/usr/eco/db) monitor.

We select the appropriate row as indicated in the above figure and then click the **Edit Monitor...** button.

This will open the Edit Sentry Monitor window as shown below.

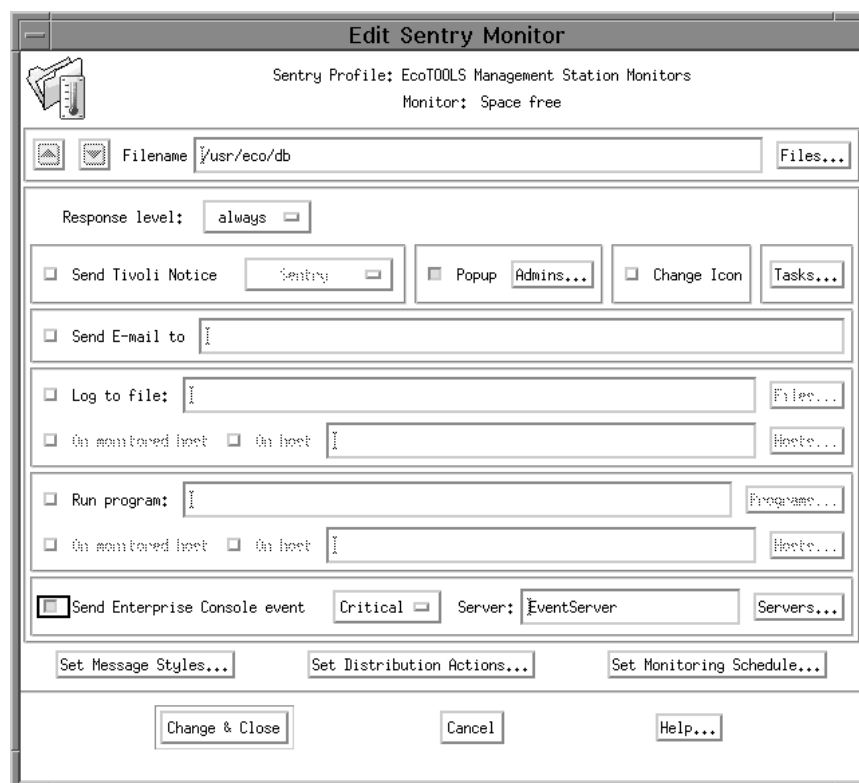


Figure 234. Edit Sentry Monitor Window

In order to check the output of the monitor we set the Response level field to always. This will trigger the specified action whenever the monitor is invoked.

In order to see the result of the monitor on the Tivoli desktop, we click on the **Popup** check box and then select the **Admins...** button. In the window that pops up we select the Tivoli Root administrator in order to display the pop-up window on that user's desktop.

We also check the **Send Enterprise Console event** box. This will send a TEC event, whenever the monitor meets the response level, which is Always in our example.

However, you should notice that these events will not show up in TEC until you have configured the event server to recognize events from EcoTOOLS as described in 4.12.3, "Event Management for EcoTOOLS Using TME 10 Enterprise Console" on page 241.

In order to test the monitor we also want to change the predefined monitoring interval from 15 minutes to 1 minute. To do so we select the **Set Monitoring Schedule...** button.

The following window appears:

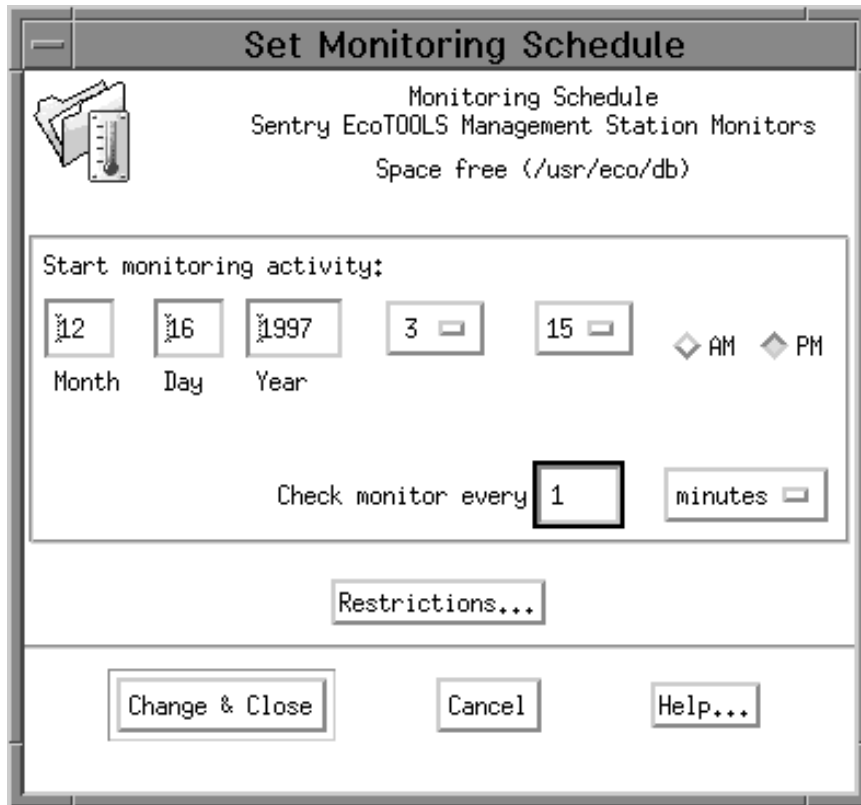


Figure 235. Set Monitoring Schedule Window

We set the current date in the Month/Day/Year fields and then enter 1 in the Check monitor every field. When finished we select the **Change & Close** button.

Back in the Edit Sentry Monitor window we select **Change & Close** again.

This returns us to the Sentry Profile Properties window.

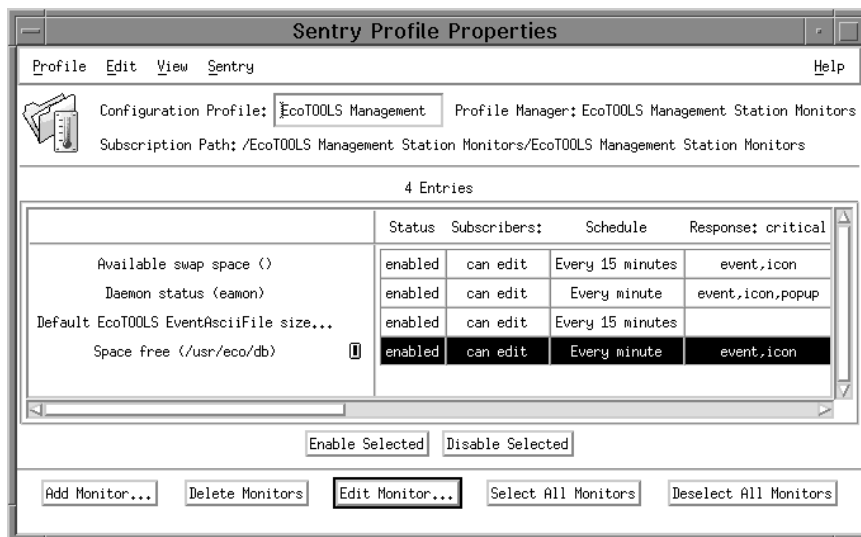


Figure 236. Sentry Profile Properties Window

The symbol next to the Space free (/usr/eco/db) monitor indicates that the profile needs to be saved.

In order to do so we select **Profile** from the menu bar and then **Save** from the pull-down menu. After that we can close the Sentry Profile Properties window.

In order to activate the monitor, we need to distribute the monitoring profile to our Management Station (rs600024). We select the profile icon with the right mouse button and then select **Distribute...** from the pull-down menu.

This will pop up the Distribute Profiles window. We select the **Distribute Now** button in order to distribute the profile. By default, this will distribute the profile to the predefined subscriber, which is rs600024 in our example. If you want to modify the subscribers before distributing the profile, you can select the **Subscribers...** option from the pull-down menu of the profile icon.

In the Tivoli desktop main window you will get a message notifying you that the profile has been distributed.

After a short period of time, the following window pops up at our Tivoli desktop.

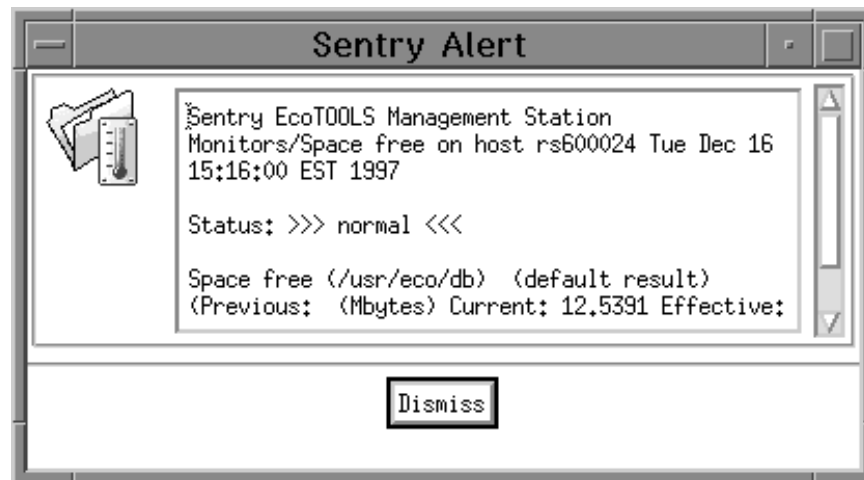


Figure 237. Sentry Alert Window

You can see that the current value for the free space in /usr/eco/db is 12.5391, which is the size in MegaBytes. As we specified a one minute interval and a Response level of always we will keep getting this message every minute until we modify and re-distribute the monitor.

Apparently this was just useful for testing the result of the monitor. You would now customize the monitor to only trigger an action when a condition is met that requires an action, for example, when the free space drops below 1 MB.

Also, we have specified to send a TEC event whenever the monitor triggers. These events will only appear in TEC after you have set up the TEC server correctly as described in 4.12.3, "Event Management for EcoTOOLS Using TME 10 Enterprise Console" on page 241.

In our example, TEC is already configured, so we get the following display in the TEC console window.

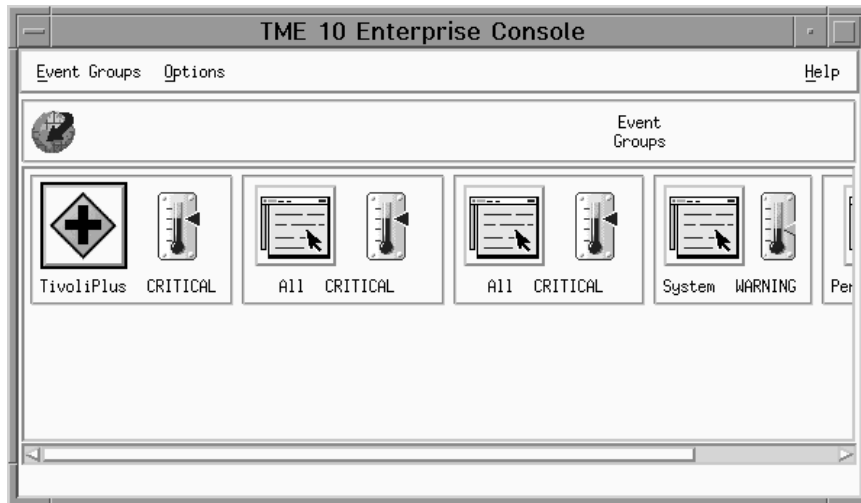


Figure 238. TME 10 Enterprise Console Window

The above window appears when you double-click on the **Root_rs600024-region** icon (the TEC console) in the Tivoli desktop main window.

The indicator displays a level of CRITICAL as this is the severity level for the TEC event we specified in the monitor.

In order to see the EcoTOOLS events, we click on the **TivoliPlus** icon. This will display the TivoliPlus window.

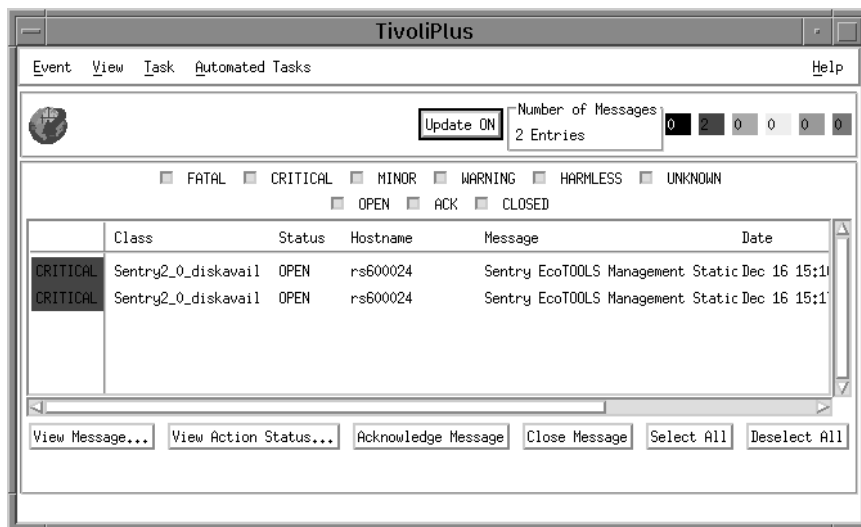


Figure 239. Tivoli Plus Window

You can see two new events, both triggered by our monitor. To see the event details we select one of the events and then click the **View Message...** button.

The Event Group Message Viewer window will appear.

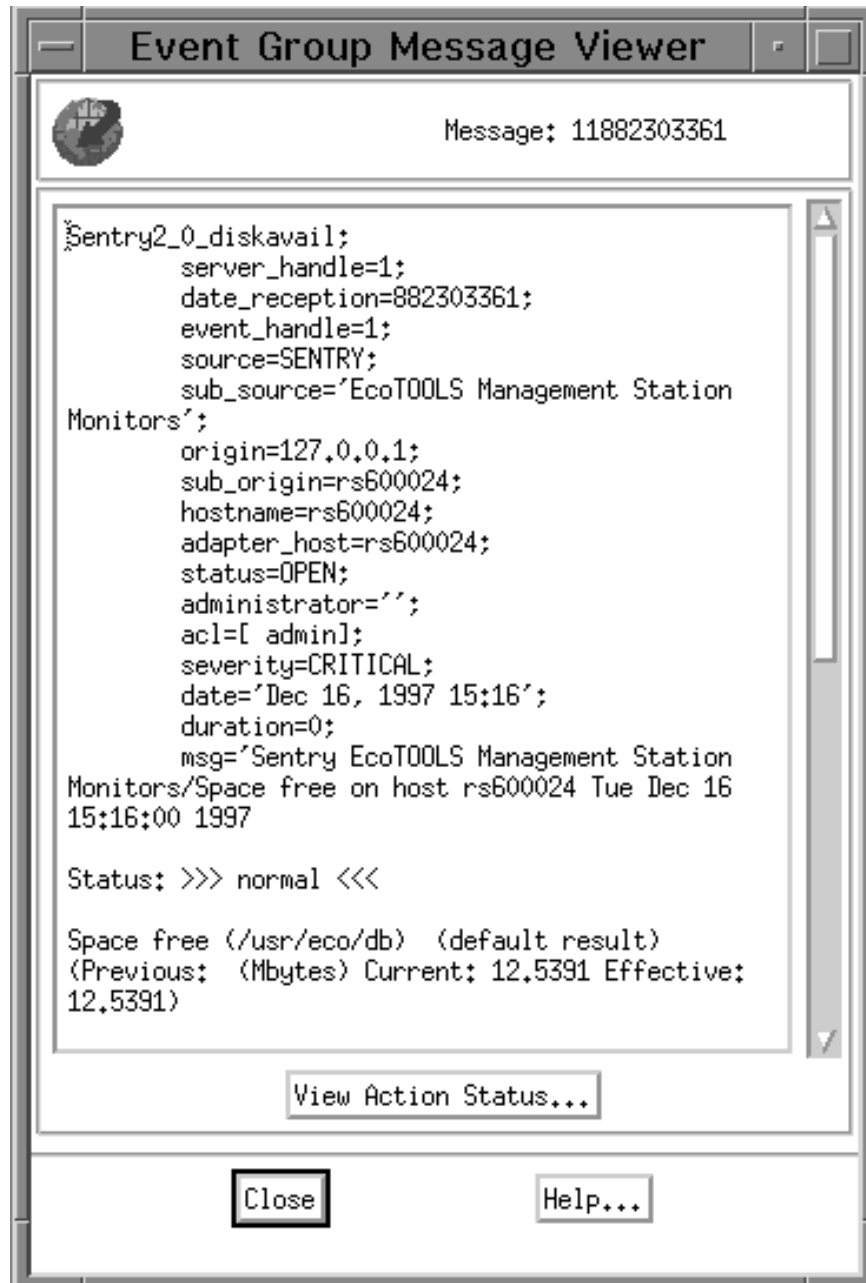


Figure 240. Event Group Message Viewer Window

As you can see, we get a TEC event with detailed information about what caused the exception.

4.12.3 Event Management for EcoTOOLS Using TME 10 Enterprise Console

In the previous section we explained the configuration of TME 10 Distributed Monitoring and the installation of the monitors on the managed nodes.

The TME 10 Enterprise Console is another possibility to perform event management from a central point. TME 10 Enterprise Console can be used to perform asynchronous monitoring, that is raise exceptions when something happens as opposed to TME 10 Distributed Monitoring which performs

synchronous monitoring, that is, a parameter is regularly checked against certain thresholds.

The advantage of TEC is the ability to display events coming from various sources. This simplifies the event management task enormously.

In this section we first describe how to configure the event server. In the second step we explain the installation of the event adapters on the Tivoli managed nodes. We then verify that each step has been accomplished properly (4.12.3.2, “Testing the Event Server” on page 247 and 4.12.3.4, “Testing the EcoTOOLS Adapter” on page 265).

4.12.3.1 Setting Up the Event Server

The EcoTOOLS Plus Module for TME 10 contains a task to set up the TEC event server correctly. That is, the task will import the necessary BAROC files into a TEC rule base and therefore enable the event server to process events coming from EcoTOOLS.



Setup T/EC EventServer for EcoTOOLS



Figure 241. Select Run Job from the Setup T/EC EventServer Icon

The setup of the TEC is initiated by selecting **Run job...** from the the context menu of the Setup TEC EventServer icon (Figure 241) or simply double-clicking on this icon.

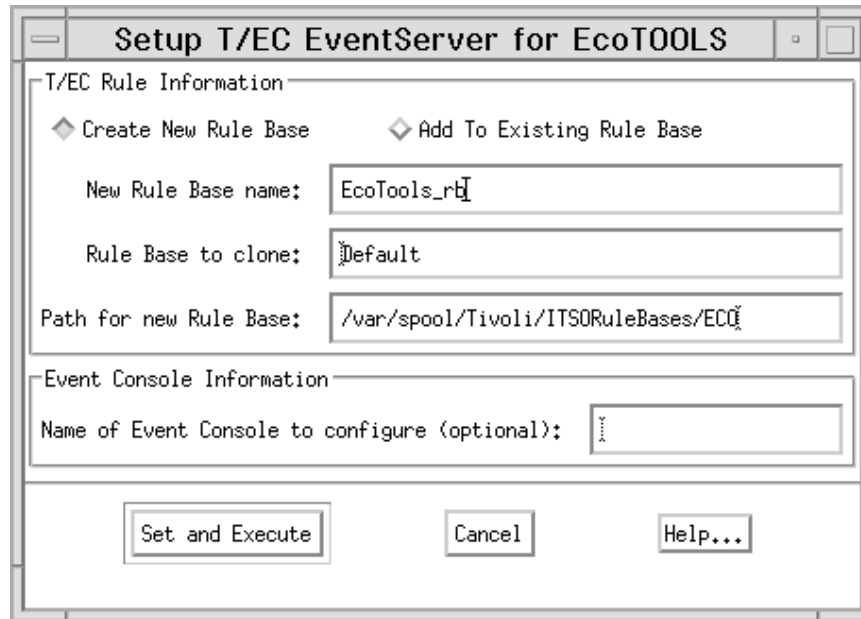


Figure 242. Setup TEC EventServer Parameters for EcoTOOLS

The Setup TEC EventServer for EcoTOOLS window appears. When configuring the event server you have the options either to create a new rule base or to utilize an existing rule base. Select one of the radio buttons at the top of the window representing the desired options. Depending on which option you choose you are asked to enter a number of parameter values.

We recommend you create a new rule base and obtain the parameter list as displayed in Figure 242. The parameters to specify are:

- The name of the new rule base.
- The name of of the rule base to be used as a template. In our case we use the Default rule base and leave this option unchanged.
- The rule base directory. All our rule bases a located under the same path. We use a new subdirectory to separate the the EcoTOOLS rule base from other rule bases.
- The name of the event console. If specified, the new event group will be assigned to this event console. This parameter is optional and the assignment of event groups can also be made later.

Select **Set and Execute** when you are finished.

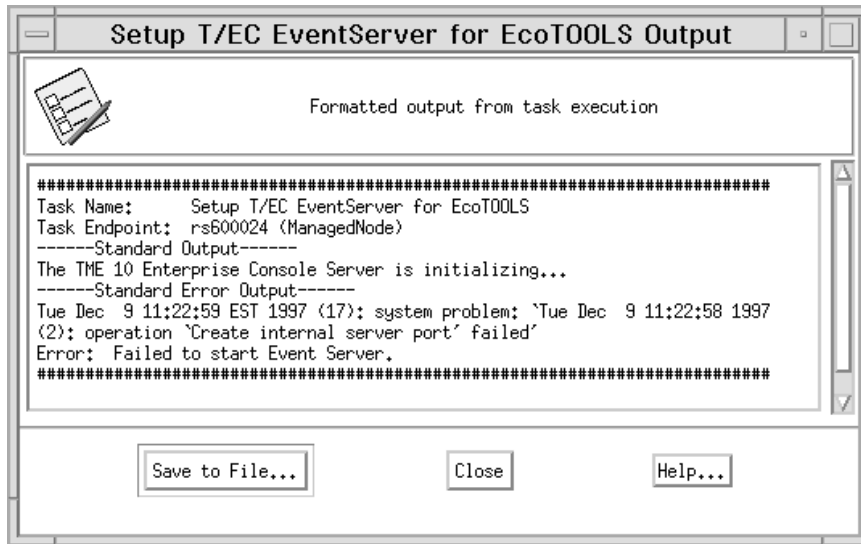


Figure 243. Setup EventServer Task Output

When the setup task terminates the output window appears as shown in Figure 243. As you see a problem is reported to stderr; the missing red arrow of the event server icon on the Tivoli desktop indicates the system is not able to restart the event server. The problem in our environment was that the Sybase server was not running when the task was executed and therefore the TEC server could not connect to it and failed to start.

Nevertheless, the import of the necessary BAROC files completes successfully, so the rule base will be ready for use.

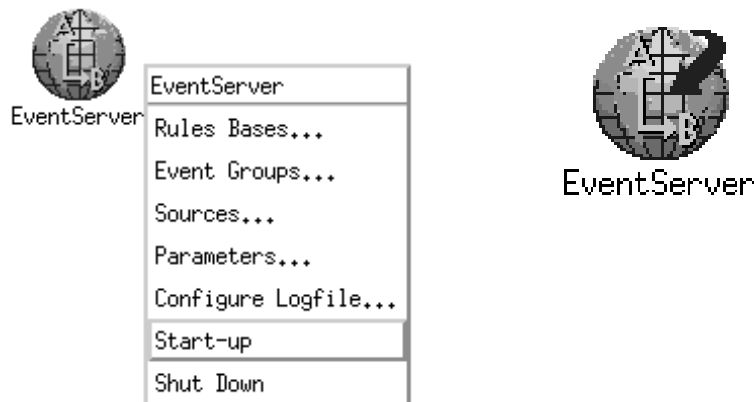


Figure 244. Manual Restart of the Event Server

In our example, we started the Sybase server and then try to start the event server manually.

To restart the event server select **Start-up** from the EventServer icon as shown in Figure 244. As soon as the event server is active again the red arrow will re-appear on the icon.

As mentioned before, despite this problem the setup process of the event server works successfully. To verify this you can check the rule base directory. The

system should have created a new subdirectory, ECO in our case, for the new rule base. Under this directory you will find three more directory entries: TEC_CLASSES, TEC_RULES and TEC_TEMPLATES.

```

root@rs600024[/var/spool/Tivoli/ITSORuleBases/ECO/TEC_CLASSES]# ls
.event_sum_slots      as400msg.baroc          tecad_nv6k.baroc
.load_classes         ecoplus.baroc           tecad_ov.baroc
.task_slots           root.baroc               tecad_snaevent.baroc
ECO_Monitors.baroc   tec.baroc                tecad_snm.baroc
ECO_via_Sentry2_0.baroc tecad_hpov.baroc        tecad_snmp.baroc
Sentry.baroc         tecad_logfile.baroc
Sentry2.0Monitors.baroc tecad_nt.baroc

```

Figure 245. BAROC Files of EcoTOOLS Rule Base

Figure 245 shows the contents of the TEC_CLASSES subdirectory which contains all event class descriptions (BAROC files) of the EcoTOOLS rule base.

```

root@rs600024[/var/spool/Tivoli/ITSORuleBases/ECO/TEC_RULES]# ls
.rule_builder.pro      tec_forward.conf
RCS                    tec_r..load_rules
aix4-r1                tec_r.ecoplus_corr.rls
ecoplus_corr.rls      tec_r.ecoplus_dup.rls
ecoplus_dup.rls       tec_r.ecoplus_mon.rls
ecoplus_mon.rls       tec_r.ecoplus_split.rls
ecoplus_split.rls     tec_r.ecoplus_split_cli.rls
ecoplus_split_cli.rls tec_r.log_default.rls
log_default.rls       tec_r.ov_default.rls
ov_default.rls        tec_r.tecad_snaevent.rls
rule_sets              tecad_snaevent.rls

```

Figure 246. Rule Files of EcoTOOLS Rule Base

The rule files stored in the TEC_RULES subdirectory (see Figure 246) contain evaluation rules for the EcoTOOLS event classes.

EcoTOOLS Plus Module for TME 10 provides a number of predefined correlation rules to automatically correlate events from EcoTOOLS. Use an ASCII editor to view the *.rls files to examine what the different rules look like.

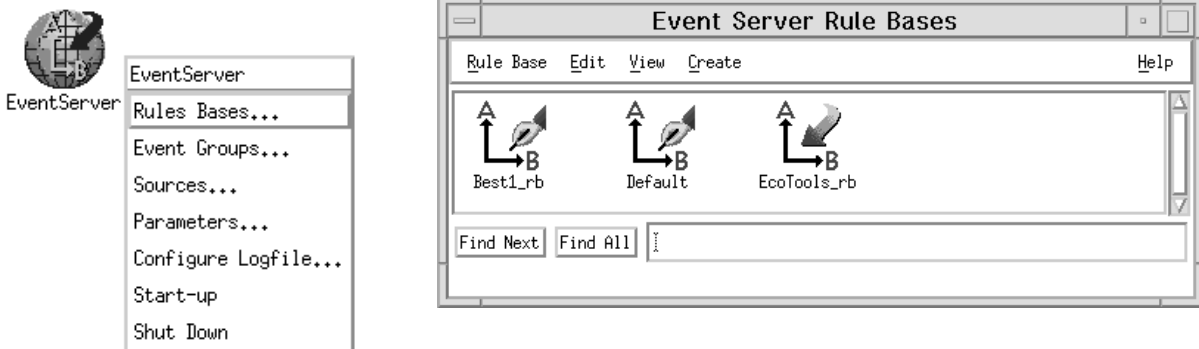


Figure 247. Rule Bases after Setup EcoTOOLS Event Server

When you select **Rule Bases...** from the pop-up menu of the EventServer icon the Event Server Rule Bases window appears as shown above. You will notice that an icon for the new EcoTOOLS rule base, named EcoTools_rb in our environment, has been added to the window.

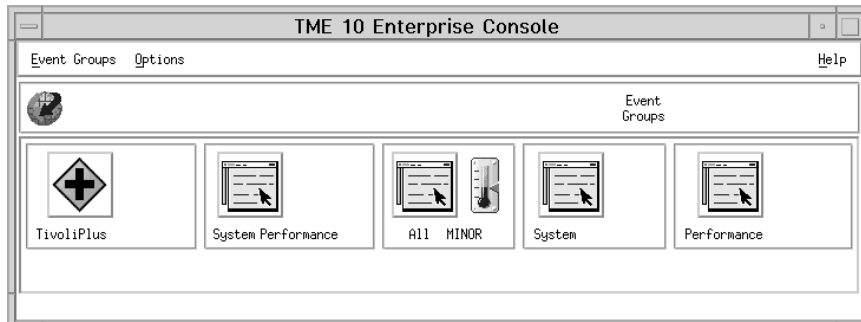


Figure 248. Event Groups after Setup EcoTOOLS Event Server

Figure 248 shows the TME 10 Enterprise Console window containing all assigned event groups. The TivoliPlus icon has been added to the console, containing filters for events from EcoTOOLS.

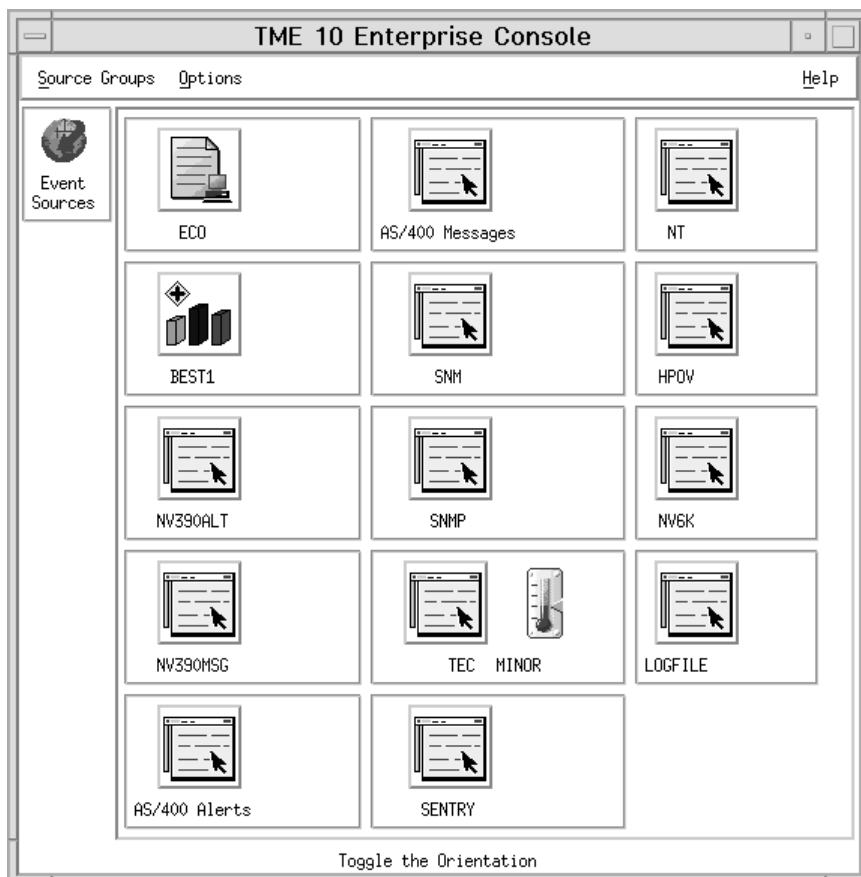


Figure 249. Event Sources after Setup EcoTOOLS Event Server

Figure 249 shows that the event source ECO has been added to the TME 10 Enterprise Console event sources.

4.12.3.2 Testing the Event Server

In order to verify that the TME 10 Enterprise Console is set up properly we use the TME 10 CLI command `wpostemsg` to send test events. The following two examples generate events for the event classes `Eco_Base_Event` and `ECO_via_Sentry2_0` and the event source `ECO`:

```
wpostemsg -m "Test EcoTOOLS rule base" Eco_Base_Event ECO
```

or

```
wpostemsg -m "Test EcoTOOLS rule base" ECO_via_Sentry2_0 ECO
```

The event classes can be determined by examining the BAROC files in the `TEC_CLASSES` subdirectory (see Figure 245 on page 245). The event class `Eco_Base_Event` is defined in the `ecoplus.baroc` file. The event class `ECO_via_Sentry2_0` is defined in the `ECO_via_Sentry2_0.baroc` file. This class is derived from the `Sentry2_0_Base` event class defined in the `Sentry.baroc` file. In addition, the file `ECO_Monitors.baroc` contains event class definitions for EcoTOOLS Plus Module for TME 10.

All three BAROC files are located in the `/usr/local/Tivoli/bin/generic_unix/TME/PLUS/ECO` directory and are shown below.

```

# $Id: ecoplus.baroc,v 1.2 1996/09/20 19:48:25 aharriso Exp $
#
# Component:    TivoliPlus for EcoTOOLS
#
# $Source: /tivoli/development/src/2.0/plus/ecotools/src/ecoplus.baroc,v$
# $Revision: 1.2 $
#
# Description:
#
# (C) COPYRIGHT Tivoli Systems, Inc. 1996
# Unpublished Work
# All Rights Reserved
# Licensed Material - Property of Tivoli Systems, Inc.
#
TEC_CLASS :
    Eco_Base_Event ISA EVENT
    DEFINES {
        source: default=ECO;
        hostname: default="N/A",dup_detect=yes;
        msg: default="N/A";
        severity: default = CRITICAL;
        orig_msg: STRING, default="N/A";
    };
END

TEC_CLASS :
    EcoLog_Event ISA Eco_Base_Event
    DEFINES {
        eco_user: STRING, default="N/A";
        eco_view: STRING, default="N/A";
        eco_agt_mach_ip_addr: STRING, default="N/A";
        eco_scenario: STRING, default="N/A";
        eco_start_date: STRING, default="N/A";
        eco_start_time: STRING, default="N/A";
        eco_interval: STRING, default="N/A";
        eco_agt: INTEGER, default=0,dup_detect=yes;
        eco_instance: INTEGER, default=0;
        eco_record_date: STRING, default="N/A";
        eco_agt_start_time: STRING, default="N/A";
        eco_record_time: STRING, default="N/A";
        eco_data_value: STRING, default="N/A";
        eco_data_unit: STRING, default="N/A";
        eco_severity: STRING, default="N/A";
    };
END

TEC_CLASS :
    EcoLog_nb_Event ISA EcoLog_Event;
END

TEC_CLASS :
    EcoLog_lb_Event ISA EcoLog_Event;
END

```

Figure 250. *ecoplus.baroc* File (Part 1 of 5)

```

TEC_CLASS :
    EcoLog_tb_Event ISA EcoLog_Event;
END
TEC_CLASS :
    EcoLog_bb_Event ISA EcoLog_Event;
END

TEC_CLASS :
    EcoLog_CLI_Event ISA Eco_Base_Event
    DEFINES {
        eco_user: STRING, default="N/A";
        eco_view: STRING, default="N/A";
        eco_agt_mach_ip_addr: STRING, default="N/A";
        eco_scenario: STRING, default="N/A";
        eco_start_date: STRING, default="N/A";
        eco_start_time: STRING, default="N/A";
        eco_interval: STRING, default="N/A";
        eco_agt: INTEGER, default=0,dup_detect=yes;
        eco_instance: INTEGER, default=0;
        eco_record_date: STRING, default="N/A";
        eco_agt_start_time: STRING, default="N/A";
        eco_record_time: STRING, default="N/A";
        eco_data_value: STRING, default="N/A";
        eco_data_unit: STRING, default="N/A";
        eco_severity: STRING, default="N/A";
    };
END

TEC_CLASS :
    EcoLog_CLI_nb_Event ISA EcoLog_CLI_Event;
END

TEC_CLASS :
    EcoLog_CLI_lb_Event ISA EcoLog_CLI_Event;
END

TEC_CLASS :
    EcoLog_CLI_tb_Event ISA EcoLog_CLI_Event;
END

TEC_CLASS :
    EcoLog_CLI_bb_Event ISA EcoLog_CLI_Event;
END

TEC_CLASS &colon,
    EcoLog_Cat_Event ISA Eco_Base_Event
    DEFINES {
        eco_user: STRING, default="N/A";
        eco_view: STRING, default="N/A";
        eco_agt_mach_ip_addr: STRING, default="N/A";
        eco_scenario: STRING, default="N/A";
        eco_start_date: STRING, default="N/A";
        eco_start_time: STRING, default="N/A";
        eco_interval: STRING, default="N/A";
        eco_agt: INTEGER, default=0,dup_detect=yes;
    };

```

Figure 251. *ecoplus.baroc* File (Part 2 of 5)

```

        eco_instance: INTEGER, default=0;
        eco_record_date: STRING, default="N/A";
        eco_agt_start_time: STRING, default="N/A";
        eco_record_time: STRING, default="N/A";
        eco_data_value: STRING, default="N/A";
        eco_data_unit: STRING, default="N/A";
        eco_severity: STRING, default="N/A";
    };
END
TEC_CLASS :
    EcoLog_Computer_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_Network_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_Oracle_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_Oracle_Financials_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_Sybase_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_Fault_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_Space_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_Performance_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_Security_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_Log_File_Management_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_General_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_System_Check_Event ISA EcoLog_Cat_Event;

```

Figure 252. *ecoplus.baroc* File (Part 3 of 5)

```

END

TEC_CLASS :
    EcoLog_Informix_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_Chargeback_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_Tuxedo_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_NoCat_Event ISA EcoLog_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_Cat_Event ISA Eco_Base_Event
    DEFINES {
        eco_user: STRING, default="N/A";
        eco_view: STRING, default="N/A";
        eco_agt_mach_ip_addr: STRING, default="N/A";
        eco_scenario: STRING, default="N/A";
        eco_start_date: STRING, default="N/A";
        eco_start_time: STRING, default="N/A";
        eco_interval: STRING, default="N/A";
        eco_agt: INTEGER, default=0,dup_detect=yes;
        eco_instance: INTEGER, default=0;
        eco_record_date: STRING, default="N/A";
        eco_agt_start_time: STRING, default="N/A";
        eco_record_time: STRING, default="N/A";
        eco_data_value: STRING, default="N/A";
        eco_data_unit: STRING, default="N/A";
        eco_severity: STRING, default="N/A";
    };
END

TEC_CLASS :
    EcoLog_CLI_Computer_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_Network_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_Oracle_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_Oracle_Financials_Event ISA EcoLog_CLI_Cat_Event;
END

```

Figure 253. *ecoplus.baroc* File (Part 4 of 5)

```

TEC_CLASS :
    EcoLog_CLI_Sybase_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_Fault_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_Space_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_Performance_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_Security_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_Log_File_Management_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_General_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_System_Check_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_Informix_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_Chargeback_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_Tuxedo_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    EcoLog_CLI_NoCat_Event ISA EcoLog_CLI_Cat_Event;
END

TEC_CLASS :
    Eco_Error_Event ISA Eco_Base_Event
    DEFINES {
        eco_error_msg: STRING, default="N/A";
    };
END

```

Figure 254. *ecoplus.baroc* File (Part 5 of 5)


```

TEC_CLASS :
    ECO_via_Sentry2_0 ISA Sentry2_0_Base;
END
TEC_CLASS :
    ECO_via_Sentry2_0_daemon ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_daemonct ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_zombies ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_host ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_oserv ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_filesize ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_filechk ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_filediff ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_fileperm ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_countstr ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_diskavail ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_diskused ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_diskusedpct ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_tivdbspace ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_inodes ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_inodesused ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_inodesusedpct ISA ECO_via_Sentry2_0;
END

```

Figure 255. ECO_via_Sentry2.0.baroc File (Part 1 of 3)

```

TEC_CLASS :
    ECO_via_Sentry2_0_netcoll ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_netcollpct ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_netinerr ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_netin ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_netouterr ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_netout ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_badnfs ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_badrpc ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_rpctmout ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_pageouts ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_swapavail ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_loadavg ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_ologins ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_ologintot ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_mailqlen ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_printstat ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_printjobs ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_printjobsize ISA ECO_via_Sentry2_0;
END

```

Figure 256. ECO_via_Sentry2.0.baroc File (Part 2 of 3)

```

END
TEC_CLASS :
    ECO_via_Sentry2_0_sasync ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_onsync ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_oscust ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_ncust ISA ECO_via_Sentry2_0;
END
TEC_CLASS :
    ECO_via_Sentry2_0_scust ISA ECO_via_Sentry2_0;
END

```

Figure 257. *ECO_via_Sentry2.0.baroc File (Part 3 of 3)*

```

TEC_CLASS :
    Eco_Custom_Monitors ISA Sentry2_0_Base;
END
TEC_CLASS :
    EcoEventfileSize ISA Eco_Custom_Monitors;
END
TEC_CLASS :
    EcoLicenses ISA Eco_Custom_Monitors;
END

```

Figure 258. *ECO_Monitors.baroc File*

As you can see from the above files, there are a large number of event classes defined for use by EcoTOOLS Plus Module for TME 10. This allows for a detailed classification of events from EcoTOOLS arriving at TEC.

There are two ways to check the arrival of the tests event. One is the TME 10 CLI command `wtdumpri`.

```

...
1 4679 0 881770495(Dec 10 11:14:55 1997)
### EVENT ###
Eco_Base_Event;
source=ECO;
msg='Test EcoTOOLS rule base';
origin=9.24.104.161;
END

### END EVENT ###
PROCESSED

1 4680 0 881770520(Dec 10 11:15:20 1997)
### EVENT ###
ECO_via_Sentry2_0;
source=ECO;
msg='Test EcoTOOLS rule base';
origin=9.24.104.161;
END

### END EVENT ###
PROCESSED

```

Figure 259. Excerpt from the WTDUMPRL Command Output

The wtdumprl command writes to stdout and can be called without any parameters. Figure 259 shows an excerpt of the output. You can see the event class in the first line of the event body. Further, the source, the message text and the origin are listed.

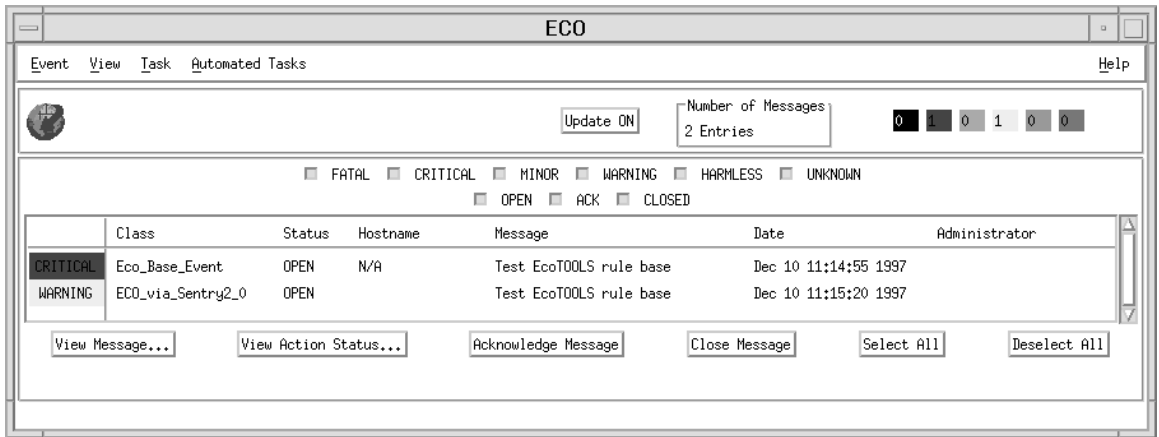


Figure 260. EcoTOOLS Test Events on the TME 10 Enterprise Console

In order to examine the result of the test events on the event console click on either the **TivoliPlus** icon in the event groups window (Figure 248 on page 246) or the ECO icon in the event source window (Figure 248 on page 246). The window in Figure 260 appears. In addition to the output of the wtdumprl command you can see the severity level assigned to the event classes.

4.12.3.3 Configuring the EcoTOOLS Event Adapter

In order to be able to forward events from EcoTOOLS to TEC, EcoTOOLS Plus Module for TME 10 utilizes the standard TEC logfile adapter. The logfile adapter is a generic TEC event adapter that can be used to parse any kind of structured ASCII logfile and generate events for TEC as a result.

The following figure illustrates how EcoTOOLS Plus Module for TME 10 and the TEC logfile adapter work together.

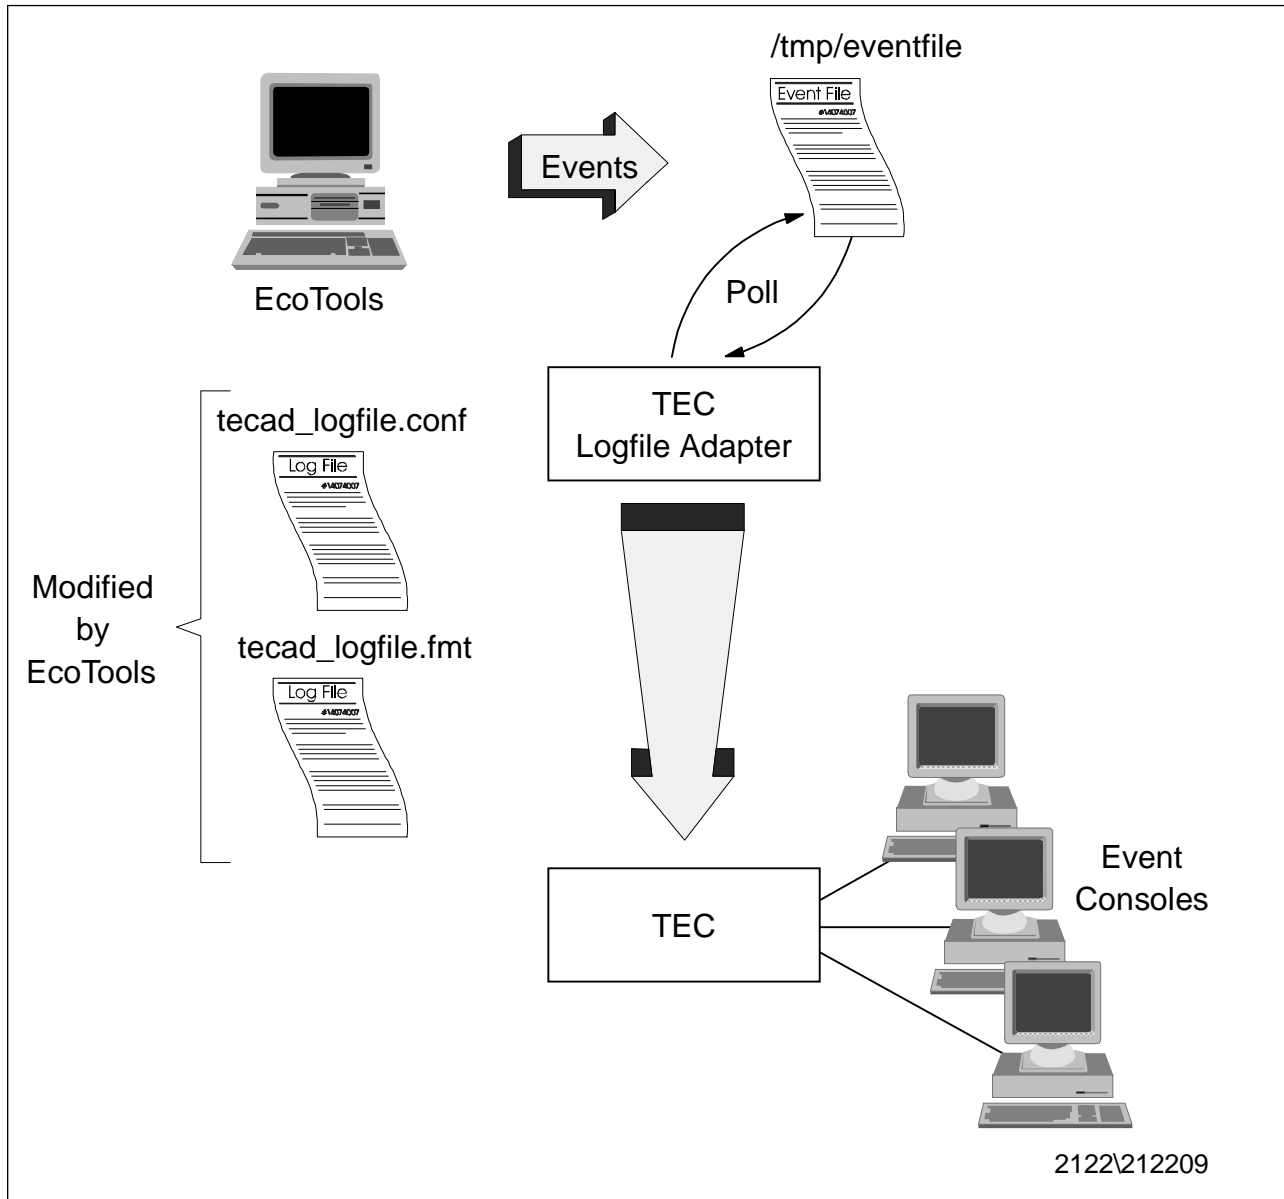


Figure 261. TEC Logfile Adapter and EcoTOOLS Plus Module for TME 10

EcoTools can write events to an event file, /tmp/eventfile by default.

Note

Be sure to uncomment the parameter for EventAsciiFile in the eco.conf file because only in this case, the event file (/tmp/eventfile) is produced.

EcoTOOLS Plus Module for TME 10 then customizes the standard TEC logfile adapter to monitor this file for events and translate them into TEC events. This customization involves the modification of the files `tecad_logfile.conf` and `tecad.logfile.fmt` which are configuration files for the TEC logfile adapter. In order to configure the EcoTOOLS Logfile adapter, double-click on the **Configure EcoTOOLS Logfile Adapter** icon in the EcoTOOLS Plus Module for TME 10 window.

This will open the following window:

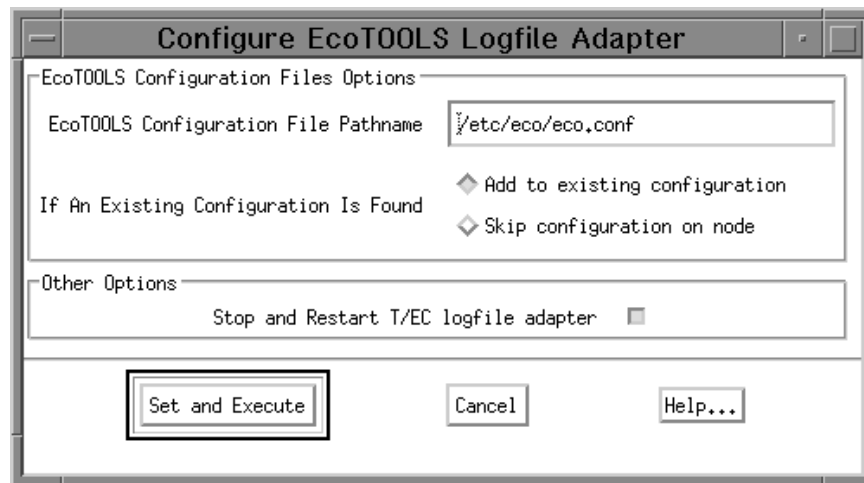


Figure 262. Configure EcoTOOLS Logfile Adapter Window

You can leave all the default options and just click the **Set and Execute** button. This will run the task on the default subscribers which are all machines in the EcoTOOLS Management Station's profile manager. If you want to execute the task on different machines, you can select the icon with the right mouse button and then select **Modify Job...** from the pull-down menu.

In our example, the tasks fails with the following message:

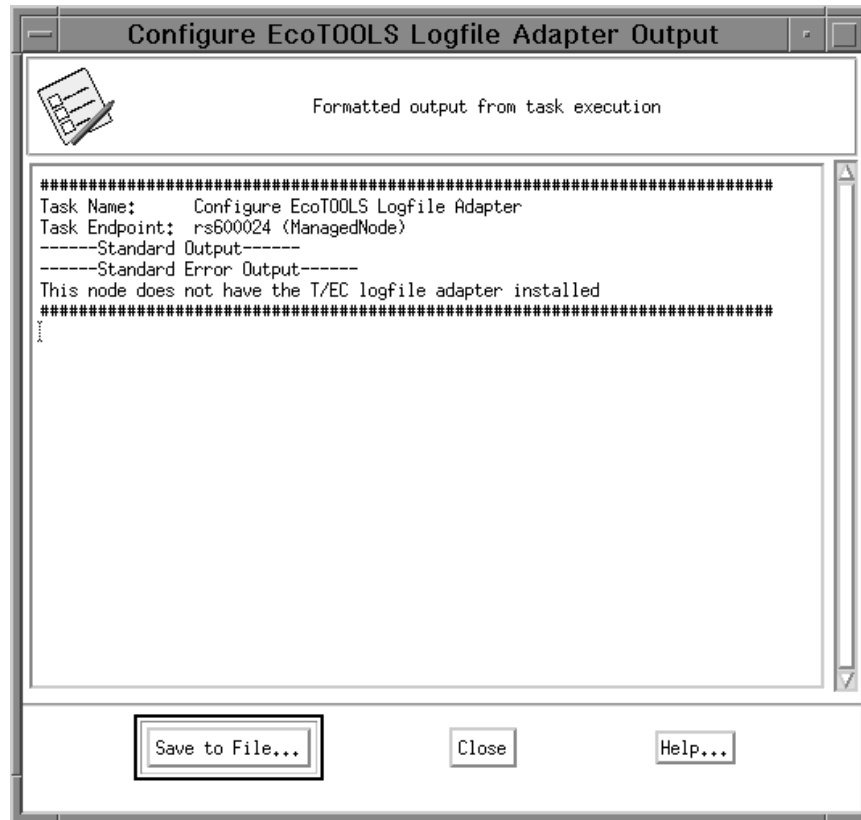


Figure 263. Error Message When Configuring Logfile Adapter

The message indicates that the TEC logfile adapter is not installed on our system. We need to install the TEC logfile adapter that comes on the TEC CD-ROM.

We insert and mount the TEC CD-ROM on rs600024 and open the Install Product window by selecting **Desktop** from the menu bar in the Tivoli desktop main window, then selecting **Install** from the pull-down menu and **Install Product...** from the submenu.

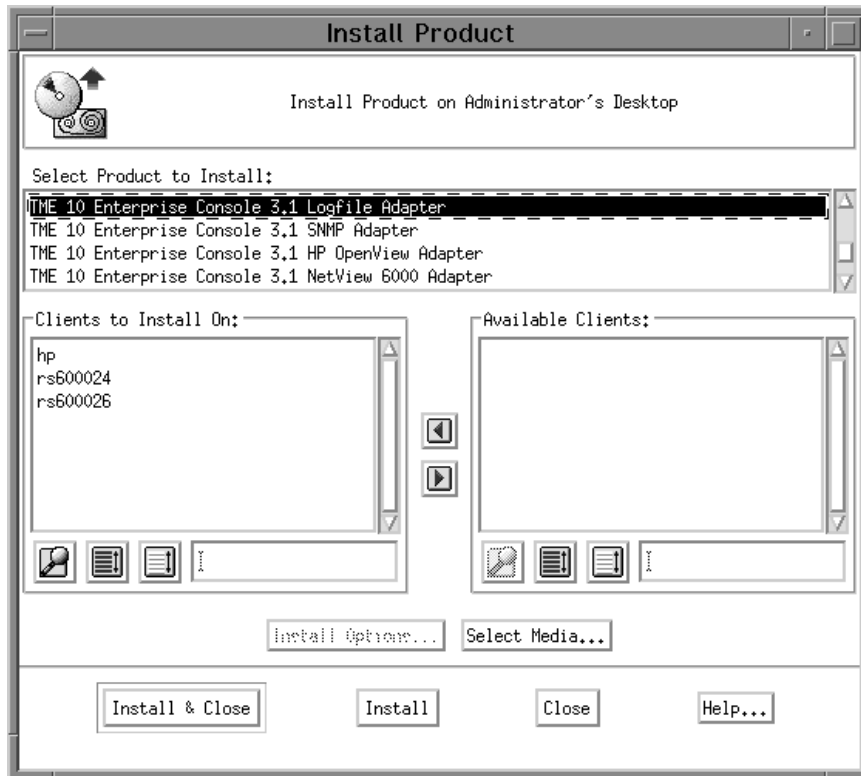


Figure 264. Install Product Window

We select the **Select Media...** button and point to the directory where the TEC CD-ROM is mounted. Back in the Install Product window, we see the TEC installation options in the Select Product to Install section.

We click on **TME 10 Enterprise Console 3.1 Logfile Adapter** and select the clients to install it on, all the managed nodes in our example. In order to start the installation we select **Install & Close**.

The Product Install window will appear.

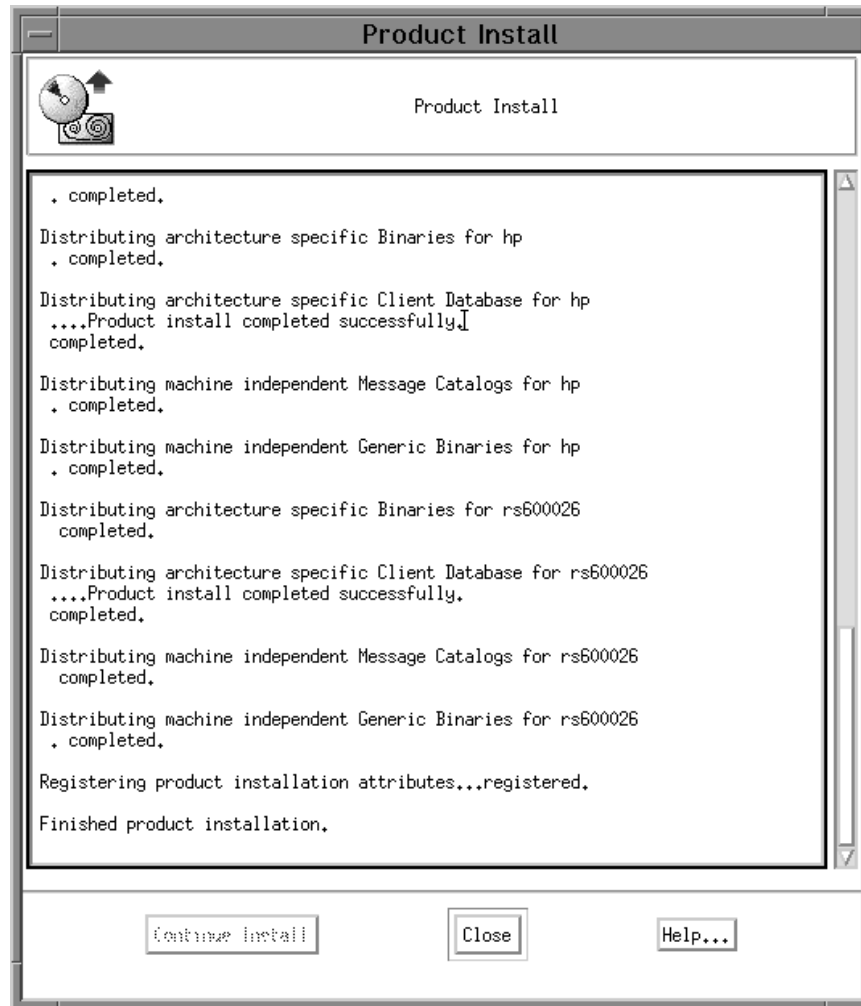


Figure 265. Product Install Window

After a short period of time, the successful completion of the install is indicated.

When starting the Configure EcoTOOLS Logfile Adapter task again, we get another error:

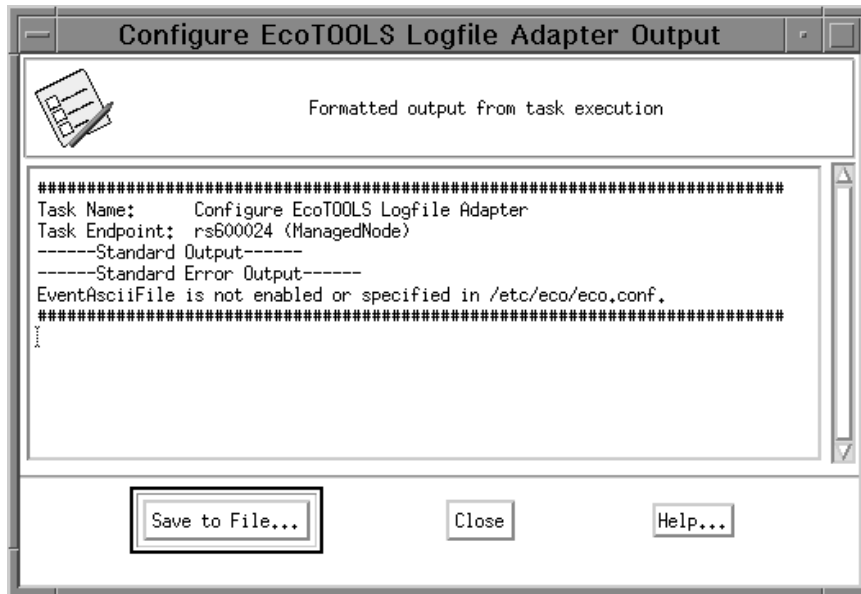


Figure 266. Another Error Message When Configuring Logfile Adapter

This time the message indicates that the `/etc/eco/eco.conf` file is not configured to write events yet. In order to do this, we need to edit the file and uncomment the line highlighted in the following figure:

```

# $Id: eco.conf,v 301.30.340.12 1997/06/02 22:23:12 padams Exp $
#
# Copyright (c) 1992,1993 EcoSystems Software, Inc.
#
# eco.conf
#
# Site-specific configuration information

homedir      /usr/eco
viewname     prod

#
# ecodb tunable parameters
#
database     /usr/eco/db      # the database directory

ECOIMDB_SERVERFMT /usr/eco/bin/ecodb.sh /usr/eco/bin/ecodb %s >/dev/nul &
# command to run ecodb
ECOIMDB_CREATETIMEOUT 60      # Create-timeout limit
ECOIMDB_IOTIMEOUT 30        # I/O timeout specification
ECOIMDB_CACHESIZE 102400    # ecodb cache size

oidpath 1.3.3 /usr/eco/oracle/agt
oidpath 1.3.4 /usr/eco/sybase/agt
oidpath 1.3.5 /usr/eco/informix/agt

...

#
# The following lines usually are commented out, since the values
# are automatically assumed by the system as defaults.
# To change any of the defaults, uncomment the line and change the
# default to the desirable value. All the time values are in seconds.
#
#data-entry 500 # up to 500 datapoints/datalog
data-entry-size 512 # max bytes per datapoint
#DataAsciiFile /tmp/datafile # external data ASCII log file
EventAsciiFile /tmp/eventfile # external event ASCII log file
#EventAsciiDir /usr/eco/db/eventlogs # external event ASCII log directory

...

```

Figure 267. Excerpt from *eco.conf* File

Uncommenting the indicated line will start logging events from EcoTOOLS to /tmp/event, which in turn will be monitored by the TEC logfile adapter.

After making the modification, we start the task again and this time get a successful completion:

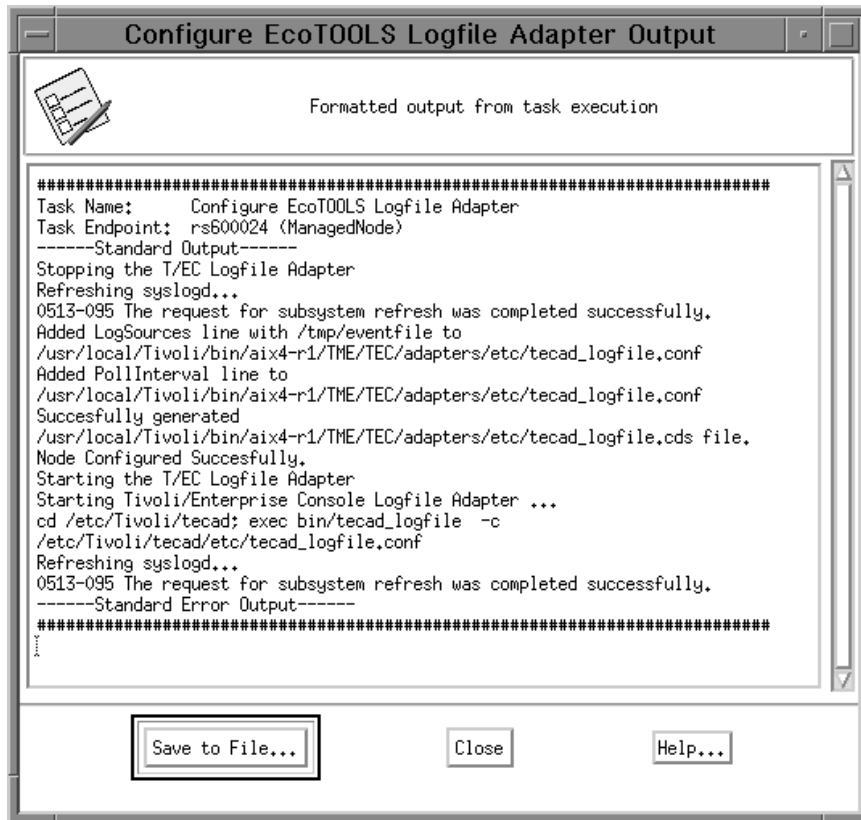


Figure 268. Successful Configuration of Logfile Adapter

It is recommended to stop and start the logfile adapter after reconfiguring it. We can do this by typing the following commands:

```

/etc/Tivoli/tecad/bin/init.tecad_logfile stop
/etc/Tivoli/tecad/bin/init.tecad_logfile start

```

Events should now be forwarded from EcoTOOLS to TEC.

The configuration task adds /tmp/eventfile as a source for the logfile adapter as shown in the following figure.

```

#
# TE/C LogFile adapter configuration file.
#
# $Date: 1995/11/27 16:23:48 $
#
# $Source: /tivoli/development/src/2.0/apps/tec/adapters/logfile/
conf/aix4-r1.conf/tecad_logfile.conf,v $
#
# $Revision: 1.1 $
#
# $Id: tecad_logfile.conf,v 1.1 1995/11/27 16:23:48 wgg Exp $
#
# Description:
#
# (C) COPYRIGHT TIVOLI Systems, Inc. 1994.
# Unpublished Work
# All Rights Reserved
# Licensed Material - Property of TIVOLI Systems, Inc.
#

ServerLocation=@EventServer
BufEvtPath=/etc/Tivoli/tec/logfile.cache
# ServerPort=9999
EventMaxSize=4096

Filter:Class=Logfile_Base
Filter:Class=Logfile_Sendmail
Filter:Class=Amd_Unmounted
Filter:Class=Amd_Mounted

LogSources=/tmp/eventfile
PollInterval=10

```

Figure 269. /etc/Tivoli/tecad/conf/tecad_logfile.conf File

You can also have a look at the tecad_logfile.fmt file in the same directory as tecad_logfile.conf to see the modifications made by the task.

4.12.3.4 Testing the EcoTOOLS Adapter

We are now ready to test the event forwarding. In order to do that we create an EcoTOOLS scenario that generates an event.

The event generated by the scenario will be written to the EcoTOOLS event logfile. As this file is monitored by the EcoTOOLS event adapter it will then be translated into a TEC event and forwarded to the event server.

In the EcoTOOLS Computer Map window we select the **rs600024** icon with the right mouse button and then select **Create Scenario...** from the pull-down menu as shown in the following figure:

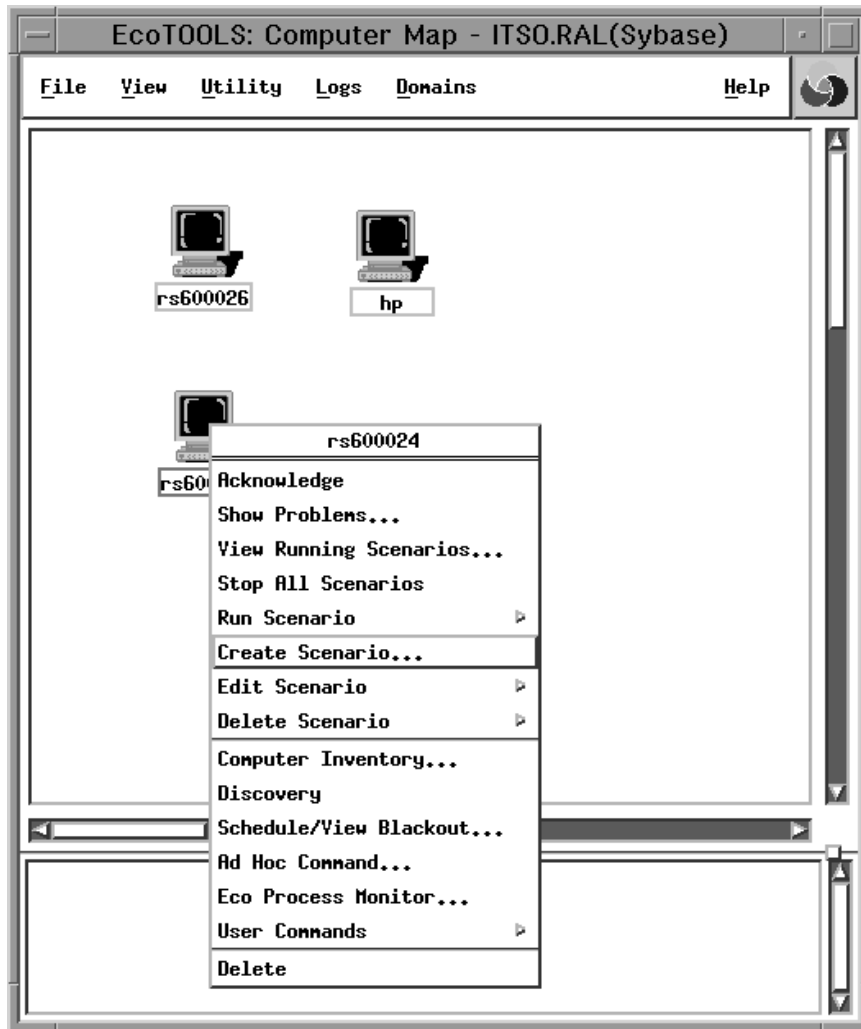


Figure 270. EcoTOOLS Computer Map Window

The Edit Scenario window will be displayed as shown in Figure 271 on page 267.

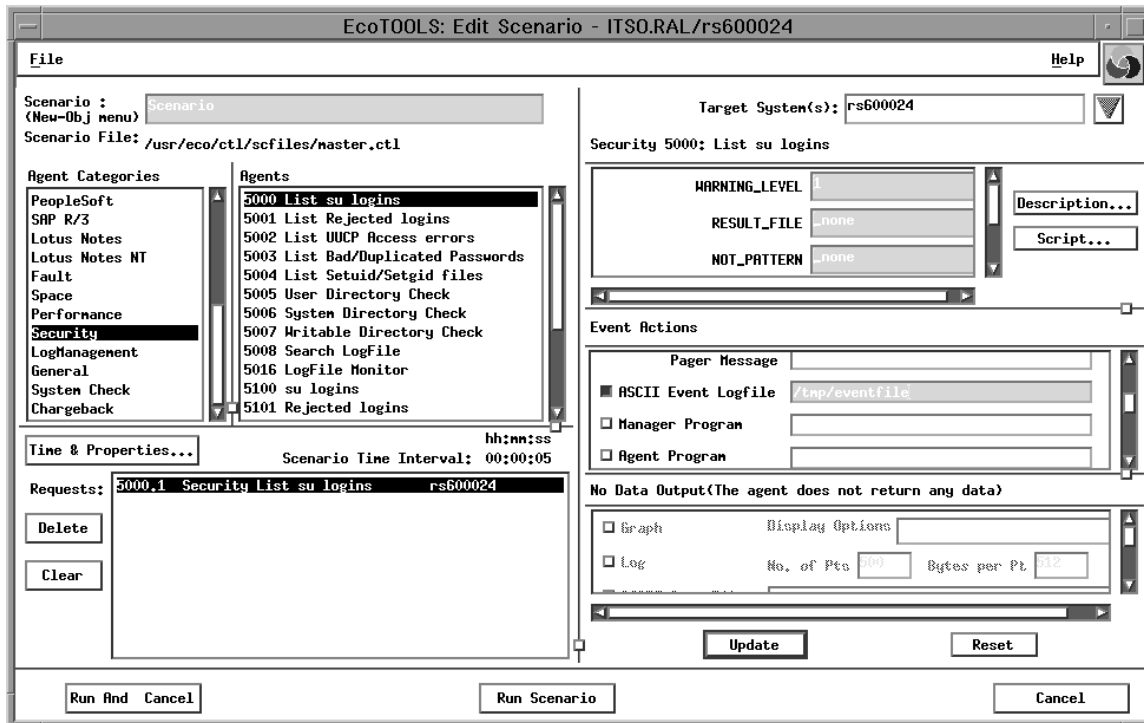


Figure 271. EcoTOOLS Edit Scenario Window

We want to create a simple scenario that watches switchuser attempts. To do so we select **Security** in the Agent Categories: section and then select **5000 List su logins** in the Agents: section.

We also check the **ASCII Event Logfile** box and enter /tmp/eventfile as the event file.

To specify the execution time of the scenario we click the **Time & Properties** button.

The following window will appear.

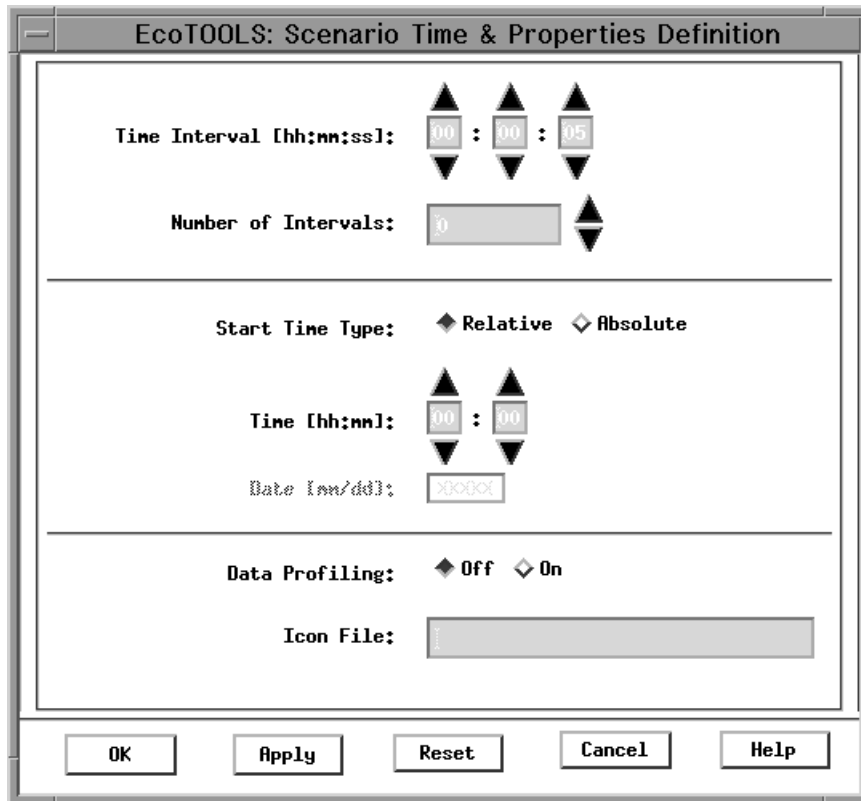


Figure 272. EcoTOOLS Time and Properties Definition Window

We enter the values as shown above and then click the **OK** button. Back in the Edit Scenario window we click on **Run and Cancel** to start our new scenario.

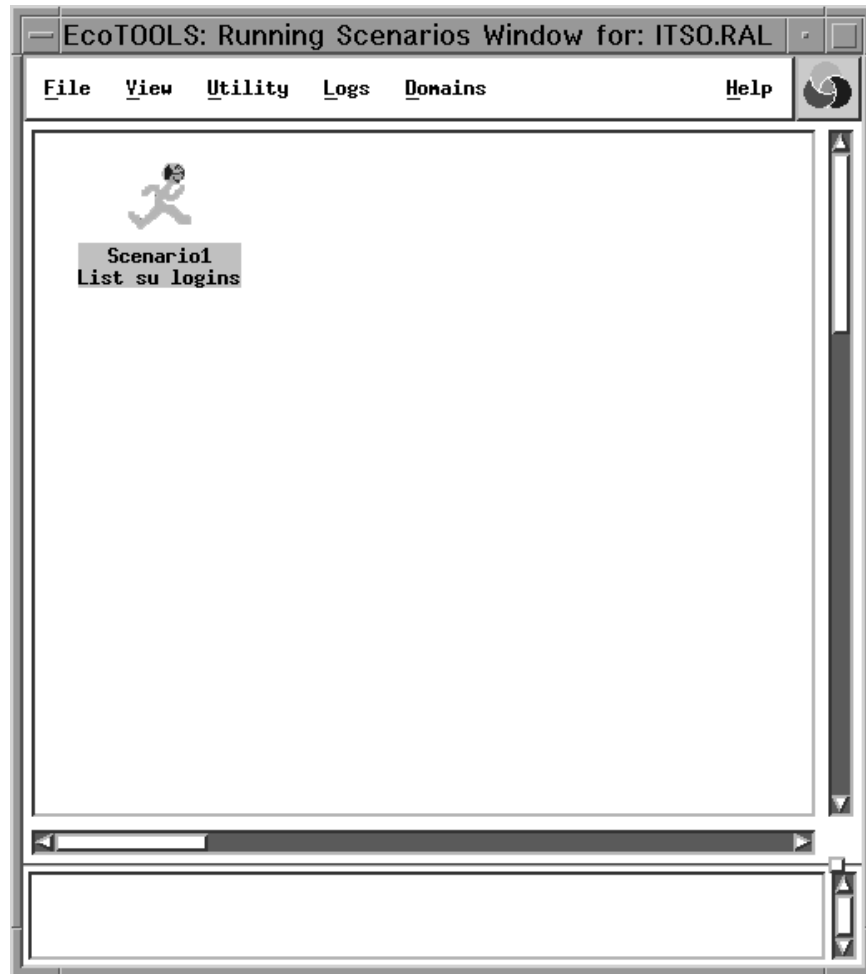


Figure 273. EcoTOOLS Running Scenarios Window

The symbol in the Running Scenarios window indicates that our scenario is running.

We can now trigger the event by typing:

```
su -
```

The switch user to root command will trigger an ecotools event. You will see a line added to the event file, /tmp/eventfile.

For example, you can use the tail command as shown below:

```

aixterm
root@rs600024[/etc/eco]# tail -3 /tmp/eventfile
V3.0 root prod rs600024 9.24.104.161 Scenario1 12/16/97 14:42:11 5 5000.1 12/16/
97 14:49:42 14:49:45 - 0.00 rs600024 "Super users logins" 2
V3.0 root prod rs600024 9.24.104.161 Scenario2 12/16/97 14:49:33 5 5000.1 12/16/
97 14:49:43 14:49:45 - 0.00 rs600024 "Super users logins" 2
V3.0 root prod rs600024 9.24.104.161 Scenario2 12/16/97 14:49:33 5 5000.1 12/16/
97 14:49:43 14:49:45 - 0.00 rs600024 "Super users logins" 2
root@rs600024[/etc/eco]# _

```

Figure 274. New Events in /tmp/eventfile

You can see that events are written to the file /tmp/eventfile when we invoke the su command.

Also, you will see events raised in the Event/Error Log window of EcoTOOLS as shown below:

The screenshot shows the EcoTOOLS: Event/Error Log window. The main window contains a table with the following data:

Time	Sev	Computer	Database	Category	Agent	Event/Error Message
15:00 12/16/97	2	rs600024		Computer	Cpu % Utilization	CPU Utilization is over 5%
14:49 12/16/97	2	rs600024		Security	List su logins	Super users logins
14:49 12/16/97	2	rs600024		Security	List su logins	Super users logins
14:49 12/16/97	2	rs600024		Security	List su logins	Super users logins
14:46 12/16/97	2	rs600024		Security	List su logins	Super users logins
14:44 12/16/97	2	rs600024		Security	List su logins	Super users logins
14:43 12/16/97	2	rs600024		Security	List su logins	Super users logins
14:43 12/16/97	2	rs600024		Security	List su logins	Super users logins
14:42 12/16/97	2	rs600024		Security	List su logins	Super users logins

Below the table, there is a detailed view of an event:

```

*****
* Computer Cpu % Utilization <1000,1>
*
* Time Started: 15:00:08 12/16/97      Domain:      ITS0,RAL
*                                     Scenario:    Scenario3
*                                     Computer:    rs600024
*
*****
Event Time:      15:00:44 12/16/97
Severity:        2

```

At the bottom of the window, it says "machine = "rs600024"

Figure 275. EcoTOOLS Event/Error Log

To see if a TEC event has occurred we open the event console and select the event group **ALL**.

We see the following window:

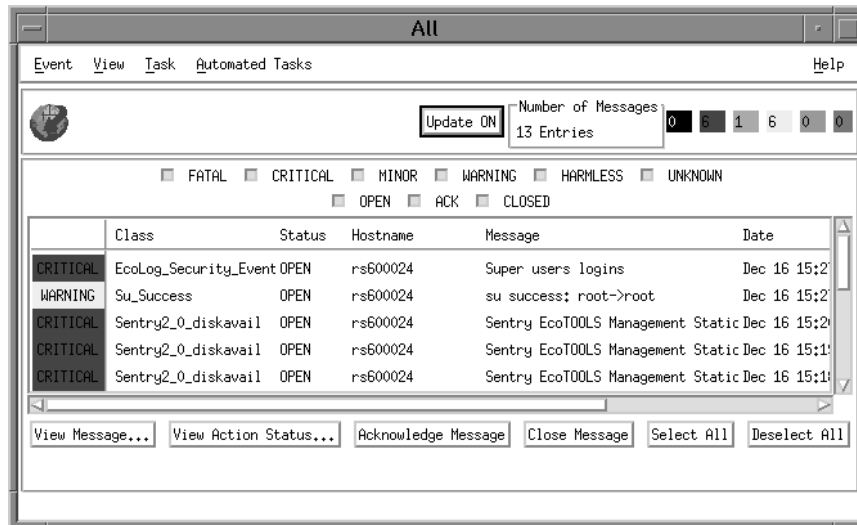


Figure 276. TEC Event Display

Interestingly, you see that two events were generated by typing the su command. The reason for this is that the TEC logfile adapter also monitors su events, therefore we get two events: one of the class `Su_Success`, which is defined by the logfile adapter itself and one of the class `EcoLog_Security_Event`, which is defined also by the logfile adapter but this time modified by EcoTOOLS Plus Module for TME 10.

We click on the **EcoTOOLS** event to display the event details:



Figure 277. Event Group Message Viewer

You can see that the message indicates a Super user login on host rs600024 with a severity of CRITICAL.

4.12.4 EcoTOOLS Plus Module for TME 10 Tasks to Manage EcoTOOLS

In this section we discuss the tasks to manage EcoTOOLS provided with the EcoTOOLS Plus Module for TME 10 in more detail.

The following tasks are described in the next subsections:

- Agent help task
- Backup the EcoTOOLS password
- Backup and restore of the EcoTOOLS database

- Starting and stopping the EcoTOOLS daemons ecolic, ecoad and eamon
- Running EcoTOOLS scenarios
- Listing user access roles of EcoTOOLS

For each task we describe how to launch the task from the desktop and we give an example of what the task output looks like.

Most tasks can be invoked by double-clicking on the corresponding icon. In order to make sure that the correct task endpoint is set, select **Modify Job...** from the pop-up menu of the icon and enter the desired execution target in the Edit Job window. The task endpoint may either be a single machine or a profile manager.

We assume that your EcoTOOLS home directory is /usr/eco.

4.12.4.1 The Task EcoTOOLS Agent Help

This task provides information about the various EcoTOOLS agents that are responsible for monitoring, gathering data, observing user-specified conditions and thresholds, and for the execution of action.




Figure 278. Executing EcoTOOLS Agent Help Task

The task must be executed on the EcoTOOLS management station. By default, the EcoTOOLS Management Station's profile manager is used to determine the target machine for execution. If the subscribers of this profile manager are set correctly, simply double-click on the task icon. This opens the EcoTOOLS Agent Help window as shown in Figure 278. The agents are identified by numbers that are four digits long. Enter the number of the agent and select the **Set and Execute** button.

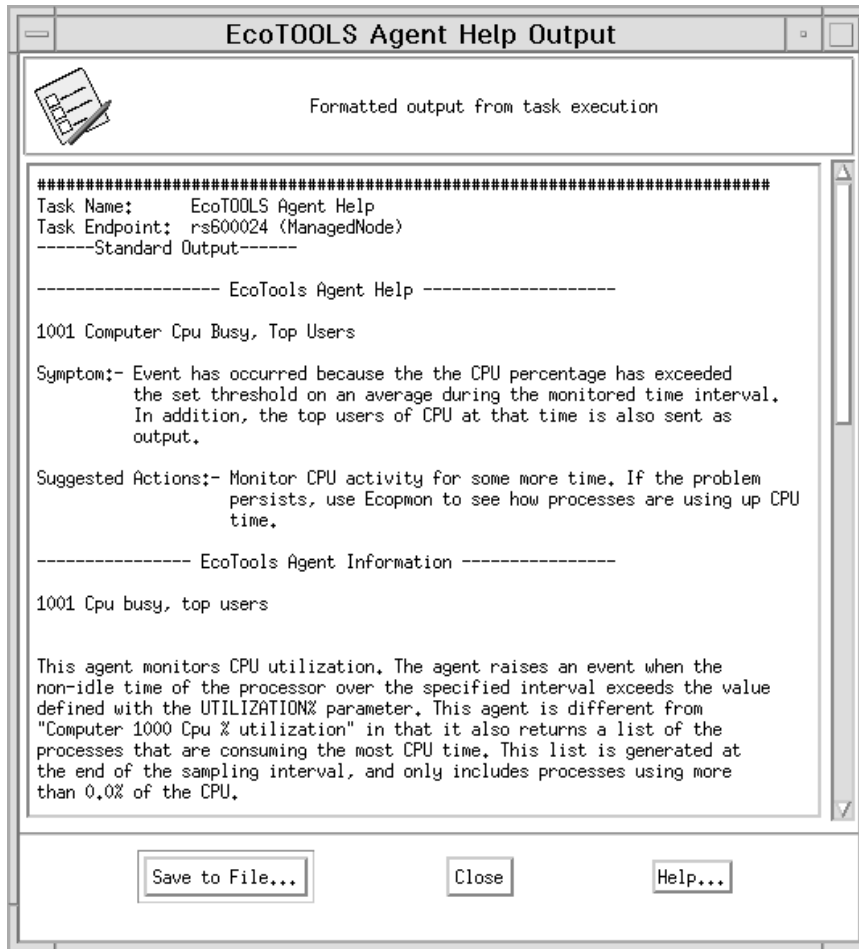


Figure 279. Output of EcoTOOLS Agent Help Task

The output of the task is displayed in the EcoTOOLS Agent Help Output window. The output consists of the two parts, EcoTOOLS Agent Help and EcoTOOLS Agent Information. The EcoTOOLS Agent Help part describes the conditions under which the agent raises an event and the recommended actions in case of an event.

The EcoTOOLS Agent Information part provides detailed information about the way of working of the agent. The description in this section contains comparisons to other agents, necessary input parameters, the name of the executed program and the recommended interval of execution.

4.12.4.2 Backing Up the EcoTOOLS Password File

This task is used to back up the EcoTOOLS password file which is located in the directory /usr/eco/etc. The file has the name ecopasswd. This task should be executed whenever the password is changed.

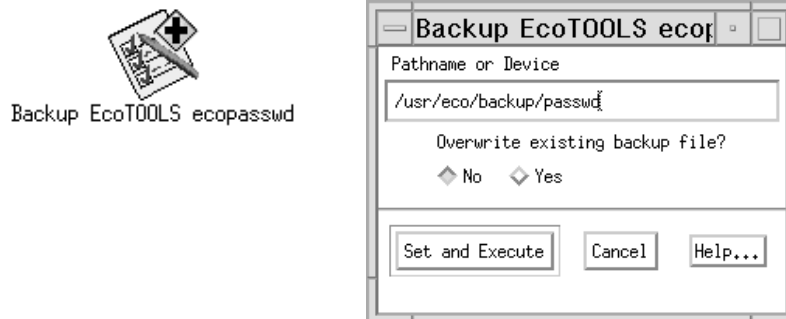


Figure 280. Backup EcoTOOLS Password File

The task is executed on the EcoTOOLS management station. Double-click on the task icon if the execution endpoints are set correctly. In the Backup EcoTOOLS ecopasswd window specify the backup target. This may be a file on a hard disk or a device, such as a tape drive.

Make sure to enter the complete path together with the file name. When you are finished click **Set and Execute**.

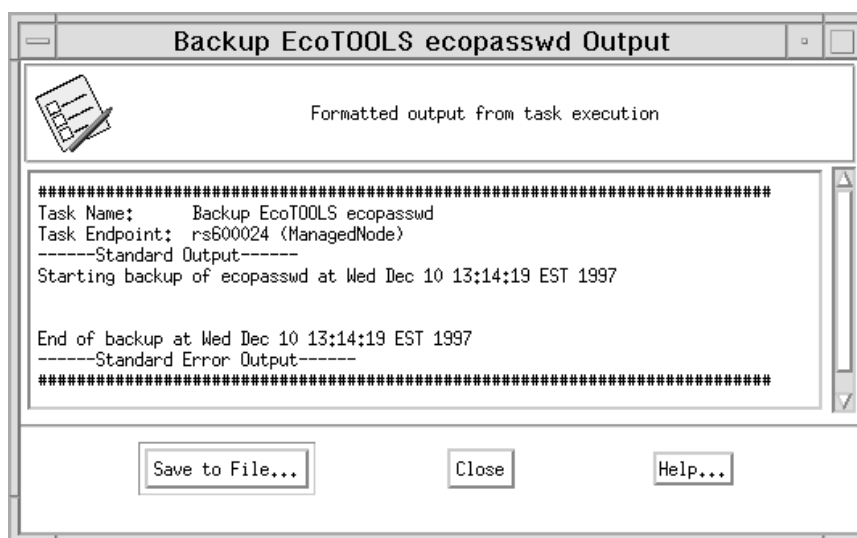


Figure 281. Output of Backup Password Task

The Backup EcoTOOLS ecopasswd Output window shows the result of the task execution. When the task completes examine the specified file. The ecopasswd file is an ASCII file so you can easily compare the contents of both files with the diff command.

4.12.4.3 Backing Up and Restoring the EcoTOOLS Database

The backup task stores the EcoTOOLS database files in a tar file. The restore task extracts the content of the tar file and restores the database to its original destination.

The EcoTOOLS database stores the domain information, performance information, the directory structure and the information gathered from running scenarios. Every user organizes his or her monitoring environment in views, a combination of

domains, log files and custom scenarios. For each user/view combination a separate database is created.

The databases are located in the a directory specified by the environment variable \$dbhome which defaults to the path /usr/eco/db. Under this directory a set of subdirectories is created containing the database and log files of the different user/view combinations.

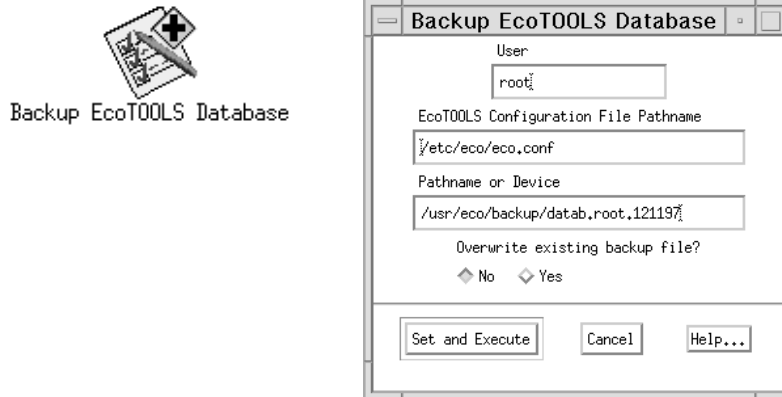


Figure 282. Executing Backup EcoTOOLS Database Task

By default the task is executed on the EcoTOOLS management station. Double-click on the task icon and enter the user name and the target tar file in the Backup EcoTOOLS Database window (Figure 282). If necessary, change the default entry for the location of the EcoTOOLS configuration file eco.conf. If you backed up the database before, you can use the same tar file. In this case, select the radio button to overwrite the existing backup file.



Figure 283. Output of the Backup EcoTOOLS Database Task

While the task is running, all files written to the tar file are listed in the Backup EcoTOOLS Database Output window (Figure 283). Besides the database file `imdb` control files and re-do log files are backed up.

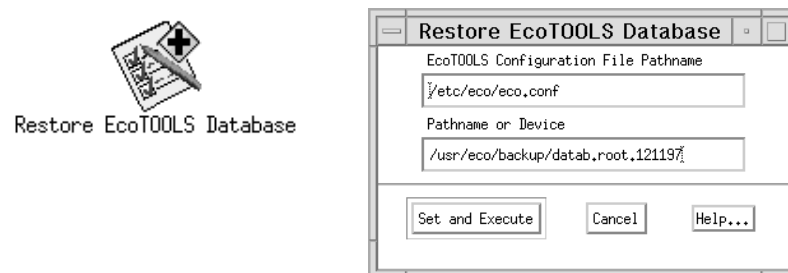


Figure 284. Executing Restore EcoTOOLS Database Task

Like the backup task, the Restore EcoTOOLS Database task is executed on the management station and the Management Station profile manager is used. Double-click on the task icon and wait for the Restore EcoTOOLS Database window to appear. Verify the correct path and name of the EcoTOOLS

configuration file and enter the name of tar file containing the database backup. Click the **Set and Execute** button when finished.

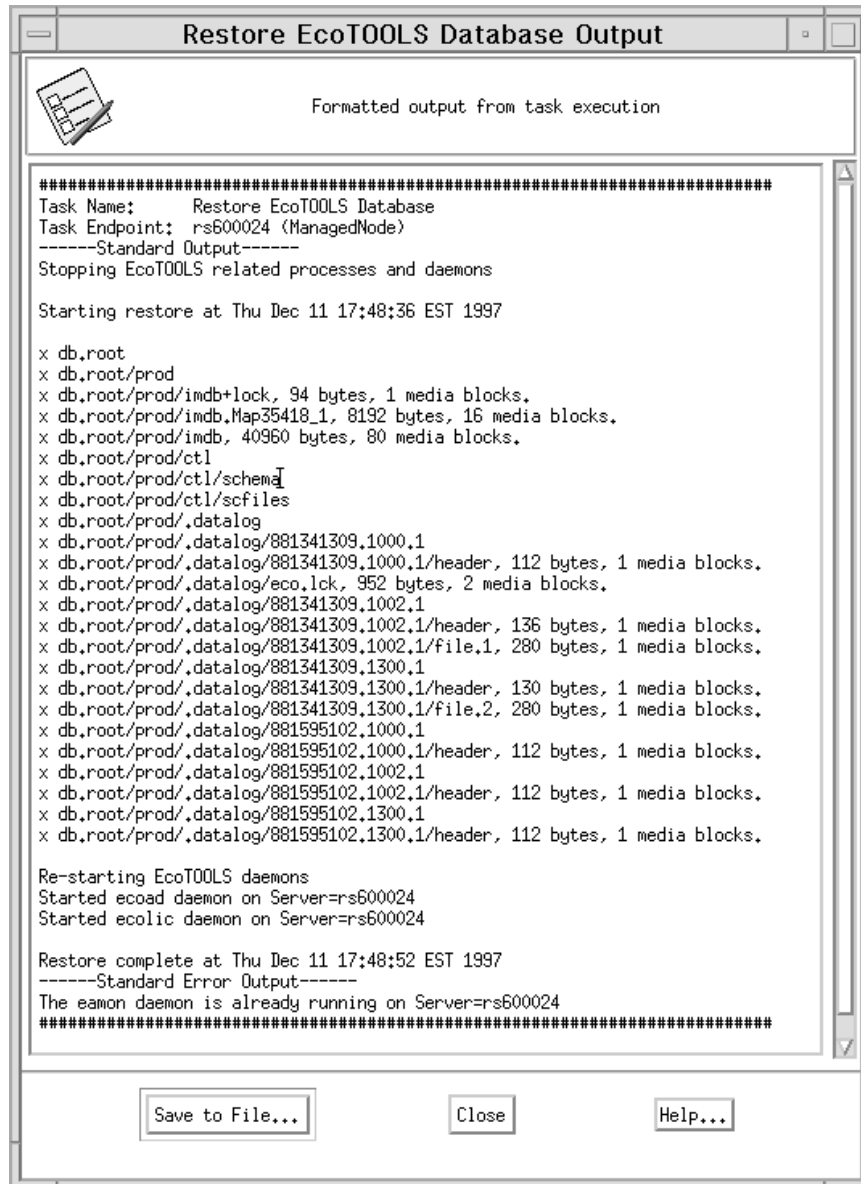


Figure 285. Output of the Restore EcoTOOLS Database Task

The restore task will read the tar file and copy all contained file to their origin. The output is logged in the Restore EcoTOOLS Database Output window.

4.12.4.4 Starting and Stopping the ecolic Daemon

This task stops and restarts the EcoTOOLS license server on a license station. This task must be executed whenever the license information has been updated. As long as the license server is not active new EcoTOOLS management consoles cannot be started.




Figure 286. Executing Start/Stop ecolic Daemon Task

By default, the task uses the EcoTOOLS License Station's profile manager to determine the task endpoint. If the execution target is set up correctly, double-click on the task icon. This will open the Start/Stop ecolic daemon window in Figure 286. Select the appropriate radio button to stop or to restart the license server.

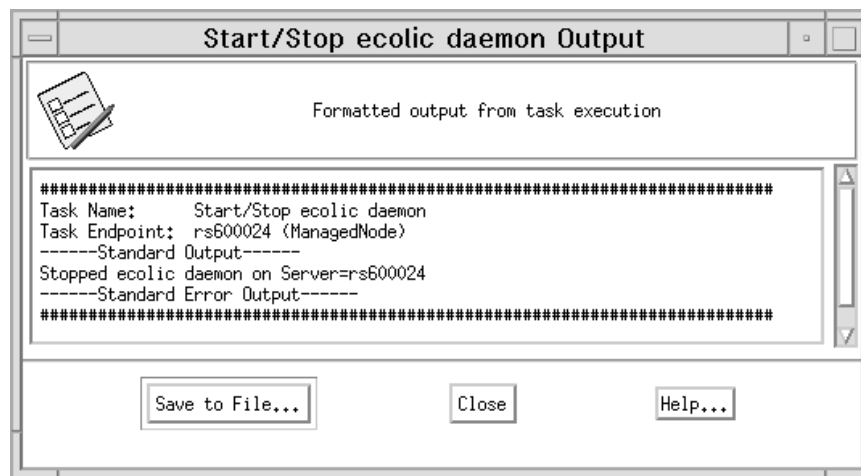


Figure 287. Output of the Start/Stop ecolic Daemon Task

The example of the Start/Stop ecolic daemon Output window in Figure 287 signals the successful execution of the task. Use the EcoTOOLS command `showlic` to verify that the license server is running.

4.12.4.5 Starting and Stopping the eamon Daemon

This task starts the agent monitoring process which is responsible for the proper execution of scenario requests. In case of scenarios not functioning properly the daemon reports error messages.



Figure 288. Executing Start/Stop eamon Daemon Task

The Start/Stop eamon Daemon window appears when you double-click on the task icon. Select the radio button representing the desired function and select **Set and Execute**.

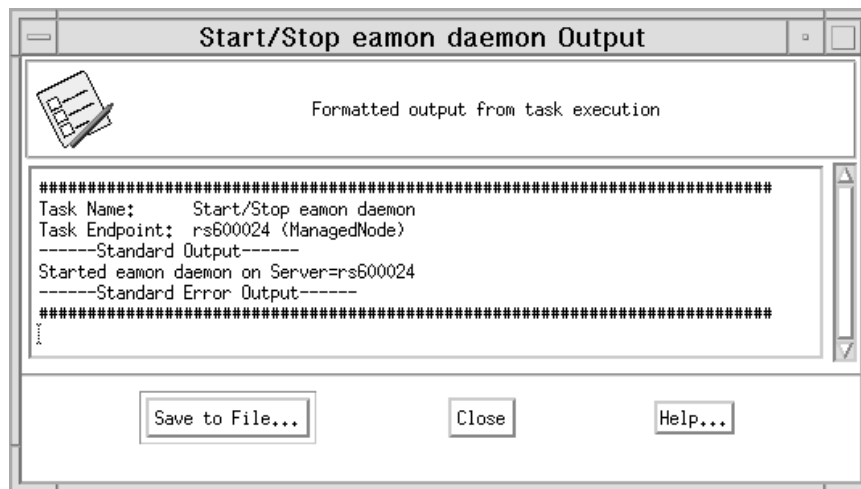


Figure 289. Output of the Start/Stop eamon Daemon Task

When restarting the daemon the task calls the executable `/usr/eco/eamon`. The Start/Stop eamon Output window logs the results of this call. As the EcoTOOLS communication is based on an RPC mechanism you can use the UNIX command `rpcinfo` to see if the daemon is alive.

4.12.4.6 Stopping and Starting the ecoad Daemon

The EcoTOOLS dispatcher `ecoad` sends and receives scenario requests and is running on the agent stations. Without this daemon the management station cannot communicate with the agent stations. As an example, we show how to start the process.



Figure 290. Executing Start/Stop ecoad Daemon Task

By default, the task employs the EcoTOOLS Agent Station's profile manager, but we recommend you launch the task by selecting **Run on selected subscribers...** to explicitly choose a task endpoint. In the Execute Task window select execution targets from the Available Task Endpoints field and click the **Execute** button. Wait for the Start/Stop ecoad Daemon window and select the appropriate radio button. Select **Select and Close** to execute the task.

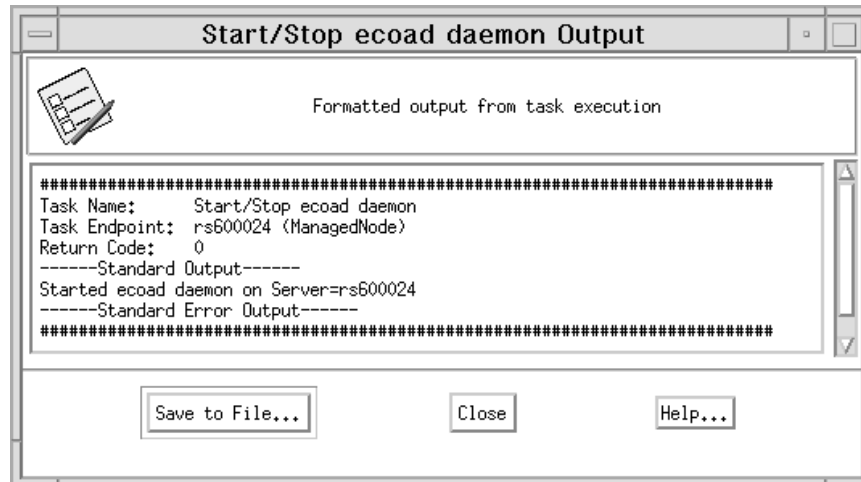


Figure 291. Output of the Start/Stop eoad Daemon Task

Examine the Start/Stop eoad Daemon Output window shown in Figure 290 on page 280 to see if the start of the daemon completes. The task calls the script /usr/eco/exec/aix6000/restartEoad (on AIX machines) which in turn executes /usr/eco/bin/eoad. On success you should see the process running when using the ps command or executing the rpcinfo command.

4.12.4.7 Running EcoTOOLS Scenarios

This task sends a request to run a custom scenario on an EcoTOOLS agent station.

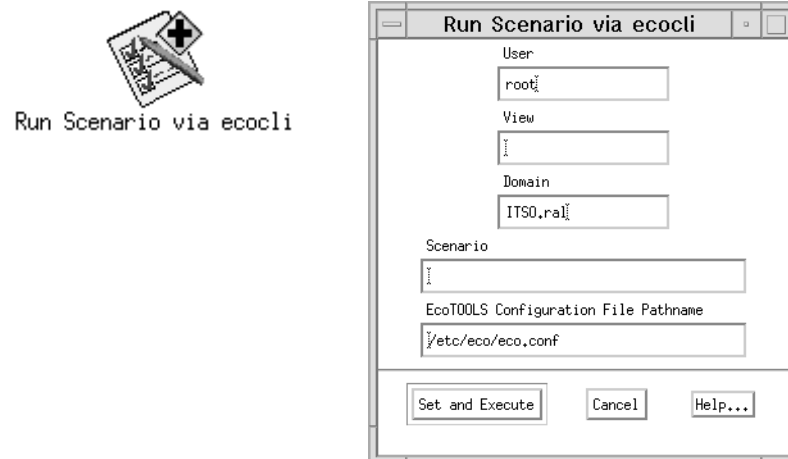


Figure 292. Executing Run Scenario via ecocli Task

Use the **Run on selected subscribers** option to initiate the task execution. In the Run Scenario via ecocli window enter a user name, the view name, the domain name and the name of of the scenario to be executed. The name of the view must be valid in the user environment. The view is a combination of domains and associated log files, custom scenarios and the EcoTOOLS database. The domain must be a valid domain within this view. The custom scenario must be defined for the domain.

The output window of the run scenario task shows the result of the attempt to run the scenario. The task fails if the scenario is not found (return code 1).

4.12.4.8 Listing User Access Roles

This task lists all users together with their roles EcoTOOLS uses to gain access to the monitored applications. The task executes the EcoTOOLS command `ecoacl -l ALL`.



Figure 293. Executing List User Access Roles Task

This command must be executed on the management station. You can use the EcoTOOLS Management Station profile manager. Double-click on the task icon and select the **Execute and Close** button.

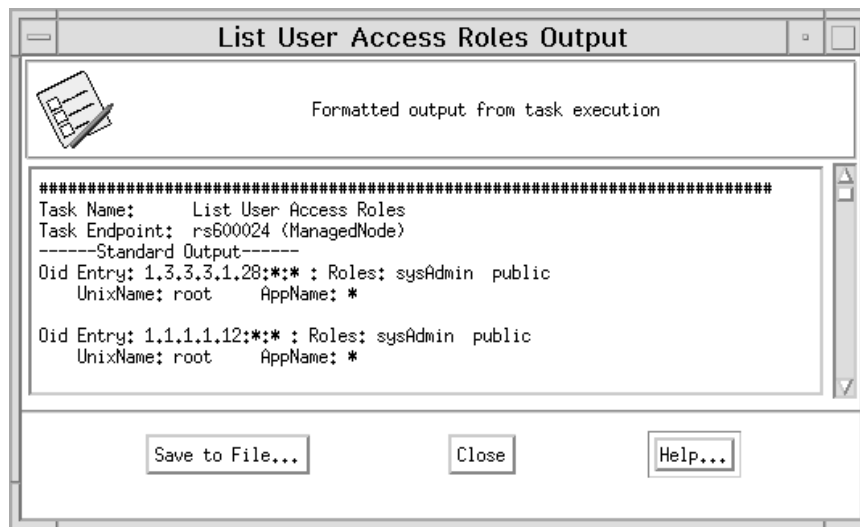


Figure 294. Output of the List User Access Roles Task

The output in Figure 294 shows an example of the List User Access Roles task. This information is derived from the oid-to-role acl database which is stored in the file `/usr/eco/acl/oidDB`. A backup is stored in `/usr/eco/acl/oidDB.OLD`. Each entry consists of the oid string, a list of roles and a user name. The oid string consists of the oid (unique object identifier), a server and an application (usually a database, for example sybase) running on this server. These access control lists are used to control if a given request is valid.

4.13 How Can EcoTOOLS Complement TME 10?

EcoTOOLS provides a large number of excellent monitors for a wide range of application components, including RDBMSs, transaction monitors, client/server applications, such as SAP R/3, Baan, Lotus Notes, Oracle applications and PeopleSoft, networks, operating system parameters and much more.

The following figure gives an impression of the monitors available.

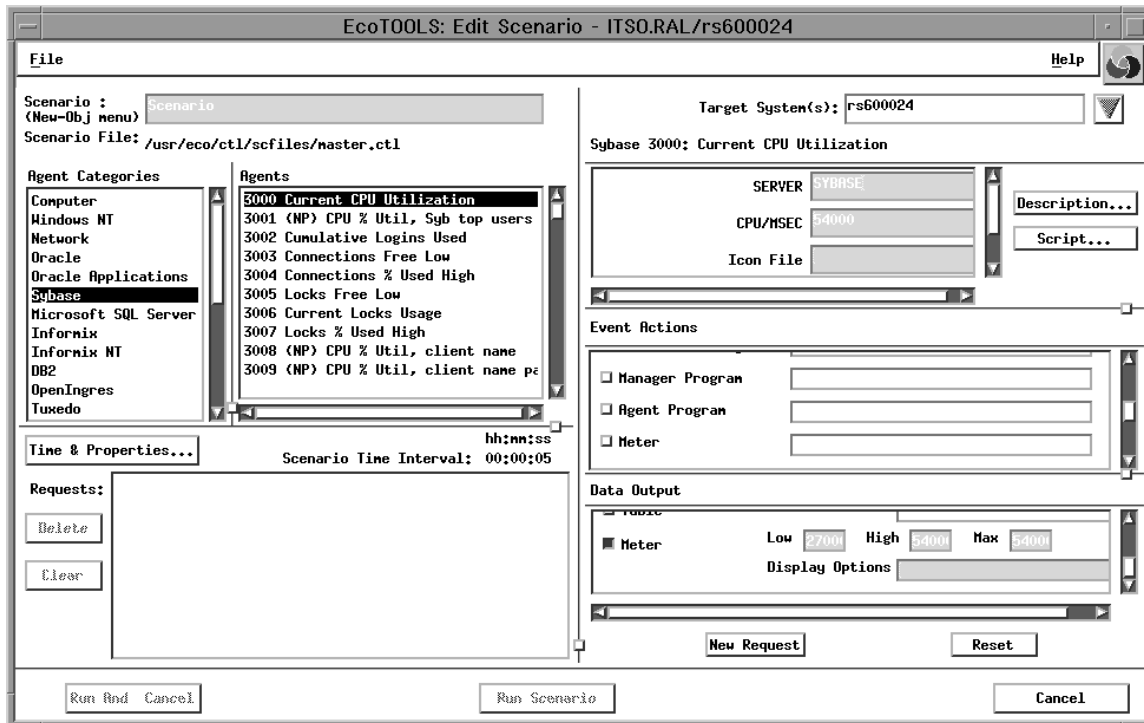


Figure 295. EcoTOOLS Monitors

The different monitor categories are displayed in the Agent Categories: field. In each category there are a number of monitors (agents) available, shown in the Agents: section.

This concept is similar to monitoring collections and monitors in TME 10 Distributed Monitoring.

All EcoTOOLS monitors are very efficient and comprehensive, so that they provide excellent availability monitoring capabilities for almost any kind of environment.

Another interesting application for the EcoTOOLS monitors is using them to monitor the health of TME 10 itself. A good example is monitoring the RDBMS in which, for example TEC stores its data. As EcoTOOLS has excellent monitors for all major RDBMSs, for example, Sybase, you can easily set up a scenario in which you monitor all vital information for the TEC database and take corrective actions in case an exception occurs.

On the other hand, when EcoTOOLS is integrated with TME 10 using EcoTOOLS Plus Module for TME 10 you combine the availability management strengths of EcoTOOLS with the excellent platform services and core applications provided by TME 10.

Chapter 5. TME 10 Solution Design Involving 10/Plus Partner Products

In this chapter we look at how TME 10 solution design is affected by the variety of modules available from Tivoli's partners.

Solution design here means creating a comprehensive solution for the systems management needs of a client, which is usually an organization needing to manage all the components of their information systems.

We look back on previous approaches to provide a comprehensive solution to the systems management challenge and then look at the partner approach of Tivoli.

5.1 The Solution Design Challenge

In today's client/server environments solution, designers such as I/T architects are confronted with the increasingly complex challenge of creating a systems management solution that involves the management of a wide variety of components, systems and networks.

For example, a typical environment will include an MVS mainframe, distributed UNIX, Windows NT or OS/2 servers, end user desktop systems, RDBMS servers, custom-written applications, TCP/IP, IPX and SNA networks, standard applications such as Microsoft Office, SAP R/3 or Lotus Notes and many, many other components.

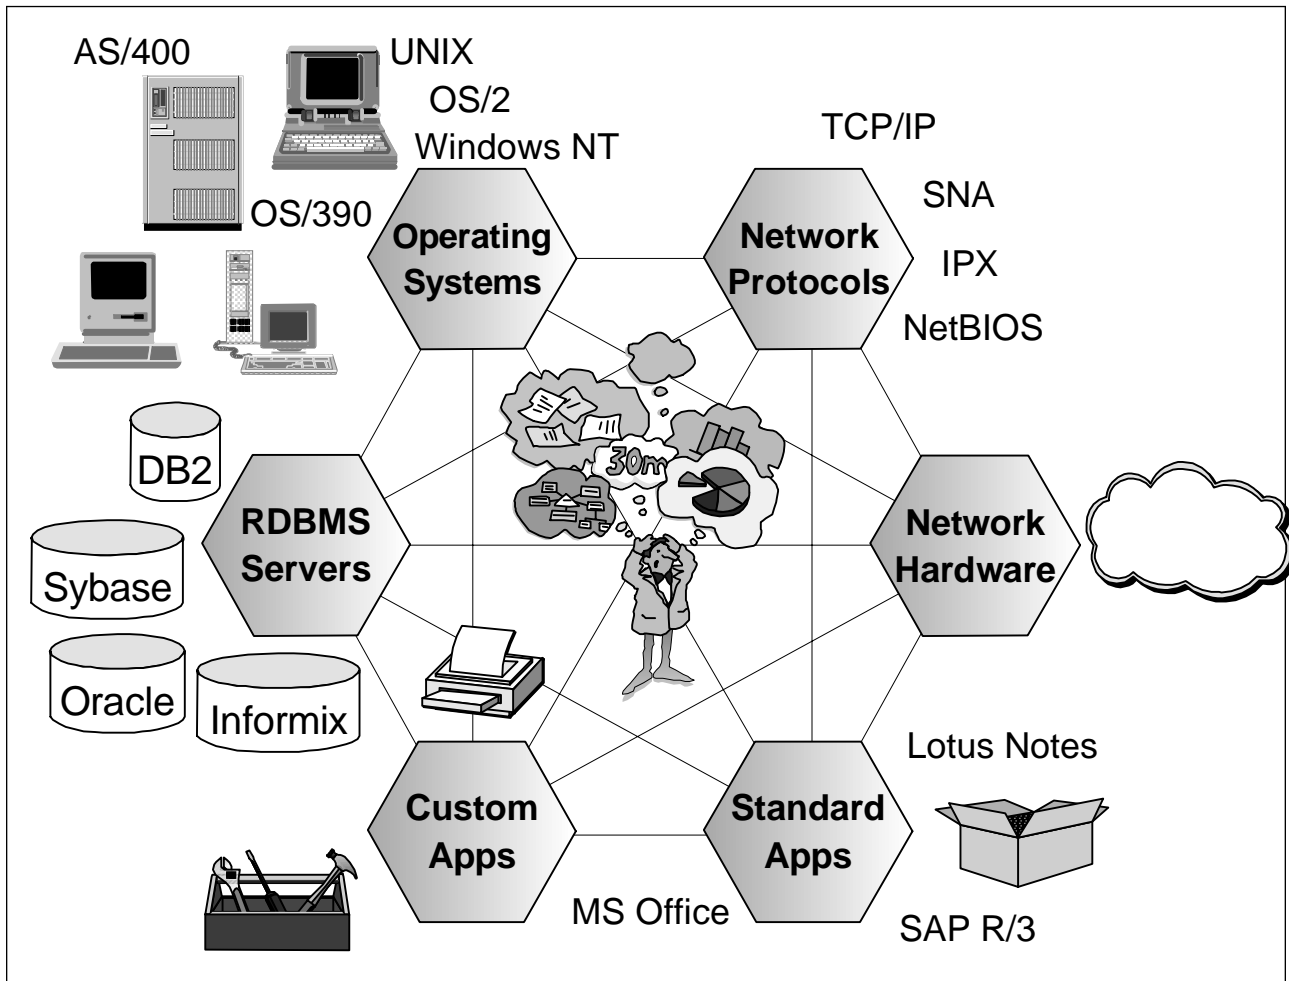


Figure 296. The Solution Design Challenge

As a result, no single vendor alone can supply a management solution that is able to manage all components in a client's organization. Nevertheless, approaches to systems management software were often based on the assumption that one vendor could actually provide a single solution.

This was mostly achieved by focusing on the most obvious problem, which is usually network management, and then either declaring the solution as achieved, or defining some interfaces between network management and systems management that allowed you to partly integrate the two. Most of these approaches were based on exploiting the SNMP protocol, using it also for systems management.

These solutions were quite popular when the client/server phenomenon first started, however, they were only successful because most of these new systems (mostly UNIX and PC) were either still used as stand-alone workstations or not used at all. For example, while several organizations started to deploy a large number of UNIX systems, in the beginning they were often used only as network or file servers while the business applications were still running on the mainframe.

The evolution of true client/server applications, including truly distributed business systems, such as SAP R/3 and distributed client/server databases such as Oracle, Sybase, Informix or DB2 changed the picture.

While in former times managing the IP network using an SNMP manager with a few added gimmicks to ensure availability of IP nodes and analyze the performance of the network may have been sufficient, in today's heterogeneous environments there are so many components to manage, that no single vendor can provide software to manage all these components.

As a result, exploiting an SNMP manager and adding component management features will not work anymore. One of the many reasons is that SNMP is a protocol designed for network management and not for systems management and can therefore not accommodate a comprehensive systems management approach.

This is where systems management frameworks come into place. Instead of expanding an existing management application, such as an SNMP manager with somewhat artificial and unnatural interfaces to enable it to manage systems, a systems management solution focuses on providing a solid platform with strong base services and applications and also open and flexible interfaces, so that component management software can easily integrate.

TME 10 is such a framework.

5.2 The Tivoli Approach

Tivoli's focus is providing the leading-edge framework for systems management and the industry-leading core applications on top of that framework.

Openness and inclusion of partners were important factors behind the design of the TME 10 Framework and result in a unique platform that is robust, scalable and provides effective and open interfaces, so that partner applications can build upon framework and core application services and easily integrate into TME 10.

Not only does Tivoli provide these interfaces, more importantly vendors are heavily using TME 10 to link their applications into a systems management framework, resulting in a community of more than 350 partners that already have written or currently write applications that link into TME 10.

For the designer of a systems management solution this opens up a completely new world of opportunities for creating a systems management solution. The dilemma in previous approaches to systems management was that they were usually linked to the product set of a single vendor. Although almost all these vendors highlighted the importance of partnerships, this was usually not reflected in the actual implementation of their products.

Let's look at an example. Customer A asks vendor B to come up with a proposal to manage the customer's information systems. The solution should include management of the corporate IP network, distribution of software, management of several RDBMS servers and applications, performance measurement and capacity planning for a number of resources and also management of the corporate-wide DCE installation.

Vendor B has a portfolio of products that include an SNMP manager, a software distribution product and a monitoring tool to manage a number of operating system parameters. All of these products use proprietary databases and their own interfaces and don't integrate well with other applications. The vendor does have no products to cope with, for example, DCE management and capacity planning.

What would vendor B usually do? As it has no products of its own to offer for a complete solution and does not want to refer the client to a competitor it can either downplay the importance of areas where no own solution is available, defer the client to a later time at which the missing components would be implemented or employ some other techniques that free it from dealing with the missing functionality.

All of these are apparently bad for the client. For example, if missing components are deferred to a later date, they are likely not to be included in the initial overall design and therefore also likely to not fit in well even if they become available some day.

What is needed is a new approach in which a customer as well as a solution provider have a solid platform and a choice of which management components they want to include.

TME 10 and the 10/Plus Association provide this choice. From a solution provider perspective, it can build the design on top of the TME 10 Framework and the TME 10 core applications. This provides a standard methodology for all scenarios as the core applications will be needed in any environment. Of course, which of the Tivoli core applications will be needed depends on the specific requirements.

All scenarios have in common that they can take advantage of TME 10 proven production strength, openness and scalability which will lead to a rapid deployment of the systems management infrastructure.

Having set the groundwork that way, the solution provider or the client can now pick additional components of their choice that are needed to included management of specific components or disciplines in the systems management solution. These additional components can be additional applications provided by Tivoli or applications provided by the members of the 10/Plus Association.

The advantages are apparent. The customer can choose whatever additional components they like for their specific environment to link into their TME 10 installation. Tivoli will continue to provide the TME 10 Framework, the core applications and selected additional systems management applications. The partners provide specialized management modules for any kind of component, adding their unique expertise in their respective area.

5.3 Experiences Working with 10/Plus Applications

To get an idea of what working with 10/Plus modules is like, let's summarize the experiences we had working with the 10/Plus applications that were investigated for this redbook.

They are examples, but are quite typical 10/Plus modules and are therefore a good representation.

The first thing that is apparent is that all 10/Plus products have the same structure and appearance, such as:

- They all have the same GUI. Usually icons are integrated into the TME 10 Desktop that are used to access the functions of the partner application.

- They are grouped into different Tivoli categories, such as profile managers, profiles, indicators, etc. These groupings are consistent. For example, most modules provide predefined TME 10 Distributed Monitoring and TME 10 Software Distribution profiles, or Sentry indicators.
- Most modules provide an event adapter for TME 10 Enterprise Console, where the configuration of this component is very similar for all modules.

All modules can be installed in the same way, that is using the TME 10 Desktop and its standard product install feature. The module will be registered in the Tivoli object database and creates resources, for example, profiles, profile managers, etc.

Also, the configuration of a 10/Plus module is always similar. TME 10 Tasks are provided to set up the module with a few mouse clicks. The Tivoli-related information needed to complete these tasks is usually identical between modules, the only thing that differs is the product-specific information related to the application that is to be linked to TME 10 using the module.

Hence, once you have completed the installation of a 10/Plus module and successfully configured it, you will be able to apply this knowledge to any other module and install and configure it in a very short period of time.

You will add a world of benefits by installing a 10/Plus module in very short time, because usually after installing and configuring the module, you can instantly start receiving TEC events from the newly linked application, monitor it with TME 10 Distributed Monitoring, or deploy client components using TME 10 Software Distribution.

5.4 Simplified Design Steps

In the following we give some example suggestions of what steps could be involved in creating a systems management solution involving TME 10 and partner modules.

This is just an example and does not constitute a structured or universal approach.

For example, you will usually:

1. Gather requirements.

This will involve gathering information about the network topology, the hardware infrastructure, involved operating systems, middleware, other infrastructure components, etc. This phase should also include the client's specific needs in regard to the applications and other components to be managed by the solution.

2. Design TME 10 Infrastructure.

This will involve the layout of the TMRs, core applications, etc. and result in deploying the TME 10 infrastructure as the foundation for the systems management solution. The solution will usually be scaled by means of the TME 10 topology; that is scalability, efficiency and openness are achieved by the TME 10 Framework. That does not mean, that, for example, a specific 10/Plus module cannot employ a different topology than TME 10; however, TME 10 guarantees the scalability and rapid deployment of the underlying management infrastructure.

3. Select/propose partner modules.

The solution provider can either propose the components, that is partner modules, to complete the solution or the client can pick the components of choice to be integrated. The initial TME 10 design will then be extended with the additional components and verified. The final result will include a design for the TME 10 infrastructure including the partner components.

4. Deploy TME 10.

Usually, the TME 10 infrastructure, that is the TME 10 Framework and the core applications will be deployed first in the client's enterprise. This deployment can take advantage of TME 10's scalability and flexibility and therefore be achieved in a short time. As deploying the TME 10 infrastructure is a standardized and proven process for which structured methodologies exist that determine the steps to be taken, this will be based on a routine process, customized for any specific environment.

5. Deploy partner module.

Once the TME 10 infrastructure is in place, the 10/Plus module can be deployed extremely fast, utilizing standard TME 10 Framework features, that allow you, for example, to install and configure software and components anywhere in the network.

6. Deploy partner product.

Usually, a 10/Plus modules will link a systems management or other application to TME 10. In a lot of cases, this includes a deployment component based on TME 10 Software Distribution. This component can be used, as soon as the 10/Plus module is installed, to automatically deploy (that is install and configure) the components of the application anywhere in the TME 10 network.

7. Manage application using TME 10 and 10/Plus module.

Once the application and the 10/Plus module are deployed, the application can be managed by TME 10 in most cases; that is it can be monitored, upgrades can be deployed, events can be received into TEC, application functions can be used from the TME 10 desktop and so on. Or, the module itself complements TME 10 by adding a specific management function that can then be accessed through TME 10, for example, managing DCE.

Appendix A. Tivoli 10/Plus Summary

In this appendix we list the 10/Plus modules that were available when this redbook was written and refer you to where you can get more information.

A.1 10/Plus Partners

The following table lists the 10/Plus partner modules currently available. To get up-to-date information on available partner modules, you can refer to the URL: www.tivoli.com.



Table 4 (Page 1 of 3). Overview of Tivoli/Plus and 10/Plus Modules

Product Name	Vendor	Level	Management Area
Legato Networker	Legato Systems, Inc.	Premiere	Backup/Restore
Alexandria	Spectra Logic, Inc.	Premiere	Backup/Restore
ADSM	IBM Corporation	Premiere	Backup/Restore
NetBackup	Veritas	Premiere	Backup/Restore
Control-M	New Dimension Software	Premiere	Workload Management
Maestro	Unison Software/Tivoli Systems	Premiere	Workload Management
AutoSys	Platinum Technology, Inc.	Premiere	Workload Management
TRUEchange	True Software	Premiere	Workload Management
AR System	Remedy Corporation	Premiere	Help Desk
Expert Advisor	Software Artistry, Inc.	Premiere	Help Desk
ServiceCenter	Peregrine Systems, Inc.	Premiere	Help Desk
Solstice Domain Manager	SunSoft	Premiere	Network Management
OpenView	Hewlett-Packard	Premiere	Network Management
Spectrum	Cabletron Systems	Premiere	Network Management
OmniGuard	Axent Technologies	Premiere	Security Management

Table 4 (Page 2 of 3). Overview of Tivoli/Plus and 10/Plus Modules

Product Name	Vendor	Level	Management Area
EcoTools	Compuware	Premiere	Performance Management
BEST/1	BGS Systems, Inc.	Premiere	Performance Management
Command Center	Candle Corporation	Premiere	Performance Management
Performance Works	Landmark Systems, Inc.	Premiere	Performance Management
Dazel	Dazel Corporation	Premiere	Output Management
PSM	IBM Corporation	Premiere	Output Management
Insight Manager	Compaq	Premiere	LAN Management
HA Plus	Qualix	Premiere	Fault Management/High Availability
XPS	Informix	Pinnacle	Database Management
DCEmgmt	Santix Software GmbH	Pinnacle	DCE Management
ESSM	Platinum Technology, Inc.	Pinnacle	Database Management
DB2 ECC	IBM Corporation	Pinnacle	Database Management
ARCserve	Cheyenne	Premiere	Backup/Restore
RAIDmanager	MTI Development Corp.	Foundation	Backup/Restore
Net Archive	Platinum Technology, Inc.	Premiere	Backup/Restore
DB Vision	Platinum Technology, Inc.	Premiere	Database Management
Roboman	Heroix	Premiere	Workload Management
DSM	Enlighten Software Solutions	Premiere	Workload Management
Enterprise Desktop Management System	Novadigm, Inc.	Premiere	Workload Management
Enterprise Systems Platform	Cybermation Inc.	Premiere	Workload Management
AutoAction	Platinum Technology, Inc.	Premiere	Workload Management
Dollar Universe	Orsyp	Premiere	Workload Management
WinInstall	Seagate	Premiere	Workload Management
Scopus	Scopus Technology	Premiere	Help Desk
ClearSupport	Clarify Inc.	Premiere	Help Desk
Gladiator 3200	MTI Technology Corp.	Foundation	Help Desk
HelpDesk	Vantive	Premiere	Help Desk
Timbukto Pro	Farallon Communications Inc.	Premiere	Help Desk
IT Service Manager	Prolin Automation BV	Premiere	Help Desk

<i>Table 4 (Page 3 of 3). Overview of Tivoli/Plus and 10/Plus Modules</i>			
Product Name	Vendor	Level	Management Area
Apriori GT	Platinum Technology, Inc.	Premiere	Help Desk
Impact	Allen Systems Group, Inc.	Premiere	Help Desk
BoKS	Security Dynamics	Premiere	Security Management
OmniGuard/URM	Axent Technologies	Premiere	Security Management
Challenger	CyberSAFE Corp.	Premiere	Security Management
Stalker Monitor Agent	Haystack Labs Inc.	Premiere	Security Management
IBM Antivirus	IBM Corporation	Premiere	Security Management
CKS MyNet	CKS	Premiere	Security Management
Firewall	Checkpoint Software	Premiere	Security Management
AutoServe	Platinum Technology, Inc.	Premiere	Security Management
SafeSuite	Internet Security Systems	Premiere	Security Management
Protegrity	Protegrity	Premiere	Security Management
COMMAND/POST	Boole & Babbage	Premiere	Performance Management
Patrol	BMC Software	Premiere	Performance Management
Server Vision	Platinum Technology, Inc.	Premiere	Performance Management
Wire Trap	Platinum Technology, Inc.	Premiere	Performance Management
NetIQ App Manager	NetIQ	Premiere	Performance Management
OMC Print	Tone Software Corp.	Premiere	Output Management
Compulert	Tone Software Corp.	Premiere	Output Management
AutoDeliver	Platinum Technolgy, Inc.	Premiere	Output Management
Telalert	Telamon	Premiere	Output Management
PowerEdge Servers	DELL	Foundation	Hardware Management

A.2 More Information about the 10/Plus Association

The easiest way to obtain current information about the 10/Plus Association, the available modules, contacts, etc. is to go to the Tivoli Web page at <http://www.tivoli.com> and follow the links to the 10/Plus Association.

A.3 Press Release BGS Systems/BMC Software

The following press release details the proposed merger between BMC Software, Inc. and BGS Systems, Inc.

For more information, contact your BMC Software or BGS Systems representative.

BMC Software to Acquire BGS Systems

Strategic Acquisition Would Combine Leading Technologies to Provide Application Service Assurance Products.

HOUSTON (Feb. 2, 1998) BMC Software, Inc. (NASDAQ: BMCS) a leading developer and vendor of computer software solutions that improve the availability, performance and recoverability of business-critical applications and data, today announced the execution of a definitive merger agreement with BGS Systems, Inc. (NASDAQ: BGSS) of Waltham, MA. This strategic acquisition combines BGS's industry-leading BEST/1* performance analysis and prediction technology and products with BMC Software's suite of PATROL* application management, data management and performance optimization solutions, providing the industry's first application service assurance products.

BMC Software will pay the equivalent of \$45 in BMC Software common stock for each share of BGS common stock, using the average price per share of BMC Software common stock during a ten day trading period preceding closing. Total value of the transaction is approximately \$285 million. The transaction is to be accounted for as a pooling of interests and is to be treated as a tax-free reorganization for federal income tax purposes. The transaction is expected to close within 60 to 90 days, subject to BGS shareholder approval, regulatory approvals and other customary closing conditions.

"BGS has the best analytical and predictive tools for enterprise information systems," said Max Watson, chairman, president and chief executive officer of BMC Software. "They have delivered solid products for over twenty years, know their market segment cold and have great people. We see a clear customer requirement for correlating and interpreting the tremendous amount of performance statistics and data that PATROL provides. BMC has a road-map for providing these additional capabilities, and the BGS BEST/1 products fit perfectly within our plans. We are very enthusiastic about the potential to leverage PATROL's market momentum and product breadth with BGS's products, people and capabilities."

Watson added that the similarities between BMC Software and BGS Systems should simplify the integration of the companies and their technologies. Both companies have extensive expertise in developing high value solutions for mainframe and open systems environments and share a direct sales distribution model targeted at large enterprise systems.

Merger Strengthens PATROL's Leadership and Expands Product Lines

BMC Software currently leads the market for monitoring and event management software with its PATROL Management Suite, which provides the broadest platform coverage with support for all of the leading operating systems, database management systems and enterprise resource planning applications.

PATROL collects and delivers more performance metrics than any comparable systems management tool and offers an unrivaled ability to collect and display performance data from large and complex distributed information systems. With this merger, BMC Software will acquire BGS's industry leading BEST/1 performance analysis and prediction products and integrate them with PATROL. The combination of BEST/1's superior performance analysis and prediction capabilities with PATROL should enable BMC Software to deliver pro-active application service assurance products that have the predictive capability to anticipate system events, such as an outage or throughput bottleneck, and take corrective action well before the problem occurs. This capability would enable internal information systems organizations to provide application service assurance to their customers and reduce the administrative costs and complexity associated with today's heterogeneous and distributed information systems.

In addition to its performance management solutions for UNIX and Microsoft NT-based distributed systems, BGS provides a full line of performance management products for mainframe systems. BMC Software believes its strong presence and reputation as a mainframe systems management vendor should provide for expanded growth opportunities for the BGS mainframe product line.

A 1997 industry analyst report by International Data Corporation (IDC) defines performance management as including performance monitoring, simulation, management and resource tracking for mainframe and open systems(1). The IDC study identifies performance management as the largest systems management market segment at approximately \$1.8 billion. The acquisition of the BGS Systems BEST/1 products will position BMC Software solutions even stronger in this market.

BGS Systems President and Chairman of the Board Dr. Harold S. Schwenk, Jr. added, "BGS is the leader in assuring high levels of application performance across the enterprise and we are thrilled at the prospect of combining our technologies with BMC Software to provide a new standard of performance management. We admire BMC Software's success in establishing PATROL as a defacto standard for monitoring and event management. BMC Software's ability to establish the PATROL franchise while maintaining impressive growth in mainframe business is just the type of business model that we want to emulate. Our performance management expertise, combined with BMC Software's recognized status as an application/data management leader makes this merger a natural fit. This strategic combination of our businesses will allow our shareholders and our employees to participate in the success of the combined companies."

The companies have scheduled two conference calls today to discuss the acquisition in more detail. Parties interested in participating in the conference calls may contact the numbers below. More information regarding the acquisition can be found at <http://www.bmc.com/bgs>.

About BMC Software

BMC Software, Inc. (NASDAQ: BMCS) is a worldwide developer and vendor of more than 160 software solutions that improve the availability, performance and recoverability of business-critical applications and data. BMC Software is the world's 12th largest independent software vendor and a Forbes 500 company, with calendar 1997 revenues exceeding \$675 million. The company is headquartered in Houston, Texas with offices, distributors and agents in 50

markets, worldwide. For more information on BMC Software, access the World Wide Web at www.bmc.com or call (800) 841-2031 or (713) 918-8800.

About BGS Systems

BGS Systems, Inc., founded in 1975, is a public company with headquarters in Waltham, MA. BEST/1 Performance Assurance products are part of the BGS Enterprise-Wide Performance Assurance framework. BEST/1 software solutions give organizations the ability to avoid risk and save money by understanding how business applications are performing today and predicting how they will perform in the future. Our consistent approach to managing performance and capacity across major computing platforms and networked environments (including UNIX, NT, MVS, VM, AS/400, OpenVMS and OS/2) is in use today at over 30,000 installations throughout the world. For more information, please contact Jim McGuire, chief operating officer of BGS Systems, Inc., One First Avenue, Waltham, MA 02254-9111 or access the World Wide Web at www.bgs.com.

(1) Copyright 1997 IDC, "Systems Management Software: 1997 Worldwide Markets and Trends"

This news release contains forward looking information as indicated by words such as will, should, would, expects and other words that contemplate future events. Numerous important factors affect the company's operating results and could cause the company's actual results from the transaction discussed in this press release to differ materially from the results indicated by this press release or by any other forward looking statements made by, or on behalf of, the company. There can be no assurance that future results including the results of the combined companies will meet expectations, estimates or projections. These factors include, but are not limited to, the following: 1) the high degree of difficulty of integrating different software products and technologies; 2) the general risks associated with mergers of high technology companies, including the loss of their personnel and cultural conflicts; 3) the extreme unpredictability of the financial and operating results of the combined entity; 4) the continually increasing levels of competition in the performance and systems management markets; 5) the additional risks and important factors described in BMC Software Inc.'s Quarterly Report on Form 10-Q for the quarter ended September 30, 1997; Annual Report to Stockholders for fiscal 1997; Annual Report to Stockholders on Form 10-K for the fiscal year ended March 31, 1997 and other filings with the SEC and in BGS Systems, Inc.'s Quarterly Report on Form 10-Q for the quarter ended October 30, 1997; Annual Report to Stockholders for fiscal 1997; Annual Report to Stockholders on Form 10-K for the fiscal year ended January 31, 1997 and other filings with the SEC.

BMC Software, the BMC Software Logo and all other product or service names are registered trademarks or trademarks of BMC Software, Inc. in the USA and in other select countries. * and TM indicate USA registration or USA trademark. Other logos and product/trade names mentioned are registered trademarks or trademarks of their respective companies. BMC is an Equal Opportunity Employer.

Appendix B. Special Notices

This publication is intended to help IT Architects, IT Specialists, IT Managers and Project Managers to understand the Tivoli 10/Plus Association and how partner products can be integrated into TME 10. The information in this publication is not intended as the specification of any programming interfaces that are provided by TME 10. See the PUBLICATIONS section of the IBM Programming Announcement for TME 10 for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

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, 500 Columbus Avenue, Thornwood, NY 10594 USA.

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, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

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

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

The following document contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples contain 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.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

ADSTAR	AIX
AS/400	BookManager
DB2	IBM
MQSeries	MVS
NetView	OS/390
OS/400	OS/2
RS/6000	SystemView
System/390	VM/ESA
IBM®	

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Incorporated.

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

Pentium, MMX, ProShare, LANDesk, and ActionMedia are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Compuware and the Compuware logo are registered trademarks of Compuware Corporation.

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

Appendix C. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

C.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 301.

- *TME 10 Cookbook for AIX*, SG24-4867
- *The TME 10 Deployment Cookbook: Courier and Friends*, SG24-4976
- *A First Look at TME 10 Distributed Monitoring 3.5*, SG24-2112
- *Migrating from Systems Monitor for AIX to TME 10 Distributed Monitoring*, SG24-4936
- *TME 10 Inventory Cookbook*, SG24-2120
- *TME 10 Inventory 3.2 New Functions and Database Support*, SG24-2135
- *A First Look at TME 10 Software Distribution 3.6*, SG24-2045 (Available at a later date.)
- *Understanding Tivoli's TME 3.0 and TME 10*, SG24-4948
- *Examples of Using the TME 10 Module for MQSeries*, SG24-2134
- *ADSM Operation and Management with TME 10*, SG24-2214
- *Managing a Notes Environment with TME 10 Module for Lotus Notes*, SG24-2104
- *Developing Plus Modules for Tivoli With the TME 10 Integration Toolkit*, SG24-2007
- *Using the Applications Management Specification*, SG24-4809
- *Tivoli TME 10 Management for DCE and DFS*, SG24-2027 (Available at a later date.)

C.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
Application Development Redbooks Collection	SBOF-7290	SK2T-8037
Personal Systems Redbooks Collection	SBOF-7250	SK2T-8042

C.3 Other Publications

These publications are also relevant as further information sources:

- *BEST/1 Performance Assurance for UNIX Installation Guide*, BGS Document Number 97081-00
- *BEST/1 Monitor for UNIX User's Guide*, BGS Document Number 97083-00
- *BEST/1 Visualizer User's Guide*, BGS Document Number 96193-00
- *BEST/1 Performance Assurance for UNIX User's Guide*, BGS Document Number 97082-00
- *Site Preparation for Using BEST/1 Visualizer in an ODBC Environment*, BGS Document Number 96168-00
- *BEST/1+ for TME 10 User Guide Version 1.0*, BGS Document Number 97095-00
- *EcoTOOLS Release 3.6 Release Notes*, Compuware Document Number 05005046360
- *EcoTOOLS Release 3.6 Installation Guide*, Compuware Document Number 05005002360
- *EcoTOOLS Release 3.6 User's Guide*, Compuware Document Number 05005005360
- *EcoTOOLS Application Categorization Guide*, Compuware Document Number 05005148360

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at <http://www.redbooks.ibm.com>.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type GOPHER.WTSCPOK.ITSO.IBM.COM
- **Tools disks**

To get LIST3820s of redbooks, type one of the following commands:

```
TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
```

To get BookManager BOOKs of redbooks, type the following command:

```
TOOLCAT REDBOOKS
```

To get lists of redbooks, type one of the following commands:

```
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE
```

To register for information on workshops, residencies, and redbooks, type the following command:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1998
```

For a list of product area specialists in the ITSO: type the following command:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
```

- **Redbooks Web Site on the World Wide Web**
<http://w3.itso.ibm.com/redbooks>
- **IBM Direct Publications Catalog on the World Wide Web**
<http://www.elink.ibm.link.ibm.com/pb1/pb1>

IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserv. To initiate the service, send an e-mail note to announce@webster.ibm.link.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (<http://www.redbooks.ibm.com/redpieces.htm>). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** — send orders to:

In United States:	IBMMAIL usib6fpl at ibmmail	Internet usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications Publications Customer Support P.O. Box 29570 Raleigh, NC 27626-0570 USA	IBM Publications 144-4th Avenue, S.W. Calgary, Alberta T2P 3N5 Canada	IBM Direct Services Sortemosevej 21 DK-3450 Allerød Denmark
--	--	--

- **Fax** — send orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States) or (+1)001-408-256-5422 (Outside USA)** — ask for:

Index # 4421 Abstracts of new redbooks
Index # 4422 IBM redbooks
Index # 4420 Redbooks for last six months

- **Direct Services** - send note to softwareshop@vnet.ibm.com

- **On the World Wide Web**

Redbooks Web Site	http://www.redbooks.ibm.com
IBM Direct Publications Catalog	http://www.elink.ibm.link.ibm.com/pbl/pbl

- **Internet Listserver**

With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibm.link.ibm.com with the keyword `subscribe` in the body of the note (leave the subject line blank).

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (<http://www.redbooks.ibm.com/redpieces.htm>). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

Invoice to customer number _____

Credit card number _____

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

Index

Special Characters

/best1 22
/etc/eco/eco.conf 186, 262
/etc/eco/msg.cat 186
/etc/eco/msg.cat.m 186
/etc/eco/msg.idx 186
/etc/hosts 196
/etc/inetd.conf 186
/etc/rc.tcpip 186
/etc/rpc 186
/etc/Tivoli/tecad/bin 143
/tmp/eco_stage 229
/tmp/eventfile 257
/usr/adm/best1_5.01 86
/usr/eco 179, 189
/usr/eco/acl/oidDB 282
/usr/eco/db 223, 236, 276
/usr/eco/eamon 280
/usr/eco/etc 226, 274
/usr/eco/src/EcoCONNECT/Plus 214
/usr/local/sybase 189
/var/spool/Tivoli 136
.an 48
.anb 57
.bestrc 26
.installed 88
.md 57
.plc 35
.udf files 15
\$BEST1_COLLECT_HOME 32
\$dbhome 276
\$DISPLAY 193

Numerics

10/Plus Association 1, 16, 174
10/Plus journal 6
10/Plus module 214

A

actions 40, 125, 204
Agent Categories 205
agent code 175
Agent Error log 213
agent monitoring process 279
agent processes 221
Agent Script 208
Agent station 175
Agent Station Configuration 182
Agent Station software 186

agent stations 221, 222
agent threshold 207
AIX 15, 177, 229
AIX 4.2 82
AIX file package 231
AIX Version 4.2 21
aix6000.TAR 231
Alarm 207, 210
alarm notifications 19
alert specifications 19
AMS 2
analytical reports 79
API 7
application 55
application component management 173
Application launch 5
Application Management Specification 2
application users 2
architecture model for client/server computing 161
AS/400 13
ASCII Data File 209
ASCII event log file 223
ASCII Event Logfile 208
asynchronous monitoring 171, 241
ATN 177
Auto Discovery 197
AutoPack 144
AutoPack file 144
Availability 1, 173, 236
Available disk space 223
available swap space 236
average CPU utilization 202

B

Baan 282
Backup EcoTOOLS Database 226
Backup EcoTOOLS ecopasswd 226
BAROC file 170
BAROC files 242
baseline 60
basic integration 5
beep 208
benchmark results 20
best-of-breed selection 9
best-of-breed solution 4
BEST/1 13
BEST/1 Agent 15
BEST/1 architecture 14
BEST/1 Central Monitors 95
BEST/1 Clients 92

- BEST/1 Collect 91
- BEST/1 Collect daemon 135
- BEST/1 Collect File File Size Monitor 126
- BEST/1 Collect File Growth 126
- BEST/1 Collect Home 32
- BEST/1 Collect Home Directory 86
- BEST/1 Collect Process CPU Usage 126
- BEST/1 daemons 125
- BEST/1 Daemons Available 126, 127
- BEST/1 Data Repository 31, 86
- BEST/1 environment variables 26
- BEST/1 event adapter 135
- BEST/1 events 89
- BEST/1 File System Full 127
- BEST/1 Home Directory 86
- BEST/1 Management Console 15
- BEST/1 Move Collect Files 91, 97
- BEST/1 Performance Assurance 13, 20, 21
- BEST/1 Port Availability 127
- BEST/1 processes 95, 127
- BEST/1 Reconfigure Collect 91
- BEST/1 Remote Monitors 95
- BEST/1 Remove Old Collect Files 91, 97
- BEST/1 Semaphore IDs Availability 127
- BEST/1 Semaphore Usage 127
- BEST/1 Servers 92
- BEST/1 Share Memory Usage 127
- BEST/1 Stop Collect 91, 97
- BEST/1-Monitor log file 25
- BEST/1-Visualizer 17, 79
- BEST/1+ for TME 10 16, 81
- best1 command 27
- BEST1_ALERT 143
- BEST1_Semaphore_IDS 143
- best1_sentry.baroc 139
- best1.tar 22
- BEST1PLUS_ALIDB_after.error 87
- BEST1PLUS_ALIDB_after.output 87
- bgs 52
- BGS Systems Inc. 13
- BGS Web page 20
- bgsadmin 52
- bgsagent 126, 127
- bgsCat.cat 89
- bgscollect 126, 143
- bgsDaemonStatus 134
- bgsfreespace 134
- bgsHelpCat.cat 89
- bgsoracollect 127
- bgsproducts.cntl 26
- bgssybcollct 127
- bibliography 299
- Blind mailing 6
- bottleneck 155
- browse interval 50

- browser-equipped desktop 79
- BS/2000 161
- business application 2

C

- capacity planning 18, 69, 98, 157
- category 205
- CD-ROM 46
- CD-ROM path 182
- cd0 46
- CDF 8
- CDS 161
- Certification 6, 8
- charts 44
- circular log 209
- client application 200
- client information 200
- client machines 198
- client/server application 13, 173
- client/server relationships 196
- Collect for Oracle 104
- Collect for Sybase 104
- collect_install_after.sh 108
- collect_ora 126
- collect_syb 126
- collect_udf 126
- collected data 125
- collection 34
- collection process 32
- collector component 15
- color coding 62
- colored indicators 201
- Component Description Files 8
- component management 174
- component-specific management 2
- computing activity 79
- computing requirements 79
- Compuware Corporation 173
- conditions 40
- Configure Agent Station Filepack 221
- Configure BEST/1 Collect 91
- Configure BEST/1 Policy File 91
- Configure BEST/1 RDBMS Collect 91
- Configure EcoTOOLS Logfile Adapter 222
- consistency 83
- control file 26
- CORBA 4
- core framework services 3
- corrective action 208
- corrective actions 174
- correlation rules 245
- CPU benchmark 75
- Cpu Load Average 201
- CPU parameters 45

- CPU power 150
- CPU queue length 202
- CPU Type 31
- CPU usage 127
- CPU utilization 210
- Critical 97, 240
- critical condition 208
- critical resources 125
- current hardware 20
- custom monitors 236
- customer-specific needs 9

D

- daemon status 223, 236
- daily operation 20
- database 79
- database backups 83
- database client software 200
- database conversion utility 188
- database scenario 210
- Database server hostname 189
- database user 200
- databases 19
- dataserver 151, 154
- DB2 174
- DB2 Enterprise Control Center 11
- DCE 159
- DCE Cell 161
- DCE Cell Directory Service 161
- DCE Distributed Time Service 161
- DCE Remote Procedure Call 161
- DCE Security Service 161
- DCE threads 161
- DCE user groups 166
- DCE users 166
- DCEmgmt 159
- DCEmgmt/Cell Manager 162
- DCEmgmt/DFS Manager 162
- DCEmgmt/Event Manager 162
- DCEmgmt/Security Manager 162
- de-install the module 89
- de-installation script 90
- DEC 161
- dedicated management applications 10
- Default rule base 137, 243
- deinstall.sh 90
- Deployment 1
- desktop display 20
- destination platform types 221
- DevKit 7
- DFS 161
- Digital Alpha 20
- Digital UNIX 20, 177
- disk space 150

- disk space requirements 178
- Distributed Computing Environment 159
- Distributed File System 161
- distributed systems environment 124
- distribution process 231
- DNS 28
- domain 28
- domain icon 197
- domain models 28
- domain name 28
- drill down 19, 79
- DTS 161
- DYNIX 177

E

- e-mail 128, 208
- e-mail message 40
- eamon 223, 279
- eco_agt_fp_after.sh 232
- eco_agt_fp_before.sh 232
- Eco_Base_Event 247
- ECO_Monitors.baroc 247
- eco_profile 231
- ECO_via_Sentry2_0 247
- ECO_via_Sentry2_0.baroc 247
- eco.conf 276
- ecoacl 282
- ecoad 223, 280
- EcoCHARGEBACK Option 188
- EcoGraph 211
- ecolic 278
- ecolic daemon 223
- EcoLog_Security_Event 271
- EcoMeter gauge 211
- EcoMeter gauges 201
- ECONTX86.EXE 190
- ecopasswd 274
- ecoplus.baroc 247
- EcoTOOLS 173
- EcoTOOLS Agent Dispatcher service 192
- EcoTOOLS Agent Help 226
- EcoTOOLS Agent Station Monitors 223, 224
- EcoTOOLS Agent Stations 222, 223
- EcoTOOLS CD-ROM 190
- EcoTOOLS Central Monitors 222, 224
- EcoTOOLS daemons 186
- EcoTOOLS discovery process 196
- EcoTOOLS dispatcher 280
- EcoTOOLS dispatcher daemon 223
- EcoTOOLS domains 194
- EcoTOOLS event adapter 222
- EcoTOOLS License Station Monitors 223, 224
- EcoTOOLS License Stations 223
- EcoTOOLS log 213

- EcoTOOLS Management Station Monitors 224
- EcoTOOLS Management Stations 223
- EcoTOOLS map 196
- EcoTOOLS Plus Module for TME 10 214
- EcoTOOLS processes 221
- EcoTOOLS Sentry Indicators 224
- EcoTOOLS Services 191
- EcoTOOLS source binaries 229
- EcoToolsDb 190
- EcoTOOLSPlus Module for TME 10 Release 3.6 216
- electronic mail 26
- enhanced problem resolution 26
- enterprise management solution 174
- enterprise systems management 16, 157
- enterprise-wide DCE installations 160
- enterprise-wide management solution 2
- error message 187
- evaluation 61
- evaluation rules 245
- event adapter for BEST/1 151
- event adapters 221, 242
- event classes 137, 255
- event console 243
- event forwarding 265
- event group 243
- Event integration 5
- Event log 213
- event management 241
- Event Server 136, 242
- event status monitors 201
- events 150
- exceptions 241
- Execution mode 103
- Execution Target 103
- exit commands 188

F

- file growth 127
- file package definition 108
- file package distribution 110
- file packages 91, 221
- file system 127
- file system utilization 236
- file systems 95
- filters 40
- flash 208
- flexibility 2
- flexible framework 2
- forecasting 79
- Foundation Members 4
- future configurations 20
- future performance 15

G

- GID 50
- GIS 177
- graph 209
- graphical representations 15
- graphs 44, 79
- Group name 50
- group of nodes 196
- guidelines 62

H

- hard disk 46
- hard disks 19
- hardware management application 5
- hardware simulation 20
- hdisk0 43
- heterogeneous environment 173
- high load 66
- historical data 79
- Host ID 187
- hostname 182
- HP 177
- HP 90000 229
- HP-UX 15, 99, 177, 229
- HP-UX 9.05 82
- HTML code 79

I

- I/T environment 2
- IBM RS6000 229
- indicator collection 92
- Informix 174
- init.tecad_logfile start 264
- init.tecad_logfile stop 264
- input device 22
- Install BEST/1 Adapter 91, 95
- INSTALL_COLLECT 116
- install_ecotools 189, 231
- install_ecotools.log 180
- install_product 85
- install.best1 22
- install.cd 180
- install.collect-aix 99
- install.collect-hp 99
- install.collect-ncr 99
- install.collect-osf 99
- install.collect-solaris 99
- install.collect-sunos 99
- installation 181
- installation errors 182
- installation media 22, 85, 179
- installation options 85

- installation password 188
- installation procedure 21
- installation program 191
- installation script 22, 104, 182
- installation scripts 99
- installation source 216
- installation status messages 184
- InstallShield 144
- integration levels 3
- integration tools 3
- inventory 213
- IRIX 178
- ITSO.RAL 194

J

- justify expenditures 79

L

- license information 186, 278
- license server 187, 194, 221, 278
- License station 175, 222
- License Station Configuration 182
- license tag 187
- licensing 26
- life expectancy 79
- List User Access Roles 227
- Log Spill Interval 119
- logical groups 35
- logical unit of work 55
- logical volumes 62
- look and feel 10
- Lotus Notes 282

M

- machine type 234
- machine-dependent software 221
- management console 206
- management disciplines 1
- Management station 175, 222
- Management Station Configuration 182
- management working groups 8
- map icon 208
- MASF 79
- memory 150
- message catalog files 89
- meters 202
- MIB integration 5
- Microsoft SQL 174
- middleware layer 161
- mission critical applications 173
- model 69
- modelling techniques 20

- models 58
- Monitor User 104
- monitoring 2
- monitoring agents 194, 204
- monitoring probes 132
- Motif 21
- mount point 180
- multiple instances 175
- multiple versions 25
- Multivariate Adaptive Statistical Filter 79
- MVS 13

N

- NCR 20, 177
- nested file package 108
- NetView 9
- NetView Association 10
- network 125
- network computing management 3
- Network Information Service 196
- network performance 60
- networked applications 173
- NIS 196
- nodes 35
- nonstop availability management 173
- Normal 97
- Notice Groups 111
- number of users 59
- NuTCRACKER Kernel 193

O

- objectives 62
- OMG 4
- online help 27, 63
- OP_SYS file 22
- open integration platform 4
- open management strategy 2
- Open Software Foundation 161
- OpenVMS 13
- operating system process 161
- operating system type 182, 227
- Operations 1
- ora_udf 126
- Oracle 20, 127, 174
- Oracle applications 282
- Oracle database servers 127
- OS/390 13, 161
- OS/400 161
- oserv 151, 154
- OSF 161
- Output Destination 103, 105
- Output Format 103
- outputs 204

P

- pager message 208
- paging space 20
- partner applications 2
- peak utilization periods 60
- Pentium 177
- PeopleSoft 282
- performance 60, 173
- performance behavior 79
- performance data 15
- performance figures 62
- performance management 13, 98, 157, 175
- Performance Management Working Group 174
- performance planning 18
- performance reports 58
- periodic tasks 71
- pie charts 62
- Pinnacle Club 4
- Pinnacle level application 159
- Pinnacle partners 3
- plan a TME 10 environment 149
- platform types 221
- point products 3, 10
- Pointcast Channel 6
- policies 35
- policy files 93
- policy region 100
- port 127
- ports 95
- postmsg.sh 89
- pre-installation check 181
- predict_syntax.doc 76
- Premier Club 4
- Premier level member 16, 174
- Premier partners 3
- preprocess.template 78
- Private Scenario 209
- process ID 120
- process name 54
- processes 20, 124
- processor model 150
- processor models 156
- processor utilization 155
- product license information 26
- profile 92
- profile manager 92, 258
- profile managers 221, 234
- progress status 184
- Public Scenario 209

R

- RDBMS 5, 58, 73, 174
- RDBMS Sample Interval 33, 74

- re-evaluate 69
- real-time chart data 19
- real-time monitors 174
- Remote Collect Home 105
- Remote Data Repository 104
- Remote Install directory 104
- remote procedure calls 176
- Remote Server Directory 105
- remote sites 124
- report file 57
- reports 47
- resource 91
- resource IDs 108
- response level 92, 237
- Responses 134
- Rest Of the World 60
- Restore EcoTOOLS Database 226
- restricted system resource 125
- RIM_Sybase_prog 155
- root 52
- ROW 60
- RPC 161, 176
- RPC mechanism 280
- RS/6000 20, 179
- RS/6000 43P 155
- rule base 243
- rule base directory 243
- rules 89, 137
- Run EcoTOOLS Process Monitor 227
- Run Scenario via ecocli 227
- run-time environment 121
- run-time libraries 21

S

- sales database 55
- sales department 55
- sample applications 7
- Santix Software GmbH 159
- SAP R/3 177, 282
- scenario 194
- SCO UNIXWARE 178
- script 40, 150, 208
- script files 89
- Security 1
- semaphore IDs 89, 125
- semaphores 21, 89, 95, 125, 127
- seminfo 89
- Sentry 224
- sentry_engine 131, 151
- Sentry.baroc 247
- Sentry2_0_Base 247
- Sequent 177
- service level objectives 62, 150
- service level threshold 154

- service levels 15
- Setup BEST/1 Collect 91
- Setup BEST/1 EventServer 94
- Setup BEST/1 Policy File 91
- Setup TEC EventServer for EcoTOOLS 222
- Severe 97
- severity 208
- severity level 92, 125, 128, 240
- SGI 178
- shared libraries 21
- shared memory 95, 125, 127
- shared memory segments 21
- shell scripts 221
- SLO 64, 150
- snapshot 144
- SNMP Manager 44
- SNMP trap 40
- software distribution 2
- software prerequisites 20
- Solaris 15, 99, 178
- Solaris 2.5.1 82
- solution design 285
- sources 242
- specialized management applications 2
- Spill Interval 33, 74
- Sponsorship opportunities 6
- SQL monitor server name 189
- SQL server 211
- SQL_Central 200
- staging directory 229
- standardization groups 4
- standards-based architecture 10
- Start/Stop eamon daemon 226
- Start/Stop ecoad daemon 226
- Start/Stop ecolic daemon 226
- stderr 244
- stress test 151
- structured ASCII logfile 257
- Su_Success 271
- subscriber list 223, 234
- subscription list 110
- sufficient capacity 125
- Sun SPARC 178
- Sun Sparc Series 20
- SunOS 178
- SuperSystem 6
- Swap space available 223
- syb_udf 126
- Sybase 20, 127, 174, 188, 189
- Sybase database server 127
- Sybase environment variable 189
- Sybase server 244
- Sybase user ID 189
- synchronous monitoring 242
- synchronous monitors 98

- system administrator (sa) 189
- system changes 66
- system performance 174
- system resources 20, 95
- System Sample Interval 119
- System Type 229
- systems management applications 2
- systems management framework 174
- systems management functions 1, 157
- systems management solution 98
- SystemView 1

T

- table 45
- tape read 21
- tar file 275
- task endpoint 228
- TCP/IP 177
- TCP/IP hostname 31
- TCP/IP port 25
- TEC 150
- TEC adapters 89
- TEC Dispatcher 151
- TEC event 128, 237, 241, 270
- TEC event adapter 89
- TEC logfile adapter 257
- TEC master process 151
- TEC process 151
- TEC Reception engine 151
- TEC rule base 242
- TEC rule bases 136
- TEC Rule engine 151
- TEC Server 81
- TEC Sybase database process 151
- TEC_CLASSES 245
- tec_dispatch 151
- tec_reception 151
- tec_rule 151
- TEC_RULES 245
- tec_server 151
- tec_task 151
- TEC_TEMPLATES 245
- tecad_best1.baroc 139
- tecad_best1.conf 89
- tecad_best1.pl 89, 151
- tecad_best1.sh 89
- tecad_collect.pl 89
- tecad.logfile.fmt 258
- telnet 189
- Terminals 67
- test criteria 8
- thermometer 225
- third parties 4
- thresholds 38, 62

- tight integration 5
- Tivoli core applications 1, 81
- Tivoli database 136, 233
- Tivoli Enterprise Console 135
- Tivoli Integration Toolkit 16, 174
- Tivoli Management Environment 1
- Tivoli Management Framework 1
- Tivoli Management Region 81
- Tivoli notice 128
- Tivoli Notice Board 234
- Tivoli object database 83, 215
- Tivoli Object Request Broker process 151
- Tivoli Ready branding 6
- Tivoli Root administrator 237
- Tivoli Systems 1
- Tivoli transaction class 154
- Tivoli/Plus Module for ADSM 11
- Tivoli/Sentry 3.0 Monitors 81
- Tivoli/Sentry 3.0.2 Monitors Patch 81
- Tivoli/Sentry Patch 3.0.2 81
- Tivoli/Sentry Version 3.0 81
- TivoliPlus collection 88
- TivoliPlus icon 215
- TME 10 1, 13, 98, 157
- TME 10 application services 3, 6
- TME 10 commands
 - objcall 88
 - odstat 132
 - wbkupdb 83, 215
 - wchkdb 83
 - wexpftfp 107, 232
 - wgetattr 89
 - wlookup 107, 232
 - wlseng 132
 - wlsmon 89
 - wlsra 134
 - wlstime 134
 - wpostemsg 141, 150, 247
 - wsprbe 134
 - wtdumpri 152, 255
- TME 10 core applications 160
- TME 10 desktop 16, 84, 111, 174, 214
- TME 10 Distributed Monitoring 1, 16, 21, 83, 94, 124, 169, 174, 221, 236
- TME 10 Distributed Monitoring profile 125
- TME 10 Enterprise Console 1, 16, 21, 89, 95, 169, 174, 214, 222
- TME 10 Enterprise Console 3.1 Logfile Adapter 260
- TME 10 Enterprise Console 3.1. Server 81
- TME 10 Enterprise Console server 150
- TME 10 environment 81
- TME 10 Framework 16, 81, 174
- TME 10 Framework 3.1 Revision C 81
- TME 10 Framework services 160
- TME 10 Framework Toolkit 7
- TME 10 infrastructure 149
- TME 10 Integration Toolkit 3
- TME 10 Inventory 1
- TME 10 Management Working Group 16
- TME 10 Module for Lotus Notes 11
- TME 10 Module for MQSeries 11
- TME 10 Module for SAP R/3 11
- TME 10 Software Distribution 1, 16, 21, 91, 174, 221, 227
- TME 10 Software Distribution notice 112
- TME 10 Software Distribution Notice Group 234
- TME 10 Software Distribution Release 3.1 81
- TME 10 Tasks and Jobs 16, 174
- TME 10 User Administration 1, 162
- TMF 1
- TMR 81, 83
- TMR role 85
- TMR server 17, 150, 174, 219
- toolkits 2
- Trade show participation 6
- traffic light icon 204
- Transaction Classes 53
- transactions 55
- trend graph 79
- turnkey performance graphs 79

U

- UDF 20
- UDF files 73, 127
- UID 50
- Universal Data Format 20
- UNIX 13, 20, 161, 175
- UNIX Agent Station 177
- UNIX commands
 - cp 99
 - grep 120
 - ls 120, 245
 - ps 120
 - rpcinfo 280
 - su 269
 - tail 269
- UNIX login name 183
- UNIX machines 81
- UNIX operating systems 99
- Unix Options 109
- UNIX server 22
- Update EcoTOOLS Agents 221
- upgrading hardware 59
- user class 50
- user name 41, 50
- user-defined text file 209
- users 20
- utilization 79

V

Visualizer 144
Visualizer files 71
Visualizer interval 77
VM 161
VM/ESA 13
VMS 161

W

Warning 97
Web 20, 79
Windows 161
Windows Explorer 190
Windows NT 13, 21, 81, 144, 175, 178, 190
Windows NT Agent Station 177
WinZip Self-Extractor 190
wizards 7
work load 20
www.bgs.com 20

X

X Users 67
X Windows session 189
X11 21
X11R5 run-time environment 177

ITSO Redbook Evaluation

An Industry Around the Tivoli Framework: Examples from the 10/Plus Association
SG24-2122-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at <http://www.redbooks.ibm.com>
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@vnet.ibm.com

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes____ No____

If no, please explain:

What other redbooks would you like to see published?

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

