AUTHORIZED DOCUMENTATION

User Guide

Novell®
Sentinel 6.1 Rapid Deployment

SP2

April, 2011

www.novell.com

Legal Notices

Novell, Inc., makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc., makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. See the Novell International Trade Services Web page (http://www.novell.com/info/exports/) for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Copyright © 1999-2011 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc.
404 Wyman Street, Suite 500
Waltham, MA 02451
U.S.A.
www.novell.com

Online Documentation: To access the latest online documentation for this and other Novell products, see the Novell Documentation Web page (http://www.novell.com/documentation).

Novell Trademarks

For Novell trademarks, see the Novell Trademark and Service Mark list (http://www.novell.com/company/legal/trademarks/tmlist.html).

Third-Party Materials

All third-party trademarks are the property of their respective owners.

Contents

About This Guide **15**

1 Managing Sentinel Rapid Deployment Through the Web Interface **17**

1.1 Accessing the Novell Sentinel Web Interface . . . 17
1.2 Applications and Installers . . . 17
1.3 Reporting . . . 19
1.3.1 Running Reports . . . 19
1.3.2 Viewing Reports . . . 22
1.3.3 Scheduling a Report . . . 24
1.3.4 Managing Reports . . . 25
1.4 Searching Events . . . 29
1.4.1 Enabling the Search Option in Web User Interface . . . 29
1.4.2 Running an Event Search . . . 30
1.4.3 Viewing Search Results . . . 33
1.4.4 Event Fields . . . 36

2 Sentinel Control Center **41**

2.1 Log In to the Sentinel Control Center . . . 41
2.1.1 Linux . . . 41
2.1.2 Windows . . . 41
2.2 About Sentinel Control Center . . . 42
2.2.1 Active Views . . . 43
2.2.2 Incidents . . . 43
2.2.3 iTRAC . . . 43
2.2.4 Analysis . . . 44
2.2.5 Advisor . . . 44
2.2.6 Admin . . . 44
2.2.7 Correlation . . . 44
2.2.8 Event Source Management . . . 45
2.2.9 Solution Packs . . . 45
2.2.10 Identity Integration . . . 45
2.3 Introduction to the User Interface . . . 45
2.3.1 Menu Bar . . . 46
2.3.2 Toolbar . . . 47
2.3.3 Tabs . . . 48
2.3.4 Frames . . . 48
2.3.5 Using the Sentinel Control Center to Navigate . . . 49
2.3.6 Changing the Appearance of the Sentinel Control Center . . . 49
2.3.7 Saving User Preferences . . . 50
2.3.8 Changing Password . . . 51
2.3.9 Configuring the Attachment Viewer . . . 51

3 Active Views Tab **53**

3.1 Understanding Active Views . . . 53
3.2 Introduction to the User Interface . . . 54
3.3 Reconfiguring Total Display Time . . . 57
3.4 Viewing Real-Time Events . . . 57
3.4.1 Resetting the Parameters and Chart Type of an Active View . . . 60

3.4.2 Rotating a 3D Bar or Ribbon Chart . . . 61
3.5 Showing and Hiding Event Details . . . 61
3.6 Sending Mail Messages about Events and Incidents . . . 62
3.7 Creating Incidents . . . 63
3.8 Viewing Events That Trigger Correlated Events . . . 64
3.9 Investigating an Event or Events . . . 65
3.9.1 Investigate: Event Query . . . 66
3.9.2 Investigate: Graph Mapper . . . 66
3.9.3 Historical Event Query . . . 67
3.9.4 Active Browser . . . 68
3.10 Viewing the Advisor Data . . . 70
3.11 Viewing the Asset Data . . . 71
3.12 Viewing Vulnerabilities . . . 73
3.13 Ticketing System Integration . . . 77
3.14 Viewing User Information . . . 77
3.15 Using Custom Menu Options with Events . . . 77
3.16 Managing Columns in a Snapshot or Navigator Window . . . 78
3.17 Taking a Snapshot of a Navigator Window . . . 79
3.18 Sorting Columns in a Snapshot . . . 79
3.19 Closing a Snapshot or Navigator . . . 79
3.20 Adding Events to an Incident . . . 80

4 Correlation Tab 83

4.1 Understanding Correlation . . . 83
4.1.1 Technical Implementation . . . 84
4.2 Introduction to the User Interface . . . 85
4.3 Correlation Rules . . . 85
4.3.1 Opening the Correlation Rule Manager . . . 86
4.3.2 Creating a Rule Folder . . . 86
4.3.3 Renaming a Rule Folder . . . 86
4.3.4 Deleting a Rule Folder . . . 86
4.3.5 Creating a Correlation Rule . . . 86
4.3.6 Creating Correlation Rules . . . 87
4.3.7 Deploying and Undeploying Correlation Rules . . . 95
4.3.8 Enabling and Disabling Rules . . . 96
4.3.9 Renaming and Deleting a Correlation Rule . . . 96
4.3.10 Sorting Correlation Rules . . . 96
4.3.11 Moving a Correlation Rule . . . 97
4.3.12 Importing a Correlation Rule . . . 97
4.3.13 Exporting a Correlation Rule . . . 98
4.4 Dynamic Lists . . . 98
4.4.1 Adding a Dynamic List . . . 99
4.4.2 Modifying a Dynamic List . . . 100
4.4.3 Deleting a Dynamic List . . . 100
4.4.4 Removing Dynamic List Elements . . . 100
4.4.5 Using a Dynamic List in a Correlation Rule . . . 100
4.5 Correlation Engine . . . 101
4.5.1 Starting or Stopping a Correlation Engine . . . 102
4.5.2 Renaming a Correlation Engine . . . 102
4.6 Correlation Actions . . . 102
4.6.1 Configuring a Correlated Event . . . 103
4.6.2 Adding to a Dynamic List . . . 104
4.6.3 Removing a Value from a Dynamic List . . . 105
4.6.4 Executing a Command . . . 106

4.6.5 Creating an Incident . . . 107
4.6.6 Sending an E-mail . . . 108
4.6.7 Imported JavaScript Action Plugins . . . 108

5 Incidents Tab 109

5.1 Understanding an Incident . . . 109
5.2 Introduction to User Interface . . . 109
5.2.1 Incident View . . . 110
5.2.2 Incident . . . 110
5.3 Manage Incident Views . . . 111
5.3.1 Adding a View . . . 111
5.3.2 Modifying a View . . . 114
5.3.3 Deleting a View . . . 115
5.3.4 Default View . . . 115
5.4 Manage Incidents . . . 115
5.4.1 Creating Incidents . . . 115
5.4.2 Viewing an Incident . . . 116
5.4.3 Attaching Workflows to Incidents . . . 116
5.4.4 Adding Notes to Incidents . . . 117
5.4.5 Adding Attachments to Incidents . . . 117
5.4.6 Executing Incident Actions . . . 118
5.4.7 E-Mailing an Incident . . . 119
5.4.8 Modifying Incidents . . . 120
5.4.9 Deleting Incidents . . . 121
5.5 Switch between Existing Incident Views . . . 121

6 iTRAC Workflows 123

6.1 Understanding iTRAC Workflows . . . 123
6.2 Introduction to the User Interface . . . 124
6.3 Template Manager . . . 125
6.3.1 Default Templates . . . 125
6.4 Template Builder Interface . . . 126
6.4.1 Creating Templates . . . 127
6.4.2 Managing Templates . . . 128
6.5 Steps . . . 129
6.5.1 Start Step . . . 130
6.5.2 Manual Step . . . 130
6.5.3 Decision Step . . . 134
6.5.4 Mail Step . . . 134
6.5.5 Command Step . . . 134
6.5.6 Activity Step . . . 135
6.5.7 End Step . . . 136
6.5.8 Adding Steps to a Workflow . . . 136
6.5.9 Managing Steps . . . 136
6.6 Transitions . . . 141
6.6.1 Unconditional Transitions . . . 141
6.6.2 Conditional Transitions . . . 142
6.6.3 Else Transitions . . . 146
6.6.4 Timeout Transitions . . . 146
6.6.5 Alert Transitions . . . 147
6.6.6 Error Transition . . . 148
6.6.7 Managing Transitions . . . 148
6.7 Activities . . . 149
6.7.1 Incident Command Activity . . . 150
6.7.2 Incident Internal Activity . . . 151

6.7.3 Eradication Activity . . . 151
6.7.4 Incident Composite Activity . . . 151
6.7.5 Creating iTRAC Activities . . . 151
6.7.6 Managing Activities . . . 154
6.8 Process Management . . . 155
6.8.1 Instantiating a Process . . . 156
6.8.2 Automatic Step Execution . . . 156
6.8.3 Manual Step Execution . . . 157
6.8.4 Displaying Status . . . 157
6.8.5 Displaying the Status of a Process . . . 157
6.8.6 Changing Views in the Process Manager . . . 158
6.8.7 Starting or Terminating a Process . . . 159

7 Work Items 161

7.1 Work Item Summary . . . 161
7.2 Processing a Work Item . . . 164
7.2.1 Accepting and Completing a Work Item . . . 164
7.3 Managing Work Items of Other Users . . . 165

8 Analysis Tab 167

8.1 Introduction to the User Interface . . . 167
8.1.1 Top Ten Dashboard . . . 167
8.2 Offline Query . . . 169
8.2.1 Creating an Offline Query . . . 169
8.2.2 Viewing, Exporting, or Deleting an Offline Query . . . 170

9 Advisor Usage and Maintenance 171

9.1 Understanding Advisor . . . 171
9.2 Understanding Exploit Detection . . . 172
9.2.1 How Exploit Detection Works . . . 172
9.2.2 Generating the Exploit Detection File . . . 174
9.2.3 Viewing the Events . . . 174
9.3 Introduction to the Advisor User Interface . . . 174
9.3.1 The Advisor Window . . . 175
9.3.2 Processing the Advisor Feed . . . 176
9.3.3 Configuring the Advisor Products for Exploit Detection . . . 177
9.4 Downloading the Advisor Feed . . . 178
9.4.1 Configuring the Sentinel Server for Automated Downloads . . . 178
9.4.2 Downloading the Advisor Feed Manually . . . 179
9.5 Viewing the Advisor Status . . . 179
9.6 Viewing the Advisor Data . . . 181
9.7 Resetting the Advisor Password . . . 182
9.8 Deleting the Advisor Data . . . 182
9.9 Advisor Audit Events . . . 182

10 Download Manager 183

10.1 Understanding the Download Manager User Interface . . . 183
10.2 Creating a Download Configuration . . . 184
10.3 Editing a Download Configuration . . . 187
10.4 Downloading the Feed Instantly . . . 187
10.5 Deleting a Download Configuration . . . 188

10.6 Audit Events for the Download Manager 188

11 Event Source Management 189

11.1 Understanding Event Source Management 189
11.1.1 Using Event Source Management 189
11.1.2 Plug-In Repository 190
11.1.3 Auxiliary Files 190
11.2 Introduction to the User Interface 190
11.2.1 Menu Bar 191
11.2.2 Toolbar 192
11.2.3 Zoom 192
11.2.4 Frames 193
11.3 Live View 198
11.3.1 Graphical ESM View 198
11.3.2 Tabular ESM View 200
11.3.3 Right-Click Menu 200
11.4 Components of Event Source Hierarchy 202
11.4.1 Component Status Indicators 203
11.4.2 Adding Components to the Event Source Hierarchy 204
11.4.3 Collectors 204
11.5 Debugging 220
11.5.1 Collector Workspace and Collector Directory 220
11.5.2 Debugging Proprietary Collectors 221
11.5.3 Debugging JavaScript Collectors 223
11.5.4 Using the Raw Data Tap to Generate a Flat File 227
11.6 Exporting a Configuration 228
11.7 Importing a Configuration 230
11.7.1 Enabling or Disabling the Import Configuration 231
11.7.2 Resetting the Layout 233
11.7.3 Undoing the Layout 233
11.7.4 Redo Layout 233
11.8 Event Source Management Scratchpad 234

12 Administration 235

12.1 Understanding the Admin Tab 235
12.2 Introduction to the User Interface 236
12.3 Servers View 237
12.3.1 Monitoring a Process 238
12.3.2 Creating a Servers View 238
12.3.3 Starting, Stopping, and Restarting Processes 239
12.4 Filters 239
12.4.1 Public Filters 239
12.4.2 Private Filters 240
12.4.3 Global Filters 240
12.4.4 Configuring Public and Private Filters 243
12.4.5 Color Filter Configuration 245
12.5 Configure Menu Options 248
12.5.1 Adding an Option to the Event Menu 250
12.5.2 Cloning an Event Menu Option 251
12.5.3 Modifying an Event Menu Option 252
12.5.4 Viewing Event Menu Option Parameters 252
12.5.5 Activating or Deactivating an Event Menu Option 252
12.5.6 Rearranging Event Menu Options 252
12.5.7 Deleting an Event Menu Option 252

12.5.8 Editing Your Event Menu Browser Settings . . . 253
12.6 DAS Statistics . . . 254
12.7 Mapping . . . 255
12.7.1 Adding Map Definitions . . . 257
12.7.2 Adding a Number Range Map Definition . . . 259
12.7.3 Editing Map Definitions . . . 262
12.7.4 Deleting Map Definitions . . . 262
12.7.5 Updating Map Data . . . 263
12.8 Event Configuration . . . 265
12.8.1 Event Mapping . . . 265
12.8.2 Renaming Tags . . . 269
12.9 Report Data Configuration . . . 270
12.9.1 Disabling or Enabling a Summary . . . 271
12.9.2 Viewing Information for a Summary . . . 272
12.9.3 Checking the Validity of a Summary . . . 272
12.9.4 Query the Event Files for a Summary . . . 273
12.9.5 Running the Event Files for a Summary . . . 274
12.10 User Configurations . . . 275
12.10.1 Opening the User Manager Window . . . 275
12.10.2 Creating a User Account . . . 275
12.10.3 Modifying a User Account . . . 280
12.10.4 Viewing Details of a User Account . . . 280
12.10.5 Cloning a User Account . . . 280
12.10.6 Deleting a User Account . . . 281
12.10.7 Terminating an Active User Session . . . 281
12.10.8 Adding an iTRAC Role . . . 281
12.10.9 Deleting an iTRAC Role . . . 282
12.10.10 Viewing the Details of a Role . . . 282

13 Sentinel Data Manager 283

13.1 Understanding the Sentinel Data Manager . . . 283
13.2 Using the SDM GUI . . . 283
13.2.1 Prerequisites . . . 283
13.2.2 Starting the SDM GUI . . . 284
13.2.3 Connecting to the Database . . . 284
13.2.4 Partitions Tab . . . 285
13.2.5 Tablespaces Tab . . . 288
13.2.6 Partition Configuration . . . 289
13.2.7 Managing Disk Space Allocation . . . 291
13.3 Using the SDM Command Line . . . 291
13.3.1 Prerequisite . . . 292
13.3.2 Syntax of the SDM command . . . 292
13.3.3 Starting the SDM GUI . . . 292
13.3.4 Saving Connection Properties for Sentinel Data Manager . . . 292
13.3.5 Adding Partitions . . . 293
13.3.6 Dropping Partitions . . . 294
13.3.7 Viewing Partition Summaries . . . 295
13.3.8 Archiving Data . . . 296
13.3.9 Importing Data . . . 297
13.3.10 Deleting Imported Data . . . 298
13.3.11 Viewing Sentinel Database Space Usage . . . 299

14 Utilities 301

14.1 Introduction to Sentinel Utilities . . . 301
14.2 Starting and Stopping a Sentinel Server . . . 301

14.2.1 Starting a Sentinel Server . . . 302
14.2.2 Stopping a Sentinel Server . . . 302
14.3 Sentinel Scripts . . . 302
14.3.1 Operational Scripts . . . 302
14.3.2 Troubleshooting Scripts . . . 304
14.4 Version Information . . . 305
14.4.1 Executable Version Information . . . 305
14.4.2 Sentinel .jar Version Information . . . 305
14.5 Database Cleanup . . . 306
14.5.1 Components . . . 306
14.5.2 Prerequisites . . . 307
14.5.3 Running Clean_Database.sh . . . 307
14.6 Connecting to PostgreSQL Database Through Command Line . . . 308
14.7 Backup and Restore Utility . . . 309
14.7.1 Parameters for the Backup and Restore Utility Script . . . 309
14.7.2 Using the Backup and Restore Utility Script . . . 310
14.8 Updating Your License Key . . . 311

15 Quick Start 313

15.1 Security Analysts . . . 313
15.1.1 Active Views Tab . . . 313
15.1.2 Exploit Detection . . . 314
15.1.3 Asset Data . . . 315
15.1.4 Event Query . . . 316
15.2 Creating Incidents . . . 317
15.3 iTRAC . . . 318
15.3.1 Instantiating a Process . . . 318
15.4 Correlation . . . 328
15.4.1 Creating a Simple Correlation Rule . . . 329
15.4.2 Deploying the Simple Correlation Rule . . . 329
15.4.3 Viewing the Events that Triggered Your Correlated Event . . . 330

16 Solution Packs 331

16.1 Solution Packs . . . 331
16.1.1 Components of a Solution Pack . . . 332
16.1.2 Permissions for Using Solution Packs . . . 333
16.2 Solution Manager . . . 334
16.2.1 Solution Manager Interface . . . 335
16.3 Managing Solution Packs . . . 336
16.3.1 Importing Solution Packs . . . 337
16.3.2 Opening Solution Packs . . . 339
16.3.3 Installing Content from Solution Packs . . . 341
16.3.4 Implementing Controls . . . 346
16.3.5 Testing Controls . . . 347
16.3.6 Uninstalling Controls . . . 347
16.3.7 Viewing Solution Pack Status . . . 349
16.3.8 Deleting Solution Packs . . . 351
16.4 Solution Designer . . . 352
16.4.1 Solution Designer Interface . . . 352
16.4.2 Connection Modes . . . 354
16.4.3 Creating a Solution Pack . . . 355
16.4.4 Managing Content Hierarchy Nodes . . . 356
16.4.5 Adding Content to a Solution Pack . . . 357
16.4.6 Documenting a Solution Pack . . . 359

16.4.7 Editing a Solution Pack . . . 360
16.5 Deploying an Edited Solution Pack . . . 361

17 Action Manager and Integrator 363

17.1 Action Manager . . . 363
17.2 Action Plug-Ins . . . 365
17.2.1 Importing JavaScript Action Plug-Ins . . . 365
17.2.2 Importing JavaScript Files . . . 368
17.3 Actions . . . 376
17.3.1 Creating Actions . . . 376
17.3.2 Editing Actions . . . 377
17.3.3 Deleting Actions . . . 377
17.3.4 Using JavaScript Actions . . . 377
17.3.5 Developing JavaScript Actions . . . 378
17.4 Integrator Manager . . . 382
17.4.1 Permissions for Using Integrators . . . 383
17.5 Integrator Plug-Ins . . . 384
17.5.1 Importing Integrator Plugins . . . 384
17.5.2 Deleting Integrator Plug-Ins . . . 384
17.6 Integrators . . . 385
17.6.1 Creating an Integrator Instance . . . 385
17.6.2 Editing an Integrator Instance . . . 385
17.6.3 Deleting an Integrator Instance . . . 385
17.6.4 Integrator Connection Status . . . 385
17.6.5 Viewing Integrator Health Details . . . 386
17.6.6 Integrator Events Query . . . 387
17.6.7 Using Integrators from Actions . . . 389

18 Identity Integration 391

18.1 Integration with Novell Identity Manager . . . 392
18.2 Identity Browser . . . 395
18.2.1 Searching Profiles . . . 395
18.2.2 Viewing Profile Details . . . 396
18.2.3 Using the Clipboard Functionality . . . 399
18.3 Reports . . . 399

A Sentinel Rapid Deployment Architecture 401

A.1 Sentinel Rapid Deployment Features . . . 401
A.2 Functional Architecture . . . 401
A.3 Architecture Overview . . . 403
A.3.1 Communication Server . . . 404
A.3.2 Sentinel Events . . . 405
A.3.3 Event Source Management . . . 409
A.3.4 Application Integration . . . 410
A.3.5 Time . . . 410
A.3.6 System Events . . . 411
A.3.7 Processes . . . 412
A.4 Logical Architecture . . . 414
A.4.1 Collection and Enrichment Layer . . . 415
A.4.2 Business Logic Layer . . . 419
A.4.3 Presentation Layer . . . 426

B System Events for Sentinel 431

B.1 Advisor Audit Events . . . 431
B.1.1 Advisor Update Successful . . . 431
B.1.2 Advisor Update Failure . . . 432
B.2 Download Manager Audit Events . . . 432
B.2.1 Download Successful . . . 432
B.2.2 Download Failed . . . 433
B.2.3 Download Config Updated . . . 433
B.2.4 Download Config Added . . . 433
B.2.5 Download Config Removed . . . 434
B.3 Authentication Events . . . 434
B.3.1 Authentication . . . 434
B.3.2 Creating Entry For External User . . . 435
B.3.3 Duplicate User Objects . . . 435
B.3.4 Failed Authentication . . . 435
B.3.5 Locked Account . . . 436
B.3.6 No Such User Event . . . 436
B.3.7 Too Many Active Users . . . 437
B.3.8 User Discovered . . . 437
B.3.9 User Logged In . . . 437
B.3.10 User Logged Out . . . 438
B.4 User Management . . . 438
B.4.1 Add Users To Role . . . 438
B.4.2 Create Role . . . 439
B.4.3 Create User . . . 439
B.4.4 Creating User Account . . . 439
B.4.5 Delete Role . . . 440
B.4.6 Deleting User Account . . . 440
B.4.7 Locking User Account . . . 440
B.4.8 Remove Users From Role . . . 441
B.4.9 Resetting Password . . . 441
B.4.10 Unlocking User Account . . . 441
B.4.11 Updating User . . . 442
B.5 Database Event Management . . . 442
B.5.1 Diskspace Usage Reached Lower Threshold . . . 443
B.5.2 Diskspace Usage Reached Upper Threshold . . . 443
B.5.3 Dropping the Oldest Partition . . . 443
B.5.4 Failing to Drop Online CurrentPartition . . . 444
B.5.5 Database Space Reached Specified Percent Threshold . . . 444
B.5.6 Database Space Reached Specified Time Threshold . . . 444
B.5.7 Database Space Very Low . . . 445
B.5.8 Error inserting events . . . 445
B.5.9 Error Moving Completed File . . . 445
B.5.10 Error Processing Event Message . . . 446
B.5.11 Error Saving Failed Events . . . 446
B.5.12 Event Insertion Is Blocked . . . 446
B.5.13 Event Insertion Is Resumed . . . 447
B.5.14 Event Message Queue Overflow . . . 447
B.5.15 Event Processing Failed . . . 447
B.5.16 No Space In The Database . . . 448
B.5.17 Opening Archive File Failed . . . 448
B.5.18 Partition Configuration . . . 448
B.5.19 Writing to Archive File failed . . . 449
B.5.20 Writing to the overflow partition (P_MAX) . . . 449
B.6 Database Aggregation . . . 449
B.6.1 Creating Summary . . . 450
B.6.2 Deleting Summary . . . 450

B.6.3 Disabling Summary . . . 450
B.6.4 Enabling Summary . . . 451
B.6.5 Error inserting Summary Data into the Database . . . 451
B.6.6 Saving Summary . . . 451
B.7 Mapping Service . . . 451
B.7.1 Error . . . 452
B.7.2 Error Applying Incremental Update . . . 452
B.7.3 Error initializing map with ID . . . 453
B.7.4 Error Refreshing Map . . . 453
B.7.5 Error Saving Data File . . . 454
B.7.6 Get File Size . . . 454
B.7.7 Loaded Large Map . . . 454
B.7.8 Long Time To Load Map . . . 455
B.7.9 Out Of Sync Detected . . . 455
B.7.10 Refreshing Map from Cache . . . 455
B.7.11 Refreshing Map from Server . . . 456
B.7.12 Save Data File . . . 456
B.7.13 Saved Data File . . . 457
B.7.14 Timed Out Waiting For Callback . . . 457
B.7.15 Timeout Refreshing Map . . . 457
B.7.16 Update . . . 458
B.7.17 Update . . . 458
B.8 Event Router . . . 458
B.8.1 Event Router is Initializing . . . 459
B.8.2 Event Router Is Running . . . 459
B.8.3 Event Router is Stopping . . . 459
B.8.4 Event Router is Terminating . . . 460
B.9 Correlation Engine . . . 460
B.9.1 Correlation Action Definition . . . 461
B.9.2 Correlation Engine Configuration . . . 461
B.9.3 Correlation Engine is Running . . . 461
B.9.4 Correlation Engine is Stopped . . . 462
B.9.5 Correlation Rule . . . 462
B.9.6 Correlation Rule Configuration . . . 462
B.9.7 Deploy Rules With Actions To Engine . . . 463
B.9.8 Disabling Rule . . . 463
B.9.9 Enabling Rule . . . 463
B.9.10 Rename Correlation Engine . . . 464
B.9.11 Rule Deployment is Modified . . . 464
B.9.12 Rule Deployment Is Started . . . 464
B.9.13 Rule Deployment is Stopped . . . 465
B.9.14 Starting Engine . . . 465
B.9.15 Stopping Engine . . . 465
B.9.16 UnDeploy All Rules From Engine . . . 466
B.9.17 UnDeploy Rule . . . 466
B.9.18 Update Correlation Rule Actions . . . 466
B.10 Event Source Management:General . . . 466
B.10.1 Collector Manager Initialized . . . 467
B.10.2 Collector Manager Is Down . . . 468
B.10.3 Collector Manager Started . . . 468
B.10.4 Collector Manager Stopped . . . 468
B.10.5 Collector Service Callback . . . 469
B.10.6 Cyclical Dependency . . . 469
B.10.7 Event Source Manager Callback . . . 469
B.10.8 Initializing Collector Manager . . . 470
B.10.9 Lost Contact With Collector Manager . . . 470
B.10.10 No Data Alert . . . 470
B.10.11 Persistent Process Died . . . 470

B.10.12 Persistent Process Restarted . . . 471
B.10.13 Port Start . . . 471
B.10.14 Port Stop . . . 471
B.10.15 Reestablished Contact With Collector Manager . . . 472
B.10.16 Restart Plugin Deployments . . . 472
B.10.17 Restarting Collector Manager (Cold Restart) . . . 473
B.10.18 Restarting Collector Manager (Warm Restart) . . . 473
B.10.19 Start Event Source Group . . . 473
B.10.20 Start Event Source Manager . . . 474
B.10.21 Starting Collector Manager . . . 474
B.10.22 Stop Event Source Group . . . 474
B.10.23 Stop Event Source Manager . . . 475
B.10.24 Stopping Collector Manager . . . 475
B.11 Event Source Management-Event Sources . . . 475
B.11.1 Start Event Source . . . 475
B.11.2 Stop Event Source . . . 476
B.12 Event Source Management-Collectors . . . 476
B.12.1 Start Collector . . . 476
B.12.2 Stop Collector . . . 476
B.13 Event Source Management-Event Source Servers . . . 477
B.13.1 Start Event Source Server . . . 477
B.13.2 Stop Event Source Server . . . 477
B.13.3 Stop Event Source Server . . . 477
B.14 Event Source Management-Connectors . . . 478
B.14.1 Data Received After Timeout . . . 478
B.14.2 Data Timeout . . . 478
B.14.3 File Rotation . . . 479
B.14.4 Process Auto Restart Error . . . 479
B.14.5 Process Start Error . . . 479
B.14.6 Process Stop . . . 480
B.14.7 WMI Connector Status Message . . . 480
B.15 Active Views . . . 480
B.15.1 Active View Created . . . 481
B.15.2 Active View Joined . . . 481
B.15.3 Active View No Longer Permanent . . . 481
B.15.4 Active View Now Permanent . . . 482
B.15.5 Idle Active View Removed . . . 482
B.15.6 Idle Permanent Active View Removed . . . 482
B.16 Data Objects . . . 483
B.16.1 Activity Definition . . . 483
B.16.2 Configuration . . . 483
B.16.3 Viewing Configuration Store . . . 484
B.16.4 Write Data . . . 484
B.17 Activities . . . 484
B.17.1 Creating an Activity . . . 485
B.17.2 Deleting an Activity . . . 485
B.17.3 Saving an Activity . . . 485
B.18 Incidents and Workflows . . . 485
B.18.1 Add Events to Incident . . . 486
B.18.2 Adding Process Definition . . . 486
B.18.3 Create Incident . . . 487
B.18.4 Creating Group . . . 487
B.18.5 Creating User . . . 487
B.18.6 Delete Incident . . . 488
B.18.7 Deleting Group . . . 488
B.18.8 Deleting Process Definition . . . 488
B.18.9 Deleting User . . . 489
B.18.10 E-Mail Incident . . . 489

B.18.11 Get Incident 489
B.18.12 Save Incident 490
B.18.13 Saving Group 490
B.18.14 Saving Process Definition 490
B.18.15 Viewing Process Definition 491
B.19 General 491
B.19.1 Configuration Service 491
B.19.2 Controlled Process is started 492
B.19.3 Controlled Process Is Stopped 492
B.19.4 Importing Auxiliary 493
B.19.5 Importing Plug-In 493
B.19.6 Load Esec Taxonomy to XML 493
B.19.7 Process Auto Restart Error 493
B.19.8 Process Restarts 494
B.19.9 Proxy Client Registration Service (medium) 494
B.19.10 Restarting Process 495
B.19.11 Restarting Processes 495
B.19.12 Starting Process 495
B.19.13 Starting Processes 496
B.19.14 Stopping Process 496
B.19.15 Stopping Processes 496
B.19.16 Store Esec Taxonomy From XML 497
B.19.17 Watchdog Process is started 497
B.19.18 Watchdog Process Is stopped 497

About This Guide

Novell Sentinel Rapid Deployment is a security information and event management solution that receives information from many sources throughout an enterprise, standardizes it, prioritizes it, and presents it to you to make threat, risk, and policy-related decisions. This guide is divided into the following sections:

- Chapter 1, “Managing Sentinel Rapid Deployment Through the Web Interface,” on page 17
- Chapter 2, “Sentinel Control Center,” on page 41
- Chapter 3, “Active Views Tab,” on page 53
- Chapter 4, “Correlation Tab,” on page 83
- Chapter 5, “Incidents Tab,” on page 109
- Chapter 6, “iTRAC Workflows,” on page 123
- Chapter 7, “Work Items,” on page 161
- Chapter 8, “Analysis Tab,” on page 167
- Chapter 9, “Advisor Usage and Maintenance,” on page 171
- Chapter 10, “Download Manager,” on page 183
- Chapter 11, “Event Source Management,” on page 189
- Chapter 12, “Administration,” on page 235
- Chapter 13, “Sentinel Data Manager,” on page 283
- Chapter 14, “Utilities,” on page 301
- Chapter 15, “Quick Start,” on page 313
- Chapter 16, “Solution Packs,” on page 331
- Chapter 17, “Action Manager and Integrator,” on page 363
- Appendix A, “Sentinel Rapid Deployment Architecture,” on page 401
- Appendix B, “System Events for Sentinel,” on page 431

Audience

This documentation is intended for information security professionals.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation or go to Novell Documentation Feedback (http://www.novell.com/documentation/feedback.html) and enter your comments there.

Additional Documentation

Sentinel technical documentation includes several different volumes:

- *Novell Sentinel Rapid Deployment Installation Guide* (http://www.novell.com/documentation/sentinel61rd/s61rd_install/data/index.html)
- *Novell Sentinel Rapid Deployment Reference Guide* (http://www.novell.com/documentation/sentinel61rd/s61rd_reference/data/index.html)
- *Novell Sentinel Rapid Deployment User Guide* (http://www.novell.com/documentation/sentinel61rd/s61rd_user/data/index.html)
- *Sentinel 6.1 SP2 Install Guide* (http://www.novell.com/documentation/sentinel61/)
- *Sentinel 6.1 SP2 User Guide* (http://www.novell.com/documentation/sentinel61)
- *Sentinel 6.1 SP2 Reference Guide* (http://www.novell.com/documentation/sentinel61)
- *Sentinel SDK* (http://www.novell.com/developer/develop_to_sentinel.html)

  The Sentinel SDK site provides the details about developing Collectors (proprietary or JavaScript) and JavaScript correlation actions.

Contacting Novell

- *Novell Web site* (http://www.novell.com)
- *Novell Technical Support* (http://support.novell.com/phone.html?sourceidint=suplnav4_phonesup)
- *Novell Self Support* (http://support.novell.com/support_options.html?sourceidint=suplnav_supportprog)
- *Patch Download Site* (http://download.novell.com/index.jsp)
- *Novell 24x7 Support* (http://www.novell.com/company/contact.html)
- *Sentinel TIDS* (http://support.novell.com/products/sentinel)
- Sentinel Community Support Forums (http://forums.novell.com/novell-product-support-forums/sentinel/)
- Sentinel Plug-in Web site (http://support.novell.com/products/sentinel/secure/sentinelplugins.html)
- Notification E-mail List: Sign up through the Sentinel Plug-in Web site

Managing Sentinel Rapid Deployment Through the Web Interface

1

This section discusses how to manage the services for Novell Sentinel by using the Sentinel Web interface.

To start Sentinel Rapid Deployment by using the Web interface, your system should have Java 1.6.0_20 or later installed. Also the JRE path should be set to launch the Sentinel applications through Webstart. Set the `JAVA_HOME` environment variable to point to the location of the `JRE 6` folder. Set the export path to point to the `bin` folder under the JRE 6 location.

- Section 1.1, “Accessing the Novell Sentinel Web Interface,” on page 17
- Section 1.2, “Applications and Installers,” on page 17
- Section 1.3, “Reporting,” on page 19
- Section 1.4, “Searching Events,” on page 29

1.1 Accessing the Novell Sentinel Web Interface

Use the Novell Sentinel Web interface to manage, run, schedule, and search reports, launch the Sentinel Control Center (SCC), the Sentinel Data Manager (SDM), and the Solution Designer, and download the Collector Manager installer and the Client installer. You can also perform full-text search on events by using the Web interface.

1. Open a Web browser to the following URL:

   `https://svrname.example.com:8443/sentinel`

   Replace *`svrname.example.com`* with the actual DNS name or IP address (such as `192.168.1.1`) of the server where Sentinel is running.

   IMPORTANT: The URL is case sensitive.

2. If you are prompted to verify the certificates, review the certificate information, then click *Yes* if it is valid.
3. Specify the username and password for the Sentinel account you want to access.
4. Use the *Languages* drop-down list to specify which language you want to use.

   This is typically the same language as the language code of the Sentinel server and your local computer. Make sure to configure your browser's Languages setting to support the desired language.
5. Click *Sign in*.

1.2 Applications and Installers

Click *Applications* in the left panel of the Novell Sentinel Rapid Deployment Web interface to download the Sentinel components.

Figure 1-1 *WebStart*

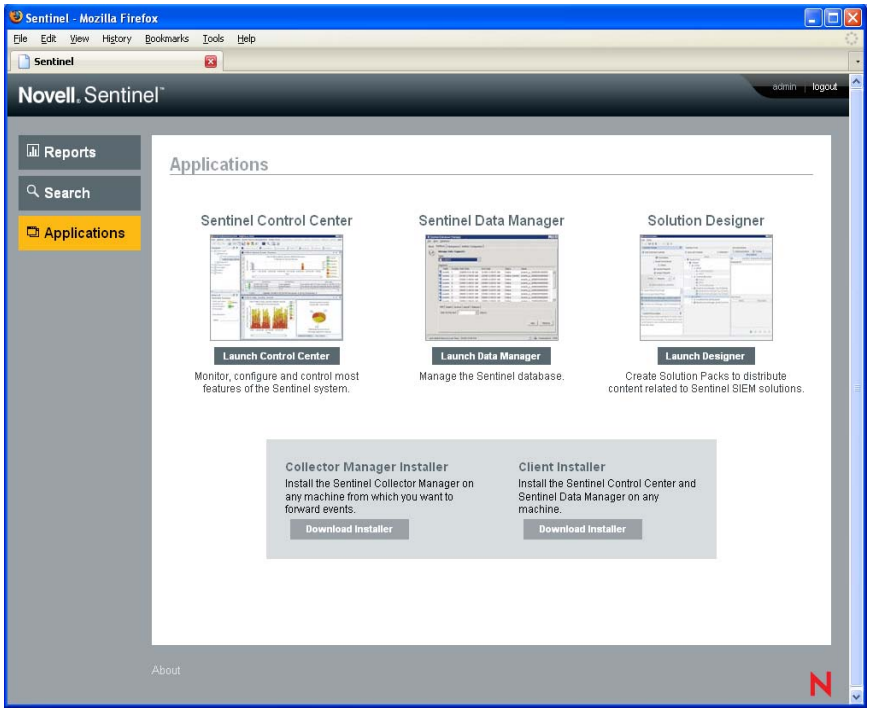

Table 1-1 *Downloading Options*

Options	Description	Action
The Sentinel Control Center (SCC)	The Sentinel Control Center allows you monitor, configure, and control most features of the Sentinel system. The SCC interface helps you manage and monitor the security information received from different network resources. It creates and deploys rules to detect suspicious or malicious patterns of events, provides real-time indication of attacks and related risks, and manages and monitors connections between Sentinel and its event sources.	1. Click *Launch Control Center.* 2. Open SCC with the Java Web Start Launcher. 3. Specify the user credentials and click *Login*.

Options	Description	Action
The Sentinel Data Manager (SDM)	The Sentinel Data Manager allows you manage the Sentinel database. You can monitor database space utilization, view and manage database partitions, configure auto-archives, and configure auto-addition of partitions.	1. Click *Launch Data Manager.* 2. Open SDM with the Java Web Start Launcher. 3. Specify the server, database, host, and port number. 4. Specify the user credentials and click *Connect*.
The Solution Designer		
Collector Manager Installer	The Collector Manager Installer allows you install the Sentinel Collector Manager on any machine where you want to forward the events from.	Click *download Collector Manager installer* and follow the on-screen instructions.
Client Installer	The Client Installer allows you install the Sentinel Control Center, Sentinel Collector Builder, Sentinel Solution Designer, and Sentinel Data Manager on any client machine.	Click *download Client installer* and follow the on-screen instructions.

1.3 Reporting

You can upload, run, view, and delete reports or report definitions by using the Sentinel Rapid Deployment Web interface. You can run a report by using the desired parameters (such as start and end date) as given in the report definition. The report results are saved with a name of your choice. After the report runs, you can retrieve the results and view them as a PDF file.

Reports are organized by category.

- Section 1.3.1, “Running Reports,” on page 19
- Section 1.3.2, “Viewing Reports,” on page 22
- Section 1.3.3, “Scheduling a Report,” on page 24
- Section 1.3.4, “Managing Reports,” on page 25

1.3.1 Running Reports

Sentinel Rapid Deployment is installed with a set of reports organized into several product categories. Reports run asynchronously, so you can continue to do other things in the application while the report is running. You can view the PDF report results after the report finishes running.

Many report definitions include parameters. You are prompted to set them before running the reports. Depending on how the report developer designed the report, the report parameters can be text, numbers, Boolean values, or dates. A parameter might have a default value or a list based on values in the Sentinel RD database.

IMPORTANT: If a report in progress is canceled by using the *Cancel* link, the query on the database is canceled.

Manually Running a Report

1 Click *Reports* to display the available reports.

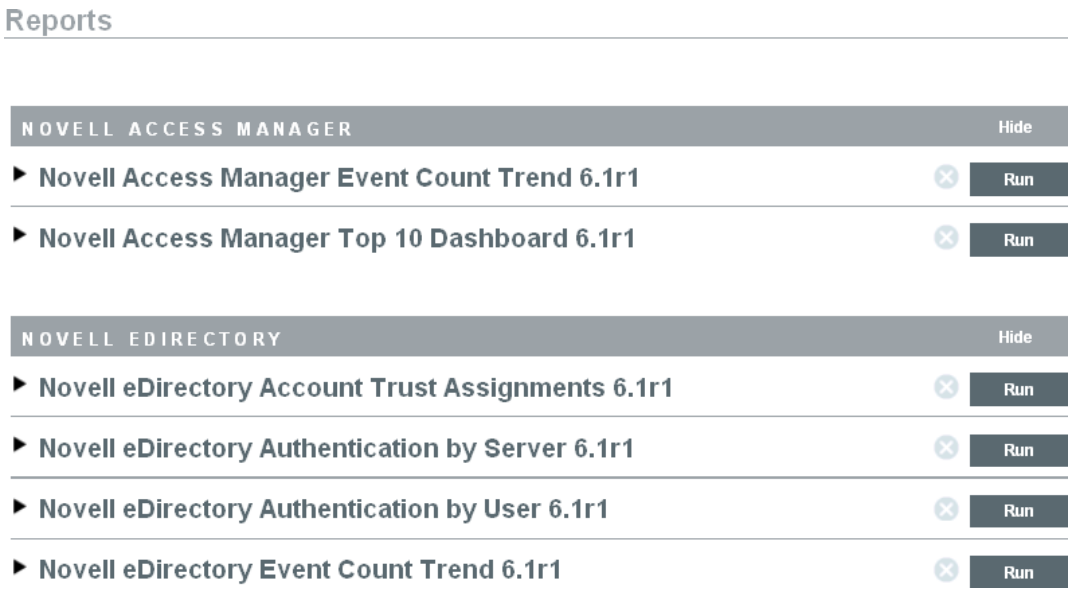

2 If desired, click a report definition to expand it. If you see a *Sample Report* link, you can click *View* to find out how the completed report looks with a set of sample data.

3 Select the report you want to run and click *Run*.

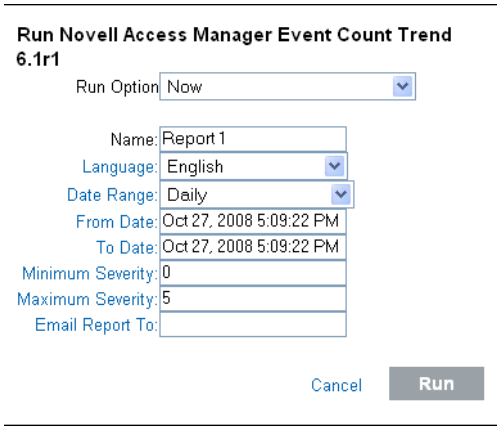

4 Specify the following:

The report parameters are specific to the report definition. Therefore, the report parameters might vary based on the report definition you select.

Report Parameters	Description
Run Option	Set the schedule for running the report. If you want the report to run later, you must also enter a start time. ◆ **Now:** This is the default. It runs the report immediately. ◆ **Once:** Runs the report once at the specified date and time. ◆ **Daily:** Runs the report once a day at the specified time. ◆ **Weekly:** Runs the report once a week on the same day at the specified time. ◆ **Monthly:** Runs the report on the same day of the month every month, starting at the specified date and time. For example, if the start date and time is October 28 at 2:00 p.m, the report will run on the 28th day of the month at 2:00 p.m every month. All time settings are based on the browser's local time. All the *Date* and *Time* fields are stored with a local time stamp and time zone. Sentinel *Date* and *Time* fields use GMT.
Name	Specify a name to identify the report results. Because the username and time are also used to identify the report results, the report name does not need to be unique.
Language	Choose the language in which the report labels and descriptions should be displayed (English, French, German, Italian, Japanese, Traditional Chinese, Simplified Chinese, Spanish, or Portuguese). The data in the report is displayed in whatever language was originally used by the event source.
Date Range	If the report includes time period parameters, choose the date range. You can also set start and end dates for all the time periods. All time periods are based on the local time for the browser. ◆ **Current Day:** Shows events from midnight of the current day until 11:59 p.m of the current day. If the current time is 8 a.m, the report shows 8 hours of data. ◆ **Previous Day:** Shows events from midnight yesterday until 11:59 p.m yesterday. ◆ **Week To Date:** Shows events from midnight Sunday of the current week until the end of the current day. ◆ **Previous Week:** Shows seven days of events, from midnight Sunday of the previous week until 11:59 p.m Saturday of the previous week ◆ **Month to Date:** Shows events from midnight the first day of the current month until the end of the current day. ◆ **Previous Month:** Shows a month of events, from midnight of the first day of the previous month until 11:59 p.m of the last day of the previous month ◆ **Custom Date Range:** For this setting only, you also need to set a start date and end date below.

Report Parameters	Description
From Date and To Date	Set the start date (From Date) and the end date (To Date) for the report.
MinSev	Specify the minimum severity of events to be included in the report. The range is 0-5.
MaxSev	Specify the maximum severity of events to be included in the report. The range is 0-5.
Email Report To	If the report should be mailed to a user or users, specify their e-mail addresses, separated by commas. To enable mailing reports, the administrator must configure the mail relay under *Rules* > *Configuration*.

5 Click *Run*.

A report results entry is created and mailed to the designated recipients.

1.3.2 Viewing Reports

You can view the reports for different applications in the Sentinel Rapid Deployment Web interface for reports. The report GUI by default shows up to 10 report results for any given report definition. The 10 report results displayed are the 10 most recent report results for that report definition.

If there are more than 10 report results for any given report definition (that is, the report has been run more than 10 times), a *Show all x reports* link is displayed after the 10th report, where *x* is the total number of results available for that given report definition.

1 To view the list of report results, click *View*.

All previously run reports are shown with the user-defined report name, the user who ran them, and the time the report was run.

IMPORTANT: The default number of report results to be displayed for each report definition is managed by the `reporting.reportResultsDisplayed` property specified in the `das_core.xml` file.

```
<obj-component id="JasperReportingComponent">
    <class>esecurity.ccs.comp.reporting.jasper.JasperReportingComponent</
class>
    <property name="reporting.reportResultsDisplayed">10</property>
</obj-component>
```

You can change the `reporting.reportResultsDisplayed` property value. After changing this value, ensure that you restart the das_core to apply the changes.

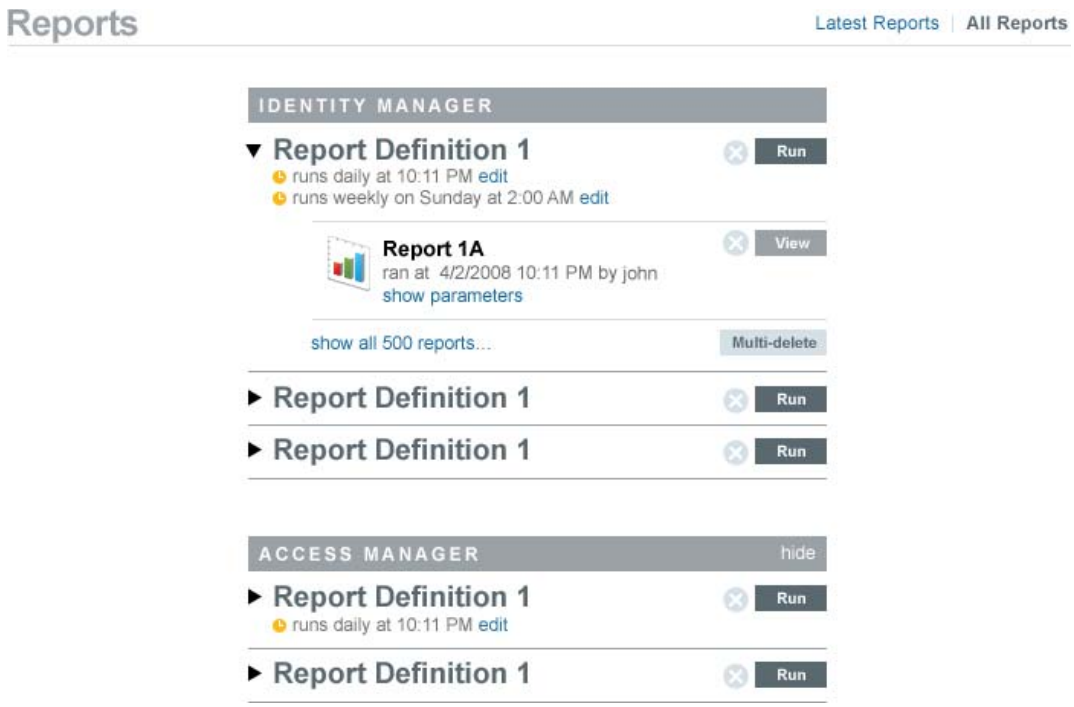

2 Click *show parameters* to see the exact values used to run the report.

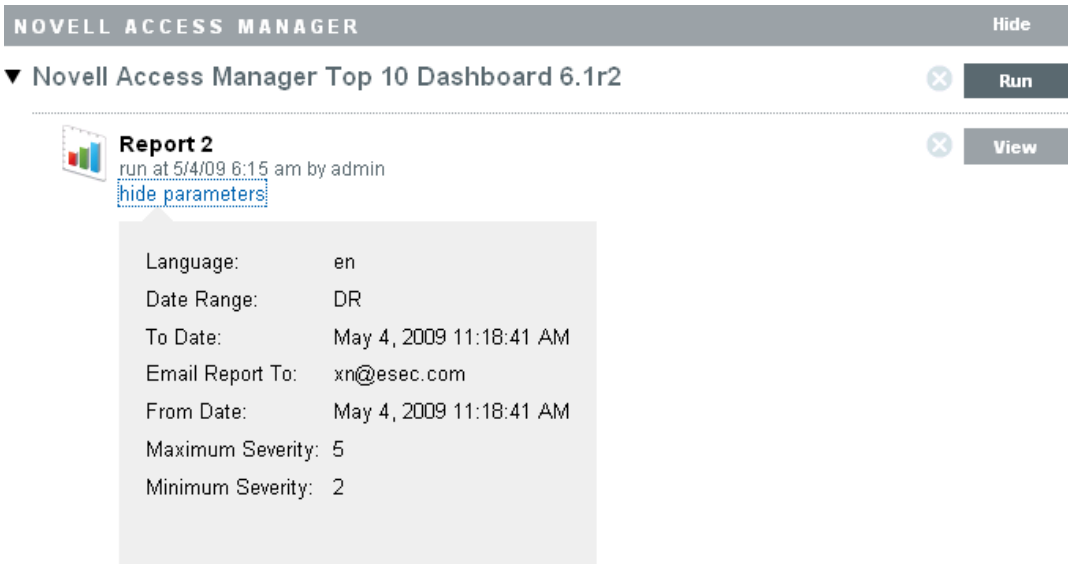

- For Date Range, D=Current Day, PD=Previous Day, W=Week To Date, PW=Previous Week, M=Month To Date, PM=Previous Month, and DR=Custom Date Range.
- For Language, en=English, fr=French, de=German, it=Italian, ja=Japanese, pt=Brazilian Portuguese, es=Spanish, zh=Simplified Chinese, and zh_TW=Traditional Chinese.

3 Click *View* for the report results you want to see. The report results are displayed in a new window in `.pdf` format.

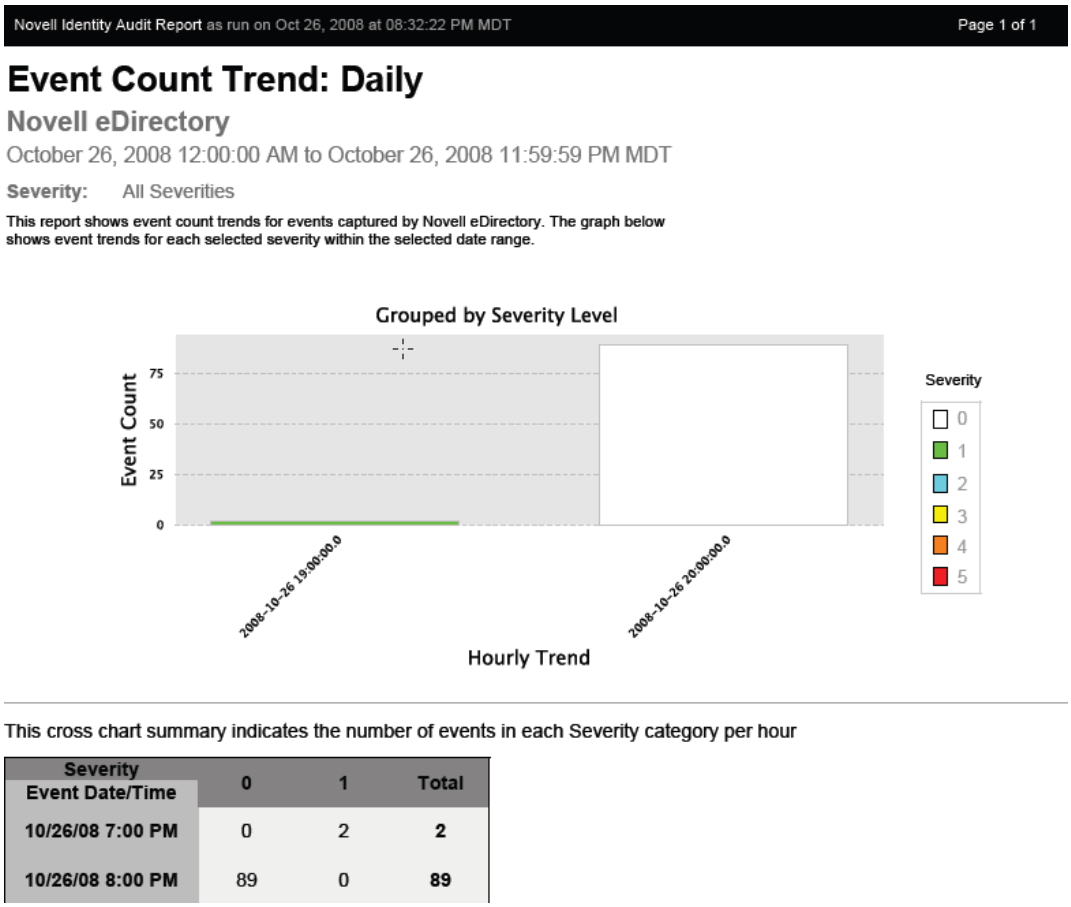

Novell Identity Audit Report as run on Oct 26, 2008 at 08:32:22 PM MDT Page 1 of 1

Event Count Trend: Daily

Novell eDirectory

October 26, 2008 12:00:00 AM to October 26, 2008 11:59:59 PM MDT

Severity: All Severities

This report shows event count trends for events captured by Novell eDirectory. The graph below shows event trends for each selected severity within the selected date range.

This cross chart summary indicates the number of events in each Severity category per hour

Severity / Event Date/Time	0	1	Total
10/26/08 7:00 PM	0	2	**2**
10/26/08 8:00 PM	89	0	**89**

TIP: Report results are organized from newest to oldest.

1.3.3 Scheduling a Report

When you run a report, you can run the report immediately or schedule it to be run later, either once or on a recurring basis. For scheduled reports, you must choose a frequency and enter a time at which the report should run.

- **Now:** This is the default. It runs the report immediately.
- **Once:** Runs the report once at the specified date and time.
- **Daily:** Runs the report once a day at the specified time.
- **Weekly:** Runs the report once a week on the same day at the specified time.
- **Monthly:** Runs the report on the same day of the month every month, starting at the specified date and time. For example, if the start date and time is October 28 at 2:00 p.m, the report runs on the 28th day of the month at 2:00 p.m.every month.

NOTE: All time settings are based on the browser's local time.

Figure 1-2 *Scheduled Reports*

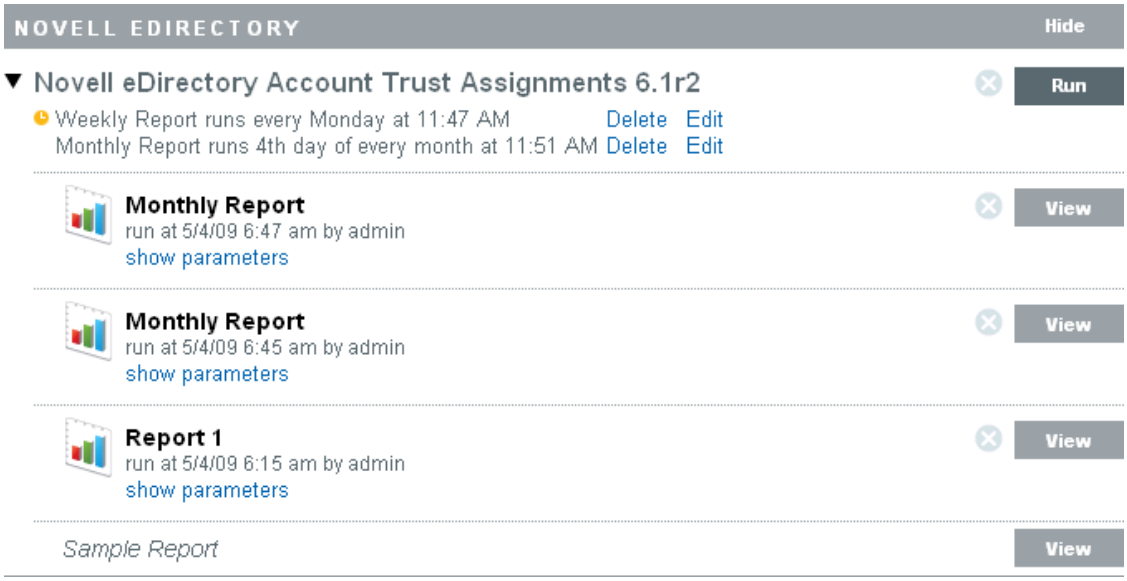

Report schedules can be removed or modified by using the *Delete* and *Edit* links.

1.3.4 Managing Reports

Sentinel Rapid Deployment users can add, delete, update, and schedule reports.

- "Adding Reports" on page 25
- "Creating New Reports" on page 27
- "Renaming Report Results" on page 27
- "Deleting Reports and Report Definitions" on page 27
- "Updating Report Definitions" on page 29

Adding Reports

Any user can add or update reports in Sentinel Rapid Deployment.

- "Downloading New or Updated Reports" on page 25
- "Adding New Reports" on page 25

Downloading New or Updated Reports

New or updated reports by Novell can be downloaded from the Novell Content Web site (http://support.novell.com/products/sentinel/secure/identityaudit.html).

Adding New Reports

Sentinel Rapid Deployment comes preloaded with reports, but new report plug-ins (special `.zip` files that include the report definition plus metadata) can be uploaded into Sentinel Rapid Deployment. If there are no reports in the system, the following screen displays:

Figure 1-3 *No Reports Loaded*

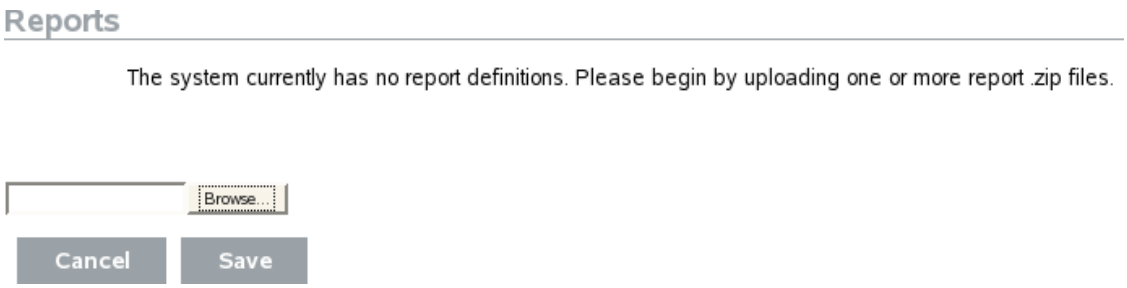

To add a report:

1. Click the *Reports* button on the left side of the screen.
2. Click the *Upload Report* button.
3. Browse and select the report plug-in `.zip` file on your local machine.
4. Click *Open*.
5. Click *Save*.
6. If the same report already exists in the report repository (based on the report's unique ID), decide whether to replace the existing report.

   Sentinel Rapid Deployment displays the details of both the report in the system and the one being imported. In the example below, the imported report is the same version as the existing report.

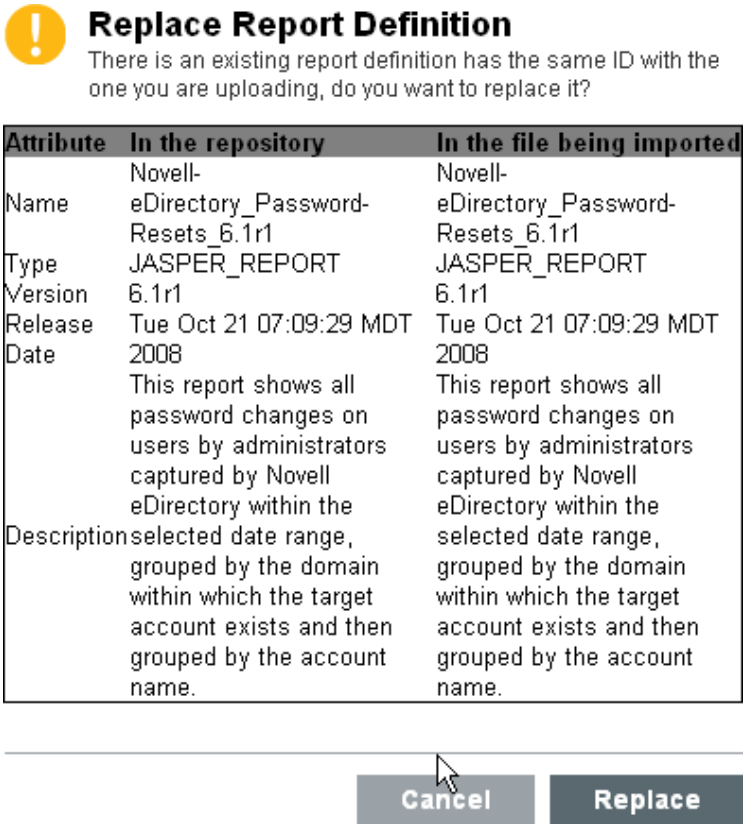

Attribute	In the repository	In the file being imported
Name	Novell-eDirectory_Password-Resets_6.1r1	Novell-eDirectory_Password-Resets_6.1r1
Type	JASPER_REPORT	JASPER_REPORT
Version	6.1r1	6.1r1
Release Date	Tue Oct 21 07:09:29 MDT 2008	Tue Oct 21 07:09:29 MDT 2008
Description	This report shows all password changes on users by administrators captured by Novell eDirectory within the selected date range, grouped by the domain within which the target account exists and then grouped by the account name.	This report shows all password changes on users by administrators captured by Novell eDirectory within the selected date range, grouped by the domain within which the target account exists and then grouped by the account name.

The new report definition is added to the list in alphabetical order and can be run immediately, if desired.

Creating New Reports

Users can modify or write reports by using JasperForge iReport. a graphical report designer for JasperReports. iReport is an open source report development tool that is available for download from JasperForge.org (http://jasperforge.org/plugins/project/project_home.php?group_id=83) (as of the time of this publication).

New or modified reports can include additional database fields that are not presented in the Sentinel Rapid Deployment Web interface. They must adhere to the file and format requirements of the report plug-ins. For more information about database fields and file and format requirements for report plug-ins, see the Sentinel SDK Web site (http://www.novell.com/developer/develop_to_sentinel.html).

Renaming Report Results

Report results (but not report definitions) can be renamed in the interface.

1. Click the *Reports* button on the left side of the screen.
2. Click a report name to expand it.
3. Click the name of the report results you want to rename.
4. Specify the new name.
5. Click *Rename*.

Deleting Reports and Report Definitions

- "Deleting Report Definitions" on page 27
- "Deleting Report Results" on page 27

Deleting Report Definitions

You can delete either a set of report results or a report definition by using the button at the right side of the report definition. If a report definition is deleted, all associated report results are also deleted.

IMPORTANT: Only the users with Manage Reports permissions can delete the report definitions.

For more information on permissions, see "Reporting" in the *Sentinel 6.1 Rapid Deployment Reference Guide*.

Deleting Report Results

There are two ways to delete report results.

- Delete a single report by using the button at the right side of the report result.

IMPORTANT: Users with the Run/View Reports or Manage Reports permission can delete the report results. For more information on permissions, see "Reporting" in the *Sentinel 6.1 Rapid Deployment Reference Guide*.

- Delete multiple report results by using the Multi-delete option at the bottom right side of the report results for each report definition.

NOTE: If the number of report results you have created for a report definition is less than or equal to the default value, you need to use the button to delete each report result.

However, you can change the default value by editing the following property of the `JasperReportingComponent` in the `config/das_core.xml` file:

```
<property name="reporting.reportResultsDisplayed">10</property>
```

After you modify this property value, restart the Sentinel services to apply the changes.

Using the Multi-delete Option

The Multi-delete option is displayed only if:

- You have either Run/View Reports or Manage Reports permissions.
- The number of report results created for a report definition is higher than the default value specified in the Jasper Reporting component.

1. Click the *Multi-delete* option to:
   - Expand the *Multi-delete* panel to list *Select all* and *delete reports* options.
   - Display a check box next to each report result.
2. Select the report results for deletion.

   You can also use the *select all* or *unselect all* options from the *Multi-delete* options panel.
3. Click *delete # reports* to delete the selected report results, where # is the number of report result selected for deletion.

   For example, if you select 3 reports for deletion, a *delete 3 reports* option displays under *Multi-delete* panel. Click *delete 3 reports* to delete all the selected reports. Click *Select all* and select *delete reports* to remove all the reports for a selected report definition.

   Click *cancel* to remove the *Multi-delete* panel and the check boxes for all the report results.

Figure 1-4 *Multi-delete*

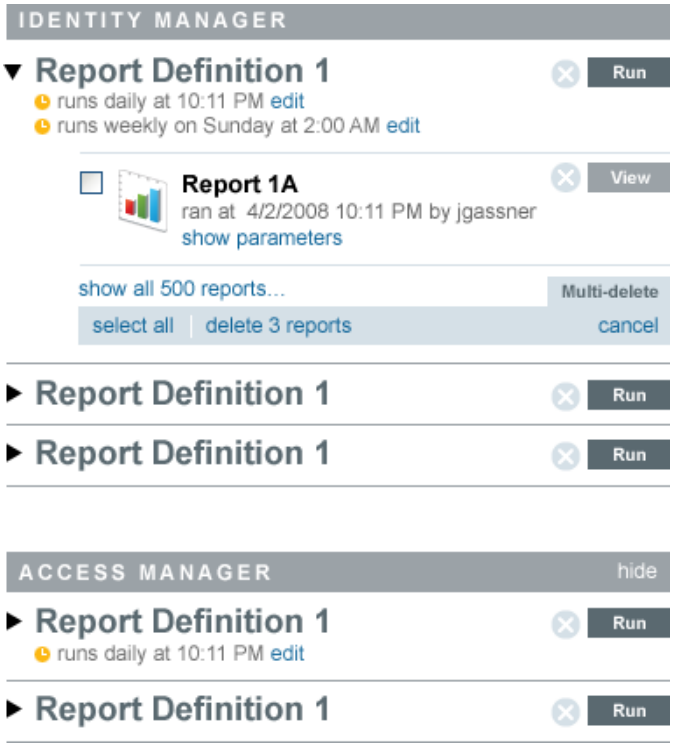

Updating Report Definitions

Users can upload updated reports to replace an existing report. For more information, see "Adding Reports" on page 25.

1.4 Searching Events

Novell Sentinel Rapid Deployment provides the ability to perform a search on events. The search includes all online data currently in the database, but internal events generated by the Sentinel system are excluded unless you select *Include System Events*. By default, events are sorted based on the search engine's relevancy algorithm.

Basic event information includes event name, source, time, severity, information about the initiator (represented by an arrow icon), and information about the target (represented by a bull's-eye icon).

1.4.1 Enabling the Search Option in Web User Interface

To enhance the stability of Sentinel Rapid Deployment, the ability to search events from the Web user interface has been disabled. The preferred methods for searching are in the Sentinel Control Center by using the following options:

- Historical Query (http://www.novell.com/documentation/sentinel61rd/s61rd_user/?page=/documentation/sentinel61rd/s61rd_user/data/investigate_event.html).
- Offline Query (http://www.novell.com/documentation/sentinel61rd/s61rd_user/?page=/documentation/sentinel61rd/s61rd_user/data/offline_query.html).

You can use the following procedure to enable the *Search* option in the Web user interface. However, under load, enabling this option might lead to das_binary crashes and even event loss:

1. Stop the Sentinel services:

```
<install_directory>/bin/sentinel.sh stop
```

2. Open the `das_binary.xml` file for editing.

```
<install_directory>/config/das_binary.xml
```

3. Uncomment the EventSearchComponent section:

```
<!--
<obj-component id="EventSearchComponent">
 <class>esecurity.ccs.comp.textsearch.EventSearchComponent</class>
 <property name="eventsearcher.sortableBatchSize">100000</property>
 <obj-component-ref>
 <name>EventProducer</name>
 <ref-id>EventStoreService</ref-id>
 </obj-component-ref>
 </obj-component>
 -->
```

4. Restart the Sentinel services:

```
<install_directory>/bin/sentinel.sh restart
```

The Search option is now enabled and you can search for events from the Web user interface.

1.4.2 Running an Event Search

You can run simple and advanced searches.

- "Basic Search" on page 30
- "Advanced Search" on page 32

Basic Search

A basic search runs against all of the event fields in Table 1-2 on page 36. Some sample basic searches include the following:

- `root`
- `127.0.0.1`
- `Lock*`
- `driverset0`

NOTE: If time is not synchronized between the end user machine and the Sentinel Rapid Deployment server (for example, one machine is 25 minutes behind), you might get unexpected results from your search. Searches such as *Last 1 hour* or *Last 24 hours* are based on the end user's machine time.

---

1. Click the *Search* link on the left.

Sentinel Rapid Deployment is configured to run a default search for non-system events with severity 3 to 5 the first time you the *Search* link. Otherwise, it defaults to the last search term you entered.

When the Search UI first comes up, by default the UI shows the search results of the last search term you entered (saved in the user preferences).These user preferences are stored in the CONFIGS table in the database as an XML string.This XML string contains the last search term, date range, from time, to time, include system events, sort by time, and results per page. If there is no search term saved in the user's preferences, it defaults to Search Term = sev >=3, Date Range = last 30 days, To Time = Present Time, From Time = Last 30 days time from now, Include system events = false, Sort By time = false results per page = 25.

2 For a different search, type a search term in the search field (for example, `admin`). The search is not case sensitive.

3 Select a time period for which the search should be performed. Most of the time settings are self-explanatory, and the default is *Last 30 Days*.

- *Custom* allows you to select a start date and time and an end date and time for the query. The start date must be before the end date, and the time is based on the browser's local time.
- *All time* searches all the data in the database.

4 Select *Include System Events* to include events that are generated by Sentinel Rapid Deployment system operations.

5 Select *Sort By Time* to arrange data with the most recent events at the beginning.

Sorting by time takes longer than sorting by relevance, which is the default.

6 Click *Search*.

All fields in the index are searched for the specified text. A spinning icon indicates that the search is taking place.

The event summaries are displayed.

Advanced Search

An advanced search can search for a value in a specific event field or fields. The advanced search criteria are based on the short names for each event field and the search logic for the index. To view the field names and descriptions, the short names that are used in advanced searches, and whether the fields are visible in the basic and detailed event views, see Table 1-2 on page 36.

To search for a value in a specific field, use the short name of the field, a colon, and the value. For example, to search for an authentication attempt to Sentinel RD by user2, use the following text in the search field:

- `evt:authentication AND sun:user2`

Other advanced searches might include:

- `pn:NMAS AND sev:5`
- `sip:123.45.67.89 AND evt:“Set Password”`

Figure 1-5 *Advanced Search Example*

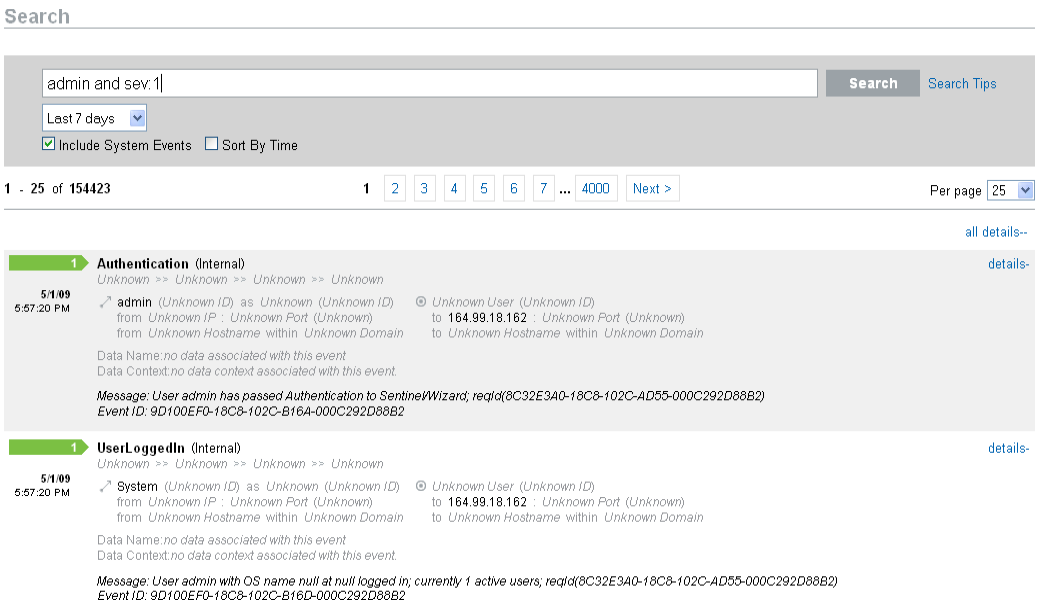

Multiple advanced search criteria can be combined by using the following Boolean operators:

- AND (must be capitalized)
- OR (must be capitalized)
- NOT (must be capitalized and cannot be used as the only search criterion)
- +
- -

Special characters must be escaped by using a \ symbol:

```
+ - && || ! ( ) { } [ ] ^ " ~ * ? : \
```

The advanced search criteria are modeled on the search criteria for the Apache Lucene open source package. More detail about the search criteria is available on the Web: Lucene Query Parser Syntax (http://lucene.apache.org/java/2_3_2/queryparsersyntax.html).

1.4.3 Viewing Search Results

Searches return a set of events. Users can view basic or detailed event information and configure the number of results per page. Search results are returned in batches. The default batch size is 25 results, but this is easily configured.

When results are sorted by relevance, only the top 100,000 events can be viewed. When they are sorted by time, this limitation does not exist.

- "Basic Event View" on page 33
- "Event View with Details" on page 34
- "Refining Search Results" on page 34

Basic Event View

The information in each event is grouped into initiator information and target information. If data isn't available for a particular event field, the fields are labeled *Unknown*.

Figure 1-6 *Basic Event View*

Occasionally, the search engine might index events faster than they are inserted into the database. If you run a search that returns events that have not been inserted into the database, you get a message that some events match the search query but could not be found in the database. If you run the search again later, the events are usually in the database and the search is successful.

Figure 1-7 Events Indexed but Not Yet in Database

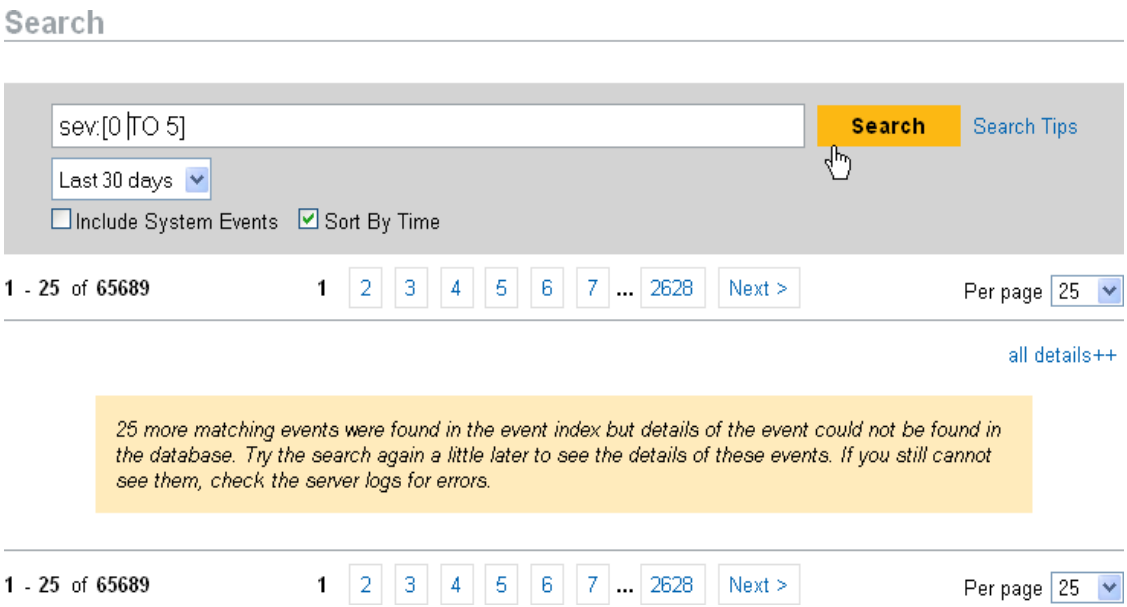

Event View with Details

You can view additional details about any event or events by clicking the *details* link on the right side of the page.The details for all events on a page can be expanded or collapsed by using the *all details* ++ or *details*-- link. This preference is retained as you scan through multiple pages of results or execute new searches.

Figure 1-8 Event View

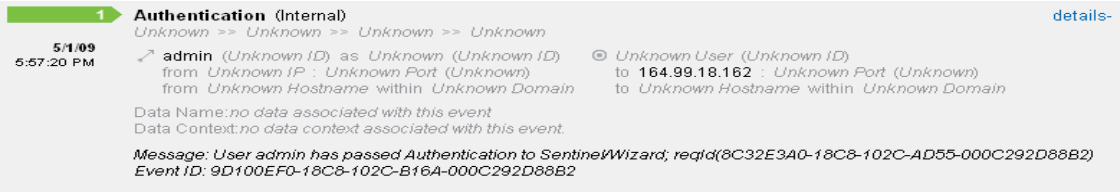

The event in Figure 1-8 shows the same event as in Figure 1-6 on page 33, but with an expanded view that shows additional data fields that might have been populated.

Refining Search Results

After viewing the results of a search, it might be necessary to refine the search results and add additional search criteria. For example, you might see one initiator user's name appear several times in the search results and want to see more events from that initiator.

To filter the search results using a specific value appearing in the search results:

1. Identify the desired filter criteria in the search results.
2. Click the value (for example, target hostname test 1900) by which you want to filter the results.

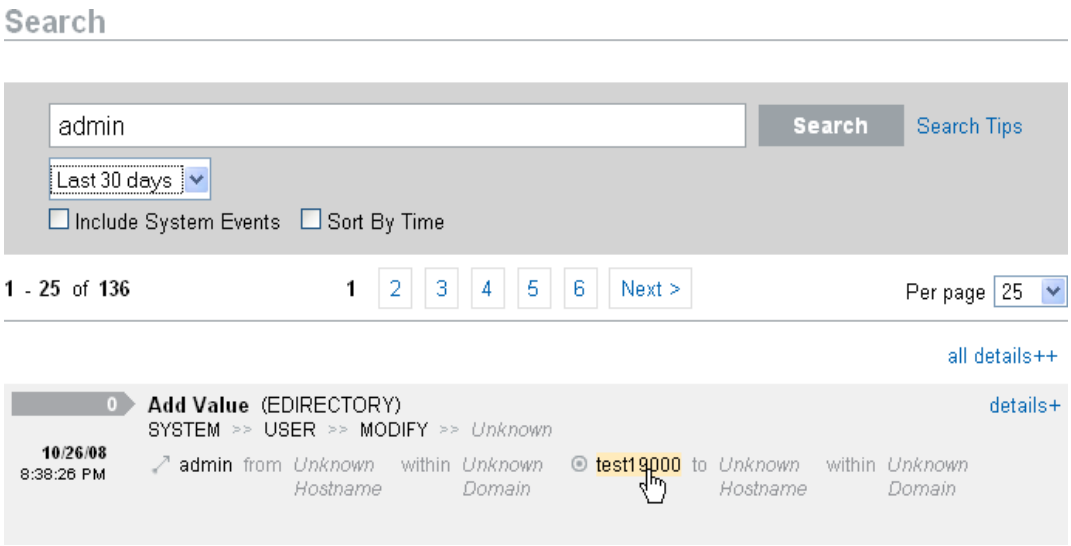

TIP: This adds the value to your filter with an AND operator. To add the value to your filter with an NOT operator, press the Alt key as you click the value.

3 Click *Search*.

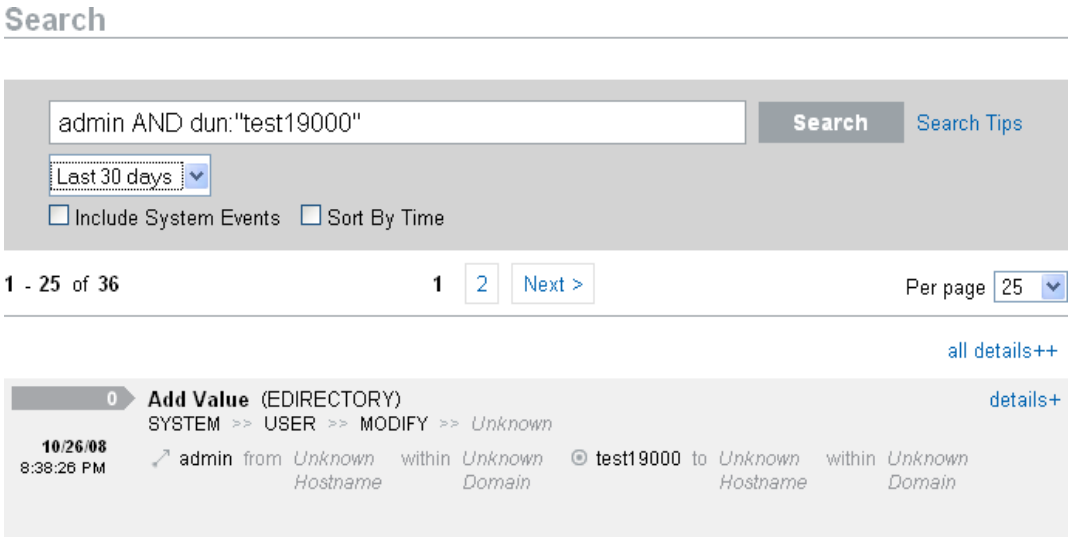

Some fields cannot be selected to refine a search this way:

- EventTime
- Message
- Any field related to the Reporter
- Any field related to the Observer
- Any field related to TargetTrust
- Any field with a value `Unknown`

1.4.4 Event Fields

Each event has fields that might or might not be populated, depending on the specific event. The values for these event fields can be viewed by using a search or running a report. Each field has a short name that is used in advanced searches. The values for most of these fields are visible in the detailed event view; other values are also visible in the basic event view.

Table 1-2 *Event Fields*

Field	Short Name	Description	Visible in Basic View	Visible in Detailed View
Severity	sev	Normalized severity of the event on a scale of 0 (informational) to 5 (critical).	X	X
EventTime	dt	Time stamp of the event. Can be the Sentinel Rapid Deployment server time stamp or the time stamp from the original event source (if *trust event time* is enabled).	X	X
EventName	evt	Short name of the event.	X	X
Message	msg	Detailed event message.	Invisible	X
ProductName	pn	Product that generated the event; the event source. Displayed after the event name.	X	X
InitUserName	sun	Username of the user who initiated the event.	X	X
InitUserID	iuid	User ID of the user who initiated the event, based on the raw data reported by the device.	Invisible	X
InitUserDomain	rv35	Domain of the user who initiated the event. Searchable but not displayed in either event view.	Invisible	Invisible
InitHostName	shn	Hostname of the machine from which the event initiated.	X	X
InitHostDomain	rv42	Domain of the machine from which the event initiated.	X	X
InitIP	sip	IP address of the machine from which the event initiated.	Invisible	X
InitServicePort	spint	Port number from which the event initiated (for example, HTTP)	Invisible	X
InitServicePortName	sp	Type of port from which the event initiated (for example, HTTP).	Invisible	X
TargetUserName	dun	Username of the user who was the target of the event.	X	X

Field	Short Name	Description	Visible in Basic View	Visible in Detailed View
TargetUserID	tuid	User ID of the user who was the target of the event, based on the raw data reported by the device.	Invisible	X
TargetUserDomain	rv45	Domain of the user who was the target of the event. Searchable but not displayed in either event view.	Invisible	X
TargetHostName	dhn	Hostname of the machine that was the target of the event.	X	X
TargetHostDomain	rv41	Domain of the machine that was the target of the event.	X	X
TargetIP	dip	IP address of the machine that was the target of the event.	Invisible	X
TargetServicePort	dpint	Port number that was the target of the event (for example, 80).	Invisible	X
TargetServicePortName	dp	Type of port that was the target of the event (for example, HTTP).	Invisible	X
TargetTrustName	ttn	Role of the user that was a target of the event (for example, FinanceAdmin). Searchable but not displayed in either event view.	Invisible	Invisible
TargetTrustID	ttid	Numerical ID representing the role of the user that was a target of the event. Searchable but not displayed in either event view.	Invisible	Invisible
TargetTrustDomain	ttd	Domain (namespace) within which the target trust exists. Searchable but not displayed in either event view.	Invisible	Invisible
EffectiveUserName	euname	Name of the user that the InitUser is impersonating (`root` using `su`, for example); follows *Initiator Username (Initiator User ID) as* in the detailed event view.	Invisible	X
EffectiveUserID	euid	Numerical ID of the user that the InitUser is impersonating (`root` using `su`, for example), based on the raw data reported by the device.	Invisible	X

Field	Short Name	Description	Visible in Basic View	Visible in Detailed View
ObserverHostName	sn	Hostname of the machine that forwarded the event to the security information event management system (for example, the hostname of a syslog server). Searchable but not displayed in either event view.	Invisible	Invisible
ObserverHostDomain	obsdom	Domain of the machine that forwarded the event to the security information event management system (for example, the domain of a syslog server). Searchable but not displayed in either event view.	Invisible	Invisible
ObserverIP	obsip	IP address of the machine that forwarded the event to the security information event management system (for example, the IP address of a syslog server). Searchable but not displayed in either event view.	Invisible	Invisible
ReporterHostName	rn	Hostname of the machine that reported the event to an observer. Searchable but not displayed in either event view.	Invisible	Invisible
ReporterHostDomain	repdom	Domain of the machine that reported the event to an observer. Searchable but not displayed in either event view.	Invisible	Invisible
ReporterIP	repip	IP address of the machine that reported the event to an observer. Searchable but not displayed in either event view.	Invisible	Invisible
SensorType	st	The single character designator for the sensor type (N=network, H=host, O=operating system, A and I=Sentinel Rapid Deployment auditing events, P=Sentinel Rapid Deployment performance events). Searchable but not displayed in either event view.	Invisible	Invisible
DataName/Filename	fn	Data object name reported in the event (for example, the file name or database table name).	Invisible	X

Field	Short Name	Description	Visible in Basic View	Visible in Detailed View
DataContext	rv36	Container for the FileName data object (for example, a directory for a file or a database instance for a database table)		X
TaxonomyLevel1	rv50	Target classification for event. Displayed under the event name in the format: TaxonomyLevel1>> TaxonomyLevel2>> TaxonomyLevel3>> TaxonomyLevel4	X	X
TaxonomyLevel2	rv51	Subtarget classification for the event. Displayed under the event name in the format: TaxonomyLevel1>> TaxonomyLevel2>> TaxonomyLevel3>> TaxonomyLevel4	X	X
TaxonomyLevel3	rv52	Action information for the event. Displayed under the event name in the format: TaxonomyLevel1>> TaxonomyLevel2>> TaxonomyLevel3>> TaxonomyLevel4	X	X
TaxonomyLevel4	rv53	Detail information for the event. Displayed under the event name in the format: TaxonomyLevel1>> TaxonomyLevel2>> TaxonomyLevel3>> TaxonomyLevel4	X	X

Some fields are tokenized. Tokenizing the fields makes it possible to search for an individual word in the field without a wildcard. The fields are tokenized based on spaces and other special characters. For these fields, articles such as “a” or “the” are removed from the search index.

- EventName
- Message
- ProductName
- FileName
- DataContext
- TaxonomyLevel1
- TaxonomyLevel2
- TaxonomyLevel3
- TaxonomyLevel4

Sentinel Control Center

2

Novell Sentinel gathers and correlates security and non-security information from across an organization's networked infrastructure, as well as third-party systems, devices, and applications. Sentinel presents the collected data in an richly functional interface, identifies security or compliance issues, and tracks remediation activities, streamlining previously error-prone processes and building a more rigorous and secure management program. The Sentinel Control Center (SCC) is the main user interface for viewing and interacting with this data.

- Section 2.1, “Log In to the Sentinel Control Center,” on page 41
- Section 2.2, “About Sentinel Control Center,” on page 42
- Section 2.3, “Introduction to the User Interface,” on page 45

2.1 Log In to the Sentinel Control Center

- Section 2.1.1, “Linux,” on page 41
- Section 2.1.2, “Windows,” on page 41

2.1.1 Linux

1. As the Sentinel Administrator (admin), change directory to:

   `<install_directory>/bin`

2. Run the following command:

   `./control_center.sh`

3. Specify your username and password, then click *OK*.

   A Certificate window displays.

4. Select *Accept*, if you want this message to display every time you start Sentinel on your system. To avoid this, you can select *Accept Permanently*.

2.1.2 Windows

1. Perform either of the following:
   - Go to *Start* > *Programs* > *Sentinel* and select Sentinel Control Center. The Sentinel Login window displays.
   - Click *Applications* in the left panel of the Novell Sentinel Rapid Deployment Web interface, then click *Launch Control Center*:

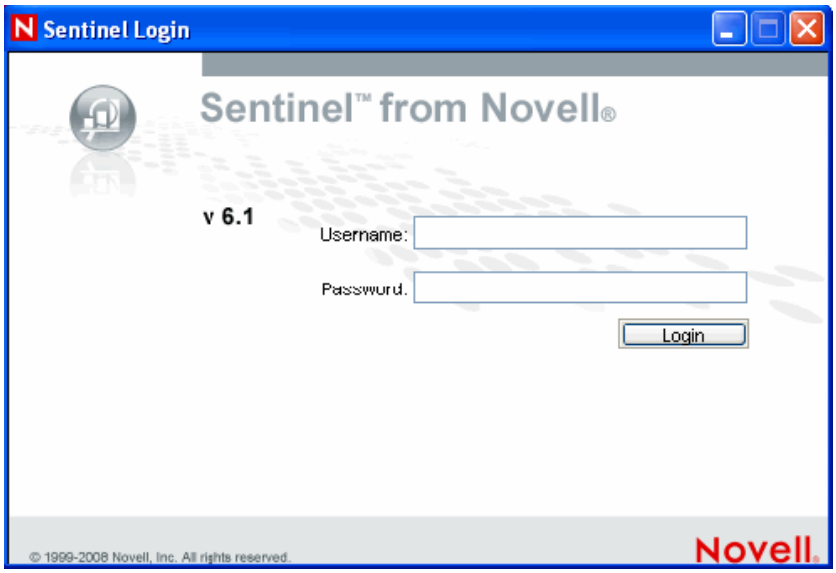

2 Specify your username and password.

3 Click *Login*.

On the first login, the following warning message displays. You must accept the certificate in order to securely log in to the Sentinel Control Center

4 Select *Accept*, if you want this message to display every time you start Sentinel on your system. To avoid this, you can select *Accept Permanently*.

2.2 About Sentinel Control Center

The Sentinel Control Center includes the following functional tabs and interfaces:

- Section 2.2.1, “Active Views,” on page 43
- Section 2.2.2, “Incidents,” on page 43
- Section 2.2.3, “iTRAC,” on page 43

- Section 2.2.4, “Analysis,” on page 44
- Section 2.2.5, “Advisor,” on page 44
- Section 2.2.6, “Admin,” on page 44
- Section 2.2.7, “Correlation,” on page 44
- Section 2.2.8, “Event Source Management,” on page 45
- Section 2.2.9, “Solution Packs,” on page 45
- Section 2.2.10, “Identity Integration,” on page 45

2.2.1 Active Views

The *Active Views* tab presents events in near-real time.

In the *Active Views* tab, you can:

- View events occurring in near-real time
- Investigate events
- Graph events
- Perform historical queries to collect data for a specified period
- Invoke right-click functions
- Initiate manual incidents and remediation workflows

2.2.2 Incidents

An incident is a set of events that require attention (for example, a possible attack). Incidents centralize the data and are typically made up of a correlated event, the associated events that triggered a correlation rule, asset details of the affected systems, vulnerability state of the affected systems, and any remediation information, if known. Incidents can be associated with a remediation workflow in iTRAC, if specified. An incident associated to an iTRAC workflow allows users to track the remediation state of the incident.

In the *Incidents* tab, you can:

- Manage incident views
- View and manage incidents and their associated data
- Switch between existing incident views

2.2.3 iTRAC

The iTRAC stateful incident remediation workflow capability allows you to incorporate your organization’s incident response processes into Sentinel.

In the *iTRAC* tab, you can:

- Create custom workflow templates
- Edit workflow templates
- Create custom activities
- Edit activities

- Associate activities with workflow steps
- Initiate and execute processes

2.2.4 Analysis

The *Analysis* tab is used to run and save an offline query for later quick retrieval of search results.

2.2.5 Advisor

Advisor is an optional module that provides real-time correlation between detected intrusion detection system attacks and vulnerability scan output in order to immediately indicate increased risk to an organization.

2.2.6 Admin

The *Admin* tab provides you access to perform the administrative actions and configuration settings in Sentinel. In the *Admin* tab, you can:

- Create and modify filters
- Use filters to format data
- Use filters to determine event routing
- View system statistics about the Data Access Service
- Start and stop system components
- Configure Sentinel event fields
- Configure the mapping service
- Create new options for right-click event menus
- Aggregate data for reporting
- Create users and assign them to roles for workflows
- Manage user sessions

2.2.7 Correlation

The *Correlation* tab provides an interface to create and deploy rules to detect suspicious or malicious patterns of events.

In the *Correlation* tab, you can:

- Create and edit rules
- Deploy/undeploy rules
- Add an action and associate it to a rule
- Configure dynamic lists

2.2.8 Event Source Management

The Event Source Management (ESM) interface is available through the Sentinel Control Center menu. It allows you to manage and monitor connections between Sentinel and its event sources by using Sentinel Connectors and Sentinel Collectors.

In the ESM, you can:

- Import/export Connectors and Collectors from and to the centralized repository available in ESM
- Add/edit connections to event sources through the configuration wizards
- View the real-time status of the connections to event sources
- Monitor data flowing through the Collectors and Connectors

Sentinel Collectors

The Collectors parse the data and deliver a richer event stream by injecting taxonomy, exploit detection, and business relevance into the data stream before events are correlated and analyzed and sent to the database.

Sentinel Connectors

The Connectors use industry standard methods to connect to the data source to get raw data.

2.2.9 Solution Packs

You can use the Solution Packs interface through the *Tools* menu in the Sentinel Control Center. Solution Packs provide a framework within which sets of content can be packaged into controls, each of which is designed to enforce a specific business or technical policy.

2.2.10 Identity Integration

The Sentinel integration framework for identity management systems provides functionality on several levels. When identity integration is implemented, you can:

- Look up the following information about a user from the Identity Browser:
  - Contact information
  - Accounts associated with that user
  - Most recent authentication events
  - Most recent access events
  - Most recent permissions changes
- Look up user information by right-clicking an event

2.3 Introduction to the User Interface

In the Sentinel Control Center user interface, you can perform the activities through the following components:

- Section 2.3.1, “Menu Bar,” on page 46

- Section 2.3.2, “Toolbar,” on page 47
- Section 2.3.3, “Tabs,” on page 48
- Section 2.3.4, “Frames,” on page 48

Sentinel Control Center provides you the “dockable” framework, which allows you to move the toolbars, tabs or frames from their default location to user-specific locations for ease of use.

Figure 2-1 *Sentinel Control Center*

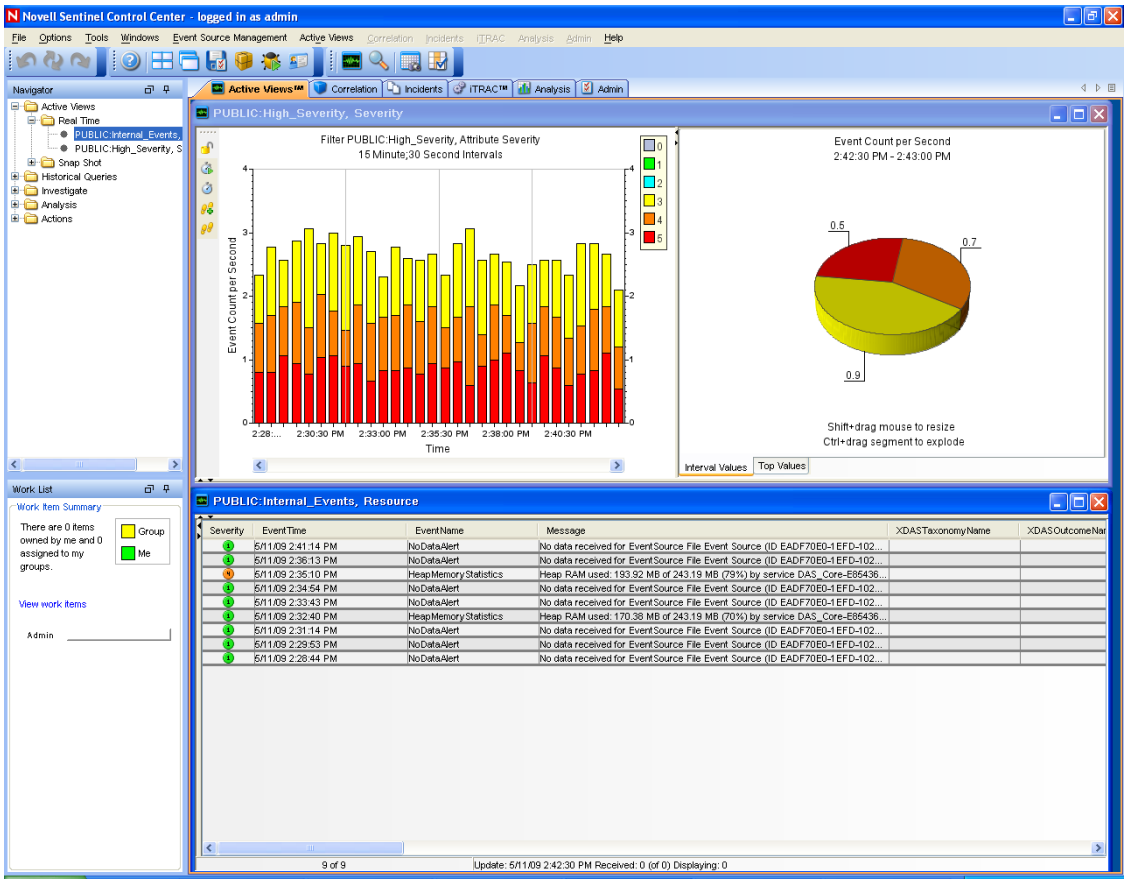

2.3.1 Menu Bar

The menu bar has the menus required to navigate, perform activities, and change the appearance of the Sentinel Control Center.

Figure 2-2 *Menu Bar*

Figure 2-3 *Menu Bar*

The *File*, *Options*, *Event Source Management*, *Windows*, and *Help* menus are always available. The availability of other menus depends on your location in the console and the permissions you have.

2.3.2 Toolbar

The toolbar allows you to perform tab-specific functions. There are four system-wide toolbar buttons that are always displayed: *View Sentinel Help*, *Cascade All Display Windows*, *Tile All Display Windows*, and *Save User Preferences*. The availability of other toolbar buttons depends on your location in the console and the permissions you have.

- "System-Wide Toolbar" on page 47
- "Tab-Specific Toolbar Buttons" on page 47

System-Wide Toolbar

The system-wide toolbar buttons are:

Figure 2-4 *Toolbar Buttons*

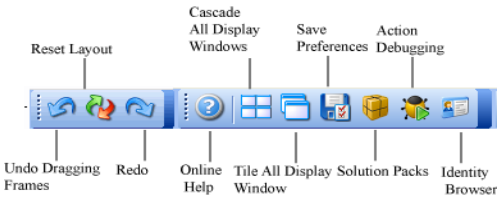

Figure 2-5 *Additional Toolbar Buttons*

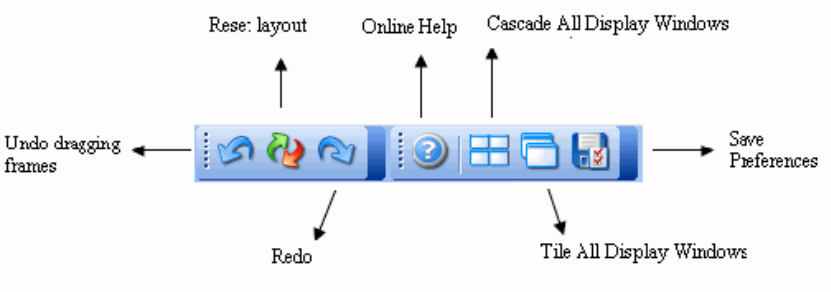

Tab-Specific Toolbar Buttons

Tab-specific toolbar buttons allows you to perform the functions related to each tab.

Table 2-1 *Tab-Specific Toolbar Buttons*

Toolbar	View
Active Views	

Toolbar	View
Correlation	
Incidents	
iTRAC	
Analysis	
Admin	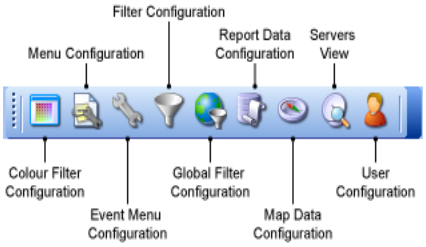

For more information on tab-specific toolbar buttons, see the sections on each of the tabs listed in Section 2.3.3, “Tabs,” on page 48.

2.3.3 Tabs

Depending on your access permissions, Sentinel Control Center displays the following tabs.

- *Active Views* tab. For more information, see Chapter 3, “Active Views Tab,” on page 53
- *Correlation* tab. For more information, see Chapter 4, “Correlation Tab,” on page 83
- *Incidents* tab. For more information, see Chapter 5, “Incidents Tab,” on page 109
- *iTRAC* tab. For more information, see Chapter 6, “iTRAC Workflows,” on page 123
- *Analysis* tab. For more information, see Chapter 8, “Analysis Tab,” on page 167
- *Admin* tab. For more information, see Chapter 12, “Administration,” on page 235

2.3.4 Frames

Sentinel provides a dockable framework that allows you to drag frames on the screen to place them in your preferred locations. The following buttons allow you to drag and/ hide frames.

- Toggle Floating
- Toggle Auto-hide

Figure 2-6 *Navigator Frame*

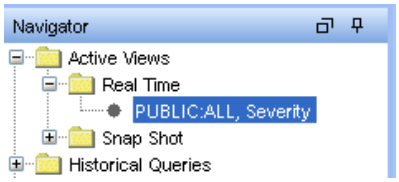

To drag a frame to any location:

1. Click the *Toggle Floating* icon on the frame or hold the frame and drag it to the desired location.

To hide a frame:

1. Click the *Toggle Auto-hide* icon.

NOTE: You can undo dragging or reset the framework to the default position by using the toolbar buttons.

2.3.5 Using the Sentinel Control Center to Navigate

To navigate by using the toolbar:

1. Click the tab you need to use.
2. Click toolbar buttons to perform the actions.

To navigate by using the menu bar:

1. Click the tab menu in the menu bar.
2. Select an action you need to perform.

NOTE: This procedure is generic for all the tabs in the Sentinel Control Center. Navigation procedures for tabs are discussed in the relevant sections.

2.3.6 Changing the Appearance of the Sentinel Control Center

You can change the Sentinel Control Center's look by:

- "Setting the Tab Position" on page 50
- "Cascading Windows" on page 50
- "Tiling Windows" on page 50
- "Minimizing Windows" on page 50
- "Restoring Windows to Original Size" on page 50
- "Closing all Open Windows" on page 50

Setting the Tab Position

1. Click *Options* > *Tab Placement*.
2. Select either *Top* or `Bottom`.

Cascading Windows

1. Click *Windows* > *Cascade All*. All open windows in the right panel cascade.

Tiling Windows

1. Click *Windows* > *Tile All*.
2. Select from the following options:
   - *Tile Best Fit*
   - *Tile Vertical*
   - *Tile Horizontal*

Minimizing Windows

1. Click *Windows* > *Minimize All*. All open windows in the right panel minimize.

Restoring Windows to Original Size

1. Click *Windows* > *Restore All*. All open windows in the right panel are restored to their original size.

NOTE: Use the Minimize and Restore options provided on the top right corner of the tab to minimize individual tabs.

Closing all Open Windows

1. Click *Windows* > *Close All*.

2.3.7 Saving User Preferences

If the user has permissions to save the workspace, they can save the following preferences:

- Permanent windows that are not dependent on data that was available at the time of their original creation.
- Active Views
- Summary displays
- Window positions
- Window sizes, including the application window
- Tab positions
- Navigator docked or floating and showing or hidden

The following preferences are not saved when the user logs out:

- Snapshots

- Historical event queries
- Secondary windows opened from one of the primary windows in the Admin Navigator
- Column widths in Active Views

To save your preferences:

1 Click *File* > *Save Preferences* or click

2.3.8 Changing Password

1 Click *Options* > *Change Password*.
2 Provide the old password.
3 Provide the new password and confirm it.
4 Click *OK*.

For more information on password security, see the *Sentinel Rapid Deployment Reference Guide*.

2.3.9 Configuring the Attachment Viewer

1 On the Tools menu, click *Attachment Viewer Configuration* or alternatively click the Configure Attachment Viewers button. The Attachment Viewer Configuration window displays.

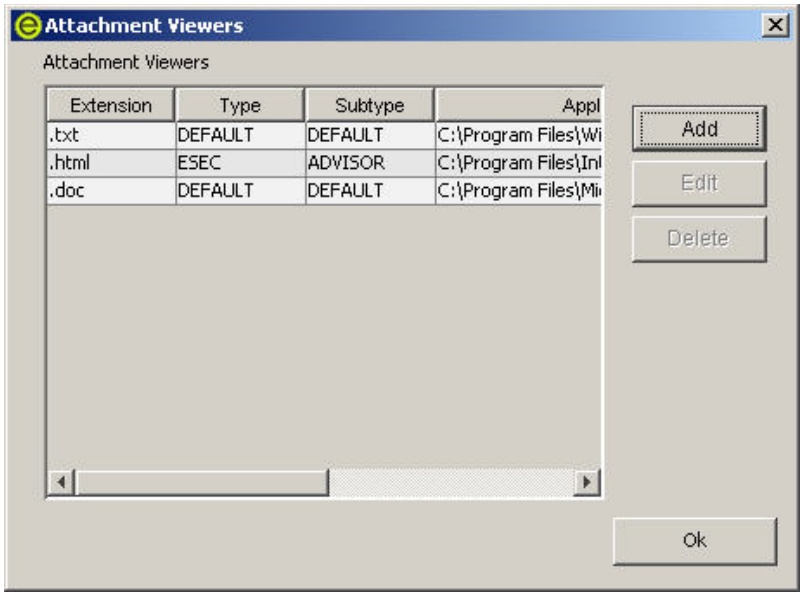

2 Click *Add*. The Attachment Identification window displays.

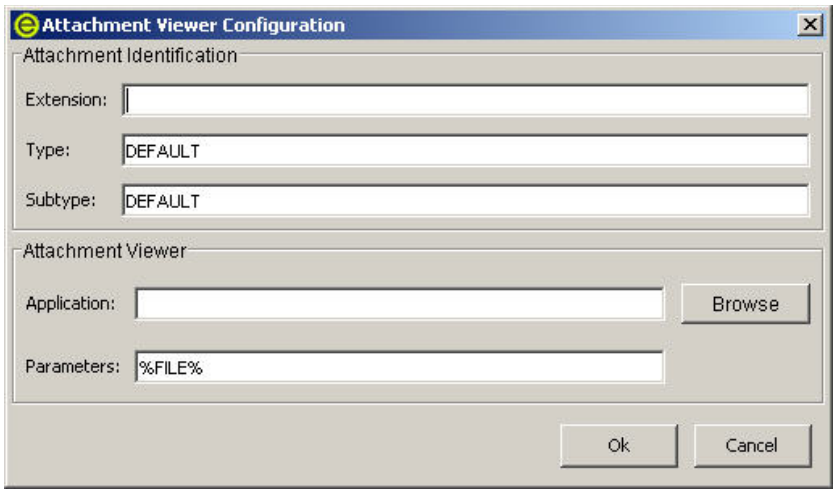

Specify the extension type (such as `.doc`, `.xls`, `.txt`, `.html` and so on) and click *Browse* or type in the application program to launch the file type (such as `notepad.exe` for Notepad).

3 Click *OK*.

Active Views Tab

3

The *Active Views* tab presents events in near-real time.

- Section 3.1, “Understanding Active Views,” on page 53
- Section 3.2, “Introduction to the User Interface,” on page 54
- Section 3.3, “Reconfiguring Total Display Time,” on page 57
- Section 3.4, “Viewing Real-Time Events,” on page 57
- Section 3.5, “Showing and Hiding Event Details,” on page 61
- Section 3.6, “Sending Mail Messages about Events and Incidents,” on page 62
- Section 3.7, “Creating Incidents,” on page 63
- Section 3.8, “Viewing Events That Trigger Correlated Events,” on page 64
- Section 3.9, “Investigating an Event or Events,” on page 65
- Section 3.10, “Viewing the Advisor Data,” on page 70
- Section 3.11, “Viewing the Asset Data,” on page 71
- Section 3.12, “Viewing Vulnerabilities,” on page 73
- Section 3.13, “Ticketing System Integration,” on page 77
- Section 3.14, “Viewing User Information,” on page 77
- Section 3.15, “Using Custom Menu Options with Events,” on page 77
- Section 3.16, “Managing Columns in a Snapshot or Navigator Window,” on page 78
- Section 3.17, “Taking a Snapshot of a Navigator Window,” on page 79
- Section 3.18, “Sorting Columns in a Snapshot,” on page 79
- Section 3.19, “Closing a Snapshot or Navigator,” on page 79
- Section 3.20, “Adding Events to an Incident,” on page 80

3.1 Understanding Active Views

In the *Active Views* tab, you can:

- View events occurring in near-real time
- Investigate events
- Graph events
- Perform historical statistical analysis
- Invoke right-click functions
- Initiate manual incidents and remediation workflows

An event represents a normalized log record reported to Sentinel from a third-party security, network, or application device or from an internal Sentinel source. There are several types of events:

- External events (event received from a security device), such as:
  - An attack detected by an intrusion detection system

    - A successful login reported by an operating system
    - A customer-defined situation such as a user accessing a file
- Internal events (an event generated by Sentinel), including:
    - A correlation rule being disabled
    - The database filling up
- Correlated events

You can monitor the events in a tabular form or you can use several different types of charts to perform queries for recent events. Access to these features can be enabled or disabled for each user.

3.2 Introduction to the User Interface

In an Active Views, you can see *Create Active View*, *Event Real Time*, and *Event Query*. You can navigate to these functions from:

Table 3-1 *Active Views User Interface*

User Interface	Description
	The *Active Views* menu in the menu bar
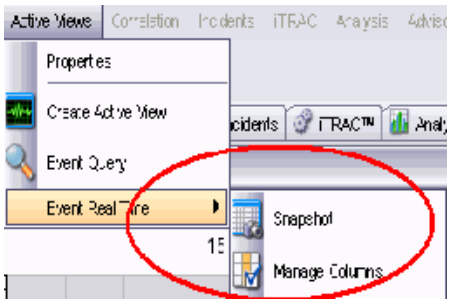	When you create a filter, the *Active Views* menu has these additional options.
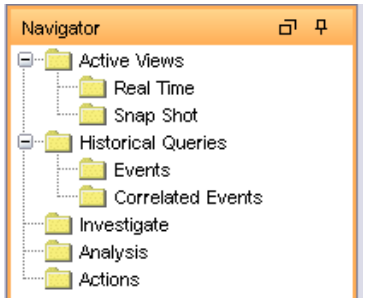	The Navigation tree in the Navigation pane

User Interface	Description
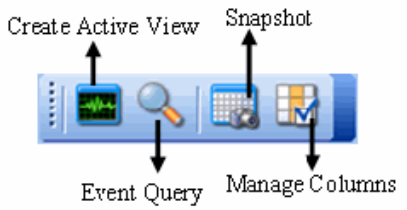	The toolbar buttons

Active Views provides two types of views that display the events in tables and graphs.

- The Table format displays the variables of the events as columns in a table. You can sort the information in the grid by clicking the column name.

Figure 3-1 *Active View Tabular Format*

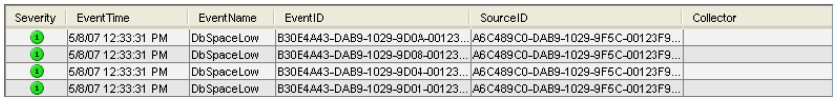

Severity	EventTime	EventName	EventID	SourceID	Collector
	5/8/07 12:33:31 PM	DbSpaceLow	B30E4A43-DAB9-1029-9D0A-00123...	A6C489C0-DAB9-1029-9F5C-00123F9...	
	5/8/07 12:33:31 PM	DbSpaceLow	B30E4A43-DAB9-1029-9D08-00123...	A6C489C0-DAB9-1029-9F5C-00123F9...	
	5/8/07 12:33:31 PM	DbSpaceLow	B30E4A43-DAB9-1029-9D04-00123...	A6C489C0-DAB9-1029-9F5C-00123F9...	
	5/8/07 12:33:31 PM	DbSpaceLow	B30E4A43-DAB9-1029-9D01-00123...	A6C489C0-DAB9-1029-9F5C-00123F9...	

- The Graphical format displays events as graphs. You can change the chart types to display other chart types.

Figure 3-2 *Active View Graphical Format*

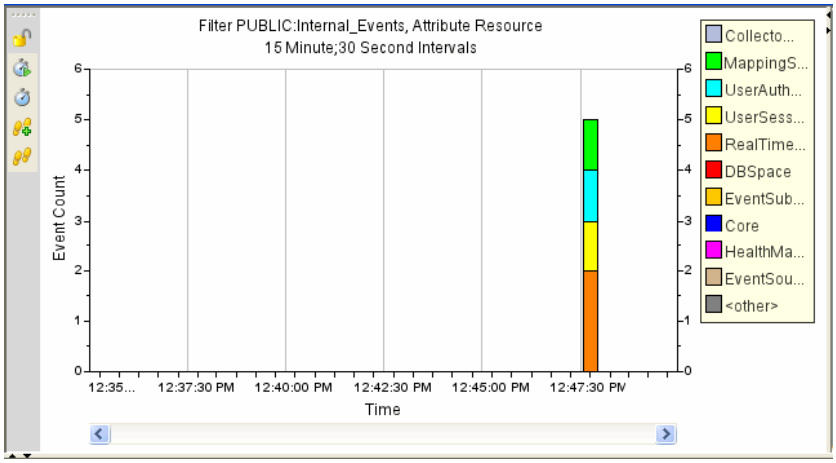

There are two types of Active Views:

- Near Real Time Event Table:
  - Holds up to 750 events per 30-second period. If there are more than 750 events, the events are displayed in the following priority order: correlated events, events that are sent to the GUI by using a global filter, and all remaining events.
  - By default, the client maintains a 24-hour period of cached events. This is configurable through Active View Properties.
  - By default, the smallest possible display interval of an active view is 30 seconds. This is represented by a gray line in the event table.

Figure 3-3 *Gray Line Smallest Possible Display Interval*

	2005.06.21 / 06:34:38 EDT			Threshold_ex
	2005.06.21 / 06:34:38 EDT	10.0.0.1	10.0.0.12	Password_ex
	2005.06.21 / 06:34:28 EDT	10.0.0.22	10.0.0.9	Program_exe

If there are more than 750 events per 30-second time period, a red separation line displays indicating that there are more events than are displayed. The other events can be viewed by using Historical Queries.

Figure 3-4 *Red Line More Events Displayed*

	2005.06.21 / 07:07:00 EDT	10.0.0.11	10.0.0.21	unsuccessfu
	2005.06.21 / 07:07:00 EDT	10.0.0.13	10.0.0.35	suspicious-fil
	2005.06.21 / 07:06:58 EDT	10.0.0.54	10.0.0.25	successful-a

- On saving user preferences, the system continues to collect data for four days. For instance, if you save your preferences, log out, and log back in the following day, your Active View displays data as if you never logged off.
- If an Active View is created and not saved, it continues to collect data for an hour. If an identical Active View is created within that hour, the Active View displays data for the last hour.

- **Snapshot:** Time-stamped views of a Real Time Event View table.

Active View provides the following unique features:

- Filter assigned to an Active View
- The z-axis attribute
- The security filter assigned to a user

The *Active Views* tab allows you to:

- Reconfigure total display time
- Add events to an incident
- Close a Snapshot or Navigator window
- Create an incident
- Custom menu options with events
- Investigate an event query
- Investigate a graph map
- View Advisor data
- Manage columnsSend messages about events by e-mail
- Show or hide event details
- Take a Snapshot of a Navigator window
- View events that triggered a correlated event
- View vulnerability visualization
- View asset data
- Integrate with the ticketing system

You can change labels (column names) to user-friendly names and the new names are populated throughout the system. For more information, see Section 3.15, “Using Custom Menu Options with Events,” on page 77.

3.3 Reconfiguring Total Display Time

Active View Properties allows you to configure the cached time in each client. The default cache time value in an Active View is 24 hours.

To configure Maximum Total Display Time:

1. Click the *Active Views* tab.
2. Click *Active Views* > *Properties*.
3. Make your changes, then click *OK*.

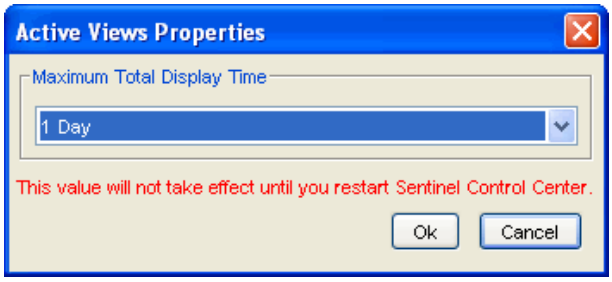

The new values do not take effect until you restart the Sentinel Control Center.

3.4 Viewing Real-Time Events

1. Click the *Active Views* tab.
2. Click *Active Views* > *Create Active View* or click the Create Active View icon .
3. In the Event Visualization Wizard window, click the down-arrows to select your Event Attribute (Z Axis), Filter, and to Display Events (Yes or No).

   In the Filter Selection window, you can build your own filter or select one of the already built filters. Selecting the All filter allows all events to display in your window. When you are creating an Active View, if the filter assigned to the Active View is changed or deleted after creation of the Active View, the Active View is unaffected.

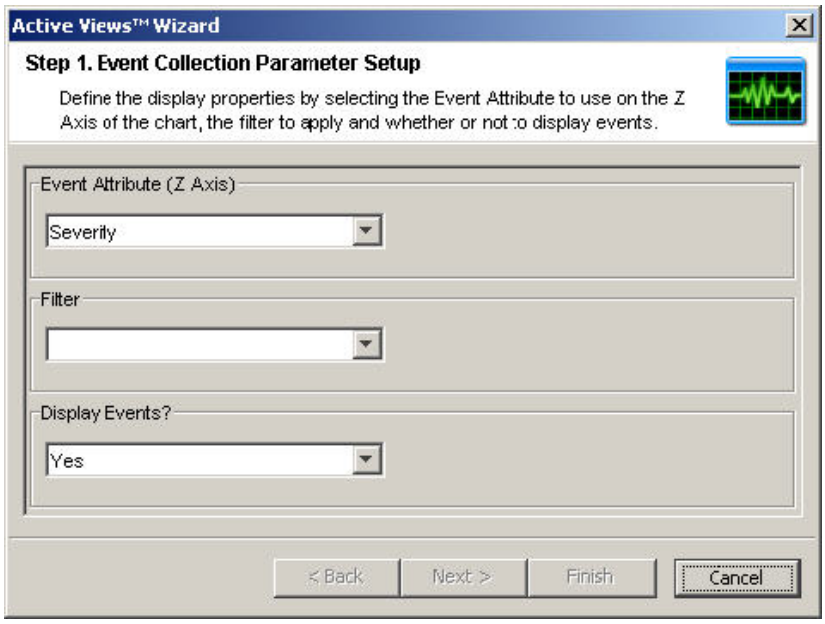

After making your selection, you can click *Next* or *Finish*. If you select *Finish*, the following default values are selected:

- Display Interval and Refresh rate of 30 seconds
- Total Display Time of 15 minutes
- Y-axis as Event Count
- Chart type of Stacked Bar 2D

4 If you click *Next*, click the down-arrows and fill in the fields:

- **Display Interval and Refresh rate:**
  - Display Interval is the time interval to display events.
  - Refresh Rate is the rate at which Active Views should refresh.
- **Total Display Time:** Amount of time to display the chart.
- **Y-axis:** Either the total Event Count or Event Count per Second.

5 Click *Next*.

6 Select your chart type from the drop-down list and click *Finish*.

Your graph looks similar to:

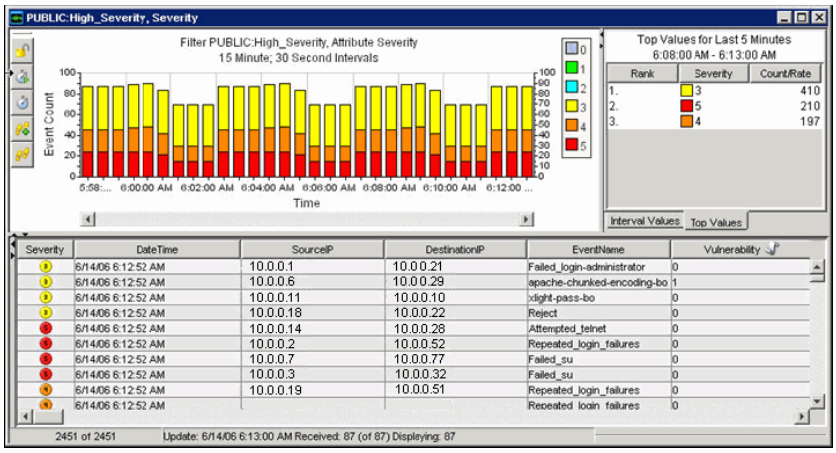

The five buttons to the left of the chart perform the following functions:

Table 3-2 *Functions of the Buttons*

Buttons	Description
Lock/Unlock the Chart	Used when performing a drill-down, zoom in, zoom out, and zoom to selection, and saving a chart as an HTML file.
Increase Display Interval	Increases the display time interval for the incoming events.
Decrease Display Interval	Decreases the display time interval for the incoming events.
Increase Display Time	Increases the time interval along the x-axis.
Decrease Display Time	Decreases the time interval along the x-axis.

When you click the *Lock* button, additional available buttons are the following:

Table 3-3 *Additional Buttons*

Buttons	Description
Lock/Unlock the Chart	Used when performing a drill-down, zoom in, zoom out, and zoom to selection, and saving a chart as an HTML file.
Zoom In	Zooms in without changing any of the time settings of the chart.
Zoom Out	Zooms out without changing any of the time settings of the chart.
Zoom to Selection	Zooms in on a selection of time intervals of events.
Snapshot Active View	Save as an HTML file with chart as images and events in a tabular format.

3.4.1 Resetting the Parameters and Chart Type of an Active View

When viewing an Active View, you can reset your chart parameters and change your chart type.

1 Within an Active View displaying a chart, right-click and select *Properties*.

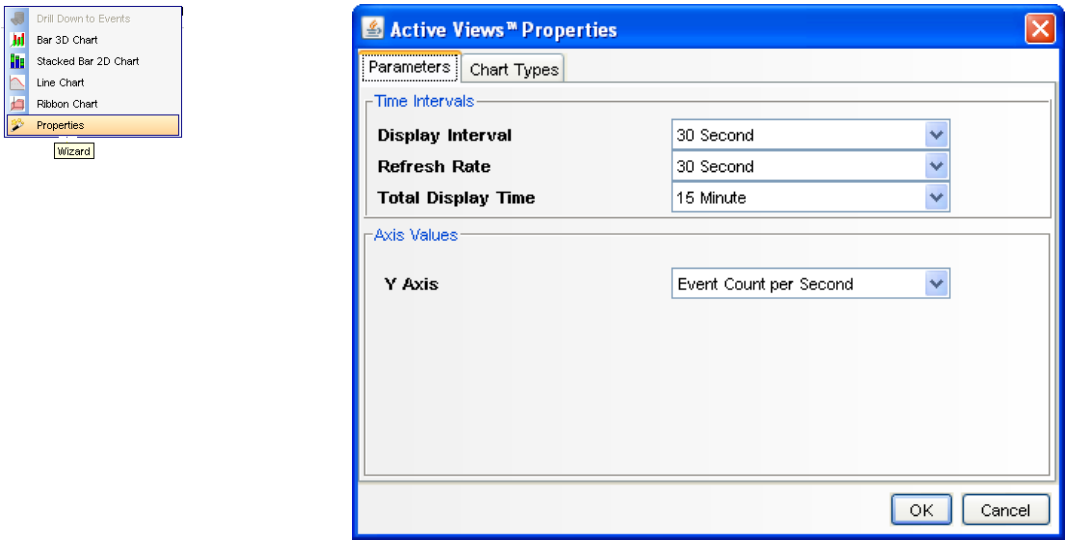

2 Under the *Parameters* tab, set the following options:

- **Display Interval:** Time between each interval.
- **Refresh Rate:** Number of seconds for the event rate to be updated.
- **Total Display Time:** Amount of time to display the chart.
- **Y-axis:** Either total Event Count or Event Count per Second.

3 Under the *Chart Types* tab, set your chart to *Stacked Bar2D*, *Bar 3D*, *Line*, or *Ribbon*.

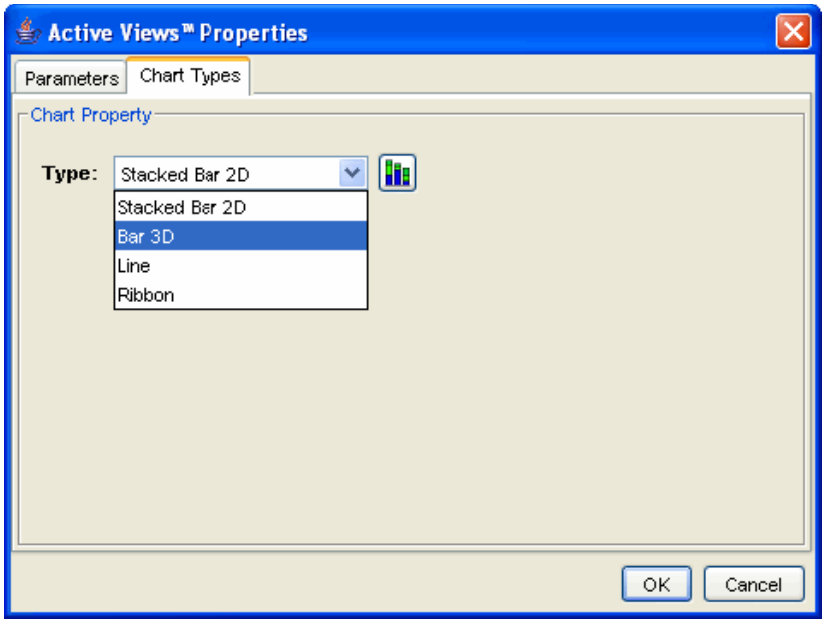

3.4.2 Rotating a 3D Bar or Ribbon Chart

1 Click anywhere on the chart and hold the mouse button.

2 Reposition the chart as desired by moving the mouse and holding the button.

3.5 Showing and Hiding Event Details

To show event details:

1 In a Real Time Event Table of the Navigator or in a Snapshot, double-click or right-click an event and click Show Details. The event details displaying the left panel of the Real Time Event Table.

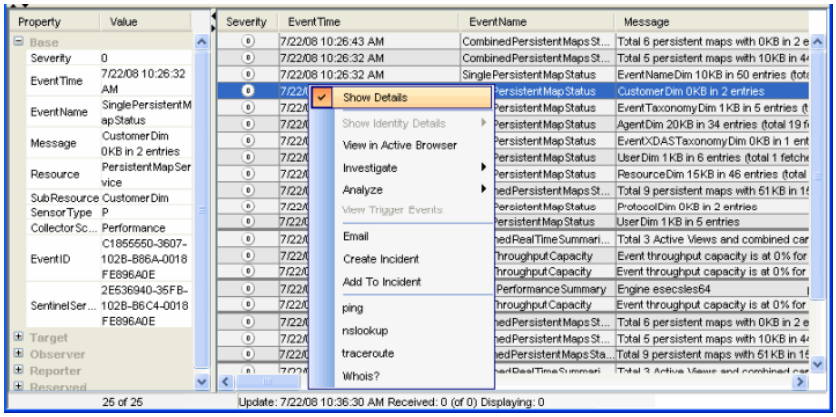

To hide event details:

1 In a Real Time Event Table of the Navigator or in a Snapshot, with event details displayed in the left panel, right-click an event and click *Show Details*. The Event Details window closes.

3.6 Sending Mail Messages about Events and Incidents

IMPORTANT: Before you send a mail by using the Sentinel Control Center, ensure that you have an SMTP Integrator configured with connection information and with the property SentinelDefaultEMailServer set to `true`.

To send an event message by e-mail:

1 In a Real Time Event Table, select an event or a group of events, then right-click and select *Email*.

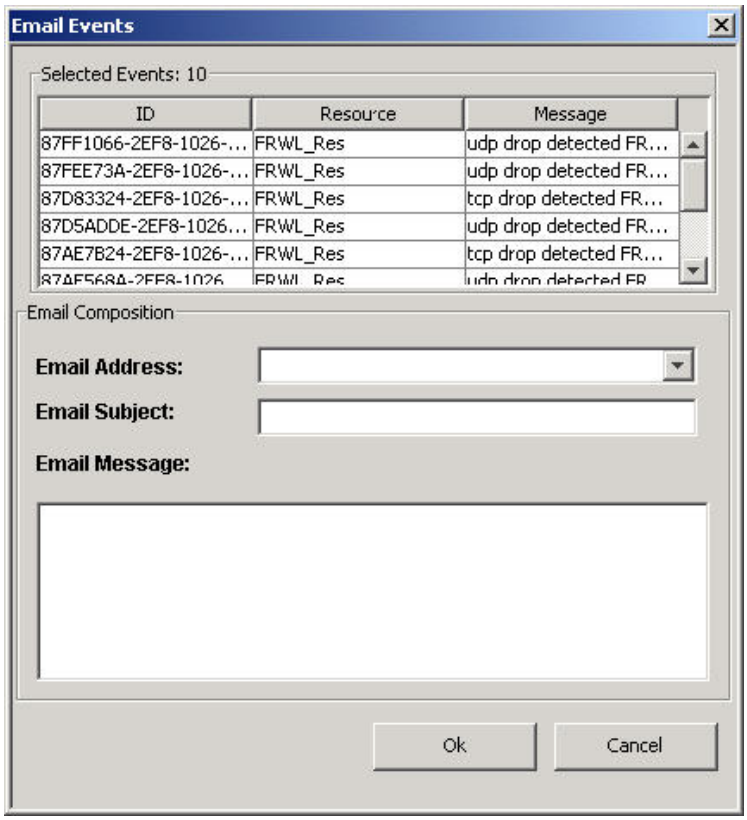

2 Provide the following information:

- Email Address
- Email Subject
- Email Message

3 Click *OK*.

To e-mail an incident:

1. After you save your incident, click the Incidents tab, *Incidents > Incidents View*.
2. Click the *All Incidents* option in the *Switch View* drop-down list located at the bottom right corner.
3. Double-click an incident.
4. Click *Email Incident* icon.
5. Provide the following information:
   - Email Address
   - Email Subject
   - Email Message
6. Click OK.

   The e-mail messages have HTML attachments that address incident details, events, assets, vulnerabilities, advisor information, attachment information, incident notes, and incident history.

3.7 Creating Incidents

To perform this function you must have user permission to create incidents.

This is useful in grouping a set of events together as a whole representing something of interest (group of similar events or set of different events that indicate a pattern of interest such an attack).

If events are not initially displayed in a newly created incident, it is probably because of a lag in the time between display in the Real Time Events window and insertion into the database. If this occurs, it takes a few minutes for the original events to be inserted into the database and display in the incident.

To create an incident:

1. In a Real Time Event Table of the Navigator or a Snapshot Real Time Event Table, select an event or a group of events, then right-click and select *Create Incident*.

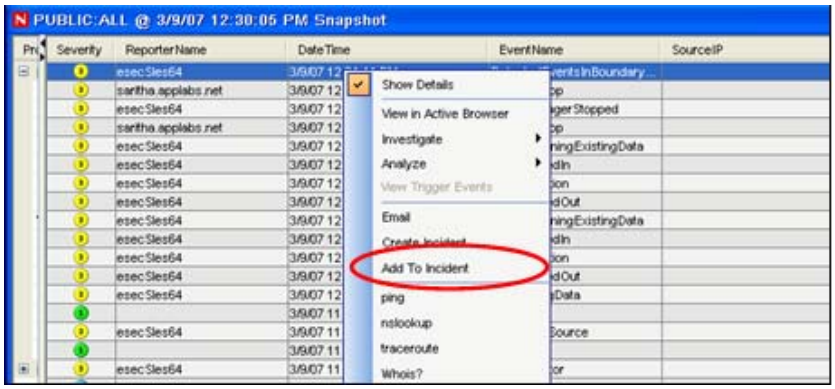

2. In the New Incident window, fill in the necessary information in the following tabs:
   - **Events:** Shows which events make up the incident
   - **Assets:** Show affected assets

- **Vulnerability:** Show related asset vulnerabilities
- **Advisor:** Asset attack and alert information
- **iTRAC:** Under this tab, you can assign a WorkFlow (iTRAC)
- **History:** Incident history
- **Attachments:** You can attach any document or text file with pertinent information to this incident
- **Notes:** You can specify any general notes regarding this incident.

3 In the Create Incident dialog box, specify:

- Title
- State
- Severity
- Priority
- Category
- Responsible
- Description
- Resolution

4 Click *Create*. The incident is added under the *Incidents* tab of the Sentinel Control Center.

3.8 Viewing Events That Trigger Correlated Events

Correlated events are determined based on the *RT2* field being set to null or any value. The *RT2* field is set to the name of the Correlation rule that triggered the Correlated event. This value is set only when the event is generated by the Correlation engine. The *Resource* field is set to `Correlation` and the *SensorType* field is set to `C`. However, the following are the exceptions:

- The *SensorType* field is set to T for the Correlated events that are routed to `gui only`.
- The *Resource* field might not be *Correlation* when the *Configure Correlated* event action is used because this action updates the correlated event Resource field.

The *View Trigger Events* option is enabled only for Correlated events.

1 In the Real-Time event table of the Navigator or Snapshot, or an Event Query table, right-click a Correlated event, and select *View Trigger Events*.

A window displays showing the events that triggered the rule and the name of the Correlation Rule.

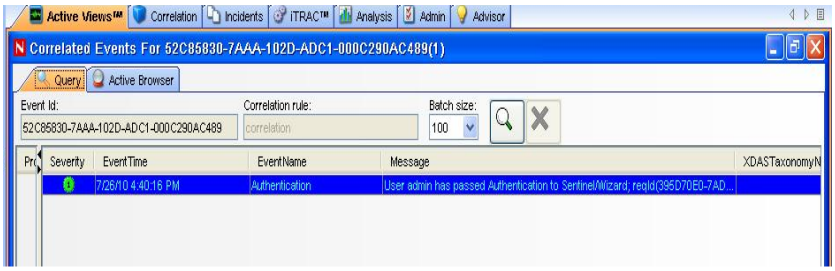

NOTE: For Correlated events, Trigger events are not available if events were routed to GUI only. However, the *View Trigger Events* option is enabled even if the Trigger events are not available.

3.9 Investigating an Event or Events

The right-click option *Investigate* allows you to:

- Perform an event query for the last hour on a single event for:
  - Other events with the same target IP address
  - Other events with the same source (initiator) IP address
  - Other targets with the same event name

  NOTE: You cannot perform a query on a null (empty) field.

- Graphically display the mappings between any two fields in the selected events. This is particularly useful to view the relationship between the initiators (IP, port, event, sensor type, Collector) and the targets (IP, port, event, sensor type, Collector name) of the selected events, but any fields can be used

Figure 3-5 is an illustration of initiator IP addresses mapped to target IP addresses.

Figure 3-5 *Graph Mapper*

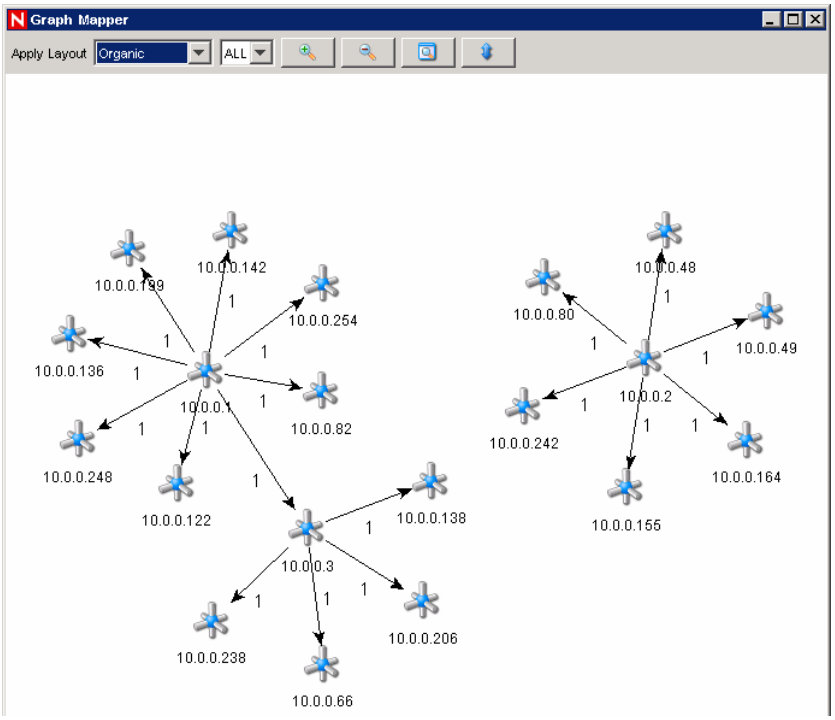

- Section 3.9.1, “Investigate: Event Query,” on page 66
- Section 3.9.2, “Investigate: Graph Mapper,” on page 66

- Section 3.9.3, “Historical Event Query,” on page 67
- Section 3.9.4, “Active Browser,” on page 68

3.9.1 Investigate: Event Query

This function allows you to perform an event query within the last hour for events similar to the selected event.

1 In a Navigator or Snapshot window, right-click an event, click *Investigate*, and select one of three options given below:

Option	Function
Show More Events to this target	Events with the same destination IP address
Show More Events from this source	Events with the same initiator IP address
What are the target objects of this event?	Events with the same event name as the selected event

An event table opens, showing the chosen event information.

3.9.2 Investigate: Graph Mapper

To create a graph map:

1 In a Real Time Event Table, right-click an event or events and select *Investigate >Show Graph*.

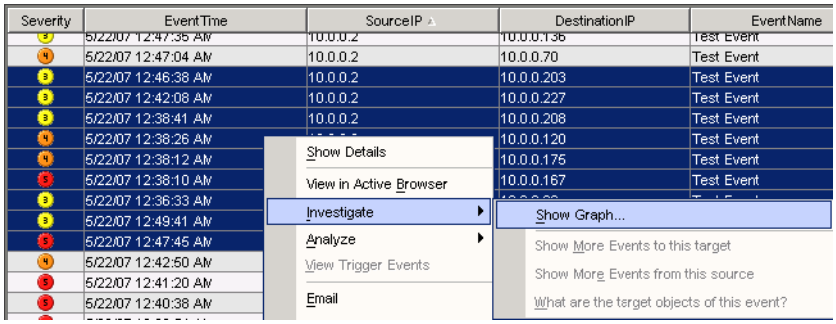

The following is a graphic depiction of Sensor Name to Event Name of severity 5 in an organic format. You can view a graphic mapping in the following formats:

- Circular
- Hierarchical
- Organic
- Orthogonal

2 You must specify the From and To fields and click *Finish*. The Graph Mapper window displays.

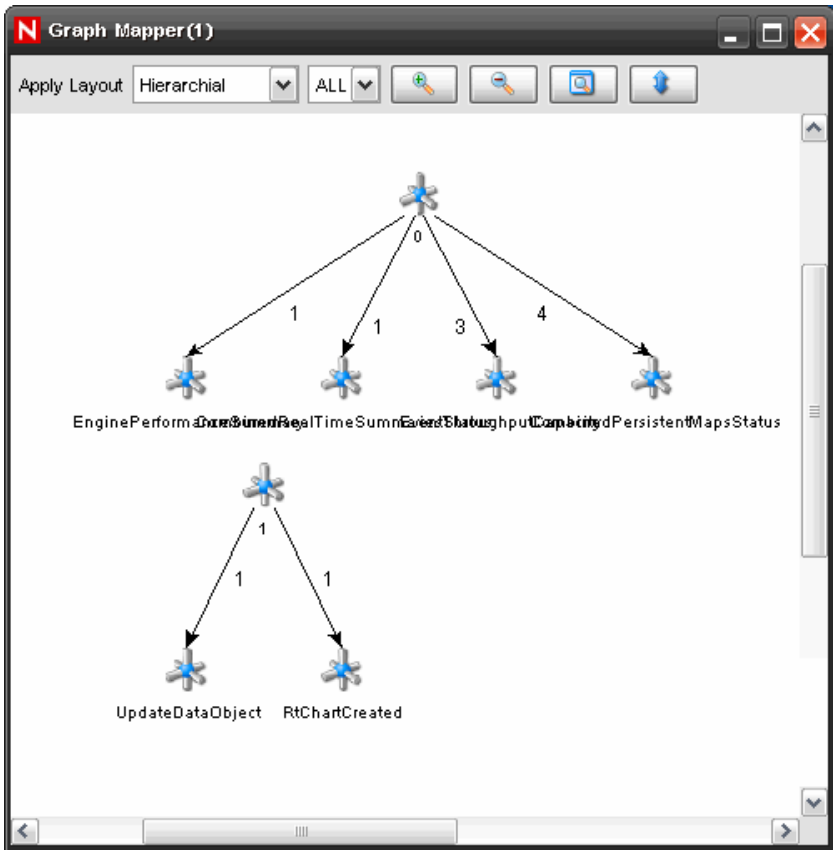

3.9.3 Historical Event Query

You can query the database for the past events through a historical event query. The events can be queried according to the filter and severity criteria in required batch size. You can export the results in HTML or CSV file format.

To query events in the Historical Event Query window:

1 In the *Active Views* tab, select *Active Views* > *Event Query*. You can also open the Historical Event Query window by clicking the *Historical Query* icon on the toolbar. The Historical Event Query window displays.

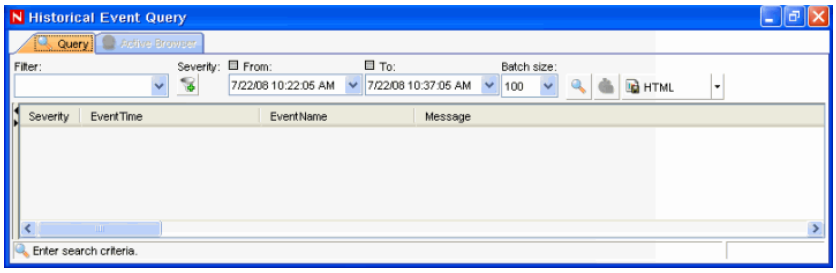

2 Click *Filter*. In Filter Selection window, select a filter from the list of available filters.

3 Click *Severity* icon. The Select Severity Values window displays.

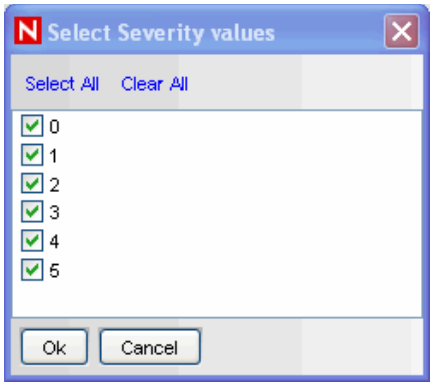

4 Select one or more values for Severity and click *OK*.

5 Select a From and To date and time.The time you select corresponds your system time.

6 Select a batch size. The events queried display in the batch size you specify.

If you select a batch size of 100, the first 100 events are displayed in the window. After the query is processed, the *Begin Searching* icon changes to the *More results* icon. You can see next 100 events along with the previous events by clicking the *More results* icon.

7 Click the *Begin Searching* icon. The query is processed. You can cancel the search by clicking the *Cancel search* icon.

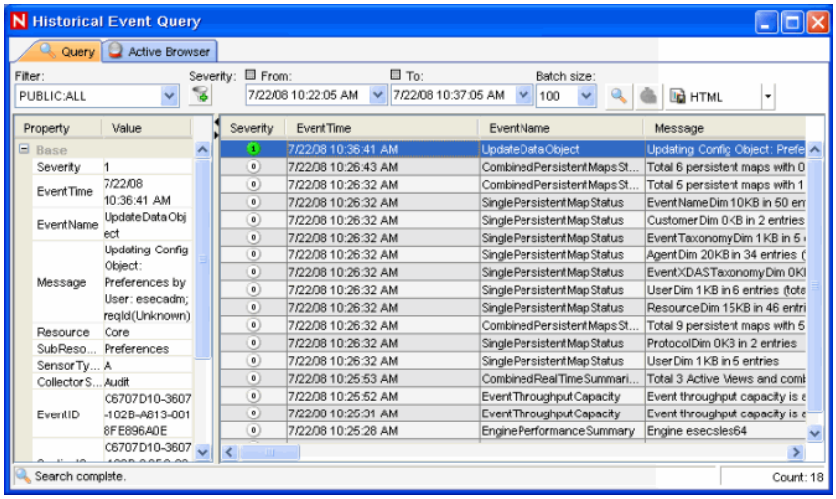

TIP: Select *HTML* or *CSV* from the drop-down list to export query results.

3.9.4 Active Browser

The Active Browser provides the ability to browse through a selected set of data to look for patterns and perform investigation. You can view the selected events in the Active Views in the Active Browser. When you open the Active Browser using *Analysis* > *Offline Query* and click *Browse* against a specific offline query, the events table is displayed only when the number of events is less than or equal to1000.

The events are grouped according to the meta tags. In these meta tags, various sub categories are defined. The numbers in the parentheses against these sub categories displays the total number of event counts corresponding to the value of the meta tag.

To view events in Active Browser:

1. In the *Active Views* tab, select the event or events you want to view in Active Browser.
2. Right-click the event or events and select *View* in the Active Browser. The selected event/s displays in the Active Browser window.

   or

   In the *Active Views* tab, select *Active Views* > *Event Query*. Historical Event Query window displays.
3. In the Historical EventQuery window, run a query and click the Active Browser tab. The selected query displays in the Active Browser window.

NOTE: The *Active Browser* tab is enabled only if the query results in at least one event display.

To view events in Active Browser in the Analysis tab:

1. In the *Analysis* tab, select the query you want to view in the Active Browser.
2. Click *Browse*. The selected query result displays in the Active Browser window.

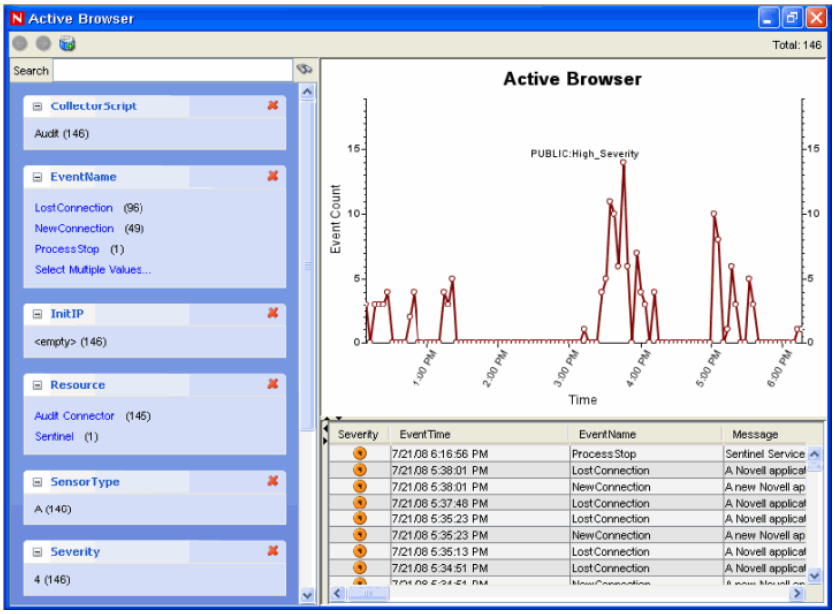

To search in the Active Browser:

1. Specify the value or text you want to search for in the *Search* field.
2. Press Enter or click the *Search* icon next to the *Search* field to search.

NOTE: You can move between the various searches by using the *Forward* and *Backward* buttons above the *Search* field.

To add attributes in Active Browser:

1 Click the *Add an attribute for categorization* icon as shown below:

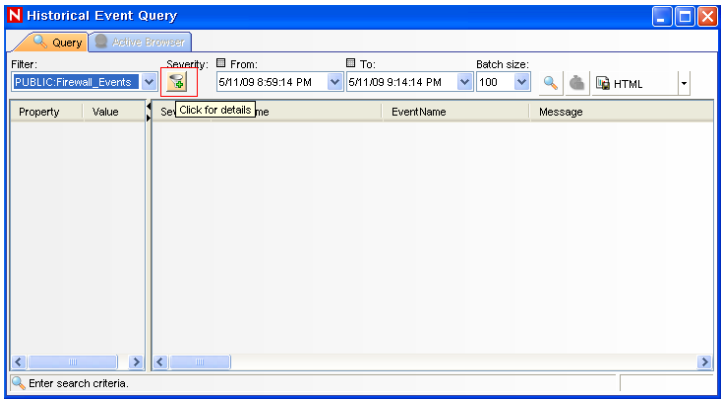

2 Select an attribute in the Add an Attribute for categorization window that displays.

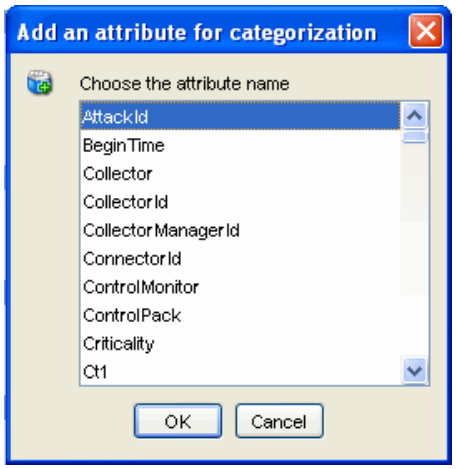

3 Click *OK*.

3.10 Viewing the Advisor Data

The Advisor provides a cross-reference between real-time intrusion detection systems attack signatures and the Advisor's knowledge base of vulnerabilities. The Advisor feed has an alert and attack feed. The alert feed contains information about vulnerabilities and viruses. The attack feed lists the exploits associated with vulnerabilities. The Advisor data is updated on a regular basis if you have opted for the optional Advisor data subscription service.

The supported intrusion detection systems are listed in Chapter 9, “Advisor Usage and Maintenance,” on page 171.

To View Advisor Data:

1 In a Real Time Event Table of the Navigator or Snapshot, right-click an event or a series of selected events, then click *Analyze* > *Advisor Data*.

If the *DeviceAttackName* field is properly populated, a report similar to the one below displays. This example is for a WEB-MISC amazon 1-click cookie theft.

Advisor Summary

Attack	Attack ID	Alert IDs
WEB-MISC amazon 1-click cookie theft	9991272,	1087, 1194, 8835, 9010
WEB-MISC amazon 1-click cookie theft	9992801	1194, 8835, 9010

Advisor Report

Microsoft Excel XLM Arbitrary Macro Execution (id 9991272) top

3 4
Urgency Severity

Microsoft Excel contains a flaw that may allow a malicious user to ru warning the user. The issue is triggered when a malicious user creat Excel macro commands, and embed commands in a spreadsheet that launch the macro without asking the user for permission. If may be p user to persuade the user to launch the file containing embedded ma loss of integrity and/or availability of data.

Scenario:

Impact:
Loss of Integrity

Safeguards:

3.11 Viewing the Asset Data

Asset data displays the asset information related to a machine or device from which you are receiving events. You can view and save the Asset data report as an HTML file. You must run your asset management Collector to view this data. The available data for viewing are:

- Hardware
  - MAC Address
  - Name
  - Type
  - Vendor
  - Product
  - Version
  - Value
  - Criticality
- Network
  - IP Address
  - Hostname
- Software
  - Name
  - Type
  - Vendor

- Product
- Version

- Contacts
  - Order
  - Name
  - Role
  - Email
  - Phone Number
- Location
  - Location
  - Address

To view Asset Data:

1. In a Real Time Event Table of the Navigator or a Snapshot window, right-click an event or multiple events.
2. Select *Analyze* > *Asset Data*.

   If both the Source IP and Destination IP are populated in an event, the asset data is displayed for both. If either of them is populated, the respective asset data is displayed.

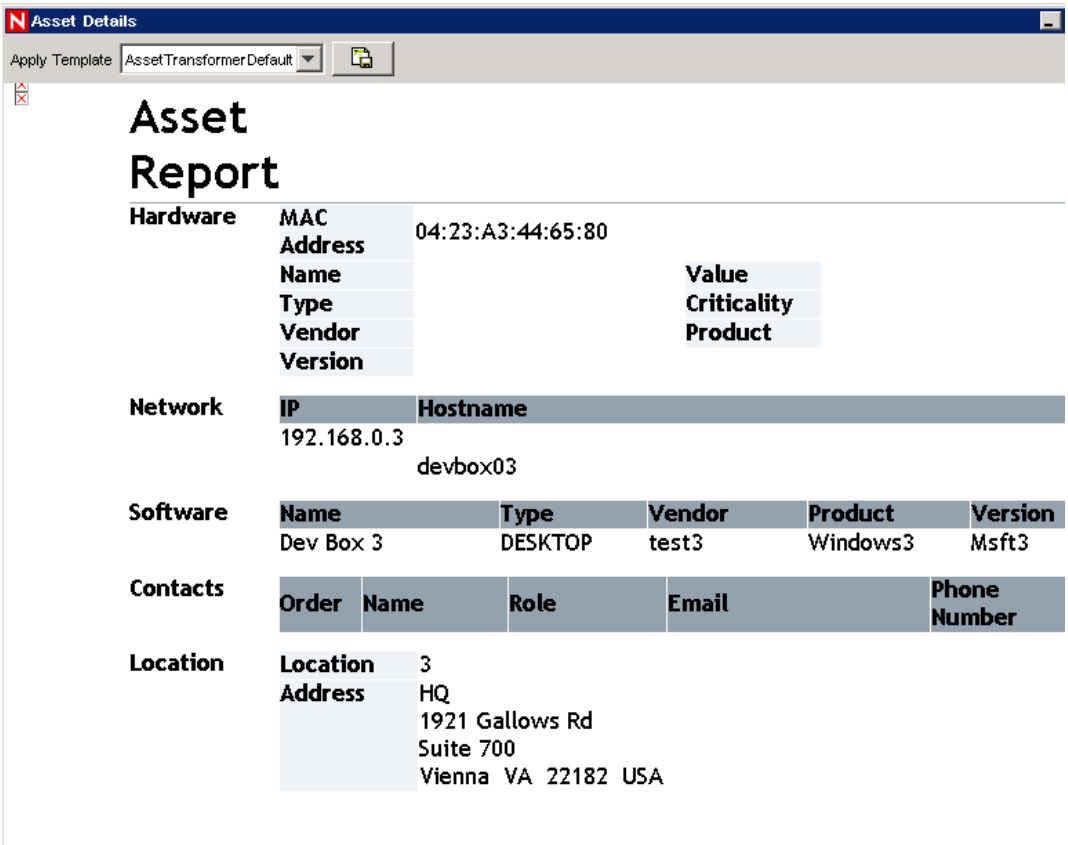

3.12 Viewing Vulnerabilities

Vulnerability Visualization provides a textual or graphical representation of the vulnerabilities of selected destination systems. Vulnerabilities for the selected destination IPs can be seen for the current time or for the time of the selected events.

Vulnerability Visualization requires that a vulnerability Collector is running and adding vulnerability scan information to the Sentinel database. The Novell Sentinel Content (http://support.novell.com/products/sentinel/secure/sentinel61.html) provides Collectors for several industry-standard vulnerability scanners, and additional vulnerability Collectors can be written by using the *Sentinel SDK* (http://www.novell.com/developer/develop_to_sentinel.html).

NOTE: Vulnerability Collectors are distinct from Event Collectors and use different commands.

There are several Vulnerability Visualization views:

- HTML
- Graphical
  - Circular
  - Organic
  - Hierarchical
  - Orthogonal

The HTML view is a report view that lists relevant fields, depending on which vulnerability scanner you have:

- IP
- Host
- Vulnerability
- Port/protocol

Figure 3-6 *Viewing Vulnerability*

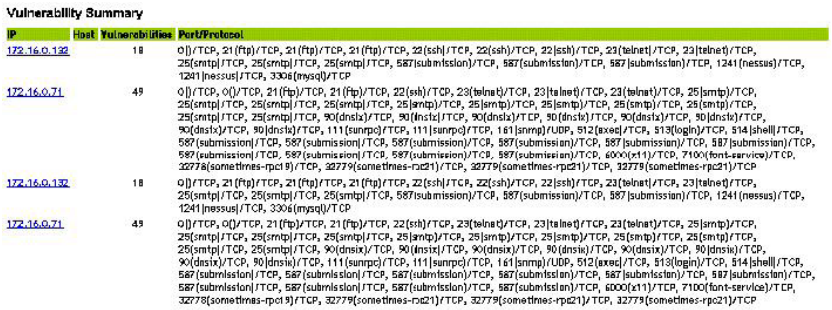

Vulnerability Summary

IP	Host	Vulnerabilities	Port/Protocol
172.16.0.132		18	0()/TCP, 21(ftp)/TCP, 21(ftp)/TCP, 21(ftp)/TCP, 22(ssh)/TCP, 22(ssh)/TCP, 22(ssh)/TCP, 23(telnet)/TCP, 23(telnet)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 1241(nessus)/TCP, 1241(nessus)/TCP, 3306(mysql)/TCP
172.16.0.71		49	0()/TCP, 0()/TCP, 21(ftp)/TCP, 21(ftp)/TCP, 22(ssh)/TCP, 23(telnet)/TCP, 23(telnet)/TCP, 23(telnet)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 90(dnsix)/TCP, 90(dnsix)/TCP, 90(dnsix)/TCP, 90(dnsix)/TCP, 90(dnsix)/TCP, 90(dnsix)/TCP, 90(dnsix)/TCP, 111(sunrpc)/TCP, 111(sunrpc)/TCP, 161(snmp)/UDP, 512(exec)/TCP, 513(login)/TCP, 514(shell)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 6000(x11)/TCP, 7100(font-service)/TCP, 32778(sometimes-rpc19)/TCP, 32779(sometimes-rpc21)/TCP, 32779(sometimes-rpc21)/TCP, 32779(sometimes-rpc21)/TCP
172.16.0.132		18	0()/TCP, 21(ftp)/TCP, 21(ftp)/TCP, 21(ftp)/TCP, 22(ssh)/TCP, 22(ssh)/TCP, 22(ssh)/TCP, 23(telnet)/TCP, 23(telnet)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 1241(nessus)/TCP, 1241(nessus)/TCP, 3306(mysql)/TCP
172.16.0.71		49	0()/TCP, 0()/TCP, 21(ftp)/TCP, 21(ftp)/TCP, 22(ssh)/TCP, 23(telnet)/TCP, 23(telnet)/TCP, 23(telnet)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 90(dnsix)/TCP, 90(dnsix)/TCP, 90(dnsix)/TCP, 90(dnsix)/TCP, 90(dnsix)/TCP, 90(dnsix)/TCP, 90(dnsix)/TCP, 111(sunrpc)/TCP, 111(sunrpc)/TCP, 161(snmp)/UDP, 512(exec)/TCP, 513(login)/TCP, 514(shell)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 587(submission)/TCP, 6000(x11)/TCP, 7100(font-service)/TCP, 32778(sometimes-rpc19)/TCP, 32779(sometimes-rpc21)/TCP, 32779(sometimes-rpc21)/TCP, 32779(sometimes-rpc21)/TCP

The graphical display is a rendering of vulnerabilities that link them to an event through common ports. Below are the examples of the four available views:

Figure 3-7 *Organic View*

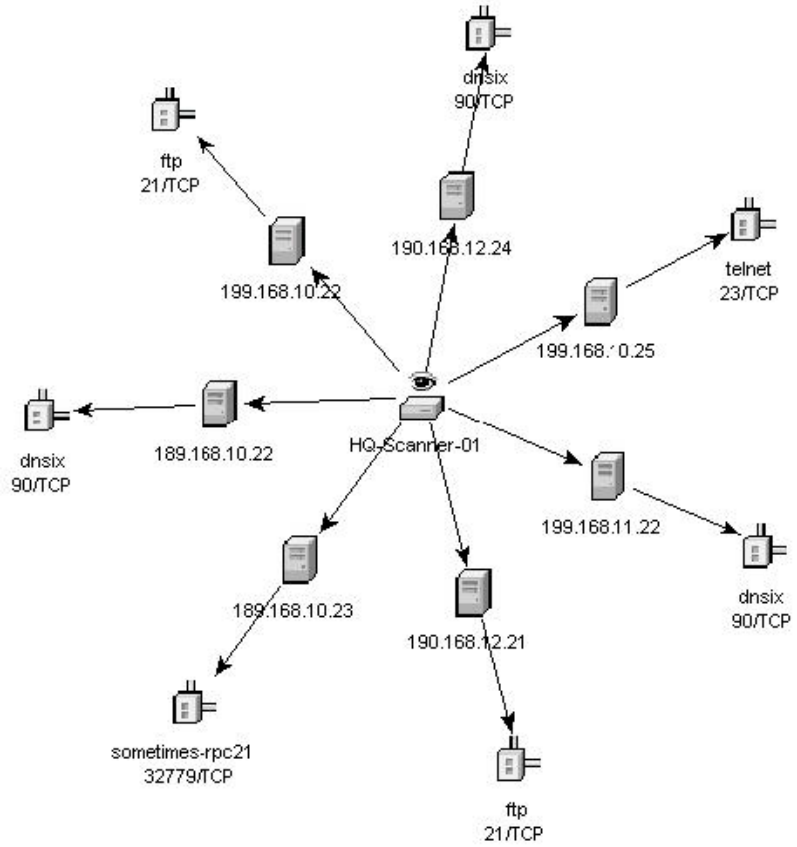

Figure 3-8 *Hierarchical View*

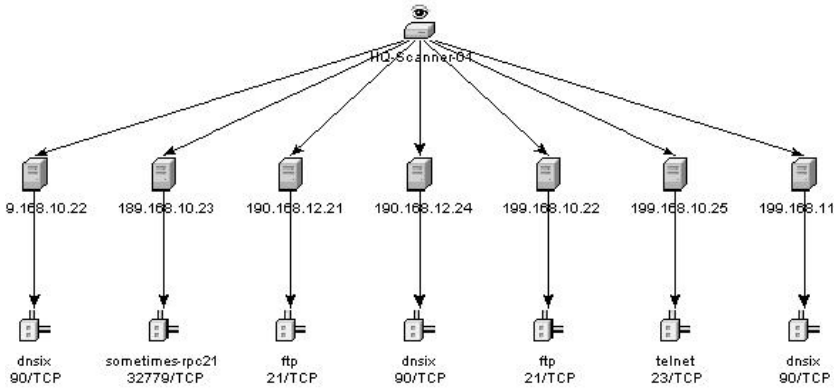

Figure 3-9 *Circular View*

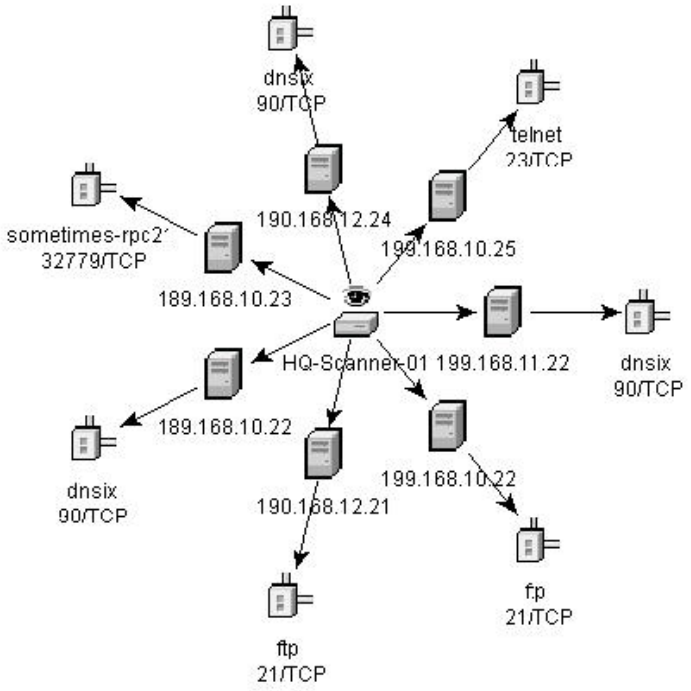

Figure 3-10 *Orthogonal View*

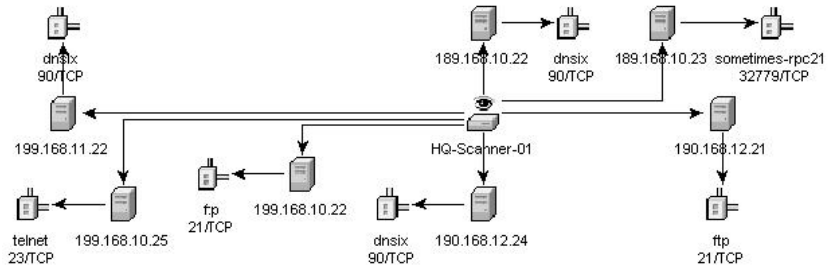

The graphical display has four panels:

- Graph panel
- Tree panel
- Control panel
- Details/events panel

The graph panel display associates vulnerabilities to a port/protocol combination of a resource (IP address). For example, if a resource has five unique port/protocol combinations that are vulnerable, there are five nodes attached to that resource. The resources are grouped together under the scanner that scanned the resources and reported the vulnerabilities. If two different scanners are used (ISS and Nessus), there are two independent scanner nodes that have vulnerabilities associated with them.

NOTE: Event mapping takes place only between the selected events and the vulnerability data returned.

The tree panel organizes data in same hierarchy as the graph. The tree panel also allows users to hide/show nodes at any level in the hierarchy.

The control panel exposes all the functionality available in the display. This includes:

- Four different algorithms to display
- The ability to show all or selected nodes which have events mapped to them
- Zooming in and out of selected areas of the graph

There are two tabs in the Details/Events panel. When you are in the *Details* tab, clicking a node displays node details. When you are in the *Events* tab, clicking an event associated with a node displays the node in tabular form as in a Real Time or Event Query window.

To run a Vulnerability Visualization:

1. In a Real Time Event Table of the Navigator or Snapshot, right-click an event or a series of selected events and click *Analysis*.
   - **Current Vulnerability:** Queries the database for vulnerabilities that are active (effective) at the current date and time.
   - **Event Time Vulnerability:** Queries the database for vulnerabilities that were active (effective) at the date and time of the selected event.

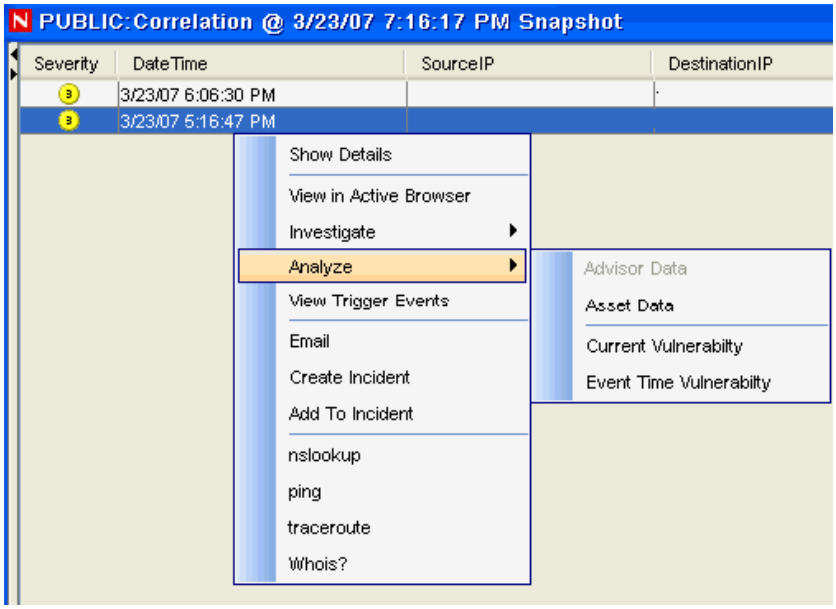

2. At the bottom the vulnerability results window, click one of the following:
   - *Event to Vulnerability Graph*
   - *Vulnerability Report*
3. (For Event to Vulnerability Graph) Adjust the display as desired:
   - Move nodes and their labels

- Use one of four different layout algorithms to display the graph
- Show all nodes or only those nodes that have events mapped to them
- Use in-line tree filtering if a large number of resources are returned as vulnerable
- Zoom in and out of selected areas

3.13 Ticketing System Integration

Novell provides optional integration modules for BMC Remedy that allow you to send events from any display screen to one of these external ticketing systems. You can also send incidents and their associated information (asset data, vulnerability data, or attached files) to Remedy.

For more information on Remedy integration, see the *Remedy Integration Guide*, available at the Novell Sentinel Content Web site (http://support.novell.com/products/sentinel/sentinel61.html) for users with a Remedy integration license.

NOTE: The permission to create Remedy incidents is controlled by the administrator on a user-by-user basis.

3.14 Viewing User Information

Novell provides optional integration with identity management systems, specifically Novell Identity Manager. With this integration, user identity information is added to incoming events when the account name matches one from Novell Identity Manager. When the *InitUserIdentity* or *TargetUserIdentity* column is populated in an event, a right-click option menu option is enabled to open the user's page in the Identity Browser.

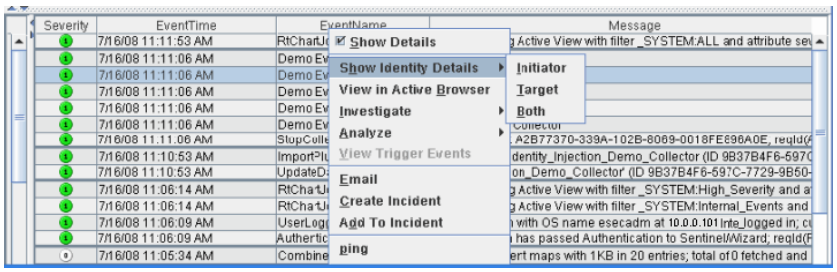

When you select *Show Identity Details*, you can choose to view the identity of the Initiator user, the Target user, or both. The Identity Browser opens and shows identifying information about the user (or users) from the identity management system, all the accounts to which the user is provisioned, and the recent activity by that user. For more information on the Identity Browser, see Chapter 18, "Identity Integration," on page 391.

3.15 Using Custom Menu Options with Events

1. In an existing Real Time Event Table of the Visual Navigator or Snapshot, right-click an event and select a menu option. The default custom menu options are as follows:
   - ping
   - nslookup

- tracert
- Whois?

The default custom menu options are available only when you right-click a single event, and are disabled when you right-click multiple events. However, custom menus with JavaScript based actions are available because JavaScript actions support multiple events. You can further assign user permissions to view vulnerability. You can add options by using the *Event Menu Configuration* option on the *Admin* tab.

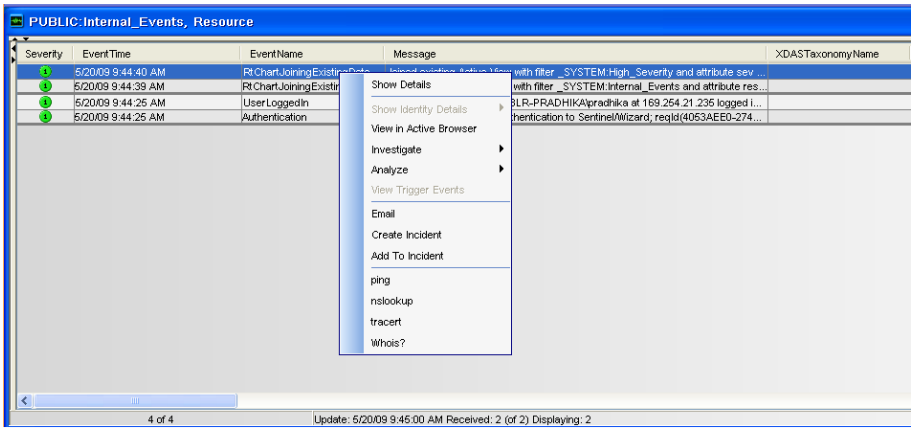

3.16 Managing Columns in a Snapshot or Navigator Window

To select and arrange columns in a Snapshot or Navigator:

1. With a Snapshot or Navigator window open, click *Active View* > *Event Real Time* > *Manage Columns* or click the Manage Columns icon of a Real Time Event Table.
2. Use the *Add* and *Remove* buttons to move column titles between the Available Columns list and the Show these columns in this order list. The Insert button can be used to insert an available column item into a specific location.

   For example, in the illustration below, clicking *Insert* places *AttackId* above *DateTime*.

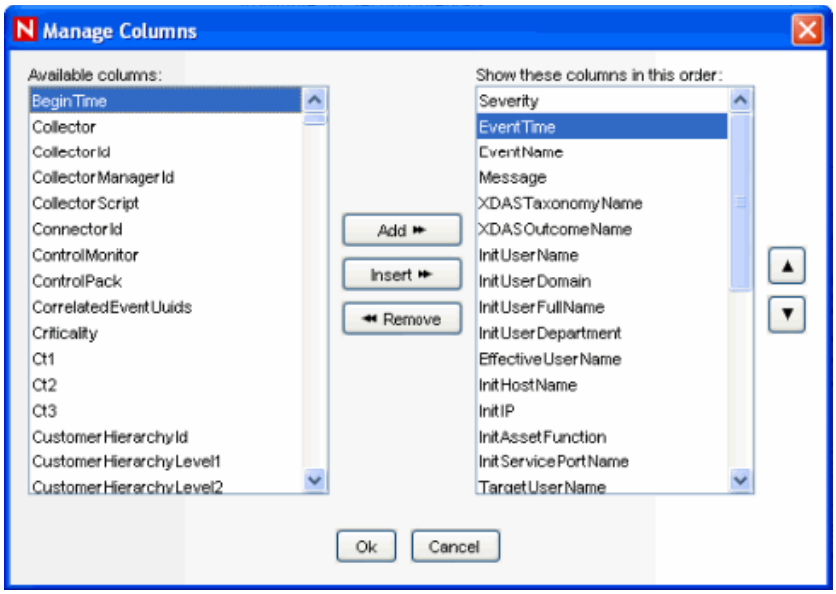

Use the up-arrow and down-arrow buttons to arrange the order of the columns as you want them to display in the Real Time Event Table. The top-to-bottom order of column titles in the Manage Column dialog box determines the left-to-right order of the columns in the Real Time Event Table.

3 In the Manage Column dialog box, click *OK*.

4 If you want your columns to display the next time you open the Sentinel Control Center, click File > Save Preferences or click the Save User Preference icon.

3.17 Taking a Snapshot of a Navigator Window

It is useful to study events this way because the Navigator refreshes automatically and the alert or alerts of interest scroll off the screen. Also, within a Snapshot, you can sort by column.

To perform this function, you must have the Snapshot user permission.

1 With a Navigator window open, click *Active Views* > *Event Real Time* > *Snapshot* or click the Snapshot Event Real Time Table icon.

A Snapshot window opens and is added to the Snap Shots folder list under Active Views in the Navigator. The graphical display is not part of the Snapshot.

3.18 Sorting Columns in a Snapshot

1 Click any column header once to sort by ascending value and twice to sort by descending value.

3.19 Closing a Snapshot or Navigator

1 When a Snapshot or Navigator is open, close it by using the *Close* button in the upper right corner.

NOTE: The view or Snapshot does not redisplay when you close and reopen the Sentinel Control Center.

3.20 Adding Events to an Incident

To perform this function you must have user permissions to Modify Incident(s) and Add to existing Incident(s).

1. In a Real Time Event Table or a Snapshot, select an event or a group of events and right-click. *Click Add To Incident*.
2. In the Add Events To Incident dialog box, click *Browse* to list the available incidents.

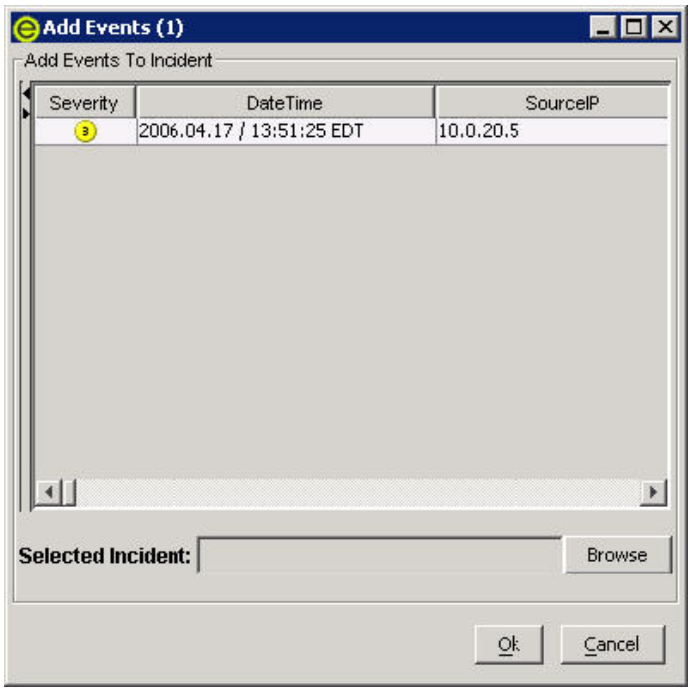

The Select Incident window displays.

3. Click *Search* to view a list of incidents with the selected criteria.

   You can define your criteria to search for a particular incident or incidents in Select Incident window.

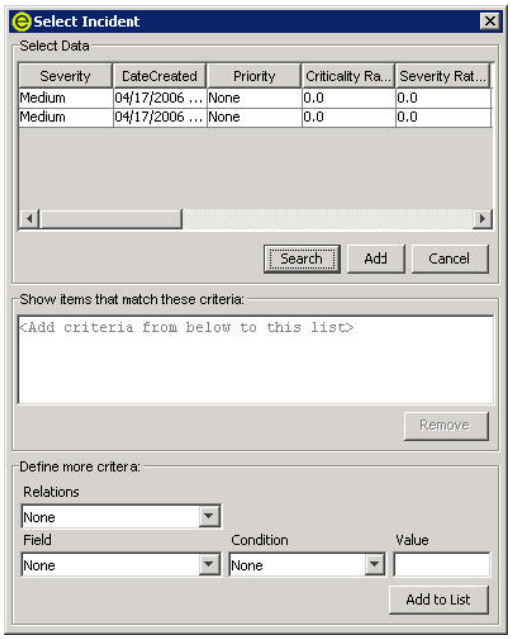

4 Select an incident and click *Add*.

5 Click *OK*. The event or events selected are added to the incident in the Incidents Navigator.

If events are not initially displayed in a newly created incident, it is probably because of a lag in the time between displaying in the Real Time Events window and insertion into the database. If this occurs, it takes a few minutes for the original events to be inserted into the database and display in the incident.

Correlation Tab

4

Sometimes, an event viewed in the system might not necessarily draw your attention. However, when you correlate a set of similar or comparable events in a given period, it might lead you to a significant event. Sentinel helps you correlate such events with the rules you create and deploy in the Correlation engine so you can take appropriate action to mitigate any alarming situation.

- Section 4.1, “Understanding Correlation,” on page 83
- Section 4.2, “Introduction to the User Interface,” on page 85
- Section 4.3, “Correlation Rules,” on page 85
- Section 4.4, “Dynamic Lists,” on page 98
- Section 4.5, “Correlation Engine,” on page 101
- Section 4.6, “Correlation Actions,” on page 102

4.1 Understanding Correlation

Correlation adds intelligence to security event management by automating analysis of the incoming event stream to find patterns of interest. Correlation allows you to define rules that identify critical threats and complex attack patterns so that you can prioritize events and initiate effective incident management and response. Starting with Sentinel 6.0, the Correlation engine is built with a pluggable framework, which allows the addition of new Correlation engines in the future.

Correlation rules define a pattern of events that should trigger, or fire, a rule. Using either the Correlation Rule Wizard or the simple RuleLG language, you can create rules that range from simple to extremely complex, for example:

- High severity event from a finance server
- High severity event from any server brought online in the past 10 days
- Five failed logins in 2 minutes
- Five failed logins in 2 minutes to the same server from the same username
- Intrusion detection event targeting a server, followed by an attempted login to root originating from that same server within 60 seconds

Two or more of these rules can be combined into one composite rule. The rule definition determines the conditions under which the composite rule fires:

- All subrules must fire
- A specified number of subrules must fire
- The subrules must fire in a particular sequence

After the rule is defined, it should be deployed to an active Correlation engine, and one or more actions can be associated with it. After the rule is deployed, the Correlation engine processes events from the real-time event stream to determine whether they should trigger any of the active rules.

NOTE: Events that are sent directly to the database or dropped by a global filter are not processed by the Correlation engine.

When a rule fires, a correlated event is sent to the Sentinel Control Center, where it can be viewed in the Active Views window.

Figure 4-1 *Active Views Window*

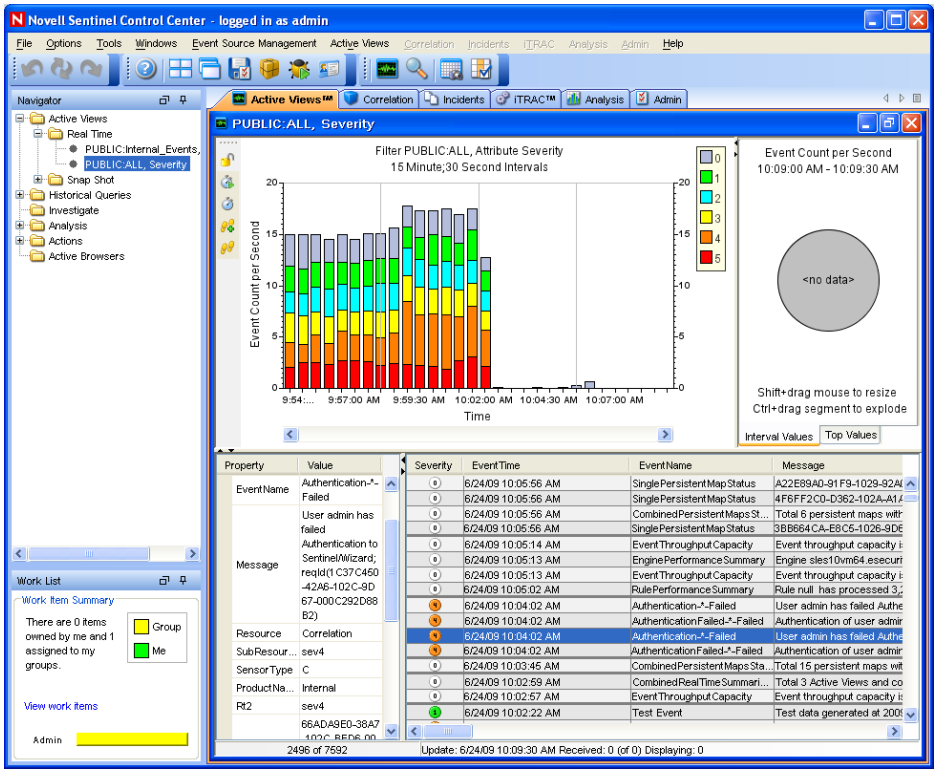

The correlated event can also trigger actions, such as sending an e-mail with the correlated event's details or creating an incident associated with an iTRAC workflow.

4.1.1 Technical Implementation

All correlation is done in-memory on the machine (or machines) that host the Correlation engine. This model allows fast, distributed processing that does not contend with database operations such as inserting events into the database.

For environments with large numbers of Correlation rules or extremely high event rates, it might be advantageous to install more than one Correlation engine and redeploy some rules to the new Correlation engine. The ability to deploy multiple Correlation engines provides the ability to scale as the Sentinel system incorporates additional data sources or as event rates increase.

Sentinel correlation is nearly real-time and depends on the time stamp for the individual events. To synchronize time, you can use an NTP (Network Time Protocol) server to synchronize the time on all devices on your network, or you can rely on the time on the Collector Manager servers and synchronize only those few machines.

Correlation relies on the data that is collected, parsed, and normalized by the Collectors, so a working understanding of the data is necessary to write rules. Many Novell Correlation rules rely on an event taxonomy that ensures that a "failed login" and an "unsuccessful logon" from two devices are classified the same.

In the *Correlation* tab, you can:

- Create/modify Correlation rules and rule folders
- Deploy Correlation rules on the Correlation engine
- Create and associate an action to a rule
- Configure dynamic lists

NOTE: Access to the correlation functions can be enabled by the administrator on a user-by-user basis.

4.2 Introduction to the User Interface

In Correlation, you can see the Correlation Rule Manager, Correlation Engine Manager, Correlation Action Manager, and dynamic lists.

You can navigate to these functions from:

Table 4-1 *Correlation User Interface*

User Interface	Description
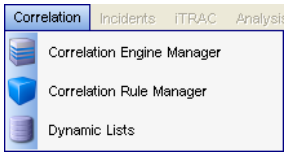	The Correlation menu in the Menu bar
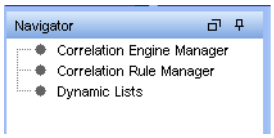	The Navigation tree in the Navigation pane
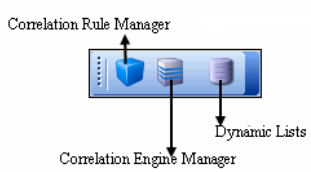	The Toolbar buttons

4.3 Correlation Rules

Correlation rules are created, modified, renamed, deployed, and undeployed in the Correlation Rule Manager. Correlation rules are organized into rule folders, which can also be managed in the Correlation Rule Manager.

NOTE: There is no limit to the number of users that can access Correlation rules. When more than one user is editing the same rule, the last person to save overwrites all previous saves.

- Section 4.3.1, “Opening the Correlation Rule Manager,” on page 86
- Section 4.3.2, “Creating a Rule Folder,” on page 86

- Section 4.3.3, “Renaming a Rule Folder,” on page 86
- Section 4.3.4, “Deleting a Rule Folder,” on page 86
- Section 4.3.5, “Creating a Correlation Rule,” on page 86
- Section 4.3.6, “Creating Correlation Rules,” on page 87
- Section 4.3.7, “Deploying and Undeploying Correlation Rules,” on page 95
- Section 4.3.8, “Enabling and Disabling Rules,” on page 96
- Section 4.3.9, “Renaming and Deleting a Correlation Rule,” on page 96
- Section 4.3.10, “Sorting Correlation Rules,” on page 96
- Section 4.3.11, “Moving a Correlation Rule,” on page 97
- Section 4.3.12, “Importing a Correlation Rule,” on page 97
- Section 4.3.13, “Exporting a Correlation Rule,” on page 98

4.3.1 Opening the Correlation Rule Manager

1. Click the *Correlation* tab.
2. In the navigator, click Correlation Rules Manager. Alternatively, click the *Correlation Rules Manager* button in the tool bar. The Correlation Rule Manager window displays.

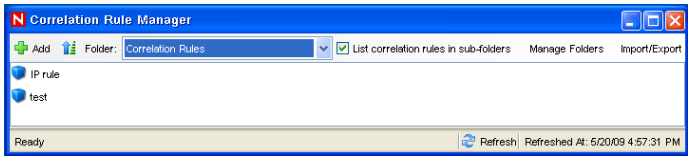

4.3.2 Creating a Rule Folder

1. Open the Correlation Rule Manager window and click Manage Folder.
2. Right-click a folder and select Add Folder.
3. Specify the Rule Folder name.

4.3.3 Renaming a Rule Folder

1. Open the Correlation Rule Manager window and click *Manage Folder*.
2. Select a folder and click *Rename*. Change the name of the folder.

4.3.4 Deleting a Rule Folder

1. Open the Correlation Rule Manager window and click *Manage Folder*.
2. Select a folder and click *Delete*. Click *Yes* when the system asks for confirmation.

4.3.5 Creating a Correlation Rule

1. Open the Correlation Rule Manager window and select a folder from the Folder drop-down list to which this rule is added.
2. Click the *Add* button located on the top left corner of the screen.

3. The Rule Wizard displays. Select one of the following rule types and follow the steps for that particular rule type:
   - Simple
   - Composite
   - Aggregate
   - Sequence
   - Custom/Freeform
4. Define the update criteria for the rule.

   If you select Continue to perform actions every time this rule fires, the rule fires every time the criteria is met. If you select *Do not perform actions every time this rule fires for the next (t) time*, the event fires only once as per user-defined time period.

   All the other events that match the Correlation rule within the specified time are grouped together with this correlated event. This user-defined time period can be a certain number of seconds, minutes, or hours.
5. Click *Next*.
6. Provide the rule name. The syntax of the rule is checked at the time it is created.
7. Under *Namespace*, select a Correlation rule folder in which to store the rule.
8. Type the description of the rule.
9. Click *Next*. The rule is created and displays in the Correlation Rule Manager window.
10. Select *Yes* if you want to create another rule or select *No* if you do not want to create another rule. Click *Next*.

The rule types and the steps to create them are described in Section 4.3.6, "Creating Correlation Rules," on page 87.

4.3.6 Creating Correlation Rules

Correlation rules can be defined in the Correlation Rule Wizard by walking through the wizard or by choosing the *Custom/Freeform* option to write the rule in the proprietary RuleLG language. All rule definitions are stored in the database in RuleLG.

Correlation rules can be defined based on any populated event field.

NOTE: When creating a rule, you can refer to a dynamic list for it. For more information, see Section 4.4.5, "Using a Dynamic List in a Correlation Rule," on page 100.

- "Simple Rule" on page 88
- "Aggregate Rule" on page 90
- "Composite Rule" on page 92
- "Sequence" on page 93
- "Custom or Freeform Correlation Rules" on page 94

Simple Rule

A simple rule is defined by specifying the events that can trigger the rule to fire (For example, firewall events, firewall events of severity 3 or higher). The filter criteria can be intersected (using the “all”option in the GUI or the “AND” operator in RuleLG) or the filter criteria can be unioned (using the “any” option in the GUI or the “OR” operator in RuleLG).

For example, a rule might be defined so that it fires anytime an event takes place on a server that is on the critical list. Another rule might be defined to fire anytime an event of severity 4 or greater takes place on a server that is on the critical list.

A simple rule requires only one event in order to fire.

For users familiar with the Correlation rule language (RuleLG), the defining operator for a simple rule is the “filter” operator. For more information about RuleLG, see “Sentinel 6.1 Rapid Deployment Correlation Engine RuleLG Language” in the *Sentinel Rapid Deployment Reference Guide*.

In Sentinel 6, filter criteria must be defined in the Correlation Rule Wizard. You cannot use existing public filters.

To create a simple rule:

1. Open the Correlation Rule Manager window and select a folder from the drop-down list to which this rule is added.
2. Click the *Add* button located on the top left corner of the screen. The Correlation Rule window displays. Select *Simple Rule*.

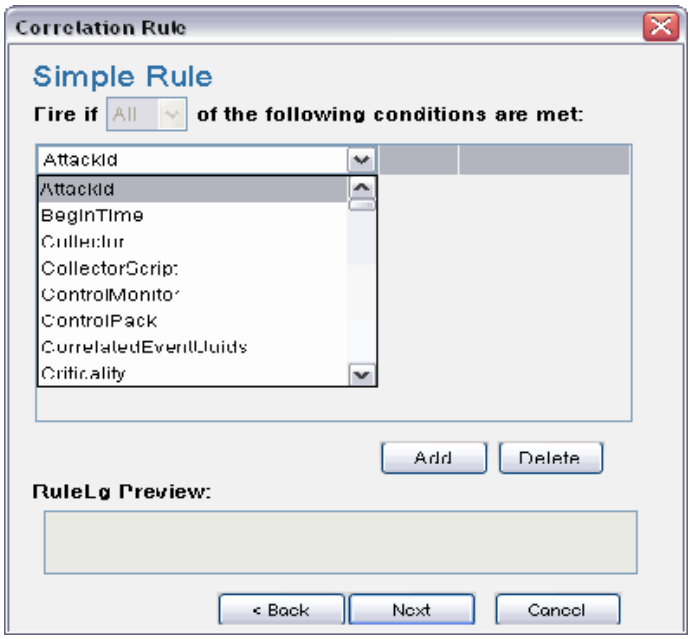

3. In the Simple Rule window, define a condition for this rule. Select the Property and Operator values from the drop-down lists and specify data in the value field.

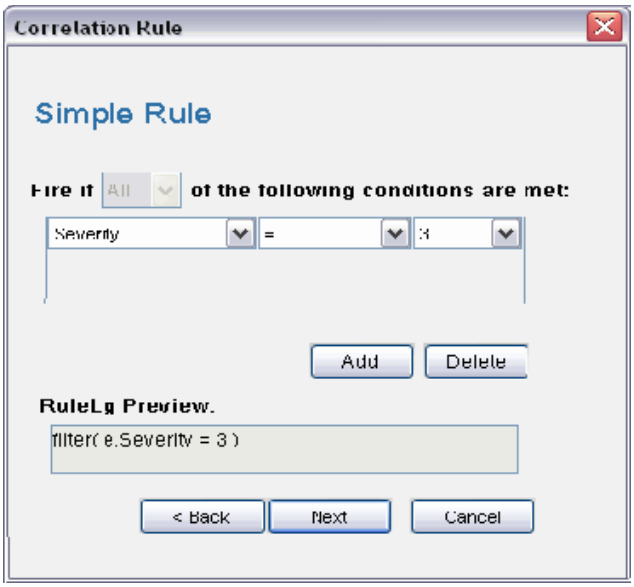

4 Click *Add* to add additional definitions for this rule.

5 Preview the rule in the RuleLG preview window. For example, `filter(e.sev=3)`.

6 Click *Next*.The Update Criteria window displays.

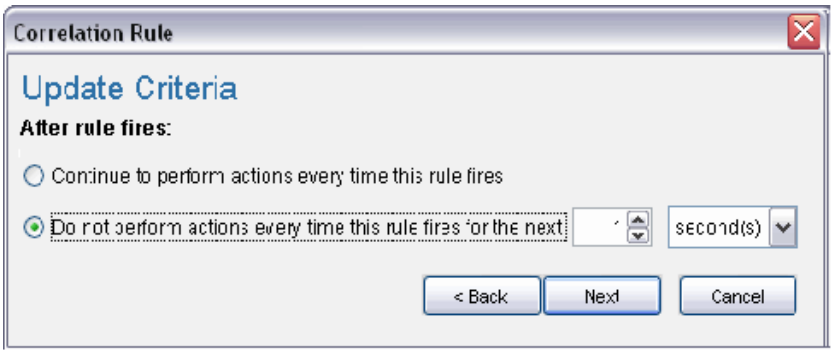

7 Enable the update criteria for the rule to fire and click *Next*. The General Description window displays.

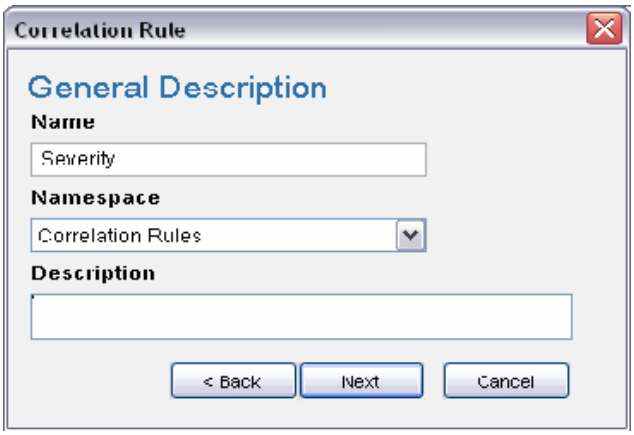

8 Provide a name for this rule. You have an option to modify the rule folder.

9 Provide rule description and click *Next*.

10 You have an option to create another rule from this wizard. Select your option and click *Next*.

Aggregate Rule

An aggregate rule is defined by specifying a subrule and the number of times the subrule must fire within a specific time window in order to trigger the aggregate rule. For example, an aggregate rule might require that a subrule fire 10 times within 5 minutes for the aggregate rule to fire.

Aggregate rules have an optional *group by* field, which can be any populated field from the events. For example, an aggregate rule might require that a subrule fire 10 times within 5 minutes where each of the 10 events has the same destination server.

NOTE: For users familiar with the Correlation rule language (RuleLG), the defining operator for an aggregate rule is the "trigger" operator. The trigger clause might also use the "discriminator" operator to define the group by field. For more information about RuleLG, see "Sentinel 6.1 Rapid Deployment Correlation Engine RuleLG Language" in the *Sentinel Rapid Deployment Reference Guide*.

To create an aggregate rule:

1 Open the Correlation Rule Manager window and select a folder from the drop-down list to which this rule is added.

2 Click the *Add* button located on the top left corner of the screen. The Correlation Rule window displays. Select *Aggregate Rule*.

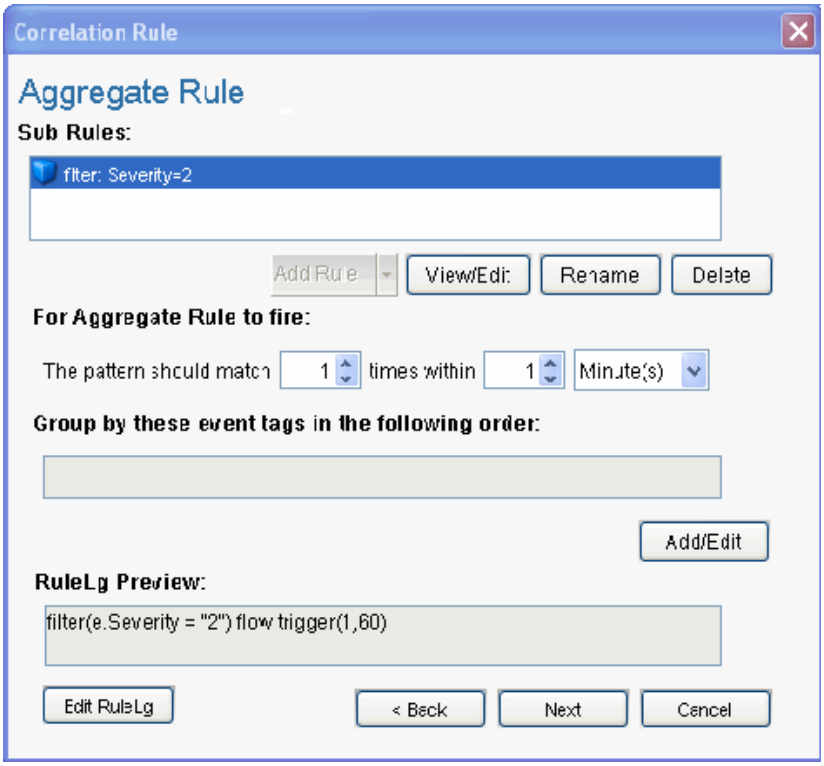

3 In Aggregate Rule window, click the *Add Rule* button to select a sub rule to create an aggregate rule. The Add Rule window displays.

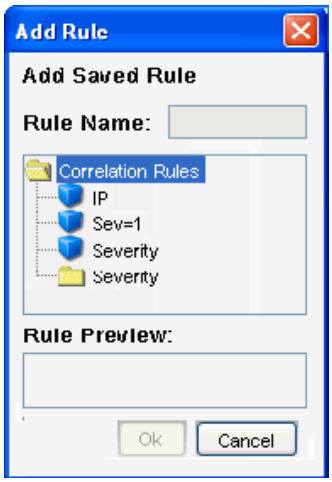

You can select only one sub rule when creating an aggregate rule.

4 Select a rule and click *OK*.

5 Set parameters for the rule to fire.

6 To group event tags according to the attributes, Click *Add/Edit*. The Attribute List window displays.

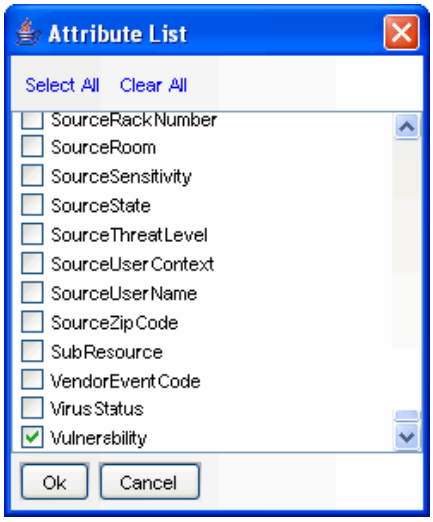

7 Select the attribute you want, then preview the rule in the RuleLG preview window.

8 Click *Next*.The Update Criteria window displays.

9 Update the criteria for the rule to fire and click *Next*. The General Description window displays.

10 Provide a name for this rule. You have an option to modify the rule folder.

11 Provide a rule description and click *Next*.

12 You have an option to create another rule from this wizard. Select your option and click *Next*.

Composite Rule

A composite rule is comprised of two or more subrules. A composite rule can be defined so that all or a specified number of the subrules must fire within the defined time frame. Composite rules have an optional *group by* field, which can be any populated field from the events.

NOTE: When a subrule is used to create a composite rule, a copy of the subrule is added to the composite rule's definition. Because a copy is added, changes to the original subrule do not affect the composite rule.

To create a composite rule:

1. Open the Correlation Rule Manager window and select a folder from the drop-down list to which this rule is added.
2. Click the *Add* button located on the top left corner of the screen. The Correlation Rule window displays. Select *Composite Rule*.

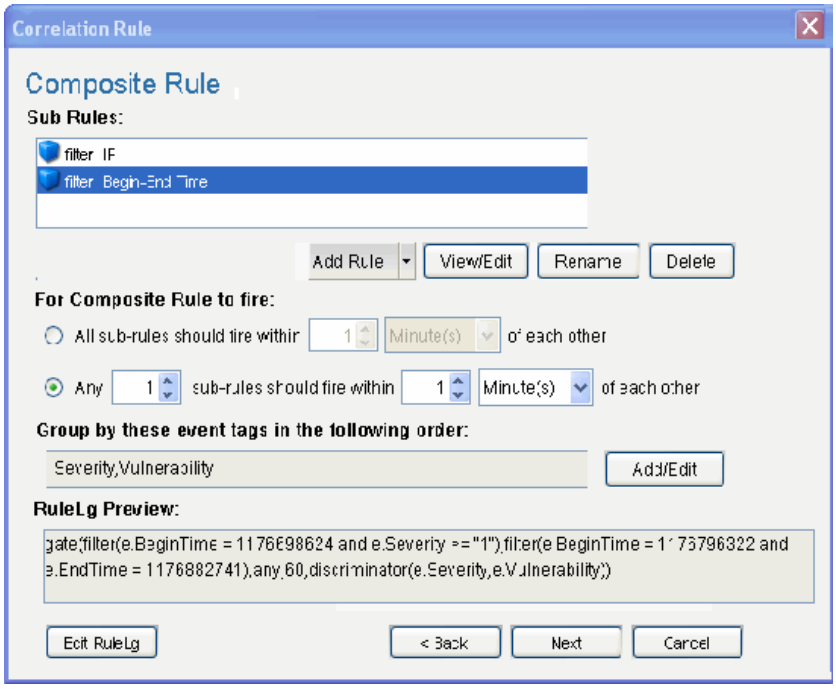

3. In the Composite Rule window, click *Add Rule* to select sub rules to create a composite rule. The Add Rule window displays.
4. Select a rule or a set of rules and click *OK*.
5. Set parameters for the rule to fire.
6. To group event tags according to the attributes, click *Add/Edit*. The Attribute window displays.
7. Select the attribute you want, then preview the rule in RuleLg preview box.
8. Click *Next*.The Update Criteria window displays.
9. Update criteria for the rule to fire and click *Next*.
10. Provide a name for this rule. You have an option to modify the rule folder.

11 Provide a rule description and click *Next*.

12 You have an option to create another rule from this wizard. Select your option and click *Next*.

Sequence

A sequence rule is comprised of two or more subrules that must be triggered in a specific order within the defined time frame. Sequence rules have an optional *group by* field, which can be any populated field from the events.

NOTE: When a subrule is used to create a sequence rule, a copy of the subrule is added to the sequence rule's definition. Because a copy is added, changes to the original subrule do not affect the sequence rule.

To create a sequence rule:

1 Open the Correlation Rule Manager window and select a folder from the *Folder* drop-down list to which this rule is added.

2 Click the *Add* button located on the top left corner of the screen. The Correlation Rule window displays. Select *Sequence Rule*.

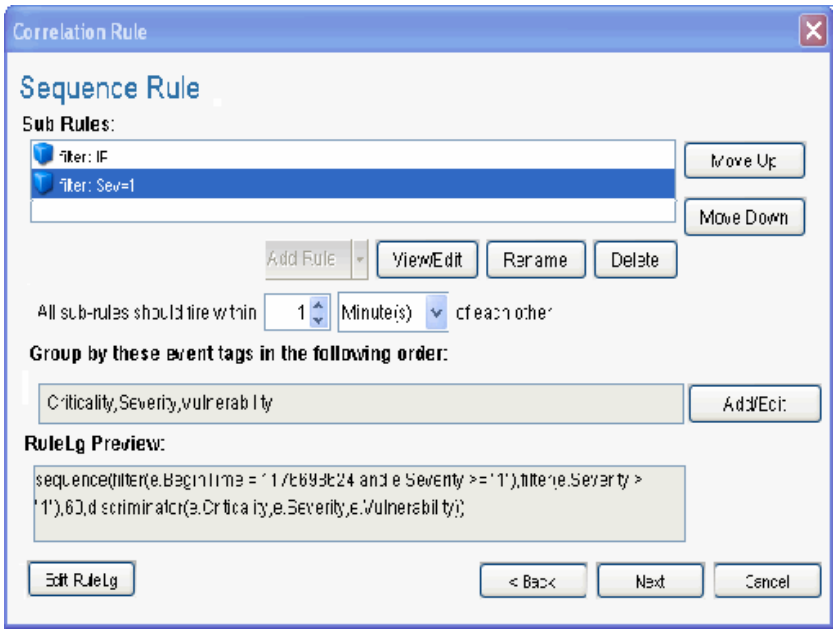

3 In the Sequence Rule window, click the *Add Rule* button to select a sub rule to create a sequence rule. The Add Rule window displays.

4 Select a rule and click *OK*.

5 Set parameters for the rule to fire. To group event tags according to the attributes, click *Add/Edit*. The Attribute List window displays.

6 Select the attribute you want, then You can preview the rule in RuleLg preview box.

7 Click *Next*.The Update Criteria window displays.

8 Update criteria for the rule to fire and click *Next*.

9 Provide a name for this rule. You have an option to modify the rule folder.

10 Provide rule description and click *Next*.

11 You have an option to create another rule from this wizard. Select your option and click *Next*.

Custom or Freeform Correlation Rules

The custom or freeform rule option is the most powerful option for creating a correlation rule. This allows the user to create any of the previous types of rules by typing the RuleLG correlation rule language directly into the Correlation Rule Wizard.

Freeform rules are the only way to include certain functionality in a correlation rule. Freeform rules give you the ability to do the following:

- Nest operations by using parentheses to specify order of operations
- Use the `inlist` operator to refer to a dynamic list
- Use the `isnull` operator to refer to unpopulated fields
- Use the `w.` prefix for a field name in the window operation to compare an incoming event's value to a set of previous events

---

TIP: You can select the functions, operators, and meta tags from the drop-down list selection. Type `e.` or `w.` in the Correlation Rule section to view the drop-down lists.

To create a custom or freeform rule:

1 Open the Correlation Rule Manager window and select a folder from the *Folder* drop-down list to which this rule is added.

2 Click the *Add* button located on the top left corner of the screen. The Correlation Rule window displays. Select *Custom/Freeform Rule*.

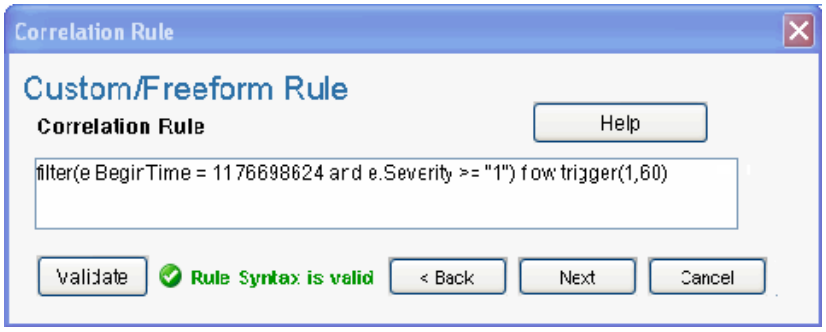

3 In the Custom/Freeform Rule window, write the condition for the rule and click *Validate* to test the validity of the rule.

4 After validation of the rule, click *Next*. The Update Criteria window displays.

5 Update the criteria for the rule to fire and click *Next*.

6 Provide a name for this rule. You have an option to modify the rule folder.

7 Provide rule description and click Next.

8 You have an option to create another rule from this wizard. Select your option and click *Next*.

4.3.7 Deploying and Undeploying Correlation Rules

Correlation rules can be deployed or undeployed from the Correlation Engine Manager or the Correlation Rule Manager. You can undeploy all rules or a single rule.

The rules can be associated with one or more actions. If no action is selected, a default correlated event is generated with the following values:

Table 4-2 *Default Correlated Event Details*

Field Name	Default Values
Severity	4
Event Name	Same as the event name for the trigger event
Message	Same as the message for the trigger event
Resource	Correlation
SubResource	<Rule Name>

Other types of actions can be configured in the Action Manager:

- Configure a Correlated Event replaces the default correlated event settings
- Add to Dynamic List adds an element to a dynamic list
- Remove from Dynamic List removes an element from a dynamic list
- Execute a Command executes a shell or batch script
- Execute a Script executes a script; only available for actions created in Sentinel 6.0
- Send an Email by using default Sentinel mail settings
- Create an Incident creates a Sentinel incident
- Configure any Action from the Action Manager that was created from an Action plug-in that takes a correlated event as input. For more information on the Action Manager, see Chapter 17, “Action Manager and Integrator,” on page 363.

To deploy correlation rules in the Correlation Engine Manager:

1. Open the Correlation Engine Manager window.
2. Right-click the engine you want to deploy the rule on and select *Deploy Rule*.
3. In the *Rules* tab, select the rule or rules you want to deploy.
4. In the *Actions* tab, select the action or actions you want to associate with the rule.
5. Click *Deploy*. Rules are deployed in an enabled state.

To deploy correlation rules in the Correlation Rule Manager:

1. Open the Correlation Rule Manager window.
2. Select a rule and click the *Deploy rules* link. The Deploy Rule window displays.
3. In the Deploy Rule window, select the engine to deploy the rule from the drop-down list.
4. (Optional) Select an action or add a new action.

If nothing is selected, a Correlated event with default values is created.

5 Click *Deploy*.

To undeploy a single rule:

1 In the Correlation Engine Manager, right-click the rule and select *Undeploy Rule*.

or

In the Correlation Rule Manager, select the rule and click the *Undeploy rule* link.

To undeploy all correlation rules:

1 Open the Correlation Engine Manager window.
2 Right-click the Correlation engine and select *Undeploy All Rules*.

4.3.8 Enabling and Disabling Rules

1 Open the Correlation Engine Manager window.
2 Right-click the rule or set of rules and select *Enable Rule* or *Disable Rule*.

4.3.9 Renaming and Deleting a Correlation Rule

To rename a correlation rule:

NOTE: You must undeploy a rule before you rename or delete the rule.

---

1 Open the Correlation Rule Manager window and select the rule you want to rename.
2 If the rule is deployed, click the *Undeploy Rule* link to undeploy the rule.
3 Click the *View/Edit* link. In the *General Description* tab, change the name of the Correlation rule.
4 Click *OK*.

To delete a correlation rule:

1 Open the Correlation Rule Manager window and select the rule you want to delete.
2 If the rule is deployed, click the *Undeploy Rule* link to undeploy the rule.
3 Click the *Delete* link. Click *Yes* when the system prompts for confirmation.

4.3.10 Sorting Correlation Rules

To sort the list of correlation rules, click the Sort button at the top left of the Correlation Rule Manager window.

4.3.11 Moving a Correlation Rule

1 Open the Correlation Rule Manager window and click *Manage Folder*.

2 Drag a correlation rule from one folder to another.

4.3.12 Importing a Correlation Rule

1 Open the Correlation Rule Manager window and click the *Import/Export Correlation Rule* icon.

The Import Export Rule window displays.

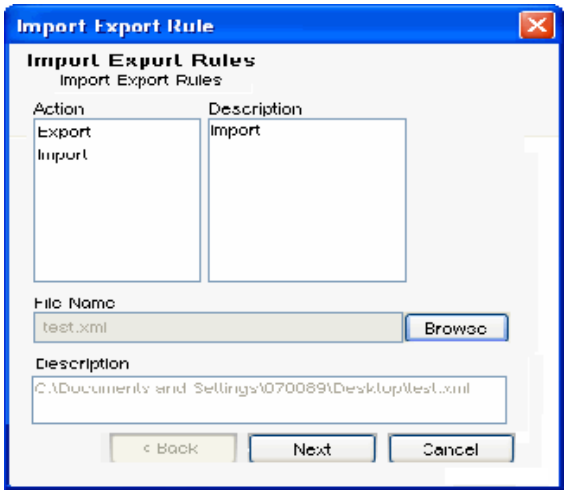

2 Select the *Import* option from the Action pane. The description in the *Description* pane changes to *Import*.

3 Click *Browse* to select the Correlation rule you want to import. Select the file and click *Import*, then click *Next*. The Import Rule window displays.

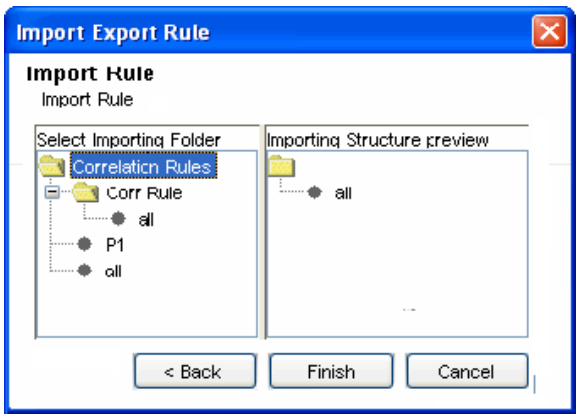

4 Select the folder you want to import the Correlation rule into, then click *Finish*.

When importing a correlation rule in a folder, if a correlation rule with the same name exists, the system displays a message and does not import the file.

IMPORTANT: If you import a correlation rule using the `inlist` operator, the dynamic list aligned to that rule must exist or you must create the dynamic list with the same name on the system to which it is imported.

4.3.13 Exporting a Correlation Rule

1. Open the Correlation Rule Manager window and click the *Import/Export Correlation Rule* icon. The Import Export Rule window displays.
2. Select the *Export* option from the Action pane. The description in the *Description* pane changes to *Export*.d
3. Click *Browse* to export the rule. Specify a filename and click *Export*, then click Next. The Export Rule window displays.

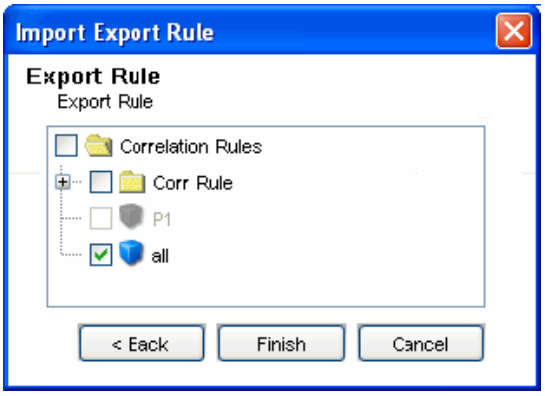

4. Select the Correlation rule you want to export. Click *Finish*.

4.4 Dynamic Lists

Dynamic lists are distributed list structures that can be used to store string elements, such as IP addresses, server names, or usernames. The lists are then used within a Correlation rule for a quick lookup to see whether an incoming event includes an element from the dynamic list. Some examples of dynamic list include:

- Terminated user lists
- Suspicious user watchlist
- Privileged user watchlist
- Authorized ports and services list
- Authorized server list

A dynamic list can be built by using the text values for any event meta tag. Elements can be added to the list manually (by an administrator) or automatically whenever a Correlation rule fires. Elements can be removed from a list manually (by an administrator), automatically whenever a correlation rule fires, when their time limit expires, or when the maximum list size is reached.

IMPORTANT: The Time To Live (TTL) must be between 60 seconds and 90 days and the maximum list size is 100,000.

Regardless of how the values were added, they can be persistent (active until manually removed or until the maximum list size is reached) or transient (active only for a specified time frame after being added to the list, also known as the Time to Live). The Time to Live can range from 60 seconds to 90 days.

NOTE: If the Time to Live period is updated on an active dynamic list, the change is not retroactive to elements already on the list. Elements that are already added to the dynamic list retain their original Time to Live.

4.4.1 Adding a Dynamic List

1. Click *Correlation* on the menu bar and select Dynamic Lists. Alternatively, you can click the *Dynamic Lists* button on the toolbar.
2. Click the *Add* button located on the top left corner of the screen. The Dynamic List Properties window displays.
3. Provide the name of the list.

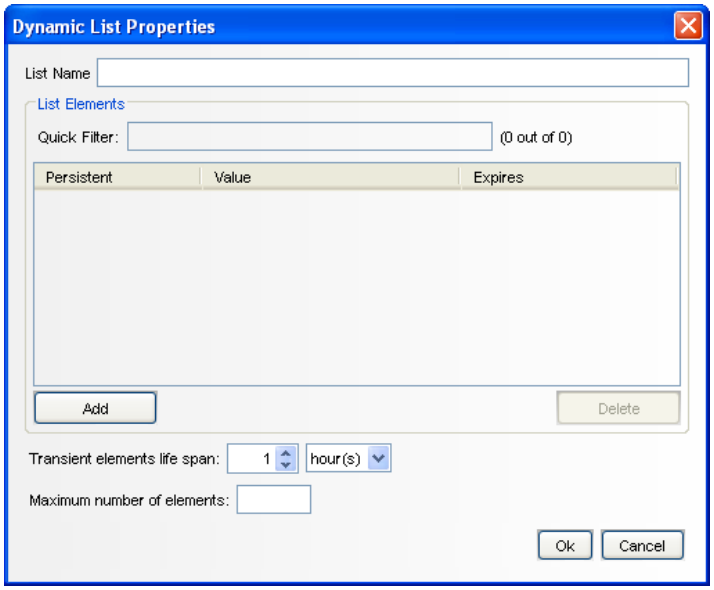

The name cannot contain special characters, such as quotations or hyphens.

4. Click *Add*. The Add Element window displays:

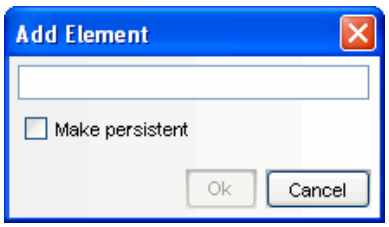

5. Provide the name of the Element. To make the Element persistent, select the *Make Persistent* check box and click *OK*.

To make an existing element persistent, select the check box next to the element name in the Dynamic Properties window.

6 Select *Transient elements life span*, then specify the time the persistent values are active in the list

7 Specify the maximum number of elements. The number defined here limits the number of elements in the list.

8 Click *OK*.

.

4.4.2 Modifying a Dynamic List

1 Click *Correlation* on the menu bar and select *Dynamic Lists*. Alternatively, you can click the *Dynamic Lists* button on the toolbar.

2 Select a dynamic list and click the *View/Edit* link.

3 The Dynamic List Properties window displays. Edit the options as required and click *OK*.

4.4.3 Deleting a Dynamic List

WARNING: Do not delete a dynamic list that is part of a correlation rule or rules.

1 Click *Correlation* on the menu bar and select *Dynamic Lists*. Alternatively, you can click the *Dynamic Lists* button on the toolbar.

2 Select a dynamic list and click the *Delete* link next to it. A confirmation message alert displays.

3 Click *Yes* to delete the list.

4.4.4 Removing Dynamic List Elements

There are several ways an element can be removed from a dynamic list:

- A user can remove it manually
- The element can be removed by a Correlation rule action
- The transient element life span can expire
- If the maximum number of elements for a dynamic list is reached, elements are removed from the list to keep the list at or below the maximum list size. The transient elements are removed (from oldest to newest) before any persistent elements are removed.

4.4.5 Using a Dynamic List in a Correlation Rule

Dynamic lists can be referenced in a Correlation rule by using the *Custom/Freeform* option of the Correlation Rule Wizard. For example:

```
filter(e.<tagname> inlist <Dynamic List Name>)
```

Where, `e.<tagname>` represents a meta tag in the incoming event, such as `e.shn` (Source Host Name) or `e.dip` (Destination IP address)
<Dynamic List Name> is the name of an existing Dynamic List, such as CriticalServerList

The following instructions assume that a dynamic list already exists.

To add a dynamic list to correlation rule:

1. Open the Correlation Rule Manager window and select a folder from the drop-down list to which this rule is added.
2. Click the *Add* button located on the top left corner of the screen. The Correlation Rule window displays. Select *Custom/Freeform Rule*.
3. In the Custom/Freeform Rule window, write the condition for the rule, including the name of the dynamic list. For example, `filter(e.sev inlist Severity)` where Severity is the dynamic list name.
4. Click *Validate* to test the validity of the rule.
5. After validation of the rule, click *Next*. The Update Criteria window displays.
6. Update the criteria for the rule to fire and click *Next*.
7. Provide a name for this rule. You have an option to modify the rule folder.
8. Provide a rule description and click *Next*.
9. You have an option to create another rule from this wizard. Select your option and click *Next*.

NOTE: Users must have the permission to Start/Stop the Correlation engine to perform these actions.

The two states of Correlation engine are:

Table 4-3 *States of the Correlation Engine*

States	Icons
Enable	
Disable	

When the Correlation engine is enabled, it processes active Correlation rules. When it is in a disabled state, all in-memory data is preserved and no new Correlation events are generated. Disabling the Correlation engine does not affect other parts of the Sentinel system.

Correlation rules are stored in the Sentinel database. When you activate the Correlation engine in the Sentinel Control Center, it requests the deployment information and rules from the database. Changes to a rule are not reflected in the Correlation engine until one of the following things happens:

- The rule is undeployed, edited, and redeployed.
- The rule is freshly deployed

4.5 Correlation Engine

- Section 4.5.1, “Starting or Stopping a Correlation Engine,” on page 102
- Section 4.5.2, “Renaming a Correlation Engine,” on page 102

4.5.1 Starting or Stopping a Correlation Engine

1. Open the Correlation Engine Manager window.
2. Right-click a correlation engine and select *Start Engine* or *Stop Engine*.

4.5.2 Renaming a Correlation Engine

A Sentinel system can have one or more correlation engines. You can rename the engines if desired.

1. Open the Correlation Engine Manager window.
2. Right-click the correlation engine and select *Rename Engine*.
3. Modify the name of the engine and click *OK*.

4.6 Correlation Actions

The Action Manager allows you to configure repeatable actions. There are several different types of actions that can be configured and then associated with a correlation rule deployment:

- Section 4.6.1, “Configuring a Correlated Event,” on page 103
- Section 4.6.2, “Adding to a Dynamic List,” on page 104
- Section 4.6.3, “Removing a Value from a Dynamic List,” on page 105
- Section 4.6.4, “Executing a Command,” on page 106
- Section 4.6.5, “Creating an Incident,” on page 107
- Section 4.6.6, “Sending an E-mail,” on page 108
- Section 4.6.7, “Imported JavaScript Action Plugins,” on page 108

NOTE: Although all of these actions can be used in Correlation rule deployments, only the JavaScript actions can be used in other areas of the Sentinel Control Center. For more information, see Chapter 17, “Action Manager and Integrator,” on page 363.

Actions associated with a Correlation rule are executed when the deployed Correlation rule fires (with the frequency of the execution determined by settings on the Update Criteria window of the Correlation Rule Wizard).

If no action is specifically selected when deploying a correlation rule, a correlated event with the following default settings is created:

Table 4-4 *Default Settings*

Field Name	Default Values
Severity	4
Event Name	Final Event Name
Message	<message>
Resource	Correlation
SubResource	<Rule Name>

4.6.1 Configuring a Correlated Event

Figure 4-2 *Configure Correlated Event*

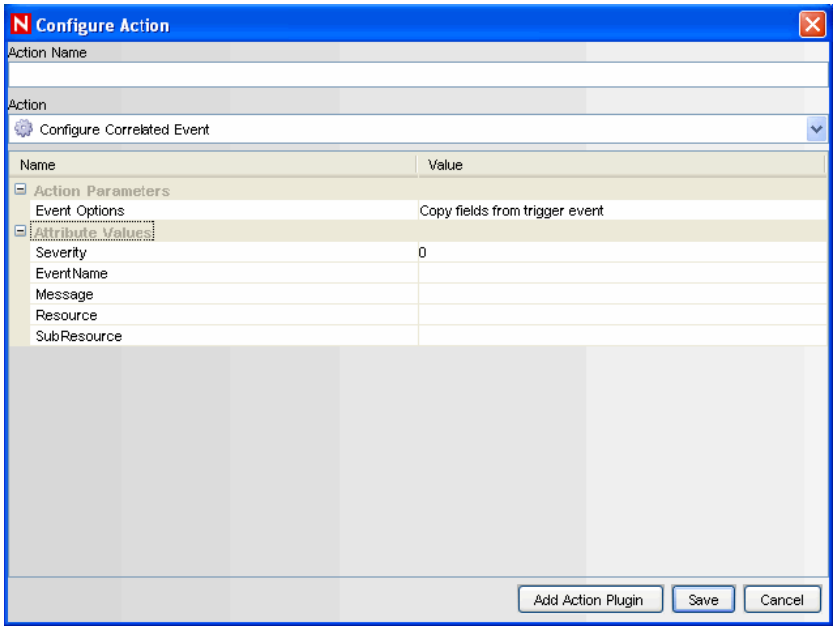

NOTE: This type of action can only be used in Correlation deployments.

To override the default values for the correlated event created when a rule fires, an action can be created to populate the following fields in the correlated event:

- Severity
- Event Name
- Message
- Resource
- SubResource

4.6.2 Adding to a Dynamic List

Figure 4-3 *Adding to a Dynamic List*

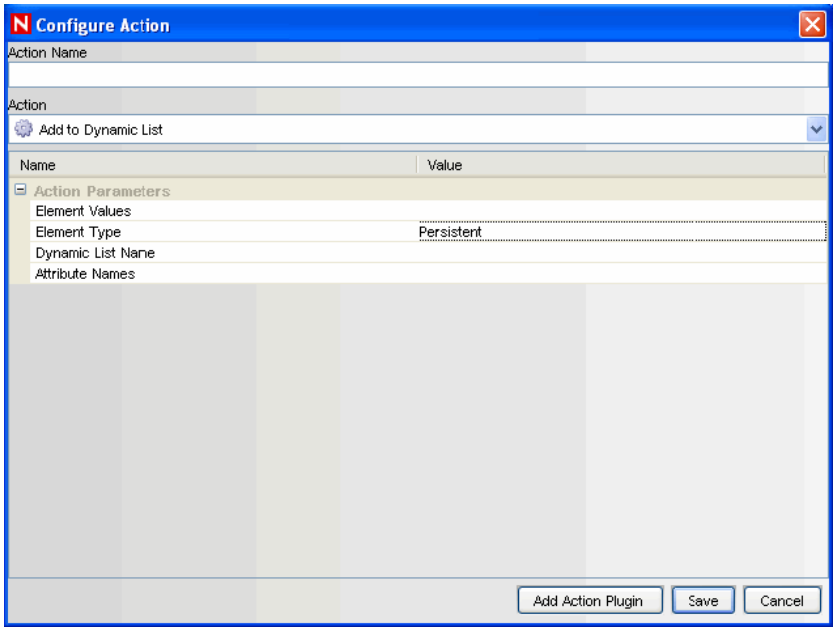

NOTE: This type of action can only be used in Correlation deployments.

This action type can be used to add a constant value or the value of an event attribute (such as Target IP or Initiator User Name) to an existing dynamic list. Any values that are repeated across multiple events are only added to the dynamic list once. The various parameters available are:

Table 4-5 *Parameters*

Option	Function
Element Values	(Optional) Specify a constant value to add to the dynamic list. If this is blank, *Attribute Name* must be populated.
Element Type	Persistent or Transient.
Dynamic List Name	Select an existing dynamic list from the drop-down menu.
Attribute Names	(Optional) For every event that is part of a correlated event, the value or values of the selected event attribute are added to the dynamic list. If this is blank, element values must be populated.

If there are entries for both *Element Values* and *Attribute Names*, both are added to the dynamic list when the rule fires. If the Element Value is filled in and the Element Type is Transient, the time stamp for the element in the dynamic list is updated each time the rule fires.

4.6.3 Removing a Value from a Dynamic List

Figure 4-4 *Removing a Value from a Dynamic List*

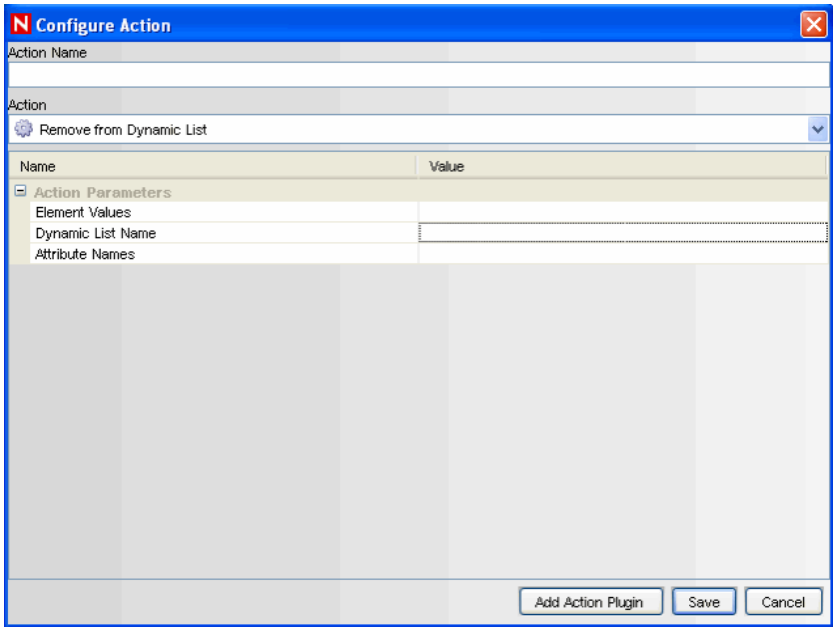

NOTE: This type of action can only be used in Correlation deployments

This action type can be used to add a constant value or the value of an event attribute (such as Target IP or Initiator User Name) from an existing dynamic list. The various parameters available are:

Table 4-6 *Parameters*

Option	Function
Element Values	Specify a constant value to remove from the list.
Dynamic List Name	Select an existing dynamic list from the drop-down menu.
Attribute Names	For every event that is part of a correlated event, the value or values of the selected event attribute are deleted from the dynamic list.

4.6.4 Executing a Command

Figure 4-5 *Executing a Command*

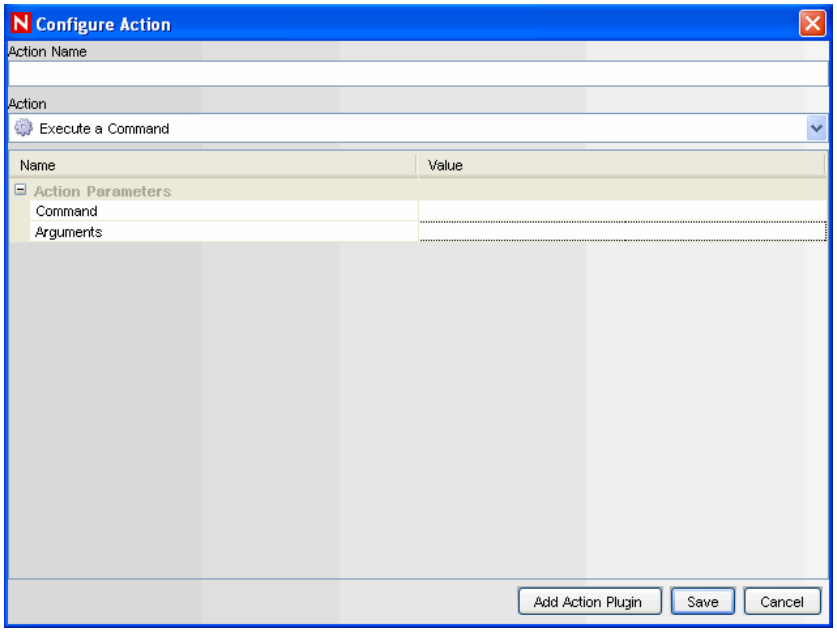

NOTE: This type of action can only be used in Correlation deployments

---

This action type can be used to execute a command when a correlated event triggers. You can set the following parameters:

Command: Arguments: This can include constants or references to an event attribute in the last event, the one that caused the rule to fire.

References to event attributes must use the values in the meta tag column enclosed in % or $ symbols. For example, %InitIP% represents the initiator IP address value from the Correlated event, except in the Configure Correlated Event action. Because the Correlated event was not created before the action is executed, the InitIP value comes from the trigger event. $InitIP$ always represents the value from the current event. Both %all% and all are the same, and they pass information (a limited set of attributes from both the trigger event and the Correlated event along with some Correlation rule data) to a Correlation action. They are provided primarily for backward compatibility with existing Correlation actions. They cannot be used in JavaScript actions or in the Configure Correlated Event action. For more information on meta tags, see "Sentinel 6.1 Rapid Deployment Event Fields" in the *Sentinel Rapid Deployment Reference Guide*.

Command actions can be created to perform a non-interactive action, such as modifying a firewall policy, entering a record in a database, or deactivating a user account. For an action that generates output, such as a command to run a vulnerability scan, the command should refer to a script that runs the command and then writes the output to a file.

NOTE: By default, the action output is stored to the working directory, `<install_directory>/data`. The action output can be written to a different directory by specifying a different storage location for the output file in the script

---

4.6.5 Creating an Incident

Figure 4-6 *Configure Action:Create Incident*

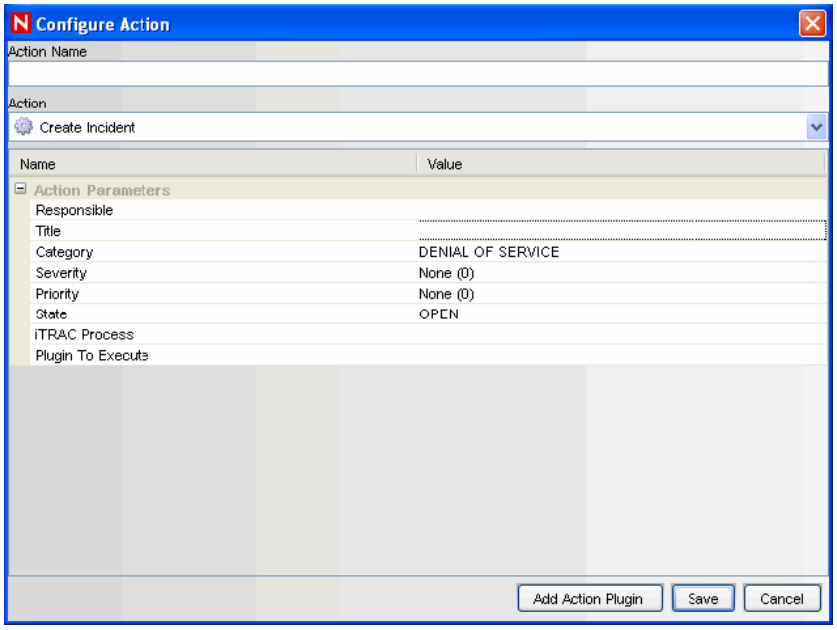

NOTE: This type of action can only be used in Correlation deployments.

This action type create an incident whenever a correlated event fires. You can also initiate an iTRAC workflow process for remediation of that incident. For more information about the values of the following parameters, see Chapter 5, “Incidents Tab,” on page 109.

- Responsible
- Title
- Category
- Severity
- Priority
- State
- (Optional) iTRAC Process list for configured iTRAC processes
- (Optional) Action Plugin to Execute list for configured JavaScript actions

---

IMPORTANT: Do not enable the Create Incident action until the correlation rule has been tuned. If the rule fires frequently, the system can create more incidents or initiate more iTRAC workflow processes than desired.

4.6.6 Sending an E-mail

Figure 4-7 *Configure Action: Send Email*

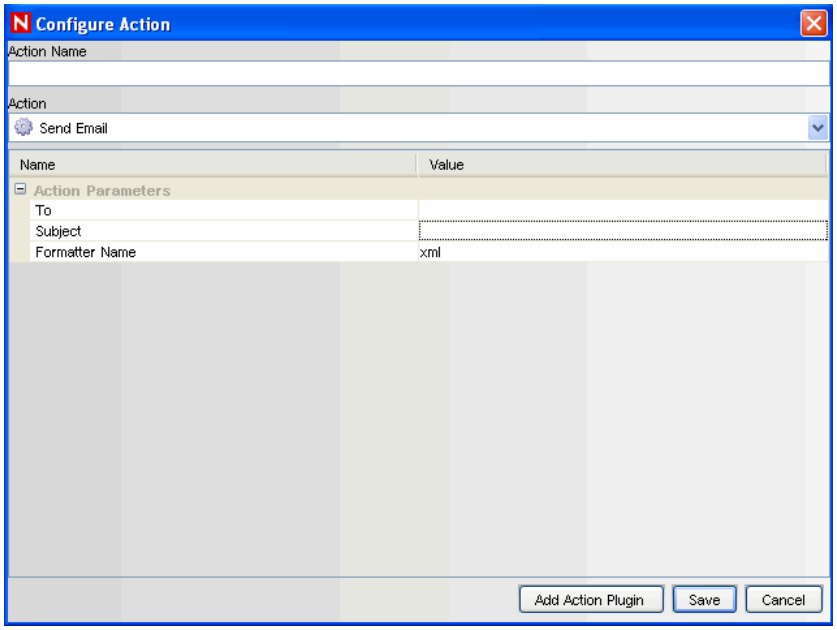

NOTE: This type of action can only be used in Correlation deployments

This action type can be used to send an e-mail when a correlated event triggers. The various parameters available are:

Table 4-7 *Parameters*

Option	Function
To	Specify the recipient e-mail address
Subject	Specify the subject of the message
Formatter Name	The format of the e-mail contains the correlated event formatted as “xml” or “Name Value Pair”, depending on what you select

4.6.7 Imported JavaScript Action Plugins

For information on JavaScript actions and how to debug them, see Section 17.1, “Action Manager,” on page 363. The JavaScript actions can be used in many places throughout the Sentinel interface.

Incidents Tab

5

In Sentinel, a set of related events (for example, a possible attack) can be grouped together to form an incident. An incident in the Open state alerts you to investigate, resolve, and close the incident. For example, the resolution to an attack might be to close a port, block a source IP, or rebuild a machine.

- Section 5.1, “Understanding an Incident,” on page 109
- Section 5.2, “Introduction to User Interface,” on page 109
- Section 5.3, “Manage Incident Views,” on page 111
- Section 5.4, “Manage Incidents,” on page 115
- Section 5.5, “Switch between Existing Incident Views,” on page 121

5.1 Understanding an Incident

Incidents can be created:

- Manually, by a security analyst monitoring incoming data or querying past data.
- Automatically, as a result of a correlation rule being triggered. For more information, see Chapter 4, “Correlation Tab,” on page 83.

In the *Incidents* tab, you can:

- Manage incident views
- Manage incidents
- Switch between existing incident views

NOTE: You need to have appropriate permissions to access this tab. Only an Administrator has controls to enable/disable access to the features of incidents for a user.

5.2 Introduction to User Interface

In the *Incidents* tab, you see the *Display Incident View*, *Create Incident*, and *Attachment Viewer Configuration*.

You can navigate to these functions from different places:

Table 5-1 *Table 4-1: Incident Tab User Interface*

User Interface	Description
	The *Incident* menu in the menu bar

User Interface	Description
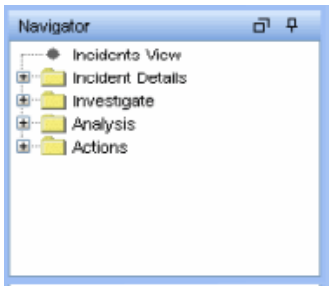	The Navigation Tree in the Navigation pane
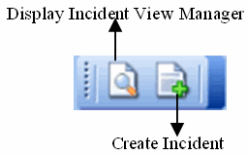	The toolbar buttons

5.2.1 Incident View

In the Incident View Manager, you can view the list of incidents and the parameters you specified when adding an incident.

To open the Incident View Manager:

1. Click *Incidents* on the menu bar and select Display Incident Views or click the *Display Incident View* button in the toolbar

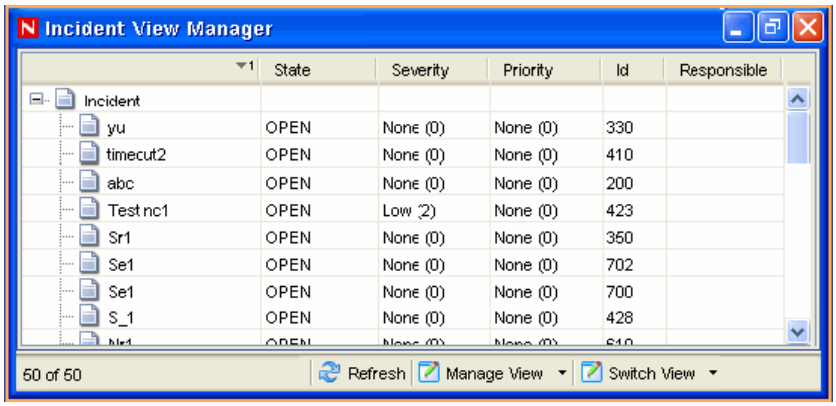

5.2.2 Incident

When you add or edit an incident, you see the tabs listed where you can perform the incident-related activities. As you investigate and remediate an incident, additional information can be added to these tabs. Except for *Events and History*, entering information on the tabs is optional.

Double click on the incident name or right-click and select modify to add/edit an incident.

Figure 5-1 *Add/Edit Incident*

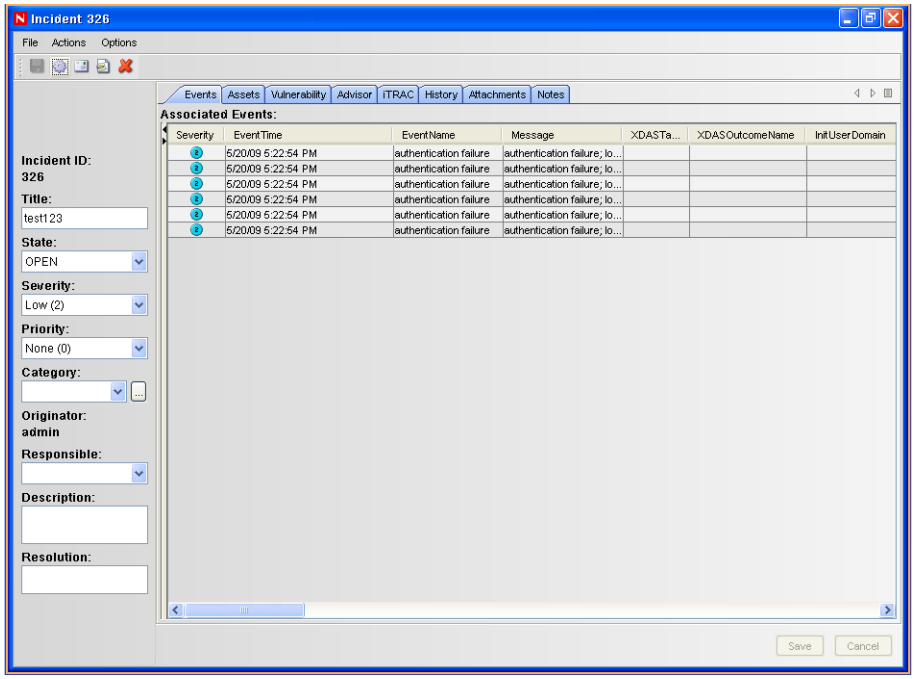

- **Events:** Lists events attached to this incident. You can attach events to incidents in an Active View.
- **Assets:** Lists assets affected by the events of this incident.
- **Vulnerability:** Lists asset vulnerabilities.
- **Advisor:** Displays asset attack and alert information.
- **iTRAC:** Allows you to add a workflow to an incident from the iTRAC tab.
- **History:** Lists activities performed on the current incident.
- **Attachments:** Allows you to add an attachment to the incident created in the system.
- **Notes:** Allows you to add notes to the incident.

5.3 Manage Incident Views

Manage View allows you to:

- Add views
- Edit views
- Delete views
- Mark a view as the default

5.3.1 Adding a View

1. Click *Incidents* > *Display Incident View Manager*. Alternatively, click the Display Incident View button on the toolbar.

2 Open the view options by doing one of the following:

- Click the down-arrow on the *Manage Views* button located in bottom right corner of the window and select *Add View*.
- Click the down-arrow on the *Manage Views* button located in the bottom right corner of the window, select *Manage Views* and then click the *Add View* button.

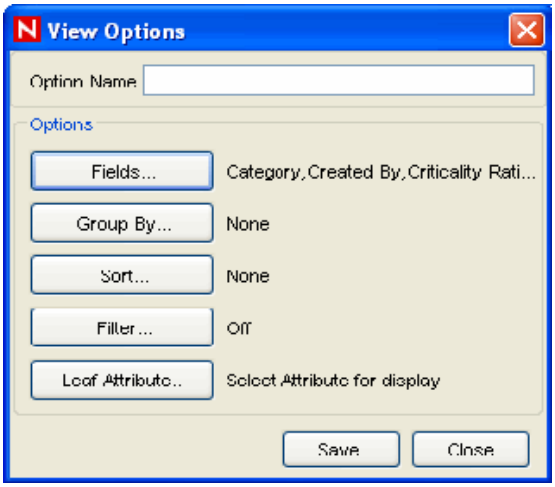

3 Provide a name in the *Option Name* field. Click the buttons listed below to specify the options.

- **Fields:** The variables of the events attached to incidents are displayed as fields. By default, all the fields are arranged as columns in the Incident View. In the Field Options window, you can add or remove columns that display and arrange the order of the columns by using the up-arrow and down-arrow.

- **Group By:** Set rules to group incidents in the display view.

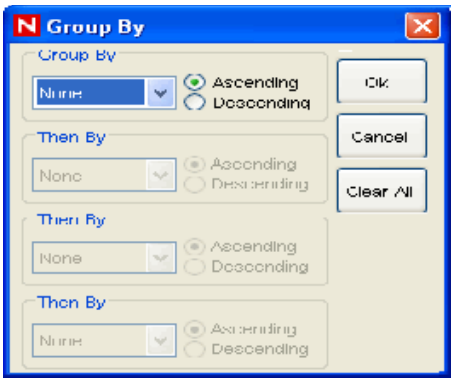

- **Sort By:** Set rules to sort the incidents in the display view.

- **Filter:** Set incident filters. Only the incidents that match your filter display in the view.

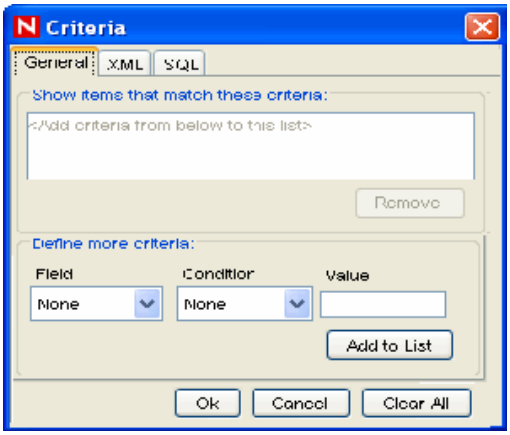

- **Leaf Attribute:** Select an attribute from the list that is displayed as the first column in the incident view.

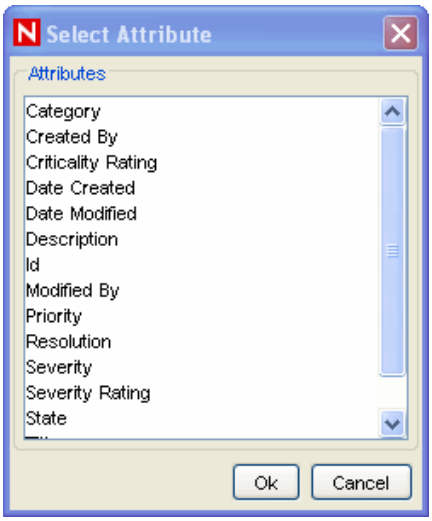

4 Click *Save*.

5.3.2 Modifying a View

1 Click *Incidents* > *Display Incident View* or click the *Display Incident View Manager* button on the toolbar.

2 Open a view by doing one of the following:

- Click the down-arrow next to the *Switch View* button in the bottom right corner, then select the view you want to edit. Click the down-arrow next to the *Manage View* button located in bottom right corner of the screen and select *Edit Current View* from the list.
- Click the down-arrow next to the *Manage Views* button located in the bottom right corner of the window, then select *Manage Views*. Select a view to edit and click *View/Edit*.

3 Edit the options as required and click *Save*.

5.3.3 Deleting a View

1. Click *Incidents* > *Incident View Manager* or click the Display Incident View button on the toolbar.
2. Click the down-arrow next to the *Manage Views* button located in bottom right corner of the screen and select *Manage View* from the list. The Manage View window displays. Select a view and click *Delete*. A confirmation message alert displays.
3. Click *Yes* to delete.

5.3.4 Default View

To mark a view as the default:

1. Click *Incidents* > *Display Incident View Manager*, or click the *Display Incident View Manager* icon on the toolbar.
2. Click the down-arrow next to the *Manage Views* button located in bottom right corner of the screen and select *Manage Views* from the list. The Incident View window displays.
3. Select the incident view you want as the default, and click *Mark as Default*.

5.4 Manage Incidents

You can perform the following activities related to incidents:

- Create an incident
- Attach workflows to incidents
- Add notes to incidents
- Add attachments to incidents
- Execute an incident action
- E-mail an incident
- Edit an incident
- Delete an incident

5.4.1 Creating Incidents

1. Click *Incidents* > *Create Incident*, or click the *Create Incident* button on the toolbar. The New Incident window displays.

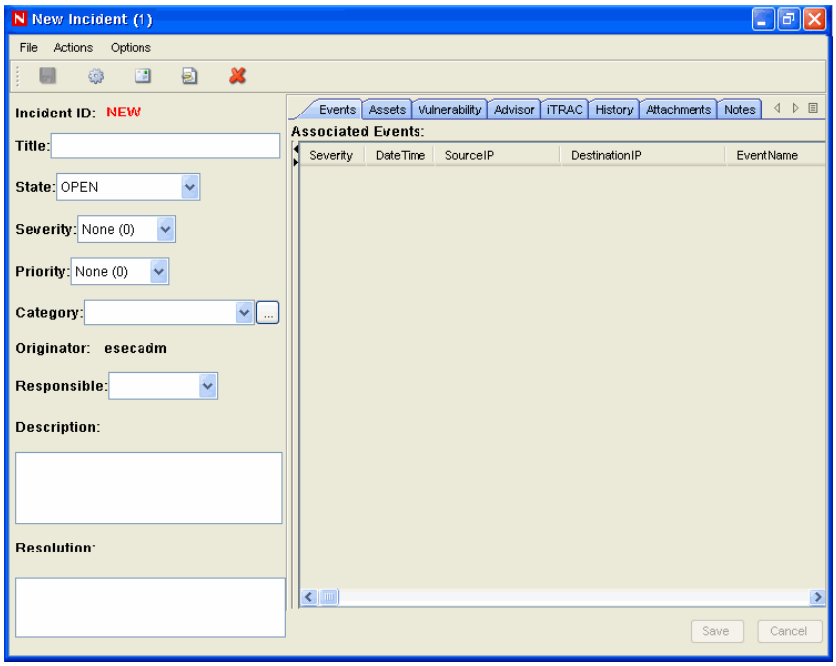

2 Specify the following information:

- **Title:** Specify the title of the incident.
- **State:** To set state of the incident, select from the drop-down list.
- **Severity:** To indicate the severity of the incident, select from the drop-down list.
- **Priority:** To indicate the priority of the incident, select from the drop-down list.
- **Category:** Specify the category of the incident.
- **Responsible:** To assign the responsibility to investigate and close the incident, select from the drop-down list.
- **Description:** Specify the description of the incident in the text area.
- **Resolution:** Specify the resolution description in the text area.

3 Click *Create*.The incident ID automatically generates after you click *Create*.

For more information on creating an incident and grouping events, see Section 3.7, “Creating Incidents,” on page 63.

5.4.2 Viewing an Incident

1 Click *Incidents* > *Display Incident View Manager* or click the *Display Incident View Manager* button on the toolbar.

2 Open an incident by doing one of the following:

- Selecting a view from the *Switch Views* button in the bottom right corner.
- Double-click an incident in the Incident View Manager window.

5.4.3 Attaching Workflows to Incidents

1 Open an incident.

2 In the Incident window, click the *iTRAC* tab.

3 Select an iTRAC process from the drop-down list.

4 Click *Save*.

NOTE: You can attach only one process to an incident.

5.4.4 Adding Notes to Incidents

1 In the Incident window, click the *Notes* tab.

2 Click *Add*. The Add Notes to Incident window displays.

3 Provide your notes and click *OK*.

4 Click *Save*.

NOTE: To edit or delete the note, right-click a note in the *Notes* tab of the Incident window and select *Edit* or *Delete*.

5.4.5 Adding Attachments to Incidents

1 In the Incident window, click the *Attachments* tab.

2 Click *Add*.The Add Attachment to Incident window displays.

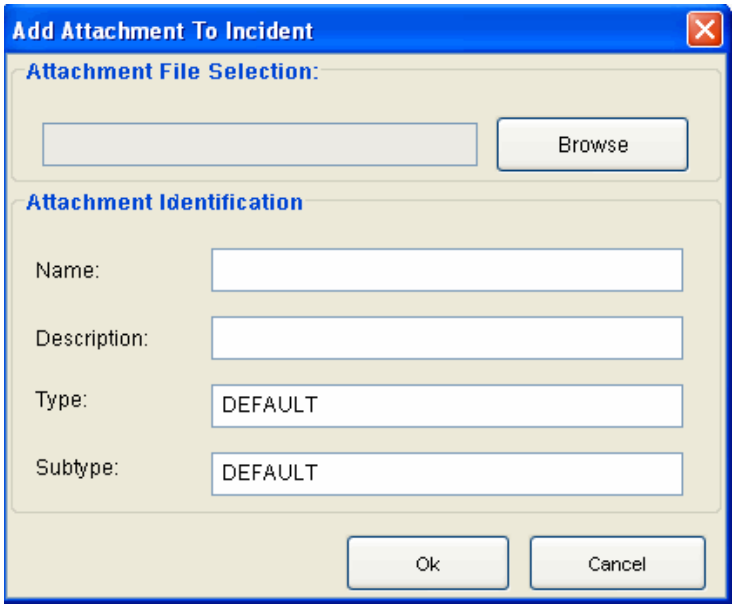

3 Click *Browse*, navigate to the attachment, and select it.

4 Provide the following information, or accept the default entries:

- Name
- Description
- Type
- Subtype

5 Click *OK*, then click Save.

Right-click the attachment to view or save.

5.4.6 Executing Incident Actions

Any configured JavasScript action or iTRAC activity can be executed on an incident.

1 Open an incident.

2 Click Actions > Execute Incident Action or click Execute Incident Action icon.

The Execute Incident Action window displays.

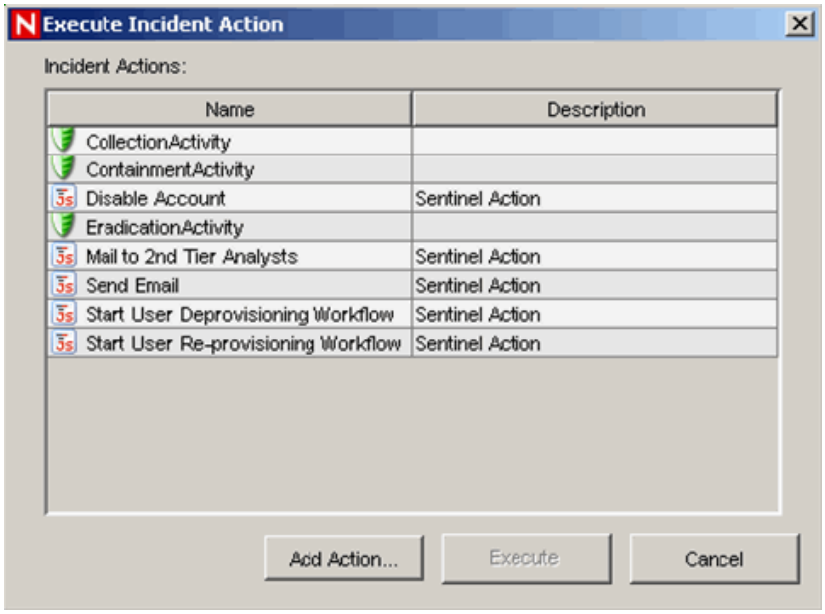

3 Select an action or click *Add Action* to create a new one.

4 Click *Execute*. If the action is a JavaScript action, a window opens to show the progress of the action.

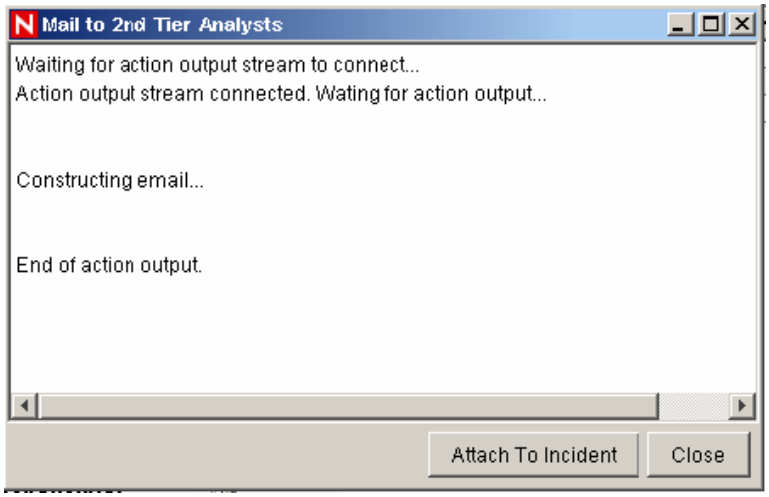

5 To add the command output to the incident, click *Attach to Incident*.

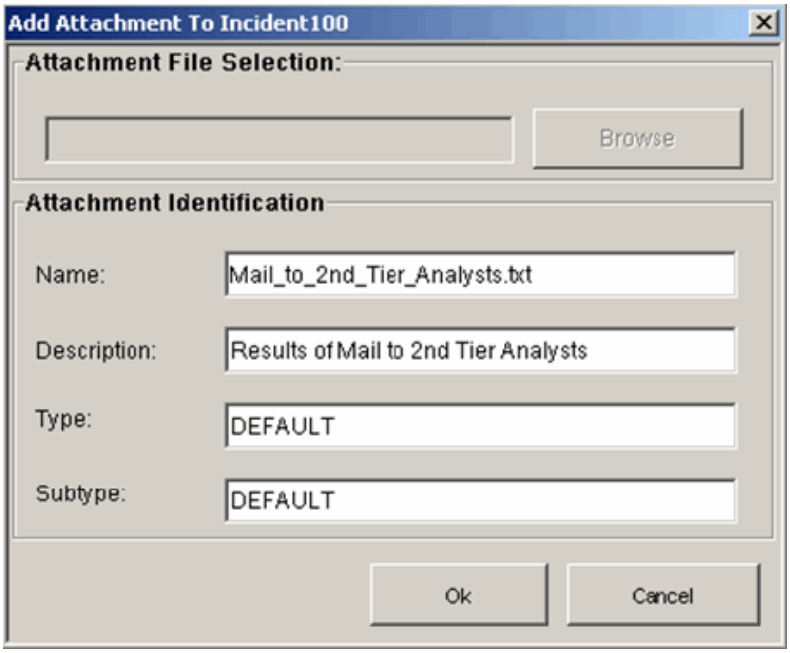

The action output is saved and can be viewed from the *Attachments* tab of the incident.

5.4.7 E-Mailing an Incident

To mail an incident by using the preinstalled Email Incident action, you must have an SMTP Integrator configured with valid connection information and with the SentinelDefaultEMailServer property set to “true”. For more information, see the SMTP Integrator documentation available at the Novell Sentinel Content Web site (http://www.novell.com/documentation/sentinel61).

1 Open an incident.

2 Click the *Email Incident* button to display the Email Incident window.

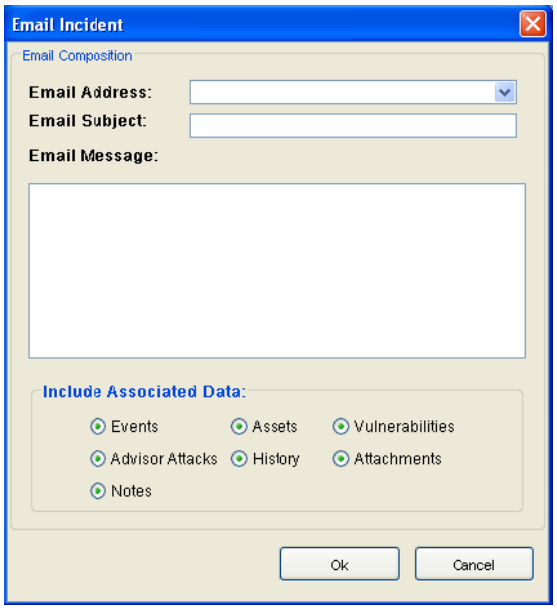

3 Provide the following:

- Email Address
- Email Subject
- Email Message

4 Select which HTML attachments should be included in the mail message, such as the events included in the incident, assets, vulnerabilities, Advisor attacks, incident history, attachments, and notes.

5 Click *OK*.

5.4.8 Modifying Incidents

1 Click the *Incident* tab, then click *Incidents* > *Display Incident View*. Alternatively, click the *Display Incident View* button on the toolbar. The Incident View window displays with the list of incidents.

2 Right-click the incident you want to edit and select *Modify*.

3 The Incident window displays. Edit the following information:

- Title
- State
- Severity
- Priority
- Category
- Responsible
- Description
- Resolution

4 Click *Save*.

Save button is active only if you modify any information in the Incidents window.

5.4.9 Deleting Incidents

1. Click the *Incident* tab, then.click *Incidents* > *Display Incident View Manager*, or click the *Display Incident View* button on the toolbar. The Incident View window displays.
2. Right-click the incident you want to delete and select *Delete*.
3. A confirmation Message displays. Select *Yes*.

5.5 Switch between Existing Incident Views

1. Click the down-arrow on the *Switch View* button on the bottom right corner of the screen to display a list of existing views.
2. Select a view.

iTRAC Workflows

6

The iTRAC workflows are designed to provide a simple, flexible solution for automating and tracking an enterprise's incident response processes. iTRAC leverages the Sentinel internal incident system to track security or system problems from identification (through correlation rules or manual identification) through resolution.

- Section 6.1, "Understanding iTRAC Workflows," on page 123
- Section 6.2, "Introduction to the User Interface," on page 124
- Section 6.3, "Template Manager," on page 125
- Section 6.4, "Template Builder Interface," on page 126
- Section 6.5, "Steps," on page 129
- Section 6.6, "Transitions," on page 141
- Section 6.7, "Activities," on page 149
- Section 6.8, "Process Management," on page 155

6.1 Understanding iTRAC Workflows

Workflows can be built using manual and automated steps. Advanced features such as branching, time-based escalation, and local variables are supported. Integration with external scripts and plug-ins allows for flexible interaction with third-party systems. Comprehensive reporting allows administrators to understand and fine-tune the incident response processes.

NOTE: Access to manage iTRAC templates, activities, and processes can be enabled on a user-by-user basis by any user with the ability to change user permissions.

The iTRAC system uses three Sentinel objects that can be defined outside the iTRAC framework:

Table 6-1 *Sentinel Objects Used by iTRAC*

Incident	Incidents within Sentinel are groups of events that represent an actionable security incident, plus associated state and meta-information. Incidents are created manually or through Correlation rules, and can be associated with a workflow process. They can be viewed on the *Incidents* tab.
Activity	An Activity is a predefined automatic unit of work, with defined inputs, command-driven activity, and outputs (for example, automatically attaching asset data to the incident or sending an e-mail). Activities can be included in a workflow template and executed during workflow processes, or they can be executed within an incident.
Role	Sentinel users can be assigned to one or more roles. Manual steps in the workflow processes can be assigned to a role.

iTRAC Workflows have four major components that are unique to iTRAC:

Table 6-2 Major Components of iTRAC

Step	A step is an individual unit of work within a workflow; there are manual steps, decision steps, command steps, mail steps, and activity-based steps. Each step displays as an icon within a given workflow template.
Transition	A transition defines how the workflow moves from one state (activity) to another. This can be determined by an analyst action, by the value of a variable, or by the amount of time elapsed.
Templates	A template is a design for a workflow that controls the flow of execution of a process in iTRAC. The template consists of a network of manual and automated steps. Activities and criteria for transition between them. Workflow templates define how an incident is responded to after a process based on that template is instantiated (see below). A template can be associated with many incidents.
Processes	A process is a specific instance of a workflow template that is actively being tracked by the workflow system. It includes all the relevant information relating to the instance, including the current step in the workflow, the associated incident, the results of steps, attachments, and notes. Each workflow process is associated to one and only one incident.

6.2 Introduction to the User Interface

Within the Sentinel Control Center, you can access the iTRAC administrative functions by selecting the *iTRAC* tab from the main screen. This tab gives you access to the Activity Manager (where you define activities), the Template Manager (where you define templates), and the Process View Manager (where you manage instantiated workflow processes).

You can navigate to these functions from:

Table 6-3 iTRAC User Interface

User Interface	Description
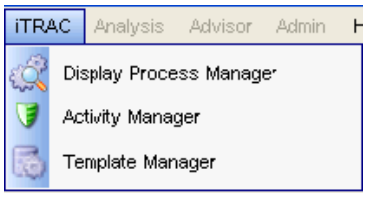	The iTRAC menu in the menu bar

User Interface	Description
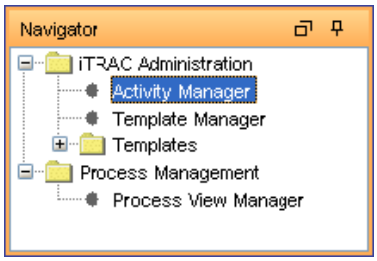	The Navigation Tree in the Navigation pane
	The toolbar buttons

6.3 Template Manager

The Template Manager can be used to create, view, modify, copy, or delete a template. Within the Template Manager you can add, delete, copy, view, and edit templates. Templates can be sorted into folders for easy management

In the Template Manager, you can:

- Create new workflow templates
- Edit or copy existing templates
- Define workflow steps
  - Manual or Automated
  - Description of step or instructions for iTRAC users
- Define transitions between steps
  - Transition type
  - Escalation procedures
  - Timeout and alert attributes

Figure 6-1 *iTRAC Workflow*

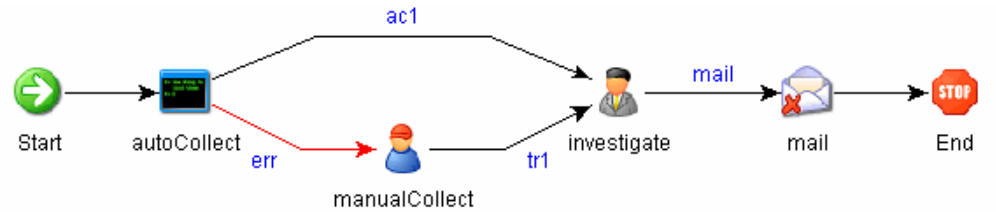

6.3.1 Default Templates

iTRAC is shipped with the following templates to use as examples. The process and activity attributes for these templates are set to predefined values. Users can modify these to suit their requirements. The default templates are:

- AlertTimeoutExample
- TwoStepSimpleExample

- ConditionalTransitionExample
- CommandExample

6.4 Template Builder Interface

Figure 6-2 *Template Builder Interface*

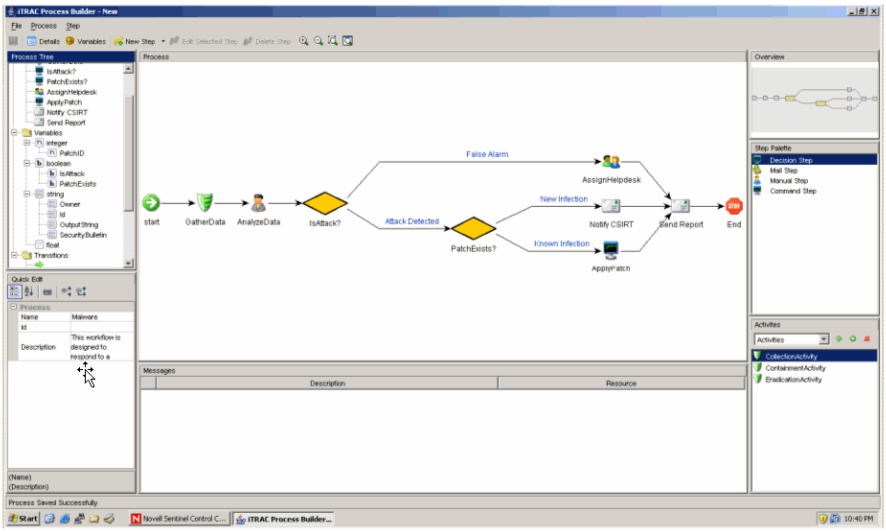

The following panes display in the Template Builder window:

- **Process Tree:** This pane displays the steps, transitions and variables added to the template. Users can add steps or variables, and edit or remove steps, variables and transitions.

  To perform an action on a step, variable or transition:

  - Expand the relevant group in the Tree.
  - Select and right-click an existing attribute.
  - Select action you want to perform.
- **Process:** This is the main GUI for viewing and creating a workflow template. For more information on creating a workflow template, see "Section 6.4.1, "Creating Templates," on page 127".
- **Quick Edit:** Select a step or transition to see its properties. This pane allows you to edit process attributes.

  To edit the details of steps using Quick Edit:

  - Click the Process Attribute value in the Quick Edit pane.
  - The attribute values are selected, indicating Edit Mode.
  - Modify the value and click anywhere outside the Quick Edit frame to save the new value.
- **Messages:** This pane displays messages if steps or transitions are incomplete. You must resolve any issues listed here before saving the template.
- **Overview:** This pane displays an overview of the entire template.

- **Step Palette:** There are four types of steps in the Step Palette. You can drag and drop the steps into the Process pane.
  - Decision Step
  - Mail Step
  - Manual Step
  - Command Step
- **Activities:** The activities added in the Activity Manager are shown in this pane and can be added to a workflow template. The user can also add, edit and remove activities. For more information, see Section 6.7.6, "Managing Activities," on page 154.

WARNING: Use caution when editing or deleting an activity that is already in use.

The following icons are used in the Template Builder to represent the steps:

Table 6-4 *Template Builder Icons*

Icon	Description
	Start Step: All workflow templates have a Start step.
	Decision Step: This step provides different execution paths, depending on the value of a variable defined in a previous step.
	Mail Step: This step sends a prewritten e-mail.
	Manual Step: This step indicates that manual work must be performed, often outside the Sentinel system (For example, telephoning the owner of the affected system or analyzing the results of a scan).
	Activity Step: This step is a predefined set of activities.
	Command Step: This step executes a command or script on the iTRAC workflow server, usually installed in the same place as the Data Access Service (DAS). The output of the command can be stored in a string variable and used as input to a decision step.
	End Step: This step signifies the completion of a workflow process.

6.4.1 Creating Templates

1. Click the *iTRAC* tab.
2. In the navigation pane, click *iTRAC Administration* > *Template Manager*.
3. Click *Add*. The iTRAC Template Builder window displays.

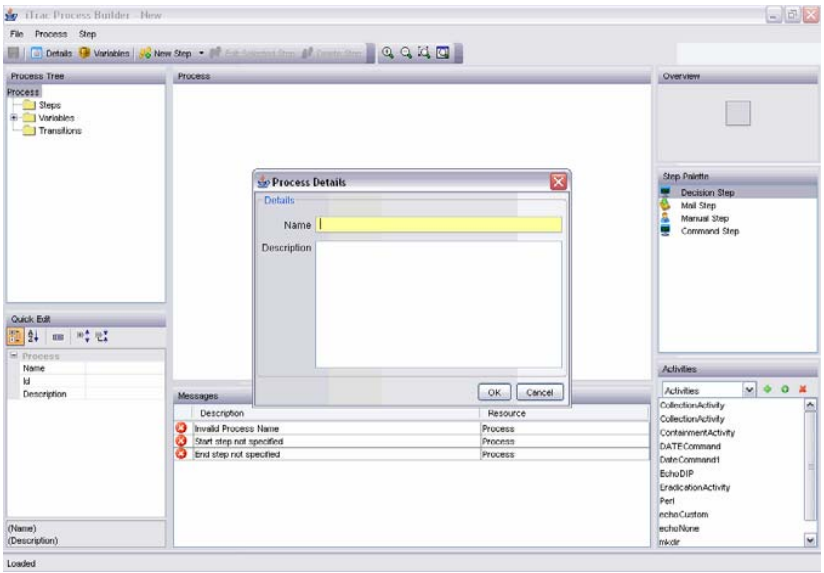

4 In the Process Details window, provide a name and description (optional) of the template and click *OK*.

5 Do one of the following:

- Drag and drop a step from the Step Palette or an activity from the Activities pane into the Process window.
- Click the New Step drop-down button in the upper left corner and select one of the following step types. Right-click *Start* step, select *Insert New* and select one of the following step types.
  - Decision Step
  - Mail Step
  - Manual Step
  - Command Step

6 Add as many steps and activities as needed to create the template.

7 Create transitions between each step. To create transitions, right-click the step after which you need to add transition and click *Add Transition*.

Any step (except for the End step) might have one or more exit transition lines. A decision step must have at least two exit lines.

8 Right-click each final step in the template and click *Add End Transition*.

On the bottom of the iTRAC Template Builder is a message pane that lists any warnings or errors about incomplete steps during the construction.

9 To save your process, go to *File>Save* or click *Save* button.

6.4.2 Managing Templates

After creating a template, you can modify, copy, or delete it.

- "Viewing/Editing Templates" on page 129

- “Copying Templates” on page 129
- “Deleting Templates” on page 129

Viewing/Editing Templates

1. In the Navigator, click *iTRAC Administration* > *Template Manager*.
2. Select a template and click *View/Edit*. The Template Builder displays.

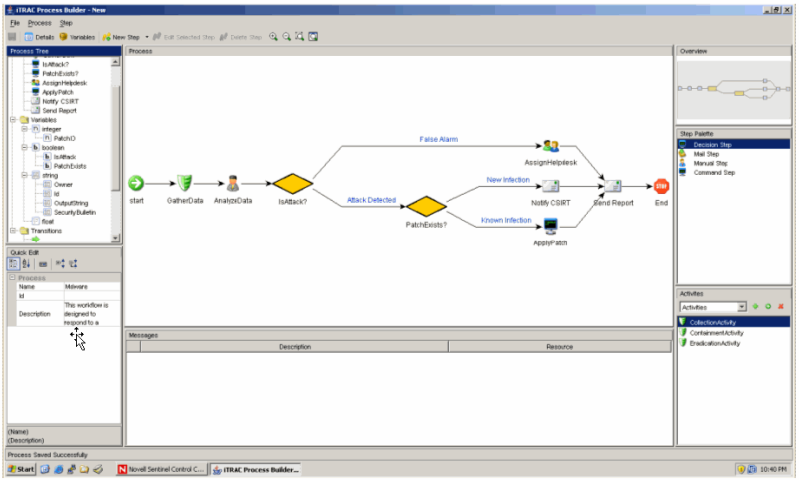

Copying Templates

One way to create a new workflow template is to copy one of the default templates and modify it.

1. Click the *iTRAC* tab.
2. In the Navigator, click *iTRAC Administration* > *Template Manager*.
3. Select a template and click *Copy*. A Template Builder with the copied template displays.
4. Provide a new name, then save and edit the template as needed.

Deleting Templates

Even if you delete a template, any instantiated workflow processes that are based on that template still completes normally.

1. Click the *iTRAC* tab.
2. In the Navigator, click *iTRAC Administration* > *Template Manager*.
3. Select a template and click *Delete*.

6.5 Steps

Steps are the basic components of a template. Every template must have a Start step and an End step. The Start step exists by default. You can also add the following types of steps to a template:

- Section 6.5.1, “Start Step,” on page 130
- Section 6.5.2, “Manual Step,” on page 130

- Section 6.5.3, “Decision Step,” on page 134
- Section 6.5.4, “Mail Step,” on page 134
- Section 6.5.5, “Command Step,” on page 134
- Section 6.5.6, “Activity Step,” on page 135
- Section 6.5.7, “End Step,” on page 136
- Section 6.5.8, “Adding Steps to a Workflow,” on page 136
- Section 6.5.9, “Managing Steps,” on page 136

6.5.1 Start Step

Every workflow template must have one and only one Start step. The transition from a Start step is always Unconditional.

6.5.2 Manual Step

This type of step indicates that manual work must be performed. Every manual step in a template must be assigned to a role. The users in that role are notified through a worklist item when an instantiated workflow process reaches the manual step. When a user accepts the worklist item, it is removed from the queue of the other users in that role. For more information about worklists and stepping through a workflow process, see Section 7.1, “Work Item Summary,” on page 161. section.

The description of the step should indicate what work needs to be performed. The user is expected to perform that work and then acknowledge completion.

A manual step includes the following attributes:

- Name of step
- Role
- Variables
  - Delete
  - Add
- Description

From a manual step, you can set Conditional, Unconditional, Timeout, or Alert transitions.

Variables

The user can also be asked to set one or more variables to appropriate values. Four variable types can be assigned to manual steps: Integer, Boolean, String, and Float. This variable can be set to an explicit default value during the step definition, or the user can set the value at run-time as part of the workflow process. The value can be optional or required.

The value of the variable can be used as part of a Conditional transition to determine the path the workflow follows. It can also be used later as part of a Conditional transition from a decision step to determine the workflow path.

NOTE: If the value is going to be used later as part of a decision step, it should be marked "Required."

For example, an integer variable can be set by the user to hold the event rate. Output transitions from the manual step can be defined so that if the event rate is greater than 500, one path is followed; otherwise, another path is followed.

To create a variable:

1. Click the *iTRAC* tab.
2. In the Navigator, click *iTRAC Administration* > *Template Manager*.
3. Click the *Add* button in upper left corner to open a new template or select an existing template, then click *View/Edit*.
4. Right-click *Variables* in the Process Tree and select the type of variable to add, or right-click the variable type and select *Add Variable*.

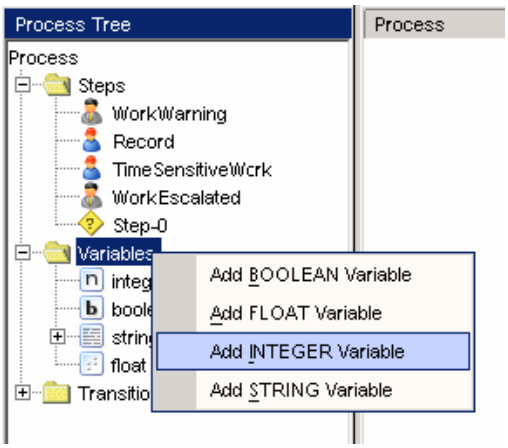

5. Give the variable a name and specify the default value, if desired.

   Boolean Variable:

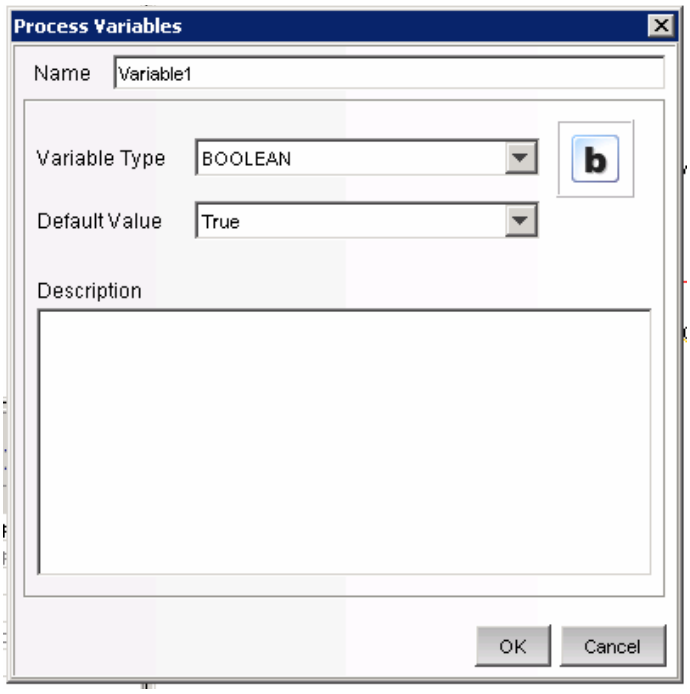

Integer Variable:

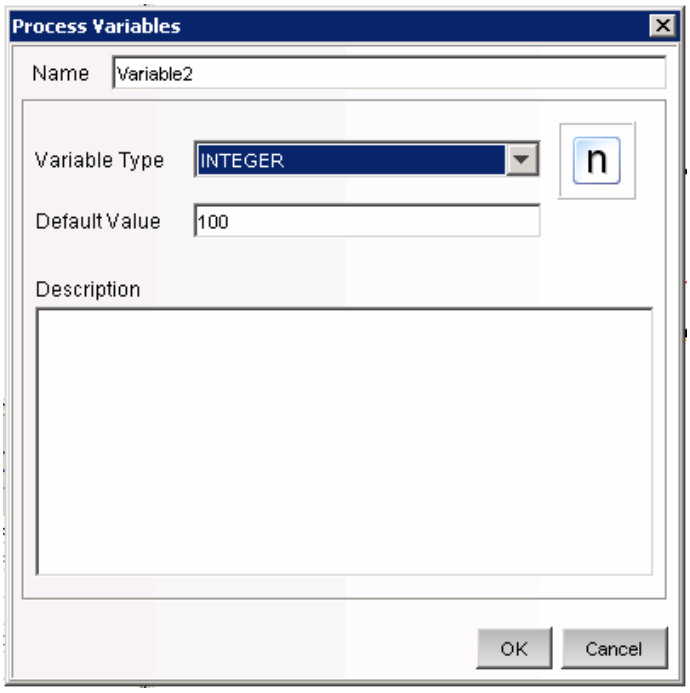

String Variable:

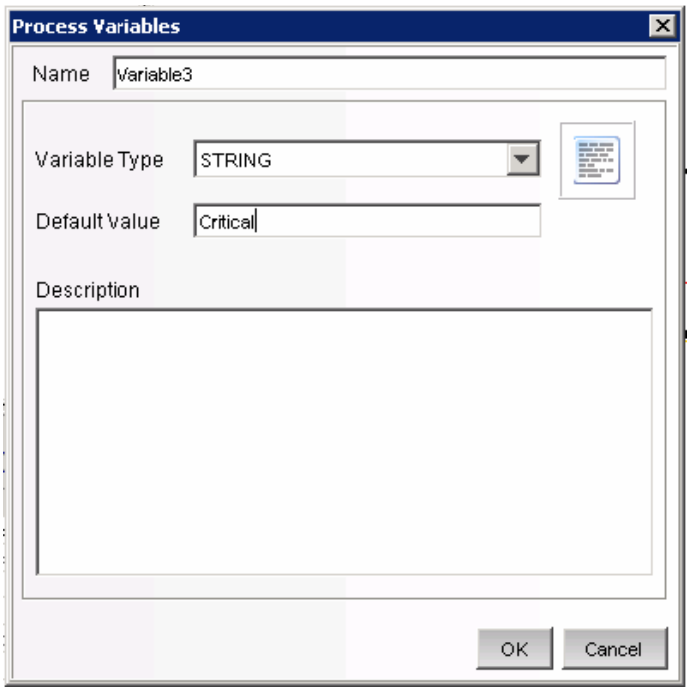

Float Variable:

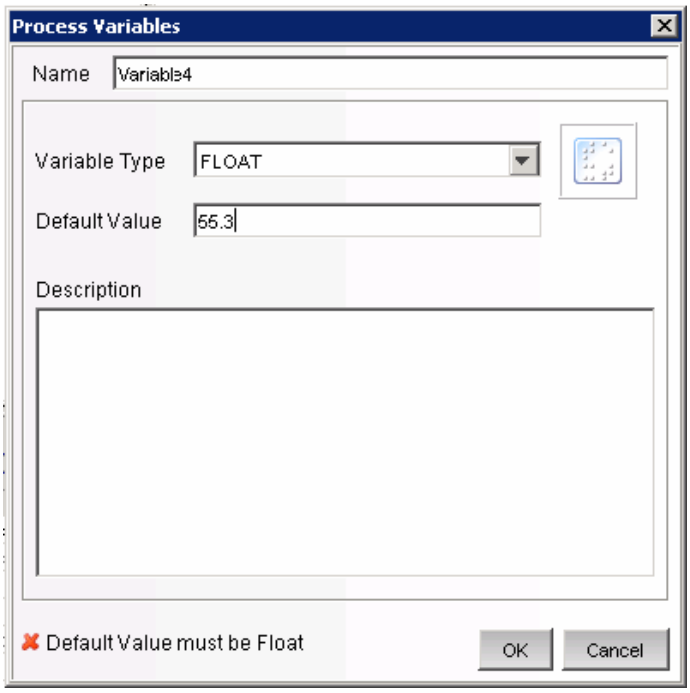

6 Click *OK*.

6.5.3 Decision Step

This type of step selects between exit transitions depending on the values of variables defined in prior steps. See Section 6.5.2, “Manual Step,” on page 130 for the available variable types. The decision step itself is very simple; you can edit only the step name and description. The workflow path is determined by the transitions.

From a decision step, you can set Conditional and Else transitions. Every decision step must have an Else transition and at least one Conditional transition. The Else transition leads to a workflow path that is followed if none of the criteria for the Conditional transitions is met.

6.5.4 Mail Step

This step sends a prewritten e-mail. A mail step includes the following attributes:

- Name of step
- To addressee
- From addressee
- Subject of email
- Body of email

From a mail step, you can set a Conditional, Unconditional, Timeout, Alert, or Error transition. An Error transition should always be included so error conditions can be handled properly.

NOTE: If the first step of a workflow fails without an error transition, the iTRAC process cannot proceed.

6.5.5 Command Step

A command step is a step in which an operating-system level command or script (shell, batch, Perl and so on) is executed. The name of the command can be provided explicitly or set as a string variable, and parameters can be passed in the same manner. Output from the command can also be placed back into a string variable.

A command step includes the following attributes:

- Name of step
- Description
- Command (Can be explicit or variable-driven)

- Arguments (Can be explicit or variable-driven)
- Output Variable

NOTE: The command must be stored in the `<install_directory>/config/exec` directory on the iTRAC workflow server. Symbolic links are not supported

---

Variables

The command output can also be used to set a variable to the appropriate values. Command steps must use String variable types.

The value of the variable can be used as part of a Conditional transition to determine the path the workflow follows. It can also be used later as part of a decision step to determine the workflow path.

For example, a Command step can return a value of 0 for failure and 1 for success. This output can be assigned to a variable, and then a Conditional transition or a decision step can use this value to determine which workflow path to take.

The command and its arguments can each be specified explicitly by the person designing the workflow or can be set as a string variable. If either one is set as a string variable, there must be a previous step in the template where the variable is set to a string value.

From a command step, you can set Conditional, Unconditional, Timeout, or Alert, or Error transitions. An Error transition should always be included so error conditions can be handled properly.

NOTE: If the first step of a workflow fails without an error transition, the iTRAC process cannot proceed.

6.5.6 Activity Step

An activity step is a type of automated step that can be used in a workflow template. The activity steps are created in the Activity Manager and can consist of internal Sentinel operations or external scripted operations. After activity steps are created, the user can select from the library of these activities and include them into in a workflow. For more information on creating each type of predefined activity, see Section 6.7.5, “Creating iTRAC Activities,” on page 151.

An activity step includes the following attributes:

- Name
- Description
- Activity Assignment

From an activity step, you can set Conditional, Unconditional, Timeout, or Alert, or Error transitions. An Error transition should always be included so error conditions can be handled properly.

NOTE: If the first step of a workflow fails without an error transition, the iTRAC process cannot proceed.

6.5.7 End Step

Every workflow template must have an End step to complete every branch of the workflow path.

6.5.8 Adding Steps to a Workflow

Steps can be added to a workflow by using the Step Palette or by using a right-click in the Process Builder. When you are adding steps to a workflow, a yellow entry field indicates an invalid entry.

To add a step from the Step Palette:

1. Drag and drop a step from the Step Palette.
2. Right-click the step and select Edit Step.
3. Edit the details of the step and click Save.

To add a step using a right-Click:

1. Right-click an existing step in the Process Builder and select *Insert New*.
2. Edit the details of the step and click *Save*.
3. Select *Manual*, *Decision*, *Mail*, *Command*, or *End Step*.
4. Edit the details of the step and click *Save*.

To add an Activity step:

1. Click and drag an activity from the Activity pane to the Process Builder.

To add an End step:

1. Right-click a step with no transition and select *Add End Transition*.

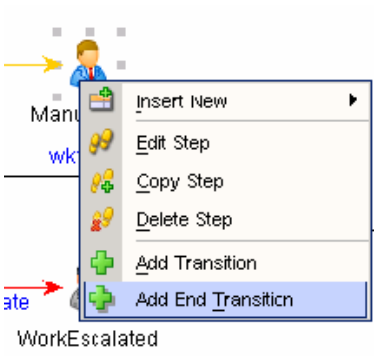

6.5.9 Managing Steps

Steps can be copied, edited, or deleted.

- "Copying Steps" on page 137

- "Modifying Steps" on page 137
- "Deleting Steps" on page 140

Copying Steps

1. Click the *iTRAC* tab.
2. In the Navigator, click *iTRAC Administration* > *Template Manager*.
3. Select an existing template, then click *View/Edit*.The iTRAC Process Builder window displays.
4. Select an existing step, right-click, and select *Copy Step*.
5. The Step window opens in edit mode with all the attributes of the selected step.
6. Specify a name for the new step.
7. Edit step attributes as required. Click *OK*.

Modifying Steps

To edit a step:

1. Click the *iTRAC* tab.
2. In the Navigator, click *iTRAC Administration* > *Template Manager*.
3. Select an existing template, then click *View/Edit*. The iTRAC Process Builder window displays.
4. Select an existing step, right-click, and select Edit Step.
5. Edit the step attributes. Click *OK*.

To edit a manual step:

1. Right-click a manual step and select *Edit Step*.

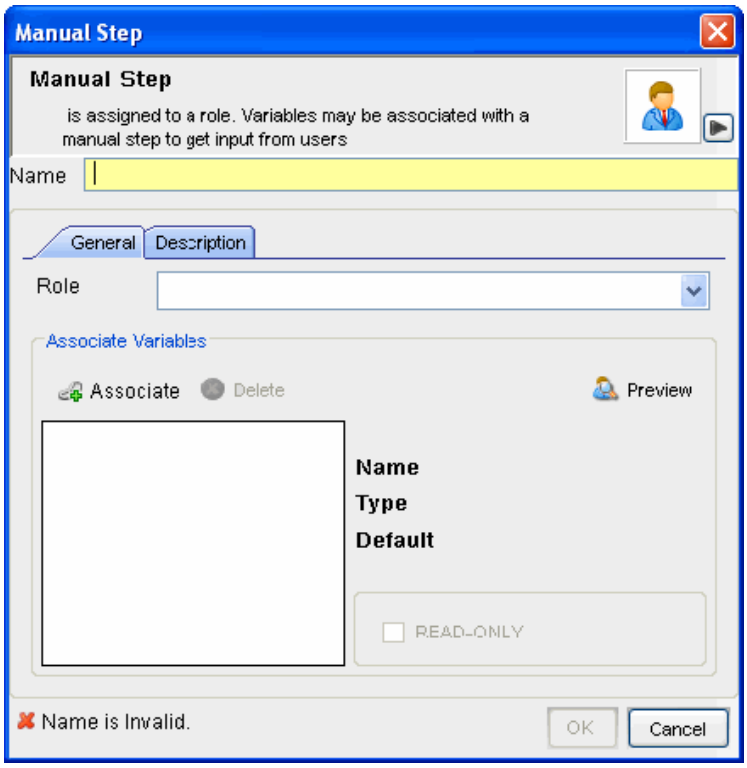

2 Provide a name for the step.

3 Attach a role to this step by selecting a role from the drop-down list. For more information on roles, see Chapter 12, “Administration,” on page 235.

4 Click *Associate* to associate a variable; select the variable from the list or create new variables to be associated. Set a default value as desired.

5 Select the *Read-Only* check box if this variable is to be forced to the default value.

6 Click the *Description* tab to provide description for this step.

7 Click *Preview* to preview the step you created.

8 Click *OK*.

To edit a decision step:

1 Right-click a decision step and select *Edit Step*.

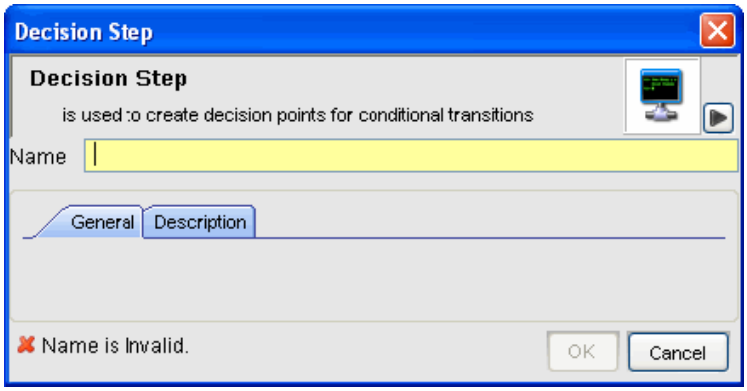

2 Provide a name.

3 Click the *Description* tab to provide a description for this step.

4 Click *OK*.

To edit a mail step:

1 Right-click a mail step and select *Edit Step*.

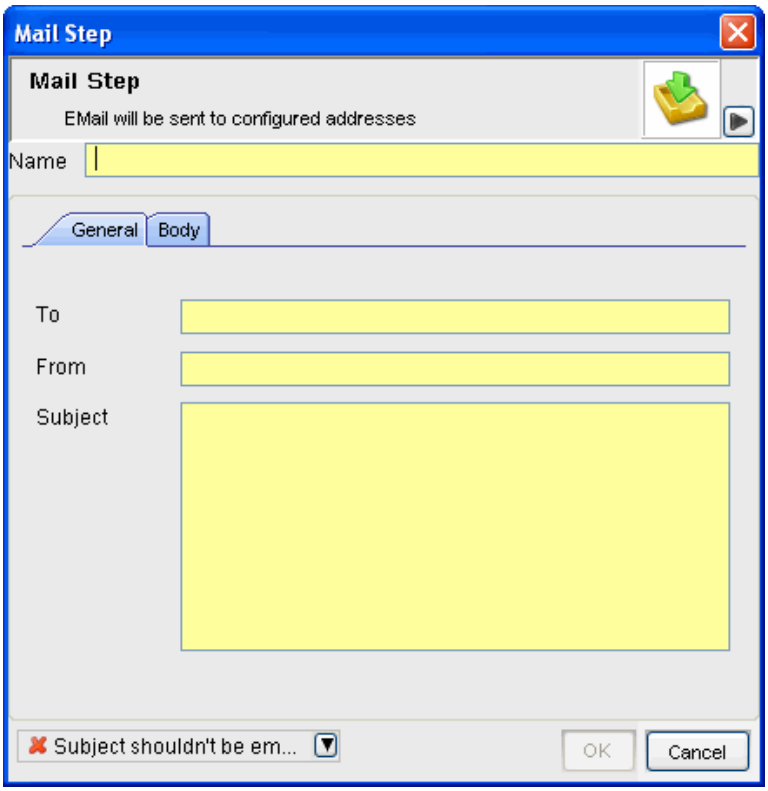

2 Provide a name for the step.

3 Provide To and From mail addresses and a Subject in the *General* tab.

4 Click the body tab and type the message.

5 Click *OK*.

To edit a command step:

1 Right-click a command step and select *Edit Step*.

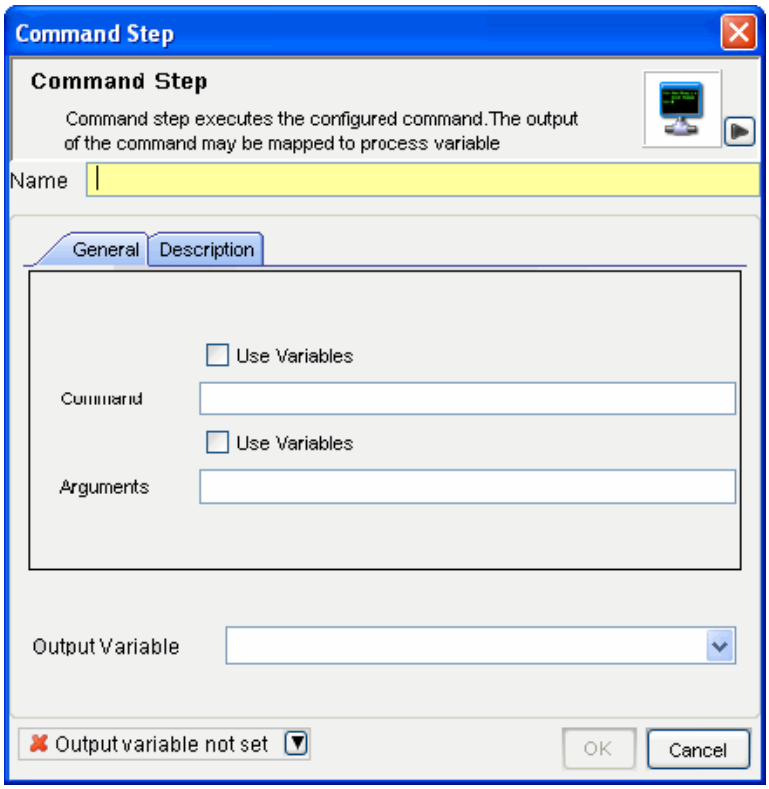

2 Provide a name for this step.

3 Specify the path and name of the command or script to execute (relative to the `<install_directory>/config/exec`)

4 If you want to run a command or script referenced in a variable that is populated during the workflow process, select the *Use Variables* check box.

5 Specify any command-line arguments to pass to the command or script. If you want to use the contents of a variable that is populated during the workflow process, select the *Use Variable* check box.

6 Specify a variable to hold output from the command or script. Any standard output is placed into these variables.

7 Click *Description* tab to provide a description for this step.

8 Click *OK*.

Deleting Steps

1 Click the *iTRAC* tab.

2 In the Navigator, click *iTRAC Administration* > *Template Manager*.

3 Select an existing template, then click *View/Edit*. The iTRAC Process Builder window displays.

4 Right-click an existing step, then click *Delete Step*.

5 In the Alert Message window, select *Yes* to delete.

6.6 Transitions

Transitions are used to connect steps. There are several types of transitions:

- Unconditional
- Conditional
- Timeout
- Alert
- Else
- Error

A transition can have the following attributes:

- Name
- Description
- Destination
- Expression
- Timeout Values

Different steps have different properties and therefore they are associated with different transition types.

Table 6-5 *Steps and Valid Transition*

Step Type	Valid Transitions
◆ Decision	◆ Conditional ◆ Else
◆ Manual	◆ Unconditional ◆ Timeout ◆ Alert
◆ Command ◆ Mail ◆ Activity	◆ Unconditional ◆ Timeout ◆ Alert ◆ Error

6.6.1 Unconditional Transitions

An unconditional transition must always be used from a start step. manual, command, activity, and mail steps can also have unconditional transitions. The only parameter for an unconditional transition is the next step.

This path is taken when the current step is completed (unless a timeout transition is configured and the timeout period elapses).

To add an unconditional transition:

1 Open the Process Builder.
2 Right-click an existing step and select *Add Transition*.
3 Specify a name for the transition.
4 Select *Unconditional* from the Transition *Type* list.

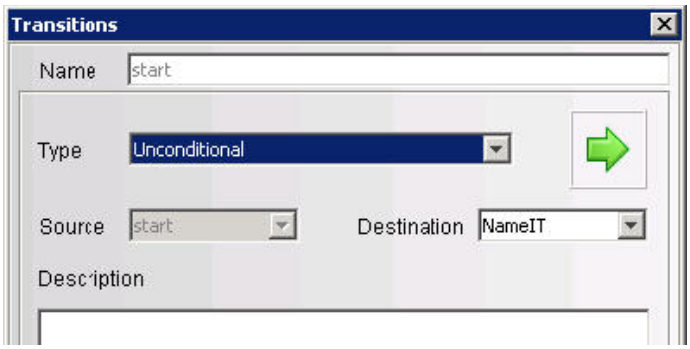

5 Click the down-arrow for the *Destination* field and select a step.

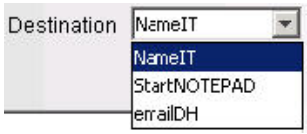

6 Provide a description for this transition and click *OK*.

6.6.2 Conditional Transitions

Select an exit path based on an expression using iTRAC variables set in a manual or command step.

NOTE: You can add Conditional transitions only from a decision step to any other step.

When creating a Conditional transition, the conditional expressions can be based on comparing a variable that is populated during the workflow process to a specific value or to another variable populated during the workflow process. Multiple conditional expressions can be combined or nested using the AND and OR operator.

To add a Conditional transition:

1 Open the Process Builder.
2 Right-click an existing decision step and select *Add Transition*.
3 Provide a name for the transition.
4 Select the *Conditional* transition type from the list.

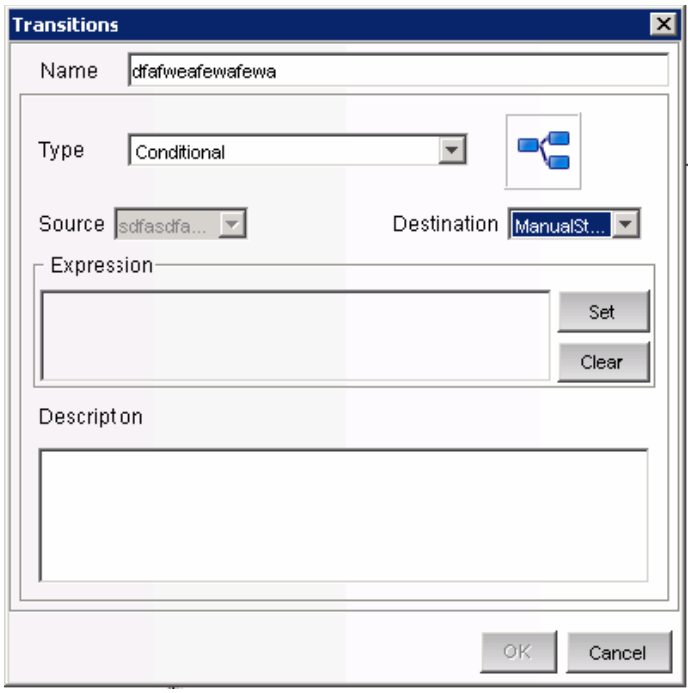

5 Specify the destination step.

6 Click *Set* to add an expression. The empty Expression window displays.

7 Click *EXP* to add the first expression. The evaluation expression is an expression that evaluates to TRUE or FALSE during the workflow process. Select the appropriate drop-down list under *Relations* to compare a variable to a constant value (Variables and Values) or to another variable (Variables and Variables).

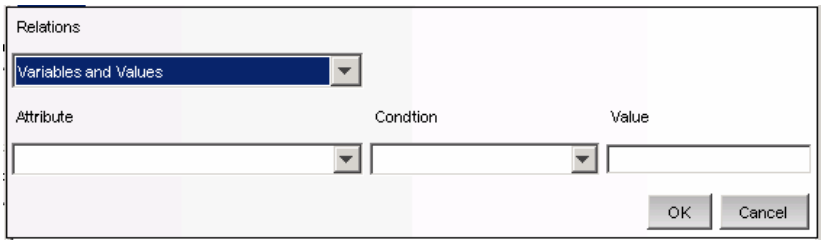

8 Select a variable from the *Attribute* drop-down list or add a new one if desired.

9 Select a condition from the *Condition* drop-down list. The condition list varies depending on the type of Attribute variable chosen.

String Variable Conditions:

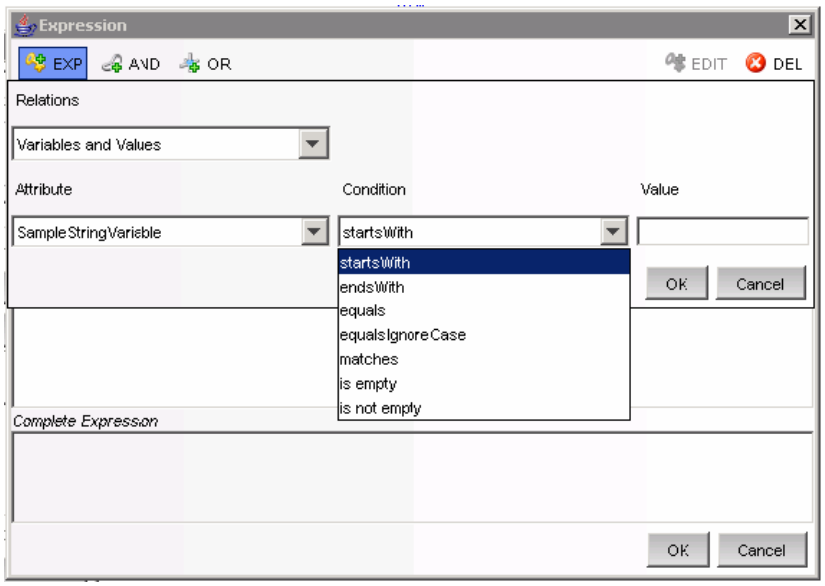

Integer and Float Variable Conditions:

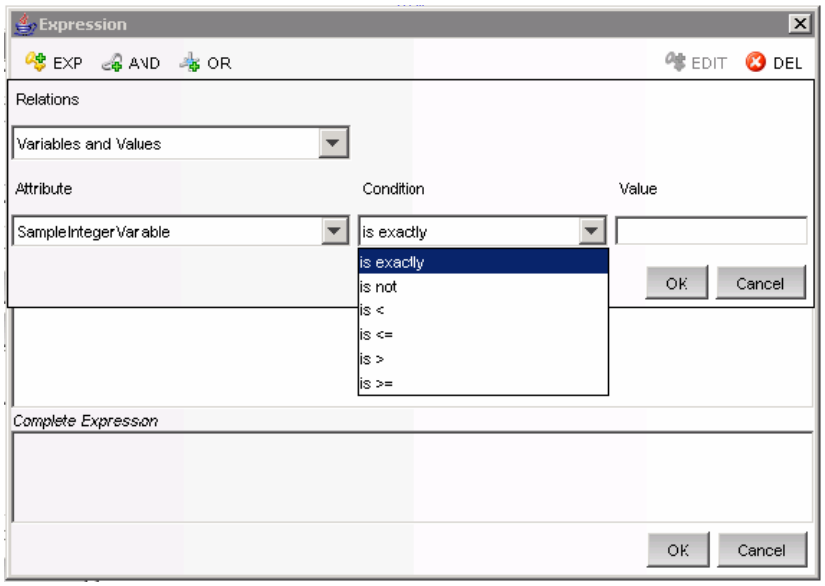

Boolean Variable Conditions:

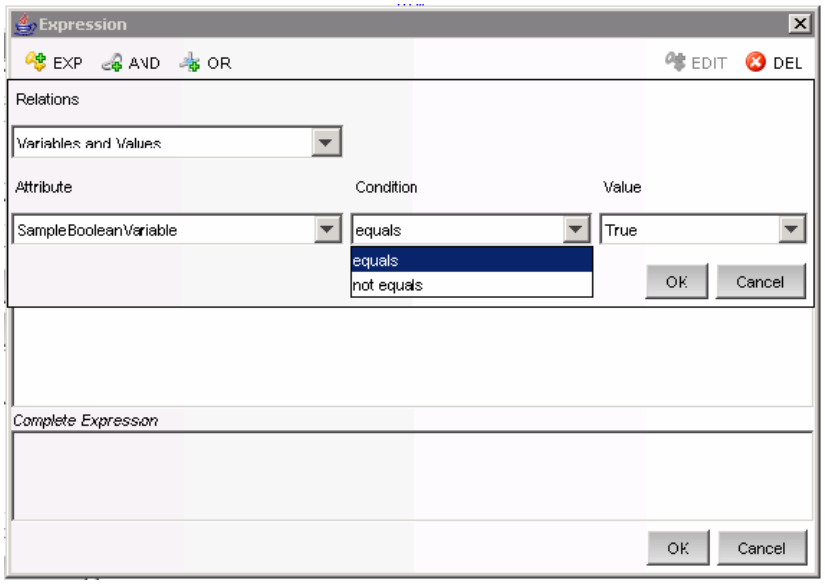

10 Set the value.

11 Click *OK*.

12 If a second expression is desired, select the root folder.

13 Repeat steps 7-12 as needed.

14 By default, all expressions at the root level are separated by AND operators. To nest expressions or to use the OR operator, click the appropriate operator button and drag and drop expressions onto that operator.

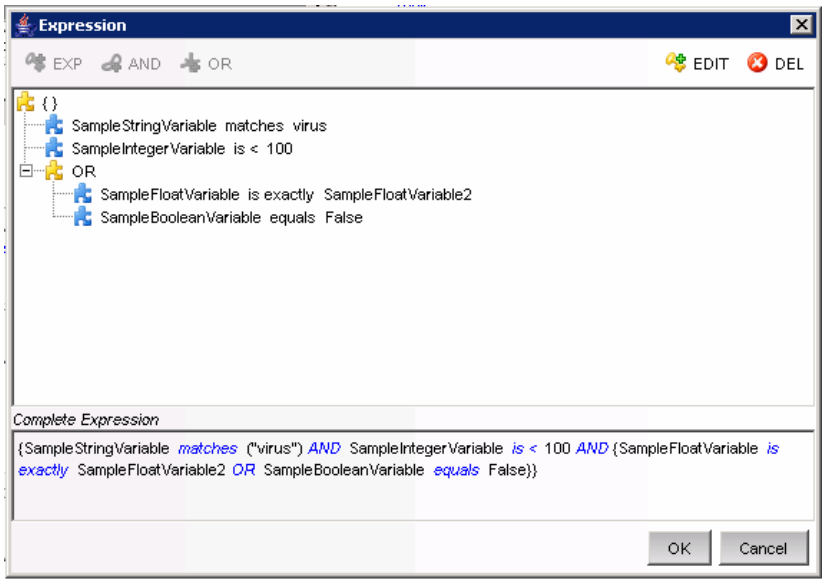

15 When the expression is complete, click `OK`.

You can edit/delete an existing expression using the *Edit* and *Delete* buttons in the Expression window.

16 Click *OK*. The expressions you provided displays in the Transition window under the Expression section.

17 Provide a description for your transition and click *OK*.

6.6.3 Else Transitions

An Else transition leads to a path that is taken from a decision step when the criteria for the Conditional transitions are not met. This transition only applies to decision steps, and every decision step must have an Else transition. The workflow path with the Else transition is only followed if none of the criteria for the Conditional transitions are met.

NOTE: You can add *Else* transitions only from a decision step to any other step.

To add an Else transition:

1 Open the Process Builder.

2 Right-click an existing decision step and select *Add Transition*.

3 Select the Transition type *Else* from the list.

4 Specify the destination step.

5 Provide a description for this step and click *OK*.

6.6.4 Timeout Transitions

A Timeout transition leads to a path that is taken when a user-specified amount of time (minutes, hours or days) elapses after a Base Time, which is either step_activated_time or step_accepted_time. Step_activated_time is the time that iTRAC activates this step within the workflow process.

Step_accepted_time is the time when a user accepts (or takes ownership) of the worklist item for this step. If the timeout time period passes without the step being completed, control moves to the next step.

Timeout transitions can be set for a manual step or a command step. Step_accepted_time is only relevant for manual steps and should not be selected for a command step.

This transition is represented by a red line.

To add a Timeout transition:

1. Open the Process Builder.
2. Right-click an existing decision step and select *Add Transition*.
3. Select the transition type timeout from the list.
4. Specify the destination step.
5. Click *Set* to specify the Timeout details. The timeout details window displays.
6. Specify the timeout value in minutes, hours, or days. Click *OK*.
7. Select *Base Time*.
8. Provide a description for your transition and click *OK*.

6.6.5 Alert Transitions

An Alert transition leads to a path that is taken when a user-specified amount of time (minutes, hours or days) elapses after step_activated_time or step_accepted_time. At this point, the workflow process is usually escalated to a user who can intervene and take action.

Step_activated_time is the time that iTRAC activates this step within the workflow process. Step_accepted_time is the time when a user accepts (or takes ownership) of the worklist item for this step.

If the alert time period passes without the step being completed, the workflow process branches into two active paths. The original step remains active for user intervention. The alert path is also initiated. For example, the alert path might escalate the workflow process to the attention of a supervisor, although the main path is still open and the original owner still has the option to complete the worklist item. Another example is that if a command is taking too long to run, you might want to alert an analyst to investigate the delay or possibly run the command manually.

Alert transitions can be set for a manual step or a command step. Step_accepted_time is only relevant for manual steps and should not be selected for a command step.

This transition is represented by a yellow line.

To add an Alert transition:

1. Open the Process Builder.
2. Right-click an existing decision step and select *Add Transition*.
3. Select the Alert transition type from the list.
4. Specify the destination step.
5. Click Set to provide the Alert details. The Alert details window displays.

6 Specify the Alert Time value, in minutes, hours, or days. Click *OK*.

7 Provide a description for your transition and click *OK*.

6.6.6 Error Transition

An Error transition leads to a path that is taken if an automated step cannot successfully complete. Error transitions can be used for command, mail, and activity steps (for example, if a command step fails to execute).

Error transitions should typically lead to some kind of notification. For example, an Error transition might lead to a manual step in which the user is instructed to manually run a process that previously failed.

NOTE: The Error transition is only taken if the iTRAC call to the command, mail, or activity step fails. If there is an internal error with the command script or the mail server fails, this does not satisfy the conditions for an Error transition.

Only the destination step can be specified, along with a description.

To add an Error transition:

1 Open the Process Builder.

2 Right-click an existing decision step and select *Add Transition*.

3 Select the Error transition type from the list.

4 Specify the destination step.

5 Provide a description for this step and click *OK*.

6.6.7 Managing Transitions

After creating a transition, you can edit or delete the transition.

- "Modifying Transitions" on page 148
- "Deleting Transitions" on page 149

Modifying Transitions

1 Click the *iTRAC* tab.

2 In the Navigator, click *iTRAC Administration* > *Template Manager*.

3 Select an existing template, then click *View/Edit*. The iTRAC Process Builder window displays.

4 Double-click an existing transition line. The Transitions window displays.

5 Edit the transition as needed.

6 If you are editing an expression from a decision step, click the button and double-click the expression.

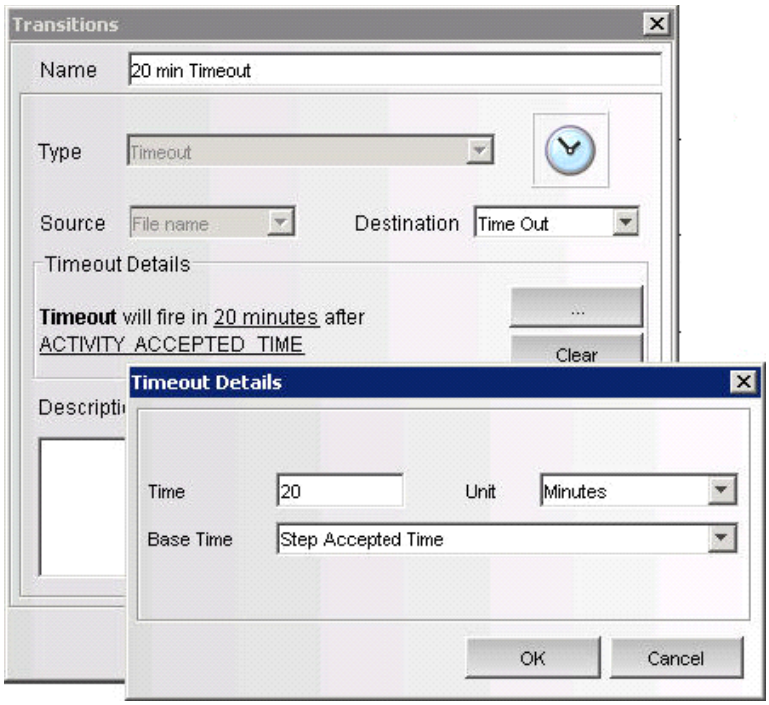

7 Edit as needed.

8 Click *OK* until you exit the Transitions window.

9 Click *Save*.

Deleting Transitions

1 Click *iTRAC* tab.

2 In the Navigator, click *iTRAC Administration* > *Template Manager*.

3 Select an existing template, then click View/Edit. The iTRAC Process Builder window displays.

4 Right-click an existing step and select *Remove Transition*.

5 In the Alert Message window, click *Yes*.

6.7 Activities

An activity is very similar to a command step, except that activities are reusable and cannot use input or output variables. The Activities pane shows a library of user-defined, reusable activities that can reduce the amount of configuration necessary when building templates.

Activities are exported or imported as XML files. These files can be exported or imported from one system to another.

- Section 6.7.1, “Incident Command Activity,” on page 150
- Section 6.7.2, “Incident Internal Activity,” on page 151
- Section 6.7.3, “Eradication Activity,” on page 151
- Section 6.7.4, “Incident Composite Activity,” on page 151

- Section 6.7.5, “Creating iTRAC Activities,” on page 151
- Section 6.7.6, “Managing Activities,” on page 154

Figure 6-3 *Activity Pane*

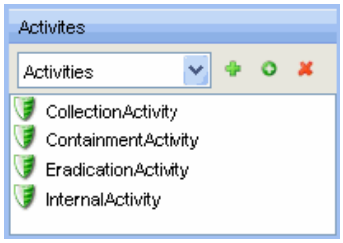

iTRAC activities can be used in iTRAC templates to define a workflow step, or they can be manually executed from within an incident. Sentinel provides three types of actions that can be used to build Activities:

- Section 6.7.1, “Incident Command Activity,” on page 150
- Section 6.7.2, “Incident Internal Activity,” on page 151
- Section 6.7.3, “Eradication Activity,” on page 151
- Section 6.7.4, “Incident Composite Activity,” on page 151
- Section 6.7.5, “Creating iTRAC Activities,” on page 151
- Section 6.7.6, “Managing Activities,” on page 154

6.7.1 Incident Command Activity

An incident command activity enables you to launch a specific command with or without arguments. The following fields from the incident associated with the workflow process can be used as input to the command:

- DIP (Target IP)
- DIP : Port
- RT1 (DeviceAttackName)
- SIP (Initiator IP)
- SIP : Port
- Text (incident information in name value pair format)

NOTE: The command must be stored in the `<install_directory>\config\exec` directory on the iTRAC workflow server, usually the same machine where the Data Access Server (DAS) is installed.

6.7.2 Incident Internal Activity

An incident internal activity enables you to mail or attach information from the Sentinel database to the incident associated with the workflow process. Each of these options has a prerequisite.

- **Vulnerability for the Initiator IP address (SIP) or the Target IP address (DIP):** This requires that you run a vulnerability scanner and bring the results of the scan into Sentinel by using a Vulnerability (or "information") Collector.
- **Advisor attack-related data:** This requires the purchase and installation of the optional Advisor data subscription service.
- **Asset data** This requires that you run an asset management tool such as NMAP and bring the results into Sentinel by using an Asset Collector.

To send mail messages from within the Sentinel Control Center, you must have an SMTP Integrator that is configured with connection information and with the SentinelDefaultEMailServer property set to true.

6.7.3 Eradication Activity

The eradication activity is used to run the `arp` command. The `arp` command displays and modifies the IP-to-Physical address translation tables used by the Address Resolution Protocol (ARP).

The `arp -a` command displays the current ARP entries by interrogating the current protocol data. If `inet_addr` option is specified, the IP and physical addresses for only the specified computer are displayed. If more than one network interface uses ARP, entries for each ARP table are displayed.

6.7.4 Incident Composite Activity

An incident composite activity enables combine one or more existing command and internal activities.

6.7.5 Creating iTRAC Activities

1. Click the *iTRAC* tab.
2. In the Navigator, click *iTRAC Administration* > *Activity Manager* or click the *Add* button in the Activity pane.
3. Select an existing activity and click the `Add` button. The Activity Wizard window displays.
4. Select an activity type: *Command*, *Internal*, or *Composite*.
5. Provide a name and description for this activity. Click *Next*.

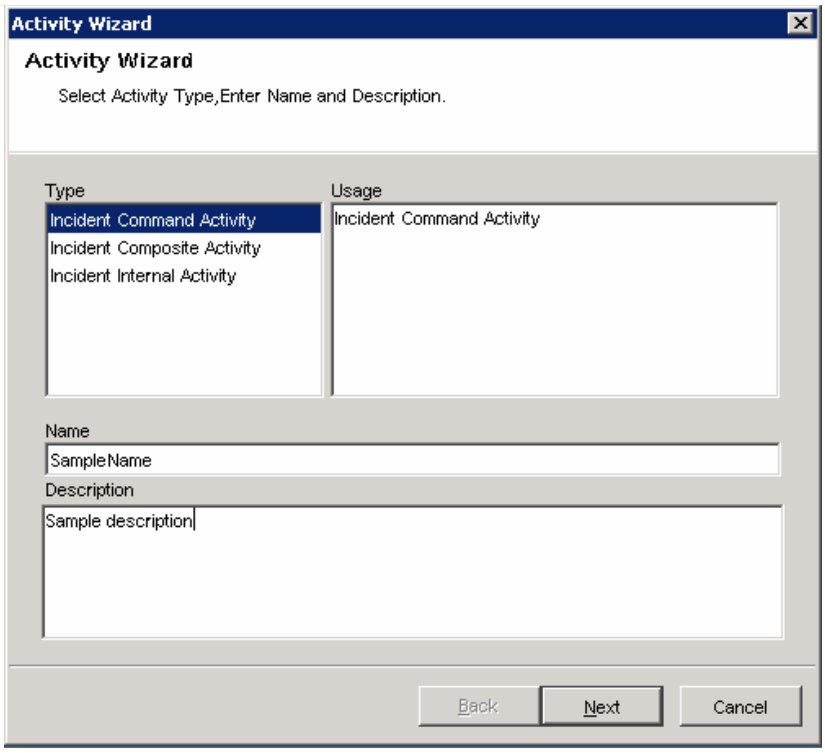

6 (Conditional) If you selected an incident command activity, configure the settings:

6a In the Command Arguments Wizard, specify the command.

6b Provide the arguments for this command. You can select *None*, *Incident Output* (Values from the Drop-down list), or specify *Custom* values.

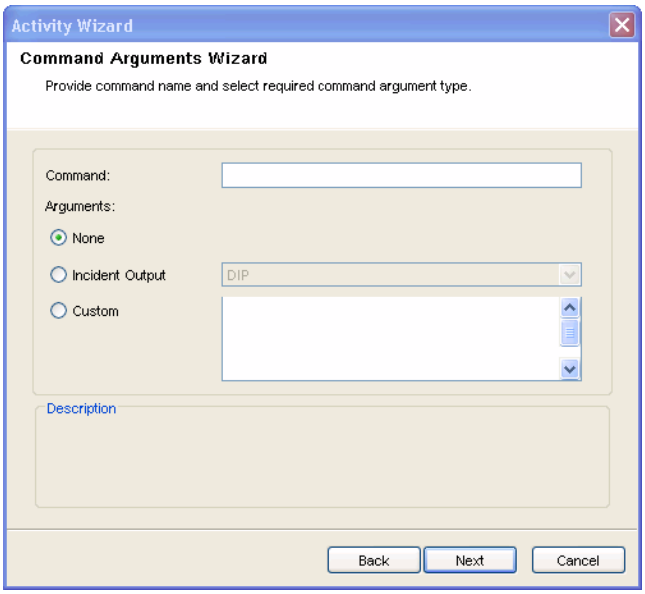

6c Click *Next*.

6d (Optional) Configure an incident command activity to e-mail the output to a specific address or attach the output to the incident associated with the workflow process in this window.

6e Select *Mail* and specify the To and From e-mail address and subject.

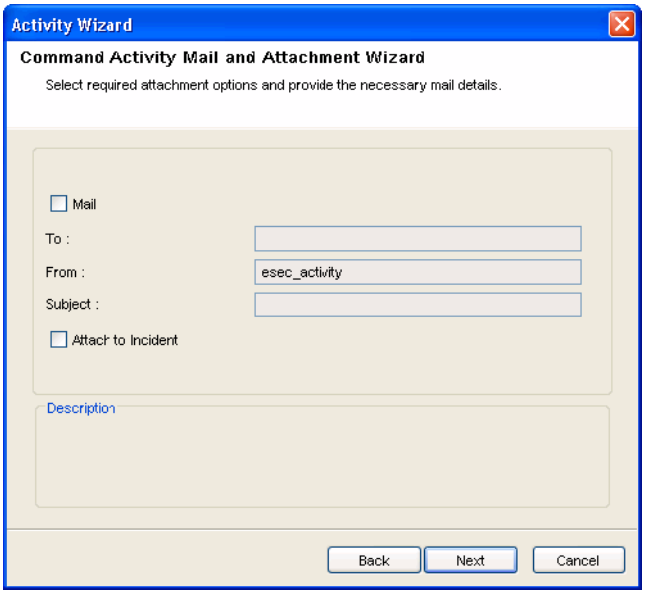

6f Select *Attach to Incident*, if required.

6g Click *Next*.

6h View and confirm the details you chose in the Summary page and click Finish.

7 Conditional) If you selected an incident internal activity, configure the settings:

7a In the Command Arguments Wizard, specify the command.

7b Provide the arguments for this command. You can select *None*, *Incident Output* (Values from the Drop-down list), or specify *Custom* values.

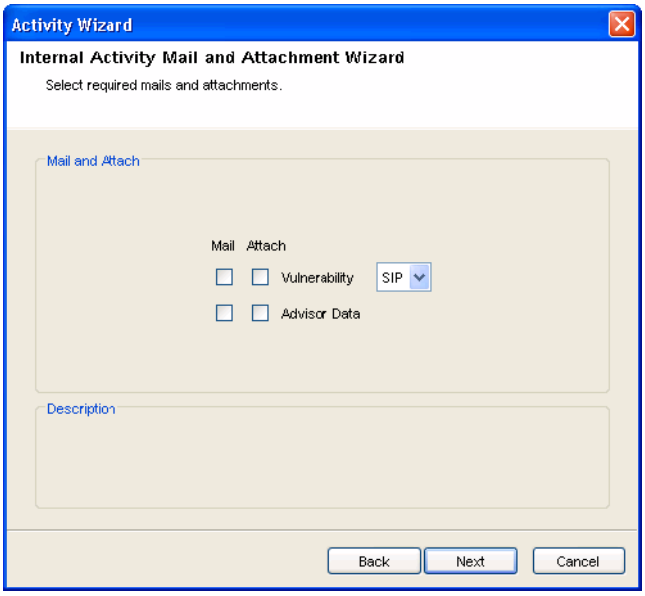

7c Click *Next*.

7d Select your options (Mail and attach).

7e If you select Mail, you are prompted to provide To and From e-mail address and subject. Provide this information and click *Next*.

View and confirm the details you chose in the Summary page and click *Finish*.

8 Conditional) If you selected an incident composite activity, configure the settings:

8a Select the activities from the list of available activities and click *Next*.

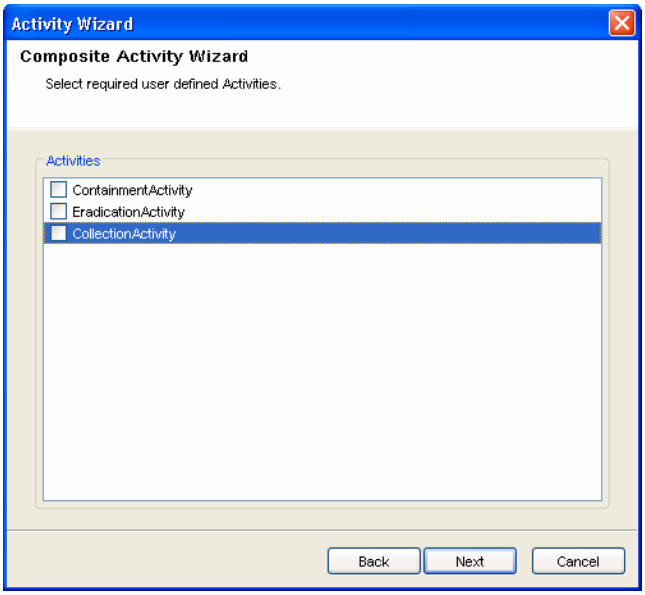

8b View and confirm the details you chose in the Summary page and click *Finish*.

6.7.6 Managing Activities

After creating an activity, you can modify, import or export it.

- "Modifying Activities" on page 154
- "Exporting Activities" on page 154
- "Importing Activities" on page 155

Modifying Activities

1. Click the *iTRAC* tab.
2. In the Navigator, click *iTRAC Administration* > *Activity Manager*.
3. Select activity that needs modification and click View/Edit. Edit Activity window displays.
4. Edit information in the *General*, *Attachment*, and *Mail* tabs.
5. Click *OK*.

Exporting Activities

1. Click the *iTRAC* tab.

2 In the Navigator, click *iTRAC Administration* > *Activity Manager*.

3 Click the *Import/Export Activity* icon. The *Import/Export* Wizard window displays.

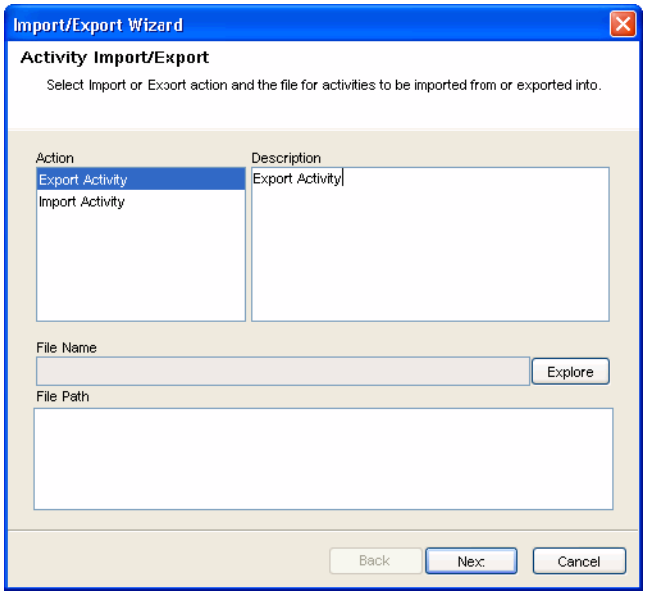

4 Select Export Activity and click *Explore*.

5 Navigate to where you want save your exported file.

6 Click *Next*.

7 Select one or more activities to be exported.

8 Click Next, then click *Finish*.

Importing Activities

1 Click the *iTRAC* tab.

2 In the Navigator, click *iTRAC Administration* > *Activity Manager*.

3 Click the *Import/Export Activity* icon. The Import/Export Wizard window displays.

4 Select Import Activity and click *Explore*.

5 Navigate to your import file. Click *Import*.

6 Click *Next*.You see a list of activities that are imported.

7 Click *Next*, then click *Finish*.

6.8 Process Management

Process management allows you to view the incident's progress in the workflow or terminate a workflow process. Process management allows you to:

- Display the status of your process
- Start your process
- Terminate your process

Process execution is the time period during which the process is operational, with process instances being created and managed.

When an iTRAC process is executed or instantiated in the iTRAC server, a process instance is created, managed, and eventually terminated by the iTRAC server in accordance with the process definition. As the process progresses towards completion or termination, it executes various activities defined in the workflow template based on the criteria for the transitions between them. The iTRAC workflow server processes manual and Automatic steps differently.

An iTRAC process must be created with a single associated incident; there is therefore a one-to-one match between iTRAC processes and incidents. Not all incidents are necessarily attached to processes.

NOTE: Only one incident can be associated to an iTRAC process instance.

- Section 6.8.1, “Instantiating a Process,” on page 156
- Section 6.8.2, “Automatic Step Execution,” on page 156
- Section 6.8.3, “Manual Step Execution,” on page 157
- Section 6.8.4, “Displaying Status,” on page 157
- Section 6.8.5, “Displaying the Status of a Process,” on page 157
- Section 6.8.6, “Changing Views in the Process Manager,” on page 158
- Section 6.8.7, “Starting or Terminating a Process,” on page 159

6.8.1 Instantiating a Process

An iTRAC process can be instantiated in the iTRAC server by associating an incident to an iTRAC process by the following three methods

- Associating an iTRAC process to the incident at the time of incident creation
- Associating an iTRAC process to incident after an incident has been created
- Associating an iTRAC process to an incident through correlation

For more information on associating a process to an incident, see Chapter 5, “Incidents Tab,” on page 109.

6.8.2 Automatic Step Execution

When the process instance executes an automatic activity step, command step, or mail step, it executes the associated activity or command defined in the template, and stores the result in process variables. It then transitions to the next step in the iTRAC template.

For example, an activity might be defined to ping a server; when this activity is executed in a workflow process the activity runs and attaches the results to the associated incident.

6.8.3 Manual Step Execution

On encountering a manual step, the iTRAC server sends out notifications in the form of work items to the assigned resource. If the step was assigned to a role, a work item is sent to all users within the role. The iTRAC server then waits for the user to complete the work item before proceeding to the next activity.

For more information, see Section 7.1, “Work Item Summary,” on page 161.

NOTE: All manual steps must be assigned to a role or to a group of users.

6.8.4 Displaying Status

The Display Status function is used to monitor the progress of a process. As the process instance progresses, you can track the progress visually by clicking the *Refresh* button. The process monitor also provides an audit trail of all the actions performed by the iTRAC server when executing the process.

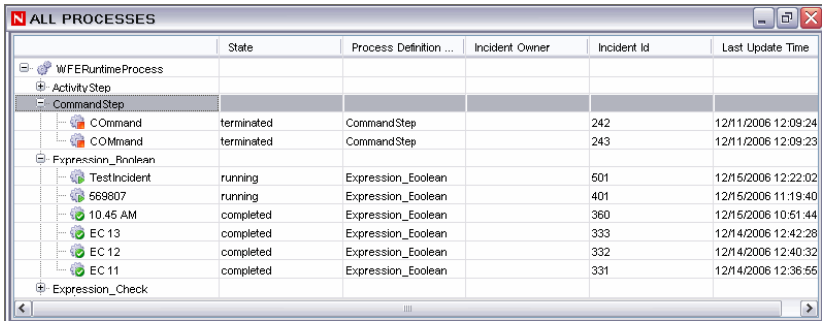

Activities that are running, completed, and terminated are represented by the following icons:

Table 6-6 *Status of an Activity*

Icon	Description
	Running
	Completed
	Terminated

6.8.5 Displaying the Status of a Process

1. Click the *iTRAC* tab.
2. Click the *Display Process Manager* icon.
3. Click the down-arrow next to the *Switch Views* button to select a view or create a new view.
4. In the Process Manager window, right-click a process and select *Actions* > *Display Status*.

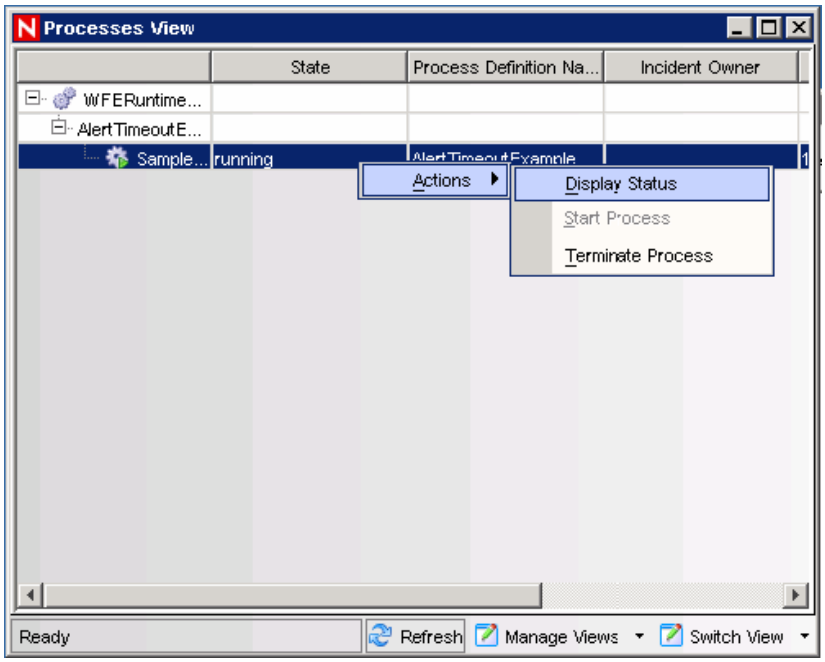

The current step is highlighted in red.

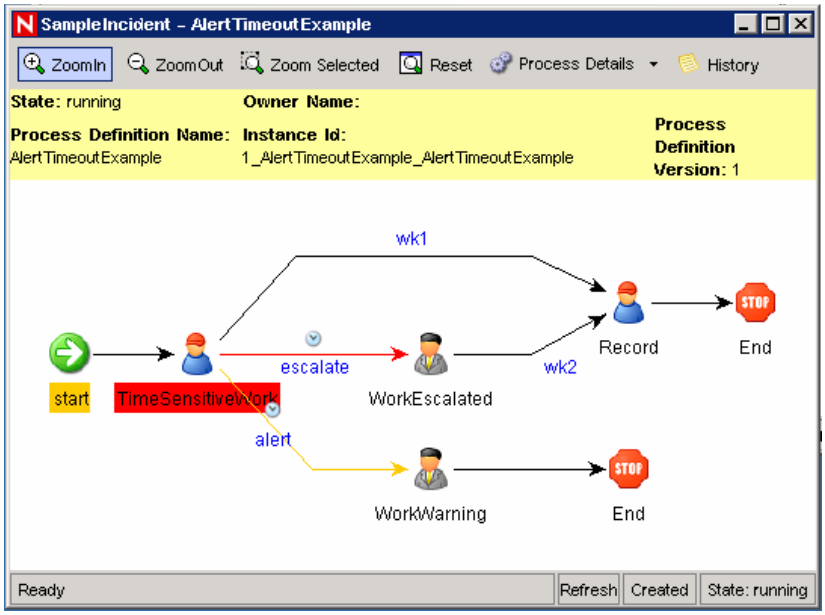

5 Close the window.

6.8.6 Changing Views in the Process Manager

1 Click the *iTRAC* tab.

2 Click the Display Process Manager icon.

3 Click the drop-down list in Manage View and select *Edit Current View* option.

4 In the View Option window, set the following options as necessary:

- Fields
- Group by
- Sort
- Filter
- Tree Display

5 Click *Apply* and *Save*.

The following is view with Tree Display set to Status (running and not started).

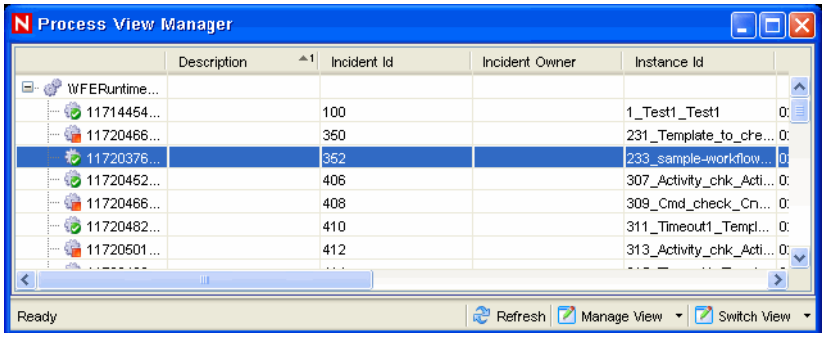

6.8.7 Starting or Terminating a Process

1 Click the *iTRAC* tab.

2 Click the Display Process Manager icon.

Alternatively, you can select *iTRAC* > *Display Process Manager*.

3 Click the drop-down arrow next to the *Switch Views* button to select a view or create a new view.

4 In the Process View Manager window, right-click a process and select *Actions* > *Start Process*, or click *Terminate Process*.

Work Items

7

A work item is a workflow task assigned to a particular user or role in the iTRAC application. The individual activities to be performed to complete an iTRAC process are listed as work items in the Work Item Summary in the Sentinel Control Center. For more information on iTRAC processes, see Chapter 6, "iTRAC Workflows," on page 123. You can access the work items from any tab in the Sentinel Control Center.

NOTE: To have access to a work item, you must assign it to you or acquire the work item management permissions. If you have Work Item management permission, you can manage work items of other users.

- Section 7.1, "Work Item Summary," on page 161
- Section 7.2, "Processing a Work Item," on page 164
- Section 7.3, "Managing Work Items of Other Users," on page 165

7.1 Work Item Summary

The Work Item Summary lists the work items allocated to a user as an individual and as a member of a group; it can be thought of as an incident workflow to-do list for a user who is a part of the incident response process. In the Work Item Summary, you can access the work items, view them, and process them to complete the task.

In the Work Item Summary, work items are grouped by current user and by other users with similar roles. The following example is for a user who is a member of the Admin, Analyst, Jr Analyst, and Admin Asst groups.

Figure 7-1 *Work Item Summary*

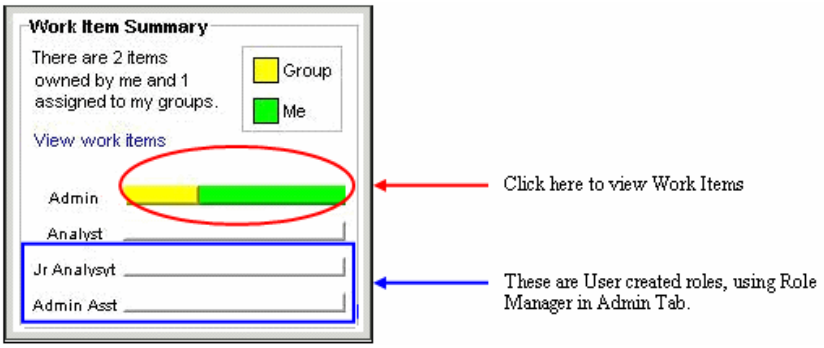

The following is an example of a user who is a member of the Analyst group who has a process assigned to his role (group).

Figure 7-2 Work Item Summary Example

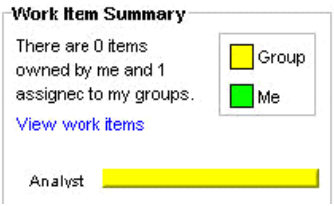

To view a work item:

1. In the Work Item Summary, click the yellow or green bar.

   A work item list for the group or the current user displays and shows the name and ID of the incident, the workflow process name, and the step name and description

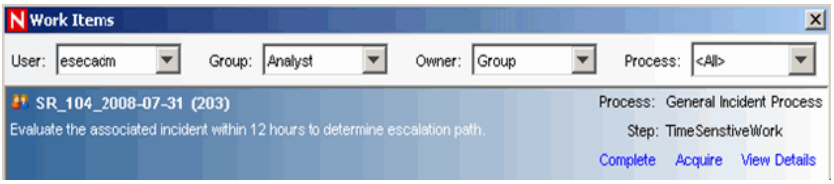

2. Double-click any work item and click *View Details*.

   The Work Item Details window displays and shows the process details, including any detailed instructions included by the iTRAC workflow developer and any variables that need to be set in the step.

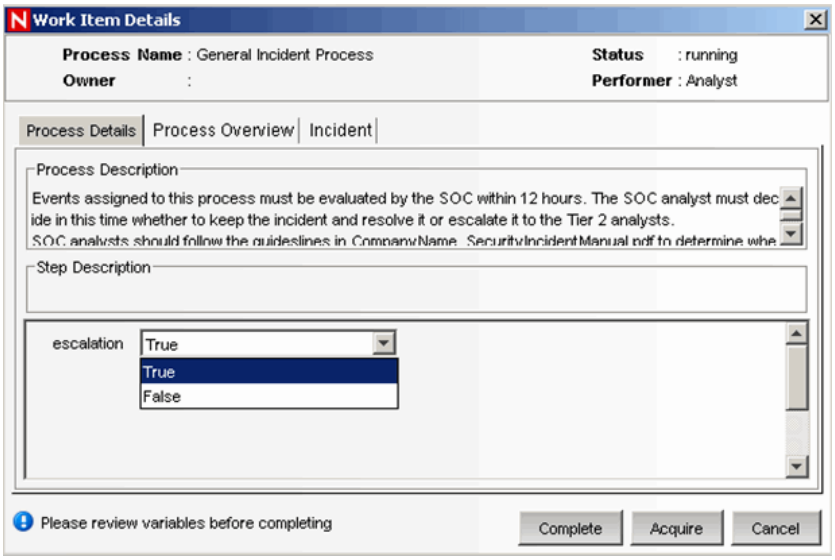

3. Click *Process Overview* to view an overview of the entire iTRAC process.

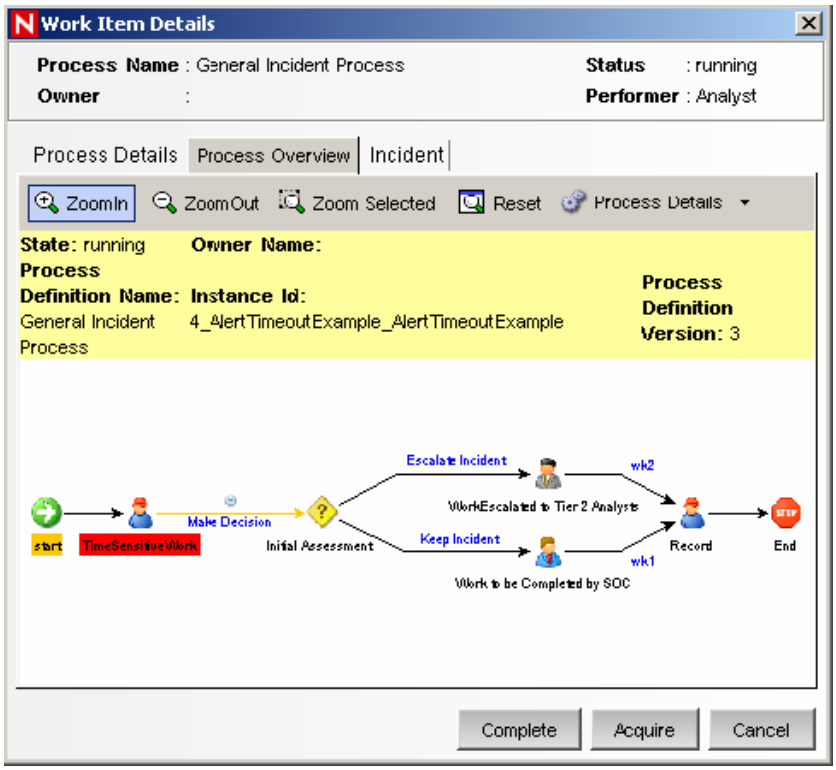

4 Click *Incident* to view the details of the associated incident.

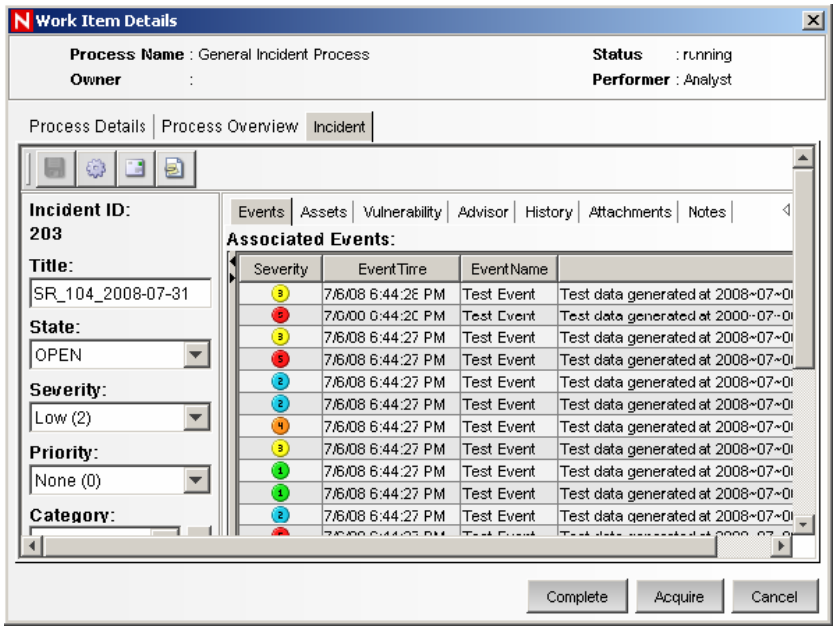

5 To take responsibility for this work item, click *Acquire*.Otherwise, click *Cancel*.

NOTE: Any changes to the incident from this screen must be saved. There is a *Save* button on the toolbar and another *Save* button at bottom of the screen.

The information on the *Process Details* and *Process Overview* tabs is defined by the iTRAC workflow designer. For more information on creating workflow templates, see Chapter 6, "iTRAC Workflows," on page 123.

7.2 Processing a Work Item

A work item can be accessed from any part of the main tabbed Sentinel Control Center interface.

- You can process a work item in a group even if you have logged in as a different user. However, you cannot acquire a step if you have logged in as a different user.
- The work item remains with the user of a group who has acquired it.
- Consecutive steps are dependent. If two steps in a row are assigned to the same role, the user who acquires the first step is also assigned the second step.
- Non-consecutive steps are independent. For example, if a workflow proceeds from steps that are assigned to the Tier 1 Analyst group, then to the Tier 2 Analyst group, then back to the Tier 1 Analyst group, the third step is available to the entire Tier 1 Analyst group; it is not assigned to the individual user who handled the first step.

7.2.1 Accepting and Completing a Work Item

1. In the Work Item Summary, click the yellow or green bar. A work item list for the group or the current user displays.

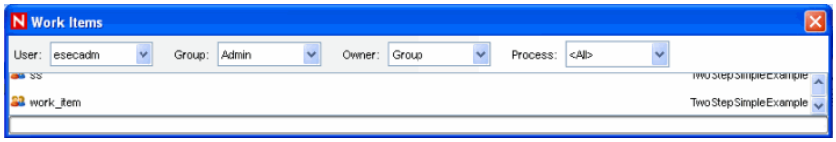

2. To assign an iTRAC process to you, select the process and click *Acquire*.

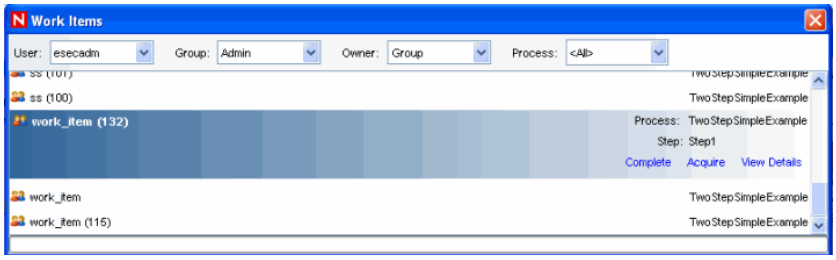

The Work Item Summary changes from yellow to green.

Work item assigned to a group (role)	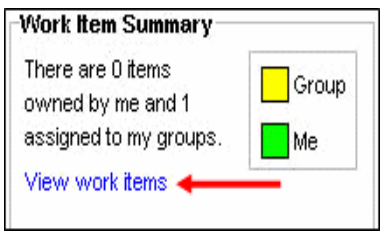

Work item assigned to the user under the Analyst role.	

When you acquire (accept) a work item, it is removed from the queue of all other users in the same role. The work item can be returned to the group by clicking *Release*.

3 Click *View Details*.

The current step within a work item is highlighted in red.

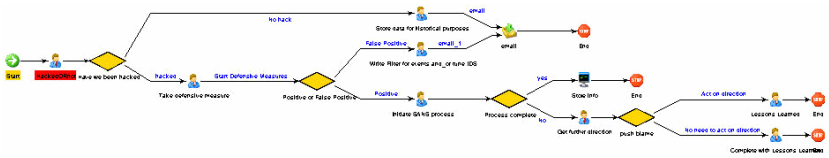

4 To take action on the step, click the *Process Details* tab.

For a manual step and depending on the type of variable (Integer, String, Boolean and Float) assigned to that step, click the down-arrow and select a decision. If necessary, you can add comments or add an attachment.

In all other cases, the steps are automatic.

5 Click *Complete* to complete the process.

Completing the work item signals the completion of the task to the iTRAC server. The updateable variables from the work item are processed by the server to move to the next step, which depends on how the workflow is defined. The work item is removed from the user's worklist and appears in the worklist of the individual or role associated with the next step in the process.

7.3 Managing Work Items of Other Users

The Administration function allows an administrative user to release a work item from a specific user back to everyone in a role. This is beneficial if a work item is in already in process but the assigned user cannot complete the work.

1 Log in to a Sentinel machine as a user with iTRAC – Manage Work Items Of Other Users user rights.

2 In the Summary pane, click *View Work Items*.

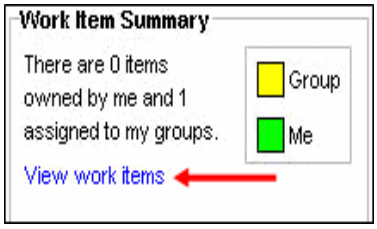

3 In the Work Items window, set the following:

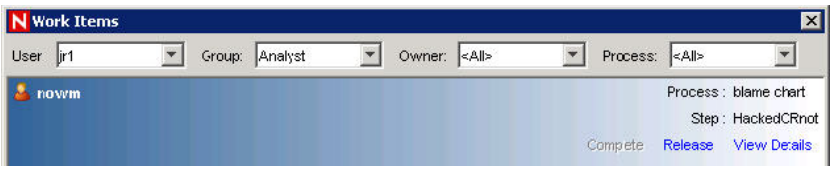

- **User:** Name of the user that has acquired the process
- **Group:** Name of the group that the user belongs to. In the above example, the user belongs to the Analyst group.
- **Owner** Select either <*All*> (all processes acquired or not), *me* (acquired processes) or *Group* (un-acquired processes).
- **Process:** Name of the process.

In the above example, all processes acquired by jr1, who belongs to the Analyst group, with all processes listed.

4 To release the work item, select the work item and click *Release*. Release changes to Acquire (not available).

In this example, only a member of the Analyst group can acquire this work item.

Analysis Tab

8

The *Analysis* tab allows you run offline queries. An offline query helps you in ad hoc reporting. In an offline query, you can save and generate the queries offline. This helps in optimizing network usage because it relieves the network from heavy processing when similar queries are triggered. You must have proper permissions to use *Analysis* tab. If this permission is not assigned, the *Analysis* tab is not displayed.

- Section 8.1, "Introduction to the User Interface," on page 167
- Section 8.2, "Offline Query," on page 169

8.1 Introduction to the User Interface

In the *Analysis* tab, you can see the *Offline Queries* options.

Table 8-1 *Analysis Tab User Interface*

User Interface	Description
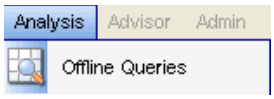	The Analysis menu in the menu bar
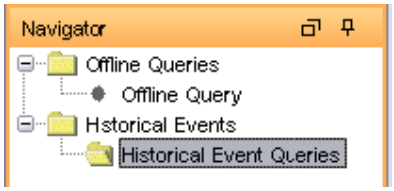	The Navigation Tree in the Navigation pane
The toolbar buttons	

8.1.1 Top Ten Dashboard

The following Top 10 dashboards are available in Sentinel and can be downloaded from the Sentinel Content page (http://support.novell.com/products/sentinel/secure/sentinel61.html):

- Top 10 Target IP Addresses
- Top 10 Initiating IP Addresses
- Top 10 Target Host Names
- Top 10 Initiating Host Names
- Top 10 Target User Names
- Top 10 Initiating User Names

- Top 10 Target Port Names
- Top 10 Event Names

The Top 10 dashboards are enabled by default, and the following summaries are turned on to enable the Top 10 dashboards:

- EventDestSummary
- EventSevSummary
- EventSrcSummary

If Top 10 dashboards are not needed, you can disable these summaries, or you can enable additional summaries in order to use them for reporting. If the summary service is not in use, you can disable it.

To enable or disable summaries:

1. In the Sentinel Control Center, go to *Admin* > *Report Data Configuration*.
2. Select the Summary to enable or disable and click the status (*Active*/*Inactive*) of that summary.
3. Select *Yes* to confirm that you want to change the status of the summary.

To enable or disable EventFileRedirectSerice:

1. At your Sentinel machine, using text editor, open:

   `<install_directory>/config/das_binary.xml`
2. For EventFileRedirectService, change the status to on or off, as appropriate. For example:

   `<property name="status">off</property>`
3. Log in to the Sentinel Control Center as the Sentinel Administrator.
4. Go to *Admin* > *Servers View*.

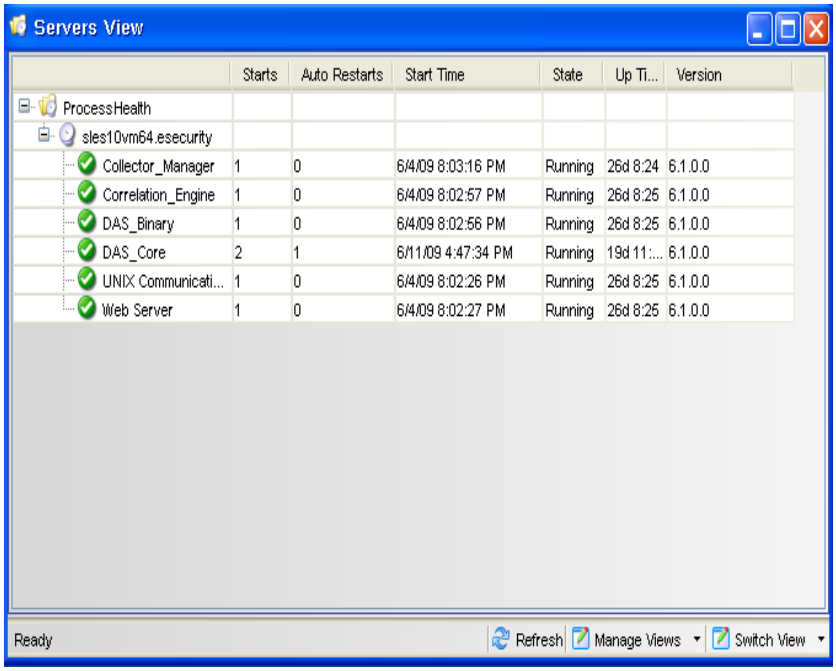

5. Right-click *DAS_Binary* and select *Restart*.

8.2 Offline Query

An offline query is most often used to run queries against large amounts of data. An offline query continues to run even after the user logs out of the Sentinel Control Center, if necessary.

NOTE: You can view the result of your query only after it is completely processed.

After the query has completely finished processing, the results are available to the user who initiated the offline query and other Sentinel users with the same security filter. When you attempt to browse or save the result as HTML or CSV, the data is transferred from the server to the local machine running the Sentinel Control Center.

For performance reasons, the result set for offline query is limited to 100,000 records. For better results, you must specify a better filter or a smaller time range when creating an offline query.

- Section 8.2.1, "Creating an Offline Query," on page 169
- Section 8.2.2, "Viewing, Exporting, or Deleting an Offline Query," on page 170

8.2.1 Creating an Offline Query

1. Click *Analysis* on the menu bar. The Offline Query window displays. Alternatively, you can click the *Offline Query* button on the toolbar.
2. In the Offline Query window, click *Add* button located on the top left corner of the page. The Add Offline Query window displays.

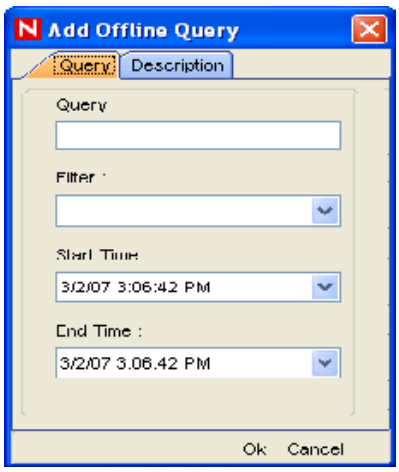

3. Provide a query name, then select an existing filter to be used for generation of offline query.

   For more information on the selection and creation of filters see Chapter 3, "Active Views Tab," on page 53.
4. Select the start date and end date for which you want to generate an offline query.
5. Specify the description in the Description tab.
6. Click *OK*. The offline query is listed in the Offline Query window.

8.2.2 Viewing, Exporting, or Deleting an Offline Query

1. Click *Analysis* on the menu bar. The Offline Query window displays. Alternatively, you can click the Offline Query button on the toolbar.
2. In the Offline Query window, select an offline query. The following options are available:
   - **Browse:** Click *Browse* to view the output of the offline query in the Active Browser window.
   - **CSV:** Click *CSV* to generate a comma separated value file with the queried information.
   - **HTML:** Click *HTML* to generate an HTML file with the queried information.
   - **Delete:** Click *Delete* to delete the offline query. A confirmation message alert displays. Click Yes to delete.
   - **Details:** Click *Details* to view the details of the offline query as specified when the query was added.

Advisor Usage and Maintenance

9

Sentinel Advisor, powered by Security Nexus, is an optional data subscription service that provides device-level correlation between real-time events, from intrusion detection and prevention systems, and from enterprise vulnerability scan results. Advisor acts as an early warning service and detects attacks against vulnerable systems by providing normalized attack information. It also provides the associated remediation information.

The Advisor subscription is optional. However, it is necessary if you want to use the Sentinel Exploit Detection or the Advisor Reporting features.

- Section 9.1, “Understanding Advisor,” on page 171
- Section 9.2, “Understanding Exploit Detection,” on page 172
- Section 9.3, “Introduction to the Advisor User Interface,” on page 174
- Section 9.4, “Downloading the Advisor Feed,” on page 178
- Section 9.5, “Viewing the Advisor Status,” on page 179
- Section 9.6, “Viewing the Advisor Data,” on page 181
- Section 9.7, “Resetting the Advisor Password,” on page 182
- Section 9.8, “Deleting the Advisor Data,” on page 182
- Section 9.9, “Advisor Audit Events,” on page 182

9.1 Understanding Advisor

The Advisor service and its corresponding Exploit Detection feature depend on the mappings between the attacks against enterprise assets and the known vulnerabilities of those assets. The Advisor and the Exploit Detection features require the following data to work with the Advisor products:

- **Vulnerability scan data:** The vulnerability scanners check enterprise assets for known vulnerabilities. The scanned data can then be loaded into the Sentinel database to serve as referential information, by using the Collectors that support Advisor.
- **Advisor mapping data:** The Advisor data contains information about known threats, including attacks and vulnerabilities. The Advisor service gathers information from various vulnerability and intrusion detection vendors, and creates mappings between abstract vulnerabilities and attacks.

  Security Nexus provides the Advisor feed data that contains information about known security vulnerabilities and threats, and also provides normalization of intrusion detection signatures and vulnerability scans. The Advisor data feed is updated on a regular basis as new attacks and vulnerabilities are reported. The updates are available at the Novell download Web site (https://secure-www.novell.com/sentinel/download/advisor/).

NOTE: With Sentinel Rapid Deployment or later, the initial Advisor data feed is installed by default on the Sentinel Rapid Deployment server at `<install_directory>/data/updates/advisor`. However, you must purchase an additional license from Novell to download the updated Advisor feed on a regular basis.

- **Real-time attack data:** Intrusion detection systems report real-time attacks against enterprise assets. However, this data does not indicate the impact of the attacks.

  The real-time attacks that are generated as events are loaded into the Sentinel database by using the intrusion detection systems or vulnerability type Collectors.

9.2 Understanding Exploit Detection

- Section 9.2.1, "How Exploit Detection Works," on page 172
- Section 9.2.2, "Generating the Exploit Detection File," on page 174
- Section 9.2.3, "Viewing the Events," on page 174

9.2.1 How Exploit Detection Works

Exploit detection instantly sends notification when an attack is attempting to exploit a vulnerable system. The Exploit Detection feature depends on the following:

- Both vulnerability scanners and the intrusion detection systems must report vulnerabilities and attacks against the same set of systems. In Sentinel, systems are identified by their IP addresses and their MSSP Customer Name. The MSSP Customer Name is a namespace identifier that prevents overlapping IP ranges from matching incorrectly.
- The vulnerability scanner and intrusion detection system products must be supported by the Advisor service. This data uses specific product identifiers to ensure proper matching.
- The specific reported attacks and vulnerabilities must be known to the Advisor service and Exploit Detection.

All Collectors shipped by Novell meet these requirements, as long as they are declared as being supported by Advisor. To write your own vulnerability or intrusion detection Collector, or to modify one of the shipped Collectors, refer to the Sentinel Plug-in SDK (http://www.novell.com/developer/develop_to_sentinel.html) for specific information about which event and vulnerability fields must be filled in to support this service.

The following table lists the supported products with their associated device type (IDS for intrusion detection system, VULN for vulnerability scanners, and FW for firewall).

Table 9-1 *Supported Products and the Associated Device Types*

Supported Products	Device Type	RV31 Value
Cisco Secure IDS	IDS	Secure
Enterasys Dragon Host Sensor	IDS	Dragon
Enterasys Dragon Network Sensor	IDS	Dragon
Intrusion.com (SecureNet_Provider)	IDS	SecureNet_Provider
ISS BlackICE PC Protection	IDS	XForce
ISS RealSecure Desktop	IDS	XForce

Supported Products	Device Type	RV31 Value
ISS RealSecure Network	IDS	XForce
ISS RealSecure Server	IDS	XForce
ISS RealSecure Guard	IDS	XForce
Sourcefire Snort/Phalanx	IDS	Snort
Symantec Network Security 4.0 (ManHunt)	IDS	ManHunt
Symantec Intruder Alert	IDS	Intruder
McAfee IntruShield	IDS	IntruShield
TippingPoint	IPS	TippingPoint
eEYE Retina	VULN	Retina
Foundstone Foundscan	VULN	Foundstone
ISS Database Scanner	VULN	XForce
ISS Internet Scanner	VULN	XForce
ISS System Scanner	VULN	XForce
ISS Wireless Scanner	VULN	XForce
Nessus	VULN	Nessus
nCircle IP360	VULN	nCircle IP360
Qualys QualysGuard	VULN	QualysGuard
Cisco IOS Firewall	FW	Secure

To enable exploit detection, the Sentinel Collectors must populate several variables as expected. Collectors built by Novell populate these variables by default.

- In intrusion detection systems and vulnerability Collectors, the RV31 (DeviceName) variable in the event must be set to the value in the RV31 column in Table 9-1. This string is case sensitive.
- In the intrusion detection systems Collector, the DIP (Destination or Target IP) must be populated with the IP address of the machine that is being attacked.
- In the intrusion detection systems Collector, RT1 (DeviceAttackName) must be set to the attack name or attack code for that intrusion detection system.
- In the intrusion detection systems and vulnerability Collectors, RV39 (MSSPCustomerName) value must be populated. For a standard corporation, the value can be anything. For a Managed Security Service Provider (MSSP), the customer name should be set for the individual customer. For either type of company, the value in the intrusion detection systems Collector must exactly match with the value in the vulnerability Collector.

These values are used by the Mapping Service to populate the VULN field in the event. This value is used to evaluate the incoming events to determine whether a vulnerability is exploited or not. When the vulnerability field (VULN) equals 1, the asset or destination device is exploited. If the vulnerability field equals 0, the asset or destination device is not exploited.

9.2.2 Generating the Exploit Detection File

When you run the intrusion detection system or vulnerability type Collectors, events from all the selected products are scanned for possible attacks and vulnerabilities, and the product name and MSSP customer name are mapped to the Advisor product name and MSSP customer name. If the events match successfully, the exploit information (IP address, Device Name, Attack Name, and MSSP Customer Name) is updated in the `exploitdetection.csv` file in the `<install_directory>/data/map_data` directory.

The initial mapping time might take up to 30 minutes. However, you can modify the time by changing the value of the `minregenerateinterval` property in the `ExploitDetectDataGenerator` component of the `das_core.xml` file. The time is given in milliseconds. For example, you can change the time from 1800000 (30 minutes) to 180000 (3 minutes).

NOTE: You must restart the das_core services after you change the time.

9.2.3 Viewing the Events

To view events that indicate a possible exploitation, create an Active View with a filter that has the Vulnerability value set to 1.

Within an event, the values in the Vulnerability field indicate the following:

- **1:** the asset or destination device is possibly exploited.
- **0:** the asset or destination device is not exploited.

NOTE: If the `exploitdetection.csv` file is not generated, the Vulnerability field is blank.

For more information on viewing events in Active Views, see Section 3.4, “Viewing Real-Time Events,” on page 57.

9.3 Introduction to the Advisor User Interface

- Section 9.3.1, “The Advisor Window,” on page 175
- Section 9.3.2, “Processing the Advisor Feed,” on page 176
- Section 9.3.3, “Configuring the Advisor Products for Exploit Detection,” on page 177

Ensure that you have Advisor Configuration permission to access the Advisor window.

You can access the Advisor user interface through one of the following methods:

Table 9-2 *Navigating to Advisor*

Location	User Interface
The *Admin* menu in the menu bar	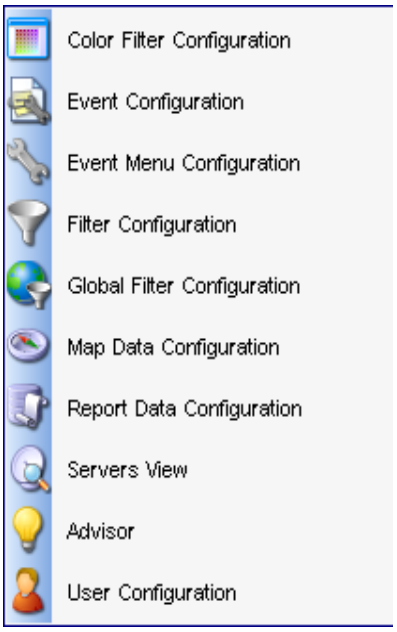
The Navigation tree in the Navigation pane	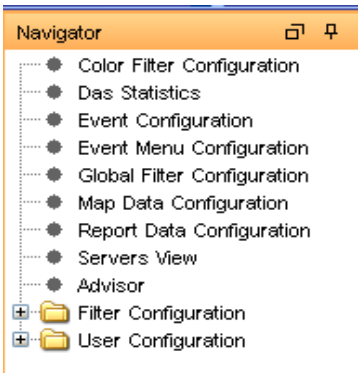
Admin Toolbar	The Advisor icon .

9.3.1 The Advisor Window

The Advisor window has two sections:

- **Download Method:** Enables you to process the Advisor feed files manually and launch the Download Manager feature to configure the Sentinel server for automated processing of the feed files. For more information on processing the Advisor feed, see Section 9.3.2, “Processing the Advisor Feed,” on page 176.
- **Exploit Detection:** Lists the vulnerable products that are included in the feed files, and enables you to configure the products for exploit detection. For more information, see Section 9.3.3, “Configuring the Advisor Products for Exploit Detection,” on page 177.

NOTE: The Exploit Detection section initially displays a blank list unless you process the initial Advisor feed that was loaded during Sentinel installation. For more information, see Section 9.3.2, "Processing the Advisor Feed," on page 176.

Figure 9-1 *Advisor Window*

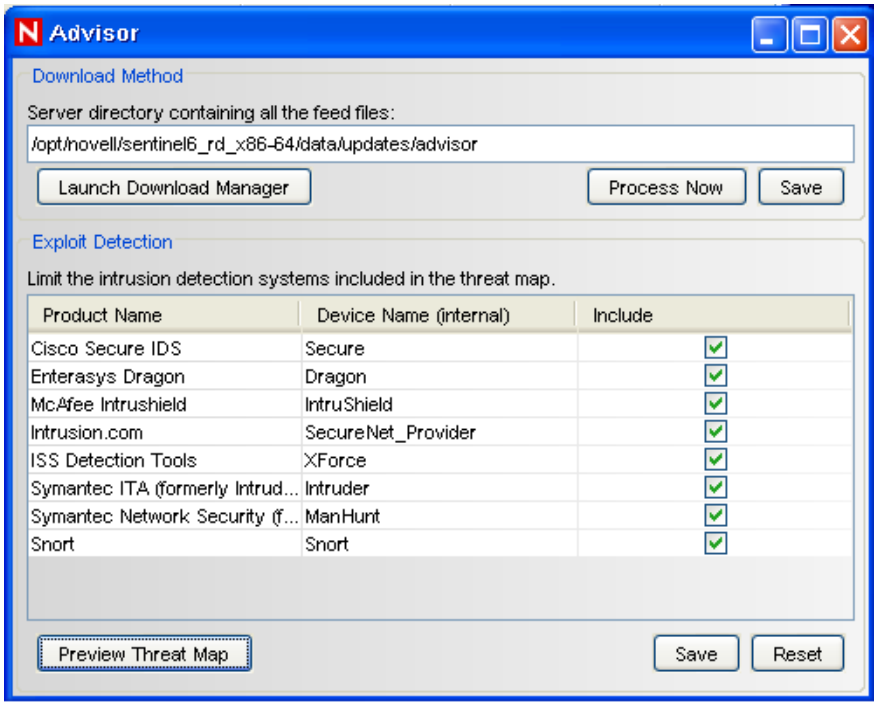

9.3.2 Processing the Advisor Feed

You can process the Advisor feed files manually or you can configure the Sentinel Rapid Deployment server to automatically process the feed files at scheduled time intervals.

- "Processing the Feed Files Manually" on page 176
- "Processing the Feed Files Automatically" on page 177

Processing the Feed Files Manually

1. In the Advisor window, select the directory where you downloaded the latest Advisor feed files.

   The initial Advisor feed is loaded at <*install_directory*>`/data/updates/advisor`.

2. Click *Process Now* to process and load the feed files into the Sentinel database.

   After the feed files are processed, the products included in the feed files are displayed in the Exploit Detection section.

3. (Optional) Click *Save* to save the location of the Advisor directory.

Processing the Feed Files Automatically

You can use the Download Manager to configure the Sentinel Rapid Deployment server to automatically process the feed files after they are downloaded.

1. In the Advisor window, click *Launch Download Manager*. For more information, see Chapter 10, "Download Manager," on page 183.

9.3.3 Configuring the Advisor Products for Exploit Detection

The Exploit Detection section of the Advisor window lists the names of the Advisor products and the device names that are included in the feed files.

1. Select the products that need to be included for Exploit Detection by selecting the corresponding check box.
2. (Conditional) To remove any product from the list, deselect the corresponding check box.
3. Click *Save* to save the changes made to the Advisor products list.

   After the product list is saved, the `exploitdetection.csv` file is updated. For more information on exploit detection, see "Generating the Exploit Detection File" on page 174.
4. (Optional) Click *Reset* to undo the changes made to the Exploit Detection products list.

For more information on exploit detection, see Section 9.2, "Understanding Exploit Detection," on page 172.

Viewing the Threat Map

The Preview Threat Map window lists the top 5000 entries of the `exploitdetection.csv` file. This list displays the attacks that attempt to exploit your machine.

To view the threat map: click *Preview Threat Map*.

NOTE: This list is blank unless the `exploitdetection.csv` file has been generated.

Figure 9-2 Preview Threat Map

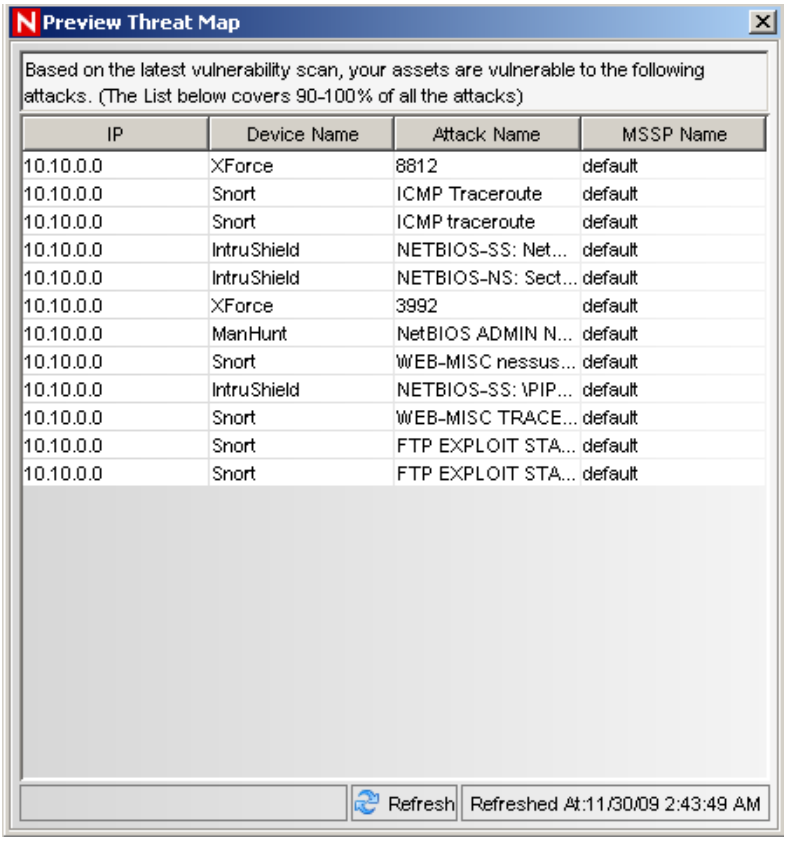

IP	Device Name	Attack Name	MSSP Name
10.10.0.0	XForce	8812	default
10.10.0.0	Snort	ICMP Traceroute	default
10.10.0.0	Snort	ICMP traceroute	default
10.10.0.0	IntruShield	NETBIOS-SS: Net...	default
10.10.0.0	IntruShield	NETBIOS-NS: Sect...	default
10.10.0.0	XForce	3992	default
10.10.0.0	ManHunt	NetBIOS ADMIN N...	default
10.10.0.0	Snort	WEB-MISC nessus...	default
10.10.0.0	IntruShield	NETBIOS-SS: \PIP...	default
10.10.0.0	Snort	WEB-MISC TRACE...	default
10.10.0.0	Snort	FTP EXPLOIT STA...	default
10.10.0.0	Snort	FTP EXPLOIT STA...	default

9.4 Downloading the Advisor Feed

To access the updated Advisor feed, you must download the feed and process the downloaded feed to load it into the Sentinel database. You can download the Advisor feed manually or you can configure the Sentinel server for automated downloads at fixed intervals.

NOTE: To download Advisor updates, you must purchase the Advisor Data Subscription and obtain the credentials.

- Section 9.4.1, “Configuring the Sentinel Server for Automated Downloads,” on page 178
- Section 9.4.2, “Downloading the Advisor Feed Manually,” on page 179

9.4.1 Configuring the Sentinel Server for Automated Downloads

You can use the Download Manager to configure the Sentinel server for automated Advisor feed downloads. For more information on Download Manager, see Chapter 10, “Download Manager,” on page 183.

To launch Download Manager from the Advisor window:

1. Open the Advisor window by doing one of the following:
   - Click *Admin* > *Advisor*.
   - Select *Advisor* in the Navigation pane.
   - Click the icon on the toolbar.
2. Click *Launch Download Manager*.

   The Download Manager window is displayed. For more information on Download Manager, see Chapter 10, "Download Manager," on page 183.

9.4.2 Downloading the Advisor Feed Manually

1. Log in to the Novell download Web site (https://secure-www.novell.com/sentinel/download/advisor/) by using your Novell eLogin username and password.

   The Novell eLogin username and password must be associated with the Advisor license.
2. Download all the `.zip` and `.md5` files.
3. Copy the downloaded feed files to the Sentinel Rapid Deployment server.

   To process the downloaded feed, you must provide the location where you have saved the feed in *Admin* > *Advisor* window. The default location is `<install_directory>/data/updates/advisor`.

9.5 Viewing the Advisor Status

The Advisor Status window lists the products Novell supports for Advisor and also displays the status of the last five feed files that have been processed or are being processed.

NOTE: Ensure that you have permissions to view the Advisor Status window.

To view the Advisor Status window: click the *Advisor* tab.

The Advisor Status window displays the Advisor information only if the feed files are loaded into the database.

Figure 9-3 *Advisor Data Status*

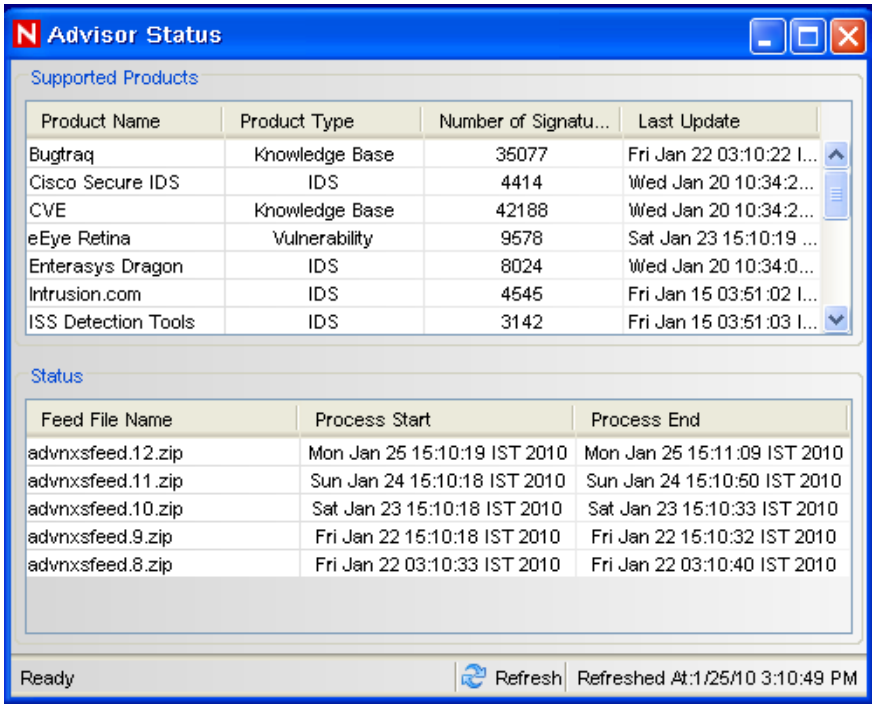

Table 9-3 *Advisor Status*

Fields	Description
Product Name	Name of the product supported by Novell for Advisor. For example: Cisco Secure IDS and Enterasys Dragon Host Sensor.
Product Type	Shows whether the product type is a Vulnerability, Intrusion Detection System (IDS), or Firewall.
Number of Signatures	Shows the number of signatures for the product by Nexus.
Last Update	Time stamp indicating when the product was last updated.
Feed File Name	Shows the name of the feed files that have been processed and are currently being processed.
Process Start	Time stamp indicating when processing the feed file started.
Process End	Time stamp indicating when processing the feed file finished. The process end time is blank if there is an error that halts processing or if processing is still in progress.

9.6 Viewing the Advisor Data

The following are the prerequisites to view the Advisor data:

- The Advisor feed must be up-to-date, processed, and loaded into the Sentinel database.
- The selected event is from a product supported by Advisor and has the Vulnerability field value set to 1.

You can view the Advisor data from the following options:

- *Active Views* tab
  1. In the Sentinel Control Center, click *Active Views*.
  2. In the real-time events table, right-click an event that has the Vulnerability field value set to 1.
  3. Click *Create Incident* > *Advisor* and double-click any of the listed attacks, or click *Analyze* > *Advisor Data*.

Figure 9-4 *Advisor Data*

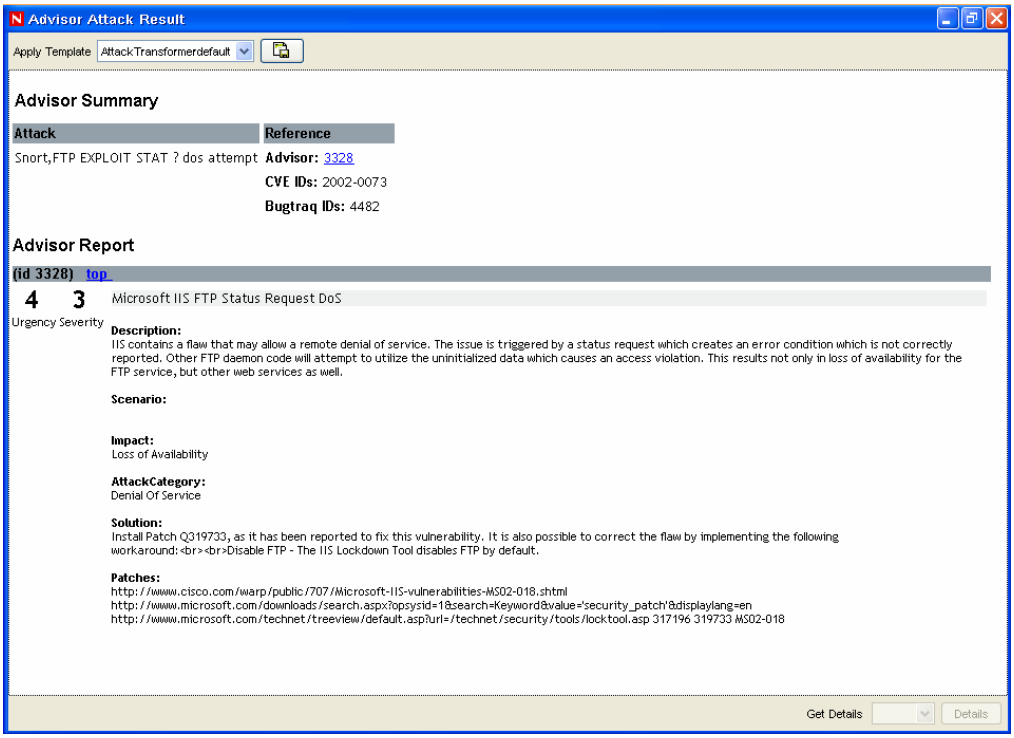

- *Analysis* tab
  1. In the Sentinel Control Center, click *Analysis* > *Offline Queries*.
  2. Add an offline query that filters events with the Vulnerability value set to1.

     For more information on adding an offline query, see Section 8.2, “Offline Query,” on page 169.
  3. Right-click on the event, click *Create Incident* > *Advisor* and double-click any of the listed attacks, or click *Analyze* > *Advisor Data*.

9.7 Resetting the Advisor Password

If you have configured automated downloads for Advisor and if your Novell eLogin password changes, then you need to change your Advisor password. You can change your Advisor eLogin password by using the Download Manager feature.

1. Open the Download Manager by doing one of the following:
   - Click *Tools* > *Download Manager*.
   - Click the Download Manager icon on the toolbar.
   - In the Advisor window, click *Launch Download Manager*.

   The Download Manager window is displayed.

   For more information on Download Manager, see Chapter 10, "Download Manager," on page 183.
2. Select the download configuration for which you want to change the password, then click *Edit*.

   The Edit window is displayed.
3. Specify the new password in the *Password* field.
4. (Optional) Click *Validate* to validate the URL and the login credentials.

   The URL and its credentials are validated and a confirmation message is displayed. If the validation fails, you must provide a valid URL and the login credentials.
5. Click *Save* to save the configuration settings.

9.8 Deleting the Advisor Data

You can delete the Advisor data from the Sentinel database by running the clean_database script. For more information on running the clean_database script, see *Database Cleanup* in the Chapter 14, "Utilities," on page 301.

9.9 Advisor Audit Events

For information on Advisor audit events, see Section B.1, "Advisor Audit Events," on page 431 in Appendix B, "System Events for Sentinel," on page 431.

Download Manager

10

The Download Manager enables you to configure the Sentinel Rapid Deployment server for automated download and processing of the downloaded files at fixed intervals. After the files are downloaded, the Download Manager notifies the Sentinel processes to process the downloaded feed and then loads the processed files to the Sentinel database.

For example, you can configure the Sentinel Rapid Deployment server to download the Advisor feed at fixed intervals. After the feed is downloaded, the Download Manager notifies the Advisor processes to process the downloaded feed and load it into the Sentinel database. However, for Advisor, you must purchase an additional license from Novell to download the updated Advisor feed.

- Section 10.1, “Understanding the Download Manager User Interface,” on page 183
- Section 10.2, “Creating a Download Configuration,” on page 184
- Section 10.3, “Editing a Download Configuration,” on page 187
- Section 10.4, “Downloading the Feed Instantly,” on page 187
- Section 10.5, “Deleting a Download Configuration,” on page 188
- Section 10.6, “Audit Events for the Download Manager,” on page 188

10.1 Understanding the Download Manager User Interface

You can access the Download Manager GUI from the *Tools* menu and also by clicking the shortcut icon in the toolbar.

NOTE: Ensure that you have permission to access the Download Manager feature. For more information on user permissions, see “Administration” in the *Sentinel 6.1 Rapid Deployment Reference Guide*.

The Download Manager window lists the download configurations and displays the status of the previous download for each download configuration. By default, a download configuration is available that is partially configured for Advisor. You can use the same configuration to download the Advisor feed, by specifying the login credentials and other details, and activate the download configuration.

Figure 10-1 *Download Manager*

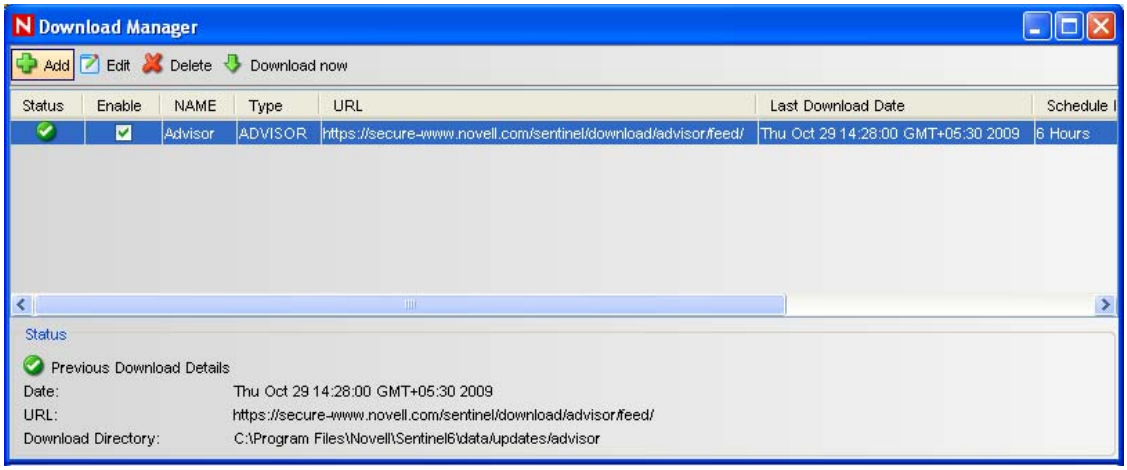

Table 10-1 *Download Configuration Status*

Status	Icon	Description
Download in progress		Indicates that the download is in progress.
Download successful		Indicates that the latest download was successful.
Download not initiated		Indicates that a download has never been initiated. This status does not display any icon.
Download failure		Indicates that all or some of the files were not downloaded. This might be because of a faulty or failed internet connection, or if you have entered invalid credentials.

The status bar displays the previous download details of the selected download configuration, such as the status, date, time, URL, download directory, and errors. If there are any errors, click *Error* to view the error details.

10.2 Creating a Download Configuration

1. Open the Download Manager window by doing one of the following:
   - Click *Tools* > *Download Manager*.
   - Click the icon on the toolbar.
2. Click *Add* to configure the download feed.

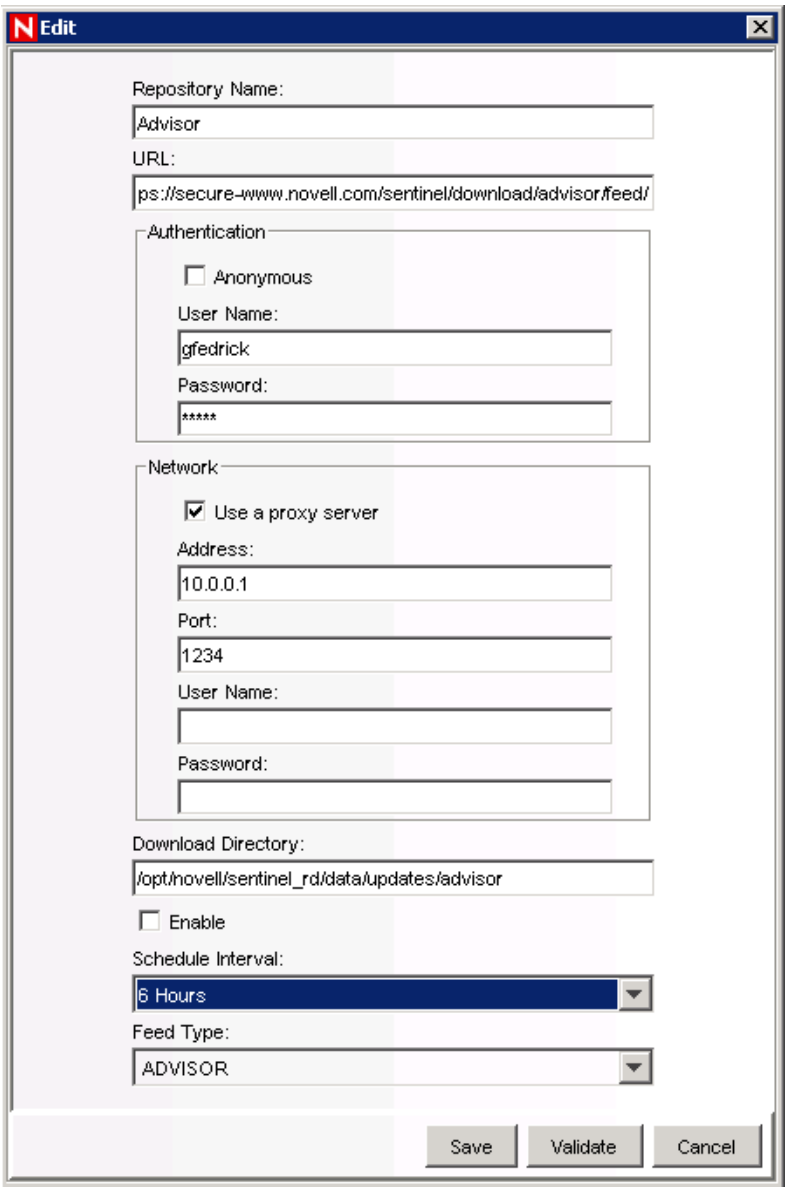

3 Specify the following information:

Option	Description
Repository Name	Name of the repository from which you are downloading the data. The repository name must be unique and meaningful to the feed that you are downloading. For example, to download Advisor data, specify the repository name as Novell Advisor.

Option	Description
URL	URL where the download feed is located. For example, to download Advisor data, specify the Advisor URL: https://secure-www.novell.com/sentinel/download/advisor/feed/ (https://secure-www.novell.com/sentinel/download/advisor/feed/). **NOTE:** Advisor feed can also be manually downloaded from the Novell Download Web site (https://secure-www.novell.com/sentinel/download/advisor/) (https://secure-www.novell.com/sentinel/download/advisor/). However, the manual download URL does not work in Download Manager.
Anonymous	Select the check box to download the information as an anonymous user. If *Anonymous* is selected, the *Username* and *Password* fields are disabled. You can only access the URLs that do not require authentication.
Username and *Password*	Specify valid credentials to log in to the URL of the repository.
Use a Proxy Server	If you do not have direct access to the server where the download feed is located, you can download the advisor feed through proxy server. Select the check box if you are using a proxy server to download the feed. **NOTE:** The proxy server through which you are connecting should have access to the server where the download feed is located.
Address	Specify the proxy server name or server IP address.
Port	Specify the port number through which you connect to the proxy server.
Username and *Password*	Specify valid credentials to log in to the proxy server.
Download Directory	Specify the location and name of the directory where you want to save the feed. Ensure that you specify the absolute path. The directory is created on the Sentinel server at the specified path while downloading the feed. **NOTE:** If the specified directory already has some existing files that are not in sync with the Sentinel server files, these files are deleted when the new download starts.
Enable	Select this check box to activate the configuration to download the feed.

Option	Description
Schedule Interval	Specify the time interval for the download by selecting one of the following options from the drop-down list: ◆ **6 Hours:** Schedules the download for every six hours. ◆ **12 Hours:** Schedules the download for every 12 hours. ◆ **Daily:** Schedules the download every day. ◆ **Weekly:** Schedules the download once a week. ◆ **Monthly:** Schedules the download once a month.
Feed Type	Select Advisor from the drop-down list to notify the Advisor processes on the Sentinel server that the Advisor feed is downloaded.

4 (Optional) Click *Validate* to validate the URL and the login credentials.

The URL and its credentials are validated and a confirmation message is displayed. If the validation fails, you must provide a valid URL and login credentials.

5 Click *Save* to save the configuration settings.

The specified directory path is validated and the download configuration that you created is displayed in the Download Manager window. If the directory path is invalid, you are prompted to specify a valid directory path.

10.3 Editing a Download Configuration

1 Open the Download Manager window by doing either of the following:

- Click *Tools* > *Download Manager*.
- Click the icon on the tool bar.

2 Select the download configuration that you want to edit, then click *Edit*.

The Edit window is displayed.

3 Edit the required information.

4 (Optional) Click *Validate* to validate the URL and the login credentials.

5 Click *Save* to save the configuration settings.

The configuration settings are updated and displayed in the Download Manager window. If the download status is changed to Enable, the Advisor feed is downloaded at the specified time interval.

10.4 Downloading the Feed Instantly

You can instantly download the feed for a selected download configuration, regardless of the time interval scheduled for the selected download configuration.

1 Open the Download Manager window by doing either of the following:

- Click *Tools* > *Download Manager*.
- Click the icon on the tool bar.

2 Click *Download Now*.

The Advisor data feed is downloaded if you have specified the appropriate URL.

NOTE: The *Download Now* button is disabled if the download is in progress for the selected configuration.

While the downloaded files for a download are being processed, other download requests are queued internally. After the downloaded files are processed, the queued download request will initiate and be reflected in the download manager.

10.5 Deleting a Download Configuration

1. Open the Download Manager window by doing either of the following:
   - Click *Tools* > *Download Manager*.
   - Click the icon on the tool bar.
2. Select the download configuration that you want to delete, then click *Delete*.

   A message is displayed to confirm whether you want to delete the selected configuration.
3. Click *Yes* to confirm deletion.

   The selected configuration is deleted.

10.6 Audit Events for the Download Manager

The Download Manager generates an audit event whenever you perform any of the following actions:

- Create a download configuration
- Edit a download configuration
- Delete a download configuration
- Download the feed

NOTE: An audit event is generated regardless of whether the download status indicates success or failure.

For more information on the audit events for the Download Manager, see Section B.2, "Download Manager Audit Events," on page 432.

You can create appropriate correlation rules to receive notifications of these audit events through e-mail. For more information on creating correlation rules, see Section 4.3, "Correlation Rules," on page 85.

Event Source Management

11

The Event Source Management (ESM) panel provides a set of tools to manage and monitor connections between Sentinel and the event sources that are providing data to Sentinel. The graphical interface shows at a glance the current event sources and the software components that are processing data from that event source. Each component can be easily deployed to quickly integrate the devices in the enterprise, and then can be monitored in real time within the ESM interface.

- Section 11.1, “Understanding Event Source Management,” on page 189
- Section 11.2, “Introduction to the User Interface,” on page 190
- Section 11.3, “Live View,” on page 198
- Section 11.4, “Components of Event Source Hierarchy,” on page 202
- Section 11.5, “Debugging,” on page 220
- Section 11.6, “Exporting a Configuration,” on page 228
- Section 11.7, “Importing a Configuration,” on page 230
- Section 11.8, “Event Source Management Scratchpad,” on page 234

11.1 Understanding Event Source Management

You need to have appropriate permissions to access this tab. Only a Sentinel Administrator has controls to enable/disable access to the ESM panel for other users.

- Section 11.1.1, “Using Event Source Management,” on page 189
- Section 11.1.2, “Plug-In Repository,” on page 190
- Section 11.1.3, “Auxiliary Files,” on page 190

11.1.1 Using Event Source Management

Through ESM, you can:

- Add/edit connections to event sources by using Configuration Wizards.
- View the real-time status of the connections to event sources.
- Import/export configuration of event sources to or from Live View/Scratchpad.
- View and configure Connectors and Collectors that are installed with Sentinel
- Import/export Connectors and Collectors from or to a centralized repository
- Monitor data flowing through the Collectors and Connectors
- Debug Collectors
- Design, configure, and create the components of the Event Source Hierarchy, and execute required actions using these components. For more information, see Section 11.3, “Live View,” on page 198.

11.1.2 Plug-In Repository

A plug-in is a package of code that provides additional functionality to Sentinel; ESM leverages two types of plug-ins called Collectors (scripts) and Connectors. Implementing these features as plug-ins allows Novell to deliver enhancements to our event collection system without the need to deliver a new version of the Sentinel platform.

- **Collector:** The Collector plug-in adds the ability to parse raw data from an event source. This is similar to the Collector in Sentinel 5; however, from Sentinel 6.x onward, the plug-in also provides additional metadata to enable the ESM panel to prompt the user for parameter values as well as enable ESM to automatically select supported connection methods that work well with the Collector. This metadata is added to the Collector plug-in by the plug-in developer. Collectors are written by using JavaScript or our legacy scripting language and as such are sometimes called scripts.
- **Connector:** In Sentinel Rapid Deployment, all Connectors are pluggable. A Connector plug-in contains both the implementation of the connection mechanism used to gather data from an event source as well as the GUI screens needed to configure the Connector. This allows for a user to easily add additional Connectors to Sentinel.
- **Hot Fixes and New Functionality:** In the future, some Sentinel enhancements and defect fixes might be available as plug-ins.
- After you import a plug-in into Sentinel, it is centrally stored in the Plug-In Repository. The appropriate Sentinel component on other machines automatically starts by using the plug-in.

11.1.3 Auxiliary Files

Some plug-ins, such as database Connectors, require one or more auxiliary files in order to function. Auxiliary files are typically files that can not be shipped by Novell within the standard plug-in, such as user-specific configuration files or third-party libraries that require specific licenses. In all cases the documentation for the plug-in includes detailed instructions about which auxiliary files are necessary and where they can be obtained.

11.2 Introduction to the User Interface

The ESM Live View and Scratchpad are independent windows. This allows you to work on other tabs in Sentinel as you work on ESM.

The Event Source Management windows include:

- A menu bar with the ESM menus
- A toolbar that helps you execute the functions of ESM
- Several different types of frames to display ESM data
- A Display Health Monitor frame with graph and table views where you can perform your activities

Figure 11-1 *Event Source Management Live View*

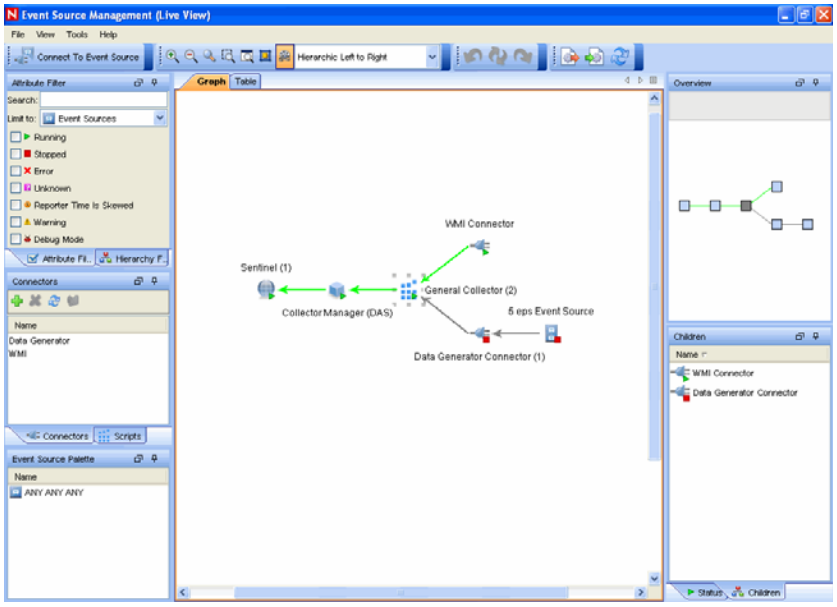

11.2.1 Menu Bar

The menu bar has *File*, *View*, *Tools*, and *Help* options.

Figure 11-2 *Event Source Management Menu Bar*

The following are the options available in the each of the menu bar options that are described in the document:

- *File*
  - *Export Configuration*
  - *Import Configuration*
  - *Save Preferences*
  - *Close*
- *View*
  - *Reset Layout*
  - *Redo Layout*
  - *Undo Layout*
- *Tools*
  - *Connect to Event Source*
  - *Import Plugin*
  - *Attachment Viewer Configuration*

- Help
  - About
  - Help

11.2.2 Toolbar

Table 11-1 *Event Source Management User Interface*

User Interface	Description
	Launch the wizard for connecting to a new event source
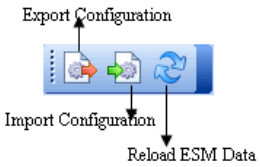	Import/Export, Reload Event Source Management configurations, and plug-ins.
	The toolbar contains several tools for displaying objects in ESM. You can zoom the entire graphical view in and out, or zoom directly to a selected region. The magnifying glass allows you to enlarge the text and icons for a small portion of the graphical view without affecting the overall zoom level. The Fit to Screen option adjusts the ESM view to fit the screen.
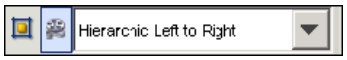	You can select from several different layouts to display the objects in ESM. You can also enable/disable animations during transition from one layout to the other in the graphical view of the Health Monitor display.
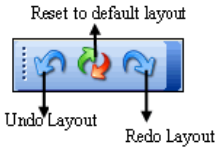	You can also reset to the default settings.

11.2.3 Zoom

In ESM, you can use magnifying glass to zoom into a region.

TIP: To enable or/ disable the magnifying glass in ESM, use the magnifying glass button on the toolbar.

You can increase or decrease the magnification factor with the following key combinations:

- **To increase the size of the size of the magnification glass cursor:** Ctrl key + backward scrolling of the mouse wheel
- **To decrease the size of the size of the magnification glass cursor:** Ctrl key + forward scrolling of the mouse wheel
- **To Zoom in:** Forward movement of the mouse wheel
- **To Zoom out:** Backward movement of the mouse wheel

NOTE: The magnifying glass is available only in the graphical view of the ESM window.

11.2.4 Frames

You can see the following frames in the Live View or Scratchpad window:

- "Attribute Filter" on page 193
- "Hierarchy Filter" on page 194
- "Connectors" on page 194
- "Scripts" on page 195
- "Event Source Palette" on page 196
- "Children" on page 196
- "Status Details" on page 197
- "Overview" on page 197

Attribute Filter

The Attribute filter allows you to display the components of ESM. You can specify the components to be displayed based on the component name and status.

Figure 11-3 *Attribute Filter Frame*

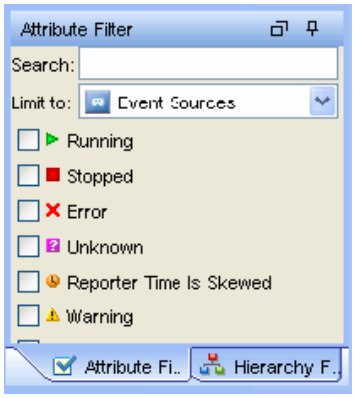

- **Text Filter:** Allows you to filter the nodes that are displayed in the graphical and tabular view based on the text they type in.
- **State Filter:** Allows you to filter the nodes that are displayed in the graphical and tabular view based on the current state of the node.

Hierarchy Filter

The Hierarchy filter sets the display based on the hierarchy you select in this frame. It allows the user to filter the nodes that are displayed in the graphical and tabular view based on the node hierarchy. All children and parents of selected nodes are shown.

Figure 11-4 *Hierarchy Filter Frame*

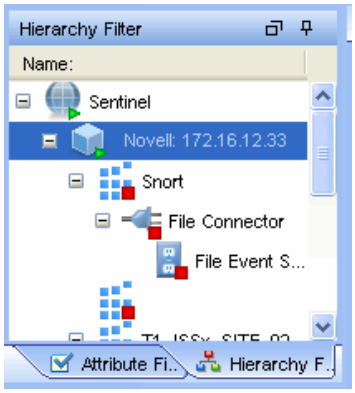

To set Hierarchy filter for displaying components:

1. In Sentinel Control Center, click *Event Source Management* in the menu bar and select *Live View or Scratch Pad*.
2. Click the Hierarchy Filter frame.
3. Select the hierarchy level to display the components.

Connectors

Connectors are plug-ins in Sentinel. Importing a Connector implements the Connector mechanism in the system. The Connectors frame allows you to add, remove, and refresh Connectors and add auxiliary files in the system.

Figure 11-5 *Connector Frame*

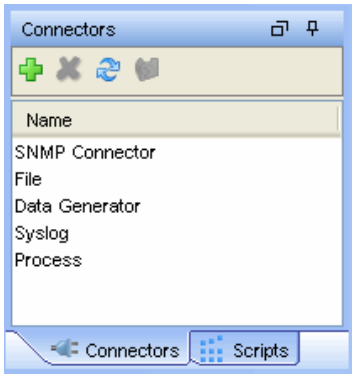

Table 11-2 Connector Frame Icons

Icon	Name	Description
	Add	Adds Connectors to the system.
	Delete	Deletes Connectors.
	Refresh	Refreshes the list.
	Add Auxiliary Files	Adds auxiliary files. For more information, see Add Auxiliary Files.

To add Connector plug-ins:

1. In Sentinel Control Center, click *Event Source Management* in the menu bar and select *Live View or Scratch Pad*.
2. Click the Script or Connectors frame. You can plug-in Connectors from here. For more information, see "Adding Connectors/Collector Plug-Ins" on page 204.

Scripts

Collectors are plug-ins in Sentinel. Collector plug-ins add the ability to parse raw data from a particular event source. The Scripts frame is used to manage the importing and updating of Collectors (also called scripts) into Sentinel.

Figure 11-6 Scripts Frame

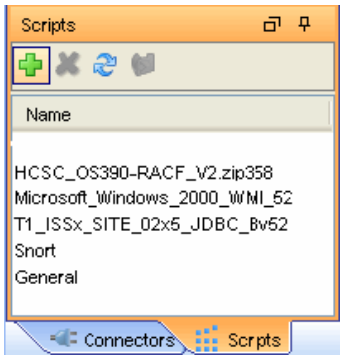

Table 11-3 Scripts Name Icons

Icon	Name	Description
	Add	Adds scripts (Collectors) to the system.
	Delete	Deletes Collectors.
	Refresh	Refreshes the list.

Icon	Name	Description
	Add Auxiliary Files	Adds auxiliary files. For more information, see Section 11.1.3, “Auxiliary Files,” on page 190.

To add Collector plug-ins:

1. In Sentinel Control Center, click *Event Source Management* in the menu bar and select *Live View or Scratch Pad*.
2. Click the Script or Connectors frame. You can import Collectors from here. For more information, see “Adding Connectors/Collector Plug-Ins” on page 204.

Event Source Palette

This frame displays the list of devices or event sources supported by the existing Collectors in the Central Repository. Each Collector ships with meta-information that describes the list of event source types supported by that Collector: This information is compiled to provide the data in this palette. The supported devices for a particular Collector might not necessarily be the same as the name of the Collector.

Figure 11-7 *Event Source Palette*

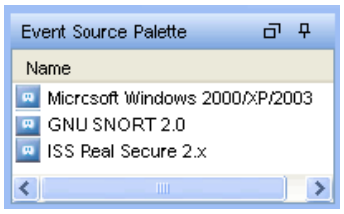

Children

This frame displays names of immediate children nodes of a parent (main) node when you click the parent node. This frame is useful for managing children of nodes that have been collapsed in the graphical view. To perform any action in ESM, right-click a component and select from options listed. For more information, see Section 11.3.3, “Right-Click Menu,” on page 200.

Figure 11-8 *Children Frame*

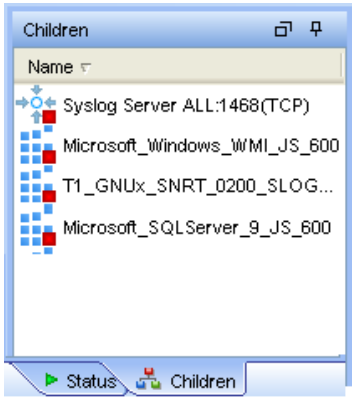

Status Details

This frame displays the status details of a selected component in the Health Monitor Display frame.

Available status information includes the current state, the number bytes processed, the number of records sent, the number of Sentinel events sent, and various other status and statistical information.

NOTE: The status information varies based on the type of component that is selected.

Figure 11-9 *Status Details Frame*

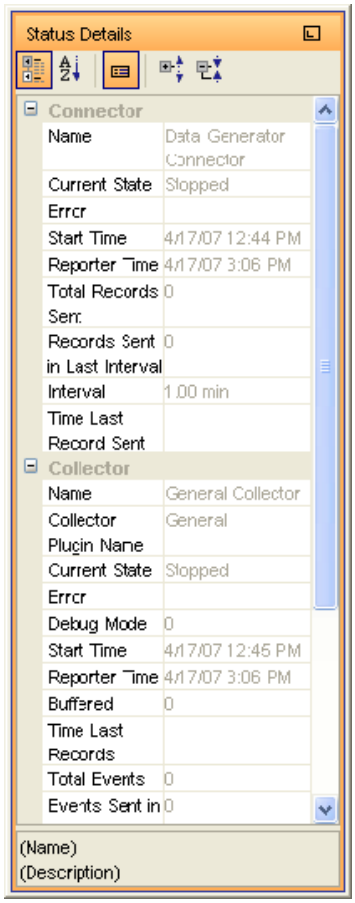

Overview

The Overview frame allows you to quickly move across the graphical view. This is particularly useful when there are many objects in the screen.

Figure 11-10 Overview Frame

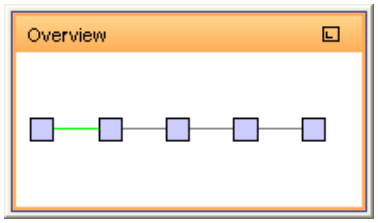

11.3 Live View

The ESM panel provides the main user interface to Event Source Management. You can view configuration data in a graphical or tabular view.

11.3.1 Graphical ESM View

The graphical view of ESM is the default view in Event Source Management. In the graphical view, you can view the status of a Collector and access the configuration settings of Collectors and Collector related objects as a graph of connected nodes.

Figure 11-11 Graphical View

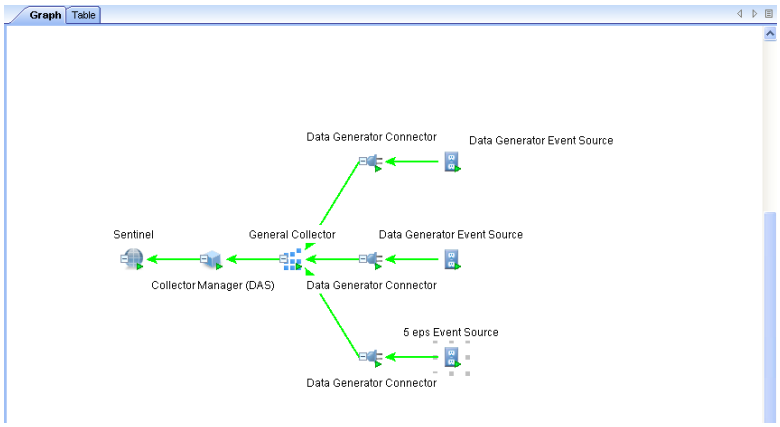

By default, the Health Monitor Display frame displays in the graphical view. The data can be displayed in seven different layouts. The default layout in graph is the “Hierarchic Left to Right” layout. You can change between these layouts by selecting the layout format from the drop-down list in the toolbar.

Figure 11-12 Layout Selection

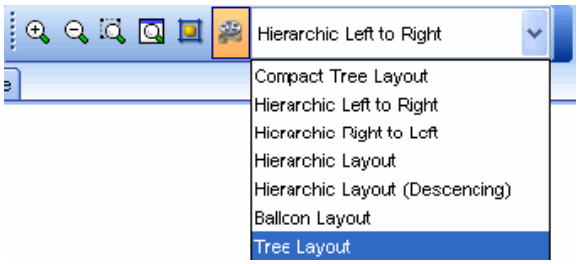

TIP: Click in the graphical ESM view and use “+” or “-” to zoom in or zoom out. Alternatively, use the mouse wheel to zoom in and zoom out.

In the graphical view, the lines connecting the components are color-coded to indicate data flow.

- **Green Line:** Indicates that data is flowing between the components.
- **Grey Line:** Indicates that the connection is not live and there is no data flow.
- **Blue dashed Line:** Indicates the logical relation of event source servers to their associated Collector Managers and event sources.

The following terminology is used for nodes:

- **Parent Node:** A node from which child nodes originate
- **Immediate Children:** The sub-nodes that are logically and functionally linked to a parent node.
- **Collapsed/Expanded nodes:** To improve the manageability and performance of the graphical display, Sentinel automatically contracts any node with 20 or more immediate children. This is especially useful for Connectors such as Syslog or Novell Audit that have the ability to automatically configure a large number of event sources.

  TIP: Collapsed nodes are identified by a “-” sign on the node and expanded nodes are identified by a “+” sign.

  Double-click a node to expand or collapse it.

In a collapsed state, a node displays the number of immediate children next to the node; for example, WMI Connector (3) [Collector name (Number of immediate children)]. The Children panel of a contracted node shows the immediate children of that node, each of which can be managed in the same way as nodes in the tabular ESM view.

NOTE: An event source server node does not have a “+” or “-” after its name even if it contains children.

Double-clicking a parent node changes the state from collapsed to expanded and vice versa. Double-clicking a node with no children displays the status details for that node. If an additional node is added to an expanded parent with over 20 children, the node is automatically collapsed. If an additional node is added to a manually expanded parent with over 20 children the node not automatically collapsed.

The parent node can take several minutes to expand if the parent node has a large enough number of child nodes to potentially cause the UI to become unresponsive; an alert message displays on the user interface to warn you about the delay in response. Click *Yes* to continue.

Figure 11-13 *Expand Selected Node Prompt*

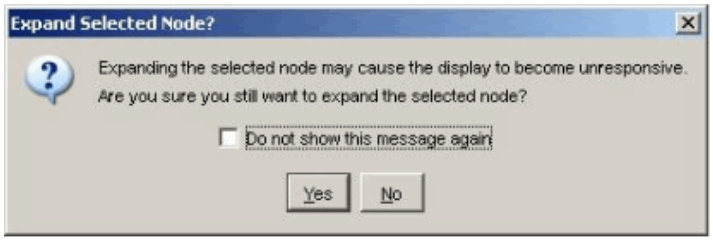

If you choose not to show this message again, the preferences are saved on that machine and any user logging into Sentinel from that machine does not get an alert again.

11.3.2 Tabular ESM View

The components visible in the graphical view of ESM can also be viewed in tabular format. In the tabular view, you can view the status of a Collector in a table and access the configuration settings of Collectors and Collector-related objects.

Figure 11-14 *Tabular View*

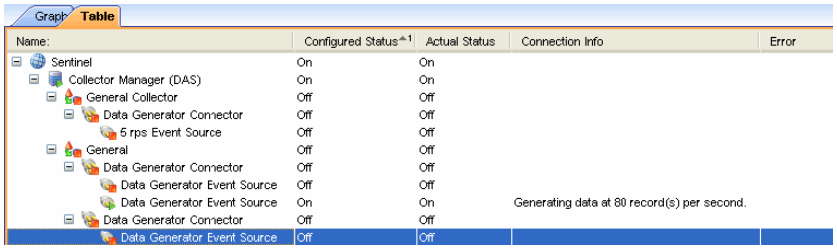

Name:	Configured Status	Actual Status	Connection Info	Error
Sentinel	On	On		
Collector Manager (DAS)	On	On		
General Collector	Off	Off		
Data Generator Connector	Off	Off		
5 rps Event Source	Off	Off		
General	Off	Off		
Data Generator Connector	Off	Off		
Data Generator Event Source	Off	Off		
Data Generator Event Source	On	On	Generating data at 80 record(s) per second.	
Data Generator Connector	Off	Off		
Data Generator Event Source	Off	Off		

The columns in the ESM tabular view are:

- **Configured Status:** The On state the object is configured to be in. This is the state that is stored in the database and it does not necessarily match the actual On state of the object. For example, the two states do not match if a parent object is turned off or if there is an error.
- **Actual Status:** The On state of the object as being reported by the actual running Collector Manager.
- **Connection Info (populated for event Sources only):** A textual description of the event source connection.
- **Error:** A textual description of an error that occurred in the running object.

TIP: Use the *Table*/*Graph* tabs to change to tabular or graphical views respectively.

11.3.3 Right-Click Menu

The Health Monitor Display view provides a set of right-click menus that help you execute a set of actions, as described below:

NOTE: The right-click actions available depend on the kind of object you clicked.

- **Status Details:** Displays all information known about the status of the selected object.
- **Start:** Sets the object to be running.

  NOTE: The selected object starts only after the parent nodes start and are running.

- **Stop:** Stops the running object.
- **Edit:** Modifies the editable information (Filter information, Object name and so on) with this option.
- **Debug:** Debugs the Collector. You must stop the running Collector before you debug it.

- **Move:** Moves the selected object from its current parent object to another parent object. You can move objects between the views; that is, move from the Live View to the Scratchpad and vice versa.
- **Clone:** Creates a new object that has its configuration information prepopulated with the settings of the currently selected object. This allows you to quickly create a large number of similar event sources without retyping the same information over and over again. You can clone objects between the views; that is, move from the Live View to the Scratchpad and vice versa. Cloning an object copies all the settings except the Run status. New objects created by using the Clone command are always in the Stopped state after creation.
- **Remove:** Deletes a selected object from the system.
- **Contract:** Collapses the child nodes into this node. This option is only available on parent nodes that are currently in an expanded state.
- **Expand:** Expands the child nodes of this node. This option is only available on parent nodes that are currently in a collapsed state.
- **Add Collector:** Opens an Add Collector Wizard that guides you through the process of adding a Collector to the selected Collector Manager.
- **Add Connector:** Opens an Add Connector Wizard that guides you through the process of adding a Connector to the selected Collector.
- **Add Event Source:** Opens an Add Event Source Wizard that guides you through the process of adding an event source to the selected Connector.
- **Open Raw Data Tap:** Displays the live stream of raw data from an event source or flowing through the selected object.
- **Open Active View:** Opens an Active View window that only displays events that have been generated by data from or flowing through the selected object.
- **Zoom:** Zooms in the graphical view display on the selected object.
- **Show in Tabular/Graphical View:** Switches over to the other view (to tabular view if you are in the graphical view, or to graphical view if you are in the tabular view) and automatically selects the object that is selected in the current view. When switching to graphical view, it also zooms in on the selected object.
- **Raw Data Filter:** Allows you to filter the raw data flowing through the selected node. The raw data filter is available on Collectors, Connectors, and event sources. If a filter is specified to drop data, the data to be dropped is not passed to the parent node and, therefore, is not converted into events.
- **Import Configuration:** Imports the configuration of ESM objects.
- **Export Configuration:** Exports the configuration of ESM objects
- **Add Event Source Server:** Allows you to add event source server to the selected Collector Manager
- **Add Collector Manager:** In Scratchpad mode, you can add a Collector Manager to the Scratchpad by using this option. In the Live view, Collector Manager objects are created automatically as each Collector Manager connects to the Sentinel system.

When you select multiple objects in the ESM panel and right-cliches following options are available:

- **Start:** Starts all the objects
- **Stop:** Stops all the objects

- **Remove selected objects:** Removes the selected object along with its children

TIP: Press Shift and click the object to select multiple objects.

11.4 Components of Event Source Hierarchy

ESM displays the information on the Collectors and other components in a hierarchy specific to ESM.

Figure 11-15 *ESM Hierarchy*

NOTE: ESM allows you to add Collectors, event sources, and Connectors.

Table 11-4 *Components of the ESM Hierarchy*

Icon	Name	Description
	Sentinel	The single Sentinel icon represents the main Sentinel server that manages all events collected by the Sentinel system. The Sentinel object is installed automatically through the Sentinel installer.
	Collector Manager	Each Collector Manager icon represents another instance of a Collector Manager process. Multiple Collector Manager processes can be installed throughout the enterprise. As each Collector Manager process connects to Sentinel, the objects are created in ESM automatically.
	Collector	Collectors instantiate the parsing logic for data from a particular event source. Each Collector icon in ESM refers to a deployed Collector script as well as the runtime configuration of a set of parameters for that Collector.
	Connector	Connectors are used to provide the protocol-level communication with an event source, using industry standards like Syslog, JDBC, and so forth. Each instance of a Connector icon in ESM represents the Connector code as well as the runtime configuration of that code.
	Event Source	An event source server (ESS) is considered part of a Connector, and is used when the data connection with an event source is inbound rather than outbound. The ESS represents the daemon or server that listens for these inbound connections. The ESS caches the received data, and one or more Connectors connects to the ESS to retrieve a set of data for processing. The Connector requests only the data from its configured event source (defined in the metadata for the event source) and that matches additional filters.

Icon	Name	Description
	Event Source Server	The event source represents the actual source of data for Sentinel. Unlike other components this is not a plug-in, but is a container for metadata, including runtime configuration, about the event source. In some cases a single event source could represent many real sources of event data, for example if multiple devices are writing to a single file.

11.4.1 Component Status Indicators

Indicators are used to represent various states as follows:

Table 11-5 *Component Status Indicators*

Icon	Name	Description
	Stopped	Indicates that the component is stopped.
	Running	Indicates that the component is running.
	Warning	Indicates that a warning is associated with the component. At this time, this warning indicator is primarily used to show when the configured state and actual state of a component differ, that is, a component is configured to be running, but the actual state of the component is stopped.
	Error	Indicates that an error is associated with the component. See the individual component's status display for details about the error.
	Reporter Time is Skewed	Indicates when the time of a component differs from the main server's time. The difference is greater than a predefined time threshold.
	Debug	Indicates that the component is in Debug mode. Only a Collector can be in Debug mode.
	Unknown	This indicator is displayed when the status of the object in the ESM panel is not yet known.

To set an attribute filter for displaying components:

1. In the Sentinel Control Center, click *Source Management* in the menu bar and select *Live View* or *Scratch Pad*.
2. Click the Attribute Filter frame.
3. Specify the *Search* and *Limit to* criteria.
4. Select the *Running* or *Stopped* check box to specify the status of the components.

To hide components based on type:

1. In the Sentinel Control Center, click *Event Source Management* in the menu bar and select *Live View* or *Scratch Pad*.
2. Click the Attribute Filter frame.

3 Specify the *Search* and *Limit to* criteria.

4 Select the component type by which to limit the view.

11.4.2 Adding Components to the Event Source Hierarchy

Although some Sentinel components are preinstalled with the Sentinel system, Novell recommends that you check the Sentinel Content Web site (http://support.novell.com/products/sentinel/sentinel61.html) for updated versions.

Collectors, Connectors and event sources can be added to the system through the right-click menus on the main ESM display.

11.4.3 Collectors

To run the Collectors and generate the events as per your requirements, you need to:

- Download Collectors
- Import and Deploy Collectors
- Generate Events

  Right-click the Collector and select Start to generate events.
- Debug Collectors

  For any errors in the output of a Collector, right-click the Collector and select *Debug*.

  For more information, see Section 11.5, “Debugging,” on page 220.
- Edit Collectors

  To troubleshoot any problems with a Collector, you can edit the Collector. The method for editing the Collector depends on the type of Collector. For proprietary (or legacy) Collectors, copy the Collector script to a Windows machine that has Collector Builder installed. For JavaScript Collectors, any standard development environment for JavaScript can be used.

  For more information on editing Collectors, see the Sentinel Collector SDK (http://www.novell.com/developer/develop_to_sentinel.html).
- Re-Import and deploy Collectors

Adding Connectors/Collector Plug-Ins

NOTE: When you use the Sentinel Control Center to browse to locate a file on the Desktop of the Collector Manager, clicking *Desktop* takes you to the desktop of the user running the Collector Manager, usually SYSTEM. Extra steps might be necessary to navigate to the correct user’s desktop.

To add a Connector:

1 Click *Tools* on the menu bar and select *Import Plugin*. The Import Plugin Wizard window displays.

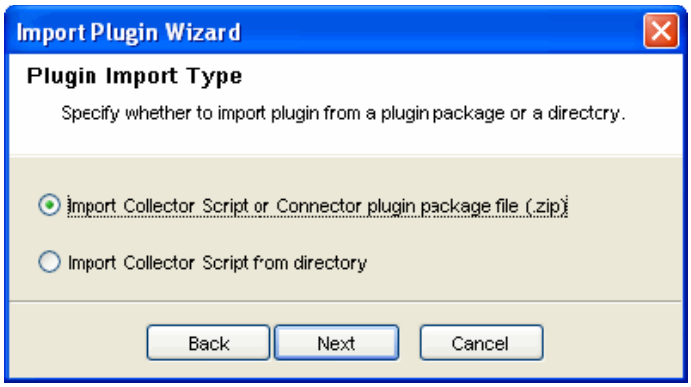

2 Select *Import Collector Script or Connector plugin package file (.zip)*. Click *Next*.

3 Browse to the location of the Connector plugin package file and click *OK*, then click *Next*.

If the file imported is not in the format specified for the Collector scripts or for the Connector plug-in package, the system displays an error message.

Plug-in details window displays.

4 Select the *Deploy Plugin* option to deploy the plug-in from this window.

For more information, see "To connect to the event sources:" on page 211.

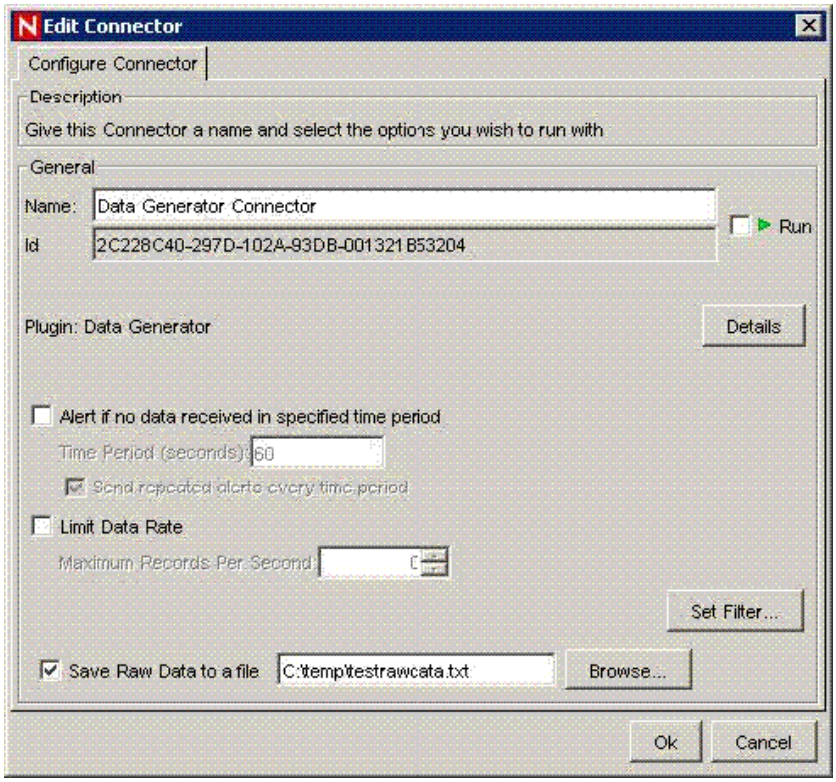

5 Click *Finish*.

When you add a plug-in into Sentinel, it is placed in the Plugin Repository, that enables Sentinel components on other machines to start using the plug-in without adding the plug-in separately.

To add a Collector plug-in:

1 Click *Tools* on the menu bar and select *Import plugin*. The Import Plugin Wizard window displays.

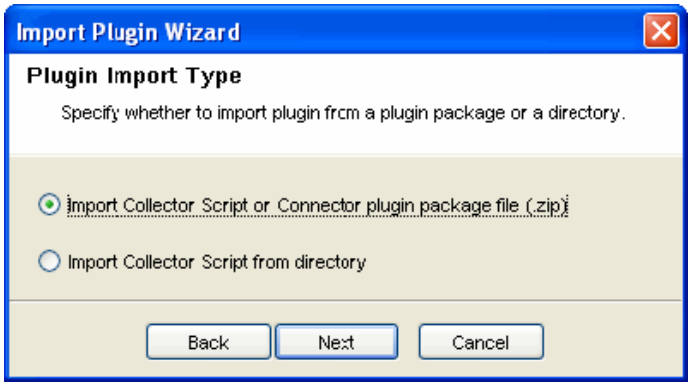

You can select from the two options available in this window.

2 Click *Next.*

3 Do one of the following:

- If you chose the first option, browse to a location of the Collector script file and click *OK*., then click Next.
- If you chose second option, you are directed to the Collector workspace. Select a Collector script directory and click *Next*.

The Collector Script Detail window displays.

4 Click the button next to the *ID* field to generate UUID.

The name and author details are displayed.

5 Edit the details as per your requirements. Specify a Version number.

6 Browse to and attach the help file.

If the help file is not in the plug-in directory, the system prompts you to copy the help file to the plug-in directory before the import. Click *Yes*.

7 Provide a description and click *Next*. The Supported Devices window displays.

You must specify at least one device.

8 Click *Add*. The Supported Devices window displays.

9 Provide a vendor, name, version, description, click *OK*, then click *Next*.

Use the *Edit* button to edit the details of a device or use the *Delete* button to delete a device from the list. The Plugin details window displays.

10 Select the *Deploy Plugin* option to deploy the plug-in from this window.

For more information on the deployment procedure, see "To connect to the event sources:" on page 211.

11 Click *Finish*.

Updating Connector/Collector Plug-Ins

If a new version of a Connector or Collector is released, you can update the Sentinel system and any deployed instances of the Connector or Collector.

NOTE: When you use the Sentinel Control Center to browse to locate a file on the desktop of the Collector Manager, clicking *Desktop* takes you to the desktop of the user running the Collector Manager, usually SYSTEM. Extra steps might be necessary to navigate to the correct user's desktop.

To update a Connector or Collector plug-in:

1. Click *Tools* and select *Import plugin*. The Import Plugin Wizard window displays.

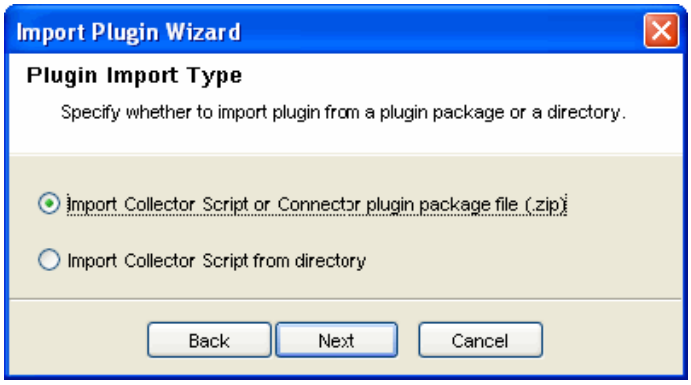

2. Select from the two options available in this window. Click *Next*.
3. Browse to a location of the Connector or Collector Plugin package file, select the file, click *OK*, then click *Next*.

   If the file imported is not in the format specified for the Collector scripts or for the Connector plug-in package, system displays an error message.
4. (Conditional) If you are updating an already-imported Connector or Collector, you are provided with the option of updating the existing plug-in, going back and selecting a different plug-in, or canceling the import. If you want to continue, click *Next*.

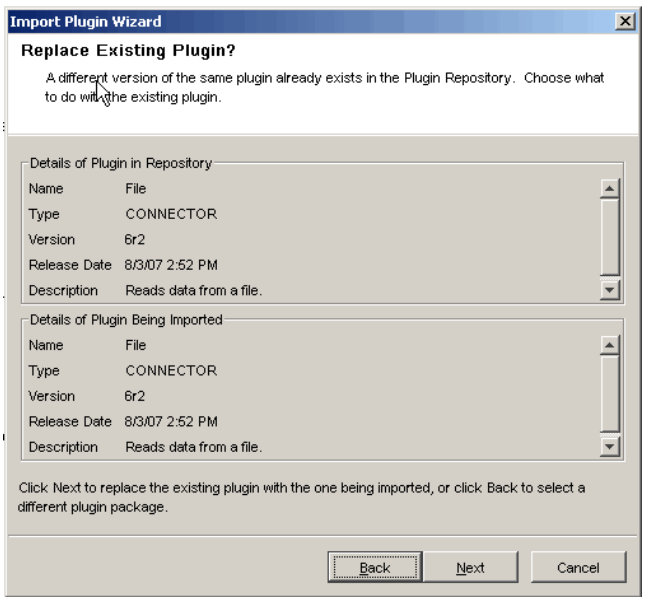

The Plugin details window displays.

5 Select the *Update Deployed Plugins* option to update any currently deployed plug-ins that use this Connector or Collector.

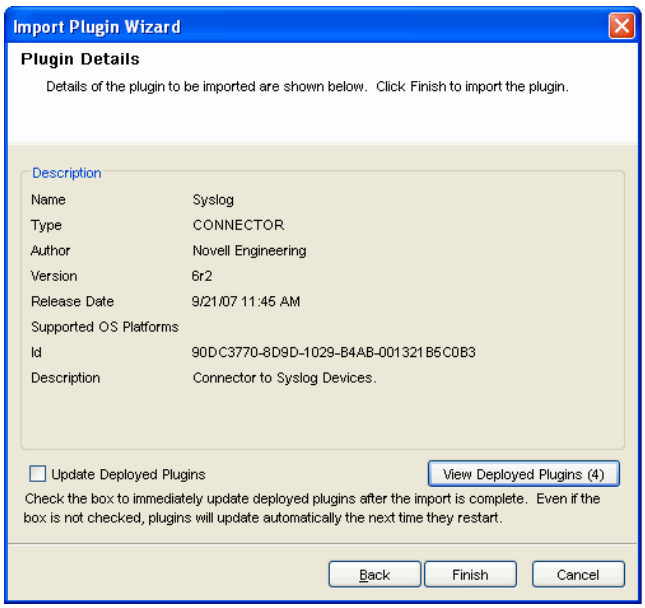

6 Click *View Deployed Plugins* to view the plug-ins deployed in the *ESM Live View*.

The number in parentheses represents the number of instances of this plug-in that are currently deployed and configured. The Deployed Plugins window displays the Affected Connectors/Event Sources/Event Source Servers or Affected Collectors. These are the components whose configuration is affected because of adding already existing Connectors/Collectors in ESM.

Description	User Interface
Affected Collectors	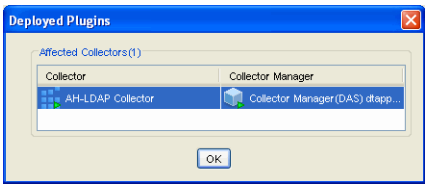
Affected Event Sources/ Connectors/ Event Source Servers:	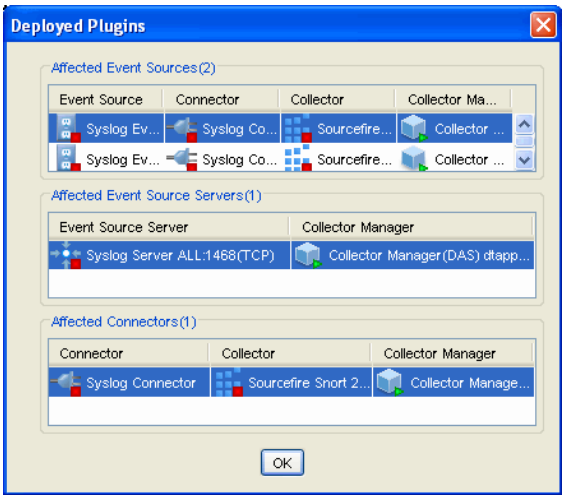

7 Click *Finish*.

NOTE: When you add a plug-in into Sentinel, it is placed in the Plugin Repository, which enables Sentinel components on other machines to start using the plug-in without adding the plug-in separately.

Deploying a Collector

1 In the main ESM display, locate the Collector Manager to which the new Collector is to be associated.

2 Right-click the Collector Manager and select the *Add Collector* menu item.

3 Follow the prompts in the Add Collector Wizard.

4 Click *Finish*.

NOTE: The Collector script enables the ESM panel to prompt you for parameter values as well as enabling ESM to automatically select supported connection methods that work well with the Collector script.

Deploying a Connector

1 In the main ESM display, locate the Collector to which the new Connector will be associated.

2 Right-click the Collector and select the *Add Connector* menu item.

3 Follow the prompts in the Add Connector Wizard.

4 Click *Finish*.

Deploying an Event Source

1 In the main ESM display, locate the Connector to which the new event source will be associated.

2 Right-click the Connector and select the *Add Event Source* menu item.

3 Follow the prompts in the Add Event Source Wizard.

4 Click *Finish*.

Deploying Event Source Servers

Certain event source Connectors (such as the Syslog Connector) require a process to collect data from the actual data source. These processes are called event source servers. They collect data from the data source and then serve it to the event source Connector. Event source servers must be added and associated to any event source Connectors that require a server.

1 In the Live View, right-click the Collector Manager and select *Add Event Source Server*. The Select Connector window displays.

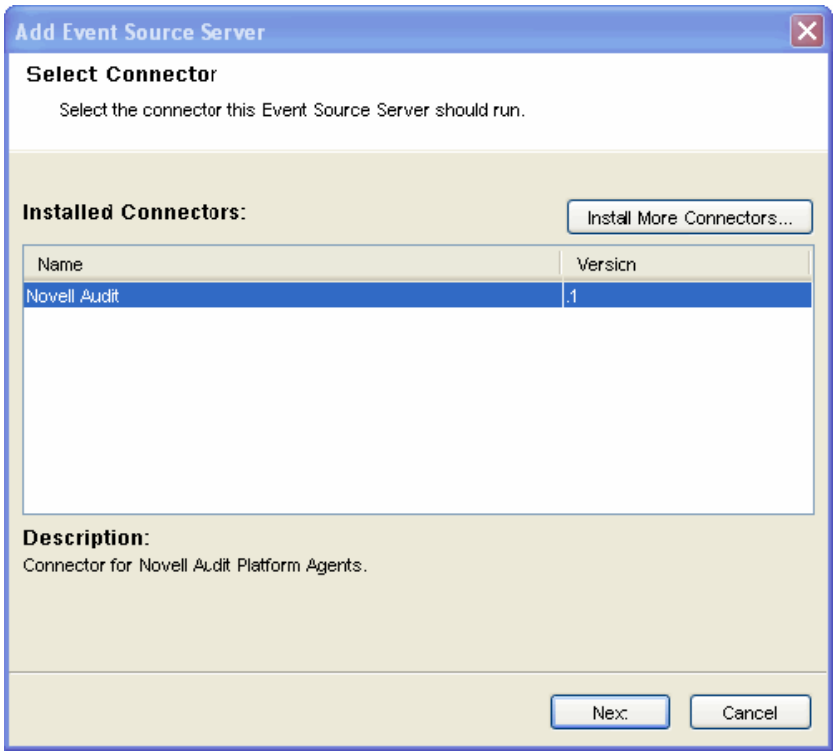

To start the Add Event Source Server Wizard, locate the Collector Manager on which the event source server process runs.

2 Select a Connector to support your device and click *Next*. If you do not have any Connectors in the list to support your device, click *Install More Connectors*.

For more information on installing a Connector, see "Adding Connectors/Collector Plug-Ins" on page 204.

3 Configure the various parameters for the server with reference to the Connector selected (For example, Syslog Connector, NAudit Connector, and so on.). The configurable parameters are different for the different Connector types.

4 Click *Next*.

5 Provide a name for the event source server. If you want this server to be running, select the *Run* check box.

6 Click *Finish*.

In the Health Monitor Display frame, the event source server added here displays with a dashed blue line showing the Collector Manager to which it is associated.

NOTE: This Add Event Source Server Wizard can also be initiated from within the Add Connector Wizard if a compatible event source server has not yet been added.

Connecting to an Event Source

There are several methods to configure an event source. Event sources can be deployed by right-clicking on an existing Collector Manager, Collector, or Connectors.

To deploy an event source, you need the following components:

- **Collector Script:** Collector scripts can be downloaded from the Sentinel Content Web site (http://support.novell.com/products/sentinel/sentinel61.html) copied from a previous Sentinel implementation (4.*x* or 5.*x)*, or built by using the Collector Builder.
- **Connector:** A Connector can also be downloaded from the Sentinel Content Web site (http://support.novell.com/products/sentinel/sentinel61.html). There are also some Connectors included in the installed Sentinel system, but there might be more recent versions on the Web site.
- Configuration information for the event source

To connect to the event sources:

1 Click *Tools* on the menu bar and select Connect to Event Source. Alternatively, click the *Connect to Event Source* button on the toolbar. The Connect to Event Source window displays.

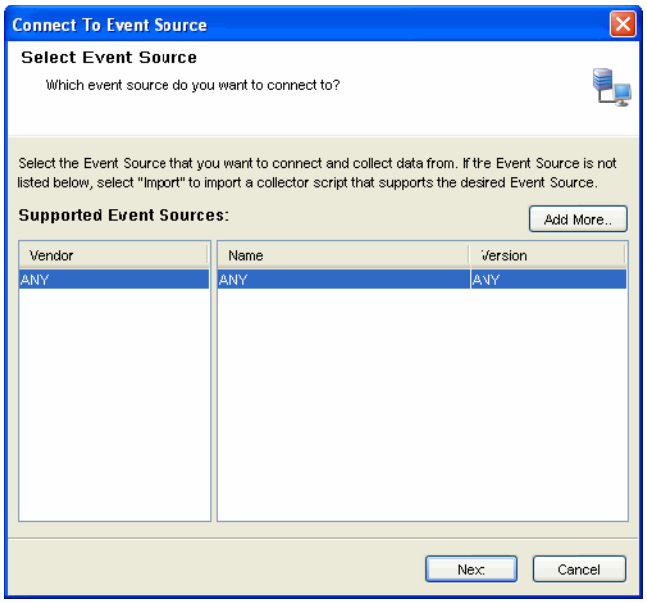

Event source types for which you currently have compatible Collector parsing scripts are listed here.

2 Select an event source from the list to which you want to connect to and collect data from. You can click *Add More* to import an event source.

3 Click *Next*. Select Collector Script window displays.

You can open the Select Collector Script window by double-clicking or dragging a selected event source from the Event Source Palette window.

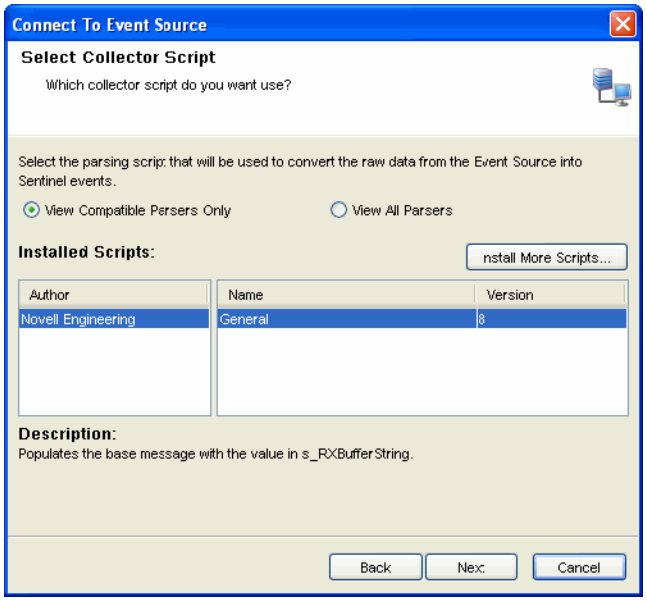

4 Select a Collector script from the list.

You can also install additional Collector scripts (click Install More Scripts) that support your event source, if it is not listed here. For more information on installing a Collector script, see "Adding Connectors/Collector Plug-Ins" on page 204.

5 Click *Next*. The Select Connection Method window displays.

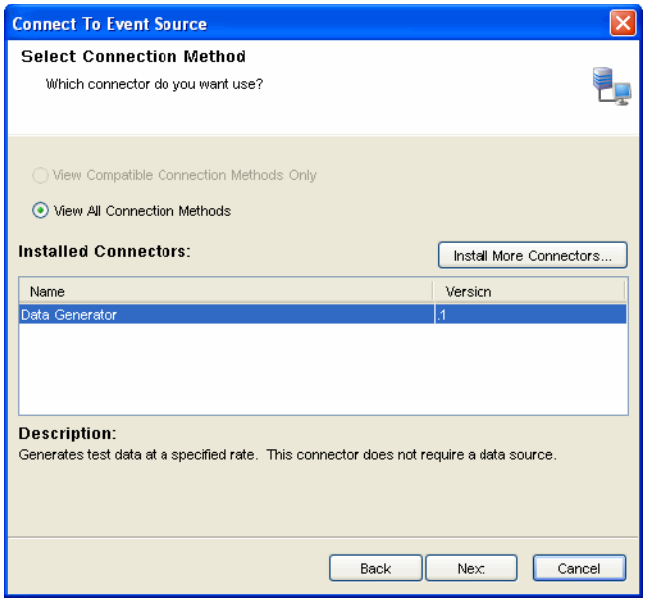

6 Select a connection method from the list.

You can also install additional Connectors by clicking on the Install More Connectors button. For more information, see "Adding Connectors/Collector Plug-Ins" on page 204 to install Connectors.

7 Click *Next*.The Event Source Management window displays.

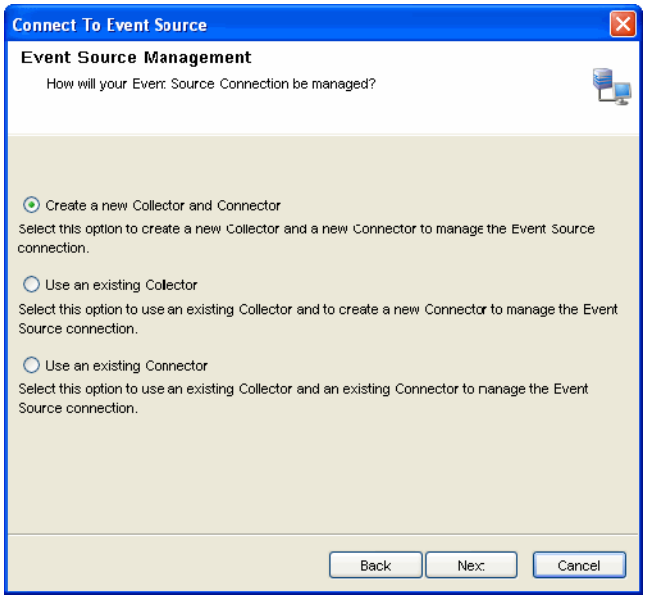

You can create a new Collector and Connector or you can use an existing Collector or Connector.

8 Select an option and click *Next*.

Based on the existing Collectors and Connectors in your system that is compatible with your new event source, one or more of these options might be unavailable.

- "Creating a new Collector and Connector" on page 216
- "Using an existing Collector:" on page 218
- "Using an Existing Connector" on page 219

9 Complete the configuration and click *Next*.

The Records Per Second window displays.

10 Set the number of records to be transferred per second and click *Next*.

The General window displays.

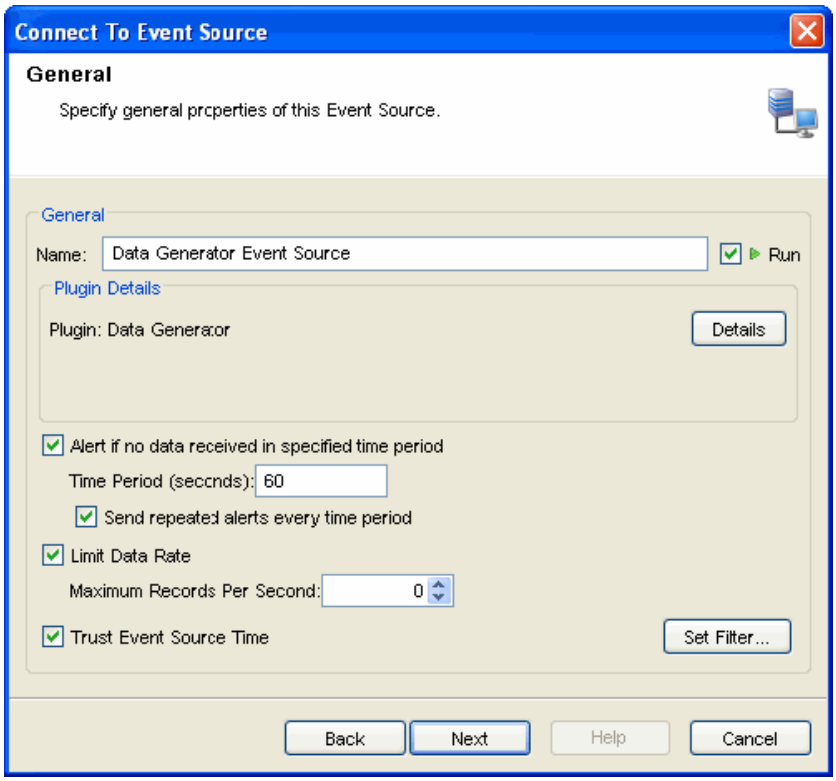

Options	Description
Name	Specify the name of the event source.
Run	Select the *Run* check box if you want to run your event source automatically.
Plugin Details	Click the *Details* button to see plug-in details.
Alert if no data is received in specified time period	Set alerts (with repeated option) indicating what to do if no data is received in a specified time interval.
Limit Data Rate	Limit the data rate as the maximum number of records per second.

Options	Description
Trust Event Source Time	Select *Trust Event Source Time* to display the Device Time (time when the event occurred) instead of the Event Source Time (time when the event was reported to the console).
Set Filter	Set the filter by using the *Set Filter* button. In the Filter window, add/edit the filters and click *OK*.

11 Click *Next*. The Summary window displays.

11a Click *Test Connection* to test the event source.

The Test Event Source window displays with Data and Error tabs. The *Error* tab displays the error message if there is any error in the configuration of event source.

After a few seconds, a sampling of raw data should be received from the event source and displayed in the tab.

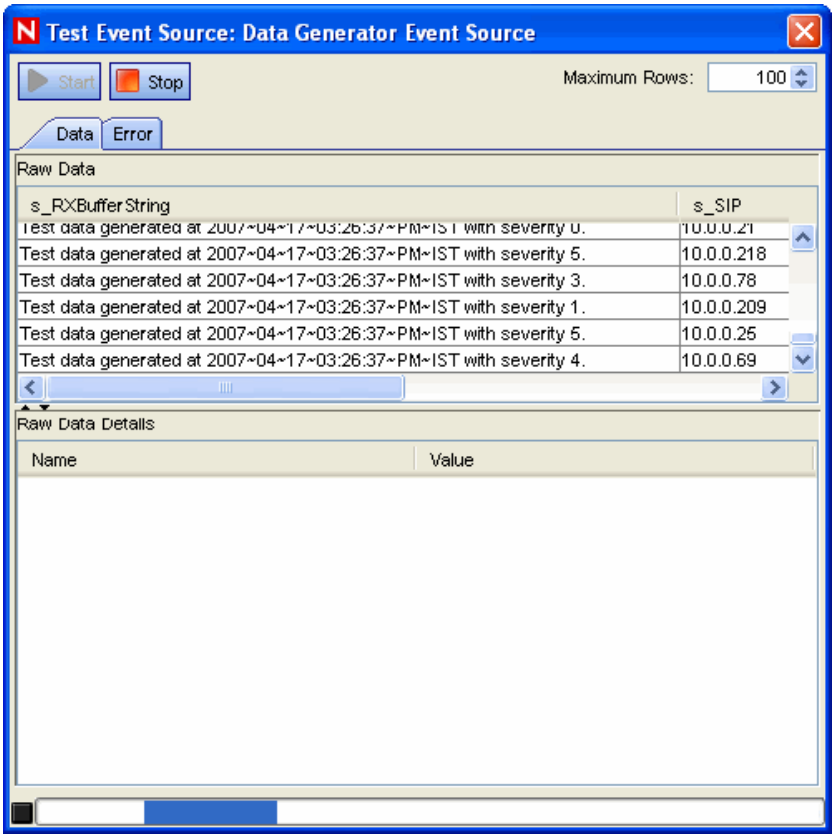

11b Click the *Start* and *Stop* buttons to start or stop the test.

Use the Maximum Rows component to control the maximum number of raw data records to obtain at once.

12 Click *Finish*.

NOTE: The Collector parsing script is executed on the same system as the Collector Manager that you select here.

1 In the Select Collector Manager window, select the Collector Manager you want to use and click *Next*.The Configure Collector Property window displays.

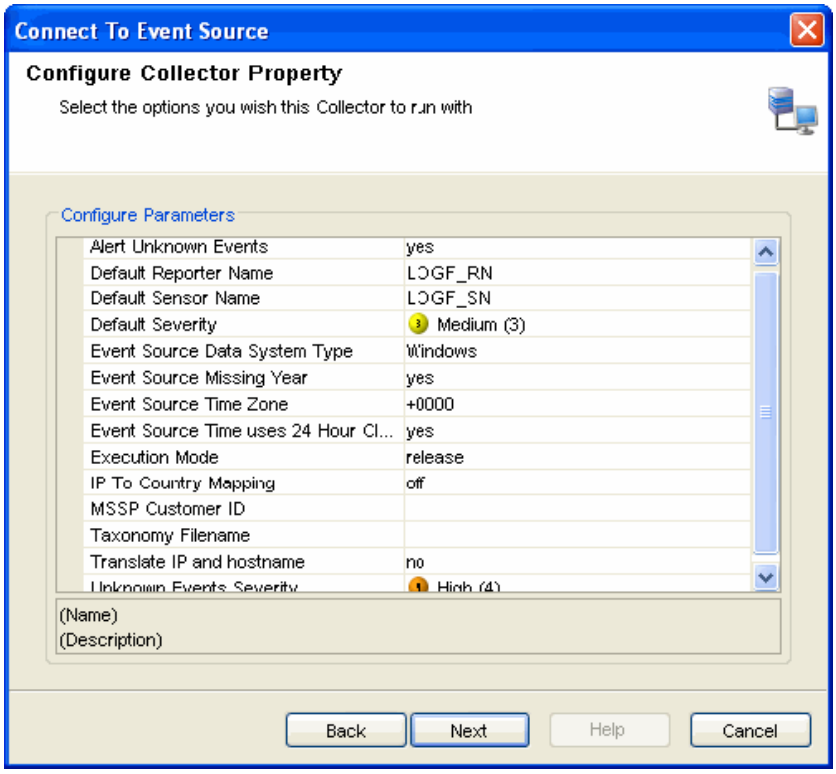

2 Configure the parameters available and click *Next*. The Configure Collector window displays.

3 Provide the name of the Collector and configure the options as desired:

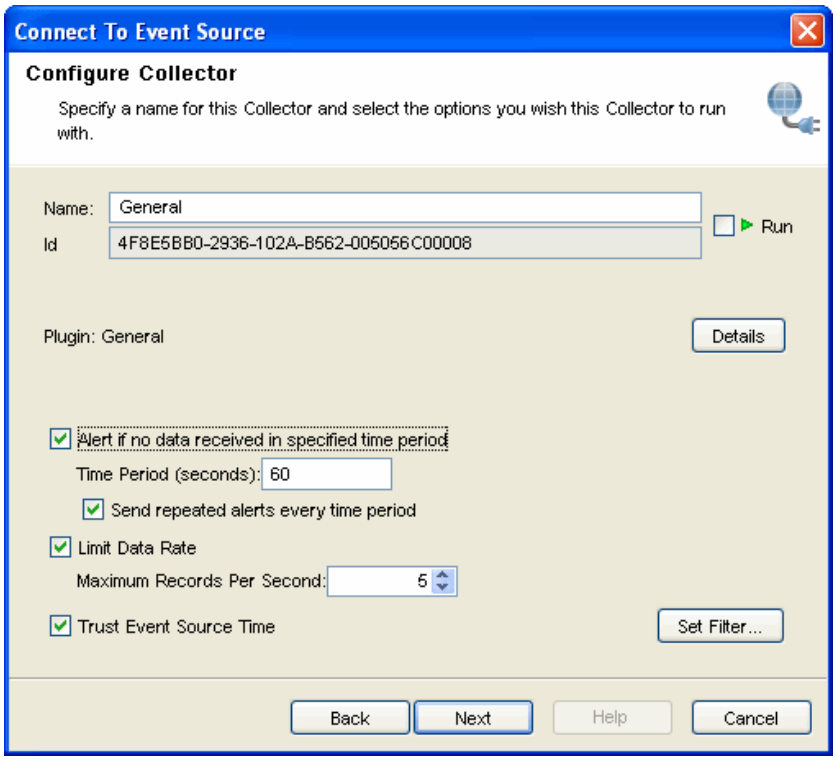

Options	Descriptions
Name	Specify the name of the event source.
Run	Select the *Run* check box if you want to run your Collector automatically.
Details	Click the *Details* button to see plug-in details.
Alert if no data is received in specified time period	Set alerts (with repeated option) indicating what to do if no data is received in a specific period.
Limit Data Rate	Limit the data rate as maximum number of records per second.
Set Filter	Set a filter by using the *Set Filter* button.
Trust Event Source Time	Select *Trust Event Source Time* to display the Device Time (time when the event occurred) instead of the Event Source Time (time when the event was reported to the console).

If the Trust Event Source Time option is selected, then all data flowing through the Collector has its Event Source Time trusted even if the event sources do not have this option selected.

4 Click *Next*. The Configure Connector window displays.

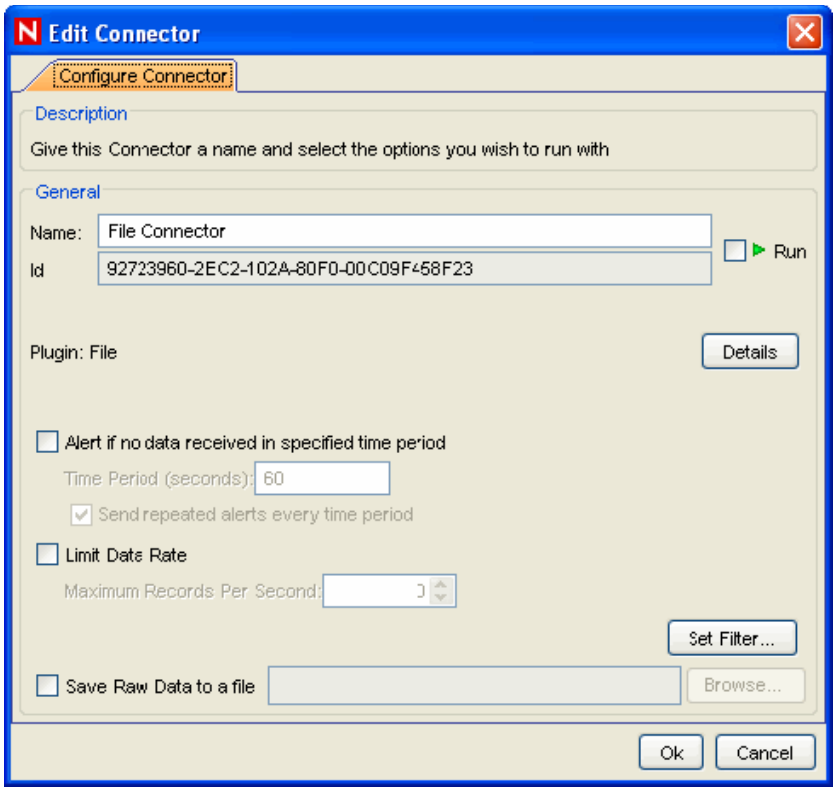

5 Provide the name of the Connector and configure the options as desired:

Options	Descriptions
Name	Specify the name of the event source.
Run	Select the *Run* check box if you want to run your Collector automatically.
Details	Click the *Details* button to see plug-in details.
Alert if no data is received in specified time period	Set alerts (with repeated option) indicating what to do if no data is received in a specific period.
Limit Data Rate	Limit the data rate as maximum number of records per second.
Set Filter	Set a filter by using the *Set Filter* button.
Trust Event Source Time	Select *Trust Event Source Time* to display the Device Time (time when the event occurred) instead of the Event Source Time (time when the event was reported to the console).

6 Click *Next*. The Event Source Configuration window displays.

7 Continue with Step 9 on page 214.

Using an existing Collector:

1 Select this option to use an existing Collector and to create a new Connector to manage the event source connection.

After you select this option and click *Next*, the Select Collector window displays.

2 Select the Collector you want to use and click *Next*. The Configure Connector window displays.

3 Provide the name of the Connector and configure the options as desired:

Options	Descriptions
Run	Select the *Run* check box if you want to run your Collector automatically.
Details	Click the *Details* button to see plug-in details.
Alert if no data is received in specified time period	Set alerts (with repeated option) indicating what to do if no data is received in a specific period.
Limit Data Rate	Limit the data rate as maximum number of records per second.
Set Filter	Set a filter by using the *Set Filter* button.
Trust Event Source Time	Select *Trust Event Source Time* to display the Device Time (time when the event occurred) instead of the Event Source Time (time when the event was reported to the console).

4 Click *Next*.The Event Source Configuration window displays.

5 Continue with Step 9 on page 214.

Using an Existing Connector

1 Select this option to use an existing Collector and an existing Connector to manage the event source connection.

After you select this option and click *Next*, the Select Connector window displays.

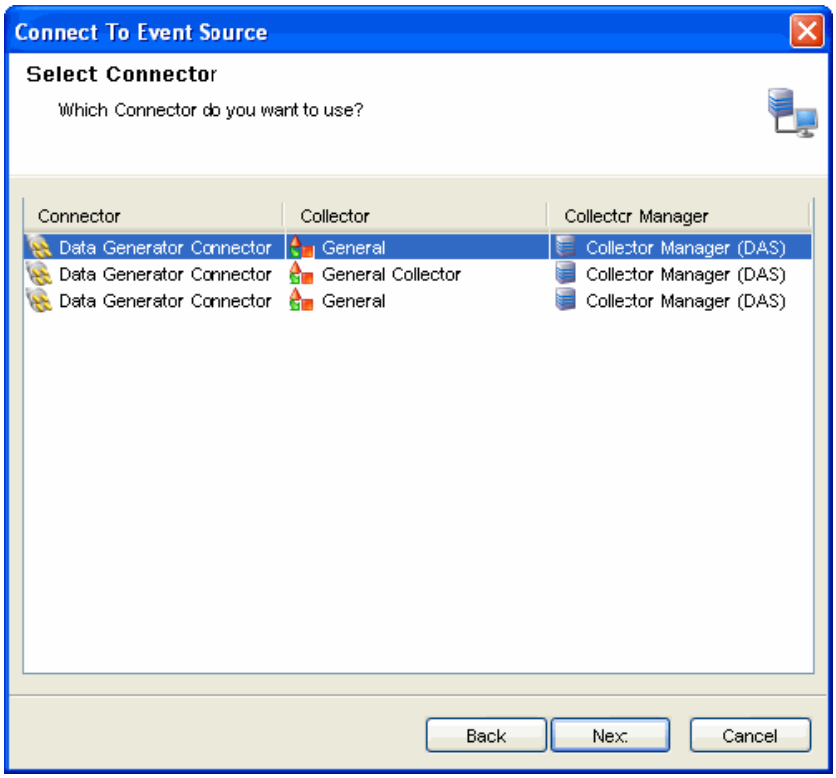

2 Select the Connector you want to use and click *Next*.

3 Continue with Step 9 on page 214.

11.5 Debugging

Sentinel's Collectors are designed to be easily customizable and to be created by customers and partners. There are two types of Sentinel Collectors: proprietary (or legacy) Collectors that are written in a language developed for Sentinel, and JavaScript Collectors. The debugging interface is slightly different for each type, and is intended to analyze the Collector code running in place on the Collector Manager For more information on customizing or creating new Collectors, obtain the Novell Developer Kit for Sentinel (http://developer.novell.com/wiki/index.php?title=Develop_to_Sentinel).

- Section 11.5.1, “Collector Workspace and Collector Directory,” on page 220
- Section 11.5.2, “Debugging Proprietary Collectors,” on page 221
- Section 11.5.3, “Debugging JavaScript Collectors,” on page 223
- Section 11.5.4, “Using the Raw Data Tap to Generate a Flat File,” on page 227

11.5.1 Collector Workspace and Collector Directory

Collectors are simple textual scripts that are run by a Collector Manager. The handling of these scripts is a bit complex:

- The code for all Collectors is stored in a Plugin Repository on the central Sentinel server when they are imported.

Location: `<Install Directory>\data\plugin_repository` on the Sentinel server.

- The runtime configuration for the Collector (when it is configured to run on a particular Collector Manager) is stored separately in the Sentinel database.
- When a Collector is actually started on the Collector Manager, the Collector plug-in is deployed to the Collector Manager in real time, the runtime configuration is applied, and the code is started. Any preexisting instance of the Collector code on that Collector Manager is overwritten.

  Location: `<Install Directory>\data\collector_mgr.cache\collector_instances` on each Collector Manager.
- In order to edit a Collector, you need to use the ESM Debuggger *Download* button, which copies the Collector to the local Collector Workspace on the client machine (the machine where you are running SCC). Edits are made against that local copy and then uploaded back into the central Plugin Repository.

  Location: `<Install Directory>\data\collector_workspace` on the client application machine.

11.5.2 Debugging Proprietary Collectors

The Debugging Collector window allows you to debug Collectors written in the Novell proprietary language. The left column on the debugger displays the commands for the current script state. The highlighted command is being executed.

The right column on the debugger displays the script's variables and their current value. The variable list expands as all the script's variables are used. The variables are color-coded to show new variables in blue, changed variables in red, and variables whose value has not changed since the last step as black.

Figure 11-16 *Debug Collector Window*

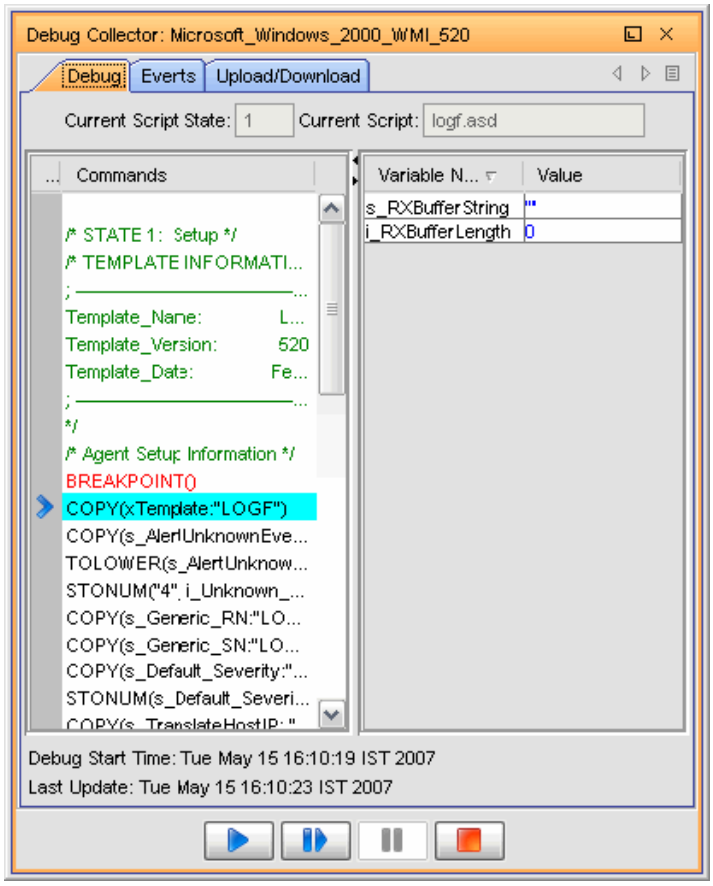

The *Events* tab displays the events generated using this Collector, and the *Upload/Download* tab allows you to upload/download another Collector script file to make modifications.

The debugger has the following four controls:

Table 11-6 *Debugger Icons*

Icon	Action	Description
	Run	Runs the script until the next breakpoint is encountered.
	Step Into	Proceeds one instruction at a time.
	Pause	Pauses the running script.
	Stop	Stops the script.

The *Command* list and the *Variable* list are not displayed in the debugger when the script is running. To see the *Command* list and the *Variable* list, the debugger must be Stepping, Paused, or Stopped.

You can view events as well as upload and download the Collector's script from the *Events* tab and the *Upload/Download* tab.

Multiple Sentinel Control Center users might connect to the same debugging session. For this reason, a Collector remains in Debug mode until one of the users specifically clicks the debugger's *Stop* button.

To debug a Collector:

1. In the main ESM display, locate the Collector that you want to debug.
2. Right-click the Collector and select *Debug*.
3. In the Debug Collector window, select a variable from the list of variables in the right pane, then click the *Run Debug* button.
4. After debugging all the variables, close the Debug window.
5. Start the Collector to generate the events.

11.5.3 Debugging JavaScript Collectors

The debugger for JavaScript Collectors can be used to debug any JavaScript Collector.

- "Accessing the Debugger" on page 223
- "Hot Keys" on page 224
- "Debugging a Collector" on page 224

Accessing the Debugger

The JavaScript debugger is launched the same way the debugger for proprietary Collectors is launched.

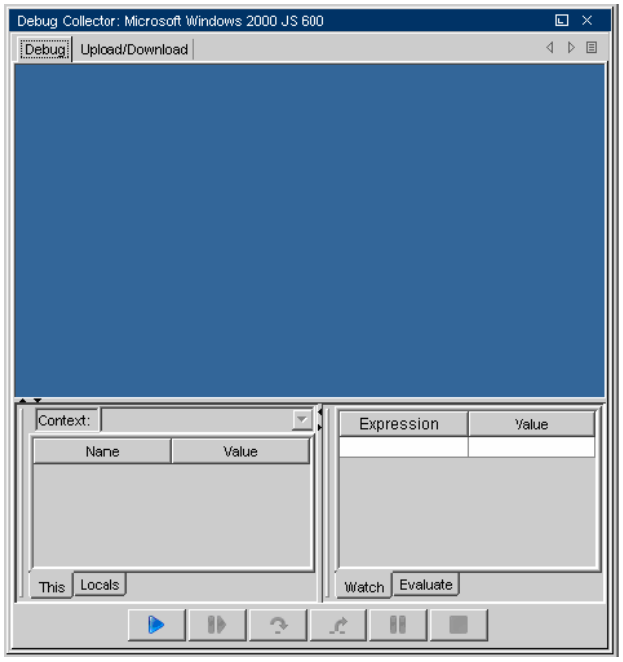

- **Debug:** Launches the JavaScript file in this window.

- **Upload/Download:** Upload/Download a JavaScript file here. You can download an existing JavaScript file, edit it, and upload it again into the system to continue debugging.
- **Context:** Displays the variable that the debugger is pointing to and its value.
- **Expression:** Displays the values of a selected parameter.

You can use the following when debugging a Collector.

Icon	Action	Description
	Run	Starts debugging.
	Pause	Pauses debugging.
	Step Into	Steps to the next line in the script.
	Step Over	Steps over a function.
	Step Out	Steps out of a function.
	Stop	Stops debugging.

Hot Keys

When the source code window is on focus in the debugger, you can use the following hot keys:

- Use Ctrl+F to find a string in the source code.
- Use Ctrl+G to go to a line number.
- Use Ctrl+M to find the parenthesis or brace that matches the highlighted one.

You can also open a script file, set a break point, step through the script code, and watch variable and method values at each step.

You can debug Collectors in Standalone or Connected modes.

Debugging a Collector

1. Log into Sentinel Control Center. On the menu bar, click *Event Source Management* > *Live View*.
2. Right-click the Collector and stop the Collector if it is running.
3. Right-click the Collector and select *Debug*.

   The Debug Mode Selection window displays.

You can choose to debug in Standalone or Live mode.

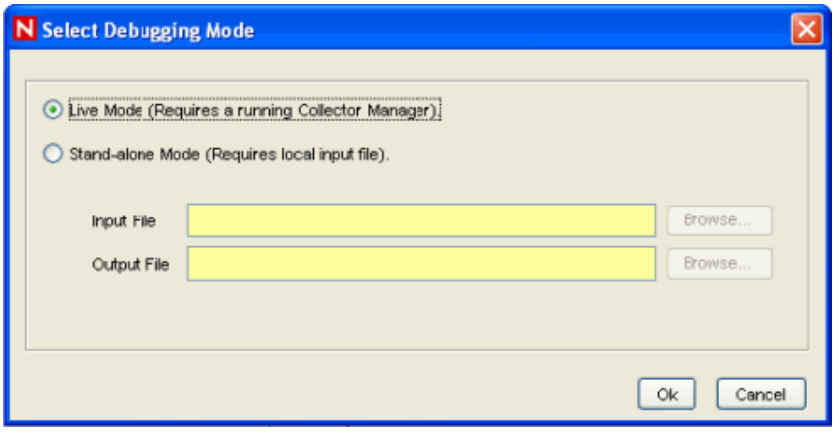

- "Standalone Mode" on page 225
- "Live Mode" on page 226

Standalone Mode

Standalone debug mode allows you to debug a Collector even if the associated Collector Manager is not running.

For standalone mode, input to the script comes from an input file rather than a live event source. Specify the path to a raw data file that will be used as input. For Collectors that use a DB Connector, the input file should be a text file with log data in NVP format. For Collectors that use a File Connector, the input should be a text file with log data in CSV format.

For standalone mode, output from the script is to an output file rather than live events. You must specify the path to the output file that the script uses for output. If you specify an output file that does not exist, the system creates the file for you

1. Select Standalone mode, then stop the Collector.

   The Collector Manager does not need to be up and running, and the events do not display in the Active Views.

2. Specify the path for the input and output files.

   If you specify an output file that does not exist, the system creates the file for you.

3. Click *OK*. The Debug Collector window displays.

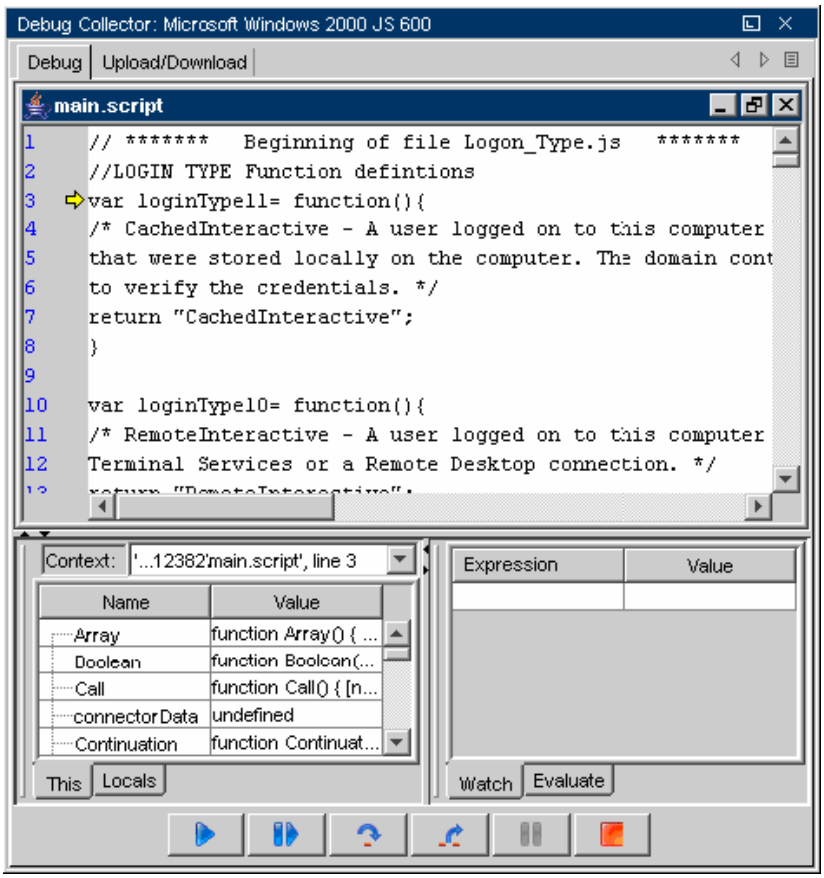

4 In the Debug Collector window, click *Run* .

In the Source text area, the source code of the Collector appears and stops at the first line of the text script.

5 Click the bar on the left and toggle a breakpoint in the script code, then click to go to the next breakpoint.

Click *Pause* to pause debugging whenever required.

6 After debugging is complete, click *Stop* to stop debugging.

7 Click the *Upload/Download* tab in the debugger window.

8 Click *Download* and specify a location to download the script file.

9 Open the file with any JavaScript editor or a text editor.

Make your edits in the code and save the file, then click *Upload*.

Debug the uploaded script to have a Collector script ready to use.

Live Mode

- Live debug mode requires that the Collector Manager associated with the Collector is running.
- In Live debug mode, Input to the script comes from actual event sources connected to the Collector. To get data from a specific event source, you must right-click and start the desired event source via the ESM display. Starting/stopping event sources can be done any time during the debug session.

NOTE: If no event source is started during the debug session, then no data is available in the buffer for the Collector and you see the Collector script's readData method blocking.

- In Live debug mode, Output from the script is via live Sentinel events. The events can be viewed on the Active Views displays.

NOTE: When in Live debug mode, the script engine is executed on the local box rather than the actual box that the associated Collector Manager is running on. The Connectors/event sources still run on the same box as the Collector Manager. When running debug mode, data is automatically routed from the event sources to the script engine running in debug mode on the local box.

11.5.4 Using the Raw Data Tap to Generate a Flat File

When debugging, it might occasionally be helpful to view Connector output data. In addition to viewing raw data from the Connector by using the *Raw Data Tap* right-click option for nodes in the Sentinel Control Center, Sentinel also includes an option to save the raw data from a Connector to a file for further analysis.

To save raw data from a deployed Connector to a file:

1. Right-click the Connector node and select *Edit*. The Edit Connector dialog box displays.

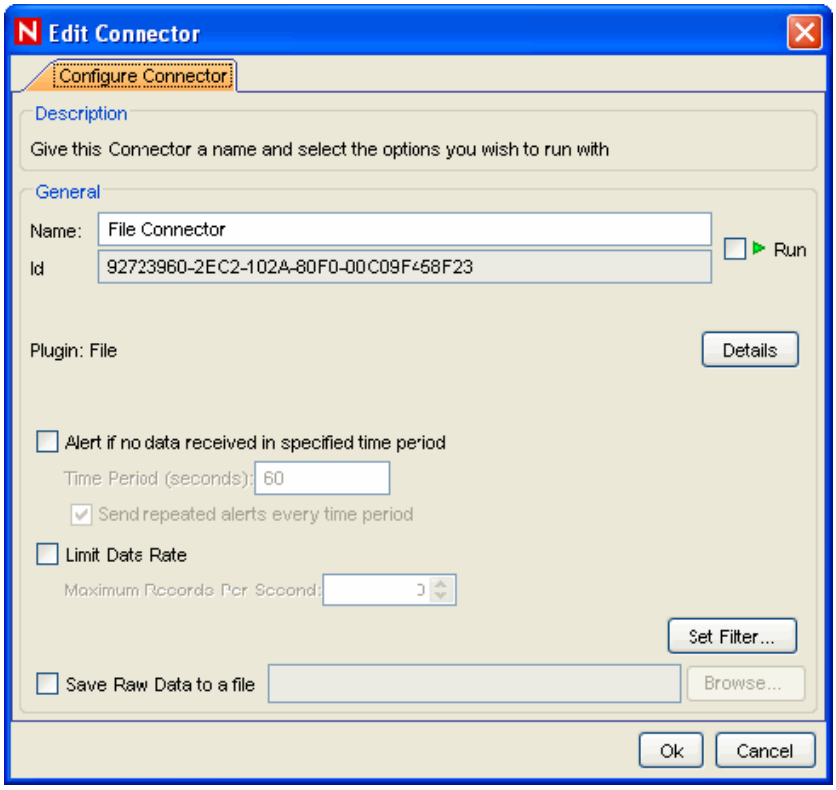

2. Select *Save Raw Data to a file*.
3. Specify a path on the Collector Manager machine where the raw data is saved.
4. Click *OK*.

IMPORTANT: The account running the Sentinel service on the Collector Manager machine must have permissions to write to the file location.

11.6 Exporting a Configuration

You can export the configuration of ESM objects along with their Collector scripts and the Connector plug-ins.

NOTE: You can export any object in the ESM panel. Depending on the object selected, all its children and parents are displayed in the Select Data window of the Export Configuration Wizard.

To export your configurations:

1. In the menu bar, click *File* > *Export Configuration* or right-click an object in the ESM panel and select *Export Configuration*.

   The Export Configuration window displays.

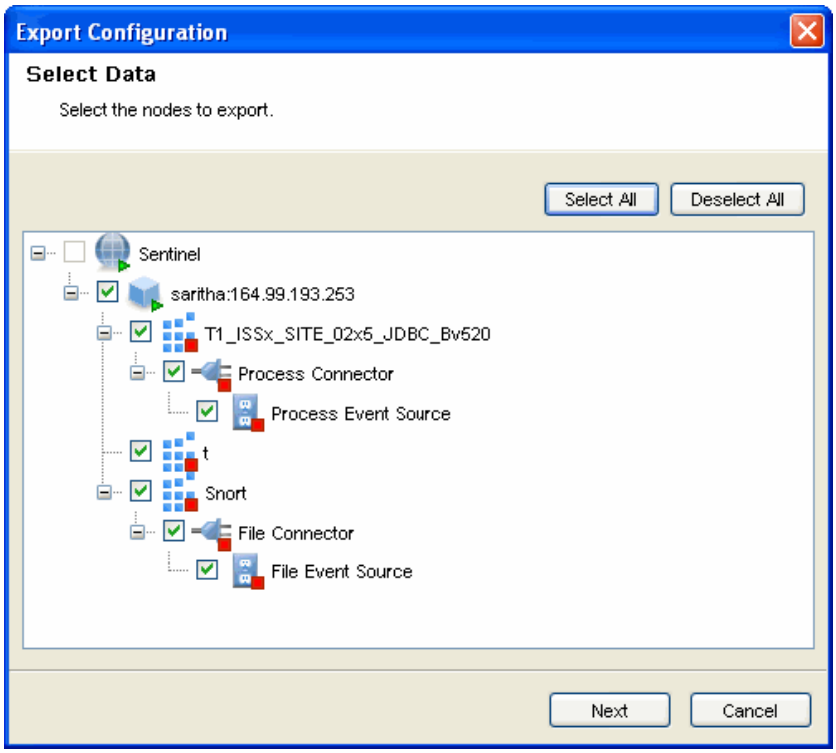

2. Select the data to export and click *Next*. The Select Collector Scripts window displays.

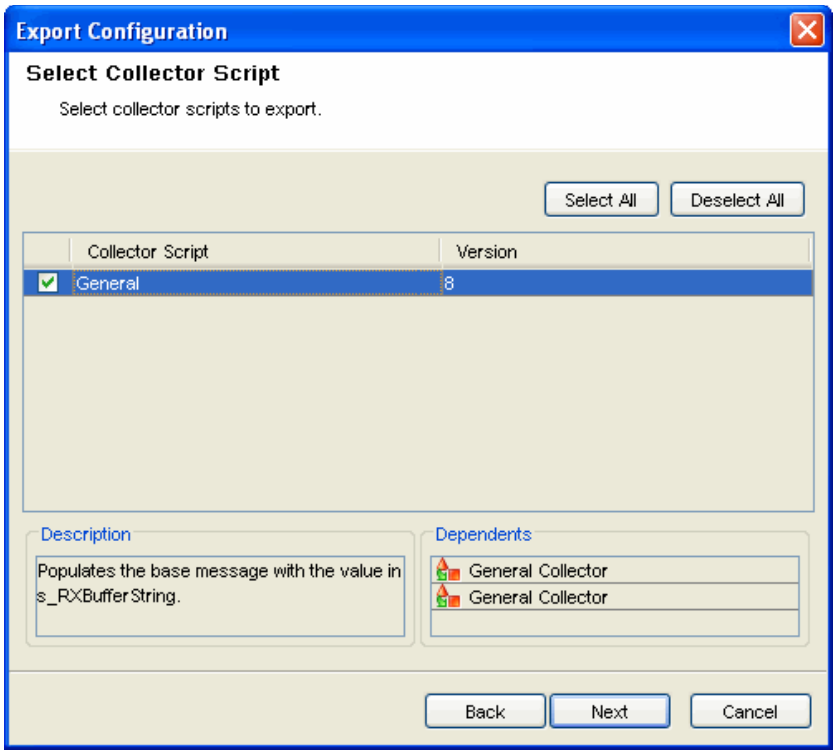

3 Select the Collector scripts from the list to export, then click *Next*. You can select or deselect all.

The Select Connectors Plugin window displays.

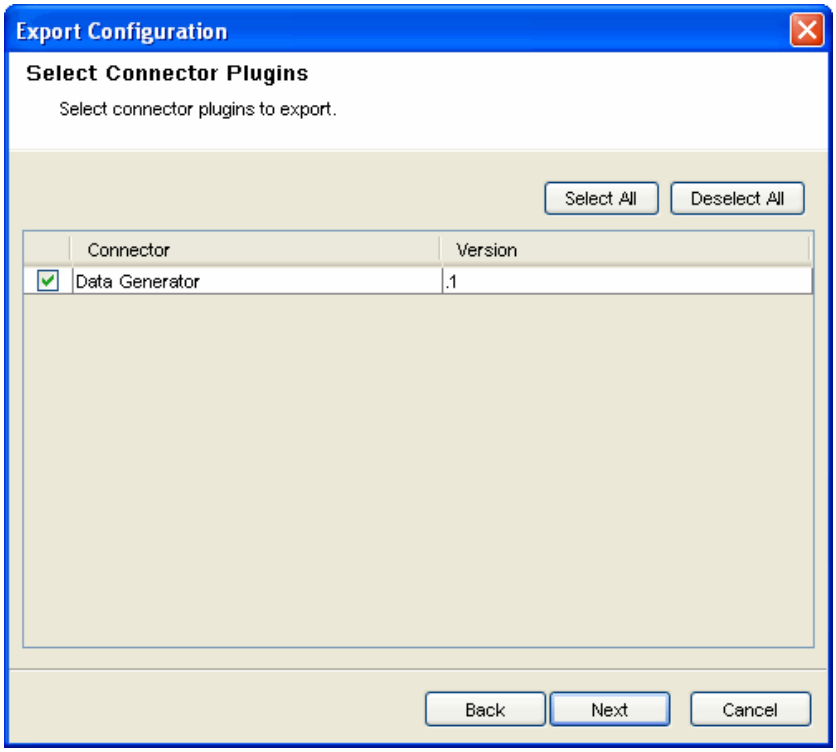

4 Select the Connector plug-ins from the list to export, then click *Next.* You can select or deselect all.

The Specify Export File window displays.

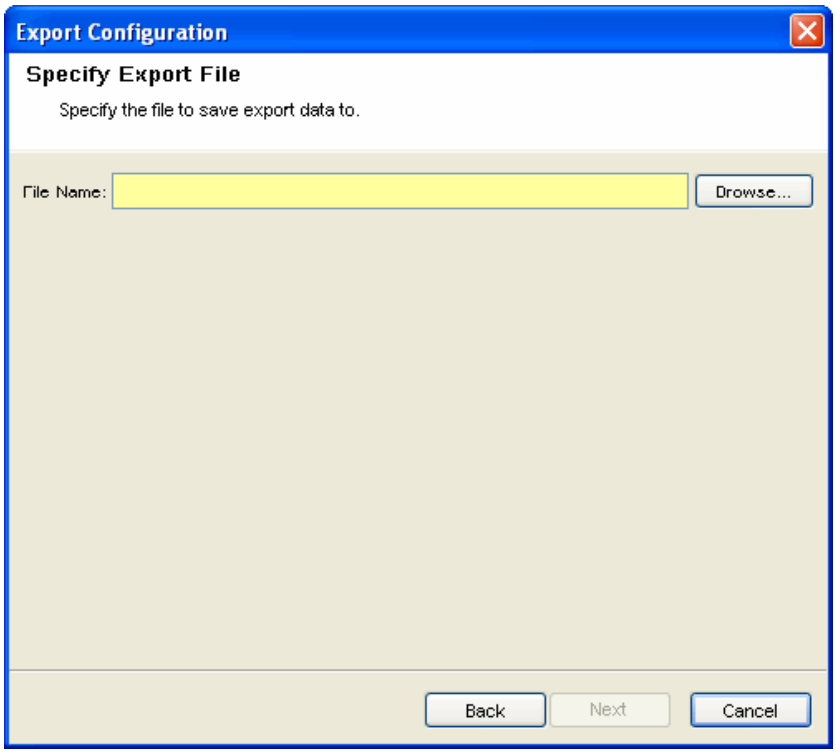

If you want to view the description and dependents of a particular plug-in in the above window, select that plug-in from the table.

5 Specify a location to save the configuration and click `Next`.

You can save the configurations only to a ZIP file.

A Summary page with the details of the configurations and plug-ins selected to export displays.

6 Click *Finish* to export. The file is exported in ZIP format.

11.7 Importing a Configuration

Importing a configuration helps you to import the configuration of ESM objects exported to a ZIP file along with the plug-ins.

- Section 11.7.1, “Enabling or Disabling the Import Configuration,” on page 231
- Section 11.7.2, “Resetting the Layout,” on page 233
- Section 11.7.3, “Undoing the Layout,” on page 233
- Section 11.7.4, “Redo Layout,” on page 233

11.7.1 Enabling or Disabling the Import Configuration

The *Import Configuration* option is enabled under the following circumstances:

- In *Live View* when you select the Collector manager, Collector, or Connector
- In the *Scratchpad* when you select any node other then the event source

Import Configuration in *Live View* and the *Scratchpad* is disabled if you do the following:

- Select Sentinel or event source nodes (only in Live View)
- Do not select any node in Live View
- Select an event source node in a child view of the graphical view
- Select multiple nodes

To import your configurations:

1. Click *File* on the menu bar and select *Import Configuration*. You can also click the *Import Configuration* button on the toolbar. The Import Configuration window displays.

   You can also import a configuration by right-clicking the object in the ESM panel. Depending on the object you have selected in the ESM panel, the node along with its child nodes is displayed in the Select Data window of the Import Configuration Wizard.

2. Browse and select the configurations file and click *Next*. The Select Data window displays.

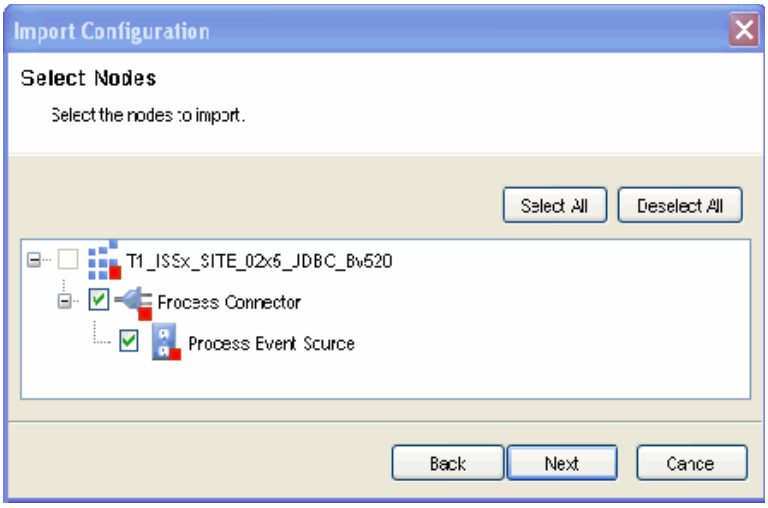

   Configurations must be saved to a ZIP file to import.

3. Select the data to import and click *Next*. The Select Collector Script window displays.

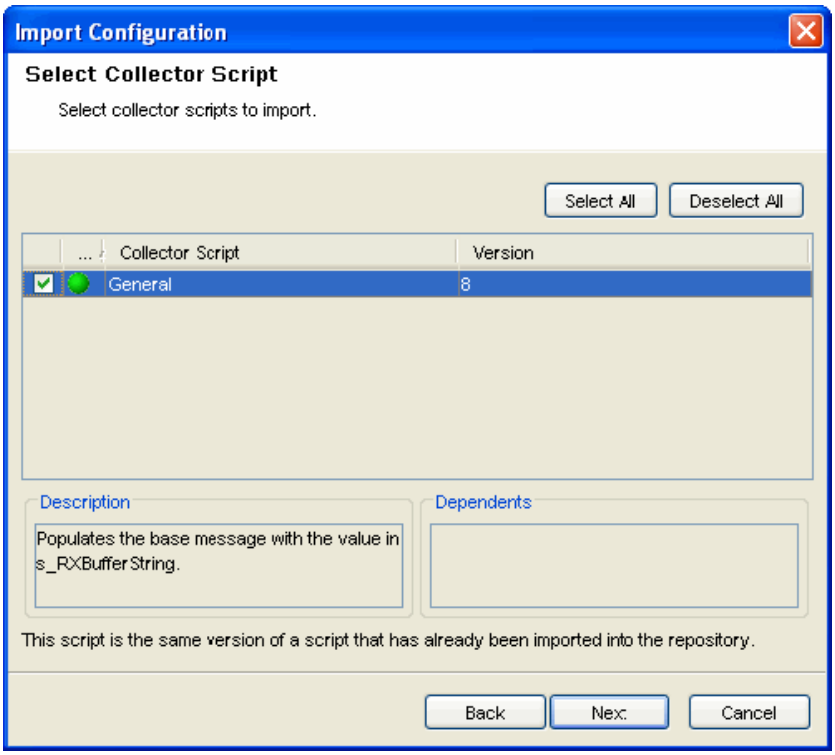

4 Select the Collector script from the list to import.

A color indicator is displayed in the Select Collector Script and Select Connector Plugins window to indicate whether the plug-in is already present in the repository or not. If the plug-in is not present in the repository, the color is displayed as red and if the same version of plug-in exists, the color is green or orange.

5 Click *Next*. The Select Connector Plugins window displays.

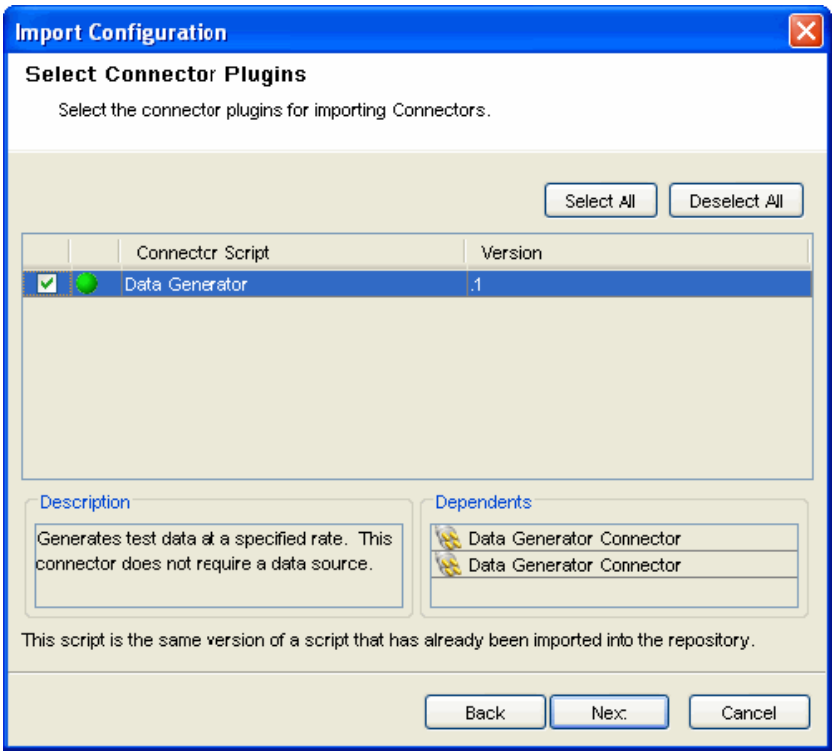

6 Select the Connector plug-ins from the list to import.

NOTE: To view the description and dependents of a particular plug-in in the above window, select that plug-in from the table. If there are any Collectors or Connectors in the ESM panel that are affected on importing the plug-in, the Affected Collectors or Affected Connectors window is displayed.

---

7 Click *Next*. A Summary page with the details of the configurations and plug-ins selected to import displays.

8 Click *Finish*.

11.7.2 Resetting the Layout

To reset to default settings:

1 Click *View* on the menu bar and select *Reset Layout*. Alternatively, click the *Reset* button on the toolbar.

11.7.3 Undoing the Layout

1 Click View on the menu bar and select *Undo Layout*. Alternatively, click the *Undo Layout* button on the toolbar.

11.7.4 Redo Layout

1 Click *View* on the menu bar and select Redo Layout. Alternatively, click the *Redo Layout* button on the toolbar.

11.8 Event Source Management Scratchpad

Scratchpad is the Design Mode of the Health Monitor. Through Scratchpad, you can design and configure various items:

- Collector Managers
- Collectors
- Event Sources
- Connectors
- Event Source Servers

You can right-click the Sentinel icon and add the components. For more information, see Section 11.4.2, “Adding Components to the Event Source Hierarchy,” on page 204.

NOTE: You cannot view the status of any object in the design mode because they are not connected to an instance of a real Collector Manager.

Administration

12

You use the *Admin* tab to configure filters and reports. You use the *User Manager* option in the *Admin* tab to create users and you can assign rights to the users.

- Section 12.1, “Understanding the Admin Tab,” on page 235
- Section 12.2, “Introduction to the User Interface,” on page 236
- Section 12.3, “Servers View,” on page 237
- Section 12.4, “Filters,” on page 239
- Section 12.5, “Configure Menu Options,” on page 248
- Section 12.6, “DAS Statistics,” on page 254
- Section 12.7, “Mapping,” on page 255
- Section 12.8, “Event Configuration,” on page 265
- Section 12.9, “Report Data Configuration,” on page 270
- Section 12.10, “User Configurations,” on page 275

12.1 Understanding the Admin Tab

The *Admin* tab allows you to access the following:

- Servers View: View the health of server components.
- Filters: Create and edit filters.
- DAS Statistics: View health statistics for DAS components.
- Color Filter Configuration: Format events based on filter criteria.
- Mapping: Configure the mapping service.
- Event Configuration: Rename event fields and configure fields to be populated by the mapping service.
- Report Data Configuration: Enable or disable the aggregation service.
- User Configurations: Create users and roles and manage active user sessions.

NOTE: You need to have appropriate permissions to access this tab. Only an Administrator has the control to enable or disable access to the features of *Admin* tab for a user.

Figure 12-1 *Sentinel Control Center*

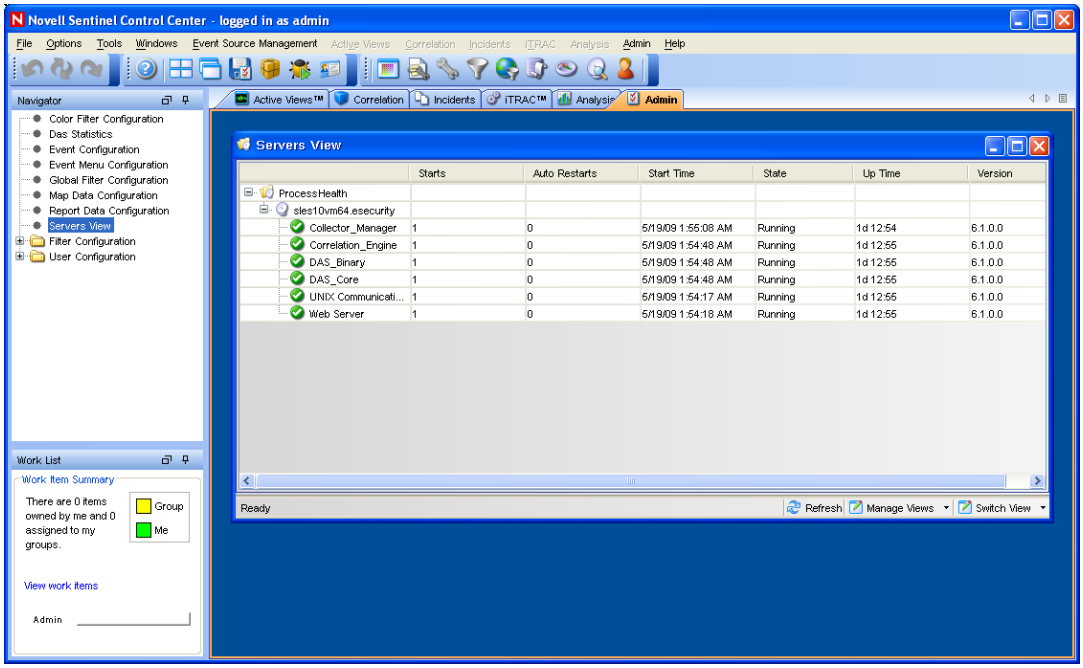

12.2 Introduction to the User Interface

In the *Admin* tab, you can see server views, filter configuration, and user configuration in the Admin Navigator.

You can navigate to these functions from:

Table 12-1 *Admin Tab User Interface*

User Interface	Description
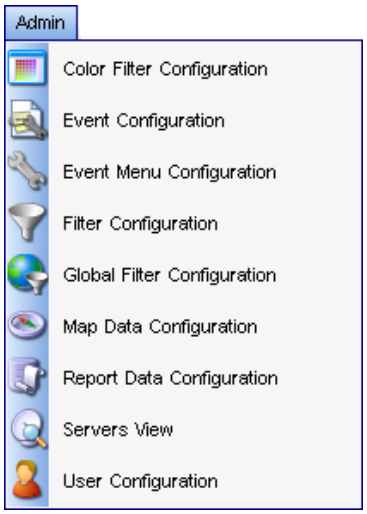	The Admin menu in the menu bar

User Interface	Description
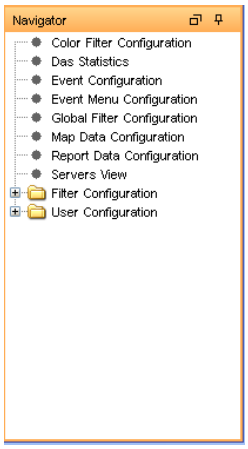	The Navigation Tree in the Navigation pane
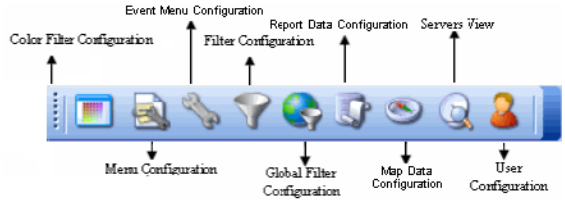	The tool bar buttons

12.3 Servers View

Through the Servers view you can start, stop, or restart processes that are installed on the product installation. Servers view also allows you to monitor the status of all Sentinel server processes across the system. The following are the Sentinel server processes:

- Collector_Manager
- Correlation_Engine
- DAS_Binary
- DAS_Core
- Web Server
- Unix Communication Server

Figure 12-2 *Servers View Window*

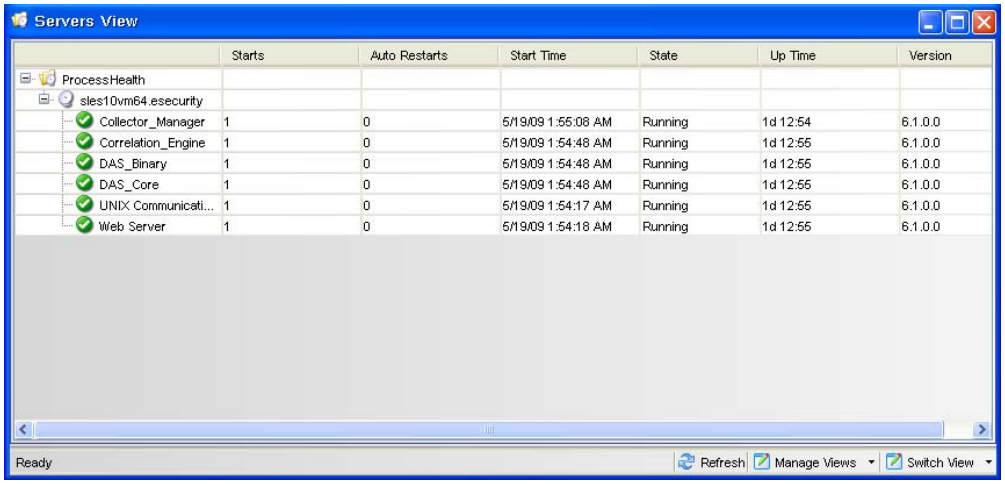

- **Start, Stop, or Restart processes:** Take these actions on a process by right-clicking the process entry.

  You cannot either stop or restart the following processes by using the right-click options *Action > Stop/Restart* in the Servers view.

  - DAS_Core
  - Web Server
  - Unix Communication Server
- **Starts:** The number of times the process was started, for whatever reason. This includes starts initiated by the user through the GUI or done automatically.
- **AutoRestarts:** The number of times the process was automatically restarted. Because this only applies to automatic restart scenarios, it does not apply to restarts initiated by a user. This field is helpful for determining if the process exited (For example, because of an error) and was automatically restarted by the Sentinel Watchdog.

12.3.1 Monitoring a Process

1. Click the Admin tab.

   Click Servers View. Alternatively, click *Servers View > Servers View* in the Navigator, or click the Servers View icon.
2. Expand the server view. All the processes are listed.

12.3.2 Creating a Servers View

1. Click the Admin tab.

   Click Servers View. Alternatively, click *Servers View > Servers View* in the Navigator, or click the *Servers View* icon.
2. To create a new view, click the Manage View drop-down arrow on the bottom right corner, then click Add View.
   - Specify your option name.

- To arrange which fields you want to be shown, click *Fields*.
- To group different attributes, click *GroupBy*.
- To sort by different attributes, click *Sort*.
- To filter, click *Filter*.
- To change the display values of the processes shown in the servers view, click *Leaf Attribute*.

3 Click *Save*.

12.3.3 Starting, Stopping, and Restarting Processes

1 Click the Admin tab.

2 Click Servers View. Alternatively, click *Servers View* > *Servers View* in the Navigator, or click the Servers View icon.

3 Expand the servers view. All the processes are listed.

4 Right-click a process, then click *Actions* and select *Start*, *Restart*, or *Stop*.

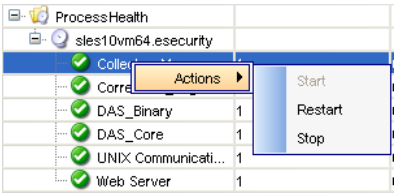

12.4 Filters

Filters allow you to process data based on specific criteria for events in real time and for users of the system. Filters enable you to manage data seen in the Sentinel Control Center. The Filter engine drives the Real Time Event windows by maintaining the data structure for each security filter. Filters prevent users from viewing unauthorized events and they drop events that users don't want to see. Filters are created in the *Admin* tab of the Sentinel Control Center.

NOTE: The following are invalid filter name characters: $ # . * & : < >.

There are three types of filters:

- Section 12.4.1, "Public Filters," on page 239
- Section 12.4.2, "Private Filters," on page 240
- Section 12.4.3, "Global Filters," on page 240
- Section 12.4.4, "Configuring Public and Private Filters," on page 243
- Section 12.4.5, "Color Filter Configuration," on page 245

12.4.1 Public Filters

Public filters are system-owned. Public filters can be used as security filters or display filters. Security filters are based on user permissions. Display filters determine which events are depicted in the real time event tables, charts, and graphs.

Figure 12-3 *Filter Manager Window*

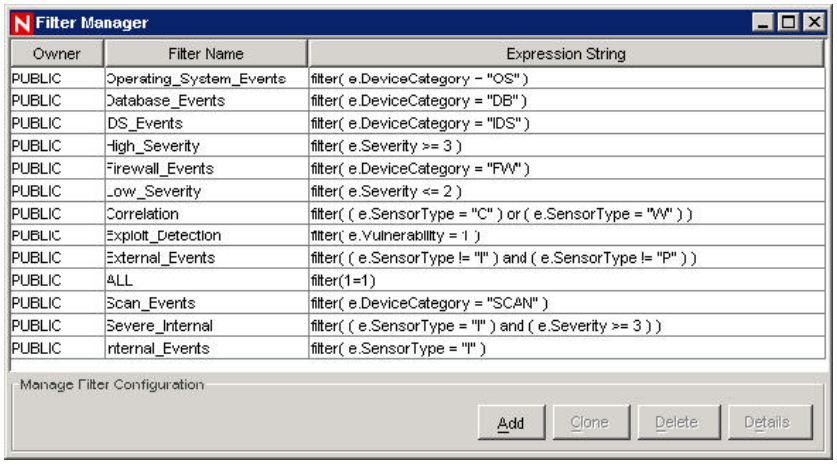

Owner	Filter Name	Expression String
PUBLIC	Operating_System_Events	filter(e.DeviceCategory = "OS")
PUBLIC	Database_Events	filter(e.DeviceCategory = "DB")
PUBLIC	IDS_Events	filter(e.DeviceCategory = "IDS")
PUBLIC	High_Severity	filter(e.Severity >= 3)
PUBLIC	Firewall_Events	filter(e.DeviceCategory = "FW")
PUBLIC	Low_Severity	filter(e.Severity <= 2)
PUBLIC	Correlation	filter((e.SensorType = "C") or (e.SensorType = "W"))
PUBLIC	Exploit_Detection	filter(e.Vulnerability = 1)
PUBLIC	External_Events	filter((e.SensorType != "I") and (e.SensorType != "P"))
PUBLIC	ALL	filter(1=1)
PUBLIC	Scan_Events	filter(e.DeviceCategory = "SCAN")
PUBLIC	Severe_Internal	filter((e.SensorType = "I") and (e.Severity >= 3))
PUBLIC	Internal_Events	filter(e.SensorType = "I")

12.4.2 Private Filters

Private filters are user-owned. Private filters are display filters and are shareable if you have the View Private Filters permission.

12.4.3 Global Filters

Global filters are classified as Public filters. Global filters are sequentially processed at the Collector Manager for each event. Once the global filter criteria are met, the evaluation stops for that event and the associated global filter action is taken for the event.

The order of evaluation of global filters is top to bottom, as shown in the console. They can be enabled or disabled as required. Global Filters enable routing actions and JavaScript actions on events. Routing actions include dropping events or routing events to database, database and GUI (SCC), or only to GUI (SCC).

Through the Global Configuration window, you can:

- Create a Global Filter
- Rearrange a Global Filter
- Delete a Global Filter

Figure 12-4 *Global Filter Configuration*

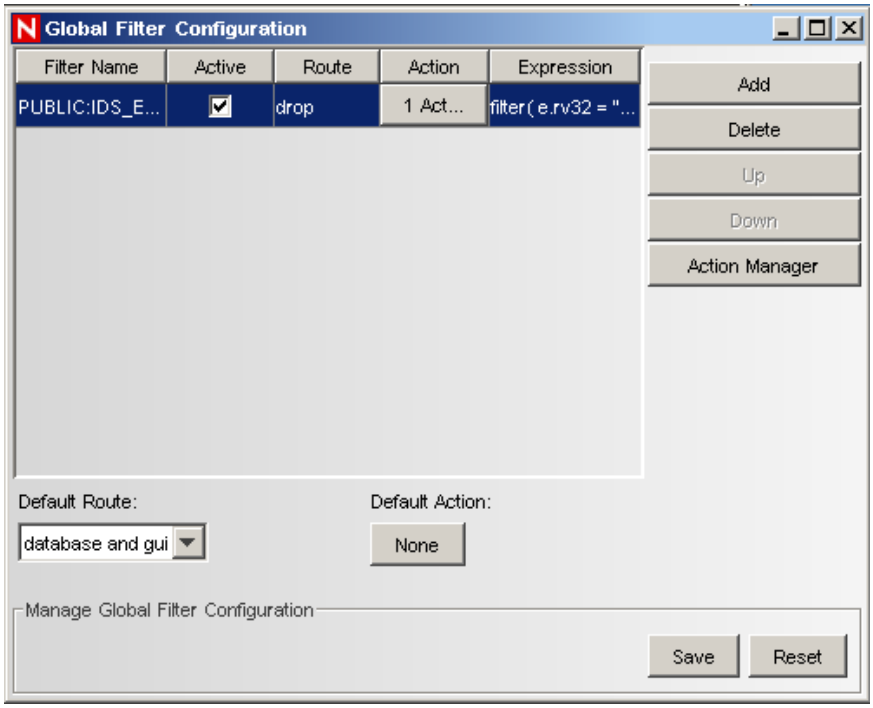

NOTE: The *Action* column and the *Action Manager* button are available only on systems that have Sentinel Rapid Deployment Hotfix 2 or later installed.

Creating a Global Filter

1. Click the *Admin* tab.
2. Click *Admin* > *Global Filter Configuration* or select *Global Filter Configuration* in the navigation tree.
3. In the *Global Configuration* window, click *Add*.
4. In the new blank row, click the *Filter Name* column.
5. In the *Filter Selection* window, highlight a relevant filter and click *Select*, or click *Add* if you need to create a filter.
6. In the *Active* column, select the checkbox to associate the filter with the options specified in the *Route* and *Action* columns.

   NOTE: If the *Active* checkbox is not selected, the options sepecified in the *Default Route* and *Default Action* will be associated to the filter. If the *Default Action* is set to None, then no action will be associated to the filter.

7. In the *Route* column, select the routing action that the global filter will have on events that pass this global filter.

   If an event does not meet any of the active global filters, the default action determines how the event is handled.

The following are the options available in the *Route* drop-down list:

- **drop:** Events are dropped and are not sent to Sentinel Control Center or the Sentinel Server database.
- **database:** Events are sent directly to the Sentinel Server database and not sent to the Sentinel Control Center.
- **database and gui:** Events are sent to the Sentinel Control Center and the Sentinel Server database.
- **gui only:** Events are sent to the Sentinel Control Center.

8 Continue adding filters until you have completed adding all the required filters.

9 In the *Action* column, select the action that needs to be performed once the filter criteria are met.

NOTE: To create new actions for the filter, click *Action Manager* or from the menu bar click *Tools* > *Action Manager*. For more information on creating actions, see Section 17.3, "Actions," on page 376.

You can associate single or multiple actions to a filter. By default, the *Action* and *Default Action* are set to None. Global Filters execute only JavaScript actions. Actions that are associated with global filters cannot be deleted from the Action Manager.

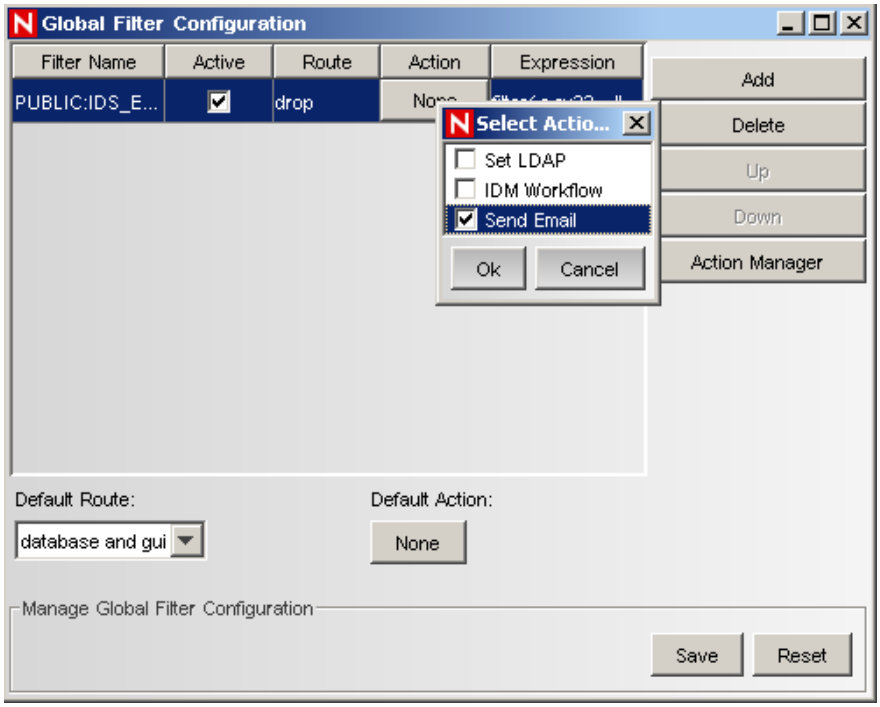

10 Continue adding filters until you have completed adding all the required filters.

11 Click *Save*.

Rearranging Global Filters

1. In the *Global Configuration* window, select a filter and click *Up* or *Down* to move it to a different location on the list.
2. Click *Save*.

Deleting a Global Filter

NOTE: When you delete a global filter, the confirmation message is not displayed.

---

1. In the *Global Configuration* window, select a filter from the list and click *Delete*.
2. Click *Save*.

12.4.4 Configuring Public and Private Filters

Configuring public and private filters allow you to:

- "Adding a Filter" on page 243
- "Cloning a Public or Private Filter" on page 245
- "Modifying a Public or Private Filter" on page 245
- "Viewing the Details of a Public or Private Filter" on page 245
- "Deleting a Public or Private Filter" on page 245

Figure 12-5 *Filter Manager Window*

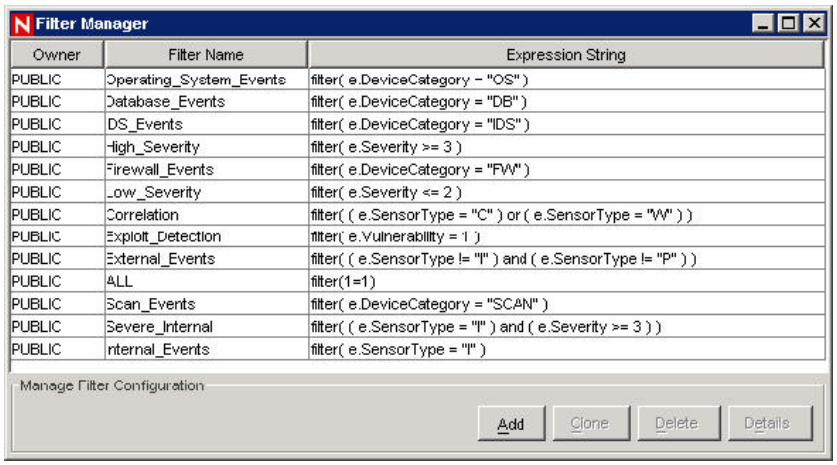

Adding a Filter

To add a public or private filter:

1. Click *Admin* > *Filter Manager* or select *File Manager* under the *Filter Configuration* folder in the Navigator; then click *Add*.
2. Select an Owner ID. You can select PUBLIC or PRIVATE (user owned).

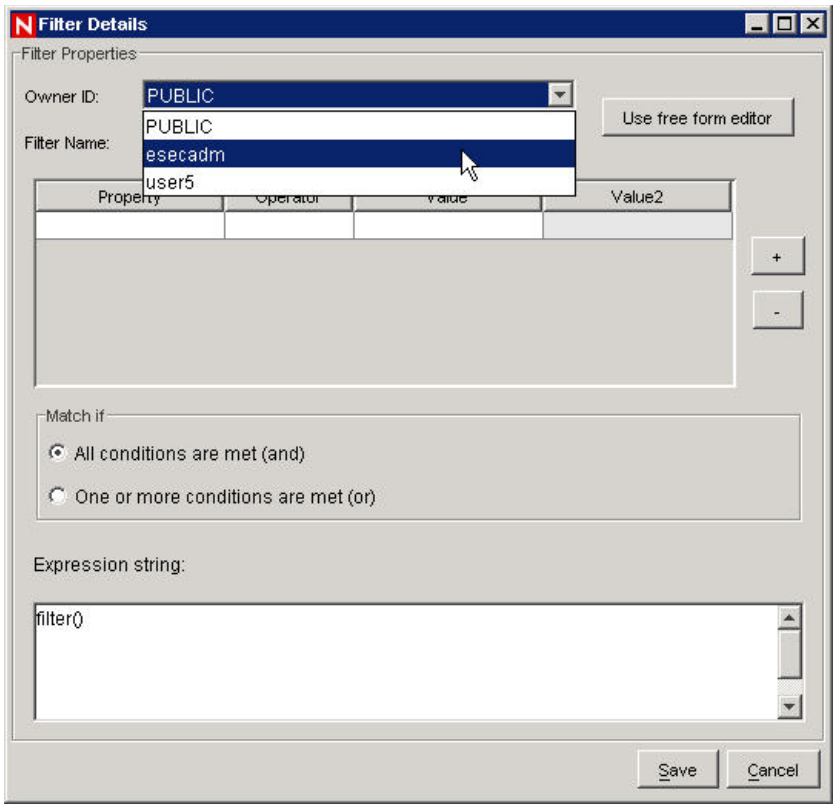

3 Specify a filter name.

The table editor is the default selection for editing the contents.

Optionally, you can click Use free form editor to display a free form editor. The free form editor allows you to create complex expressions not possible with the table editor. However, after the expression is modified with the free form editor, the table editor cannot be used with the expression.

4 Select the criteria for the following columns:

- Property
- Operator
- Value columns.

NOTE: To include special characters in the *Value* column, you should provide the hexadecimal value (character code) of the special character. For example, if the value is “10.1.1.1”, you should enter \x2210.1.1.1\x22 to embed the double quote in a string value.

The *Expression string* box displays the filters that you created in RuleLg language.

5 In the Match if box, click one of the following:

- All conditions are met (and)
- One or more conditions are met (or)

6 To create another filter expression, click Create a New Filter Expression (+) to add another row to the filter expression table.

7. To remove a filter expression, select a filter expression from the table and click Remove the Selected Expression (-).
8. Click Save.

Cloning a Public or Private Filter

Cloning is a convenient way to duplicate a filter to assure consistency of criteria among a group of filters or users.

1. Open the Filter Manager window.
2. Click *Clone*.
3. Provide a new filter name.
4. Change any original filter's criteria.
5. Click *Save*.

Modifying a Public or Private Filter

1. Open the Filter Manager window.
2. Select a filter and click *Details*.
3. Change any of the criteria as desired. You cannot change the Owner ID and the Filter Name.
4. Click Save.

Viewing the Details of a Public or Private Filter

1. Open the Filter Manager window.
2. Select a filter and click *Details*.

Deleting a Public or Private Filter

1. Open the Filter Manager window.
2. Select a filter and click *Delete*.

   A confirmation window displays.
3. Click *Yes* in delete confirmation dialog box.

12.4.5 Color Filter Configuration

The Color Filter Configuration feature allows you to assign background and text colors to events in the Sentinel Control Center based on filter criteria. The background and text colors assigned to a filter apply to all Sentinel tables, including Active Views, event tables associated with incidents, offline queries and historical event queries.

On applying a color filter, all the event tables are updated.

Figure 12-6 *Color Filter Configuration*

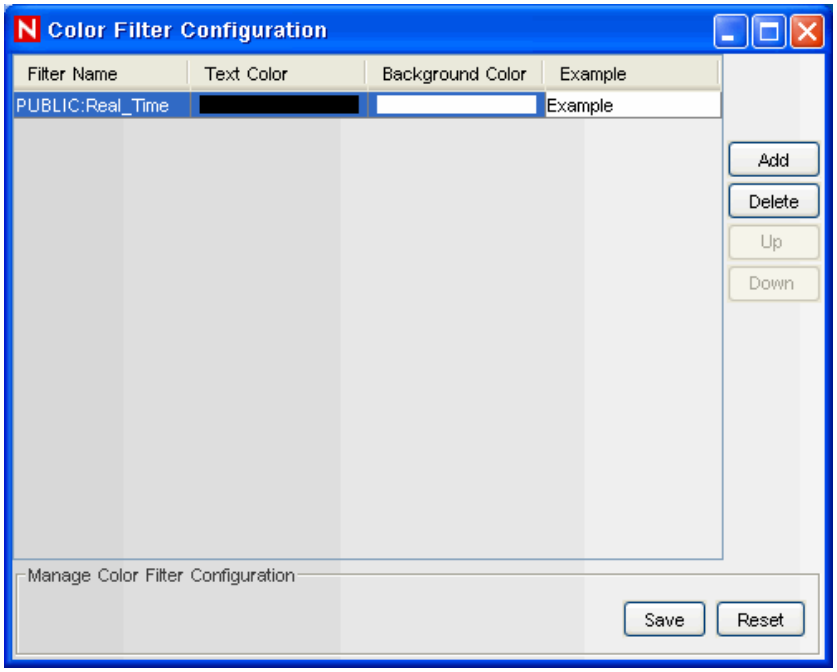

The Color Filter Configuration GUI displays a list of all the color filters that are defined in the order in which they should be applied. If an event meets the criteria for more than one of the color filters, the first color filter configuration is applied. For example, the following filter configurations are created and attached to color filter configuration:

- Color filter configuration 1: sev=2 (with background color red and text color yellow)
- Color filter configuration 2: sev>1 (with background color white and text color black)

Any event with severity=2 will meet the criteria for both color filters, but because the sev=2 color filter configuration is at the top, all the events with sev=2 are coded according to color filter configuration 1. All the other events with sev>1 (For example, sev=3, 4, 5 and so on) follow color filter configuration 2.

- "Adding a Color Filter" on page 246
- "Deleting a Color Filter" on page 248
- "Setting Color Filter Priorities" on page 248

Adding a Color Filter

1. Click Color Filter Configuration in the navigation pane or click the *Color Filter Configuration* button.
2. Click Add. A new Color Filter Configuration row is created as shown below.

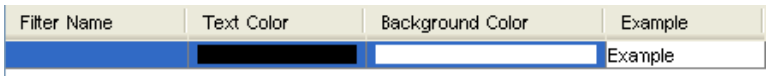

3. Click the Filter Name drop-down list. The Filter Selection window displays.

4 From the list, select a filter to which you want to apply the color filter configuration and click *Select*, or click *Add* to create a new filter.

For more information on configuring filters, see Section 12.4.4, “Configuring Public and Private Filters,” on page 243.

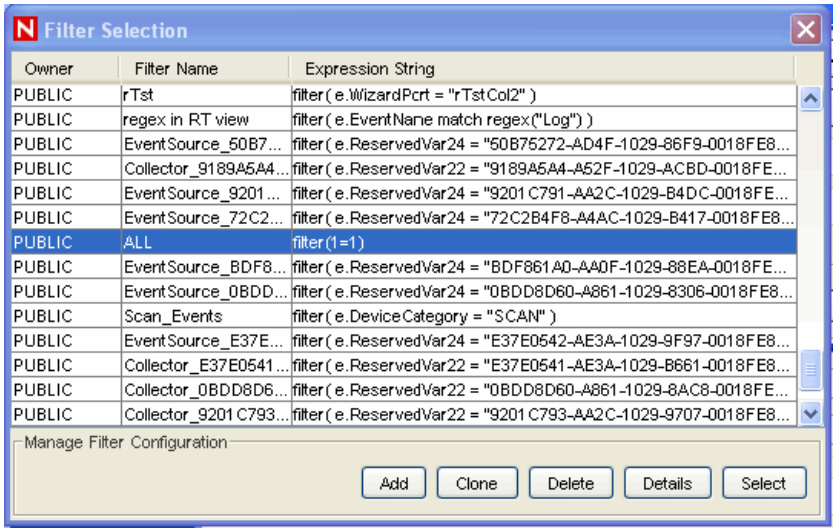

5 In the Color Filter Configuration window, click *Text Color*. The Pick a Color window displays.

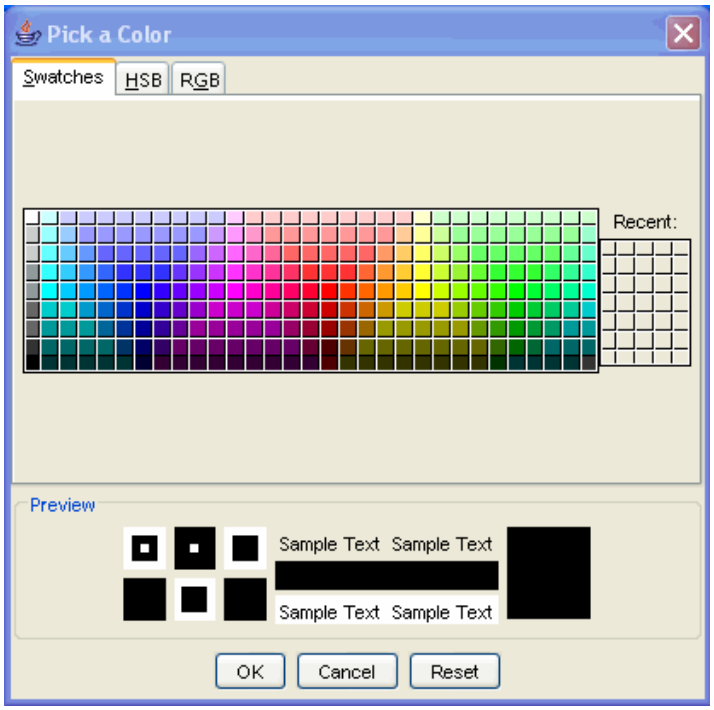

6 Select a color from the *Swatches* tab. Alternatively, click the *HSB* or *RGB* tab and specify the HSB or RGB color value in the respective tab.

7 Click *OK*.

8. In the Color Filter Configuration window, click *Background Color*. The Pick a Color window displays.
9. Select a color from the Swatches tab. Alternatively, click the *HSB* or *RGB* tab and specify the HSB or RGB color value in the respective tab.
10. Click *OK*.
11. Click *Save*.

NOTE: The order of the color filter configuration row in the Color Filter Configuration window matters. If more than one color filter definition applies to an event, the formatting for the first color filter takes precedence.

Deleting a Color Filter

1. Click *Color Filter Configuration* in the navigation pane.
2. Select a Color Filter Configuration row and click *Delete*.

Setting Color Filter Priorities

1. Click Color Filter Configuration in the navigation pane or click the Color Filter Configuration button.
2. Select a Color Filter Configuration row.
3. Click the *Up* or *Down* button to set the priority.

NOTE: The Up and Down buttons are active only when there is more than one color filter configuration row available in the Color Filter Configuration window.

12.5 Configure Menu Options

NOTE: To use this feature, you must have the Event Menu Configuration user permission.

Use the Event Menu Configuration window to create the menu items that appear on the *Event* menu, which is available by right-click when an event (or set of events, if the action is written in JavaScript) is selected in any event table (for example, an Active View window, Snapshot window, Incidents Events window, or Offline Query window). Sentinel has the following default Event Menu Configuration items that you can clone, activate, or deactivate:

Figure 12-7 Event Menu Configuration

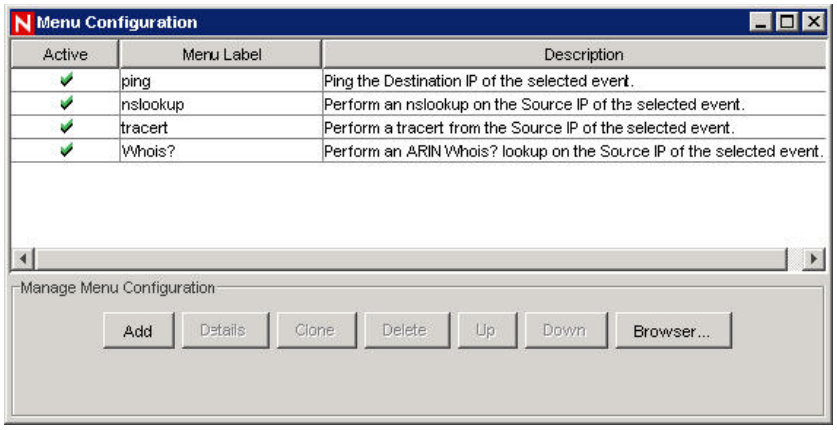

- **Ping:** Ping the destination (or target) IP of the selected event
- **nslookup:** Perform an nslookup on the Source (or initiator) IP of the selected event
- **tracert:** Perform a traceret from the Source (or initiator) IP of the selected event to the Sentinel Server
- **Whois?:** Perform an ARIN Whois? lookup on the Source (or initiator) IP of the selected event

To view the configuration details for any of these options, select the item and click *Details*.The following is the nslookup configuration.

Figure 12-8 Menu Item

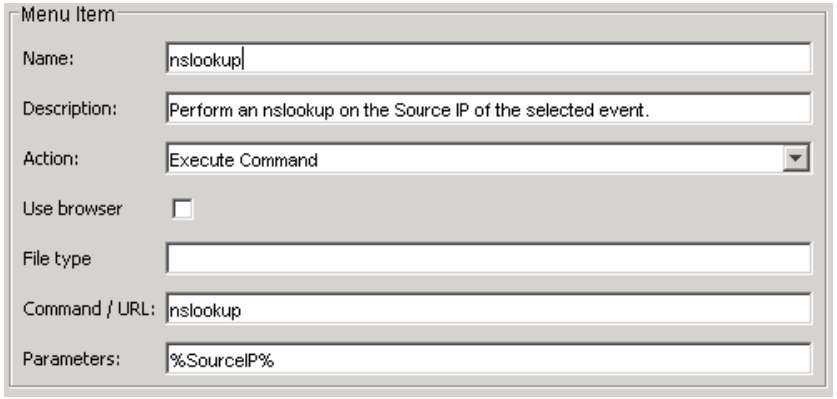

In addition, new options can be customized to execute a command, open a Web browser, or execute a JavaScript Action configured through the Action Manager.

NOTE: The Execute Command scripts, commands, or applications must be available in `<install_directory>/config/exec`.

Event Menu Configuration allows you to perform the following activities:

- Section 12.5.1, “Adding an Option to the Event Menu,” on page 250
- Section 12.5.2, “Cloning an Event Menu Option,” on page 251
- Section 12.5.3, “Modifying an Event Menu Option,” on page 252

- Section 12.5.4, “Viewing Event Menu Option Parameters,” on page 252
- Section 12.5.5, “Activating or Deactivating an Event Menu Option,” on page 252
- Section 12.5.6, “Rearranging Event Menu Options,” on page 252
- Section 12.5.7, “Deleting an Event Menu Option,” on page 252
- Section 12.5.8, “Editing Your Event Menu Browser Settings,” on page 253

12.5.1 Adding an Option to the Event Menu

Users with the appropriate permissions can add new actions to the event menu that appears when users right-click an event or events in any event table. There are three types of actions that can be configured for the event menu:

- **Execute Command:** Executes a script or an application, and opens the output in a specified application. This action can take the value of a field or fields as input, and can only be executed on a single event.
- **Launch a Web Browser:** Launches a Web browser with a specified URL. This action can take the value of a field or fields as input, and can only be executed on a single event.
- JavaScript Actions configured through the Action Manager. JavaScript actions can be executed on a single event or multiple events.

NOTE: Some JavaScript action plug-ins require a correlated event or incident as input. Actions configured from these plug-ins are excluded from the Event Menu Configuration list. This Action Plugin property is defined by the developer.

To add a command to the right-click menu:

1. Click the *Admin* tab.
2. In the Admin Navigator, click *Admin* > *Event Menu Configuration*.
3. Click *Add*. The Event Menu Configuration window opens.

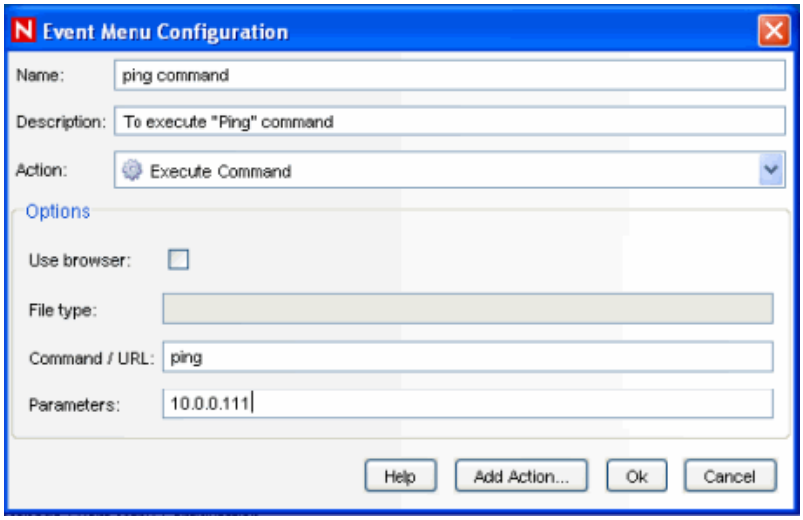

4. Specify a name and description.

   To place the command in a folder, provide folder name/command name in the *Name* field.

5 Select an action from the drop-down menu or click *Add Action* to configure a new JavaScript action. The available settings vary based on which action is chosen:

Option	Description
Use browser	Displays the output of your command by using the defaults configured for the Web browser, based on the file type. This is only available with the Execute Command action.
File Type	If you selected the Execute Command Action, if your browser settings are set up to use the default browser, and if you selected the *Use the following commands* option to launch a browser, you have the option of setting the file type for the output of this command (such as .pdf). This is only available with the Execute Command action if Use browser is selected.
Command/URL	The script or URL that the browser should open or the script or application name to invoke. This is only available with the Execute Command and Launch Web Browser actions.
Parameters	Parameters to represent information from the selected event must be enclosed by percent signs (for example, %InitIP%). For a list of available tags you can use when specifying parameters, click *Help* in the Event Menu Configuration dialog box or see "Sentinel 6.1 Rapid Deployment Event Fields" in the *Sentinel Rapid Deployment Reference Guide*.

This option is only available if your menu configuration browser settings are set to Use Default Browser. For more information, see Section 12.5.8, "Editing Your Event Menu Browser Settings," on page 253.

NOTE: The script or application for Execute Command must be located in `<install_directory>/config/exec`.

6 Click *OK*. The new option is added to the list of menu items when users right-click an event or events.

12.5.2 Cloning an Event Menu Option

1 Open the Event Menu Configuration window.

2 Select a menu item from the table and click Clone.

3 In the Event Menu Configuration dialog box, edit the following as necessary:

- Name
- Description
- Action
- To use a browser or not. For information, see Section 12.5.8, "Editing Your Event Menu Browser Settings," on page 253.
- Command/URL
- Parameters
- Select an action:
  - Execute Command

- Launch Web Browser.
- Any JavaScript action configured in the Action Manager

For a list of available tags you can use when specifying parameters, click Help on the Event Menu Configuration dialog box or see "Sentinel 6.1 Rapid Deployment Event Fields" in the *Sentinel 6.1 Rapid Deployment Reference Guide*.

4 Click OK. The new option is added to the list of menu items in the Event Menu Configuration window.

12.5.3 Modifying an Event Menu Option

1 Open the Event Menu Configuration window.

2 Double-click a menu option.

3 Type your desired changes and click *OK*.

12.5.4 Viewing Event Menu Option Parameters

1 Open the Event Menu Configuration window.

2 Select a menu item and click *Details*.

12.5.5 Activating or Deactivating an Event Menu Option

1 Open the Event Menu Configuration window.

Right-click a menu option and select either *Activate* or *Deactivate*.

12.5.6 Rearranging Event Menu Options

1 Open the Event Menu Configuration window.

2 Select a menu option and click *Up* or *Down*.

12.5.7 Deleting an Event Menu Option

1 Open the Event Menu Configuration window.

2 Select a menu option and click *Delete*.

- Click *Yes* to delete the menu option
- Click *No* to retain the menu option

12.5.8 Editing Your Event Menu Browser Settings

This option allows you to send your Event Menu output to an external browser. The external browser can be any application. It is not restricted to Internet browsers. By changing the file extension, you can launch whatever application is associated with that extension. For example, txt is often associated with Notepad. You can also select to launch a specific program (for example, you can set txt files to be opened by Wordpad or another editor).

1. Open the Event Menu Configuration window.
2. Click Browser.

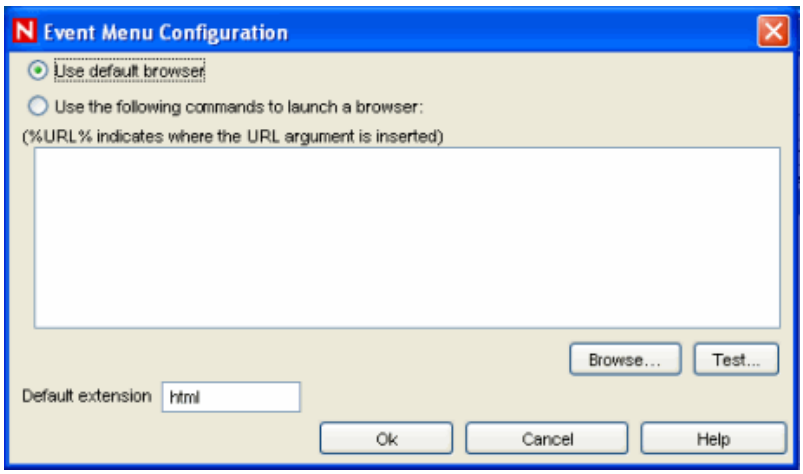

3. Select one of the following two options:
   - **Use default browser:** Uses the default browser set for that particular machine.
   - **Use the following commands to launch a browser:** Allows you to specify a specific application to launch. When you are using a browser other than the default browser, your command line must be followed by a %URL%. For example:

     `C:\Program Files\Internet Explorer\IEXPLORE.EXE %URL%`
   - **Default extension:** This file extension is assumed if the file type in a configured action is blank.

   The following is an example where the output of the Menu Option launches into Internet Explorer.

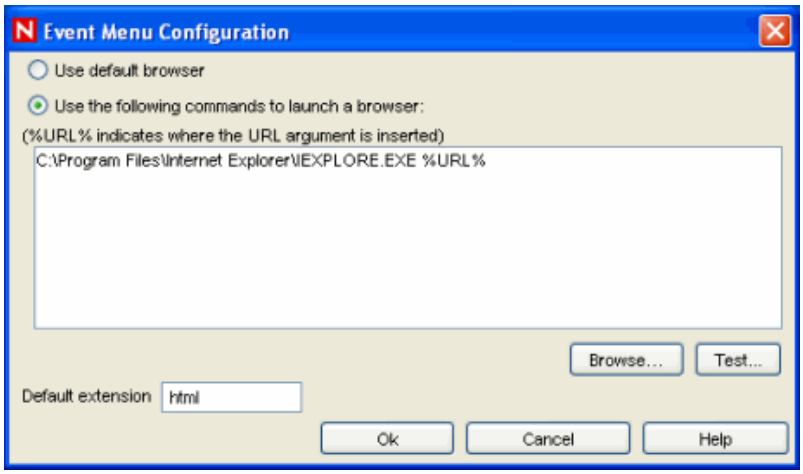

4 After you set your configuration, click *OK*.

12.6 DAS Statistics

This feature is for internal monitoring of your system. It is not intended for the average user. DAS Statistics monitors the following:

- DAS_Binary
- DAS_Core
- Unix Communication Server
- Collector_ Manager
- Correlation _Engine
- Web Server

Statistics includes the following:

- **Service:** Name of the service, such as DAS_Core
- **Time:** Time since the last update
- **num:** Number of requests processed for this entry
- **WaitTime:** Average wait time in seconds for a request before its processing starts
- **Runtime:** Average time to process a request (in seconds)
- **#wait:** Average size of the wait queue
- **#run:** Average size of the run queue

The information is divided into three sections:

- Requests
- Services
- ThreadPools

Requests keeps all the requests by channel (such as services.CorrelationService). Services does the same by service. Sometimes the information provides a breakdown by appending a category under the name, such as Services.CorrelationService or Services.RemoteObjectService.EMap.getMapPK.

For Services, the remote method calls from user-defined services (your XML services) are all under services.RemoteObjectService. Under that it puts the name of the service (such as EMap in the above example) and if asked, the name of the method (getMapPK in the above example).

When a request such as a DAS query is received by a server, a task is created and scheduled. The task is then assigned to a thread pool for execution. There can be more than one thread pool and a thread pool can service multiple services. For that reason, a request needs to wait for an available thread even if the service is not heavily used. If the statistics indicate that the wait time for a request is long and the number of requests for that service is low, check the information about the thread pools.

The numbers next to an entry are the sum for all its children. For example, requests 15 means that there are 15 requests for all requests method calls. Under that, requests.configurations 1 means that 1 of the 15 are to configurations, requests.esecurity.correlation.config 2 means that 2 of the 15 are to esecurity.correlation.config, and so on.

Figure 12-9 *DAS Statistics Window*

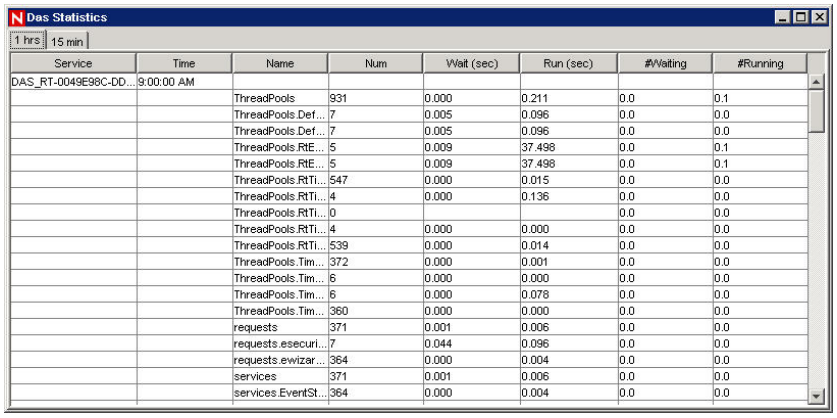

Service	Time	Name	Num	Wait (sec)	Run (sec)	#Waiting	#Running
DAS_RT-0049E98C-DD...	9:00:00 AM						
		ThreadPools	931	0.000	0.211	0.0	0.1
		ThreadPools.Def...	7	0.005	0.096	0.0	0.0
		ThreadPools.Def...	7	0.005	0.096	0.0	0.0
		ThreadPools.RtE...	5	0.009	37.498	0.0	0.1
		ThreadPools.RtE...	5	0.009	37.498	0.0	0.1
		ThreadPools.RtTi...	547	0.000	0.015	0.0	0.0
		ThreadPools.RtTi...	4	0.000	0.136	0.0	0.0
		ThreadPools.RtTi...	0			0.0	0.0
		ThreadPools.RtTi...	4	0.000	0.000	0.0	0.0
		ThreadPools.RtTi...	539	0.000	0.014	0.0	0.0
		ThreadPools.Tim...	372	0.000	0.001	0.0	0.0
		ThreadPools.Tim...	6	0.000	0.000	0.0	0.0
		ThreadPools.Tim...	6	0.000	0.078	0.0	0.0
		ThreadPools.Tim...	360	0.000	0.000	0.0	0.0
		requests	371	0.001	0.006	0.0	0.0
		requests.esecuri...	7	0.044	0.096	0.0	0.0
		requests.ewizar...	364	0.000	0.004	0.0	0.0
		services	371	0.001	0.006	0.0	0.0
		services.EventSt...	364	0.000	0.004	0.0	0.0

The number of requests is especially useful, because you can see where requests are going or where they are concentrated. The # waiting information is useful because it shows how busy the server is. That number should be small. If it is large, new requests (even for simple tasks) need to wait for potentially slow ones. The average run time is very important because it shows which requests are actually taking all the time, as opposed to waiting for others.

12.7 Mapping

A map is a collection of values and keys defined in a CSV or text file. You can enrich your data by using maps to add additional information to the incoming events from your source device. This additional information can be used for correlation and reporting.

You can create your custom maps in addition to the default maps available. You can use event mapping, which allows you to add additional data to an event by using data already present in the event and by referencing and pulling data from an outside source. For more information, see Section 12.8, “Event Configuration,” on page 265 and Section 12.8.1, “Event Mapping,” on page 265.

NOTE: In order to do mapping, your `configuration.xml` file must be pointing to a communication server that has DAS_Binary and DAS_Core connected to it. This is normally the case by default, as long as the communication server and DAS processes are running.

The *Mapping* tab allows you to:

- Add new map definitions
- Edit map definitions
- Delete map definitions
- Update map data

Mapping works together with the Referenced from Map Data Source setting for individual fields under Section 12.8, "Event Configuration," on page 265. You can map by using a string or number range. The following are the default maps available:

- **AccountIdentity:** Contains information about identities and the accounts associated with them. The keys are UserName, UserDomain, and CustomerName (for MSSPs). This map is populated from information in the Account and Identity tables in the Sentinel database.
- **Asset:** Contains the data from the map data source file `asset.csv`. The `asset.csv` is automatically generated from asset data from Sentinel Database when an asset Collector is run. This file can also be populated manually. The keys are PhysicalAssetName and CustomerName (for MSSPs).
- **AssetToRegulation:** Contains the data from the map data source file `AssetToRegulation.csv`. This file must be populated manually.
- **CustomerHierarchy:** Generally used for Managed Security Service Providers (MSSPs). This file can be used to organize customers into a four-level hierarchy. It contains data from the `customerhierachy.csv`. This file must be populated manually. The key is `CustomerName`.
- **IpToCountry:** Contains the data from the `IpToCountry.csv` map data source file. This file must be populated manually.
- **IsExploitWatchlist:** Contains the data from the `exploitDetection.csv` map data source file. (vulnerabilities and threats). The `exploitDetection.csv` file is automatically generated from Advisor and Vulnerability data from the Sentinel Database when either an Advisor feed is completed or a vulnerability Collector is run. The keys are IP, AttackName, DeviceName, and CustomerName (for MSSPs).

To view maps in the GUI:

1. Navigate to the Admin tab and select Map Data Configuration from the Navigation pane or click the Map Data Configuration button .

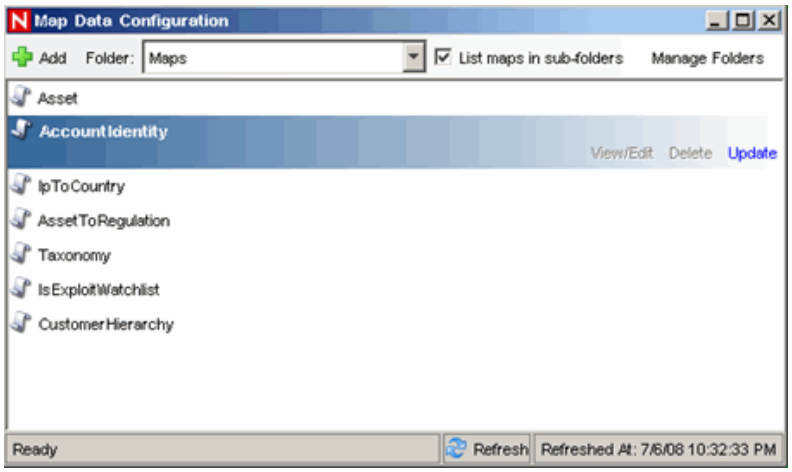

The main Mapping GUI displays a listing of all of the maps that have been defined for the system.

NOTE: Default Sentinel maps cannot be edited or deleted.

12.7.1 Adding Map Definitions

1. Navigate to the *Admin* tab and select *Map Data Configuration* from the navigation pane or click the Map Data Configuration button.
2. Click *Add*.
3. If you are creating a new map folder, click New Directory. Specify a folder name.
4. Ensure that the folder you want to provide your map definition into is selected. (that is, the folder indicates that it is open).
5. Specify your map name.
6. Click Next.

   The Map Type field box is disabled.
7. Select either *Local File* or *Remote File*.
   - **Local File:** Allows you to browse for your file on your local file system on the machine where the Sentinel Control Center was launched.
   - **Remote File:** Allows you to select from existing map source data files on the server where DAS is running. Remote file points to `<install_directory>/data/map_data`.

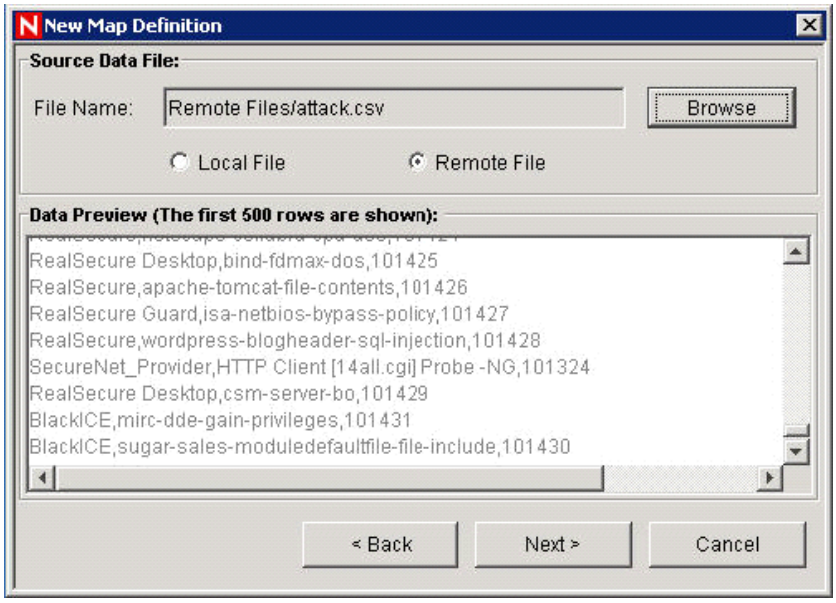

8. Select your map definition file, then click *Next*.

   Only the first 500 rows of the map appear in the interface.
9. In the New Map Definition window, set the following:
   - **Delimiter:** The option are Pipe, Comma, Semicolon, Tab and Other. Specify the delimiter of the data in rows of the map data source file.
   - **Start at row:** Specify the number of rows to skip from the top of the map data source file.

- **Column names:** Specify the column name.
- **Column types:** The currently supported column types are:
  - **String:** A group of characters used as a single object by a computer. A string might consist of a single letter, word, or number. The word FINANCE or IP address 192.168.2.40 might be a string. A string can also consist of a combination of words, spaces, and numbers. The street address of 1313 LION DOG TOWER could be a string.
  - **Number Range:** A range of numbers. For example, 10 to 200 are represented as 10-200. To use the range map functionality, a map definition must have exactly one key column and the key column must be of type NumberRange. If there are any other key columns, or if the key column is of a different type, the mapping service does not consider the map to be a range map.
- **Active columns:** When a column is marked as active, the data in the column is distributed to processes by using maps. All key columns must be active. Only active columns (but not key columns) can be selected as the Map Column under the Event Configuration tab.
- **Key columns:** A unique identifier for the row of data in the map data. If more than one column is selected as a key, the overall key of the map includes all of the columns selected as keys.
- **Column filtering:** A row can be explicitly included or excluded based on matching criteria for a particular column. This can be used to exclude rows from the map source data that are not needed or will interfere with your mapping.

As you configure each setting and filter, the data table automatically updates to allow you to preview your data and to ensure that your data is being parsed as expected.

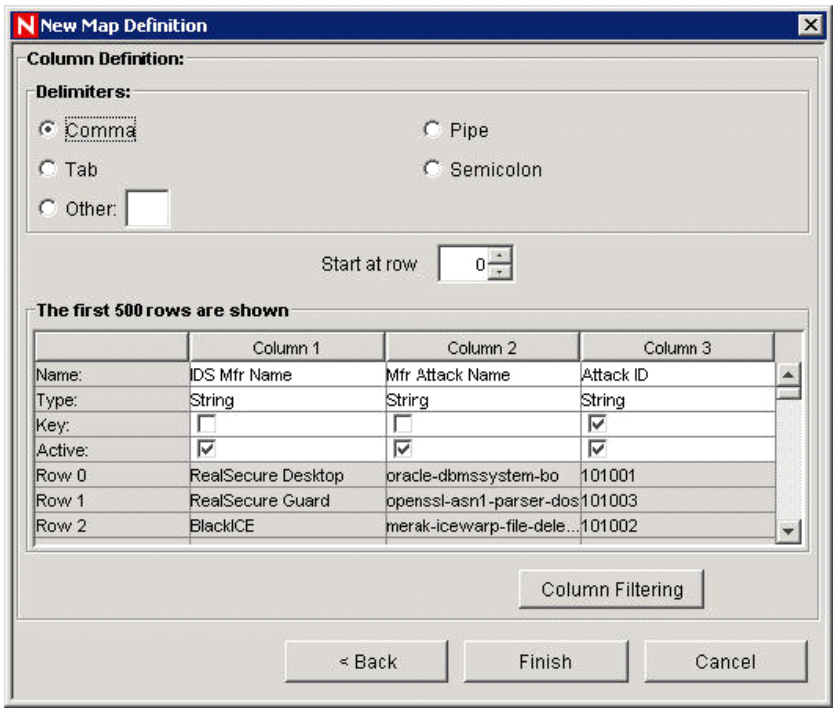

10 After you finish configuring all parameters and filters for the definition, click Finish.

11 If you selected Local File in Step 7, you are prompted to upload your file to the Remote Files virtual folder located at `<install_directory>\data\map_data`.

12 Specify a filename and click *OK*.

12.7.2 Adding a Number Range Map Definition

To use the range map functionality, a map definition must have exactly one key column and the key column must be of type NumberRange. If there are any other key columns, or if the key column is of a different type, the mapping service does not consider the map to be a range map.

To create a range map, select a single column to be the key of the map and select NumberRange as the type of the column. The format of the data in a column of type NumberRange must be m-n, where m is the minimum number in the range and n is the maximum number in the range (that is, 10-200). The maximum number in the range is not included in the range (that is, [m,n)). This means a range of 10-200 only keys off numbers equal to 10 to 199. An example set of data is with the first column as the key:

```
1-2,AA
2-4,AA
4-12,BB
10-20,BB
30-31,BB
100-200,AA
110-120,CC
```

Figure 12-10 *Number Range Map Definition*

The first 500 rows are shown

	Column 1	Column 2
Name:	Range	Value
Type:	NumberRange	String
Key:	☑	☐
Active:	☑	☑
Row 0	1-4	AA
Row 1	4-20	BB
Row 2	30-31	BB
Row 3	100-110	AA
Row 4	110-120	CC
Row 5	120-200	AA

The example table is transformed to:

Figure 12-11 *Table Transformation*

FROM	TO:
1-2,AA	1-4,AA
2-4,AA	4-20,BB
4-12,BB	30-31,BB
10-20,BB	100-110,AA
30-31,BB	110-120,CC
100-200,AA	120-200,AA
110-120,CC	

An example event configuration on the above map might look like:

Figure 12-12 *Event Configuration*

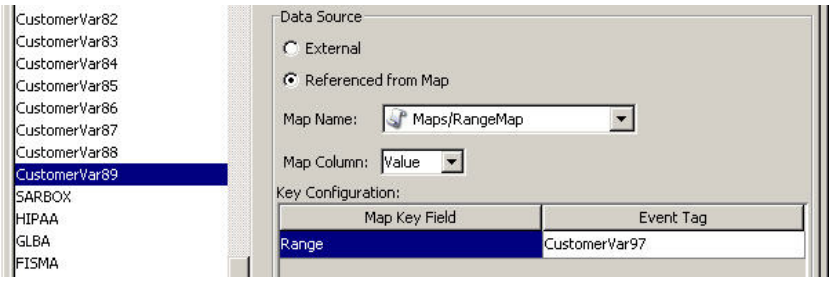

In this example, CustomerVar97 is expected to contain a numeric value or is of a type that can be converted to a numeric value, such as an IP or Date.

When you look into the example range map, the value in CustomerVar97 takes the range map and searches for the range that the value belongs in (if any). Some examples and their results are:

```
CustomerVar97 = 1; CustomerVar89 will be set to AA
CustomerVar97 = 4; CustomerVar89 will be set to BB
   CustomerVar97 = 300; CustomerVar89 will not be set
```

Internally, Sentinel converts IP addresses and dates to an integer for tags of the type IPv4 and Date.

IPv4 tags are:

- TargetIP (dip)
- InitIP (sip)

Date tags are:

- CustomerVar11 to CustomerVar20 (cv11 to cv20)
- DateTime (dt)
- ReservedVar11 to ReservedVar20 (rv11 to rv20)
- DeviceEventTime
- SentinelProcessTime
- BeginTime
- EndTime

For more information on meta tags, see "Sentinel 6.1 Rapid Deployment Event Fields" in the *Sentinel 6.1 Rapid Deployment Reference Guide*.

For example, for the table below, column 1 is numerical range equivalent to an IP range of 10.0.0.0 to 10.0.2.255.

```
167772160-167772415,AAA
167772416-167772671,BBB
167772672-167772927,CCC
```

Using the same setup as the previous example, if:

- The Event Tag is set to TargetIP and key column set to column 1 (range)
- Map Column is set to column 2 (value). The output values are for CustomerVar89.

Figure 12-13 *Number Range Map Definition*

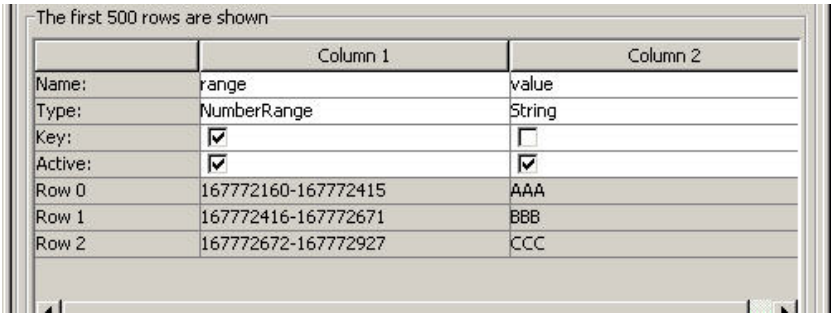

Figure 12-14 *Event Configuration*

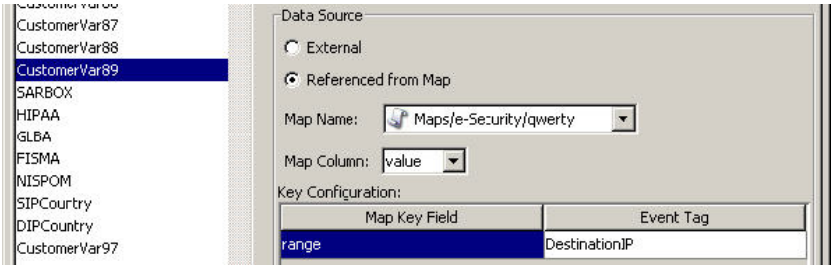

If an event contains a target IP of 10.0.1.14 (equivalent to a numerical value of 167772430), the output for the CustomerVar89 column within the event is BBB.

Sentinel supports the following number ranges:

- Range from negative number to negative number (for example, “-234—34”)
- Range from negative number to positive number (for example, “-234-34”)
- Range from positive number to positive number (for example, “234-236”)
- Single number range (negative) (for example, “-234”). In this case, the minimum and the maximum are both “ -234”.
- Single number range (positive) (for example, “234”). In this case, the minimum and the maximum are both “234”.
- Range from negative number to max number (for example, “-234-”). In this case, the minimum is “-234” and the maximum is (2^63 - 1).
- Range from positive number to max number (for example, “234-”). In this case, the minimum is “234” and the maximum is (2^63 - 1).

NOTE: In all cases, the min must be less than or equal to the max (for example, “-234- -235” is not valid).

12.7.3 Editing Map Definitions

1 Navigate to the *Admin* tab and select *Map Data Configuration* from the navigation pane or click the *Map Data Configuration* button.

2 Expand the folder of interest.

3 Select a map definition and click *Edit*.

The editing function is disabled for map definitions that are under the UNMANAGED ITEMS folder.

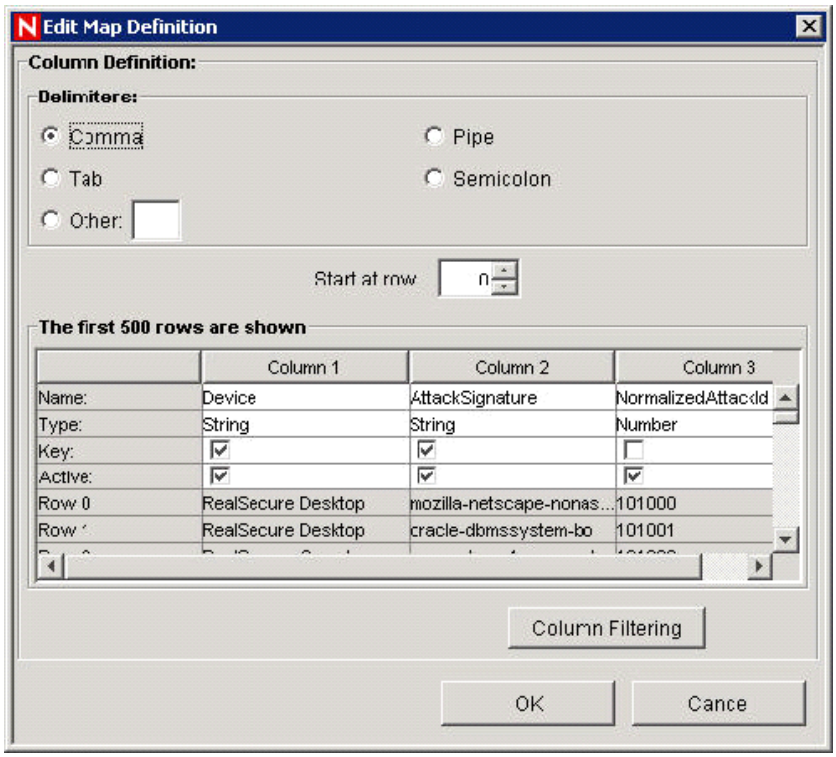

The edit function allows you to:

- Set your delimiters
- Activate or deactivate a column
- Set your column keys
- Set a column filter
- Set which row to start your map
- Rename your columns

4 After making your changes, click *OK*.

12.7.4 Deleting Map Definitions

1 Navigate to the *Admin* tab and select *Map Data Configuration* from the navigation pane or click the *Map Data Configuration* button.

2 Expand the folder of interest.

3 Select the map definition to be deleted.

4 Click *Delete*.

NOTE: Default Sentinel maps cannot be edited or deleted.

12.7.5 Updating Map Data

Updating allows you to replace the map source data file of a map on the server running DAS with another file. Your new map source data file must have the same delimiter, number of columns, and overall structure as the existing map data source file in order for the map to function properly after the update. The new map source data file should only differ from the existing file by the values that appear in the columns. If the new map source data file has a different structure than the existing file, use the Edit feature to update the map definition.

Map updates can be performed on demand from the Sentinel Control Center. To set up an automated process to update map data, you can run an equivalent process from the command line using `map_updater.sh` or `map_updater.bat`.

There are two map locations: the location referenced by the Event Map Configuration (which is a user-defined location) and the location where Sentinel stores its internal representation of the map (`<install_directory>/data/map_data`). The internal representation of the map should never be manually updated.

- "Updating Map Data from the Sentinel Control Center" on page 263
- "Updating Map Data by Using the Command Line" on page 264

Updating Map Data from the Sentinel Control Center

1 If you have not already done so, create a file containing the new map source data.

This file can be generated (for example, from a data dump script), created manually from scratch, or be an edited version of the existing map data source file. If needed, you can obtain the existing map data source file from one of the following locations:

`<install_directory>/data/map_data`

2 Navigate to the *Admin* tab and select *Map Data Configuration* from the navigation pane or click the *Map Data Configuration* button.

3 Expand the folder of interest, select the mapping to be updated, then click *Update*.

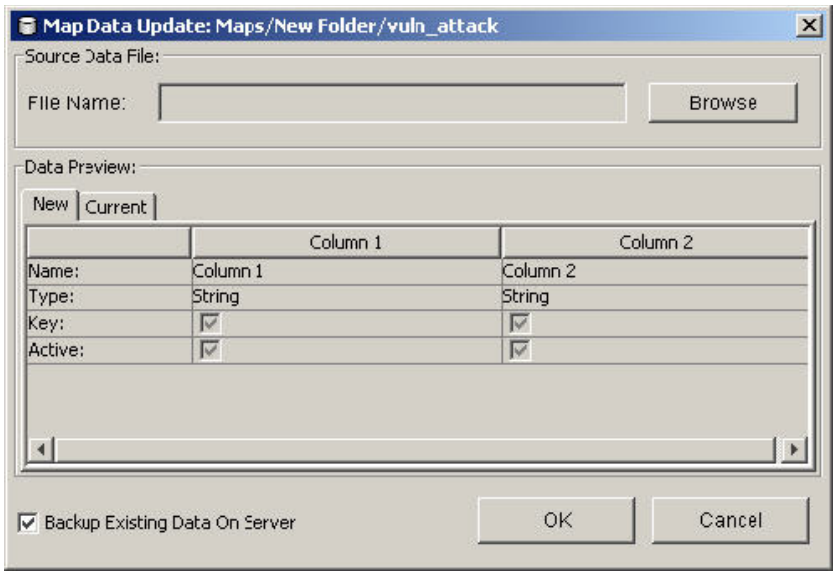

4 Select the new map data source file by clicking Browse and selecting the file with the new map data.

After you select the file, the data from the new map data source file displays under the New tab. The map data you are replacing is under the Current tab.

5 Deselect or leave the default setting for Backup Existing Data On Server.

Enabling this option results in a backup of the existing map data source file being put in the `<install_directory>/data/map_data` folder. The prefix of the name of the backup map data source file is the name of the existing map data source file. The end of the filename contains a set of random numbers followed by the `.bak` suffix. For example: `vuln_attacks10197.bak`.

6 Click *OK*.

The data from the new map data source file is uploaded to the server, replacing the contents of the existing map data source file. After the source data is completely uploaded, the map data is regenerated and distributed to map clients (For example, Collector Manager).

Updating Map Data by Using the Command Line

1 If you haven't already done so, create a file containing the new map source data.

This file can be generated (for example, from a data dump script), created manually from scratch, or be an edited version of the existing map data source file. If needed, you can obtain the existing map data source file from one of the following locations

`<install_directory>/data/map_data`

2 Log into the Sentinel database.

3 Find the UUID for the map in the MD_CONFIG table (refer to the CONFIG_ID column for the appropriate map listed in the VALUE column).

4 On the Sentinel Server machine, log in as esecadm.

5 Run the following command:

```
map_updater.sh <uuid> <source path> [nobackup]
```

6 The data from the new map data source file is uploaded to the server, replacing the contents of the existing map data source file. After the source data is completely uploaded, the map data is regenerated and distributed to map clients (for example, Collector Manager).

Unless the optional -nobackup argument is added, the previous map data is saved in a backup file on the server. Enabling this option results in a backup of the existing map data source file being put in the `<install_directory>/data/map_data` folder. The prefix of the name of the backup map data source file is the name of the existing map data source file. The end of the filename contains a set of random numbers followed by the `.bak` suffix. For example: `vuln_attacks10197.bak`.

12.8 Event Configuration

- Section 12.8.1, “Event Mapping,” on page 265
- Section 12.8.2, “Renaming Tags,” on page 269

12.8.1 Event Mapping

Event Mapping is a mechanism that allows you to add data to an event by using data already in the event to reference and pull in data from an outside source. The outside data source is a map, which is defined by using Map Data Configuration. The data already in the event that should be used as the reference into the map and the data to be pulled from the map into the event are specified by using the *Events* tab.

Because virtually any data set can be made into a map, Event Mapping is useful for incorporating data from elsewhere in your organization into the event stream. Some opportunities Event Mapping provides are:

- Regulatory compliance monitoring
- Policy compliance
- Response prioritization
- Enabling security data to be analyzed related to business operations
- Enhancing accountability

When an Event Mapping is defined, it is applied system-wide to all events from all Collectors. Additionally, Sentinel automatically distributes map data to all processes that perform event mappings as well as keeping the map data in these processes up-to-date. For these reasons, Event Mapping provides significant capabilities to support enterprise deployments.

Event Mapping is made up of four main parts:

- **Controller:** Stores all map information
- **Distributor:** Automatically redistributes modified maps to those processes that registered for the map
- **Monitor:** A monitor to detect changes in map source data
- **Generator:** Generates maps from source data

One application of Event Mapping is Sentinel's Asset Data functionality. For example, asset information is collected and stored in the Sentinel Database asset schema and is represented by a Physical Asset Entry. Soft assets, such as services and applications, are represented by an entry that

is linked to a physical asset. The primary automated update mechanism for asset data is through an asset Collector reading data from a scanner such as Nmap. The asset Collector automates the retrieval of asset information by reading asset data from the scanner and populating the asset schema tables with this data. For Event Mapping, asset information is mapped from the destination IP and source IP.

There are two types of data sources:

- **External:** A Collector populates the value in the event tag.
- **Referenced from Map:** Data is retrieved from a map to populate the tag.

Figure 12-15 *Data Sources*

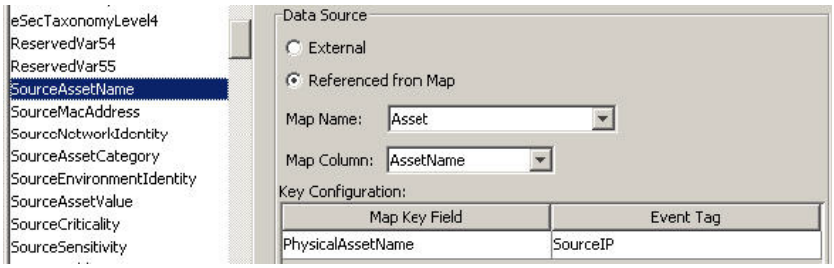

In the above illustration, the SourceAssetName tag is populated from the map called `Asset` (which has `asset.csv` as its map data source file). The specific value for SourceAssetName is taken from the AssetName column from the Asset map. The PhysicalAsssetName column is set as the key. When the InitIP tag of the event matches one of the source IP values in the PhysicalAsssetName column of the map, the row with the matching key is used to intersect the AssetName Column. For instance, in the following example the IP corresponds to AssetName Finance35.

NOTE: When a column is set as a key, it does not appear in the Column drop-down field.

Figure 12-16 *Physical Assent Name Corresponds to the Asset Name*

PhysicalAssetName	CustomerID	MacAddress	AssetName
198.168.1.91			Marketing01
198.168.1.95			Marketing02
198.168.1.96			ProgramMgmt03
198.168.1.98			Finance34
198.168.1.100			Finance35

You can have more than one column set as a key if you do not want the map to be a range map (range maps can only have one key column, with that column type set to NumberRange). For instance (with the column type set to String) the AttackId tag has the DeviceName (name of the security device) and DeviceAttackName columns set as keys and uses the NormalizedAttackID column in the AttackNormalization map for its value. In a row where the DeviceName event tag matches the data in the Device map column and the DeviceAttackName matches the data in the AttackSignature map column, the value for AttackId is the value in the NormalizedAttackID column. The configuration for Event Mapping just described is as follows:

Figure 12-17 *Event Mapping Configuration*

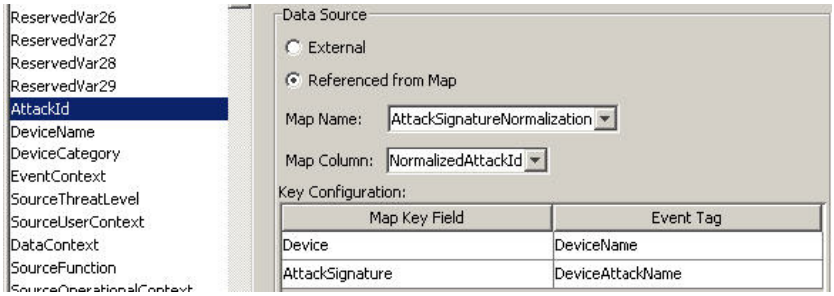

Figure 12-18 *Device and Attack Signature Corresponds to the Asset Name*

Device	AttackSignature	NormalizedAttackId	
Secure	BackDoorProbe (TCP 1234)	3	Trojan: Backdoor.SubSeven
Secure	BackDoorProbe (TCP 1999)	3	Trojan: Backdoor.SubSeven
Dragon	RWALLD:SYLOG-FORMAT	4	Sun Microsystems Solaris rwall Elevated F
Snort	RPC TCP rwalld request	4	Sun Microsystems Solaris rwall Elevated F
Snort	RPC UDP rwalld request	4	Sun Microsystems Solaris rwall Elevated F
Snort	WEB-IIS foxweb.dll access	12	Microsoft Exchange Server Arbitrary Code
RealSecure	SMTP_Exchange_Verb_DoS	12	Microsoft Exchange Server Arbitrary Code

To configure event tags (columns) to use mapping:

1. Navigate to the Admin tab and click Event Configuration in the navigation pane or click the Event Configuration button.
2. Select an event tag entry from the Event Columns list.

   The original Event Tag name displays above the Label field. In addition, the description of the event column is provided.

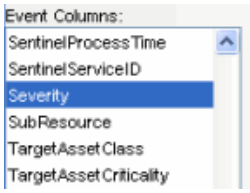

3. Click Referenced from Map to configure the event tag to be populated with data from a map.

   or

   Click External to keep whatever value the Collector put in the event tag (if any).
4. Click the Map Name field down-arrow.

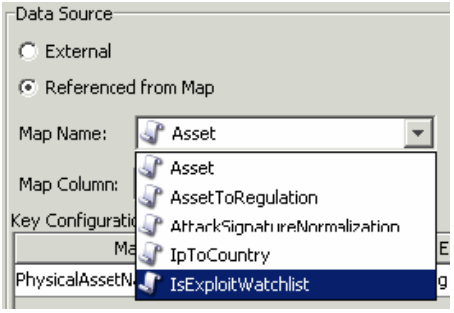

Select one of the available default maps or select a map you have created.

5 Click the Map Column field down-arrow and select a Map Column name. Depending on your Map Name choice in the previous step, these values vary.

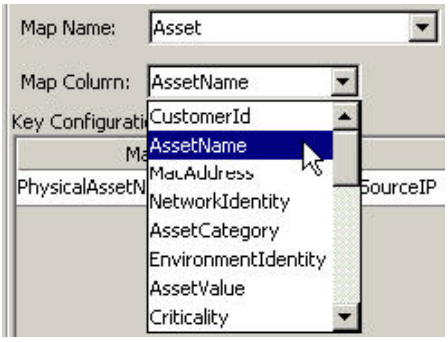

- **_EXIST_ :** This is a special map column that exists in every map. If this map column is selected, a “1” is placed in the event tag if the key is in the map data. If the key is not in the map data, a “0” is placed in the event tag.
- **All other choices:** Names of active columns within the map definition that are not set as a key (for example, the CustomerId column in Asset or the NormalizedAttackId column in AttackNormalization)

6 In the key configuration, select the event tag for each row in the table in the Event Tag column that will be matched against the map key column specified in the corresponding Map Key Field column. The rows in the Key Configuration table depend on the Map Name selected.

A key is a unique identifier for the row of data in the map data.

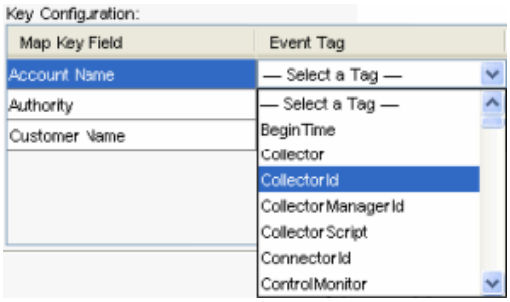

7 Click *Apply*.

Clicking Apply saves the changes you made for the currently selected event column in a temporary buffer. If you don't click Apply, the changes you made to the previously selected event column are lost when you select a different event column. Changes won't be saved to the server until you click Save.

8 If you want to edit the event mapping of another Event column, repeat Step 2 through Step 7.

Remember to click Apply after editing the Event Mapping of each Event column.

9 Click Save.

Clicking Save saves your changes to the server. The save function saves all changes stored in the temporary buffer.

12.8.2 Renaming Tags

The Event Configuration window also allows you to assign names to existing event tag labels. For example, you can rename the label for event tag Ct2 to City. Doing this results in the event tag that formerly appeared in the Sentinel Control Center as Ct2 to now appear as City. Event tags appear in the Sentinel Control Center in places such as filters, correlation rules, and Active Views.

Renaming tags does not change the name of the variable in Collector scripts or in internal Sentinel representations of the tag. For example, even if the event tag labeled Ct2 is renamed to City, the variable that must be used in a Collector script to reference this meta tag is still s_CT2. Any references to this variable in correlation or filters still work, even if they were originally written using Ct2.

Below is a before and after illustration of this feature in an Active View.

Figure 12-19 *Active View Window: Before*

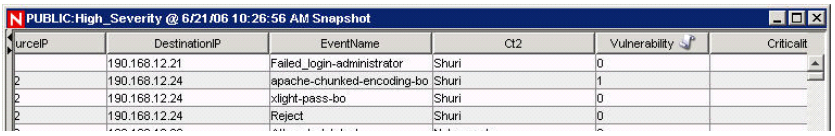

Figure 12-20 *Active View Window: After*

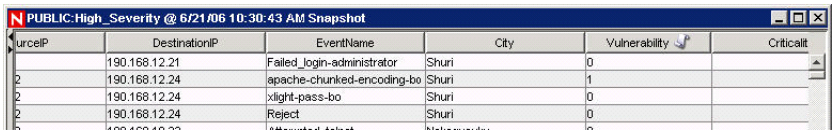

To rename an event column:

1 Click Event Configuration in the navigation pane or click the Event Configuration button.

The original Event Column name displays above the Label field. In addition, the description of the event column is provided.

2 Select an event column entry.

3 Specify a new value for your Event Column in the *Label* field.

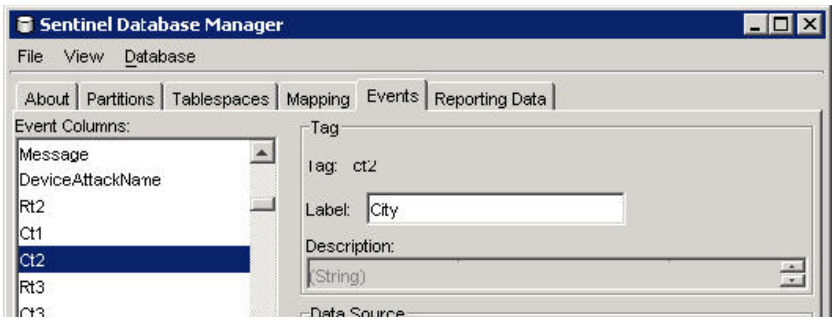

4 Click *Apply*.

Clicking Apply saves the changes you made for the currently selected event tag in a temporary buffer. If you don't click Apply, the changes you made to the previously selected event tag are lost when you select a different event tag. Changes aren't saved to the server until you click Save.

5 Click *Save*.

Clicking *Save* saves the changes to the server. The save function saves all changes stored in the temporary buffer.

6 In order for changes to be visible in Sentinel Control Center, close and reopen any Sentinel Control Centers that are running.

12.9 Report Data Configuration

The Report Data Configuration option allows you to enable and disable summaries or aggregate tables in the Sentinel database. Enabling a summary allows aggregation to start computing the counts for that particular summary and shortens the execution time for any report that uses the summary table. Sentinel Top 10 reports use summary tables.

A summary is a defined set of attributes that make up the key for which to compute the number of unique occurrences (event count) by each hour time period (event time). For EventSevDestPortSummary, it saves the count of events for each unique combination of destination port and severity for an hour. These saved computations of the event data allow for quicker summary reporting and querying. Certain summaries need to be active in order for the summary reports to be accurate.

Aggregation is the process of calculating the running count for all active summaries as events flow through the system. These running counts are saved to the database in the summary tables.

Summaries Benefits:

- Greatly reduced event data set
- Conformed dimensions that allow the ability to drill down, roll up and drill across on event data
- Summary reports run much faster with precomputed summaries

Aggregation Benefits:

- Only processes active summaries
- Does not affect event insertion into the real-time database.

Report Data Configuration tab allows you to:

- Enable/disable any predefined summaries
- View attributes of each summary
- See the validity of a summary for a period of time
- Query which Event files need to be run so that the summary is complete

The following are all summaries already defined in the system.

Table 12-2 *Summary Name Description*

Summary Name	Table/Description
EventSrcSummary	EVT_SRC_SMRY_1 Sums the event count by source IP, source asset information, source port, source user, taxonomy, event_name, resource, Collector, protocol, severity, and event time by hour.
EventDestSummary	EVT_DEST_SMRY_1 Sums the event count by destination IP, destination asset information, destination port, destination user, taxonomy, event_name, resource, Collector, protocol, severity, and event time by hour.
EventSevDestTxnmySummary	EVT_ DEST_TXNMY_SMRY_1 Sums the event count by destination IP, destination asset information, taxonomy, severity, and event time by hour.
EventSevDestEvtSummary	EVT_DEST_EVT_NAME_SMRY_1 Sums the event count by destination IP, destination event asset, taxonomy, event name, severity, and event time by hour.
EventSevDestPortSummary	EVT_PORT_SMRY_1 Sums the event count by destination port, severity, and event time by hour.
EventSevSummary	EVT_SEV_SMRY_1 Sums the event count by severity and event time by hour.

- Section 12.9.1, “Disabling or Enabling a Summary,” on page 271
- Section 12.9.2, “Viewing Information for a Summary,” on page 272
- Section 12.9.3, “Checking the Validity of a Summary,” on page 272
- Section 12.9.4, “Query the Event Files for a Summary,” on page 273
- Section 12.9.5, “Running the Event Files for a Summary,” on page 274

12.9.1 Disabling or Enabling a Summary

1. Click *Report Data Configuration* in the navigation pane or click the *Report Data Configuration* button.

2 To disable a summary, click *Active* in the Status column until it changes to say *Inactive*.

3 To enable a summary, click *InActive* in the Status column until it changes to say *Active*.

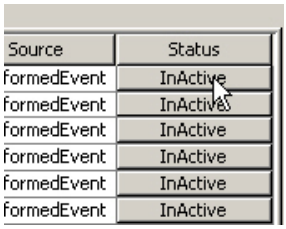

12.9.2 Viewing Information for a Summary

1 Click Report Data Configuration in the navigation pane or click the Report Data Configuration button.

2 Click the … button in the Attributes column to see the attributes that makeup a summary.

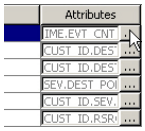

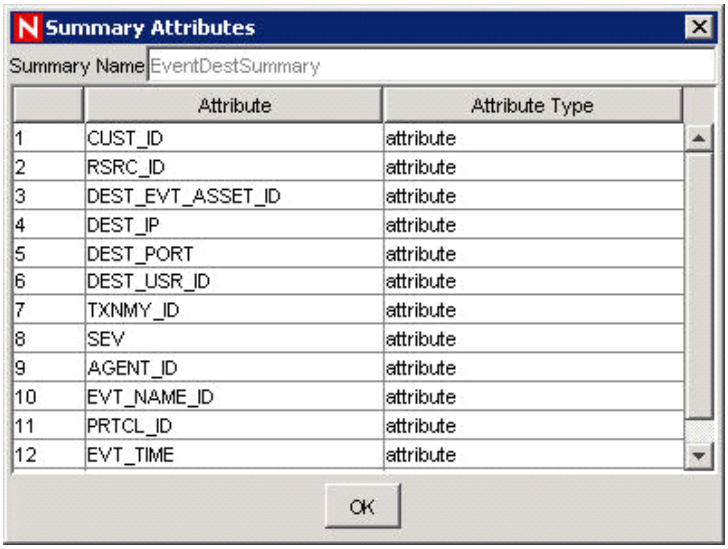

12.9.3 Checking the Validity of a Summary

1 Click *Report Data Configuration* in the navigation pane or click the *Report Data Configuration* button.

2 Select *Status*.

3 Select the summary or summaries you want to query.

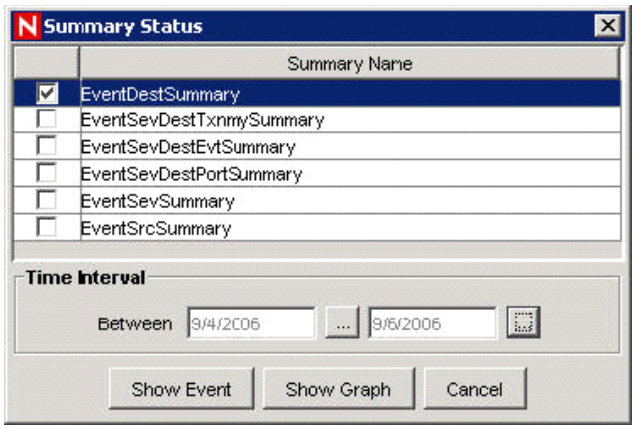

4 Select a time interval.

5 Click *Show Graph*.

The green bars signify that the summary is complete for that time frame. The red sections signify that the summary is missing data during that time period.

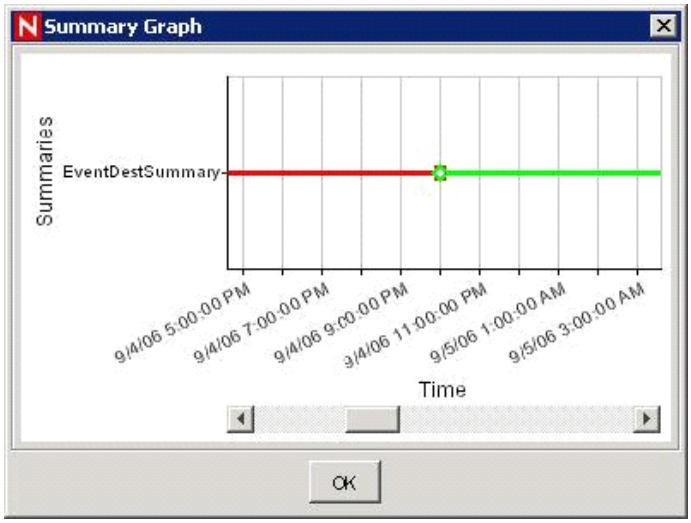

12.9.4 Query the Event Files for a Summary

1 Click *Report Data Configuration* in the navigation pane or click the *Report Data Configuration* button.

2 Select *Status*.

3 Select the summary or summaries you want to query.

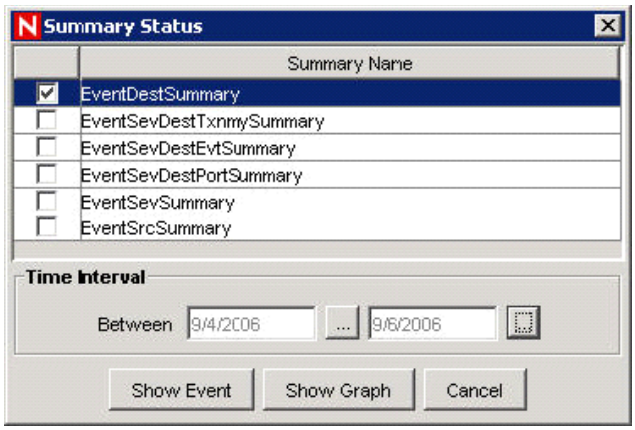

4 Select a time interval.

5 Click *Show Event*.

6 The event files needed to complete the summary display in a list format.

To complete summaries, see Section 12.9.5, "Running the Event Files for a Summary," on page 274.

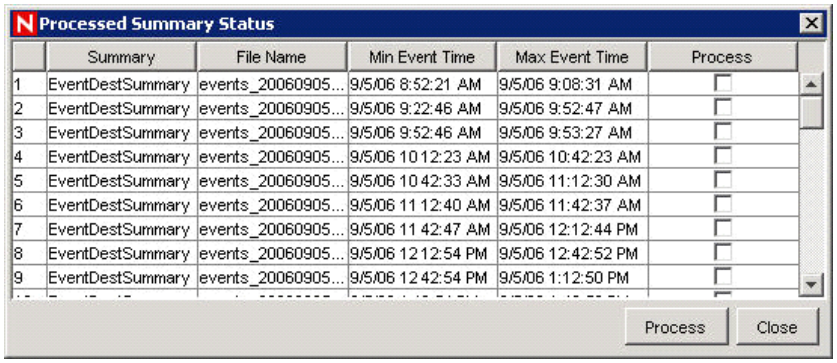

12.9.5 Running the Event Files for a Summary

1 Click *Report Data Configuration* in the navigation pane or click the *Report Data Configuration* button.

2 Select *Status*.

3 Select the summary or summaries you want to query.

4 Select a time interval.

5 Click *Show Event*.

The event files needed to complete the summary display in a list format.

6 Select the event files that you want to run so that the summary is complete.

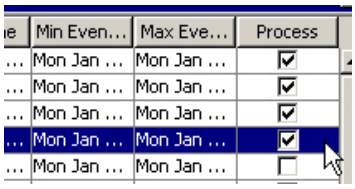

7 Click *Process*.

12.10 User Configurations

You must have the user permission in order to work in the User Configuration window.

User configuration allows you to:

- Section 12.10.1, “Opening the User Manager Window,” on page 275
- Section 12.10.2, “Creating a User Account,” on page 275
- Section 12.10.3, “Modifying a User Account,” on page 280
- Section 12.10.4, “Viewing Details of a User Account,” on page 280
- Section 12.10.5, “Cloning a User Account,” on page 280
- Section 12.10.6, “Deleting a User Account,” on page 281
- Section 12.10.7, “Terminating an Active User Session,” on page 281
- Section 12.10.8, “Adding an iTRAC Role,” on page 281
- Section 12.10.9, “Deleting an iTRAC Role,” on page 282
- Section 12.10.10, “Viewing the Details of a Role,” on page 282

NOTE: The Sentinel Database Administrator, Sentinel Administrator, Sentinel Application User, and Sentinel Report User are created during installation.

12.10.1 Opening the User Manager Window

1 Click the *Admin* tab.

2 Click *Admin* > *User Configuration*.

12.10.2 Creating a User Account

In order to meet stringent security configurations required by Common Criteria Certification, Sentinel requires a strong password with the following characteristics:

- Select passwords of at least 8 with characters in length that includes at least one uppercase letter, one lower case letter, one special symbol (!@#$%^&*()_+), and one numeral (0-9).
- Your password should not contain your e-mail name or any part of your full name.
- Your password should not be a common word. For example, it should not be a word in the dictionary or slang in common use.

- Your password should not contain words from any language, because numerous password-cracking programs exist that can run through millions of possible word combinations in seconds.
- You should select a password you can remember that is still complex. For example, Msi5!YOld (My Son is 5 years old) or IhliCf5#yN (I have lived in California for 5 years now).

NOTE: Do not use \ and ‘ in the username and password because the database does not allow these characters.

To use this feature, you must have the User Management user permission. For more information, see the *Sentinel 6.1 Rapid Deployment Reference Guide*.

- “Creating a Local User Account for Sentinel” on page 276
- “Creating an LDAP User Account for Sentinel” on page 277
- “Creating a Domain User Account for Sentinel” on page 279

Creating a Local User Account for Sentinel

1. Select the *Admin* tab.
2. Open the *User Configuration* folder.
3. Open the *User Manager* window.
4. Click *Add a new User*

or

Right-click any user and select *Add User*.

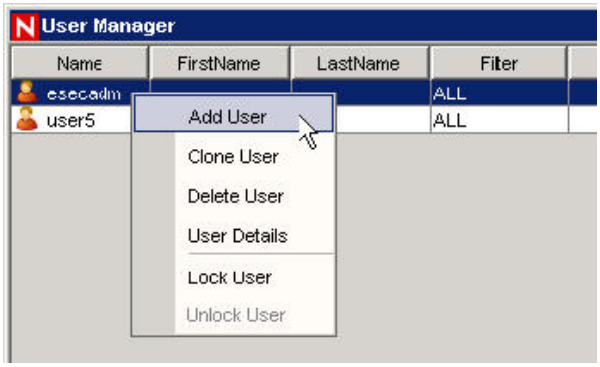

5. Under Authorization:
   - Select Local for Authentication.
   - Specify the username.
   - Specify the password.
   - Confirm the password.

6 For a Security Filter, click the down-arrow. The Filter Selection window displays and shows all public filters.

7 Select a filter and click *Select* or click *Add* to create, then select a new filter.

After assigning a security filter to a user, you cannot delete that filter.

8 (Optional) Under *Details*, specify:

- First Name
- Last Name
- Department
- Phone
- Email

9 Click the *Permissions* tab and assign user permissions.

10 Click the *Roles* tab and select an iTRAC workflow role for the user.

11 Click *OK*.

Creating an LDAP User Account for Sentinel

NOTE: This option is applicable only for Sentinel Rapid Deployment SP1 and later. By default, this option is disabled. Configure the server as given in "LDAP Authentication" in the *Sentinel Rapid Deployment Install Guide* to enable this option.

---

1 Select the *Admin* tab.

2 Expand the *User Configuration* folder in the navigation tree.

3 Select *User Manager*.

The User Manager window is displayed.

4 Click *Add User* or right-click any user and select *Add User*.

The Add User window is displayed.

5 In the Add user window, perform the following:

5a Select *LDAP* for authentication.

5b Specify the LDAP username based on the value you specified for "Anonymous searches on LDAP directory:" parameter while configuring LDAP authentication.

- **y:** The *User Name* must be the same as the eDirectory username or Active Directory sAMAccountName.
- **n:** The *User Name* need not be the same as the eDirectory username or Active Directory sAMAccountName.

5c Click the drop-down arrow on the *Security Filter* drop-down list.

The Filter Selection window is displayed that lists all the public filters.

5c1 Select a filter, and click *Select* or click *Add* to create a filter, then select the new filter.

After assigning a security filter to a user, you cannot delete that filter.

5d Specify the fully qualified Distinguished Name of the LDAP user in the *LDAP USER DN* field. Do not leave the *LDAP User DN* field empty.

For example:

eDirectory User: cn=sentinel_ldap_user,o=novell

Active Directory User: cn=sentinel_ldap_user,cn=users,dc=test,dc=com

This field is available only if you have specified n for "Anonymous searches on LDAP directory:" parameter while configuring LDAP authentication. For more information, see "LDAP Authentication" in the *Sentinel Rapid Deployment Installation Guide*.

NOTE: If you had opted to perform anonymous searches when you had last run the `ldap_auth_config` script, and now you do not want to perform anonymous searches:

Run the script `ldap_auth_config` script again, and specify n for "Anonymous searches on LDAP directory:". For each existing LDAP user, right-click and select *User Details* and specify the fully qualified DN of the LDAP user in the *LDAP User DN* field.

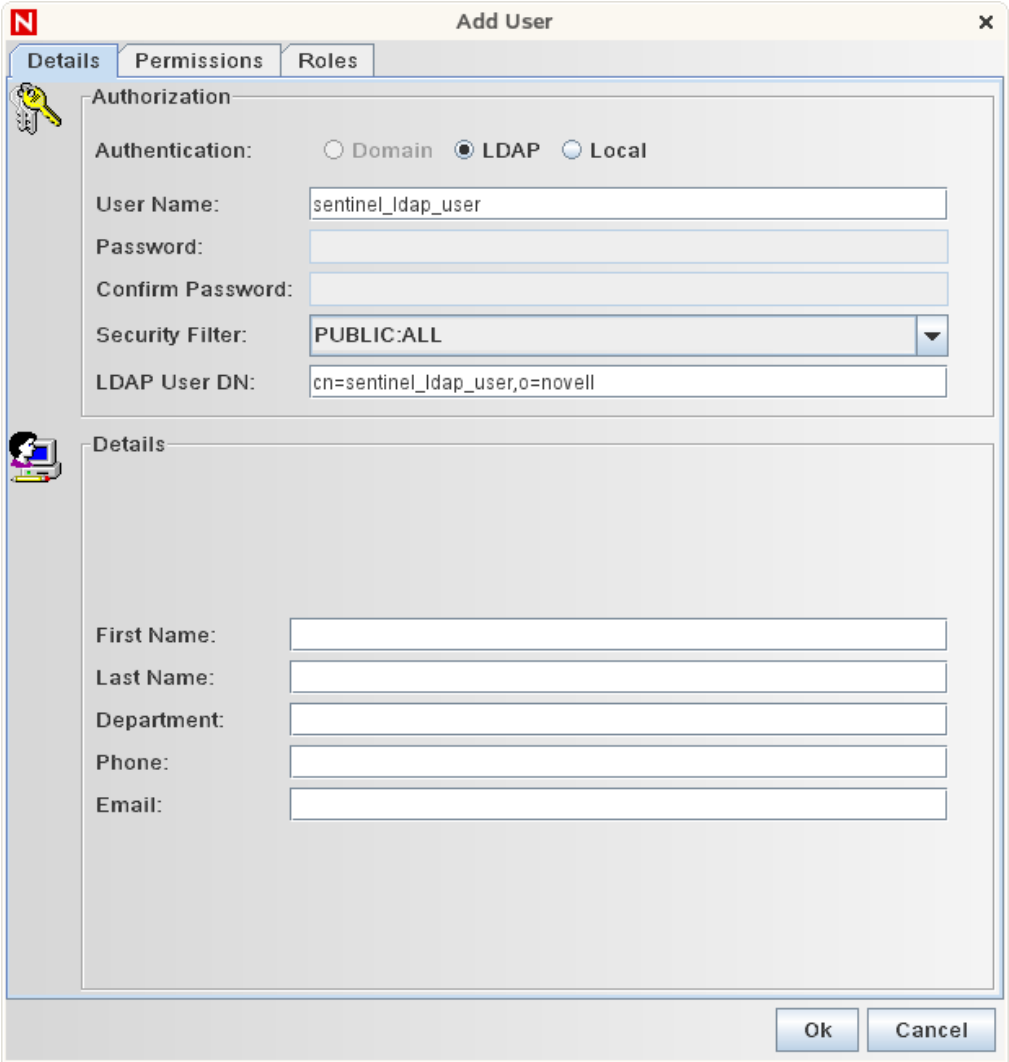

5e (Optional) Under *Details*, specify the following:

- *First Name*
- *Last Name*
- *Department*

- *Phone*
- *Email*

5f Click the *Permissions* tab and assign user permissions. For more information about permissions, see "Sentinel 6.1 Rapid Deployment Control Center User Permissions" in the *Sentinel Rapid Deployment Reference Guide.*.

5g Click the *Roles* tab and select an iTRAC workflow role for the user. This affects what work items appear in the user's work list.

5h Click *OK*.

You can now log in to Sentinel Rapid Deployment Web user interface, Sentinel Control Center, and Sentinel Solution Designer by using your LDAP username and password.

Creating a Domain User Account for Sentinel

NOTE: This option is applicable only in Sentinel Rapid Deployment Hotfix 2 and is used to create LDAP user accounts.

---

1. Select the *Admin* tab.
2. Open the *User Configuration* folder.
3. Open the User Manager window.

   Click *Add* a new User.

   or

   Right-click any user and select *Add User*.

   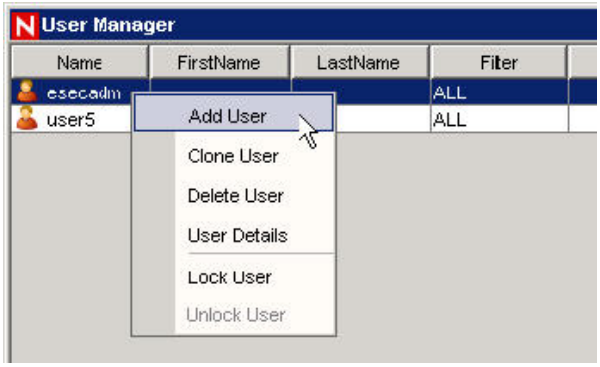

4. Under Authorization:
   - Select Domain authentication.
   - Specify an existing User Name in the form Domain\Username.
5. For Security Filter, click the down-arrow. The Filter Selection window displays and shows all public filters.
6. Select a filter and click *Select* or click *Add* to create and then select a new filter.

   After assigning a security filter to a user, you cannot delete that filter.

(Optional) Under Details, specify:

- First Name
- Last Name
- Department
- Phone
- Email

7 Click the Permissions tab and assign user permissions. For more information about permissions, see "Sentinel 6.1 Rapid Deployment Control Center User Permissions" in the *Sentinel 6.1 Rapid Deployment Reference Guide*.

8 Click the Roles tab and select an iTRAC workflow role for the user. This affects what work items appear in the user's work list.

9 Click *OK*.

NOTE: PostgreSQL does not allow the creation of users named the same as one of the PostgreSQL Reserved words. Also, Sentinel does not allow you to use these names.

12.10.3 Modifying a User Account

To use this feature, you must have the User Management permission.

NOTE: The Sentinel Database Administrator, Sentinel Administrator, Sentinel Application User, and Sentinel Report User are created during installation.

---

1 Open the User Manager window.

2 Double-click a user account or right-click it, then click *User Details*.

3 Modify the account.

4 Click *OK*.

12.10.4 Viewing Details of a User Account

To use this feature, you must have the User Management permission.

1 Open the User Manager window.

2 Double-click a user account or right-click it, then click *User Details*.

Review the details of the user account and close the window.

12.10.5 Cloning a User Account

1 Open the User Manager window.

2 Right-click a user account, then click *Clone User*.

3 Change the user information and the user permissions.

4 Click *Save*.

12.10.6 Deleting a User Account

To use this feature, you must have the User Management permission.

1. Open the User Manager window.
2. Right-click a user account, then click *Delete User*.
3. A Delete box displays. Click *Yes* to delete the User.

12.10.7 Terminating an Active User Session

1. Open the Active User Sessions window.
2. Right-click an active session you want to terminate, then click *Kill Session*.

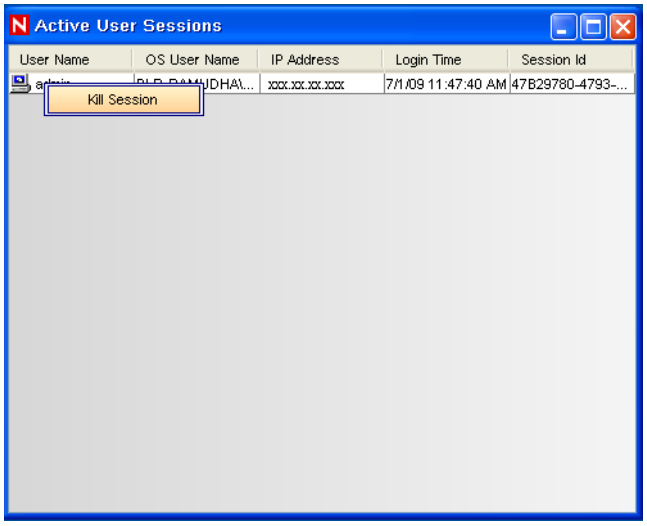

You are prompted for a termination message. This option is provided so that you can inform the user why you are killing the session.

3. Provide a message, then click *OK*.

   or

   Close the window to terminate the session without sending a message.

NOTE: If the client machine has multiple network interfaces, the IP Address displayed in the Active User Sessions window might not be the desired IP address, as the non-loop back IP address of the first NetworkInterface returned by the system is displayed.

12.10.8 Adding an iTRAC Role

1. Open the Role Manager window.
2. Right-click a role, then click *Add New Role*.

12.10.9 Deleting an iTRAC Role

1. Open the Role Manager window.
2. Right-click a role, then click *Delete Role*.

12.10.10 Viewing the Details of a Role

1. Open the Role Manager window.
2. Right-click a role, then click *Role Details*.

Sentinel Data Manager

13

The Sentinel Data Manager (SDM) is a tool by which users can manage the Sentinel database.

- Section 13.1, “Understanding the Sentinel Data Manager,” on page 283
- Section 13.2, “Using the SDM GUI,” on page 283
- Section 13.3, “Using the SDM Command Line,” on page 291

13.1 Understanding the Sentinel Data Manager

The SDM allows users to perform the following operations:

- Monitor Database Space Utilization
- View and Manage Database Partitions
- Configure Auto-Archives
- Configure Auto-Addition of Partitions

Monitor Database Space Utilization, View and Manage Database Partitions, and Configure Auto-Archives operations can be accessed by using the Sentinel Data Manager GUI or by using a command line interface to the SDM.

NOTE: Some SDM functionality has been moved to the Sentinel Control Center, including Event Mapping, Summary Data, and Reporting Data.

13.2 Using the SDM GUI

- Section 13.2.1, “Prerequisites,” on page 283
- Section 13.2.2, “Starting the SDM GUI,” on page 284
- Section 13.2.3, “Connecting to the Database,” on page 284
- Section 13.2.4, “Partitions Tab,” on page 285
- Section 13.2.5, “Tablespaces Tab,” on page 288
- Section 13.2.6, “Partition Configuration,” on page 289
- Section 13.2.7, “Managing Disk Space Allocation,” on page 291

13.2.1 Prerequisites

There are several prerequisites to running the SDM GUI:

- The user must know the following information:
  - Name and password for the Sentinel Database User (dbauser by default)
  - Database host server
  - Database (instance) name
  - Port used for database communications (the default port number is 5432)

13.2.2 Starting the SDM GUI

- "Using the Command Line Option" on page 284
- "Using the Web Interface" on page 284

Using the Command Line Option

1. Log in to the machine as dbauser.
2. Go to `<install_directory>/sdm`
3. Enter the following command:

   ```
   ./sdm
   ```

Using the Web Interface

1. Log in to the Sentinel Rapid Deployment Web interface, then click *Applications*.
2. For more information, see Section 1.1, "Accessing the Novell Sentinel Web Interface," on page 17.
3. Click *Launch Data Manager.*
4. Open the SDM with the Java Web Start Launcher.
5. Specify the server, database, host, and port number.
6. Specify the user credentials and click *Connect*.

   To run the SDM from the command line, see Section 13.3, "Using the SDM Command Line," on page 291.

13.2.3 Connecting to the Database

1. Log into the machine that has the SDM installed.

   If the Sentinel Database Administrator account uses Windows Authentication, you must log into the SDM machine by using the Sentinel Database Administrator account.
2. Start the SDM GUI, using the appropriate procedure:
   - "Using the Command Line Option" on page 284
   - "Using the Web Interface" on page 284
3. Select the database type.
4. Specify the database instance name used during the Sentinel database installation.
5. Specify the database host (hostname or IP address).
6. Specify the port used for database communications.
7. If you are using PostgreSQL Server authentication, specify the Sentinel Database Administrator username and password.

Database	Interface
PostgreSQL	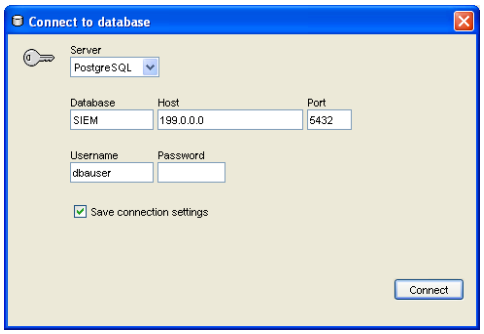

If you select to save your connection settings, the settings are saved to the local `sdm.connect` file. By default the `sdm.connect` file is located in `<install_directory>/bin`. Next time you start the GUI, the connection settings are repopulated from the `sdm.connect` file. This file can be used when you run the SDM from the command line.

8 Click *Connect*. The SDM is now ready for use.

13.2.4 Partitions Tab

The Sentinel database is partitioned by time to simplify maintenance and improve the performance of the database. The *Partitions* tab in the SDM allows users to view and manage database partitions for the tables that hold event data, correlated event data, and summary data.

To view partitions in the GUI:

1 Click the *Partitions* tab.

2 In the drop-down list, select the table you want to see.

The SDM displays the partitions of the currently selected database table.

Each row in the Segments table displays the related database table, time range, status and the name of the partition.

The status of each of the partitions shown in the stegments table has one of the following states:

Table 13-1 *Partition States*

Status	Description
Online	Partition with data that is available for access
Online Current	Partition to which events are currently being inserted
Online Archived	Partition with data that has been archived but is still accessible because the partition has not been dropped
Offline Archived	Partition with data that has been archived and then dropped from the database

Status	Description
Online Archived Imported	Partition with data that has been archived, dropped from the database, and then re-imported into the database

NOTE: If you delete a partition without archiving it, it is deleted from the partition list in the GUI.

Figure 13-1 *Sentinel Data Manager*

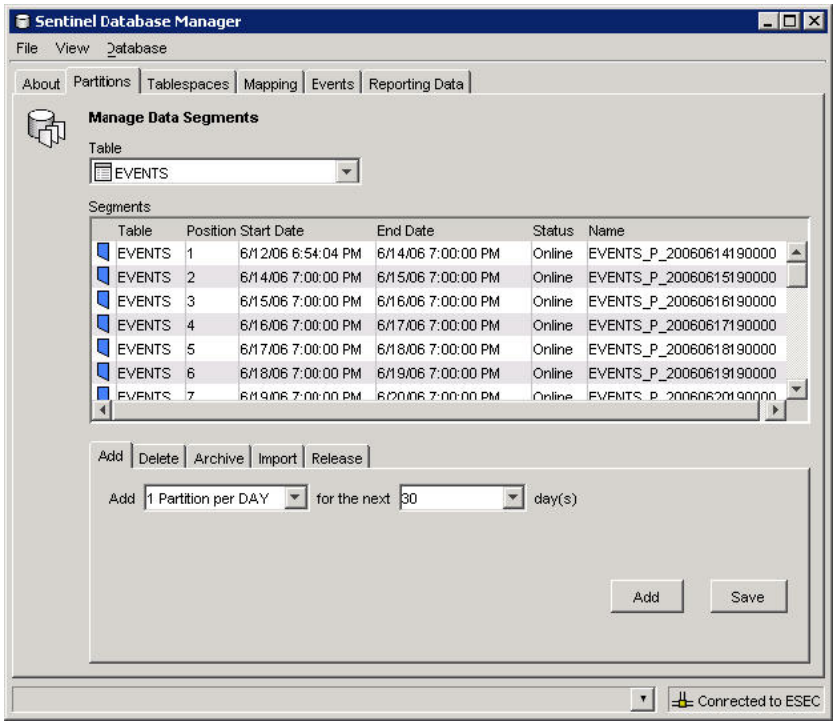

At the bottom of the Partitions page, there are several smaller tabs that allow the user to perform the following operations:

- Add empty partitions to the database
- Delete partitions from the database
- Archive data from partitions to flat files in a specified, preexisting directory
- Import partitions
- Drop partitions

Many of these operations can be executed automatically in the database by using stored procedures, but this page allows the administrator to perform these tasks manually.

To manage partitions:

1. Click the *Partitions* tab.
2. Select the table in the drop-down list.

Sentinel partitioned tables are organized into two groups. One is the EVENTS table group, which includes EVENTS and CORRELATED_EVENTS; the other is the summary table group, which includes all summary, or aggregate, tables. If any one of the tables in the group is selected, the changes apply to all the tables in the group.

3 At the bottom of the window, select the tab that relates to the operation that you want to perform : *Add*, *Delete*, *Archive*, *Import*, or *Release*.

To add partitions

1 Select the *Add* partitions tab.
2 Specify the number of days to use for adding the partitions.

   You can specify the number of partitions in *Partition Configuration* in the SDM GUI.

3 Click *Add*.

To delete partitions:

1 Select the *Delete partitions* tab.
2 Specify the number of days after which older partitions will be deleted.
3 Click *Delete*.

To import partitions:

1 Select the *Import partitions* tab.
2 Select the partition in the Segment table into which the data will be imported.

   You can specify the input directory in the *Archive Destination* field in the *Partition Configuration* tab in the SDM GUI.

3 Click *Import*.

To release imported partitions:

1 Select the *Release partitions* tab.
2 In the Segment table, select the partitions that need to be released.
3 Click *Release*.

Archiving Partitions

Events, correlated events, and aggregation (or summary) tables can all be archived by using the SDM. There are several requirements for archiving:

- The directory to which the partitions are archived must already exist on the database server (not the machine running the SDM); the SDM does not create the directory.
- You cannot archive the data to the `/root` directory.
- You must have permissions to write to the archive directory.

To archive partitions:

1 Select the *Archive* partitions tab.
2 Specify the number of days the older partitions are archived for.

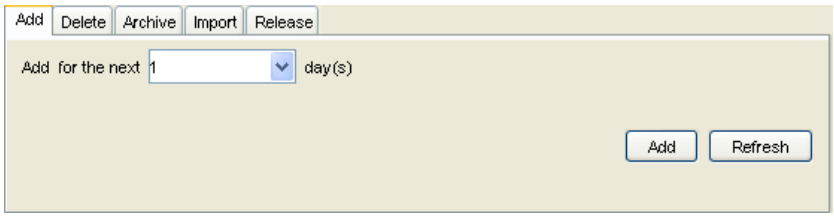

You can specify the archive directory in the Archive Destination field in the Partition configuration tab in the SDM GUI.

3 Click *Archive*.

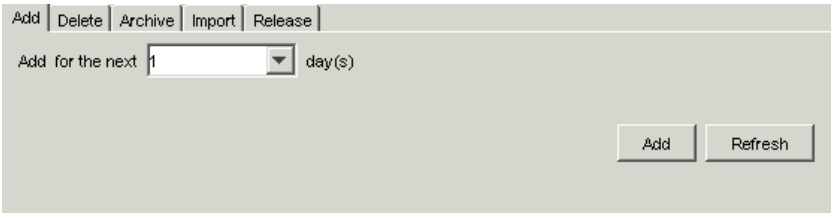

13.2.5 Tablespaces Tab

The *Tablespaces* tab in the SDM allows users to view the current database space utilization, including:

- Total space allocated for each tablespace
- Space used by each tablespace
- Space available (free) for each tablespace.

NOTE:

- PostgreSQL does not allocate a maximum size for a tablespace. Typically tablespaces can grow up to the maximum free space available on a file system. Therefore, Sentinel Rapid Deployment allocates 100% of the free disk space for tablespaces at the time of installation, and is represented as the total space allocated for each tablespace.
- The space used by each tablespace that is displayed is not correct. This is because, the space used by the tablespaces is only considered and does not include other files within that partition.
- The free space available for each tablespace that is displayed is not correct. This is because the free space is calculated on the total available space with the space used only by the tablespace and does not include other files within that partition.

Figure 13-2 *Sentinel Data Manager*

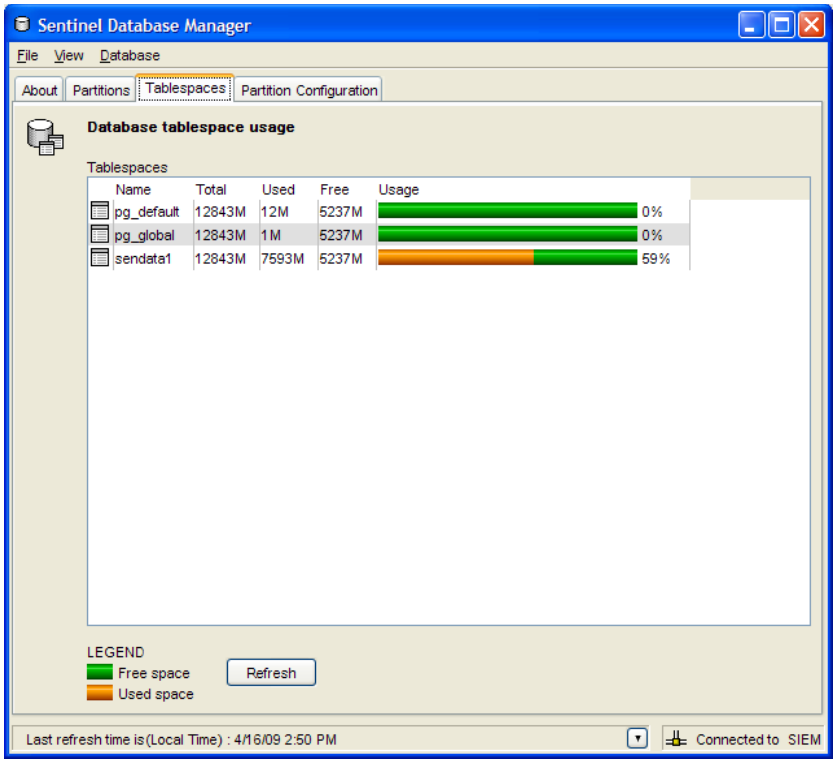

Color-coded bar graphs help to visualize the total space allocated for each tablespace and the percent used of each tablespace.

13.2.6 Partition Configuration

The *Partition Configuration* tab in the SDM allows you to set parameters to auto-archive partitions. It also allows you to auto-add partitions.

To configure auto-archive parameters:

1. Click the *Partition Configuration* tab. The Partition Configuration window displays.

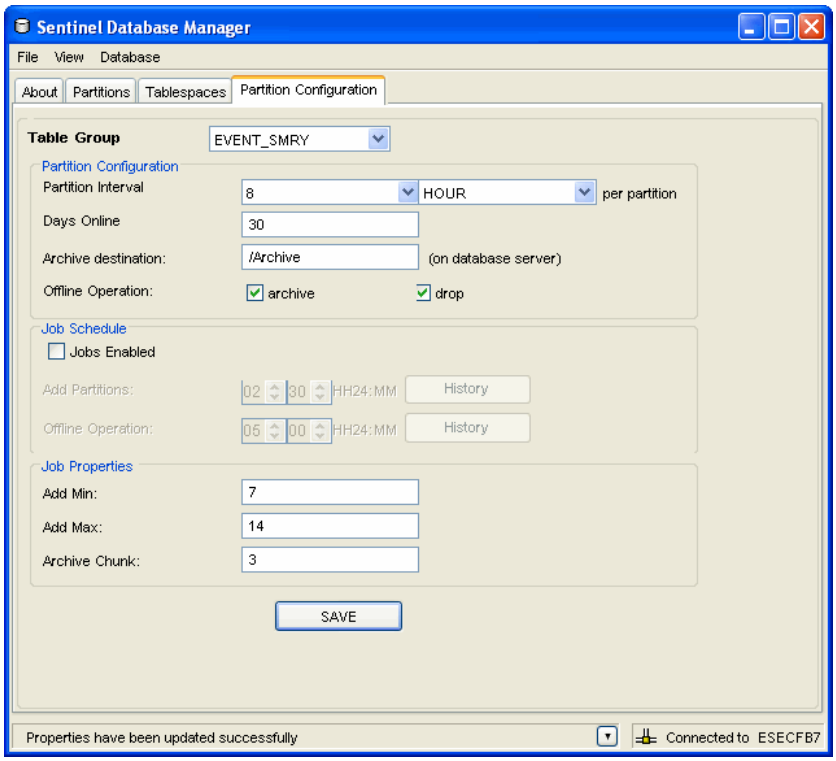

2 Select the table group from the drop-down list.

3 Specify the following partition configuration information:

- **Partition Interval:** Specify the number of partitions that should be created per day or per hour.
- **Days Online:** Number of days of data to keep online in the database.
- **Archive destination:** Specify the destination to store the automatically archived data and the manually archived data.
- **Offline operation:** Select to archive or drop the data.

Data that is dropped without archiving cannot be retrieved by using the SDM. You should almost always select the archive option.

4 Specify the Job Schedule parameters:

- Select *Jobs Enabled* check box if it's not selected. By default, the *Jobs Enabled* check box is selected if you selected this feature during installation.
- Schedule adding partitions and offline operation parameters. Click *Save*.
- Click *History* to view the Job History.

Partition Job scheduling through the SDM is reflected only after the partition job refresh interval.The default partition job refresh interval is 5 minutes.

To change the partition job refresh interval, edit the `partitionJobRefreshInterval` option specified in the `<install_directory>/config/das_core.xml` file. The `partitionJobRefreshInterval` option is provided as part of the `Scheduler` component in the `DAS_Core` container.

After you update the `partitionJobRefreshInterval`, restart the Sentinel service in order for the new refresh interval to take effect.

5 Specify the Job Properties:

- **Add Min:** Minimum number of days of partitions for future data that should exist in the database at any time
- **Add Max:** Maximum number of days of partitions for future data that should exist in the database at any time
- **Archive Chunk:** Minimum number of days of partitions that apply to the total number of days of partitions for the Archive.

NOTE: If the fewer than Add Min days partitions exist in the database, partitions are added until there are enough partitions for Add Max days. Archiving is done in chunks of days so that these database operations are not necessary every day.

6 Click Save.

13.2.7 Managing Disk Space Allocation

The Sentinel Rapid Deployment installation allocates 70% of the free disk space available at the time of installation for database, and the value is specified in the `diskSpaceAllocated` property of the `das_core.xml`. If the database consumes more than what is allocated, the Sentinel services might halt. To monitor such incidents, a scheduler job has been created in the Sentinel services. The scheduler job monitors the disk space for threshold. There are two thresholds for the disk space. The value for the lower threshold is 85% and for the upper threshold it is 95% of the disk space allocated.

The scheduler job runs based on the value specified in the dbStatsInterval property of the `das_core.xml`. When the database/tablespace size reaches the threshold limits, Sentinel Rapid Deployment system warns you of the disk space limit and behaves as follows:

Lower Threshold: When database/tablespace size reaches 85% of the disk space allocated, the Sentinel Rapid Deployment system warns you with an internal audit event indicating the limit. These are logged as internal audit events with severity 4.

Upper Threshold: When the database/tablespace size reaches 95% of the disk space allocated, Sentinel Rapid Deployment system removes the oldest partitions for each partition group until database/tablespace size falls below the threshold level (85%), and also sends an internal audit event for each partition dropped.

NOTE: Sentinel Rapid Deployment does not attempt to remove the online current partition.

13.3 Using the SDM Command Line

The SDM command line functions can be used instead of the GUI. The command line can be used to create a batch file or cron job for SDM operations, but Novell recommends using auto-archiving instead. Auto-archiving can be configured on the *Partition Configuration* tab of the SDM GUI.

- Section 13.3.1, “Prerequisite,” on page 292
- Section 13.3.2, “Syntax of the SDM command,” on page 292
- Section 13.3.3, “Starting the SDM GUI,” on page 292
- Section 13.3.4, “Saving Connection Properties for Sentinel Data Manager,” on page 292
- Section 13.3.5, “Adding Partitions,” on page 293

- Section 13.3.6, “Dropping Partitions,” on page 294
- Section 13.3.7, “Viewing Partition Summaries,” on page 295
- Section 13.3.8, “Archiving Data,” on page 296
- Section 13.3.9, “Importing Data,” on page 297
- Section 13.3.10, “Deleting Imported Data,” on page 298
- Section 13.3.11, “Viewing Sentinel Database Space Usage,” on page 299

13.3.1 Prerequisite

The first step to using the SDM command line is to create a file that stores the connection properties for the database.

13.3.2 Syntax of the SDM command

```
[path to SDM] –action [actionname] [action-specific flags] [path to database
connection file]
```

The specific flags for each action are described below.

13.3.3 Starting the SDM GUI

```
startGui (DEFAULT)
-action startGui [-connectFile <filePath>]
```

13.3.4 Saving Connection Properties for Sentinel Data Manager

The `saveConnection` command saves the database connection details to a specified file. These connection details are necessary for all other SDM command line operations.

If you run the SDM GUI with *Save connection settings* selected, the `saveConnection` command is not necessary. You can use the `sdm.connect` file located in `<install_directory>/sdm`.

The `saveConnection` command uses the following flags:

Table 13-2 *saveConnection command Flags*

Command	Command Flags
-action	saveConnection
-server	<postgresql>
-host	<database host IP Address or host name to connect to>
-port	<database port number to connect to >
-database	<database name/SID>
-driverProps	<Properties File>
-dbuser	<database username>
-password	<database password>

Command	Command Flags
winAuth	Used for Windows authentication. When using this option, -user and –password are not needed.
connectFile	<filenameToSaveConnection>

The application saves all the above connection details along with the encrypted password to the `sdm.connect` file. All other SDM command line commands refer to the specified file. This step should be completed the first time you use the SDM command line on a machine and every time you want to change the connection details the application uses.

To run saveConnection:

1. Execute the command as follows:

   ```
   -action saveConnection -server <postgresql> -host <hostIpaddress/hostName>
   -port <portnum> -database <databaseName/SID> [-driverProps
   <propertiesFile] {-user <dbUser> -password <dbPass> | -winAuth} -
   connectFile <filenameToSaveConnection>
   ```

   The following example saves connections for a host with an IP address of 10.0.0.1 at port 5432.

   - PostgreSQL Example:

     ```
     -action saveConnection -server postgresql -host 10.0.0.1 -port 5432 -
     database SIEM -user dbauser -password xxxxxx -connectFile sdm.connect
     ```

   This saves the connection details to the `sdm.connect` file. the rest of the commands take this filename as input to connect to the designated database and to perform their actions.

13.3.5 Adding Partitions

The addPartitions action adds the required number of partitions in the following tables according to the partition configuration settings:

- PostgreSQL:
  - EVENTS
  - AUDIT_RECORD
  - CORRELATED_EVENTS
  - EVT_DEST_EVT_NAME_SMRY_1
  - EVT_DEST_SMRY_1
  - EVT_DEST_TXNMY_SMRY_1
  - EVT_PORT_SMRY_1
  - EVT_SEV_SMRY_1
  - EVT_SRC_SMRY_1

NOTE: Partitions are added in database both for events and correlated events if you select any one of these two. Partitions are added for all the summary tables if you select any one of them.

If you have configured the database to have 10 days worth of partitions, every time you run addPartitions it checks to see if you have 10 days of partitions available. If you have enough partitions for the next 10 days it does nothing. If not, it adds the required number of partitions.

This action uses the following flags:

Table 13-3 *Adding Partition Flags*

Command	Command FLags
-action	addPartitions
-connectFile	<filePath>
-tableName	<table name>
-keepDays	<days to add>

To run addPartitions:

1. Execute this command as follows:

   ```
   -action addPartitions -connectFile <filePath> -tableName <table name> -
   keepDays <days to add>
   ```

   ```
   ./sdm -action addPartitions -connectFile sdm.connect -tableName EVENTS -
   keepDays 10
   ```

13.3.6 Dropping Partitions

The dropPartition action drops all the partitions older than the flag keepDays from the following tables:

- EVENTS
- AUDIT_RECORDS
- CORRELATED_EVENTS
- EVT_DEST_EVT_NAME_SMRY_1
- EVT_DEST_SMRY_1
- EVT_DEST_TXNMY_SMRY_1
- EVT_PORT_SMRY_1
- EVT_SEV_SMRY_1
- EVT_SRC_SMRY_1

To prevent unintentional loss of data, this action does not drop any partitions that are not archived. If you want to delete unarchived partitions, use the forceDelete flag.

WARNING: If - forceDelete is used, the deleted data cannot be recovered, so use this option with caution.

This action uses the following flags:

Table 13-4 Dropping Partition Flags

Command	Command Flags
-action	dropPartitions
-keepDays	<number of days to keep>
-forceDelete (optional)	<either “true” or “false”>
	This defaults to false if not specified, meaning that only the partitions that are older than keepDays and are already archived are dropped.
	If this is set to true, all partitions older than keepDays are dropped, even if they have not been archived.
-connectFile	<filePath>
-tableName	<table name>

NOTE: Sentinel partitioned tables are organized into two groups. One is the EVENTS table group, which includes EVENTS and CORRELATED_EVENTS; the other is the summary table group, which includes all summary, or aggregate, tables. If any one of the tables in the group is specified by the –tableName parameter, the dropPartition operation is applied to all tables in that group.

To run dropPartition:

1. Execute this command as follows:

   ```
   -action dropPartitions -keepDays <numberofDaysToKeep> -tableName <table
   name> [-forceDelete <true/false>] -connectFile <filePath>
   ```

   The following examples drops all the partitions older than 30 days, making sure all the partitions are archived. All partitions that were skipped (not removed) because they have not been archived are listed when the operation completes.

   PostgreSQL Example:

   ```
   ./sdm –action dropPartitions –keepDays 30 –tableName CORRELATED_EVENTS –
   forceDelete false –connectFile sdm.connect
   ```

13.3.7 Viewing Partition Summaries

The viewPartitions action displays the partition summary of the following supported tables:

- EVENTS
- AUDIT_RECORDS
- CORRELATED_EVENTS
- EVT_DEST_EVT_NAME_SMRY_1
- EVT_DEST_SMRY_1
- EVT_DEST_TXNMY_SMRY_1
- EVT_PORT_SMRY_1
- EVT_SEV_SMRY_1
- EVT_SRC_SMRY_1

NOTE: You need to have the SDM installed in order to view the partition summary.

This command uses the following flags:

Table 13-5 *Viewing Partition Summaries Flags*

Command	Command Flags
-action	viewPartitions
-tableName	<table name>
-connectFile	<filePath>

To View Partition Summaries:

1. Execute this command as follows:

   ```
   -action viewPartitions -tableName <table name> -connectFile <filePath>
   ```

   The following example, displays the list of partitions of the EVENTS table and status of each partition.

   ```
   ./sdm –action viewPartitions –tableName EVENTS –connectFile sdm.connect
   ```

13.3.8 Archiving Data

Run the archiveData action after you set your archive configuration (configured in the *Partition Configuration* tab in the SDM GUI). This action archives the data from the given table name according to the archive configuration. It archives data from:

- EVENTS
- AUDIT_RECORDS
- CORRELATED_EVENTS
- EVT_DEST_EVT_NAME_SMRY_1
- EVT_DEST_SMRY_1
- EVT_DEST_TXNMY_SMRY_1
- EVT_PORT_SMRY_1
- EVT_SEV_SMRY_1
- EVT_SRC_SMRY_1

NOTE: Sentinel partitioned tables are organized into two groups. One is the EVENTS table group, which includes EVENTS and CORRELATED_EVENTS; the other is the summary table group, which includes all summary, or aggregate, tables. If any one of the table in the group is specified by the –tableName parameter, the archiveData operation is applied to all tables in that table group.

This command uses the following flags:

Table 13-6 Archiving Data Flags

Command	Command Flags
-action	archiveData
-connectFile	<filePath>
-tableName	<table name>
-keepDays	<numberOfDaysToKeep>

To run archiveData:

1. Execute this command as follows:

   ```
   -action archiveData -connectFile <filePath> -tableName <table name> -
   keepDays <numberOfDaysToKeep>
   ```

   The following examples archive events and correlated events from the EVENTS and CORRELATED_EVENTS tables according to the value set during archive configuration.

   ```
   ./sdm -action archiveData -connectFile sdm.connect –tableName EVENTS –
   keepDays 30
   ```

13.3.9 Importing Data

The importData action imports data between the given dates into the Sentinel database so it can be used for historical reporting or other purposes. The data is imported into the following tables:

- EVENTS
- AUDIT_RECORDS
- CORRELATED_EVENTS
- EVT_DEST_EVT_NAME_SMRY_1
- EVT_DEST_SMRY_1
- EVT_DEST_TXNMY_SMRY_1
- EVT_PORT_SMRY_1
- EVT_SEV_SMRY_1
- EVT_SRC_SMRY_1

NOTE: The tables are imported in Oracle with the same name they are archived with.

If the data has already been imported or there is no archived data found between the specified dates, the command returns a notification.

The application imports data from each file into a table and builds the historical view on all the historical tables. The report view joins on the original table and historical view. All Sentinel reports use the report view, so they see any imported data.

This command uses the following flags:

Table 13-7 Importing Data Flags

Command	Command Flags
-action	importData
-tableName	<table name>
-startDate	<mm/dd/yyyy hh24:mi:ss>
-endDate	<mm/dd/yyyy hh24:mi:ss>
-connectFile	<filePath>

hh24 is hours represented in 24-hour format. For example, 1:15:00 p.m. is 13:15:00 and 3:00:00 a.m. is 03:00:00.

NOTE: The files to be imported must exist in the directory with their originalfile names.

To run importData:

1. Place all the files you want to import in a specific directory (that is, dirPath - <directory to import files from>) and execute the following command

   ```
   -action importData -startDate <mm/dd/yyyy hh24:mi:ss> -endDate <mm/dd/yyyy
   hh24:mi:ss> -tableName <table name> -connectFile <filePath>
   ```

   The following example imports the archived files from the *tmp*directory containing the data between dates 09/25/2007 00:00:00 (Sep 25 midnight) and 09/26/2007 00:00:00 (Sep 26 midnight).

   ```
   ./sdm –action importData –startDate 09/25/2007 00:00:00 –endDate 09/26/
   2007 00:00:00 -tableName Events –connectFile sdm.connect
   ```

13.3.10 Deleting Imported Data

The droImported action deletes the imported data between the given dates from the following supported tables:

- EVENTS
- AUDIT_RECORDS
- CORRELATED_EVENTS
- EVT_DEST_EVT_NAME_SMRY_1
- EVT_DEST_SMRY_1
- EVT_DEST_TXNMY_SMRY_1
- EVT_PORT_SMRY_1
- EVT_SEV_SMRY_1
- EVT_SRC_SMRY_1

NOTE: The tables are imported in Oracle with the same name they are archived with.

If there is no data imported between two specified dates, the command returns a notification.

This command uses the following flags:

Table 13-8 *Deleting Imported Data Flags*

-action	dropImported
-startDate	<mm/dd/yyyy hh24:mi:ss>
-endDate	<mm/dd/yyyy hh24:mi:ss>
-tableName	<table name>
-connectFile	<filePath>

NOTE: hh24 is hours represented in 24-hour format. For example, 1:15:00 p.m. is 13:15:00 and 3:00:00 a.m. is 03:00:00.

To run dropImported:

1. Execute this command as follows:

   ```
   -action dropImported -startDate <mm/dd/yyyy hh24:mi:ss> -endDate <mm/dd/
   yyyy hh24:mi:ss> -tableName <table name> -connectFile <filePath>
   ```

   The following example deletes the imported data between the given dates from the tables.

   ```
   ./sdm –action dropImported –startDate 09/25/2007 00:00:00 –endDate 09/26/
   2007 00:00:00 -tableName Events –connectFile sdm.connect
   ```

13.3.11 Viewing Sentinel Database Space Usage

In tablespace management, the command line option allows you to view Sentinel database space usage

The dbstats action displays the Sentinel database usage for all Sentinel tablespaces in Oracle and Sentinel file groups in MS SQL.

This command uses the following flags:

Table 13-9 *Viewing Sentinel Database Space Usage Flags*

Command	Command Flags
-action	dbstats
-connectFile	<filePath>

To view Sentinel Database Space Usage (Command Line):

1. Execute the following command:

   ```
   -action dbStats -connectFile <filePath>
   ```

The following example displays the tablespaces of Sentinel database with their total space, used space and free space available.

```
./sdm –action dbStats –connectFile sdm.connect
```

Utilities

14

This section helps you to understand the utilities provided by Sentinel.

- Section 14.1, “Introduction to Sentinel Utilities,” on page 301
- Section 14.2, “Starting and Stopping a Sentinel Server,” on page 301
- Section 14.3, “Sentinel Scripts,” on page 302
- Section 14.4, “Version Information,” on page 305
- Section 14.5, “Database Cleanup,” on page 306
- Section 14.6, “Connecting to PostgreSQL Database Through Command Line,” on page 308
- Section 14.7, “Backup and Restore Utility,” on page 309
- Section 14.8, “Updating Your License Key,” on page 311

14.1 Introduction to Sentinel Utilities

You can use these utilities for the following purposes:

- Starting or stopping certain Sentinel services.
- Modifying Sentinel configuration.
- Determining the version of a Sentinel library.
- Troubleshooting.
- Configuring Sentinel e-mail.

14.2 Starting and Stopping a Sentinel Server

- Section 14.2.1, “Starting a Sentinel Server,” on page 302
- Section 14.2.2, “Stopping a Sentinel Server,” on page 302

A Sentinel server is made up of the following components:

- Communication Server
- Correlation Engine
- DAS
- Collector Manager
- Reporting Engine
- Advisor
- Web Server

When a Sentinel server is started or stopped, all components installed in that Sentinel server are also started or stopped. To start or stop a particular component on a Sentinel server, use the *Servers View* under the Admin tab in Sentinel Control Center.

You need to start or stop a Sentinel server because of the following routine maintenance:

- Upgrades
- Patches
- Hotfixes

- Section 14.2.1, “Starting a Sentinel Server,” on page 302
- Section 14.2.2, “Stopping a Sentinel Server,” on page 302

14.2.1 Starting a Sentinel Server

1. Log in to the machine where the Sentinel server you want to start as the Sentinel Administrator operating system user.
2. Go to the `<install_directory>/bin` directory.
3. Run the following command:

   ```
   ./sentinel.sh start
   ```

14.2.2 Stopping a Sentinel Server

1. Log in to the machine where the Sentinel server you want to stop is installed. Use the Sentinel Administrator operating system user credentials (by default admin).
2. Go to the `<install_directory>/bin` directory.
3. Run the following command:

   ```
   ./sentinel.sh stop
   ```

14.3 Sentinel Scripts

Depending upon which components are installed, `<install_directory>/bin` might contain some or all of the scripts below. The operational scripts are appropriate for use during normal operations of Sentinel. The troubleshooting scripts should only be used when troubleshooting an issue.

For most scripts that require arguments, running the scripts without arguments provides details about the arguments and usage of the script.

- Section 14.3.1, “Operational Scripts,” on page 302
- Section 14.3.2, “Troubleshooting Scripts,” on page 304

14.3.1 Operational Scripts

The scripts below can be used during the normal operation of Sentinel.

Table 14-1 *Operational Scripts*

Script File	Description
`adv_change_passwd.sh`	Resets the encrypted Advisor password stored in the Advisor configuration files.

Script File	Description
`advisor.sh`	Starts the Internet download and processing of Advisor feed data. This script is scheduled to run automatically when Advisor is installed.
`BackupIncidentData.sh`	Used to back up incident-related data before running the delete incident utilities. For more information, contact Novell Support (http://support.novell.com/phone.html?sourceidint=suplnav4_phonesup).
`Clean_Database.sh`	Used to delete incidents or Identity information from the database.
`control_center.sh`	Launches the Sentinel Control Center graphical user interface.
`dbconfig`	Configures the database connection settings stored in the DAS container XML files. For more information, see "Sentinel 6.1 Rapid Deployment Data Access Service" in the *Sentinel 6.1 Rapid Deployment Reference Guide*.
`runadvisor_client.sh`	Launches the client to download Advisor data.
`sdm`	Launches the Sentinel Data Manager application. For more information, see Chapter 13, "Sentinel Data Manager," on page 283.
`sentinel.sh`	Starts or stops the Sentinel server. For more information, see Section 14.2, "Starting and Stopping a Sentinel Server," on page 301.
`setadvenv.sh`	Used by the Advisor scripts to set some local environment variables.
`setenv.sh`	Used by many of the Sentinel scripts to set some local environment variables.
`softwarekey.sh`	Resets the Sentinel license key. For more information, see Section 14.8, "Updating Your License Key," on page 311.
`solution_designer.sh`	Starts the Solution Designer application.
`uninstallcron.sh`	Removes the Advisor feed download and processing scheduled jobs. This script is run automatically by the uninstaller.
`versionreader.sh`	Displays the version information stored in a Sentinel jar file. For more information, see Section 14.4.2, "Sentinel .jar Version Information," on page 305.
`AnalyzePartitions.sh`	Analyzes only the event partitions
javac	
Agentengine	
`Install cron.sh`	
`Physicalmemory.sh`	
`setmemory.sh`	
wrapper	

14.3.2 Troubleshooting Scripts

The scripts in this section are useful when you are troubleshooting an issue you are experiencing. They provide finer -grained control of certain components in Sentinel, allowing you to drill down to the root cause of the issue.

NOTE: These scripts should not be used during normal operation of Sentinel. They are intended for troubleshooting purposes.

Table 14-2 *Troubleshooting Scripts*

Script File	Description
start_broker.sh	Starts the message bus component of the communication server. This script is useful if you are having problems starting the message bus. This script is automatically run by the installer. For more information, see "Starting the Communication Server in Console Mode" on page 304.
stop_broker.sh	Stops the message bus component of the communication server. For more information, see "Stopping the Communication Server in Console Mode" on page 304.

- "Starting the Communication Server in Console Mode" on page 304
- "Stopping the Communication Server in Console Mode" on page 304

Starting the Communication Server in Console Mode

These scripts start the communication server on the command line in console mode. The scripts are useful for debugging the communication server without requiring you to run the rest of Sentinel server.

NOTE: During normal operations, you should not use these scripts. Instead, follow the procedures in Section 14.2.1, "Starting a Sentinel Server," on page 302.

1. Log in as the Sentinel Administrator operating system user.
2. Go to:

```
<install_directory>/bin
```

3. Enter:

```
./start_broker.sh
```

Stopping the Communication Server in Console Mode

These scripts stop the communication server on the command line in console mode. The scripts are useful for troubleshooting the communication server without forcing you to stop the rest of Sentinel server.

NOTE: During normal operations, you should not use these scripts. Instead, follow the procedures in Section 14.2.2, "Stopping a Sentinel Server," on page 302.

1. Log in as the Sentinel Administrator operating system user.
2. Go to:

```
<install_directory>/bin
```

3. Enter:

```
./stop_broker.sh
```

14.4 Version Information

The following processes provide information about versions:

- Section 14.4.1, "Executable Version Information," on page 305
- Section 14.4.2, "Sentinel .jar Version Information," on page 305

14.4.1 Executable Version Information

Sentinel has a command line option to display the version information of the `agentengine` executable:

1. Go to:

   `<install_directory>/bin`
2. At the command line, enter:

```
./<process> -version
```

For example:

```
./agentengine –version
```

14.4.2 Sentinel .jar Version Information

The following procedure describes how to gather the version information of Sentinel `.jar` files:

1. Log in to the machine where Sentinel is installed by using the Sentinel Administrator operating system user credentials (the default is admin).
2. Go to:

   `<install_directory>/bin`
3. At the command line, enter either of the following:

```
./versionreader.sh <path/jar file name>
./sentinel.sh version
```

14.5 Database Cleanup

The `Clean_Database.sh` scripts are used to purge incidents, identities, assets, advisor data, and vulnerabilities from the Sentinel database. For example, an improperly configured Correlation rule might create hundreds of unwanted incidents in the database. It's also possible that the identity information might encounter an error. For example, if someone attempts to delete the `IdentityAccountMap.csv` file.

WARNING: Because these scripts are designed to delete information from your database, they should be used very carefully and only after understanding the implications.

- Section 14.5.1, "Components," on page 306
- Section 14.5.2, "Prerequisites," on page 307
- Section 14.5.3, "Running Clean_Database.sh," on page 307

14.5.1 Components

`<install_directory>/bin/Clean_Database.sh`	Main database cleanup script. This calls the other scripts.
`<install_directory>/bin/BackupIncidentData.sh`	Script used to back up Incident data.
`delete_incidents_by_query`	Stored procedure used to delete incidents specified by an SQL query.
`delete_incidents_by_rule`	Stored procedure used to delete incidents created by a specified correlation rule.
`delete_incidents_by_id`	Stored procedure used to delete an incident with a specified ID.
`identity_cleanup`	Stored procedure used to delete identity-related data.
`<install_directory>/bin/BackupAdvisor.sh`	Script used to back up Advisor data.
delete_advisor_all	Stored procedure used to delete Advisor data.
`<install_directory>/bin/BackupAsset.sh`	Script used to back up Asset data.
delete_assets_all	Stored procedure used to delete all Asset data.
delete_assets_by_id	Stored procedure used to delete Asset data based on Asset ID.
`<install_directory>/bin/BackupVuln.sh`	Script used to back up Vulnerability data.
delete_vuln_all	Stored procedure used to delete Vulnerability data.

14.5.2 Prerequisites

There are several prerequisites for running the `Clean_Database` script.

- The user running the script must be a `novell` user, and each script must have the permission set so that only the `novell` user is allowed to execute the cleanup script.
- The user running the PostgreSQL script must have permission to access/execute all of the database tools and utilities. Run the script as a `dbauser`.
- (Identity Cleanup only) The database must be in a healthy state and in good running condition because the Identity cleanup stored procedure disables and enables foreign key constraints.
- (Identity Cleanup only) All Identity/Account loaders and Collectors, such as the Identity Vault Collector, should be stopped.
- (Identity Cleanup only) Reports that are running queries against the Identity tables should be stopped.

The Identity cleanup DDL operations are atomic, so if one DDL statement execution fails, the script exits with errors written to the specified log file.

WARNING: If identity information is cleaned out of the database and then reloaded, the new identity information is not synchronized with any past events that had identity information injected. Therefore, attempts to perform identity lookups on past events (received before the cleanup) or run reports on past events with identity information is not successful.

Use this option with extreme caution.

14.5.3 Running Clean_Database.sh

1. Open a console, go to `<install_directory>/bin` and enter `Clean_Database.sh` to start the script.

   NOTE: You can cancel the execution of the cleanup script at any time by entering `q` at any prompt.

   ---

2. At the prompt, indicate which objects you want to remove from the database:

   ```
   Which objects would you like to cleanup?
   (1) Incidents
   (2) Identities
   (3) Assets
   (4) Advisor
   (5) Vulnerabilities
   (6) Incidents and Identities
   (7) All
   ```

3. At the prompts, enter the following information to connect to the PostgreSQL database:

   ```
   Database server hostname (Press ENTER for default localhost)=>
   Database name (Press ENTER for default SIEM) =>
   Database username (press ENTER for default dbauser)
   ```

   The database connection is verified before proceeding to the next step.

4 (Conditional) If you selected to clean incidents:

The following prompt displays:

```
Would you like to backup Incidents first? (y or n) =>
```

4a If you select `y` to back up the incidents, enter the destination directory (a full path or a path relative to the location of the cleanup script) for the backup files.

The user running the script must have permission to write to this directory.

4b Select an incident cleanup option:

- **Delete Incidents By Query:** You are prompted to enter a custom SELECT query. For example:

  ```
  select inc_id from incidents where inc_id=500
  ```

  The SELECT statement cannot include quotation marks.

- **Delete Incidents By Rule:** You are prompted to enter the name of the Correlation rules that created the incidents. For example:

  ```
  My Test Rule
  ```

- **Delete Incidents By Id:** You are prompted to enter the ID of a specific incident. For example:

  ```
  101
  (q) Quit without action
  ```

4c At the Incident Cleanup Confirmation prompt, enter `start` to start the incident cleanup or enter `abort` to quit without performing any cleanup.

The results of the incident cleanup are written to the specified log file.

You should review the log file for any errors before continuing.

5 Conditional) If you selected to clean identity:

5a At the Identity Cleanup Confirmation prompt, enter `start` to start the Identity cleanup or enter `abort` to quit without performing the identity cleanup.

The results of the Identity Cleanup are written to the specified log file.

You should review the log file for any errors before continuing.

5b In addition to deleting the Identity information from the database tables, the script attempts to delete the Identity Account Map file (`identityAccountMap.csv`).

If you have a distributed Sentinel install, you might need to manually connect to the main Sentinel server to delete the `identityAccountMap.csv` file.

5c At the prompt, enter the novell user's password.

14.6 Connecting to PostgreSQL Database Through Command Line

Sentinel Rapid Deployment has a command line option to connect to the PostgreSQL database:

1 Connect to the Sentinel Rapid Deployment server as the user who owns the Sentinel Rapid Deployment installation files.

2 Change to the Rapid Deployment bin directory.

```
cd <RD_install_home>/bin
```

3 Run the `setenv.sh` script to set the environment variables.

```
. setenv.sh
```

4 Change to the PostgreSQL bin directory.

```
cd <RD_install_home>/3rdparty/postgresq/bin
```

5 Run psql to connect to the PostgreSQL database.

```
./psql SIEM -U dbauser
```

14.7 Backup and Restore Utility

The backup and restore utility performs a back up of the system data and also restores the data at any given point in time without a considerable amount of effort.This utility backs up and restores data only for the Sentinel Rapid Deployment server and can not be used for Collector Manager systems.

You can back up the following data:

- **Configuration data:** Data stored in the `config` directory and other directories, and in the Sentinel database. This data includes configuration files, property files, and keystore files. The Sentinel database contains various configuration information related to users, plug-ins, Collectors, Connectors, and filters.
- **Event data:** Event data stored in the database. The event data includes the aggregated events files and any temporary event files that were created when the events failed to be inserted into the database.
- **Runtime data:** Dynamic data stored on the file system. This data includes the Lucene search indexes used in queries by the Sentinel Rapid Deployment Web UI.
- **Advisor data:** Advisor data stored in the Sentinel database.

NOTE: You need not back up the Advisor data, because the default Advisor data is loaded into the Sentinel database when you install Sentinel Rapid Deployment. Also, if you purchase the Advisor license, a backup is not required because you can download the updated Advisor data from the Novell download Web site (https://secure-www.novell.com/sentinel/download/advisor/) (https://secure-www.novell.com/sentinel/download/advisor/).

The backup and restore script is controlled by various command line parameters that are described in Table 14-3.

14.7.1 Parameters for the Backup and Restore Utility Script

The following table lists the various command line parameters that you can use with the `backup_util.sh` script:

Table 14-3 *Backup and Restore Script Parameters*

Parameters	Description
-m backup	Takes a backup of the specified data.

Parameters	Description
-m restore	Restores the data from the specified backup file. The restore mode of the script is interactive and allows you to specify the data to be restored from the backup file. The restore parameter can be used in the following scenarios: ◆ **System Failure:** In the event of a system failure, you must first reinstall Sentinel Rapid Deployment and then use the `backup_util.sh` script with the restore parameter to restore the most recent data that you had backed up. ◆ **Data Loss:** In the event of data loss, use the `backup_util.sh` script with the restore parameter to restore the most recent data that you had backed up. You must restart the server after you restore any data because the script might make several modifications to the database.
-m info	Displays the information for the specified backup file.
-a	Takes a backup of the Advisor data. By default, this parameter is disabled in the script, because the data is available from other sources. To enable this parameter, edit the script manually and change FLAG_A_DISABLED = "disabled" to FLAG_A_DISABLED = "enabled".
-e	Takes a backup of all the online event data. If the backup is performed on the Sentinel server, the current online partition is not backed up unless you shut down the server.
-c	Takes a backup of the configuration data.
-dN	Takes a backup of the event data for the specified number of days. By default, just specifying the -e option backs up all the online event data. Based on the current *Partition Configuration* settings specified in the SDM, event data partitions can be kept online for up to last 90 days. Backing up all 90 days of event data with every backup might not be essential. **NOTE:** If you specify only -d, no event data will be backed up.You must specify this parameter along with the -e parameter. For example: `backup_util.sh -m backup -e -d5 -f <install_directory>/data/<events_5days_backup.tar.gz>`
-s	Shuts down the Sentinel server. You must use this command if you want to backup the current online partitions in the database and/or to backup the dynamic runtime data because the server must be shutdown before taking a backup of these data. After the backup is complete, the server restarts automatically.
-f	Enables you to specify the location and name of the backup file. If this option is not used, then the backup file gets a random name based on the current date.
-l	Includes the log files in the backup.

14.7.2 Using the Backup and Restore Utility Script

1. Open a console, and navigate to the `<install_directory>/bin` directory as the `novell` user.

2. Enter `backup_util.sh`, along with the necessary parameters for the data that you want to back up or restore.

For more information on the different parameters, see Table 14-3. The following table gives examples of how to specify the parameters:

Syntax	Action
`backup_util.sh -m backup -a -c -e -l -s -f <install_directory>/data/<full_backup.tar.gz>`	Shuts down the server and takes a backup of the complete system data. You should shut down the Sentinel server only when you back up the event data of the current online partition or when you back up the dynamic runtime data.
`backup_util.sh -m backup -c -f <install_directory>/data/<config_backup.tar.gz>`	Takes a backup of the current configuration data.
`backup_util.sh -m backup -e -f <install_directory>/data/<events_backup.tar.gz>`	Takes a backup of the event data.
`backup_util.sh -m backup -e -d5 -f <install_directory>/data/<events_5days_backup.tar.gz>`	Takes a backup of the event data for the last five days.
`backup_util.sh -m info -f <install_directory>/data/<config_backup.tar.gz>`	Displays the backup information for the specified backup file.
`backup_util.sh -m restore -f <config_backup.tar.gz>`	Restores the data from the specified filename.

3 (Conditional) If you have restored any data, restart the server because the script might make several modifications to the database.

14.8 Updating Your License Key

If your Sentinel license key has expired and Novell has issued you a new one, run the software key program to update your license key.

1 Log into the Sentinel Server machine as the Sentinel Administrator operating system user.

2 Go to `<install_directory>/bin`

3 Enter the following command:

```
./softwarekey.sh
```

4 Specify the number 1 to set your primary key, then press Enter.

Quick Start

15

This section assumes that your security administrator has built the necessary filters and configured Collectors for your system.

- Section 15.1, “Security Analysts,” on page 313
- Section 15.2, “Creating Incidents,” on page 317
- Section 15.3, “iTRAC,” on page 318
- Section 15.4, “Correlation,” on page 328

15.1 Security Analysts

- Section 15.1.1, “Active Views Tab,” on page 313
- Section 15.1.2, “Exploit Detection,” on page 314
- Section 15.1.3, “Asset Data,” on page 315
- Section 15.1.4, “Event Query,” on page 316

15.1.1 Active Views Tab

In the *Active Views* tab, you can monitor events as they happen, performing queries on these events. You can monitor them in a table form or through a 3-D graphical representation.

1. Select the *Active View* tab.
2. Click *Active Views* > *Create an Active View*.

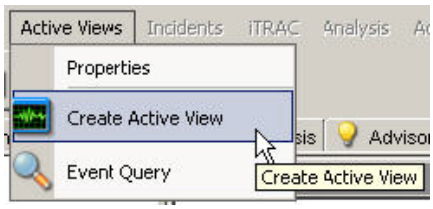

3. Select a filter from the *Filter* drop-down menu, then click *Select*.

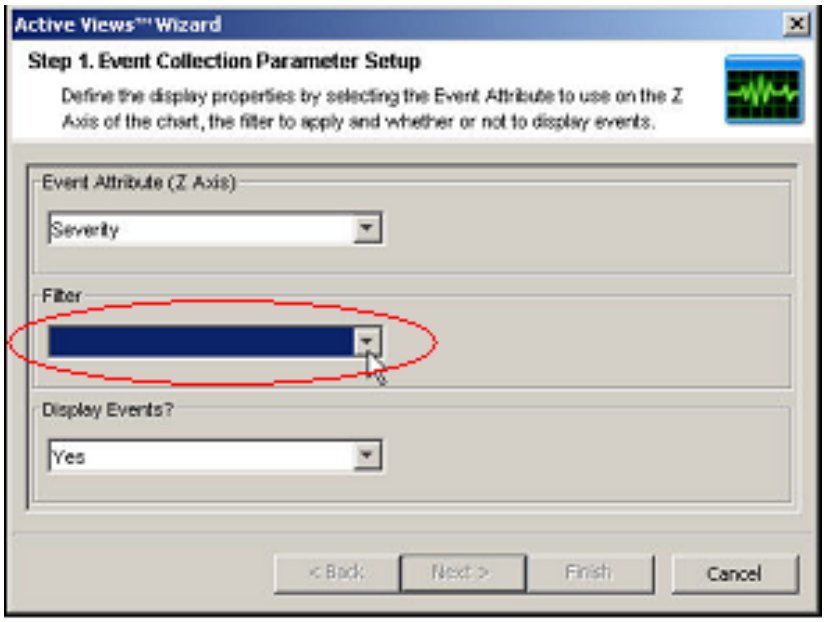

4 Click *Finish*. If you have an active network, you might see something similar to:

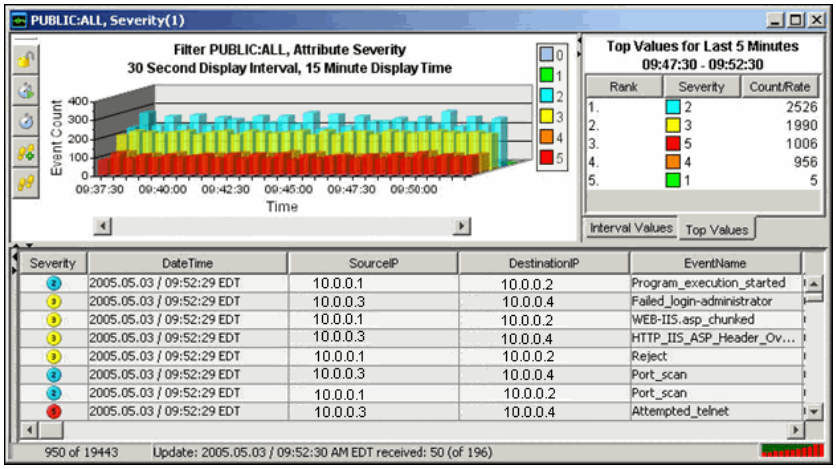

NOTE: To display a 3-D graph without real-time events, click the Display Events down-arrow and select No.

15.1.2 Exploit Detection

To view any events indicating a possible exploitation, you must have the following:

- Advisor Feed
- Intrusion detection
- Vulnerability scanning

Figure 15-1 *Severity, Vulnerability, and AttackId Columns*

Severity	Vulnerability	AttackId
2	0	
3	0	

Within an event, the values in the *Vulnerability* field convey the following:

- When the *Vulnerability* field equals 1, the asset or destination device is possibly exploited.
- When the *Vulnerability* field equals 0, the asset or destination device is not being exploited.
- When the *Vulnerability* field is blank, the exploit detection feature of Sentinel is not enabled.

To view events that indicate a possible exploitation, create an Active View with a filter where Vulnerability equals 1. For example, if you have Nmap and have run the Nmap Collector, you can view asset information on the exploited asset or any asset.

For more information on how exploit detection works and which intrusion detection systems and vulnerability scanners are supported, see Chapter 2, "Sentinel Control Center," on page 41.

15.1.3 Asset Data

To view Asset information for any event, right-click an event or events, then select *Analysis* > *Asset Data*.

A window similar to the one below displays:

Figure 15-2 *Asset Report*

Asset Report

Hardware	**MAC Address**	04:23:A3:44:65:87		
	Name		**Value**	UNKNOWN
	Type	DESKTOP	**Criticality**	UNKNOWN
	Vendor	UNKNOWN	**Sensitivity**	UNKNOWN
	Product		**Environment**	UNKNOWN
	Version		**Location**	UNKNOWN

Network	**IP**	**Hostname**
	192.168.0.10	
		devbox10

Software	**Name**	**Type**	**Vendor**	**Product**	**Version**

Contacts	**Order**	**Name**	**Role**	**Email**	**Phone Number**
		OwnerFirstName10 OwnerLastName10	ASSET_OWNER	OwnerEmail10	OwnerPhoneNumber10
		MaintainerFirstName10 MaintainerLastName10	ASSET_MAINTAINER	MaintainerEmail10	MaintainerPhoneNumber10
		BusinessUnit10	BUSINESS_UNIT		
		LineOfBusiness10	LINE_OF_BUSINESS		
		Division10	DIVISION		
		Department10	DEPARTMENT		

Location	**Room**	709
	Rack	10
	Address	HQ 1921 Gallows Rd Suite 700 Vienna VA 22182 USA

Hardware	**MAC Address**	04:23:A3:44:65:78		
	Name		**Value**	AssetValue
	Type	DESKTOP	**Criticality**	Criticality
	Vendor	Vendor	**Sensitivity**	Sensitivity
	Product	ProductName	**Environment**	EnvironmentIdentity
	Version	ProductVersion	**Location**	NetworkIdentity

Network	**IP**	**Hostname**
	192.168.0.1	

15.1.4 Event Query

You can use an event query to find out if your system has been attacked. For example, during monitoring, you see numerous Telnet attempts from source IP 10.0.0.1 Telnet attempts could be an attack. Telnet potentially allows an attacker to remotely connect to a remote computer as if they were locally connected. This can lead to unauthorized configuration changes, installation of programs, viruses, and so on.

You can use an event query to determine how often this possible attacker has attempted a Telnet attack by setting up a filter to query for this particular attacker. For example, you know the following:

- Source IP: 10.0.0.1
- Destination IP: 10.0.0.2
- Severity: 5
- Event Name: Attempted_telnet
- Sensor Type: H (Host Intrusion Detection)

To perform an event query:

1. In the Sentinel Control Center, click *Event Query* (*Magnifying Glass* icon) and click the *Filter* drop-down menu.

   A window with a list of filters displays.

2. Click *Add*; specify a filter name of Telnet SIP 10.0.0.1. In the field below the filter, specify:
   - SourceIP = 10.0.0.3
   - EventName = Attempted_telnet
   - Severity = 5
   - SensorType = H
   - DestinationIP = 10.0.0.4
   - Match if, select All conditions are met (and)
3. Click *Save*. Select your filter and click *Select*.
4. Provide your time period of interest, then click *Search* (*Magnifying Glass* icon).

   The result of your query displays. If your event query makes a match, you see a result similar to the following illustration.

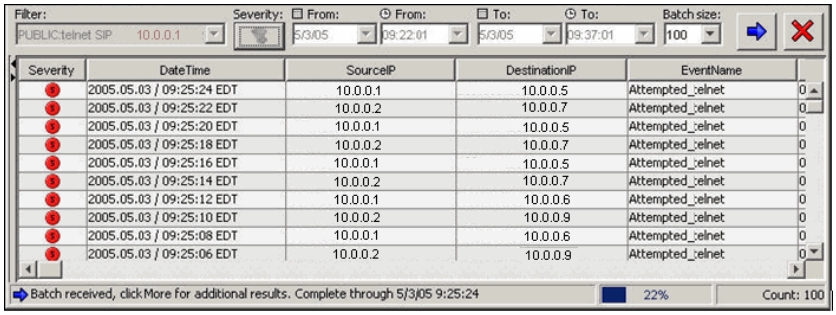

If you want to see how often in general this user is attempting a Telnet, remove DestinationIP, SensorType and, Severity from your filter or create a new filter. The results show all the destination IPs this user is attempting to Telnet to.

If any of your events are correlated events, you can right-click *View Trigger Events* to find what events triggered that correlated event.

NOTE: Correlated events have the SensorType column populated with a C.

15.2 Creating Incidents

Creating an incident is useful in grouping a set of events together as a whole representing something of interest (a group of similar events or set of different events that indicate a pattern of interest such as an attack).

If events are not initially displayed in a newly created incident, it is probably because of a lag in the time between display in the Real Time Events window and insertion into the database. If this occurs, it might take a few minutes for the original events to finally be inserted into the database and display in the incident.

NOTE: It is possible to create an incident that does not contain any events. Events can always be added to incidents.

---

1 In a Real Time Event Table of the Visual Navigator or a Snapshot Real Time Event Table, right-click an event or a group of events and select *Create Incident*.

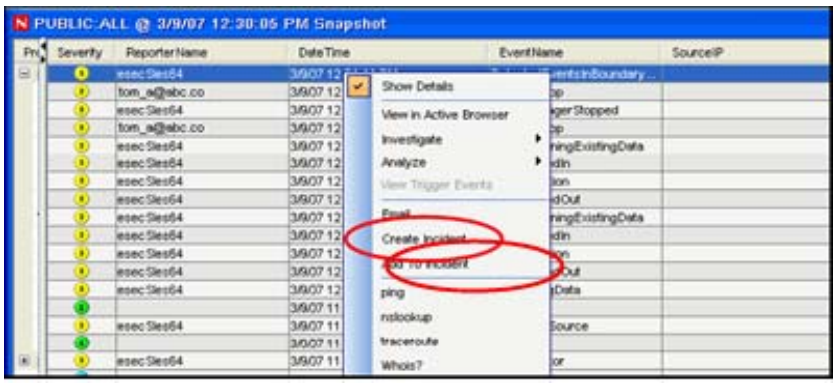

In the Incident Window are the following tabs:

- **Events:** Shows which events make up the incident.
- **Assets:** Show affected assets.
- **Vulnerability:** Show related asset vulnerabilities.
- **Advisor:** Asset attack and alert information.
- **iTRAC:** Use this tab to assign an iTRAC process.
- **History:** Incident history.
- **Attachments:** Use this tab to attach any document or text file with pertinent information to this incident.
- **Notes:** Specify any general notes regarding this incident.

2 In the Create Incident dialog box, provide the following information:
   - Title
   - State
   - Severity
   - Priority
   - Category
   - Responsible
   - Description
   - Resolution
3 Click Create. The incident is added to the Incidents page of the Sentinel Control Center.

To do this, you must have user permission to create incidents.

15.3 iTRAC

This section gives and idea relevant to iTRAC.

- Section 15.3.1, “Instantiating a Process,” on page 318

15.3.1 Instantiating a Process

An iTRAC process can be instantiated on the iTRAC server by using one of the following methods to associate an iTRAC process to an incident:

- Associating an iTRAC process to the incident at the time of incident creation
- Associating an iTRAC process to the incident after the incident is created
- Associating an iTRAC process to an incident as an action when deploying a correlation rule

For more information on associating a process to an incident, see Chapter 4, “Correlation Tab,” on page 83 and Chapter 5, “Incidents Tab,” on page 109.

NOTE: If you want to perform all of the iTRAC scenarios, you must go through them in the order they are presented.

- “Example Scenario: Creating a Simple Two-Tiered iTRAC Process for a Possible Network Attack” on page 318
- “Example Scenario: Running an iTRAC Process for a Possible Network Attack” on page 325

Example Scenario: Creating a Simple Two-Tiered iTRAC Process for a Possible Network Attack

This process is a series of steps that you can take if there is a possible attack on your system.

The example procedure does the following:

- Asks the user to decide if a preliminary look indicates that the network has been attacked. This leads to a decision step.

NOTE: All decision steps provide different execution paths, depending on the value of the variable defined in the previous step.

- The Collect Data step reviews the data to make a better determination if there has been an attack.
- If there has been an attack, iTRAC takes measures to prevent another attack and sends an e-mail to the supervisor indicating that proper measures have been taken. If there is no attack, iTRAC sends an e-mail to the supervisor indicating that there is not an attack.

Figure 15-3 *iTRAC Process*

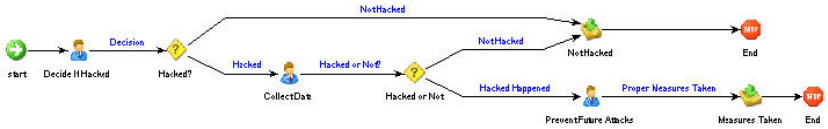

To create this iTRAC process:

1. Click the *iTRAC* tab.
2. In the navigation pane, click *iTRAC Administration* > *Template Manager*.
3. In the Template Manager window, click *Add*.

   The iTRAC Process Builder displays with a Process Details window.
4. Use the name iTRAC Tutorial. Optionally, add a description.

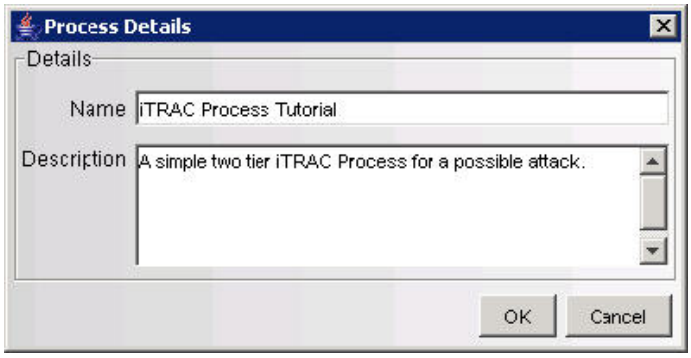

5. From the Step Palette pane, drag and drop three manual steps, two mail steps, and two Decision Steps. Rename and the attributes to the steps as follows by right-clicking and selecting Edit Step.

   5a Manual Step-0 to Decide If Hacked.

   5a1 Set the Role to Analyst.

   5a2 Click Associate, then click *Add*.

   5a3 Specify Hacked in the `Name` field.

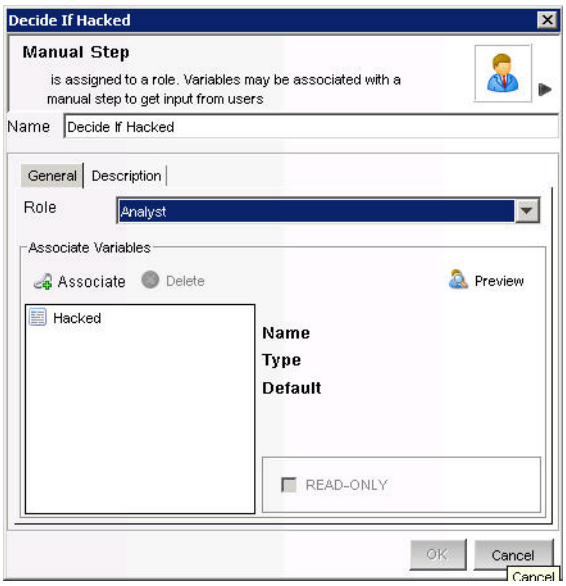

5a4 In the Process Variables window, select the *Variable Type* as *String*.

5a5 Set the *Default Value* to *yes*.

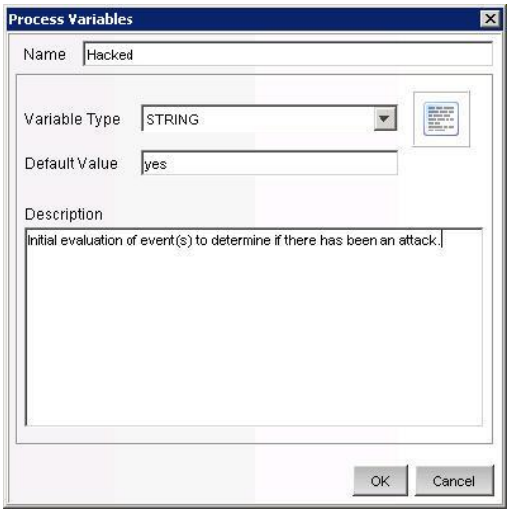

5a6 (Optional) Under the Description tab, specify Initial evaluation of events to determine if there has been an attack.

5a7 Click *OK*.

5a8 Select the newly created association, then click `OK` until the step is renamed.

5b Manual Step-1 to Collect Data:

5b1 Set the *Role* to *Analyst*.

5b2 Click *Associate*.

5b3 Select *Hacked*, then click *OK*.

5b4 (Optional) Under the Description tab, specify To further evaluate after collecting of events to determine if there has been an attack.

5b5 Click *OK* to rename the step.

5c Manual Step-2 to Prevent Future Attacks:

5c1 Set *Role* to *Analyst*.

5c2 (Optional) Under the *Description* tab, specify Take measures to stop the attack. (firewall, router or other intrusion protection method). Also, if possible, determine how the attacked was done.

5c3 Click *OK* to rename the step.

5d Mail Step-3 to Not Hacked:

5d1 In the *To* field (because this is for a tutorial), provide your e-mail address. When this step finishes, sends you an e-mail.

5d2 In the *From* field, provide a made up address such as me@nowhere.com.

5d3 In the Subject field, specify *We have not been hacked*.

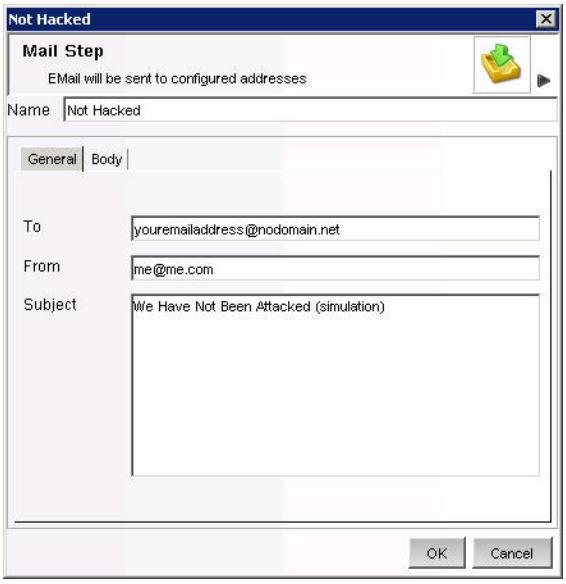

5d4 (Optional) Under the *Body* tab, specify This e-mail is generated from a tutorial (simulation) iTRAC process.

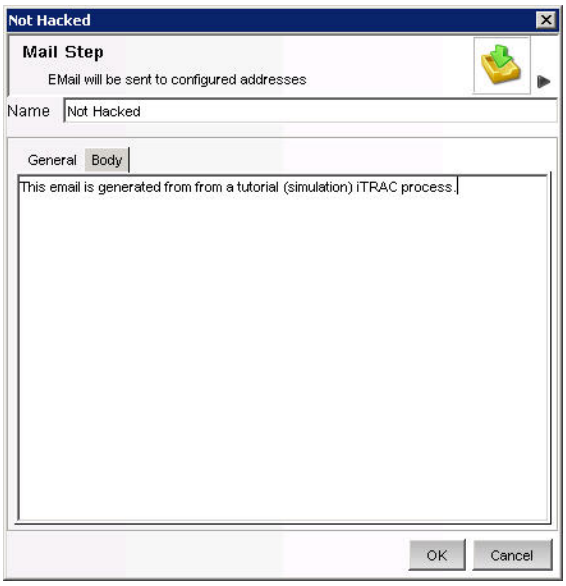

5d5 Click *OK*.

5e Mail Step-4 to Prevent Future Attacks:

5e1 In the *To* field, specify your e-mail address.

5e2 In the *From* field, specify a made up e-mail address.

5e3 In the Subject field, specify *Proper Attack Measures Taken*.

5e4 (Optional) Under the *Body* tab, specify This e-mail is generated from a tutorial (simulation) iTRAC process.

5f (Optional) Decision Step-5 to Hacked:

Under the Description tab, provide a description such as Preliminary decision if there has been an attack or not.

5g (Optional) Decision Step-6 to Hacked or Not:

Under the Description tab, provide a description such as Decision if there has been an attack or not.

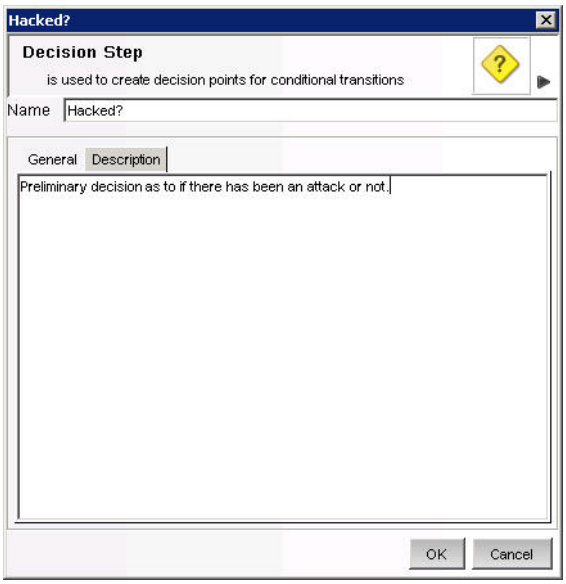

6 Right-click *Start* and select *Add Start Transition*. Select *Decide If Hacked* as the destination.

7 Right-click *Decide If Hacked* and select *Add Transition*. Specify the following:

- Name: Specify Decision.
- Type: Select Unconditional.
- Destination: Hacked.

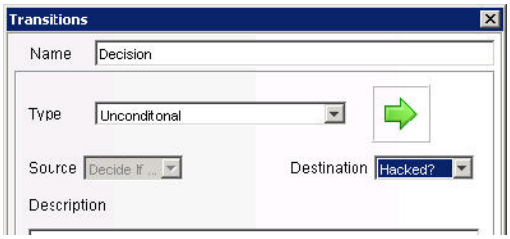

8 Click *OK*

9 Right-click *Hacked?* and select *Add Transition*. Specify the following:

- Name: Not Hacked.
- Type: Select else.
- Destination: Not Hacked.

10 Click *OK*.

11 Right-click *Not Hacked* and select *End Transition*.

12 Right-click *Hacked?* and select *Add Transition*. Specify the following:

- Name: Specify Hacked.
- Type: Select Conditional.
- Destination: Collect Data.

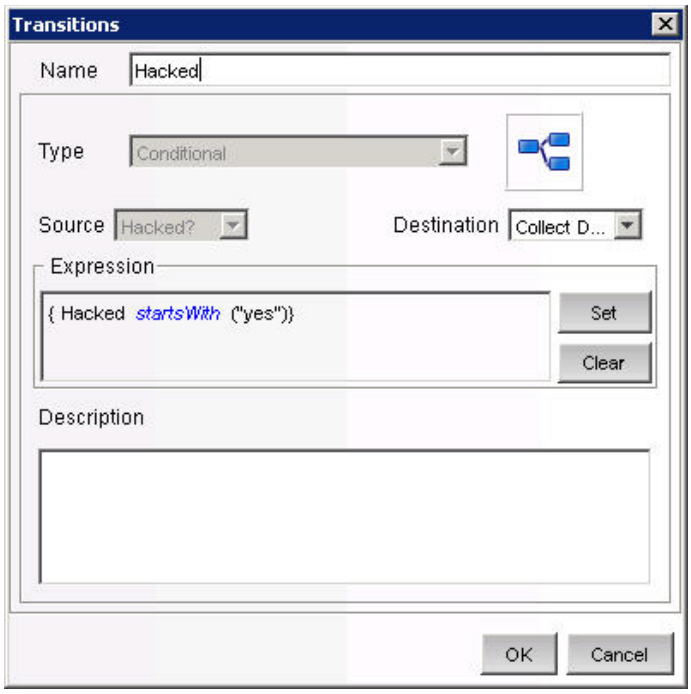

13 Click *Set* > *EXP*.

13a Select *Variables* and *Values*.

13b Select *Attribute Hacked*.

13c Select *Condition equals*.

13d Specify a value of yes.

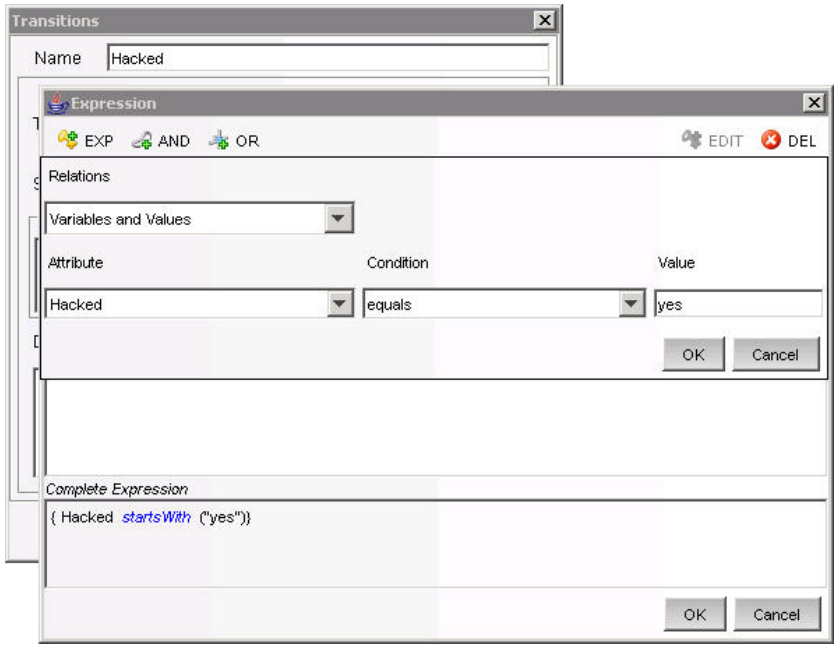

13e Click *OK* until the transition is complete.

14 Right-click *Collect Data* and select *Add Transition*. Select and specify the following:

- Name: Hacked or Not?
- Type: Unconditional
- Destination: Hacked or Not

15 Right-click *Hacked or Not* and select *Add Transition*. Specify the following:

- Name: Not Hacked.
- Type: Else.
- Destination: Not Hacked.

16 Right-click *Hacked or Not* and select *Add Transition*. Specify the following:

- Name: Hack Happened.
- Type: Conditional.
- Destination: Prevent Future Attacks.

17 Click *Set* > *EXP*.

17a Select *Variables* and *Values*.

17b Select *Attribute Hacked*.

17c Select *Condition equals*.

17d Specify Value of yes.

17e Click *OK* until the transition is complete.

18 Right-click *Prevent Future Attacks* and select *Add Transition*. Specify the following:

- Name: Proper Measures Taken.
- Type: Unconditional.
- Destination: Measures Taken.

19 Right-click *Measures Taken* and select *Add End Transition*.

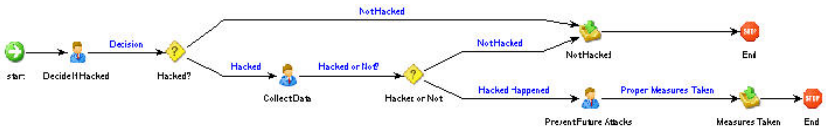

20 Click *Save*. Your new process should appear in the Template Manager.

Example Scenario: Running an iTRAC Process for a Possible Network Attack

The following example assumes the following:

- A process named iTRAC Process Tutorial has been assigned to your role (analyst)

  This is a process created in "Example Scenario: Creating a Simple Two-Tiered iTRAC Process for a Possible Network Attack" on page 318.

- All steps within the process belong to the Analyst group

NOTE: If you assign steps to other roles, you need to log out and then log in as a user assigned to that role and accept the process. For simplicity, the following example is assigned to one role.

To run this process, this process must first be assigned to an incident.

To start or terminate a process:

1. Click the *Incident* tab.
2. Click *Incidents* > *Create Incidents*.
3. Specify the following:
   - Title: iTRAC Tutorial.
   - Category: Other.
   - Responsible: assign this incident to yourself.
4. Click the *iTRAC* tab, then select *iTRAC Process Tutorial*.
5. Click Create.

   Because this is a tutorial incident and not a true incident, it can be deleted without negatively affecting your Sentinel setup.
6. From anywhere in the Sentinel GUI, click the Analyst group (yellow bar) under View Work Items.

   Your bar might already be partially green, indicating that you have accepted (acquired) an iTRAC Process.

   All of the processes assigned to the Analyst role display.

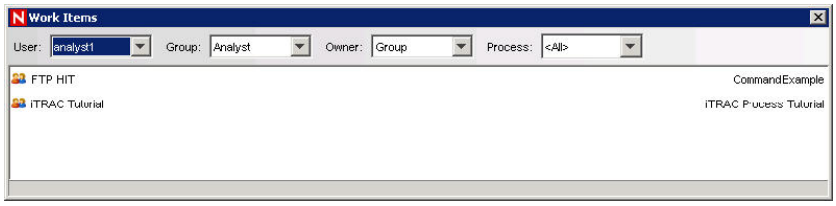

7. To accept a work item, select *iTRAC Tutorial* and click *Acquire*.

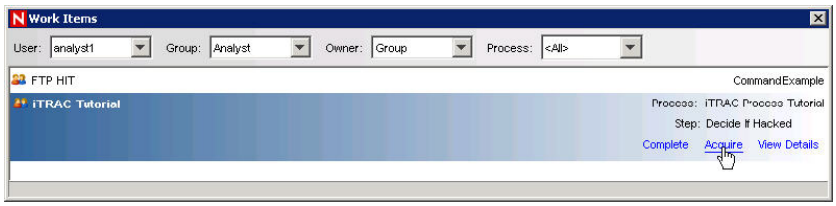

   If the View Work Item list bar was yellow as illustrated above, it changes with an addition of a green bar.

8. Click the green bar under View Work Items. In the Work Items window, click *View Details*.

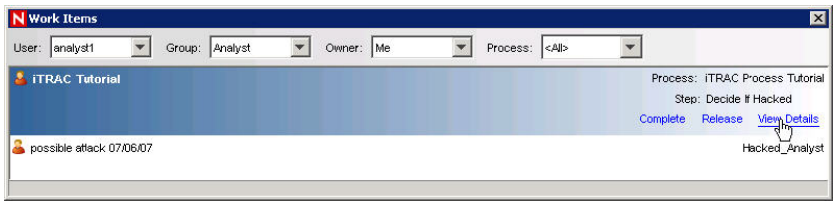

The red highlighted step indicates what step this process is currently in.

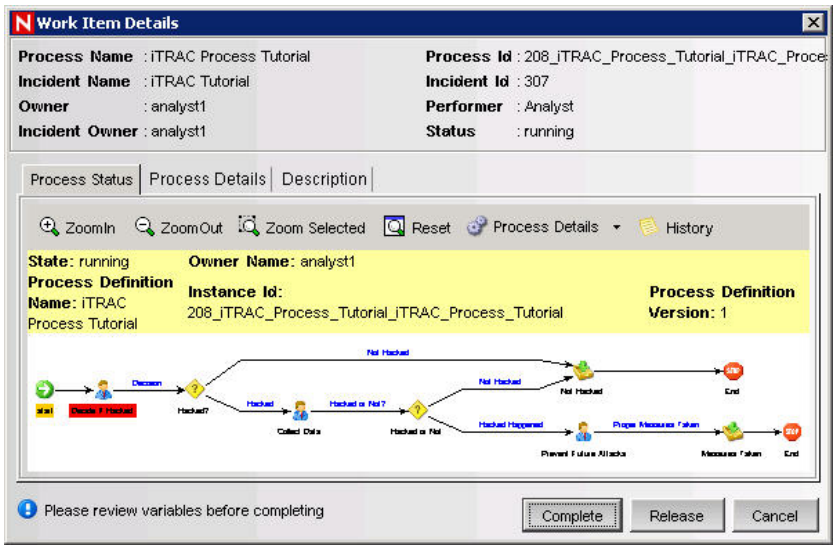

9 To start the steps within this process, click the *Process Details* tab.

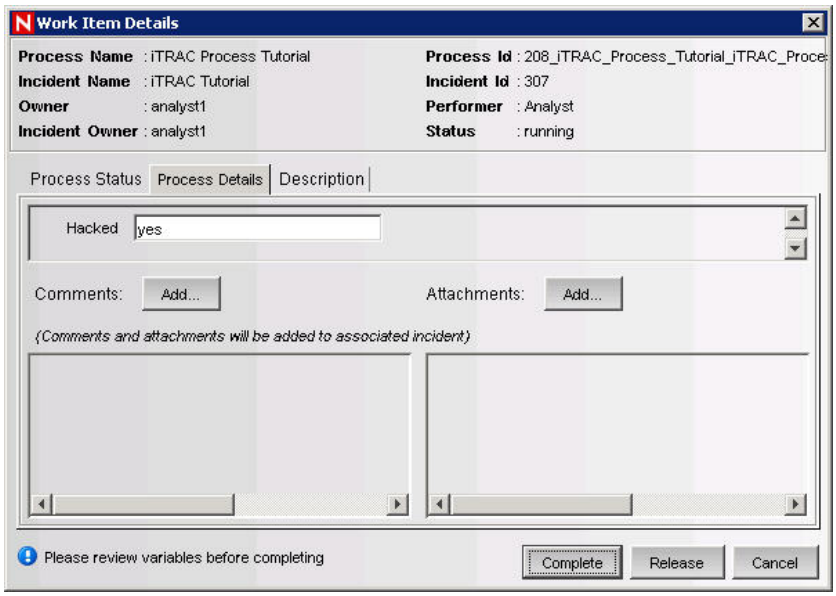

For this manual step, the variable yes is specified. Providing another value such as no or else (no attack) results in an e-mail that completes the process. For example, if initial assessment is that there is an attack and the hacked variable is equal to yes, you click Complete to complete this step.

10 In the Work Items window, select the process and click *View Details*. The Collect Data step should be highlighted in red. As before, this is a manual step.

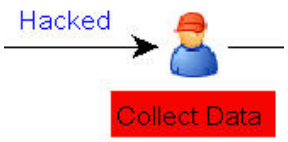

11 Click the *Process Details* tab.

12 Again, the variable page displays. In the previous step of the iTRAC Process, Collect Data is a step to further determine by analyzing the events of interest if an attack has occurred. For example, assume that an attack has occurred. Leave the default value of yes. If this were a real attack, it is beneficial to add clear notes or attachments as to the information about this attack.

13 Click *Complete*.

14 In Work Items window, select the process and click *View Details*. The Prevent Future Attacks step should be highlighted in red. As before, this is a manual step.

15 In this manual step, take measures to harden the network to prevent future attacks. When this is done, you should add notes and attachments as to the information about this attack.

16 Click *Complete*.

The next step is an automatic e-mail step indicating that proper anti-attack measures have been taken. The iTRAC Process is removed from the Work Items window.

If you go to the Process View window or if you double-click this process, it appears as Complete.

15.4 Correlation

Correlation is the process of analyzing security events to identify potential relationships between two or more events. Correlation allows quick association of priority attacks based on common elements of event data.

The following example is written for the Data Generator Connector that comes installed in Sentinel as a test event generator.

NOTE: Anytime the Data Generator Connector is running, it adds data into your database. Using a correlation rule that is associated with the Data Generator Connector also adds additional data to your database.

- Section 15.4.1, “Creating a Simple Correlation Rule,” on page 329
- Section 15.4.2, “Deploying the Simple Correlation Rule,” on page 329
- Section 15.4.3, “Viewing the Events that Triggered Your Correlated Event,” on page 330

15.4.1 Creating a Simple Correlation Rule

1. Click the *Correlation* tab and select *Correlation Rule Manager* in the navigation bar.
2. In the Correlation Rule Manager window, click *Add*.
3. Click *Simple to create a simple rule*.
4. Select *All* in the *Fire if* drop-down menu.

5. Specify the following
   - SourcePort = 10025
   - DestinationPort = 25

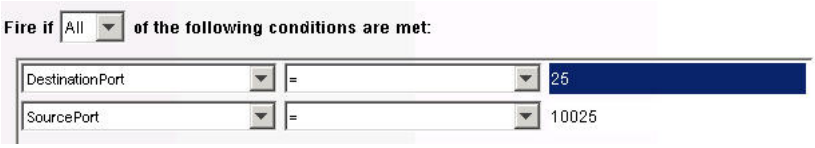

6. Click *Next*.
7. To have this rule fire as many times as possible, select *Continue to perform actions every time this fires*.

8. Click *Next*.
9. In the General Description window, specify a name. A name and description that indicates that this is tutorial rule that does not apply to the network.

10. Click *Next*.
11. Select not to create another rule, then click *Next*.

15.4.2 Deploying the Simple Correlation Rule

1. Click the *Correlation* tab and select *Correlation Rule Manager* in the navigation bar.
2. Click *Tutorial_SourcePort_DestinationPort* > *Deploy Rule*.

3 (Optional) In the Deploy Rule window, add an action. This allows you to:

- Configure Correlated Event
- Add to Dynamic List
- Remove from Dynamic List
- Execute a Command
- Send Email
- Create Incident

4 Click *Next*. The rule indicates deployed by the color green.

15.4.3 Viewing the Events that Triggered Your Correlated Event

1 Right-click the correlated event.

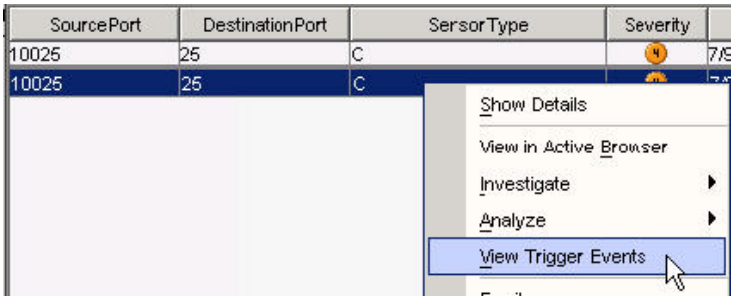

2 Select *View Trigger Events* to see how many events triggered this correlation rule.

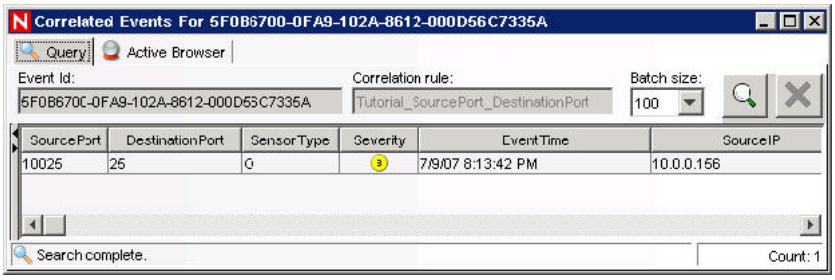

Solution Packs

16

Solution Packs allow Novell partners, and customers to create and easily manage solutions to specific business problems. They provide a framework within which sets of content can be packaged into controls, each of which is designed to enforce a specific business or technical policy. The control can use any of the detection, filtering, alerting, and response features of Sentinel, as well as provide documentation on control status and enforcement. By managing the set of content as a unit within the control, the Solution Pack solves dependency problems and simplifies implementation.

- Section 16.1, “Solution Packs,” on page 331
- Section 16.2, “Solution Manager,” on page 334
- Section 16.3, “Managing Solution Packs,” on page 336
- Section 16.4, “Solution Designer,” on page 352
- Section 16.5, “Deploying an Edited Solution Pack,” on page 361

16.1 Solution Packs

Controls within a Solution Pack can include the following types of content:

- Correlation rule deployments, including deployment status and associated correlation rules, correlation actions, including JavaScript plug-ins and integrators, and dynamic lists
- Reports
- iTRAC workflows, including associated roles
- Event enrichment, including map definitions and event meta tag configuration
- Other associated files added when the Solution Pack is created, such as documentation, example report PDFs, or sample map files.

Although Solution Packs have many uses, one of the most important use is to package content related to governance and regulatory compliance into a comprehensible and easily enforceable framework that is easy to deploy. Novell and its partners offer and extend Solution Packs around such regulations or other customer needs.

Solution Packs are created with Solution Designer application. Using this tool, a user creates the Solution Pack, associated controls and documentation (including implementation and testing steps), and then associates Sentinel content with each control. The entire package is then exported as a ZIP file.

The ZIP file containing the Solution Pack is imported and deployed into an existing Sentinel system by using the Solution Manager in the Sentinel Control Center. The Solution Manager displays implementation and testing steps in the Solution Pack and tracks the status of each control. At any time, users can generate a detailed document with implementation status for each control.

- Section 16.1.1, “Components of a Solution Pack,” on page 332
- Section 16.1.2, “Permissions for Using Solution Packs,” on page 333

16.1.1 Components of a Solution Pack

Solution Packs consist of categories, controls, content, and content groups. These components are represented in a hierarchy. The following image depicts the hierarchy in a Solution Pack:

Figure 16-1 *Solution Pack Hierarchy*

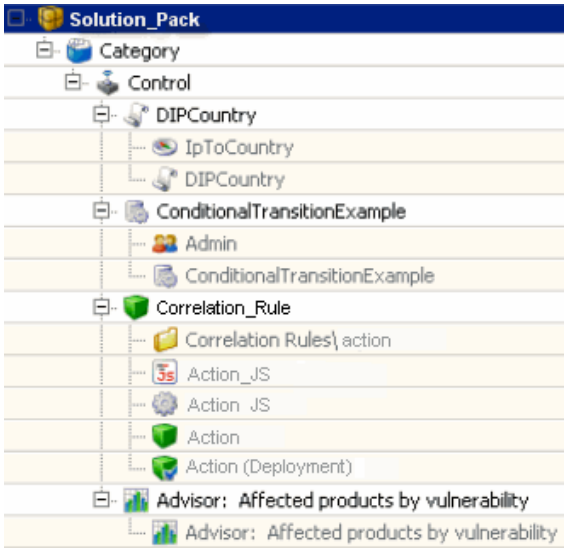

The table below describes each level in a Solution Pack hierarchy.

Table 16-1 *Solution Pack Hierarchy Levels*

Icon	Name	Description
	Solution Pack	Solution Pack is the root node in the content hierarchy. Each Solution Pack can contain one or multiple category nodes.
	Category	Category is a conceptual classification. Each category can contain one or multiple controls.
	Control	Control is another level of classification, which often corresponds to a particular control defined by a set of regulations. Each control can contain one or multiple content groups.
N/A	Content Group	A content group is a set of related content. There are several types of content groups, such as reports, correlation rules, and event configurations, each with its own icon.

The table below describes the types of content groups and the content that they contain.

Table 16-2 *Table 14-2: Types of Content Group*

	Event Configuration	A content group that contains a map definition and the configuration of one or more related Sentinel meta tags. This icon is also used for the meta tag configuration definition.
	Map	Indicates the map definition instance.
	Workflow	A content group that contains an iTRAC workflow template and any associated roles. This icon is also used for the iTRAC workflow template itself.
	Role	Indicates a role used in a workflow.
	Correlation Rule	A content group that contains a correlation rule, the namespace in which it is stored, and any associated correlation actions or dynamic lists. This icon is also used for the correlation rule definition.
	Namespace	Indicates a namespace Instance in which the correlation rule is stored.
	JavaScript Action Plugin	Indicates a JavaScript Action plug-in.
	JavaScript Action	Indicates a configured JavaScript Action instance.
	Integrator Plugin	Indicates an Integrator plug-in.
	Integrator	Indicates a configured Integrator instance.
	Action	Indicates an Action configuration for a correlation action.
	Correlation Rule Deployment	Indicates the correlation rule deployment.
	Report	A content group that contains a JasperReport. This icon is also used for the `.rpt` report file.
	Dynamic List	Indicates a dynamic list.

16.1.2 Permissions for Using Solution Packs

To use the Solution Manager or Solution Designer, a user must be assigned the necessary permissions in the User Manager.

1. Log into the *Sentinel Control Center* as a user with permissions to use the User Manager.
2. Go to the *Admin* tab.
3. Open the *User Configuration* folder.

4 Open the *User Manager* window.

5 Click the *Permissions* tab.

6 Select Solution Designer, Solution Manager, or Solution Pack, which automatically selects both child permissions. The new permissions are applied the next time the user logs in.

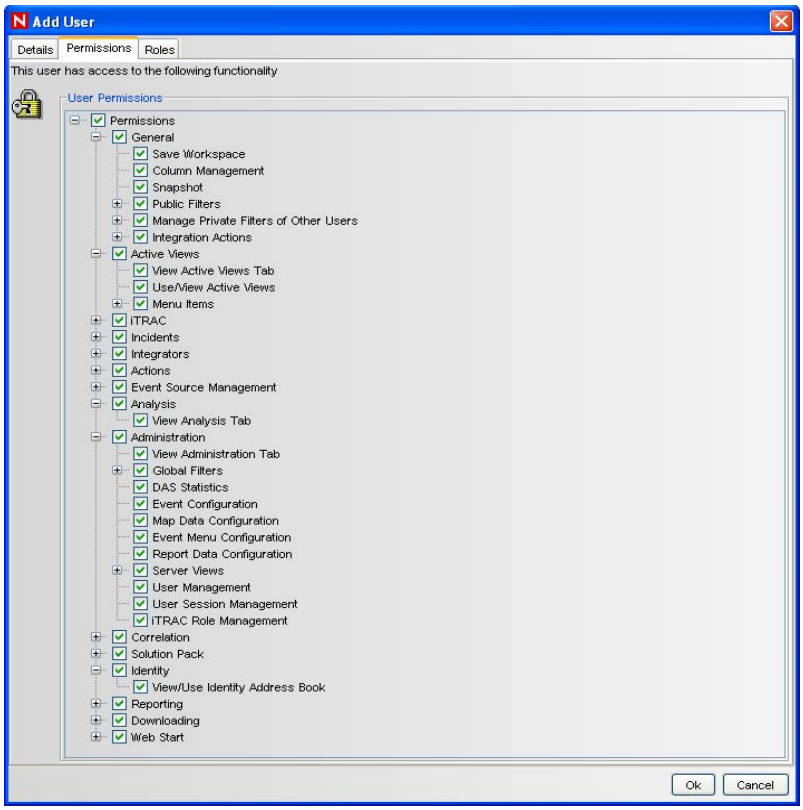

16.2 Solution Manager

After a Solution Pack is imported, the Solution Manager in the Sentinel Control Center is used to install, implement and test each control.

- Installing a control installs the child content for the control into the Sentinel system. When the content is initially installed, its status is Not Implemented.
- Implementing a control is the process to configure event source systems and Sentinel to use the content associated with the control. Novell Solution Packs include detailed documentation describing implementation steps. The user should change the status of the control to Implemented after following all of these steps.
- Testing a control is the process to verify the content associated with the control. Novell Solution Packs include detailed documentation describing testing steps. The user should change the status of the control to Tested after following all of these steps.

To use the Solution Manager, a user must be assigned Solution Manager permissions under Solution Pack. For more information, see Section 16.1.2, “Permissions for Using Solution Packs,” on page 333.

16.2.1 Solution Manager Interface

The Solution Manager window is divided into two frames: Content and Documentation.

- "Content Frame" on page 335
- "Documentation Frame" on page 336

Content Frame

A content frame provides Solution Pack extracted information in ZIP format. The Content frame displays a hierarchical view of the category, control, content group, and various types of content. All parent nodes reflect the overall state of the controls they contain. This means that parent nodes have an inherited status based on their child content.

The Content frame consists of the following columns:

- **Name:** Displays the name of the node.
- **Installed:** Indicates whether the content is installed in the target Sentinel system. If not, this column is blank.
- **State:** This column is available for the control node. This column contains a drop-down box with the following values:
  - **Not Implemented:** The default state when the control is first deployed.
  - **Implemented:** Indicates that the content is fully implemented using the associated documentation.
  - **Tested:** Indicates that you have fully tested the content for this control using the associated documentation.

NOTE: Because of the regulatory significance of implementing controls, status changes for each control are tracked for auditing purposes.

Figure 16-2 *Content Frame*

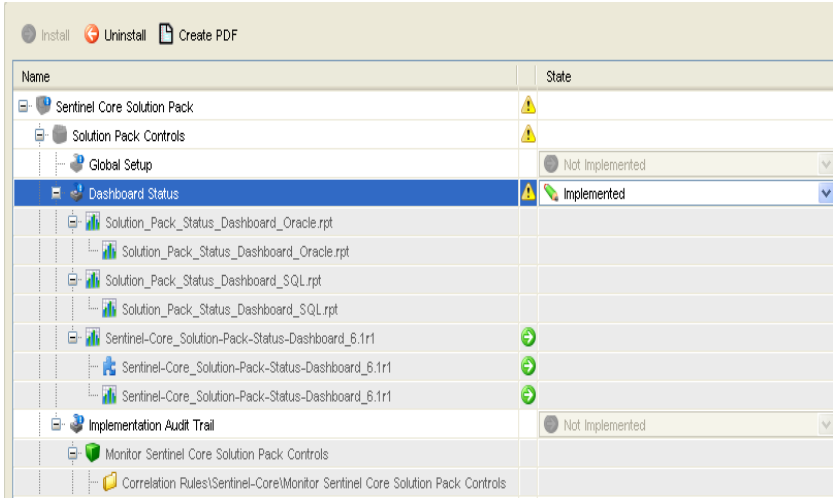

Documentation Frame

The Documentation frame provides a description of selected node.The information was provided when you created the Solution Pack by using the Solution Designer. For more information on the Solution Designer, see Section 16.4, “Solution Designer,” on page 352.

The following informational tabs, which are populated and edited by using the Solution Designer, are available in Documentation frame:

- **Description:** Displays the description of selected node. An additional Attachment panel is included on this tab. You can view attachments and their description in the Description tab.

  The user can add text to the *External ID* field to refer to specific regulations or corporate IDs.
- **Implementation:** This tab, which is associated with the control nodes, displays the instructions for implementing the selected control.
- **Testing:** This tab, which is associated with the control nodes, displays the instructions for testing the selected control.
- **Notes:** This tab, which is associated with the control nodes, is editable. It can be used for any notes related to the control, including user comments on the testing or implementation process.

Figure 16-3 *Documentation Frame*

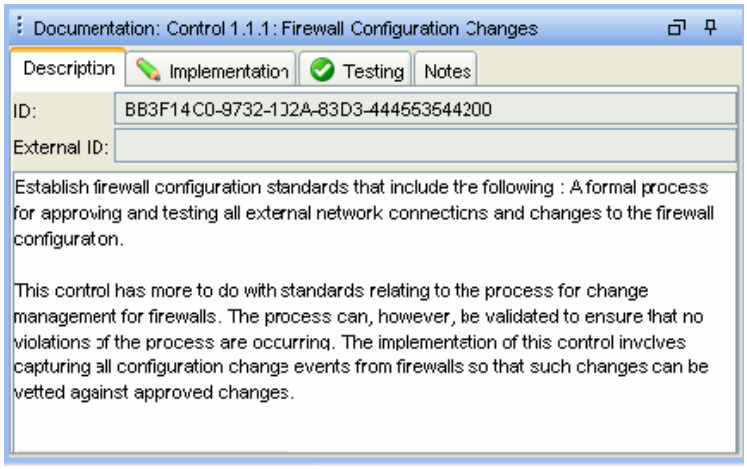

16.3 Managing Solution Packs

- Section 16.3.1, “Importing Solution Packs,” on page 337
- Section 16.3.2, “Opening Solution Packs,” on page 339
- Section 16.3.3, “Installing Content from Solution Packs,” on page 341
- Section 16.3.4, “Implementing Controls,” on page 346
- Section 16.3.5, “Testing Controls,” on page 347
- Section 16.3.6, “Uninstalling Controls,” on page 347
- Section 16.3.7, “Viewing Solution Pack Status,” on page 349
- Section 16.3.8, “Deleting Solution Packs,” on page 351

16.3.1 Importing Solution Packs

Solution Packs are available from several sources. They can be downloaded from the Sentinel product page (http://support.novell.com/products/sentinel/sentinel61.html) (an additional license might be needed). Solution Pack can also be provided by one of Novell's partners, or they can be created from content in your own Sentinel system.

The first step in using a Solution Pack is to import the `.zip` file into the system by using the Import Plugin Wizard. When a Solution Pack is imported, the `.zip` file is copied to the server where the DAS (Data Access Service) components are installed. The actual contents of the Solution Pack are not available in the target Sentinel system until the controls are installed through the Solution Manager.

If you import an updated version of a Solution Pack, you are prompted to replace the existing plug-in.

To import a Solution Pack

1. Click the *Tool* menu and select *Solution Packs*. The *Solution Packs* window displays.

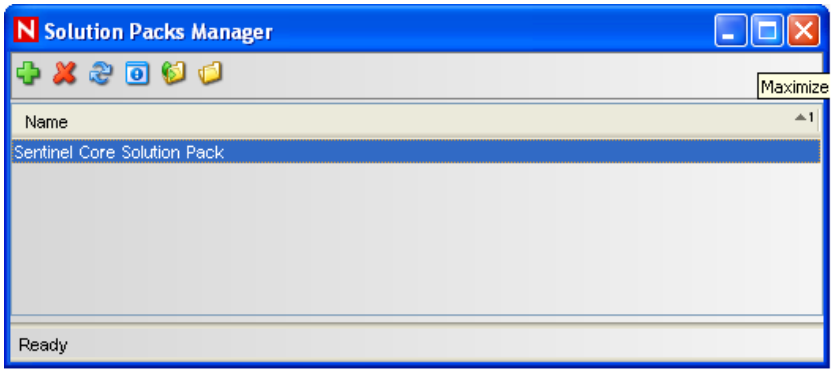

2. Click the *Import* icon in the *Solution Packs* window. The *Import Plugin Type* window is displayed.

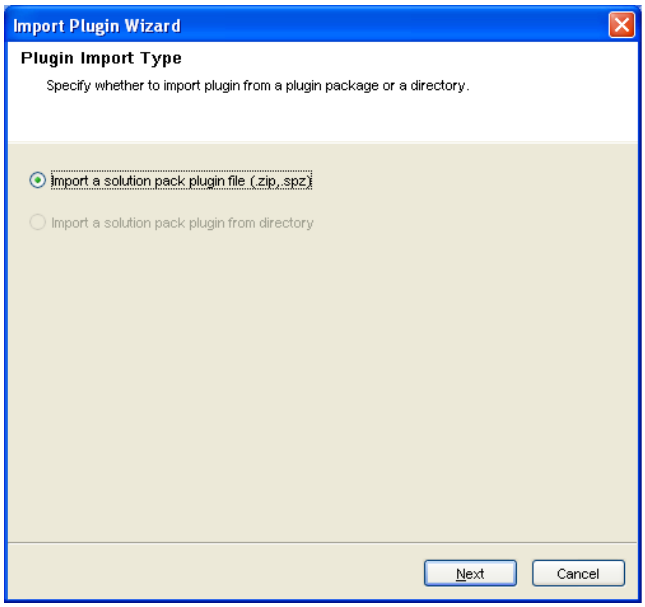

3 Select Import Solution package plug-in file (.zip), then click *Next*. The *Choose Plugin Package File* window displays.

4 Use the *Browse* button to the locate Solution Pack to import to the plug-in repository. Select a ZIP file and click *Open*.

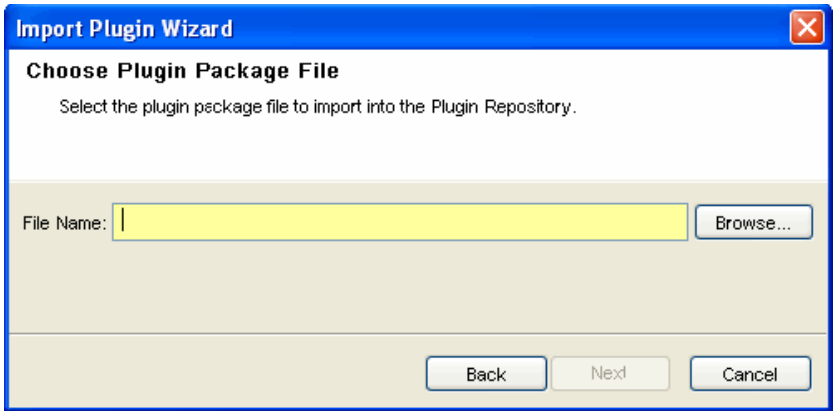

If you have selected a Solution Pack that already exists, the Replace Existing Plugin window displays.

5 Click *Next* if you want to replace the existing plug-ins

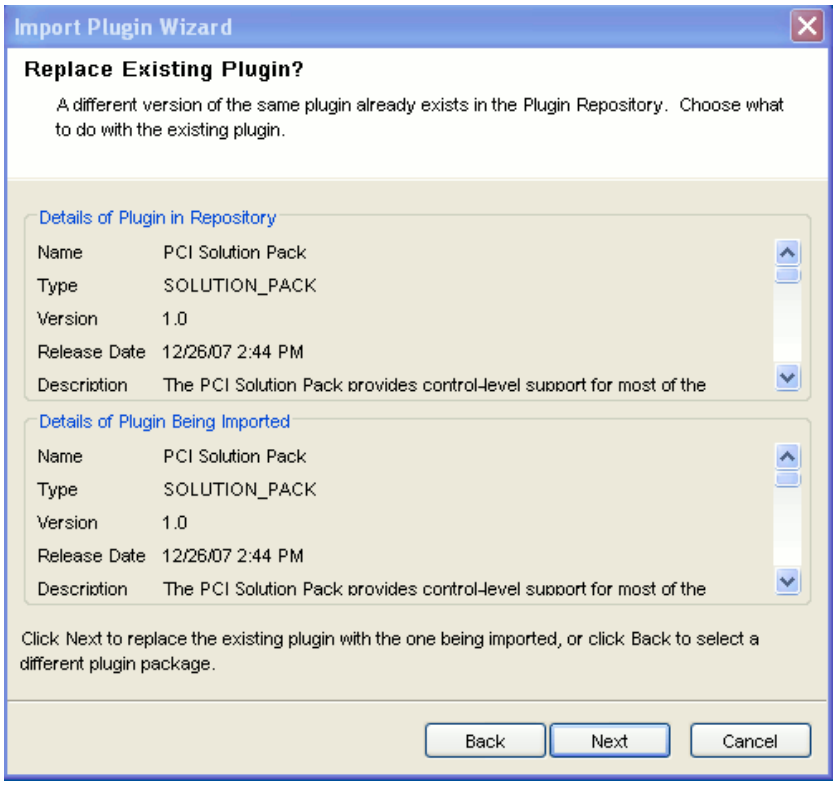

6 Click *Next*. *The Plugin Detail* window displays, including the details of the plug-in to be imported.

7 Select the Launch Solution Manager check box if you want to deploy the plug-in after importing the Solution Pack.

If you select the *Launch Solution Manager* check box, the Solution Manager displays.

8 Click *Finish*.

16.3.2 Opening Solution Packs

To use the Solution Manager and view the contents of a Solution Pack, a user must be assigned Solution Manager permissions. For more information, see Section 16.1.2, “Permissions for Using Solution Packs,” on page 333.

- “Opening a Solution Pack in the Solution Manager” on page 339
- “Content Comparison” on page 340
- “Out Of Sync Status” on page 340

Opening a Solution Pack in the Solution Manager

1 Click the *Tool* menu and select Solution Packs. The Solution Package window displays:

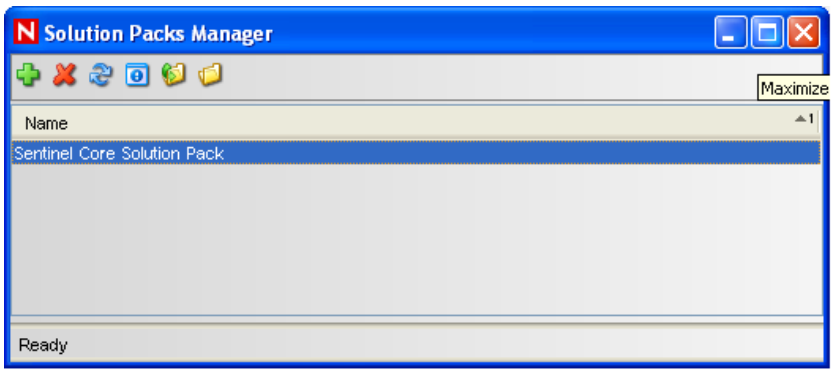

2 Double-click a Solution Pack in the Solution Packs window. The Solution Manager window is displayed.

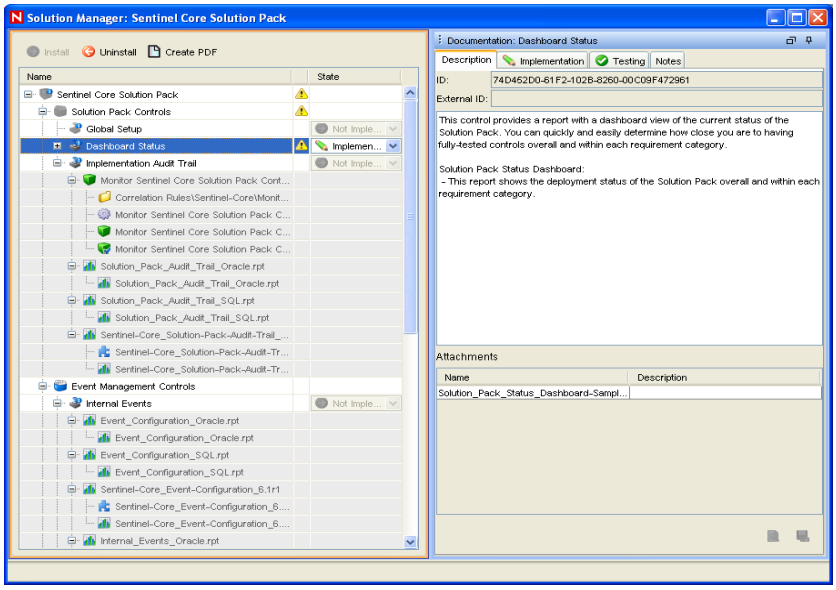

Content Comparison

When the Solution Pack is opened, the Solution Manager compares the contents of the Solution Pack to other Solution Pack content from different Solution Packs or previous versions of the same Solution Pack.

Table 16-3 *Content Status*

Icon	Name	Description
	Installed	Indicates that the content is already installed in the target Sentinel system. The version is the same in the opened Solution Pack and the previously installed Solution Pack.
	Out of Sync	Indicates that a different version of the content is already installed in the target Sentinel system. A difference in name, definition, or description could trigger an Out of Sync status.

Out Of Sync Status

The *Out of Sync* icon indicates that content in the newly opened Solution Pack differs from a version that was previously installed by another Solution Pack (either a different Solution Pack or a previous version of the same Solution Pack). The name, definition, or description of the content might be different.

NOTE: The Solution Manager only compares content from different Solution Packs (or different versions of the same Solution Pack) for installed content. It does not compare content that has not yet been installed. It also does not compare Solution Pack content to content in the target system; manual changes to content in the Sentinel Control Manager are not reflected in Solution Manager.

When you right-click a Solution Pack, you can select *Expand Only Out of Sync Nodes*. This option expands all controls that are out of sync and collapses all controls that are either uninstalled or in sync. This makes it easy to find the out of sync content in a large Solution Pack.

To resolve out of sync content:

1. Select the out of sync content (not the control or category) in the Solution Manager.
2. Right-click and select *Out of sync content details*.

   A message displays with information about which Solution Pack is the source of the out of sync content
3. Compare the description of content item in the two Solution Packs to determine which version you want to keep.
4. Uninstall the out of sync control from all Solution Packs.

   Ideally you should resolve the out of sync issue before installing the new Solution Pack.
5. Reinstall the control with the content you want to keep.
6. Implement and test as required.

16.3.3 Installing Content from Solution Packs

To use the content of a Solution Pack in the Sentinel Control Center, you must install the Solution Pack or selected controls in a Sentinel system (also known as the “target” Sentinel system).

- “Installing the Contents of a Solution Pack” on page 341
- “Correlation Rules and Actions” on page 342
- “JasperReports” on page 344
- “Default Reports” on page 344
- “Content Placeholders” on page 344
- “Duplicate Content within a Solution Pack” on page 345
- “Content with the Same Name in the Target Sentinel System” on page 346

When you install either a Solution Pack or an individual control, all of the child nodes are installed.

Installing the Contents of a Solution Pack

1. Go to *Tools* > *Solution Packs*.
2. Double-click a Solution Pack to open Solution Manager. Alternatively, you can click the *Open with Solution Manager* icon. The Solution Manager window displays.
3. Select a Solution Pack or a control you want to install, then click Install.

   Alternatively, right-click a Solution Pack or control and select *Install*. The Install Control Wizard displays. If you select a Solution Pack, all the controls in that Solution Pack display. If you select an individual control, that control is displayed in the Install Control Wizard window.

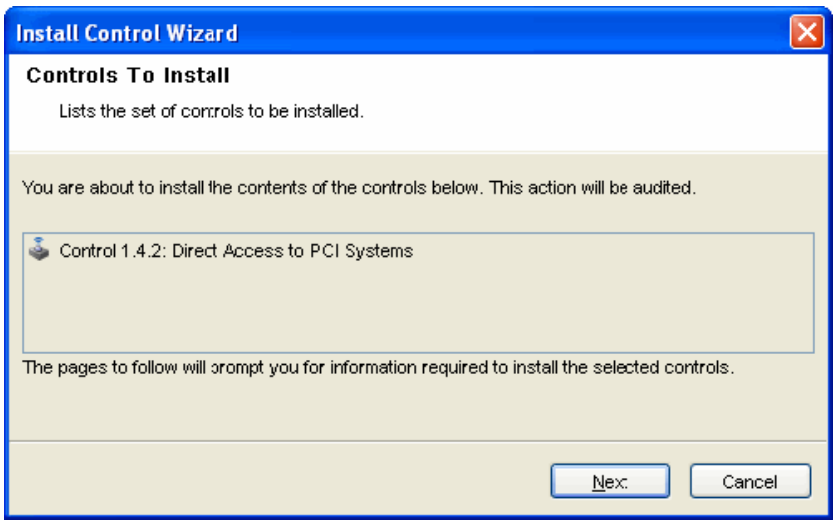

4. Click *Next*. If correlation rules or reports are included in the Solution Pack, you need to proceed through several additional screens until you reach the Install Content window.

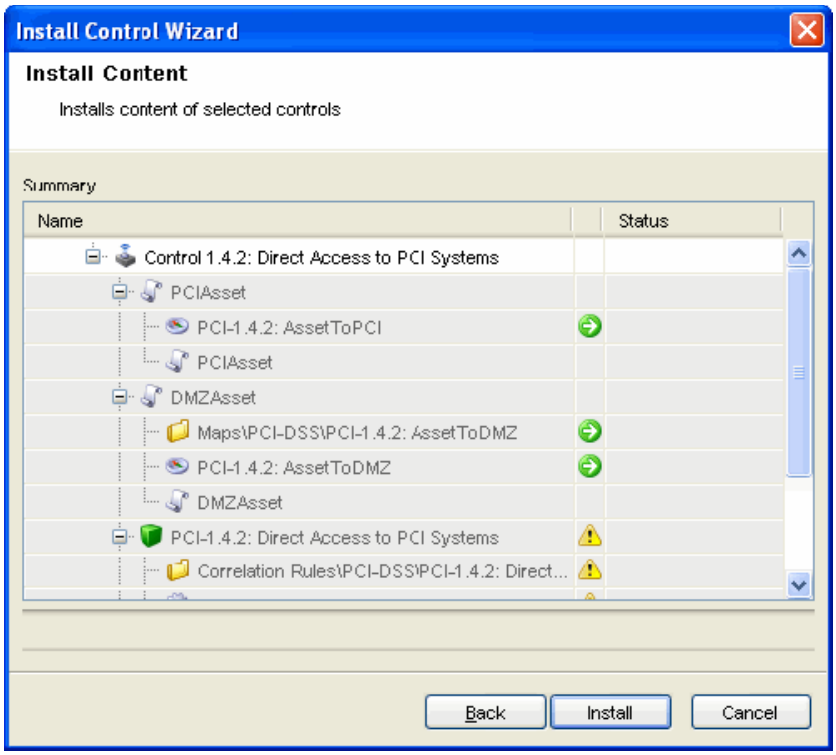

5 Click *Install*.

After installation the *Finish* button displays

6 Click *Finish*.

If the installation fails for any content item in the control, the Solution Manager rolls back all the contents in that control to uninstalled.

There are special considerations for installing certain types of content, including correlation rules and reports; these issues are described below.

Correlation Rules and Actions

Correlation rules are deployed to a specific correlation engine. During the control installation, Figure 16-1 on page 332 shows the correlation engines in the target Sentinel system and the rules that are already running on those engines. Based on the number and complexity of the rules running on the engines, you can decide which correlation engine to deploy the correlation rule to.

Correlation rules deploy in an Enabled or Disabled state, depending on their status in the source Sentinel system when the Solution Pack was created.

If an Execute Script Correlation action (created in Sentinel 6.0) is associated with the correlation rule, the Solution Manager attempts to install the associated JavaScript code on all correlation engines. If any of the correlation engines is unavailable, a message displays.

Figure 16-4 *Install Control Wizard: Select Correlation Engine*

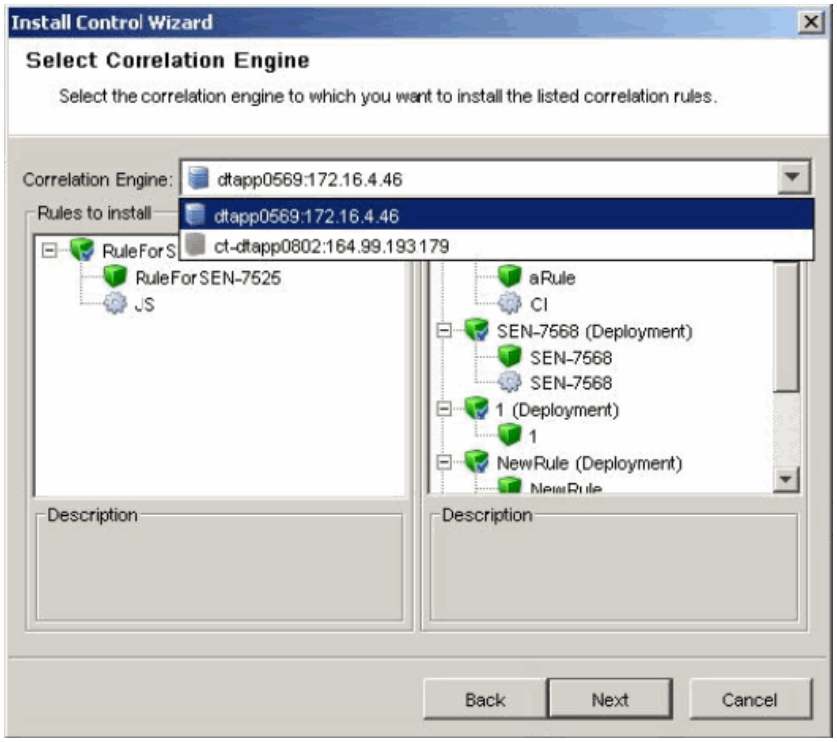

You can cancel the control's installation and fix the problem or continue installation on only the available correlation engines.

Figure 16-5 *Unavailable Correlation Engines*

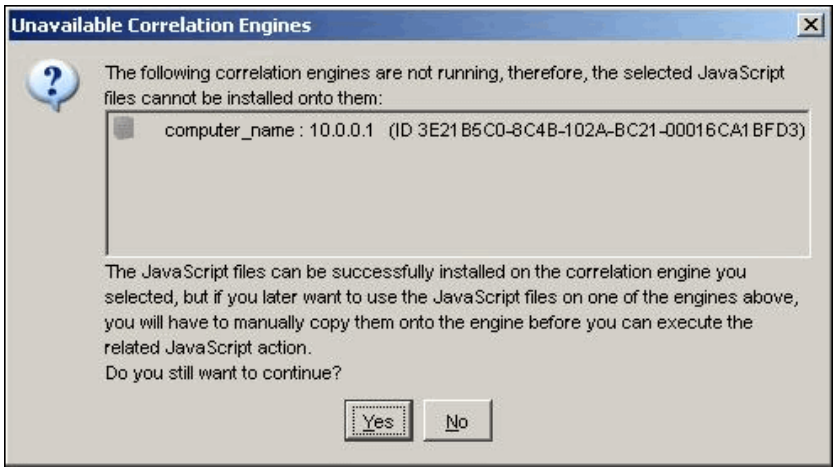

The Execute Script Correlation action (created in Sentinel 6.0) cannot run on a particular correlation engine if the installation of the JavaScript code fails for that correlation engine. The `.js` file can be manually copied to the proper directory on the correlation engine. In a default installation, the proper directory is `<install_directory>/config/exec`.

If an Execute Command correlation action is associated with the correlation rule, the Solution Manager installs the command and its arguments, but the script, batch file, or utility must be manually configured on the correlation engines. This might require installing the utility, configuring permissions, or manually copying a script or batch file to the proper directory on the correlation engines.

In a default installation, the proper directory for the script file is `<install_directory>/config/exec`.

If a JavaScript Action is associated with the correlation rule, the Solution Manager installs the Action configuration, the Action plug-in, and the associated Integrator configuration and Integrator plug-in if needed.

JasperReports

Sentinel Rapid Deployment uses JasperReports for report generation. There are two options to add JasperReports to the Solution Pack. They can either be added from the local machine (`.zip` or `.rpz` files) or from the Sentinel server you are connected to.

Sentinel Rapid Deployment does not support Crystal Reports. However, existing Solution Packs containing Crystal Reports can still be opened/edited/saved in the Solution Designer. When you attempt to install a control that also contains the Crystal Report along with other non-Crystal content such as JasperReports, Correlation rules, Action plug-ins, and Integrator plug-ins, all other contents except the Crystal Report are installed. If you attempt to open a control that contains only Crystal Reports, it stops you with an error message. In both scenarios, a log message is entered to the Sentinel Control Center log.

Default Reports

Sentinel Rapid Deployment bundles the following reports with the Sentinel Core solution pack.

- Sentinel Core Event Configuration
- Sentinel Core Event Source List
- Sentinel Core Event Source Overview
- Sentinel Core Incident Management Dashboard
- Sentinel Core Incident Status Summary
- Sentinel Core Internal Events
- Sentinel Core Solution Pack Audit Trail
- Sentinel Core Solution Pack Status Dashboard

Content Placeholders

Only fully defined controls can be installed. For controls that contain placeholders, the *Install* option is disabled:

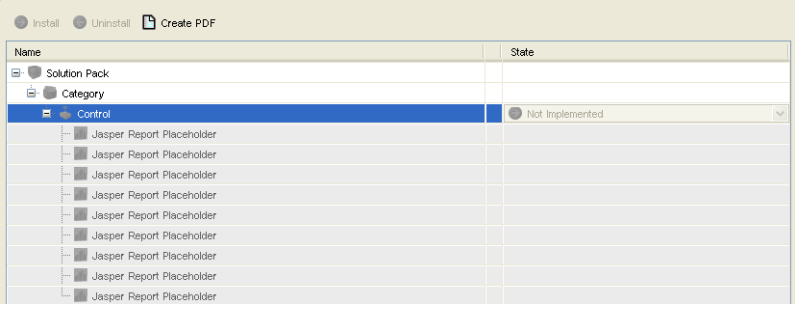

The following warning displays in the Description frame:

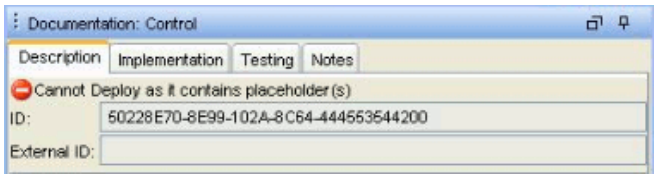

Duplicate Content within a Solution Pack

If two separate controls contain identical content and one control is deployed successfully, the status of the duplicate content in the other control is changed to Installed.The remaining child nodes in the second control stay uninstalled.

Each content item is only installed once. If the same content item (for example, a correlation rule) is included in more than one control, it is only installed once. Therefore, if you install one of those controls, the content displays with an installed status in the other control. In this scenario, the Solution Manager might show that the content for the second control is only partially installed. See Control 1.4.2 in the example below:

Figure 16-6 *Duplicated Content with in a Solution Pack*

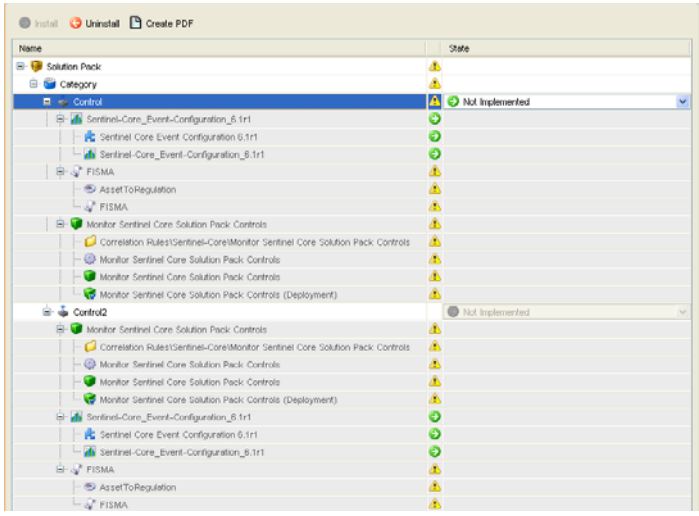

Content with the Same Name in the Target Sentinel System

If the Solution Manager detects content with the same name but a different unique identifier in the target Sentinel system, the Solution Manager installs the content with a unique ID appended to the name. For example, the rule from the Solution Pack might be named Unauthorized Firewall Change (1). The existing rule in the Sentinel system is unchanged.

NOTE: To prevent confusion for end users, Novell recommends that one of these rules be renamed.

16.3.4 Implementing Controls

After the content installation, additional steps might be necessary to fully implement a control, such as the following examples:

- Populating a `.csv` file that is used by the mapping service for event enrichment.
- Scheduling automatic report execution in the Crystal Reports Server.
- Enabling auditing on source devices.
- Copying an attached script for the Execute Command correlation action to the appropriate location on the correlation engines.

These steps should be added when the Solution Pack is created in Solution Designer.

To implement a control:

1. Open a Solution Pack in the Solution Manager.
2. Select a control.
3. Click the *Implementation* tab in the *Documentation* frame.
4. Follow all of the instructions in the *Implementation* tab.
5. Add notes to the *Notes* tab of the Documentation frame as necessary to document progress or necessary deviations from the recommended implementation steps.
6. When the implementation is complete, select the control and change the status drop-down to Implemented.

An audit event is generated and sent to the Sentinel Control Center.

Because of potential legal and regulatory implications, the status for a control should only be changed after all of the implementation steps have been successfully completed.

NOTE: A control must be installed before it can be implemented.

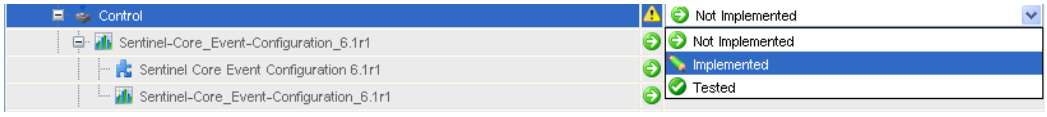

16.3.5 Testing Controls

After the content implementation, the content should be tested to verify that it is working as expected. Testing might require steps such as the following:

- Run a report.
- Generate a failed login on a critical server and verify that a correlated event is created.

These steps should be added when the Solution Pack is created in Solution Designer.

To test a control:

1. Open a Solution Pack in Solution Manager.
2. Select a control.
3. Click the *Testing* tab in the *Documentation* frame.
4. Follow all of the instructions in the *Testing* tab.
5. Add notes to the *Notes* tab of the Documentation frame as necessary to document progress or necessary deviations from the recommended testing steps.
6. When the testing is complete, select the control and change the status drop-down to Tested.

An audit event is generated and sent to the Sentinel Control Center.

Because of potential legal and regulatory implications, the status for a control should only be changed after all of the testing steps have been successfully completed.

NOTE: A control must be installed and should be implemented before it can be tested.

16.3.6 Uninstalling Controls

Controls are often used to meet legal or regulatory requirements. After they are implemented and tested, controls should be uninstalled only after careful consideration.

When a control is uninstalled, the status for the control reverts to Not Implemented and child content is deleted from the Sentinel system. There are a few exceptions and special cases:

- Dependencies are checked to ensure that no content that is still in use is deleted. Some examples of this include a dynamic list that is used by a correlation rule created in the target Sentinel system, a report that is used in a control that is still installed, an iTRAC workflow template that is used in a Solution Pack that is still installed, or a folder that still contains other content.
- Reports copied to a local system cannot be removed if the uninstall is performed from a Sentinel Control Center on a different machine.
- JavaScript files associated with Execute Script Correlation actions remain on the correlation engines.
- Maps (`.csv` files) and the data they contain are not deleted.
- Roles associated with workflows are not deleted.

To uninstall a Control:

1 Right-click the control you want to uninstall and select *Uninstall*. Alternatively, you can click the *Uninstall* icon. The Controls To Uninstall window displays

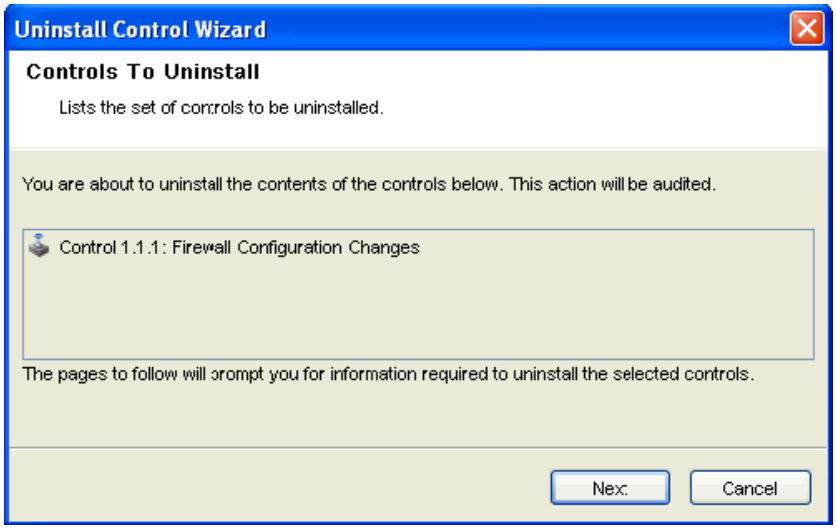

2 Click *Next*

If the control you are uninstalling includes one or more reports, you are prompted whether to uninstall the reports from the local server or the Crystal Reports Server. Ideally, this information was recorded on the *Notes* tab when the reports were installed.

3 Click *Next*. The Uninstall Content window displays.

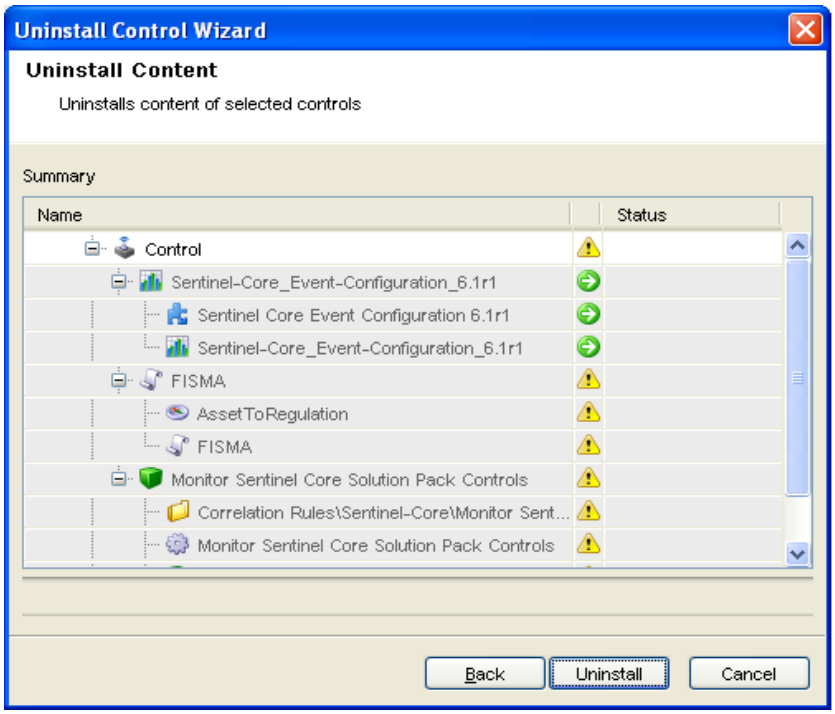

4 Click *Uninstall*. The selected contents are uninstalled.

You cannot uninstall local reports from a different Sentinel Control Center machine than the one that they were installed on or if the files were copied to a new location after installation. If the Solution Manager cannot find the `.rpt` files in the expected location, a message is logged in the Sentinel Control Center log file.

5 Click *Finish*.

16.3.7 Viewing Solution Pack Status

There are several sources of information about the status of a Solution Pack.

- "Viewing Status in the Solution Manager" on page 349
- "Generating Status Documentation" on page 349
- "Audit Events in the Sentinel Control Center" on page 351

Viewing Status in the Solution Manager

You can view the status of Solution Pack contents in the Solution Manager:

- **None/Blank:** No status indicator for a control indicates that the associated content has not been installed yet.
- **Not Implemented:** When none or some of the contents of a control are installed, the control is in the Not Implemented state. If the same content is installed by another control, a control might be Not Implemented even if some of its child content is Installed.
- **Implemented:** This status indicates that a user has completed all of the implementation steps and manually set the control status to Implemented.
- **Tested:** This status indicates that a user has completed all of the testing steps and manually set the control status to Tested.
- **Out of Sync:** This status indicates that a different version of the content in the Solution Pack is deployed in the Sentinel target system by another Solution Pack or a previous version of the same Solution Pack.

Generating Status Documentation

The information about the Solution Pack can be exported in PDF format. The report contains details about every node in the Solution Pack, including category, control, and content group. You can select the following available options:

- **Show status:** Select this option to show deployment status for each control (Not Installed, Not Implemented, Implemented, or Tested) and whether it's Out of Sync.
- **Show individual content:** Select this option to include information about the child content for each control in the documentation.

Figure 16-7 *Status Document*

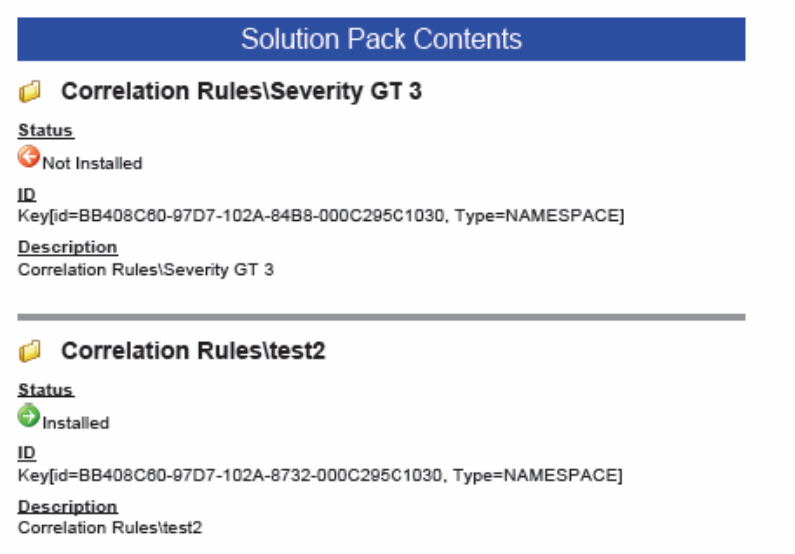

To generate Solution Pack documentation:

1. Open the Solution Pack for which you want to generate a status report.
2. Click *Create PDF*. The Report Options window displays.
3. Select Show status and Show individual content if desired.
4. To view the documentation, click *Preview*. If this is the first time a PDF has been opened from your Sentinel Control Center, you might need to locate Acrobat Reader.

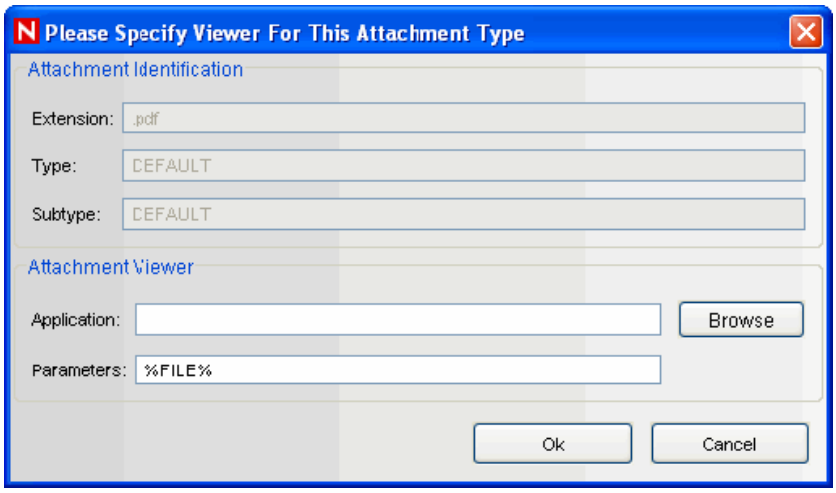

5 To save the PDF, click *Browse*. Navigate the location where you want to save the PDF and specify a filename. Click *Save*.

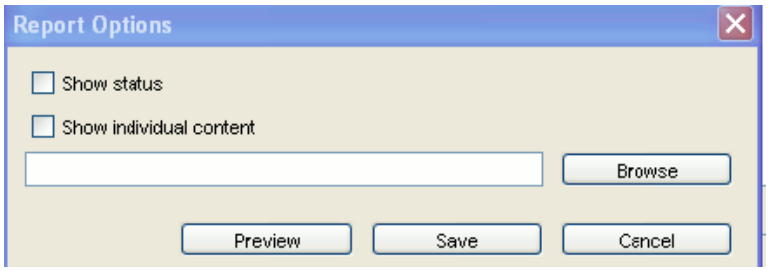

Audit Events in the Sentinel Control Center

All major actions related to Solution Packs and controls are audited by the Sentinel system, with information about which user performed the action. The following events are visible in the Sentinel Control Center and are stored in the Sentinel database:

- Solution Pack is imported.
- Control is installed.
- Control status is changed to Implemented.
- Control status is changed to Tested.
- Control status is changed to Not Implemented.
- Control is uninstalled.
- Notes are modified for a control
- Solution Pack is deleted.

16.3.8 Deleting Solution Packs

Solution Packs are often used to meet legal or regulatory requirements. After they are implemented and tested, Solution Packs should be deleted only after careful consideration.

All deletions are audited by the Sentinel system and sent to both the Sentinel Control Center and the Sentinel database.

1 Click the *Tool* menu and select Solution Packs. The Solution Packs window displays.
2 Select the Solution Pack you want to delete and click the *Open* icon on the toolbar.
3 Select the Solution Pack node and click *Uninstall*. All controls are uninstalled.
4 Close the Solution Manager
5 With the same Solution Pack selected, click *Remove plugin*. Click Yes when you are prompted to delete the Solution Pack.

NOTE: If you attempt to delete a Solution Pack without uninstalling the content first, you are notified that content is still deployed. You have the option to open the Solution Pack in the Solution Manager and uninstall the content.

16.4 Solution Designer

You can use the Solution Designer to package and export different contents, for example, a correlation rule with associated Actions and Dynamic lists and JasperReports. These contents can be selected and packaged in a ZIP file with their respective configuration. You can then view or select the content of the file by using the Solution Manager. For more information on the Solution Manager, see Section 16.2, “Solution Manager,” on page 334.

To use the Solution Designer, a user must be assigned Solution Designer permissions under Solution Pack. For more information, see Section 16.1.2, “Permissions for Using Solution Packs,” on page 333.

- Section 16.4.1, “Solution Designer Interface,” on page 352
- Section 16.4.2, “Connection Modes,” on page 354
- Section 16.4.3, “Creating a Solution Pack,” on page 355
- Section 16.4.4, “Managing Content Hierarchy Nodes,” on page 356
- Section 16.4.5, “Adding Content to a Solution Pack,” on page 357
- Section 16.4.6, “Documenting a Solution Pack,” on page 359
- Section 16.4.7, “Editing a Solution Pack,” on page 360

16.4.1 Solution Designer Interface

The Solution Designer is divided into several frames: Content Palette, Content Description, Solution Pack, and Documentation. The Content Palette includes several sections that can be expanded, including Correlation Deployment, Event Enrichment, Workflow Templates, and Reports. The displayed contents are populated from the Sentinel server and can be exported into a Solution Pack.

Table 16-4 *Table 14-4: Solution Designer - User Interface*

Frames	Image
Content Palette	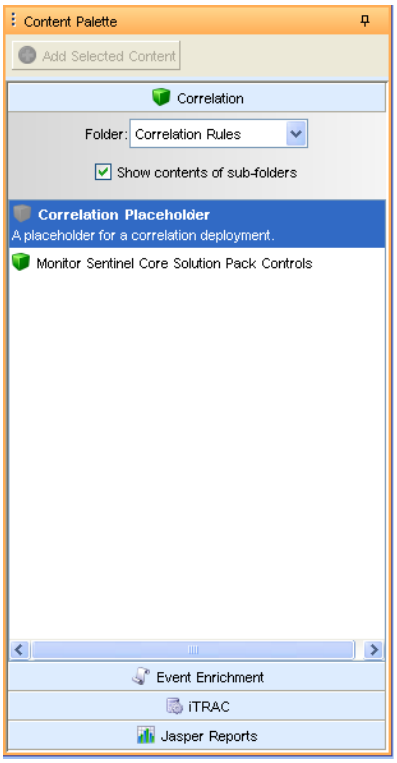
Content Description	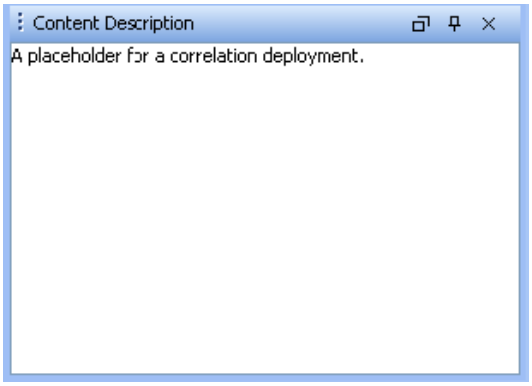

Frames	Image
Solution Pack	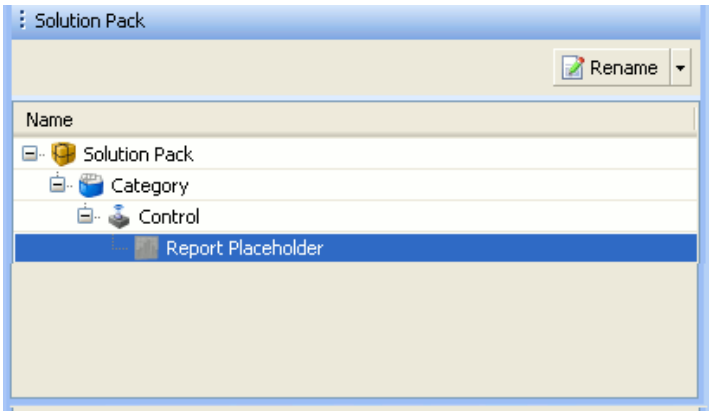
Documentation	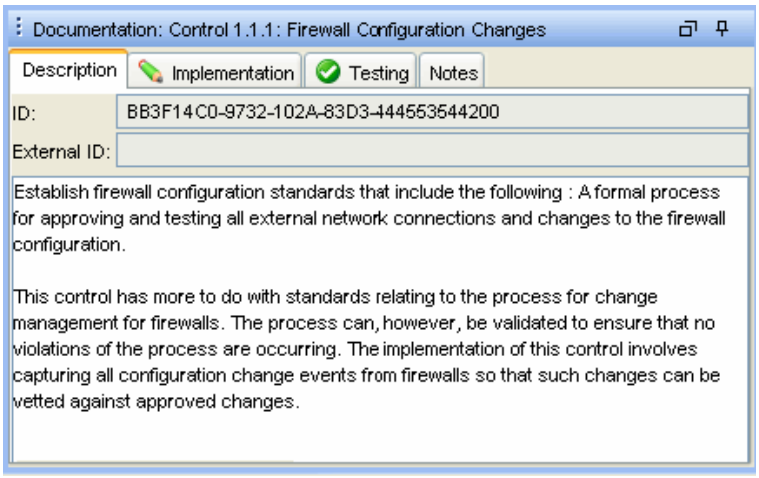

16.4.2 Connection Modes

Solution Packs can be created or edited in the Solution Designer in connected or offline modes.

In offline mode, there is no connection to an active Sentinel server or its content (such as event enrichment or correlation rules). However, you can perform the following actions:

- Define the structure of the Solution Pack (including Categories, controls, and content placeholders).
- Write implementation documentation.
- Write testing documentation.
- Add JasperReports available in your local system.
- Add attachments to any node of the Solution Pack.

In connected mode, all content in the Sentinel system is available. In addition to all of the actions that are available in offline mode, you can also perform the following actions:

- Add Sentinel content (such as correlation rules and Maps).
- Replace placeholders with Sentinel content.

To open the Solution Designer in offline mode:

1 Start the Solution Designer by executing the following command:

```
<install_directory>/bin/solution_designer.sh
```

The Sentinel Solution Designer login window is displayed.

2 Provide your login credentials. Select the *Work Offline* check box if desired, then click *Login*. The Solution Designer is displayed.

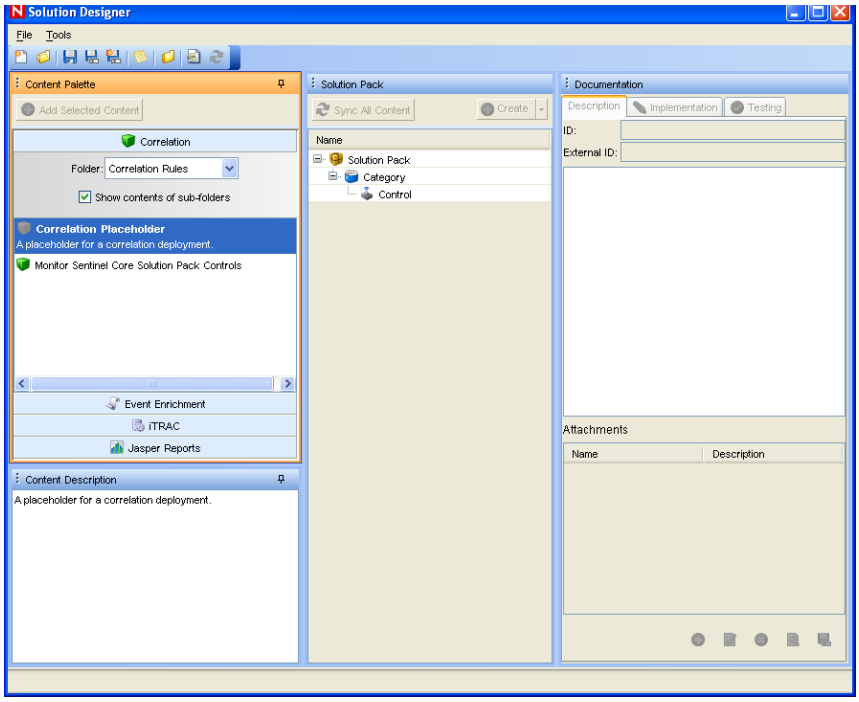

3 Open or create a Solution Pack.

For instructions on creating a Solution Pack see Section 16.4.3, “Creating a Solution Pack,” on page 355.

16.4.3 Creating a Solution Pack

Using Solution Designer, you can create a Solution Pack using existing content objects (for example, correlation rules or dynamic lists) from Sentinel. The Solution Designer analyzes the dependencies for a content object and includes all necessary components in the Solution Pack. For example, a correlation rule deployment includes a correlation rule definition and can also include one or more actions and the ability to create an incident using a workflow. The Solution Designer includes the correlation rule, the associated correlation actions, the iTRAC template, and the roles associated with the iTRAC template in the Solution Pack.

NOTE: To add a content object to a Solution Pack, it must already exist in Sentinel. Content objects cannot be created through the Solution Designer.

To create a new Solution Pack:

1 Open the Solution Designer in either connected or offline mode.

2. Click *File* > *New*. An empty Solution Pack displays in the Solution Pack frame.
3. Add Categories, controls, content groups, and content placeholders, using the proper procedures for each.
4. Add file attachments to the hierarchy nodes as desired.
5. Click *File* > *Save*. The Save window displays.

   Provide a name and click *Save*. The Solution Pack is saved in a `.zip` or `.spz` format.

NOTE: Although you can save a Solution Pack with empty placeholders, you cannot install controls in Solution Manager unless all placeholders have been filled with content.

16.4.4 Managing Content Hierarchy Nodes

All content in a Solution Pack is hierarchically organized into categories, controls, and content groups. These nodes in the hierarchy can be added, deleted, renamed, or reordered.

Table 16-5 *Adding, Deleting, Renaming, and Reordering the Content Hierarchy*

Function	Description
Create	Add a node to the existing control. Right-click an existing node and select Create, or click Create in the Solution Pack frame. Specify the details and click Create.
Rename	Rename an existing node. Right-click an existing node and select *Rename*, or click *Rename* in the Solution Pack frame. Provide the new name and click *OK*.
Delete	Delete a category, control, or content group object. Right-click an existing node and select *Delete*, or click the *Delete* option in the Solution Pack frame. The Delete Selected Objects? message displays. Click *OK*.
View or Edit Properties	View or edit the properties of a Solution Pack, such as the creator. Click *File* > *Properties* from the menu bar or right-click the Solution Pack node and select Properties.
Expand or Collapse Nodes	Expand or collapse all child nodes. Select the Solution Pack or any category, control, or content group level. Right-click a node and select *Expand All* or *Collapse All*.
Move Nodes	category, control, and content group nodes can be created in any order and then reordered or moved to a different parent in the hierarchy. To move a node to another branch in the hierarchy, drag and drop a node to its new parent node. A control can be moved to a new category. A content group can be moved to a new control. To reorder a node, drag and drop it on top of the node where it should appear in the Solution Pack.

16.4.5 Adding Content to a Solution Pack

A vital part of creating a Solution Pack is adding content to the controls. Each control can have one or more types of content associated with it.

- "Sentinel Content" on page 357
- "JasperReports" on page 358
- "Placeholders" on page 358
- "File Attachments" on page 358

Sentinel Content

The same general procedure is used to add all types of Sentinel content to a Solution Pack. The Sentinel content options include the following:

- Correlation rule deployments, including their deployment status (enabled or disabled) and associated correlation rules, correlation actions, and dynamic lists
- Reports
- iTRAC workflows, including associated roles
- Event enrichment, including map definitions and event meta tag configuration
- Other associated files added when the Solution Pack is created, such as documentation, example report PDFs, or sample map files.

The general steps for Sentinel content are described below.

NOTE: Because dynamic list elements and map data are often highly dependent on the system environment, this data is not included as part of the dynamic list or map definition in the Solution Pack. However, this data can be attached to the Solution Pack as a `.csv` file.

To add Sentinel content to a control:

1. Log into Solution Designer in connected mode.
2. Open or create a Solution Pack.
3. Click the appropriate panel to display the available reports from the Content Palette: Solution Pack, category, control, content group and contents.
4. Select the specific content group you want to add.
5. Select the appropriate control or placeholder and click Add Selected Content. Alternatively, drag and drop the selected content group to the appropriate control or placeholder in the Solution Pack frame.

NOTE: If you try to add preexisting content in Solution Designer by dragging and dropping, the existing content is highlighted. After you drop the content, a message prompt displays, stating existence of similar content.

JasperReports

You can add a JasperReport (`.jpz` file) from a local file system. Adding a JasperReport is similar to adding other types of contents.

1. Log into Solution Designer in connected mode or offline mode, then open or create a Solution Pack.
2. Click the Jasper Report panel in the Content Palette. The Jasper Report Panel expands.

   You are prompted about the availability of Jasper Report file on your local machine.
3. Select *Local Jasper Report Plugin*.
4. In the browser window, browse to the location on your local drive where the report is located.
5. Select the file (`.zip` or `.jpz` file) and click *Open*.

Placeholders

If the user is working in offline mode or is not ready to associate content with a control, an empty placeholder can be used instead.

To add a placeholder

1. Click a button in the Content Palette to open the panel for the type of placeholder you want to add: Correlation, Event Enrichment, iTRAC workflow, or report.
2. Drag and drop the placeholder to the appropriate control in the Solution Pack frame.
3. Rename it if desired.

To replace a placeholder with content:

1. Click a button in the Content Palette to open the panel for the type of placeholder you want to replace: Correlation, Event Enrichment, iTRAC workflow, or report.
2. Drag and drop the appropriate content group from the Content Palette to the placeholder in the Solution Pack frame.

File Attachments

You can attach a file or files to any node in the hierarchy, and they are included in the Solution Pack. These files can include anything useful for a user who must deploy the Solution Kit, such as a PDF view of a report, sample map data for event enrichment, or a script for an Execute Command correlation action. These files can be added, deleted, viewed, renamed, or saved to the local machine.

Table 16-6 *File Attachment*

Icon	Name	Description
	Add File	Adds an attachment to a node. The system prompts for another file if you attempt to add one that is already attached. Select a node. Click *Add a new attachment* icon in the Attachments panel. Locate the file, provide a description, and save.

Icon	Name	Description
	View	Views an attachment. Select a node, right-click the attachment in the Attachment panel, then select *View File*. The file displays in the associated application.
N/A	Rename	Renames an attachment. Select a node, right-click the attachment in the Attachment panel, then select *Rename*. Specify the new name and click *OK*.
	Delete	Deletes an attachment. Select a node, right-click the attachment in the Attachment panel, then select *Delete*. Click *OK* to delete.
	Save	Save a copy of the attachment to the local system. Select a node, right-click the attachment in the Attachment panel, then and select *Save As*. Select a file location and click *Save*.

16.4.6 Documenting a Solution Pack

- "Implementation Steps" on page 359
- "Testing Steps" on page 359

Implementation Steps

You need to add the steps required to implement the content in the target Sentinel system to the Implementation tab of the Documentation frame. The steps might include instructions for the following types of implementation actions:

- Populating a `.csv` file that is used by the mapping service for event enrichment.
- Enabling auditing on source devices.
- Copying an attached script for an Execute Command correlation action to the appropriate location on the correlation engines.

After the content implementation, the content should be tested to verify that it is working as expected.

Testing Steps

You need to add the steps required to test the content in the target Sentinel system to the Testing tab of the Documentation frame. The steps can include instructions for the following types of testing activities:

- Run a report and verify that data is returned.
- Generate a failed login on a critical server and verify that a correlated event is created and assigned to an iTRAC workflow.

16.4.7 Editing a Solution Pack

A saved Solution Pack can be edited by using the Solution Designer. For information about deploying the changes into an existing system, see Section 16.5, “Deploying an Edited Solution Pack,” on page 361.

When an existing Solution Pack is saved, the user has several options:

- **Save:** Saves an updated version of the original Solution Pack. If the Solution Pack is re-imported into a Sentinel system, it replaces the old version.
- **Save As:** Saves a renamed version of the original Solution Pack. If the Solution Pack is re-imported into a Sentinel system, it replaces the old version.
- **Save As New:** Saves a Solution Pack with a new unique identifier. If the Solution Pack is imported into a Sentinel system, it does not impact any previously imported Solution Packs.

To edit a Solution Pack:

1. Start the Solution Designer by executing one of the following commands:

   `solution_designer.sh (in <install_directory>/bin)`

   The Sentinel Solution Designer login window displays.
2. Provide your login credentials. Select the *Work Offline* check box if desired, then click *Login*. The Solution Designer displays.

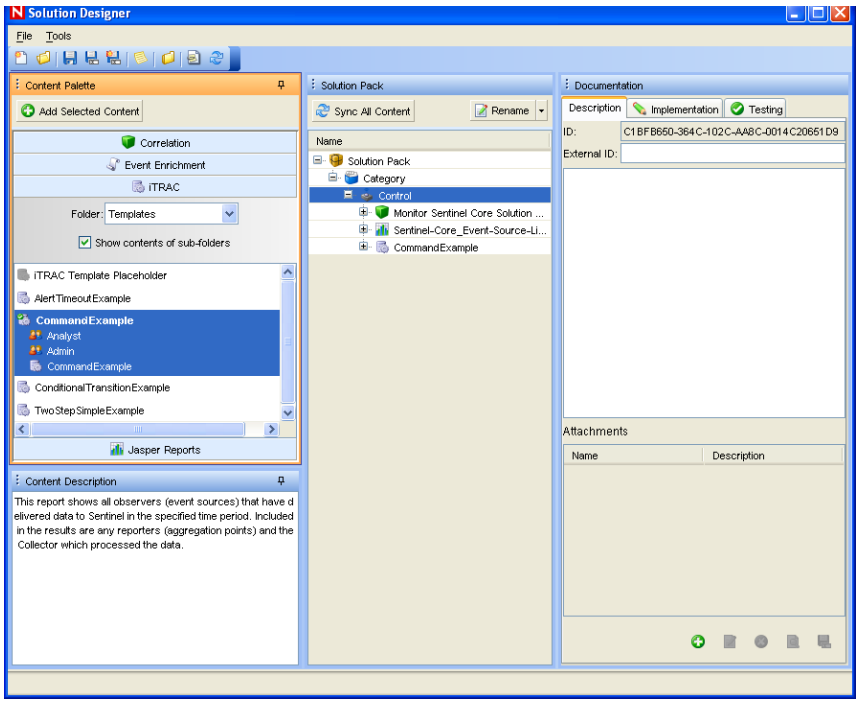

3. To edit a Solution Pack, click *File* > *Open*. Browse and select the existing Solution Pack ZIP file. Click *Open*.
4. To update the Solution Pack with modified content from the source Sentinel system, drag and drop the content from the Content Palette to the appropriate control.
5. Add or delete controls as necessary.

6 Click File > Save, Save As, or Save As New, and save the file to the location you want.

If you selected Save or Save As and some of the content is out of sync, you are prompted to synchronize. See "Out Of Sync Status" on page 340for instructions on how to synchronize content.

Out of Sync Content

If the content in the source system is modified, the content in the source system and the content in the original Solution Pack can be out of sync.

- You can drag and drop the content from the Content Palette onto the control.
- For simple content with no dependencies, the modified content is immediately updated. For example, a report has no dependencies.
- For content with dependencies, the dependencies are checked and updates are made when you click Sync All Content or when you save the Solution Pack.

NOTE: In the special case in which an action uses the Send Email action that is included in all Sentinel systems by default, the Send Email action always appears as Out of Sync. This is expected and does not cause an error.

16.5 Deploying an Edited Solution Pack

When a Solution Pack is modified and saved by using the *Save* or *Save As* options in the Solution Designer, it is considered to be a new version of the original Solution Pack. When it is imported, it replaces any older versions of the original Solution Pack. There is no immediate impact on any installed content in the target Sentinel system.

After the Solution Pack is installed, its behavior varies depending on the status of the original Solution Pack's content.

- If the content from the original Solution Pack was not installed yet, the content is simply replaced. When a user installs content, the new content is installed to the target Sentinel system.
- If the content from the original Solution Pack was installed (Not Implemented), Implemented, or Tested, the original content is compared to the new content.
- If the content version is the same, the original content is still valid and no action is necessary.
- If the content version is different, the content status is set to Out of Sync. The user must decide how to resolve the synchronization issue. For more information, see "Out Of Sync Status" on page 340.
- If the content did not exist in the original Solution Pack, it is displayed in Solution Manager as Not Installed. You can install, implement, and test the new content.
- If the content existed in the original Solution Pack but has been deleted from the modified Solution Pack, it does not appear in the Solution Manager.

NOTE: The Solution Manager only handles differences in the contents of Solution Packs. It does not recognize manual content changes that are performed after content is installed.

Action Manager and Integrator

17

Actions are used to execute some type of action in Sentinel, either manually or automatically. An action plug-in framework was introduced in Sentinel 6.1. This framework consolidates several different ways of executing actions in Sentinel 6.0. The same Action framework is now used to execute actions in all of the following contexts:

- When a deployed correlation rule fires (automatic)
- When a user chooses the action from within an incident
- When a user chooses a right-click menu option using an action in an Active View or other event table

The plug-in framework has several advantages over the method for using JavaScript actions in previous versions of Sentinel.

- There is no need to place the JavaScript file in a particular directory. The plug-in is placed in a central repository.
- There is no need to manually distribute the file to multiple machines in a distributed environment. The plug-ins are downloaded as needed.
- Importing the updated plug-in from one Sentinel Control Center machine is sufficient to update the plug-in everywhere it is used.

One or more configured action instances can be created from an action plug-in by using different parameters.

An action can be executed on its own, or it can make use of an Integrator instance, configured from an Integrator plug-in. Integrators provide the ability to connect to an external system, such as an LDAP, SMTP, or SOAP server, to execute an action.

- Section 17.1, “Action Manager,” on page 363
- Section 17.2, “Action Plug-Ins,” on page 365
- Section 17.3, “Actions,” on page 376
- Section 17.4, “Integrator Manager,” on page 382
- Section 17.5, “Integrator Plug-Ins,” on page 384
- Section 17.6, “Integrators,” on page 385

17.1 Action Manager

The Action Manager allows you to configure repeatable actions that can be executed in various contexts throughout the Sentinel system. The Action Manager allows you to configure the following types of actions:

- Configure a Correlated Event
- Add to Dynamic List
- Remove from Dynamic List
- Execute a Command

- Send an Email
- Create an Incident
- Execute JavaScript Action Plug-ins

NOTE: Except for JavaScript actions, the actions above can only be used in the context of a correlation rule deployment. For more information about correlation-only actions, see the Correlation section. This section focuses exclusively on JavaScript action plug-ins and actions.

Using the Action Manager you can import, create, and manage action plug-ins (`.zip` files) and configure specific action instances.

To use action plug-ins, a user must be assigned the necessary permissions in the User Manager. By default these permissions are assigned to admin user.

1. Log into the Sentinel Control Center as a user with permissions to use the User Manager.
2. Go to the Admin tab.
3. Open the User Configuration folder.
4. Open the User Manager window. Double-click the desired user. The User Details window displays.
5. Click the Permissions tab.

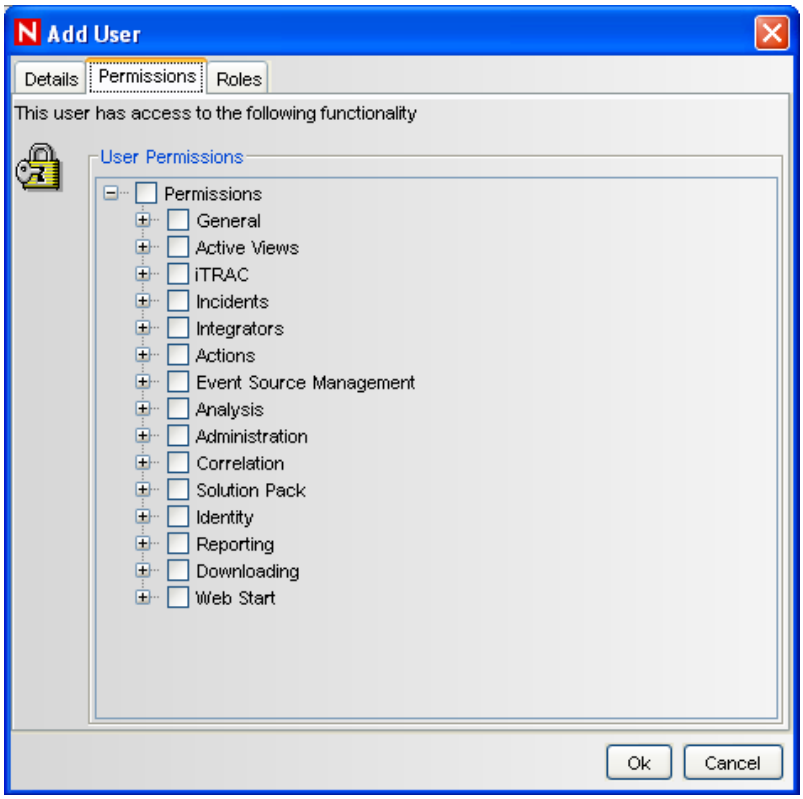

6. Select View Actions, Manage Actions, or Manage Action Plugins (which automatically selects all child permissions). The new permissions are applied the next time the user logs in. For more information, see "Sentinel 6.1 Rapid Deployment Control Center User Permissions" in the *Sentinel 6.1 Rapid Deployment Reference Guide*.

17.2 Action Plug-Ins

You can download action plug-ins from the Sentinel Plug-in Web site (http://support.novell.com/products/sentinel/secure/sentinelplugins.html).

Action plug-ins are frequently included in Solution Packs. Also, JavaScript actions used in Execute Script actions in versions of Sentinel before Sentinel Rapid Deployment can be converted to action plug-ins by using the Action Manager.

- Section 17.2.1, “Importing JavaScript Action Plug-Ins,” on page 365
- Section 17.2.2, “Importing JavaScript Files,” on page 368

17.2.1 Importing JavaScript Action Plug-Ins

JavaScript plug-ins from Novell or other sources can be imported into Sentinel.

1 Click *Tool* > *Action Manager*. The Action Manager window displays.

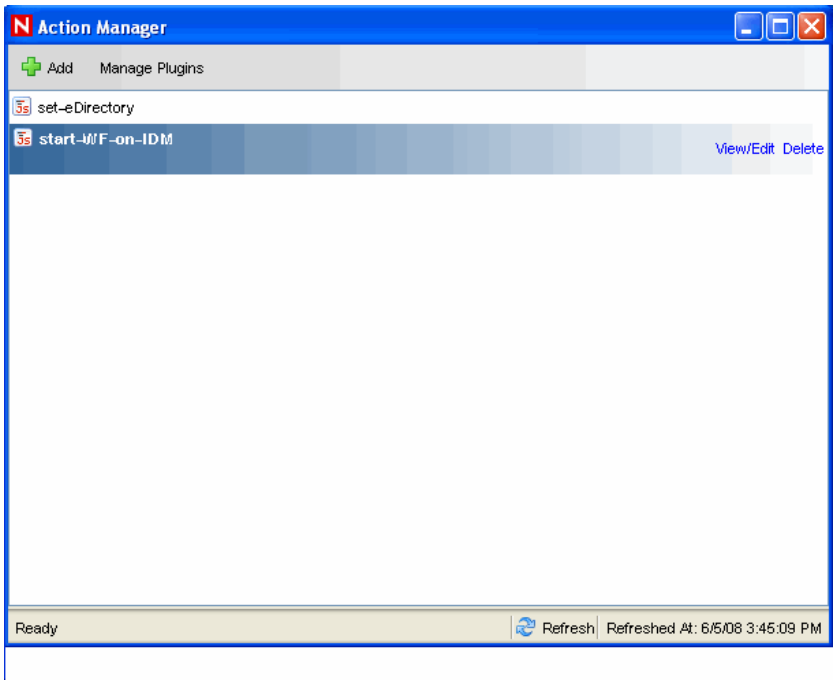

2 Click *Manage Plugins*. The Action Plugin Manager window displays.

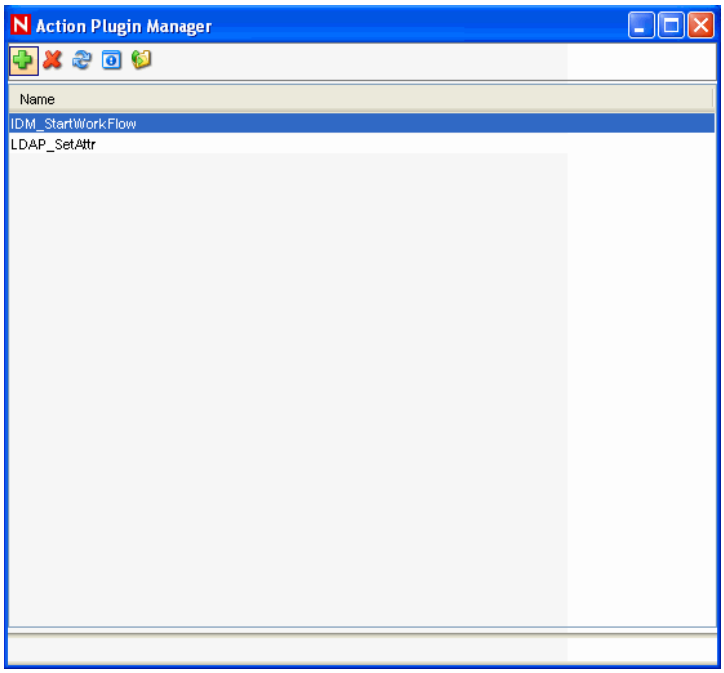

3 Click the *Add* icon on the top left corner to import plug-ins. The *Plugin Import Type* window displays.

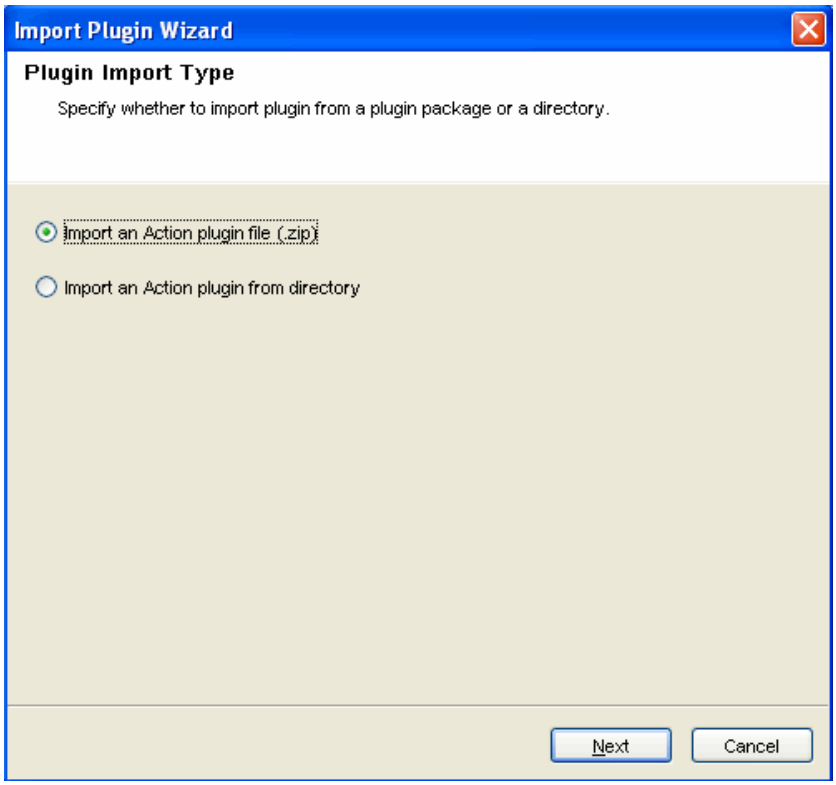

4 Select Import an Action plugin file (`.zip`). Click *Next*.

The Choose Plugin Package File window displays.

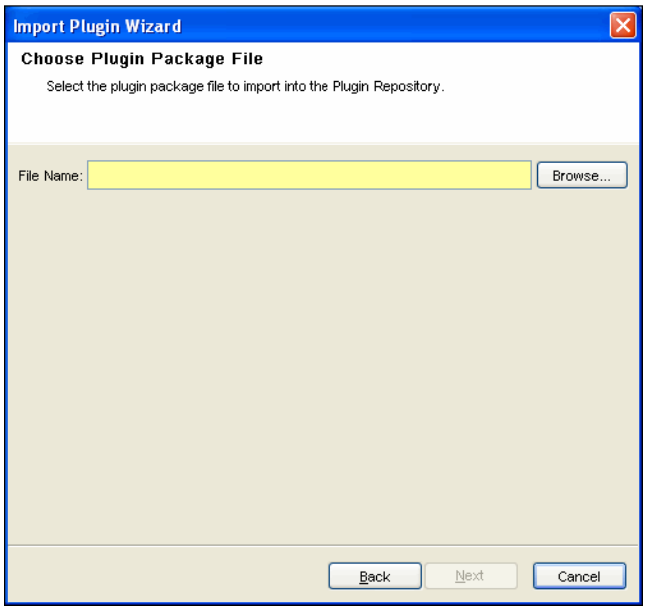

5 Browse to a location of the plug-in package file and click OK, then click *Next*.

If the file you have selected is not the proper format, the Next button does not activate.

If you are updating an already-imported plug-in file, you are provided with the option of updating the existing plug-in, going back and selecting a different plug-in, or canceling the import.

6 If you want to continue, click *Next*.

The Plugin Details window displays. Details of the plug-ins to be imported are displayed.

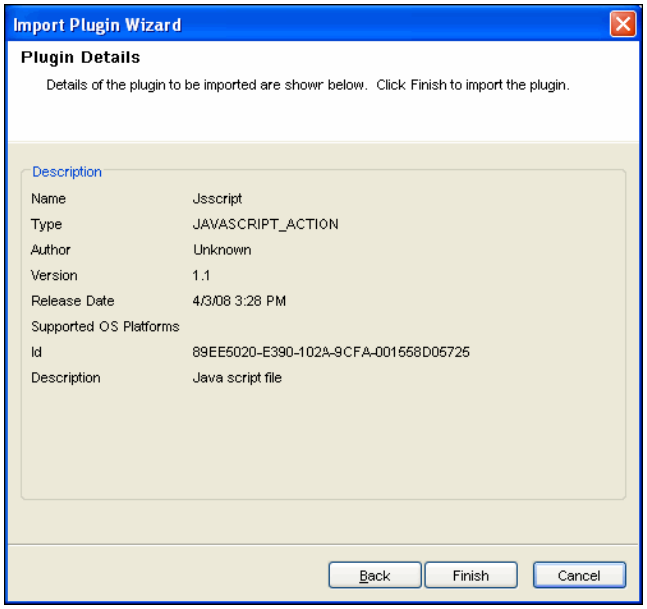

7 Click *Finish*.

17.2.2 Importing JavaScript Files

Although JavaScript action plug-ins can be obtained from Novell, it is also possible to create and manage your own JavaScript action plug-ins. Plug-ins can be created by using JavaScript files that were used in the Execute Script command in versions prior to Sentinel Rapid Deployment, or they can be created using any JavaScript file written by using the Sentinel JavaScript API.

NOTE: For information about the API for developing JavaScript scripts for Sentinel correlation, see Sentinel JavaScript Action API on the Novell Developer Community Web site (http://developer.novell.com/wiki/index.php?title=Develop_to_Sentinel).

After you import a JavaScript file into Sentinel, a JavaScript action plug-in is created and stored in the central plug-in repository. Then the action plug-in can be used to configure an action instance. Unlike the Sentinel 6.0 Execute Script command, the JavaScript file does not need to be manually moved to a specific directory location.

When you import a JavaScript file from a directory, it is important to define the required objects correctly so the JavaScript actions that use the plug-in are available in the right parts of the Sentinel Control Center interface. The following table shows the Required Objects options in the import wizard and where the actions are available if those options are selected.

Table 17-1 *Required Objects*

	Actions Available for Selection in these Contexts:			
Required Object	**Event Menu Configuration**	**Deploy Correlation Rule**	**Associate with Create Incident Correlation.Action**	**Execute Incident Action**
None	Yes	Yes	Yes	Yes
Event	Yes	Yes	Yes	Yes
Correlation Rule	No	Yes	Yes	Yes
Incident	No	No	Yes	Yes

To import JavaScript files:

1. Click *Tool* > *Action Manager*. The Action Manager window displays.

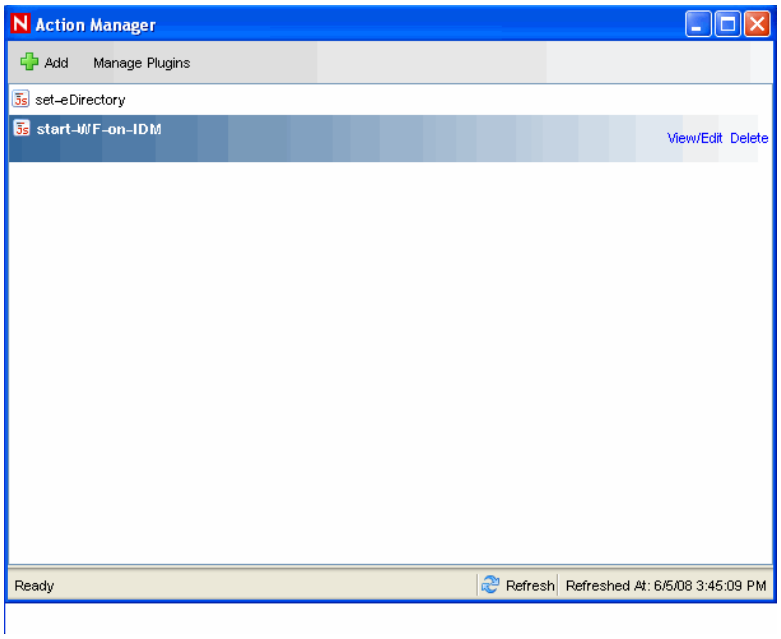

2 Click *Manage Plugins*. The Action Plugin Manager window displays.

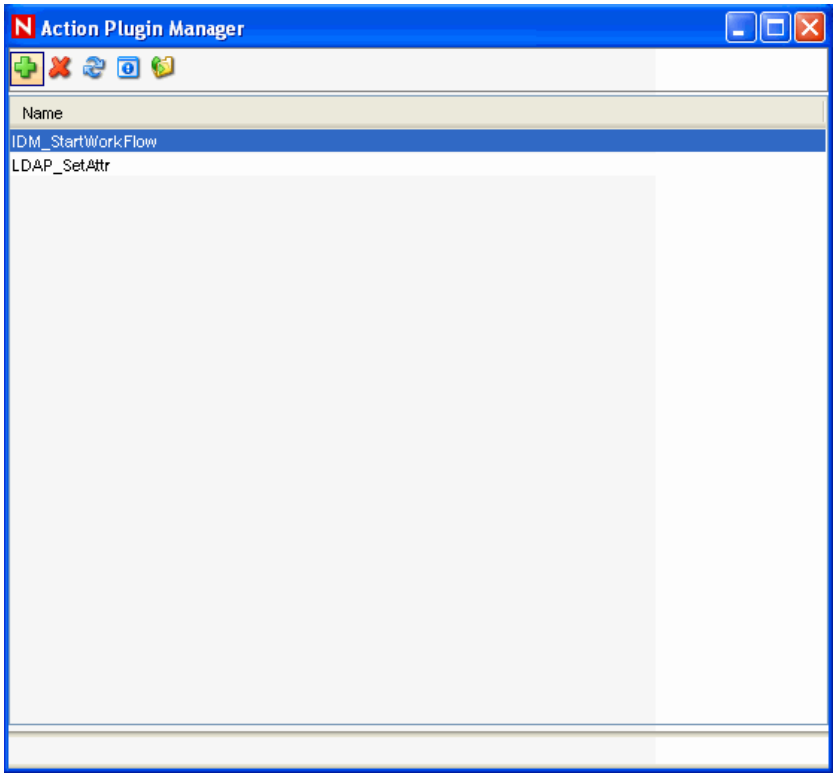

3 Click the *Add* icon on the top left corner to Import plug-ins. The Plugin Import Type window displays.

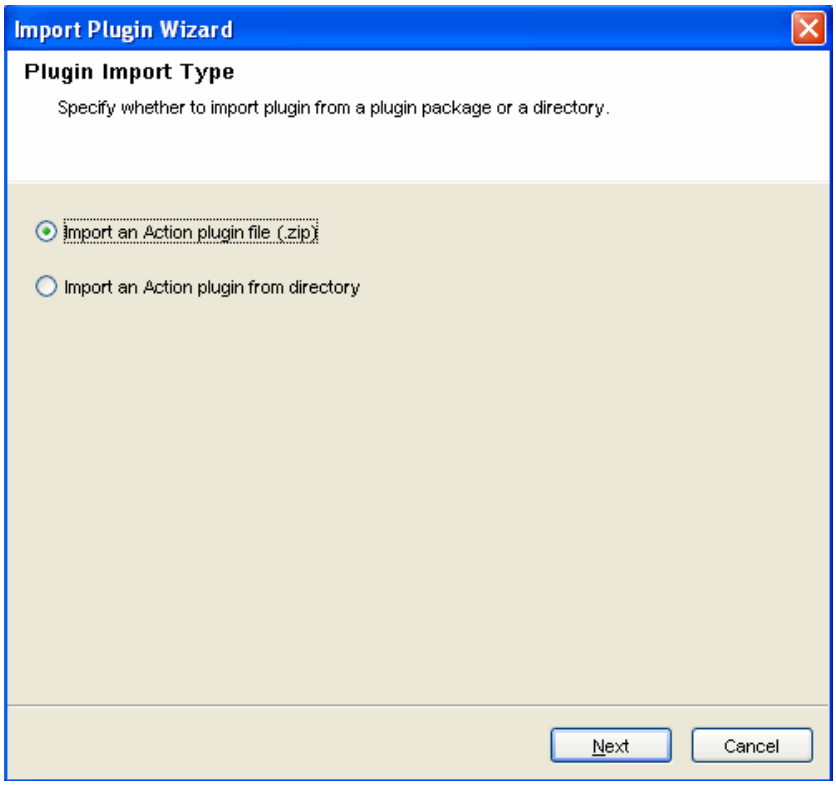

4 Select *Import an Action plugin from directory*. The Choose JavaScript Directory window displays.

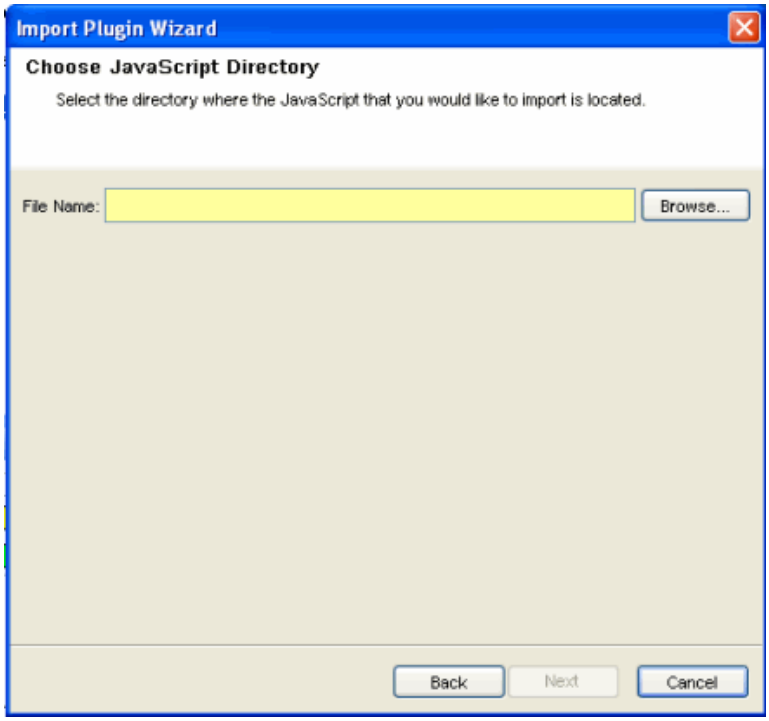

5 Browse to a location of the JavaScript Plug-in directory and click *OK*, then click *Next*.

6 The Action Plugin Detail window displays. Provide the required information. Attach a main JavaScript file and a help file.

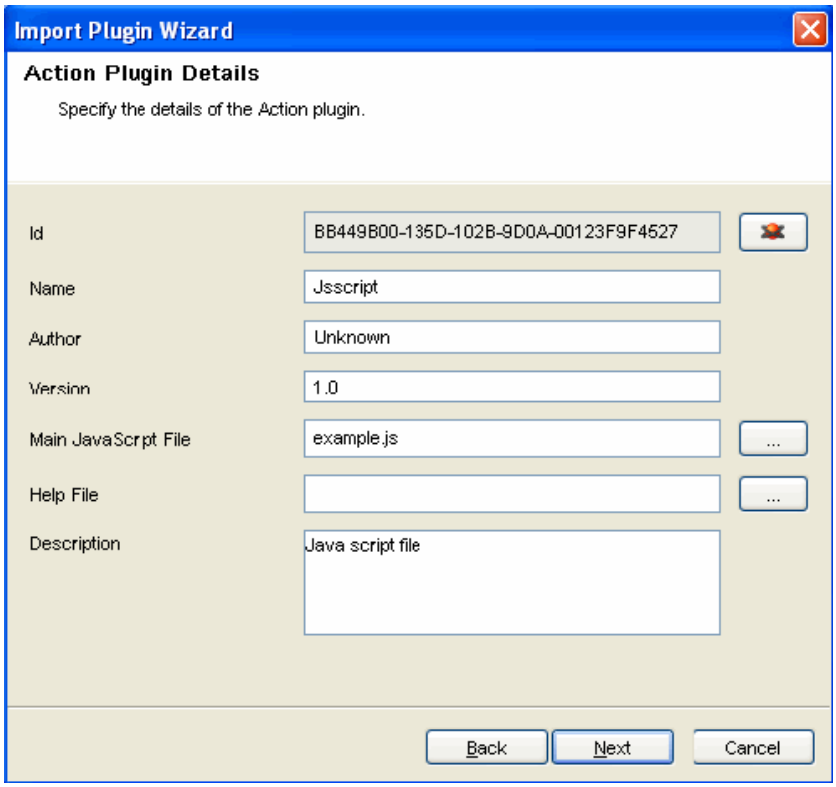

If the file you have selected is not the correct format, the Next button does not activate.

When you are updating an already-imported JavaScript file, you are provided with the option of updating the existing plug-in, going back and selecting a different plug-in, or canceling the import.

7 If you want to continue, click *Next*.

The Required Input window displays.

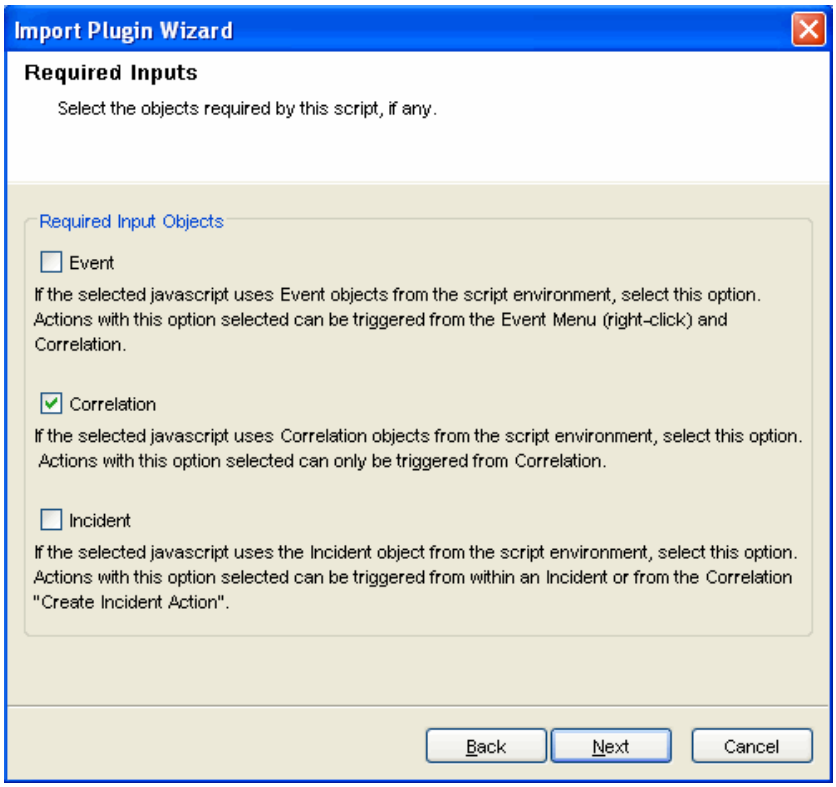

8 Select the objects that the JavaScript action requires.

This affects where the action is available in the interface. For more information, see the Table 17-1 on page 368.

9 Click *Next*. The Plugin Parameters window displays.

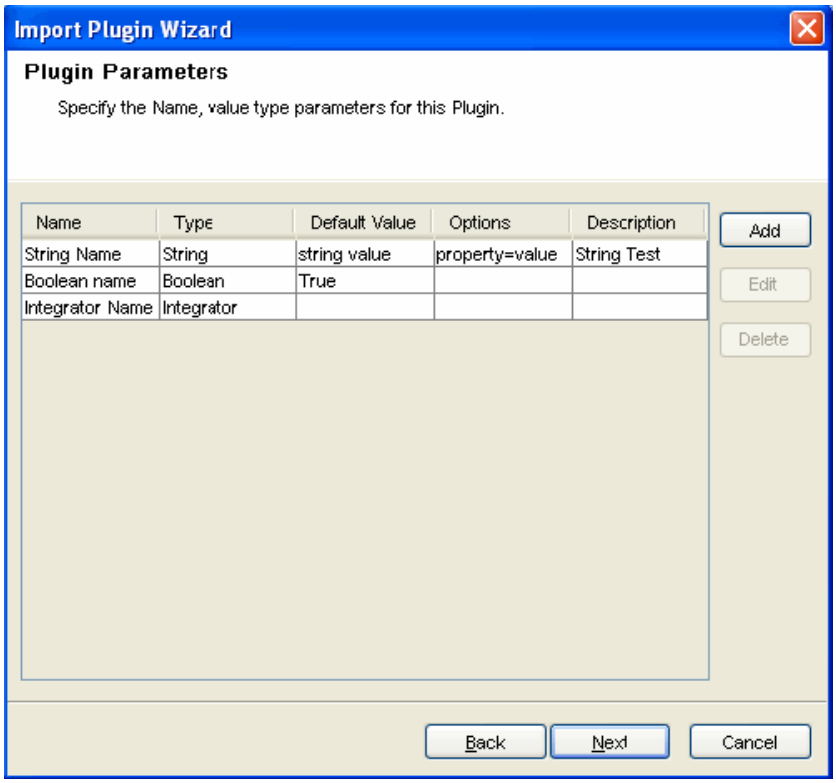

10 [Optional] Click the Add button to add parameters that can be set when an action is configured.

This option should be used for any JavaScript files that expect to receive parameterized information. The Parameter Definition window displays.

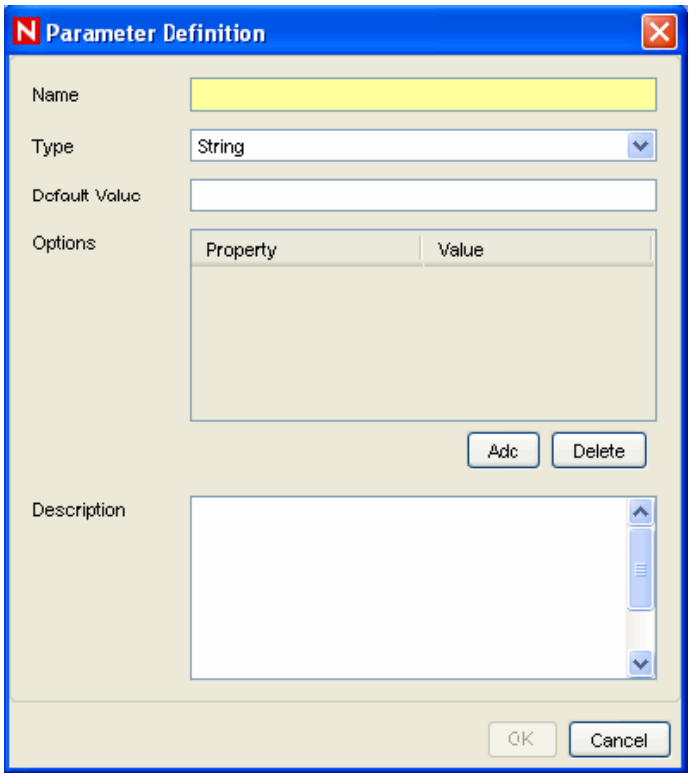

10a Specify the parameter name.

The name used here should be identical to one used in the JavaScript API method `scriptEnv.getParameter` in the script that is being imported.

10b Select parameter name from Type drop-down list.

The various parameter types available are:

- **String:** Accepts the sting values for the parameters.
- **Boolean:** The parameter can take a True or False value.
- **Integrator:** Select Integrator name for the parameters.
- **Event Tag:** Select an Event Tag for the parameters.
- **Severity:** Select the Severity for the parameters.

NOTE: The Options area is only available for string-type parameters.

10c (Optional) Specify a description.

10d Click *Next*.

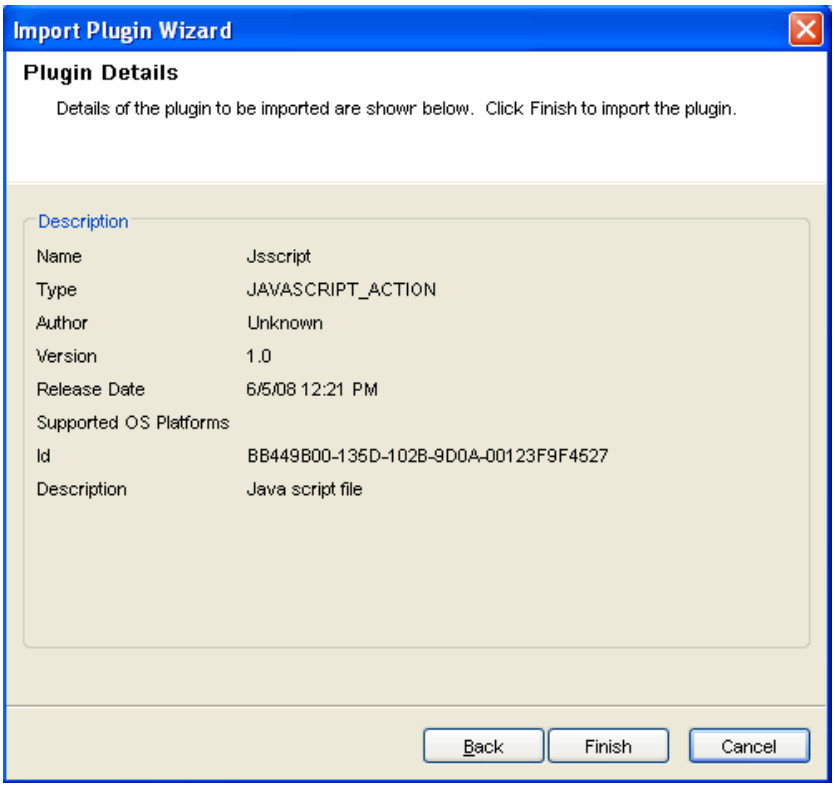

The Plugin Details window displays. Details of the plug-ins to be imported are displayed.

11 Click *Finish*.

If the directory from which the JavaScript file is imported contains a `package.xml` file, the system updates the `package.xml` file with the information defined in the wizard. If no `package.xml` file exists in the directory, a new `package.xml` file is automatically created.

An action plug-in is also created from the JavaScript file. The `package.xml` file is zipped as part of the JavaScript plug-in along with other files in the specified directory.

NOTE: When a plug-in is created from a directory, the original contents of the directory are stored in a backup `.zip` file located on the same directory level as the directory being zipped. The name of the backup file is in the format `<Directory Name>_<Randomly Generated Number>_bak.zip` where `<Directory Name>` is the directory in which the plug-in is created.

The following is the example of `package.xml` file:

```
<?xml version="1.0" encoding="UTF-8"?>
<JavaScriptActionPackage>
 <ID>FA6944D0-DC43-102A-976F-001321B5C0B3</ID>
 <Name>Example JavaScript Plugin</Name>
 <Type>JAVASCRIPT_ACTION</Type>
 <DisplayName>Example JavaScript Plugin</DisplayName>
 <Author>Novell Engineering</Author>
 <Version>61r1</Version>
 <ReleaseDate>1206414663439</ReleaseDate>
 <MainScriptFile>example.js</MainScriptFile>
 <Description>An example JavaScript Action plugin.</Description>
</JavaScriptActionPackage>
```

NOTE: When a plug-in is created from a JavaScript file and an existing `package.xml` file, the `package.xml` file is updated with the list of files contained in the package, hash codes, current dates, and so on.

17.3 Actions

There are many types of actions, many of which are intended only to be used with correlation rules. For more information about the correlation rule actions, see Chapter 4, "Correlation Tab," on page 83. This section focuses on JavaScript actions, which can be used in correlation rule deployments, within an incident, or in a right-click menu action.

17.3.1 Creating Actions

The Action Manager allows you to manage action instances, which are individual configurations of an action plug-in.

1. Click the Tools menu and select Action Manager.
2. Click the *Add* button located on the top left corner of the screen. The Configure Action window displays.

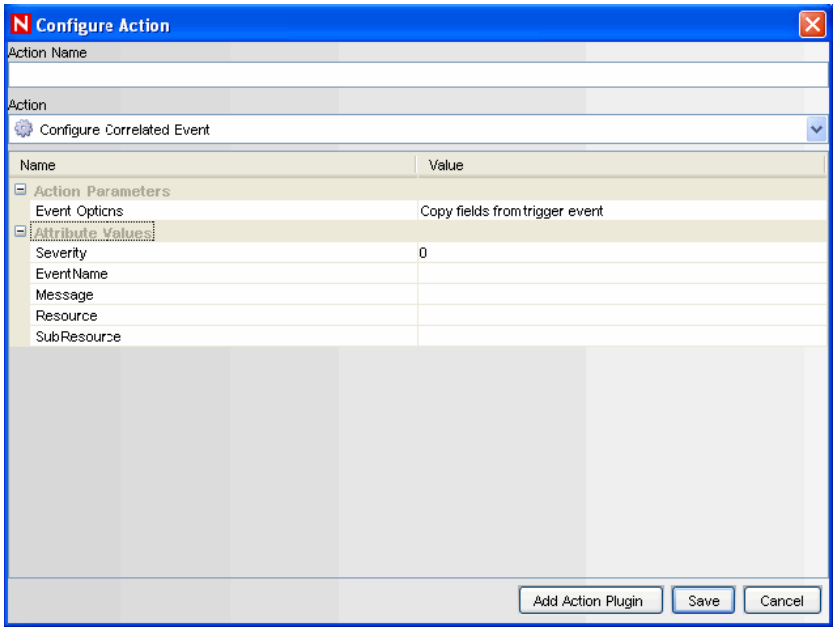

3 To create a JavaScript action, select an already imported JavaScript action plug-in from the available action types in the Action drop-down list.

   Alternatively, you can import another plug-in by clicking the *Add Action Plugin* button.

   If you select an action plug-in that is configured to use an Integrator to connect to an external system, the Add Integrator button displays.

   The parameters for the selected plug-in display. For actions provided by Novell, more information about configuration and the available parameters are available in the help file for the action.

4 Specify the attribute values for the type of action selected.

5 Click *Save*.

17.3.2 Editing Actions

If you edit an action that is associated with a deployed rule, the changes take effect the next time the correlation rule fires.

1 Click the Tools menu and select Action Manager.

2 Select an action and click the View or Edit link next to it. The Configure Action window displays.

3 Edit the options as required and click *Save*.

17.3.3 Deleting Actions

You cannot delete an action that is associated with a deployed correlation rule or Event Menu Configuration item.

To delete an action:

1 Click *Tools* > *Action Manager*.

2 Select an action and click *Delete*.

3 Click *Yes* to confirm.

17.3.4 Using JavaScript Actions

After an action instance is configured, it can be selected in one or more of the following locations:

- *Event Menu Configuration* on the *Admin* tab to create right-click menu actions
- *Actions* tab when deploying a correlation rule (to be executed when a correlation rule fires)
- Execute Incident Action within an incident (to be executed within an incident)

However, not all JavaScript actions are available in all contexts. The developer who creates the JavaScript action plug-in can define the required inputs for a JavaScript action, which determines what type of input it requires and in what contexts it can be used. For more information, see Table 17-1 on page 368. For more information on using these actions, see Chapter 4, “Correlation Tab,” on page 83, Chapter 5, “Incidents Tab,” on page 109, and Chapter 12, “Administration,” on page 235.

17.3.5 Developing JavaScript Actions

The information below is very basic development information about developing JavaScript actions. For more information, see Novell Developer Community web site (http://developer.novell.com/wiki/index.php?title=Develop_to_Sentinel).

- "Creating a JavaScript Action" on page 378
- "Debugging JavaScript Actions" on page 379

Creating a JavaScript Action

JavaScript actions take advantage of the functionality and flexibility of the JavaScript language and can be used to execute actions using Sentinel system methods to do things such as:

- Start/stop the Collectors
- Add/remove from dynamic lists
- Get a current event
- Get a correlated event
- Get a correlation event collection
- Get an incident
- Execute actions by using Integrators

The code sample below starts or stops a Collector based on the information in the correlated event.

```
importPackage(java.lang);
var CollectorName = "TC_5";
var evt = scriptEnv.getCurrentEvent();
var collNm = evt.getPort();
var outfile = new java.io.PrintWriter(new java.io.FileWriter("/opt/jaya/
strtcoll.txt", true));
if(collNm && collNm.equals(CollectorName))
{
   var collist = ESM.collectorsForName(collNm);
   if (collist.size() > 0)
   {
      var coll = collist.get(0);
      outfile.println("Stopping " + CollectorName);
      coll.stop();
      Thread.sleep(60000);
      outfile.println("starting " +CollectorName);
      coll.start();
   }
}
else
{
   outfile.println("JSTest collector does not exist");
}
outfile.close();
```

Debugging JavaScript Actions

You can debug JavaScript files from the Sentinel Control Center with the help of the JavaScript debugger. The JavaScript debugger is a local debugger that executes scripts with respect to the machine on which the Sentinel Control Center is running. The JavaScript debugger instantiates a debug session from the Data Access Service (DAS) machine.

A JavaScript Correlation action can only be debugged after it is associated with a fired correlation rule. Therefore, a prerequisite to debugging is to create a correlation rule that is guaranteed to fire, then associate the JavaScript correlation action with that rule.

The debugger has the following controls:

Table 17-2 *Debugger Controls*

Icon	Name	Description
	Run	Runs the script until the next breakpoint is encountered.
	Step Into	Steps into a function, one line at a time.
	Pause	Pauses the running script.
	Stop	Stops the script.
	Step Over	Steps over a function to the next line in the script.
	Step Out	Steps out of the function to the next line in the script.

To open a JavaScript Debugger:

1. Click Correlation on the menu bar and select Correlation Engine Manager. Alternatively, you can click the Correlation Engine Manager button on the toolbar.

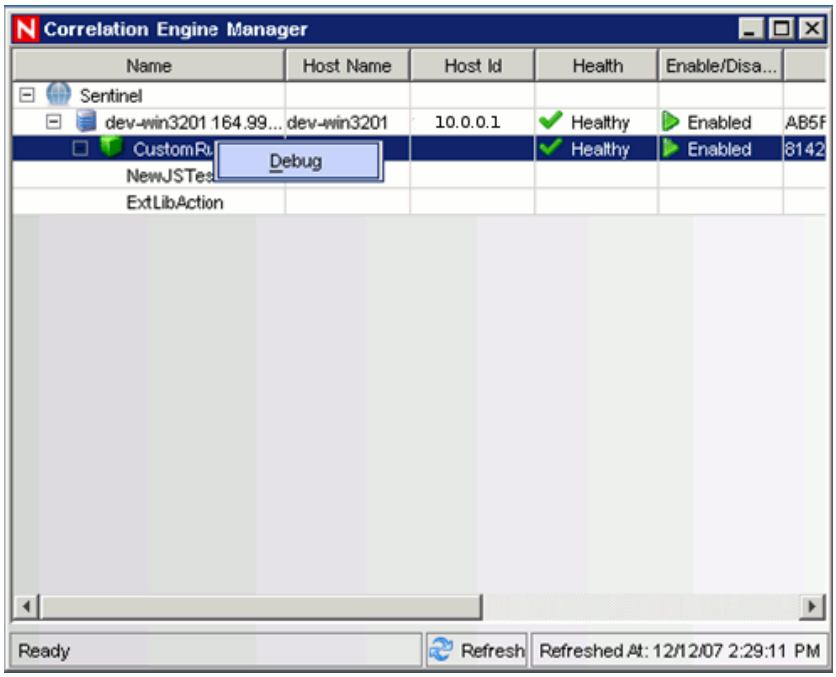

2 Right-click a JavaScript action associated with a correlation rule and select *Debug*. The Debug JavaScript Correlation Action window displays.

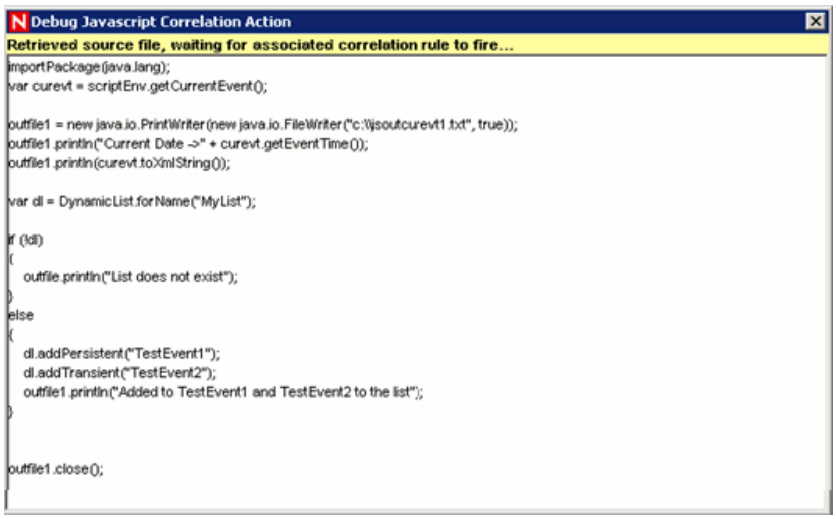

The screen displays the following message: `Retrieved source file, waiting for associated correlation rule to fire`.

The correlation rule must fire (and a correlated event or incident must be created) before you can debug the script. After the rule fires, this text panel is replaced by a debug panel and the actual debugging session begins. The following JavaScript Correlation Action window displays.

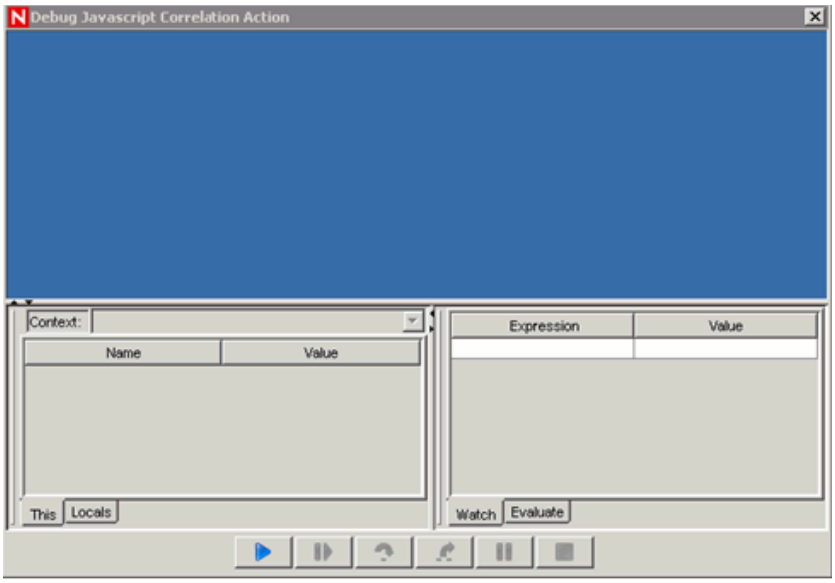

3 Click *Run*. The debugger panel displays the source code and positions the cursor on the first line of the script.

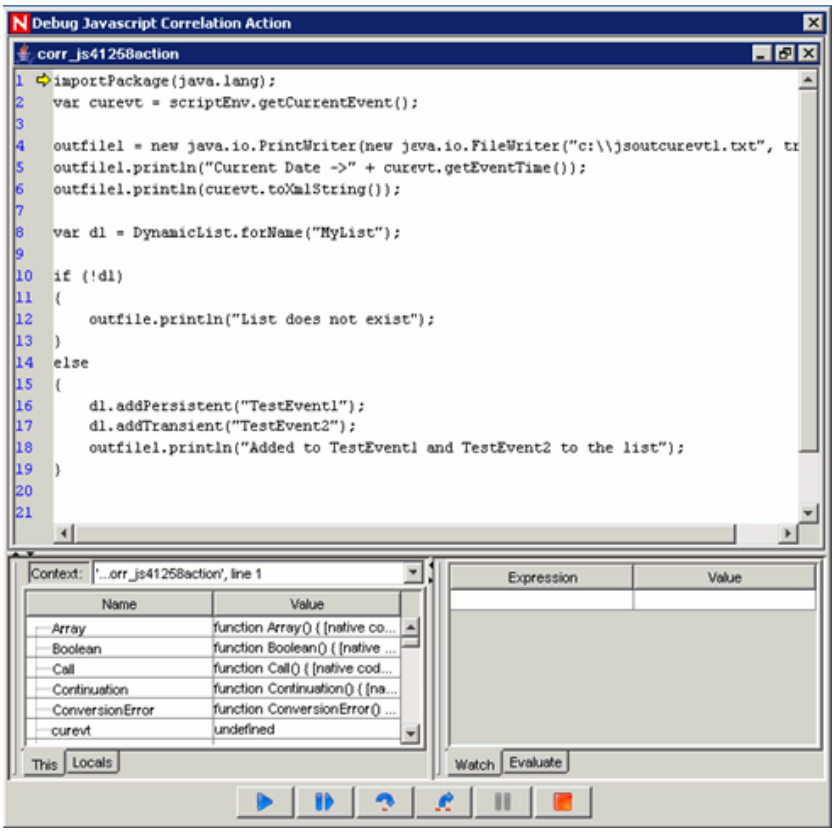

You can debug the script as many times as needed (without requiring a new correlation rule to fire). After the debugger gets to the end of the script (or after you click the Stop button), click Run again.

To debug the script by using a different rule, different correlated event, or different incident, close the Debug JavaScript Correlation Action window and repeat the debugging process.

17.4 Integrator Manager

Integrators are plug-ins that can be used in Sentinel Rapid Deployment to extend the features and functionality of Sentinel remediation actions. The Sentinel system is loaded with several Integrators by default, but you can download updates and additional Integrators from the Sentinel Plug-in Web site (http://support.novell.com/products/sentinel/secure/sentinelplugins.html).

Integrators allow Sentinel to connect to other external systems, for example, an LDAP server, SMTP server, or SOAP server. JavaScript actions can use Integrators to interact with other systems. For example, you can set the attribute in Novell eDirectory (an LDAP server) to enable or disable a user, edit details and so on. You could also start an Identity Manager workflow, such as a provisioning request, by using SOAP calls.

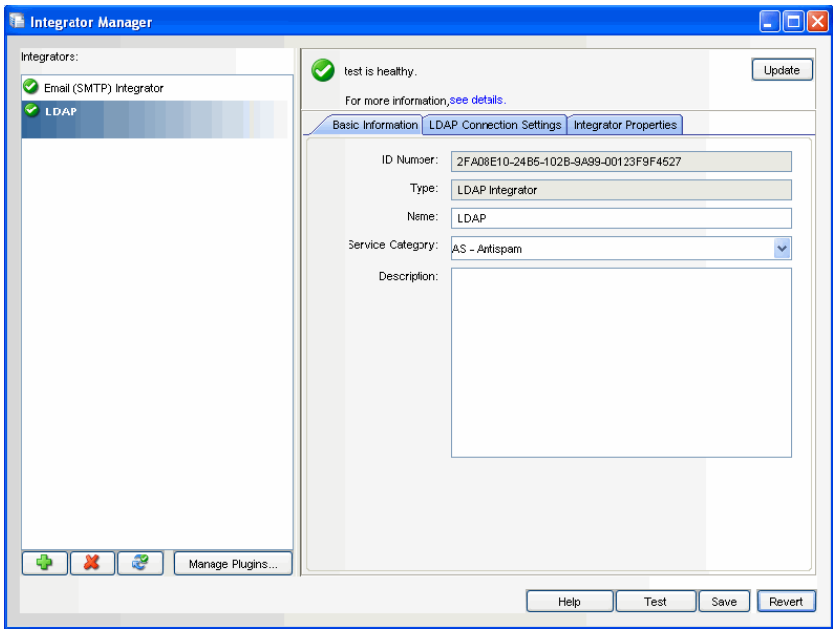

The general process for using an Integrator to perform remediation actions includes the following steps:

1. Determine the best type of Integrator to access the external system with which you want to interact.
2. Import and configure the appropriate Integrator to connect to the external system.
3. Write a JavaScript action to be executed through the Integrator. This script makes calls to methods specific to the Integrator in order to execute actions on the external system.
4. Import and configure the JavaScript action by using the Action Manager.

5 Perform additional configuration, if desired, to associate the action with a deployed correlation rule or an event menu action.

6 Execute the action when a correlation rule fires, manually by a user from the event menu, or from the Execute Incident Action menu in an incident.

For more information on specific Integrators, see the documentation that is available with the Integrators. You can download the updated Integrators from the Sentinel documentation Web site (http://www.novell.com/documentation/sentinel61). Alternatively, you can view the Integrators documentation by clicking the *Help* button in Integrator Manager after configuring the Integrator.

17.4.1 Permissions for Using Integrators

To use the Integrator Manager, a user must be assigned the necessary permissions in the User Manager. By default these permissions are not assigned to the users.

1 Log in to the Sentinel Control Center as a user with permissions to use the *User Manager*.

2 Go to the *Admin* tab.

3 Open the *User Configuration* folder.

4 Open the User Manager window. Double-click *User*. The User Details window displays.

5 Click the *Permissions* tab.

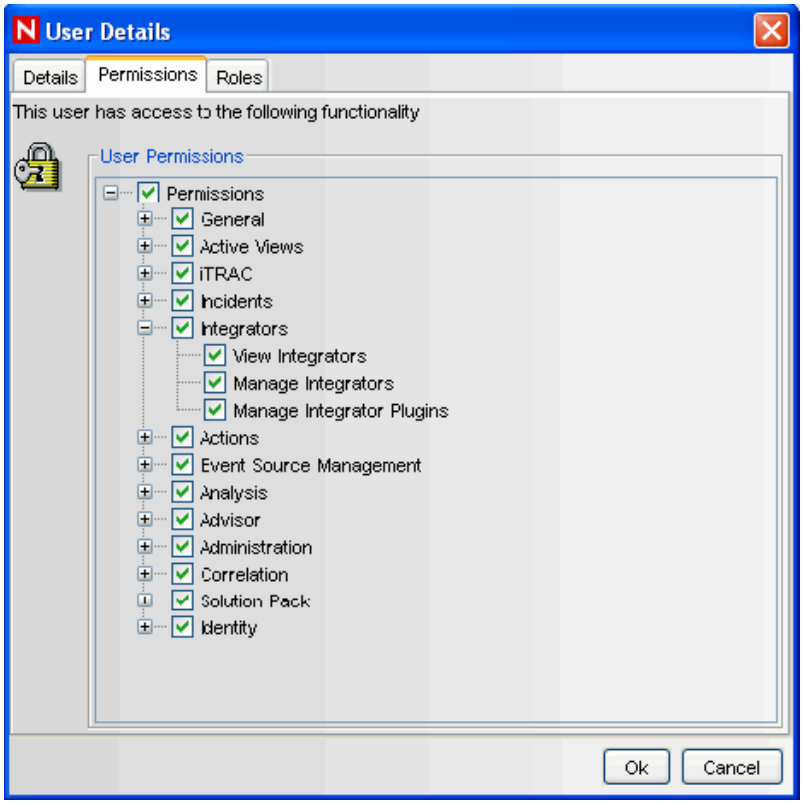

6 Select View Integrators, Manage Integrators, Manage Integrator Plugins, or Integrators (which automatically selects all child permissions).

The new permissions are applied the next time the user logs in. For more information, see "Sentinel 6.1 Rapid Deployment Control Center User Permissions" in the *Sentinel 6.1 Rapid Deployment Reference Guide*.

17.5 Integrator Plug-Ins

- Section 17.5.1, "Importing Integrator Plugins," on page 384
- Section 17.5.2, "Deleting Integrator Plug-Ins," on page 384

17.5.1 Importing Integrator Plugins

1 Click *Tools* > *Integrator Manager*. The Integrator Manager window displays.

2 Click the *Manage Plug-Ins* button.

The Integrator Plugin Manager window displays, where you can add, delete, refresh, view Integration plug-in details, configure Integrators and add auxiliary files.

3 Click the*Import* icon in the Integrator Plugin Manager window. The Plugin Import Type window displays.

4 Select *Import an Integrator plugin file (.zip)*.

5 Click *Next*. The Choose Plugin Package File window displays

6 Use the *Browse* button to locate an Integrator file to import to the plugin repository.

7 Select a zip file and Click *Open*.

If you have selected an Integrator file that already exists, the Replace Existing Plugin window displays.

8 Click *Next* if you want to replace the existing plug-ins.

9 Click *Next*. The Plugin Detail window displays.

The details of the plug-in to be imported are displayed.

10 Select the *Launch Integrator Configuration Wizard* check-box if you want to deploy the plug-in after importing the Integrator Plug-in.

11 Click *Finish*.

17.5.2 Deleting Integrator Plug-Ins

1 Click *Tools* > *Integrator Manager*. The Integrator Manager window displays.

2 Click the `Manage Plug-Ins` button. The Integrator Plugin Manager window displays.

3 Select an *Integrator Plug-in* and click the *Delete* icon. A confirmation message displays.

4 Click *Yes*.

NOTE: You can delete an Integrator plug-in only if there are no Integrators configured to use it.

17.6 Integrators

- Section 17.6.1, “Creating an Integrator Instance,” on page 385
- Section 17.6.2, “Editing an Integrator Instance,” on page 385
- Section 17.6.3, “Deleting an Integrator Instance,” on page 385
- Section 17.6.4, “Integrator Connection Status,” on page 385
- Section 17.6.5, “Viewing Integrator Health Details,” on page 386
- Section 17.6.6, “Integrator Events Query,” on page 387
- Section 17.6.7, “Using Integrators from Actions,” on page 389

17.6.1 Creating an Integrator Instance

An Integrator is a configured instance of an Integrator plug-in. There can be one or more Integrator instances with different parameters or settings using an Integrator plug-in.

The specific steps to configure an Integrator instance depend on the type of Integrator, and those steps are described in detail in documents that come with the Integrators. Documentation for installed plug-ins can be viewed by selecting an Integrator in the Integrator Manager and clicking *Help*.

17.6.2 Editing an Integrator Instance

1. Click *Tools* > *Integrator Manager*. The Integrator Manager window displays.
2. Select an Integrator from the left panel. Edit the Integrator instance information by using the Basic Information, Connection Settings, and the *Integrator Properties* tab.
3. Click *Save* after you edit the information.

17.6.3 Deleting an Integrator Instance

An Integrator instance cannot be deleted if it is currently associated with an action. To delete an Integrator instance, you must first delete or modify any actions that are associated with it.

1. Click *Tools* > *Integrator Manager*. The Integrator Manager window displays.
2. Select an Integrator from the left panel, then click the *Delete* icon to delete an Integrator instance.

17.6.4 Integrator Connection Status

1. Click *Tools* > *Integrator Manager*. The Integrator Manager window displays.
2. Click the *Refresh health of all Integrators* button. The Integrator Connection Status window displays.

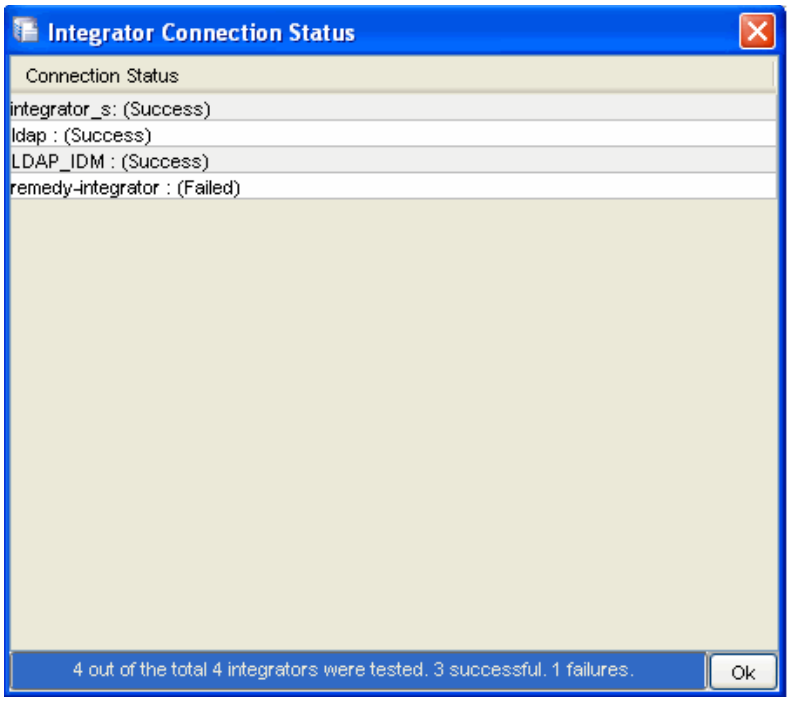

The server performs a test of the Integrators in the actual service where the Integrators are used when actions are executed.

3 Click *OK*.

17.6.5 Viewing Integrator Health Details

1 Click *Tools* > *Integrator Manager*. The Integrator Manager window displays.

2 Select an Integrator from the left pane.

3 Click *See Details*. The Refresh Health Information window displays.

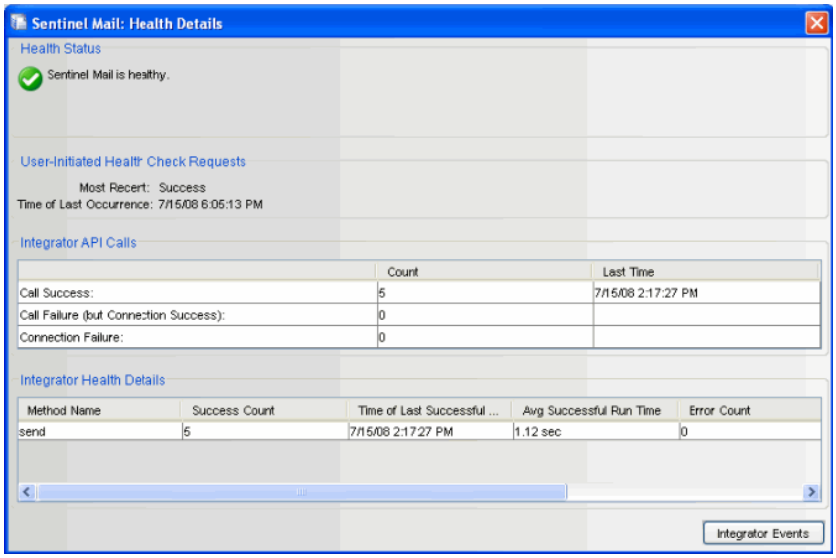

The Health screen displays the Refresh Health State, Time of last occurrence, its method calls, and the related events of the selected Integrator configuration.

- **Integrator API Calls:** Indicates the status of count and time of both the connection and the method calls used from the API of the selected Integrator. For more information on JavaScript plug-ins, see Section 17.1, "Action Manager," on page 363.
  - **Call Success Count:** Displays the count for the number of times the connection was established successfully and the methods were called successfully from the API. T*ime of Last Occurrence* displays the time when the connection and the method call were successful.
  - **Call Failure (but Connection Success) Count:** Displays the count for the number of times the connection was established successfully but the methods call failed. *Time of Last Occurrence* displays the last time when the connection was successful and the method call failed.
  - **Connection Failure Count:** Displays the count for the number of times the connection failed. `Time of Last Occurrence` displays the last time when the connection and method call failed.

  NOTE: The most recent time among *Connection Success* and *Call Success Count*, *Connection Success* and *Call Failure Count*, and *Connection Failure* and *Call Failure Count* is reflected in the overall health status for the configured Integrator.

- **Integrator Health Details:** The health details are displayed in Integrator Health Details pane. It provides information about the success of the API methods called in the JavaScript action files associated with the Integrator. It provides information specific to the methods called:
  - **Method Name:** Name of the API method used in the JavaScript.

    Success Count: Number of times the API method executed successfully.

    Time of Last Successful Call: The time at which the method was last successfully executed.

    Average Successful Run Time: Average time to make a successful method call.

    Error Count: Number of times the API method failed.

    Time of Last Error Call: The time at which the method call failed.

    Average Error Run Time: Average time to make a failed method call.

  NOTE: The most recent time among *Time of Last Successful Call* and *Time of Last Error Call* is reflected in the overall health status of the method.

17.6.6 Integrator Events Query

When an Integrator faces connection failures, it generates internal audit events. If you want to query these events, you can use an integrator events query to automatically create a filter for the selected Integrator and process a query.

1. Click T*ools* > *Integrator Manager*. The Integrator Manager window displays.
2. Click *See Details*. The Refresh Health Information window displays.

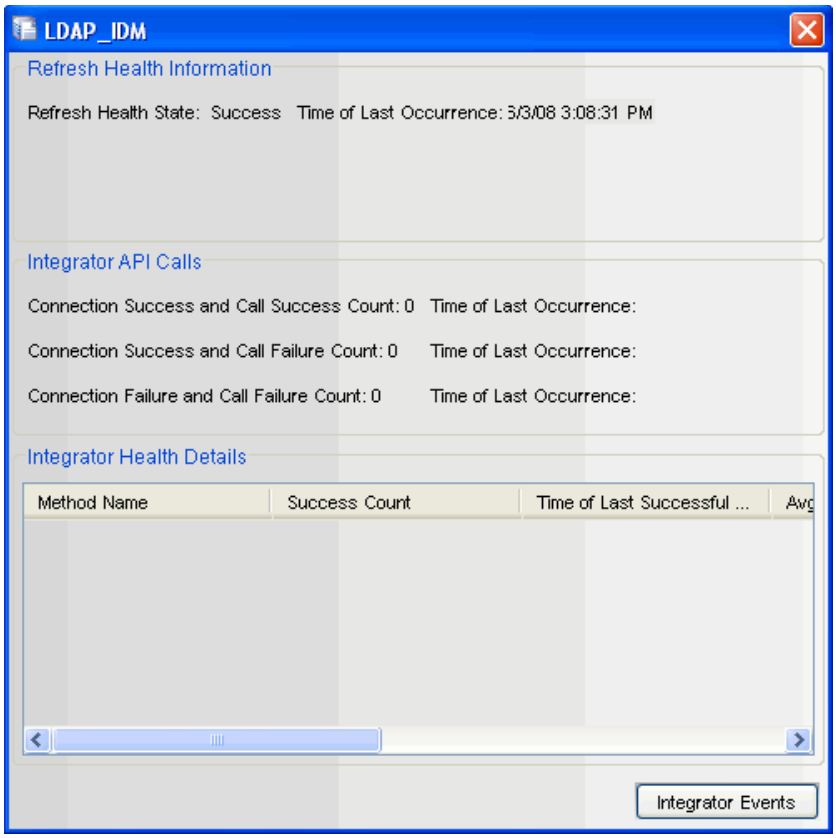

3 Click the *Integrator Events* button. The Query window displays.

All the events related to the configured Integrator automatically display in the Query window. You can filter the displayed events by using the filter criteria. For more information, see Section 3.9.3, “Historical Event Query,” on page 67.

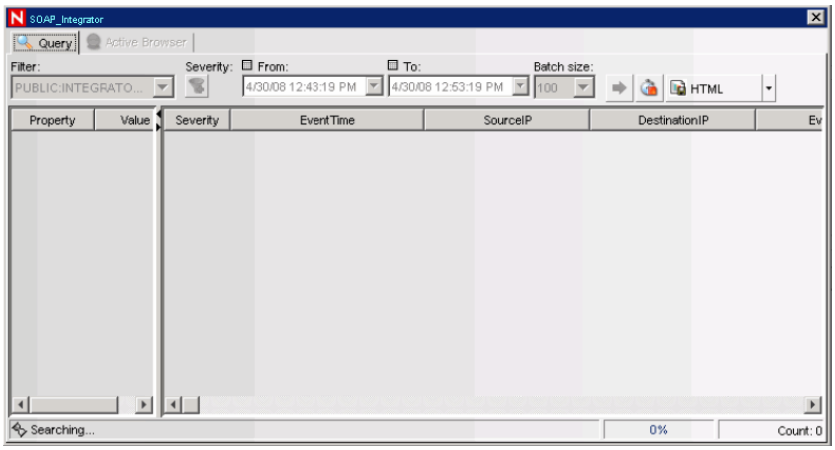

17.6.7 Using Integrators from Actions

Some actions might require an Integrator in order to make a connection to an external system. You can write or customize JavaScript code that connects to an external system by using the Integrator and executes methods appropriate for the external system. Because all the connection and other configuration information is already configured as part of the Integrator, the code only needs to perform a task on the system with which it integrates.

When writing code that needs to access an Integrator, you must determine how to locate a specific Integrator. You can locate an Integrator in the following ways:

- Look up an Integrator by its name
- Look up an Integrator by its ID.
- Look up a set of Integrators by their service category
- Retrieve a set of Integrators that have a specific property name or value
- Retrieve all Integrators and iterate through them to find the required one based on custom logic

After you retrieve the Integrator, you can access the API for the external system to make programmatic calls to achieve the required integration.

Identity Integration

18

Sentinel Rapid Deployment provides an integration framework for identity management systems. This integration provides functionality on several levels:

- Identity Browser provides the ability to look up the following information about a user:
  - Contact information
  - Accounts associated with that user
  - Most recent authentication events
  - Most recent access events
  - Most recent permissions changes
- Identity Browser lookup from events
- Reports and correlation rules provide an integrated view of a user's true identity, even across multiple systems on which that user has separate accounts. For example, accounts like NOVELL\testuser; > cn=testuser,ou=engineering,o=novell, and TUser@novell.com can be mapped to the actual person who owns the accounts.

By displaying information about the people initiating a given action or people affected by an action, incident response times are improved and behavior-based analysis is enabled.

Novell provides an optional integration with Novell Identity Manager. The screenshots and descriptions in this section are based on Novell Identity Manager.

Sentinel Rapid Deployment synchronizes identity information with major identity management systems and stores local copies of key information about each Identity. The following table summarizes the commonly used information provided:

Table 18-1 *Identity Information*

Name	Description
AccountGUID	Auto-generated internal ID.
Name	Username that references the account, generally provided by the user to log in.
ID	The numeric or other identifier that represents the account in the event source. This ID is used for resolution when the username is not available.
Authority	The realm within which this account is unique. Collectors calculate the realm based on event information.
Status	The status of the account.
IdentityGUID	A reference to the identity that owns this account.

The Identities stored by Sentinel are then linked with accounts created on endpoint systems by the identity management system. This helps Sentinel associate the correct identity information with the native events from those endpoint system. Some identity information is injected directly into the inbound event by using the mapping service. The remaining identity information, such as photograph and contact information, is accessible through the Identity Browser.

Figure 18-1 *Accessing the Identity Browser*

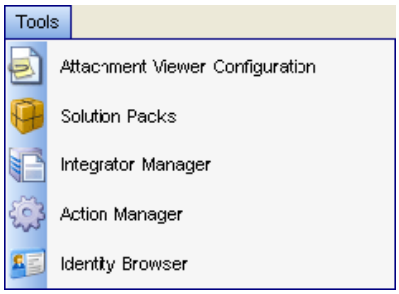

The identity information injected into the event can be used for correlation and for performing actions on the identities that are associated with detected activity. For example, Sentinel is able to see multiple failed logins from a given person and not just an account. A detected violation could trigger disabling activities for all accounts associated with an identity.

Figure 18-2 *Identity Details*

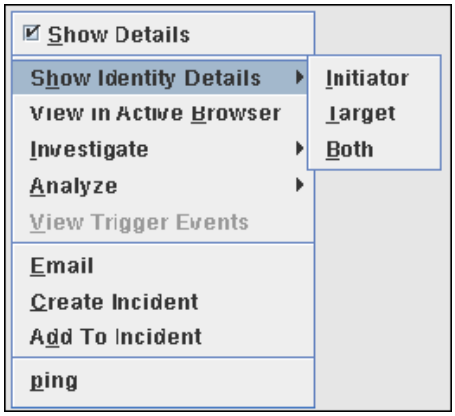

18.1 Integration with Novell Identity Manager

Integration with Novell Identity Manager is available as part of the Novell Compliance Management Platform 1.0.1 and Novell Compliance Management extension for SAP environments 1.0.1, which includes the following components:

- Sentinel Rapid Deployment
- eDirectory 8.8.5
- Identity Manager 3.6.1
- Access Manager 3.1
- Identity Tracking Solution Pack 6.1r3
- Analyzer for Identity Manager 1.1

- Identity Manager Resource Kit 1.2
- Identity Manager Driver for Sentinel 3.6

For more information, see Novell Compliance Management Platform (http://www.novell.com/documentation/ncmp10/) and Novell Compliance Management Platform extension for SAP environments 1.0 (http://www.novell.com/documentation/ncmp_sap10/).

The Solution also requires identity-enabled Collectors, which are available for download at the Standard Sentinel Content download Web site (http://support.novell.com/products/sentinel/sentinel61.html).

After Sentinel and Identity Manager are installed, the Sentinel Driver for Identity Manager sends identity and account information from the Identity Vault to the Sentinel Identity Vault Collector, which populates the Sentinel database. The information is inserted into two new tables in Sentinel Rapid Deployment. These two tables are the Identity table (USR_IDENTITY) and the Account table (USR_ACCOUNT). For more information, see "Sentinel 6.1 Rapid Deployment Database Views for PostgreSQL" in the *Sentinel 6.1 Rapid Deployment Reference Guide*.

The time required to initially populate the Sentinel database depends on the amount of data in the Identity Vault; identity information including photographs requires significantly more time to load.

The Sentinel Driver for Identity Manager and Identity Vault Collector also keep the identity information synchronized as information is updated in the Identity Vault during normal Identity Manager operations.

After the identity information and account information are loaded in their respective tables with a link between them, a map named `IdentityAccountMap` is generated automatically in the location `<install_directory>/data/map_data`. The map contains the following information:

- Account Name
- Authority
- Customer Name
- Identity GUID
- Full Name
- Department
- Job Title
- Manager GUID
- Account Status

IMPORTANT: An identity can have multiple accounts but one account cannot be assigned to multiple identities.

The identity map is automatically applied to all events from Collectors to look for an identical match between the information in the event and key fields in the map. The table below shows the fields that are populated if all of the map key fields and event data exactly match. These mappings are automatically configured and are not editable.

Label	Populated by which Column from IdentityAccount Map	Map Key Field : Event Label
InitUserDepartment	Department	Account Name : InitUserName
		Authority : InitUserDomain
		Customer Name : MSSPCustomerName
InitUserFullName	Full Name	Account Name : InitUserName
		Authority : InitUserDomain
		Customer Name : MSSPCustomerName
InitUserIdentity	Identity GUID	Account Name : InitUserName
		Authority : InitUserDomain
		Customer Name : MSSPCustomerName
TargetUserDepartment	Department	Account Name : TargetUserName
		Authority : TargetUserDomain
		Customer Name : MSSPCustomerName
TargetUserFullName	Full Name	Account Name : TargetUserName
		Authority : TargetUserDomain
		Customer Name : MSSPCustomerName
TargetUserIdentity	Identity GUID	Account Name : TargetUserName
		Authority : TargetUserDomain
		Customer Name : MSSPCustomerName

NOTE: To find a match, the event fields and map key fields must match exactly. This might require modifications to existing Collectors to “identity enable” them to parse or concatenate data to make these fields match the data from the Identity Vault.

Once added to the event by the mapping service, these fields are used by correlation rules, remediation actions, and reports in the Identity Tracking Solution Pack. In addition to using the content included in the Solution Pack, users can also perform the following actions:

- Create correlation rules based on identity in addition to account name. This allows you to look for similar events from a single user, which provides a more comprehensive view than looking at events from a single account
- Create reports that show identity, including all accounts associated with a user
- Use the Identity Browser to get more information about users and their activity

NOTE: For other identity systems, similar integration can be achieved by writing an identity synchronization Collector that uses the Identity API.

18.2 Identity Browser

The Identity Browser in Sentinel allows you to search and view user profiles of the identities in the Sentinel database that have been synchronized from the identity management system. In addition to information from the identity management system, the Identity Browser also shows recent activity for the user that has been collected using the Sentinel Collectors.

To open the Identity Browser:

1. Click the *Tools* menu and select *Identity Browser*. The Identity Browser window displays.

   Alternatively, you can launch the Identity Browser through the icon that appears when you launch the Sentinel Control Center.

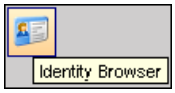

18.2.1 Searching Profiles

1. Click the *Tools* menu and select *Identity Browser*. The Identity Browser window displays.

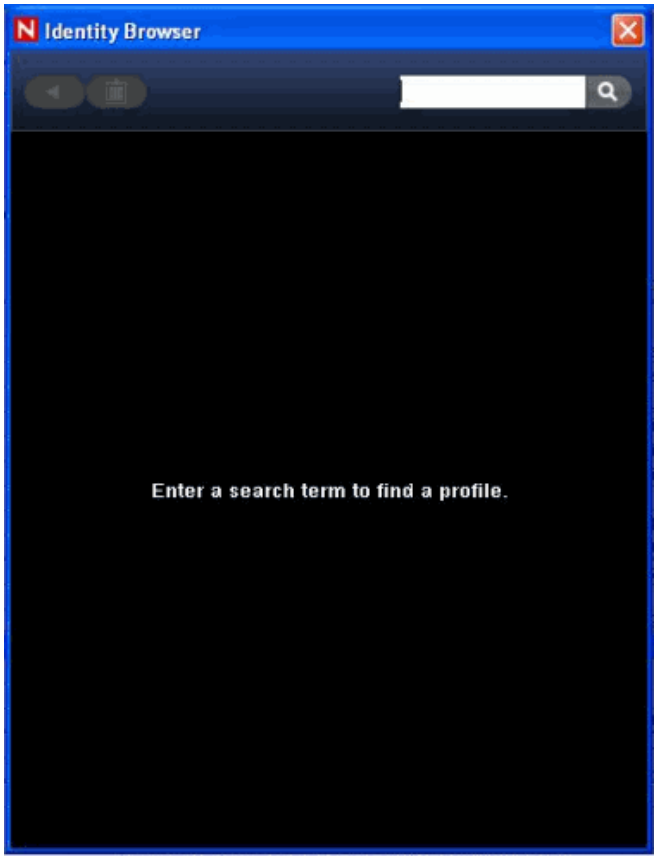

2. Enter the first name or last name or first character of either name for the profile in the Search box.

TIP: You can input letters to view all the identities whose first or last name starts with the letters. For example, if the user enters the letters “a,b” the names Abraham, Abdullah and so on are matched.

If the search is broad, the results show the first 100 names with a Load <x> More Records button, where <x> depends on the number of records remaining and can be up to 100.

3 Click Search Icon. The searched profile displays:

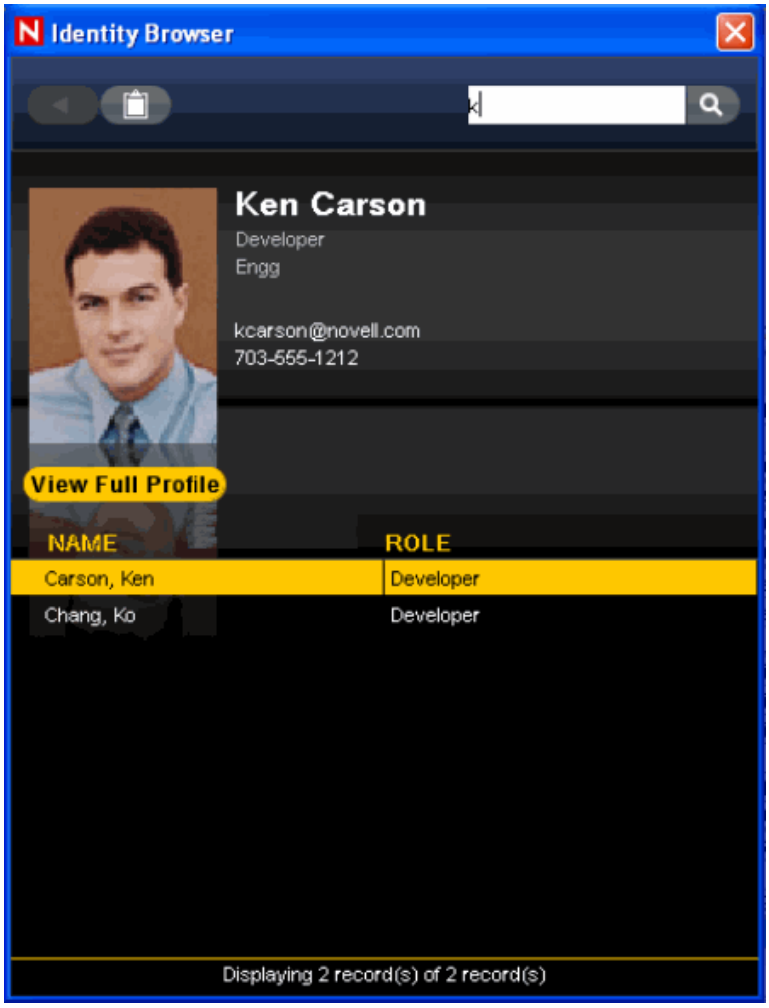

4 Select a user and click View Full Profile to see more information.

Alternatively, you can right-click a user name (identity) and select Open New Window. It opens a new Identity Browser window. It is similar to the parent Identity Browser window and you view the full profile in a new window.

5 Use the back arrow icon to navigate to the previous profile.

18.2.2 Viewing Profile Details

1 Click T*ools* > *Identity Browser*. The Identity Browser window displays.

2 Type the first name or first character of the profile in the Search box. Click the *Search* icon. The searched profile displays.

3 Click the *View Full Profile* button. The user profile displays:

Using the view profile window, you can view User Profile, Accounts, and Recent Activities performed by the user. By default, the User Profile displays.

4 Select Accounts. The details of the account are displayed:

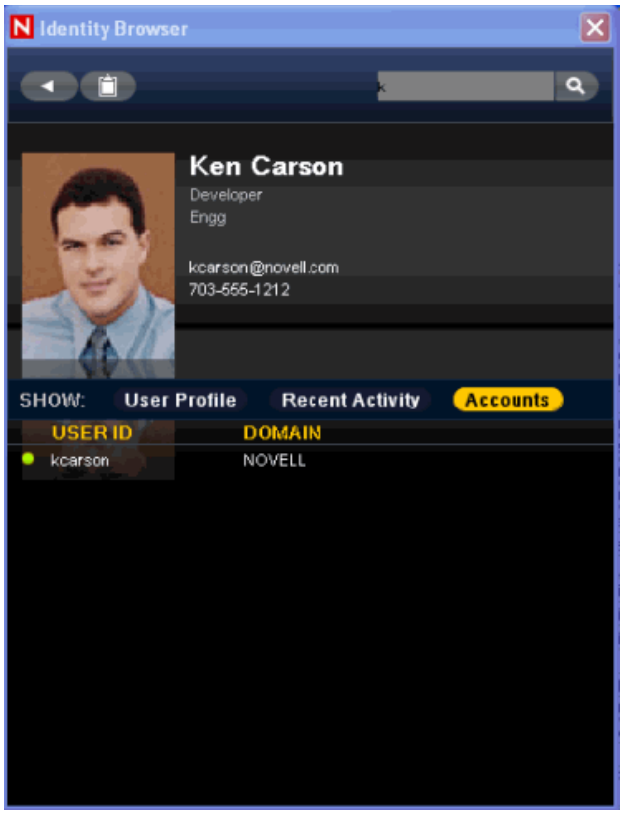

You can access Accounts in Active View by right-clicking an event generated by the Identity Collector and by selecting the *Show Identity Details* option. Select the Initiator, Target, or Both option. The account details of the associated Identity in that event displays in a pop-up window.

5 Select Recent Activity.

The contextual event information such as Authentication, Access, and Permission change events for that identity are displayed. The events displayed are limited to last 10 events in each category as shown below:

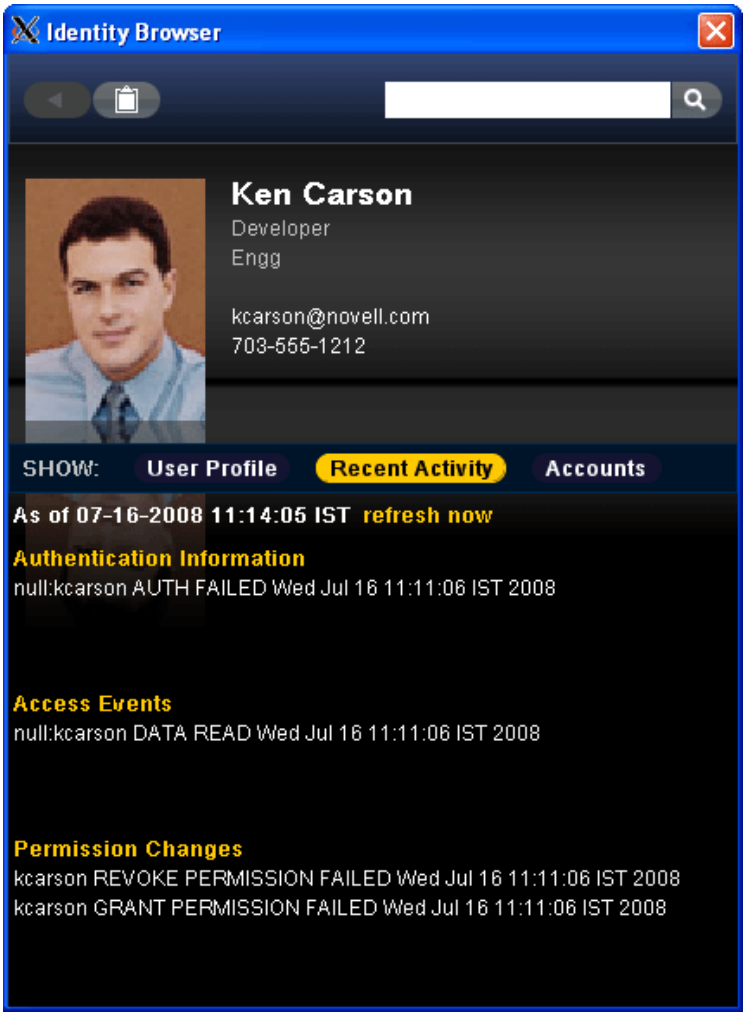

18.2.3 Using the Clipboard Functionality

You can use the clipboard functionality to copy the data of *User Profile*, *Recent Activity*, or the *Account* tabs. On any of the three tabs, you can click the clipboard option and copy the information to the clipboard. For example, if you are viewing the Recent Activity tab and click the clipboard, the recent activity data is copied to the clipboard. You can then choose to paste the information on the clipboard to a Notepad and save.

You can also copy the information to the clipboard after you have visited the tabs. For example, you can visit the *User Profile* and *Recent Activity* tabs, then use the clipboard to copy the data. Similarly, you can visit all the three tabs and then use the clipboard to copy the data.

18.3 Reports

Sentinel Rapid Deployment reports include identity information. If identity management integration is configured, this information appears on the reports. For example, see the report below.

Figure 18-3 *Reports*

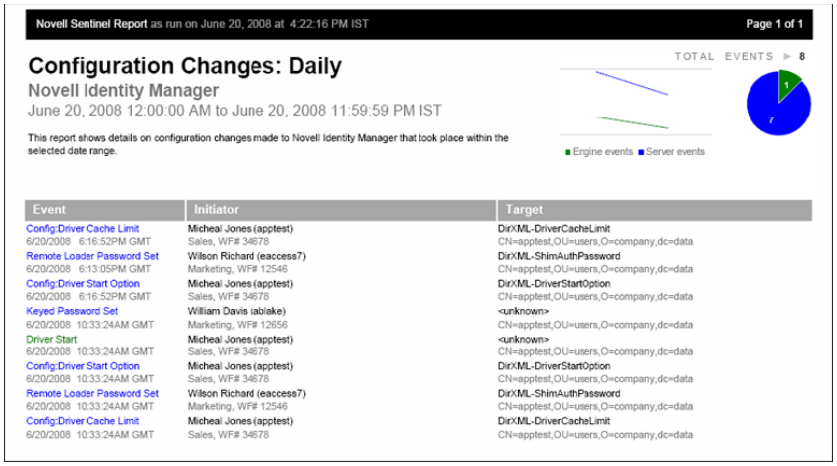

Novell Sentinel Report as run on June 20, 2008 at 4:22:16 PM IST — **Page 1 of 1**

Configuration Changes: Daily

Novell Identity Manager

June 20, 2008 12:00:00 AM to June 20, 2008 11:59:59 PM IST

This report shows details on configuration changes made to Novell Identity Manager that took place within the selected date range.

TOTAL EVENTS ▶ 8

Engine events ■ Server events

Event	Initiator	Target
Config:Driver Cache Limit 6/20/2008 6:16:52PM GMT	Micheal Jones (apptest) Sales, WF# 34678	DirXML-DriverCacheLimit CN=apptest,OU=users,O=company,dc=data
Remote Loader Password Set 6/20/2008 6:13:05PM GMT	Wilson Richard (eaccess7) Marketing, WF# 12546	DirXML-ShimAuthPassword CN=apptest,OU=users,O=company,dc=data
Config:Driver Start Option 6/20/2008 6:16:52PM GMT	Micheal Jones (apptest) Sales, WF# 34678	DirXML-DriverStartOption CN=apptest,OU=users,O=company,dc=data
Keyed Password Set 6/20/2008 10:33:24AM GMT	William Davis (ablake) Marketing, WF# 12656	<unknown> CN=apptest,OU=users,O=company,dc=data
Driver Start 6/20/2008 10:33:24AM GMT	Micheal Jones (apptest) Sales, WF# 34678	<unknown> CN=apptest,OU=users,O=company,dc=data
Config:Driver Start Option 6/20/2008 10:33:24AM GMT	Micheal Jones (apptest) Sales, WF# 34678	DirXML-DriverStartOption CN=apptest,OU=users,O=company,dc=data
Remote Loader Password Set 6/20/2008 10:33:24AM GMT	Wilson Richard (eaccess7) Marketing, WF# 12546	DirXML-ShimAuthPassword CN=apptest,OU=users,O=company,dc=data
Config:Driver Cache Limit 6/20/2008 10:33:24AM GMT	Micheal Jones (apptest) Sales, WF# 34678	DirXML-DriverCacheLimit CN=apptest,OU=users,O=company,dc=data

Sentinel Rapid Deployment Architecture

A

Sentinel Rapid Deployment is a simplified version and an alternate platform for Novell Sentinel that provides security information and an event management (SIEM) solution that automates the collection, analysis, and reporting of system network, application, and security logs to help organizations manage IT risks.

Sentinel Rapid Deployment provides full Sentinel functionality in a single-box SUSE Linux package. It features an easy-to-install SIEM solution that uses open source components such as PostgreSQL, ActiveMQ, and JasperReports for the database, messaging, and reporting.

This section discusses the functional and technical architecture of Sentinel.

- Section A.1, “Sentinel Rapid Deployment Features,” on page 401
- Section A.2, “Functional Architecture,” on page 401
- Section A.3, “Architecture Overview,” on page 403
- Section A.4, “Logical Architecture,” on page 414

A.1 Sentinel Rapid Deployment Features

Sentinel allows you to monitor and manage a variety of functions. Some of the main functions include:

- Real-time views of large streams of events
- Reporting capabilities based on real-time and historical events, through the Web interface
- Managing users and what they are able to see and do by permission assignment
- Managing access to events for different users
- Organizing events into incidents for efficient response management and tracking
- Detecting patterns in events and streams of events
- An intuitive and flexible rule-based language for correlation
- Rules compiled for high performance
- Embedded Sentinel database, based on the open source PostgreSQL database engine
- Web-based search tool to quickly search for strings and patterns within the Sentinel event database
- Web-based client application launch and installation

Sentinel processes communicate with each other through message-oriented middleware (MOM).

A.2 Functional Architecture

Sentinel Rapid Deployment is composed of the following component subsystems, which form the core of the functional architecture:

Table A-1 Sentinel Rapid Deployment Components

Components	Description
Sentinel Rapid Deployment Server	The Sentinel Rapid Deployment server runs the core back-end components of the software. There are a number of subcomponents that performs the key functions. ◆ **ActiveMQ Message Bus:** The JMS-based message bus over which the other components communicate with each other. ◆ **Data Access Services (DAS):** Data storage, query, display, and processing components. ◆ **Correlation Engine:** Performs real-time event analysis. ◆ **iTRAC:** A role-based incident-response workflow engine. ◆ **Jasper Reporting Engine:** Open source reporting engine.
Event Source Management (ESM)	An extensible framework built to manage and monitor connections between Sentinel and third-party event sources, by using Sentinel Connectors and Sentinel Collectors. In addition to ESM, there are a number of subcomponents that are hosted by a distributable service called the Collector Manager. This service can be installed on a number of systems to balance the processing load or for scalability. The data collection components are downloaded from the Novell Sentinel Content page and are installed to the Collector Managers via a central ESM interface.
Event Source	An event source can be a device, an operating system, a database, or an application. The actual event sources are represented in ESM and can be configured with certain meta information.
Connector	Connectors perform protocol-based communications with the event source. For example, over JDBC, Syslog, WMI, file reads, etc.
Collectors	Collectors are used to parse data from a specific event source and normalize the data into Sentinel's standard event schema.
Advisor	A key vulnerability or attack information service that helps you enhance your security posture. For example, the Exploit Detection feature of Advisor reduces false positives from intrusion detection systems.
Solution Packs	The Solution Pack framework provides the ability to group various types of content, such as reports, rules, data enrichment, remediation actions, and workflows. The content is grouped into a familiar control framework. Solution Packs can be built around specific business issues like PCI compliance, and partners can extend and customize them for industry-specific solutions.
User Applications	Sentinel includes the following three key user applications: ◆ Sentinel Control Center (SCC) An SCC interface includes the Event Source Management and Solution Manager interfaces. ◆ Solution Designer that creates Solution Packs. ◆ Sentinel Database Manager

Components	Description
Collector Builder	The Collector Builder helps you develop new Collectors from scratch by using the proprietary language. It is similar to an IDE. Sentinel Rapid Deployment provides the ability to develop Collectors in Java Script by using the third-party tools like Eclipse.
PostgreSQL Server	Sentinel requires a back-end database component to store the data. Sentinel Rapid Deployment uses a PostgreSQL database that is installed with Sentinel Rapid Deployment installation. The database can be used with all the required schema.
Tomcat Server	For generating reports and event search features on Web UI. It provides Sentinel Applications to launch and install through the Web interface.

A.3 Architecture Overview

The Sentinel Rapid Deployment system is responsible for receiving events from the Collector Manager. The events are then displayed in real-time in an Active View and logged into a database for historical analysis.

At a high level, the Sentinel system uses a PostgreSQL database and is comprised of Sentinel processes and a reporting engine. The system accepts events from the Collector Manager as its input. The Collector Manager interfaces with third-party products and normalizes the data from these products. The normalized data is then sent to the Sentinel processes and database.

Historical analysis and reporting can be done by using the Sentinel integrated JasperReports reporting engine. The reporting engine extracts data from the database and integrates the report displays in the Web interface by using HTML documents over an HTTP connection.

Figure A-1 *Sentinel Architecture*

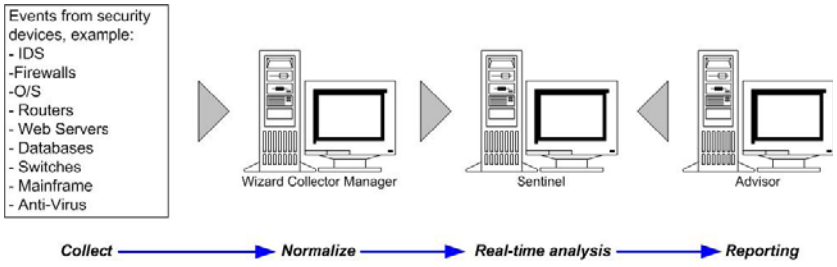

- Section A.3.1, “Communication Server,” on page 404
- Section A.3.2, “Sentinel Events,” on page 405
- Section A.3.3, “Event Source Management,” on page 409
- Section A.3.4, “Application Integration,” on page 410
- Section A.3.5, “Time,” on page 410
- Section A.3.6, “System Events,” on page 411
- Section A.3.7, “Processes,” on page 412

A.3.1 Communication Server

Sentinel Rapid Deployment's Apache ActiveMQ is an open source message broker. The architecture is built around the Java Message Oriented Middleware (JMOM), which supports asynchronous calls between the client and server applications. Message queues provide temporary storage when the destination program is busy or not connected. MOM reduces the complexity of the master-slave nature of the client/server mechanism.

ActiveMQ Message Bus

The ActiveMQ message bus allows for independent scaling of individual components while also allowing for standards-based integration with external applications. The key to scalability is that, unlike other distributed software, no two peer components communicate with each other directly. All components communicate through the message bus, which is capable of moving thousands of message packets per second.

Leveraging the unique features of the message bus, the high-throughput communication channel can maximize and sustain a high data throughput rate across the independent components of the system. Events are compressed and encrypted on the wire for secure and efficient delivery from the edge of the network or collection points to the hub of the system, where real-time analytics are performed.

The ActiveMQ message bus employs a variety of queuing services that improve the reliability of the communication beyond the security and performance aspects of the platform. Using a variety of transient and durable queues, the system offers unparalleled reliability and fault tolerance. For instance, important messages in transit are saved (by being queued) in case of a failure in the communication path. The queued message is delivered to the destination after the system recovers from the failure state.

Figure A-2 *Message Bus*

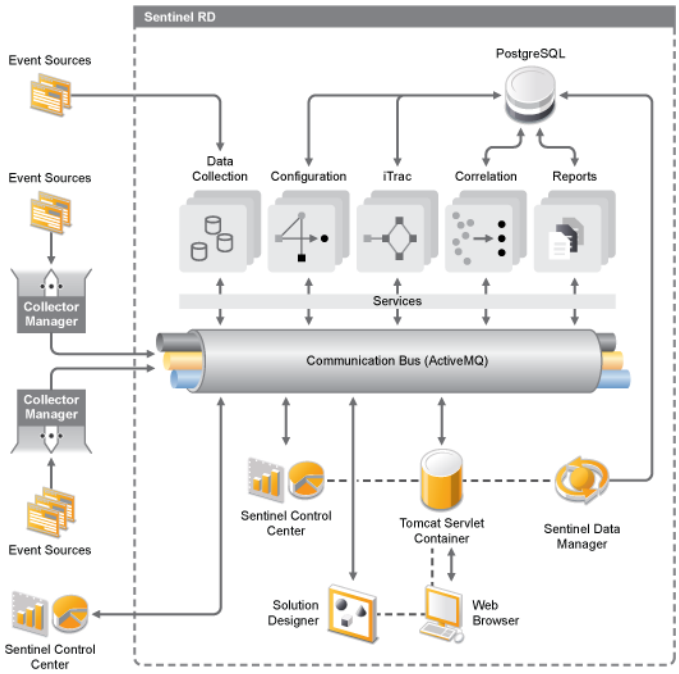

ActiveMQ takes advantage of an independent, multi-channel environment, which virtually eliminates contention and promotes parallel processing of events. These channels and sub channels work not only for event data transport but also offer fine-grained process control for scaling and load balancing the system under varying load conditions. Using independent service channels such as control channels and status channels, in addition to the main event channel, allows sophisticated and cost-effective scaling of event-driven architecture.

A.3.2 Sentinel Events

Sentinel receives information from devices, normalizes this information into a structure called a Sentinel event, and sends the event for processing. Events are processed by the real-time display, correlation engine, and the back-end server.

An event is made up of more than 200 tags. Tags are of different types and have different purposes. There are some predefined tags such as severity, criticality, destination IP, and destination port. There are two sets of configurable tags: reserved tags are for Novell internal use to allow future expansion and customer tags are for customer extensions.

Tags can be repurposed by renaming them. The source for a tag can either be external, which means that it is set explicitly by the device or the corresponding Collector, or referential. The value of a referential tag is computed as a function of one or more other tags using the mapping service. For example, a tag can be defined to be the building code for the building containing the asset mentioned as the destination IP of an event. Or, a tag can be computed by the mapping service by using a customer-defined map with the destination IP from the event.

- "Map Service" on page 405
- "Streaming Maps" on page 406
- "Exploit Detection" on page 406

Map Service

The Map Service allows a sophisticated mechanism to propagate business relevance data throughout the system. This facility aids scalability and provides an extensibility advantage by enabling intelligent data transfer between different nodes of the distributed system.

The Map Service cross-references vulnerability scanner data with intrusion detection system signatures and more (for example, asset data and business-relevant data). This allows immediate notification when an attack is attempting to exploit a vulnerable system. Three separate components provide this functionality:

- Collection of real-time events from an intrusion detection source
- Comparing those signatures to the latest vulnerability scans
- Cross-referencing an attack feed through Sentinel Advisor (an optional product module, which cross-references between real-time intrusion detection system attack signatures and the user's vulnerability scanner data).

The Map Service dynamically propagates information through out the system without impacting the system load. When important data sets (that is, "maps" such as asset information or patch update information) are updated in the system, the Map Service propagates the updates across the system.

Streaming Maps

The Map Service employs a dynamic update model and streams the maps from one point to another, avoiding the buildup of large static maps in dynamic memory. The value of this streaming capability is particularly relevant in a mission-critical real-time system such as Sentinel where there must be a steady, predictive, and agile movement of data independent of any transient load on the system.

Exploit Detection

Sentinel provides the ability to cross-reference event data signatures with vulnerability scanner data.You are notified automatically and immediately when an attack is attempting to exploit a vulnerable system. This is accomplished through:

- The Advisor feed
- Intrusion detection
- Vulnerability scanning
- The firewall

Advisor provides a cross-reference between event data signatures and vulnerability scanner data. The Advisor feed has both an alert feed and an attack feed. The alert feed contains information about vulnerabilities and threats. The attack feed is a normalization of event signatures and vulnerability plug-ins.

The supported systems are:

- Intrusion Detections Systems
  - Cisco Secure IDS
  - Enterasys Dragon Host Sensor
  - Enterasys Dragon Network Sensor
  - Intrusion.com (SecureNet_Provider)
  - ISS BlackICE
  - ISS RealSecure Desktop
  - ISS RealSecure Network
  - ISS RealSecure Server
  - ISS RealSecure Guard
  - Snort
  - Symantec Network Security 4.0 (ManHunt)
  - Symantec Intruder Alert
  - McAfee IntruShield
- Vulnerability Scanners
  - eEYE Retina
  - Foundstone Foundscan
  - ISS Database Scanner
  - ISS Internet Scanner
  - ISS System Scanner

- ISS Wireless Scanner
- Nessus
- nCircle IP360
- Qualys QualysGuard

You need at least one vulnerability scanner and either an intrusion detection system, IPS, or firewall from each category above. The intrusion detection system and Firewall DeviceName (rv31) must appear in the event as shown above. Also, the intrusion detection system and the firewall must properly populate the DeviceAttackName (rt1) field (for example, WEB-PHP Mambo uploadimage.php access).

The Advisor feed is sent to the database and then to the Exploit Detection Service. The Exploit Detection Service generates one or two files, depending upon what kind of data has been updated.

Figure A-3 *Exploit Detection*

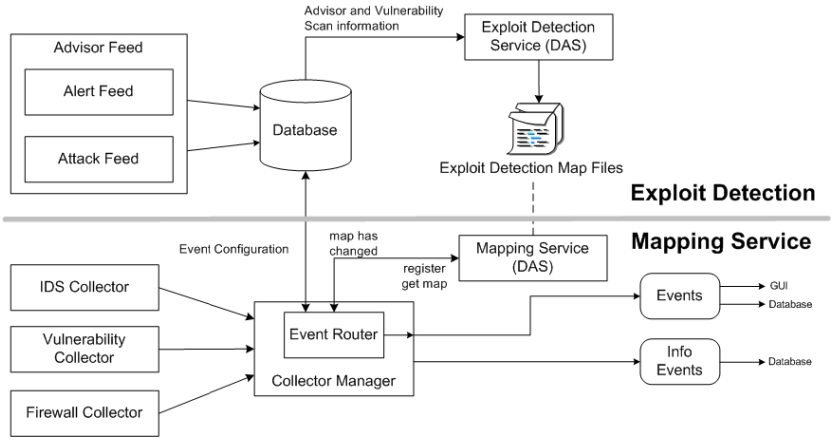

The Exploit Detection map files are used by the Mapping Service to map attacks to exploits of vulnerabilities.

Vulnerability scanners scan for system (asset) vulnerable areas. Intrusion detection systems detects attacks (if any) against these vulnerable areas. Firewalls detect if any traffic is against any of these vulnerable areas. If an attack is associated with any vulnerability, the asset has been exploited.

The Exploit Detection Service generates two files located in:

`<install_directory>/bin/map_data`

The two files are `attackNormalization.csv` and `exploitDetection.csv`.

The `attackNormalization.csv` is generated after:

- Advisor feed
- DAS Startup (if enabled in `das_core.xml`; disabled by default)

The `exploitDetection.csv` is generated after one of the following:

- Advisor feed
- Vulnerability scan
- Sentinel server startup (if enabled in `das_core.xml`; disabled by default)

By default, there are two configured event columns used for exploit detection and they are referenced from a map (all mapped tags have the Scroll icon).

- Vulnerability
- AttackId

Figure A-4 *Event Columns*

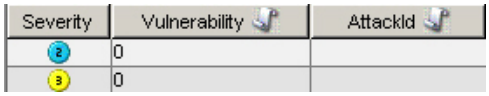

When the Vulnerability field (vul) equals 1, the asset or destination device is exploited. If the Vulnerability field equals 0, the asset or destination device is not exploited.

Sentinel comes preconfigured with the following map names associated with `attackNormalization.csv` and `exploitDetection.csv`.

Table A-2 *Map Name and csv Filename*

Map Name	csv Filename
AttackSignatureNormalization	attackNormalization.csv
IsExploitWatchlist	exploitDetection.csv

There are two types of data sources:

- **External:** Retrieves information from the Collector
- **Referenced from Map:** Retrieves information from a map file to populate the tag.

The AttackId tag has the Device (type of the security device, such as Snort) and AttackSignature columns set as Keys and uses the NormalizedAttackID column in the `attackNormalization.csv` file. In a row where the DeviceName event tag (an intrusion detection system device such as Snort, with information filled in by Advisor and Vulnerability information from the Sentinel database) is the same as Device and where the DeviceAttackName event tag (attack information filled in by Advisor information in the Sentinel Database through the Exploit Detection Service) is the same as AttackSignature, the value for AttackId is where that row intersects with the NormalizedAttackID column.

Figure A-5 *AttackId and Data Source Information*

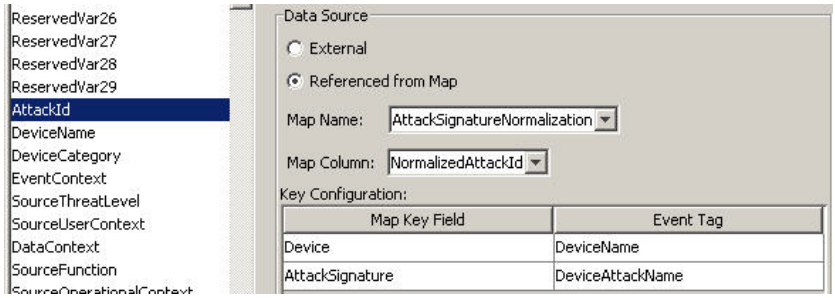

Figure A-6 *attackNormalization.csv Sample*

Device	AttackSignature	NormalizedAttackId	
Secure	BackDoorProbe (TCP 1234)	3	Trojan: Backdoor.SubSeven
Secure	BackDoorProbe (TCP 1999)	3	Trojan: Backdoor.SubSeven
Dragon	RWALLD:SYLOG-FORMAT	4	Sun Microsystems Solaris rwall Elevated F
Snort	RPC TCP rwalld request	4	Sun Microsystems Solaris rwall Elevated F
Snort	RPC UDP rwalld request	4	Sun Microsystems Solaris rwall Elevated F
Snort	WEB-IIS foxweb.dll access	12	Microsoft Exchange Server Arbitrary Code
RealSecure	SMTP_Exchange_Verb_DoS	12	Microsoft Exchange Server Arbitrary Code

The Vulnerability tag has a column entry `_EXIST_`, which means that the map result value is 1 if the key is in IsExploitWatchlist (`exploitDetection.csv` file) or 0 if it is not. The key columns for the vulnerability tag are IP and NormalizedAttackId. When an incoming event with a DestinationIP event tag that matches the IP column entry and an AttackId event tag that matches the NormalizedAttackId column entry in the same row, the result is one (1). If no match is found in a common row, the result is zero (0).

Figure A-7 *Vulnerability and Data Source*

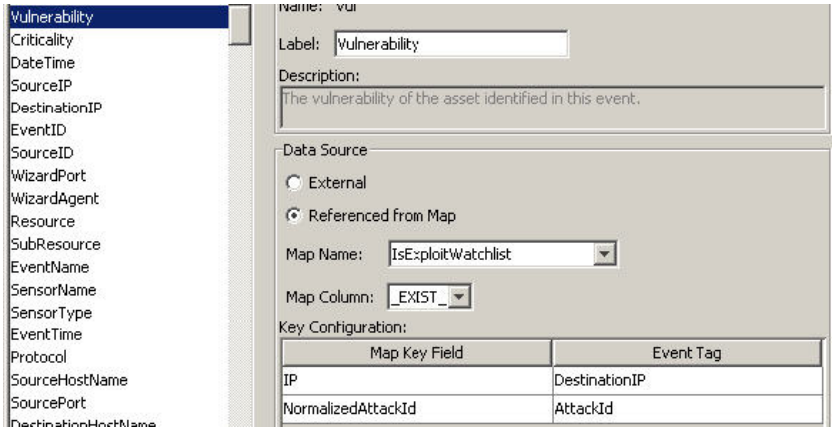

A.3.3 Event Source Management

Sentinel Rapid Deployment delivers a centralized event source management framework to facilitate data source integration. This framework enables all aspects of configuring, deploying, managing and monitoring data Collectors for a broad set of systems, which include databases, operating systems, directories, firewalls, intrusion detection/prevention systems, antivirus applications, mainframes, Web and application servers, and many more.

Using adaptable and flexible technology is central to Sentinel's event source management strategy, which is achieved through interpretive Collectors that parse, normalize, filter and enrich the events in the data stream.

These Collectors can be modified as needed and are not tied to a specific environment. An integrated development environment allows for interactive creation of Collectors by using a "drag and drop" paradigm from a graphical user interface. Non-programmers can create Collectors, ensuring that both current and future requirements are met in an ever-changing IT environment. The command and control operation of Collectors (for example, starting, stopping, and so on) is performed centrally from the Sentinel Control Center. The event source management framework

takes the data from the source system, performs the transformations, and presents the events for later analysis, visualization, and reporting purposes. The framework delivers the following components and benefits:

- **Collectors:** Parse and normalize events from various systems.
- **Connectors:** Connect to the data source to get raw data.
- **Taxonomy:** Allows data from disparate sources to be categorized consistently.
- **Filtering:** Eliminates irrelevant data at the point of collection, saving bandwidth and disk space.
- **Business relevance:** Offers a way to enrich event data with valuable information.
- **Collector Builder:** An integrated development environment for building custom Collectors to collect from unique or proprietary systems.
- **Live view:** User interface for managing live event sources.
- **Scratch pad:** User interface for offline design of event source configuration.

A.3.4 Application Integration

External application integration through standard APIs is central to Sentinel. For example, when dealing with a third party trouble-ticketing system, Sentinel 6 can open an initial ticket in its own iTRAC workflow remediation system. Sentinel then uses bidirectional API to communicate with the other trouble-ticketing systems, such as Remedy and HP OpenView's ServiceDesk, allowing straightforward integration with external systems.

The API is Web Services-based and therefore allows any external systems that are SOAP-aware to take advantage of pervasive integration with the Sentinel system.

A.3.5 Time

The time of an event is very critical to its processing. It is important for reporting and auditing purposes as well as for real time processing. The correlation engine processes time-ordered streams of events and detects patterns within events as well as temporal patterns in the stream. However, the device generating the event might not know the real time when the event is generated. In order to accommodate this, Sentinel allows two options in processing alerts from security devices: trust the time the device reports and use that as the time of the event, or do not trust the device time and instead stamp the event at the time it is first processed by Sentinel by the Collector.

Sentinel is a distributed system and is made up of several processes that can be in different parts of the network. In addition, there can be some delay introduced by the device. In order to accommodate this, the Sentinel processes reorder the events into a time ordered stream before processing.

The following illustration explains the concept of Sentinel time.

Figure A-8 *Time*

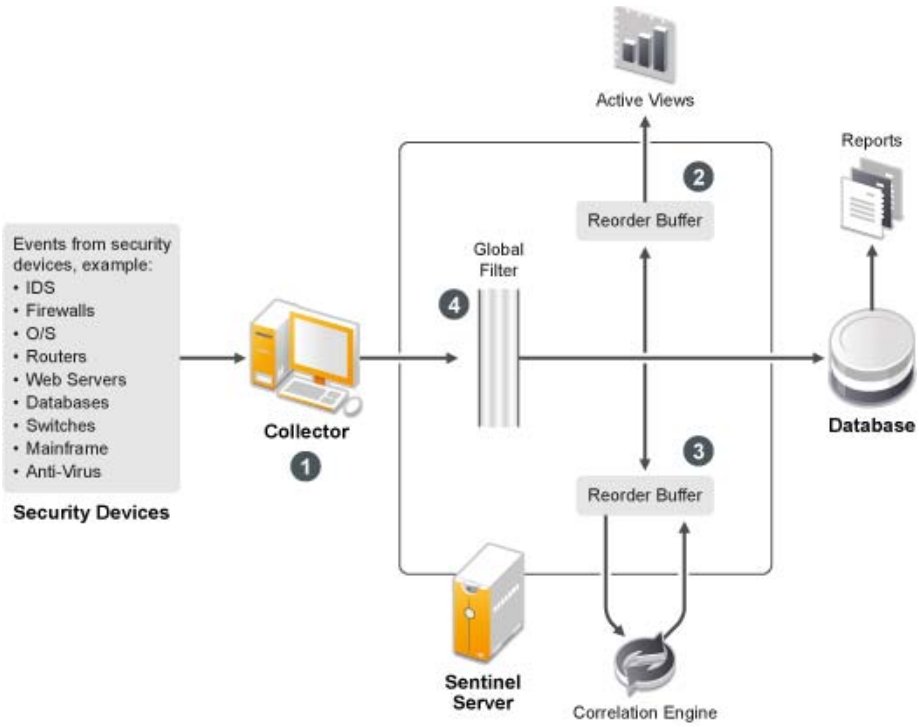

1. By default, the event time is set to Collector Manager time. The ideal time is the device time. Therefore it is best to set the event time to the device time if the device time is available, accurate, and properly parsed by the Collector.
2. Events are sorted into 30 second buckets so that they can be viewed in Active Views. By default, the events that have a timestamp within a 5 minute range from the DAS Core server time (in the past or future) are processed normally. Events that have timestamps more than 5 minutes in the future do not show in the Active Views, but are inserted into the database. Events that have timestamps more than 5 minutes and less than 24 hours in the past are still shown in the charts, but are not shown in the event data for that chart. A drill down operation is necessary to retrieve those events from the database.
3. If the event time is more than 30 seconds older than the server time, the correlation engine does not process the events.
4. If the event time is older than 5 minutes than the Collector Manager time (correct time), events are directly routed to the database.

A.3.6 System Events

System events are a means to report on the status and status changes of the system. There are three types of events generated by the internal system:

- "Internal Events" on page 412
- "Performance Events" on page 412
- "Audit Events" on page 412

Internal Events

Internal events are informational and describe a single state or change of state in the system. They report when a user logs in or fails to authenticate, when a process is started, or when a correlation rule is activated.

Performance Events

Performance events are generated on a periodic basis and describe average resources used by different parts of the system.

Audit Events

Audit events are generated internally. Each time an audited method is called or an audited data object is modified, the audit framework generates audit events. There are two types of Audit events: one that monitors user actions such as user login/out, add/delete user and another that monitors system actions and health, such as process start/stop.

Some of these events were formerly called internal events (mainly for system actions/health monitoring), so the functionality of Audit events is similar to internal events. Audit events can be logged into log files, saved into database, and sent out as Audit events simultaneously (internal events are only sent out as events.).

All System events populate the following attributes:

- **Sensor Type (ST) field:** For internal events this field is set to `I`, for Audit events it is set to `A`, and for performance events it is set to `P`.
- **Event ID:** A unique UUID for the event.
- **Event Time:** The time the event was generated.
- **Source:** The UUID of the process that generated the event.
- **Sensor Name:** The name of the process that generated the event (for example, DAS_Binary).
- **RV32 (Device Category):** Set to `.ESEC`.
- **Collector:** .Performance. for performance events, `Audit` for Audit events, and `Internal` for internal events.

In addition to the common attributes, every system event also sets the resource, sub-resource, the severity, the event name, and the message tags. For internal events, the event name should be specific enough to identify the exact meaning of the event (for example, UserAuthenticationFailed). The message tags add some specific detail; for UserAuthenticationFailed, the message tag contains the name of the user, the OS name if available, and the machine name). For performance events the event name is generic, describing the type of statistical data and the data itself is in the message tag.

Performance events are sent directly to the database. To view them, do a quick query.

For more information, see Appendix B, “System Events for Sentinel,” on page 431.

A.3.7 Processes

The following processes and the services communicate with each other through the ActiveMQ message bus.

- “Sentinel Service (Watchdog)” on page 413

- "Data Access Service (DAS) Process" on page 413
- "Correlation Engine Process (correlation_engine)" on page 414
- "Collector Manager" on page 414
- "ActiveMQ" on page 414

The following illustration shows the architecture for the Sentinel server.

Figure A-9 *Sentinel Server Architecture*

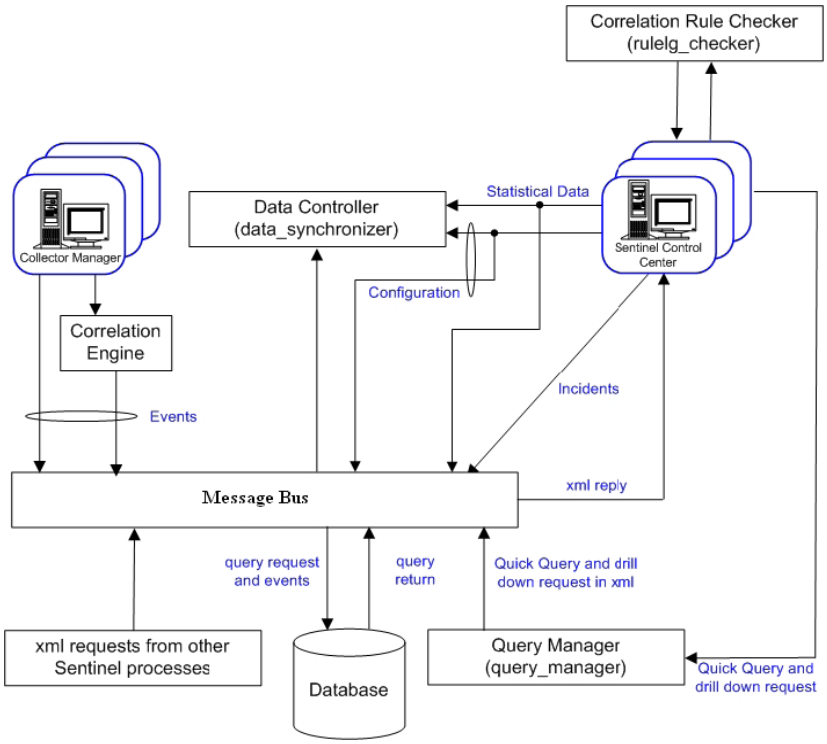

Sentinel Service (Watchdog)

Watchdog is a Sentinel process that manages other Sentinel processes. If a process other than Watchdog stops, Watchdog reports this and then restarts that process.

If this service is stopped, it stops all Sentinel processes on that machine. It executes and reports the health of other Sentinel processes. This process is launched by the Sentinel service.

Data Access Service (DAS) Process

The Data Access Service (DAS) is Sentinel server's persistence service and provides an interface to the database. It provides data-driven access to the database back-end.

DAS is a container composed of two different processes. Each process is responsible for different types of database operations.

- **DAS Core:** DAS core container, which performs the following functions:
  - General Sentinel Service operations including login and historical queries.
  - Provides the server-side functionality for Active Views.

    - Calculates event data summaries that are used in reports.
    - Provides the server-side functionality for Sentinel iTRAC.
    - Provides the server side of the SSL proxy connection to Sentinel server.
- **DAS Binary:** Performs event database insertion.

These processes are controlled by the following configuration files:

- **das_binary.xml:** Used for event and correlated event insertion operations
- **das_core.xml:** All other database operations

DAS receives requests from the different Sentinel processes, converts them to a query against the database, processes the result from the database, and converts it back to a reply. It supports requests to retrieve events for Quick Query and Event Drill Down, in order to retrieve vulnerability information and advisor information and to manipulate configuration information. DAS also handles logging of all events being received from the Collector Manager and requests to retrieve and store configuration information.

Correlation Engine Process (correlation_engine)

The correlation engine (correlation_engine) process receives events from the Collector Manager and publishes correlated events based on user-defined correlation rules.

Collector Manager

The Collector Manager services, processes, and sends events.

ActiveMQ

An open source communication server built around the message-oriented middleware (MOM) architecture. It provides the communication platform for all other Sentinel processes.

A.4 Logical Architecture

Sentinel is composed of three logical layers: the collection and enrichment layer, the business logic layer, and the presentation layer.

- Section A.4.1, “Collection and Enrichment Layer,” on page 415
- Section A.4.2, “Business Logic Layer,” on page 419
- Section A.4.3, “Presentation Layer,” on page 426

Figure A-10 *Sentinel Logical Layers*

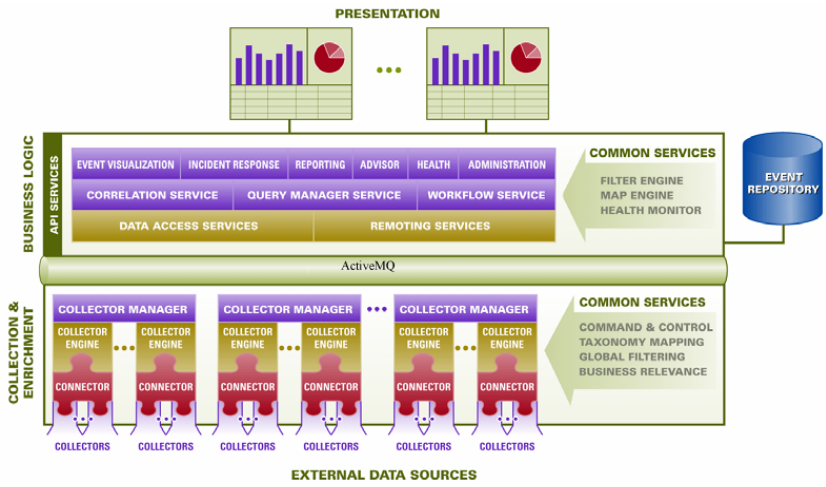

- The collection and enrichment layer aggregates the events from external data sources, transforms the device-specific formats into Sentinel format, enriches the native events source with business-relevant data, and dispatches the event packets to the message bus. The key component orchestrating this function is the Collector, aided by a taxonomy mapping and global filter service.
- The business logic layer contains a set of distributable components. The base component is a Remoting service that adds messaging capabilities to the data objects and services to enable transparent data access across the entire network and Data Access service that is an object management service to allow users to define objects using metadata. Additional services include Correlation, Query Manager, Workflow, Event Visualization, Incident Response, Health, Advisor, Reporting, and Administration.
- The presentation layer renders the application interface to the end user. A comprehensive dashboard called the Sentinel Control Center offers an integrated user workbench consisting of an array of seven different applications accessible through a single common framework. This cross-platform framework is built on Java 1.4 standards and provides a unified view into independent business logic components: real-time interactive graphs, actionable incident response, automated enforceable incident workflow, reporting, incident remediation against known exploits and more.

Each of the layers are illustrated in Figure A-10 and subsequently discussed in detail in the following sections.

- Section A.4.1, “Collection and Enrichment Layer,” on page 415
- Section A.4.2, “Business Logic Layer,” on page 419
- Section A.4.3, “Presentation Layer,” on page 426

A.4.1 Collection and Enrichment Layer

Event Source Management (ESM) provides tools to manage and monitor connections between Sentinel and third-party event sources. Events are aggregated by using a set of flexible and configurable Collectors, which collect data from a myriad of sensors and other devices and sources. User can use prebuilt Collectors, modify existing Collectors or build their own Collectors to ensure that the system meets all requirements.

Data aggregated by the Collectors in the form of events is subsequently normalized and transformed into XML format, enriched with a series of metadata (that is, data about data) using a set of business relevance services, and propagated to the server side for further computational analysis through the message bus platform. The collection and enrichment layer consists of the following components:

- "Connectors and Collectors" on page 416
- "Collector Manager and Engine" on page 416
- "Collector Builder" on page 416
- "Common Services" on page 418

Connectors and Collectors

A Connector is a concentrator or multiplexed adapter that connects the Collector Engine to the actual monitored devices.

Collectors are the component-level aggregators of event data from a specific source. Sentinel primarily supports remote "Collector-less" connections to sources; however, Collectors can be deployed on specific devices where a remote approach is less efficient.

Collectors are controlled from the Sentinel Control Center, which orchestrates the communication between the Collectors and the Sentinel platform for real time analysis, correlation computation and incident response.

Collector Manager and Engine

Collector Manager manages the Collectors, monitors system status messages, and performs event filtering as needed. The main functions of the Collector Manager include transforming events, adding business relevance to events through taxonomy, performing global filtering on events, routing events, and sending health messages to the Sentinel server.

A Collector Engine is the interpreter component that parses the Collector code.

Collector Builder

The Collector Builder is a standalone application that is used to build, configure, and debug Collectors. This application serves as an integrated development environment that allows the user to create new Collectors to parse data from source devices, using a special-purpose interpretive language designed to handle the nature of network and security events.

ESM introduces a new hierarchy of deployment objects that allows users to group multiple connections into sets. The hierarchy is as follows:

Figure A-11 *ESM Hierarchy*

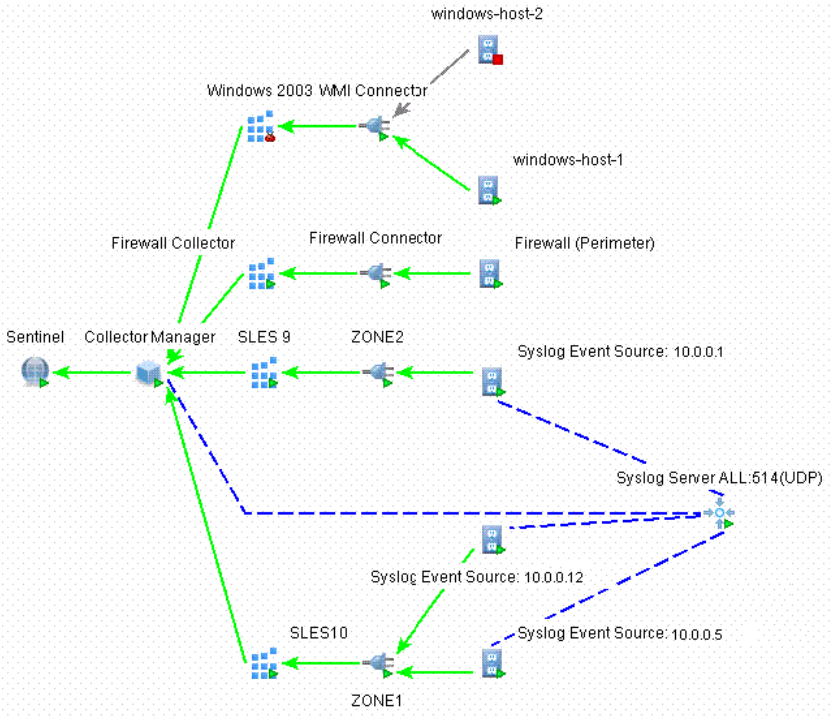

The event source, event source server, Collector, and Connector are configuration-related objects that can be added through the ESM user interface.

- **Event Source:** This node represents a connection to a specific source of data, such as a specific file, firewall, or Syslog relay, and contains the configuration information necessary to establish the connection. The health of this node represents the health of the connection to the data source. This node sends raw data to its parent Connector node.
- **Event Source Server:** This node represents a deployed instance of a server-type Connector plug-in. Some protocols, such as Syslog UDP/TCP, NAudit, and others, push their data from the source to a server that is listening to accept the data. The event source server node represents this server and can be configured to accept data from protocols that are supported by the selected Connector plug-in. This node redirects the raw data it receives to an event source node that is configured to receive data from it.
- **Collector:** This node represents a deployed instance of a Collector script. It specifies which Collector script to use as well as the parameter values with which the Collector should run. This node sends Sentinel events to its parent Collector Manager node.
- **Connector:** This node represents a deployed instance of a Connector plug-in. It includes the specification of which Connector plug-in to use as well as some configuration information, such as auto-discovery. This node sends raw data to its parent Collector node.

Common Services

All of the components in this Collection and Enrichment layer are driven by a set of common services. These utility services form the fabric of the data collection and data enrichment and assist in filtering the noise from the information (through global filters), applying user-defined tags to enrich the events information (through business relevance and taxonomy mapping services), and governing the data Collectors' functions (through command and control services).

- "Taxonomy" on page 418
- "Business Relevance" on page 418
- "Exploit Detection" on page 419

Taxonomy

Nearly all security products produce events in different formats and with varying content. For example, Windows and Solaris report a failed login differently.

Sentinel's taxonomy automatically translates heterogeneous product data into meaningful terms, which allows for a real-time homogeneous view of the entire network security. Sentinel taxonomy formats and filters raw security events before adding event context to the data stream. This process formats all the security data in the most optimal structure for processing by the Sentinel Correlation engine, as you can see in the following diagram.

Figure A-12 *Sentinel Taxonomy*

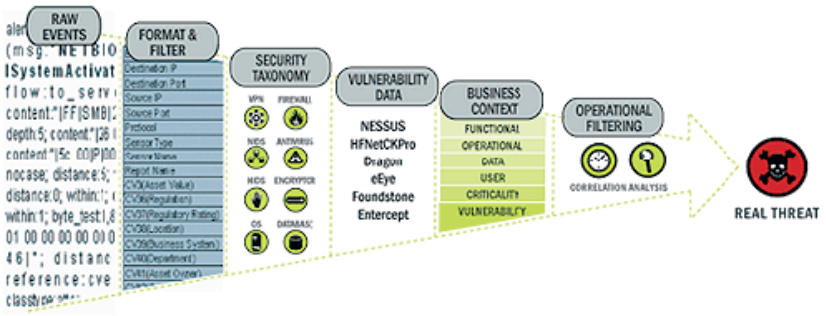

Business Relevance

Sentinel injects business-relevant contextual data directly into the event stream. It includes up to 135 customizable fields where users can add asset specific information such as business unit, owner, asset value, and geography. After this information is added into the system, all other components can take advantage of the additional context.

Figure A-13 *Injecting Business Relevance*

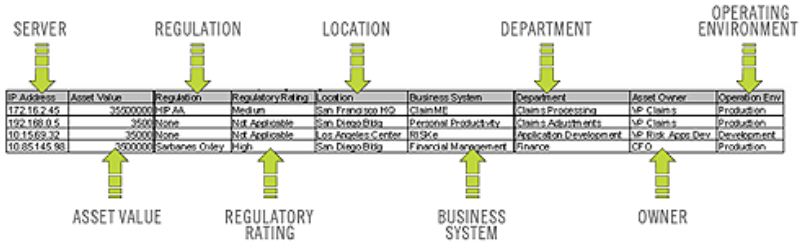

Exploit Detection

Exploit Detection enables immediate, actionable notification of attacks on vulnerable systems. It provides a real-time link between intrusion detection system signatures and vulnerability scan results, notifying users automatically and immediately when an attack attempts to exploit a vulnerable system. This dramatically improves the efficiency and effectiveness of incident response.

Exploit Detection provides users with updates of mappings between intrusion detection systems and vulnerability scanner product signatures. The mappings include a comprehensive list of intrusion detection systems and vulnerability scanners. Users simply upload vulnerability scan results into Sentinel. Exploit Detection automatically parses them and updates the appropriate intrusion detection system Collectors. It uses the embedded knowledge of vulnerability status to efficiently and effectively prioritize responses to security threats in real time.

When an attack is launched against a vulnerable asset, Exploit Detection alerts users with the corresponding severity level of the exploited vulnerability. Users can then take immediate action on high-priority events. This takes the guesswork out of alert monitoring and increases incident response efficiency by focusing reaction on known attacks against vulnerable assets.

Exploit Detection also enables users to map or “un-map” signatures and vulnerabilities to tune out false positives and negatives and to leverage custom signatures or vulnerability scans.

A.4.2 Business Logic Layer

Sentinel services run in specialized containers and allow unparalleled processing and scaling because they are optimized for message-based transport and computation. The key services that make up the Sentinel server include:

- “Remoting Service” on page 419
- “Data Access Service” on page 420
- “Query Manager Service” on page 420
- “Correlation Service” on page 420
- “Dynamic Lists” on page 421
- “Workflow Service (iTRAC)” on page 421
- “Event Visualization” on page 421
- “Incident Response Through iTRAC” on page 423
- “Reporting Service” on page 425
- “Advisor” on page 425
- “Health” on page 426
- “Administration” on page 426
- “Common Services” on page 426

Remoting Service

Sentinel’s Remoting Service provides the mechanism by which the server and client programs communicate. This mechanism is typically referred to as a distributed object application.

The Remoting Service provides the following capabilities:

- **Locating remote objects:** This is achieved through metadata that describes the object name or registration token, although the actual location is not required, because the iSCALE message bus allows for location transparency.
- **Communicating with remote objects:** Details of communication between remote objects are handled by the iSCALE message bus.
- **Object streaming and chunking:** When large amounts of data need to pass back and forth from the client to the server, these objects are optimized to load the data on demand.
- **Callbacks:** Another pattern and layer of abstraction built into the Remoting Service that allows for PTP remote object communication.
- **Service monitoring and statistics:** Provides performance and load statistics for using these remote services.

Data Access Service

Data Access Service (DAS) is an object management service that allows users to define objects using metadata. DAS manages the object and access to objects and automates transmission and persistence. DAS also serves as a facade for accessing data from any persistent data store such as databases, directory services, or files. The operations of DAS include uniform data access through JDBC, and high-performance event insert strategies using native Connectors.

Query Manager Service

The Query Manager Service orchestrates drill-down and event history requests from the Sentinel Control Center. This service is an integral component for implementing the paging algorithm used in the Event History browsing capability. It converts user-defined filters into valid criteria and appends security criteria to it before events are retrieved. This service also ensures that the criteria do not change during a paged event history transaction.

Correlation Service

Sentinel's correlation algorithm computes correlated events by analyzing the data stream in real time. It publishes the correlated events based on user-defined rules before the events reach the database. Rules in the correlation engine can detect a pattern in a single event of a running window of events. When a match is detected, the correlation engine generates a correlated event describing the found pattern and can create an incident or trigger a remediation workflow through ActiveMQ. The correlation engine works with a rules checker component that computes the correlation rule expressions and validates the syntax of filters. In addition to providing a comprehensive set of correlation rules, Sentinel's correlation engine provides specific advantages over database-centric correlation engines.

- By relying on in-memory processing rather than database inserts and reads, the correlation engine performs during high steady-state volumes as well as during event spikes when under attack, which is the time when correlation performance is most critical.
- The correlation volume does not slow down other system components, so the user interface remains responsive, especially with high event volumes.

- Organizations can deploy multiple correlation engines, each on its own server, without the need to replicate configurations or add databases. Independent scaling of components provides cost-effective scalability and performance.
- The correlation engine can add events to incidents after an incident has been determined.

Users are encouraged to use a metric called Event Rules per Second (ERPS). ERPS is the measure of the number of events that can be examined by a correlation rule per second. This measure is a good performance indicator because it estimates the impact on performance when two factors intersect: events per second and number of rules in use.

Dynamic Lists

Dynamic lists are distributed list structures that can be used for storing elements and performing fast lookups on those elements. These lists can store a set of strings such as IP addresses, server names, or usernames. Examples of dynamic lists include:

- Terminated user list
- Suspicious user watch list
- Privileged user watch list
- Authorized ports and services list
- Authorized server list

In all cases, correlation rules might reference named dynamic lists to perform lookups on list members. For example, a rule can be written to identify a file access event from a user who is not a member of the Authorized Users list. Additionally, correlation actions integrate with the dynamic list module to add or remove elements from a list. The combination of lookups and automated actions on the same list provides a powerful feedback mechanism used to identify complex situations.

Workflow Service (iTRAC)

The Workflow Service receives triggers on incident creation and initiates workflow processes based on predefined workflow templates. It manages the life cycle of these processes by generating work items or executing activities. This service also maintains a history of completed processes that can be used for auditing incident responses.

Event Visualization

Active Views, the interactive graphical user interface for event visualization, provides an integrated security management dashboard with a comprehensive set of real-time visualization and analytical tools to facilitate threat detection and analysis. Users can monitor events in real time and perform instant drill-downs from seconds to hours in the past. A wide array of visualization charts and aids allow monitoring of information through 3D bar, 2D stacked, line and ribbon chart representation and others. Additional valuable information can be viewed from the Active Views dashboard, including notification of asset exploits (exploit detection), viewing asset information, and graphical associations between pertinent source IPs and destination IPs.

Because Active Views uses the ActiveMQ architecture, analysts can quickly drill down for further analysis because Active Views provides direct access to the real-time memory-resident event data, which easily handles thousands of events per second without any performance degradation. Data is

kept in memory and written to the database as needed (Active Views can store up to 8 hours of data in memory with typical event loads). This uninterrupted, performance-oriented real-time view is essential when under attack or in a steady state.

Figure A-14 *Active Views*

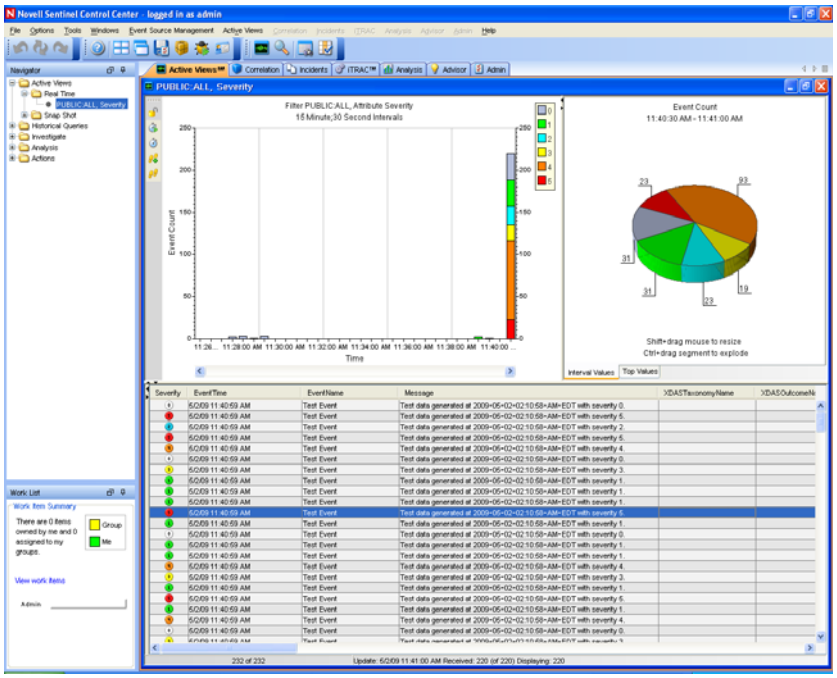

Figure A-15 *Network*

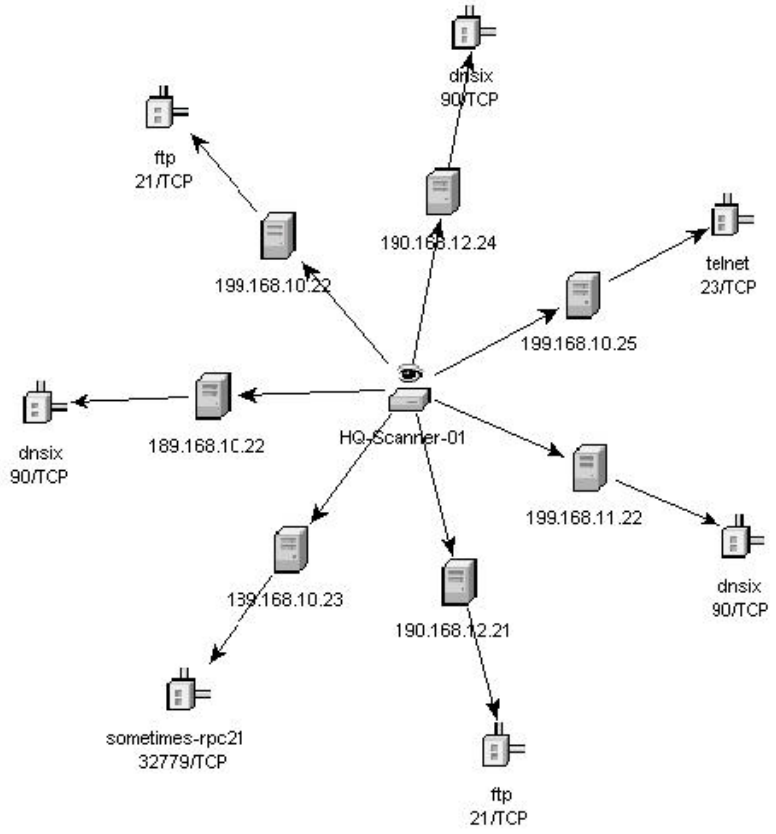

Incident Response Through iTRAC

Sentinel iTRAC transforms traditional security information management from a passive alerting and viewing role to an actionable incident response role by enabling organizations to define and to document incident resolution processes and then guide, enforce and track resolution processes after an incident or violation has been detected.

Sentinel comes with “out-of-the-box” process templates that use the SANS Institute’s guidelines for incident handling. Users can start with these predefined processes and configure specific activities to reflect their organization’s best practices. These processes can be automatically triggered from incident creation or correlation rules or manually engaged by an authorized security or audit professional. iTRAC keeps an audit trail of all actions to support compliance reporting and historical analysis.

Figure A-16 *Process Template*

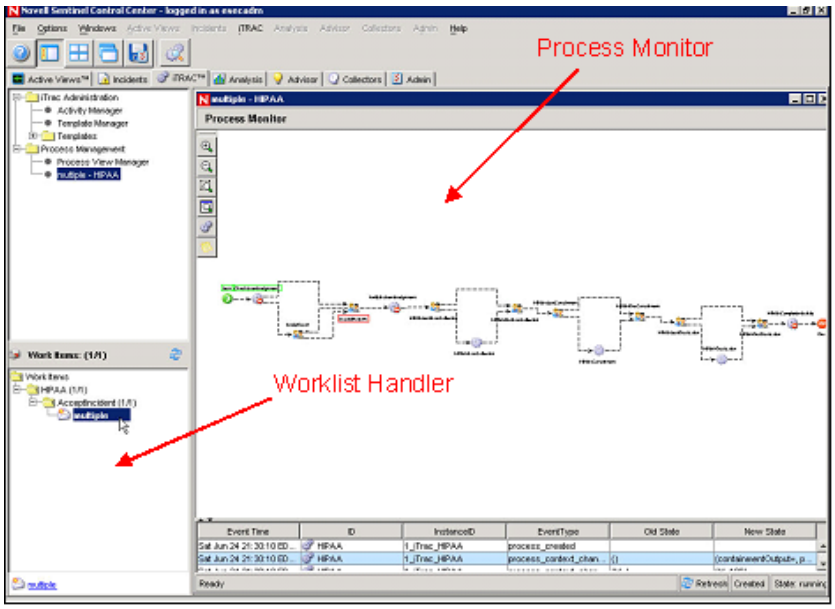

A worklist provides the user with all tasks that have been assigned to the user and a process monitor provides real-time visibility into process status during a resolution process life cycle.

iTRAC's activity framework enables users to customize automated or manual tasks for specific incident-resolution processes. The iTRAC process templates can be configured by using the activity framework to match the template with an organization's best practices. Activities are executed directly from the Sentinel Control Center.

iTRAC's automation framework works using two key components:

- **s container:** Automates the activity's execution for the specified set of steps, based on input rules
- **Workflow container:** Automates the workflow execution based on activities through a worklist.

The input rules are based on the XPDL (XML Processing Description Language) standard and provide a formal model for expressing executable processes in a business enterprise. This standards-based approach to the implementation of business-specific rules and rule sets ensures future-proofing of process definitions for customers.

The iTRAC system uses three Sentinel Rapid Deployment objects that can be defined outside this framework:

- **Incident:** Incidents within Sentinel 6 are groups of events that represent an actionable security incident, associated state, and meta-information. Incidents are created manually or through correlation rules, and can be associated with a workflow process. They can be viewed on the *Incidents* tab.
- **Activity:** An activity is a predefined automatic unit of work, with defined inputs, command-driven activity and outputs such as automatic attachment of asset data to the incident or generation of an e-mail. Activities can be used within workflow templates, triggered by a correlation rule, or executed by a right-click when viewing events.

- **Role:** Users can be assigned to one or more roles, such as Analyst, Admin, and so on. Manual steps in the workflow processes can be assigned to a role.

Sentinel workflows have four major components that are unique to iTRAC:

- **Step:** A step is an individual unit of work within a workflow; there are manual steps, decision steps, command steps, mail steps, and activity-based steps. Each step displays as an icon within a given workflow template.
- **Transition:** A transition defines how the workflow moves from one state (activity) to another and can be determined by an analyst action, by the value of a variable, or by the amount of time elapsed.
- **Templates:** A template is a design for a workflow that controls the execution of a process in Sentinel iTRAC. The template consists of a network of manual and automated steps, activities and criteria for transition between them. Workflow templates define how to respond to an incident when a process based on that template is instantiated. A template can be associated with many incidents.
- **Processes:** A process is a specific instance of a workflow template that is actively being tracked by the workflow system. It includes all the relevant information relating to the instance, including the current step in the workflow, the associated incident, and the results of the steps, attachments and notes. Each workflow process is associated with one incident.

Figure A-17 *iTRAC Workflow*

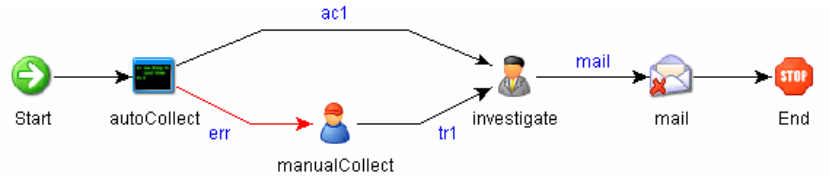

Reporting Service

The Reporting service allows for reporting, including historical and vulnerability reports. Sentinel comes with out-of-the-box reports and enables users to configure their own reports using Jasper Reports. Some examples of reports included with Sentinel are:

- Trend analysis
- Security status of lines of business or critical assets
- Attack types
- Targeted assets
- Response times and resolution
- Policy compliance violations

Advisor

Sentinel Advisor cross-references Sentinel's real-time alert data with known vulnerabilities and remediation information, bridging the gap between incident detection and response. With Advisor, organizations can determine if events exploit specific vulnerabilities and how these attacks impact their assets. Advisor also contains detailed information on the vulnerabilities that attacks intend to

exploit, the potential effects of the attacks if successful and necessary steps for remediation. Recommended remediation steps are enforced and tracked by using iTRAC incident response processes.

Health

The Health service enables users to get a comprehensive view of the distributed Sentinel platform. It aggregates health information from various processes that are typically distributed on various servers. The health information is periodically displayed on the Sentinel Control Center for the end user.

Administration

The Administration facility allows for user management and settings facilities typically needed by application administrators of Sentinel.

Common Services

All of the components in this business logic layer of the architecture are driven by a set of common services. These utility services assist in fine-grained filtering (through the filter engine) of events to users, continuous monitoring of system health statistics (through the Health Monitor) and dynamic updates of system wide data (through the Map Service). Together, these utility services form the fabric of the loosely coupled services that allow for unparalleled processing and scaling over the message bus-based transport for real-time analytics and computation.

A.4.3 Presentation Layer

The presentation layer renders the application interface to the end user. The Sentinel Control Center, the Sentinel Rapid Deployment Web interface are the two comprehensive dashboards that present information to the user, and the Active Browser helps in viewing the selected events.

- "Sentinel Rapid Deployment Web Interface" on page 426
- "Sentinel Control Center" on page 427
- "Active Browser" on page 428

Sentinel Rapid Deployment Web Interface

With the Novell Sentinel Rapid Deployment Web interface, you can manage and search reports and launch the Sentinel Control Center (SCC), the Sentinel Data Manager (SDM), and the Solution Designer. You can also download the Collector Manager installer and the Client installer from the *Application* tab of the Sentinel Rapid Deployment Web interface.

The Web console used for Sentinel Rapid Deployment reporting and full text search also includes the option to launch and install the Sentinel client applications. You can now launch the Sentinel Control Center, Sentinel Solution Designer, and Sentinel Data Manager from a Web browser without installing these client applications locally. The Web console also includes the option to install the client applications and the Sentinel Collector Manager without manually retrieving the installation package.

Sentinel Control Center

In the SCC interface, event presentation is possible through Active Views that display the events in a tabular form or by using different types of charts. The table format displays the variables of the events as columns in a table. Sorting information is possible by clicking on the column name in the grid.

Figure A-18 *Active Views Tabular Format*

Severity	EventTime	EventName	EventID	SourceID	Collector
	5/8/07 12:33:31 PM	DbSpaceLow	B30E4A43-DAB9-1029-9D0A-00'23...	A6C489C0-DAB9-1029-9F5C-00123F9...	
	5/8/07 12:33:31 PM	DbSpaceLow	B30E4A43-DAB9-1029-9D08-00123...	A6C489C0-DAB9-1029-9F5C-00123F9...	
	5/8/07 12:33:31 PM	DbSpaceLow	B30E4A43-DAB9-1029-9D04-00123...	A6C489C0-DAB9-1029-9F5C-00123F9...	
	5/8/07 12:33:31 PM	DbSpaceLow	B30E4A43-DAB9-1029-9D01-00123...	A6C489C0-DAB9-1029-9F5C-00123F9...	

The graphical format displays events as graphs. Stacked Bar 2D, Bar, 3D, Line, and Ribbon graphs are available for representation of information.

Figure A-19 *Active Views Graphical Format Stacked Bar 2D Graph*

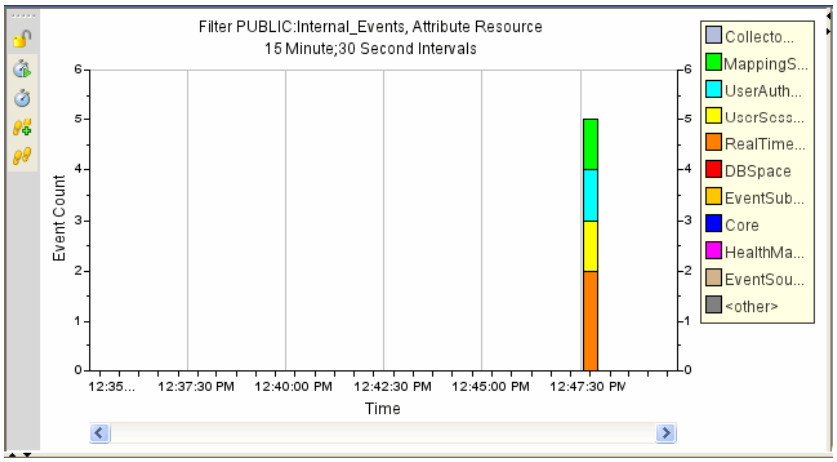

Figure A-20 *Active Views Graphical Format Bar Graph*

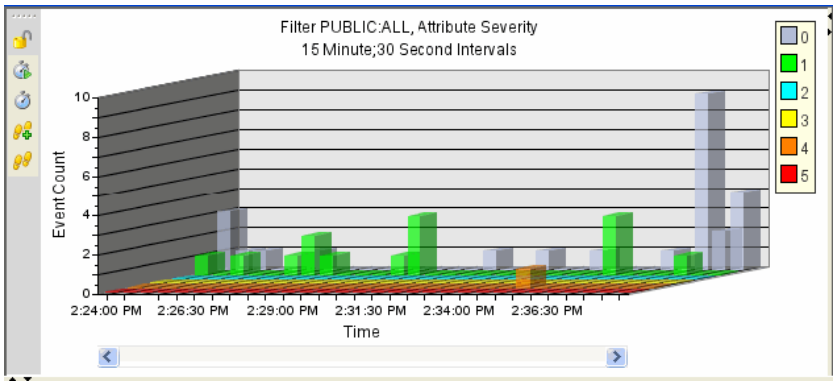

Figure A-21 *Active Views Graphical Format Line Graph*

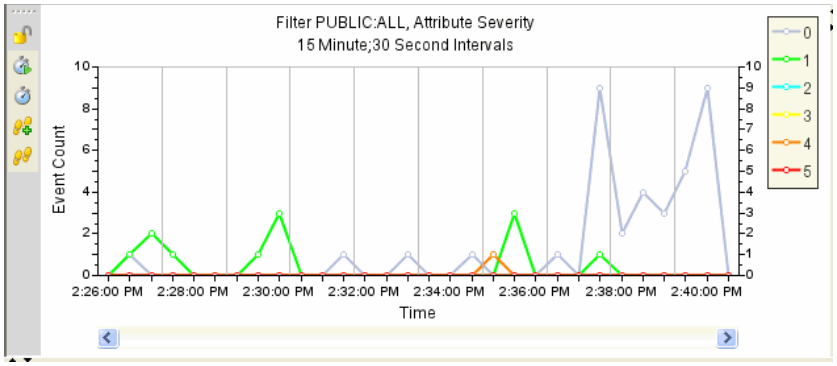

Figure A-22 *Active View Graphical Format Ribbon Graph*

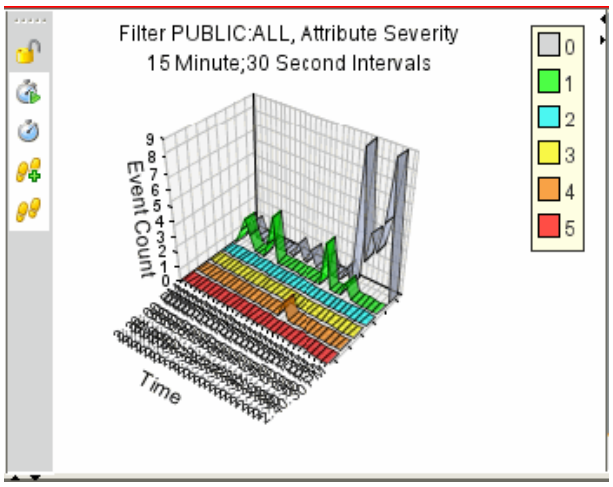

Active Browser

The Active Browser facility helps in viewing the selected events. In the Active Browser, the events are grouped according to the meta tags. In these meta tags, various subcategories are defined. The numbers in the parentheses against these subcategories display the total number of event counts corresponding to the value of the meta tag.

Figure A-23 *Active Browser*

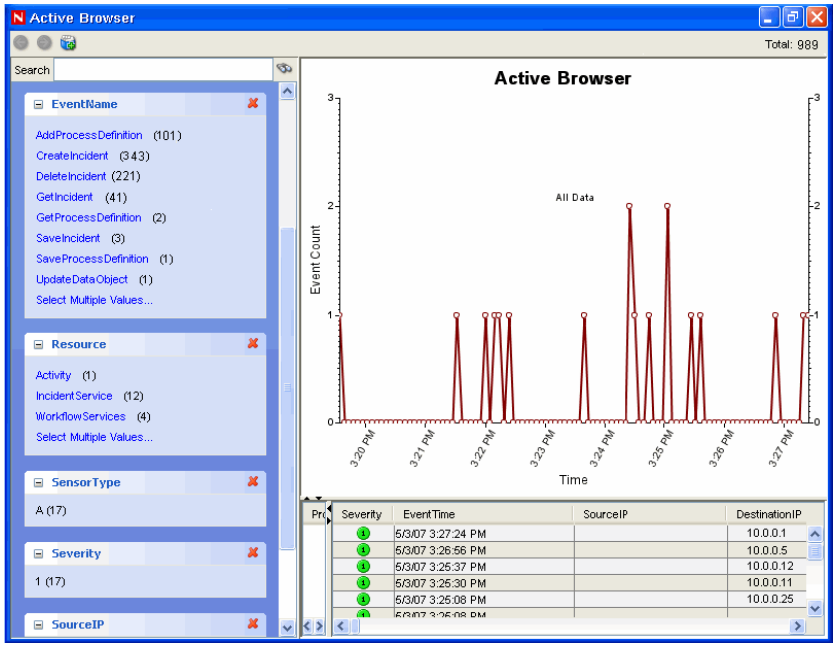

In the Active Browser, the query manager service retrieves a list of events taken from any part of the system and performs a statistical analysis of these events to break them down into ranges of values for each desired attribute of the event. Using single clicks through a Web browser interface, you can select ranges to quickly drill down on a large set of events. Individual event details can be viewed or exported to an HTML or CSV file. Additional event attributes for analysis can be added dynamically at any time, and the interface provides an interactive way to drill down on events in a given time range.

System Events for Sentinel

B

In the tables below, words in italics surrounded by <…> are replaced by relevant values in the real messages.

- Section B.1, “Advisor Audit Events,” on page 431
- Section B.2, “Download Manager Audit Events,” on page 432
- Section B.3, “Authentication Events,” on page 434
- Section B.4, “User Management,” on page 438
- Section B.5, “Database Event Management,” on page 442
- Section B.6, “Database Aggregation,” on page 449
- Section B.7, “Mapping Service,” on page 451
- Section B.8, “Event Router,” on page 458
- Section B.9, “Correlation Engine,” on page 460
- Section B.10, “Event Source Management:General,” on page 466
- Section B.11, “Event Source Management-Event Sources,” on page 475
- Section B.12, “Event Source Management-Collectors,” on page 476
- Section B.13, “Event Source Management-Event Source Servers,” on page 477
- Section B.14, “Event Source Management-Connectors,” on page 478
- Section B.15, “Active Views,” on page 480
- Section B.16, “Data Objects,” on page 483
- Section B.17, “Activities,” on page 484
- Section B.18, “Incidents and Workflows,” on page 485
- Section B.19, “General,” on page 491

B.1 Advisor Audit Events

- Section B.1.1, “Advisor Update Successful,” on page 431
- Section B.1.2, “Advisor Update Failure,” on page 432

B.1.1 Advisor Update Successful

Table B-1 *Advisor Update Successful Event Details*

Tag	Value
Severity	1
Event Name	Advisor update succeeded
Resource	Advisor Processor

Tag	Value
SubResource	Advisor Processor
Message	If the feed file is not available, the message displayed is: 'No new feed available to process' If the feed file is available, the message displayed is: "Number of records inserted: (value) Number of records updated (value) Processing started: (datetime) Processing ended: (datetime)'

B.1.2 Advisor Update Failure

For all types of failures the event would be similar, except that the Message field will have the actual cause of error.

Table B-2 *Advisor Update Failure Event Details*

Tag	Value
Severity	4
Eevnt Name	Advisor update failed
Resource	Advisor Processor
SubResource	Advisor Processor
Message	Advisor feed file `advnxsfeed.1.zip` could be corrupt

B.2 Download Manager Audit Events

- Section B.2.1, "Download Successful," on page 432
- Section B.2.2, "Download Failed," on page 433
- Section B.2.3, "Download Config Updated," on page 433
- Section B.2.4, "Download Config Added," on page 433
- Section B.2.5, "Download Config Removed," on page 434

B.2.1 Download Successful

Table B-3 *Download Successful Event Details*

Tag	Value
Severity	1

Tag	Value
Event Name	Download Success
Resource	DownloadFeedService
SubResource	DOWNLOAD
Message	Download successful for config:<Displays download configuration>

B.2.2 Download Failed

Table B-4 *Download Failed Event Details*

Tag	Value
Severity	1
Event Name	Download Failed
Resource	DownloadFeedService
SubResource	DOWNLOAD
Message	Exception for which the download failed

B.2.3 Download Config Updated

Table B-5 *Download Config Updated Event Details*

Tag	Value
Severity	1
Event Name	Update Download Config
Resource	DownloadFeedService
SubResource	DOWNLOAD
Message	Successfully Updated Download Configuration

B.2.4 Download Config Added

Table B-6 *Download Config Added Event Details*

Tag	Value
Severity	1
Event Name	AddDownloadConfig
Resource	DownloadFeedService

Tag	Value
SubResource	DOWNLOAD
Message	Successfully saved Download Configuration

B.2.5 Download Config Removed

Table B-7 *Download Config Added Event Details*

Tag	Value
Severity	1
Event Name	RemoveDownloadConfig
Resource	DownloadFeedService
SubResource	DOWNLOAD
Message	Successfully removed Download Configuration

B.3 Authentication Events

- Section B.3.1, “Authentication,” on page 434
- Section B.3.2, “Creating Entry For External User,” on page 435
- Section B.3.3, “Duplicate User Objects,” on page 435
- Section B.3.4, “Failed Authentication,” on page 435
- Section B.3.5, “Locked Account,” on page 436
- Section B.3.6, “No Such User Event,” on page 436
- Section B.3.7, “Too Many Active Users,” on page 437
- Section B.3.8, “User Discovered,” on page 437
- Section B.3.9, “User Logged In,” on page 437
- Section B.3.10, “User Logged Out,” on page 438

B.3.1 Authentication

When a user is authentic, the following event is generated:

Table B-8 *Authentication Events Authentication*

Tag	Value
Severity	0
Event Name	Authentication
Resource	UserAuthentication

Tag	Value
SubResource	Authenticate
Message	User <name> has passed Authentication to Sentinel/Wizard

B.3.2 Creating Entry For External User

When creating an external user, the following event is generated:

Table B-9 *Authentication Events: Creating Entry For External User*

Tag	Value
Severity	1
Event Name	CreatingEntryForExternalUser
Resource	UserAuthentication
SubResource	Authentication
Message	No existing local user entry with name <name> found, creating one

B.3.3 Duplicate User Objects

When there is an unexpected second active user object, the following event is generated. This is an internal error.

Table B-10 *Authentication Events: Duplicate User Objects*

Tag	Value
Severity	4
Event Name	TooManyActiveUsers
Resource	UserAuthentication
SubResource	Authenticate
Message	Error in user table : Multiple users with the name <name> found

B.3.4 Failed Authentication

When a user authentication fails, the following event is generated:

Table B-11 *Authentication Events:Failed Authentication*

Tag	Value
Severity	4

Tag	Value
Event Name	AuthenticationFailed
Resource	UserAuthentication
SubResource	Authenticate
Message	Authentication of user <name> with OS name <domUser> from <IP> failed

B.3.5 Locked Account

When a locked user account is attempting to log in, the following event is generated:

Table B-12 *Authentication Events: Locked Account*

Tag	Value
Severity	4
Event Name	LockedUser
Resource	UserAuthentication
SubResource	Authentication
Message	Attempt to login using locked account <acct>

B.3.6 No Such User Event

When a user attempts to log in to the application and authentication succeeds, but the user is not an Sentinel user, the following event is generated:

Table B-13 *Authentication Events: No Such User Event*

Tag	Value
Severity	4
Event Name	NoSuchUser
Resource	UserAuthentication
SubResource	Authenticate
Message	No existing user with name <name> found

B.3.7 Too Many Active Users

Table B-14 *Authentication Events: Too Many Active Users*

Tag	Value
Severity	
Event Name	
Resource	
SubResource	
Message	

B.3.8 User Discovered

If the server restarts, it loses the session information. It then reconstructsthe session when it receives messages from active users. When it discovers a connected user, the following internal event is generated:

Table B-15 *Table B-8: Authentication Events:User Discovered*

Tag	Value
Severity	1
Event Name	UserLoggedIn
Resource	UserSessionManager
SubResource	User
Message	Discovered active user <user> with OS name <osName> at <IP> logged in; currently <number> active users

B.3.9 User Logged In

When a user logs in, the following internal event is generated:

Table B-16 *Authentication Events:User Logged In*

Tag	Value
Severity	1
Event Name	UserLoggedIn
Resource	UserSessionManager
SubResource	User

Tag	Value
Message	User <user> with OS name <osName> at <IP> logged in; currently <number> active users

B.3.10 User Logged Out

When a user logs out, the following internal event is generated:

Table B-17 *Authentication Events : User Logged Out*

Tag	Value
Severity	1
Event Name	UserLoggedOut
Resource	UserSessionManager
SubResource	User
Message	Closing session for <user> OS name <osName> from <IP> was on since <date>; currently <number> active users

B.4 User Management

- Section B.4.1, “Add Users To Role,” on page 438
- Section B.4.2, “Create Role,” on page 439
- Section B.4.3, “Create User,” on page 439
- Section B.4.4, “Creating User Account,” on page 439
- Section B.4.5, “Delete Role,” on page 440
- Section B.4.6, “Deleting User Account,” on page 440
- Section B.4.7, “Locking User Account,” on page 440
- Section B.4.8, “Remove Users From Role,” on page 441
- Section B.4.9, “Resetting Password,” on page 441
- Section B.4.10, “Unlocking User Account,” on page 441
- Section B.4.11, “Updating User,” on page 442

B.4.1 Add Users To Role

Table B-18 *User Management: Add Users To Role*

Tag	Value
Severity	

Tag	Value
Event Name	createRole
Resource	WorkflowServices
SubResource	WorkflowAdminService
Message	Adding users <name> to role <role>

B.4.2 Create Role

Table B-19 *User Management: Create Role*

Tag	Value
Severity	
Event Name	createRole
Resource	WorkflowServices
SubResource	WorkflowAdminService
Message	Creating role with name <name> and description <description>

B.4.3 Create User

Table B-20 *User Management: Create User*

Tag	Value
Severity	
Event Name	createUser
Resource	WorkflowServices
SubResource	WorkflowAdminService
Message	Creating user {0} Name {1} {2} belonging to roles <roles>

B.4.4 Creating User Account

Table B-21 *User Management:Creating User Account*

Tag	Value
Severity	
Event Name	createUser
Resource	Config

Tag	Value
SubResource	UserManagementService
Message	Creating User Account: {0} with Last Name: <lastName>, First Name: <firstName>, State: <state>

B.4.5 Delete Role

Table B-22 *User Management: Delete Role*

Tag	Value
Severity	
Event Name	deleteRole
Resource	WorkflowServices
SubResource	WorkflowAdminService
Message	Deleting role with name <name>

B.4.6 Deleting User Account

Table B-23 *User Management: Deleting User Account*

Tag	Value
Severity	
Event Name	deleteUser
Resource	Config
SubResource	UserManagementService
Message	Deleting User Account: {0}

B.4.7 Locking User Account

Table B-24 *User Management:Locking User Account*

Tag	Value
Severity	
Event Name	lockUser
Resource	Config
SubResource	UserManagementService
Message	Locking User Account: {0}

B.4.8 Remove Users From Role

Table B-25 *User Management:Remove Users From Role*

Tag	Value
Severity	
Event Name	removeUsersFromRole
Resource	WorkflowServices
SubResource	WorkflowAdminService
Message	Removing users <name> from role <role>

B.4.9 Resetting Password

Table B-26 *Resetting Password*

Tag	Value
Severity	
Event Name	setPassword
Resource	Config
SubResource	UserManagementService
Message	Resetting password for User Account {0}

B.4.10 Unlocking User Account

Table B-27 *User Management:Unlocking User Account*

Tag	Value
Severity	
Event Name	unlockUser
Resource	Config
SubResource	UserManagementService
Message	Unlocking User Account: {0}

B.4.11 Updating User

Table B-28 *User Management:Updating User*

Tag	Value
Severity	
Event Name	updateUser
Resource	Config
SubResource	UserManagementService
Message	Updating user: {0} Last Name:<lastName>, First Name: <firstName>, State: <state>

B.5 Database Event Management

- Section B.5.1, “Diskspace Usage Reached Lower Threshold,” on page 443
- Section B.5.2, “Diskspace Usage Reached Upper Threshold,” on page 443
- Section B.5.3, “Dropping the Oldest Partition,” on page 443
- Section B.5.4, “Failing to Drop Online CurrentPartition,” on page 444
- Section B.5.5, “Database Space Reached Specified Percent Threshold,” on page 444
- Section B.5.6, “Database Space Reached Specified Time Threshold,” on page 444
- Section B.5.7, “Database Space Very Low,” on page 445
- Section B.5.8, “Error inserting events,” on page 445
- Section B.5.9, “Error Moving Completed File,” on page 445
- Section B.5.10, “Error Processing Event Message,” on page 446
- Section B.5.11, “Error Saving Failed Events,” on page 446
- Section B.5.12, “Event Insertion Is Blocked,” on page 446
- Section B.5.13, “Event Insertion Is Resumed,” on page 447
- Section B.5.14, “Event Message Queue Overflow,” on page 447
- Section B.5.15, “Event Processing Failed,” on page 447
- Section B.5.16, “No Space In The Database,” on page 448
- Section B.5.17, “Opening Archive File Failed,” on page 448
- Section B.5.18, “Partition Configuration,” on page 448
- Section B.5.19, “Writing to Archive File failed,” on page 449
- Section B.5.20, “Writing to the overflow partition (P_MAX),” on page 449

B.5.1 Diskspace Usage Reached Lower Threshold

Table B-29 *Diskspace Usage Reached Lower Threshold*

Tag	Value
Severity	4
EventName	DBLowSpace
Resource	DBSPace
Message	Diskspace usage reached lower threshold. Its current size is {0} MB as against allocated size {1} MB.

B.5.2 Diskspace Usage Reached Upper Threshold

Table B-30 *Diskspace Usage Reached Upper Threshold*

Tag	Value
Severity	4
EventName	DBNoSpace
Resource	DBSPace
Message	Diskspace usage reached upper threshold. Its current size is {0} MB as against allocated size {1} MB.

B.5.3 Dropping the Oldest Partition

Table B-31 *Dropping the Oldest Partition*

Tag	Value
Severity	4
EventName	DBNoSpace
Resource	DBSPace
Message	Diskspace usage reached upper threshold. Dropping the oldest partition {0}.

B.5.4 Failing to Drop Online CurrentPartition

Table B-32 *Failing to Drop Online CurrentPartition*

Tag	Value
Severity	4
EventName	DBNoSpace
Resource	DBSPace
Message	Diskspace usage reached upper threshold. Failing to drop online current partition {0}.

B.5.5 Database Space Reached Specified Percent Threshold

When event insertion is resumed after being blocked, the following event is sent.:

Table B-33 *Database Event Management: Database Space Reached Specified Percent Threshold*

Tag	Value
Severity	0
Event Name	DbSpaceReachedPercentThrshld
Resource	Database
SubResource	Database
Message	Tablespace <string> has current size of <number> MB with a max size of <number> MB and has reached the percentage threshold of <number> %

B.5.6 Database Space Reached Specified Time Threshold

When event insertion is resumed after being blocked, the following event is sent:

Table B-34 *Database Event Management: Database Space Reached Specified Time Threshold*

Tag	Value
Severity	0
Event Name	DbSpaceReachedTimeThrshld
Resource	Database
SubResource	Database
Message	Tablespace <string> has <number> MB left and growing <number> bytes per second and will run out space within the time threshold specified <number> seconds

B.5.7 Database Space Very Low

When event insertion is resumed after being blocked, the following event is sent:

Table B-35 *Database Event Management: Database Space Very Low*

Tag	Value
Severity	5
Event Name	DbSpaceVeryLow
Resource	Database
SubResource	Database
Message	Tablespace <string> has current size of <number> MB and has reached the physical threshold of <number> MB

B.5.8 Error inserting events

When inserting events into the database fails,the following internal event is generated:

Table B-36 *Database Event Management: Error inserting events*

Tag	Value
Severity	5
Event Name	InsertEventsFailed
Resource	EventSubsystem
SubResource	Events
Message	Error inserting events into the Database—the events might be permanently lost. Please check the Database and backend server logs<Exception>

B.5.9 Error Moving Completed File

When an event file is completed it is moved to the output directory. If that move fails, the following internal event is generated:

Table B-37 *Database Event Management : Error Moving Completed File*

Tag	Value
Severity	3
Event Name	MoveArchiveFileFailed
Resource	<DAS name>
SubResource	ArchiveFile

Tag	Value
Message	Error moving completed archive file <fileName> to <directory>

B.5.10 Error Processing Event Message

Table B-38 *Database Event Management:Error Processing Event Message*

Tag	Value
Severity	
Event Name	ErrorProcessingEventMessage
Resource	EventSubsystem
SubResource	EventStore
Message	Error processing event message, events may be lost; check the log file for more details: {0}

B.5.11 Error Saving Failed Events

Table B-39 *Database Event Management : Error Saving Failed Events*

Tag	Value
Severity	
Event Name	ErrorSavingFailedEvents
Resource	EventSubsystem
SubResource	EventStore
Message	Error inserting failed events to cache; {0} events may be permanently lost. Check the logs for more detail and correct the problem immediately: {1}

B.5.12 Event Insertion Is Blocked

If DAS is writing into the overflow partition and the user attempts to add partitions SDM sends a request to DAS to temporarily stop inserting events into the database. When this happens, DAS sends internal events every time it attempts to insert events into the database.

Table B-40 *Database Event Management : Event Insertion Is Blocked*

Tag	Value
Severity	4
Event Name	EventInsertionIsBlocked
Resource	EventSubSystem

Tag	Value
SubResource	Events
Message	Event insertion is blocked, waiting <number> sec

B.5.13 Event Insertion Is Resumed

When event insertion is resumed after being blocked, the following event is sent:

Table B-41 *Database Event Management : Event InsertionIs Resumed*

Tag	Value
Severity	2
Event Name	EventInsertionResumed
Resource	EventSubSystem
SubResource	Events
Message	Event insertion has resumed after being blocked

B.5.14 Event Message Queue Overflow

Table B-42 *Database Event Management : Event Message Queue Overflow*

Tag	Value
Severity	
Event Name	EventMessageQueueOverflow
Resource	EventSubsystem
SubResource	EventStore
Message	In the previous {0}ms, failed to execute event store task for {1} events because task queue is ful-Events were stored to file for later insertion. Check the log files and the database " "for more information. The error occurred {2} times in this time range: {3}";

B.5.15 Event Processing Failed

Table B-43 *Database Event Management : Event Processing Failed*

Tag	Value
Severity	
Event Name	EventProcessingFailed

Tag	Value
Resource	EventSubsystem
SubResource	EventStore
Message	In the previous {0}ms, failed to process {1} events--Events were stored for later insertion. Check the log files and the database for more information. The error occurred {2} times in this time range: {3}, cause {4}";

B.5.16 No Space In The Database

Table B-44 *Database Event Management : No Space In The Database*

Tag	Value
Severity	
Event Name	DbNoSpace
Resource	DBSpace
SubResource	tableSpace
Message	

B.5.17 Opening Archive File Failed

When opening an archive file for storing the events for aggregation fails, the following internal event is generated.

Table B-45 *Database Event Management : Opening Archive File Failed*

Tag	Value
Severity	3
Event Name	OpenArchiveFileFailed
Resource	<Das name>
SubResource	ArchiveFile
Message	Error opening archive file <fileName> in <directory>

B.5.18 Partition Configuration

Table B-46 *Database Event Management : Partition Configuration*

Tag	Value
Severity	

Tag	Value
Event Name	New/Update/Remove
Resource	
SubResource	PartitionConfig
Message	ableName=<name> PartTimeUnit={1} PartTimeFactor={2} NumberOfUnits={3}

B.5.19 Writing to Archive File failed

When opening an archive file for storing the events for aggregation fails, the following internal event is generated.

Table B-47 *Database Event Management : Writing to Archive File failed*

Tag	Value
Severity	3
Event Name	WriteArchiveFileFailed
Resource	<Das name>
SubResource	ArchiveFile
Message	Error writing newly received events to aggregation archive file <fileName>

B.5.20 Writing to the overflow partition (P_MAX)

An event is sent approximately every 5 minutes, notifying the user when events are being written to the overflow partition (P_MAX). When this occurs, the administrator needs to use SDM and add more partitions or performance starts to degrade.

Table B-48 *Database Event Management : Writing to the overflow partition (P_MAX)*

Tag	Value
Severity	5
Event Name	InsertIntoOverflowPartition
Resource	EventSubSystem
SubResource	Events
Message	Error: currently inserting into the overflow partitions (P_MAX), add more partitions

B.6 Database Aggregation

- Section B.6.1, “Creating Summary,” on page 450
- Section B.6.2, “Deleting Summary,” on page 450

- Section B.6.3, “Disabling Summary,” on page 450
- Section B.6.4, “Enabling Summary,” on page 451
- Section B.6.5, “Error inserting Summary Data into the Database,” on page 451
- Section B.6.6, “Saving Summary,” on page 451

B.6.1 Creating Summary

Table B-49 *Database Aggregation : Creating Summary*

Tag	Value
Severity	
Event Name	createSummary
Resource	
SubResource	
Message	Creating summary: <summaryDescription>

B.6.2 Deleting Summary

Table B-50 *Database Aggregation : Deleting Summary*

Tag	Value
Severity	
Event Name	deleteSummary
Resource	
SubResource	
Message	Deleting summary: <summaryDescription>

B.6.3 Disabling Summary

Table B-51 *Database Aggregation : Disabling Summary*

Tag	Value
Severity	
Event Name	disableSummary
Resource	
SubResource	
Message	Disabling summary: <summaryDescription>

B.6.4 Enabling Summary

Table B-52 *Database Aggregation : Enabling Summary*

Tag	Value
Severity	
Event Name	enableSummary
Resource	
SubResource	EventAggregationAdminService
Message	Enabling summary: <summaryDescription>

B.6.5 Error inserting Summary Data into the Database

If an error is encountered while writing aggregation data into the database, the following internal event is generated:

Table B-53 *Database Aggregation : Error inserting Summary Data into the Database*

Tag	Value
Severity	4
Event Name	SummaryUpdateFailure
Resource	Aggregation
SubResource	Summary
Message	Error saving summary batch to the database for summary <summaryName>

B.6.6 Saving Summary

Table B-54 *Database Aggregation : Saving Summary*

Tag	Value
Severity	
Event Name	saveSummary
Resource	
SubResource	
Message	Saving summary: <summaryDescription>

B.7 Mapping Service

- Section B.7.1, “Error,” on page 452

- Section B.7.2, “Error Applying Incremental Update,” on page 452
- Section B.7.3, “Error initializing map with ID,” on page 453
- Section B.7.4, “Error Refreshing Map,” on page 453
- Section B.7.5, “Error Saving Data File,” on page 454
- Section B.7.6, “Get File Size,” on page 454
- Section B.7.7, “Loaded Large Map,” on page 454
- Section B.7.8, “Long Time To Load Map,” on page 455
- Section B.7.9, “Out Of Sync Detected,” on page 455
- Section B.7.10, “Refreshing Map from Cache,” on page 455
- Section B.7.11, “Refreshing Map from Server,” on page 456
- Section B.7.12, “Save Data File,” on page 456
- Section B.7.13, “Saved Data File,” on page 457
- Section B.7.14, “Timed Out Waiting For Callback,” on page 457
- Section B.7.15, “Timeout Refreshing Map,” on page 457
- Section B.7.16, “Update,” on page 458
- Section B.7.17, “Update,” on page 458

B.7.1 Error

Table B-55 *Database Aggregation : Error*

Tag	Value
Severity	
Event Name	error
Resource	
SubResource	
Message	Error while updating map data: {0}

B.7.2 Error Applying Incremental Update

This event is sent when the mapping service fails to apply an update to an existing client map.

Table B-56 *Database Aggregation : Error Applying Incremental Update*

Tag	Value
Severity	4
Event Name	ErrorApplyingIncrementalUpdate
Resource	MappingService

Tag	Value
SubResource	ReferentialDataObjectMap
Message	The error <error> occurred while applying updates to map <mapName> (ID <mapId>) v.<version>. Rescheduling a refresh to complete map update.

B.7.3 Error initializing map with ID

This internal event is generated from the client side of the mapping service (the one that is part of the Collector Manager). This error is generated when the Collector Manager attempts to retrieve a map that does not exist. This can happen if maps are created and deleted.

Table B-57 *Database Aggregation : Error initializing map with ID*

Tag	Value
Severity	4
Event Name	ErrorNoSuchMap
Resource	MappingService
SubResource	ReferentialDataObjectMap
Message	Error initializing map with id <ID>: no such map

B.7.4 Error Refreshing Map

This internal event is generated from the client side of the mapping service (the one that is part of the Collector Manager). When the Collector Manager is told to refresh the map because it has been modified or its definition has changed, it sends an internal event. This means that there was some unexpected non-transient error while trying to refresh a map. The Collector Manager waits 15 minutes and tries again. If this happens during initialization, the initialization proceeds and this map is ignored until it can be successfully loaded.

Table B-58 *Database Aggregation : Error Refreshing Map*

Tag	Value
Severity	4
Event Name	ErrorRefreshingMapData
Resource	MappingService
SubResource	ReferentialDataObjectMap
Message	Error refreshing map <mapName>: <exc>

B.7.5 Error Saving Data File

Table B-59 *Database Aggregation : Error Saving Data File*

Tag	Value
Severity	
Event Name	ErrorSavingDataFile
Resource	MappingService
SubResource	MapService
Message	The error <error> occurred while saving data to file <fileName> (no) backup

B.7.6 Get File Size

Table B-60 *Database Aggregation : Get File Size*

Tag	Value
Severity	
Event Name	getFileSize
Resource	
SubResource	
Message	Retrieving size for file <fileName>

B.7.7 Loaded Large Map

This internal event is an information event sent by the mapping service indicating that a large map was loaded to the Collector Manager. A map is considered large if the number of rows exceeds 100,000.

Table B-61 *Database Aggregation : Loaded Large Map*

Tag	Value
Severity	0
Event Name	LoadedLargeMap
Resource	MappingService
SubResource	ReferentialDataObjectMap
Message	Finished loading map <name> with id <ID> and <number> entries and total size <#>Kb in <##>sec

B.7.8 Long Time To Load Map

This internal event is an information event sent by the mapping service informing that loading a map took an unusually long time (greater than one minute).

Table B-62 *Database Aggregation : Long time To load Map*

Tag	Value
Severity	0
Event Name	LongTimeToLoadMap
Resource	MappingService
SubResource	ReferentialDataObjectMap
Message	It took <##>sec to load map <name> with id <ID> and <number> entries and total size <##>Kb

B.7.9 Out Of Sync Detected

This event is sent when the mapping service detects that a map is out-of-date. The mapping service automatically schedules a refresh.

Table B-63 *Database Aggregation : Out Of Sync Detected*

Tag	Value
Severity	2
Event Name	OutOfsyncDetected
Resource	MappingService
SubResource	ReferentialDataObjectMap
Message	Map <mapName> detected the map data is out-of-sync, probably because of a missed update notification--scheduling a refresh

B.7.10 Refreshing Map from Cache

This internal event is generated from the client side of the mapping service (the one that is part of the Collector Manager). When the Collector Manager is told to refresh the map because it has been modified or its definition has changed, it sends an internal event. This means that its cache is up-to-date and is refreshing the map from cache.

Table B-64 *Database Aggregation : Refreshing Map from Cache*

Tag	Value
Severity	1

Tag	Value
Event Name	LoadingMapFromCache
Resource	MappingService
SubResource	ReferentialDataObjectMap
Message	Loading from cache v<version> of map <mapName> (ID <id>)

B.7.11 Refreshing Map from Server

This internal event is generated from the client side of the mapping service (the one that is part of the Collector Manager). When the Collector Manager is told to refresh the map because it has been modified or its definition has changed, it sends an internal event. This means that the map was either not in the cache or the version in the cache was not up-to-date and the Collector Manager is retrieving the map from the server.

Table B-65 *Database Aggregation : Refreshing Map from Server*

Tag	Value
Severity	1
Event Name	RefreshingMapFromServer
Resource	MappingService
SubResource	ReferentialDataObjectMap
Message	Refreshing from server map <name> with id <ID>

B.7.12 Save Data File

Table B-66 *Database Aggregation : Save Data File*

Tag	Value
Severity	
Event Name	saveDataFile
Resource	
SubResource	MapService
Message	Saving data file {0}, backup? {1}

B.7.13 Saved Data File

Table B-67 *Database Aggregation : Saved Data File*

Tag	Value
Severity	
Event Name	SavedDataFile
Resource	MappingService
SubResource	MapService
Message	Saved "+fileSize+" bytes to file <fileName> with original backed up to "+backupFile:"no backup of original

B.7.14 Timed Out Waiting For Callback

When the Collector Manager needs to refresh a map, it sends a request to the back end. This request contains a callback. The back-end generates the map and when it is ready it sends the map to the Collector Manager, using the callback. If it takes too long for the response to arrive (more than ten minutes) the Collector Manager submits a second request assuming the first was lost. When this occurs, the following internal event is generated:

Table B-68 *Database Aggregation : Timed Out Waiting For Callback*

Tag	Value
Severity	2
Event Name	TimedoutWaitingForCallback
Resource	MappingService
SubResource	ReferentialDataObjectMap
Message	Map <name> timed out waiting for callback with new map data--retrying

B.7.15 Timeout Refreshing Map

This internal event is generated from the client side of the mapping service (the one that is part of the Collector Manager). When the Collector Manager is told to refresh the map because it has been modified or its definition has changed it sends an internal. This means that the Collector Manager attempted to retrieve the map from the server and the server never acknowledged the request and timed out. This error is considered transient and the Collector Manager retries.

Table B-69 *Database Aggregation : Timeout Refreshing Map*

Tag	Value
Severity	4

Tag	Value
Event Name	TimeoutRefreshingMap
Resource	MappingService
SubResource	ReferentialDataObjectMap
Message	Request timed out while refreshing map <name>: <exception>

B.7.16 Update

Table B-70 *Database Aggregation : Update*

Tag	Value
Severity	
Event Name	update
Resource	
SubResource	MapDataCallback
Message	Updating map data

B.7.17 Update

Table B-71 *Database Aggregation : Update*

Tag	Value
Severity	
Event Name	update
Resource	
SubResource	(low)
Message	Updating map data (ser)

B.8 Event Router

- Section B.8.1, “Event Router is Initializing,” on page 459
- Section B.8.2, “Event Router Is Running,” on page 459
- Section B.8.3, “Event Router is Stopping,” on page 459
- Section B.8.4, “Event Router is Terminating,” on page 460

B.8.1 Event Router is Initializing

This event is sent when an event router starts its initialization. The event router starts initializing when it has established a connection with the back end.

Table B-72 *Event Router : Event Router is Initializing*

Tag	Value
Severity	1
Event Name	EventRouterInitializing
Resource	CollectorManager
SubResource	EventRouter
Message	Event router is initializing in standalone mode; reqId(1EEAD430-E790-1029-93AC-000C296FC5D4)

B.8.2 Event Router Is Running

The Event router is the main component of the Collector Manager (the one that performs the maps, applies global filters, and publishes the events). This internal event is sent when the event router is ready during initialization. When the Collector Manager is restarted, another event is sent when it is ready.

This event is not sent until the event router successfully loaded all the global filters and map information.

Table B-73 *Event Router : Event Router is Running*

Tag	Value
Severity	1
Event Name	EventRouterIsRunning
Resource	CollectorManager

B.8.3 Event Router is Stopping

This event is sent when a request is received by the event router to stop during shutdown.

Table B-74 *Event Router : Event Router is Stopping*

Tag	Value
Severity	2
Event Name	EventRouterStopping
Resource	CollectorManager

Tag	Value
SubResource	EventRouter
Message	Event router is stopping; reqId(B408EC15-F4D2-1029-A795-000C296FC5D4)

B.8.4 Event Router is Terminating

This event is sent when a request is received by the event router to stop during shutdown.

Table B-75 *Event Router : Event Router is Terminating*

Tag	Value
Severity	2
Event Name	EventRouterTerminating
Resource	CollectorManager
SubResource	EventRouter
Message	Event router is terminating; reqId(B408EC15-F4D2-1029-A797-000C296FC5D4)

B.9 Correlation Engine

- Section B.9.1, “Correlation Action Definition,” on page 461
- Section B.9.2, “Correlation Engine Configuration,” on page 461
- Section B.9.3, “Correlation Engine is Running,” on page 461
- Section B.9.4, “Correlation Engine is Stopped,” on page 462
- Section B.9.5, “Correlation Rule,” on page 462
- Section B.9.6, “Correlation Rule Configuration,” on page 462
- Section B.9.7, “Deploy Rules With Actions To Engine,” on page 463
- Section B.9.8, “Disabling Rule,” on page 463
- Section B.9.9, “Enabling Rule,” on page 463
- Section B.9.10, “Rename Correlation Engine,” on page 464
- Section B.9.11, “Rule Deployment is Modified,” on page 464
- Section B.9.12, “Rule Deployment Is Started,” on page 464
- Section B.9.13, “Rule Deployment is Stopped,” on page 465
- Section B.9.14, “Starting Engine,” on page 465
- Section B.9.15, “Stopping Engine,” on page 465
- Section B.9.16, “UnDeploy All Rules From Engine,” on page 466
- Section B.9.17, “UnDeploy Rule,” on page 466
- Section B.9.18, “Update Correlation Rule Actions,” on page 466

B.9.1 Correlation Action Definition

Table B-76 *Correlation Engine : Correlation Action Definition*

Tag	Value
Severity	
Event Name	New/Update/Remove
Resource	Correlation
SubResource	CorrelationActionDefinition
Message	Action Name: <name> with Id: <ID>

B.9.2 Correlation Engine Configuration

Table B-77 *Correlation Engine : Correlation Engine Configuration*

Tag	Value
Severity	
Event Name	New/Update/Remove
Resource	Correlation
SubResource	CorrEngineConfig
Message	Correlation Engine ID: <ID> Name: <name> Active: {2}

B.9.3 Correlation Engine is Running

The correlation engine process can be idled by the user. Its running state determines whether the active process is processing events or not. The process starts in the idle (stopped) state and waits to retrieve its configuration from the database. This event is sent when the engine changes state from stopped to running.

Table B-78 *Correlation Engine : Correlation Engine is Running*

Tag	Value
Severity	1
Event Name	EngineRunning
Resource	CorrelationEngine
SubResource	CorrelationEngine
Message	Correlation Engine is processing events.

B.9.4 Correlation Engine is Stopped

This event is sent out when the engine changes state from running to stopped.

Table B-79 *Correlation Engine : Correlation Engine is Stopped*

Tag	Value
Severity	1
Event Name	EngineStopped
Resource	CorrelationEngine
SubResource	CorrelationEngine
Message	Correlation Engine has stopped processing events.

B.9.5 Correlation Rule

Table B-80 *Correlation Engine : Correlation Rule*

Tag	Value
Severity	
Event Name	New/Update/Remove
Resource	Correlation
SubResource	CorrRule
Message	Rule Name: <name> Type: <type> Rule Id: <ID>

B.9.6 Correlation Rule Configuration

Table B-81 *Correlation Engine : Correlation Rule Configuration*

Tag	Value
Severity	
Event Name	New/Update/Remove
Resource	Correlation
SubResource	CorrRuleConfig
Message	Correlation Rule Config ID: <ID> Rule Definition ID: {1} Name: <name> Active: {3}

B.9.7 Deploy Rules With Actions To Engine

Table B-82 *Correlation Engine : Deploy Rules With Actions To Engine*

Tag	Value
Severity	
Event Name	deployRulesWithActionsToEngine
Resource	CorrelationManagementService
SubResource	CorrelationManagementService
Message	Deploy Rules With Actions To Engine <enginId>: Rules: <ruleID> Actions: <actionID>

B.9.8 Disabling Rule

Table B-83 *Correlation Engine : Disabling Rule*

Tag	Value
Severity	
Event Name	disableRule
Resource	CorrelationManagementService
SubResource	CorrelationManagementService
Message	Disable Rule: {ruleCfgId}

B.9.9 Enabling Rule

Table B-84 *Correlation Engine : Enabling Rule*

Tag	Value
Severity	
Event Name	enableRule
Resource	CorrelationManagementService
SubResource	CorrelationManagementService
Message	Enable Rule: {ruleCfgId}

B.9.10 Rename Correlation Engine

Table B-85 *Correlation Engine : Rename Correlation Engine*

Tag	Value
Severity	
Event Name	renameCorrEngine
Resource	CorrelationManagementService
SubResource	CorrelationManagementService
Message	Rename Engine to: <name> with EngineId: <ID>

B.9.11 Rule Deployment is Modified

This event is sent when an engine successfully reloads a rule deployment. This message is sent regardless of the engine‘s running state.

Table B-86 *Correlation Engine : Rule Deployment is Modified*

Tag	Value
Severity	1
Event Name	DeploymentModified
Resource	CorrelationEngine
SubResource	Deployment
Message	Deployment <name> modified

B.9.12 Rule Deployment Is Started

This event is sent when an engine successfully loads a rule deployment. This message is sent regardless of the engine‘s running state.

Table B-87 *Correlation Engine : Rule Deployment is Started*

Tag	Value
Severity	1
Event Name	DeploymentStarted
Resource	CorrelationEngine
SubResource	Deployment
Message	deployment <name> started

B.9.13 Rule Deployment is Stopped

This event is sent when an engine successfully unloads a rule deployment. This message is sent regardless of the engine‘s running state.

Table B-88 *Correlation Engine : Rule Deployment is Stopped*

Tag	Value
Severity	1
Event Name	DeploymentStopped
Resource	CorrelationEngine
SubResource	Deployment
Message	deployment <name> stopped

B.9.14 Starting Engine

Table B-89 *Correlation Engine : Starting Engine*

Tag	Value
Severity	
Event Name	startEngine
Resource	CorrelationManagementService
SubResource	CorrelationManagementService
Message	Start engine: <engineID>

B.9.15 Stopping Engine

Table B-90 *Correlation Engine : Stopping Engine*

Tag	Value
Severity	
Event Name	stopEngine
Resource	CorrelationManagementService
SubResource	CorrelationManagementService
Message	Stop engine: <engineID>

B.9.16 UnDeploy All Rules From Engine

Table B-91 *Correlation Engine : UnDeploy All Rules From Engine*

Tag	Value
Severity	
Event Name	undeployAllRulesFromEngine
Resource	CorrelationManagementService
SubResource	CorrelationManagementService
Message	Undeploy all rules from Engine:

B.9.17 UnDeploy Rule

Table B-92 *Correlation Engine : UnDeploy Rule*

Tag	Value
Severity	
Event Name	undeployRule
Resource	CorrelationManagementService
SubResource	CorrelationManagementService
Message	Undeploy Rule: {ruleCfgId}

B.9.18 Update Correlation Rule Actions

Table B-93 *Correlation Engine : Update Correlation Rule Actions*

Tag	Value
Severity	
Event Name	updateCorrRuleActions
Resource	CorrelationManagementService
SubResource	CorrelationManagementService
Message	Update Rule Config {0} by deleting Actions: <actionID> and adding Actions: <actionID>

B.10 Event Source Management:General

- Section B.10.1, “Collector Manager Initialized,” on page 467
- Section B.10.2, “Collector Manager Is Down,” on page 468

- Section B.10.3, “Collector Manager Started,” on page 468
- Section B.10.4, “Collector Manager Stopped,” on page 468
- Section B.10.5, “Collector Service Callback,” on page 469
- Section B.10.6, “Cyclical Dependency,” on page 469
- Section B.10.7, “Event Source Manager Callback,” on page 469
- Section B.10.8, “Initializing Collector Manager,” on page 470
- Section B.10.9, “Lost Contact With Collector Manager,” on page 470
- Section B.10.10, “No Data Alert,” on page 470
- Section B.10.11, “Persistent Process Died,” on page 470
- Section B.10.12, “Persistent Process Restarted,” on page 471
- Section B.10.13, “Port Start,” on page 471
- Section B.10.14, “Port Stop,” on page 471
- Section B.10.15, “Reestablished Contact With Collector Manager,” on page 472
- Section B.10.16, “Restart Plugin Deployments,” on page 472
- Section B.10.17, “Restarting Collector Manager (Cold Restart),” on page 473
- Section B.10.18, “Restarting Collector Manager (Warm Restart),” on page 473
- Section B.10.19, “Start Event Source Group,” on page 473
- Section B.10.20, “Start Event Source Manager,” on page 474
- Section B.10.21, “Starting Collector Manager,” on page 474
- Section B.10.22, “Stop Event Source Group,” on page 474
- Section B.10.23, “Stop Event Source Manager,” on page 475
- Section B.10.24, “Stopping Collector Manager,” on page 475

B.10.1 Collector Manager Initialized

Table B-94 *Event Source Management (General) : Collector Manager Initialized*

Tag	Value
Severity	
Event Name	CollectorManagerInitialized
Resource	CollectorManager
SubResource	Internal
Message	Initialized Collector Manager…

B.10.2 Collector Manager Is Down

Table B-95 *Event Source Management (General) : Collector Manager Is Down*

Tag	Value
Severity	
Event Name	CollectorManagerDown
Resource	HealthManager
SubResource	CollectorManagerHealth
Message	

B.10.3 Collector Manager Started

Table B-96 *Event Source Management (General) : Collector Manager Started*

Tag	Value
Severity	
Event Name	CollectorManagerStarted
Resource	CollectorManager
SubResource	Internal
Message	Started Collector Manager...

B.10.4 Collector Manager Stopped

Table B-97 *Event Source Management (General) : Collector Manager Stopped*

Tag	Value
Severity	
Event Name	CollectorManagerStopped
Resource	CollectorManager
SubResource	Internal
Message	Stopped Collector Manager...

B.10.5 Collector Service Callback

Table B-98 *Event Source Management (General) : Collector Service Callback*

Tag	Value
Severity	
Event Name	Restart
Resource	
SubResource	CollectorServiceCallback
Message	Restart Collector with Id: <ID>

B.10.6 Cyclical Dependency

The Event Service sends this event when it detects a cycle in the Event Definition (in dependencies among tags because of referential map assignments). Check the event configuration in SDM and resolve the dependency.

Table B-99 *Event Source Management (General) : Cyclical Dependency*

Tag	Value
Severity	5
Event Name	CyclicalDependency
Resource	EventService
SubResource	ObjectAttrInfos
Message	Cyclical dependency detected in event transformations. Check event configuration.

B.10.7 Event Source Manager Callback

Table B-100 *Event Source Management (General) : Event Source Manager Callback*

Tag	Value
Severity	
Event Name	restart
Resource	
SubResource	EventSourceManagerCallback
Message	Restart node with Id: <ID>

B.10.8 Initializing Collector Manager

Table B-101 *Event Source Management (General) : Initializing Collector Manager*

Tag	Value
Severity	
Event Name	CollectorManagerInitializing
Resource	CollectorManager
SubResource	Internal
Message	Initializing Collector Manager…

B.10.9 Lost Contact With Collector Manager

Table B-102 *Event Source Management (General) : Lost Contact With Collector Manager*

Tag	Value
Severity	
Event Name	LostContactWithCollectorManager
Resource	HealthManager
SubResource	CollectorManagerHealth
Message	Lost contact with collector manager <name> UUID {1}--down for {2} days {3} hrs {4} min

B.10.10 No Data Alert

Table B-103 *Event Source Management (General) : No Data Alert*

Tag	Value
Severity	
Event Name	NoDataAlert
Resource	CollectorManager
SubResource	objectName
Message	No data received for {7} {0} (ID {1}) for last {2} days {3} hrs {4} min {5} sec (threshold {6} ms)

B.10.11 Persistent Process Died

The Collector Engine sends this event when the persistent process Connector detects that its controlled process has died.

Table B-104 Event Source Management (General) : Persistent Process Died

Tag	Value
Severity	5
Event Name	PersistentProcessDied
Resource	AgentManager
SubResource	AgentManager
Message	Persistent Process on port <port ID> has died.

B.10.12 Persistent Process Restarted

The Collector Engine sends this event when the persistent process Connector is able to restart the controlled process that had died.

Table B-105 Event Source Management (General) : Persistent Process Restarted

Tag	Value
Severity	1
Event Name	PersistentProcessRestarted
Resource	AgentManager
SubResource	AgentManager
Message	Persistent Process on port <port ID> has restarted.

B.10.13 Port Start

TheCollector Manager sends this event when a port is started.

Table B-106 Event Source Management (General) : Port Start

Tag	Value
Severity	1
Event Name	PortStart
Resource	AgentManager
SubResource	AgentManager
Message	Processing started for port_<port ID>

B.10.14 Port Stop

The Collector Manager sends this event when a port is stopped.

Table B-107 Event Source Management (General) : Port Stop

Tag	Value
Severity	1
Event Name	PortStop
Resource	AgentManager
SubResource	AgentManager
Message	Processing stopped for port_<port ID>

B.10.15 Reestablished Contact With Collector Manager

Table B-108 Event Source Management (General) : Reestablished Contact With Collector Manager

Tag	Value
Severity	
Event Name	ReestablishedContactWithCollectorManager
Resource	HealthManager
SubResource	CollectorManagerHealth
Message	Reestablished contact with collector manager {0} UUID {1} after {2} days {3} hrs {4} min

B.10.16 Restart Plugin Deployments

Table B-109 Event Source Management (General) : Restart Plugin Deployments

Tag	Value
Severity	
Event Name	restartPluginDeployments
Resource	EventSourceManagement
SubResource	EventSourceManagerService
Message	Restart deployments of plugin: {0}

B.10.17 Restarting Collector Manager (Cold Restart)

Table B-110 *Event Source Management (General) : Restarting Collector Manager (Cold Restart)*

Tag	Value
Severity	
Event Name	CollectorManagerRestart
Resource	CollectorManager
SubResource	Internal
Message	Restarting Collector Manager (Cold restart)

B.10.18 Restarting Collector Manager (Warm Restart)

Table B-111 *Event Source Management (General) : Restarting Collector Manager (Warm Restart)*

Tag	Value
Severity	
Event Name	CollectorManagerRestart
Resource	CollectorManager
SubResource	Internal
Message	Restarting Collector Manager (Warm restart)

B.10.19 Start Event Source Group

Table B-112 *Event Source Management (General) : Start Event Source Group*

Tag	Value
Severity	
Event Name	startEventSourceGroup
Resource	EventSourceManagement
SubResource	EventSourceManagerService
Message	Start Connector: {0}

B.10.20 Start Event Source Manager

Table B-113 *Event Source Management (General) : Start Event Source Manager*

Tag	Value
Severity	
Event Name	startEventSourceManager
Resource	EventSourceManagement
SubResource	EventSourceManagerService
Message	Start Collector Manager: <eventSourceManagerID>

B.10.21 Starting Collector Manager

Table B-114 *Event Source Management (General) : Starting Collector Manager*

Tag	Value
Severity	
Event Name	CollectorManagerStarting
Resource	CollectorManager
SubResource	Internal
Message	Starting Collector Manager

B.10.22 Stop Event Source Group

Table B-115 *Event Source Management (General) : Stop Event Source Group*

Tag	Value
Severity	
Event Name	stopEventSourceGroup
Resource	EventSourceManagement
SubResource	EventSourceManagerService
Message	Stop Connector: {0}

B.10.23 Stop Event Source Manager

Table B-116 *Event Source Management (General) : Stop Event Source Manager*

Tag	Value
Severity	
Event Name	StopEventSourceManager
Resource	EventSourceManagement
SubResource	EventSourceManagerService
Message	Stop Collector Manager: <eventSourceManagerID>

B.10.24 Stopping Collector Manager

Table B-117 *Event Source Management (General) : Stopping Collector Manager*

Tag	Value
Severity	
Event Name	CollectorManagerStopping
Resource	CollectorManager
SubResource	Internal
Message	Stopping Collector Manager...

B.11 Event Source Management-Event Sources

- Section B.11.1, “Start Event Source,” on page 475
- Section B.11.2, “Stop Event Source,” on page 476

B.11.1 Start Event Source

Table B-118 *Event Source Management (Event Sources) : Start Event Source*

Tag	Value
Severity	
Event Name	startEventSource
Resource	EventSourceManagement
SubResource	EventSourceManagerService
Message	Start EventSource: {0}

B.11.2 Stop Event Source

Table B-119 *Event Source Management (Event Sources) : Stop Event Source*

Tag	Value
Severity	
Event Name	stopEventSource
Resource	EventSourceManagement
SubResource	EventSourceManagerService
Message	Stop EventSource: {0}

B.12 Event Source Management-Collectors

- Section B.12.1, “Start Collector,” on page 476
- Section B.12.2, “Stop Collector,” on page 476

B.12.1 Start Collector

Table B-120 *Event Source Management (Collectors) : Start Collector*

Tag	Value
Severity	
Event Name	startCollector
Resource	EventSourceManagement
SubResource	EventSourceManagerService
Message	Start Collector: {0}

B.12.2 Stop Collector

Table B-121 *Event Source Management (Collectors): Stop Collector*

Tag	Value
Severity	
Event Name	stopCollector
Resource	EventSourceManagement
SubResource	EventSourceManagerService
Message	Stop Collector: {0}

B.13 Event Source Management-Event Source Servers

- Section B.13.1, “Start Event Source Server,” on page 477
- Section B.13.2, “Stop Event Source Server,” on page 477
- Section B.13.3, “Stop Event Source Server,” on page 477

B.13.1 Start Event Source Server

Table B-122 *Event Source Management (Event Source Servers): Start Event Source Server*

Tag	Value
Severity	
Event Name	startEventSourceServer
Resource	EventSourceManagement
SubResource	EventSourceManagerService
Message	Start EventSourceServer: <eventSourceServerID>

B.13.2 Stop Event Source Server

Table B-123 *Event Source Management (Event Source Servers): Stop Event Source Server*

Tag	Value
Severity	
Event Name	stopEventSourceServer
Resource	EventSourceManagement
SubResource	EventSourceManagerService
Message	Stop EventSourceServer: <eventSourceServerID>

B.13.3 Stop Event Source Server

Table B-124 *Event Source Management (Event Source Servers): Stop Event Source Server*

Tag	Value
Severity	
Event Name	stopEventSourceServer
Resource	EventSourceManagement
SubResource	EventSourceManagerService

Tag	Value
Message	Stop EventSourceServer: <eventSourceServerID>

B.14 Event Source Management-Connectors

- Section B.14.1, “Data Received After Timeout,” on page 478
- Section B.14.2, “Data Timeout,” on page 478
- Section B.14.3, “File Rotation,” on page 479
- Section B.14.4, “Process Auto Restart Error,” on page 479
- Section B.14.5, “Process Start Error,” on page 479
- Section B.14.6, “Process Stop,” on page 480
- Section B.14.7, “WMI Connector Status Message,” on page 480

B.14.1 Data Received After Timeout

When the File Connector is configured with a data timeout greater than 0 in the `package.xml` file, the data timeout period is reached without reading any data, then new data is read from the file, the following internal event is generated:

Table B-125 *Event Source Management (Connectors): Data Received After Timeout*

Tag	Value
Severity	4
Event Name	FileUpdatedAfterTimeout
Resource	FileConnector
SubResource	FileConnector
Message	After Event source<File Event Source ID> reached time out of<Timeout Period>, file<File Location> received new data.

B.14.2 Data Timeout

When the File Connector is configured with a data timeout greater than 0 in the `package.xml` file and no data is read from the file in the data timeout period, the following internal event is generated:

Table B-126 *Event Source Management (Connectors): Data Timeout*

Tag	Value
Severity	4
Event Name	FileTimeout
Resource	FileConnector

Tag	Value
SubResource	FileConnector
Message	Event source <File Event Source ID> reached time out of <Timeout Period> when processing file <File Location>.

B.14.3 File Rotation

When the File Connector is configured to use file rotation and the Connector changes from one file to the next, the following internal event is generated:

Table B-127 *Event Source Management (Connectors): File Rotation*

Tag	Value
Severity	4
Event Name	RotatingFile
Resource	FileConnector
SubResource	FileConnector
Message	File rotated for event source <File Event Source ID>. Rotating file from <Previous File Location> to <New File Location>.

B.14.4 Process Auto Restart Error

Table B-128 *Event Source Management (Connectors): Process Auto Restart Error*

Tag	Value
Severity	4
Event Name	ProcessAutoRestartError
Resource	ProcessConnector
SubResource	ProcessConnector
Message	Process <{0}> [command: {1}] was automatically restarted more than the allowed {2} automatic restart(s) in {3} min. The process will no longer be automatically restarted. Please check process configuration.

B.14.5 Process Start Error

Table B-129 *Event Source Management (Connectors): Process Start Error*

Tag	Value
Severity	1

Tag	Value
Event Name	ProcessStartError
Resource	ProcessConnector
SubResource	ProcessConnector
Message	Error starting command: {0}

B.14.6 Process Stop

Table B-130 *Event Source Management (Connectors) : Process Stop*

Tag	Value
Severity	1
Event Name	ProcessStop
Resource	ProcessConnector
SubResource	ProcessConnector
Message	Process <{0}> exited [command: {1}]

B.14.7 WMI Connector Status Message

Table B-131 *Event Source Management (Connectors) : WMI Connector Status Message*

Tag	Value
Severity	4
Event Name	WMIConnectorStatusMessage
Resource	WMIConnector
SubResource	WMIConnector
Message	<Exception>

B.15 Active Views

- Section B.15.1, “Active View Created,” on page 481
- Section B.15.2, “Active View Joined,” on page 481
- Section B.15.3, “Active View No Longer Permanent,” on page 481
- Section B.15.4, “Active View Now Permanent,” on page 482
- Section B.15.5, “Idle Active View Removed,” on page 482
- Section B.15.6, “Idle Permanent Active View Removed,” on page 482

B.15.1 Active View Created

DAS_Binary sends this event when an Active View is created.

Table B-132 *Active View : Active View Created*

Tag	Value
Severity	1
Event Name	RtChartCreated
Resource	RealTimeSummaryService
SubResource	ChartManager
Message	Creating new Active View with filter <filter> and attribute <attribute> for users with security filter <security filter>. Currently <n> Active View(s) Collecting.

B.15.2 Active View Joined

DAS_Binary sends this event when a user connects to an existing Active View.

Table B-133 *Active View : Active View Joined*

Tag	Value
Severity	1
Event Name	RtChartJoiningExistingData
Resource	RealTimeSummaryService
SubResource	ChartManager
Message	Joining existing Active View with filter <filter> and attribute <attribute> for users with security filter <security filter>. Currently <n> Active View(s) Collecting.

B.15.3 Active View No Longer Permanent

DAS_Binary sends this event when it detects a formerly permanent Active View that is no longer permanent. This check happens periodically, so it can be several minutes after an Active View is removed from preferences before this event is generated.

Table B-134 *Active View : Active View No Longer Permanent*

Tag	Value
Severity	1
Event Name	RtChartNotPermanent
Resource	RealTimeSummaryService

Tag	Value
SubResource	ChartManager
Message	Active View with filter <filter> and attribute <attribute> for users with security filter <security filter> is no longer permanent.

B.15.4 Active View Now Permanent

DAS_Binary sends this event when it detects an Active View as newly permanent. This check happens periodically, so it can be several minutes after an Active View is saved to preferences before this event is generated.

Table B-135 *Active View : Active View Now Permanent*

Tag	Value
Severity	1
Event Name	RtChartIsNowPermanent
Resource	RealTimeSummaryService
SubResource	ChartManager
Message	Active View with filter <filter> and attribute <attribute> for users with security filter <security filter> is now permanent.

B.15.5 Idle Active View Removed

DAS_Binary sends this event when a non:permanent Active View is removed because of inactivity.

Table B-136 *Active View : Idle Active View Removed*

Tag	Value
Severity	1
Event Name	RtChartInactiveAndRemoved
Resource	RealTimeSummaryService
SubResource	ChartManager
Message	Removed idle Active View with filter <filter> and attribute <attribute> for users with security filter <security filter>. Currently <n> Active View(s) Collecting.

B.15.6 Idle Permanent Active View Removed

DAS_Binary sends this event when a permanent Active View is removed because of inactivity. Permanent Active Views are saved in user preferences, and they time out after several days of inactivity by default.

Table B-137 Active View : Idle Permanent Active View Removed

Tag	Value
Severity	1
Event Name	RtPermanentChartRemoved
Resource	RealTimeSummaryService
SubResource	ChartManager
Message	Removed idle permanent Active View with filter <filter> and attribute <attribute> for users with security filter <security filter>. Currently <n> Active View(s) Collecting.

B.16 Data Objects

- Section B.16.1, “Activity Definition,” on page 483
- Section B.16.2, “Configuration,” on page 483
- Section B.16.3, “Viewing Configuration Store,” on page 484
- Section B.16.4, “Write Data,” on page 484

B.16.1 Activity Definition

Table B-138 Data Objects : Activity Definition

Tag	Value
Severity	
Event Name	New/Update/Remove
Resource	
SubResource	ActivityDefinition
Message	Activaty Name: <name> Description: <description>

B.16.2 Configuration

Table B-139 Data Objects : Configuration

Tag	Value
Severity	
Event Name	New/Update/Remove
Resource	Core

Tag	Value
SubResource	FilterConfig, GlobalFilterConfig, MenuConfig, OptionsConfig, IncidentActionConfig, AnalyzeDefaultConfig, AnalyzeReportConfig, AdvisorDefaultConfig and AdvisorReportConfig
Message	Updating Config Object: <name> by User: _SYSTEM

B.16.3 Viewing Configuration Store

Table B-140 *Data Objects : Viewing Configuration Store*

Tag	Value
Severity	
Event Name	New/Update/Remove
Resource	
SubResource	ViewConfigurationStore
Message	name <name> type <type> description <description>

B.16.4 Write Data

Table B-141 *Data Objects : Write Data*

Tag	Value
Severity	
Event Name	WriteData
Resource	ListService
SubResource	ListUpdater
Message	Could not write data for list

B.17 Activities

- Section B.17.1, “Creating an Activity,” on page 485
- Section B.17.2, “Deleting an Activity,” on page 485
- Section B.17.3, “Saving an Activity,” on page 485

B.17.1 Creating an Activity

Table B-142 *Activities : Creating an Activity*

Tag	Value
Severity	
Event Name	createActivity
Resource	
SubResource	ActivityNamespace
Message	Creating iTRAC Activity <name>

B.17.2 Deleting an Activity

Table B-143 *Activities : Deleting an Activity*

Tag	Value
Severity	
Event Name	deleteActivity
Resource	
SubResource	ActivityNamespace
Message	Deleting iTRAC Activity <name>

B.17.3 Saving an Activity

Table B-144 *Activities : Saving an Activity*

Tag	Value
Severity	
Event Name	saveActivity
Resource	
SubResource	ActivityNamespace
Message	Saving changes for iTRAC Activity <name>

B.18 Incidents and Workflows

- Section B.18.1, “Add Events to Incident,” on page 486
- Section B.18.2, “Adding Process Definition,” on page 486
- Section B.18.3, “Create Incident,” on page 487

- Section B.18.4, “Creating Group,” on page 487
- Section B.18.5, “Creating User,” on page 487
- Section B.18.6, “Delete Incident,” on page 488
- Section B.18.7, “Deleting Group,” on page 488
- Section B.18.8, “Deleting Process Definition,” on page 488
- Section B.18.9, “Deleting User,” on page 489
- Section B.18.10, “E-Mail Incident,” on page 489
- Section B.18.11, “Get Incident,” on page 489
- Section B.18.12, “Save Incident,” on page 490
- Section B.18.13, “Saving Group,” on page 490
- Section B.18.14, “Saving Process Definition,” on page 490
- Section B.18.15, “Viewing Process Definition,” on page 491

B.18.1 Add Events to Incident

Table B-145 *Incidents and Workflow : Add Events to Incident*

Tag	Value
Severity	
Event Name	addEventsToIncident
Resource	IncidentService
SubResource	IncidentService
Message	User: <name> adding <number> events to incident with ID: <ID>

B.18.2 Adding Process Definition

Table B-146 *Incidents and Workflow : Adding Process Definition*

Tag	Value
Severity	
Event Name	addProcessDefinition
Resource	WorkflowServices
SubResource	WorkflowObjectMgrService
Message	reading iTRAC Template <name>

B.18.3 Create Incident

Table B-147 *Incidents and Workflow : Create Incident*

Tag	Value
Severity	
Event Name	createIncident
Resource	IncidentService
SubResource	IncidentService
Message	User: <name> created incident with name: <incidentName>, state: <state>, severity: <severity>, resolution: <resolution>

B.18.4 Creating Group

Table B-148 *Incidents and Workflow : Creating Group*

Tag	Value
Severity	
Event Name	createGroup
Resource	WorkflowServices
SubResource	WorkflowObjectMgrService
Message	Creating iTRAC Role {0} : description : <description>

B.18.5 Creating User

Table B-149 *Incidents and Workflow : Creating User*

Tag	Value
Severity	
Event Name	createUser
Resource	WorkflowServices
SubResource	WorkflowObjectMgrService
Message	Creating User in WorkFlow: {0} with firstname: <firstName> lastname : <lastName>

B.18.6 Delete Incident

Table B-150 *Incidents and Workflow : Delete Incident*

Tag	Value
Severity	
Event Name	deleteIncident
Resource	IncidentService
SubResource	IncidentService
Message	Delete incident with ID: <ID>

B.18.7 Deleting Group

Table B-151 *Incidents and Workflow : Deleting Group*

Tag	Value
Severity	
Event Name	deleteGroup
Resource	WorkflowServices
SubResource	WorkflowObjectMgrService
Message	Deleting iTRAC Role {0} : description : <description>

B.18.8 Deleting Process Definition

Table B-152 *Incidents and Workflow : Deleting Process Definition*

Tag	Value
Severity	
Event Name	deleteProcessDefinition
Resource	WorkflowServices
SubResource	WorkflowObjectMgrService
Message	Deleting iTRAC Template <ID>

B.18.9 Deleting User

Table B-153 *Incidents and Workflow : Deleting User*

Tag	Value
Severity	
Event Name	deleteUser
Resource	WorkflowServices
SubResource	WorkflowObjectMgrService
Message	Deleting User in WorkFlow: {0} with firstname: <firstName> lastname : <lastName>

B.18.10 E-Mail Incident

Table B-154 *Incidents and Workflow : E-mail Incident*

Tag	Value
Severity	
Event Name	emailIncident
Resource	IncidentService
SubResource	IncidentService
Message	User: <name> emailed incident with name: <incidentName>, state: <state>, severity: <severity>{2}, resolution: <resolution> to email address: <e-mailID>

B.18.11 Get Incident

Table B-155 *Incidents and Workflow : Get Incident*

Tag	Value
Severity	
Event Name	getIncident
Resource	IncidentService
SubResource	IncidentService
Message	Get incident with ID: <ID>

B.18.12 Save Incident

Table B-156 *Incidents and Workflow : Save Incident*

Tag	Value
Severity	
Event Name	saveIncident
Resource	IncidentService
SubResource	IncidentService
Message	Save incident with name: <name>, state: <state>, severity: <severity>, resolution: <resolution>

B.18.13 Saving Group

Table B-157 *Incidents and Workflow : Saving Group*

Tag	Value
Severity	
Event Name	saveGroup
Resource	WorkflowServices
SubResource	WorkflowObjectMgrService
Message	Saving iTRAC Role {0} : description : <description>

B.18.14 Saving Process Definition

Table B-158 *Incidents and Workflow : Saving Process Definition*

Tag	Value
Severity	
Event Name	saveProcessDefinition
Resource	WorkflowServices
SubResource	WorkflowObjectMgrService
Message	Saving iTRAC Template <name>

B.18.15 Viewing Process Definition

Table B-159 *Incidents and Workflow : Viewing Process Definition*

Tag	Value
Severity	
Event Name	getProcessDefinition
Resource	WorkflowServices
SubResource	WorkflowObjectMgrService
Message	Viewing iTRAC Template <ID>

B.19 General

- Section B.19.1, “Configuration Service,” on page 491
- Section B.19.2, “Controlled Process is started,” on page 492
- Section B.19.3, “Controlled Process Is Stopped,” on page 492
- Section B.19.4, “Importing Auxiliary,” on page 493
- Section B.19.5, “Importing Plug-In,” on page 493
- Section B.19.6, “Load Esec Taxonomy to XML,” on page 493
- Section B.19.7, “Process Auto Restart Error,” on page 493
- Section B.19.8, “Process Restarts,” on page 494
- Section B.19.9, “Proxy Client Registration Service (medium),” on page 494
- Section B.19.10, “Restarting Process,” on page 495
- Section B.19.11, “Restarting Processes,” on page 495
- Section B.19.12, “Starting Process,” on page 495
- Section B.19.13, “Starting Processes,” on page 496
- Section B.19.14, “Stopping Process,” on page 496
- Section B.19.15, “Stopping Processes,” on page 496
- Section B.19.16, “Store Esec Taxonomy From XML,” on page 497
- Section B.19.17, “Watchdog Process is started,” on page 497
- Section B.19.18, “Watchdog Process Is stopped,” on page 497

B.19.1 Configuration Service

Table B-160 *General : Configuration Service*

Tag	Value
Severity	

Tag	Value
Event Name	saveConfig
Resource	
SubResource	ConfigService
Message	Saving configuration, unit {0} app {1} userId {2}

B.19.2 Controlled Process is started

Watchdog is run as a service. Its main purpose is to keep Sentinel processes running. If a process dies, Watchdog automatically restarts that process. This event is sent when a process is started.

Table B-161 *General : Controlled Process is started*

Tag	Value
Severity	1
Event Name	ProcessStart
Resource	Sentinel
SubResource	Process
Message	Process <ProgramName> spawned (command <pID>)

B.19.3 Controlled Process Is Stopped

This event is sent when a process is stopped. The severity is set to 5 if the process was set to respawn (that is, it is not expected to die). The severity is set to 1 if the process was set to run once.

Table B-162 *General : Controlled Process Is Stopped*

Tag	Value
Severity	1/5
Event Name	ProcessStop
Resource	Sentinel
SubResource	Process
Message	Process <ProgramName> exited (command <exit_code>)

B.19.4 Importing Auxiliary

Table B-163 *General : Importing Auxiliary*

Tag	Value
Severity	
Event Name	importAuxiliary
Resource	
SubResource	PluginRepositoryService (Medium)
Message	Import auxiliary file <auxiliaryJarName> into plugin <pluginID>.

B.19.5 Importing Plug-In

Table B-164 *General : Importing Plugin*

Tag	Value
Severity	
Event Name	importPlugin
Resource	
SubResource	PluginRepositoryService
Message	Import plugin <name> (ID <ID>) of type <type>.

B.19.6 Load Esec Taxonomy to XML

Table B-165 *General : Load Esec Taxonomy to XML*

Tag	Value
Severity	
Event Name	loadEsecTaxonomyToXML
Resource	
SubResource	EsecTaxonomyNodeService
Message	Loading Esecurity taxonomy Info to an xml format:

B.19.7 Process Auto Restart Error

This event is sent when a process is stopped. The severity is set to 5 if the process was set to respawn (that is, it is not expected to die). The severity is set to 1 if the process was set to run once.

Table B-166 General : Process Auto Restart Error

Tag	Value
Severity	1/5
Event Name	ProcessAutoRestartError
Resource	Sentinel
SubResource	Process
Message	Process <{0}> [command: {1}] was automatically restarted more than the allowed {2} automatic restart(s) in {3} min. The process will no longer be automatically restarted. Please check process configuration.

B.19.8 Process Restarts

Table B-167 General : Process Restarts

Tag	Value
Severity	
Event Name	ProcessRestart
Resource	Sentinel
SubResource	Process
Message	Process <ProgramName> spawned (command <pID>)

B.19.9 Proxy Client Registration Service (medium)

Table B-168 General : Proxy Client Registration Service (medium)

Tag	Value
Severity	
Event Name	registerClient
Resource	
SubResource	ProxyClientRegistrationService (medium)
Message	Registering new client

B.19.10 Restarting Process

Table B-169 *General : Restarting Process*

Tag	Value
Severity	
Event Name	restartProcess
Resource	SentinelHealth
SubResource	SentinelHealthService
Message	Restarting process <name> on Sentinel server <name> UUID {2}

B.19.11 Restarting Processes

Table B-170 *General : Restarting Processes*

Tag	Value
Severity	
Event Name	restartProcesses
Resource	SentinelHealth
SubResource	SentinelHealthService
Message	Restarting <number> processes: <number> name <name> server <name> server ID <ID>;

B.19.12 Starting Process

Table B-171 *General : Starting Process*

Tag	Value
Severity	
Event Name	startProcess
Resource	SentinelHealth
SubResource	SentinelHealthService
Message	Starting process <name> on Sentinel server <name> UUID {2}

B.19.13 Starting Processes

Table B-172 *General : Starting Processes*

Tag	Value
Severity	
Event Name	startProcesses
Resource	SentinelHealth
SubResource	SentinelHealthService
Message	Starting <number> processes: <number> name <name> server <name> server ID <ID>;

B.19.14 Stopping Process

Table B-173 *General : Stopping Process*

Tag	Value
Severity	
Event Name	stopProcess
Resource	SentinelHealth
SubResource	SentinelHealthService
Message	Stopping process <name> on Sentinel server <name> UUID {2}

B.19.15 Stopping Processes

Table B-174 *General : Stopping Processes*

Tag	Value
Severity	
Event Name	stopProcesses
Resource	SentinelHealth
SubResource	SentinelHealthService
Message	Stopping <number> processes: <number> name <name> server <name> server ID <ID>;

B.19.16 Store Esec Taxonomy From XML

Table B-175 *General : Store Esec Taxonomy From XML*

Tag	Value
Severity	
Event Name	storeEsecTaxonomyFromXML
Resource	
SubResource	EsecTaxonomyNodeService
Message	Storing Esecurity taxonomy Info :

B.19.17 Watchdog Process is started

As the Watchdog process starts, the following internal event is generated:

Table B-176 *General : Watchdog Process is started*

Tag	Value
Severity	1
Event Name	ProcessStart
Resource	WatchDog
SubResource	WatchDog
Message	WatchDog Service Starting

B.19.18 Watchdog Process Is stopped

When the Watchdog service is stopped, the following internal event is generated:

Table B-177 *General : Watchdog Process is stopped*

Tag	Value
Severity	5
Event Name	ProcessStop
Resource	WatchDog
SubResource	WatchDog
Message	WatchDog Service Ended