Bluebird Administrator Reference Guide

Open Source Network Management for the Next Millenium

OpenNMS.org

Table of Contents

B File and Directory Structure

a sect1 title

a sect2 title

another sect1 title

another sect2 title

Preface
Acknowledgments, Copyrights and Terms

Copyrights   Copyright (c) 1999-2000 PlatformWorks, Inc.

Copyright (c) 1996-2000 PostgresSQL, Inc.

Copyright (c) 1994 Regents of California, Inc.

Copyright (c) 2000 Sun Corporation

Copyright (c) 1999-2000 Apache Software Foundation

Useful links for for information

XSLFO Spec, Oct 2000

XSLFO Bible

XSL Info Site

XML Organzation

Docbook Organization

O\'Reilly's DocBook: The Definitive Guide

OpenNMS.org - Bluebird web site

apache.org - home of Xerces, XML4J etc

Sun - home of JAVA

W3C Organization - home of XSLFO and XML

PostgreSQL - home of PostgreSQL

P.1 Compendium of Terms

To help understand the scope of breadth of the Bluebird project, this document defines the high level objectives and philosophy of Bluebird. Rather than reading through reams of paper to try and understand the various pieces, and how they are assembled, we have put this document together.

One of the challenges of a project is developing an internal language used amongst team members which is concise and efficient without confusing new users with lingo. As any complex project, Bluebird has its share of these "buzzwords". Rather than change the established culture, we present a compendium of terms:

P.1.2 Geneal Terminology

Bluebird   The name of the project to build a network and service management application in the open source model.

OpenNMS.org   A project supported by PlatformWorks, Inc. chartered with designing, building and supporting the various projects including Bluebird and JoeSNMP.

PlatformWorks	The incorporated company organized by the founders to fund and support the Bluebird project and OpenNMS.org.
JoeSNMP	A JAVA SNMP library developed by OpenNMS.org and available as LGPL open source code.

P.2 Bluebird Modules

P.2.1 Bluebird Configuration

Poller Package	A poller package is a bundle of configuration information which a distributed poller uses to poll devices and services within its domain. Poller packages are built and stored on the master station and pulled by the distributed poller at startup (or reconfig).
	Poller packages are ASCII text files stored in XML in a well-defined format. Any application which can generate this XML is available to build poller packages, however, graphical tools are available to allow simple drag/drop configurations.
Iceberg	Same as poller package. The name is only of significance to historians of the OpenNMS.org project.

P.2.2 Bluebird Functional Areas

Master Station	The central computer in the network where all the network information resides. It is expected, in this release, that the Master Station is a single computer capable of handling all the load of this information. Master stations can be redundant, but are not distributed.
	Users and administrators log into the master station. They cannot access distributed poller information except via the master station.
Distributed Poller	A remote computer in the network which contains the Bluebird distributed poller programs and one or more poller packages for configuration information. The assumption is that the distributed poller is a Unix, NT or other computer type device. Since Distributed Poller logic is not implemented in network elements, specific computers are selected for the role of distributed poller. The fundamental idea of a distributed poller is autonomy. A poller should be able to operate completely disconnected from the master station after receiving a poller package of information. Distributed pollers report information to the master station in a push/ pull arrangement. Configuration information is pulled from the master station by the distributed poller. Should the master station need to tell a distributed poller of configuration changes, the master station will push an "alert" to the distributed poller telling it to pull down the most current poller information.
Discovery	The process of automatically finding IP addressable devices in the network. Discovery only takes place on a distributed poller. In first release, discovery is implemented as ping sweeps to insure a baseline discovery method.

Filters	When a distributed poller discovers a device, it applies filters to the device to see if it matches. Filters allow unimportant devices to be skipped and important devices to pass through.
capsd	When a device is discovered, it is tested for capabilities against a filter. The process responsible for this test is called capsd or the capability checker.
Capability Checker	Same as capsd.
Service poller	A service poller is responsible for testing the availability of a service on a target device in the network. Examples of service pollers are HTTP, SMTP, DNS, FTP, ICMP and any other service which can be tested externally.
Synthetic Transaction	A service poller tests a network device by "faking" a transaction from a bogus user. For example, for HTTP, a synthetic transaction would test port 80 and send an HTTP GET command. Service pollers are specially built to mimic the dialog of a user service.
eventd	eventd is the event daemon responsible for receiving events from the distributed pollers. eventd uses a push/pull arrangement where events are periodically pulled from a distributed poller. In the vent of a major event, a distributed poller can inform the master station to perform an unscheduled (or asynchronous) upload.
Event Browser	Events in the system (node down, threshold exceeded, etc) are available via the event browser. The event browser is accessible via the real-time console.
Real-time Console	When an operator logs in, access to network management information is via the Real-Time Console (RTC). The RTC amalgamates information from the event stream, the view rules and the category configuration into a simple, histogram presentation of live data.
Rule Builder	Distributed pollers use rules to determine which devices to filter. Since rules are typically complex grammars to decode and understand, the rule builder allows a graphical drag and drop metaphor for building rules. Rules are dragged from the template area and connections are established by drawing lines.

P.2.3 Users, Groups and Views

User	A user is a particular operator who has capabilities assigned by the administrator. A user cannot log into the system without a userID and password.
User Group	Users may be collected together into user groups for ease of administration. User groups can be assigned capabilities which then effect all the users within that group.

Business Views

Bluebird administrators can partition the network into smaller subsets called views. For a large network outsourcer, a view might be a customer or a type of customer. For a smaller Bluebird installation a view might correspond to a sales region or a geographical area. Views can be broken up into smaller pieces called categories.

View Categories

A view can be broken up into smaller pieces called categories. For a view defining a particular region of the network, the categories might be "routers", "ISDN equip" or any other set of devices which can be categorized by rules.

P.2.4 Distributed Poller Configuration

Packages

Packages are a bundle of configuration information stored in an XML structure. Packages are sent to distributed pollers to inform the poller of the desired behaviors, parameters and actions. Poller packages are additive; i.e. you can send more than one poller package to a distributed poller. A distributed poller must have at least one poller package assigned to it do do any useful work.

Calendars

Calendars control the time when a poller ignores a managed device in its domain. Calendars are schedule outages. There are two different types of calendars; repeating and one-time. Repeating is for periodic scheduled down-time, such as backups. One-time calendars are used for special outages which are unplanned but do not want to be included in the down-time calculations.

Service

A service is a type of protocol test; e.g. http, smtp, pop3 etc. Services are assigned to a poller package along with target rules and domain information.

Service Pollers

A service poller is a software component on a distributed poller which is responsible for performing a service test. The service poller will probe a device, check for service availability, retry failures and report problems to the master station.

P.2.5 Data Storage

XML

XML is the Extensible Markup Language. XML can define any data structure using tags. HTML is a poorly defined implementation of XML. Eventually, most of the web will migrate from ill-structured HTML to XML with styles.

Bluebird uses XML almost exclusively for communication between and within the various components of the product. To configure Bluebird, you can use the Graphical Tools or generate the appropriate XML structure programmatically or by hand.

ODBC

ODBC is a method of accessing data from a (typically) SQL database. Applications written to ODBC have a level of independence which improves the ability to change database systems without changing existing code.

P.2.6 Data Presentation and Reporting

XSL   XSL is the Extensible Style Language. XML does not define the look and feel of data. Instead, XML uses styles, defined by XSL, to render the data into a particular presentation. This allows data to reside independently from the visual representation.

Bluebird uses XSL to generate reporting and Web presentation. Eventually, when web browser support full implementations of XSL, the need for the Bluebird Data Presenter will go away or be reduced.

FOT   When XSL renders an XML document to apply styles, it generates an intermediate file called an Formatted Object Tree or FOT. Bluebird generates FOTs when building reports against XML report data and applies the user defined styles.

- Page 12 -

Chapter 1
Theory of Operation

1.1 Fundamental Tenets

Bluebird is a service and network management platform for automatic node discovery, network services monitoring, operator notification of problems, events consolidation, automatic action launching and service level performance monitoring.

Scalability

The architecture must be scalable to handle the small to medium companies on a single computer and handle multiple poller approaches in a large international environment. To accomplish this, Bluebird uses a distributed architecture of a master station and one or more distributed pollers. Distributed Pollers can reside on the same processor as the master station.

Resilient and Autonomous

The architecture aspires to the goal of allowing a distributed poller to communicate with the master station only once a day. Though most people won't do this, the architecture becomes resilient to long term disconnects from the master station to a distributed poller.

For example, assume a master station process requests new nodes from a distributed poller every 5 minutes. Should the connection from the master station to the distributed poller fail, the first time it will ask for 5 minutes of data. The next time the master station talks to the distributed poller, it will ask for 10 minutes of new nodes. The third time, it will ask for 15 minutes of data etc. When the connection between the two is established, it will download all the missing data and they will be back in synch.

Open Architecture

Using JAVA, XML, SOAP, JSDT and servlets to communicate between the various components allows an open approach where any application which can understand XML can communicate with the most intimate areas of the product. For example, if an application would like to get events from a distributed poller, it merely queries the servlet on the distributed poller using a well defined XML data stream. The poller returns the requested events in XML. This approach opens the architecture to non-programmers and scripters. The only requirement is an XML parser.

Service Management

Rather than focus solely on ICMP for determining availability of the network, Bluebird uses "synthetic transactions" to probe various services on a device. Bluebird views a network device as a series of services, including ICMP, SMTP, DNS, HTTP etc.

Service Level Orientation

Bluebird eschews the idea of topological presentation of the network for a rule-based, statistical view of the network. Rather than watching red and green icons blink on and off, Bluebird presents problems as histograms effecting service levels. If a machine fails but is not of interest, the service levels are not effected and notification is suppressed.

Deterministic   For limiting traffic generated by a network management system, Bluebird deploys "Bandwidth Trolls" to provide deterministic control of critical polling functions. Bandwidth Trolls throttle back Bluebird processes to pre-determined, user-defined traffic levels.

Discovery   Network devices are automatically discovered using IP, (or SNMP or other protocols) queried for applicability, and added to an object database.

Thread Technology   Pollers and other real-time components take advantage of threaded technology to eliminate queuing overhead, improve responsiveness and take advantage of multi-processor hardware technology.

Graphical   Bluebird administration utilizes graphical metaphors for all configuration and maintenance tasks to reduce the time to learn and understand the product.

Filters   Bluebird administration is rule based to reduce administration and maintenance of the platform.
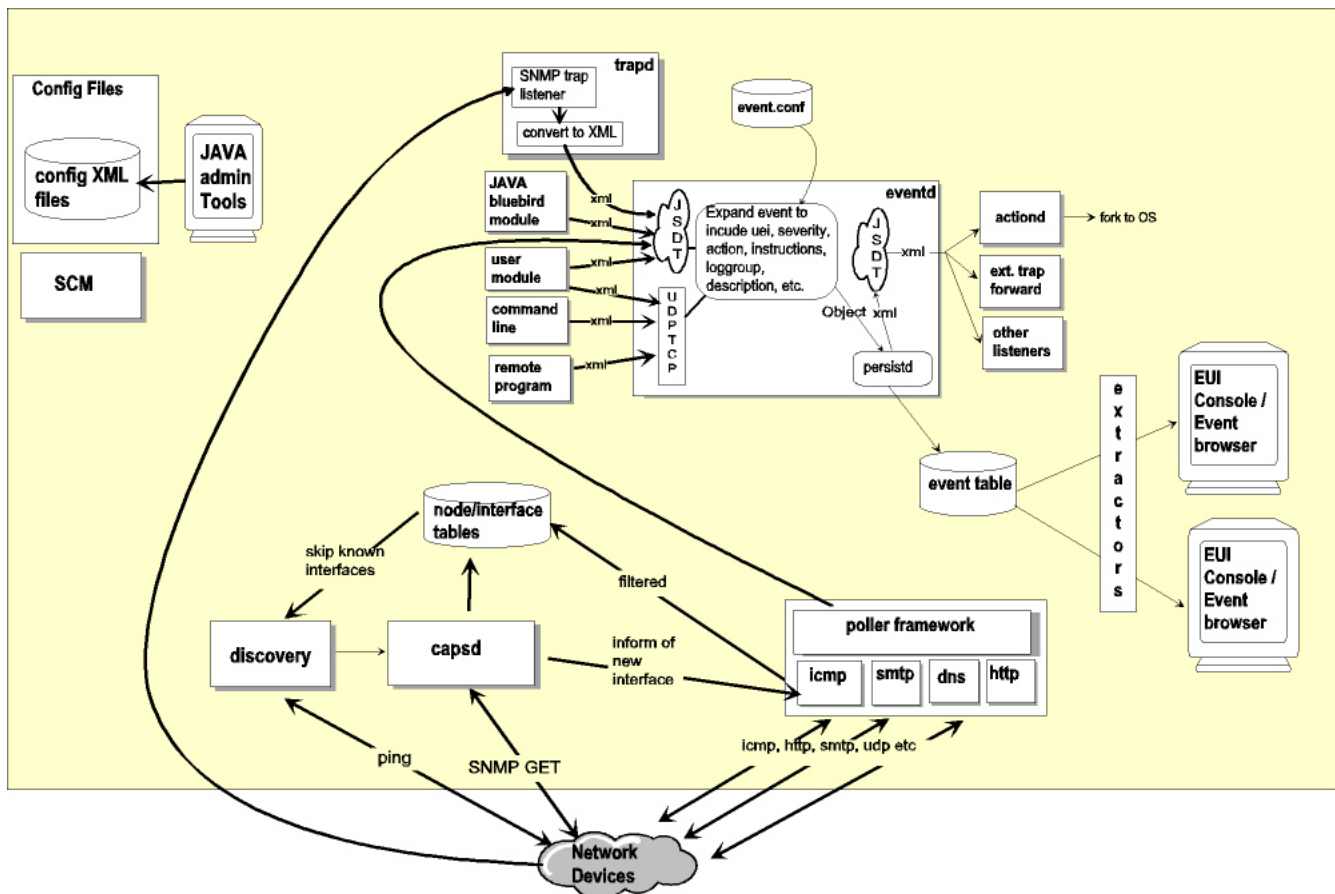
1.2 Bluebird Functional Overview

Figure: Bluebird Overview

Each functional area of the architecture is described below:

Configuration Files	Configuration files used by the Bluebird system control the behavior and actions of the various parts of Bluebird. Configuration files are maintained and stored in XML format on the file system of the Master Station but are pushed to the Distributed Pollers as necessary.
	Configuration files are assembled into bundles called poller packages. These packages are assigned to specific distributed pollers to control how the distributed poller operates. Packages are re-usable and allow for redundancy and ease of administration.
SCM (Service Control Manager)	SCM is responsible for starting, stopping and controlling the various Bluebird processes (services) on the Master Station and the distributed pollers. SCM uses the serviceconfig.xml file to determine the services to run, how to run them and dependencies between services.
JAVA Admin Tools	The Administrator tools are used to graphically manipulate the configuration files. One could edit the config files directly with editors or perl, but the admin tools are designed to be user friendly ways to configure the system.
Discovery	Discovery performs an advanced ICMP sweep of devices in a discovery range. All responding devices in the range are then tested against discovery filters to see if they are "interesting" to Bluebird. Discovery is a threaded system allowing pools to pollers. Discovery is limited by Bandwidth trolls which control how much traffic is consumed by discovery.
capsd	capsd is the capability checker. When a device is found by discovery, capsd checks it against the discovery filter. If it passes through the filter, it is added to the object database and added to the known node list.
Service Poller/Poller Framework	Service pollers provide the actual probing of the service under test. Initial service pollers include ICMP, HTTP, SNMP, SMTP, DNS, FTP and others. Additional pollers can be bolted into the architecture as needed.
EUI Extractors	When the Real-time Console (EUI) and event browser need information about the network, they register with an extractor channel to receive network updates. The extractors deliver a "tree" of information in XML format so that additional EUIs can be built and integrated.
	Extractors are shared by multiple EUIs so that additional overhead is not incurred for users viewing similar information.
trapd	Trapd is the SNMP trap listener which waits on a UDP port form messages. When received, trapd converts the event into XML and sends it to the eventd process.

eventd   Events are processed and expanded by eventd. An event is received by eventd from one of several sources; trapd, a third party application, a Bluebird module or from a TCP/UDP port from a different computer. Once received, the event is "expanded"; i.e. additional information for the the event is appended. Additional information would typically be event description, operator instructions, automatic actions, log groups and many other fields.

After the event is expanded, it is sent to persistd for committing to the database.

persistd   persistd takes a fully expanded event from eventd, writes it to the database and broadcasts it to the various event listeners for processing. persistd insures that the event gets sent to the database before event listeners receive it.

actiond   actiond is one of the event listeners. actiond takes an event and looks for automatic actions. If enabled for that event, the action is launched and tracked. actiond uses threads to allow processes to run in parallel.

Chapter 2
Installing and Removing Bluebird NMS

The current version of Bluebird is a single computer version of the Bluebird distributed architecture and requires only a single computer to run. Installation is significantly simplified in an environment where the distributed poller and master station run on the same computer.

2.1 Installation Roadmap

The installation process involves identifying the software and hardware pre-requisites, installing the software dependant packages, installing the Bluebird package, configuring the Bluebird poller packages and starting up Bluebird's SCM (Service Control Manager).

Bluebird contains small portions of code that are specific to an operating system and are not in JAVA. These code modules are installed during the installation process and cannot be moved between different operating system architectures.
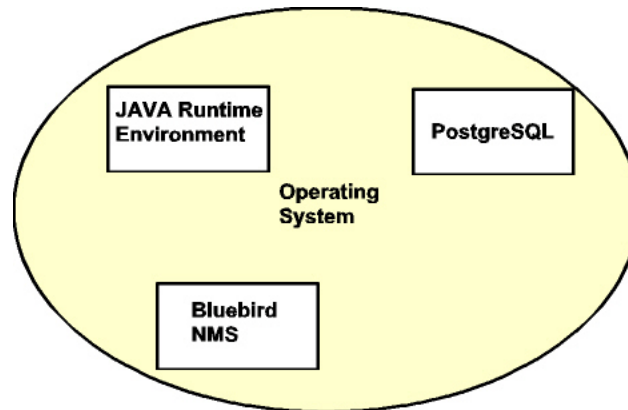

Figure: Installation Roadmap

2.2 Installation Requirements

2.2.1 Shipped Distribution Packages

Some software packages required by Bluebird are not shipped on the distribution package due to licensing or practical considerations. The following table shows the packages shipped on the Bluebird distribution media:

Software	Version
Bluebird	.8
PostgreSQL for NT	7.02

2.2.2 Software Dependencies

Bluebird requires additonal software packages to operate. Depending upon the base operating system of the target computer, these packages may already be installed or may need to be installed or updated. The following table shows all the software packages required by Bluebird. If the installation target computer does not have all of the required packages, contact the provider referenced in the

"Source" column to get the latest version.

Software	Source	Version	Patches
Operating System	Redhat Linux	6.2	
Operating System	SUSE Linux	7	
Operating System	Windows NT	4.0	SP 6
JAVA Runtime Environment	Sun Microsystems	1.3	
JAVA Runtime Environment	IBM	1.3	
Bluebird NMS	OpenNMS.org	0.8	
PostgreSQL	Postgres.org	7.02	

2.2.3 Hardware Requirements

Hardware requirements are completely dependant upon the scope and scale of the network being managed. A low end laptop with 64 meg could potentially manage a small workgroup or 20-50 devices. A more practical environment would involve a significantly large computer with more memory, disk and CPU.

Resource	Minimum	Recommended Min
Memory	128 MB	512 MB
Disk	256 MB	10 GB +
CPU	Pentium II	Multi Pentium PIIIs
Network Interface	Ethernet	Fast Ethernet

2.3 Installation Procedures

Bluebird is installed differently depending upon the target computer. For Linux and Unix based operating systems, Bluebird installs from the user shell. For Windows/NT based systems, Bluebird is installed from the NT command line.

2.3.1 Linux based distributions

Bluebird on Linux distributions is in RPM format. To install, perform the following steps:
1.   Acquire the install RPM file either from installation media or across the Internet.

2.   Move the RPM file to a temporary location on disk having 20-25 MB of free space.

3.   Run the installation procedure:
      rpm -i RPMfilename.rpm

2.3.2 Windows/NT based distributions

Bluebird for Windows/NT is in Windows exe format. To install, perform the following steps:
1.   Mount the media containing the install exe file or copy the exe file across the Internet.

2.   Run the installation procedure:
      <installdrive>:\installbb.exe

Chapter 3
Configuring Users, User Groups and Views

Users, user groups and views are configured by the administrator. A user cannot access information in Bluebird unless that user is authorized by the administrator to access a view. A view is a way to visualize the network.

3.1 Configuring Users

An operator of Bluebird must be assigned a user id and password prior to using the interface. The user id and password are assigned by the administrator using the Configure User/Group/View interface.

3.2 Running the Configure User/Group/Views Panel

The administrator builds new users, assigns them to user groups and enables users to access views. To configure users, groups and views:

1. Log in as the administrator. The Administrator Main Panel appears:

Figure: Administrator Main Panel

2. Select the "Configure User/Groups/Views" icon from the Administrator Main Panel. The Configure User/Group/View Panel appears
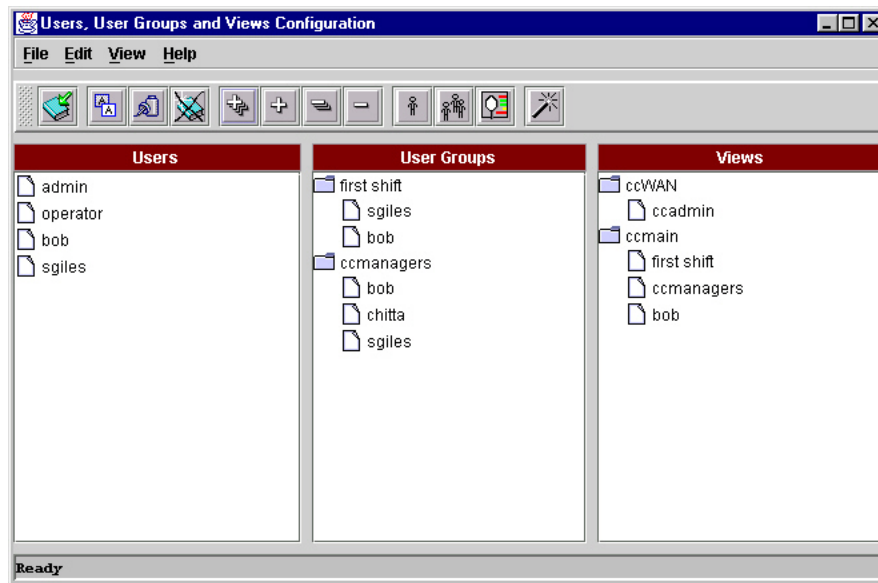
Figure: Configure User/Group/Views panel

3.3 Understanding Users, User Groups and Views

Users, user groups and views control the end user interface into the Bluebird system. Efficient organization allows rapid control and maintenance of the various users and what they have permission to access.

3.3.1 Users

Before a person can use Bluebird, they must be assigned a unique user name and password. The user and password are required for login and determine what a particular user has permission to access. The administrator is the only user who has permission to add users and change their views.

When a user is initially built by the administrator, that user has no access to any information in the Bluebird system. Users can only access information via data windows called "views". When the administrator assigns a view to a user, any information provided via the view is available to the user.

A user can be assigned to one or more views which provides a superset of information of each view. For example, if the "East Coast" view permits a user to see nodes in New York and the "West Coast" view permits a user to see nodes in Los Angeles, then if both views are assigned to a user, that user will see devices for both New York and Los Angeles.

A user cannot be logged onto the system more than once. Therefore, each concurrent user must be assigned a unique user id and login. Sharing of user IDs amongst people in the same group is not permitted.

3.3.2 User Groups

To make it easier to configure and administer large groups of users, users can be assembled into collections called groups. Groups can be organized in almost any fashion. The following are some simple examples of ways to organize users:

·**Time** - if users perform different tasks on third shift than first, user groups are built called "first shift", "second shift" etc. Users are then assigned to a particular shift group. The entire group is then assigned permissions and access to views.

·**Network Region** - if users perform different tasks for a specific area of the network, user groups are built called "east coast", "west coast", "asia" etc. Users are then assigned to an area of the network by assigning them to the appropriate user group. The group is then assigned permission and access to appropriate views.

·**Job Function** - if users perform different tasks depending upon the job function of the user, user groups are built called "administrators", "managers", "staff", "contractors" etc. Users are then assigned a set of views by assigning them to the appropriate user group. The group is then assigned to the appropriate views.

·**Device Type** - if users perform the same type of task for a set of devices, user groups are built called "routers", "hubs", "ATM switches" etc. Users are then assigned to the appropriate groups by the administrator. The group is then assigned a set of views which the users can access.

3.3.3 Views

A view is a data window into a subset of the network and services. When a view built, it typically defines an organized area of the network. Since views are rule driven, sophisticated relationships can be built which accurately describe the desired information.

Views are assigned to a user via the administrator; a user cannot remove or add views to their user interface. The user can switch between views as desired so the administrator can assign views to a user which overlap or are supersets of each other. For example, a view can be built for all the devices on the east coast as well as a view for New York, Boston, Atlanta and Washington DC. The user can then see the entire east coast in one view and a specific city in one of the other views.

The following are some simple examples of views:

·**Geography** - A view can depict an area of the network, typically using IP Address. This is useful for viewing devices within a "management domain" where responsibility are assigned.

·**Filtered** - Often you want to view an area of the network but without specific types of devices. A view of this type allows users to focus on only the critical devices of interest.

·**Organizational** - Often, the network is broken up by company organizations; engineering, finance, research, manufacturing etc. Viewing the network this way allows different users to focus on only those devices within that group.

·**Device Function** - Often, the network contains core backbone equipment and regional equipment. Some users need to see only backbone WAN equipment, others need to see only region equipment and yet other users need to see both.

3.4 Using the Configure Users, User Groups and Views Panels

Users, user groups and views are configured using the Configure Users, Groups and Views tools. This tool allows new users to be added and maintained and assigned to views.

When selected the following panel appears:
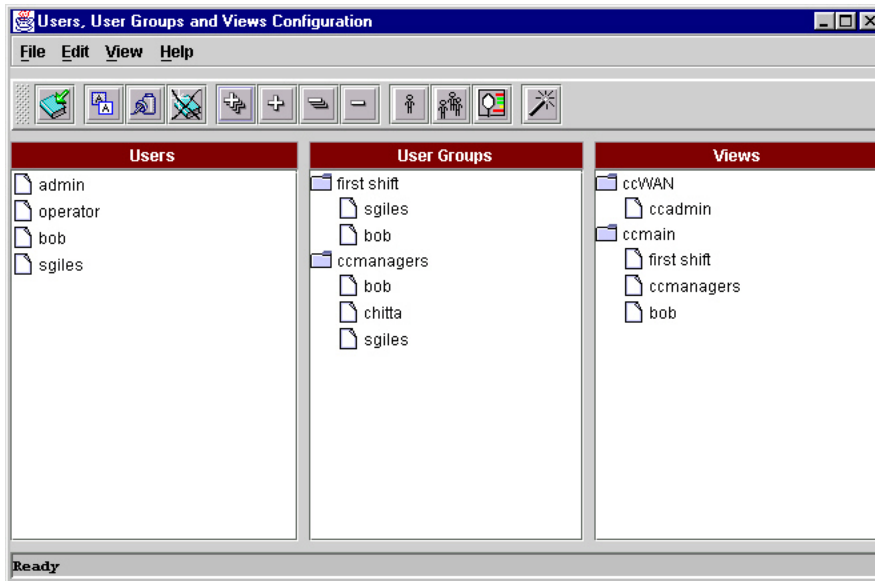
Figure: Configure User/Group/Views panel

The Configure User, Groups and View Panel contains the following areas:

Menu Bar	The menu bar allows operations from the keyboard and mouse as well as view appearance and help.
Icon Bar	The icon bar contains commonly used tools and shortcuts. Adding new users, groups, views, expanding the trees and the wizards are contained here.
"Users" Area	The Users area contains all of the configured users which have access to the system. Users are shown in a scroll list in the order in which they were configured. The label of the user is the valid user ID used at login time.
"User Groups" Area	The User Groups area contains all the of the user groups currently configured on the system. Each user group contains a branch below it which shows the members of that group. Only a single level of users can be contained in a user group.
"Views" Area	The Views area contains all of the views configured on the system. Each view contains a single branch showing which users or user groups have access that that view. If a user is not contained in a user group under a view, then that user will not see that view when they use Bluebird.
Status Bar	The status bar show information about the panel. The last operation is shown in the status when a normal operation has occurred.

3.4.1 Icon Bar

The icon bar contains the common operations for manipulating users, user groups and views.

Icons in the icon bar will display textual hints when the mouse moves over the icon. Icons can also be selected using the keyboard using the TAB and arrow keys. There are not accelerator keyboard keys for icons.

The icon bar is detachable and can float above the rest of the panel as desired. To detach the iconbar,

simply drag the icon bar from the panel. This action may differ slightly depending upon the "Look and Feel" settings. The floating icon bar can be reattached to the panel by closing the floating icon bar.

Figure: Floating Icon Bar

Icon Bar Tools

- Save the current configuration. Allows the changed configuration to be make permanent. As changes are made, they are stored in a temporary location until a save if performed. Exiting the panel without saving changes will generate a confirmation message warning about the loss of data.

- Copy the selected user, group or view to the clipboard.

- Paste a user, group or view from the clipboard.

- Delete the selected user, group or view.

- Expand the tree branches for the entire panel.

- Expand the selected branch.

- Collapse the tree branches for the entire panel.

- Collapse the selected branch.

- Add a new user to the "Users" area.

- Add a new user group to the "User Groups" area.

- Add a new view to the "Views" area.

- Launch the New User Wizard.

3.4.2 "Users" Area
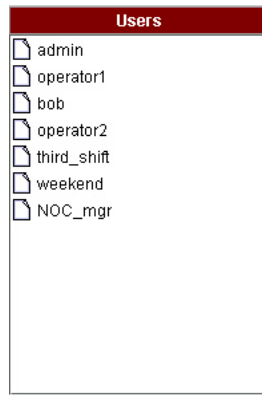
Figure: "Users" Area

An individual must have a valid user ID and password to get access to Bluebird. If a user is not listed in the "Users" area of the panel, they do not have access to Bluebird information.

Users are added to the "Users" area using the "Add a New User" icon in the icon bar. When selected, the "Add User" Panel appears:

Figure: "Add User" Panel

"Add User" Fields

The following fields are defined in the "Add User" Panel:

User ID

This is the string used in the login panel to identify the user at login time. The ID can be any string of characters, however, it is recommended that user names be limited to alphanumeric A-Za-z0-9   characters.

Full Name

The full name of the user. This is a free-form field which can contain any textual data.

Password

The password of the user. This field is used at login time to allow access to the assigned user views. A password is required and cannot be blank.

Confirm Password

The confirmation password of the user. This field is compared to the password field to insure that they match. If the passwords do not match, the user cannot be entered.

Comments

This is a free form text field for any notes, comments or other information about the user.

Modifying Users

Users are changed by double clicking on the user in the "Users" area or right clicking the mouse on the user. When right clicked, the following menu appears:
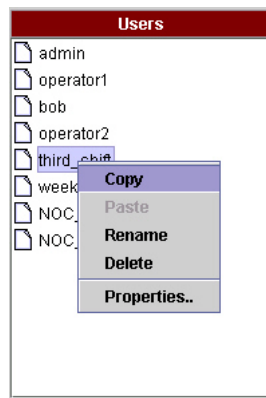
Figure: "Users" Area Right Click

The following options are defined in the right click:

|	
Delete	The "Delete" operation removes the user. See *Deleting Users* below for more information.
Copy	The "Copy" operation takes the user information and adds it to the clipboard. If another copy operation was previously performed, the previous information is lost without warning.
Paste	The "Paste" operation takes the user information from the clipboard and pastes it to the "Users" panel. The "Paste" option is unavailable unless there is a valid user in the clipboard.
Rename	The "Rename" operation allows the User ID to be changed without going into the "Properties" operation. The other fields are not effected by the rename operation.
Properties...	The "Properties" operation allows the User information to be modified. When selected, the "Modify User" panel appears.

Deleting Users

A user can be a part of one or more user groups and be assigned to one or more views. Deleting a user removes the user, any membership in user groups and assignments to views.

When a delete operation is performed, the following confirmation message appears verifying the delete:
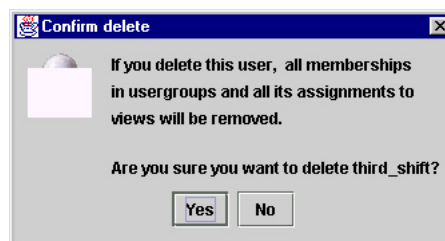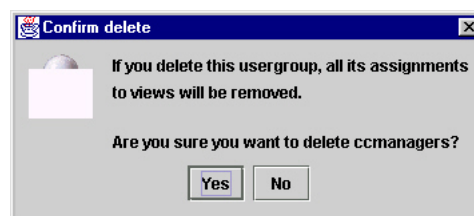

Figure: Delete Confirmation

User Drag Operations

Users are dragged from the "Users" area onto user groups and views. When dropped onto a user group, the user becomes a member of that group. Any privledges and access granted to that group become available to the users in that group.

3.4.3 "User Groups" Area
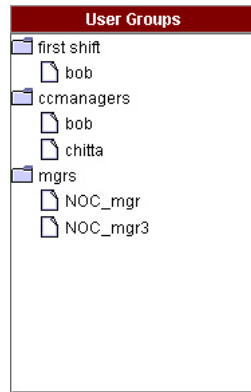
Figure: "User Groups" Area

Users can be collected into groups called user groups. User groups are not required for configuration, however, it is strongly recommended that similar users be organized and assigned views.

When the "Add User Group" icon is selected form the icon bar, the "Add User Group" panel appears. Since a user group is a container for organizing users, the fields in the panel are simply descriptive and are not used elsewhere.
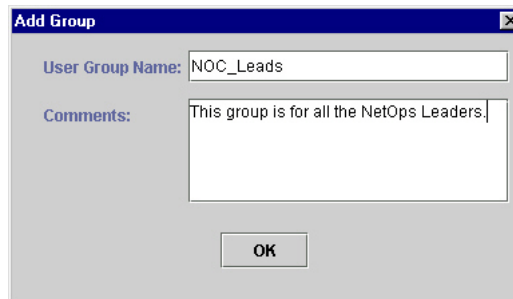
Figure: "Add User Group" Panel

"Add User Group" Fields

The following fields are defined in the "Add User Group" Panel:

User Group Name	The name of the user group. This name is used in the view area to identify the users who have access to that view.
Comments	Textual comments about the user group. These comments are informational and for documentation purpose. Comments are not used elsewhere in the system.

Modifying User Groups

User groups are changed by double clicking on the user group in the "User Groups" area or right clicking the mouse on the user group. When right clicked, the following menu appears:
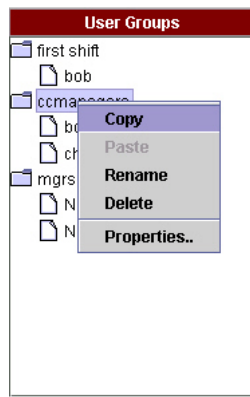
Figure: "User Groups" Area Right Click

The following options are defined in the right click:

Delete   The "Delete" operation removes the user group. See *Deleting User Groups* below for more information.

Copy   The "Copy" operation copies the user group to the clipboard.

Paste   The "Paste" operation pastes the user group information from the clipboard to the "User Groups" panel. The "Paste" option is unavailable unless there is a valid user group in the clipboard.

Rename   The "Rename" operation allows the User Group Name to be changed without going into the "Properties" operation. The other fields are not effected by the rename operation.

Properties...   The "Properties" operation allows the user group information to be modified. When selected, the "Modify User Group" panel appears.

Deleting User Groups

A user can be a part of one or more user groups and be assigned to one or more views. Deleting a user removes the user, any membership in user groups and assignments to views.

When a delete operation is performed, the following confirmation message appears verifying the delete:

Figure: Delete Confirmation

User Group Drag Operations

User Groups are dragged from the "User Groups" area onto views. When dropped onto a view, the user group becomes a member of that view. All access granted by that view become available to the users in that user group.

3.4.4 "Views" Area
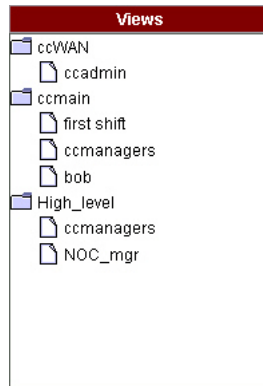
Figure: "Views" Area

A view is a rule based window into a subset of the network. Views are assigned by the administrator to a user.

When the "Add New View" icon is selected form the icon bar, the "Add View" panel appears.

Figure: "Add View" Panel

"Add View" Fields

The following fields are defined in the "Add View" Panel:

View Name	The name of the user group. This name is used in the view area to identify the users who have access to that view.
View Title	The full description of the view.
Comments	Textual comments about the view. These comments are informational and for documentation purpose. Comments are not used elsewhere in the system.

"Configure View Categories, Thresholds and Rules" Button	Since a view is a subset of the managed network, a view has a series of rules which define what devices to show and where. To configure the rules for the current view, select the "Configure View Categories, Thresholds and Rules" button from the panel. The rule builder will launch. For more information, refer to the section on the rule builder.

Modifying Views

Views are changed by double clicking on the view in the "Views" area or right clicking the mouse on the view. When right clicked, the following menu appears:
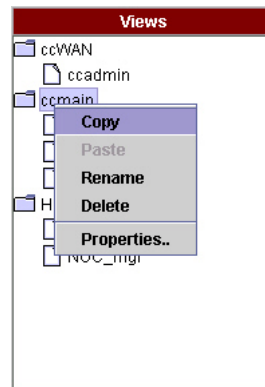
Figure: "Views" Area Right Click

The following options are defined in the right click:

Delete	The "Delete" operation removes the view. See *Deleting User Groups* below for more information.
Copy	The "Copy" operation copies the view and it's children to the clipboard.
Paste	The "Paste" operation pastes the view information from the clipboard to the "Views" panel. The "Paste" option is unavailable unless there is a valid view in the clipboard.
Rename	The "Rename" operation allows the View Name to be changed without going into the "Properties" operation. The other fields are not effected by the rename operation.
Properties...	The "Properties" operation allows the view information to be modified. When selected, the "Modify View" panel appears.

3.5 Building Rules with the View Rule Builder

When a view is built, it consists of one or more categories. Each category has a rule assigned which determines which devices are included in that category. The collection of all categories and rules is called a view.

The Rule Builder is a graphical tool for generating syntactically and semantically correct rule grammar. Sophisticated users can type rules directly into the rule area or use the drag/drop capabilities to let the rule builder create the grammar.

Nodes managed by Bluebird enter the Rule Builder at the Source and must leave at the Sink to be included in the View. Rules are represented as flowing through a series of pipes which have filter at various junctions. A series of tests one after another performs the logical boolean "AND" function; i.e. they must pass through this pipe and that pipe. Branches in the pipe represent boolean logical "OR"s; i.e. they may pass through this pipe or that pipe.

3.5.1 Understanding the View Rule Builder

Views and view categories are used to organize devices, interfaces and services in the network into more managable groups.

View Architecture

The structure of how views are organized is shown below:
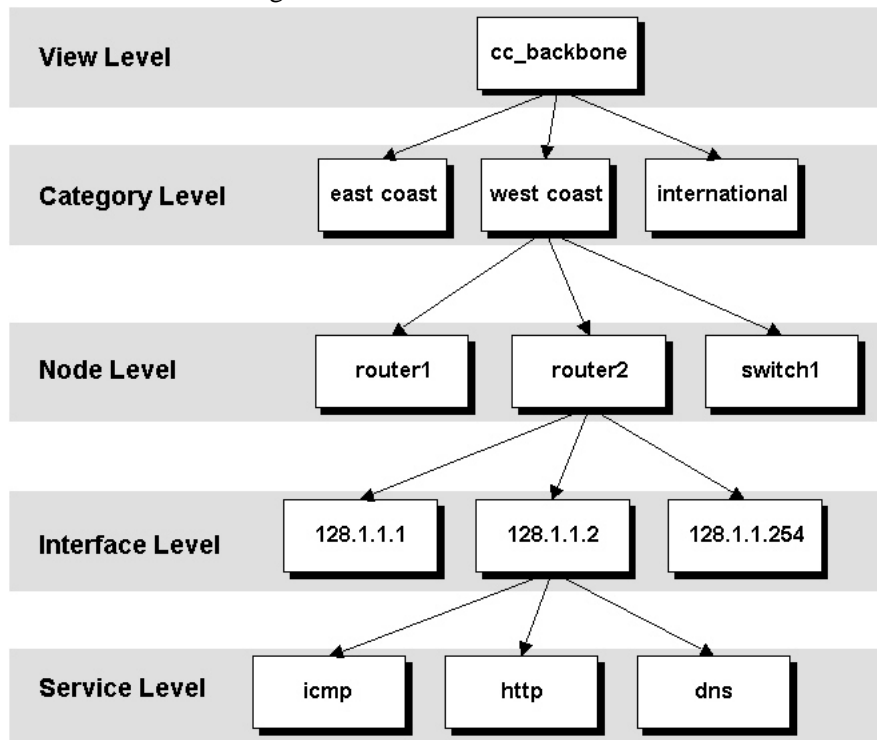
Figure: View/Category Structure

Each level of this structure is defined:

View Level	The views in the view level are configured by the administrator to organize what an operator sees. Views are (typically) shared between users to allow teams to work on areas of the network together. Views are assigned to operators or operator groups.
Category Level	Categories are containers within Views which subdivide the nodes into smaller, more manageable groups. For a node to appear in a catgory, the node must have passed through the view filter.
Node Level	Nodes for that category appear at this level. A node is defined as a set of interfaces.

Interface Level	Interfaces are addressable entities in a node; in fact, a node may have many interfaces (e.g. a router). An interface can have many different services running on it.
Service Level	Services are pollable protocols which run on an interface of a node. Services are reachable or unreachable.

Relationship between the Real-time Console and the View Rule Builder

Using information configured in Users/Groups/View and the Rule Builder by the administrator, the Real-time Console (RTC) dynamically builds the view folder tabs and category histograms. In the following example, the administrator has built 3 views for this user. For the currently selected view folder tab, there are six categories configured. Each category has a set of filters which determine which nodes appear and which ones do not.
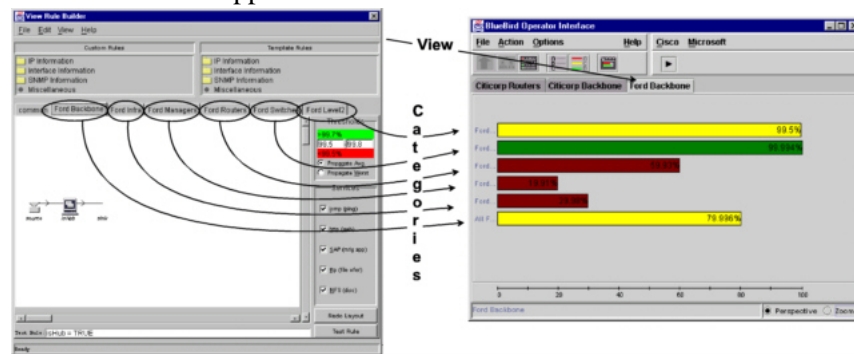
Figure: Real-time Console Structure

3.5.2 Launching the Rule Builder

The View Rule Builder is launched by:

1. Log in as administrator

2. Select "Configure User/Group/View" from the Administrator Panel

3. Right click on the desired View in the View Area

4. Select [Properties...]

5. Click on the "Configure View Categories, Thresholds and Rules" Button

When selected, the Rule Builder Appears:
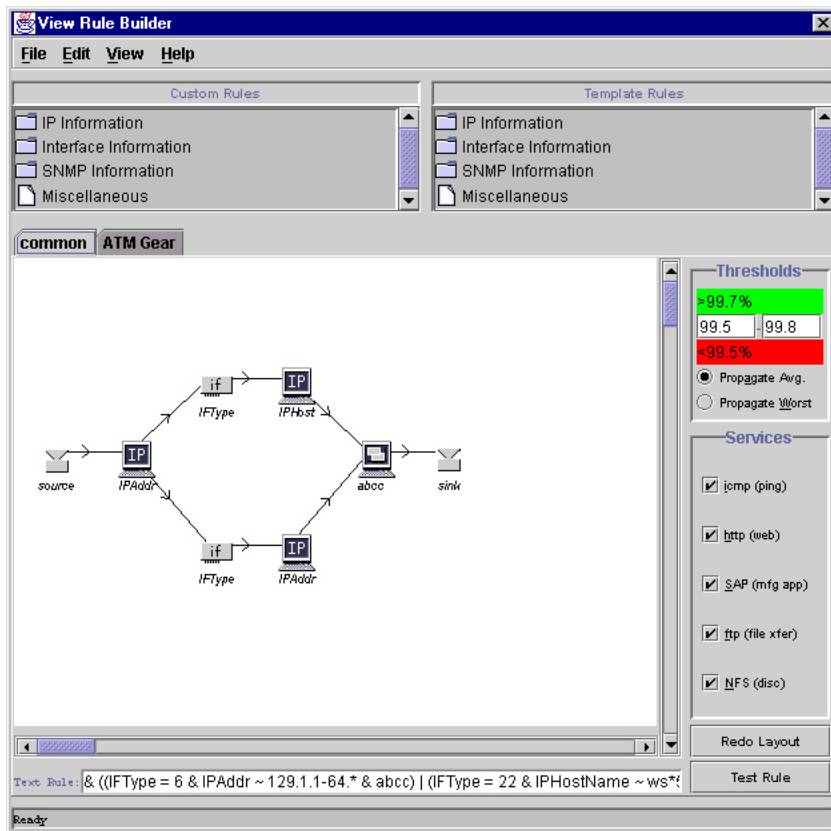
Figure: Rule Builder

3.5.3 Using the View Rule Builder

The Rule Builder Panel consists of seven areas; Custom Rules area, Template Rules area, Category Folder Tabs, Rule Draw area, Thresholds Area, Services area and the Text Rule field. Each area is described below:
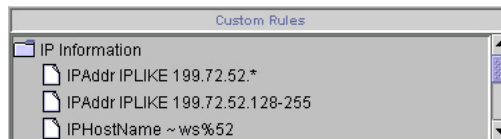
**Custom
Rules Area**

Figure: Custom Rules

The Custom Rules area holds specific rules which have been customized by the user to the environment. Custom Rules can be dragged onto the rule pallette from the Custom Rules.

To modify Custom Rules in the Custom Rules Area, right click on the rule in the Custom Rules area and the following pop-up will appear:
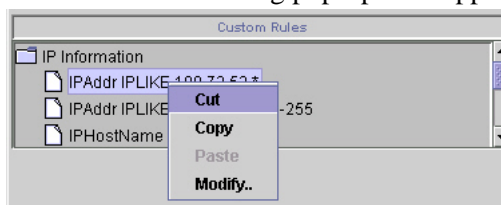
Figure: Modifying Custom Rules

**Template
Rules Area**

Figure: Template Rules

The Template Area holds template rules which can be dragged onto the Custom Rules Area. Template rules are used as basic building blocks for custom rules. Everyone shares the same template rules.

When dragging a template rule onto the Custom Rules area, be sure to drag onto the folder for the same type of Rule; i.e. drag an "IP Information" rule onto the "IP Information" folder in the Custom Rules. For example:
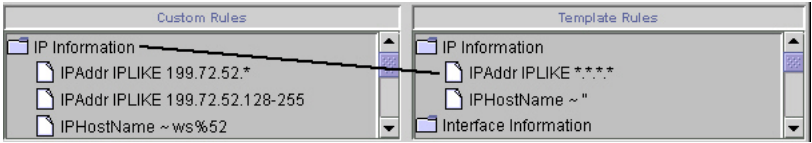

Figure: Dragging Template Rules

**Category
Folder
Tabs**

Figure: Category Folder Tabs

Nodes are organized into groups called categories. Categories are useful for managing large numbers of devices into smaller, more managable categories. A view must contain at least one category to allow nodes to appear to users of that view.

To delete, add or modify View Categories, right click on a folder tab. The following pop-up panel will appear:
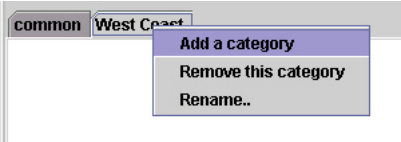
Figure: Modifying Folder Tabs
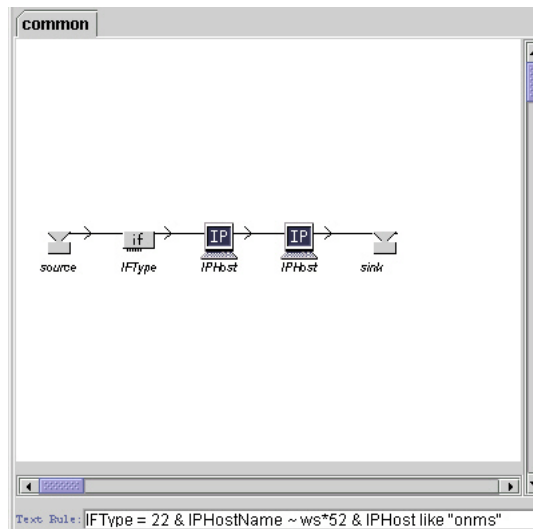
**Rule Draw
Area**

Figure: Rule Draw Area

The Rule Area is a free-form area where rules are dragged and connected to form complex expressions of rules.

Each icon in the Draw Area is an expression in the form of "variable operator value"; e.g. "IPAddr iplike 199.72.52.*". To conserve space on the screen, long expressions are shortened. To view the entire expression, flyt he cursor over the icon and the full rule will appear.

Categories are shown in the folder tabs at the top of the panel. Every folder tab has a different Rule Draw Area and Text Rule.

**Threshold
Area**
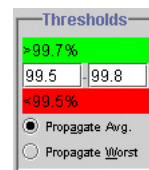
Figure: Thresholds Area

Thresholds define the status of the category based upon a three tiered threshold system. Values greater than the higher threshold number are in the normal range. Values below the lower threshold number are in the Critical range. Values between the higher and lower threshold numbers are in the warning range.

The radio button for "Propagate Average" and "Propagate Worst" define whether the histogram value reflects either the average availability for the node or the worst availability.

"Propagate Average" and "Propogate Most Critical" are currently undefined.

Services Area

Figure: Services Area

The Services area defines which services are to be included in the evaluation of uptime. If a device has multiple services, but only one service is of interest, the services radio buttons allow the administrator to restrict the services the category contains.

The options in the Services area are currently unused.

Text Rule Field

Figure: Text Rule Field

The Text Rule is the actual grammar for the rule drawn in the Rule Draw area. Grammar can be entered into the Text Rule area or built using the Rule Draw area. The Text Rule is only updated when the "Redo Layout" button is pressed.

Rule grammar is defined in Chapter 6.

3.5.4 The "Common" Category Folder Tab

Often, all the categories within a view use the same base rule. For example, for a view of all Cisco Routers, all the categories build upon a rule of ((SNMPsysDescr ~ 'cisco'). Rather than entering that base rule with a logical "and" for every category, the "Common" folder tab allows a rule to apply to all categories.

In the following figure, the common rule defines that categories "East Coast" and "West Coast" will only contain devices which match the common rule of SNMPsysDescr ~ 'cisco'.
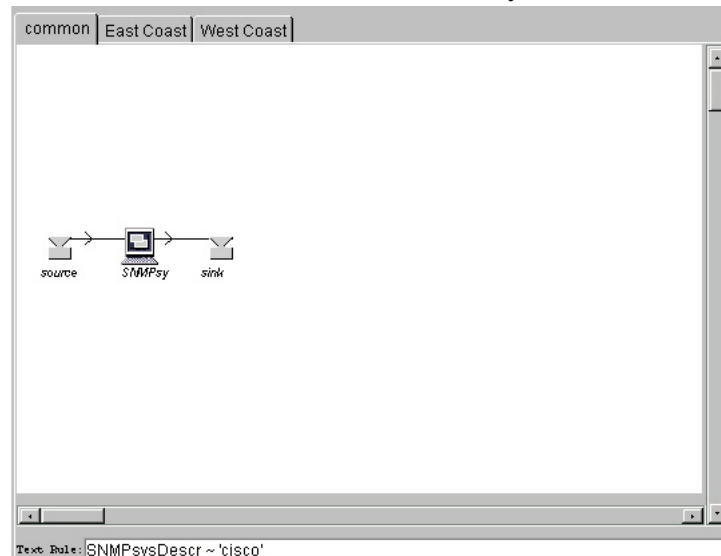

Figure: Common Rule

3.5.5 Step-by-step Examples of Building Rules

The process of building rules is an accretive one. Rules build upon each other by using the template rules to build more and more complex expressions.

Step 1 - Nodes enter the test area at the Source and (may) leave at the Sink. If a node can successfully pass through any paths to the Sink icon, then the node will be part of that Category.

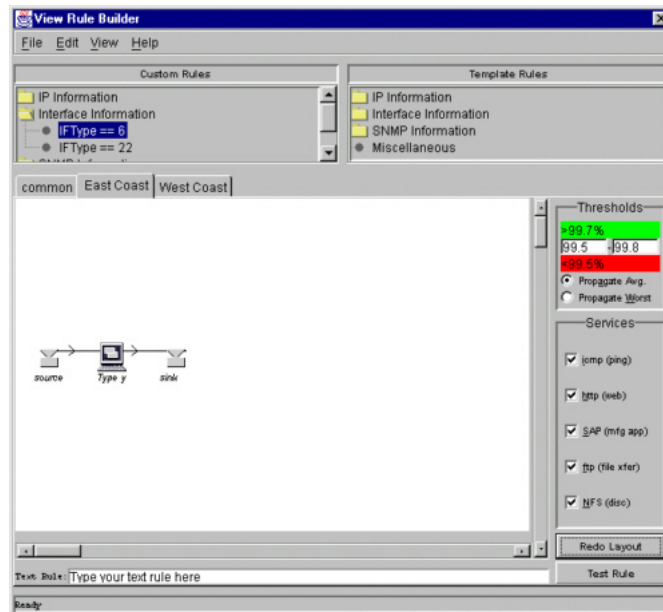When the rule builder is first launched, the rule is set to "type your text rule here". The figure is shown below:

Figure: Step 1 - The Default Rule

Step 2 - We have right clicked on the default icon. A popup panel displays the option to join, cut, copy or modify the rule. The following shows a right click on the default rule
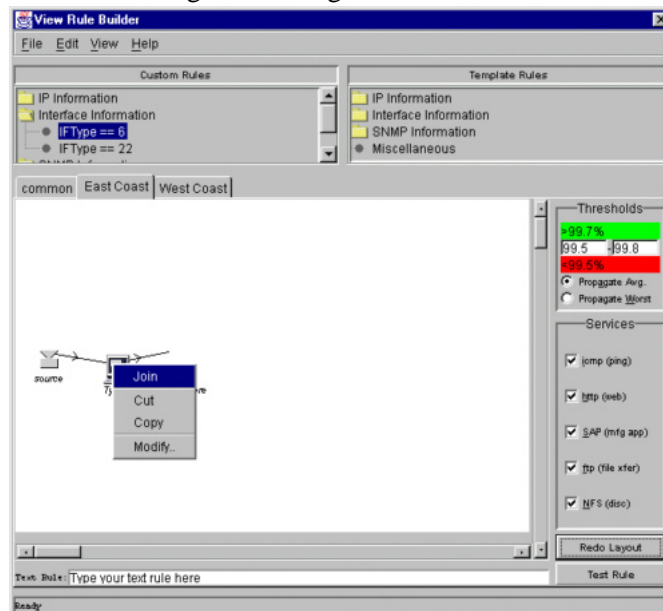
Figure: Step 2 - Right Clicking a Rule

Step 3 - We have chosen to modify the rule and the following rule panel appears.

Figure: Step 3 - Viewing the Default Rule

Step 4 - We typed in a new rule to limit the IP addresses of devices in the category to only 199.72.52 network number. The resultant rule is shown in the figure below:
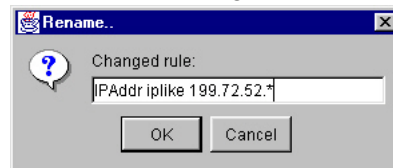

Figure: Step 4 - Changing the Rule

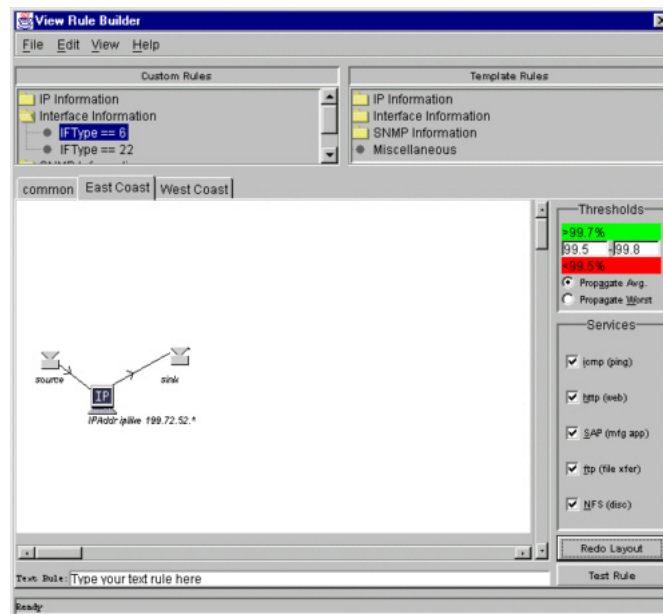**Step 5** - After clicking on the [OK] button, we now view the changed icon. The changed icon is shown in the figure below:

Figure: Step 5 - Viewing the Changes

Step 6 - After clicking on the [Redo Layout] button, we now view the changed rule in the Text Rule field at the bottom of the panel:
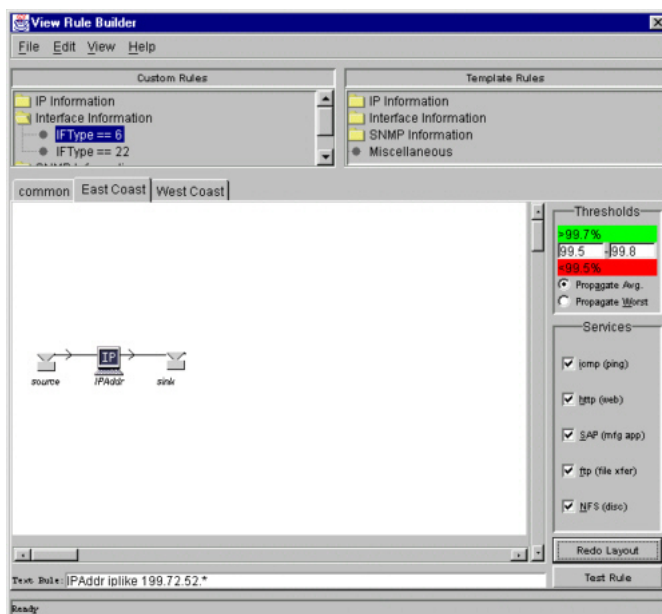
Figure: Step 6 - Clicking on Redo Layout

Step 7 - We drag a rule called "ifType ==22" from the Custom Rules area and drop it onto the Drawing Area. The new rule now floats on the panel. The following figure shows the new floating rule:
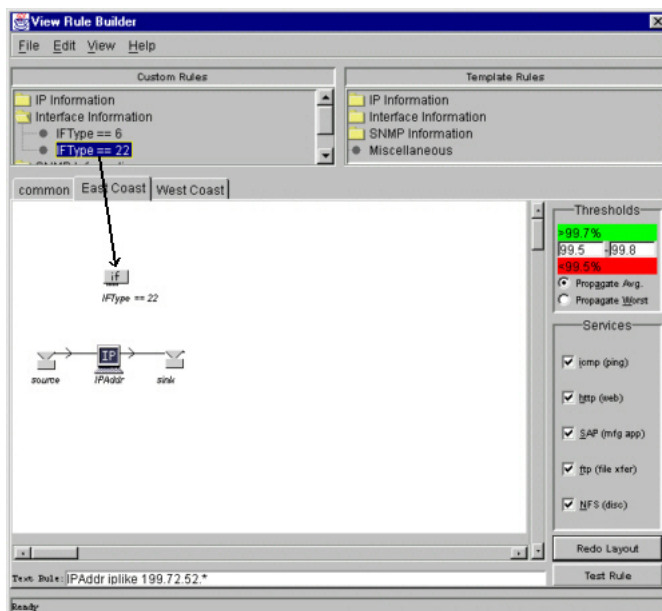

Figure: Step 7 - Dropping a Rule on the Drawing Area

Step 8 - To connect a path from the Source to the new rule, a right click on the Source pops up a panel. "Join" will allow a connection form the selected icon. The following figure shows the pop-up when a right click occurs on an icon:
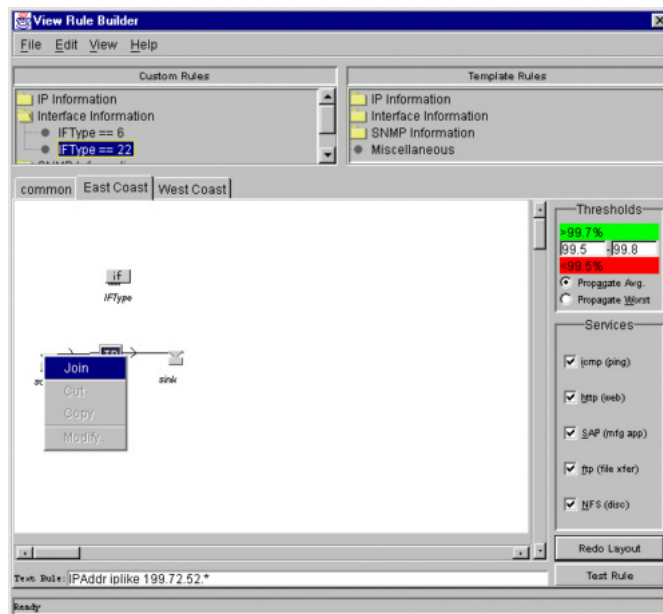
Figure: Step 8 - Right Click to Join

Step 9 - After "Join" is selected, the cursor becomes a stretchable line. A click on the destination icon will complete the line. The following figure shows the line being dragged around the panel:
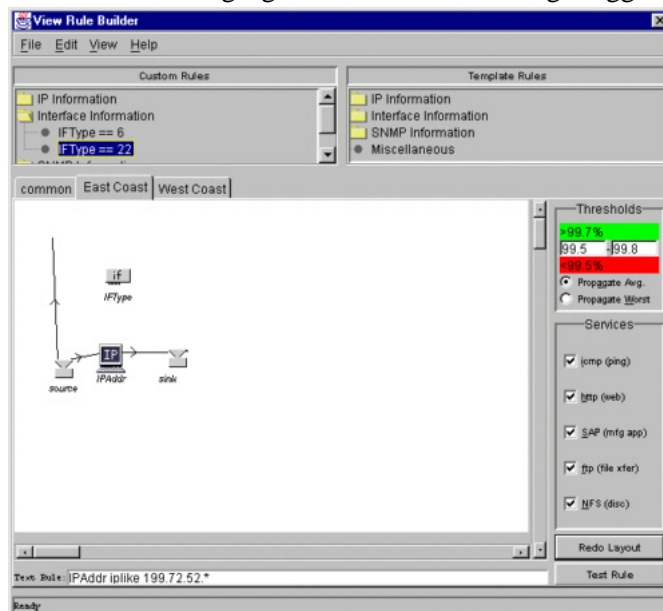

Figure: Step 9 - Drawing a New Path

Step 10 - The Join command has now been completed by clicking on the "ifType ==22" icon. The two icons are now related to each other. The following figure shows the new line between the the icons:
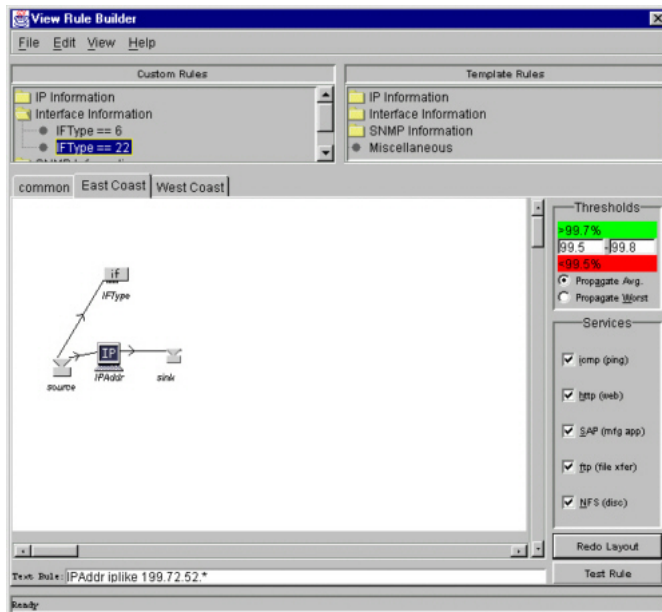
Figure: Step 10 - After the Drawn Line

Step 11 - In a similar fashion, a line drawn from the "ifType ==22" icon to the "IPAddr iplike 199.72.52.*" icon. The following figure shows the two new lines:
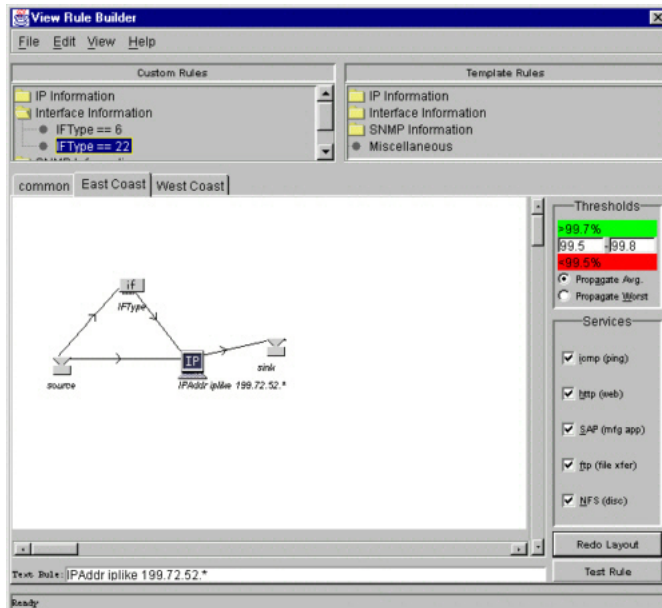

Figure: Step 11 - Drawing Another Relationship

Step 12 - Since we have a closed loop, the relationship between the Source and the "IPAddr iplike 199.72.52.*" icon needs to be deleted. The following figure show the pop-up panel when the line between the two icons is right clicked:
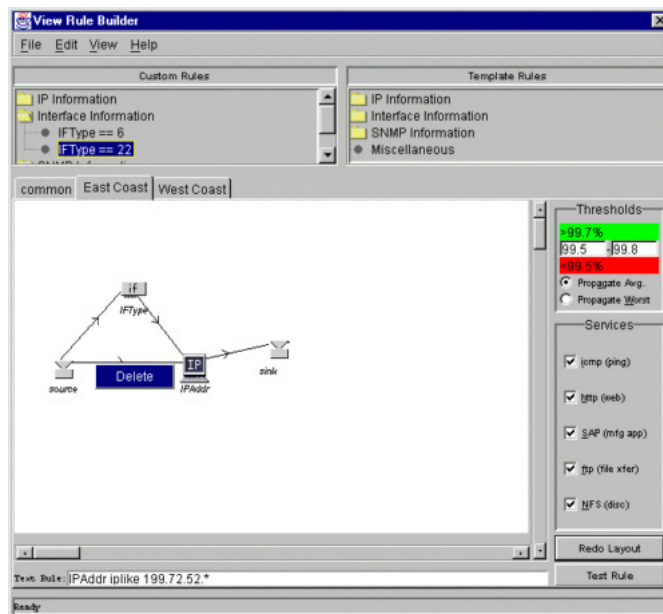
Figure: Step 12 - Removing a Relationship

Step 13 - The following figure shows the panel after "Delete" is selected from the pop-up panel:
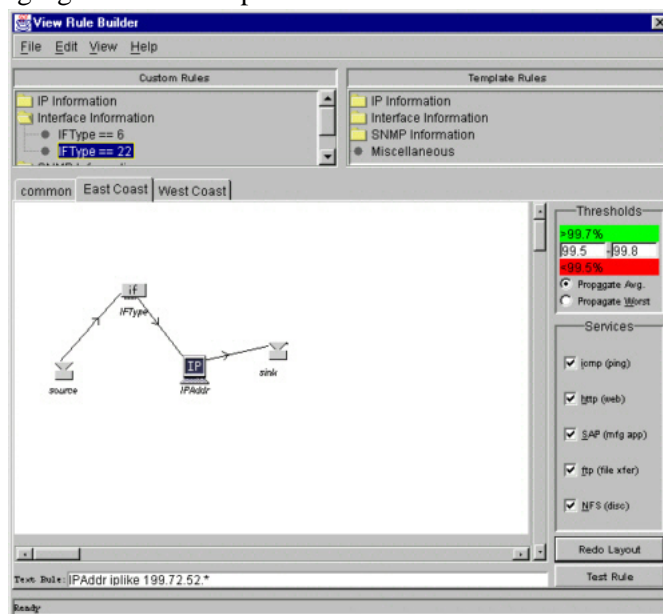

Figure: Step 13 - After the delete

Step 14 - After clicking on the [Redo Layout] button, the icons are reordered and the Text Rule now shows the grammar of the new "AND" relationship. The following figure shows the reordered icons and the new Text Rule:
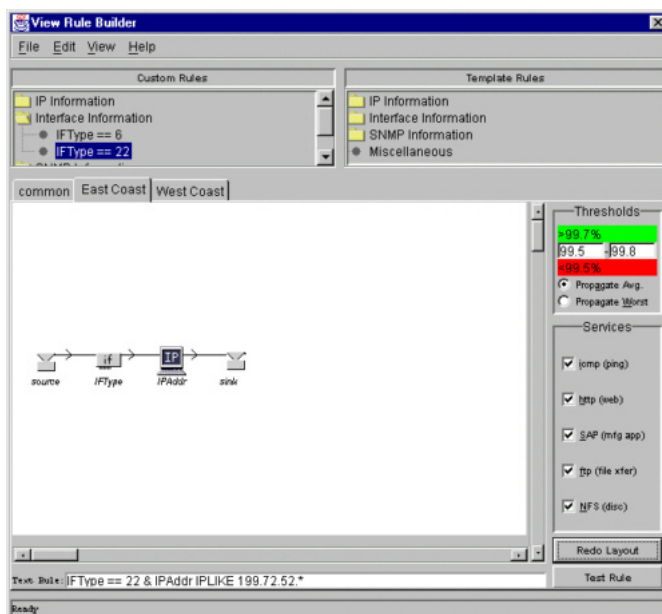
Figure: Step 14 - After Redoing the Layout

Step 15 - The following figure show the panel after we have dropped a new rule onto the Drawing Area:
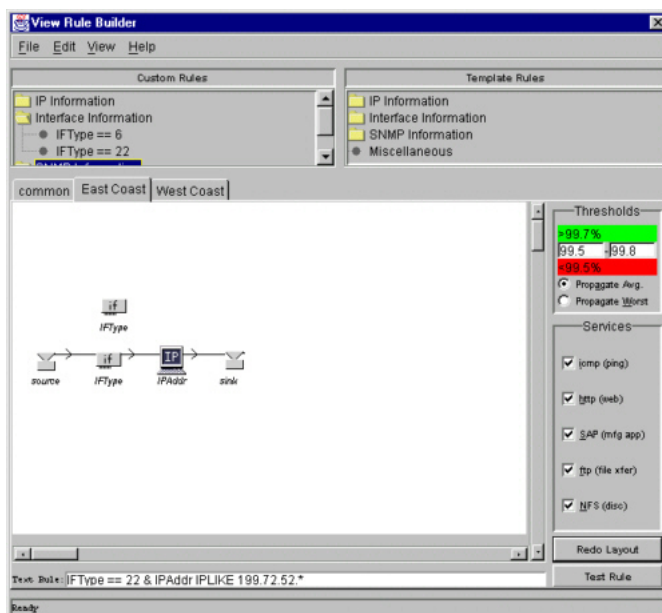

Figure: Step 15 - Adding an "OR" relationship

Step 16 - After adding two new connections and clicking on the "Redo Layout] button, the new rules and Text Rules look like the following:.
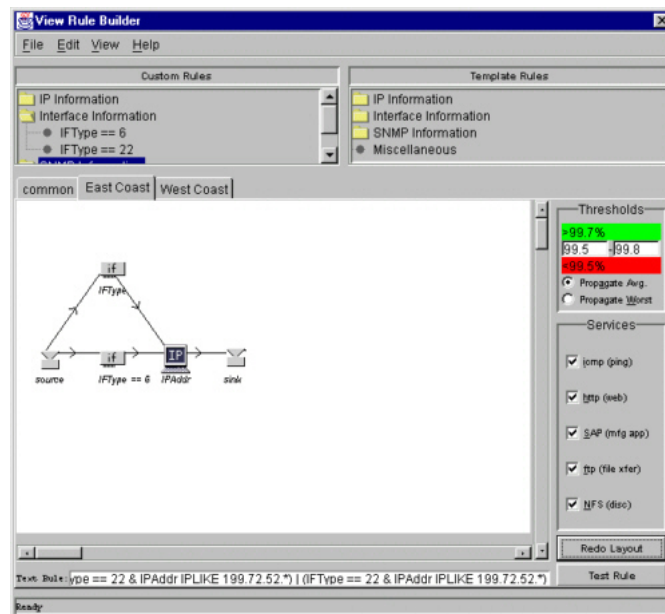
Figure: Step 16 - Completing the "OR"

Chapter 4
Configuring Distributed Pollers

Distributed pollers are configured by the administrator on the master station. A poller is a computer in the network which can talk to local devices and report the result to the master station. A distributed poller can be any machine capable of communicating with the local nodes via an interface.

4.1 Understanding Distributed Pollers

When configuring Distributed Pollers, a strong understanding of how pollers and poller packages are built will help reduce configuration time and eliminate errors.

4.1.1 Understanding Poller Packages

Poller packages are ways to organize configuration information into reusable containers of information. A poller package contains four basic items:

Ranges	defines the boundaries of the discovery domain. Devices not listed in a range will be ignored by the distributed poller.
Filters	defines the types of devices to discover. Filters are rules which are evaluated TRUE or FALSE when a device is interrogated. If the rule evaluates TRUE, the device will be added to the database of the distributed poller. If the rule evaluates FALSE, the device will be ignored. Filters are combined with ranges to define the boundaries of the management domain and the interesting devices within that domain.
Services	defines the types of services to probe on a device. Services include icmp, http, smtp, or any currently supported protocols.
Calendars	defines the hours of operation for the package; both one-time and repeating calendars. Periodic backup, preventative maintenance and other scheduled outages are flagged here. (This functionality is curently disabled).

4.1.2 Poller Packages and Pollers

A Distributed Poller can have one or more poller package assigned to it. When a poller package is assigned to a poller, the poller uses the information in the package to control the behavior, parameters and methods of the poller. The package is downloaded to the poller at initialization time. This architecture allows several kinds of configurations:

|	
Overlapping Pollers	If a poller package is assigned to more than one poller (currently not supported), the domains defined in that poller package overlap. Should one of the pollers fail, the other poller will continue to operate. Overlapping pollers will generate multiple traffic since both pollers operate against the same set of devices.

Independent Pollers   If a poller package is assigned to only one poller, then that poller is solely responsible for polling services. Traffic is minimized in this configuration, however, a single poller has a single point of failure.

4.2 Running the Configure Distributed Pollers Panel

The administrator builds new pollers, assigns poller packages and (in the near future) sets calendars of operation. To launch the configuration panel for the distributed pollers:

1.   Log in as the administrator. The Administrator Main Panel appears:
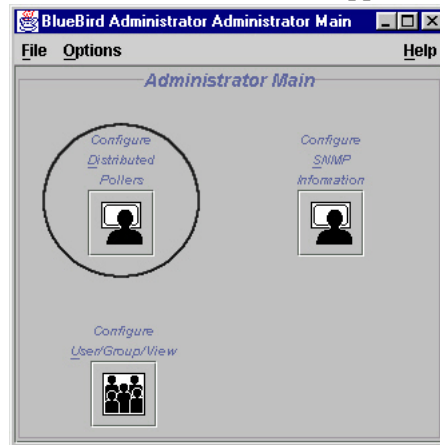
Figure: Administrator Main Panel

2.   Select the "Configure Distributed Pollers" icon from the Administrator Main Panel. The Configure Distributed Pollers panel appears:
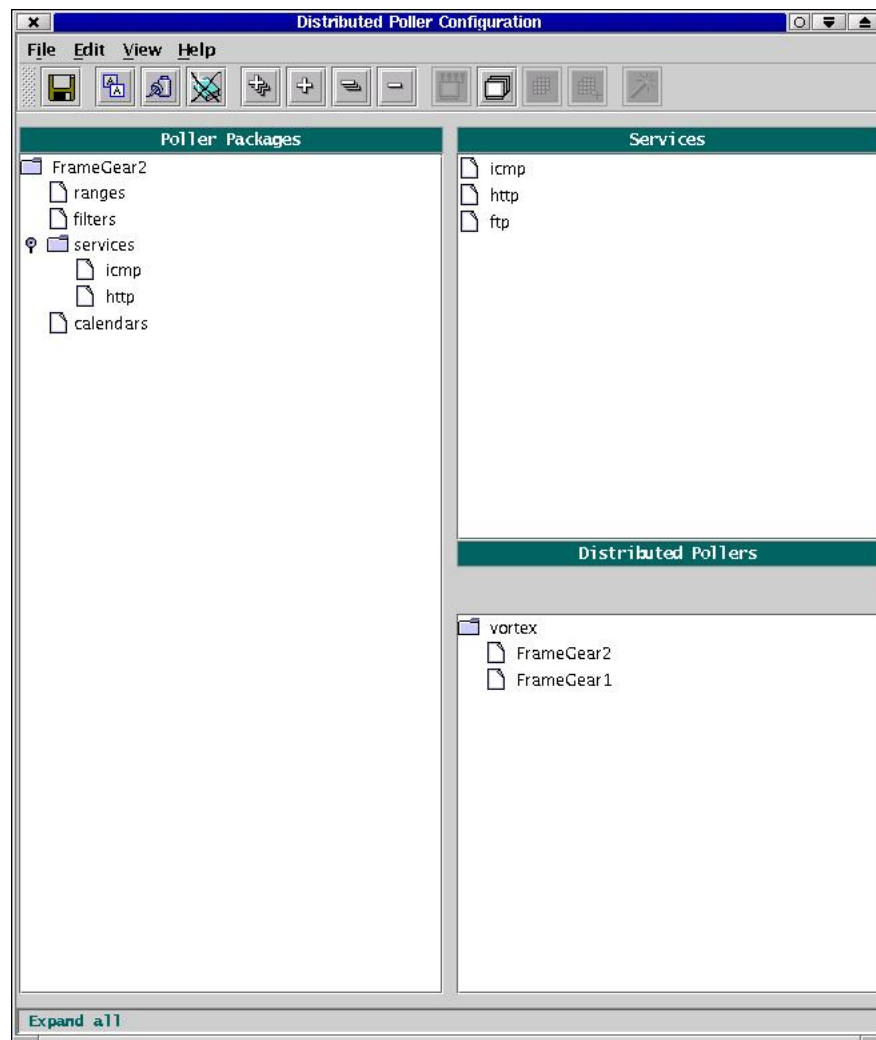
Figure: Configure Distributed Pollers Panel

4.3 Using the Configure Distributed Pollers Panel

The "Configure Distributed Pollers" panel is used to build or modify distributed pollers, poller packages, and assign packages to pollers. The panel has a menu bar at the top, icon bar for tools, configuration areas and a status bar.

4.3.1 Configure Distributed Pollers Configuration Areas

The Configure Distributed Pollers panel is divided into 3 separate areas:

Poller Packages	Defines the packages configured on the system. Each package has four branches below it representing the ranges, filters, services and calendars assigned to that package. Packages are dragged onto pollers using drag and drop operations.
Services	The services area contains all of the supported services which a poller could potentially probe. Since a service requires poller code specific to that service, additional services can only be added to the list by adding new service modules to Bluebird.

Distributed Pollers   Pollers are shown in this area of the panel. If a poller is not listed here, it is not polling. Pollers must have at least one poller package assigned to them. Packages are dragged from the Poller Package area and dropped onto a Distributed Poller.

4.3.2 Icon Bar

The icon bar holds the more commonly used tools for configuring pollers. Icons exist for copy, paste, expand/contract, wizards and adding new pollers.

Figure: Distributed Poller Icon Bar

- Save the current configuration. Allows the changed configuration to be make permanent. As changes are made, they are stored in a temporary location until a save if performed. Exiting the panel without saving changes will generate a confirmation message warning about the loss of data.

- Copy the selected item to the copy buffer. A paste is necessary to insert the copied information.

- Paste the last copied information into the selected item. Pasting one type of information into a different type of information is not allowed and will be ignored or generate an error message.

- Expand all the folder icons to view the contents.

- Expand the selected folder icon and show all the contents. Double clicking on the folder icon will accomplish the same task.

- Collapse all folder icons. All contents of folder tabs will be hidden.

- Collapse the selected folder icon. Double clicking on an open folder icon will accomplish the same task.

- Add a new poller icon. This function is disabled until distributed management capability is enabled.

- Add a new poller package. Selecting this icon will pop up a panel allowing the poller information to be configured.

- Add a one-time calendar. This function is disabled until calendar functionality is available.This

- Add a repeating calendar. This function is disabled until calendar functionality is available.

- Launch the Wizard. Currently disabled until distributed management capability is available.

4.3.3 Menu Bar

The menu bar allows menu operation of basic functions such as copying and pasting. In addition to mouse clicks, keyboard accelerators allow keyboard operation of basic functions. The various menu bar options are shown below:

File Menu Bar

The file menu allows the current configuration to be saved and the application to be exited. If modification to the distributed poller configuration has not been saved, a warning message will appear allowing the configuration to be saved or discarded.
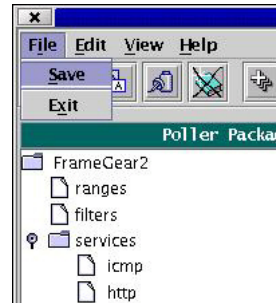
Figure: File Menu Bar

Edit Menu Bar

The Edit menu allows basic delete,copy and paste functions. The selected item in the node tree is either deleted, copied into the copy storage area or pasted from the copy storage area.
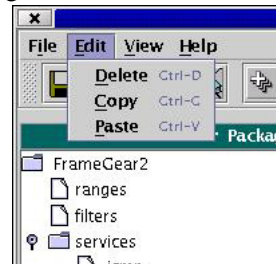
Figure: Edit Menu Bar

View Menu Bar

View allows the look and feel of the panel to be chosen. Personal attributes of the panel are saved on a user by user basis, including look and feel.
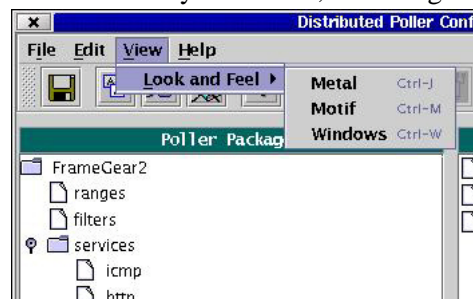
Figure: View Menu Bar

Help Menu Bar

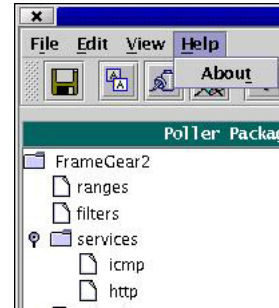The Help menu bar provides basic "About" functionality and entry point into the help subsystem.

Figure: Help Menu Bar

4.3.4 Right Click Functionality

The right mouse click is used extensively for changing the properies of the various configurations in the poller packages and pollers. Right click functionality is sensitive:

Ranges Properties

Right clicking on the Ranges branch in the Poller Packages section allows ranges to be copied, pasted and properties changed.
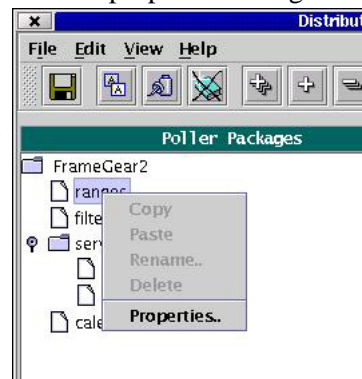

Figure: Right click on "Ranges"

Filters Properties

Right clicking on the Filters branch in the Poller Packages section allows filters to be copied, pasted and properties changed. Under Properties, the graphical rule builder is used to define the rules for that filter.
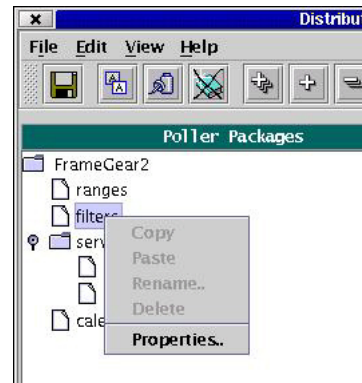

Figure: Right click on "Filters"

Services Properties   Right clicking on the Services branch in the Poller Packages section allows services to be copied, pasted and properties changed.
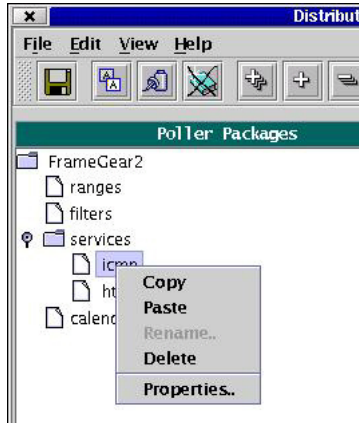

Figure: Right click on "Services"

4.4 Defining Ranges

Ranges defined the "responsibility domain" of a poller package. This domain is bounded by the start and stop IP addresses of devices within it's responsibility. If a device is not defined within a range, Bluebird will not manage it. Ranges are additive and subtractive; i.e. you can define multiple ranges and all the ranges are added together to form a superset. In addition, ranges can be excluded to never manage devices within a range.

Careful definition of ranges can allow multiple distributed pollers to overlap responsibility. Packages should be defined to be reusable by different pollers.

4.4.1 Default Range Properties Panel

When the "Properties" option is selected from a right click on a ranges folder, the following panel appears:
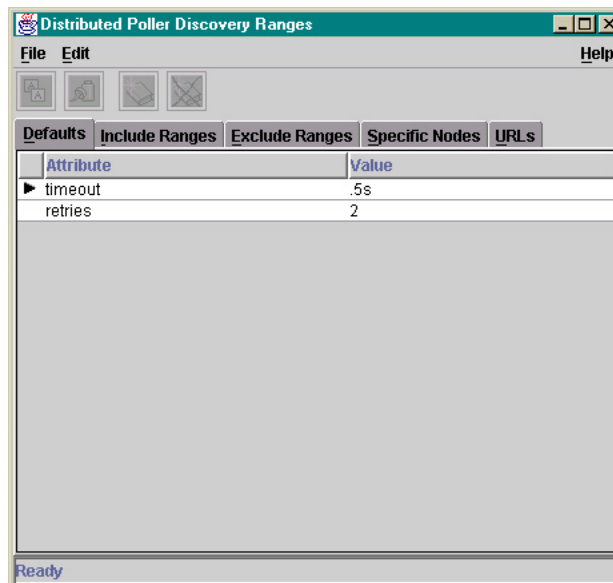

Figure: Ranges Default Panel

The Ranges Default panel defines the defaults used by the various other panels in the Ranges configuration. If a value is defined as <default>, then the values from this panel are used.

Fields defined within this panel include:

Timeout   The timeout determines how long Bluebird will wait for a device to respond to an ICMP message before considering the device unreachable.

Retries   The retries determine the number of times Bluebird will retry a node which has timed out (see Timeout above) and is unreachable. After the "Retries" number of retries, the node is considered unreachable.

4.4.2 Include Range Properties Panel

When the "Include Ranges" folder tab is selected, the following panel appears:
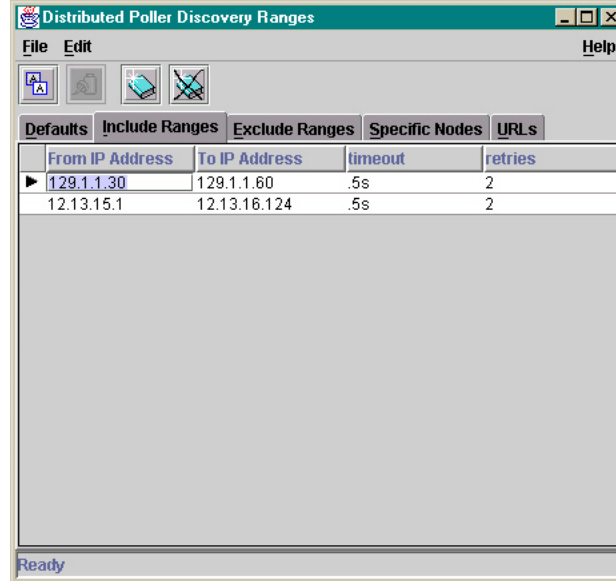

Figure: Include Ranges Panel

The "Include Ranges" panel starting and stopping addresses in the management domain. As many ranges as necessary can be added to define the bondaries of this domain. Wildcards are not allowed in the from and to fields. The timeout and retries fields define the timeouts for devices within that range.

Fields defined within this panel include:

From IP Address The beginning IP address in the management range.

To IP Address The ending IP address in the management range.

timeout   The timeout determines how long Bluebird will wait for a device to respond to an ICMP message before considering the device unreachable.

retries   The retries determine the number of times Bluebird will retry a node which has timed out (see Timeout above) and is unreachable. After the "retries" number of retries, the node is considered unreachable.

4.4.3 Exclude Range Properties Panel

When the "Exclude Ranges" folder tab is selected, the following panel appears:
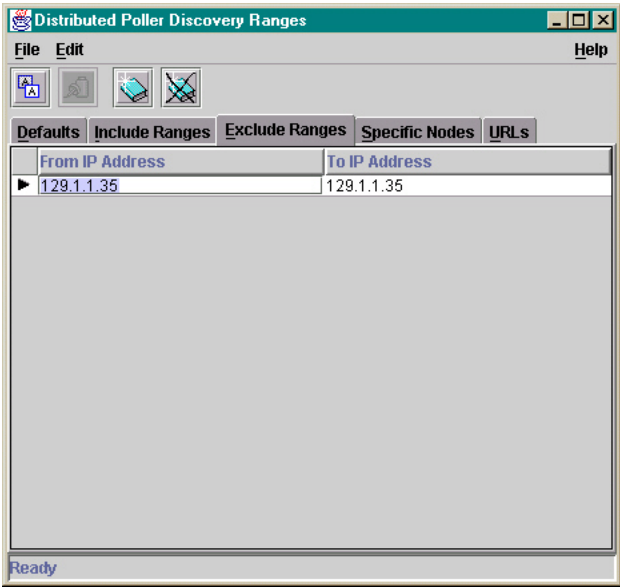
Figure: Exclude Ranges Panel

The "Exclude Ranges" panel defines the range of devices to ignore within the management domain. A range of excluded devices allows for blocks to be ignored.

Fields defined within this panel include:

From IP Address   The beginning IP address to be excluded from the management range.

To IP Address   The ending IP address to be excluded from the management range.

4.4.4 Specific Node Properties Panel

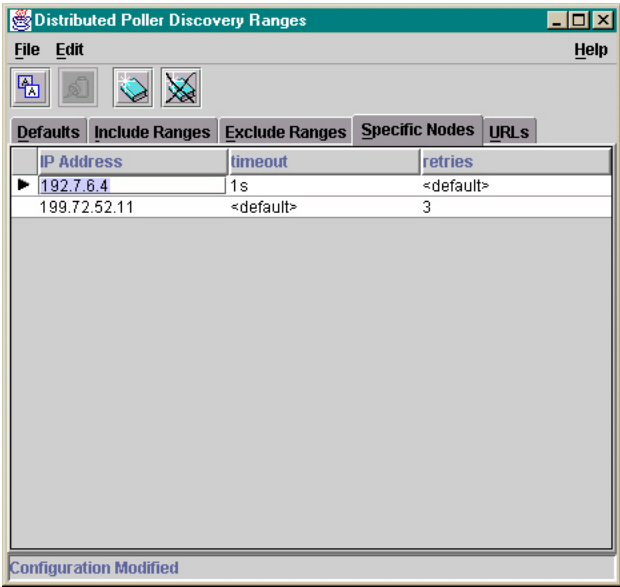When the "Specific Nodes" folder tab is selected, the following panel appears:

Figure: Specific Node Panel

The "Specific Nodes" panel allows individual devices to be entered and the attributes assigned.
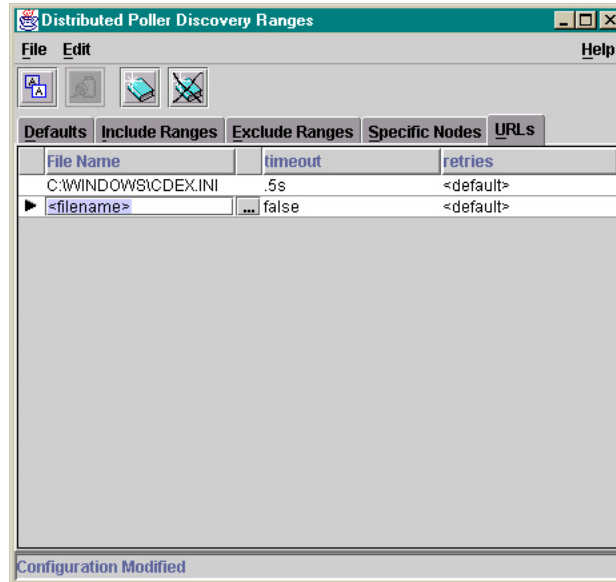
Fields defined within this panel include:

IP Address   The IP address to be specifically included from the management range.

timeout   The timeout determines how long Bluebird will wait for the device to respond to an ICMP message before considering the device unreachable.

retries   The retries determine the number of times Bluebird will retry the node which has timed out (see Timeout above) and is unreachable. After the "retries" number of retries, the node is considered unreachable.

4.4.5 URL Properties Panel

When the "URLs" folder tab is selected, the following panel appears:

Figure: URLs Panel

The "URLs" panel provides the same functionality as specific nodes, however, the addresses of the specific nodes as contained in a file located at that URL location. This allows files to be build with desired nodes and the file referenced here.

Fields defined within this panel include:

File Name   The URL of a file containing a list of nodes to be included in the management domain.

...   A file browser allowing you to pick the URL.

timeout   The timeout determines how long Bluebird will wait for a device to respond to an ICMP message before considering the device unreachable.

retries   The retries determine the number of times Bluebird will retry a node which has timed out (see Timeout above) and is unreachable. After the "retries" number of retries, the node is considered unreachable.

Chapter 5
Configuring SNMP Information

When a device is discovered by a distributed poller, information about the device is retrieved and stored in the device object database. In order for the distributed poller to properly access SNMP information on SNMP agents, basic SNMP configuration information must be configured.

5.1 Running the SNMP Configuration Panel

The administrator uses the "Configure SNMP Information" option from the Administrator main panel to set up SNMP community information, timeouts, retries and backoff methods.

To run the SNMP Configuration Tool, select the "Configure SNMP Information" icon from the administrator main panel:

1. Log in as the administrator. The Administrator Main Panel appears:

Figure: Administrator Main Panel

2. Select the "Configure SNMP Information" icon from the Administrator Main Panel. The SNMP Configuration Panel appears.
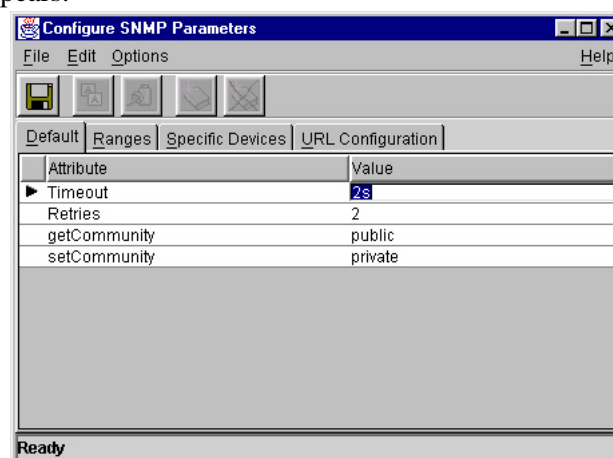

Figure: Configure SNMP Information Main Panel

5.2 The SNMP Configuration Panel Areas

The SNMP Configuration Panel contains the following areas:

Folder Tabs   The four folder tabs allow SNMP information to be configured for defaults, IP address ranges, specific IP address or a URL (file) containing a list of IP addresses.

Icon Bar   The icon bar allows entries in the display area to be manipulated. Copy, paste, save, new entry and delete entry icons are available.

Configuration Display Area   the display area shows the configured entries and their settings. Items are evaluated from top to bottom in the panel such that if there are ambiguous entries, the entry at the top takes precedent. For example, if there is an entry for 129.1.1.1-129.1.1.256 and an entry for 129.1.1.1-129.1.1.64, then the entry at the top of the list takes precedent.

Menu Bar   provides program options, file operations, copy, paste etc.

Status Bar   shows the status of the SNMP panel and state of the configuration file. The status bar can be either "Ready" or "Configuration Modified". Saving the configuration will toggle the state of the status from "Configuration Modified" to "Ready".

5.2.1 Icon Bar

Tools in the Icon bar allow quit modification and manipulation operations.

Figure: Icon Bar

Icon Bar Tools

- Save the current configuration. Allows the changed configuration to be make permanent. As changes are made, they are stored in a temporary location until a save if performed. Exiting the panel without saving changes will generate a confirmation message warning about the loss of data.

- Copy the selected row. The arrow indicator on the left side of the table shows the selected row. When the "Copy" button is selected, the contents of the row is stored in the copy buffer. The "copy" icon is disabled in the "default" folder tab since defaults cannot be copied and added.

- Paste the selected row. When a paste is selected, the contents of the copy buffer is then pasted back into the table as a new row. New rows are always added at the end of the list. The "paste" icon is disabled in the "default" folder tab since new defaults cannot be added. The "paste" icon is disabled unless a row has been copied using the "copy" option.

- Add a new row. When the "New" \ button is selected, a new row is added to the table at the bottom of the table. The cell contents contain default information which refers back to the global SNMP defaults. When a new row is built which contains cells that have edit requirements (e.q. IP Address must be xxx.xxx.xxx.xxx), a cell is built which contains "placeholder" information.

These placeholder cells contain syntactically correct entries which pass edits but do not contain "real" information (e.q. IP address will be "255.255.255.255). The "new" icon is disabled in the "default" folder tab since new defaults cannot be added.

- Delete the selected row. The arrow indicator on the left column of the table shows the selected row which will be deleted. There is no confirmation message before the delete occurs. There is no way to "undo" a deleted row (except for exiting the panel without saving). The delete icon is disabled in the "default" folder tab since defaults cannot be deleted. The delete icon is disabled if there is not a row selected in the table.

5.2.2 Folder Tabs

The folder tabs select the type of configuration for setting SNMP information. Information in folder tabs can be combined to build complex sets of wildcard and specific node information. For example, a range configuration can be used to set the timeout for an entire section of the network and combined with a specific node configuration to override the range for that one device.

Configurations are evaluated from most specific to most general; i.e. a specific node configuration overrides a range configuration which overrides the default configuration

If a node is not in the specific, url or range configuration, it is polled using the default SNMP configuration information in the "default" folder tab.
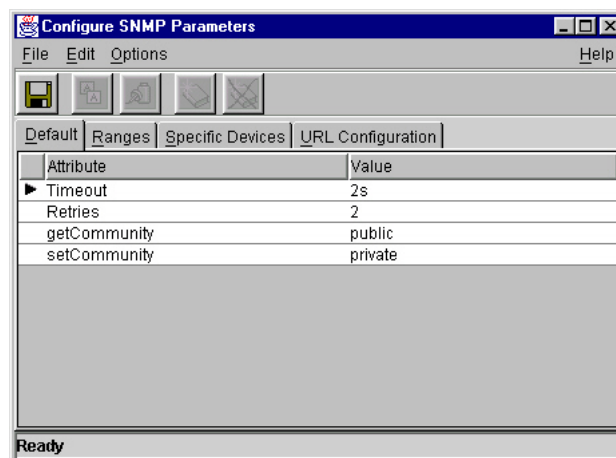
"Default" Folder Tab

Figure: Configure SNMP Information Main Panel

The default folder holds the information used when no other SNMP configuration applies. Configuration information includes community strings, retries, timeouts and backoffs.

timeout   Amount of time in seconds which the poller will wait for an SNMP response after sending an SNMP request. Time is specified in "x.ys" format where x is the number of seconds, y is the number of tenths of seconds and the letter s indicates seconds. Currently, the only valid suffix is "s".

Retries   the number of additional times that the poller will retry the device if the response is not received within "timeout" seconds.

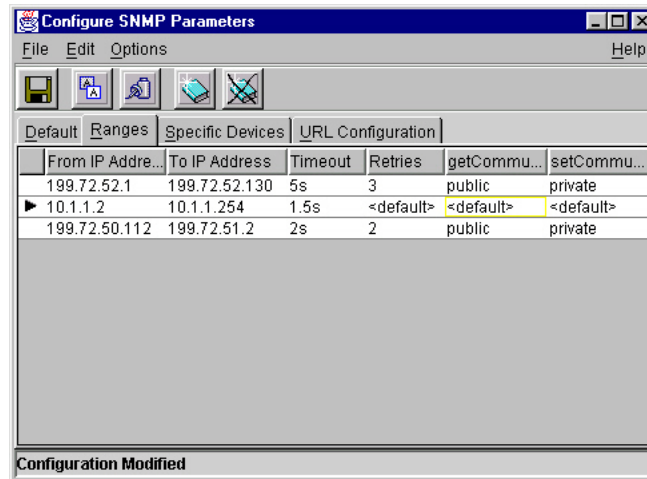|	
Get Community	when an SNMP request is sent to the device, this string is used for the GET request.
Set Community	when an SNMP request is sent to the device, this string is used for the SET request.

"Ranges" Folder Tab

Figure: Ranges Folder Tab

The "ranges" folder tab holds the information used for configuring large blocks of the network. Configuration information includes community strings, retries, timeouts and backoffs.

|	
From IP	The starting IP address within the IP address range. The IP address entered here is inclusive; i.e. the "From IP" address is considered in the range.
To IP	The ending IP address within the IP address range. The IP address entered here is inclusive; i.e. the "To IP" address is considered in the range.
Timeout	Amount of time in seconds which the poller will wait for an SNMP response after sending an SNMP request to a device in the range. Time is specified in "x.ys" format where x is the number of seconds, y is the number of tenths of seconds and the letter s indicates seconds. Currently, the only valid suffix is "s".
Retries	the number of additional times that the poller will retry a device in the range if a response is not received within "timeout" seconds.
Get Community	when an SNMP request is sent to a device in this range, this string is used for the GET request.
Set Community	when an SNMP request is sent to a device in this range, this string is used for the SET request.

"Specific Node" Folder Tab
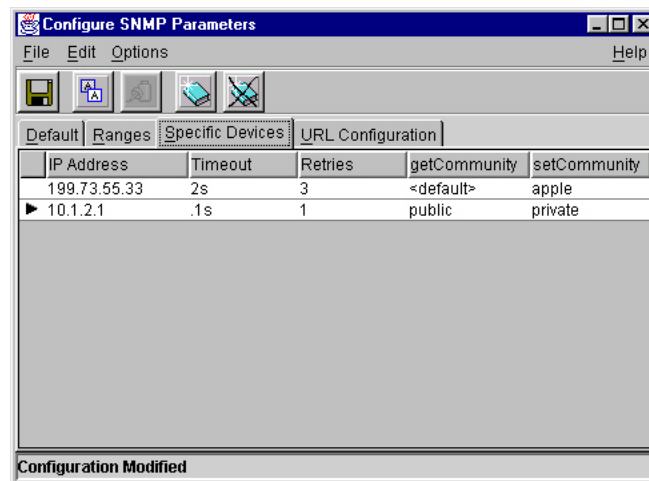
Figure: Specific Node Folder Tab

The "specific node" folder tab holds the information about a specific node which is different from every other node. Configuration information includes community strings, retries, timeouts and backoffs.

|	
Specific IP	The IP address of the node.
Timeout	Amount of time in seconds which the poller will wait for an SNMP response after sending an SNMP request. Time is specified in "x.ys" format where x is the number of seconds, y is the number of tenths of seconds and the letter s indicates seconds. Currently, the only valid suffix is "s".
Retries	Retries - the number of additional times that the poller will retry the device if the response is not received within "timeout" seconds.
Get Community	when an SNMP request is sent to this device, this string is used for the GET request.
Set Community	when an SNMP request is sent to this device, this string is used for the SET request.

URL Folder Tab
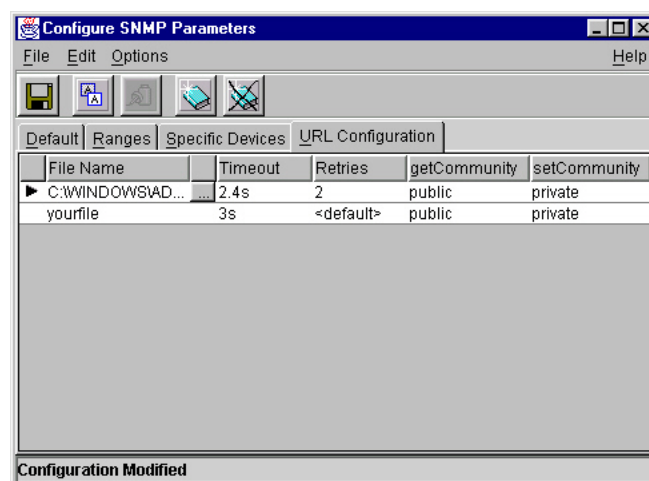
Figure: URL Node Folder Tab

The "URL" folder tab holds the information about a list of specific nodes which is different from every other node. The URL is typically a file location which contains a list of IP addresses. Configuration information includes community strings, retries, timeouts and backoffs.

URL	The location where the list of nodes can be obtain. The is is a file of IP addresses of the node.
Timeout	Amount of time in seconds which the poller will wait for an SNMP response after sending an SNMP request. Time is specified in "x.ys" format where x is the number of seconds, y is the number of tenths of seconds and the letter s indicates seconds. Currently, the only valid suffix is "s".
Retries	the number of additional times that the poller will retry the device if the response is not received within "timeout" seconds.
Get Community	when an SNMP request is sent to devices in the list, this string is used for the GET request.
Set Community	when an SNMP request is sent to devices in the list, this string is used for the SET request.

When the URL file browser button is selected for a row, the file chooser appears. This allows the user to selected a file from the local file system rather than type in a local file name and location.
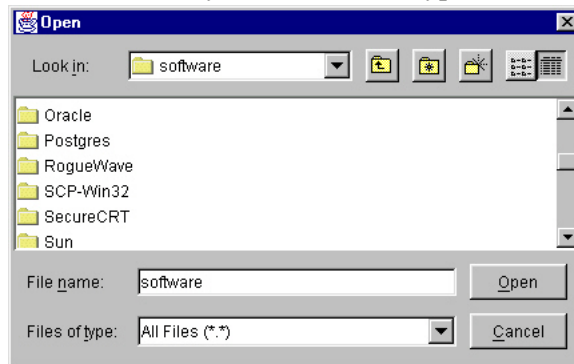
Figure: URL File Chooser

5.2.3 Configuration Display Area

The configuration display area is a table of information related to the particular type of folder tab. Columns of variables define the allowable parameters which can be configured. The rows define the configured variables for a range, specific node or URL list of nodes.

New Row   When a new row is added to the Configuration Display Area, a placeholder row is built which contains default information. After the row is added, non-default information is entered into the row to configure those IP addresses:
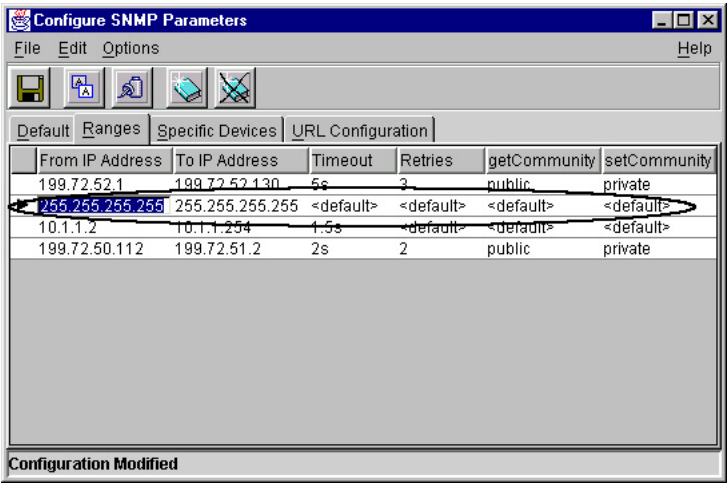

Figure: New Row

Incorrect   Cells cannot be left in an inconsistent state during editing. When a row which
Cell   is in the process of being edited loses focus on another part of the panel, a
Syntax   warning message is displayed:

Figure: Error Message Box

5.2.4 Menu Bar

The menu bar allows mouse and keyboard control of basic functions in the SNMP Configuration panel. The menus can be accessed using accelerator keys from the keyboard as well as the mouse.
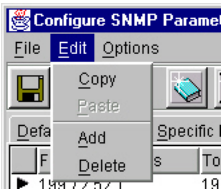

Figure: Edit Menu Bar

Chapter 6
Rules and Filters

6.1 Understanding Rules

Rules are used in two distinct places in Bluebird

1. Buusiness Views - filter what and operator sees in the Real-time console.

2. Discovery Filters - filter the devices discovered by a distributed poller.

Appendix A
Care and Feeding of Bluebird

This is text after the appendix starts and before the section.

Files

This is the text of the first section in the appendix A

Appendix B
File and Directory Structure

Appendix text after title but before section.

a sect1 title

some text. some text. some text. some text. some text. some text.

a sect2 title

some text. some text. some text. some text. some text. some text.

another sect1 title

Section 1 text.

another sect2 title

sSection 2 text.