

WANJet® Appliance Administrator Guide

version 4.2.11

MAN-0229-01



Product Version

This manual applies to product version 4.2.11 of the WANJet® appliance.

Publication Date

This manual was published on February 27, 2007.

Legal Notices

Copyright

Copyright 2003-2007, F5 Networks, Inc. All rights reserved.

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described by applicable user licenses. F5 reserves the right to change specifications at any time without notice.

Trademarks

F5, F5 Networks, the F5 logo, BIG-IP, 3-DNS, iControl, Internet Control Architecture, IP Application Switch, iRules, OneConnect, Packet Velocity, SYN Check, Control Your World, ZoneRunner, uRoam, FirePass, TrafficShield, Swan, WANJet, WebAccelerator, and TMOS are registered trademarks or trademarks, and Ask F5 is a service mark, of F5 Networks, Inc. in the U.S. and certain other countries. All other trademarks mentioned in this document are the property of their respective owners. F5 Networks' trademarks may not be used in connection with any product or service except as permitted in writing by F5.

Patents

This product protected by U.S. Patents 6,327,242 and 7,126,955. Other patents pending.

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

RF Interference Warning

This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.

FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user, at his own expense, will be required to take whatever measures may be required to correct the interference.

Any modifications to this device, unless expressly approved by the manufacturer, can void the user's authority to operate this equipment under part 15 of the FCC rules.

Canadian Regulatory Compliance

This Class A digital apparatus complies with Canadian ICES-003.

Standards Compliance

This product conforms to the IEC, European Union, ANSI/UL and Canadian CSA standards applicable to Information Technology products at the time of manufacture.

Acknowledgments

This product includes software developed by the University of California, Berkeley and its contributors.

This product includes software developed by the NetBSD Foundation, Inc. and its contributors.

This product includes software developed by Bob Withers.

This product includes software developed by Jean-Loup Gailly and Mark Adler.

This product includes software developed by Markus FXJ Oberhumer.

This product includes software developed by Gullaume Fihon.



Table of Contents

I

Introducing the WANJet Appliance

Overview of the WANJet appliance	I-1
About this guide	I-2
Stylistic conventions in this guide	I-2
Using the solution examples	I-2
Identifying new terms	I-2
Identifying references to objects, names, and commands	I-3
Identifying references to other documents	I-3
Identifying command syntax	I-3
Finding additional information and technical support	I-4

2

Overview of WANJet Appliance Features

About the WANJet appliance	2-1
Transparent Data Reduction	2-2
TDR-2: Repeat transfer data reduction	2-2
TDR-1: First transfer data reduction and networking adaptivity	2-5
Application QoS	2-6
Type of Service	2-6
Simple Network Management Protocol Support	2-7
Remote Monitoring support	2-8
System Log protocol support	2-9
Connection Intercept	2-10

3

Installation

WANJet appliance deployment	3-1
Types of deployment	3-1
One-arm deployment	3-4
Firewall guidelines	3-8
Hardware installation	3-8
Site information worksheet	3-9

4

Initial Configuration

Logging on to the WANJet Web UI	4-1
Using the Web UI	4-2
Activating the license	4-6
Basic WANJet appliance configuration	4-9
Naming the first WANJet appliance	4-10
Configuring multiple subnets	4-10
Defining the second WANJet appliance as a remote WANJet appliance	4-12
Naming the second WANJet appliance	4-13
Configuring multiple subnets on the second WANJet appliance	4-13
Defining the first WANJet appliance as a remote WANJet appliance	4-13
Testing connectivity	4-14
Additional configuration tasks	4-15
Troubleshooting	4-16

5

Managing the WANJet Appliance

Configuring authentication settings	5-1
Changing the Web UI passwords	5-1
Configuring remote authentication	5-2
Changing the WANJet LCD PIN code	5-3
Granting Web UI access	5-4
Using the Management port	5-5
Configuring time settings	5-6
Setting the time zone	5-6
Synchronizing time automatically	5-7
Setting the time manually	5-7
Shutting down and restarting the WANJet appliance	5-8
Booting from an alternate image	5-10
Backing up and restoring settings	5-10
Restoring factory default values	5-12
Upgrading the WANJet appliance software	5-13

6

Advanced Configuration

Creating optimization policies	6-1
Managing subnets	6-1
Managing port settings	6-4
Enabling Connection Intercept	6-8
Setting operational modes	6-11
Adjusting tuning settings	6-13
Updating a configuration	6-14
Modifying a local WANJet appliance network configuration	6-15
Replicating configuration changes on remote WANJet appliances	6-16
Managing virtual LANs	6-17
Managing remote WANJet appliances	6-20
Changing the interface speed	6-22
Managing static routes	6-23
Configuring Syslog and SNMP settings	6-24
Configuring email alerts	6-25
Configuring redundant peers	6-26
Configuring one-arm topology	6-28
Introducing high-availability features	6-30
WANJet appliance bridging functionality	6-30
WANJet appliance fail-to-wire feature	6-31
WANJet appliance fail-close feature	6-32
WANJet appliance peer port	6-33
WANJet appliance remote peer load balancing and failure detection	6-35
Alternate high-availability configurations	6-36

7

Configuring Service Policies

Defining IT service policies	7-1
Adding, editing, or removing an IT service policy	7-1
Creating Application QoS policies	7-3
Adding, editing, or removing an Application QoS policy	7-3
Managing WAN links	7-6
Adding, editing, or removing WAN links	7-6
Adding a subnet to a WAN Link	7-7

8

Monitoring Performance

Introducing reports	8-1
Status report	8-2
Real Time Traffic report	8-3
Comparative Throughput reports	8-5
Diagnostics reports	8-11
Monitoring	8-11
Connectivity	8-20
General	8-23
Administration Tools	8-24
Diagnostic Log	8-27
System Snapshots	8-27
Third-party reporting systems	8-28
Syslog reports	8-28
SNMP reports	8-28
RMON2 Reports	8-29

9

Configuration Examples

Basic configuration	9-1
Mesh configuration	9-3
Hub and spoke configuration	9-5
Redundant system configuration	9-7
LAN router configuration	9-9

10

Working from the Command Line

Command line overview	10-1
Attaching a computer to the WANJet appliance	10-2
Getting to the command line	10-2
Using commands	10-4
Configuring the WANJet appliance from the command line	10-5
Running diagnostic tools from the command line	10-7
Shutting down or restarting the WANJet appliance from the command line	10-9
Replacing the console cable	10-9

A

WANJet Appliance Messages

WANJet appliance messages and codes	A-1
---	-----

B

WANJet Appliance Private MIB File

Ethernet card information	B-1
Introducing MIB files	B-2
WANJet appliance MIB file	B-2

C

RMON2 Tree

RMON2 tree overviewC-1

MIB treeC-1

Protocol directory treeC-1

Network layer matrixC-2

Application data matrixC-2

D

WANJet 200 Specifications

WANJet 200 platform specificationsD-1

E

WANJet 400 Specifications

WANJet 400 platform specificationsE-1

Glossary

Index



I

Introducing the WANJet Appliance

- Overview of the WANJet appliance
- About this guide
- Stylistic conventions in this guide
- Finding additional information and technical support

Overview of the WANJet appliance

F5® Networks WANJet® appliance increases distributed application performance by optimizing, thus reducing, the amount of data that is transferred over the WAN. As a result, the WANJet appliance accelerates applications, such as file transfer, email, client-server applications, and data replication, resulting in increased performance for all WAN users.

Various WANJet hardware platforms are available for large, medium, and small corporations, data centers, and branch offices. Some platforms, such as the WANJet 500, can optimize over 20,000 connections. WANJet appliances feature fault tolerance, and work seamlessly across all wide-area networks, including dedicated links, frame relays, and satellite connections.

Operating at Layer 5 of the OSI reference model, the WANJet appliance can gather application knowledge, and analyze data streams to determine what data to optimize. The WANJet appliance combines several technologies to accomplish the optimization, including Transparent Data Reduction (TDR), adaptive TCP optimization, Application QoS, and site-to-site encryption.

TDR technology reduces the amount of bandwidth that repeated data transfers consume across a WAN link, and compresses the data. Adaptive TCP optimization enables the WANJet appliance to adapt, in real time, to the latency, packet loss, and congestion characteristics of WAN links, and then to accelerate application traffic. Application QoS policies let you assign more bandwidth to critical network traffic. The WANJet appliance uses SSL encryption to protect the traffic moving from site to site.

About this guide

This guide describes how to configure and use the WANJet appliance. Its intended audience consists of network administrators, information system engineers, and network managers responsible for the configuration and ongoing management of the WANJet appliance.

This guide provides information about:

- Installing and configuring the WANJet appliance
- Administering and managing the WANJet appliance
- Performing advanced configuration tasks involving subnets, hubs, static routes, and VLANs
- Configuring remote WANJet appliances
- Managing IT service policies and application QoS policies
- Monitoring the WANJet appliance's performance
- Working from the command line
- Troubleshooting issues
- WANJet 200 and 400 specifications

Stylistic conventions in this guide

To help you easily identify and understand certain types of information, this document uses the following stylistic conventions.

Using the solution examples

All examples in this documentation use only private class IP addresses. When you set up the solutions we describe, you must use valid IP addresses suitable to your own network in place of our sample IP addresses.

Identifying new terms

When we first define a new term, the term is shown in ***bold italic*** text.

For example, after you have completed the hardware configuration, using either the LCD panel or a console connected to the F5 appliance's serial port, you can configure the WANJet appliance using the browser-based utility, called the ***Web UI***.

Identifying references to objects, names, and commands

We apply bold formatting to a variety of items to help you easily pick them out of a block of text. These items include web addresses, IP addresses, utility names, most controls in the Web UI, and portions of commands, such as variables and keywords.

For example, if the IP address of the appliance is **192.168.168.102**, type **https://192.168.168.102:10000** in the web browser to log in to the WANJet appliance.

Identifying references to other documents

We use italic text to denote a reference to a specific section, or another document. In references where we provide the name of a book as well as a specific chapter or section in the book, we show the book name in bold, italic text, and the chapter/section name in italic text to help quickly differentiate the two.

For example, see Chapter 6, *Reviewing Hardware Specifications*, in the ***Platform Guide: WANJet® 500*** for details about the WANJet 500 appliance.

Identifying command syntax

We show actual, complete commands in bold Courier text. Note that we do not include the corresponding screen prompt, unless the command is shown in a figure that depicts an entire command line screen.

For example, the following command traces the route from the WANJet appliance you are working on to the device at IP address **10.1.102.204**:

```
traceroute -v 10.1.102.204
```

Table 1.1 explains additional special conventions used in command line syntax.

Item in text	Description
\	Continue to the next line without typing a line break.
< >	You enter text for the enclosed item. For example, if the command has <your name> , type in your name.
	Separates parts of a command.
[]	Syntax inside the brackets is optional.
...	Indicates that you can type a series of items.

Table 1.1 *Command line conventions used in this manual*

Finding additional information and technical support

In addition to this guide, there are other sources of documentation that you can use to work with the WANJet appliance. The information is available in the guides and documents described below.

◆ ***WANJet® Appliance Quick Start Card***

The WANJet platform includes a printed *Quick Start Card* written for the specific platform that you purchased. It provides basic instructions for a quick setup and initial configuration of the WANJet appliance.

◆ ***Platform Guide: WANJet® 500***

This guide describes the WANJet 500 platform and includes detailed instructions on how to install the WANJet 500.

In addition to the documentation included with the platform, you can find additional technical documentation by using the following resources.

◆ **F5 Networks Technical Support web site**

The F5 Networks Technical Support web site, <http://tech.f5.com>, provides the latest documentation for the product, including:

- Release notes for the WANJet appliance, current and past
- Updates for guides (in PDF form)
- Technical notes
- Answers to frequently asked questions
- The Ask F5SM natural language question and answer engine.

◆ **Note**

To access this site, you need to register at <http://tech.f5.com>.



2

Overview of WANJet Appliance Features

- About the WANJet appliance
- Transparent Data Reduction
- Application QoS
- Type of Service
- Simple Network Management Protocol Support
- System Log protocol support
- Connection Intercept

About the WANJet appliance

The WANJet appliance is designed to improve the performance of your networks, reducing the bandwidth consumed when transmitting data. For the WANJet appliance to reduce the bandwidth consumed in data transmission, it processes data on one side and reverses this process on the other. This process requires installing at least two WANJet appliances, one to process data on one side and another to reverse data processing on the other side.

WANJet appliance optimization works by identifying redundancy patterns in input data and replacing those redundant patterns with symbols (encoding). When data arrives at its destination, symbols are replaced with the original patterns (decoding). The WANJet appliance stores a list of all identified redundancy patterns and their equivalent symbols, enabling it to handle both sent and received data at the same time.

The WANJet appliance's compression technology operates at Layer 5, the session layer, of the OSI (Open System Interconnection) reference model. This technology enables the WANJet appliance to recognize the redundancies in data traffic.

Some compression products operate at Layer 3 of the OSI model. They wait until individual application data streams merge before searching for redundancies. Merged data streams yield fewer redundancies than data streams that are not merged, so the Layer 3 approach is less than optimal.

Other compression products operate at Layer 7 of the OSI model, the application layer. These products do a great job for specific applications (such as Telnet and FTP), but other traffic crosses the WAN uncompressed, so overall bandwidth savings are limited.

Operating at Layer 5, as the WANJet appliance does, is more efficient than operating at any other layer in the OSI model. Unlike data compression based on Layer 3, the WANJet appliance compresses data streams before data merge, so it finds and removes more redundancies than Layer 3 methods. Unlike Layer 7 techniques, the WANJet appliance's data reduction technology examines all applications and compresses all traffic types.

Transparent Data Reduction

F5 Networks' Transparent Data Reduction (TDR) technology dramatically reduces the amount of bandwidth consumed across a WAN link for repeated data transfers. For example, without TDR, a 1 MB file transferred across a WAN link by 100 different users would consume 100 MB of bandwidth. With TDR, the same transfer would consume less than 10 MB of bandwidth. This is a reduction of more than 90% in WAN traffic volume.

With TDR, files are not stored or cached, so data is never out of date and it does not need to be refreshed. Every request for a piece of data is sent to the server that actually has that data (even across the WAN link).

In other words, unlike traditional caching algorithms, requests are never served from a local WANJet appliance without the file actually being sent by the server that has the data. As a result, a user can change the name of a file and still experience the same dramatic reduction with TDR.

The WANJet appliance implements TDR technology as a two-stage compression process to maximize bandwidth savings while minimizing processing latency. The first step of the process, called **TDR-2**, examines the transmitted data to determine if any part of it has been previously sent. If so, the WANJet appliance replaces the previously transmitted data with references. The second step, called **TDR-1**, further compresses the data through the use of dictionary-based compression and advanced encoding schemes.

TDR-2: Repeat transfer data reduction

TDR-2 data reduction routines identify and remove all repetitive data patterns on the WAN. As data flows through the two WANJet appliances, each one records the byte patterns and builds a synchronized dictionary. If an identical pattern of bytes traverses the WAN more than once, the WANJet appliance nearest the sender replaces the byte pattern with a reference to it, compressing the data. When the reference reaches the remote WANJet appliance, it replaces the reference with the data, restoring the data to its original format.

Following is an illustrated example of how TDR-2 works.

In Figure 2.1, Client A requests a file named **antivirus.dat**.

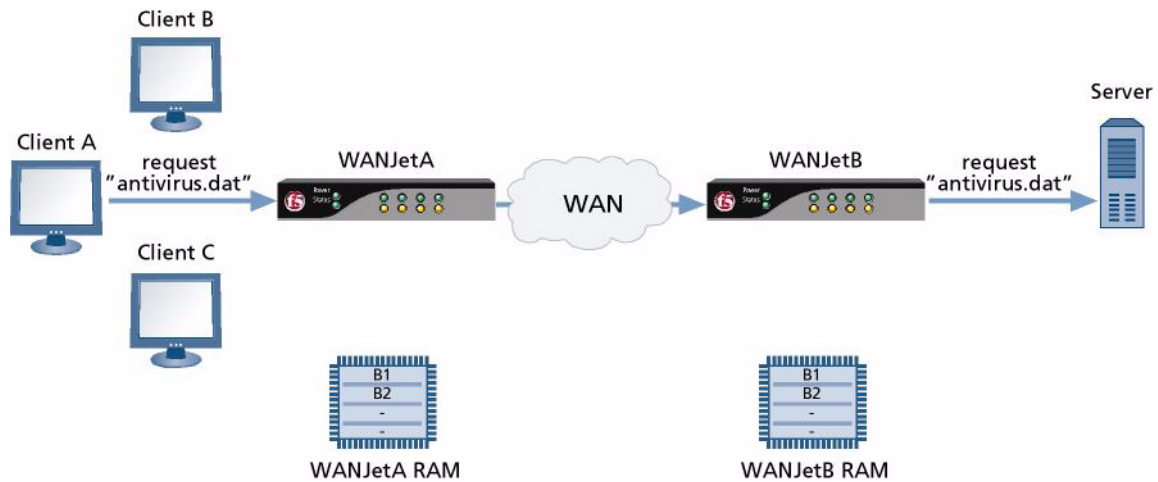


Figure 2.1 Client requests a file

In Figure 2.2, the server on which the file is stored returns the **antivirus.dat** file. WANJetA and WANJetB copy the data to RAM.

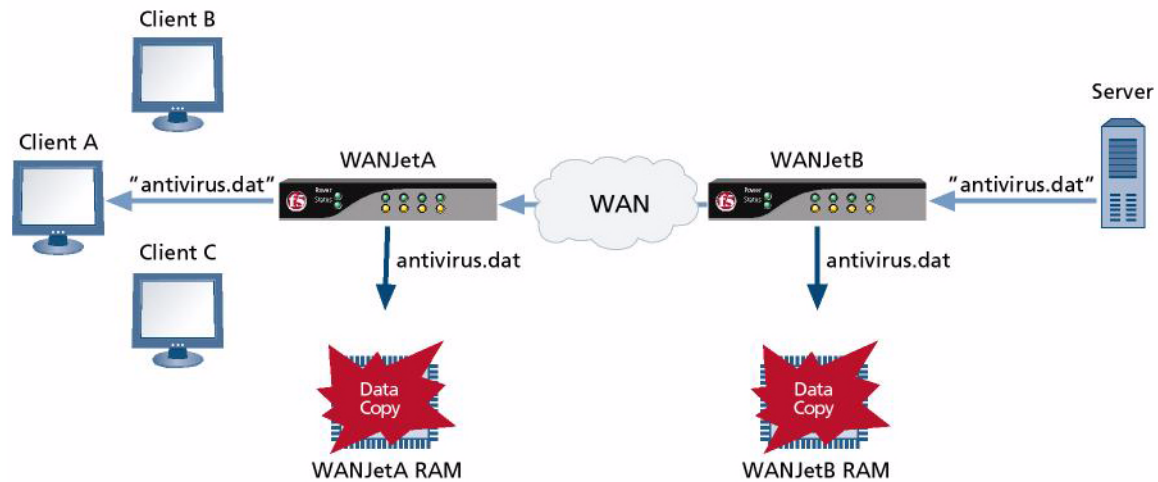


Figure 2.2 Server returns the file

In Figure 2.3, Client B requests the same **antivirus.dat** file.




Figure 2.3 Second client requests the same file

In Figure 2.4, WANJetB compares the **antivirus.dat** file with the data in its RAM to see if the data has changed, confirming that the data in its RAM is still current.




Figure 2.4 WANJetB compares the file to the file in its RAM

Finally, in Figure 2.5, WANJetB sends a message to WANJetA to use the local data instead of resending the file, because the data has not changed. WANJetA sends Client B the **antivirus.dat** file from its local RAM, saving bandwidth over the WAN.




Figure 2.5 WANJetA sends file from RAM

TDR-1: First transfer data reduction and networking adaptivity

After TDR-2 has removed all previously transferred byte patterns, the WANJet appliance applies a second level of data reduction routines called TDR-1. While TDR-2 compression focuses on repeat transfer performance, **TDR-1** improves first transfer performance by examining smaller repetitive patterns and, at the same time, by adapting to changing networking conditions and application requirements.

During periods of high congestion, TDR-1 increases compression levels to reduce congestion and networking queuing delay. During periods of low congestion, TDR-1 reduces compression levels to minimize compression-induced latency. The adaptive nature of TDR-1 ensures that the appropriate compression strategy is applied without degrading application performance.

TDR-1 compresses the remaining network data through intelligent network and application-aware routines that encode the remaining data in as few bytes as possible, improving performance for WAN users.

Application QoS

The goal of Application QoS (Quality of Service) is to provide better service for specific data flows by raising the priority of a specific type of traffic and limiting the priority of other traffic. Accordingly, Application QoS provides complex networks with a guaranteed level of performance for different applications and traffic types. Your network's data transmission is optimized, providing more control over network resources, and ensuring the delivery of mission-critical data.

Utilizing Application QoS policies enables you to downsize the bandwidth consumed over less important network activities and, at the same time, prioritize important and critical data transfer. This way, your bandwidth is used optimally for the transfer of the data that is most important to you.

In addition, the WANJet appliance provides high quality of service for applications that are sensitive to delays by supporting the Voice over Internet Protocol (VoIP).

See *Creating Application QoS policies*, on page 7-3 for information on how to add, edit, or remove Application QoS policies.

Type of Service

The Type of Service (TOS) feature helps to provide the highest quality of data delivery by prioritizing the delivery of one data stream over another. The WANJet appliance deploys the Type of Service methodologies, giving you control over your data streams. You decide which data stream will get to the receiver first by using the Type of Service feature to assign a priority to data traffic using a specific port.

You can assign TOS priorities from 0 to 7, where 0 is the lowest priority, and 7 is the highest. This means that the data using a specific port is transferred according to its priority. For example, you can decide to give the HTTP traffic the lowest priority while giving the FTP traffic the highest priority. You can also assign the same priority, such as priority 7, to multiple protocols. See *Configuring ports and services*, on page 6-4 for instructions on setting TOS priorities.

Simple Network Management Protocol Support

Simple Network Management Protocol (SNMP) governs the management and monitoring of network devices. SNMP sends messages to SNMP-compliant servers, where users can retrieve these messages using SNMP-compliant software. SNMP data is stored in a data structure called a Management Information Base (MIB). An *SNMP trap* provides notification of a significant event (such as a power outage, an error, a fault, or a security violation) that occurred on the network.

The WANJet appliance sends SNMP traps to the SNMP server you specify. The traps you view on the SNMP server are errors for troubleshooting purposes. See *WANJet appliance messages and codes*, on page A-1 for error codes and descriptions.

The WANJet appliance also stores more detailed SNMP reports that you can access using SNMP-compliant software. For the SNMP-compliant software to access the WANJet appliance, it should authenticate itself using a community string you specify. The machine on which the SNMP-compliant software resides should have access to the SNMP data in the WANJet Web UI. See *Granting Web UI access*, on page 5-4.

Figure 2.6 illustrates the interaction between the WANJet appliance and the SNMP traps.

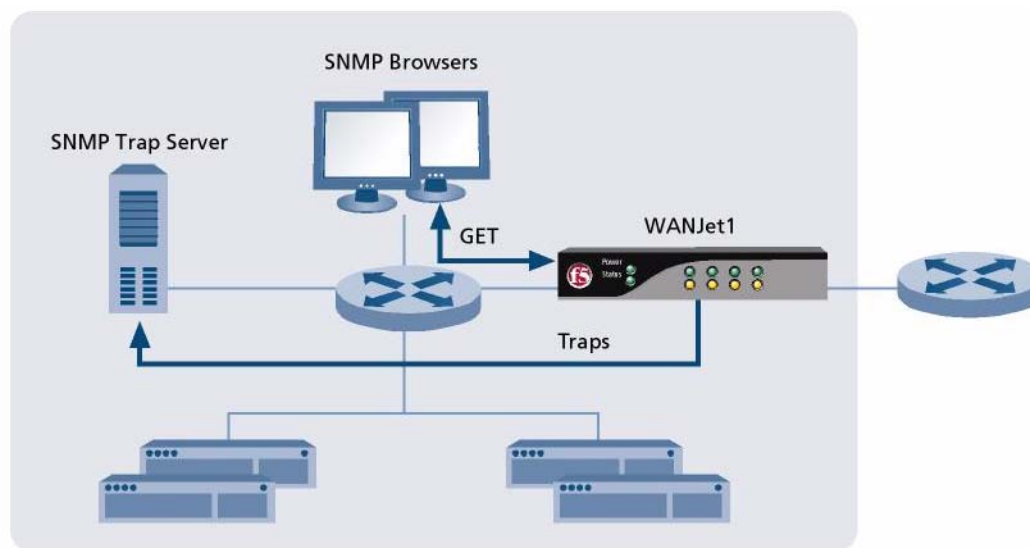


Figure 2.6 WANJet appliance and SNMP data

The Management Information Base (MIB) that stores the SNMP data contains details about the network cards like the network card type, physical address, the card speed, the packets sent and received through each card, the bytes sent and received through each card, and the errors of each card.

In addition, the SNMP reports include detailed information about the WANJet appliance such as total bandwidth saved for sent data and for received data.

For more information about configuring SNMP settings, see *Configuring Syslog and SNMP settings*, on page 6-24.

Remote Monitoring support

Remote Monitoring (RMON) is an extension to SNMP that provides comprehensive network monitoring capabilities. It is a network management protocol that monitors different types of data traffic passing through the network. Unlike SNMP, RMON can gather network data from multiple types of MIBs. Thus, RMON provides much richer data about network usage. For RMON to work, network devices, such as hubs and switches, must be designed to support it.

RMON1 is the Remote Network Monitoring MIB that was developed so that network administrators could see the traffic and collect information about remote network segments for troubleshooting and performance monitoring. RMON1 focuses on Layer 1 and Layer 2 of the OSI model.

RMON2 is an extension of RMON1 that includes open, comprehensive network fault diagnosis, planning, and performance-tuning features. In addition, RMON2 includes monitoring of packets on the higher layers of the OSI model, from Layer 3 to Layer 6. Therefore, RMON2 provides data about traffic on all network layers for network and application monitoring.

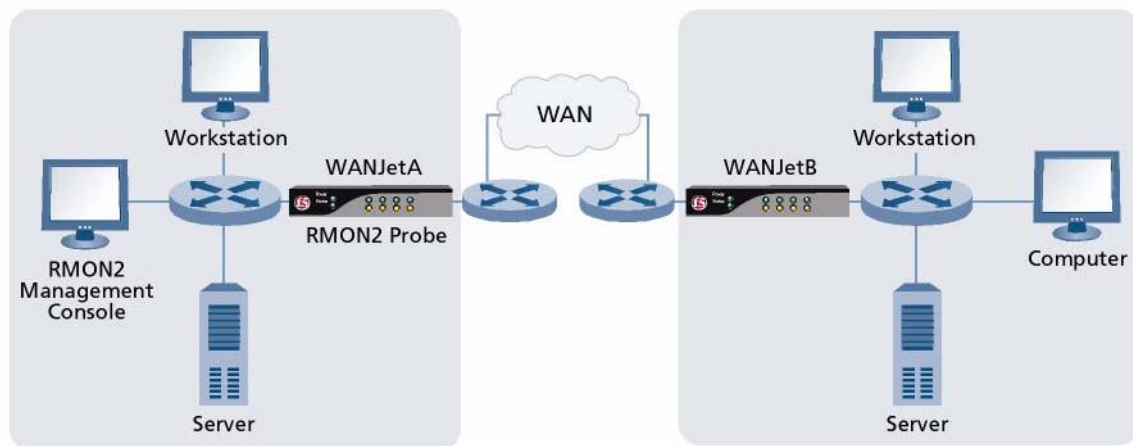


Figure 2.7 WANJet appliance and RMON2

The WANJet appliance supports RMON2 to help administrators gather and analyze detailed information about network traffic, before or after the WANJet appliance processes it, including:

- Data sent and received between two nodes
- IP addresses of the nodes
- Port used to send and receive data
- Data size before and after the WANJet appliance processes the traffic
- Time stamp
- Number of connections

The WANJet appliance supports the RMON2 groups listed in Table 2.1.

RMON2 group	Description
Protocol Directory	Provides a way for an RMON2 application to determine a list of protocols for which the WANJet appliance monitors and maintains statistics.
Network Layer Matrix	Stores and retrieves network layer (IP layer) statistics for conversations between pairs of network addresses.
Application Layer Matrix	Stores and retrieves application layer statistics for conversations between pairs of network layer addresses.

Table 2.1 Supported RMON2 groups

To see where these RMON2 groups fit into the MIB tree, see Appendix C, *RMON2 Tree*. For more information about configuring RMON2, see *Configuring Syslog and SNMP settings*, on page 6-24.

System Log protocol support

System Log (Syslog) protocol provides a way to send event notification messages across IP networks to centralized event message collectors called Syslog servers. Messages are sent at the start or end of a process, or to transmit the current status of a process. The WANJet appliance can send system event messages to the Syslog server you specify. The data log sent by the WANJet appliance includes the sent data, and the received data. In addition, the WANJet appliance can send warning logs to the Syslog server, when necessary.

For more information on how to configure the Syslog settings, see *Configuring Syslog and SNMP settings*, on page 6-24.

Connection Intercept

Connection Intercept (CI) intercepts and resets connections that were initiated before the WANJet appliance became active on the network. If set, this feature ensures that the WANJet appliance resets then optimizes existing connections. As usual, the WANJet appliance optimizes new connections starting after the appliance is up and running.

You can enable Connection Intercept for specific services or ports when creating optimization policies. For details, see *Enabling Connection Intercept*, on page 6-8.



3

Installation

- WANJet appliance deployment
- Firewall guidelines
- Hardware installation
- Site information worksheet

WANJet appliance deployment

This chapter provides conceptual guidelines concerning WANJet appliance installation and configuration. The *Quick Start Card* included in the shipping box with your WANJet appliance provides the initial hardware installation and setup instructions. You can also find the *Quick Start Card* on the F5 Networks Technical Support web site, <http://tech.f5.com>.

Types of deployment

Following are the primary ways to deploy a WANJet appliance within a corporate network:

- ◆ Inline deployment, using one of the following configurations:
 - Point-to-point
 - Point-to-multi-point
- ◆ One-arm deployment, using one of the following configurations:
 - Static routing (non-transparent)
 - Static transparent proxy
 - Transparent proxy with WCCP v2 protocol
 - Transparent proxy with GRE encapsulation

The way you choose to deploy the WANJet appliance depends on your current network topology and requirements.

Inline deployment

Inline deployment is the most common way to deploy WANJet appliances. In this configuration, you place WANJet appliances directly in the path of traffic, or *inline*, between a WAN router and LAN switch.

You can scale inline deployment from a simple point-to-point configuration to a more complex point-to-multi-point configuration.

Point-to-point configuration

Point-to-point configuration is a simple one-to-one topology where you place WANJet appliances at each end of the WAN between their respective WAN routers and LAN switches.

Each WANJet appliance is configured to search for traffic that matches specified source and destination subnets, and ports. If the local WANJet appliance detects a match, it processes the traffic and sends it through a tunnel to the remote WANJet appliance, which, in turn, reverses the process and delivers the packets exactly as they originally were. If there is no match, the local WANJet appliance acts as a bridge, and passes the packets unaltered to the WAN.

Figure 3.1 shows inline deployment with two WANJet appliances in a point-to-point configuration, connecting a corporate data center and one remote office.




Figure 3.1 Inline deployment in point-to-point configuration

Point-to-multipoint configuration

Point-to-multipoint configuration is more complex and involves three or more WANJet appliances. Figure 3.2 illustrates a point-to-multipoint deployment that consists of five appliances that connect to each other across intranets and the Internet.

As with the point-to-point configuration, the WANJet appliance processes traffic that matches user-specified source and destination subnets and ports, and then delivers the traffic across the WAN through a tunnel to the appropriate WANJet appliance.

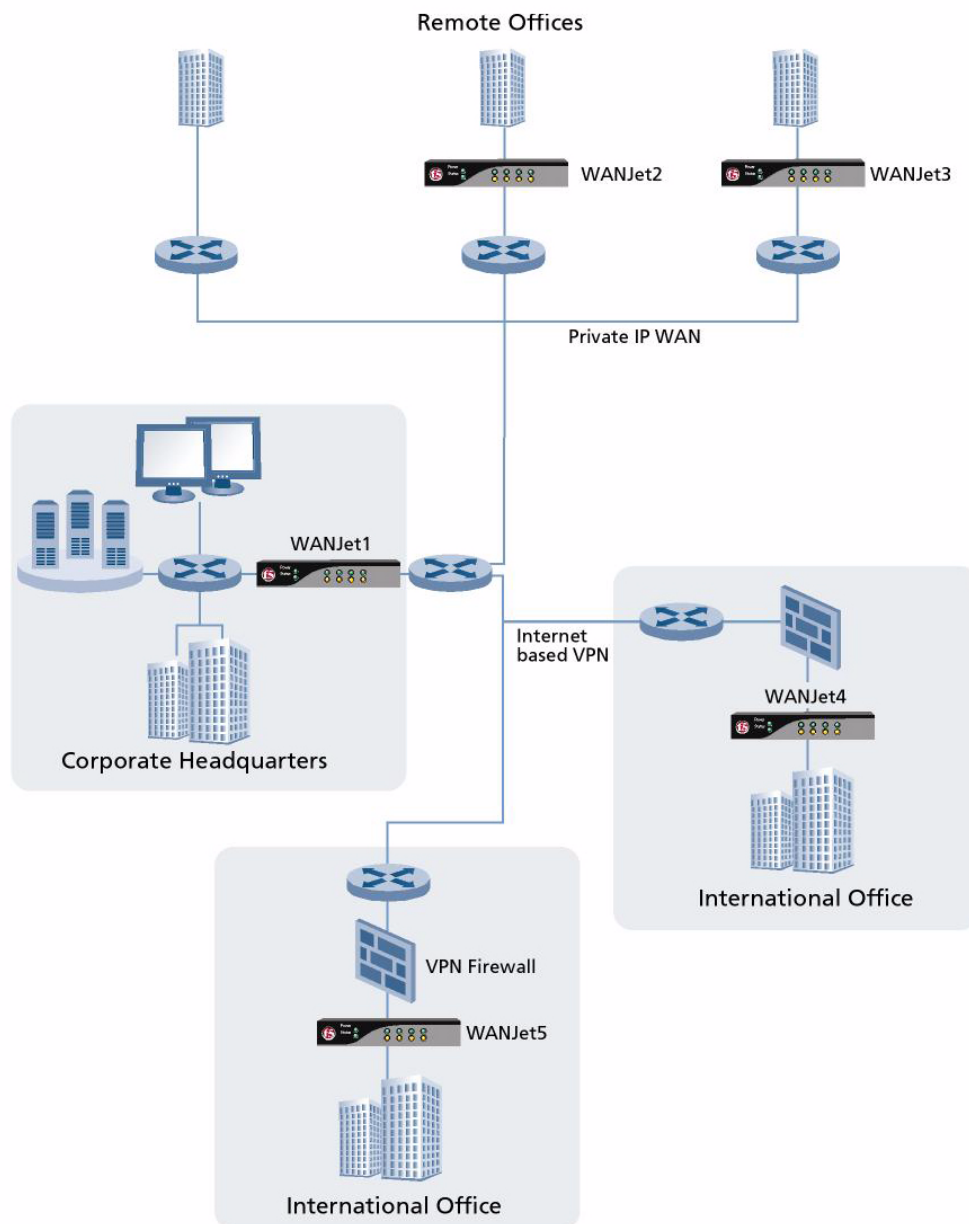


Figure 3.2 *Inline deployment in a point-to-multi-point configuration*

One-arm deployment

In certain cases, it is not desirable or even possible to deploy the WANJet appliance inline. For example, in the case of a collapsed backbone where the WAN router and LAN switch are in one physical device, you may not be able to deploy the WANJet appliance inline.

When inline deployment is not an option, you can use *one-arm deployment*. In this deployment, the WANJet appliance has a single (hence the term *one-arm*) connection to the WAN router (or LAN switch) and has all relevant traffic redirected to it by the WAN router. Figure 3.3 shows a simple one-arm deployment in a corporation that has two networks. Network 1 includes the servers, and network 2 is where the clients are located.




Figure 3.3 One-arm deployment of WANJet appliance

Figure 3.4 shows the basic topology and traffic flow for a one-arm deployment.




Figure 3.4 One-arm topology and traffic flow

The traffic flow sequence shown in Figure 3.4 is as follows.

- **Step 1:** The client initiates a session.
- **Step 2:** A WAN router redirects traffic to the WANJet appliance.

- **Step 3:** The WANJet appliance processes traffic and sends it back to the WAN router.
- **Step 4:** The WAN router forwards traffic across the WAN.

For more information on how to configure one-arm deployment, refer to *Configuring one-arm topology*, on page 6-28.

One-arm deployment methods

You deploy the WANJet appliance using a one-arm configuration in a transparent or non-transparent manner. The transparent method is most common, but there may be certain circumstances when the non-transparent method is more appropriate. This decision depends on your network configuration and specific needs.

As the name implies, the transparent method is totally transparent on the network and requires no modification to any client settings (such as the default gateway). However, you must reconfigure the WAN router.

You can configure the WAN router to redirect traffic by using one of the following methods:

- Static routing (non-transparent)
- Static transparent proxy
- Transparent proxy with WCCP v2 protocol
- Transparent proxy with GRE encapsulation

Using static routing

Static routing is a non-transparent one-arm deployment most suitable for smaller offices. If using static routing, the WANJet appliance connects to a LAN switch, and the LAN switch connects to all of the clients on the network, as well as to the router. Every client on the LAN uses the WANJet appliance as its default gateway. In this deployment, all client traffic is routed to the WANJet appliance. The WANJet appliance can optimize specific traffic, apply different services to specific traffic, and leave other traffic untouched.

This method is non-transparent because, for this to work, all clients have to be reconfigured to use the WANJet appliance IP address as their default gateway. You can reconfigure the clients by either individually modifying each client's default gateway address or, more typically, by updating the DHCP server to provide the WANJet appliance IP address as the default gateway for all of its DHCP clients. All outbound traffic from any client is then first sent to the WANJet appliance for optimization (or passthrough), and the WANJet appliance, in turn, forwards the traffic to the WAN router.

Static routing supports only one subnet (all clients must be in the same subnet as the WANJet appliance), and there is no redundancy. If the WANJet appliance were to fail, clients would no longer have a way to forward traffic to the WAN, just as if a WAN router failed. In a branch office, where support for multiple subnets and redundancy are not as crucial, this deployment mode may be ideal because its principle benefit is its

simplicity. All WAN-bound traffic is automatically sent to the WANJet appliance, which processes it according to defined policies, and sends the traffic on to the WAN. You do not have to reconfigure the WAN router.

In the WANJet appliance Web UI, you set up static routing on the Operational Mode screen by selecting **Static Routing** as the redirection method.

Using static transparent proxy

If using static transparent proxy, the WANJet appliance connects directly to the router and is transparent to the rest of the LAN clients.

The router (by means of a configured routing rule) directs to the WANJet appliance only traffic that the WANJet appliance is configured to process (optimize or applying specific services). You configure the router not to send passthrough traffic to the WANJet appliance. Otherwise, the WANJet appliance drops the passthrough traffic. In this deployment, the WANJet appliance optimizes traffic according to specified policies and then sends all traffic back to the router.

Using transparent proxy with WCCP v2 protocol

Web Cache Communication Protocol (WCCP) was originally developed by Cisco Systems® to specify interactions between one or more routers (or Layer 3 switches) and one or more devices, such as a web cache. The purpose of the interaction is to establish and maintain the transparent redirection of selected types of traffic flowing through a group of routers. WCCP v2 supports traffic redirection to other devices, such as the WANJet appliance. For detailed specifications about the WCCP protocol, see <http://www.faqs.org/rfcs/rfc3040.html>.

The WANJet appliance can use the WCCP protocol to advertise itself to a LAN router as a web cache. Local routers and web caches together form a service group. Routers redirect traffic to the group-member web caches, for example, the local WANJet appliance, in accordance with an algorithm defined for the service group.

The advantage to this deployment method is that it is more tolerant of a failure. If the WANJet appliance fails, the router detects that and handles the traffic properly without sending it back to the WANJet appliance.

If using transparent proxy with WCCP v2 protocol, the WANJet appliance connects to the router directly and is transparent to the LAN clients. You route all LAN traffic to the WANJet appliance just as you do for static transparent proxy.

The difference is that the WANJet appliance communicates with the router using WCCP v2 protocol (the router must support WCCP v2 and you must configure it on the router). In accordance with configured optimization policies, the WANJet appliance determines which traffic to optimize, and which traffic to apply services to. The rest of the traffic is sent back to the router for proper handling.

Using transparent proxy with generic routing encapsulation tunneling

If your network topology uses generic routing encapsulation (GRE), the WANJet appliance may need to process encapsulated traffic within a GRE tunnel. Typically, GRE tunneling connects private IP networks over an Internet connection using two routers (or switches) that support GRE encapsulation.

If you are using GRE tunneling, you can deploy the WANJet appliances using a one-arm configuration so they connect to the routers on both ends of the GRE tunnel. Each router is configured to forward GRE traffic to the WANJet appliance which either optimizes the traffic or sends it through as pass-through traffic.

Figure 3.5 illustrates GRE tunneling and shows how the WANJet appliance optimizes traffic through a TCP tunnel but sends passthrough traffic through the GRE tunnel.

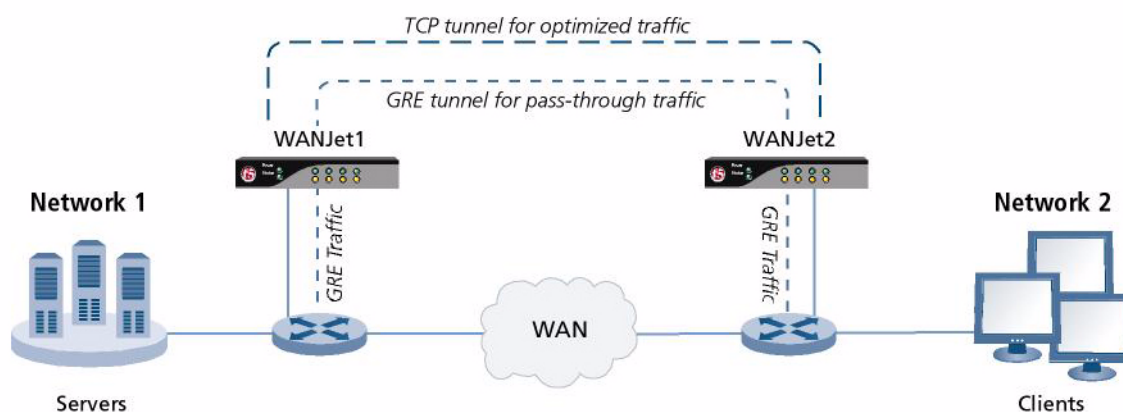


Figure 3.5 WANJet appliance configuration with GRE tunnel

To use this method, on the Operational Mode screen, you select **Transparent Proxy** as the redirection method and **Static** as the discovery method on both WANJet appliances. If you are using GRE tunneling, you also need to configure local (source) and remote (destination) GRE IP addresses on both appliances. You set the **Local GRE IP** (the local end of the GRE tunnel) and **Remote GRE IP** (the remote end of the GRE tunnel) from the Remote WANJets screen. See *Configuring one-arm topology*, on page 6-28 for more details.

When the WANJet appliance receives a GRE packet and recognizes the local and remote GRE IP addresses included in the header, the appliance processes the original encapsulated packet in accordance with its optimization policies. Traffic that the WANJet appliance can optimize uses the TCP tunnel, and pass-through traffic uses the GRE tunnel.

Firewall guidelines

If the WANJet appliance is placed behind a firewall, you must open certain ports for the WANJet appliance to operate properly. Table 3.1 lists the ports that you must open to allow the traffic to pass through the firewall.

Port Number	Used for
53	A UDP port used for DNS.
161	A UDP port used for SNMP.
162	An optional UDP port used for SNMP traps.
22	A TCP port used for SSH.
10000	A TCP port that you use to connect to the Web UI for managing the WANJet appliance remotely.
10001	The default port that the WANJet appliance uses to send real-time chart data.
3701	The default port that the WANJet appliance uses for managing connections.
3702	The default port that the WANJet appliance uses for TCP data tunnels.
3703	The default port that the WANJet appliance uses to proxy UDP over TCP.

Table 3.1 Ports to open when the WANJet appliance is behind a firewall

You must also allow the ICMP protocol to pass through the firewall, so that you can ping the WANJet appliance.

Hardware installation

See the **WANJet® Appliance Quick Start Card** for the WANJet 200, 400, or 500 appliance for instructions on installing WANJet appliances and connecting them to your network. If you have a WANJet 500, refer also to the **Platform Guide: WANJet® 500** for additional details on hardware installation.

Site information worksheet

Use the following site information worksheet to capture relevant site data. When you complete the site information sheet, we recommend that you attach a detailed network diagram for each WANJet appliance site.

Site:	Name:		
	Address:		
	City:		
	State/Province, Country:		
Contact Person:	Name/Title:		
	Email:		
	Work phone:	Cell phone:	
Link:	Type:		
	Speed in Kb/s:		
	Latency:		
	Utilization %:	Peak	Average
Router:	Make:	Model:	
	IP address:		
	Routing protocols used:		
	Static routing table rules:		
Switch:	Make:	Model:	
	IP address:		
WANJet Appliance:	Alias:	IP address:	
	Subnet mask:		
	Default gateway:		
Local Networks:	Alias:	IP address:	Subnet:
	Alias:	IP address:	Subnet:
	Alias:	IP address:	Subnet:
Remote Networks:	Alias:	IP address:	Subnet:
	Alias:	IP address:	Subnet:
	Alias:	IP address:	Subnet:



4

Initial Configuration

- Logging on to the WANJet Web UI
- Activating the license
- Basic WANJet appliance configuration
- Testing connectivity
- Troubleshooting

Logging on to the WANJet Web UI

After you complete the initial hardware configuration, using the LCD panel, a computer connected to the WANJet appliance's serial port, or a secure shell (SSH), you can set up the WANJet appliance using a browser-based interface, called the **Web UI**. By default, you can access the Web UI from any computer that is connected to the network, and can run a web browser.

This chapter describes how to log on to the WANJet Web UI and perform the basic configuration required for the WANJet appliance to start processing traffic. This basic configuration is also covered on the **Quick Start Card** that ships in the box with the WANJet appliance. If you have already performed the basic configuration steps on the **Quick Start Card**, you do not need to repeat them.

After you finish installing and configuring the WANJet appliance, you use the Web UI to administer the appliance and perform additional configuration. You need to log on to the Web UI of each WANJet appliance using the **admin** account to fully configure it.

Another account called **roadmin** is available to log on to the Web UI with read-only access to the configuration settings. You log on the same as you do for the **admin** account using **roadmin** as the user name and default password.

◆ Note

*If your web browser cannot access the Web UI, it may be because Web UI access is restricted to specific IP addresses. You can grant access through the console by specifying the IP address of the machine on which your browser runs. Once you have access, you can use the Web UI to change the list of addresses. See **Granting Web UI access**, on page 5-4.*

To log on to the Web UI

1. In a web browser, access the Web UI using HTTPS and port **10000** in this format:

https://<WANJet_IP_address>:10000

For example, if the IP address of the appliance is **192.168.168.102**, type **https://192.168.168.102:10000** in the web browser. (If you set up the Management port, use the Management IP address instead of the IP address of the WANJet appliance.)

The Welcome screen opens where you can log on.

2. Type the user name and password.
The default user name is **admin** and the default password is **admin** (unless it was changed by a local administrator).

***Note:** F5 Networks recommends that you change the default password to something more secure at your earliest opportunity. See **Changing the Web UI passwords**, on page 5-1, for details.*

3. Click the **Log On** button.
The WANJet appliance Web UI opens.

To log on to the Web UI as a read-only administrator

1. In a web browser, access the Web UI using HTTPS and port **10000** in this format:

`https://<WANJet_IP_address>:10000`

For example, if the IP address of the appliance is **192.168.168.102**, type **`https://192.168.168.102:10000`** in the web browser.

The Welcome screen opens where you can log on.

2. Type the user name and password.
The default user name is **roadmin** and the default password is **roadmin** (unless it was changed by a local administrator).

***Note:** F5 Networks recommends that you change the default password to something more secure at your earliest opportunity. See **Changing the Web UI passwords**, on page 5-1 for details.*

3. Click the **Log On** button.
The WANJet appliance Web UI opens.

Using the Web UI

When you log on to the Web UI of a WANJet appliance, that appliance is considered to be the local WANJet appliance. All other WANJet appliances are remote WANJets appliances in relation to the one you are working on.

When you are logged on to the WANJet appliance, it remains available as long as you are using the Web UI. The WANJet appliance automatically logs you off after 30 minutes of inactivity.

The first screen that you see when you log on to the Web UI is the WANJet Status screen, which displays in the main browser frame. This screen displays a brief summary of the status, IP address, alias, and software version of each remote WANJet appliance. Figure 4.1 shows the parts of the Web UI.

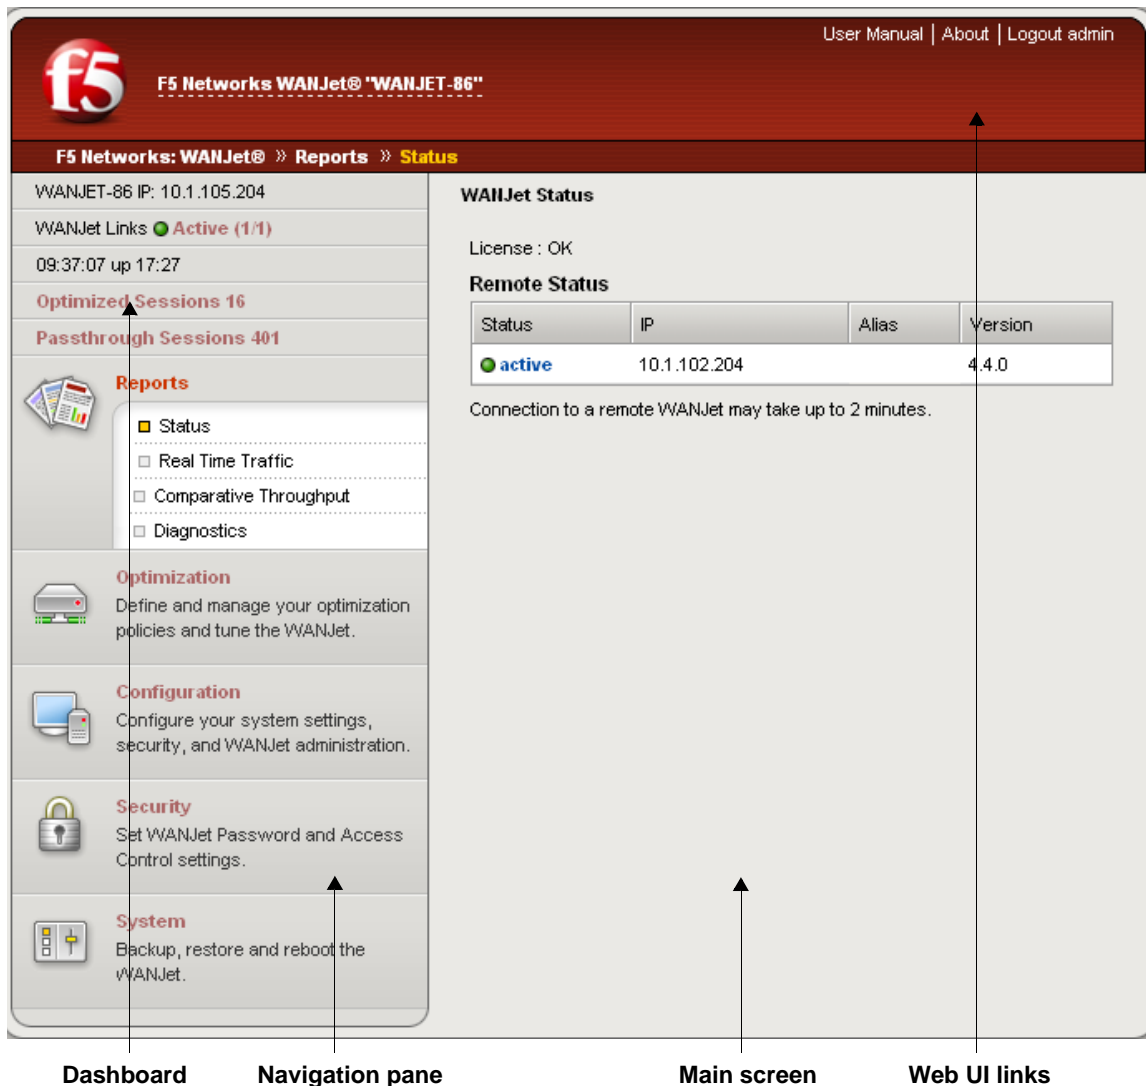


Figure 4.1 WANJet appliance Web UI

For instructions on obtaining additional remote status information, refer to *Status report*, on page 8-2.

Dashboard

The Web UI, below the F5 logo, displays a variety of status indicators and shortcuts. This area, called the *dashboard*, is always visible, regardless of where you are in the Web UI.

The dashboard displays the following information:

- ◆ **WANJet appliance name and IP address**

The alias (name) and the IP address of the local WANJet appliance.

- ◆ **WANJet links**

The number of links to remote WANJet appliances, and a light showing the status of remote WANJet appliances:

- A green light indicates that all links are active.
- A red light displays if no links are active.
- A yellow light displays if only some links are active.

For more information about each link, click the word **Active** on the screen to display the Remote Status report. For more information, see *Status report*, on page 8-2.

- ◆ **Current time and time active**

The current time on the local WANJet appliance and how long (in days, hours, and minutes) it has been active.

- ◆ **Optimized sessions**

The number of WAN sessions that the WANJet appliance is currently optimizing (this number includes all established sessions plus connections that are in the process of being optimized). This links to the Optimized Sessions report (this number lists only established sessions). For more information, see *Optimized Sessions diagnostics*, on page 8-13.

- ◆ **Passthrough sessions**

The number of WAN sessions that are passing through the WANJet appliance, without optimization. This links to the Passthrough Sessions. For more information, see *Passthrough Sessions diagnostics*, on page 8-14.

Navigation pane

The *navigation pane* is the area on the left of the screen, below the dashboard. It includes five sections that you can expand:

- ◆ **Reports**

Provides links for you to view reports and diagnostic logs depicting WANJet appliance performance.

- ◆ **Optimization**

Provides screens to set up the operational mode, optimization policies, application QoS, IT service policies, and WANJet appliance tuning.

- ◆ **Configuration**

Provides screens to allow you to configure local and remote WANJet appliances, interfaces, third party monitoring through SMTP, email alerts, and time.

- ◆ **Security**

Provides screens for changing the administrator passwords, setting up remote authentication, adding a PIN code for accessing the LCD (WANJet 400 and 500 only), and restricting access to the WANJet appliance based on IP addresses.

- ◆ **System**

Provides screens for you to view licensing information, backup and restore configuration settings, initiate an upgrade or boot from a different software image, or shut down the WANJet appliance.

To view other Web UI screens, expand a section in the navigation pane, on the left side of the screen, and click an option. Information displays in the main area of the screen. For example, if a step says to go to the Optimization Policy screen, expand **Optimization** and click **Optimization Policy**. The WANJet Optimization Policy screen replaces the WANJet Status screen in the main browser frame.

Main screen

The main screen is the area of the Web UI that contains reports showing information about WANJet appliance operations, or fields where you can configure how the WANJet appliance works.

Web UI links

The following links always appear at the top right of the Web UI:

- ◆ **User Manual**

Displays the current version of the *WANJet® Appliance Administrator Guide* in PDF format.

- ◆ **About**

Displays an informational screen that contains:

- Information about the WANJet appliance hardware platform
- WANJet appliance version and build number (required when contacting F5 Networks Technical Support)
- Space to add your chassis serial number
- Contact details for F5 Networks Technical Support

- ◆ **Logout admin**

Logs you off the Web UI. It changes to **Logout roadmin** if you are logged in as the read-only administrator.

Logging off the Web UI

To maintain the security of the WANJet appliance, you should log off when you are done using it. Figure 4.2 shows the location of the **Logout** button. The WANJet appliance automatically logs you off after 30 minutes of inactivity for added security.

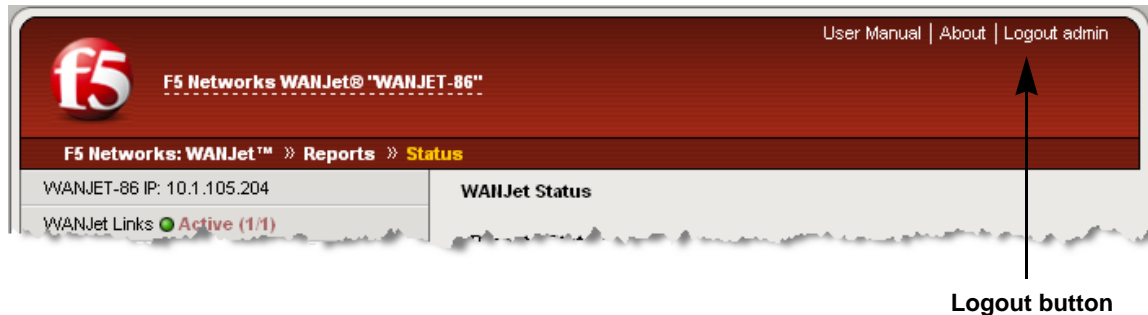


Figure 4.2 Logout button

To log off the Web UI

1. Click the **Logout admin** button in the upper right corner of the Web UI. If you are logged in as **roadmin**, the button reads **Logout roadmin**.
A popup screen verifies that you really want to log off.
2. Click **OK**.

Activating the license

You must activate the license associated with your WANJet appliance after setting the addresses. The WANJet appliance cannot optimize traffic until the license is activated. If you followed the instructions in the *Quick Start Card*, you may have already activated the license. (If the license is already activated, the WANJet Status screen displays **License: OK**.)

You can activate the license automatically or by using the manual procedure. You only need to activate the license once. When you update the WANJet appliance in the future, the license information is retained.

The license purchased for the WANJet appliance is associated with the bandwidth of the WAN link. To increase the bandwidth of that link, you need to contact F5 to obtain a new license, then activate it.

Automatic activation is the easiest method (and is the default method) because the WANJet appliance directly contacts the F5 licensing server and handles the activation. However, in certain cases, you may need to manually activate the license. For example, follow the manual procedure if the WANJet appliance does not have a direct connection to the Internet, or if it resides behind a firewall that does not allow for a direct Internet connection. You can try the automatic method first, and if you receive a message concerning a connection failure, then try the manual method.

To manually activate the license, you need an administrative workstation with a connection to the WANJet appliance and the Internet.

To activate the license automatically

1. Log on to the WANJet appliance. (For details, see *Logging on to the WANJet Web UI*, on page 4-1.)
The WANJet Status screen opens.
2. On the WANJet Status screen, next to **License**, click the **Not entered** link.
The WANJet License Details screen opens.
3. For **Base Registration Key**, you should see your registration key (filled in at the factory).
4. Click **Next**.
The EULA (End User License Agreement) screen opens.
5. Read the EULA, and then click **Accept** if you agree to the conditions.
The WANJet appliance automatically activates the license. When license activation is complete, the WANJet License Details screen opens and shows the current license date.

◆ Note

If automatic license activation does not work, you can use manual activation instead, described in the next section.

To activate the license manually

1. Log on to the WANJet appliance. (For details, see *Logging on to the WANJet Web UI*, on page 4-1.)
The WANJet Status screen opens.
2. On the WANJet Status screen, next to **License**, click the **Not entered** link.
The WANJet License Details screen opens.
3. For **Base Registration Key**, you should see your registration key (filled in at the factory).
4. Next to **Activation Method**, click **Manual**.
5. Click **Next**.
The Manual Activation screen opens.

6. Select and copy the entire contents of the Dossier box (Ctrl + A, Ctrl + C).
7. Click the **Click here to access F5 Licensing Server** link, located below the dossier.
In a separate browser, the Activate F5 License screen opens.
8. On the Activate F5 License screen, paste the dossier that you copied and click **Next**.
The Activate F5 Product displays, still in a separate browser.
9. Read the license, then select **I have read and agree to the terms of this license** and click **Next**.
The license information is displayed on the Activate F5 Product screen.
10. Select the entire license (Ctrl + A, Ctrl + C), and paste it into the Manual Activation screen after **Step 3: License**.
11. Click **Next**.
The WANJet License Details screen displays with the current license date.
12. In the navigation pane, expand **Reports** and click **Status**.
The license should now display as **OK**.

Basic WANJet appliance configuration

You must set up WANJet appliances in pairs, with one appliance on each side of the WAN link. You can perform the configuration steps for both appliances either on each physical appliance, or from a single computer by logging on to the Web UI remotely.

Figure 4.3 shows two WANJet appliances that are deployed in a point-to-point configuration.




Figure 4.3 Inline deployment in point-to-point configuration

The WANJet appliances in this example are connected as follows:

- WANJet1 is in the data center and connects with local IP address **175.16.2.1**.
- WANJet2 is in a remote office and connects to the remote end of the private IP WAN link with IP address **10.2.0.1**.

For this example, basic WANJet appliance configuration includes the following steps:

- Naming the first WANJet appliance
- Configuring multiple subnets (if required)
- Defining the second WANJet appliance as a remote WANJet appliance on the first WANJet appliance
- Naming the second WANJet appliance
- Defining the first WANJet appliance as a remote WANJet appliance on the second WANJet appliance

Naming the first WANJet appliance

You configure WANJet appliances in pairs to optimize the traffic that flows between them. A pair of WANJet appliances consists of a local WANJet appliance and a remote WANJet appliance, one on either side of a WAN link. A typical configuration might include one WANJet appliance in a data center where company servers reside, and a second WANJet appliance on the other side of the WAN in an office where employees work.

You can start by naming the first WANJet appliance in the pair.

To name the first WANJet appliance

1. Into a browser, type the address and port for the first WANJet appliance. For this example, you type the following URL in the browser for WANJet1:
`https://175.16.2.1:10000`
2. Log on to the Web UI.
The user name is **admin**. The default password is **admin**.
3. In the navigation pane, expand **Configuration** and click **Local WANJet**.
The Local WANJet screen opens.
4. In the **WANJet Alias** box, type a name for the WANJet appliance.
For example:
`WANJet1`
5. Click the **Save** button.

Configuring multiple subnets

If your local area network has multiple subnets connected through a router, you need to configure the local router IP address and add the local subnets that you want to optimize to the WANJet appliance. You can add specific subnets, or you can optimize all local subnets.

Once the WAN link between the WANJet appliance pair is up, subnet specifications are automatically exchanged between the appliances. So, for example, the local subnets specified on WANJet1 appear as remote subnets on WANJet2, and local subnets on WANJet 2 appear as remote subnets on WANJet 1.

Before performing the following steps, verify that you require additional subnets, and decide whether you want to optimize all of them or selected subnets.

To configure multiple subnets

1. In the navigation pane, expand **Configuration** and click **Local WANJet**.
The Local WANJet screen opens.

2. In the **LAN Router** box, type the router's IP address.
This address is the next-hop router in your LAN.
3. Click the **Save** button.
4. In the navigation pane, expand **Optimization** and click **Optimization Policy**.
The Optimization Policy screen opens.
5. Check the **Include WANJet Subnet** box, if it is not already checked (by default, it is checked).
Note: If you do not check this box, you must have a specific reason for not optimizing the traffic from the subnet that includes the first WANJet appliance.
6. Specify how you want to optimize subnets.
 - If you want to optimize all subnets, check the **Optimize All Subnets** box and skip to step 13.
 - If you want to optimize selected subnets, continue with steps 7-12 to add them.
7. Click the **Add** button beneath **Local Subnet**.
The Add Local Subnet screen opens in a separate browser window.
8. In the **Local Subnet** box, type the IP address for the subnet.
You can use the shorthand address format of, **xxx.xxx.xxx.xxx/nn**, to provide both the subnet address and the subnet mask. For example:
175.16.2.0/24
Where **/24** means that the first 24 bits of the address must match the local subnet address and the address of any host in the subnet is defined by the last 8 bits of the address. For example, **175.16.2.6** is a valid address for the subnet defined in this configuration example.
9. In the **Netmask** box, type the subnet mask. For example:
255.255.255.0
Note: If you entered the subnet address in the /nn format, as described in the previous step, the system automatically populates the corresponding subnet mask box.
10. In the **Alias** box, type a string to serve as a name for the subnet.
For example:
Subnet A
11. For **Operational Status**, click the **Enabled** button.
12. Click the **OK** button.
The Optimization Policy screen displays with the new subnet in the list of local subnets.
13. Click the **Save** button.
14. Repeat steps 7 through 13 to add as many subnets as required.

Defining the second WANJet appliance as a remote WANJet appliance

After you finish adding subnets to the first WANJet appliance, define the second appliance as a remote WANJet appliance of the first one.

To define the second WANJet appliance as a remote WANJet appliance

1. In the navigation pane, expand **Configuration** and click **Remote WANJets**.
The Remote WANJets screen opens.
2. Click the **Add** button.
The Manage Remote WANJet screen displays in a new browser window.
3. Leave the **WANJet Type** set to **Single**.
*Note: For information about configuring redundant WANJet appliances, refer to **Configuring redundant peers**, on page 6-26.*
4. In the **WANJet IP** box, type the IP address for the remote WANJet appliance. For example:
10.2.0.1
5. In the **WANJet Alias** box, type a name for the remote WANJet appliance. For example:
WANJet2
6. Leave the **WANJet Port** setting.
7. In the **Shared Key** box, type the shared key that was assigned by the network administrator. The only requirement for the key is that it matches the key added for its partner on the corresponding system pair. For this example, you must use the same key when adding WANJet1 as a remote WANJet appliance to WANJet2.
8. Leave the **Local GRE IP** and **Remote GRE IP** addresses blank unless you are using GRE tunneling between two routers with a one-arm WANJet appliance configuration.
9. Click the **OK** button.
The browser window closes.
10. On the Remote WANJet screen, click the **Save** button.
The new remote WANJet appliance appears in the Remote WANJet appliance list.

Naming the second WANJet appliance

After you finish configuring the first WANJet appliance, you can configure the second WANJet appliance in the pair. The second WANJet appliance must already be installed as described in the *Quick Start Card* included in the shipping box.

To name the second WANJet appliance

1. Into a browser, type the address and port for the second WANJet appliance. For example, you could type the following URL in the browser for WANJet2:

https://10.2.0.1:10000

2. Log on to the Web UI.
The user name is **admin**. The default password is **admin**.
3. In the navigation pane, expand **Configuration** and click **Local WANJet**.
The Local WANJet screen opens.
4. In the **WANJet Alias** box, you may type a name for the WANJet appliance. For example:
WANJet2
5. Click the **Save** button.

Configuring multiple subnets on the second WANJet appliance

If you have defined a LAN router and added subnets for WANJet1, you must do the same for WANJet2, unless WANJet2 is on a simpler LAN. Refer to steps 1-11 in *Configuring multiple subnets*, on page 4-10, for instructions.

Defining the first WANJet appliance as a remote WANJet appliance

After you have named the second WANJet appliance, define the first appliance as a remote WANJet appliance on the second WANJet appliance.

To define the first WANJet appliance as a remote WANJet appliance

1. In the navigation pane, expand **Configuration** and click **Remote WANJet**.
The Remote WANJets screen opens.
2. On the Remote WANJets screen, click **Add**.
The Manage Remote WANJet screen displays in a new browser window.

3. In the **WANJet IP** box, type the IP address for the remote WANJet appliance. For example:
175.16.2.1
4. In the **WANJet Alias** box, type a name for the remote WANJet appliance. For example:
WANJet1
5. In the **Shared Key** box, type the shared key.
This key is assigned by the network administrator. The only requirement for the key is that it matches the key added for its partner on the corresponding system pair. For this example, you must use the same key when adding WANJet2 as a remote WANJet appliance to WANJet1.
6. Leave the settings as they are for **WANJet Type** and **WANJet Port**.
7. Leave the **Local GRE IP** and **Remote GRE IP** addresses blank unless you are using GRE tunneling between two routers with a one-arm WANJet appliance configuration.
8. Click the **Logoff** button.
9. Close the browser window.

Testing connectivity

When the WAN link is established between the WANJet pair, the two WANJet appliances automatically exchange subnet specifications. For example, the local subnets that you specify for WANJet A become remote subnets for WANJet A in WANJet B's Remote WANJet appliance configuration information.

You can test the connectivity between the local and remote WANJet appliances by viewing the following details on each:

- Status of remote WANJet appliance(s)
- Traffic passing through network
- Diagnostics

For additional information about WANJet appliance reports, such as those described in the following procedures, see Chapter 8, *Monitoring Performance*.

To view the status of the remote WANJet appliance

In the navigation pane, expand **Reports** and click **Status**.

A green light displays next to the IP address for remote WANJet appliances that are enabled and connected.

To view traffic passing through the network

1. In the navigation pane, expand **Reports** and click **Comparative Throughput**.
2. Click **Total Throughput**, **Sent Throughput**, and **Received Throughput** to view the various reports.

To view diagnostics

1. In the navigation pane, expand **Reports** and click **Diagnostics**. The Diagnostics screen opens.
2. From the Connectivity menu, choose Remote WANJets. The Diagnose Remote WANJet report opens.
3. Review the status of each remote WANJet appliance. The status should be **Active**.

Additional configuration tasks

The initial configuration steps described in this chapter are only the minimal steps you need to take to establish a WAN link between two WANJet appliances and start optimizing traffic between the two.

When you have completed the initial configuration steps, we recommend (but do not require) that you perform additional administrative tasks, such as the following:

- Change the Web UI password, see page 5-1.
- Use the Management port, see page 5-5.
- Synchronize the time automatically, see page 5-7.
- Configure email alerts, see page 6-25.
- Adjust tuning settings, see page 6-13.

You can also fine-tune the optimization policies for the WANJet appliances. Refer to *Creating optimization policies*, on page 6-1.

Troubleshooting

One of the first steps we recommend for troubleshooting the WANJet appliance is to create a system snapshot immediately. It provides detailed information about the WANJet appliance, including:

- Date and time of the snapshot
- WANJet appliance version and build number
- Connection counts
- System and network status
- Recent errors
- Status and configuration settings of the interfaces

Refer to *System Snapshots*, on page 8-27 for information on taking a system snapshot. You can provide the system snapshot to the F5 Networks Technical Support team to help resolve technical issues.

Some common problems are listed in Table 4.1. If you are experiencing an issue that is not included in the following table, contact http://www.f5.com/customer_support/ for assistance.

Issue	Suggested actions
I cannot ping the WANJet appliance.	<p>Verify that the computer from which you are pinging has a valid network connection.</p> <p>Try pinging other known devices.</p> <p>Verify that you are using the correct IP address for the appliance, by reading it from the LCD display.</p>
I can ping the WANJet appliance, but I cannot ping the WAN gateway.	<p>Verify that the cabling is connected properly, as described in the Quick Start Card.</p> <p>Make sure that you connected the gateway router to the WANJet appliance's WAN port, using the supplied crossover cable.</p>
I cannot see that the WANJet appliance is optimizing traffic, or the optimization is extremely low.	<p>Review your configuration of local and remote subnets at both appliances. You might have heavy traffic on a subnet that is not included in the WANJet appliance's configuration. You must include all subnets for which traffic should be optimized.</p>
My browser connection times out when I attempt to access the Web UI.	<p>Check to see that you are accessing the correct URL for the Web UI. If you enter just http:// followed by the WANJet appliance's IP address, it will not work. You must connect to port 10000 using the secure HTTPS protocol. For example: https://123.123.123.123:10000/</p> <p>See <i>Logging on to the WANJet Web UI</i>, on page 4-1.</p>

Table 4.1 Troubleshooting suggestions

Issue	Suggested actions
When I attempt to access the Web UI, I get a Page Not Found error.	<p>If you are certain that you entered the URL correctly and the WANJet appliance appears to be running, it may indicate that the computer from which you are running your web browser does not have access to the Web UI. Although the default setting grants access to all machines, that setting can be changed to limit access based on IP address.</p> <p>Use the LCD to add your computer's IP address to the list for access. After that, use the Web UI to change the access settings. For instructions, see <i>Granting Web UI access</i>, on page 5-4.</p>
I can access the Login screen for the Web UI, but my browser connection times out when I try to log on.	<p>This issue can occur when the WANJet appliance is not able to access the RADIUS authentication server or when the Timeout and NRetry variables are set too high. See <i>To configure the WANJet appliance for remote RADIUS authentication</i>, on page 5-2.</p> <p>Log on as a local user, using the admin user name and a default password of admin (note that the local administrator may have changed the default password). After you are logged in, in the navigation pane, expand Security and click Remote Authentication, and verify that RADIUS authentication is enabled.</p> <p>Review the Timeout and NRetry values. F5 Networks recommends a value of 3 for each of these settings. If these settings are too high, authentication might take a long time to fail, causing the connection to time out. For information, see <i>Configuring remote authentication</i>, on page 5-2.</p>
The Link LED (for the WAN or LAN port) does not light up.	<p>Verify that the cables are installed properly on the WANJet appliance.</p> <p>Check to see if the ports on the WAN router and the LAN switch connected to the WANJet appliance are set to autonegotiate. If either port is forced to a specific link speed and duplex value, you must set the WANJet port to match this value. For information about resetting the NIC configuration (link speed and duplex value) for a WANJet port, see <i>Changing the interface speed</i>, on page 6-22</p> <p>F5 Networks strongly recommends that if you force the link for one of the WANJet ports, you force the link for both ports. This prevents link problems in pass-through mode if power to the WANJet device is lost.</p>

Table 4.1 Troubleshooting suggestions (Continued)



5

Managing the WANJet Appliance

- Configuring authentication settings
- Granting Web UI access
- Using the Management port
- Configuring time settings
- Shutting down and restarting the WANJet appliance
- Booting from an alternate image
- Backing up and restoring settings
- Restoring factory default values
- Upgrading the WANJet appliance software

Configuring authentication settings

To maintain the security of the WANJet appliance settings, the Web UI is password-protected. Some WANJet appliances, including the 400 and 500, have LCDs on the front of the unit and can be PIN-protected.

The **admin** account is the primary account that you use to access the Web UI. Administrators who log on as **admin** can view and change all settings on the WANJet appliance. (See *To log on to the Web UI*, on page 4-1, for details on how to log on.) There is only one **admin** password for all users of the Web UI.

Administrators can also log on to the Web UI using a restricted account called **roadmin** (for read-only admin). The **roadmin** is a read-only user who can view the WANJet Web UI, but cannot reconfigure the appliance. There is no command line interface to the **roadmin** account. The user name and the default password for the account are both **roadmin**. (See *To log on to the Web UI as a read-only administrator*, on page 4-2, for details on how to log on.)

Changing the Web UI passwords

An administrator logged in as **admin** can change the password and/or PIN code at any time. F5 recommends that you immediately change the administrator account passwords and the PIN code from the defaults, and then change them regularly (once a month, for example) thereafter.

You can also log on as **admin** from the command line interface on a computer connected through the serial interface. You can change the password using the **passwd** command. Only an administrator who is logged in as **admin** can change the **roadmin** password in the Web UI, or at the command line using the **passwd** command.

You can also use remote accounts to access the Web UI; however, you cannot change the passwords for remote accounts from the WANJet Password screen. For more details, refer to *Configuring remote authentication*, following.

Important

*Since there is only one **admin** password for the Web UI, be sure to warn all other users that you are changing the password (unless they are using remote authentication).*

To change the WANJet Web UI password

1. In the navigation pane, expand **Security** and click **Password**. The WANJet Password screen opens.
2. In the **Old Password** box, type the old password.

***Note:** If you did not change the default password during the initial configuration, leave this box blank.*

3. In the **New Password** box, type the new password.
As a general rule, passwords should consist of at least 6 characters, and include a mixture of lowercase and uppercase letters, numbers, and punctuation marks. A blank password is not allowed.
4. In the **Confirm Password** box, retype the new password.
This must match the password that you typed in the **New Password** box.
5. Click the **Save** button to save the new password.
A popup confirmation screen opens.
6. Click **Yes** to confirm the new password.

Configuring remote authentication

You can choose to authenticate WANJet appliance administrators using the WANJet appliance's local database, or using remote authentication on a RADIUS server.

To configure the WANJet appliance for local authentication

1. In the navigation pane, expand **Security** and click **Remote Authentication**.
The WANJet Remote Authentication screen opens.
2. Select **No Remote Authentication**.
3. Click the **Save** button.

To configure the WANJet appliance for remote RADIUS authentication

1. In the navigation pane, expand **Security** and click **Remote Authentication**.
The WANJet Remote Authentication screen opens.
2. Select **RADIUS**.
The screen displays additional settings.
3. In the **Server IP** box, type the IP address of the RADIUS server.
4. In the **Secret** box, type the server's shared secret.
This is the key that authenticates RADIUS transactions between the local WANJet appliance and the RADIUS server.
5. In the **Timeout** box, type the number of seconds that the WANJet appliance should wait after sending a RADIUS request. After this time has expired, WANJet appliance stops waiting for a response. F5 Networks recommends a value of **3**.

6. In the **NRetry** box, type the number of times that you want the WANJet appliance to send a RADIUS request to the server before giving up. F5 Networks recommends a value of **3**.

*Note: If you type a value in the **Timeout** box, you must also enter a value in the **NRetry** box. If you set the values too high, it could take a long time to determine that the server is not responding to a login attempt. This problem is compounded if you are using more than one RADIUS server.*

7. Click the **Add** button to store the new information.
8. Repeat Steps 2 through 7 for any additional RADIUS servers.
9. Click the **Save** button.
The Remote Authentication screen refreshes with the RADIUS server details that you added.

Now that you have configured the WANJet appliance to use remote authentication, you can view diagnostic RADIUS reports. For more information, see **Connectivity**, on page 8-20 for details about this report. For information about RADIUS protocol, refer to <http://www.ietf.org/rfc/rfc2865.txt>.

If you later want to edit the remote authentication settings, you need to delete the information (click **Delete** next to the server information) and add it again.

Changing the WANJet LCD PIN code

There is no default PIN code set for the Liquid Crystal Display (LCD) on the WANJet appliance. If you assign a pin code, you need to enter it before you can change the WANJet appliance configuration using the LCD.

To create or change the LCD PIN code

1. In the navigation pane, expand **Security** and click **LCD PIN**.
The LCD PIN screen opens.
2. In the **Old PIN** box, type the old PIN.
This is a four-digit number.
3. In the **New PIN** box, type the new PIN.
This must be a four-digit number.
4. In the **Confirm PIN** box, retype the new PIN.
It must match the PIN that you typed in the **New PIN** box.
5. Click the **Save** button.
6. Restart the WANJet appliance. See *To restart the WANJet appliance using the Web UI*, on page 5-9.

Granting Web UI access

You can restrict or allow access to the WANJet appliance's Web UI, and the SNMP reports residing on it, to specific WANJet appliances or subnets as follows:

- **Allow all addresses** (default)
Permit all IP addresses on the network to access the Web UI.
- **Allow Listed Addresses**
Permit only a specified list of IP addresses on the network to access the Web UI.
- **Deny Listed Addresses**
Prevent a specified list of IP addresses on the network from accessing the Web UI.

Once this is configured, if an administrator tries to log on to the Web UI from a restricted IP address (that is, an IP address that is not allowed or is denied access), the browser returns a **404: Page Not Found** error. Restrictions apply to both the **admin** and **roadmin** administration accounts.

◆ **Note**

*To ensure that only specific users access the Web UI, you can create a password for the Web UI and provide this password only to approved personnel. See **Changing the Web UI passwords**, on page 5-1.*

To allow specific IP addresses access to the Web UI

1. In the navigation pane, expand **Security** and click **IP Access Control**.
The WANJet IP Access Control screen opens.
2. Select **Allow Listed Addresses**.
3. In the box, type the IP addresses where an administrator can log on to the Web UI. At a minimum, specify the IP addresses for the following:
 - The SNMP server, so that you can view SNMP and RMON2 reports. (See *Configuring Syslog and SNMP settings*, on page 6-24.)
 - The Syslog server, so that you can view Syslog data. (See *Configuring Syslog and SNMP settings*, on page 6-24.)
 - The WANJet appliance from which you are currently accessing the Web UI through a browser, and any other WANJet appliance from which you and other administrators want to access the Web UI.
4. Click the **Save** button.

To deny specific IP addresses access to the Web UI

1. In the navigation pane, expand **Security** and click **IP Access Control**.
The WANJet IP Access Control screen opens.
2. Select **Deny Listed Addresses**, and in the box, type the IP addresses of users who do not have permission to log on to the Web UI.
3. Click the **Save** button.

Using the Management port

The WANJet appliance has a port called the *Management port* that you can use for out-of-band management. *Out-of-band management* provides a dedicated management channel (separate from the data channel) that is used for administration only. Use of the Management port is optional. If you use the Management port, you will want to log on to the Web UI using its IP address rather than the WANJet IP address (although it is possible to use either).

The advantage of using the Management port is that it provides a way to separate the WANJet appliance management data from the data that is being optimized. You can connect the Management port to a separate subnet dedicated to a management network, for example, where only administrators have access. You can set up the Management port from the LCD or from the Web UI.

The following procedure describes how to set up the Management port from the Web UI if you did not configure it when you initially installed the WANJet appliance hardware.

To set up the Management port

1. On the WANJet appliance, plug an Ethernet cable into the Management port and connect the other end to your management network.
2. Access the Web UI using the WANJet IP address:
`https://<WANJet_IP_address>:10000`
3. Log on to the WANJet appliance using the **admin** user name and password.
4. In the navigation pane, expand **Configuration** and click **Local WANJet**.
The Local WANJet screen opens, and the Local WANJet Management section appears at the bottom of the screen.
5. In the **Management IP** box, type the IP address you want to use to manage the WANJet appliance.

6. In the **Management Netmask** box, type the netmask for the Management port.
7. In the **Management Gateway** box, type the IP address for the gateway to the management network.
8. Click the **Save** button.
9. Access the Web UI using the Management IP address:
`https://<Management_IP_address>:10000`

Configuring time settings

Time management for the WANJet appliance involves setting the time zone and synchronizing all linked WANJet appliances. Synchronizing the time settings is one of the most frequent administrative management tasks that you perform.

The settings include:

- **Time zone**
When you initially configure a WANJet appliance, you must set the time zone and the first day of the week.
- **Time server**
With this option, you can choose a time server to use for automatic time synchronization for all WANJet appliances.
- **Time**
With this option, you can set the current time manually for the WANJet appliance.

Setting the time zone

Use the following procedure to set the time zone and the first day of the week for the WANJet appliance.

To set the time zone

1. In the navigation pane, expand **Configuration** and click **Time**. The WANJet Time Settings screen opens.
2. In the Timezone section, from the **Current location** list, select the closest geographical location to your site.
3. From the **First Day of the Week** list, select a day.
4. Click the **Change timezone** button to save the changes.

Repeat these steps on every WANJet appliance in your network.

Synchronizing time automatically

You can use a specific time server to synchronize the WANJet appliance's time automatically. The IP addresses of several commonly used time servers are provided, or you can specify the address of another time server.

◆ **Note**

For information about time servers, refer to www.eecis.udel.edu/~mills/ntp/clock2a.html.

To use a time server to synchronize time automatically

1. In the navigation pane, expand **Configuration** and click **Time**. The WANJet Time Settings screen opens.
2. In the Time Server section, from the **Host/Address** list, select the IP address of the time server you want to use. Alternatively, select **User Specified**, and in the box, type the IP address of the preferred time server.
3. Click the **Sync time** button to save the changes.

Repeat these steps for every WANJet appliance in your network.

Setting the time manually

You can adjust the time on your WANJet appliances manually through the Web UI, instead of synchronizing with a time server.

To set the date and time manually

1. In the navigation pane, expand **Configuration** and click **Time**. The WANJet Time Settings screen opens.
2. In the Time section, use the **Day**, **Month**, **Year**, **Hour**, **Minute**, and **Second** settings and select the appropriate options.
3. Click the **Set time** button to save the changes.

Repeat these steps for every WANJet appliance in your network.

Shutting down and restarting the WANJet appliance

Shutting down WANJet appliance stops all data processing. You can shut down or restart the WANJet appliance from the Web UI or the LCD panel (WANJet 400 and 500 appliances only). Because the LCD panels on the WANJet 400 and 500 are different, the procedures for restarting or shutting them down differ slightly.

◆ Important

Notify your users before you shut down or restart a WANJet appliance, because network performance is affected.

To shut down the WANJet appliance using the Web UI

1. In the navigation pane, expand **System** and click **Shutdown & Restart**.
The WANJet Shutdown & Restart screen opens.
2. Click **Shutdown WANJet**.
A confirmation message appears on the LCD screen.
3. Click the **OK** button to shut down the WANJet appliance.

To shut down the WANJet 400 appliance using the LCD

1. On the front LCD panel of the WANJet 400 appliance, press **X** (Cancel) to activate the main menu.
2. Press **✓** (Enter) to display the Setup menu.
3. From the menu, choose Shutdown.
4. Press **✓** (Enter).
A confirmation message appears on the LCD screen.
5. Press **✓** (Enter) to shut down the appliance, or press **X** to cancel and escape the menu sequence.
The system shuts down.
6. Turn off the WANJet 400 appliance completely by pressing the On/Off button located on the back of the appliance.

To shut down the WANJet 500 appliance using the LCD

1. On the LCD keypad of the WANJet 500 appliance, press **►** twice.
The LCD displays **Menu**, followed by **>Configure**.
2. Press **▼** three times to go to **Shutdown**.
3. Press **►** once.
The LCD displays **Shutdown Now?**
4. Press **►** once to shut down the WANJet appliance.
The system shuts down, and then power turns off.

To restart the WANJet appliance using the Web UI

1. In the navigation pane, expand **System** and click **Shutdown & Restart**.
The WANJet Shutdown & Restart screen opens.
2. Click the **Restart WANJet** button.
A confirmation message appears on the LCD screen.
3. Click the **OK** button to restart the WANJet appliance.

To restart the WANJet 400 appliance using the LCD

1. On the front LCD panel of the WANJet 400 appliance, press the **X** (Cancel) button to activate the main menu.
2. Press the **✓** (Enter) button to display the **Setup** menu.
3. From the menu, choose **Restart**.
4. Press the **✓** (Enter) button.
A confirmation message appears on the LCD screen.
5. Press **✓** to restart the WANJet appliance, or press **X** to cancel and escape the menu sequence.

To restart the WANJet 500 appliance using the LCD

1. On the LCD keypad of the WANJet 500 appliance, press **►** twice.
The LCD displays **Menu**, followed by **>Configure**.
2. Press **▼** twice to go to **Restart**.
3. Press **►** once.
The LCD displays **Restart Now?**
4. Press **►** once to restart the WANJet appliance.

Booting from an alternate image

You typically have two software images on the flash memory card of the WANJet appliance. One image is active and the other is inactive. If something goes wrong with the first installation, you can boot from the alternate image.

Also, having two images lets you upgrade one yet maintain the current installation for cases when you want to test an upgrade without losing the previous version. When you perform an upgrade, the WANJet appliance copies the configuration settings from the current installation to the upgraded image.

◆ Important

*If you boot from the second image without upgrading, the WANJet appliance initially starts up with the default settings. Before you boot the alternate image, you can create a backup file of the current image settings, boot the alternate image, then restore the settings from the backup file. See **Backing up and restoring settings**, following.*

To boot the WANJet appliance from the alternate WANJet appliance image

1. In the navigation pane, expand **System** and click **Upgrade & Boot Menu**.
The WANJet Boot Menu screen displays the WANJet appliance's version and build number for each image. The active version has a green button next to it, and the inactive image has a red button next to it.
2. Click the **Make Active** button next to the image you want to activate.
A popup confirmation request opens.
3. Click the **Yes** button.

Backing up and restoring settings

F5 recommends that you create backups of your current WANJet appliance settings on a regular basis. You should also perform a backup before making any major changes to the settings. This makes it easy to restore the system in the event of a failure. Backing up your current content is one of the most frequent administrative management tasks that you perform.

To create a backup file of the current WANJet appliance settings

1. In the navigation pane, expand **System** and click **Backup & Restore**.
The WANJet Configuration Backup & Restore screen opens.
2. Click the word **here**.
The browser opens a File Download window for you to save the backup file to your local computer. This default backup file is **Settings-<ServerName>.NTCL**.
3. Save the file to your local hard drive.
4. Rename the backup file to identify the specific WANJet appliance you are backing up, and the current date.

To restore a saved backup of WANJet appliance settings

1. In the navigation pane, expand **System** and click **Backup & Retstore**.
The WANJet Configuration Backup & Restore screen opens.
2. From the WANJet Configuration Restore section, click the **Browse** button to locate the backup file that you want to upload. The WANJet appliance's backup files end with the extension **.NTCL**.
3. Click the **Upload** button.
The Web UI refreshes to the home page, and the backup settings are in effect.

Restoring factory default values

You can restore the factory default values on the WANJet appliance. The values in all fields on all screens are reset to the default values. You should restore the defaults only in extreme cases, when you want to discard all configuration changes that you have made, and start configuration all over again. The WANJet appliance retains the licensing information, therefore, you do not have to revalidate the license. You do have to reconfigure the rest of the settings and policies.

Important

*Before restoring factory default values, you should back up the current configuration in case you decide that you want to use the policies and settings that you have already entered. Refer to **To create a backup file of the current WANJet appliance settings**, on page 5-11, for the procedure on how to save the current settings.*

To restore factory default values

1. In the navigation pane, expand **System** and click **Backup & Restore**.
The WANJet Configuration Backup & Restore screen opens.
2. Click the **Factory Defaults** button.
A message warns you that you are about to reset the WANJet appliance configuration to the factory default values, and asks if you want to continue.
3. Click the **OK** button.
The WANJet appliance stops and reboots automatically using the default values.
4. Perform the initial hardware configuration of the appliance including setting the WANJet appliance IP address, netmask, and WAN gateway, and if using out-of-band management, also set the Management port IP address, netmask, and gateway. Refer to the **Quick Start Card** for your WANJet appliance platform for details on how to configure the appliance.
5. Reconfigure the appliance by setting up the remote WANJet appliance and local subnets as described in Chapter 4, *Initial Configuration*.
6. Set up optimization policies (see Chapter 6, *Advanced Configuration*), service policies (see Chapter 7, *Configuring Service Policies*), and perform other configuration, as needed (for example, see *Configuring time settings*, on page 5-6, to correctly set the time zone).

Upgrading the WANJet appliance software

You can upgrade the software on the WANJet appliance from the Web UI in one of two ways:

- **Normal upgrade:** You need to have a copy of the new software.
- **Web upgrade:** The WANJet appliance must have Internet access to retrieve the new software.

When you upgrade the WANJet appliance, a new version of the software and the configuration settings from the current installation are installed onto the alternate image. After the upgrade, that image boots and becomes the active image.

Important

During the upgrade process, the WANJet appliance stops processing traffic for approximately 5-10 minutes and resets all connections. F5 Networks recommends that you upgrade during a time that is the least disruptive to network users.

To upgrade the WANJet appliance

1. Verify that a disk image of the new version of the WANJet software is accessible from the local computer on which you are viewing the Web UI (on CD-ROM, for example).
2. In the navigation pane, expand **System** and click **Upgrade & Boot Menu**.
The WANJet Boot Menu screen opens.
3. Click the **Upgrade** button.
A popup confirmation request opens.
4. Click the **OK** button to continue.
5. Click the **Browse** button and locate the upgrade file on your computer.
6. Upload the upgrade file to the WANJet appliance.
7. Click **Upgrade WANJet**.
A blue Upgrading screen opens and shows the status of the upgrade. The WANJet appliance restarts automatically when the upgrade process is complete. At that point, you need to log on to the Web UI again to continue working on the upgraded WANJet appliance.

Note

For you to perform a web upgrade (described in the following procedure), the WANJet appliance must be able to access the Internet and be running release 4.2.4 or later.

To perform a web upgrade of the WANJet appliance

1. In the navigation pane, expand **System** and click **Upgrade & Boot Menu**.
The WANJet Boot Menu screen opens.
2. Click the **Go** button.
The WANJet appliance checks for available upgrades. If one is available, it displays a message stating the upgrade number.
3. Click **Upgrade** if you want to upgrade the WANJet appliance to the available release.
A confirmation message states that if you continue, this action will stop the WANJet appliance functionality.
4. Click **OK** to continue.
The Remote Upgrade screen opens.
5. Under **Upgrade from the Web**, click **Go**.
A blue Upgrading screen opens and shows the status of the upgrade. The WANJet appliance restarts automatically when the upgrade process is complete. At that point, you need to log on to the Web UI again to continue working on the WANJet appliance.



6

Advanced Configuration

- Creating optimization policies
- Adjusting tuning settings
- Updating a configuration
- Managing virtual LANs
- Managing remote WANJet appliances
- Changing the interface speed
- Managing static routes
- Configuring Syslog and SNMP settings
- Configuring email alerts
- Configuring redundant peers
- Configuring one-arm topology
- Introducing high-availability features

Creating optimization policies

You use optimization policies to specify where and how you want the WANJet appliance to optimize traffic, including the local and remote subnets on the local and remote appliances. Optimization policies specify the TCP and UDP ports on which the WANJet appliance applies its Transparent Data Reduction (TDR) optimization algorithms.

To develop optimization policies, you set up subnets and ports:

- ◆ Specify the subnets for which you want to optimize traffic.
See the next section, *Managing subnets*, for details on how to determine the subnets for which the WANJet appliance will optimize traffic.
- ◆ Specify how you want traffic to be optimized on ports.
See *Managing port settings*, on page 6-4, for how to optimize ports.

Managing subnets

The procedures to add, remove, or modify subnets are different for local and remote WANJet appliances.

Adding, editing, or removing subnets on a local WANJet appliance

You can optimize all subnets connected to the local WANJet appliance or create a list of subnets for which you want to optimize traffic. You can add, update, or remove subnets from the list.

When you first open the Optimization Policies screen, the subnet in which the WANJet appliance resides is the only subnet on the list of local subnets to optimize, and the **Include WANJet Subnet** setting is enabled. The **Optimize all Subnets** check box is cleared by default.

To optimize all traffic on all subnets including the subnet in which the local WANJet resides, check the **Optimize all Subnets** box. When you check this option, the tables of local and remote subnets are removed from the screen.

To optimize traffic only on specific subnets, clear the **Optimize all Subnets** check box, and add the subnets for which you want to optimize traffic.

To optimize all subnets connecting to the local WANJet appliance

1. In the navigation pane, expand **Optimization** and click **Optimization Policy**.
The WANJet Optimization Policy screen opens.
2. Check the **Optimize all Subnets** box.
Other local and remote WANJet options are removed from the screen.
3. Click the **Save** button at the bottom of the WANJet Optimization Policy screen.

To add a new subnet to the local WANJet appliance

1. In the navigation pane, expand **Optimization** and click **Optimization Policy**.
The WANJet Optimization Policy screen opens.
2. Clear the **Optimize all Subnets** check box, if it is checked.
3. Click the **Add** button below the Local Subnets list.
The Add Local Subnet popup screen opens.
4. In the **Local Subnet** box, type the IP address of the subnet. For example:
10.8.0.0
5. In the **Netmask** box, type the netmask of the subnet. For example:
255.255.0.0
6. In the **Alias** box, type a name for the subnet. For example:
Subnet B
7. Set the **Operational Status** to one of the options:
 - **Enabled**
To have the WANJet appliance optimize the traffic for the subnet.
 - **Disabled**
To prevent the WANJet appliance from optimizing the traffic for the subnet.
8. Click the **OK** button.
The window closes.
9. Click the **Save** button at the bottom of the WANJet Optimization Policy screen.
Traffic on subnets that you added will be optimized for all new connections.

To update or remove a subnet on the local WANJet appliance

1. In the navigation pane, expand **Optimization** and click **Optimization Policy**.
The WANJet Optimization Policy screen opens.
2. Clear the **Optimize all Subnets** check box.
3. In the list of local subnets, click the link of the subnet that you want to remove or edit.
The Edit Local Subnet screen opens.
4. Edit the settings, or click **Remove** to remove this subnet.
5. Click the **OK** button.
The screen closes.

6. Click the **Save** button at the bottom of the WANJet Optimization Policy screen.
Changes to subnets take effect for all new connections.

◆ **Note**

You cannot update or remove the IP address of the subnet in which the WANJet appliance is located.

Adding, editing, or removing subnets on a remote WANJet appliance

◆ **Important**

Always add the gateway of any remote WANJet appliance as one of its subnets, and confirm that the status of this subnet is disabled.

To add a new subnet to a remote WANJet appliance

1. In the navigation pane, expand **Optimization** and click **Optimization Policy**.
The Optimization Policy screen opens.
2. From the **Remote WANJet** list, select the remote WANJet appliance to which you want to add subnets.
3. Click the **Add** button, located below the Remote Subnets table.
The Add Remote Subnet screen opens.
4. In the **Supported Subnet** box, type the IP address of the machine/subnet that you want to make visible to the remote WANJet appliance.
5. In the **Netmask** box, type the netmask of the remote subnet.
6. In the **Machine(s) Alias** box, type a name for the machine/subnet.
7. If you do not want the WANJet appliance to process the traffic for this subnet at this time, click **Disabled**. Otherwise, leave it at the default of **Enabled**.
8. Click the **OK** button.
The window closes.
9. Click the **Save** button at the bottom of the WANJet Optimization Policy screen.
The WANJet appliance optimizes traffic for all new connections on the subnets that you added.

To update or remove a subnet on a remote WANJet appliance

1. In the navigation pane, expand **Optimization** and click **Optimization Policy**.
The WANJet Optimization Policy screen opens.
2. From the **Remote WANJet** list, select the IP address of the remote WANJet appliance that you want to modify.
3. In the list of remote subnets, click the subnet that you want to remove or edit.
The Edit Remote Subnet screen opens.
4. Edit the settings, or click **Remove** to remove this subnet.
5. Click the **OK** button.
The screen closes.
6. Click the **Save** button at the bottom of the WANJet Optimization Policy screen.
Changes to subnets take effect for new connections.

Managing port settings

You can adjust the processing mode and the Type of Service (ToS) priority that are assigned to packets for each port on a remote WANJet appliance. You can assign these values separately for TCP and UDP packets so that you can, for example, optimize TCP traffic on a port while allowing UDP traffic to pass through untouched.

It is typical to have optimization enabled on commonly used ports such as those used for Active FTP, SMTP, HTTP, POP3, IMAP, and HTTPS. You can also consider enabling TDR-1 compression on these ports, except **443** (HTTPS). You can edit the settings for ports that have been added by clicking the corresponding link.

◆ Note

It is difficult to optimize Passive FTP sessions because the server port that Passive FTP uses varies from session to session. However, if you need to optimize Passive FTP, enable optimization for all TCP ports and disable optimization for ports that do not require it (typically ports used by real-time applications such, as VoIP telephony).

Configuring ports and services

You can customize the optimization policy for ports, or services that use specific ports. To do this, you need to add the ports (or services) and indicate how you want the system to handle connections through that port. For example, for any port or service, you can specify whether data is optimized or passed through, set its ToS priority, and select the type of optimization.

To add ports or services

1. In the navigation pane, expand **Optimization** and click **Optimization Policy**.
The WANJet Optimization Policy screen opens.
2. Click the **Add** button located beneath the Protocol Optimization Policies table.
The Add Port/Service Name popup screen opens.
3. From the **Service Name** list, select the service or application for which you want to customize the optimization policy. The default port used by the service appears in the **From Port** box.

If you do not know the name of the service, in the **From Port** box, type the port number. To specify a range of ports, type the first port in the range in the **From Port** box, and the last port in range in the **To** box.

Note: Refer to <http://www.iana.org/assignments/port-numbers> for a list of commonly assigned TCP/UDP port numbers and the services and applications that use them. Keep in mind that these may differ on your system.

4. From the **Processing Mode** list, select one of the options:
 - **Passthrough**
Leave traffic over this port in its raw state.
 - **Optimized**
Apply WANJet appliance optimization to traffic over this port.
5. From the **TOS Priority** list, select a priority for the port(s):

7 - Network Control
6 - Internet Control
5 - Critical
4 - Flash Overdrive
3 - Flash
2 - Immediate
1 - Priority
0 - Routine

Note: Refer to <http://www.ietf.org/rfc/rfc0791.txt> for more information about ToS priority levels.

6. Select a WANJet appliance optimization option by checking one of the optimization option check boxes.

The following options are available only if you have selected **Optimized** as the processing mode.

Check box	Optimization description
TDR-1	Check this box to compress network traffic on the specified port. This is not necessary if the traffic would not benefit from compression, for instance if it consists largely of JPEG or ZIP files.
TDR-2	Check this box to apply the WANJet appliance's TDR-2 intelligent caching algorithm.
Encryption	Check this box if network traffic on the specified port is encrypted to use SSL.
Connection Intercept	Check this box to reset any connection over the specified port that was opened before the new settings were applied.

7. Click the **OK** button.
The window closes and the WANJet Optimization Policy screen displays with a new row in the Protocol Optimization Policies table with the details that you entered. You can click the port number (in the Service Name column) to edit these settings.
8. Click the **Save** button at the bottom of the WANJet Optimization Policy screen to apply the new port settings.
The new port settings take effect immediately for all new connections.

Configuring all other ports

In addition to defining optimization policies for specific ports, you can change the default policies that have been set up for all TCP and UDP ports. (Any policies defined for individual ports will override these default policies.) Instead of listing specific ports or services in the list of optimization policies, the table shows **All ports** or **All other ports**.

To set the default processing mode for all TCP/UDP ports

1. In the navigation pane, expand **Optimization** and click **Optimization Policy**.
The WANJet Optimization Policy screen opens.
2. From the Remote WANJet box, choose the IP address of the remote WANJet appliance to which you want to connect.
3. In the Protocol Optimization Policies table, in the Service Name column, for TCP or UDP protocol, click **All Ports**. (This reads **All other ports** if optimization policies are defined for specific ports.)
The Edit Port/Service Name popup screen opens.
4. From the **Processing Mode** list, select one of the options:
 - **Passthrough**
Leave traffic over this port in its raw state.
 - **Optimized**
Apply WANJet appliance optimization to traffic over this port.
5. From the **TOS Priority** list, select a priority for the port(s):
 - 7 - Network Control
 - 6 - Internet Control
 - 5 - Critical
 - 4 - Flash Overdrive
 - 3 - Flash
 - 2 - Immediate
 - 1 - Priority
 - 0 - Routine

Note: Refer to <http://www.ietf.org/rfc/rfc0791.txt> for more information about ToS priority levels.

6. If you selected **Optimized** as the processing mode in step 5, check the optimization options you want to enable.

Check box	Optimization description
TDR-1	Check this box to compress network traffic on the specified port. This is not necessary if the traffic would not benefit from compression, for instance if it consists largely of JPEG or ZIP files.
TDR-2	Check this box to apply the WANJet appliance's TDR-2 intelligent caching algorithm.
Encryption	Check this box if network traffic on the specified port is encrypted to use SSL.
Connection Intercept	Check this box to reset any connection over the specified port that was opened before the new settings were applied.

7. Click **OK**.
The Optimization Policy screen displays with a new row in the third table that contains the details that you entered. You can click the port number (in the Service Name column) to edit these settings.
8. Click the **Save** button to apply the new port settings.
The new port settings take effect immediately for all new connections.

Enabling Connection Intercept

When you start the WANJet appliance, some connections may have already been established. Connection Intercept lets you reset connections that were initiated before you started the WANJet appliance. You can use Connection Intercept to reset connections for specific ports or services, without having to reboot the relevant servers or restart those services.

Using the **Connection Intercept** option is particularly effective when performing any of the following tasks:

- Installing the WANJet appliance on your network
- Upgrading the WANJet appliance
- Changing the WANJet appliance's mode from inactive to active
- Restarting the WANJet appliance

The ports on which you implement Connection Intercept require the following settings:

- **Optimized** as the processing mode
- **Connection Intercept** option enabled

◆ **Note**

You can use the following process to optimize any port. The best usage for Connection Intercept is when you want to reset connections on a range of ports, without having to either reboot the relevant servers or restart a whole range of services. The WANJet appliance allows you to reset connections automatically, without having to restart the server or manually reset the connections.

To enable connection intercept

1. In the navigation pane, expand **Optimization** and click **Optimization Policy**.
The WANJet Optimization Policy screen opens.
2. In the Protocol Optimization Policies section, in the Service Name column, click the service or port for which you want to enable Connection Intercept.
The Edit Port/Service Name popup screen opens.
***Note:** If the service is not listed, you need to add it. Refer to **Configuring ports and services**, on page 6-4.*
3. For the **Processing Mode**, select **Optimized**.
4. Check the **Connection Intercept** box.
5. Click **OK**.
The Edit Port/Service Name screen closes, and you see the WANJet Optimization Policy screen.
6. Verify that the WANJet appliance operational mode is set to **Active**.
7. Click the **Save** button to apply the changes.
The WANJet appliance enables Connection Intercept on all configured ports, and resets existing connections on these ports so that data transfers are optimized.

For additional details about using the **Connection Intercept** option, refer to *Configuring ports and services*, on page 6-4.

Example: Connection Intercept implementation

One of the uses of Connection Intercept is for client systems that use the CIFS (Common Internet File System) protocol to request file services from server systems over a network. Here we provide an example of how to use Connection Intercept to automatically reset CIFS connections.

In this example, the customer is concerned that they may have existing CIFS connections, already in progress, that are not being optimized after starting the WANJet appliance. It shows how to enable the Connection Intercept option on the CIFS ports (typically ports **139** and **445**). This causes the WANJet appliance to automatically reset connections that are not being optimized, without having to restart each of the connections manually.

To automatically reset CIFS connections

1. In the navigation pane, expand **Optimization** and click **Optimization Policy**.
The WANJet Optimization Policy screen opens.
2. In the Protocol Optimization Policies section, in the Service Name column, click **139 (Netbios-ssn)**.
The Edit Port/Service Name popup screen opens.
*Note: If the Netbios-ssn service is not listed, you need to add it. Refer to **Configuring ports and services**, on page 6-4.*
3. For the **Processing Mode**, select **Optimized**.
4. Check the **Connection Intercept** box.
5. Click **OK**.
The Edit Port/Service Name screen closes, and you see the WANJet Optimization Policy screen.
6. In the Protocol Optimization Policies section, in the Service Name column, click **445 (Microsoft-ds)**.
The Edit Port/Service Name popup screen opens.
7. For the **Processing Mode**, select **Optimized**.
8. Check the **Connection Intercept** box.
9. Click **OK**.
The Edit Port/Service Name screen closes, and you see the WANJet Optimization Policy screen again.
10. In the navigation pane, expand **Optimization** and click **Operational Mode**.
The Operational Mode screen opens.
11. Verify that **Mode** is set to **Active**.
12. Click the **Save** button to apply the changes.
This implements Connection Intercept on ports **139** and **445**. The next time you restart the WANJet appliance, it resets connections on these ports, and then optimizes the traffic.

Setting operational modes

You can set operational modes on the WANJet appliance. From the Operational Mode screen, you can:

- Specify the operating mode of the WANJet appliance (whether it is active or inactive).
- Determine how you want to handle traffic in case of failure (on WANJet 400 and 500 only).
- Specify how to deploy the WANJet appliance in your network topology (inline or one-arm).

To configure the operational mode settings

1. In the navigation pane, expand **Optimization** and click **Operational Mode**.
The Operational Mode screen opens.
2. For the **Mode** setting, select one of the following options:
 - **Active**
Enables optimization.
 - **Inactive**
Optimization does not occur and the WANJet appliance is completely transparent to network traffic.
3. For the **Failure Mode** setting (WANJet 400/500 only), select one of the following options:
 - **Fail to Wire** (default)
If the WANJet appliance fails for any reason, network traffic continues to flow and bypasses the WANJet appliance. On the WANJet 500, if the power is off on the WANJet appliance, this option is always in effect even if you select **Fail Close**.
 - **Fail Close**
If the WANJet appliance fails for any reason, the appliance breaks the link and stops traffic from passing through.

***Note:** If you select **Fail Close** on a WANJet 400, you must also make a hardware adjustment on the appliance. Refer to **To enable Fail Close on the WANJet 400 hardware**, following, for instructions on how to open the unit and change the setting on the NIC.*

4. For the **Topology** setting, specify the way the WANJet appliance is connected to the network by clicking one of the options:
 - **In-Line**
This is the most common network topology. ***Inline*** means that the WANJet appliance is located between the LAN (or the LAN switch) and the WAN gateway (or the LAN router).
 - **One-Arm**
Select this option if your WANJet appliance is located on a separate, independent link. Refer to *Configuring one-arm topology*, on page 6-28, for additional instructions.
5. Click the **Save** button.

To enable Fail Close on the WANJet 400 hardware

1. Set the **Failure Mode** setting to **Fail Close** as described in the previous procedure, *To configure the operational mode settings*. (Do not forget to click **Save** to save the changed setting.)
2. Shut down the WANJet 400 appliance. See *Shutting down and restarting the WANJet appliance*, on page 5-8.
3. Turn the WANJet appliance upside down. On the bottom of the unit, unscrew the four screws on the left and right edges of the unit.
4. Slide the cover off the top of the WANJet 400 appliance.
5. Facing the front of the WANJet 400 appliance, locate the PXG2BP NIC card on the right near the front of the unit.
6. Tip the WANJet 400 appliance over onto the left side so you can see the buttons on the NIC card better.
7. On the upper right of the NIC card, locate the two switches (labeled **BYPASS MODE**).

The ENB switch is on the left (towards the front of the card), and is turned off by default. The DIS switch is on the right (towards the back of the card), and is turned on by default.
8. Turn the appliance on. If a warning beep sounds, press the red reset button on the back of the unit next to the power supplies.
9. On the NIC card, press the ENB switch on the left (the one towards the front of the appliance). You hear an audible click.
10. Turn the appliance off and replace the cover.

To re-enable Fail to Wire on the WANJet 400 hardware

1. Set the Failure Mode setting to **Fail to Wire** as described in the procedure, *To configure the operational mode settings*, on page 6-11. (Do not forget to click **Save** to save the changed setting.)
2. Shut down the WANJet 400 appliance.
3. Turn the WANJet appliance upside down. On the bottom of the unit, unscrew the four screws that are on the left and right edges of the unit.
4. Slide the cover off the top of the WANJet 400 appliance.
5. Facing the front of the WANJet 400 appliance, locate the PXG2BP NIC card on the right near the front of the unit.
6. Tip the WANJet 400 appliance over onto the left side so you can see the buttons on the NIC card better.
7. On the upper right of the NIC card, locate the two switches (labeled **BYPASS MODE**).

The ENB switch is on the left (towards the front of the card), and was previously turned on. The DIS switch is on the right (towards the back of the card), and is off.

8. Turn the appliance on. If a warning beep sounds, press the red reset button on the back of the unit next to the power supplies.
9. On the NIC card, press the DIS switch on the right (the one towards the center of the appliance). You hear an audible click.
10. Turn the appliance off and replace the cover.

Adjusting tuning settings

From the Tuning screen, you can guarantee maximum throughput by specifying the link bandwidth and the Round Trip Time (RTT) for the WAN link. The maximum bandwidth value is a global setting that relates to the WANJet appliance license that your company purchased. You should only modify the tuning settings when initially setting up the WANJet appliance (after licensing), or if the bandwidth of your WAN link changes. Once they are set, you rarely need to change the tuning settings.

To modify tuning settings

1. In the navigation pane, expand **Optimization** and click **Tuning**. The WANJet Tuning screen opens.
2. In the **Bandwidth** box, type a value for your WAN link bandwidth. You can set it to the bandwidth for which the appliance is licensed or lower. The default bandwidth varies depending on the license purchased and the platform. You can adjust the value lower than the

default (but not higher), and use the list to change the units to kilobits per second for lower-bandwidth links.

F5 Networks does not recommend changing this value.

3. In the **RTT** box, type the value for the average round trip time for the WAN link. You determine the RTT by using the ping utility to send a request to a device on the other side of the WAN link and reviewing the command output. The default RTT is **300** milliseconds.
4. Check the **Congestion Control** box to have the WANJet appliance handle traffic if congestion occurs in the case of packet loss. The **Congestion Control** box is checked by default.
5. Review the value in the **Queue Size** box. It contains the maximum number of outgoing packets to keep in the queue before dropping them (in case of network problems). The WANJet automatically calculates the **Queue Size** based on the values specified for **Bandwidth** and **RTT**. F5 Networks does not recommend changing this value.
6. Click the **Save** button.
The WANJet Tuning screen refreshes, and the WANJet appliance saves the changes.

Updating a configuration

When you initially configure the local WANJet appliance (as described in Chapter 4, *Initial Configuration*) you specify the network settings for the WANJet appliance, such as IP address, ports, subnets, redundant peers, and connected remote WANJet appliances.

From the Local WANJet appliance screen, you can edit the network information for the local WANJet appliance, such as defining redundant peers, adding subnets, and defining VLANs to the local WANJet appliance. The initial values displayed on the Local WANJet appliance screen are the ones that you specified during initial configuration.

Important

*You must replicate any changes that you make to the WANJet appliance's IP address, port, or subnet address on each remote WANJet appliance to which the local WANJet appliance is connected. See **Replicating configuration changes on remote WANJet appliances**, on page 6-16.*

Modifying a local WANJet appliance network configuration

If you need to modify the local WANJet appliance configuration, perform the following steps.

To modify the local WANJet appliance network configuration

1. In the navigation pane, expand **Configuration** and click **Local WANJet**.
The Local WANJet appliance screen opens.
2. Modify the values as required. The values are defined as follows:
 - **WANJet Alias**
The name that is used for the local WANJet appliance. This name is displayed at the upper-left corner of the home when you log onto the Web UI.
 - **WANJet IP**
The IP address that is assigned to the local WANJet appliance on your network. If you change this value, you must change it on each remote WANJet appliance that accesses the local appliance. See *Replicating configuration changes on remote WANJet appliances*, on page 6-16.
 - **WANJet Netmask**
Subnet mask assigned to the WANJet appliance on your network.
 - **WAN Gateway**
The gateway the WANJet appliance uses to reach the WAN.
 - **LAN Router**
The gateway that the WANJet appliance uses to reach the LAN.
 - **WANJet Port**
The main port number that the local WANJet appliance uses to communicate with remote WANJet appliance. The default port is **3701**. If you change this value, you must change it on each remote WANJet appliance that accesses the local WANJet appliance. See *Replicating configuration changes on remote WANJet appliances*, on page 6-16.
 - **Redundant Peer IP**
IP address of the redundant WANJet appliance. If you check the **Redundant Peer IP** check box, the **IP address** box displays.
3. Click the **Save** button.

Configuring delayed acceptance

The WANJet appliance generally accepts incoming connections from the LAN, then attempts to connect with the server on the remote LAN. If the server is unreachable, the WANJet appliance closes the original client-side LAN connection. A delayed connection acceptance feature, enabled by default, postpones acceptance of LAN requests coming from ports **445** and **139** until the server connection is verified.

You can configure the ports that will delay accepting requests, or disable the setting, as needed. This setting is particularly useful for ports that use CIFS (that is, ports **445** and **139**).

To configure delayed connection acceptance settings

1. In the navigation pane, expand **Configuration** and click **Local WANJet**.
The Local WANJet screen opens.
2. Click **Settings for Delayed Acceptance**.
3. Check the **Enable** box to enable delayed connection acceptance (if it is not selected already).
4. In the **Ports** box, type the numbers of any ports for which you want to delay the acceptance of a connection until verifying that the server is reachable. Separate multiple ports with colons (for example, **139:445**).
5. Click **Save** to make the changes.

Replicating configuration changes on remote WANJet appliances

If you make any changes to the IP address, port setting, or subnet address on a local WANJet appliance, you must replicate the changes everywhere they appear, including to connected remote WANJet appliances.

For example, if you have four connected WANJet appliances named B1, B2, B3, and B4, and you log on to the Web UI for B1, the Web UI shows B1 as the local WANJet appliance and B2, B3, and B4 as its remote WANJet appliances. Therefore, if you change the IP address for B1, you must also change the IP address for B1 for the remote WANJet appliances (B2, B3, and B4) so that it matches.

To update the remote WANJet appliance settings from the local WANJet appliance

1. Log on to the Web UI of the WANJet appliance.
2. In the navigation pane, expand **Configuration** and click **Remote WANJets**.
The Local WANJet screen opens.

3. Click the IP address of the remote WANJet appliance that you want to edit.
The Manage Remote WANJet popup screen opens.
4. Edit the settings as required.
5. Click the **OK** button.
The Manage Remote WANJet screen closes.
6. Click the **Save** button at the bottom of the Remote WANJets screen.
7. Repeat steps 3 through 6 for each remote WANJet appliance that connects to the local WANJet appliance for which you changed settings.

Once complete, the local WANJet appliance should be able to communicate with all connected remote WANJet appliances.

◆ **Note**

Alternatively, you can change the settings for the connected WANJet appliances by logging on to each WANJet appliance's Web UI.

Managing virtual LANs

A **virtual LAN** (VLAN) is a segment of a computer network that has logically defined (rather than physically defined) boundaries. VLANs provide a way to structure a large network for increased security, separating systems with sensitive data, special projects, or separate departments. You need to configure the WANJet appliance so it can optimize traffic from those VLANs.

Unless you want to have separate WANJet appliances on every VLAN, you must use the Web UI to add the WANJet appliance to any VLANs that are linked to your network and for which you want to optimize traffic. When added, the WANJet appliance becomes part of the VLAN and can optimize traffic on the VLAN. You need to add the WANJet appliance to the VLAN because VLANs are often implemented by adding tags to Ethernet frames. These tags must be preserved during optimization. See *To add a VLAN to a WANJet appliance* in the next section for how to add WANJet appliances to VLANs.

Traffic emanating from VLANs is tagged with a **VLAN tag** (also called a VLAN ID). By default, the WANJet appliance handles only untagged traffic. If you want to optimize traffic on VLANs in your networking environment, you need to configure the VLANs on the WANJet appliance. By configuring the VLANs and their VLAN tags, you include the WANJet appliance in those VLANs. The WANJet appliance recognizes the tags and can optimize traffic from those VLANs.

The WANJet appliance IP address is automatically listed in the WANJet VLAN settings table with a VLAN tag of **0** (meaning no tag is assigned).

You can:

- ◆ Add VLANs. The WANJet appliance optimizes traffic tagged with those VLAN tags; it also optimizes untagged traffic that meets the optimization policy guidelines.
- ◆ Configure the local WANJet appliance IP address with a VLAN tag. The WANJet appliance optimizes traffic tagged with that VLAN tag. If you add additional VLANs, it optimizes traffic from those VLANs as well, but it does not optimize untagged traffic.

You can also modify VLAN information or delete VLANs no longer configured. The following procedures describe how to perform these VLAN management tasks.

To add a VLAN to a WANJet appliance

1. In the navigation pane, expand **Configuration** and click **Local WANJet**.
The Local WANJet appliance screen opens.
2. Click the **VLAN Settings** link beneath the table.
The VLAN Setting screen displays with all of the currently defined VLANs.
3. Click the **Add** button.
The Add VLAN screen opens in a separate window.
4. In the **WANJet Virtual IP** box, type the virtual IP address assigned to the local WANJet appliance on this VLAN. That is, the IP address that other machines on the VLAN use to communicate with the local WANJet appliance.
5. In the **VLAN Netmask** box, type the subnet mask for the VLAN.
6. In the **VLAN Gateway** box, type the virtual IP address of the gateway machine for the VLAN.
7. In the **VLAN Tag** box, type the VLAN ID that the WANJet appliance uses to preserve tagged Ethernet frames that pass to and from the VLAN. Valid values are **2** through **4094** (**0** means assign no VLAN tag, and **1** is reserved).
8. Click the **OK** button.
The Add VLAN screen closes.
9. Click the **Save** button.
The VLAN is automatically added as a local subnet as part of the optimization policy on the local WANJet appliance. It is also added as a remote subnet on any remote WANJet appliances that are linked to the local appliance.

To configure the local WANJet appliance IP address with a VLAN tag

1. In the navigation pane, expand **Configuration** and click **Local WANJet**.
The Local WANJet appliance screen opens.
2. Click the **VLAN Settings** link beneath the table.
The VLAN Setting screen opens.
3. Click the IP address of the local WANJet appliance (it is automatically listed first in the table).
The Edit VLAN Tag Associated with WANJet IP popup screen opens in a separate window. You can edit only the **VLAN Tag**.
4. In the **VLAN Tag** box, type the ID of the VLAN to which the local WANJet appliance belongs.
5. Click the **OK** button.
The popup screen closes.
6. Click the **Save** button.
The WANJet appliance optimizes only tagged traffic from this point on. To be optimized, the traffic needs to be tagged with the VLAN tag of the local WANJet appliance or with the tags of other VLANs that you added on the WANJet VLAN Settings screen.

To edit or remove a VLAN on the WANJet appliance

1. In the navigation pane, expand **Configuration** and click **Local WANJet**.
The Local WANJet appliance screen displays
2. Click the **VLAN Settings** link beneath the table.
The VLAN Setting screen displays with all of the currently defined VLANs.
3. Click the IP address of the VLAN configuration you want to edit or remove.
The Edit VLAN popup screen opens in a separate window.
4. Edit the VLAN information, or click the **Remove** button to remove it.

***Note:** On the local WANJet appliance, you can edit only the **VLAN Tag**, not the **WANJet IP** address, **WANJet Netmask**, or the **WANJet Gateway**. You cannot remove the local WANJet appliance.*
5. Click the **OK** button.
The Edit VLAN screen closes.
6. Click the **Save** button.

◆ Important

If you remove a VLAN from a local WANJet appliance, you must also remove it from the list of subnets supported by that WANJet appliance.

Managing remote WANJet appliances

To optimize the data that is sent over a network link, you need a pair of WANJet appliances, each running the WANJet appliance software. A remote WANJet appliance reverses the optimization process for data that is sent from the local WANJet appliance. For this configuration to work, the local WANJet appliance must be aware of the remote WANJet appliance. If you do not specify a remote WANJet appliance to receive the processed data, network traffic passes through the local WANJet appliance without being optimized.

To add a remote WANJet appliance

1. In the navigation pane, expand **Configuration** and click **Remote WANJet**.
The Remote WANJet screen opens.
2. Click the **Add** button.
The Manage Remote WANJet appliance screen opens.
3. From the **WANJet Type** list, select **Single**.
Or, if you have two connected WANJet appliance peers on the same remote LAN, select **Redundant**. (See *Configuring redundant peers*, on page 6-26 for an explanation about these node types.)
4. In the **WANJet IP** box, type the IP address for the remote WANJet appliance.
5. If you selected **Redundant** in Step 3, for **Redundant Peer**, type the IP address for the peer WANJet appliance in the **Node 2** box.
Otherwise, the field is not available, and you can skip to Step 6.
Note: The Node 2 box appears only if you select Redundant from the WANJet appliance type menu.
6. In the **WANJet Alias** box, type a name for the remote WANJet appliance. The name must have fewer than 14 characters.
7. In the **WANJet Port** box, type the main port number on which the remote WANJet appliance listens for data from the local WANJet appliance. The default port number is **3701**.
Note: If you change the WANJet appliance port number, you must change it for all connected WANJet appliances.
8. In the **Shared Key** box, type the shared key that authenticates between the local and remote WANJet appliances. You can set a unique shared key for every pair of WANJet appliances.
9. Leave the **Local GRE IP** and **Remote GRE IP** addresses blank unless you are using GRE tunneling between two routers with a one-arm WANJet appliance configuration (see *Using transparent proxy with generic routing encapsulation tunneling*, on page 3-7 for details on this configuration).

If you are using GRE tunneling, for **Local GRE IP**, type the IP address of the local end of the GRE tunnel, and for **Remote GRE IP**, type the IP address of the remote end of the GRE tunnel.

10. If you specified an IP address in the **LAN Router** field on the Local WANJet screen, you can select an MTU (Maximum Transmission Unit) type. The MTU is the maximum packet size in bytes that can be transmitted across a link. For **MTU**, select one of the following MTU types:
 - **Direct**
The default value for this type is 1500 bytes, and is the most common MTU type used for the IP protocol.
 - **VPN**
The default MTU for this option is 1400 bytes.
 - **Other**
You can specify the MTU value required by your network.
11. Click the **OK** button.
The Manage Remote WANJet screen closes.
12. Click the **Save** button.

You now need to add the gateway of the remote WANJet appliance as a disabled subnet. For more information, see *Managing subnets*, on page 6-1.

To edit or remove a remote WANJet appliance

1. In the navigation pane, expand **Configuration** and click **Remote WANJet**.
The Remote WANJet screen opens.
2. Click the IP address for the WANJet appliance that you want to edit or remove.
The Manage Remote WANJet appliance screen opens.
3. Edit the information or click the **Remove** button to remove the remote WANJet appliance.
Note: If you edit a port number, you must change that port number on all connected WANJet appliances. If you remove a WANJet appliance, you remove all associated subnets and ports.
4. Click the **OK** button.
The Manage Remote WANJet appliance screen closes.
5. Click the **Save** button.

Important

If you remove a remote WANJet appliance, the local WANJet appliance no longer recognizes it, and any data sent to the removed remote WANJet appliance's network passes through without being optimized.

Changing the interface speed

The WANJet appliance supports both half-duplex and full-duplex data transmission for the LAN (**eth0**), WAN (**eth1**) and Management port (**eth3**) interfaces. The available speeds (for example, 100 BaseTX or 10 BaseT) vary depending on the WANJet platform you are using.

By default, the interface speeds of all WANJet platforms are set to **Auto Negotiate**, causing them to negotiate the interface speeds automatically; however, you can use the following procedure to manually specify the speed of the network interfaces that the WANJet appliance uses to communicate with the LAN, the WAN, and the management network.

To change network interface settings

1. In the navigation pane, expand **Configuration** and click **Interfaces**. The NIC Configuration screen opens.
2. From the **Media Type** list for **eth0**, select the interface speed and duplex setting that corresponds with the link between the LAN switch and the WANJet appliance (the default is **Auto Negotiate**). The speed and duplex values for the LAN and the WAN interfaces must match.
3. From the **Media Type** list for **eth1**, select the interface speed and duplex setting that corresponds with the link between the WAN router and the WANJet appliance (the default is **Auto Negotiate**). The speed and duplex values for the LAN and the WAN interfaces must match.
4. From the **Media Type** list for **eth3**, select the interface speed and duplex setting that corresponds with the link between the management network and the WANJet appliance (the default is **Auto Negotiate**). The speed and duplex values do not need to match those of the LAN and WAN interfaces.
5. Click the **Save** button.

Managing static routes

The Static Routes table contains information about the gateway (router) that you specify to route the data for a specific network. Data packets sent to the defined gateway use the relevant static route to identify their destination.

When you specify a LAN router for your local WANJet appliance, all subnets configured for the local WANJet appliance use it to identify the destinations of packets.

◆ **Note**

*To specify a gateway for each subnet, remove the IP address from the **LAN Router** box on the Local WANJet appliance page. See **Updating a configuration**, on page 6-14, for specific instructions.*

To add a static route

1. In the navigation pane, expand **Configuration** and click **Routes**. The WANJet Routes screen opens.
2. In the **Network** box, type the subnet's IP address for which you want to route data to a specific gateway.
3. In the **Netmask** box, type the netmask for the network.
4. In the **Next Hop** box, type the IP address for the gateway to which the data should be routed. Data packets use this gateway to send them to their destination.
5. In the **MTU** box, type the maximum packet size of datagrams that you want transferred through this route.
6. Click the **Save** button.

To edit or remove an existing static route

1. In the navigation pane, expand **System Settings** and click **Routes**. The WANJet Routes screen opens.
2. Modify the **Network** and/or **Netmask** settings as required, or clear the **Network** settings for the route that you want to remove.
3. Click the **Save** button.

Configuring Syslog and SNMP settings

You can configure the WANJet appliance to retrieve Syslog, SNMP, and RMON2 reports from specific servers, and specify whether RMON2 data is gathered before (raw data) or after the WANJet appliance processes it (WANJet data). You can explicitly specify which IP address to use for SNMP: either the Management IP (the default) or the WANJet data IP (also called the *bridge IP*). You can also define the community string for viewing SNMP reports.

To configure Syslog and SNMP settings

1. In the navigation pane, expand **Configuration** and click **Monitoring**.
The **WANJet Syslog and SNMP** screen opens.
2. Check the **Syslog Server IP** check box and type the IP address of the server that receives Syslog data from the WANJet appliance.
3. Specify which log to store:
 - **Application**
Stores only the application error log on the server that you specified.
 - **All**
Stores all error logs on the server that you specified.
4. Check the **SNMP Server IP** check box and type the IP address of the SNMP server to which the WANJet appliance sends error messages. (For more information about viewing SNMP reports, see *SNMP reports*, on page 8-28.)
5. To view RMON2 data, check the **Enable RMON2 Logs** check box and select an option:
 - **Raw Data**
To view RMON2 logs before the WANJet appliance processes traffic.
 - **WANJet Data**
To view RMON2 logs after the WANJet appliance processes traffic.

*Note: For information about RMON2 data, refer to **RMON2 Reports**, on page 8-29. For details on the Raw Data and WANJet Data settings, see **RMON2 configuration settings**, on page 8-30.*
6. For **SNMP IP**, select which IP address you want SNMP to use as the source address in response to SNMPv1 GET requests, and for sending SNMP traps. The choices are **Management IP** (set by default) or **WANJet IP**.

7. In the **Community String** box, type the shared community string that enables the SNMP server to access SNMP reports on the WANJet appliance. The community string is a text string (up to 32 characters including a-z, A-Z, 0-9, hyphen, and underscore) set on the SNMP server to authenticate access.
8. Click the **Save** button.
The Syslog and SNMP page refreshes, and the changes are committed to the WANJet appliance.

Configuring email alerts

You can configure the WANJet appliance to send an email that includes a system snapshot (containing current system information) to a specific email address in the event of system failure.

◆ Note

*For information about how to download system snapshots directly, refer to **Diagnostic Log**, on page 8-27.*

To configure email alerts

1. In the navigation pane, expand **Configuration** and click **Email Alert**.
The WANJet Email Alert screen opens.
2. In the **Email address** box, type the email address to which you want the system snapshot sent. If you want the email alert to go directly to F5 Networks, type **WANJetSupport@f5.com**.
3. In the **From Email address** box, type the email address from which you want the email to appear to be sent.

This does not need to be a valid email address, but it should look like a valid address to pass through spam filters. F5 Networks recommends that you use the alias of the WANJet appliance from which the snapshot was taken as the first part of the address (before the @ symbol), and your company's domain name as the second part of the address. For example, **WJ_NewYork@company.com**.
4. In the **SMTP Server IP** box, type the IP address (not the domain name) of an SMTP mail server that is accessible from the WANJet appliance.

5. In the **SMTP Server Port** box, type the port number for the mail server to which the SMTP request for the email alert will be sent.
Note: Typically, the port for SMTP is 25; however, the default port that the WANJet appliance uses for email alerts is 443 (which is normally used by SSL traffic). The WANJet appliance uses port 443, because it is more likely to be allowed through by a firewall. Verify that the mail server specified in the SMTP Server IP box is set up to forward traffic on port 443 to port 25.
6. To automatically email system snapshots, check the **Enabled** box.
Email alerts are disabled by default, but F5 Networks recommends that you enable them after you configure the settings on the WANJet Email Alert screen.
7. Click the **Test Me** button to confirm that the WANJet appliance can access the mail server and send the email. You can use the test feature to send a simple test message, create a new system snapshot to send, or send all past system snapshots. F5 Networks recommends that you send a test message, because the WANJet appliance does not attempt to resend failed emails.
8. After you have confirmed that the email alert that you configured works, click the **Save** button.

Configuring redundant peers

The WANJet appliance supports high availability through redundant pairs, or *peers*. Redundancy offers a continuous mode of operation and eliminates a central point of failure for LAN switching and routing. The WANJet appliance supports redundancy using a second WANJet appliance on a LAN, connected to a redundant router. The second WANJet appliance is known as a redundant peer. If one of the LAN's routers fail, the corresponding WANJet appliance detects that the router is down and continues service through the remaining active router and WANJet appliance.

Not only does this redundant system offer you a continuous mode of operation, but it also provides load balancing under normal network conditions by distributing network traffic over two WANJet appliances. Figure 6.1 shows an example of redundant peer configuration.




Figure 6.1 Redundant peer configuration

To access redundant peers through the Web UI of a remote WANJet appliance, you must add both the primary peer (WANJet1) and the redundant peer (WANJet2) to the remote WANJet appliance (WANJet3). For example, to configure the WANJet appliances shown in Figure 6.1, you need to perform the following tasks:

- On WANJet1, set up WANJet2 as a redundant peer, and WANJet3 as a remote WANJet appliance.
- On WANJet2, set up WANJet1 as a redundant peer, and WANJet3 as a remote WANJet appliance.
- On WANJet3, add both WANJet1 and WANJet2 as remote WANJet appliances, set its type as **Redundant**, then specify WANJet2's IP address as the **Redundant Peer**.

For information about how to add remote WANJet appliances, see *Managing remote WANJet appliances*, on page 6-20.

To set up a local WANJet appliance as a redundant peer

1. In the navigation pane, expand **Configuration** and click **Local WANJet**.
The Local WANJet screen opens.
2. Check the **Redundant Peer IP** box.
A new field appears next to the check box.
3. In the new **Redundant Peer IP** field, type the IP address of the WANJet appliance that is the redundant peer of this WANJet appliance.
4. Click the **Save** button.
The Local WANJet screen refreshes, and the changes are committed to the WANJet appliance.

Configuring one-arm topology

You can deploy the WANJet appliance out-of-line in a one-arm topology, with one physical connection to the LAN and no direct connection to the WAN.




Figure 6.2 One-arm deployment

For more information about this configuration, see *One-arm deployment*, on page 3-4.

To configure the WANJet appliance in a one-arm configuration

1. Perform initial configuration of the WANJet appliance as described on the *Quick Start Card*.
2. Position the WANJet appliance so that it connects to a router that is inline, connecting the WAN port of the WANJet appliance to the router (see Figure 6.2).
3. Log on to the WANJet appliance, as described in *Logging on to the WANJet Web UI*, on page 4-1.
4. In the navigation pane, expand **Optimization** and click **Operational Mode**.
The Operational Mode screen opens.
5. For the **Topology** setting, select **One-Arm**.
When you select **One-Arm** topology for the operational mode setting, a new section entitled Redirection Method appears.
6. For the **Redirection Method** setting, select one of the following options:
 - **Static Routing**
Use this option if each client on your LAN is configured to route network traffic through the WANJet appliance. Go to step 8.
 - **Transparent Proxy**
Use this option if LAN traffic designated for optimization is directed to the WANJet appliance by a router (also for GRE tunneling). Go to step 7.

- **Non-Transparent Proxy**
Use this option if you want the WANJet appliance to act as the default gateway for all clients in the LAN. In this configuration, every client on the LAN must be configured to use the WANJet appliance's IP address as its default gateway. Go to step 8.
7. If you select **Transparent Proxy** in step 3, a new section entitled **Discovery Method** appears. From this section, select one of the following options:
- **Static**
Use this option if passthrough traffic is not routed to the WANJet appliance. (Use this option also for GRE tunneling.) When you choose **Static** as the Discovery Method, only network traffic that is scheduled for optimization is routed through the WANJet appliance. This traffic is lost if the WANJet appliance is not running. If you select this option, skip to step 9.
 - **WCCP v2**
Use this option if the WANJet appliance communicates with your network router using the Web Cache Coordination Protocol (WCCP). In this case, all network traffic is routed through the WANJet appliance, but the router by-passes the appliance if the WANJet appliance is not running. If you select this option, proceed to step 8.
8. If you select **WCCP v2** as the discovery method in step 7, configure the following settings:
- a) In the **Service ID** box, type the service group identifier. This must be a number between 51 and 100, and must match the service ID configured on the LAN router.
 - b) In the **Priority** box, type the priority assigned by the router to the service group. This number determines the order in which redirection rules are followed. This must be a number between 0 and 255, and must match the priority configured on the LAN router.
 - c) In the **Router** box, type the IP address that the LAN router uses to communicate with the WANJet appliance.
 - d) Check the **Authenticate** box.
 - e) If WCCP is configured to require authentication between the WANJet appliance and the LAN router, type a password in the **Password** box.
9. Click the **Save** button.

Introducing high-availability features

When combined with a redundant network topology, the WANJet appliance's high-availability features ensure that a failure in a networking device connected to WANJet appliance peers does not result in loss of connectivity or loss of application optimization.

The following WANJet appliance functions work with the high-availability features of the network:

- Bridging of traffic not configured for optimization
- Fail-to-wire mode for all traffic
- Peer port for redundant data path and peer health checks
- Load balancing and health checks of remote WANJet appliance peers

The following sections discuss each feature in detail and provide information on integrating these features into the network architecture. Additionally, standard and alternate network configurations are discussed.

WANJet appliance bridging functionality

When deployed in an inline configuration (LAN and WAN ports connected), the WANJet appliance acts as a Layer 2 bridge for network traffic that is not configured for WAN optimization. Ethernet frames with unoptimized traffic are bridged between the LAN and WAN interfaces.

The ability to act as a bridge for traffic that is not optimized allows the WANJet appliance to be incorporated into redundant network topologies and to support the high-availability features of other network devices. Protocols such as the Address Resolution Protocol (ARP), the Spanning Tree Protocol (STP), the Virtual Router Redundancy Protocol (VRRP) and the Hot Standby Redundancy Protocol (HSRP) function normally in the presence of a WANJet appliance.

Redundancy protocols typically create a shared Virtual IP address (VIP). The VIP is the default gateway for the hosts on the LAN. One router uses the VIP to actively pass traffic, while the other router acts as a standby. The redundancy protocol sends multicast packets between the active and standby routers to indicate that the active router is healthy and continues to pass traffic. These packets are bridged through the active router's WANJet appliance and LAN switches, and bridged back through the peer WANJet appliance to the standby router.

If a failure in a network component (other than the WANJet appliance) prevents the multicast packets from reaching the standby router, the standby router becomes the active router by sending out an ARP packet indicating that it now owns the VIP (this process is often called gratuitous ARP). The gratuitous ARP packet is a Layer 2 broadcast packet, which is bridged by

the WANJet appliance to the LAN hosts. LAN hosts then begin using the new router (but with the same IP address, namely the VIP) as their default gateway to send traffic to other networks.

◆ **Note**

WANJet appliances themselves can use the VIP as their default gateway IP address. If WANJet appliances connect directly from their WAN ports to their routers, both WANJet appliances must use the non-virtual IP address of their connected router's interface.

To use the VIP as the WANJet appliances' default gateway (to achieve redundant default gateways for the WANJet appliances), both of the WANJet appliance WAN ports must connect to switches or other Layer 2 devices that then connect to both routers. Depending on the details of the topology and configuration of your Layer 2 devices and routers, this may introduce Layer 2 loops that require resolution through the Spanning Tree Protocol or other means.

WANJet appliance fail-to-wire feature

A core feature of WANJet appliance high availability is its **fail-to-wire** feature (set by default). Fail-to-wire functionality guarantees that a failure of a WANJet appliance does not block data traveling between the LAN and WAN ports when the WANJet appliance is deployed in an inline topology (as opposed to one-armed topology). When a failure in WANJet appliance occurs, the WANJet appliance network interface hardware opens a path that connects the LAN and WAN ports directly. Refer to *Setting operational modes*, on page 6-11, for how to configure the fail-to-wire setting.

A WANJet appliance in fail-to-wire state acts effectively as a patch panel connecting two Ethernet cables. In the event of a WANJet appliance failure, data continues to flow between the two connected devices (such as switches, routers, or another WANJet appliance) on either side of the WANJet appliance. By allowing data to pass between connected devices in this manner, WANJet appliance failure does not result in the loss of network connectivity for clients, servers, and other networking devices.

You can configure fail-to-wire to occur regardless of the type of failure in the WANJet appliance, including software bugs, hardware bugs, or hardware failures in components, such as memory chips or hard disks (except physical damage to the WANJet appliance's fail-to-wire hardware components), and loss of power to the WANJet appliance.

The fail-to-wire feature requires that the Ethernet parameters (that is, duplex and speed) of the connected devices' network interfaces are the same, as they would be if cabled directly together.

Duplex and speed

You must set the duplex and speed appropriately for the ports on the connected devices. F5 Networks recommends configuring the WANJet appliance interfaces and the interfaces of connected devices to autonegotiate duplex and speed.

After you configure the connected to autonegotiate, F5 Network recommends checking the Diagnostics report (see *Diagnostics reports*, on page 8-11) to determine whether both the LAN and WAN interfaces have autonegotiated the same settings. If so, fail-to-wire will work correctly in case of failure. If duplex, speed, or both settings have different values, you need to manually set the parameters on all devices to the same values.

Cable type

Cabling two network devices together may require use of an Ethernet cable with standard wiring (often called a straight-through cable), or may require an Ethernet cable with pinouts 1, 2, 3, and 6 of one connector wired to pinouts 3, 6, 1 and 2 (respectively) of the connector on the other end (often called a crossover cable). The WANJet appliance Gigabit Ethernet network interfaces can automatically sense which cable type is present (auto-sensing MDI/MDI-X), so during normal operation cable type should not be an issue.

However, in fail-to-wire mode, the effective cable type (that is, the combination of the two cable types) may or may not be appropriate for the two connected devices. As per the Gigabit Ethernet specification, Gigabit Ethernet network interfaces perform auto-sensing of the crossover cable, and configure themselves appropriately. If one or both devices possess Gigabit Ethernet interfaces, you can use any combination of the two cable types for the two cables connected to the WANJet appliance. If neither connected device possesses a Gigabit Ethernet network interface, you must choose the cable type based on the type of devices that effectively connect during fail-to-wire mode.

WANJet appliance fail-close feature

An alternative configuration to fail-to-wire exists. You can configure the WANJet appliance to *fail-close*, which breaks the connectivity between connected devices. This may be desirable if an administrator wishes to create a redundant network architecture in which all traffic is routed to the peer WANJet appliance when a WANJet appliance failure occurs. When used with the redundancy features of the other network components, fail-close can prevent the creation of an unoptimized path through the network. Fail-close requires a hardware modification on the WANJet 400 appliance.

If you cannot use fail-close, but requirements do not permit a path in the network that does not have optimization, you can use the router connected to the WAN port to perform policy-based routing of unoptimized traffic, directing it to the active peer WANJet appliance for optimization.

Consult the documentation for your routing device, and contact F5 Networks support for additional information on high-availability configuration of WANJet appliances with policy routing.

WANJet appliance peer port

Every WANJet appliance model has an Ethernet port labeled **Peer**. You connect the peer ports of two WANJet appliances using a crossover cable when the WANJet appliances are deployed as a redundant pair. The peer network has two functions: passing network traffic being optimized by a peer, and sending heartbeat packets between the peers.

Passing network traffic being optimized by a peer

The peer network provides an alternate path for network traffic that is being optimized by a given WANJet appliance, but due to a failure in the network the normal path to the WANJet appliance is not available. When a failure on the network prevents traffic from reaching a WANJet appliance, redundant paths in a network should permit this traffic to take a path to the peer WANJet appliance. However, this peer has no knowledge of established optimized sessions belonging to the other WANJet appliance, because WANJet appliances do not pass information about existing sessions to each other.

Instead of passing state information, the WANJet appliances pass actual traffic from one WANJet appliance to the other (and possibly receive it back) over the peer network. The WANJet appliance passes packets to the peer if it does not recognize the packets as part of one of its own existing optimized sessions, or if the packet has not yet been identified as part of an unoptimized session that should be bridged between LAN and WAN ports.

The peer network uses multiple VLANs; a WANJet appliance sending packets to its peer selects the VLAN based on the interface from which the packet was received (LAN or WAN), and whether the packet is being sent to a peer (first transit of peer network) or returned from a peer (second transit of peer network). In this way, the VLAN tag represents the meta-information associated with each packet that is needed to make decisions on packet handling.

Packets arriving on the LAN or WAN ports

The WANJet appliance checks all packets arriving on the LAN or WAN ports to determine how to handle them. A ***SYN packet*** is a type of packet that TCP uses when initiating a connection to another computer. The WANJet appliance handles SYN packets differently from other packets and applies the following logic when deployed in a peer configuration:

If the packet is a SYN packet

The WANJet appliance checks the source IP address and destination port against the optimization rules:

- If a matching rule is found, the WANJet appliance that received the packet optimizes the packet's session.
- If a matching rule is not found, the WANJet appliance that received the packet bridges the packet.

If the packet is not a SYN packet

The WANJet appliance checks the source and destination IP addresses and ports against the table of existing optimized sessions:

- If a matching record is found, the WANJet appliance that received the packet optimizes the packet

If a matching record is not found, the WANJet appliance then checks the source and destination IP addresses, and the ports against the table of connections to be bridged:

- If a matching record is found, the WANJet appliance that received the packet bridges the packet.
- If a matching record is not found, the WANJet appliance forwards the packet over the peer network to the peer WANJet appliance.

Packets arriving on the peer port

If the packet is in the VLAN representing the first transit of the peer network, the WANJet appliance checks the source and destination IP addresses and ports against the table of existing optimized sessions:

- If a matching record is found, the WANJet appliance that received the packet on the peer port optimizes the packet.
- If a matching record is not found, the WANJet appliance returns the packet over the peer network to the WANJet appliance that first received the packet.

If the packet is in the VLAN representing the second transit of the peer network (no matching record found by the peer after the first transit of peer network), the packet's source and destination IP addresses and ports are added to the table of connections to be bridged, and the packet is bridged in the direction indicated by the VLAN in which it was sent.

Unlike the LAN and WAN ports, the peer port does not perform Layer 2 bridging. Connecting WANJet appliance peer ports does not create bridge loops.

Heartbeat packets between peers

The WANJet appliance also uses the peer port to send heartbeat packets to its peer. These heartbeats are standard ICMP echo requests and responses. Failure to obtain a response after a request causes the requesting WANJet appliance to mark its peer as down, and to stop sending across the peer connection those packets which may be part of previously optimized sessions.

WANJet appliance remote peer load balancing and failure detection

To configure a pair of WANJet appliances as peers, you must add the IP addresses of the two peers to all remote WANJet appliances. When a remote WANJet appliance starts, it attempts to establish optimized TCP tunnels to both peers. If the network topology allows packets from the remote WANJet appliance to both of the peer IP addresses, the remote WANJet appliance passes optimized sessions to both peers, achieving a basic form of load balancing.

If the remote sites themselves have peers, you configure both remote peers with the IP addresses of the peers. For example, Site 1 has WANJet appliance peers A and B. Site 2 has WANJet appliance peers C and D. WANJet appliances A and B each attempt to establish tunnels to WANJet appliances C and D, and WANJet appliances C and D each attempt to establish tunnels to WANJet appliances A and B.

A WANJet appliance that establishes two tunnels to remote peers constantly monitors the optimized TCP tunnel connection. If one connection fails, either due to failure of a remote WANJet appliance or failure in the network path, the WANJet appliance considers the remote WANJet appliance to be down, and redirects new optimized traffic to the remaining remote peer until the tunnel can be reestablished. In our example, if a networking device at Site 2 fails, causing WANJet appliance C to be inaccessible, then WANJet appliances A and B send new connections only to WANJet appliance D.

When a remote WANJet appliance experiences a failure, existing optimized sessions to the failed WANJet appliance are lost, and must be reestablished at the application level. Many applications perform this step without end-user interaction (for example, CIFS file sharing). Some applications require end-user interaction. A web browser connected to an HTTP server might require the user to relick on a link or hit the refresh button, for example. The WANJet appliance does not exchange state information about existing connections with its peer. See *WANJet appliance peer port*, on page 6-33, for more information.

Alternate high-availability configurations

Most high-availability configurations use redundant WANJet appliance peers inline in a topology with parallel data paths. However, you can also deploy redundant WANJet appliances using other topologies, such as the one-arm configuration or an advanced inline configuration.

High-availability and one-arm configuration

When you deploy WANJet appliances in a one-arm topology, you have several options to direct traffic for optimization. The most common of these options is using the Web Cache Control Protocol (WCCP). WCCP contains load-balancing and fault tolerance mechanisms, which permit the use of multiple WANJet appliances.

In this scenario, peer functionality using the peer port and peer network is unnecessary. The WCCP-enabled router or other device that directs traffic to WANJet appliances performs all needed tasks to handle failures of the WANJet appliances or failures in the intervening network.

The number of WANJet appliances employed for redundancy is not limited to two; you can place multiple WANJet appliances, up to the limits of the WCCP device.

In addition to WCCP, you can use policy-based routing to direct traffic to the WANJet appliances in a one-arm deployment. The devices performing policy-based routing may provide health-checking mechanisms to verify the routing through the WANJet appliances.

Alternatively, you can use the peer functionality in combination with fail-to-wire to allow a WANJet appliance to seamlessly optimize traffic if its peer fails. For further information on redundancy in policy-based routing scenarios, consult the documentation for your routing device, and contact F5 Networks support for additional information on high-availability configuration of WANJet appliances with policy routing.

Alternate inline configuration for high availability

Additionally, there is another configuration for redundant WANJet appliances that are deployed inline. You can deploy two WANJet appliances in sequence (with the WAN port of one connecting to the LAN port of the second). You configure both to optimize the same network traffic. The WANJet appliance closer to the clients or servers performs the optimizations, while the WANJet appliance behind it bridges all traffic.

If the optimizing WANJet appliance fails in this configuration, the fail-to-wire feature passes unoptimized traffic to the second WANJet appliance, which performs the optimization. Sequential deployment eliminates the potential drawback to the basic inline topology, that a WANJet appliance in the fail-to-wire state can create a network path with no optimization. This deployment scenario is attractive when the network topology itself does not contain redundant paths (often the case with a branch office network), but you want redundancy of WANJet appliances.



7

Configuring Service Policies

- Defining IT service policies
- Creating Application QoS policies
- Managing WAN links

Defining IT service policies

With the IT service policy feature, you define services that you can use to achieve specific QoS standards. You can group ports, machines, and subnets under the heading of an IT service policy. By assigning a minimum and a maximum amount of bandwidth to this service (in an Application QoS policy), you treat this group of ports, machines, and subnets as one entity. This is simpler than creating many different services, each of which handles a single type of traffic.

You can use the WANJet appliance to define IT service policies and Application Quality of Service (QoS) policies for your various applications, and apply them to optimally allocate bandwidth. An ***IT service policy*** specifies a named group of ports, machines, and subnets. When you define an Application QoS policy, you can specify an IT service group as well as the bandwidth you want to allocate to particular applications, such as:

- Mission-critical applications
- Video and voice streaming
- Interactive video or voice
- Data transfers
- Web-based applications

These individual classes of applications have very different network requirements. The challenge is to align the network services to the application's requirements from a performance perspective.

Adding, editing, or removing an IT service policy

You can add, edit, or remove an IT service policy from the IT Service Policy screen.

To add an IT service policy

1. In the navigation pane, expand **Optimization** and click **IT Service Policy**.
The WANJet IT Service Policies screen opens.
2. Click the **Add** button.
The Add IT Service Policy popup screen opens.
3. In the **Policy Name** box, type a name for the policy.
4. In the **From** box, type the IP address of the subnet that sends the data, for which you want to specify an IT service policy.
5. In the **Netmask** box, type the full netmask of the subnet that sends the data, for which you want to specify an IT service policy.
6. In the **To** box, specify the subnet that receives the data, for which you want to specify an IT service policy.

7. In the **Netmask** box, type the full netmask of the subnet that receives the data, for which you want to specify an IT service Policy.
8. You can specify a port in one of the following ways:
 - From the **Ports** list, select a port.
 - In the **From Port** and **To** boxes, specify a range of ports.
9. From the **Protocol** list, select a protocol type for the ports that you specified.
10. Click the **OK** button.
The Add IT Service Policy screen closes, and the WANJet IT Service Policies screen refreshes with the new IT service policy displayed.
11. Click the **Save** button to save the changes.

To edit or remove an IT service policy

1. In the navigation pane, expand **Optimization** and click **IT Service Policy**.
The WANJet IT Service Policies screen opens.
2. Click the name of the IT service policy that you want to edit or remove.
The Edit IT Service Policies screen displays in a separate browser window.
3. Edit the policy settings, or click the **Remove** button to delete the policy.
The Edit IT Service Policies screen closes and the WANJet IT Service Policies screen refreshes with the new IT Service Policy displayed.
4. Click the **Save** button to save the changes.

Creating Application QoS policies

Application QoS policies help you to obtain better network performance by dedicating bandwidth to specific network traffic traveling between two WANJet appliances. At the same time, you can ensure that providing sufficient bandwidth to one or more data flows does not handicap the transmission of other data flows. **Application QoS** is a per endpoint setting that you can use to override the tuning page bandwidth setting in a multi-node network.

The Application QoS policies handle two types of services:

- **Fundamental services**

The basic protocols supported by your network, such as FTP, HTTP, HTTPS, Pop3, and so on.

- **IT service policies**

Tailored services that include different types of traffic. If you want to create an Application QoS policy to handle tailored services, you need to have created or imported the IT service policy already.

(See *Defining IT service policies*, on page 7-1.)

Before you create an Application QoS policy, you need to plan how you want to allocate bandwidth.

Adding, editing, or removing an Application QoS policy

You can add, edit, or remove an Application QoS policy from the screen, Manage the Application QoS Settings of a Remote WanJet.

To add an Application QoS policy to a remote WANJet appliance

1. In the navigation pane, expand **Optimization** and click **Application QoS**.
The Application QoS screen opens.
2. In the table, click the IP address of the remote WANJet appliance to which you want to apply an Application QoS policy.
The Manage the Application QoS Settings of a Remote WANJet appliance popup screen opens.
3. In the **Link Bandwidth** box, type the bandwidth size of the link between the local WANJet and the remote WANJet, and from the **Link Bandwidth** list, select a unit (Kb/s or Mb/s).
4. Click the **Add** button.
The Application QoS Policy popup screen opens.
5. In the **Alias** box, type a name for the policy.

6. In the **Bandwidth** box, type the percentage of bandwidth that you want to guarantee that the policy can use. For example, if you specify 50%, the connections associated with the policy are guaranteed 50% of the available bandwidth if needed.
Note: The bandwidth that you allocate to all of the Application QoS policies should not exceed 100%.
7. In the **Maximum** box, type the maximum percentage that the policy can borrow from unused additional bandwidth. For example, if you specify 90%, the connections associated with the policy can use up to 90% of the additional bandwidth that is available.
8. In the **Services** box, add the services or IT service policies to use for the Application QoS policy. For each service that you add:
 - a) From the **Services** list, select the service or IT service policy to add to the Application QoS policy.
 - b) From the adjacent service type list, select the associated protocol (TCP or UDP), if applicable.
*Note: You can configure some ports for both TCP and UDP protocols. To do this, select the service port (for example, **FTP**) and then select **TCP**. Then on a new line, select service **FTP** again, and service type **UDP**. If you select **VoIP**, it uses only the UDP protocol. If you choose a defined IT Service policy from the menu, the adjacent service type menu disappears.*
 - c) Click the **OK** button.
The Manage the Application QoS Settings of a Remote WANJet appliance popup screen opens.
9. Repeat steps 4-8 to add as many other Application QoS policies as you need.
10. Click the **OK** button.
The Manage the Application QoS Settings of a Remote WANJet appliance popup screen closes.
11. Click the **Save** button.
The Application QoS screen refreshes.

To edit or remove an Application QoS policy from a remote WANJet appliance

1. In the navigation pane, expand **Optimization** and click **Application QoS**.
The Application QoS screen opens.
2. Click the link for the remote WANJet appliance for which you want to edit or remove an Application QoS policy.
The Manage the Application QoS Settings of a Remote WANJet appliance popup screen opens.
3. Click the link for the Application QoS policy that you want to edit or remove.
4. Edit the settings to alter an existing policy, or click the **Remove** button to delete the policy.
5. Click the **OK** button.
The Application QoS screen opens.
6. Click the **Save** button.
The Application QoS screen refreshes.

Managing WAN links

Using WAN links, you can add an Application QoS Policy to the traffic passing through the local WANJet appliance and going to a remote network, whether or not the remote network has a WANJet appliance installed. In this way, you can manage and manipulate the bandwidth size for all the traffic transferred through the local WANJet appliance, regardless of the traffic's processing mode.

Adding, editing, or removing WAN links

To add a new WAN link

1. In the navigation pane, expand **Optimization** and click **Application QoS**.
The Application QoS screen opens.
2. Click the **Add WAN Link** button.
The Manage the Application QoS Settings of a WAN Link popup screen opens.
3. In the **WAN Link Alias** box, type a name.
4. In the **Link Bandwidth** box, type the size of the bandwidth between the local WANJet appliance and the WAN network.
5. From the **Link Bandwidth** list, select a unit (Kb/s or Mb/s).
6. Click the **OK** button.
The Manage the Application QoS Settings of a WAN Link screen closes, and the Application QoS screen refreshes with the new WAN link displayed.
7. Click the **Save** button to save the changes.

To edit or remove a WAN link

1. In the navigation pane, expand **Optimization** and click **Application QoS**.
The Application QoS screen opens.
2. From the Alias column, click the name of the link that corresponds to the WAN Link that you want to edit or remove.
The Manage the Application QoS Settings of a WAN Link popup screen opens.
3. Edit the WAN Link settings, or click the **Remove** button to delete the WAN Link.
4. Click the **OK** button.
The Manage the Application QoS Settings of a WAN Link screen closes, and the Application QoS screen refreshes.
5. Click the **Save** button.

Adding a subnet to a WAN Link

You can add subnets or machines to any existing WAN Link. In doing so, you can make use of the Application QoS policies with more nodes (computers, subnets, or networks).

◆ **Note**

In addition to the following procedure, you can also add a subnet when you add a WAN Link.

To add a subnet to a WAN Link

1. In the navigation pane, expand **Optimization** and click **Application QoS**.
The Application QoS screen opens.
2. From the Alias column, click the name of the link that corresponds to the WAN Link to which you want to add a subnet.
The Manage the Application QoS Settings of a WAN Link popup screen opens.
3. Click the **Add** button beneath the Support Subnet table.
The Add Subnet popup screen opens.
4. In the **Supported Subnet** box, type the IP address of the machine or subnet that you want to add.
5. In the **Netmask** box, type the netmask of the new machine or subnet.
6. In the **Machine(s) Alias** box, type a name for the new machine or subnet.
7. Click the **OK** button.
The Application QoS screen opens, and the new subnet appears in the Support Subnet column.
8. Click the **OK** button.
The Manage the Application QoS Settings closes, and the Application QoS screen refreshes with the new changes.
9. Click the **Save** button.

To edit or remove a subnet from a WAN Link

1. In the navigation pane, expand **Optimization** and click **Application QoS**.
The Application QoS screen opens.
2. From the Alias column, click the name of the link that corresponds to the WAN Link from which you want to edit or remove a subnet.
The Manage the Application QoS Settings of a WAN Link popup screen opens.
3. From the Supported Subnet column, click the name that corresponds to the subnet that you want to edit or remove.
The Edit Subnet popup screen opens.
4. Edit the subnet settings, or click the **Remove** button to delete the subnet.
5. Click the **OK** button.
The Edit Subnet screen closes and the Manage the Application QoS Settings of a WAN Link screen displays.
6. Click the **OK** button.
The Manage the Application QoS Settings of a WAN Link screen closes and the Application QoS screen displays with the new changes.
7. Click the **Save** button.



8

Monitoring Performance

- Introducing reports
- Status report
- Real Time Traffic report
- Comparative Throughput reports
- Diagnostics reports
- Third-party reporting systems

Introducing reports

The WANJet appliance provides reports that you can use to monitor the status, connectivity, and performance of your WANJet appliance. You can display the reports in the Web UI by expanding **Reports** in the navigation pane and clicking one of the following report options:

- Status
- Real Time Traffic
- Comparative Throughput
- Diagnostics

It is easier to view the reports if your Web UI browser window is full-screen size.

Note

*To ensure accurate reports, we suggest that you frequently synchronize the time setting on all WANJet appliances. For more information, see **Configuring time settings**, on page 5-6.*

This chapter describes how to display the reports and also explains other ways of obtaining information about performance, including using network diagnostic tools, reviewing operational logs, and integrating with third-party reporting tools.

Status report

The Status report provides the status and details of remote WANJet appliances. If the remote WANJet appliance has a redundant peer, the Status report also displays details about the peer appliance. The Status report is the first screen displayed when you log on to the WANJet Web UI.

To view the Status report

In the navigation pane, expand **Reports** and click **Status**.

The initial WANJet Status screen displays the following information for all remote WANJet appliances:

- Status (active or inactive)
- IP address of the remote WANJet appliance
- Alias
- Version of the WANJet appliance software

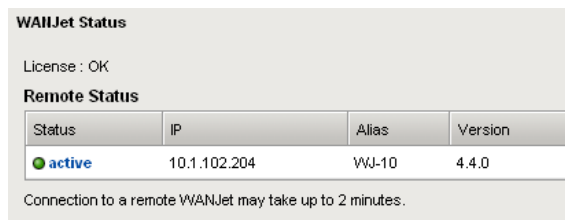
The WANJet Status screen displays the following information for the local WANJet appliance:

- License key status (not entered, not valid, expired, or OK)

◆ Note

If you change the settings of a remote WANJet appliance, it takes a couple of minutes for the two WANJet appliances to communicate the changes and display an updated Status report.

Figure 8.1 shows a sample Status report for a WANJet appliance that has one active remote WANJet appliance.



The screenshot shows the 'WANJet Status' page. At the top, it says 'License : OK'. Below that is a section titled 'Remote Status' containing a table with four columns: Status, IP, Alias, and Version. The table has one row with the values: active (with a green dot icon), 10.1.102.204, VWJ-10, and 4.4.0. Below the table, a note states: 'Connection to a remote WANJet may take up to 2 minutes.'

Status	IP	Alias	Version
● active	10.1.102.204	VWJ-10	4.4.0

Figure 8.1 WANJet Appliance Status report

To test connectivity with the remote WANJet appliance

Click the information in the Status column of any remote WANJet appliance. The Remote WANJet Link Status screen opens and shows the results of running the **ping** and **traceroute** commands to the remote WANJet appliance. For details on the **ping** command, see *Ping*, on page 8-24; for details on the **traceroute** command, see *Traceroute*, on page 8-25.

Real Time Traffic report

The Real Time Traffic report displays a graph of all network traffic, in real time, over both the LAN and the WAN. This provides an at-a-glance overview of the network traffic that is passing through the WANJet appliance.

To view a graph of network traffic in real time

In the navigation pane, expand **Reports** and click **Real Time Traffic**. The Real Time Traffic report opens as shown in Figure 8.2. You may need to reply to Security Information questions, and press Enter or spacebar to view the report.




Figure 8.2 Real Time Traffic report (standard inline configuration)

In a Real Time Traffic report for a standard inline configuration:

- The vertical axis indicates the amount of network traffic, in bits per second.
- The horizontal axis shows the time (using a 24-hour clock) in hours, minutes, and seconds, in 10-second intervals.
- The blue line (**LAN In**) represents the raw data that is destined for the WAN passing into the local WANJet appliance from the LAN.
- The yellow line (**LAN Out**) represents optimized data passing out of the local WANJet appliance on its LAN interface.
- The red line (**WAN In**) represents optimized data passing into the local WANJet appliance from its remote partner.

- The green line (**WAN Out**) represents reconstituted data passing out of the local WANJet appliance on its WAN interface.

If you run the Real Time Traffic report on a WANJet appliance set up using a one-arm configuration, the report changes. Because there are no LAN and WAN connections in that configuration, you see only two lines showing the **Data In** and **Data Out**, as shown in Figure 8.3.




Figure 8.3 Real Time Traffic report (one-arm configuration)

In a Real Time Traffic report for a one-arm configuration:

- The vertical axis indicates the amount of network traffic, in bits per second.
- The horizontal axis shows the time (using a 24-hour clock) in hours, minutes, and seconds, in 10-second intervals.
- The blue line (**Data In**) represents the raw data that is coming into the local WANJet appliance.
- The yellow line (**LAN Out**) represents optimized data passing out of the local WANJet appliance.

For more information on one-arm configuration, refer to *One-arm deployment*, on page 3-4.

Comparative Throughput reports

You can generate a Comparative Throughput report based on any combination of traffic direction, data type, and time period. Comparative Throughput reports refresh automatically every two minutes.

At the top of each report, there is a summary of the amount of data handled before and after compression, and the compression ratio achieved (expressed as a percentage). These figures vary according to the time period selected and the direction of traffic.

You can download any of the reports as text files with comma-separated values (CSV). Then, you can import CSV reports into a database or spreadsheet package.

To generate a Comparative Throughput report and save it to a file

1. In the navigation pane, expand **Reports** and click **Comparative Throughput**.
2. Near the top of the main screen, click one of the following to select the direction of traffic and display the associated report:
 - **Total Throughput**
Shows all traffic that the WANJet appliance processes.
 - **Sent Throughput**
Shows the outgoing (sent) data that was optimized.
 - **Received Throughput**
Shows the incoming (received) data that was optimized.
3. To determine how to display the data, click one of the following options within the report:
 - **Performance Increase report**
Shows the performance increase by comparing the bandwidth before and after optimization. See *Performance Increase report*, on page 8-6.
 - **Actual Bandwidth Expansion report**
Shows the actual bandwidth amount that the WANJet appliance freed during optimization. See *Actual Bandwidth Expansion report*, on page 8-7.
 - **Optimized Data report**
Lets you compare the amount of raw data with the amount of optimized data. See *Optimized Data report*, on page 8-8.
 - **Overall Data report**
Shows the amount of passthrough data, raw data, and compressed data. See *Overall Data report*, on page 8-9.
 - **Link Utilization report**
Displays the average amount of bandwidth used, compared to what would have been used without optimization. See *Link Utilization report*, on page 8-10.

- Under the report, click the time period for which you want to view collected data (hour, day, week, month, quarter, or year). The default is hour.

Note: The WANJet appliance saves all of the reports generated for the last hour, every hour. If you stop or restart the WANJet appliance, or any external termination occurs, you can view the last set of saved reports when you restart the WANJet appliance.

- If you want to save the report to a file, in the Download Report box, click the **Download** button.
The report is saved into a text file with a CSV extension.

Performance Increase report

The Performance Increase report displays the percentage increase in bandwidth due to using the WANJet appliance.

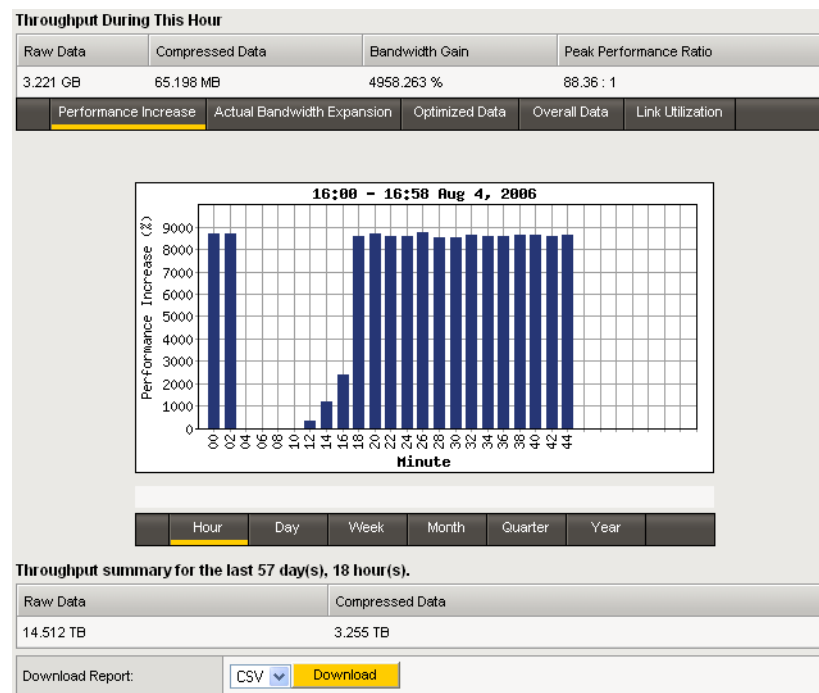


Figure 8.4 Performance Increase report

In this report, the vertical axis indicates the percentage increase in bandwidth. This is calculated by comparing the bandwidth freed up by the WANJet appliance with the bandwidth used after optimization. This is calculated as follows:

$$(\text{Freed Bandwidth} / \text{Bandwidth after optimization}) \times 100 = \text{Percentage Performance Increase}$$

For example, if your bandwidth before the optimization was 100 MB, and the bandwidth used by data after the optimization is 25 MB, then the amount of bandwidth freed up by the WANJet appliance is 75 MB. Using these values in the equation results in the following performance increase:

$$(75 \text{ MB}/25 \text{ MB}) \times 100 = 300\% \text{ performance increase}$$

Actual Bandwidth Expansion report

The Actual Bandwidth Expansion report displays the actual bandwidth amount that the WANJet appliance has freed by optimizing network data.

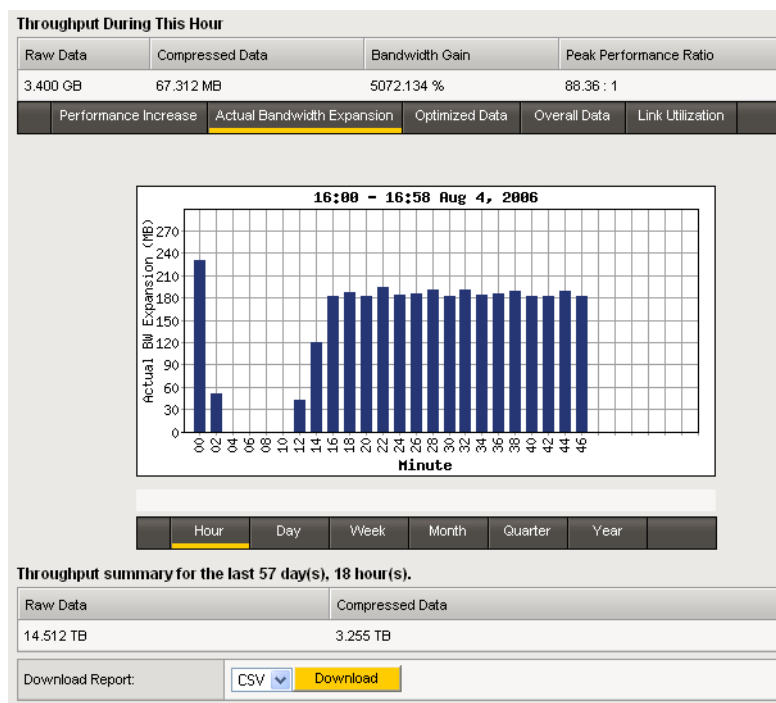


Figure 8.5 Actual Bandwidth Expansion report

In this report, the vertical axis represents the bandwidth expansion in kilobytes, megabytes, and so forth. (The unit used depends on the extent to which the bandwidth has expanded over the selected time period.)

Optimized Data report

The Optimized Data report displays the difference in the amounts of network traffic before and after the WANJet appliance processes the data.

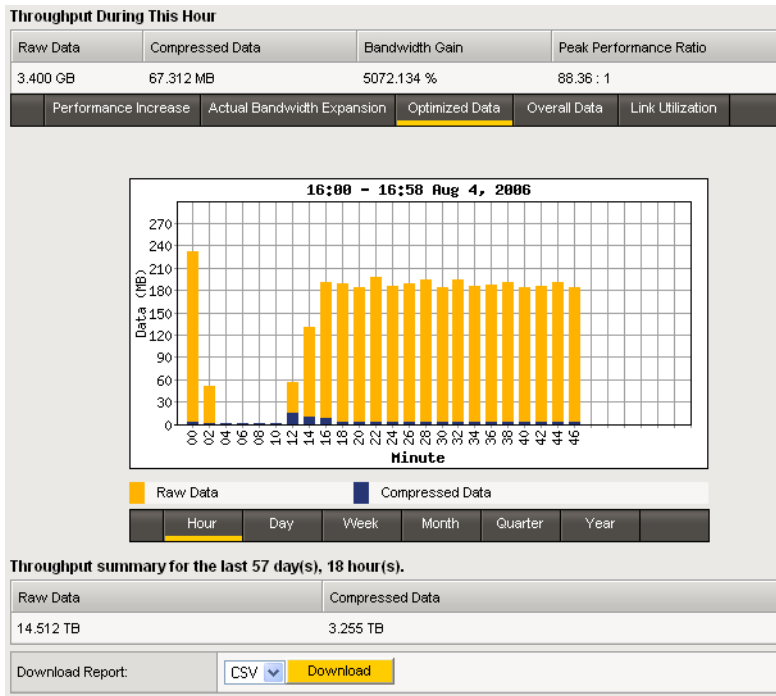


Figure 8.6 Optimized Data report

In this report:

- The vertical axis indicates the amount of network traffic before and after optimization (in kilobytes, megabytes, and so forth).
- The blue bar represents the amount of traffic before optimization.
- The yellow bar represents the amount of freed bandwidth.

Overall Data report

The Overall Data report allows you to view and compare the amounts of passthrough data, raw data, and optimized data.




Figure 8.7 Overall Data report

In this report:

- The vertical axis indicates the amount of data passing through the link (in kilobytes, megabytes, gigabytes, and so forth).
- The green bars represent the amount of passthrough data.
- The blue bars represent the amount of compressed (optimized) data.
- The yellow bars represent the amount of freed bandwidth.
- The bars as a whole represent the total amount of data passing through the WANJet appliance.

Link Utilization report

The Link Utilization report is similar to the *Optimized Data report*, on page 8-8. However, instead of showing the total amount of data optimized over a given time period, this Link Utilization report displays the average amount of bandwidth used per second, compared to what would have been used if network traffic had not been optimized.

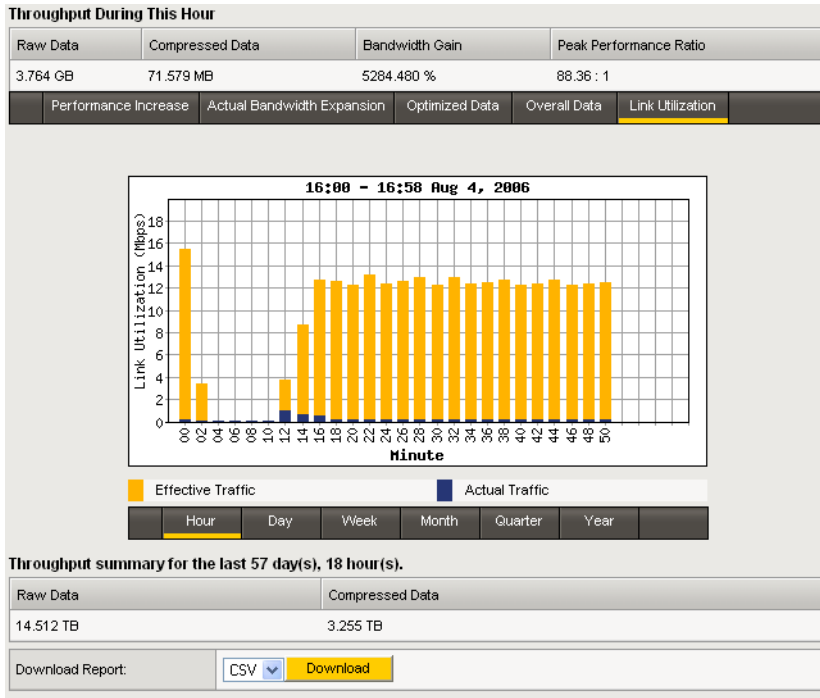


Figure 8.8 Link Utilization report

- In this report:
- The vertical axis indicates the amount of bandwidth (in kilobits per second, megabits per second, and so forth).
 - The blue bars represent the actual bandwidth used.
 - The bars as a whole represent the amount of bandwidth that would have been used if network traffic had not been optimized; therefore, the yellow bars represent the amount of bandwidth saved.

Diagnostics reports

Diagnostics reports provide you access to a range of useful information, such as IP addresses, error log files, and the results of popular network analysis tools.

To view diagnostics information

1. In the navigation pane, expand **Reports** and click **Diagnostics**. The initial diagnostics screen opens.
2. On the menu bar, click one of the following options to display detailed reports:
 - Monitoring
 - Connectivity
 - General

Each menu bar option provides several additional reports.

Monitoring

The following monitoring diagnostic reports are available:

- Interfaces
- Optimized Sessions
- Passthrough Sessions
- WANJet Links
- RADIUS
- TCP Statistics
- TDR Statistics
- QoS
- VLANs

Interfaces diagnostics

A WANJet appliance typically has at least two active network interfaces: one for the connection to the LAN and one for the connection to the WAN. In addition, if a redundant peer WANJet appliance is present on your LAN, there is an interface for that connection. For more information, see *Configuring redundant peers*, on page 6-26.

To view diagnostics for interfaces

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.
2. From the Monitoring menu, choose Interfaces.
The Interfaces report opens with the following information for each network interface:
 - MAC address (a unique identifier attached to most forms of networking equipment).
 - Maximum speed (in Mbits per second) and duplex setting (Full Duplex or Half Duplex).
 - Current status (Link OK/Link error).
 - Reception (RX) errors raised by the interface, including dropped packets, overruns, and frame errors.
 - Transmission (TX) errors raised by the interface, including dropped packers, overruns, carrier errors, and collisions

Interfaces

Reset

eth0 (LAN)

MAC:	00:E0:ED:08:69:55
Speed:	N/A
Status:	Link error
RX: errors: 0 dropped:0 overruns:0 frame:0	
TX: errors: 0 dropped:0 overruns:0 carrier:0	
collisions: 0	

eth1 (WAN)

MAC:	00:E0:ED:08:69:54
Speed:	1000 Full Duplex
Status:	Link ok
RX: errors: 0 dropped:0 overruns:0 frame:0	
TX: errors: 0 dropped:0 overruns:0 carrier:0	
collisions: 0	

eth2 (Peer)

MAC:	00:04:23:B9:23:16
Speed:	N/A
Status:	Link error

TCP Passthrough

Total	0 Packets/ 0 Bytes
-------	--------------------

UDP Passthrough

Total	0 Packets/ 0 Bytes
-------	--------------------

Figure 8.9 Interfaces diagnostics report

Optimized Sessions diagnostics

The Optimized Sessions report displays all of the network connections at the application layer that the WANJet appliance is currently optimizing. In contrast to the number of optimized sessions shown on the dashboard, the report shows established sessions only, and does not include those in the process of being set up or torn down. Therefore, the number of optimized sessions shown in the dashboard may not match the number in the Optimized Sessions report.

To view diagnostics for Optimized Sessions

1. In the navigation pane, expand **Reports** and click **Diagnostics**. The initial diagnostics screen opens.
2. From the Monitoring menu, choose Optimized Sessions. The Optimized Sessions screen opens.

*Note: Optionally, you can view the Optimized Sessions report by clicking **Optimized Sessions** from the navigation pane on any screen in the Web UI. The counter beside the **Optimized Sessions** link in the dashboard displays the number of optimized sessions including those in the process of being set up.*

The Optimized Sessions report has two sections: one for TCP and another for UDP traffic.

The TCP section contains the following information:

- **Processing mode**
How the traffic is processed, optimized or passthrough.
- **Local IP**
IP address and port for the local machine.
- **Direction**
Direction of optimized data traffic flow. A right arrow indicates that the direction is from the local machine to the remote machine. A left arrow indicates that the direction is from the remote machine to the local machine.
- **Remote IP**
IP address and port for the remote WANJet appliance.
- **WANJet IP**
IP address for the remote WANJet appliance handling the optimized session.

The UDP section contains two columns with the IP address and port number for each UDP session's source (from) and destination (to).

◆ Note

*For information about how to specify connections for optimization, see **Creating optimization policies**, on page 6-1.*

Passthrough Sessions diagnostics

A *passthrough session* is a network connection (at the application layer) for traffic that the WANJet appliance does not optimize, but allows that particular type of traffic to pass through the appliance untouched.

To view diagnostics for Passthrough Sessions

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.
2. From the Monitoring menu, choose Passthrough Sessions.
The Passthrough Sessions screen opens.

The Passthrough Sessions report has two sections: one for TCP traffic and another for UDP traffic, with specific information in each section.

From this screen, you can view the following reports:

- **All Passthrough Connections**
Displays a detailed list of all passthrough connections.
- **Optimize Eligible Connections**
Displays connections that were set up before the WANJet appliance was last activated. If the protocol and software allow it, you can intercept and reset these connections so that from this point on, they will be optimized. This is most useful for connections that need to be live for a long time so that they can transfer large amounts of data, such as replication processes.
- **Autopass**
Displays a list of connections that pass through automatically when the destination server is refusing connections.
- **Realtime**
Displays passthrough traffic throughput in real time.

◆ Note

*For information about how to specify connections for optimization, see **Creating optimization policies**, on page 6-1.*

WANJet Links diagnostics

The WANJet Links report displays information about other WANJet appliances that connect to the one you are working on.

To view diagnostics for WANJet Links

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.
2. From the Monitoring menu, choose WANJet Links.
The WANJet Links report opens and displays the following information:
 - **Remote IP**
IP address of the remote WANJet appliance.
 - **#Retrans**
Number of retransmitted packets to the remote WANJet appliance.
 - **#ACM5**
Number of network connections to the remote WANJet appliance that are being optimized.
 - **#ACM5 without compression**
Number of passthrough network connections (not optimized).

◆ Note

*For additional information about links to remote WANJet appliances, refer to **Managing remote WANJet appliances**, on page 6-20.*

RADIUS status diagnostics

The RADIUS report displays information about RADIUS authentication servers known to the local WANJet appliance. Remote authentication through the RADIUS protocol is an alternative to local authentication with a user name and password stored on the WANJet appliance.

◆ Note

*For information about how to configure WANJet appliance to use RADIUS authentication, see **Configuring remote authentication**, on page 5-2. For technical details about the RADIUS protocol, refer to <http://www.ietf.org/rfc/rfc2865.txt>.*

To view diagnostics for RADIUS status

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.
2. From the Monitoring menu, choose RADIUS.
The RADIUS report shows the RADIUS status, the number of RADIUS servers defined, and displays the following information for each RADIUS server:
 - **IP address**
 - **Secret**
The key that is used to authenticate RADIUS transactions between client and server.
 - **Timeout period, in seconds**
 - **Number of times to retry a connection**

◆ **Note**

The WANJet appliance displays a warning message if the settings for both the timeout and number of retries are too high; this could cause a delay in determining whether the RADIUS server is responding to a login attempt.

TCP Statistics diagnostics

The TCP Statistics menu provides the following reports for TCP connectivity activity:

- Connection States
- Packet retransmissions
- Receive queued packets pruned

The Connections States report is displayed by default.

To view diagnostics for TCP Statistics

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.
2. From the Monitoring menu, choose TCP Statistics.
The TCP Statistics screen displays with the Connections States report, by default.
3. Click the options above the report to view the following reports:
 - **Connections States**
 - **Packet Retransmissions**
 - **Receive Queue Packets Pruned**

Connection States

The Connection States report displays a graph of current state for each TCP connection that is visible to the WANJet appliance, for both optimized and passthrough connections. In this report, three lines represent the number of connections in the following states:

- **ESTABLISHED**
Established connections have been successfully opened and are working normally.
- **TIME-WAIT**
Connections in the TIME-WAIT status are waiting to see that the remote TCP received the acknowledgment of a connection termination request. This can take up to four minutes.
- **Other**
Other possible connection states include:
 - LISTEN
 - SYN-SENT
 - SYN-RECEIVED
 - FIN-WAIT-1
 - FIN-WAIT-2
 - CLOSE-WAIT
 - CLOSING
 - LAST-ACK

For more information about these states, see IETF RFC #793 at <http://www.ietf.org/rfc/rfc793.txt>.

Packet Retransmissions

TCP segments that time out without being acknowledged by a destination host are retransmitted by the source host. A high number of these retransmitted segments can indicate network problems. Therefore, the Web UI includes a report that tracks those numbers and their trends.

The Packet Retransmissions report consists of a graph with a blue line. The blue line indicates the number of TCP segments (which often correspond to IP packets) that had to be retransmitted per second.

Receive Queue Packets Pruned

The Receive Queue Packets Pruned report provides a graphic representation of the number of segments pruned from the TCP receive queue due to socket overrun. Pruning can occur if the TCP receive buffer on the receiving host is too large. The optimal buffer size is twice the product of the bandwidth and the delay.

For more information about TCP tuning background, see <http://www.didc.lbl.gov/TCP-tuning/background.html>.

TDR Statistics diagnostics

Transparent Data Reduction (TDR) further enhances network optimization by storing the contents of frequently accessed files in memory. You can view information about TDR statistics through the WANJet appliance Web UI. The report shows TDR-2 statistics that relate to repeat transfer data reduction. For more information about TDR, refer to *Transparent Data Reduction*, on page 2-2.

◆ Note

The WANJet appliance updates the statistics in the report at the completion of a session. For example, if you transfer a 1 GB file, the updated TDR-2 statistics are available when the file transfer is complete.

To view diagnostics for TDR Statistics

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.
2. From the Monitoring menu, choose TDR Statistics.
The TDR Statistics screen opens and displays the following TDR-2 information:
 - **WANJetIP**
IP address of the remote WANJet appliance.
 - **Sent Bytes (TDR)**
The amount of data sent in bytes, to which TDR-2 has been applied since the WANJet Link became active.
 - **Sent Bytes (other)**
Amount of data in bytes to which TDR has not been applied.
 - **Received Bytes (TDR)**
Amount of received data in bytes to which TDR-2 has been applied.
 - **Received Bytes (other)**
Amount of received data in bytes to which TDR has not been applied.
 - **TDR efficiency %**
Percentage of data sent across the link to which TDR-2 has been applied. The bold number at the bottom of the report is the average for all remote WANJet appliance links.

QoS diagnostics

Quality of Service (QoS) policies can help to improve network performance by dedicating bandwidth to specific network traffic.

To view diagnostics for QoS policies for remote networks

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.

2. From the Monitoring menu, choose QoS.
The QoS report opens and displays the following information:
 - **Remote**
The Remote network that has a QoS policy assigned to it.
 - **Policy**
Name of the QoS policy assigned to the remote network.
 - **Rate**
Actual bandwidth assigned to each QoS policy.
 - **Bytes Sent**
Number of bytes sent for each QoS policy.
 - **Packets Sent**
Number of packets successfully sent for each QoS policy.
 - **Dropped**
Number of packets dropped for each QoS policy.

◆ **Note**

*For additional information about QoS, refer to **Creating Application QoS policies**, on page 7-3.*

VLAN diagnostics

A Virtual LAN (VLAN) is a computer network which has its boundaries defined logically, rather than physically. VLANs must be explicitly added to the WANJet Web UI, since they are often implemented by adding tags to Ethernet frames, and these tags must be preserved during optimization.

To view VLANs supported by the WANJet appliance

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.
2. From the Monitoring menu, choose VLANs.
The VLANs Information report opens and displays the following information.
 - **Packets/Bytes**
Number of packets and total size in bytes of the network traffic exchanged with the VLAN.
 - **Aware**
Indicates whether the WANJet appliance can identify this virtual LAN.

◆ **Note**

*For information about configuring VLANs to work with the WANJet appliance, refer to **Managing virtual LANs**, on page 6-17.*

Connectivity

Connectivity diagnostic information includes the following reports:

- **All**
- **Ethernet**
- **IP**
- **Bridge**
- **Remote WANJet appliance**

All connectivity diagnostics

The Diagnose Connectivity report displays details about all types of connectivity (Ethernet, IP, bridge, and remote WANJet appliance).

To view diagnostics for all connectivity

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.
2. From the Connectivity menu, choose All.
The Diagnose Connectivity screen opens.

For details on the information included for each type of connectivity, refer to the following sections.

Ethernet diagnostics

The Diagnose Ethernet screen displays details about the Ethernet interfaces for the local WANJet appliance. For WANJet appliances to work correctly, the speed and duplex settings for the LAN and WAN interfaces should be the same. The Diagnose Ethernet screen confirms whether that is the case, and displays a warning if it is not.

◆ Note

*For information about configuring the speed and duplex settings for Ethernet interfaces, see **Changing the interface speed**, on page 6-22.*

To view diagnostics for Ethernet connectivity

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.
2. From the Connectivity menu, choose Ethernet.
The Diagnose Ethernet report opens and includes a section for each of the Ethernet interfaces on the WANJet appliance:
 - **eth2** (PEER)
 - **eth3** (ETH3)
 - **eth1** (WAN)
 - **eth0** (LAN)

The following information displays for each interface:

- **Speed**
- **Transmitted**
- **Received**
- **Receive errors**
- **Collisions**

Application QoS does not work unless the Ethernet interfaces are connected as follows:

- The **eth0** interface must connect to the LAN switch or router.
- The **eth1** interface must connect to the WAN gateway.

◆ **Note**

*If a redundant pair is present, the **eth2** interface (also labeled **Peer** on some WANJet appliances) must be connected to the redundant peer. For more information, see **Configuring redundant peers**, on page 6-26.*

IP diagnostics

The Diagnose IP screen displays technical details about the local WANJet appliance's IP configuration.

To view diagnostics for IP connectivity

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.
2. From the Connectivity menu, choose IP.
The Diagnose IP report opens and displays the following information:
 - The IP address of the local WANJet appliance.
 - The netmask of the local subnet.
This determines how much of the address identifies the subnetwork on which the WANJet appliance host resides, and how much identifies the host itself.
 - The IP address of the WAN gateway used by the local WANJet appliance.
 - The results of the local gateway ping.

◆ **Note**

*Addresses must adhere to the Internet Protocol standards. For more information about configuring addresses, see **Updating a configuration**, on page 6-14.*

Bridge diagnostics

The Diagnose Bridge screen displays details of the internal connectivity, or *bridge*, between Ethernet interfaces between the two WANJet appliances.

To view diagnostics for bridge connectivity

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.
2. From the Connectivity menu, choose Bridge.
The Diagnose Bridge report opens and displays the following information:
 - The IP address and MAC address of the WAN gateway that the local WANJet appliance uses
 - The Ethernet interfaces that are linked by the bridge

The WANJet appliance bridge does not work unless the Ethernet interfaces are connected as follows:

- The **eth0** interface must connect to the LAN switch or router.
- The **eth1** interface must connect to the WAN gateway.

Remote WANJet appliance diagnostics

The Diagnose Remote WANJet screen displays details about the remote WANJet appliances that are connected to the local WANJet appliance.

◆ Note

*For information about how to configure remote WANJet appliances, see **Managing remote WANJet appliances**, on page 6-20.*

To view diagnostics for remote WANJet appliance connectivity

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.
2. From the Connectivity menu, choose Remote WANJets.
The Diagnose Remote WANJet report opens and displays the following information for each remote WANJet appliance:
 - The software version number, which is compared to the local version number
 - The status of the local WANJet appliance
 - The number of remote WANJet appliances
 - The IP address for the remote WANJet appliance
 - The WANJet appliance type, which is Single if there is no redundant peer at the remote end
 - Whether the remote WANJet appliance is responding to pings from the local WANJet appliance

- Whether the local WANJet appliance can connect to the remote WANJet appliance on the ports that WANJet appliances use to communicate with each other. These ports are **3701**, **3702**, and **3703**, by default.

General

General diagnostic information includes the following three reports:

- **Bridge Forwarding Database**
- **Administration Tools**
- **Diagnostic Log**

To view general diagnostic information

From the General menu, choose the option that corresponds to the information that you want to view.

Bridge Forwarding Database diagnostics

The Bridge Forwarding Database report lists all of the network devices, by Media Access Control (MAC) Address, that have sent traffic through the local WANJet appliance bridge.

To view diagnostics for Bridge Forwarding Database

1. In the navigation pane, expand **Reports** and click **Diagnostics**. The initial diagnostics screen opens.
2. From the General menu, choose Bridge Forwarding Database. The Bridge Forwarding Database report opens and displays the following information for each network device:
 - **MAC Address**
A unique identifier attached to most networking devices, and used by many network protocols.
 - **IP Address**
Available only if the device has communicated directly with the WANJet.
 - **Interface**
The interface is defined as **eth0** if the device is connected to the local WANJet appliance through the LAN and as **eth1** if the device is connected through the WAN.
 - **Local**
This column displays **Yes** if the interface is one of the WANJet appliance's internal network interface cards.

Administration tools

The WANJet appliance provides a browser-based user interface for the following three network administration diagnostic tools:

- Ping
- Traceroute
- Packet Capture

To use the administration tools

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.
2. From the General menu, choose Administration Tools.
The Administration Tools screen opens. Each tool includes a box where you can specify command-line parameters, and a button to initiate the program.
3. Click the button for the tool that you want to run.
The lower half of the screen shows the command results:
 - The full path including parameters to the process, as it appears on the command line.
 - The process number allocated by the operating system. You can stop a process by clicking the process number before it has finished running.
 - The process output, which is similar to what you would see in the shell after running the program from the command line.
 - The return code, which is **0** if the program completes successfully.

Ping

The **ping** tool provides a simple test to confirm that a target host is online and reachable through a TCP/IP network. It works by sending ICMP request packets to the target and listening for response packets in return. The percentage of packets lost and the time taken to send and receive them indicate how well the connection is working.

◆ Note

If a ping is unable to reach a target host, such as when the statistical summary shows a 100% packet loss, it does not necessarily mean that there is no working network connection between source and target. For example, a firewall might be blocking ICMP requests from reaching the target host, but allowing some other network traffic through.

By default, WANJet appliance provides the following parameters for ping:

```
-R -c 5 -w 10 <IP address of target host>
```

The default target is the gateway machine for the subnet on which the WANJet appliance resides. You can change the parameters by typing new parameters in the associated text box.

◆ Important

F5 Networks recommends that only advanced users change parameters.

The WANJet appliance displays the following output for ping:

- The IP addresses of both the target host and the source host (the server on which ping is running)
- A line for each ICMP response packet received back from the target showing the packet's sequence number, time-to-live value, and round-trip time (request time + response time)
- A statistical summary showing:
 - The number of request packets transmitted
 - The number of response packets received back
 - The percentage of lost packets
 - The minimum, average, and maximum round-trip times

Traceroute

The **traceroute** tool plots the route that packets take to a target host. It can be helpful in determining the location of any network disruption.

Traceroute works by incrementing the time-to-live (TTL) value of successive packets sent out. TTL values are decremented as packets pass through intermediate hosts (known as hops). When the TTL reaches a value of **1**, a time exceeded message is sent back to the source host (the host on which traceroute is running). By examining the origins of these messages, you can reconstruct the path that packets take to the target host.

◆ Note

*Traceroute sends out UDP datagram packets by default. If UDP probes are being blocked by a firewall, you can use ICMP echo requests instead (as ping does) by specifying the **-I** option. Packets are normally sent to port **33434**, which should not be in use. If the target host is listening on port **33434**, you can specify a different port using the **-p** option.*

By default, the WANJet appliance provides the following parameters for traceroute:

```
-v <IP address of target host> -c 10 (not port 10000)
```

As with the ping tool, the default target is the gateway for the local subnet. You can change the parameters by typing new parameters in the associated text box.

◆ **Important**

F5 Networks recommends that only advanced users change parameters.

The WANJet appliance displays the following output for traceroute:

- The IP address of the target host, the maximum number of hops (that is, the maximum TTL), and the size of the packets sent.
- A list of hosts through which packets are passing together with the round-trip time taken for each of the three packets (packets are sent out in threes, by default) to travel from the source host, to the intermediate host, and back again.

Packet Capture

You can use the **tcpdump** utility to intercept and display the contents of TCP/IP packets on the network. This is useful for debugging your network configuration, because it allows you to isolate the source of a problem by determining if all routing is working correctly. The utility saves the data in a PCAP file.

◆ **Note**

*You need a specialized application, such as **Ethereal** (a network protocol analyzer that runs on both Linux and Windows) to read PCAP files produced by **tcpdump**. You can download **Ethereal** and its documentation for free from <http://www.ethereal.com/>.*

By default, the WANJet appliance provides the following parameters for **tcpdump**:

-c 10 (not port 1000)

Packets sent to port **10000** are ignored, since this is the port that the Web UI uses to communicate with the local WANJet appliance. You can change the parameters by typing new parameters in the associated text box.

◆ **Important**

F5 Networks recommends that only advanced users change parameters.

When the **tcpdump** process has finished running, the Tools screen displays a link to the PCAP file that is produced. If you have an application that can read PCAP files, you can open the PCAP file directly, or you can save the file to disk. The PCAP file is also stored on the server where **tcpdump** is running, at the following location:

/usr/local/NetOptimizer/logs/dump.pcap

Diagnostic Log

The Diagnostic Log contains status information and errors that the WANJet appliance records during a session. This log keeps you informed and helps resolve problems that you might encounter while working with the WANJet appliance. You can clear the data in the Diagnostic Log at any time.

To view the Diagnostic Log

1. In the navigation pane, expand **Reports** and click **Diagnostics**.
The initial diagnostics screen opens.
2. From the General menu, choose Diagnostic Log.
The Diagnostic Log opens.

To clear the Diagnostic Log

1. At the top of the Diagnostic Log screen, click the **Clear Logs** button.
A warning message displays to inform you that this action will delete all data in the error and report logs.
2. Click **OK** to delete the logs.

System Snapshots

You can create a system snapshot and download it as a zipped text file to your hard disk. You can provide this zipped text file to the F5 Networks Technical Support team to help resolve technical issues.

To create a system snapshot

1. At the top the Diagnostic Log screen, next to **System Snapshot**, click **now** to get the current system status.
The browser opens a download window for you to save the snapshot file to your local disk.

*Note: Next to **System Snapshot**, you can also click **historic** to get a tar file containing previously taken snapshots.*

2. Save the snapshot file. The system snapshot file is named **snapshot.txt.gz**. This is a compressed plain text file.

*Note: To view the snapshot file, you first need to extract it using a tool such as **gunzip**, which is available at www.gzip.org.*

3. Rename the compressed file using the following format:

snapshot-<yourcompanyname-yyyy-mm-dd>

For example:

snapshot-acme-2005-04-22

You can provide this file to F5 Networks Technical Support for assistance with troubleshooting issues.

Third-party reporting systems

You can configure the WANJet appliance to work with several third-party reporting systems, including Syslog, SNMP, and RMON2.

Syslog reports

With the WANJet appliance, you can view syslog reports from an external syslog server. These reports include data, such as the amount of sent and received data that the WANJet appliance has processed.

◆ **Note**

*You must configure the IP address of the syslog server in the **Syslog Server IP** box of the **Syslog and SNMP** screen, to view syslog data. For more information, see **Configuring Syslog and SNMP settings**, on page 6-24.*

SNMP reports

With the WANJet appliance, you can use an external computer as a management station for viewing Simple Network Management Protocol (SNMP) logs that the WANJet appliance produces on the local appliance. The SNMP data trees reside in a Management Information Base (MIB).

The SNMP data on the WANJet appliance includes information about the network cards, the total bandwidth saved for sent and received data, and the amount of sent and received data that was optimized.

For the private MIB file, see Appendix B, *WANJet Appliance Private MIB File*.

To configure the WANJet appliance to use an SNMP server

1. In the navigation pane, expand **Configuration** and click **Monitoring**.
The WANJet Syslog and SNMP screen opens.
2. For **SNMP Server IP**, select the check box, then type the IP address of the SNMP server.
For additional instructions, see *Configuring Syslog and SNMP settings*, on page 6-24.
3. For **SNMP IP**, select which IP address you want SNMP to use as the source address in response to SNMPv1 GET requests, and for sending SNMP traps. The choices are **Management IP** (set by default) or **WANJet IP**.
4. For **Community String**, type the shared community string used on both the SNMP server and the WANJet appliance.
5. Click **Save** to save your settings.

6. In the navigation pane, expand **Security** and click **IP Access Control**.
The IP Access Control screen opens.
7. Verify that the SNMP server can access the WANJet appliance. The default setting is **Allow all addresses**, but an administrator may have changed it to allow or deny listed addresses. If so, you must make sure that the SNMP server is listed (if allowing listed addresses) or not listed (if denying listed addresses).

For additional instructions, see *Granting Web UI access*, on page 5-4.
8. Click **Save** if you change any information on the WANJet IP Access Control screen.

To view SNMP tables

SNMP data is stored in tables in the MIB. To view the SNMP tables, you need to use SNMP-compliant software. In the SNMP-compliant software, you must configure the IP address of the WANJet appliance and the community string that you specified on the WANJet Syslog and SNMP screen.

◆ Note

*For a list of WANJet appliance SNMP errors and descriptions, see Appendix A, **WANJet Appliance Messages**.*

RMON2 reports

The WANJet appliance provides monitoring of network and application performance through RMON2 (Remote Monitoring). You can configure the WANJet appliance so that you can view RMON2 reports (data trees), which are part of the SNMP data trees that are stored on the WANJet appliance. The RMON2 data is stored in a MIB. For information on the WANJet MIB file, see Appendix B, *WANJet Appliance Private MIB File*.

The RMON2 data on the WANJet appliance includes data sent and received between two nodes, the IP addresses of these nodes, the port used to send and receive data, data size before and after the WANJet appliance processes it, times at which data was sent, and the numbers of connections. You can configure the WANJet appliance to report these statistics as **Raw Data** (prior to being optimized) or **WANJet Data** (optimized).

To enable RMON2 Logs

1. In the navigation pane, expand **Configuration** and click **Monitoring**.
The initial WANJet Syslog and SNMP screen opens.
2. Check the **Enable RMON2 Logs** box.
3. Click either **Raw Data** or **WANJet Data** to control the type of data that the RMON2 tables will export: **Raw Data** exports unoptimized data and **WANJet Data** exports optimized data. For more details on these options, see *RMON2 configuration settings*, following.
4. For **Community String**, type the shared community string that enables the SNMP server to access RMON2 data on the WANJet appliance.
5. Click the **Save** button.

You access RMON2 data the same way that you access SNMP data. Before accessing RMON2 data, you must specify a community string and IP address for the SNMP server as discussed in the previous section for SNMP reports. Set the RMON2 preferences on the Syslog and SNMP screen.

For additional instructions, see *Configuring Syslog and SNMP settings*, on page 6-24. Note that the SNMP server must have access to the WANJet, as described in *Granting Web UI access*, on page 5-4.

To view RMON2 reports

To view the RMON2 data tree, you must use SNMP-compliant software. You need to provide SNMP-compliant software with the IP address of the WANJet appliance and the community string that you specified on the WANJet Syslog and SNMP screen.

RMON2 configuration settings

RMON2 provides network and application protocol statistics (bytes transmitted and bytes received) that include both the unoptimized protocol statistics from the LAN side, and the optimized (compressed) protocol statistics from the WAN side. When enabling RMON2 logs as described in the previous section, you can configure the WANJet appliance to report this information as:

- **Raw Data**
- **WANJet Data**

Which setting you choose depends on your RMON2-based monitoring software. Some RMON2 probes recognize only a fixed list of standard protocols, and do not recognize F5 Networks-added table entries. Thus, the software can monitor only one of the two types of data.

If your software recognizes a fixed list of entries, choose the setting that reports the type of data you require. If you want to monitor the unoptimized data, enable the **Raw Data** setting. If you want to monitor the optimized data, enable the **WANJet Data** setting.

If your RMON2-based monitoring software recognizes the F5 Networks-added table entries, the choice of setting is not as important. Select the setting that allows you to use both types of data in the way best suited to the configuration of your network management software.

The following sections provide further technical details concerning the implications of selecting either option.

Raw data

By selecting the **Raw Data** setting on the WANJet Syslog and SNMP screen, you instruct the WANJet appliance to augment the set of standard protocols reported through RMON2 with a second set of standard protocols, representing those protocols after optimization. For any table in the RMON2 MIB that contains network and application protocol statistics, the standard protocols contain the unoptimized data byte counts from the LAN, and each protocol is paired with a new instance that contains the optimized data byte counts from the WAN.

For example, the **protocolDirTable** object (OID 1.3.6.1.2.1.16.11.2 in the RMON2 MIB) includes the following entries:

Standard RMON2 protocols (LAN/unoptimized)

any.ip
any.ip.udp
any.ip.tcp
any.ip.tcp.22

F5 Networks-created pairing (WAN/optimized)

any.IPcompressed
any.IPcompressed.udp
any.IPcompressed.tcp
any.IPcompressed.tcp.22

In this example, the first three protocols in both lists are permanent entries. The fourth protocol in both shows a protocol that is added to the tables at runtime. The fourth, **any.ip.tcp.22**, contains statistics for the Secure Shell (SSH) protocol.

WANJet appliance data

By selecting the **WANJet Data** setting on the WANJet Syslog and SNMP screen, you instruct the WANJet appliance to place optimized WAN data byte counts into the standard protocols, and insert new instances to contain the unoptimized data byte counts from the LAN.

For example, the **protocolDirTable** object in this configuration might have the following entries:

F5 Networks-created pairing (LAN/unoptimized)

```
any.IPuncompressed
any.IPuncompressed.udp
any.IPuncompressed.tcp
any.IPuncompressed.tcp.22
```

Standard RMON2 protocols (WAN/optimized)

```
any.ip
any.ip.udp
any.ip.tcp
any.ip.tcp.2
```



9

Configuration Examples

- Basic configuration
- Mesh configuration
- Hub and spoke configuration
- Redundant system configuration
- LAN router configuration

Basic configuration

Figure 9.1 shows an example of a basic configuration of WANJet appliances.




Figure 9.1 Basic WANJet appliance configuration

In this example:

- This configuration includes two WANJet appliances.
- The network includes two LANs (LAN1 and LAN2) that are connected.
- LAN2 is a remote network of LAN1, and LAN1 is the remote network of LAN2.
- LAN1 includes WANJet1, and LAN2 includes WANJet2.
- Configure WANJet2 as a remote WANJet appliance on WANJet1, and configure WANJet1 as a remote WANJet appliance on WANJet2.
- WANJet1 sends optimized data to WANJet2, while WANJet2 sends optimized data to WANJet1.

	WANJet1	WANJet2
IP Address	192.168.150.100	192.168.100.100
Local Network	192.168.150.0/24	192.168.100.0/24
WAN Gateway	192.168.150.2	192.168.100.2
Remote WANJet appliance	192.168.100.100	192.168.150.100

Table 9.1 Basic configuration specifications

Mesh configuration

Figure 9.2 shows an example of a mesh configuration including three WANJet appliances.

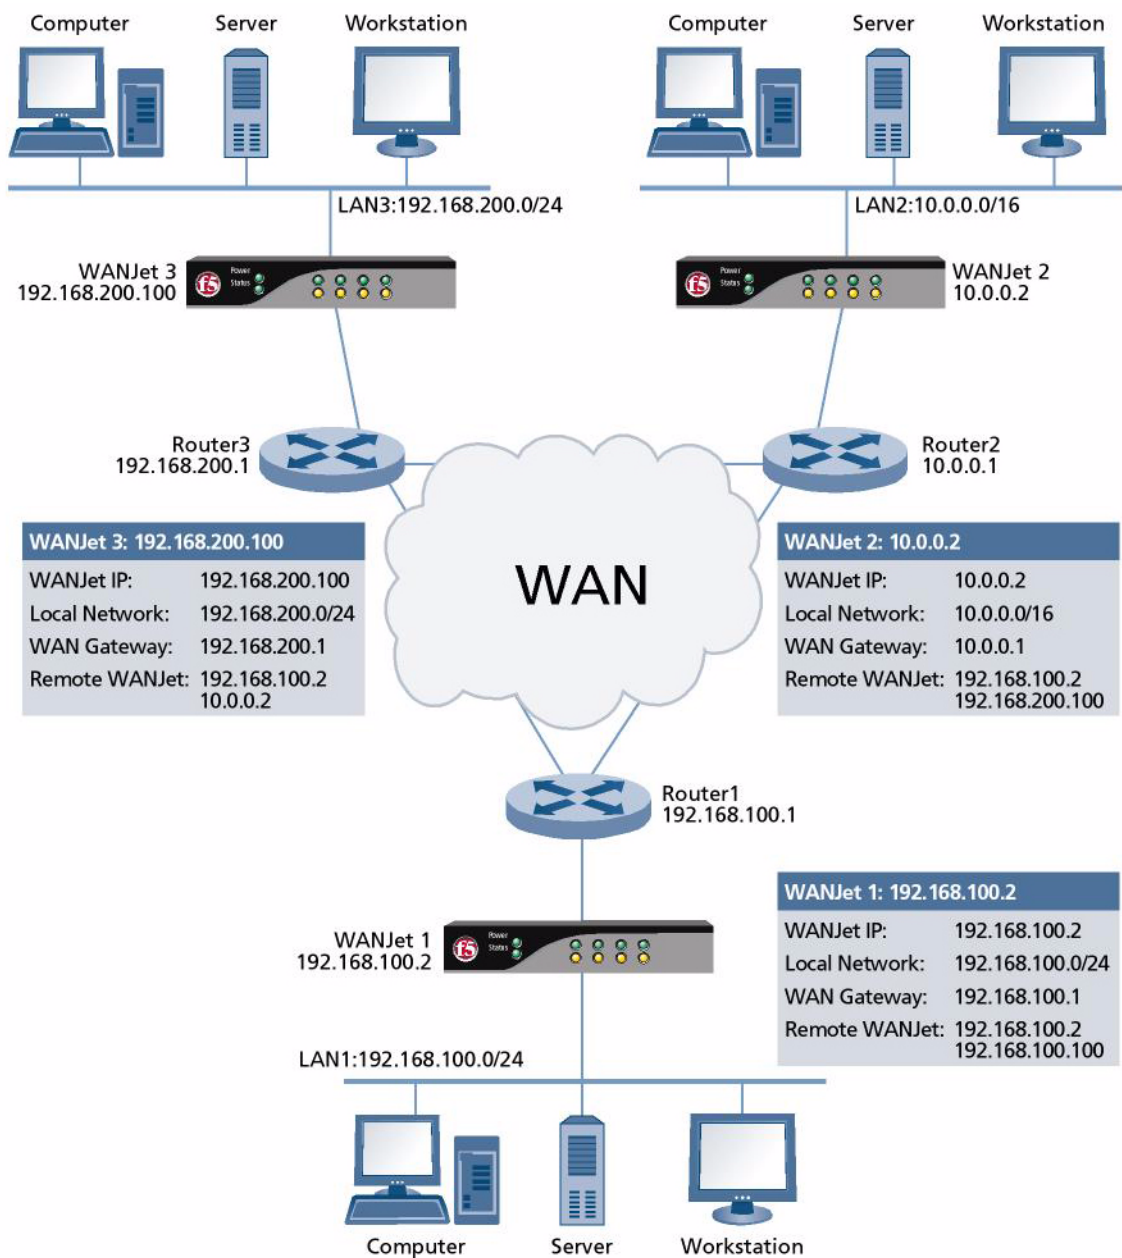


Figure 9.2 Mesh configuration

In this example:

- This configuration includes three WANJet appliances.
- The network includes three LANs (LAN1, LAN2, and LAN3) that are connected.
- LAN1 includes WANJet1, LAN2 includes WANJet2, and LAN3 includes WANJet3.
- Configure WANJet2 and WANJet3 as remote WANJet appliances on WANJet1; configure WANJet1 and WANJet3 as remote WANJet appliances on WANJet2; and configure WANJet1 and WANJet2 as remote WANJet appliances on WANJet3.
- WANJet1 sends optimized data to WANJet2 and WANJet3; WANJet2 sends optimized data to WANJet1 and WANJet3; and WANJet3 sends optimized data to WANJet1 and WANJet2.

	WANJet1	WANJet2	WANJet3
IP Address	192.168.100.2	10.0.0.2	192.168.200.100
Local Network	192.168.100.0/24	10.0.0.0/16	192.168.200.0/24
WAN Gateway	192.168.100.1	10.0.0.1	192.168.200.1
Remote WANJet appliance 1	10.0.0.2	192.168.200.100	192.168.100.2
Remote WANJet appliance 2	192.168.200.100	192.168.100.2	10.0.0.2

Table 9.2 Example mesh configuration specifications

Hub and spoke configuration

Figure 9.3 shows an example of a hub and spoke configuration.

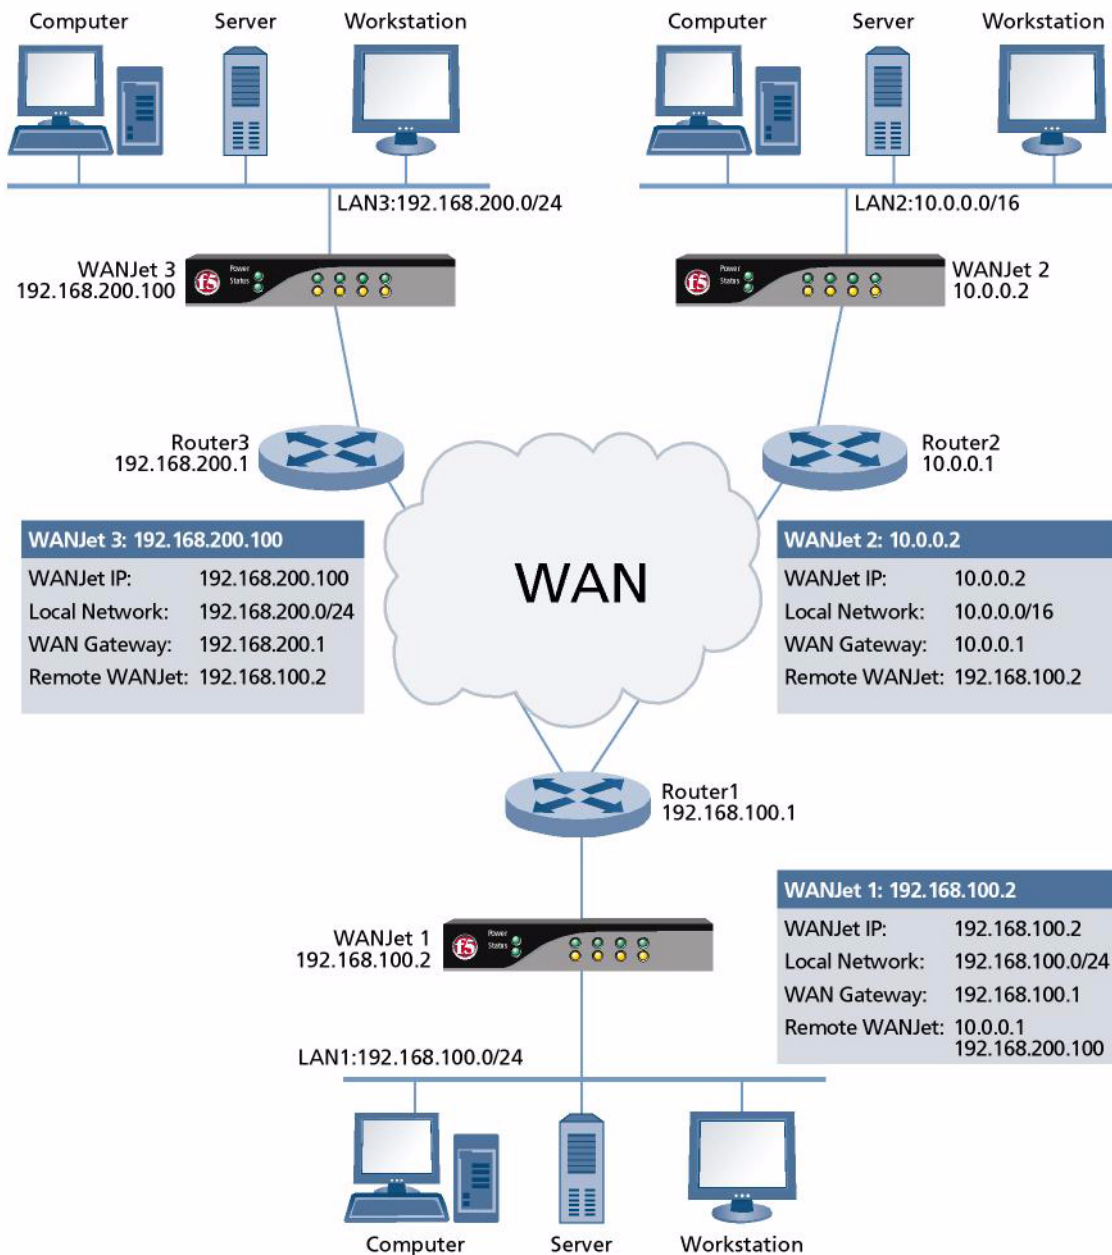


Figure 9.3 Hub and spoke configuration

In this example:

- This configuration includes three WANJet appliances.
- The network includes three LANs (LAN1, LAN2, and LAN3) that are connected.
- LAN1 connects to LAN2 and LAN3, and LAN2 and LAN3 connect to LAN1, but LAN2 and LAN3 do not connect to each other.
- Configure WANJet2 and WANJet3 as remote WANJet appliances on WANJet1; configure WANJet1 as a remote WANJet appliance on both WANJet2 and WANJet3.
- LAN1 includes WANJet1, LAN2 includes WANJet2, and LAN3 includes WANJet3.
- WANJet1 sends optimized data to both WANJet2 and WANJet3, WANJet2 sends optimized data to WANJet1 only, and WANJet3 sends optimized data to WANJet1 only.

	WANJet1	WANJet2	WANJet3
IP Address	192.168.100.2	10.0.0.2	192.168.200.100
Local Network	192.168.100.0/24	10.0.0.0/16	192.168.200.0/24
WAN Gateway	192.168.100.1	10.0.0.1	192.168.200.1
Remote WANJet appliance 1	10.0.0.2	192.168.100.2	192.168.100.2
Remote WANJet appliance 2	192.168.200.100		

Table 9.3 Example hub and spoke configuration specifications

Redundant system configuration

Figure 9.4 shows an example of a redundant system configuration.




Figure 9.4 Redundant system configuration

In this example:

- Two LANS are connected, and one of the LANs has WANJet appliances installed in a redundant system configuration.
- LAN1 includes two WANJet appliances, WANJet1A and WANJet1B, and LAN2 includes WANJet2. WANJet1B is the redundant peer of WANJet1A. In case of router or WANJet appliance failure, the other router and its corresponding WANJet appliance resume function.
- WANJet1B processes the half of the data in LAN1, and WANJet1A processes the other half of the data in LAN1 (load balancing).
- WANJet1A sends optimized data to WANJet2, and WANJet1B sends optimized data to WANJet2.

	WANJet1A	WANJet1B	WANJet2
IP Address	10.55.55.3	10.55.55.4	192.168.200.100
Local Network	10.55.55.0/24	10.55.55.0/24	192.168.200.0/24
Gateway	10.55.55.1	10.55.55.2	192.168.200.1
Remote WANJet appliance	192.168.200.100	192.168.200.100	10.55.55.3
Subnet			10.55.55.0/24
Remote WANJet appliance	192.168.200.100	192.168.200.100	10.55.55.4

Table 9.4 Example redundant configuration specifications

LAN router configuration

Figure 9.5 shows an example of a LAN router configuration.




Figure 9.5 LAN router configuration

In this example:

- A LAN router connects two networks to a WANJet appliance, and the WANJet appliance connects to the outside WAN through another router (WAN gateway).
- LAN1 includes WANJet1, and LAN2 includes WANJet2.
- LAN1 comprises three networks VLAN 100, VLAN 200, and WANJet1's network (**192.168.1.0/24**). The LAN router connects all three networks. You configure the LAN router on WANJet1.
- You configure VLAN 100 and VLAN 200 as local subnets on WANJet1.
- LAN1 and WANJet1 connect to the WAN through the WAN gateway.

	WANJet1	WANJet2
IP Address	192.168.1.100	10.10.20.100
Local Network	192.168.1.0/24	10.10.20.0/24
Local Subnets	VLAN 100: 192.168.100.0/24 VLAN 200: 192.168.200.0/24	10.10.20.0/24
LAN Router	192.168.1.1	
WAN Gateway	192.168.1.2	10.10.20.1

Table 9.5 LAN router configuration specifications



10

Working from the Command Line

- Command line overview
- Attaching a computer to the WANJet appliance
- Getting to the command line
- Using commands
- Replacing the console cable

Command line overview

You can use the command line to configure the WANJet appliance and perform several commands. While it is not usually necessary to use the command line interface, some administrators may prefer to use the restricted shell in some cases.

For example, you may find that it is easier to perform initial configuration on the WANJet appliance at the command line than it is using the LCD included on some platforms. You can configure the WANJet appliance IP address, subnet mask, gateway, and so on using the **config** command.

All of the commands available at the command line are also available through the Web UI, after you have completed the initial configuration. Refer to *Using commands*, on page 10-4, for a list of all available commands including where to find them in the Web UI.

If you are having problems with the WANJet appliance (such as getting to the Web UI), you can reboot the appliance. From the command line on the console, you can attempt to boot from the alternate image (rather than the default), bringing up the Web UI and reviewing the WANJet appliance status.

In addition, if you do not have physical access to the WANJet appliance, you can use an SSH program to connect to the command line interface. You can perform initial configuration, as described in *Configuring the WANJet appliance from the command line*, on page 10-5. You can also use SSH for troubleshooting the WANJet appliance. For example, if your corporate firewall is blocking port **10000**, you can log on to the command line using the console to troubleshoot that fact. (Port **10000** is the port that the Web UI uses to connect to the WANJet appliance.)

Attaching a computer to the WANJet appliance

You need to attach a computer to the WANJet appliance to access the command line through a serial console. You use a terminal emulator program, such as HyperTerminal, to connect to the WANJet appliance's command line interface.

Different WANJet platforms require different console cables to connect the computer. Refer to *Replacing the console cable*, on page 10-9, for more information on the cable provided for each platform.

To connect a computer to a WANJet appliance

1. On a computer or laptop, plug the console cable into the serial port on the computer and tighten the cable screws.
2. Plug the other end of the cable into the port labeled **Console** on the WANJet appliance. The console port is an RJ-45 port on the front of the WANJet 400 and 500, and a serial port on the back of the WANJet 200.

Getting to the command line

You can access the command line in the following ways:

- Using a terminal emulator on a serial console (local only)
- Using an SSH program (local or remote)

You can use a terminal emulator program to access the command line if working locally using a serial console connected to the WANJet appliance.

If you are using a computer running Microsoft® Windows®, you can use HyperTerminal™ or other terminal emulator program. You can also use any SSH program, such as PuTTY, AlphaCom, or SecureCRT® to get to the command line if working locally or remotely.

The following procedure uses HyperTerminal; other terminal emulators require similar setup steps. The procedure *To use PuTTY to get to the command line*, on page 10-3, describes how to use an SSH client to get to the command line.

To use HyperTerminal to get to the command line

1. On the connected computer, from the Start menu, choose All Programs, then Accessories, then Communications, and then click HyperTerminal.
The Connection Description popup screen opens.
2. In the **Name** box, type a name for the connection, then click **OK**.
The Connect To popup screen opens.

3. From the **Connect using** list, select the port where the WANJet appliance is connected to the computer.
The Properties screen for the port opens.
4. From the **Bits per second** list, select **9600**.
5. From the **Data bits** list, select **8**.
6. From the **Parity** list, select **None**.
7. From the **Stop bits** list, select **1**.
8. From the **Flow control** list, select **Xon/Xoff**.
9. Click **OK**.
The connection session opens.
If the WANJet appliance login prompt does not appear, press **Enter** once or twice.
The login prompt appears in the console window.
10. At the **WANJet login** prompt, type the user name you use to log on to the WANJet appliance.
11. At the **Password** prompt, type the password.
The Welcome message is displayed, as shown in Figure 10.1.

```

Welcome to F5 WANJet
The following commands are available:
  help - this help message
  config - set WANJet's network configuration and password
  ping [ip address] - the classic ping tool
  traceroute [ip address] - map the path to an ip address
  diagnose-connectivity - help diagnose connectivity issues
  reboot - reboot this WANJet
  halt - halt this WANJet
  exit - Disconnect from this session
WANJet Version: x.x.x build xxxxx

```

Figure 10.1 Welcome message using HyperTerminal

To use PuTTY to get to the command line

1. From any computer on a network that can connect to the WANJet appliance, start PuTTY.
The PuTTY Configuration popup screen opens.
2. In the **Host Name (or IP address)** box, type the IP address of the WANJet appliance.
3. For the **Protocol**, select **SSH**.
4. Click **Open**.
5. When connecting to a WANJet appliance for the first time, a message asking you to check the host key appears. To continue, you must click **Yes** to indicate that you trust the host.
A console window opens.
6. At the **login as** prompt, type the administrator user name (default is **admin**).

7. At the **admin@<IP address> password** prompt, type the administrator password (default is **admin**).
The Welcome message is displayed, as shown in Figure 10.2.

```

Welcome to F5 WANJet
The following commands are available:
  help - this help message
  config - set WANJet's network configuration and password
  ping [ip address] - the classic ping tool
  traceroute [ip address] - map the path to an ip address
  diagnose-connectivity - help diagnose connectivity issues
  reboot - reboot this WANJet
  halt - halt this WANJet
  exit - Disconnect from this session
WANJet Version: x.x.x build xxxxx

```

Figure 10.2 Welcome message using PuTTY

To log off the command line

When you are done working at the command line, you can log off by typing **Exit** or **Logout**.

Using commands

After you have logged on to the command line through your terminal emulator or SSH client, you can use the commands listed in Table 10.1. All of the commands listed in the table are also available through the Web UI.

Command	Description
config	<p>Configure network settings for the WANJet appliance including setting the IP address, NetMask, Gateway, Media Type for eth0 (LAN) and eth1 (WAN), the Management IP address, Management Netmask, Management Gateway, and the IP address of the only computer that can access the WANJet appliance.</p> <p>In the Web UI, you can configure network settings through the Configuration section of the navigation pane.</p>
diagnose-connectivity	<p>Investigate connectivity issues and provide information including Ethernet configuration, LAN and WAN duplex and speed settings, WANJet IP configuration, software version, state, and remote WANJet appliances.</p> <p>In the Web UI, you can display the command output in the Diagnostic reports; on the Connectivity menu, choose All.</p>

Table 10.1 WANJet appliance commands

Command	Description
exit	Log off the WANJet appliance. In the Web UI, click Logout admin to log off the WANJet appliance.
halt	Shut down the WANJet appliance. In the Web UI, you can shut down the WANJet appliance in the System section of the navigation pane. On the Shutdown & Restart screen, click Shutdown WANJet .
help	List available commands and a brief description of each. In the Web UI, you can get help by clicking User Manual .
logout	Log off the WANJet appliance. In the Web UI, click Logout admin to log off the WANJet appliance.
ping <IP address>	Test whether an IP address is reachable across an IP network. In the Web UI, you can display the command output in the Diagnostic reports; on the General menu, choose Administration Tools.
reboot	Restart the WANJet appliance. In the Web UI, you can restart the WANJet appliance in the System section of the navigation pane. On the Shutdown & Restart screen, click Restart WANJet .
traceroute <IP address>	Determine the route that packets take across an IP network to the specified IP address. In the Web UI, you can display the command output in the Diagnostic reports; on the General menu, choose Administration Tools.

Table 10.1 WANJet appliance commands (Continued)

The following procedures describe how to use the commands after having logged on to the WANJet appliance in the terminal emulator.

Configuring the WANJet appliance from the command line

You can perform initial configuration of the WANJet appliance from the command line. On platforms that include a liquid crystal display, or LCD, you can configure the addresses there instead. The *Quick Start Card* included in the shipping box with your WANJet appliance describes the initial hardware installation and setup instructions using the LCD.

To configure the WANJet appliance from the command line

1. In the terminal emulator, type **config**.
You are prompted for new configuration values. Press Enter to keep the current value.
2. At the **Set WANJet IP** prompt, type the IP address of the WANJet appliance (default is **192.168.168.100**) and press Enter. This is the IP address you will use to log on to the Web UI (unless you want to use out-of-band management as described in step 6).
3. At the **Set WANJet NetMask** prompt, type the netmask you want to assign to the WANJet appliance (default is **255.255.255.0**) and press Enter.
4. At the **Set WANJet Gateway** prompt, type the IP address of the gateway (default is **192.168.168.001**) and press Enter.
The media types are displayed.

Type	0	for	Auto Sense
Type	100F	for	100BaseTX Full-Duplex
Type	100H	for	100BaseTX Half-Duplex
Type	10F	for	10BaseT Full-Duplex
Type	10H	for	10BaseT Half-Duplex

5. To set the media type, type the appropriate alphanumeric code that represents the required speed and duplex setting for the **eth0** (LAN) and **eth1** (WAN) media types, then press Enter. You can use the default value of **Auto Sense** (called **Auto Negotiate** in the Web UI) unless you are familiar with the settings required in your network environment and want to change them.
6. If you want to manage the WANJet platform from a separate management subnet, you can optionally set the IP addresses for the **WANJet Management IP**, **WANJet Management NetMask**, and **WANJet Management Gateway**. (These options are blank, by default.) If you set the Management IP address, this is the address you use to log on to the Web UI (instead of the WANJet IP).
7. Determine how you want to allow access to the Web UI.
 - To allow access from one IP address, at **Set the IP of the first machine allowed to access WANJet UI**, type the IP address of the only computer that you want to access the Web UI.
 - To allow access from multiple IP addresses, leave this blank (the default). In the navigation pane, expand **Security** and click **IP Access Control**. See *Granting Web UI access*, on page 5-4.
 - To allow access from any browser on any computer in your network (using the correct user name and password), leave this field blank (the default).

When you complete this step, the command line prompt returns and the new configuration information is applied to the WANJet appliance immediately. Values that you changed on the command line are also changed in the Web UI.

Running diagnostic tools from the command line

You can use the **ping** command to test whether a destination address is reachable across an IP network. For example, if you are having trouble connecting to a second WANJet appliance, you could ping the gateway IP address or the IP address of the other WANJet appliance.

To see the complete syntax of the **ping** or **traceroute** commands, type the command with no arguments.

To use ping from the command line

In the terminal emulator, type **ping <destination>**, where **<destination>** is the IP address or host name of the computer you are trying to reach. **Ping** output shows an overview of the command, followed by a list of responses received. It continues to show responses until you stop the command by typing Ctrl + C.

At the end of the command output, you can see **ping** statistics including the number of packets transmitted and received between the WANJet appliance on which you are working and the destination computer, the percentage of packet loss, the total time it took, and the round trip time between the two computers.

For example, the following command checks to see whether or not the computer at the IP address **192.168.72.254** is responsive. The **-c 5** option specifies the count and stops sending requests after five replies are received (or when the deadline is reached). The **-w 10** option indicates a deadline of 10 seconds.

```
ping -c 5 -w 10 192.168.72.254
```

Type **ping** with no arguments to see the complete syntax of the command. Refer to documentation on the Linux **ping** command for details on using all of the options. See *Administration Tools*, on page 8-24, for how to run the command in the Web UI.

To use traceroute from the command line

In the terminal emulator, type **traceroute <host>**, where **<host>** is the IP address or host name at the end of the route you are tracing. **Traceroute** output prints the route that packets take to the host, specifying the IP addresses of all of the computers (hops) in between the WANJet appliance and the host computer.

For example, the following command traces the route from the WANJet appliance you are working on to the computer at IP address **10.1.102.204**. In the output, you see a summary of the command, then notice that it takes three hops to get to the IP address, as shown in Figure 10.3.

```
traceroute -v 10.1.102.204
traceroute -v 10.1.102.204
traceroute to 10.1.102.204 (10.1.102.204), 30 hops max, 38 byte
packets
 1 10.1.105.1 (10.1.105.1) 36 bytes to 10.1.105.204 67.116 ms
   60.849 ms 60.594 ms
 2 10.1.101.251 (10.1.101.251) 36 bytes to 10.1.105.204 60.332 ms
   60.360 ms 60.024 ms
 3 10.1.102.204 (10.1.102.204) 36 bytes to 10.1.105.204 60.924 ms
   60.721 ms 60.035 ms

[Finished (returned 0)]
```

Figure 10.3 Traceroute command example

Type **traceroute** with no arguments to see the complete syntax of the command. Refer to documentation on the Linux **traceroute** command for details on using all of the options. See *Administration Tools*, on page 8-24, for how to run the command in the Web UI.

To use diagnose-connectivity from the command line

In the terminal emulator, type **diagnose-connectivity** to investigate connectivity issues and display the following information:

- Ethernet configuration including details on all four interfaces
- LAN and WAN duplex and speed settings
- WANJet appliance IP configuration
- WANJet appliance software version and build number
- WANJet appliance state (Active or Inactive)
- Remote WANJet appliances

Refer to *All connectivity diagnostics*, on page 8-20, for information on running a report from the Web UI that includes this information.

Shutting down or restarting the WANJet appliance from the command line

You can shut down the WANJet appliance or restart it from the command line. Refer to *Shutting down and restarting the WANJet appliance*, on page 5-8, for details on performing these tasks from the Web UI.

To shut down the WANJet appliance from the command line

Type **halt**.

The WANJet appliance initiates its shutdown sequence. On some platforms, the appliance powers off when the sequence completes. On others, you may need to press the power switch to turn the power off after shutting it down.

To restart the WANJet appliance from the command line

Type **reboot**.

The WANJet appliance stops, then restarts automatically. You need to log on again to access the command line.

Replacing the console cable

The WANJet appliance includes a console cable for attaching a console to the appliance. If this cable is lost, or becomes damaged, you may need to create a new cable. This section describes the included cable, which you can use as a model to create or purchase a new cable.

The WANJet 200 uses a DB-9 null-modem female-to-female serial cable. This type of cable can be readily purchased, if needed.

Figure 10.4 illustrates the DB9 to RJ45 cable connector provided on the WANJet 400 and 500 platforms.




Figure 10.4 DB9 to RJ45 cable diagram (WANJet 400 and 500)

The numbers indicated in Figure 10.4 correspond to the mappings specified in Table 10.2.

DB9 pin	Name	RJ45 pin
1	Carrier Detect (CD)	NC
2	Receive Data (RXD)	3
3	Transmit Data (TXD)	6
4	Data Terminal Ready (DTR)	7
5	GND	5
6	Data Set Ready (DSR)	4
7	Request to Send (RTS)	8
8	Clear to Send (CTS)	2
9	Ring Indicator (RI)	NC

Table 10.2 Pinouts for DB9 to RJ45 cable (WANJet 400 and 500)



A

WANJet Appliance Messages

- WANJet appliance messages and codes

WANJet appliance messages and codes

Table A.1 lists information about messages, such as errors, warnings or diagnostic log messages, that may occur on the WANJet appliance. If your company uses a centralized message server, you can configure the WANJet appliance to send messages to an associated SNMP server or Syslog server. Refer to *Configuring Syslog and SNMP settings*, on page 6-24, for more information.

You can also view messages that have occurred on your WANJet appliance from the Web UI. In the navigation pane, expand **Reports** and click **Diagnostics**; then from the General menu, choose Diagnostic Log.

Code	Message and Description	WANJet Appliance Component
1000 to 1002	Error: Configuration error	Optimization Engine
1003 to 1005	Error: Initialization error	Optimization Engine
1006 to 1007	Error: Internal error Description: Internal errors related to TCP compression.	Optimization Engine
1100 to 1103	Error: Internal error	Packet Processor
1150	Maximum number of optimized connections reached	Packet Processor
1151	Connection reset: Source IP:Port <IP address:port #> Destination IP:Port <IP address:port #> Description: The WANJet appliance reset the connection because Connection Intercept is turned on.	Packet Processor
1200 to 1201	Error: Configuration error	Optimization
1202 to 1203	Error: Initialization error	Optimization
1204 to 1207	Error: Internal error	Optimization
1209	Link down with <Proxy IP address> Description: The local WANJet appliance has no communication with the remote WANJet appliance.	Optimization
1210	Link up with <Proxy IP address> Description: The WANJet appliance resumed communication with the remote WANJet appliance.	Optimization

Table A.1 WANJet appliance messages

Code	Message and Description	WANJet Appliance Component
1211	Authentication failed with <Proxy IP address>	Optimization
1212	Error: Connection from unauthorized proxy <Proxy IP address>	Optimization
1213	Error: Internal error	Optimization
1214	Error: The version <version #> is incompatible with <Proxy IP address> version <version #>	Optimization
1215	Error: License expired on <date>	Optimization
1216	Error connecting to remote client <Proxy IP> from <Initiator_IP> Description: The remote WANJet appliance was unable to connect to a server from the initiator IP address.	Optimization
1217	Error connecting to local client <Proxy IP> error (<error spec>) Description: where <error spec> is a standard socket error. The local WANJet appliance was unable to connect to a server.	Optimization
1218	Error: Diagnostic log test error Description: For internal use only.	Optimization
1219	Connection (<server IP> <-> <client IP>) not optimized: SMB signing required (code 1219) Description: The connection requires SMB signing so it was not optimized.	Optimization
1221	GP Session Up Description: A diagnostic log message that occurs when the generic proxy session starts.	Optimization
1222	GP Session Down Description: A diagnostic log message that occurs when the generic proxy session fails.	Optimization
1250	Version (<version #>) up and running	Optimization
1251	Error: Internal error	Optimization
1252	Warning: License Limit Exceeded	Optimization

Table A.1 WANJet appliance messages (Continued)

Code	Message and Description	WANJet Appliance Component
1253	Warning: Invalid license key - Bandwidth optimization off	Optimization
1254	Warning: License key not entered - Bandwidth optimization off or Warning: License expired - Bandwidth optimization off or Warning: License invalid <reason> - Bandwidth optimization off	Optimization
1255	Warning: <#> days remain for evaluation license to expire	Optimization
1256	Warning: WANJet is activated for evaluation for <#> days	Optimization
1257	Warning: Evaluation license key expired	Optimization
1258	License violation: Bandwidth optimization stopped	Optimization
1259	Cannot complete the remote upgrade. Not enough free space.	Optimization
NA	<SSL> certs invalid (error : <error>) - backing up to /tmp/certs_bak and recreating Description: The SSL Certificates are invalid.	Optimization
NA	<NOProxy> has waited too long for GPStart Description: A problem occurred while starting the optimization service; the WANJet appliance was unable to become active.	Optimization
1260	Error: Upgrade failed. Please contact F5 technical support team.	WANJet Management
1300	Error: Logging error Description: The logging facility could not get error messages from the WANJet appliance.	Logs
1420	WCCP ServiceGroup (TCP) is up	WCCP
1421	WCCP ServiceGroup (UDP) is up	WCCP
1422	WCCP ServiceGroup (TCP) is down	WCCP

Table A.1 WANJet appliance messages (Continued)

Code	Message and Description	WANJet Appliance Component
1423	WCCP ServiceGroup (UDP) is down	WCCP
1424	WCCP Configuration Error	WCCP
1425	WCCP Runtime Error	WCCP
1426	WCCP is not enabled on the router	WCCP

Table A.1 WANJet appliance messages (Continued)



B

WANJet Appliance Private MIB File

- Ethernet card information
- Introducing MIB files

Ethernet card information

The WANJet appliance provides Ethernet card information by means of the standard SNMP.

The Ethernet card-related information path is:

`.iso.org.dod.internet.mgmt.mib-2.interfaces = .1.3.6.1.2.1.2`

The Ethernet card-related information description is:

`IfNumber`

`ifTable.ifEntry.ifIndex`

`ifTable.ifEntry.ifDescr`

`ifTable.ifEntry.ifEnter`

`ifTable.ifEntry.ifMtu`

`ifTable.ifEntry.ifSpeed`

`ifTable.ifEntry.ifPhysAddress`

`ifTable.ifEntry.ifInOctets`

`ifTable.ifEntry.ifInUcastPkts`

`ifTable.ifEntry.ifInDiscards`

`ifTable.ifEntry.ifInErrors`

`ifTable.ifEntry.ifOutOctets`

`ifTable.ifEntry.ifOutUcastPkts`

`ifTable.ifEntry.ifOutDiscards`

`ifTable.ifEntry.ifOutErrors`

Introducing MIB files

You can use WANJet Private Management Information Base (MIB) file if you need to compile the MIB file in order to browse the MIB using a standard MIB browser. You can download the MIB file from the WANJet appliance, and use it with your SNMP-compliant software.

To download the private MIB file for the WANJet appliance

1. In the navigation pane, expand **Configuration** and click **Monitoring**.
The WANJet Syslog and SNMP screen opens.
2. Click **WANJet Private MIB** to display the file in a web browser.
3. Save the file from the browser.
4. Copy the MIB file into your SNMP-compliant software and compile it.

Refer to the documentation of your SNMP-compliant software for specific instructions.

WANJet appliance MIB file

Following is the WANJet Private Management Information Base (MIB) file.

```
F5NETWORKS-GLOBAL-REG DEFINITIONS ::= BEGIN
IMPORTS
enterprises FROM SNMPv2-SMI;
F5 OBJECT IDENTIFIER
 ::= { enterprises 3375 }
WANJet OBJECT IDENTIFIER
 ::= { F5 11 }
Statistics OBJECT IDENTIFIER
 ::= { WANJet 2 }
SnmptTraps OBJECT IDENTIFIER
 ::= { WANJet 3 }
-- ***** Statistics
TotalSentBandwidthSavingPercent OBJECT-TYPE
SYNTAX INTEGER
ACCESS read-only
STATUS current
DESCRIPTION "Percentage of bandwidth saved on the traffic sent
to other WANJet appliances today."
 ::= { Statistics 1 }
```



```
TotalRecvBandwidthSavingPercent OBJECT-TYPE
SYNTAX INTEGER
ACCESS read-only
STATUS current
DESCRIPTION "Percentage of bandwidth saved on the traffic
received from other WANJet appliances today."
::= { Statistics 2 }

TotalSentBeforeWANJet OBJECT-TYPE
SYNTAX INTEGER
ACCESS read-only
STATUS current
DESCRIPTION "Effective traffic sent from this WANJet appliance
to other WANJet appliances
today in MB (before WANJet)."
::= { Statistics 3 }

TotalSentAfterWANJet OBJECT-TYPE
SYNTAX INTEGER
ACCESS read-only
STATUS current
DESCRIPTION "Optimized traffic sent from this WANJet appliance
to other WANJet appliances
today in MB (after WANJet)."
::= { Statistics 4 }

TotalRecvBeforeWANJet OBJECT-TYPE
SYNTAX INTEGER
ACCESS read-only
STATUS current
DESCRIPTION "Effective traffic received
from other WANJet appliances
today in MB (before WANJet)."
::= { Statistics 5 }

TotalRecvAfterWANJet OBJECT-TYPE
SYNTAX INTEGER
ACCESS read-only
STATUS current
DESCRIPTION "Optimized traffic received
from other WANJet appliances
today in MB (after WANJet)."
::= { Statistics 6 }
```

```

LastSentBandwidthSavingPercent OBJECT-TYPE
SYNTAX INTEGER
ACCESS read-only
STATUS current
DESCRIPTION "Percentage of bandwidth saved on the traffic sent
to other WANJet appliances during the last five minutes.
This value may be plotted to create a chart."
::= { Statistics 7 }
LastRecvBandwidthSavingPercent OBJECT-TYPE
SYNTAX INTEGER
ACCESS read-only
STATUS current
DESCRIPTION "Percentage of bandwidth saved on the traffic
received from other WANJet appliances during the last five
minutes.
This value may be plotted to create a chart."
::= { Statistics 8 }
LastSentBeforeWANJetRate OBJECT-TYPE
SYNTAX INTEGER
ACCESS read-only
STATUS current
DESCRIPTION "The rate of effective traffic sent from this WANJet
appliance to other WANJet appliances in Kbps (before WANJet).
This value may be plotted to create a chart."
::= { Statistics 9 }
LastSentAfterWANJetRate OBJECT-TYPE
SYNTAX INTEGER
ACCESS read-only
STATUS current
DESCRIPTION "The rate of optimized traffic sent from this WANJet
appliance to other WANJet appliances in Kbps (after WANJet).
This value may be plotted to create a chart."
::= { Statistics 10 }
LastRecvBeforeWANJetRate OBJECT-TYPE
SYNTAX INTEGER
ACCESS read-only
STATUS current
DESCRIPTION "The rate of effective traffic received
from other WANJet appliances in Kbps
(before WANJet).
This value may be plotted to create a chart."
::= { Statistics 11 }

```

```
LastRecvAfterWANJetRate OBJECT-TYPE
SYNTAX INTEGER
ACCESS read-only
STATUS current
DESCRIPTION "The rate of optimized traffic received
from other WANJet appliances in Kbps
(after WANJet).
This value may be plotted to create a chart."
::= { Statistics 12 }
-- ***** SnmpTraps
SnmpTrapObjs OBJECT IDENTIFIER
::= { SnmpTraps 1 }
SnmpTrapID OBJECT-TYPE
SYNTAX INTEGER
ACCESS read-only
STATUS current
DESCRIPTION "Holds the ID of the SNMP Trap."
::= { SnmpTrapObjs 1 }
SnmpTrapDescription OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Holds the description of the SNMP Trap."
::= { SnmpTrapObjs 2 }
SnmpTrapList OBJECT IDENTIFIER
::= { SnmpTraps 2 }

-- Optimization Engine Traps
Trap1000 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Configuration error."
::= { SnmpTrapList 1000 }
Trap1001 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Configuration error."
::= { SnmpTrapList 1001 }
```

```

Trap1002 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Configuration error."
::= { SnmpTrapList 1002 }
Trap1003 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Initialization error."
::= { SnmpTrapList 1003 }
Trap1004 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Initialization error."
::= { SnmpTrapList 1004 }
Trap1005 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Initialization error."
::= { SnmpTrapList 1005 }
Trap1006 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Internal error."
::= { SnmpTrapList 1006 }
Trap1007 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Internal error."
::= { SnmpTrapList 1007 }
-- Packet Processor Traps
Trap1100 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Internal error."
::= { SnmpTrapList 1100 }

```

```
Trap1101 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Internal error."
::= { SnmpTrapList 1101 }
Trap1102 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Internal error."
::= { SnmpTrapList 1102 }
Trap1103 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Internal error."
::= { SnmpTrapList 1103 }
Trap1150 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Maximum number of optimized connections reached."
::= { SnmpTrapList 1150 }
-- ACM5 Traps
Trap1200 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Configuration error."
::= { SnmpTrapList 1200 }
Trap1201 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Configuration error."
::= { SnmpTrapList 1201 }
Trap1202 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Initialization error."
::= { SnmpTrapList 1202 }
```

```

Trap1203 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Initialization error."
::= { SnmpTrapList 1203 }
Trap1204 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Internal error."
::= { SnmpTrapList 1204 }
Trap1205 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Internal error."
::= { SnmpTrapList 1205 }
Trap1206 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Internal error."
::= { SnmpTrapList 1206 }
Trap1207 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Internal error."
::= { SnmpTrapList 1207 }
Trap1209 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Link down with (Proxy IP).".
::= { SnmpTrapList 1209 }
Trap1210 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Link up with (Proxy IP).".
::= { SnmpTrapList 1210 }

```

```
Trap1211 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Authentication failed with (Proxy IP)."
```

```
 ::= { SnmpTrapList 1211 }

Trap1212 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Connection from unauthorized Proxy (Proxy
IP)."
```

```
 ::= { SnmpTrapList 1212 }

Trap1213 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Internal error."
```

```
 ::= { SnmpTrapList 1213 }

Trap1214 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: This version (%f) is incompatible with
(Proxy IP) version
(%f)."
```

```
 ::= { SnmpTrapList 1214 }

Trap1215 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: License expired on (%f)."
```

```
 ::= { SnmpTrapList 1215 }

Trap1216 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Connecting to remote client (%f)."
```

```
 ::= { SnmpTrapList 1216 }

Trap1217 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Connecting to local client (%f) error (%f)."
```

```
 ::= { SnmpTrapList 1217 }
```

```

Trap1250 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Version (%f) up and running."
::= { SnmpTrapList 1250 }
Trap1251 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Internal error."
::= { SnmpTrapList 1251 }
Trap1252 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Warning: License limit exceeded."
::= { SnmpTrapList 1252 }
Trap1253 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Warning: Invalid license key - Bandwidth
optimization off."
::= { SnmpTrapList 1253 }
Trap1254 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Warning: License key not entered - Bandwidth
optimization off."
::= { SnmpTrapList 1254 }
Trap1255 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Warning: Evaluation license key expires in x days."
::= { SnmpTrapList 1255 }
Trap1256 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Warning: WANJet evaluation license is activated for
x days."
::= { SnmpTrapList 1256 }

```



```
Trap1257 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Warning: Evaluation license key expired."
::= { SnmpTrapList 1257 }

Trap1258 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: License violation - Bandwidth optimization
stopped."
::= { SnmpTrapList 1258 }

Trap1259 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Cannot complete the remote upgrade. Not
enough free space."
::= { SnmpTrapList 1259 }

Trap1260 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Upgrade failed. Please contact F5 technical
support team."
::= { SnmpTrapList 1260 }

-- Logging Traps
Trap1300 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "Error: Logging error."
::= { SnmpTrapList 1300 }

-- WCCP Traps
Trap1420 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "WCCP ServiceGroup TCP is up."
::= { SnmpTrapList 1420 }
```

```
Trap1421 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "WCCP ServiceGroup UDP is up."
::= { SnmpTrapList 1421 }
Trap1422 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "WCCP ServiceGroup TCP is down."
::= { SnmpTrapList 1422 }
Trap1423 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "WCCP ServiceGroup UDP is down."
::= { SnmpTrapList 1423 }
Trap1424 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "WCCP Configuration Error."
::= { SnmpTrapList 1424 }
Trap1425 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "WCCP Runtime Error."
::= { SnmpTrapList 1425 }
Trap1426 OBJECT-TYPE
SYNTAX OCTET STRING
ACCESS read-only
STATUS current
DESCRIPTION "WCCP is not enabled on the router."
::= { SnmpTrapList 1426 }
END
```



C


RMON2 Tree

- RMON2 tree overview
- MIB tree
- Protocol directory tree
- Network layer matrix
- Application data matrix

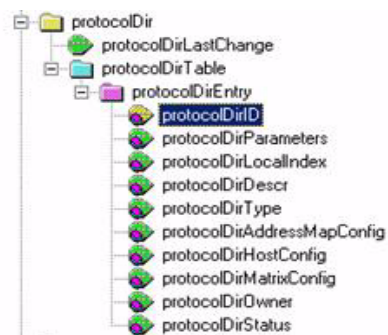
RMON2 tree overview

This appendix shows the standard RMON MIB tree and all the groups for RMON1 and RMON2. You can use it to locate the RMON2 groups that the WANJet appliance supports.

MIB tree



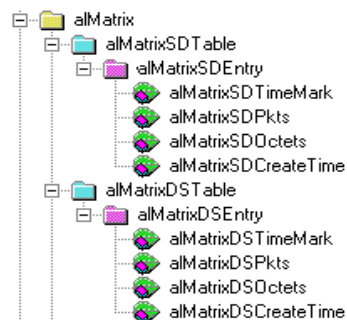
Protocol directory tree



Network layer matrix



Application data matrix





D

WANJet 200 Specifications

- WANJet 200 platform specifications

WANJet 200 platform specifications

The following specifications apply to the WANJet 200 platform.

Item	Specification
Dimensions	3.5 cm H x 22.4 cm W x 15.6 cm D (1.3"H x 8.81"W x 6.14"D) per unit
Weight	1.7 kg (3.8 lbs.) per unit
Power Supply	External
Input Voltage	100–240 VAC, +/-10%
Maximum Power Consumption	20 W
Maximum Heat Output	70 BTUs
Network Interfaces	Four 10/100 BaseT Fast Ethernet, with two of them bridged fail-to-wire
LCD module	None
Memory	PC133 SDRAM 256 MB
Operating Temperature	0° to 60° C (32° to 140° F)
Non-operating ambient temperature range	Temperature -40° to 70° C (-40° to 158° F) at a relative humidity of 10% to 90%
Relative Humidity	10% to 90% noncondensing
Hazardous Substance Compliance	RoHS compliant
Safety Agency Approval	CE/ETC
Electromagnetic Emissions Certifications	FCC CB Scheme

Table D.1 WANJet 200 platform specifications

◆ Important

Specifications are subject to change without notification.



E

WANJet 400 Specifications

- WANJet 400 platform specifications

WANJet 400 platform specifications

The following specifications apply to the WANJet 400 platform.

Item	Specification
Dimensions	8.9 cm H x 43.2 cm W x 54.0 cm D (3.5" H x 17" W x 21.25" D) per unit 2U industry standard rack-mount chassis
Weight	8.6 kg (19 lbs.) per unit
Power Supply	2 x 300 W redundant hot swappable power supplies
Input Voltage	100–240 VAC, +/-10%
Maximum Power Consumption	300 W
Maximum Heat Output	1025 BTUs
Network Interfaces	Four 10/100 BaseT Fast Ethernet, with two of them bridged fail-to-wire
LCD module	6-button keypad, 2 x 20 LCD
Memory	DDR-400 1GB with ECC x 4, Total 4 GB
Operating Temperature	0° to 60° C (32° to 140° F)
Non-operating ambient temperature range	-40° to 70° C (-40° to 158° F) at a relative humidity of 10% to 90%
Relative Humidity	10% to 90% noncondensing
Safety Agency Approval	CE/ETL
Electromagnetic Emissions Certifications	FCC CB Scheme

Table E.1 WANJet 400 platform specifications

◆ Important

Specifications are subject to change without notification.



Glossary

active unit

In a redundant system, the active unit is the system that currently optimizes connections. If the active unit in the redundant system fails, the standby unit assumes control and begins to optimize connections. See also *redundant system*.

Application QoS

Application QoS provides better service for data flows by configuring policies that can adjust the bandwidth consumed by specific types of network traffic (also referred to as *traffic shaping*). This way, you can use the network bandwidth optimally for critical network traffic. See also *Quality of Service (QoS) level*.

bandwidth

Bandwidth specifies the amount of data that can be transferred over a network connection in a fixed amount of time. Bandwidth is usually stated in bits per second (bps), kilobits per second (Kbps), or megabits per second (Mbps).

Connection Interception (CI)

Connection Interception (CI) intercepts and resets connections that were initiated before the WANJet appliance became active on the network.

dashboard

The dashboard is the area at the top of the navigation pane in the Web UI. It displays a variety of status indicators and shortcuts, and is always visible in the Web UI.

failover

Failover is the process whereby a standby unit in a redundant system takes over when a software failure or a hardware failure is detected on the active unit.

generic routing encapsulation

Generic routing encapsulation (GRE) tunneling typically connects private IP networks over an Internet connection using two routers (or switches) that support GRE encapsulation.

Layer 1 through Layer 7

Layers 1 through 7 refer to the seven layers of the Open System Interconnection (OSI) model. Thus, Layer 1 represents the physical layer, Layer 2 represents the data-link layer, Layer 3 represents the IP layer, and Layer 4 represents the transport layer (TCP and UDP). Layer 5 manages connections between applications, and Layer 6 represents the presentation layer. Layer 7 represents the application layer, handling traffic such as HTTP and SSL.

LCD

LCD stands for liquid crystal display. An LCD panel is available on the front of WANJet 400 and 500 platforms. You can use the LCD and its associated keypad to configure the LAN, WAN, and Management ports on the unit and perform basic administration tasks.

LED indicators

The LED indicators on the front of the WANJet appliance are lights that show the status of the system.

Management port

The Management port on the WANJet appliance can connect to the management network, if your organization is using out-of-band management.

NIC

A network interface card (NIC) is an expansion board used to connect a computer to a network.

OSI model

See *Layer 1 through Layer 7*.

Peer port

The Peer port on the WANJet appliance is an Ethernet port that can connect using a crossover cable to a second WANJet appliance to form a redundant system. See also *redundant system*.

Quality of Service (QoS) level

The Quality of Service (QoS) level is a means by which network equipment can identify and treat traffic differently based on an identifier. Essentially, the QoS level specified in a packet enforces a throughput policy for that packet. See also *Application QoS*.

redundant system

Redundant system refers to a pair of units that are configured for failover. In a redundant system, there are two units, one running as the active unit and one running as the standby unit. If the active unit fails, the standby unit takes over and manages connection requests.

router

A router is a Layer 3 networking device. If no VLANs are defined on the network, a router defines a broadcast domain.

SSL (Secure Sockets Layer)

SSL is a network communications protocol that uses public-key technology as a way to transmit data in a secure manner.

standby unit

A standby unit in a redundant system is a unit that is always prepared to become the active unit if the active unit fails.

Transparent Data Reduction™ (TDR)

TDR technology provides a dramatic reduction in the amount of bandwidth consumed by repeated data transfers across a WAN link.

Type of Service (ToS) level

The Type of Service (ToS) level is another means, in addition to the Quality of Service (QoS) level, by which network equipment can identify and treat traffic differently based on an identifier.

VLAN (Virtual LAN)

VLAN stands for virtual local area network. A VLAN is a logical grouping of network devices. VLANs provide a way to structure a large network for increased security, separating systems with sensitive data, for special projects, or into separate departments.

VLAN tag

A VLAN tag (also known as a VLAN ID) is a unique number between 2 and 4094 that identifies a particular VLAN.

WAN (wide area network)

A WAN is a computer network that spans a large geographic area, and typically consists of two or more local area networks (LANs). A WAN may also include public or shared user networks. The most well-known example of a WAN is the Internet.

Web UI

The Web UI is the interface to the WANJet appliance through which you can configure it and monitor network activity.



Index

A

About box 4-5

access

- allowing access by IP address 5-4
- denying access by IP address 5-5
- granting access to SNMP server 5-4

Active mode setting 6-11

Actual Bandwidth Expansion report 8-7

admin account 4-1, 5-1

administration tools 8-24

All Passthrough Connections report 8-14

application data matrix C-2

Application QoS policy

- adding 7-3
- and WAN Links 7-6
- description 2-6, 7-3
- editing or removing 7-5

ARP 6-30

ARP packets 6-30

Ask F5 search engine 1-4

authentication settings 5-1

Auto Negotiate setting 6-22, 10-6

Auto Sense setting 10-6

autonegotiation 6-32

Autopass report 8-14

B

backup files, creating 5-10

bandwidth

- reporting increased 8-6
- setting size for WAN link 6-13

bandwidth used, report 8-10

basic configuration example 9-1

boot images 5-10

bridge

- and high-availability 6-30
- description 8-22
- diagnosing connectivity 8-22

Bridge Forwarding Database report 8-23

buffer size, application 6-13

C

cable connector 10-9

cable, console 10-9

cables, Ethernet 6-32

CIFS protocol

- and Connection Intercept 6-10

command line

- attaching a computer 10-2
- overview 10-1
- using commands 10-3, 10-4

command syntax, conventions 1-3

command-line commands 10-4

community string 6-25, 8-30

Comparative Throughput report 8-5

config command 10-5

configuration

- adding subnets 4-10
- and deployment options 3-1
- deploying basic 9-1
- deploying hub and spoke 9-5
- deploying LAN router 9-9
- deploying mesh 9-3
- deploying redundant system 9-7
- modifying local network 6-15
- replicating changes 6-16
- setting up first appliance 4-10
- testing 4-14
- using the command line 10-5

congestion control 6-13

Connection Intercept 6-8

Connection States report 8-17

connectivity diagnostics 8-20, 10-8

console cable 10-9

CSV, and reports 8-5

D

data traffic, setting priority 2-6

DB9 to RJ45 cable pinouts 10-10

DB9 to RJ45 connector

- and pinouts 10-10

delayed acceptance settings 6-16

deployment

- configuring in-line 3-1
- configuring one-arm 3-4
- configuring point-to-multi-point 3-2
- configuring point-to-point 3-1

Diagnose Bridge report 8-22

Diagnose Connectivity report 8-20

Diagnose Ethernet report 8-20

Diagnose IP report 8-21

Diagnose Remote WANJet report 8-22

diagnose-connectivity command 10-8

diagnosing problems

- using error codes A-1
- using reports 8-11

diagnostic tools

- running at the command line 10-7
- running in the Web UI 8-24

diagnostics

- and logs 8-27
- downloading 8-27
- monitoring information 8-11
- viewing 4-15

Diagnostics reports 8-11

documentation 1-4

duplex setting 6-22, 6-32, 10-6

E

- email alerts 6-25
- error messages and codes A-1
- eth0 port 6-22
- eth1 port 6-22
- eth3 port 6-22
- Ethernet cables 6-32
- Ethernet cards
 - and SNMP information B-1
 - viewing diagnostics 8-20
- event messages, syslog 2-9
- Exit command 10-3, 10-4

F

- fail-close feature 6-11, 6-12, 6-32
- fail-to-wire feature 6-11, 6-31
- failure detection 6-35
- Failure Mode setting 6-11
- firewall ports 3-8
- flash memory card 5-10
- freed bandwidth, report 8-6

G

- gateway, specifying a static route 6-23
- Generic routing encapsulation tunneling 3-7
- green light 4-14
- guaranteed performance
 - providing for networks 2-6

H

- halt command 10-9
- hardware specifications D-1, E-1
- heartbeat packets 6-35
- high-availability features 6-30
- HSRP 6-30
- hub and spoke configuration example 9-5
- HyperTerminal 10-2

I

- ICMP protocol 3-8
- Inactive mode setting 6-11
- initial configuration, verifying connectivity 4-14
- in-line deployment
 - and high availability 6-36
- inline deployment
 - description 3-1
 - topology 6-12
- interface diagnostics
 - viewing 8-12
- interface speed 6-22, 10-6

IP address

- accessing Web UI 5-4
- and multiple subnets 4-10
- diagnosing connectivity 8-21
- setting from the command line 10-5

IT service policy

- adding 7-1
- description 7-1
- editing 7-2

L

- LAN port 6-22
- LAN router 4-13, 6-15, 6-23
- LAN router configuration example 9-9
- Layer 5 1-1
- LCD keypad
 - restarting the unit 5-9
 - shutting down the unit 5-8
- LCD PIN code 5-3
- license
 - upgrading 5-13
- Link Utilization report 8-10
- links 4-4, 7-6, 8-15
- Liquid Crystal Display. See LCD keypad.
- load balancing 6-35
- local routers
 - and service groups 3-6
- local subnets 6-1
- Logout button 4-5
- Logout command 10-3, 10-4
- logs
 - and diagnostics 8-11
 - downloading 8-27

M

- Management Information Base. See MIB.
- Management IP address 10-6
- Management port 5-5, 6-22
- matrix
 - viewing application data C-2
 - viewing network layer C-2
- media types 6-22, 10-6
- mesh configuration example 9-3
- MIB 2-7
- MIB file B-2
- MIB tree C-1
- modes, processing 6-4
- monitoring, diagnostics information 8-11

N

- navigation in user interface 4-2
- network interface settings 6-22
- network layer matrix C-2
- non-transparent proxy 6-28

O

- one-arm deployment
 - and high-availability 6-36
 - description 3-4
 - running Real Time Traffic report 8-4
 - setting one-arm option 6-12
- operational mode settings 6-11
- optimization 2-1
- optimization policies 6-1
- Optimize Eligible Connections report 8-14
- optimized connections 1-1
- Optimized Data report 8-8
- Optimized Sessions report 4-4, 8-13
- OSI reference model 1-1, 2-1, 2-8
- outgoing packets
 - setting queue size 6-13
- Overall Data report 8-9

P

- Packet Retransmissions report 8-17
- packet size
 - configuring outgoing 6-13
- packets
 - directing through a gateway 6-23
 - sending heartbeat 6-35
- passthrough sessions 8-14
- password
 - setting for Web UI 5-1
- path
 - viewing MIB tree C-1
- peer port 6-33
- peers, redundant 6-26
- performance
 - providing guaranteed level 2-6
- Performance Increase report 8-6
- PIN code, setting 5-3
- ping command 8-2, 8-24, 10-7
- pinout table 10-10
- planning worksheet 3-9
- platform
 - reviewing hardware specifications D-1, E-1
- point-to-multipoint configuration 3-2
- point-to-point configuration 3-1
- port settings 6-4
- ports
 - adding 6-5
 - and WANJet appliance 3-8
 - configuring delayed acceptance 6-16
 - opening behind a firewall 3-8
- power off 5-8
- private MIB file B-2
- problems
 - diagnosing 4-16
 - using diagnostic reports 8-11
 - viewing error codes A-1

- processing modes

- assigning to ports 6-5
- protocol directory tree C-1
- proxy, non-transparent 6-28
- proxy, transparent 3-5, 3-7, 6-28
- PuTTY 10-3

Q

- QoS diagnostics 8-18
- queue size
 - configuring for outgoing packets 6-13
- Quick Start Cards 1-4, 3-1, 3-8, 10-5

R

- RADIUS authentication 5-2
- RADIUS report 8-15
- Raw Data setting 6-24, 8-31
- read-only administrator 4-1
- Real Time Traffic report 8-3
- Realtime Connections report 8-14
- reboot command 10-9
- Receive Queue Packets Pruned report 8-17
- redundant configuration example 9-7
- redundant peers
 - description 6-26
- remote authentication 5-2
- remote links 4-4
- remote monitoring. See RMON.
- remote user accounts
 - accessing Web UI 5-1
- remote WANJet appliance
 - adding to local appliance 6-20
 - defining 4-13
 - editing or removing 6-21
 - testing connectivity 8-2
- reporting systems
 - configuring RMON2 8-30
 - configuring third-party 8-28
- reports
 - and actual bandwidth expansion 8-7
 - and comparative throughput 8-5
 - and diagnostic log 8-27
 - and link utilization 8-10
 - and optimized data 8-8
 - and overall data 8-9
 - and performance increase 8-6
 - and real time traffic 8-3
 - and RMON2 8-29
 - and SNMP 8-28
 - and status 8-2
 - and syslog 8-28
 - and traffic reduction 8-8
 - overview 8-1
 - saving to CSV 8-5
- restart process 5-8, 5-9, 10-9

retransmitted packets 8-17

RMON

- supporting 2-8

RMON1

- description 2-8

- See also RMON2.

RMON2

- configuring reports 6-24

- description 2-8

- listing supported groups 2-9

- understanding configuration 8-30

- viewing application data matrix C-2

- viewing MIB tree C-1

- viewing network layer matrix C-2

- viewing protocol directory tree C-1

- viewing reports 8-29, 8-30

RMON2 groups 2-9

RMON2 logs 8-29

roadadmin account 4-1, 5-1

round trip time

- specifying for WAN Link 6-13

S

security

- setting PIN code for LCD 5-3

- setting Web UI password 5-1

service groups

- and local routers 3-6

- and web caches 3-6

- definition 3-6

service policies

- adding 7-1

- and Application QoS policy 2-6

- editing 7-2

services

- adding 6-5

sessions

- diagnosing optimized 8-13

- diagnosing passthrough 8-14

settings

- creating backup file for 5-10

- restoring from backup 5-11

Shared Key setting 4-14

shut down process 5-8

shutdown process 5-8, 10-9

Simple Network Management Protocol. See SNMP.

site information worksheet 3-9

snapshot of system 4-16, 8-27

SNMP

- accessing reports 5-4

- and Ethernet card information B-1

- and MIB 2-7, B-2

- configuring reports 6-24

- description 2-7

- viewing reports 8-28

SNMP reports

- configuring 8-28

- configuring access to 2-7

- viewing 6-24

SNMP tables 8-29

SNMP traps 2-7

software images 5-10

specifications, platform D-1, E-1

speed setting 6-32

SSH 10-3

static routes

- adding 6-23

- and one-arm deployment 6-28

- removing 6-23

Static Routes table 6-23

static routes, specifying 3-5, 6-23

Status report 4-2, 4-14, 8-2

STP 6-30

stylistic conventions 1-2

subnets

- adding on a local WANJet appliance 6-1

- adding on a remote WANJet appliance 6-3

- adding to a WAN Link 7-7

- configuring multiple 4-10

- editing on a local WANJet appliance 6-1

- editing or removing from a WAN Link 7-8

- removing from a local WANJet appliance 6-1

support

- downloading diagnostics 8-27

SYN packet 6-34

synchronizing time 5-7

syslog

- and event messages 2-9

- and reports 8-28

- configuring settings 6-24

- description 2-9

syslog server

- configuring 6-24

- sending messages to 2-9

System Log protocol. See syslog.

system snapshot 4-16, 6-25, 8-27

T

TCP ports 3-8, 6-1

TCP statistics diagnostics 8-16

TCP/UDP ports

- setting default processing mode 6-7

tcpdump utility 8-26

TDR Statistics diagnostics 8-18

terminal emulator 10-2

time

- displaying 4-4

- setting manually 5-7

- synchronizing automatically 5-7

time management 5-6

- time server
 - description 5-6
 - synchronizing 5-7
- time zone, setting 5-6
- topology settings 6-11
- ToS
 - assigning priorities to data traffic 2-6
 - description 2-6
- ToS priority 6-5, 6-7
- traceroute command 8-2, 8-25
- traffic optimized report 8-8
- traffic priority 2-6
- Transparent Data Reduction 1-1, 2-2
- transparent proxy 3-5, 3-7, 6-28
- trees
 - viewing MIB C-1
 - viewing protocol directory C-1
- troubleshooting
 - testing connectivity 4-14
 - understanding common problems 4-16
 - using error codes A-1
- tuning
 - specifying bandwidth and RTT for WAN Link 6-13
- Type of Service. See ToS.

U

- UDP ports 3-8, 6-1
- upgrading software 5-13

V

- version, upgrading 5-13
- Virtual IP address 6-30
- VLAN diagnostics 8-19
- VLANs, managing 6-17
- VRRP 6-30

W

- WAN Gateway setting 6-15
- WAN Links
 - adding 7-6
 - adding subnets 7-7
 - description 7-6
 - editing or removing 7-6
 - setting bandwidth size 6-13
- WAN port 6-22
- WAN router 3-4
- WAN sessions 4-4
- WANJet Alias setting 6-15
- WANJet Data setting 6-24, 8-32
- WANJet Links report 8-15
- WCCP v2 protocol 3-6, 6-36
- web caches
 - and service groups 3-6

Web UI 10-6

- authenticating users 5-1
- granting access 5-4
- logging off 4-6
- logging on 4-1
- setting password 5-1
- viewing screens 4-5

Web upgrade 5-13

worksheet, site information 3-9

