

BusinessCom Sentinel PF Operating System Handbook

BUSINESSCOM[®]
NETWORKS

BusinessCom Networks Limited
Glacis Road, Portland House, Suite 2
Gibraltar, GX11 1AA

December 2014, Revision D

BusinessCom Networks Limited has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more U.S. patents or pending patent applications in the U.S. and in other countries.

Portions Copyright © 2005 - 2013 Peter N. M. Hansteen, Portions Copyright © 1995, 1996, 1997, 1998, 1999, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013 The FreeBSD Documentation Project. Portions Copyright © Wikipedia. Portions Copyright © OpenVPN Technologies, Inc.

U.S. Government Rights – Commercial software. Government users are subject to the BusinessCom Networks Limited standard license agreement and applicable provisions of the FAR and its supplements.

This distribution may include materials developed by third parties. Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd. Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries. Intel Inside is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries. NVIDIA is a trademark or registered trademark of NVIDIA Corporation or its subsidiaries in the United States and other countries. LSI is a trademark or registered trademark of LSI Corporation or its subsidiaries in the United States and other countries.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this document, and the FreeBSD Project was aware of the trademark claim, the designations have been followed by the “™” or the “®” symbol.

FreeBSD is a registered trademark of the FreeBSD Foundation.

Heidelberg, Helvetica, Palatino, and Times Roman are either registered trademarks or trademarks of Heidelberger Druckmaschinen AG in the U.S. and other countries.

IEEE, POSIX, and 802 are registered trademarks of Institute of Electrical and Electronics Engineers, Inc. in the United States.

Intel, Celeron, EtherExpress, i386, i486, Itanium, Pentium, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

Motif, OSF/1, and UNIX are registered trademarks and IT DialTone and The Open Group are trademarks of The Open Group in the United States and other countries.

SAP, R/3, and mySAP are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world.

Products covered by and information contained in this publication are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical or biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

THIS DOCUMENTATION IS PROVIDED "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL BUSINESSCOM NETWORKS LIMITED BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS DOCUMENTATION, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Do not discard electronic products in household trash. All waste electronics equipment should be recycled according to local regulations. Information for the recycler: Please cut off Lithium battery, if present, for separate recycling.

Contents

| | |
|---|-----------|
| Preface | 11 |
| Typographic Conventions | 11 |
| Naming Conventions | 11 |
| About Sentinel PF | 12 |
| Deployment | 12 |
| Deployment: Branch Office | 13 |
| Deployment: Wireless ISP | 14 |
| Customer Service | 15 |
| BusinessCom Welcomes Your Comments | 15 |
| Warranty Policy | 15 |
| Limitations of Warranty | 16 |
| Exclusive Remedies | 16 |
| Electrical Safety | 17 |
| Change History | 18 |
| Installation Overview | 20 |
| Process Flow | 20 |
| Required Experience | 21 |
| Required Equipment and Software | 21 |
| Blackbird and Westwind: Checking Package Contents | 22 |
| D2, D2W and Sentinel 3: Checking Package Contents | 22 |
| Precautions | 22 |
| Physical Connections | 23 |
| Blackbird and Westwind: Front Panel | 23 |
| Blackbird and Westwind: Motherboard | 24 |

| | |
|---|-----------|
| Blackbird: LAN and WAN Ethernet Ports | 24 |
| Westwind: LAN and WAN Ethernet Ports | 25 |
| Blackbird and Westwind: Ethernet Indicator Lights | 26 |
| D2, D2W: Rear Panel Motherboard | 27 |
| Sentinel 3: Rear Panel Motherboard | 28 |
| Powering Sentinel On and Off | 29 |
| To Power On Sentinel | 29 |
| To Power Off Sentinel | 29 |
| Command Line Interface | 30 |
| Introduction to CLI | 30 |
| CLI Prompt | 31 |
| CLI Functionality Overview | 31 |
| Logout (SSH Only) | 32 |
| Assign Interfaces | 32 |
| To Assign Interfaces | 33 |
| Set Interface(s) IP Address | 34 |
| To Set Interface IP Address | 34 |
| Reset GUI Password/Reset To Factory Defaults | 34 |
| To Reset Sentinel | 34 |
| Ping host | 35 |
| CLI | 35 |
| View Network Connections | 36 |
| To Show Connections With Most Packet Rate | 38 |
| To Show Connections With Most Traffic | 38 |
| To Show Oldest Connections | 38 |
| Filter Logs | 38 |
| Restart GUI | 38 |
| Web GUI | 39 |

| | |
|--|-----------|
| Introduction to Web GUI | 39 |
| To Connect to Web GUI | 39 |
| Initial Configuration Steps | 41 |
| First Steps | 41 |
| Web GUI: System | 42 |
| Cert Manager | 42 |
| To Create New Certificate Authority (CA) | 42 |
| To Create New User or Server Certificate | 43 |
| General Setup | 44 |
| Packages | 45 |
| Routing | 47 |
| To Add New Gateway | 47 |
| To Add New Static Route | 48 |
| To Create New Gateway Group | 49 |
| User Manager | 49 |
| To Create New User | 50 |
| To Create New Group | 50 |
| To Assign Privileges To A Group | 51 |
| To Use LDAP or RADIUS As Authentication | 51 |
| Web GUI: Interfaces | 52 |
| (assign) | 52 |
| To Create Interface Group | 52 |
| To Use Sentinel As WPA Access Point | 54 |
| To Use Sentinel As WPA Client | 55 |
| To Create a VLAN | 56 |
| To Create QinQ Interface | 57 |
| To Create PPP Interface | 58 |
| To Create GRE Tunnel | 59 |

| | |
|---|-----------|
| To Create GIF Tunnel | 60 |
| To Create Bridge | 61 |
| LAN, WAN, OPT1 and Other Interface Menus | 65 |
| Web GUI: Firewall | 67 |
| Aliases | 67 |
| To Create Alias | 67 |
| NAT | 68 |
| Rules | 72 |
| To Create Firewall Rule | 74 |
| To Move Rule | 74 |
| Schedules | 78 |
| To Create Schedule | 79 |
| Traffic Shaper | 80 |
| To Create Limiter | 89 |
| Use Limiters To Prioritize LAN Users | 91 |
| To Shape Traffic via Layer 7 | 93 |
| Virtual IPs | 95 |
| Web GUI: Services | 97 |
| To Show Bandwidth Chart For Individual IP | 98 |
| Captive Portal | 100 |
| To Start Using Vouchers In Your LAN | 106 |
| DHCP Relay and DHCPv6 Relay | 107 |
| To Enable DHCP Relay | 107 |
| DHCP Server and DHCPv6 Server/RA | 107 |
| To Add Static Mapping | 109 |
| DNS Forwarder | 110 |
| To Enable DNS Forwarder | 110 |
| To Add DNS Host Override | 111 |

| | |
|--|------------|
| To Add DNS Domain Override | 111 |
| Dynamic DNS | 112 |
| To Use Dynamic DNS | 112 |
| IGMP Proxy | 113 |
| To Enable IGMP Proxy | 113 |
| Load Balancer | 114 |
| To Create Internal Server Pool | 115 |
| To Create Virtual Server | 115 |
| NTP | 116 |
| OLSR | 116 |
| For Transparent PEP HTTP Redirect | 118 |
| PPPoE Server | 119 |
| For PEP Acceleration with HTTP Proxy Cache | 129 |
| RIP | 132 |
| SNMP | 132 |
| To Enable SNMP Daemon | 134 |
| To Create New NIDS Suppress List | 137 |
| To Edit NIDS Suppress List | 137 |
| UPnP & NAT-PMP | 145 |
| Wake on LAN | 146 |
| Web GUI: VPN | 147 |
| IPSec | 147 |
| L2TP | 151 |
| OpenVPN | 152 |
| PPTP | 157 |
| Web GUI: Status | 159 |
| CARP (failover) | 159 |
| Dashboard | 159 |

| | |
|--|------------|
| DHCP Leases and DHCPv6 Leases | 159 |
| Filter Reload | 159 |
| Gateways | 159 |
| Interfaces | 160 |
| IPSec | 160 |
| Load Balancer | 161 |
| Network Monitor | 161 |
| NTP | 161 |
| OpenVPN | 161 |
| Package Logs | 161 |
| Queues | 162 |
| RRD Graphs | 163 |
| Services | 163 |
| System Logs | 163 |
| Traffic Graph | 164 |
| UPnP & NAT-PMP Status | 164 |
| Web GUI: Diagnostics | 165 |
| ARP Table | 165 |
| Authentication | 165 |
| Backup/Restore | 166 |
| To Backup Current Configuration | 166 |
| To Restore Configuration From File | 166 |
| To View Changes Between Configurations | 167 |
| Command Prompt | 167 |
| DNS Lookup | 167 |
| Edit File | 167 |
| Factory Defaults | 168 |
| To Reset Sentinel | 168 |

| | |
|---|------------|
| Halt System | 168 |
| Limiter Info | 168 |
| Packet Capture | 168 |
| pfInfo | 169 |
| pfTop | 170 |
| Ping | 170 |
| Reboot | 170 |
| Routes | 170 |
| SMART Status | 171 |
| States | 172 |
| States Summary | 173 |
| System Activity | 174 |
| Tables | 174 |
| Traceroute | 174 |
| NanoBSD | 174 |
| NDP Table | 174 |
| Sockets | 174 |
| Test Port | 175 |
| Redundancy | 176 |
| Sentinel Redundancy Overview | 176 |
| Example CARP Redundant Configuration | 177 |
| Redundancy Configuration Example | 178 |
| Multiple WAN Connections | 181 |
| Mutli-WAN Environment | 181 |
| Example Redundant Multi-WAN Configuration | 182 |
| State Filter Expressions | 183 |
| Netmask/CIDR Translation Table | 185 |
| Factory Default Settings | 186 |

| | |
|--|------------|
| Additional Support | 187 |
| DiffServ Classification & Marking | 188 |
| NIDS Incidents Classification | 190 |
| Sentinel D2 RS-232 Console Cable | 195 |
| Troubleshooting | 196 |
| Boot | 196 |
| Run-Time | 196 |
| Services | 197 |

Preface

The *BusinessCom Sentinel PF Operating System Handbook* provides instructions for network administrators and field engineers deploying Sentinel PF operating system based hardware servers (Sentinel thereafter); as well as any other personnel involved into the installation, operation or monitoring of networks featuring Sentinel.

Typographic Conventions

The following table describes the typographic conventions that are used in this book.

TABLE P-1 Typographic Conventions

| Typeface | Meaning | Example |
|------------------|---|--|
| AaBbCc123 | The names of commands, files, and directories, and onscreen computer output | Edit your <code>.login</code> file. |
| AaBbCc123 | What you type, contrasted with onscreen computer output | <code>machine_name% su</code> |
| <i>aabbcc123</i> | Placeholder: replace with a real name or value | The command to remove a file is <code>rm filename</code> . |
| <i>AaBbCc123</i> | Book titles, new terms, and terms to be emphasized | A <i>cache</i> is a local copy. |

Chapter names with ❖ mark designate external software packages.

Naming Conventions

The following table describes the naming conventions that are used in this book.

TABLE P-2 Naming Conventions

| Names | Meaning |
|---|---|
| Sentinel PF | The operating system software. |
| Sentinel Blackbird Sentinel Westwind Sentinel Sierra Sentinel 3 Sentinel D2 and D2W | Hardware servers. |
| Sentinel | In this book, refers to the hardware server running Sentinel PF operating system. |

About Sentinel PF

Sentinel PF is a specialized operating system designed for bandwidth management and optimization purposes. It is built on a solid foundation of the FreeBSD and features rich UNIX® heritage. The code is based on U.C. Berkeley's "4.4BSD-Lite" release, with some "4.4BSD-Lite2" enhancements. It is also based indirectly on William Jolitz's port of U.C. Berkeley's "Net/2" to the i386™, known as "386BSD", though very little of the 386BSD code remains.

The "PF" part in the name derives from the OpenBSD Packet Filter developed by Daniel Hartmeier and a number of OpenBSD developers. PF is a packet filter, that is, code which inspects network packets at the protocol and port level, and decides what to do with them. PF operates in a world which consists of packets, protocols, connections and ports. Based on where a packet is coming from or where it's going, which protocol, connection or port it is designated for, PF is able to determine where to lead the packet, or decide if it is to be let through at all.

One important feature of the packet filter and similar software, perhaps the most important feature, is that it is able to identify and block traffic which you do not want to let into your local network or let out to the world outside. The packet filter is a very flexible tool which is extremely useful when you want to take control of what happens in your network.

Sentinel PF is much more than a simple firewall. It provides a wealth of advanced networking features, such as Layer 7 traffic shaping, captive portal, NAT, DHCP server, HTTP and DNS caching, failover and load balancing, NIDS (Network Intrusion Detection System), VPN tunneling, CARP redundancy, PEP WAN optimization and traffic compression, and much more. Our primary goal for developing Sentinel PF is to enable network administrators to take full control of their networks.

Deployment

BusinessCom provides some initial deployment scenarios to help illustrate how Sentinel PF based servers can provide bandwidth management and advanced networking solutions. The initial deployment scenarios are based on actual environments configured during the product testing phase. They include common combinations of different networking environments.

These deployment scenarios represents specific examples of how this product can be deployed and are intended to augment the information provided in the product documentation. There are many other supported deployment scenarios that are possible. Feel free to share your deployment scenarios and show off what you have done with Sentinel PF.

Deployment: Branch Office

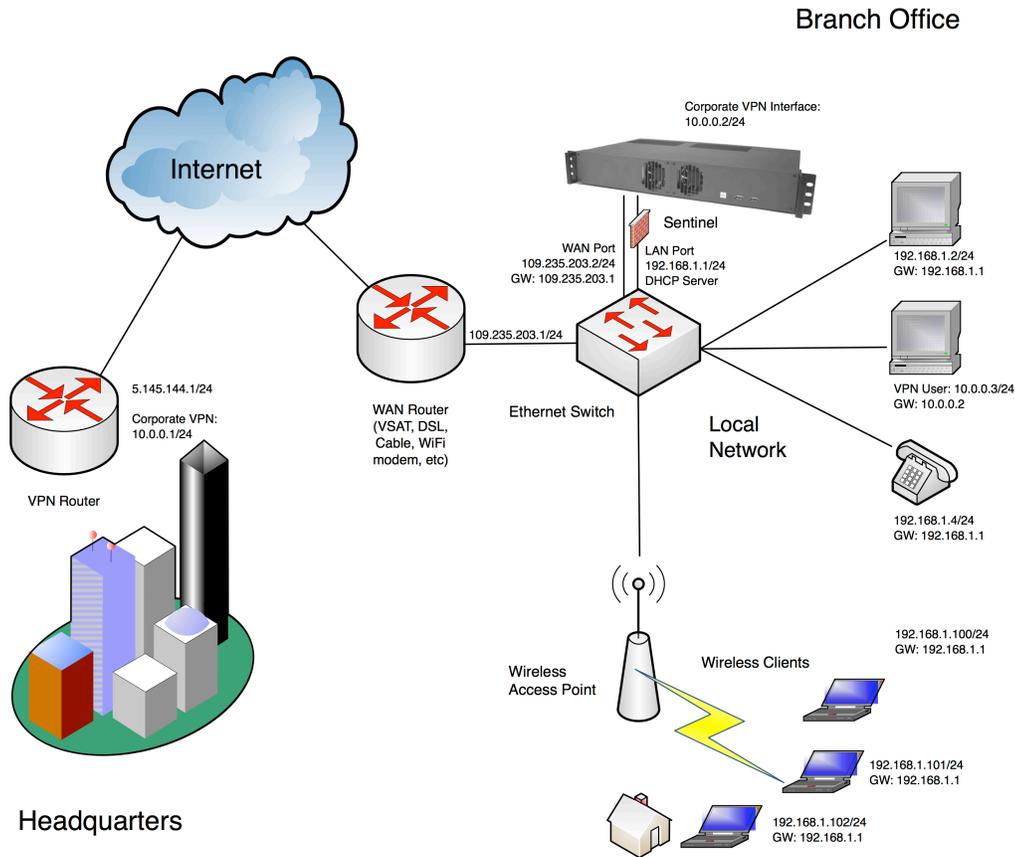


FIGURE P-1 Branch Office Example.

With the example diagram above, Sentinel is deployed to serve as the bandwidth management appliance for a branch office. The Sentinel is configured with a Public IP interface on the WAN port and provides routing, NAT and firewall services to the local network. Sentinel acts as the DHCP server for the network and assigns private IP addresses in the 192.168.1.100 to 192.168.1.102 range to wireless clients. Office computers and Voice over IP telephones are configured with static private IP addresses in the same subnet.

Additionally, Sentinel acts as the VPN client and has a 10.0.0.2/24 interface to the router at the headquarters for encrypted communications. The local VPN user (10.0.0.3) uses Sentinel as its gateway to the corporate HQ. In addition to essential networking features, Sentinel provides Layer 7 traffic shaping, monitoring and reporting, Network Intrusion Detection System (NIDS) functionality to the office network and also acts as the captive portal with user database and management back-end for clients connecting via wireless links.

Deployment: Wireless ISP

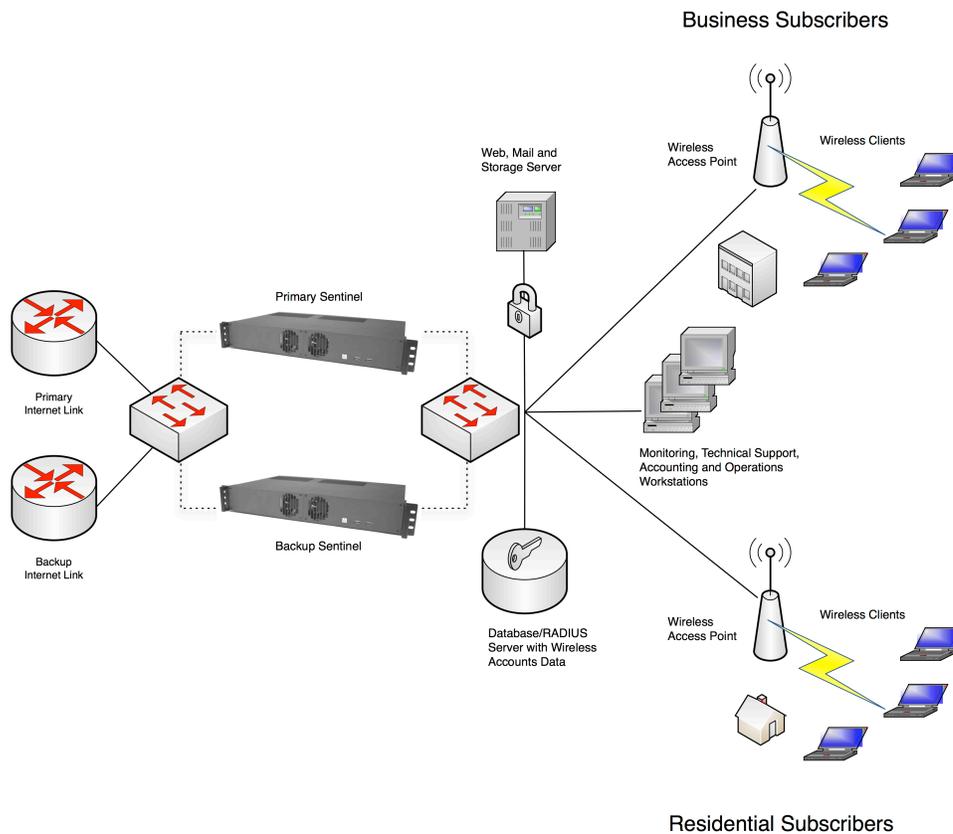


FIGURE P-2 Wireless ISP Example.

Sentinel can be deployed to manage bandwidth and provide critical services in Wireless Hot Spot and ISP networks. In the example above, a pair of Sentinel servers are configured for redundancy purposes. Both servers are synchronized via an internal protocol to ensure configuration is always kept the same. Each Sentinel server is configured with two WAN interfaces, connected to Primary and Backup Internet links. The WAN interfaces are configured in the fail-over mode whereby Sentinel automatically switches all the traffic to the backup link if primary link fails. Sentinel can also be set up to load balance both links, so bandwidth on the backup link won't be wasted.

Sentinels also serve as the captive portal for both residential and business subscribers connected via two wireless access points. Sentinels are setup to communicate with the external account database via 3rd party RADIUS server that also provides authentication to Web, Mail and storage server deployed in the Hot Spot's local network. In addition to standard DHCP, NAT and Firewall functionality with real-time reporting of malicious activity to network's administrator, Sentinels enforce Layer 7 traffic shaping with different QoS profiles setup for business and residential users. Business traffic is prioritized during business hours, and mission-critical traffic is prioritized on a 24 x 7 basis. Peer to Peer and bandwidth intensive applications are de-prioritized to ensure service consistency, however they are allowed to consume idle bandwidth. Sentinels can be configured to allow unauthorized users to access a "walled garden" of web sites or payment portals.

Customer Service

Contact BusinessCom Networks Support for product support or training, information on upgrading or returning a product, reporting comments or suggestions concerning manuals.

BusinessCom Networks
Attention: Customer Support Department
Cernochova 1291/4, Prague, Czech Republic, 158 00

E-mail: support@bcsatellite.net
WWW: <http://www.bcsatellite.net>

BusinessCom Welcomes Your Comments

BusinessCom is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments via e-mail. Please include the title of your document with your feedback: *BusinessCom Sentinel PF Operating System Handbook*.

Warranty Policy

BusinessCom Networks products are warranted against defects in material and workmanship for a specific period from the date of shipment, and this period varies by product. In most cases, the warranty period is one year. During the warranty period, BusinessCom Networks will, at its option, repair or replace products that prove to be defective. Repairs are warranted for the remainder of the original warranty. Contact BusinessCom Networks for the warranty period specific to the product purchased.

For equipment under warranty, the owner is responsible for freight to BusinessCom Networks and all related customs, taxes, tariffs, insurance, etc. BusinessCom Networks is responsible for the freight charges only for return of the equipment from the factory to the owner. BusinessCom Networks will return the equipment by the same method (i.e., Air, Express, Surface) as the equipment was sent to BusinessCom Networks.

All equipment returned for warranty repair must have a valid RMA number issued prior to return and be marked clearly on the return packaging. BusinessCom Networks strongly recommends all equipment be returned in its original packaging.

BusinessCom Networks' obligations under this warranty are limited to repair or replacement of failed parts, and the return shipment to the buyer of the repaired or replaced parts.

Limitations of Warranty

The warranty does not apply to any part of a product that has been installed, altered, repaired, or misused in any way that, in the opinion of BusinessCom Networks, would affect the reliability or detracts from the performance of any part of the product, or is damaged as the result of use in a way or with equipment that had not been previously approved by BusinessCom Networks. The warranty does not apply to any product or parts thereof where the serial number or the serial number of any of its parts has been altered, defaced, or removed. The warranty does not cover damage or loss incurred in transportation of the product. The warranty does not cover replacement or repair necessitated by loss or damage from any cause beyond the control of BusinessCom Networks, such as lightning or other natural and weather related events or wartime environments. The warranty does not cover any labor involved in the removal and or reinstallation of warranted equipment or parts on site, or any labor required to diagnose the necessity for repair or replacement. The warranty excludes any responsibility by BusinessCom Networks for incidental or consequential damages arising from the use of the equipment or products, or for any inability to use them either separate from or in combination with any other equipment or products. A fixed charge established for each product will be imposed for all equipment returned for warranty repair where BusinessCom Networks cannot identify the cause of the reported failure.

Exclusive Remedies

BusinessCom Networks' warranty, as stated is in lieu of all other warranties, expressed, implied, or statutory, including those of merchantability and fitness for a particular purpose. The buyer shall pass on to any purchaser, lessee, or other user of BusinessCom Networks' products, the aforementioned warranty, and shall indemnify and hold harmless BusinessCom Networks from any claims or liability of such purchaser, lessee, or user based upon allegations that the buyer, its agents, or employees have made additional warranties or representations as to product preference or use. The remedies provided herein are the buyer's sole and exclusive remedies. BusinessCom Networks shall not be liable for any direct, indirect, special, incidental, or consequential damages, whether based on contract, tort, or any other legal theory.

Electrical Safety

Please refer to the specifications sheet of the Sentinel hardware server you are using for the information on nominal voltage operating range and maximum power consumption.



Caution – Sentinel hardware servers contain a Lithium Battery. **DANGER OF EXPLOSION EXISTS** if the battery is incorrectly replaced. Replace only with the same or equivalent type recommended by the manufacturer. Dispose of used batteries in accordance with local and national regulations.



Caution – **PROPER GROUNDING PROTECTION IS REQUIRED:** The installation instructions require that the integrity of the protective earth must be ensured and your equipment shall be connected to the protective earth connection at all times. Therefore, it is imperative during installation, configuration, and operation that the user ensures that the unit has been properly grounded using the ground stud provided on the rear panel of the unit.

Sentinel servers are designed for connection to a power system that has separate ground, line and neutral conductors. The equipment is not designed for connection to a power system that has no direct connection to ground.

Sentinel servers must be operated with cover on at all times. If it becomes necessary to remove the cover, the user should ensure that the cover is correctly re-fitted before normal operation commences.

Sentinel servers are shipped with a line inlet cable suitable for use in the country of operation. If it is necessary to replace this cable, ensure the replacement has an equivalent specification. Examples of acceptable ratings for the cable include HAR, BASEC and HOXXX-X. Examples of acceptable connector ratings include VDE, NF-USE, UL, CSA, OVE, CEBEC, NEMKO, DEMKO, BS1636A, BSI, SETI, IMQ, KEMA-KEUR and SEV.

Change History

The following changes have been made to the documentation set.

- 05/2013, Rev A initial documentation was published for Sentinel PF version 2.0.3 build 04/24/2013.
- 09/2013, Rev B. Updated troubleshooting section. Added Sentinel D2W information. Added Huginn, iDirect Monitor section. Reviewed PEP service parameters. Added notices about logging on NanoBSD based Sentinel models for PEP and Proxy services. Added proxy server (TP) package.
- 04/2014, Rev C. Added transparent PEP redirect how-to. Added Network Monitor. Updated Westwind NIC information.
- 12/2014, Rev D. Modified limiter priorities description. Low-number priorities are higher. Added Sentinel 3 information. Published for Sentinel PF version 2.1.5.

Installation Overview

This overview provides procedures for preparing Sentinel for deployment into your existing networking environment.

Process Flow

Figure 1–1 shows the process flow for getting Sentinel ready for the initial configuration.

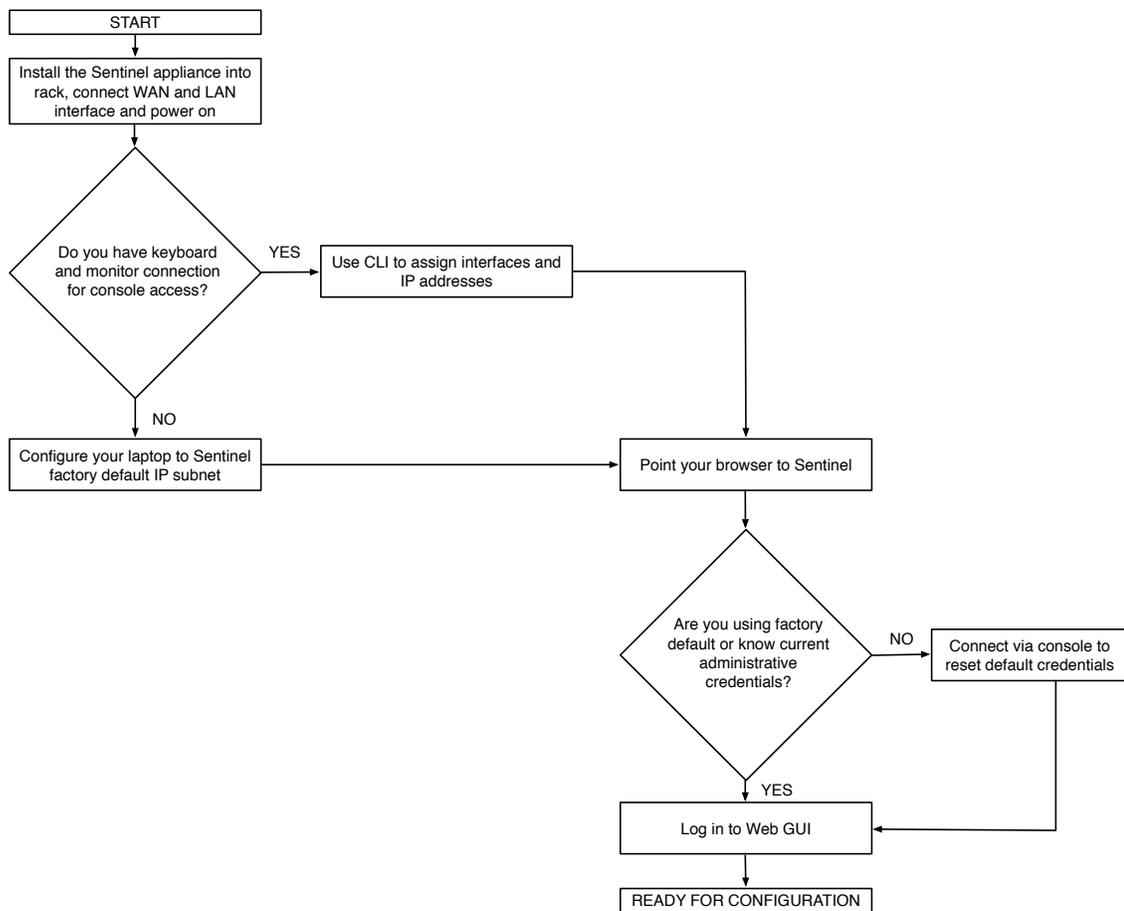


FIGURE 1–1 Installation Flow Diagram

Required Experience

It is recommended that you read through this book in order to get an overall insight on how Sentinel servers need to be configured for your particular network topology. Some basic knowledge of TCP/IP concepts, operating systems and networking is expected. Background experience with FreeBSD, Linux or similar UNIX® based operating systems is recommended for advanced configuration tasks, although it is not required for all of the Sentinel PF operating system features documented in this book.

Required Equipment and Software

You will need the following equipment and software in order to successfully deploy Sentinel into your network:

- PC with web browser software, such as Chromium, Safari, Firefox, Internet Explorer or similar. We recommend having your PC firewall switched off to make sure you have full networking functionality. SSH client software is desired.
- For Sentinel Blackbird, Westwind, Sierra and Cirrus: USB standard PC keyboard, display with DVI-I or VGA connector and cable, for console access to Sentinel. Sentinel Blackbird and Westwind servers use DVI-I connector by default.
- For Sentinel D2, D2W and Sentinel 3: RS-232 null-modem cable for console access. (See *Physical Connections* chapter and *Appendix G - Sentinel D2 RS-232 Console Cable*.)
- Two Ethernet Category 5 cables with RJ45 connectors, straight or cross-over, depending on your network topology. Having both types in advance of the installation may be useful. For 100 Mbps operation (100BaseTX), use Category 5 wiring or better. For 1000 Mbps operation (1000BaseT), use Category 5 or better (must be 4-pair wiring). Make sure you use Category 5 cable that complies with the TIA-568 wiring specification. For more information on this specification, see the Telecommunications Industry Association's website: www.tiaonline.org
- A single computer or a LAN with multiple computers to perform final routing and performance tests during the bandwidth management and advanced configuration stages.

Note: You may use factory default IP address to access Sentinel Web GUI interface via an Ethernet connection, bypassing the console access step.

To insure compliance with CISPR 24 and the EU's EN55024, devices should be used only with CAT 5E shielded cables that are properly terminated according to the recommendations in EN50174-2. If you are using Sentinel in a residential environment (at any speed), use Category 5 or better wiring. If the cable runs between rooms or through walls and/or ceilings, it should be plenum-rated for fire safety.

Blackbird and Westwind: Checking Package Contents

Carefully unpack all Sentinel Blackbird (or Westwind) components from the packing cartons. The following items are contained in the package.

| | |
|---------------------|--|
| Hardware | <ul style="list-style-type: none">▪ Server▪ Rack mount kit (attached to the server)▪ Power cable |
| Miscellaneous Items | <ul style="list-style-type: none">▪ Spare bolts (carton box) |

D2, D2W and Sentinel 3: Checking Package Contents

The following items are contained in the Sentinel D2, D2W or Sentinel 3 package.

| | |
|----------|---|
| Hardware | <ul style="list-style-type: none">▪ Server▪ Wireless Antennae (2) – optional▪ Power Supply Unit |
|----------|---|

Precautions

When unpacking and handling Sentinel, please take all precautions against electrostatic discharge as this may cause permanent damage the equipment. Any damage due to electrostatic discharge will void the product warranty. The Sentinel server must not be operated in an environment where the unit is exposed to precipitation; condensation; humid atmospheres, altitudes in excess of values specified in the hardware specification sheets, excessive dust or vibration; flammable gases, corrosive or explosive atmospheres; or extremes of temperature outside the ambient range specified. Operation in vehicles or other transportable installations that are equipped to provide a stable environment is permitted. If such vehicles do not provide a stable environment, safety of the equipment to EN 60950 may not be guaranteed.

◆ ◆ ◆ CHAPTER 2

Physical Connections

This chapter provides overview of physical connections available on different Sentinel servers.

Blackbird and Westwind: Front Panel

Sentinel Blackbird and Westwind are rack mountable servers built on the same 1.5U form factor hardware platform. Figures below provide an overview of physical connections these servers provide.



FIGURE 2-1 Sentinel Blackbird and Westwind Front Panel

Blackbird and Westwind: Motherboard

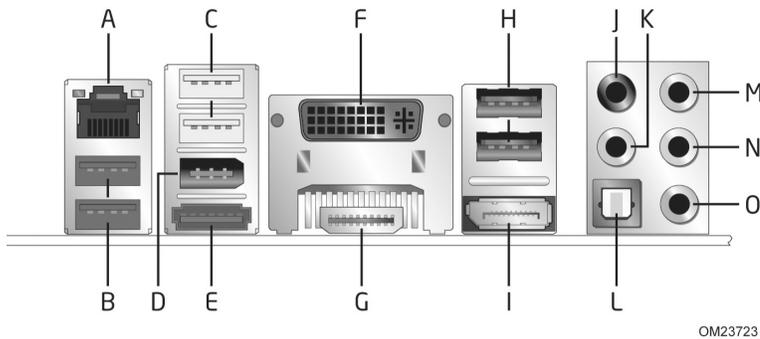


FIGURE 2-2 Sentinel Blackbird and Westwind Server Real Panel Motherboard Connectors.

Sentinel Blackbird and Westwind servers provide the following connections on the rear panel via standard motherboard:

- A – Motherboard integrated Ethernet RJ45 port, designated *em0* (1)
- B and C – USB 3.0 ports (2) and USB 2.0 high current ports (2)
- D and E – IEEE 1394a connector (1) and eSATA connector (1)
- F and G – DVI-I connector (1) and HDMI connector (1)
- H – USB 2.0 ports (2)
- I – DisplayPort connector (1)
- J, K, L, M, N and O – Sound connectors (6)

Please note that some models may use alternative connection layout. Please consult with BusinessCom representative if you are unsure.

Blackbird: LAN and WAN Ethernet Ports

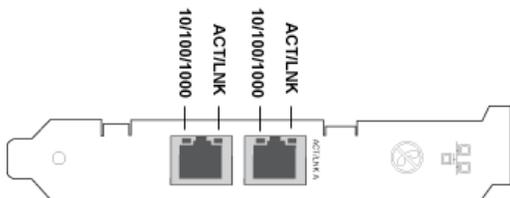


FIGURE 2-3 Sentinel Blackbird Server Intel® Dual Gigabit Ethernet NIC.

Sentinel Blackbird server features Intel® Dual Gigabit Ethernet Network Interface Card on the back panel. In the Sentinel PF operating system, the inner port, situated closer to the motherboard, is designated *igb1*, and the outer port is designated *igb0*. These ports are used to connect your LAN and WAN to Sentinel. We recommend using *igb0* for WAN and *igb1* for LAN connections, however this is optional.

Westwind: LAN and WAN Ethernet Ports

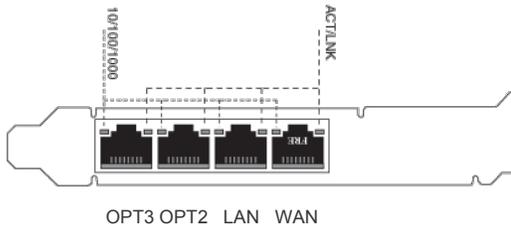


FIGURE 2-3-WESTWIND Sentinel Westwind Server Intel® Quad Gigabit Ethernet NIC.

Sentinel Westwind server features Intel® Quad Gigabit Ethernet Network Interface Card on the back panel. The ports are assigned as following by default, from left (closest to the motherboard) to right (closest to the outer edge):

- Interface designated *igb3* (OPT3)
- Interface designated *igb2* (OPT2)
- Interface designated *igb1* (LAN)
- Interface designated *igb0* (WAN)

We recommend using *igb0* for WAN and *igb1* for LAN connections, however this is optional.

Blackbird and Westwind: Ethernet Indicator Lights

The Intel® Ethernet Network Interface Cards and Motherboard Integrated Ethernet port installed in Sentinel Blackbird and Westwind models have the following indicator lights:

TABLE 2-1 Ethernet Indicator Lights

| Label | Indication | Meaning |
|-------------|----------------|--------------------------|
| ACT/LNK | Green on | The adapter is connected |
| | Green flashing | Data activity. |
| | Off | No link. |
| 10/100/1000 | Off | 10 Mbps link. |
| | Green | 100 Mbps link. |
| | Yellow | 1000 Mbps link. |

In the default mode, the Intel Ethernet NIC using copper-based connections will attempt to auto-negotiate with its link partner to determine the best setting. If the adapter cannot establish link with the link partner using auto-negotiation, you may need to manually configure the adapter and link partner to identical settings to establish link and pass packets. This should only be needed when attempting to link with an older switch that does not support auto-negotiation or one that has been forced to a specific speed or duplex mode. Your link partner must match the setting you choose.



Caution - Only experienced network administrators should force speed and duplex manually. The settings at the switch must always match the adapter settings. Adapter performance may suffer or your adapter may not operate if you configure the adapter differently from your switch.

D2, D2W: Rear Panel Motherboard

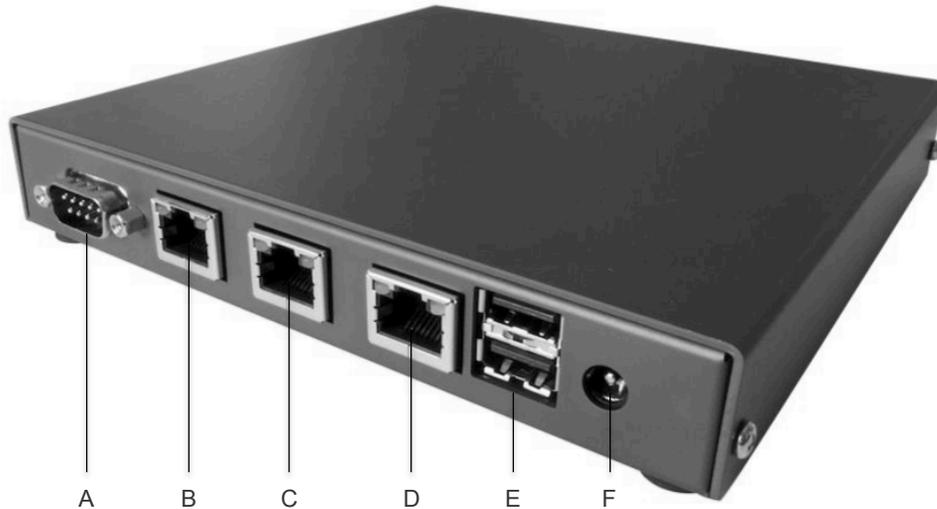


FIGURE 2-4 Sentinel D2 Real Panel Motherboard Connectors.

Sentinel D2 server provides the following connections on the rear panel via standard motherboard:

- A – RS-232 Serial Console Port, DB9 male (RXD/TXD only) (1)
- B – OPT1 (Optional) Ethernet RJ45 port, designated *vr2* (1)
- C – WAN Ethernet RJ45 port, designated *vr1* (1)
- D – LAN Ethernet RJ45 port, designated *vr0* (1)
- E – USB 2.0 ports (2)
- F – Power connector (1)

Note: Power connector F has center pin positive, sleeve is ground. Diameter is 2.1mm. Recommended power supply is 18 V DC, 15 W. 7-20 V DC range is accepted.

Serial console port settings are 9,600 baud, 8 data bits, no parity bit and 1 stop bit (8-N-1). Set flow control to *none* or *XON/XOFF* at client.

Sentinel D2W with Wireless module installed will have a pair of additional WiFi antennae SMA-type connectors not shown above. The OPT2 wireless interface is designated *ath0*.

Sentinel 3: Rear Panel Motherboard

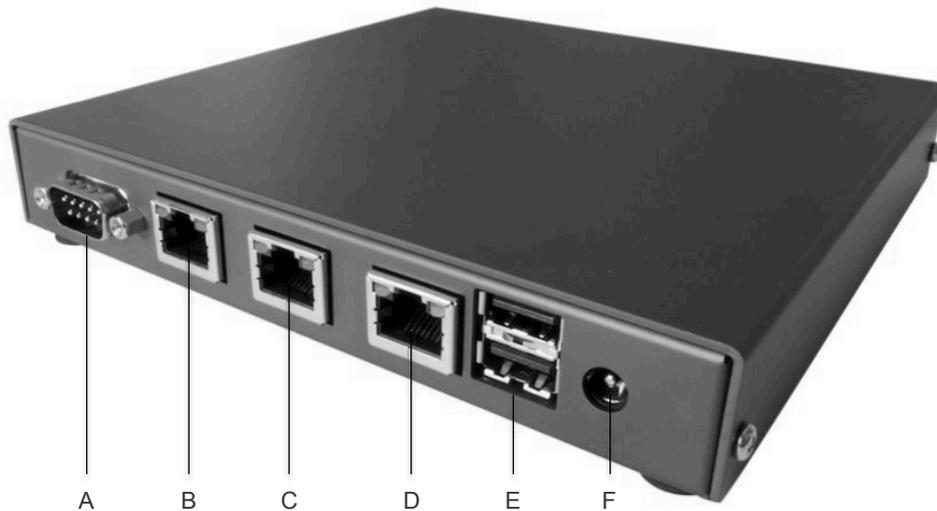


FIGURE 2-4 Sentinel 3 Real Panel Motherboard Connectors.

Sentinel D2 server provides the following connections on the rear panel via standard motherboard:

- A – RS-232 Serial Console Port, DB9 male (RXD/TXD only) (1)
- B – OPT1 (Optional) Ethernet RJ45 port, designated *re0* (1)
- C – WAN Ethernet RJ45 port, designated *re1* (1)
- D – LAN Ethernet RJ45 port, designated *re2* (1)
- E – USB 2.0 ports (2)
- F – Power connector (1)

Note: Power connector F has center pin positive, sleeve is ground. Diameter is 2.1mm. Recommended power supply is 18 V DC, 15 W. 7-20 V DC range is accepted.

Serial console port settings are 9,600 baud, 8 data bits, no parity bit and 1 stop bit (8-N-1). Set flow control to *none* or *XON/XOFF* at client.

Sentinel 3 with Wireless module installed will have a pair of additional WiFi antennae SMA-type connectors not shown above. The OPT2 wireless interface is designated *ath0*.

Powering Sentinel On and Off

▼ To Power On Sentinel

- 1 Turn on the power to the monitor and to all external devices, if necessary.
- 2 Press and release the power button on the front panel.
- 3 After several seconds, verify that the power light on the power button is lit.
The Power LED lights up once Sentinel begins the internal booting process.

On Sentinel D2 servers, a single LED lights up for approximately 1 minute which indicates a booting process. Then, three LEDs start flashing which indicates progress, and after the booting process is complete, a single LED light comes back on.

▼ To Power Off Sentinel

- 1 If you have done any changes in the configuration via Web GUI, save your settings.
- 2 Use one of the following power off options:
 - a. Use the Halt menu option in the Sentinel PF operating system CLI.
 - b. Use Halt System option in the Diagnostics menu in the Web GUI.



Caution: To avoid data loss or operating system corruption, use the steps described in a or b whenever possible.

- c. If the first step does not shut off Sentinel power, press and release the Power button and wait approximately one minute.
With Sentinel D2 and Sentinel 3 units, disconnect power supply adapter from the outlet.
- d. If the power is still on, press and hold the Power button for approximately four seconds to force shut down.
- e. If the preceding options do not power off Sentinel, disconnect the power cable from the line.

Note: After powering off Sentinel, make sure the power button light is not lit and wait at least 10 seconds before powering on Sentinel again.

Command Line Interface

This chapter provides overview of console access to Sentinel PF operating system *CLI (Command Line Interface)*.

Introduction to CLI

The CLI is a text based access to Sentinel PF operating system. It is used to assign IP addresses and roles to Ethernet NICs, halt or shut down the server, access Sentinel low-level interface for system troubleshooting and maintenance purposes, and similar tasks. The CLI can be accessed using a physical console via keyboard and display ports on the Sentinel server, or via SSH (Secure Shell) connection over TCP/IP.

The SSH connection uses public-key cryptography to authenticate the remote computer and allow it to authenticate the user, if necessary. A matching pair of different keys - public and private - is required to establish an SSH connection. The public key is placed on Sentinel that must allow access to the owner of the matching private key. The owner keeps the private key secret. While authentication is based on the private key, the key itself is never transferred through the network during authentication.

For SSH access to CLI, you will need an SSH client software installed on your computer. Many modern UNIX-like operating systems, including Apple OS X, have `ssh` client from the OpenSSH suite already pre-installed. For your convenience, a list of SSH client software is provided in the table below.

TABLE 3-1 SSH Client Software

| Name | URL | Operating Systems |
|--------------|---|-------------------------------|
| OpenSSH | http://www.openssh.com/ | UNIX-like systems, Apple OS X |
| ZOC Terminal | http://www.emtec.com/zoc/ | Apple OS X, Windows |
| PuTTY | http://www.putty.org/ | Windows |

Once Sentinel server is powered on, the CLI is automatically enabled on the physical console after the boot process is complete. The default SSH port is 22.

CLI Prompt

Once accessed via physical console or authenticated SSH connection, the Sentinel PF operating system provides the following CLI prompt (example):

```
Sentinel PF 2.0.2-RELEASE (amd64)
Tue Feb 26 16:55:00 UTC 2013

WAN (wan)           -> igb0           -> 192.168.57.6
LAN (lan)           -> igb1           -> 192.168.58.5
OPT1 (opt1)        -> em0            -> NONE

0) Logout (SSH only)
1) Assign Interfaces
2) Set interface(s) IP address
3) Reset GUI password
4) Reset to factory defaults
5) Reboot system
6) Halt system
7) Ping host
8) CLI
9) View network connections
10) Filter logs
11) Restart GUI

>
```

CLI Functionality Overview

The table below provides an overview of Sentinel PF operating system functionality available via CLI.

TABLE 3-2 CLI Menu

| CLI Menu Item | Function |
|--------------------------------|--|
| 0) Logout (SSH only) | Log out and close CLI connection. Works via SSH only. |
| 1) Assign Interfaces | Allows to set up VLANs and assign different roles to interfaces. |
| 2) Set interface(s) IP address | Set IP address for interface(s). |
| 3) Reset GUI password | Reset Web GUI password to factory default. |
| 4) Reset to factory defaults | Reset global Sentinel configuration files to factory default. |
| 5) Reboot system | Safely reboot Sentinel server. |
| 6) Halt system | Safely halt Sentinel server and power off. |
| 7) Ping host | ICMP ping tool. |
| 8) CLI | Low-level shell access to Sentinel PF operating system. |
| 9) View network connections | Use pfTop functionality to view established connections. Press q to quit the pfTop view. |
| 10) Filter logs | View Sentinel filter logs in real-time. Press ctrl+C to quit the filter log view. Stopping the filter output may take a few seconds. |
| 11) Restart GUI | Safely restart Web GUI. |

Logout (SSH Only)

This menu option closes the current SSH connection. It does not apply to physical console access.

Assign Interfaces

Assign Interfaces menu allows to assign roles to physical network interfaces and, optionally, define VLANs. In the most common Sentinel deployment scenario, one physical network interface is used for the LAN connection, and another physical interface is used for the WAN connection. The menu will guide you through assigning roles to each physical interface available in your Sentinel server.

Please consider *Chapter 2: Physical Connections* of this book for a review of physical interfaces available on your Sentinel server as well as recommended roles.

When the Assign Interfaces menu is selected, Sentinel provides a list of physical interfaces available in the server, as well as their MAC addresses, state (up or down), and corresponding NIC hardware installed, with a prompt to set up VLANs:

```
valid interfaces are:
igb0 90:e2:ba:01:7c:5a (up) Intel(R) PRO/1000 Network Connection
igb1 90:e2:ba:01:7c:5b (up) Intel(R) PRO/1000 Network Connection
em0 4c:72:b9:98:a1:67 (down) Intel(R) PRO/1000 Network Connection

Do you want to set up VLANs first?

If you are not going to use VLANs, or only for optional interfaces, you
should say no here and use the webConfigurator to configure VLANs later, if
required.

Do you want to set up VLANs now [y|n]?
```

Valid interfaces shown on your Sentinel may be different from those indicated above. You can skip the VLAN set up by pressing **n** at this stage, otherwise you will be guided through defining parent interfaces per each VLAN and assigning VLAN tags. Please refer to *Chapter 6 – Web GUI: Interfaces, (assign) section* for more information on VLANs.

Note: After VLANs are setup, only VLAN interfaces will appear in the Assign Interfaces menu.

Sentinel will guide you through assigning physical interfaces to the WAN and LAN roles first, as well as offer auto-detection functionality if **a** is entered as the physical interface name. More physical interfaces may be assigned to the *Optional* roles at this stage, such as motherboard integrated NICs or any external NICs installed. These optional interfaces can be used for advanced networking functionality, such as load balancing or fail-over, which is configured through the Web GUI. Once you have assigned all the physical interfaces in your system, simply press enter when asked to enter next interface name.

▼ To Assign Interfaces

- 1 Enter the Assign Interfaces CLI menu.
- 2 Choose whether you want to set up VLANs at this stage. Press **y** to set up VLANs. If you do not wish to set up VLANs, press **n** and skip to the step 5.
- 3 For each new VLAN, enter the parent interface name, e.g. **i gb0**. Press enter for parent interface name when all VLANs are defined.
- 4 Enter the VLAN tag.
- 5 Enter the WAN interface name, e.g. **i gb0**.
- 6 Enter the LAN interface name, e.g. **i gb1**.
- 7 Enter optional interface name, if any. Press enter for interface name when all optional interfaces are defined.
- 8 Review the interface assignments and press **y** if you would like to apply new settings, otherwise press **n**.

Set Interface(s) IP Address

This menu option is used to assign IP addresses to physical interfaces. DHCP can be used on the interface to obtain an IP address automatically.

▼ To Set Interface IP Address

- 1 Enter the Set Interface(s) IP Address CLI menu.
- 2 Select the interface to assign IP address to from the list.
- 3 Choose whether to obtain IP address via DHCP for this interface.
- 4 If DHCP is not selected, enter IPv4 address manually for this interface.
- 5 Enter subnet mask as bit counts (as in CIDR notation). E.g. entering 24 will result in 255.255.255.0 netmask selected. See *Appendix B – Netmask/CIDR Translation Table*.



Caution – If you are using a remote SSH session to assign IP address to the network interface that is used for this session then you will be disconnected once Sentinel will be restarting its filter.

Reset GUI Password/Reset To Factory Defaults

This menu option resets the Web GUI password back to the default password. Please consider *Appendix C – Factory Default Settings* for reference. The Reset To Factory Defaults resets the Sentinel PF operating system configuration back to the factory default settings.

▼ To Reset Sentinel

- 1 Enter the Reset GUI Password or Reset To Factory Defaults CLI menu.
- 2 Confirm reset by pressing y.



Caution – With Reset to factory defaults, all your configuration will be lost.

Ping host

The ping utility uses the ICMP protocol's mandatory ECHO_REQUEST datagram to elicit an ICMP ECHO_RESPONSE from a host or gateway. When using ping for fault isolation, it should first be run on the local host, to verify that the local network interface is up and running. Then, hosts and gateways further and further away should be “pinged”. Round-trip times and packet loss statistics are computed. If duplicate packets are received, they are not included in the packet loss calculation, although the round trip time of these packets is used in calculating the round-trip time statistics. When the specified number of packets have been sent (and received) or if the program is terminated, a brief summary is displayed, showing the number of packets sent and received, and the minimum, mean, maximum, and standard deviation of the round-trip times.

CLI

This menu option provides low-level shell access to Sentinel PF operating system based on the FreeBSD kernel. When entered, the CLI presents the user with the # prompt indicating *root* level access to the operating system.

Type **exit** to return to standard user-level CLI.



Caution – Low-level CLI access is not required and shall not be used for any bandwidth management purposes. Sentinel provides low-level access for remote troubleshooting by BusinessCom staff only.

View Network Connections

This menu option provides text based overview of network connections generated by the *pfTop* utility. *pfTop* displays the active packet filter states and rules, and periodically updates this information. If standard output is an intelligent terminal then as many states as will fit on the terminal screen are displayed by default. Otherwise, a good number of them are shown (around 20). If number is given, then the top number states will be displayed instead of the default. The displayed states are filtered according to the filter specification. *pfTop* reads commands from the terminal and acts upon them accordingly. A character will be processed as soon as it is typed. The command will be processed and the display will be updated immediately thereafter (reflecting any changes that the command may have triggered). If a key is pressed while *pfTop* is in the middle of updating the display, it will finish the update and then process the command. These commands are currently recognized:

TABLE 3-3 *pfTop* Commands

| Character | Command |
|-----------|--|
| c | Enable/disable state caching (enabled by default). |
| f | Set the state filter expression. |
| h,? | Display a summary of the commands (help screen). |
| n | Set number of lines to display. |
| o | Select next sorting Order. |
| p | Pause/resume display updates. |
| q | Quit <i>pfTop</i> . |
| r | Reverse current sorting order. |
| s | Set display update interval in Seconds. |
| v | Select next View. |
| 0-8 | Select one of the views directly. |
| Cursor | Scroll display (up/down), and switch views (left/right). |
| SPACE | Update display immediately. |
| CTRL-L | Refresh display. |
| CTRL-G | Clear command entry line. |

The following keys are shortcuts for sorting the display – note the upper case:

| | |
|---|---|
| A | Sort states by Age. |
| B | Sort states by number of Bytes. |
| D | Sort by Destination port. |
| E | Sort states by Expiry time. |
| F | Sort by source address (From). |
| K | Sort by peak speed when caching is enabled. |
| N | No ordering. |
| P | Sort states by the number of Packets. |
| R | Sort by instantaneous speed (Rate) when caching is enabled. |
| T | Sort by destination address (To). |

Please refer to the *Appendix A – State Filter Expressions* for advanced filtering in the *pfTop* output.

The pfTop tool offers multiple views into network connections. You can select pfTop views by pressing any number from 0 to 8. The default view provides information about protocol, direction, source and destination IP addresses, connection state, age, expiration, as well as passed packets and bytes.

Views 1 to 4 and 7 provide the same information as the default view, with different column layouts in the output for your convenience. Views 5 and 6 provide the same traffic overview in the context of filtering rules defined, and their labels. View 8 provides information about queues.

The following table provides a description of the pfTop labels used.

TABLE 3-4 pfTop View Lables

| Label | Meaining |
|-------|-----------------------------------|
| PR | Protocol |
| DIR | Direction |
| SRC | Source IP Address |
| DEST | Destination IP Address |
| GW | Gateway |
| STATE | Connection State |
| AGE | Connection Age |
| EXP | Connection Expires |
| PKTS | Packets |
| BYTES | Bytes |
| RATE | Rate (packets per second) |
| PEAK | Peak rate (packets per second) |
| AVG | Average rate (packets per second) |

The following labels are specific to filtering rules views 5 and 6

| | |
|--------|--|
| RULE | Applicable filtering rule number |
| ACTION | Filtering action undertaken |
| LOG | Set to «Log» if filtering rule is logged |
| Q | Queued |
| IF | Physical interface |
| STATES | Amount of states/connections per rule |
| INFO | Filtering rule expression |
| LABEL | Rule label/description |

The following labels are specific to queues view 8

| | |
|--------|--------------------|
| BW | Bandwidth |
| SCH | Schedule |
| PRIO | Priority |
| DROP_P | Dropped packets |
| DROP_B | Dropped bytes |
| QLEN | Queue length |
| BORROW | Borrowed bandwidth |
| SUSPEN | Suspends |
| P/S | Packets per second |
| B/S | Bytes per second |

The following example provides some of the popular pfTop usage scenarios.

▼ To Show Connections With Most Packet Rate

- 1 Enter the View Network Connections CLI menu.
- 2 Enter **R** to sort connections by packet rate. Note you need to press **Shift+r**.
- 3 When finished, press **q** to return to the CLI menu.

▼ To Show Connections With Most Traffic

- 1 Enter the View Network Connections CLI menu.
- 2 Enter **B** to sort connections by number of bytes. Note you need to press **Shift+b**.
- 3 When finished, press **q** to return to the CLI menu.

▼ To Show Oldest Connections

- 1 Enter the View Network Connections CLI menu.
- 2 Enter **A** to sort connections by age. Note you need to press **Shift+a**.
- 3 When finished, press **q** to return to the CLI menu.

Filter Logs

This CLI menu enables real-time output of filtering logs. Once enabled, the CLI will start output of logs for all filtering rules that have logging enabled. Press **Ctrl+c** to return back to CLI.

Restart GUI

This CLI menu restarts Sentinel Web GUI. You may find it useful in case the Web GUI becomes unresponsive.

Web GUI

This chapter provides overview of Web GUI access to Sentinel PF operating system.

Introduction to Web GUI

The Web GUI is the primary tool for advanced Sentinel PF operating system configuration. It provides much more extensive and user-friendly configuration environment, if compared to CLI. Most of the bandwidth management and advanced network features are configured via Web GUI.

The Web GUI is automatically enabled by default.

▼ To Connect to Web GUI

- 1 **Make sure you have physical connection to Sentinel (either LAN or WAN ports).**

Note: Sentinel may not reply to ICMP pings for security purposes. If the below access procedure does not work, make sure Sentinel LAN and WAN interfaces are assigned and configured, as described in *Chapter 3 – Command Line Interface*.

- 2 **Point your web browser to Sentinel IP address and Web GUI port.**
- 3 **Authenticate with your username and password.**
If this is the initial Sentinel configuration, please consider *Appendix C – Factory Default Settings* for the default username and password.

Once you are logged in, you will be presented with the Sentinel Dashboard that provides an overview of system resources, traffic graphs, Sentinel PF network interface and services status and Network Intrusion Detection System (NIDS) alarms.

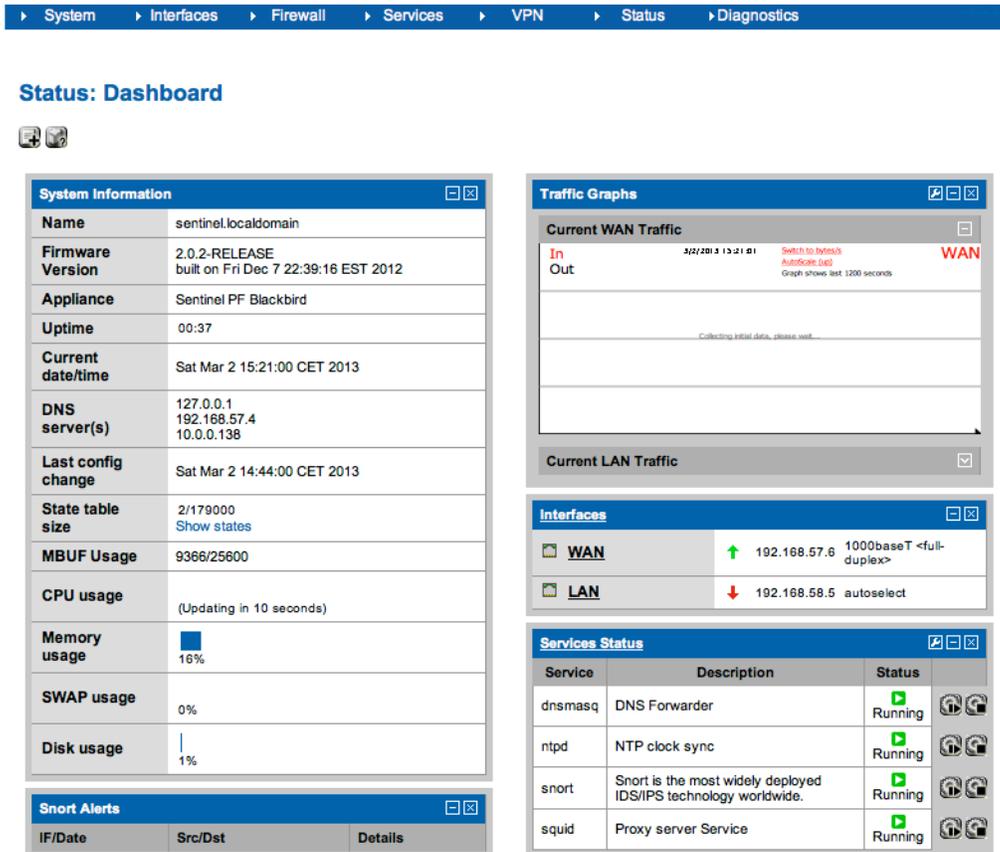


FIGURE 4-1 Sentinel Blackbird Web GUI Dashboard.

These are the standard dashboard sections:

System Information

Provides information on the state of the Sentinel PF operating system and hardware usage.

Snort Alerts

Latest Snort Network Intrusion Detection System (NIDS) Alerts

Traffic Graphs and Interface

Real-time traffic graphs on WAN and LAN interfaces, as well as interface status.

Services Status

Shows which Sentinel PF services are running.

Initial Configuration Steps

If this is your initial Sentinel deployment, we advise to walk through the following Web GUI menus and configure essential Sentinel functionality prior to any advanced networking and bandwidth management configuration.

▼ First Steps

- 1 **In System / General Setup Menu:**
 - a. **Setup DNS servers to use.**
 - b. **Set the correct time zone.**

See *Chapter 5 – General Setup*.



Caution: With an incorrect time zone setup, your bandwidth charts and NIDS alarms will have wrong time frames.

- 2 **In System / Routing Menu, add the new WAN gateway that will be used as the default gateway to the WAN.**

See *Chapter 5 – To Add New Gateway*.
- 3 **In Interfaces / WAN, make sure your default WAN gateway is selected in the *Static IP Configuration* section for this interface.**

Web GUI: System

System menu items are used to configure Sentinel system-level settings, such as hostname, domain, DNS servers, time zone, default gateways and routing tables, as well as manage software packages, user accounts and PKI certificates.

Cert Manager

The Cert Manager menu allows to manage Sentinel PF public key infrastructure (PKI). It is based on the X.509, an ITU-T standard for PKI and Privilege Management Infrastructure (PMI). X.509 specifies, amongst other things, standard formats for public key certificates, certificate revocation lists, attribute certificates, and a certification path validation algorithm. The most common use for X.509 keys within Sentinel PF environment would be authentication for VPN connections. With PKI, a *Certificate Authority (CA)* is created that is used to sign individual *certificates*. The CA's certificate is used on VPN connections to verify authenticity of certificates used. Signing individual certificates requires knowledge of CA's certificate private key. This private key is the cornerstone of your PKI, thus it must be kept secure.

▼ To Create New Certificate Authority (CA)

- 1 In System / Cert Manager Menu, select CAs tab and click on the  button in order to create new CA.
- 2 Enter descriptive name for the CA (e.g. PrimaryCA).
- 3 Select «Create an internal Certificate Authority» as the method.
- 4 Specify key length and lifetime of the CA.
- 5 Provide Distinguished Name data for this CA.
This includes country code, state or province, city, organization, e-mail address for the administrator and CA common name.
- 6 Click «Save».

Individual user and server certificates are created in a similar fashion:

▼ To Create New User or Server Certificate

- 1 In System / Cert Manager Menu, select Certificates tab and click on the  button in order to create a new certificate.
- 2 Enter descriptive name for the certificate (e.g. VPNServerCert).
- 3 Select «Create an internal Certificate » as the method.
- 4 Select the Certificate Authority (CA) that will be signing this certificate.
- 5 Specify Certificate Type.
- 6 Specify key length and lifetime of the CA.
- 7 Provide Distinguished Name data for this CA if it is different from the CA data.
This includes country code, state or province, city, organization, e-mail address for the administrator and CA common name.
- 8 Click «Save».

It is also possible to re-use existing certificates by specifying *Import an existing Certificate* (or *CA*) as the method when creating new certificates in Sentinel PF. The certificate and private key data can then be pasted in X.509 PEM format, typically found in *.crt* files. Example of X.509 PEM certificate data is as following:

```
-----BEGIN CERTIFICATE-----
base-64 encoded data follows
-----END CERTIFICATE-----
```

Private key files are usually found in *.key* files, with format looking as following:

```
-----BEGIN RSA PRIVATE KEY-----
base-64 encoded data follows
-----END RSA PRIVATE KEY -----
```

The *Certificate Revocation* tab allows to create CRL (Certificate Revocation Lists) and point to certificates that have been compromised. Revoking a certificate will cause it to be considered untrusted, as long as the application using CA is also using CRL. All CRLs must be signed by a CA.

General Setup

The General Setup menu allows configuration of essential Sentinel features.

| Menu Item | Description |
|---|---|
| Hostname | Name of the Sentinel host, without the domain part. |
| Domain | Name of the domain. Do not use 'local' as a domain name. It will cause local hosts running mDNS (avahi, bonjour, etc.) to be unable to resolve local hosts not running mDNS. <i>e.g. mycorp.com, home, office, private, etc.</i> |
| DNS Servers | Enter IP addresses to be used by the system for DNS resolution. These are also used for the DHCP service, DNS forwarder and for PPTP VPN clients. In addition, optionally select the gateway for each DNS server. When using multiple WAN connections there should be at least one unique DNS server per gateway. |
| Allow DNS server list to be overridden by DHCP/PPP on WAN | |
| If this option is set, Sentinel PF will use DNS servers assigned by a DHCP/PPP server on WAN for its own purposes (including the DNS forwarder). However, they will not be assigned to DHCP and PPTP VPN clients. | |
| Do not use the DNS Forwarder as a DNS server for the firewall | |
| By default localhost (127.0.0.1) will be used as the first DNS server where the DNS forwarder is enabled, so system can use the DNS forwarder to perform lookups. Checking this box omits localhost from the list of DNS servers. | |
| Timezone | Time zone configuration. Set to the location closest to you. |
| NTP Time Server | Configure which NTP (Network Time Protocol) servers to use for NTP synchronization. |

Packages

The Sentinel PF operating system uses external software packages to provide some advanced network services:

TABLE 5-1 Sentinel PF External Packages

| Package Name | Description |
|-------------------------|--|
| Network Monitor | Network Monitor tracks usage of TCP/IP network subnets and builds HTML files with graphs to display utilization. Charts are built by individual IPs, and by default display utilization over 2 day, 8 day, 40 day, and 400 day periods. Furthermore, each ip address's utilization can be logged out at intervals of 3.3 minutes, 10 minutes, 1 hour or 12 hours in cdf format, or to a backend database server. HTTP, TCP, UDP, ICMP and VPN traffic are color coded. |
| Dashboard Widget: Snort | Dashboard widget for displaying Snort NIDS (Network Intrusion Detection System) alerts. |
| Huginn | Huginn is BusinessCom VSAT modem monitoring tool. Currently compatible with iDirect Evolution series modems. (<i>Status / iDirect Monitor</i> menu). |
| LightSquid | High performance web proxy report (LightSquid). Proxy realtime statistics (SQStat). Requires Squid HTTP proxy. |
| PEP | BusinessCom PEP (Performance Enhancing Proxy) client. PEP is a WAN optimization solution. The current protocol version implemented is 4.2.0. PEP client establishes a connection with BusinessCom PEP gateway and opens a local HTTP proxy or TCP SOCKS port to route traffic through. |
| snort | Snort® is an external open source package developed by Sourcefire, Inc. available for Sentinel PF. It provides NIDS (Network Intrusion Detection System) and IPS (Intrusion Prevention System) functionality and enables real-time traffic analysis of IP networks for malicious activity. |
| squid | High performance web proxy cache. |
| Tinyproxy | A light-weight HTTP/HTTPS proxy daemon. |
| | Note: Tinyproxy is recommended for Sentinel D2 and D2W servers instead of the default Squid, especially for demanding network environments. This is non-caching proxy that is conservative on RAM usage. |

Standard Sentinel PF distribution comes with no pre-installed external packages. External packages can be installed at any time from the *Available Packages* tab. Sentinel PF downloads the software package from a software repository hosted by BusinessCom Networks, and the installation process is automatic.

Additional actions can be taken on the packages from the *Installed Packages* tab:

Click  button to remove package from Sentinel PF system.

Click  button to re-install a package.

Click  button to re-install Web GUI package components only. These include Web GUI XML files.

Note: BusinessCom PEP service is *not* included in Sentinel PF license, technical support plans or hardware costs. Please contact BusinessCom Networks Sales Department if you require PEP service.

Routing

For one machine to be able to find another over a network, there must be a mechanism in place to describe how to get from one to the other. This is called routing. A “route” is a defined pair of addresses: a “destination” and a “gateway”. The pair indicates that if you are trying to get to this destination, communicate through this gateway. There are three types of destinations: individual hosts, subnets, and “default”. The “default route” is used if none of the other routes apply. When the local system needs to make a connection to a remote host, it checks the routing table to determine if a known path exists. If the remote host falls into a local subnet attached to an interface, then the system checks to see if it can connect along that interface. If all known paths fail, the system has one last option: the “default” route. This route is a special type of gateway route (usually the only one present in the system). For hosts on a local area network, this gateway is set to whatever machine has a direct connection to the outside world (whether via PPP link, DSL, cable modem, T1, or another network interface). If you are configuring the default route for Sentinel which itself is functioning as the gateway to the outside world, then the default route will be the gateway device at your Internet Service Provider's (ISP) site, or your satellite modem for example. The Sentinel Web GUI routing menu allows to set up gateways, define static routes for selected networks as well as create gateway groups in order to enable load balancing, failover, or policy-based routing.

▼ To Add New Gateway

- 1 In System / Routing Menu, select Gateways tab and click on the  button in order to create new gateway.
- 2 Select interface used to reach the new gateway.
- 3 Enter gateway name (e.g. WANGW) and IP address.
- 4 If this gateway will be used as the default gateway for this Sentinel server, select the *Default Gateway* checkbox.
- 5 Check whether you want to disable gateway monitoring for this gateway.

Note: Gateway monitoring will trigger Sentinel to send periodic ICMP ping requests in the background in order to measure packet loss and latency between Sentinel and the gateway host. This is used for load balancing and failover purposes. Gateway monitoring will result in small amounts of background traffic.

- 6 Enter optional gateway description (e.g. satellite Modem).

When editing the gateway, you can specify advanced settings by clicking on the *Advanced* button. This includes specifying weight, latency and packet loss thresholds, monitoring probe frequency and the number of bad probes for the interface to be considered down.

Note: When using gateway monitoring to probe gateways located over long distance WAN links, such as geostationary satellite connections, we recommend setting latency thresholds to be from 600 to 1,000 ms. Some TDMA based satellite links may exceed the 1,000 ms mark, especially when ICMP traffic is not prioritized over the satellite link.

▼ To Add New Static Route

- 1 In System / Routing Menu, select Routes tab and click on the  button in order to create new route.
- 2 Enter destination network.
- 3 Select gateway.
- 4 Enter optional route description. (e.g. Private Network 1)
- 5 Click Save.

Note: Do not enter static routes for networks assigned on any interface of Sentinel. Static routes are only used for networks reachable via a different router, and not reachable via your default gateway.

The groups tab allows to create a group of gateways for load balancing and failover purposes. The group will consist of multiple individual gateways that can be assigned a tiered priority structure. The priority selected defines in what order failover and balancing of the gateway links will be done. Multiple links of the same priority will balance connections until all links in the priority will be exhausted. If all links in a priority level are exhausted, Sentinel will use the next available link(s) in the next priority level.

Additionally, a trigger level is defined in order to trigger exclusion of a gateway member from the group by selected criteria, such as packet loss, high latency, or both. These criteria are defined in the *Gateway Monitoring* setup of each individual gateway.

▼ To Create New Gateway Group

- 1 In System / Routing Menu, select Groups tab and click on the  button in order to create new gateway group.
- 2 Enter group name.
- 3 Assign priorities to individual gateways. Use «Never» if you do not want any gateway to be used in this group. Choose trigger level.
- 4 Enter the optional gateway group description (e.g. All Internet Links).
- 5 Click Save

User Manager

Sentinel PF allows multiple users to use it at the same time. While only one user can sit in front of the physical screen and use the keyboard at any one time, any number of users can log in to the system through the network. To use the system, every user must have a user account. Since all access to the Sentinel PF system is achieved via accounts and all processes are run by users, user and account management is important. Every account on a Sentinel PF system has certain information associated with it to identify the account:

User name

The user name is typed at the username: prompt. User names must be unique on the system as no two users can have the same user name. Typically user names consist of eight or fewer all lower case characters.

Password

Each account has an associated password.

Full Name

Optional field – for informative purposes only.

Expiration Date

By default Sentinel PF does not expire accounts. When creating accounts that need a limited lifespan, such as student accounts, specify the account expiry date. After the expiry time has elapsed, the account cannot be used to log in to the system, although the account's directories and files will remain.

Group Memberships

Used to uniquely identify primary group that the user belongs to. Groups are a mechanism for controlling access to resources based on a user's group membership. A user may also be a member of more than one group.

Authorized Keys

If this user will be accessing the system with SSH, and you want them to use key-based authentication instead of password-based, then paste their public ssh key into the *IPsec Pre-Shared Key* box.

By default, there are two user groups defined – *all* and *admins*. The *all* group does not provide any privileges to a user. It is simply a placeholder for any user. The *admins* group, on the other hand, grants access to all Sentinel PF functionality. The factory default user *admin* is in the *admins* group, however you may want to create custom groups to access only specific Sentinel PF functionality. This could be used, for example, to create observer level access to your Sentinel PF for remote monitoring purposes by outsourced employees.

The Users tab allows to create, modify and delete users in the Sentinel PF operating system.

▼ To Create New User

- 1 In System / User Manager, select Users tab and click on the  button in order to create new user.
- 2 Enter user name, password and, optionally, full name and expiration date for this account.
- 3 Assign this user to one or multiple groups by clicking  and  buttons.
- 4 If this user will be using IPsec pre-shared key for SSH authentication, paste the key to the IPsec pre-shared key textbox.
- 5 Click Save.

▼ To Create New Group

- 1 In System / User Manager, select Groups tab and click on the  button in order to create new group.
- 2 Enter group name, (e.g. `observers`) and description (e.g. `Diagnostics Access Only`).
- 3 Assign users to this group by clicking  and  buttons.
- 4 Click Save.

▼ To Assign Privileges To A Group

- 1 In System / User Manager, select Groups tab and click on the  button in order to edit a group.
- 2 In «Assigned Privileges» section, click the  button to add a privilege.
- 3 Select one or multiple privileges and click Save.
- 4 In the group menu, click Save to apply new settings.

The Settings tab allows to change miscellaneous settings for user authentication. By default, Sentinel PF expires all management sessions in 4 hours after the log in. You can change this by entering value in minutes in the *Session Timeout* page.

In the factory default configuration, Sentinel PF uses local database to authenticate all users. Sentinel PF can be configured to use LDAP (*Lightweight Directory Access Protocol*) and RADIUS (*Remote Authentication Dial In User Service*) as the authentication mechanism.

▼ To Use LDAP or RADIUS As Authentication

- 1 In System / User Manager, select Servers tab and click on the  button in order to add a new authentication server.
- 2 Enter your authentication server data and click Save.
- 3 In System / User Manager, select Settings tab and select your newly created server in the *Authentication server* drop-down menu.

Web GUI: Interfaces

Interfaces menu items are used to assign and configure Sentinel network interfaces, interface groups, VLANs, QinQ tagging, PPP, GRE, GIF functionality, bridges and link aggregation, also known as *LAGG*.

(assign)

The interface assignments tab in the (assign) menu allows to assign roles to physical LAN and WAN interfaces and configure additional interfaces. The interface assignments tab is a GUI frontend to the interface assignment, as described in the *Chapter 3 – CLI, Assign Interfaces* section, thus we will avoid duplicating description of this tab here.

Interface groups tab allow you to create rules that apply to multiple interfaces without duplicating the rules. If you remove members from an interface group, the group rules no longer apply to that interface.

▼ To Create Interface Group

- 1 In Interfaces / (assign), select Interface Groups tab and click on the  button in order to create a new interface group.
- 2 Enter interface group name, e.g. WAN interfaces.
- 3 Enter optional interface group description.
- 4 In the members section, click on the  button to add interfaces to the group.
- 5 Click Save.

The Sentinel PF operating system also supports wireless interfaces and can act as both wireless client and access point. The wireless tab in the (assign) menu allows to configure clones of wireless interfaces, which can be assigned as separate independent interfaces.

Most wireless networks are based on the IEEE® 802.11 standards. A basic wireless network consists of multiple stations communicating with radios that broadcast in either the 2.4GHz or 5GHz band (though this varies according to the locale and is also changing to enable communication in the 2.3GHz and 4.9GHz ranges).

802.11 networks are organized in two ways: in infrastructure mode one station acts as a master with all the other stations associating to it; the network is known as a BSS and the master station is termed an access point (AP). In a BSS all communication passes through the AP; even when one station wants to communicate with another wireless station messages must go through the AP. In the second form of network there is no master and stations communicate directly. This form of network is termed an IBSS and is commonly known as an ad-hoc network.

802.11 networks were first deployed in the 2.4GHz band using protocols defined by the IEEE 802.11 and 802.11b standard. These specifications include the operating frequencies, MAC layer characteristics including framing and transmission rates (communication can be done at various rates). Later the 802.11a standard defined operation in the 5GHz band, including different signalling mechanisms and higher transmission rates. Still later the 802.11g standard was defined to enable use of 802.11a signalling and transmission mechanisms in the 2.4GHz band in such a way as to be backwards compatible with 802.11b networks.

Separate from the underlying transmission techniques 802.11 networks have a variety of security mechanisms. The original 802.11 specifications defined a simple security protocol called WEP. This protocol uses a fixed pre-shared key and the RC4 cryptographic cipher to encode data transmitted on a network. Stations must all agree on the fixed key in order to communicate. This scheme was shown to be easily broken and is now rarely used except to discourage transient users from joining networks. Current security practice is given by the IEEE 802.11i specification that defines new cryptographic ciphers and an additional protocol to authenticate stations to an access point and exchange keys for doing data communication. Further, cryptographic keys are periodically refreshed and there are mechanisms for detecting intrusion attempts (and for countering intrusion attempts). Another security protocol specification commonly used in wireless networks is termed WPA. This was a precursor to 802.11i defined by an industry group as an interim measure while waiting for 802.11i to be ratified. WPA specifies a subset of the requirements found in 802.11i and is designed for implementation on legacy hardware. Specifically WPA requires only the TKIP cipher that is derived from the original WEP cipher. 802.11i permits use of TKIP but also requires support for a stronger cipher, AES-CCM, for encrypting data. (The AES cipher was not required in WPA because it was deemed too computationally costly to be implemented on legacy hardware.)

Some Sentinel servers with wireless interfaces support networks that operate using 802.11a, 802.11b, and 802.11g. The WPA and 802.11i security protocols are likewise supported (in conjunction with any of 11a, 11b, and 11g) and QoS and traffic prioritization required by the WME/WMM protocols are supported for a limited set of wireless devices.

Note: Not all Sentinel hardware servers have wireless interfaces. See the *Appendix D – Additional Support* for information on supported servers.

▼ To Use Sentinel As WPA Access Point

- 1 In Interfaces / (assign), select Wireless tab and click on the  button in order to create a wireless interface clone
- 2 Select physical wireless interface as the parent interface. Set mode to Access Point.
- 3 Set SSID for your access point.
- 4 Check «Enable WPA»
- 5 Enter the PSK (8-63 characters).
- 6 Set WPA mode to WPA.
- 7 Set WPA Key Management Mode to Pre-shared Key
- 8 Set Authentication to Open System Authentication
- 9 Set WPA Pairwise to AES
- 10 Click Save.

▼ To Use Sentinel As WPA Client

- 1 In Interfaces / (assign), select Wireless tab and click on the  button in order to create a wireless interface clone
- 2 Select physical wireless interface as the parent interface. Set mode to Infrastructure.
- 3 Set SSID to the SSID of access point to connect.
- 4 Check «Enable WPA»
- 5 Enter the PSK used by the access point.
- 6 Set WPA mode to WPA.
- 7 Set WPA Key Management Mode to Pre-shared Key
- 8 Set Authentication to Open System Authentication
- 9 Set WPA Pairwise to AES
- 10 Click Save.

The VLANs tab is used to create VLANs. In computer networking, a single Layer 2 network may be partitioned to create multiple distinct broadcast domains, which are mutually isolated so that packets can only pass between them via one or more routers; such a domain is referred to as a Virtual Local Area Network, Virtual LAN or VLAN.

Simpler devices only support partitioning on a port level (if at all), so sharing VLANs across devices requires running dedicated cabling for each VLAN. More sophisticated devices can mark packets through tagging, so that a single interconnect (trunk) may be used to transport data for various VLANs.

Grouping hosts with a common set of requirements regardless of their physical location by VLAN can greatly simplify network design. A VLAN has the same attributes as a physical local area network (LAN), but it allows for end stations to be grouped together more easily even if they are not on the same network switch. VLAN membership can be configured through software instead of physically relocating devices or connections. Most enterprise-level networks today use the concept of virtual LANs. Without VLANs, a switch considers all interfaces on the switch to be in the same broadcast domain.

To physically replicate the functions of a VLAN would require a separate, parallel collection of network cables and equipment separate from the primary network. However, unlike physically separate networks, VLANs share bandwidth, so VLAN trunks may require aggregated links and/or quality of service prioritization.

Network architects set up VLANs to provide the segmentation services traditionally provided by routers in LAN configurations. VLANs address issues such as scalability, security, and network management. Routers in VLAN topologies provide broadcast filtering, security, address summarization, and traffic-flow management. By definition, switches may not bridge IP traffic between VLANs as doing so would violate the integrity of the VLAN broadcast domain.

VLANs can also help create multiple Layer 3 networks on the same Layer 2 switch. For example, if a DHCP server is plugged into a switch it will serve any host on that switch that is configured to get its IP from a DHCP server. By using VLANs you can easily split the network up so some hosts won't use that DHCP server and will obtain link-local addresses, or obtain an address from a different DHCP server. Hosts may also use a DNS server if a DHCP is not available.

VLANs are Layer 2 constructs, compared with IP subnets, which are Layer 3 constructs. In an environment employing VLANs, a one-to-one relationship often exists between VLANs and IP subnets, although it is possible to have multiple subnets on one VLAN. VLANs and IP subnets provide independent Layer 2 and Layer 3 constructs that map to one another and this correspondence is useful during the network design process. By using VLANs, one can control traffic patterns and react quickly to relocations. VLANs provide the flexibility to adapt to changes in network requirements and allow for simplified administration.

Note: Not all Sentinel hardware servers support 802.1Q VLAN tagging properly. On cards that do not explicitly support it, VLAN tagging will still work, but the reduced MTU may cause problems. See the *Appendix D – Additional Support* for information on supported servers.

▼ To Create a VLAN

- 1 In **Interfaces / (assign)**, select **VLANs** tab and click on the  button in order to create a new VLAN.
- 2 Select parent interface.
- 3 Enter the VLAN tag.
- 4 Enter optional VLAN description, e.g. **Office VLAN**.
- 5 Click **Save**.

The QinQ tab is used to create QinQ interfaces that have nested 802.1Q VLAN tags. The IEEE 802.1QinQ standard is an amendment to IEEE standard IEEE 802.1Q-1998. The technique is also known as provider bridging, Stacked VLANs or simply QinQ.

The original 802.1Q specification allows a single VLAN header to be inserted into an Ethernet frame. QinQ allows multiple VLAN headers to be inserted into a single frame, an essential capability for implementing Metro Ethernet network topologies. In a multiple VLAN header context, out of convenience the term "VLAN tag" or just "tag" for short is often used in place of "802.1Q VLAN header". QinQ allows multiple VLAN tags in an Ethernet frame; together these tags constitute a tag stack. When used in the context of an Ethernet frame, a QinQ frame is a frame that has 2 VLAN 802.1Q headers (double-tagged). 802.1QinQ specifies architecture and bridge protocols to provide separate instances of the MAC services to multiple independent users of a Bridged Local Area Network in a manner that does not require cooperation among the users, and requires a minimum of cooperation between the users and the provider of the MAC service.

The idea is to provide, for example, the possibility for customers to run their own VLANs inside service provider's provided VLAN. This way the service provider can just configure one VLAN for the customer and customer can then treat that VLAN as if it was a trunk.

▼ To Create QinQ Interface

- 1 In Interfaces / (assign), select QinQ tab and click on the  button in order to create a new interface.
- 2 Select parent interface.
- 3 Enter first level VLAN tag.
- 4 Enter optional QinQ description, e.g. Primary QinQ.
- 5 In the members section, click on the  button to add tags or tag ranges to this QinQ.
- 6 Click save.

The PPPs tab allows to configure PPP (Point-to-Point Protocol) interfaces. This will commonly be used to drive 3G modems for wireless WAN access, but can also be used for controlling dialup modems. The parent interface of a PPP interface is a serial device. There are some pre-defined connection profiles available for cellular modems from AT&T and Sprint, click the links to fill in the information automatically.

Please note that Sentinel PPP support also includes PPPoE (Point-to-Point Protocol over Ethernet), PPTP (Point to Point Tunneling Protocol) and L2TP (Layer 2 Tunneling Protocol) with parent interfaces being Ethernet physical network interfaces.

▼ To Create PPP Interface

- 1 **In Interfaces / (assign), select PPPs tab and click on the  button in order to create a new interface.**
- 2 **Select link type and parent interface.**
- 3 **Enter optional PPP interface description, e.g. 3G Connection.**
- 4 **Enter interface data.**
For L2TP and PPTP connections, you will need to enter username, password as well as local and gateway IP addresses.

Note: PPP interface configuration, if used in Serial PPP and PPPoE modes should be requested from your Internet Service Provider.

- 5 **Click save.**

Note: Not all Sentinel hardware servers have serial interfaces for PPP connections. See the *Appendix D – Additional Support* for information on supported servers.

The GRE tab allows to configure GRE tunnels. Generic Routing Encapsulation (GRE) is a tunneling protocol developed by Cisco Systems that can encapsulate a wide variety of network layer protocols inside virtual point-to-point links over an Internet Protocol network. Outgoing datagrams are encapsulated by an IP header of protocol type 47, and a GRE header specifying the type of the encapsulated datagram (currently only IP). This mode is described in *RFC 1702*. It's also the default mode on Cisco routers. With mobile encapsulation, outgoing IP datagrams are encapsulated by a smaller header, and the original IP header is modified. This mode is described in *RFC 2004*.

▼ To Create GRE Tunnel

- 1 In Interfaces / (assign), select GRE tab and click on the  button in order to create a new GRE tunnel.
- 2 Select parent interface.
- 3 Enter IP addresses for: GRE remote address, GRE tunnel local address and GRE tunnel remote address.
Please consider figure 5-1 below for additional details on addressing.
- 4 Check additional options, if required, such as Mobile tunnel, Route search type and WCCP version and add optional tunnel description.
- 5 Click Save.

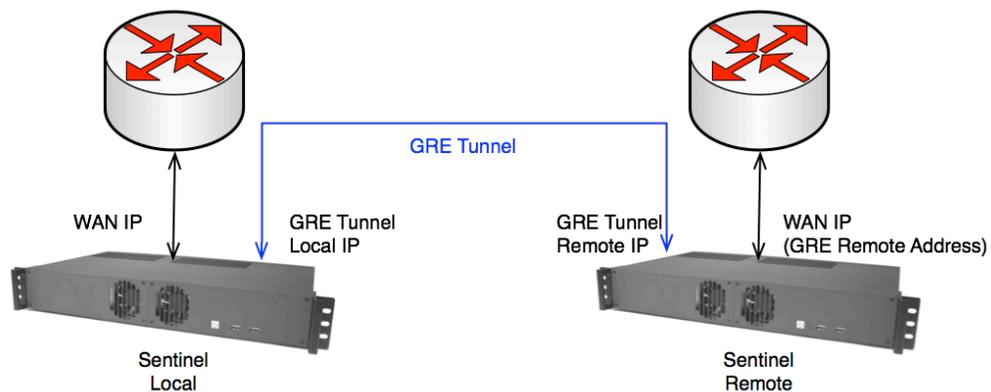


FIGURE 6-1 GRE Tunnel Addressing.

The GIF tab allows to configure GIF (Generic Tunnel Interface) tunnel interfaces. The GIF interface is a generic tunnelling device for IPv4 and IPv6 traffic. It can tunnel IPv4 or IPv6 traffic over IPv4 or IPv6 connection. Therefore, there can be four possible configurations.

The behavior of GIF is mainly based on *RFC2893 IPv6-over-IPv4 configured tunnel*.

▼ To Create GIF Tunnel

- 1 **In Interfaces / (assign), select GIF tab and click on the  button in order to create a new GIF tunnel.**
- 2 **Select parent interface.**
- 3 **Enter IP addresses for: GIF remote address, GIF tunnel local address and GIF tunnel remote address.**
Please consider figure 5-1 that refers to the GRE tunnels, however the same logic applies to the GIF tunnels.
- 4 **Check additional options, if required.**
- 5 **Click Save.**

A GIF tunnel can be configured with a number of additional options:

The *Route caching* option controls whether or not the route to the remote endpoint is cached. If your path to the remote peer is static, setting this can avoid one route lookup per packet. However if the path to the far side can change, this option could result in the GIF traffic failing to flow when the route changes.

The *ECN friendly behaviour* option controls whether or not Explicit Congestion Notification (ECN)-friendly practice of copying the TOS bit into/out of the tunnel traffic is done. By default the TOS bit on the packets is cleared or set to 0, depending on the direction of the traffic. With this option set, the bit is copied as needed between the inner and outer packets to be more friendly with intermediate routers that can perform traffic shaping.

The bridging tab allows to configure interface bridging. It is sometimes useful to divide one physical network (such as an Ethernet segment) into two separate network segments without having to create IP subnets and use a router to connect the segments together. A device that connects two networks together in this fashion is called a “bridge”. A Sentinel server with two network interface cards can act as a bridge. The bridge works by learning the MAC layer addresses (Ethernet addresses) of the devices on each of its network interfaces. It forwards traffic between two networks only when its source and destination are on different networks. There are many common situations in which a bridge is used today:

Connecting Networks

There are many reasons to use a host based bridge over plain networking equipment such as cabling constraints, firewalling or connecting pseudo networks such as a Virtual Machine interface. A bridge can also connect a wireless interface running in access point mode to a wired network and act as an access point.

Filtering/Traffic Shaping Firewall

A common situation is where firewall functionality is needed without routing or network address translation (NAT). An example is a small company that is connected via satellite link, DSL or ISDN to their ISP. They have a 13 globally-accessible IP addresses from their ISP and have 10 PCs on their network. A bridge-based firewall can be configured and dropped into the path just downstream of their DSL/ISDN router without any IP numbering issues.

Network Tap

A bridge can join two network segments and be used to inspect all Ethernet frames that pass between them.

Layer 2 Redundancy

A network can be connected together with multiple links and use the Spanning Tree Protocol to block redundant paths. For an Ethernet network to function properly only one active path can exist between two devices, Spanning Tree will detect loops and put the redundant links into a blocked state. Should one of the active links fail then the protocol will calculate a different tree and reenable one of the blocked paths to restore connectivity to all points in the network.

▼ To Create Bridge

- 1 In Interfaces / (assign), select Bridges tab and click on the  button in order to create a new bridge.
- 2 Select interfaces participating in the bridge.
- 3 Click Save.

There is a variety of advanced options that can be configured for a bridge:

| Option | Description |
|----------------|---|
| Protocol | <p>Protocol used for spanning tree. Supported options are RSTP and STP. The Spanning Tree Protocol (STP) is a network protocol that ensures a loop-free topology for any bridged Ethernet local area network. The basic function of STP is to prevent bridge loops and the broadcast radiation that results from them. Spanning tree also allows a network design to include spare (redundant) links to provide automatic backup paths if an active link fails, without the danger of bridge loops, or the need for manual enabling/disabling of these backup links. Spanning Tree Protocol (STP) is standardized as IEEE 802.1D. As the name suggests, it creates a spanning tree within a network of connected Layer 2 bridges (typically Ethernet switches), and disables those links that are not part of the spanning tree, leaving a single active path between any two network nodes. STP is based on an algorithm that was invented by Radia Perlman while she was working for Digital Equipment Corporation. In 2001, the IEEE introduced Rapid Spanning Tree Protocol (RSTP) as 802.1w. RSTP provides significantly faster spanning tree convergence after a topology change, introducing new convergence behaviors and bridge port roles to do this. RSTP was designed to be backwards-compatible with standard STP.</p> |
| STP Interfaces | Enable Spanning Tree Protocol on interface |
| Valid Time | Set the time that a Spanning Tree Protocol configuration is valid. The default is 20 seconds. The minimum is 6 seconds and the maximum is 40 seconds. |
| Forward Time | Set the time that must pass before an interface begins forwarding packets when Spanning Tree is enabled. The default is 15 seconds. The minimum is 4 seconds and the maximum is 30 seconds. |
| Hello Time | Set the time between broadcasting of Spanning Tree Protocol configuration messages. The hello time may only be changed when operating in legacy STP mode. The default is 2 seconds. The minimum is 1 second and the maximum is 2 seconds. |
| Priority | Set the bridge priority for Spanning Tree. The default is 32768. The minimum is 0 and the maximum is 61440. Additionally, you can set the Spanning Tree priority of interface to value. The default is 128. The minimum is 0 and the maximum is 240. These values are specified in separate entry boxes. |
| Hold Count | Set the transmit hold count for Spanning Tree. This is the number of packets transmitted before being rate limited. The default is 6. The minimum is 1 and the maximum is 10. |
| Path Cost | Set the Spanning Tree path cost of interface to value. The default is calculated from the link speed. To change a previously selected path cost back to automatic, set the cost to 0. The minimum is 1 and the maximum is 20000000. |

(Table continued)

| Option | Description |
|-------------------------|--|
| Cache Size | Set the size of the bridge address cache to size. The default is 100 entries. |
| Cache Entry Expiry Time | Set the timeout of address cache entries to this number of seconds. If seconds is zero, then address cache entries will not be expired. The default is 240 seconds. |
| SPAN Port | <p>Add the interface named by interface as a span port on the bridge. SPAN ports transmit a copy of every frame received by the bridge. This is most useful for snooping a bridged network passively on another host connected to one of the span ports of the bridge.</p> <hr/> <p>Note: the SPAN interface cannot be part of the bridge member interfaces.</p> |
| Edge Ports | Set interface as an edge port. An edge port connects directly to end stations and cannot create bridging loops in the network; this allows it to transition straight to forwarding. |
| Auto Edge Ports | <p>Allow interface to automatically detect edge status. This is the default for all interfaces added to a bridge.</p> <hr/> <p>Note: This will disable the autoedge status of interfaces.</p> |
| PTP Ports | Set the interface as a point-to-point link. This is required for straight transitions to forwarding and should be enabled on a direct link to another RSTP-capable switch. |
| Auto PTP Ports | <p>Automatically detect the point-to-point status on interface by checking the full duplex link status. This is the default for interfaces added to the bridge.</p> <hr/> <p>Note: The interfaces selected here will be removed from default autoedge status.</p> |
| Sticky Ports | Mark an interface as a "sticky" interface. Dynamically learned address entries are treated as static once entered into the cache. Sticky entries are never aged out of the cache or replaced, even if the address is seen on a different interface. |
| PTP Ports | Mark an interface as a "private" interface. A private interface does not forward any traffic to any other port that is also a private interface. |

The LAGG tab allows you to bond ("team") multiple network interfaces together into a virtual LAGG interface for fault tolerance or increased bandwidth.

LAGG interfaces can be enabled with different protocols:

Failover

Sends and receives traffic only through the master port. If the master port becomes unavailable, the next active port is used. The first interface added is the master port; any interfaces added after that are used as failover devices.

FEC (Cisco® Fast EtherChannel®)

Balances outgoing traffic across the active ports based on hashed protocol header information and accepts incoming traffic from any active port. This is a static setup and does not negotiate aggregation with the peer or exchange frames to monitor the link. The hash includes the Ethernet source and destination address, and, if available, the VLAN tag, and the IP source and destination address.

This protocol is based on the Cisco EtherChannel. Fast EtherChannel is a technology-leveraging, standards-based Fast Ethernet used in parallel to provide the additional bandwidth network backbones require today. It provides flexible, scalable bandwidth with resiliency and load sharing across links for switches, router interfaces and servers. Supports up to eight links per channel.

LACP

Supports the IEEE 802.3ad Link Aggregation Control Protocol (LACP) and the Marker Protocol. LACP will negotiate a set of aggregable links with the peer in to one or more Link Aggregated Groups (LAGs). Each LAG is composed of ports of the same speed, set to full-duplex operation. The traffic will be balanced across the ports in the LAG with the greatest total speed, in most cases there will only be one LAG which contains all ports. In the event of changes in physical connectivity, Link Aggregation will quickly converge to a new configuration.

Load Balance

This is an alias for FEC.

Round Robin

Distributes outgoing traffic using a round-robin scheduler through all active ports and accepts incoming traffic from any active port.

None

This protocol is intended to do nothing: it disables any traffic without disabling the LAGG interface itself.

Note: Only unassigned interfaces can be added to LAGG.

LAN, WAN, OPT1 and Other Interface Menus

These menu items are used to configure LAN, WAN, OPT1 and other network interfaces in the Sentinel system. The following configuration options are generally available per each interface:

| Option | Description |
|-------------|--|
| Enable | Whether interface is enabled in the system. |
| Description | Interface description. |
| Type | Specify interface type, from one of the following: <p><i>Static</i> This interface is manually configured with the IP address and CIDR mask. If this will be a WAN interface, you can either select a gateway from the list or create a new gateway. If you are creating a gateway, you can check the box to select it as a default gateway, enter a name, gateway IP address, and a description.</p> <p><i>DHCP</i> This interface will obtain its IP address via DHCP.</p> <p><i>PPP, L2TP, PPTP and PPPoE</i> These interfaces take a username and password, and optionally a service name, dial-on-demand setting, idle timeout, and optional periodic reset (PPPoE only).</p> |
| Description | Leave blank to use default hardware MAC address. Otherwise, this field can be used to modify ("spoof") the MAC address of this interface. Enter a MAC address in the following format: xx:xx:xx:xx:xx:xx. |
| MTU | In computer networking, the maximum transmission unit (MTU) of a communications protocol of a layer is the size (in bytes) of the largest protocol data unit that the layer can pass onwards. A larger MTU brings greater efficiency because each packet carries more user data while protocol overheads, such as headers or underlying per-packet delays, remain fixed; the resulting higher efficiency means a slight improvement in bulk protocol throughput. A larger MTU also means processing of fewer packets for the same amount of data. In some systems, per-packet-processing can be a critical performance limitation. However, this gain is not without some downside. Large packets can occupy a slow link for some time, causing greater delays to following packets and increasing lag and minimum latency. If you leave this field blank, the adapter's default MTU will be used. This is typically 1500 bytes but can vary on some hardware. |

(Table continued)

| Option | Description |
|--------------------------------------|--|
| MSS | The maximum segment size (MSS) is a parameter of the TCP protocol that specifies the largest amount of data, specified in bytes, that a computer or communications device can receive in a single TCP segment, and therefore in a single IP datagram. If you enter a value in this field, then MSS clamping for TCP connections to the value entered above minus 40 (TCP/IP header size) will be in effect. Otherwise, leave blank for default MSS. |
| Speed and Duplex | Click on «Advanced» button to edit. Here you can explicitly set speed and duplex mode for this interface. <u>Note: You must leave this set to autoselect (automatically negotiate speed) unless the port this interface connects to has its speed and duplex forced.</u> |
| Static IP Configuration / IP Address | IP address of the interface. Netmask is specified in CIDR notation. E.g. selecting 24 will result in 255.255.255.0 netmask selected. See <i>Appendix B – Netmask/CIDR Translation Table</i> . |
| Static IP Configuration / Gateway | If this interface is an Internet connection, select an existing Gateway from the list or add one using the «Add a new one» link. |
| DHCP / Hostname | The value in this field is sent as your DHCP client identifier and hostname when requesting a DHCP lease from the server. Some ISPs may require this (for client identification). |
| DHCP / Alias IP Address | The value in this field is used as a fixed alias IP address by the DHCP client. |
| Block Private Networks | When set, this option blocks traffic from IP addresses that are reserved for private networks as per <i>RFC 1918</i> (10/8, 172.16/12, 192.168/16) as well as loopback addresses (127/8). You should generally leave this option turned on, unless your WAN network lies in such a private address space, too. |
| Block Bogon Networks | When set, this option blocks traffic from IP addresses that are reserved (but not <i>RFC 1918</i>) or not yet assigned by IANA. Bogons are prefixes that should never appear in the Internet routing table, and obviously should not appear as the source address in any packets you receive. <u>Note: IP addresses that are currently in the bogon space may not be bogons at a later date because IANA and other registries frequently assign new address space to ISPs. Announcements of new assignments are often published on network operators' mailing lists (such as NANOG) to ensure that operators have a chance to remove bogon filtering for addresses that have become legitimate. For example, addresses in 49.0.0/8 were not allocated prior to August 2010, but are now used by APNIC. As time goes on, the IPv4 address exhaustion will mean there are fewer and fewer IPv4 bogons.</u> |

Web GUI: Firewall

Firewall menu items are used to configure Sentinel PF packet filter: firewall, NAT, traffic shaping rules and schedules and Virtual IP addresses.

Aliases

Aliases act as placeholders for real hosts, networks or ports. They can be used to minimize the number of changes that have to be made if a host, network or port changes. You can enter the name of an alias instead of the host, network or port in all fields that have a red background in this menu.

▼ To Create Alias

- 1 In Firewall / Aliases, click on the  button in order to create a new alias.
- 2 Enter alias name. The name of the alias may only consist of the characters "a-z, A-Z, 0-9 and _", e.g. localnetwork.
- 3 Enter optional alias description, e.g. Local Network.
- 4 In the final section, add objects (hosts, networks, ports, URLs or URL tables) to an alias, by pressing the  button.
- 5 Click Save.

A common usage of aliases would be, for example, to define a «local network» alias with a single subnet. If later your local network expands to multiple subnets, updating all the complex filtering rules would be a matter of simply adding a new subnet to the alias, instead of reconfiguring every filtering rule to span any additional subnets. The same concept can also be applied to ports that you may want to be filtered or processed in a special way in your network. Instead of multiplying different rules per each port, you can have a single filtering rule that uses an alias consisting of multiple ports. For your convenience, Sentinel PF comes with a set of pre-configured port aliases for different applications, beginning with *ports_* prefix.



Caution – If an alias cannot be resolved (e.g. because you deleted it), the corresponding filtering element - e.g. filter/NAT/shaper rule - will be considered invalid and skipped.

NAT

Note: By default the Sentinel PF operating system has outbound NAT enabled and NAT rules are automatically created and applied for all LAN-type interfaces. If all you need is to simply NAT outbound connections from LAN-type interfaces to WAN-type interface(s) then nothing extra should be configured.

In computer networking, network address translation (NAT) is the process of modifying IP address information in IP packet headers while in transit across a traffic routing device. In the mid-1990s NAT became a popular tool for alleviating the consequences of IPv4 address exhaustion. It has become a common, indispensable feature in routers for Internet connections. Most systems using NAT do so in order to enable multiple hosts on a private network to access the Internet using a single public IP address. This is also known as *masquerading*.

Sentinel PF supports features NAT configuration components represented as tabs in the Firewall / NAT menu:

Port Forward

The basic NAT technique enables communication only when the conversation originates in the masqueraded network, since this establishes the translation tables. For example, a web browser in the masqueraded (private) network can browse a website outside on the Internet, but a web browser outside could not browse a web site in the masqueraded network. However, Sentinel PF allows the network administrator to configure translation table entries for permanent use. This feature is often referred to as *static NAT* or *port forwarding*, and allows traffic originating in the "outside" network to reach designated hosts in the masqueraded network.

1:1 NAT

This is the simplest type of NAT that provides a one-to-one translation of IP addresses or subnets. *RFC 2663* refers to this type of NAT as basic NAT, which is often also called a 1:1 NAT. In this type of NAT only the IP addresses, IP header checksum and any higher level checksums that include the IP address are changed. The rest of the packet is left untouched, at least for basic TCP/UDP functionality; some higher level protocols may need further translation. Basic NATs can be used to interconnect two IP networks that have incompatible addressing. The 1:1 NAT binds a specific internal address or subnet to a specific external address or subnet. Incoming traffic from the Internet to the specified IP will be directed toward the associated internal IP. Outgoing traffic to the Internet from the specified internal IP will originate from the associated external IP.

Outbound NAT

This controls how traffic leaving Sentinel will be translated. The factory default scenario is to have all traffic that enters from a LAN or LAN-type interfaces to have NAT applied so it is translated to the WAN IP address before it leaves your network. This also includes most VPN-connected networks. You can disable the automatic outbound NAT and use so called *AON (Advanced Outbound NAT)* with all NAT rules specified manually.

When a client on the internal network contacts a machine on the Internet, it sends out IP packets destined for that machine. These packets contain all the addressing information necessary to get them to their destination. NAT is concerned with these pieces of information:

Source IP address (for example, 192.168.1.35)

Source TCP or UDP port (for example, 2132)

When the packets pass through the NAT gateway they will be modified so that they appear to be coming from the NAT gateway itself. The NAT gateway will record the changes it makes in its state table so that it can a) reverse the changes on return packets and b) ensure that return packets are passed through the firewall and are not blocked. For example, the following changes might be made:

Source IP: replaced with the external address of the gateway
(for example, 24.5.0.5)

Source port: replaced with a randomly chosen, unused port on the gateway
(for example, 53136)

Neither the internal machine nor the Internet host is aware of these translation steps. To the internal machine, the NAT system is simply an Internet gateway. To the Internet host, the packets appear to come directly from the NAT system; it is completely unaware that the internal workstation even exists.

When the Internet host replies to the internal machine's packets, they will be addressed to the NAT gateway's external IP (24.5.0.5) at the translation port (53136). The NAT gateway will then search the state table to determine if the reply packets match an already established connection. A unique match will be found based on the IP/port combination which tells PF (Sentinel's Packet Filter) the packets belong to a connection initiated by the internal machine 192.168.1.35. PF will then make the opposite changes it made to the outgoing packets and forward the reply packets on to the internal machine. Translation of ICMP packets happens in a similar fashion but without the source port modification.

We will begin with the Outbound NAT configuration tab in the NAT menu, as this is the common NAT scenario used on most Sentinel systems. The outbound NAT tab provides the following modes of operation:

Automatic Outbound Rule Generation

With automatic outbound NAT enabled, a mapping is automatically created for each interface's subnet (except WAN-type connections) and the rules on the outbound NAT page are ignored.

Manual Outbound NAT rule generation (AON - Advanced Outbound NAT)

If manual outbound NAT is enabled, outbound NAT rules will not be automatically generated and only the mappings you specify on this tab will be used. If there are no mappings specified then the outbound NAT will be disabled.

Note: To disable NAT in Sentinel, switch to the Manual Outbound NAT rule generation and delete all mappings.

The Advanced Outbound NAT Entry can be configured with a number of options:

| Option | Description |
|-------------|--|
| Do not NAT | Enabling this option will disable NAT for traffic matching this rule and stop processing Outbound NAT rules. |
| Interface | Choose interface the rule applies to. In most cases this will be WAN interface. |
| Protocol | Choose which protocol the rule will match. Specifying <i>any</i> will match all protocols. |
| Source | Specify source network for outbound NAT mapping. |
| Destination | Specify destination network for outbound NAT mapping. |
| Translation | <p>Packets matching this rule will be mapped to the IP address given here. If you want this rule to apply to another IP address rather than the IP address of the interface chosen above, select it here (you will need to define Virtual IP addresses on the interface first). You will also need to specify the source port for the outbound NAT mapping.</p> <p>Alternatively, you can map to a different subnet by choosing «Other subnet» in the Address dropdown menu. You can select how IP addresses are selected from the subnet pool is for mapping purposes:</p> <p><i>Round Robin:</i> Loops through the translation addresses. <i>Random:</i> Selects an address from the translation address pool at random. <i>Source Hash:</i> Uses a hash of the source address to determine the translation address, ensuring that the redirection address is always the same for a given source. <i>Bitmask:</i> Applies the subnet mask and keeps the last portion identical; 10.0.1.50 -> x.x.x.50. <i>Sticky Address:</i> The Sticky Address option can be used with the Random and Round Robin pool types to ensure that a particular source address is always mapped to the same translation address.</p> |

(Table continued)

| | |
|--|--|
| Note: Only <i>Round Robin</i> works with Host Aliases. Any type can be used with a Subnet. | |
| No XMLRPC Sync | Enabling this option prevents this NAT rule to be synchronized with other Sentinel servers connected via CARP for redundancy purposes. |
| Description | Optional NAT rule description. |

The 1:1 NAT Entry can be configured with the following options:

| Option | Description |
|--------------------|---|
| Disabled | Set this option to disable this rule without removing it from the list. |
| Interface | Choose interface the rule applies to. In most cases this will be WAN interface. |
| External Subnet IP | Enter the external (usually on a WAN) subnet's starting address for the 1:1 mapping. The subnet mask from the internal address below will be applied to this IP address. This is generally an address owned by the router itself on the selected interface. |
| Internal IP | Enter the internal (LAN) subnet for the 1:1 mapping. The subnet size specified for the internal subnet will be applied to the external subnet. |
| Destination | The 1:1 mapping will only be used for connections to or from the specified destination. This is usually set to <i>any</i> . |
| Description | Optional NAT rule description. |
| NAT Reflection | See <i>NAT Reflection</i> description below. |

NAT Reflection Option:

With NAT rules often used to forward incoming connections from the WAN to a local server with a private IP address, the mechanism does not facilitate a connection to the local server's public IP address from within the LAN. The reason is that NAT rules apply only to packets that pass through the specified interface, and this is often the external, WAN interface. Connecting to the external address of the firewall from a host on the LAN, however, does not mean the packets will actually pass through its external interface.

The TCP/IP stack on the firewall compares the destination address of incoming packets with its own addresses and aliases and detects connections to itself as soon as they have passed the internal interface. Such packets do not physically pass through the external interface, and the stack does not simulate such a passage in any way. Thus, Sentinel's PF never sees these packets on the external interface, and the NAT redirection rule, specifying the external interface, does not apply. This situation may occur, for example, if you wanted to access the local web server internally via its public IP address.

With additional NAT rules, the lacking passage can be achieved. Enabling NAT reflection will cause packets from the client to be translated, replacing the client's source address with the firewall's internal address. The internal server will reply back to the firewall, which can reverse all translations when forwarding traffic to the local client. This construct is rather complex as it creates multiple separate states for each reflected connection. Care must be taken to prevent the NAT rule from applying to other traffic, for instance connections originating from external hosts (through other redirections) or the firewall itself.

Rules

The Rules menu allows to configure firewall rules in the Sentinel PF operating system. Firewall rules control what traffic is allowed to enter an interface on your firewall. Once traffic is passed on the interface it enters, an entry in the *state* table is created, which allows through subsequent packets that are part of that connection.

Firewall rules are processed from the top down, and the first match wins. The default on all interfaces is to deny traffic, and only what is explicitly allowed via firewall rules will be passed.

On the Firewall Rules menu, there will be a tab for each interface, plus a tab for each active VPN type (IPsec, OpenVPN, PPTP), and a tab for "Floating Rules" which contains rules that apply to multiple interfaces, to avoid repetition. When editing a rule, many of the options are explained in detail on the rule editor screen. Be mindful of the default settings on the rule editor, especially the protocol.

Note: Rules are evaluated on a first-match basis, i.e. the action of the first rule to match a packet will be executed. This means that if you use block rules, you'll have to pay attention to the rule order. Everything that isn't explicitly passed is blocked by default.

Sentinel PF operates with three types of packet filter rules:

LAN

These are rules applicable for packets arriving to LAN interface.

WAN

These are rules applicable for packets arriving to WAN interface.

Floating

Floating rules are parsed before rules on other interfaces. Thus, if a packet matches a floating rule and the *Quick* option is active on that rule, Sentinel PF will not attempt to filter that packet against any rule on any other interface. The most common use of floating rules are traffic shaper rules to match and queue traffic. Other common uses are to ensure that no traffic can exit into a secure network by blocking the outbound toward a secure network. Additionally, you can use floating rules to enact state timeouts and tag & match operations. Floating rules are advanced firewall rules which can apply in any direction to any (or multiple) interfaces. Some more advanced/low-level options are available on floating rules than exist for the normal per-interface rules. By using floating rules you can control and restrict traffic from the firewall itself, you can have rules which apply to multiple interfaces in the same way, you can have traffic shaping rules which match traffic but do not affect it's pass/block action, and much more. Many firewalls do not need any floating rules, or may only have them for the traffic shaper.

You can switch between Floating, LAN and WAN rules by clicking on corresponding tabs in the Rules menu. You may have already noticed that the default Sentinel PF installation includes WAN rules setup to allow Web GUI and SSH access on WAN, and an Anti-Lockout Rule to allow the same on LAN. The Anti-Lockout Rule allows access to the SSH and Web GUI on the LAN interface, regardless of the user-defined firewall rule set.

Each firewall rule can be configured with one of the following actions:

TABLE 7-1 Sentinel PF Firewall Rule Actions

| Action | Meaning |
|--|--|
|  Pass | Allow packets to pass. |
|  Block | Silently drop packets. |
|  Reject | Reject packets. |
|  Queue | Assign packets to queue and pass. (Floating rules only.) |

The difference between *block* and *reject* is that with *reject*, a packet (TCP RST or ICMP port unreachable for UDP) is returned to the sender, whereas with *block* the packet is dropped silently. In either case, the original packet is discarded.

▼ To Create Firewall Rule

- 1 In Firewall / Rules, select the Floating, WAN or LAN tab depending on which interface(s) this rule has to apply to, and click on the  button in order to create a new rule.
- 2 Configure your rule.
- 3 Click Save.
- 4 Click Apply Changes to reload Sentinel PF packet filter immediately for the rule to take effect.

Note: Newly created firewall rules are not active unless the filter is reloaded. This can be done either by clicking «Apply Changes» when offered, or in the Status / Filter Reload menu manually.

Because rules are evaluated on a first-match basis, you may want to review your firewall rules and move the newly created rule above or below existing rules, depending on your preference.

▼ To Move Rule

- 1 In Firewall / Rules, select the Floating, WAN or LAN tab depending on which interface(s) your rule applies to, and click on the  button in order to create a new rule.
- 2 Select the rule you want to move by clicking on a checkbox on the left side of the rule table.
- 3 Click on the  button next to any existing rule you would like your selected firewall rule to appear on top of.
You can't move anything on top of anti-lockout rules on LAN interface.
- 4 Click Apply Changes to reload Sentinel PF packet filter immediately for the new rule sequence to take effect.

Firewall rules can be configured with a number of options:

| Option | Description |
|------------------------|---|
| Action | Choose what to do with packets that match the criteria specified below. This is one of the firewall actions specified in <i>Table 7-1</i> . |
| Disabled | Set this option to disable this rule without removing it from the list. <u>Note: Once a rule is disabled, its action icons will appear lighter than normal in firewall rules tables.</u> |
| Interface | The Interface drop down specifies the interface on which the rule will be applied. Remember that on interface and group tab rules, traffic is only filtered on the interface where the traffic is initiated. Traffic initiated from your LAN destined to the Internet or any other interface on your firewall is filtered by the LAN ruleset. |
| Protocol | This is where you specify the protocol this rule will match. Most of these options are self-explanatory. <i>TCP/UDP</i> will match both TCP and UDP traffic. Specifying <i>ICMP</i> will make another drop down box appear where you can select the ICMP type. Several other common protocols are also available. <u>Note: This field defaults to <i>TCP</i> for a new rule because it's a common default and it will display the expected fields for that protocol. If you wish the rule to apply to any protocol, you will need to change this field to <i>any</i>. One of the most common mistakes in creating new rules is accidentally creating a TCP rule and then not being able to use ping, DNS, etc.</u> |
| Source | Specify source packets must come from to match this rule. You can specify <i>any</i> type for the rule to match any sources. Otherwise, you can specify a single host or alias, or a subnet, or PPP type clients. <u>Note: There are separate types for WAN and LAN addresses and subnets. It is easier to use these types rather than entering your WAN and LAN IP addresses manually. For example, if your LAN IP addressing changes in future, the rule with a <i>LAN Subnet</i> type will match new LAN addresses automatically.</u> |
| Source Port Range | Click <i>Advanced</i> to specify ports at the <i>source</i> packets must come from to match this rule. |
| Destination | Specify destination packets must come to, in order match this rule. You can specify <i>any</i> type for the rule to match any destination. Otherwise, you can specify a single host or alias, or a subnet, or PPP type clients. |
| Destination Port Range | Specify ports at the <i>destination</i> packets must come to, in order match this rule. |

(Table continued)

| | |
|----------------------------|---|
| Log | <p>Log packets that are handled by this rule.</p> <hr/> <p>Note: Sentinel PF, like any other server, has limited local log space. Don't turn on logging for everything. If you want to do a lot of logging, consider using a remote <i>syslog</i> server (see the <i>Diagnostics / System logs / Settings</i> menu).</p> |
| Description | Optional description of this firewall rule. |
| Source OS | <p>Sentinel can determine some operating system types of the packet sender using its TCP «fingerprints». This may be useful to build rules to block vulnerabilities known to exist with specific operating systems only.</p> <hr/> <p>Note: TCP fingerprinting is not 100% accurate, so use this technique with care.</p> |
| DiffServ Code Point (DSCP) | <p>Differentiated services or DiffServ is a computer networking architecture that specifies a simple, scalable and coarse-grained mechanism for classifying and managing network traffic and providing quality of service (QoS) on modern IP networks. The DS field contains a 6-bit Differentiated Services Code Point (DSCP) value you can use to match packets. See <i>Appendix E - DiffServ Classification & Marking</i> for additional information.</p> |
| Advanced Options | <p>Click on the button to open <i>Advanced Options</i> section that allows to specify the following options:</p> <p>The first checkbox allows packets with IP options to pass. Otherwise they are blocked by default. This is usually only seen with multicast traffic. Please consider</p> <p>http://www.iana.org/assignments/ip-parameters/ip-parameters.txt</p> <p>for additional information on IP options.</p> <p>The second checkbox allows to disable auto generated <i>reply-to</i> for this rule. This prevents replies to be automatically routed to the same network interface packets have been received on.</p> <p>Additionally, you can mark a packet matching this rule and use this mark to match on other NAT/filter rules; and you can match packet on a mark placed before by another rule.</p> <p>Finally, you can specify maximum state entries this rule can create, maximum number of unique source hosts, established connections per host, state entries per host and new connections per second, as well as the state timeout in seconds.</p> <p>Advanced options above will remain out of the scope of this handbook and will be discussed in supplemental Sentinel PF documentation available from BusinessCom.</p> |
| TCP Flags | Choose TCP flags (also known as control bits) that must be set or cleared for this rule to match. |

(Table continued)

| | |
|----------------|---|
| State Type | Select which type of state tracking mechanism you would like to use. See state tracking description below. If in doubt, use keep state. <hr/> Note: This only affects traffic in the inbound direction. Use another floating rule for the outbound direction. <hr/> |
| No XMLRPC Sync | This prevents the rule from automatically syncing to other CARP members. |
| Schedule | Select schedule when this rule applies. |
| Gateway | Leave as <i>default</i> to use the system routing table. Or choose a gateway to utilize policy based routing. |
| In/Out | These selections let you pick from your defined <i>limiters</i> to apply a bandwidth limit to the traffic entering this interface (<i>In</i>) and leaving this interface (<i>Out</i>). Please remember that <i>in</i> and <i>out</i> are from the perspective of that interface on the firewall. If you're choosing limiters on the LAN interface, <i>out</i> is download speed (traffic from the LAN interface into the LAN) and <i>in</i> is upload speed (traffic from the LAN into the LAN NIC). The <i>Out</i> selection is applied to traffic leaving the interface where the rule is created, <i>In</i> is applied to traffic coming into the chosen interface. <hr/> Note: Choose the <i>Out</i> queue/Virtual interface only if you have also selected In. <hr/> |
| AckQueue/Queue | These options define which traffic shaper queues are applied to traffic entering and exiting this interface. |
| Layer 7 | Selecting an entry for Layer 7 will redirect traffic into a Layer 7 inspection instance. <hr/> Note: Rules for Layer 7 should always use the <i>pass</i> action. The decision to block or queue is made by the Layer 7 inspection instance, the firewall rule merely passes the traffic into the inspection daemon so it can be acted upon later. <hr/> |

State tracking:

One of Sentinel PF's important abilities is "keeping state" or "stateful inspection". Stateful inspection refers to Sentinel PF's ability to track the state, or progress, of a network connection. By storing information about each connection in a state table, Sentinel PF is able to quickly determine if a packet passing through the firewall belongs to an already established connection. If it does, it is passed through the firewall without going through ruleset evaluation.

Keeping state has many advantages including simpler rulesets and better packet filtering performance. Sentinel PF is able to match packets moving in either direction to state table entries meaning that filter rules which pass returning traffic don't need to be written. And, since packets matching stateful connections don't go through ruleset evaluation, the time PF spends processing those packets can be greatly lessened.

When a rule creates state, the first packet matching the rule creates a "state" between the sender and receiver. Now, not only do packets going from the sender to receiver match the state entry and bypass ruleset evaluation, but so do the reply packets from receiver to sender. This is the default *keep state* mechanism.

The *sloppy state* mechanism uses a sloppy TCP connection tracker that does not check sequence numbers at all, which makes insertion and ICMP teardown attacks way easier. This is intended to be used in situations where one does not see all packets of a connection, e.g. in asymmetric routing situations.

The *synproxy state* option can be used to cause Sentinel PF itself to complete the handshake with the active endpoint, perform a handshake with the passive endpoint, and then forward packets between the endpoints. No packets are sent to the passive endpoint before the active endpoint has completed the handshake, hence so-called SYN floods with spoofed source addresses will not reach the passive endpoint, as the sender can't complete the handshake. The proxy is transparent to both endpoints; they each see a single connection from/to the other endpoint. Sentinel PF chooses random initial sequence numbers for both handshakes. Once the handshakes are completed, the sequence number modulators are used to translate further packets of the connection.

The *none* option tells Sentinel PF not to use state mechanisms to keep tracking. This is only useful if you're doing advanced queuing in certain situations.

Schedules

Schedules act as placeholders for time ranges to be used in firewall rules. A common use for schedules would be activate and deactivate firewall rules depending on the time of the day. This can be used, for example, to prioritize or block certain applications during business hours.

▼ To Create Schedule

- 1 In Firewall / Schedules, click on the  button in order to create a new schedule.
- 2 Enter schedule name. The name of the schedule may only consist of the characters "a-z, A-Z, 0-9 and _", e.g. `businesshours`.
- 3 Enter optional schedule description, e.g. `Business Hours`.
- 4 Select month of the year schedule will apply to.
- 5 Click on the days to select days.
You can select certain week days by clicking on the calendar column title. E.g. clicking on *Mon* will select all Mondays in the month. Selections are marked red.
- 6 Select Start time and Stop time for the schedule.
- 7 Enter optional time range description, e.g. `Business Hours`.
- 8 Click «Add Time».
- 9 Repeat steps 4 to 8 until you have added all time ranges in the schedule.
- 10 Click «Save».

Note: Schedules are automatically repeated every month. For example, if you have defined a schedule for a certain month then the same schedule will be automatically created for all following months. Thus, in order to create, for example, a schedule to highlight business hours, you can pick current month, select all days and enter the business hours time range (e.g. 07:00 to 18:00). Once you save the time range, it will automatically apply to all the other months in this schedule.

| September 2013 | | | | | | |
|----------------|-----|-----|-----|-----|-----|-----|
| Mon | Tue | Wed | Thu | Fri | Sat | Sun |
| | | | | 1 | 2 | 3 |
| 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| 11 | 12 | 13 | 14 | 15 | 16 | 17 |
| 18 | 19 | 20 | 21 | 22 | 23 | 24 |
| 25 | 26 | 27 | 28 | 29 | 30 | 31 |

FIGURE 7-1 Example of “Every Friday” Schedule.

Traffic Shaper

This is the primary configuration menu of the Sentinel PF traffic shaper. Traffic shaping is one of the Quality of Service (QoS) techniques used to control bandwidth allocation to specific hosts or subnets in your network as well as prioritize certain traffic types over another. Without QoS, Sentinel PF defaults to the FIFO (First In First Out) packets processing, with all traffic assigned to a default priority level.

Sentinel PF shapes traffic where it can control the flow. The traffic coming from WAN to LAN is shaped on the LAN interface, as it is actually coming *out* of the LAN interface. Similarly, traffic coming from the LAN to WAN is shaped on the WAN interface, as it is going *out* from the WAN interface:

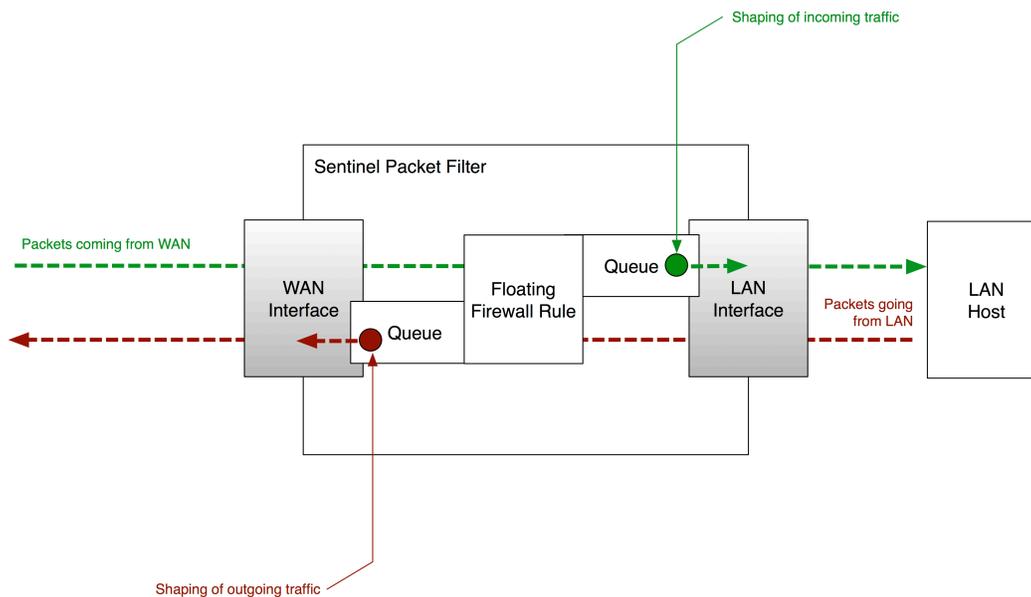


FIGURE 7-2 Operation of Sentinel PF Traffic Shaper.

Sentinel PF traffic shaper uses a combination of *queues* and *firewall rules* to work. The queues are where bandwidth and priorities are actually allocated. There can be multiple queues defined to enforce QoS for different traffic types. Firewall rules control how traffic is assigned into those queues. Once a packet matches the firewall rule, it is assigned into a queue by that rule.

Note: Firewall rules for traffic shaping are usually set up as *floating* firewall rules configured with the *queue* action that is used to assign traffic to queues.

Note: Sentinel uses a separate queue for LAN-to-LAN traffic, so traffic shaping will not affect local traffic.

Sentinel PF features multiple scheduler types you can use for queuing. These scheduler types are summarized in the table below for your convenience.

TABLE 7-2 Sentinel PF Scheduler Types

| Feature | Hierarchical Fair Service Curve (HFSC) | Class-Based Queueing (CBQ) | Priority Queueing (PRIQ) |
|---------------------------|--|----------------------------|--------------------------|
| Orientation | Bandwidth | Bandwidth & Priorities | Priorities |
| Queue Structure | Hierarchical | Hierarchical | Flat |
| Non-Linear Service Curves | Yes | No | No |
| CIR Guarantees | Yes | No | No |
| MIR Limits | Yes | Yes | No |
| Bandwidth Sharing | Yes | Yes | Yes |
| Priorities | No | Yes | Yes |
| Configuration Complexity | High | Medium | Low |

Hierarchical Fair Service Curve (HFSC)

This is an extension of service curve based QoS model proposed in “*A Hierarchical Fair Service Curve Algorithm for Link-Sharing, Real-Time and Priority Services*” (1997) by Ion Stoica, Hui Zhang and T. S. Eugene Ng. In HFSC, the queues are arranged in a hierarchy, or a tree, with root queues for each interface, parent queues underneath, and child queues nested under the parent queues. HFSC is very effective if you need to guarantee a certain amount of minimum bandwidth (CIR) for hosts or applications in your network, yet still allow idle bandwidth to be shared among queues. In addition to CIR guarantees, HFSC can impose maximum information rate (MIR) limits on queues.

HFSC, as proposed in the original paper by Stoica et al, achieves three performance goals for realtime, hierarchical linksharing, and priority services:

- Guarantee the service curves of all leaf classes – CIR guarantees.
- Minimize the short-term discrepancy between the amount of services provided to an interior class and the amount specified by the Fair Service Curve linksharing model.
- Allocate the excess bandwidth to sibling classes with bounded fairness, also known as *bandwidth sharing*.

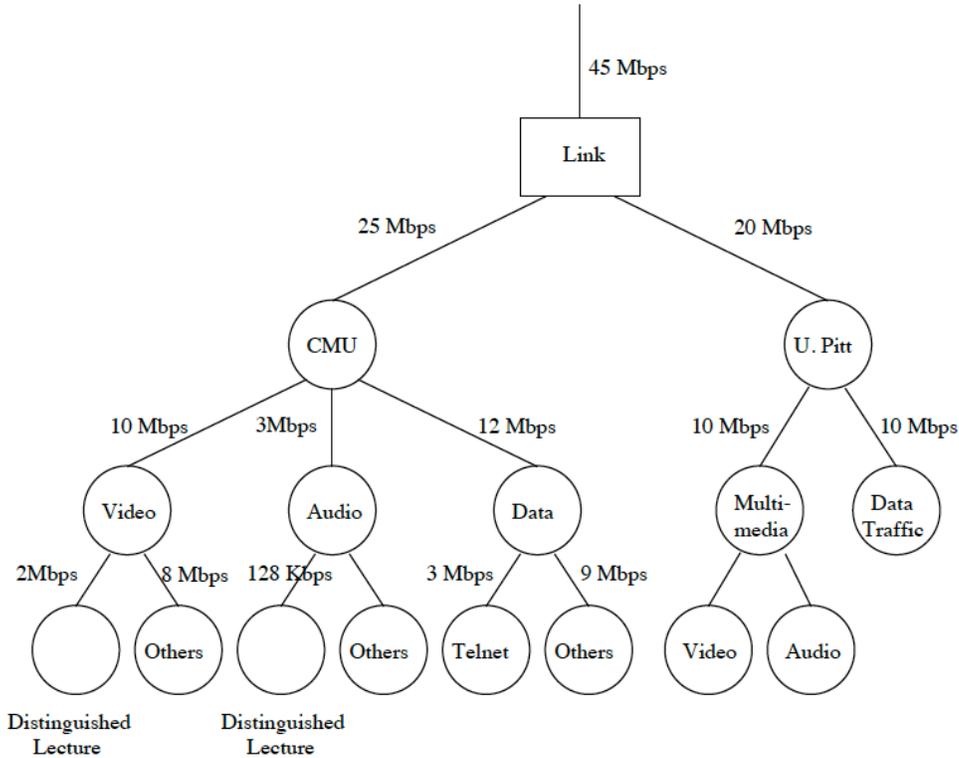


FIGURE 7-3 Example of link-sharing hierarchy used by HFSC and CBQ.

One of the most important aspects of the HFSC is the ability to define *Non-Linear Service Curves*, or simply *Service Curves*. With service curves, you can configure HFSC to allocate an increase amount of bandwidth to a connection for initial pre-defined time period, after which the connection will be throttled down to default throughput level. This may be useful to allow lightweight content, such as HTML code and images, to traverse in bursts at higher throughput, if compared to large multimedia files and downloads.

Service curves can improve performance of delay-sensitive bursty applications without negatively affecting the performance of throughput-oriented applications like FTP.

Consider the following service curve example:

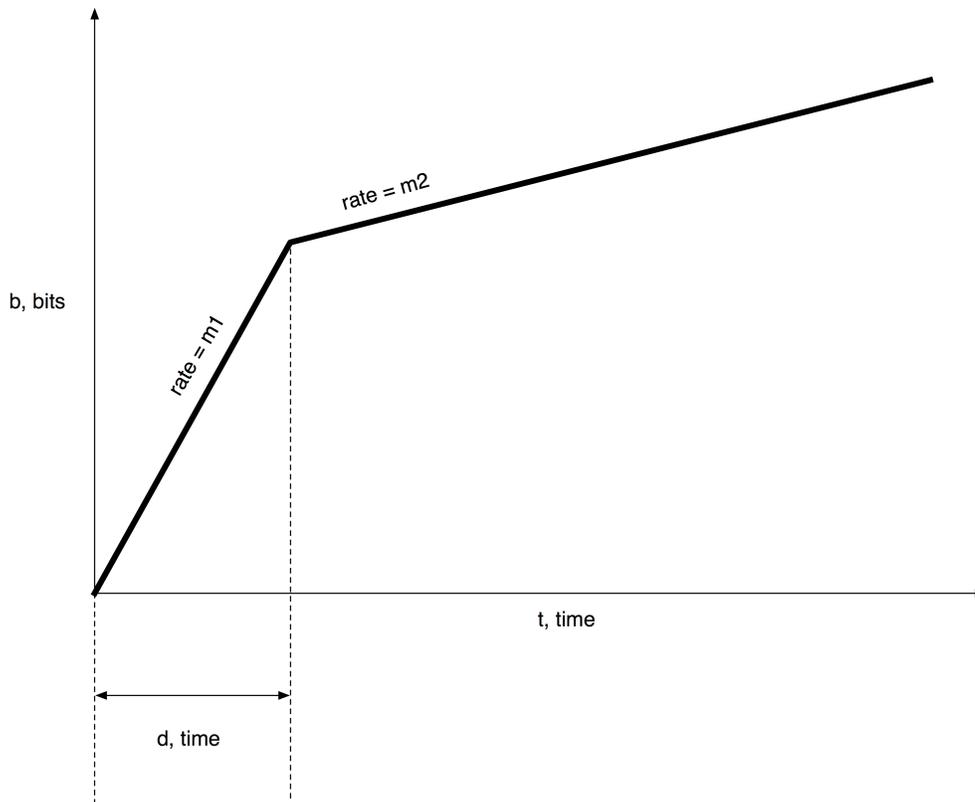


FIGURE 7-4 Example of non-linear convex service curve.

An HFSC service curve is based on $m1$ and $m2$, which are the initial and the normal throughput rates respectively. A connection is allocated $m1$ bandwidth for the initial time of d , after which it is throttled down or allowed to burst at the $m2$ level. With the convex curve above, the connection is throttled down after d ; that is, the rate of the bits flow decreases. It is possible to define concave curves whereby the initial amount of bandwidth allocated to a connection will be less than the normal $m2$ rate.

Within the initial d time, the $m2$ throughput rate is not checked. After d has expired, if the traffic is still above the $m2$ rate then it will be shaped down to $m2$. It is also possible to define a linear service curve with $m1$ and d left blank and only $m2$ specified.

Note: Current Sentinel PF HFSC implementation does not support priorities. Web GUI HFSC priority entries will be ignored.

These values can be used to define a service curve for the following purposes in HFSC:

- *Upper Limit*. This is the MIR (maximum information rate) bandwidth allowed for the queue. It will do hard limiting.
- *Real Time*. This is the CIR (committed information rate) bandwidth guaranteed for the queue. This is only valid for child queues. The *m1* parameter will always be satisfied in timeframe *d*, and *m2* is the maximum CIR this discipline will allow to be used.

Note: The value for *m2* cannot exceed 30% of the parent's queue's available bandwidth.

- *Link Share*. Any additional amount of bandwidth this queue gets from sharing bandwidth with other queues after all *Real Time* requirements are satisfied.

By combining these factors, a queue will get the bandwidth specified by the Real Time factors, plus those from Link Share, up to a maximum of Upper Limit.

Consider the following non-linear service curve example:

TABLE 7-3 HFSC Example

| Parameter | m1 | d | m2 |
|-------------|----------|-----------|----------|
| Upper Limit | <blank> | <blank> | 500 kbps |
| Real Time | 300 kbps | 10,000 ms | 0 kbps |
| Link Share | 50 kbps | 10,000 ms | 200 kbps |

This service curve will result in a connection being allocated 300 kbps of guaranteed CIR bandwidth for the initial period of 10,000 ms (10 seconds). It will also be able to borrow up to 50 kbps of unused (shared) bandwidth from other queues. After the 10 seconds period, the CIR rate drops to zero, however the connection will be able to borrow up to 200 kbps of shared bandwidth from other queues.

Note: With the example above, the *m2* parameter is specified for both the Upper Limit and the Link Share. In these situations, Sentinel PF enforces the *m2* from the Link Share setting, and it will ignore the *m2* parameter from the Upper Limit. Upper Limit *m2* could have as well been blank in this example.

Class-Based Queueing (CBQ)

CBQ is a queueing algorithm that divides a network connection's bandwidth among multiple queues or classes. Each queue then has traffic assigned to it based on source or destination address, port number, protocol, and other criteria. A queue may optionally be configured to borrow bandwidth from its parent queue if the parent is being under-utilized.

A CBQ queue can be assigned a priority such that those containing interactive traffic, such as SSH, can have their packets processed ahead of queues containing bulk traffic, such as FTP. CBQ queues, similar to HFSC, are arranged in an hierarchical manner. At the top of the hierarchy is the root queue which defines the total amount of bandwidth available. Child queues are created under the root queue, each of which can be assigned some portion of the root queue's bandwidth.

CBQ supports setting MIR bandwidth limits, however it does not support CIR guarantees like HFSC. Because of the simpler queue configuration, it can be a good alternative to HFSC, especially if you do not need to guarantee CIR in the local network.

With CBQ, queue priorities range from 0 to 7, higher numbers indicating higher priority. Queues of an equal priority are processed in a round-robin fashion.

Note: The MIR limit in CBQ may sound to operate like a *limiter*, however with CBQ, the sum of child queues cannot exceed the total bandwidth available in the parent queue. CBQ is not an alternative to using Sentinel PF *limiters* to shape hosts on a per-IP basis.

Priority Queueing (PRIQ)

Priority Queueing (PRIQ) assigns multiple queues to a network interface with each queue being given a priority level. A queue with a higher priority is always processed ahead of a queue with a lower priority. If two or more queues are assigned the same priority then those queues are processed in a round-robin fashion. The queueing structure in PRIQ is flat - you cannot define queues within queues. The root queue is defined, which sets the total amount of bandwidth that is available, and then sub queues are defined under the root. The PRIQ values for priority go from 0 to 15, and the higher the priority number, the more likely the queue is to have its packets processed.

It is important to note that when using PRIQ you must plan your queues very carefully. Because PRIQ always processes a higher priority queue before a lower priority one, it's possible for a high priority queue to cause packets in a lower priority queue to be delayed or dropped if the high priority queue is receiving a constant stream of packets.

In addition to schedulers, Sentinel PF provides *Random Early Detection (RED)* congestion avoidance algorithm. Its job is to avoid network congestion by making sure that the queue doesn't become full. It does this by continually calculating the average length (size) of the queue and comparing it to two thresholds, a minimum threshold and a maximum threshold. If the average queue size is below the minimum threshold then no packets will be dropped. If the average is above the maximum threshold then *all* newly arriving packets will be dropped. If the average is between the threshold values then packets are dropped based on a probability calculated from the average queue size. In other words, as the average queue size approaches the maximum threshold, more and more packets are dropped. When dropping packets, RED randomly chooses which connections to drop packets from. Connections using larger amounts of bandwidth have a higher probability of having their packets dropped. RED is useful because it avoids a situation known as global synchronization and it is able to accommodate bursts of traffic. Global synchronization refers to a loss of total throughput due to packets being dropped from several connections at the same time. For example, if congestion occurs at a router carrying traffic for 10 FTP connections and packets from all (or most) of these connections are dropped (as is the case with FIFO queueing), overall throughput will drop sharply. This isn't an ideal situation because it causes all of the FTP connections to reduce their throughput and also means that the network is no longer being used to its maximum potential. RED avoids this by randomly choosing which connections to drop packets from instead of choosing all of them. Connections using large amounts of bandwidth have a higher chance of their packets being dropped. In this way, high bandwidth connections will be throttled back, congestion will be avoided, and sharp losses of overall throughput will not occur. In addition, RED is able to handle bursts of traffic because it starts to drop packets before the queue becomes full. When a burst of traffic comes through there will be enough space in the queue to hold the new packets.

RED should only be used when the transport protocol is capable of responding to congestion indicators from the network. In most cases this means RED should be used to queue TCP traffic and *not* UDP or ICMP traffic.

Explicit Congestion Notification (ECN) works in conjunction with RED to notify two hosts communicating over the network of any congestion along the communication path. It does this by enabling RED to set a flag in the packet header instead of dropping the packet. Assuming the sending host has support for ECN, it can then read this flag and throttle back its network traffic accordingly.

For more information on RED and ECN, please refer to *RFC 3168* and Floyd, S., and Jacobson, V., *Random Early Detection gateways for Congestion Avoidance V.1 N.4*, August 1993, p. 397-413.

Setting up schedulers, queues and firewall rules can be a complex task. To simplify the process, Sentinel PF provides multiple wizards to create initial traffic shaping frameworks you can build a more advanced model on top of. These wizards are:

TABLE 7-4 Sentinel PF Traffic Shaping Wizards

| Wizard | Description | Schedulers | Per-host MIR and CIR | Bandwidth Sharing in LAN | Multiple LAN | Multiple WAN | LAN/WAN Pairings |
|------------------------------|---|---------------------|----------------------|--------------------------|--------------|--------------|------------------|
| Sentinel Shaper (under dev.) | Shaping framework for single LAN and single WAN connection that provides per-IP MIR and CIR limits with bandwidth sharing capability. | CBQ | Yes | Yes | No | No | No |
| Single LAN, Multi-WAN | Basic traffic shaping framework for a single LAN and multiple WAN connections. | HFSC CBQ PRIQ | No (queue only) | Yes | No | Yes | No |
| Single WAN, Multi-LAN | Basic traffic shaing framework for a single WAN and multiple LAN connections. | HFSC CBQ PRIQ | No (queue only) | Yes | Yes | No | No |
| Multiple LAN/WAN | Basic traffic shaing framework for multiple LAN and WAN connections. | HFSC CBQ PRIQ | No (queue only) | Yes | Yes | Yes | No |
| Dedicated Links | Basic traffic shaing framework for multiple LAN and WAN connections whereby each WAN connection is dedicated to each LAN connection. | HFSC CBQ PRIQ | No (queue only) | Yes | Yes | Yes | Yes |

All wizards are capable to create multiple queues for different types of traffic, and if used with CBQ and PRIQ schedulers, assign different priorities to these queues. The *Sentinel Shaper* wizard allows to specify MIR and CIR per each IP in the LAN, as well as enable bandwidth sharing among different IPs.

In addition to traffic shaping, Sentinel PF provides a concept of *limiters* that are used to impose hard MIR bandwidth limits on IPs, or group of IPs. You can also assign priorities inside traffic controlled by limiters. Limiters are handled differently to traffic shaping. The bandwidth limits are imposed using regular *pass* action firewall rules as it enters and leaves interface.

Note: For environments with multiple WAN connections, limits for non-default gateways will need to be applied using floating firewall rules, with the *out* direction and appropriate gateway set.

Similar to HFSC and CBQ, limiters may be nested with *queues* inside other *queues*. Root-level limiters (*pipes*), may have bandwidth limits and delays, while child limiters (*queues*) belonging to a pipe, may have different packet order priorities (*weights*). Bandwidth limits can be optionally masked by either the source IP or the destination IP, so that the limits can be applied per-IP instead of as a group. Limiters are almost always used in pairs, one for incoming traffic and one for outgoing traffic.

There are a few applications for limiters within bandwidth management tasks:

- Limit maximum amount of bandwidth (specify MIR) per protocol, or limit all the traffic *except* a specific protocol.
- Limit maximum amount of bandwidth (specify MIR) per IP, or group of IPs or subnets.

Similar to traffic shaping, limiters hold traffic to a certain point by dropping or delaying packets to achieve a specific throughput. Usually a protocol congestion algorithms will detect the packet loss and will throttle back to a limit it can sustain. If multiple connections are queued under the same limiter pipe then their weights will be considered when ordering packets before they are sent. Unlike priorities in traffic shaping, the weight of a queue in a limiter will never cause it to starve for bandwidth.

There are certain drawbacks with limiters:

- Limiters cannot provide any CIR guarantees.
- Limiters cannot share bandwidth from other pipes.
- Queues do not have MIR values and use pipe's MIR setting. You can not limit a queue inside a pipe, however you can specify weights to prioritize packets.

▼ To Create Limiter

- 1 In Firewall / Traffic Shaper, select the «Limiter» tab and click on the  button in order to create a new limiter.
- 2 Click on «Enable limiter».
- 3 Enter limiter name. The name of the schedule may only consist of the characters "a-z, A-Z, 0-9 and _", e.g. `penaltylimiter-in`.
- 4 Enter optional limiter description, e.g. `Penalty Limiter at 200 kbit/s: In direction`.
- 5 Enter bandwidth amount that is used as MIR setting for the limiter.
- 6 **Specify masking option.**
If *none* is specified then the MIR bandwidth limit will apply to all the traffic passing through the limiter, as the whole. With *source address* and *destination address* the MIR limit will apply on a per-IP basis. If you are creating a limiter on the LAN interface then, usually, you want to use the *source address* mask when creating an *in* direction limiter that will handle upload traffic coming from your LAN (as if traffic goes from your LAN *into* the LAN interface); and use *destination address* mask for *out* direction limiters used to limit downloads going *out* to your LAN. Swap the logic if you are building a limiter on the WAN side.
- 7 Click «Show advanced options» button to specify any advanced options, if necessary.
- 8 Click «Save».
At this point, the parent limiter (pipe) is created.
- 9 Click «Add New Queue» to add a child queue into the limiter.
- 10 Specify child queue Name, bandwidth, mask, advanced options (optional) and description, similar to how the parent limiter was created.
- 11 Click «Save».
- 12 Repeat steps 9 to 11 for any additional child queues you need in this limiter.
- 13 **After the limiter is created, create a firewall rule on the desired interface with the In/Out parameter set in order to apply the limiter to the matched traffic.**
Please refer to the *Rules* section of this Chapter on how to create rules.

Note: Limiters are usually used in pairs, one limiter per *in* and *out* directions.

Limiters can be configured with a number of advanced options:

| Option | Description |
|------------------|--|
| Delay | Used only on limiter pipes. Introduces artificial delay (latency), specified in milliseconds. Delay is typically left blank to transmit packets as fast as possible, however it can be used to simulate high-latency WAN environment, such as satellite links for experimental purposes. |
| Packet Loss Rate | Another method for link quality degradation, typically used for experimental purposes. If set to 0.01, the limiter will drop 1% of all packets passing through the pipe. Normally left empty. |
| Queue Size | In most cases, you should leave the field empty. All packets in this pipe are placed into a fixed-size queue first, then they are delayed by value specified in the Delay field (defaults to 50 slots), and then they are delivered to their destination. |
| Bucket Size | In most cases, you should leave the field empty. Sets the hash table sized used for queue storage (defaults to 64 slots). Must be numeric value between 16 and 65536. |

Limiter's child queues allow to specify the *weight* parameter which defines the priority of a queue. Higher values assign higher priority to packets that are in a given queue. Limiter priorities range from 1 (highest) to 100 (lowest).

Once a limiter is created, traffic should be assigned to it using a firewall rule, using the In/Out parameter. Please consider the *Rules* section of this chapter. You can specify a range of criteria to match desired traffic and assign it to a limiter. Note the In/Out direction is designated from the perspective of Sentinel PF itself, e.g. inbound traffic on the LAN interface is actually going toward the WAN interface. For your convenience, limiter directions are shown in the figure below.

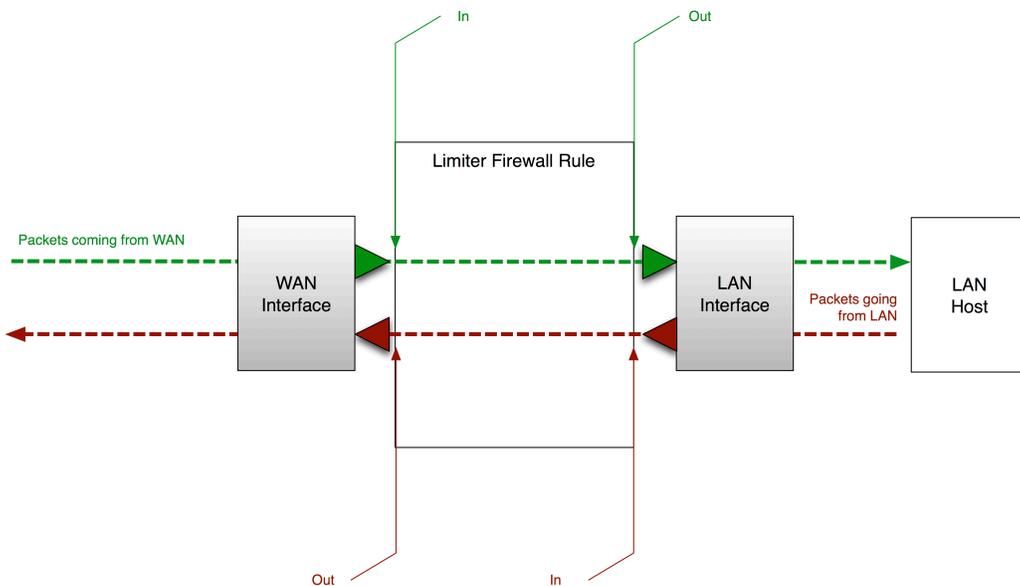


FIGURE 7-5 Sentinel PF Limiter Directions.

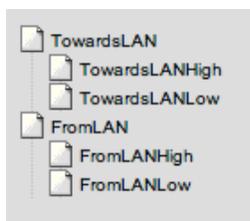
Limiters can be used to prioritize traffic from/to specific hosts or users in the local network. One way to do this would be to create a firewall alias (see *Chapter 7 - Web GUI: Firewall, Aliases* section) with IP addresses of prioritized hosts and then assign their traffic to a limiter with higher *weight*.

▼ Use Limiters To Prioritize LAN Users

- 1 In Firewall / Aliases, create an alias holding IP addresses of hosts to prioritize, e.g. alias named `vipusers`.**

This implies these LAN hosts have static IPs assigned to them, i.e. no DHCP. In this example other, non-prioritized hosts may have their addresses assigned dynamically.
- 2 In Firewall / Traffic Shaper, create two limiters for the In and Out direction.**

A suitable name for these limiters may be **TowardsLAN** and **FromLAN** respectively. The **TowardsLAN** limiter would be set up with WAN downlink bandwidth and **FromLAN** would use WAN uplink bandwidth:



- 3 In each limiter, create a «high» and «low» weight queue. An example of suitable queue names may be: `TowardsLANHigh`, `TowardsLANLow`, `FromLANHigh`, `FromLANLow`.**

High priority queues would have *weight* parameter set up in the *Advanced* section, a value of **30** would be a good start. Low priority queues would have a higher weight value, e.g. **70**.
- 4 In Firewall / Rules, LAN section, create two rules to assign traffic to limiters depending on the source.**

Here is an example of the rule to assign traffic to high weight queues:
 Action: **Pass**
 Source: **single host or alias `vipusers`**
 Destination: **any**
 In/Out (Advanced features): **`FromLANHigh` / `TowardsLANHigh`**

Here is an example of the rule to assign other traffic to low weight queues:
 Action: **Pass**
 Source: **single host or alias `!vipusers`** (everything except `vipusers`)
 Destination: **any**
 In/Out (Advanced features): **`FromLANLow` / `TowardsLANLow`**
- 5 Use Diagnostics, Limiter Info section to double check traffic is assigned to correct queues.**

Sentinel PF traffic shaper offers a method of shaping traffic based on the Layer 7 Packet Classifier. Unlike most other approaches, the Layer 7 method doesn't just look at values such as port numbers. Instead, it does regular expression matching on the application layer data to determine what protocols are being used. Since this classifier is much more processor and memory intensive than others, we recommend that you only use it if you have reason to believe that matching by standard firewall rules (by port, protocol, IP address or netmask) is insufficient for your purposes.

Layer 7 shaping is right for you if you need to:

- match any protocol that uses unpredictable ports (e.g. P2P file sharing),
- match traffic on non-standard ports (e.g. HTTP on non-80 port),
- distinguish between protocols that share a common port (e.g. P2P file sharing that uses port 80).

Sentinel PF Layer 7 Packet Classifier compares an IP packet against a known “fingerprint” that expresses how the protocol traffic should look and applies an action to the connection if it finds the match. There are some limitations to the Layer 7 approach, namely:

- Layer 7 Packet Classifier can not match encrypted traffic. Many P2P protocols, for example, are using SSL today and will evade the classifier.
- Some protocols vary too frequently to be reliably identified. A good examples of this are Bittorrent and Skype. Because these protocols change very often, the classifier’s fingerprint library must also be updated.
- Layer 7 Packet Classifier is very CPU intensive and will decrease the maximum throughput your Sentinel server will be able to process.
- Layer 7 inspection can not be used in multi-WAN environments. The inspection happens after the connection has been established, so it cannot be re-routed via an alternate path.

The way Layer 7 traffic shaping works in Sentinel PF is different from traffic shaper or limiter firewall rules. A firewall rule is created to direct the traffic to a *Layer 7 container*. Such a firewall rule must always *pass* the traffic to the container which then takes decisions on whether to *block*, *queue* or *limit* the packet matched. That is, even if you want to block the traffic via Layer 7, you use the pass action on the firewall rule, and the container is then set up to do the block action. The figure below provides an example of how Sentinel PF queues traffic via Layer 7.

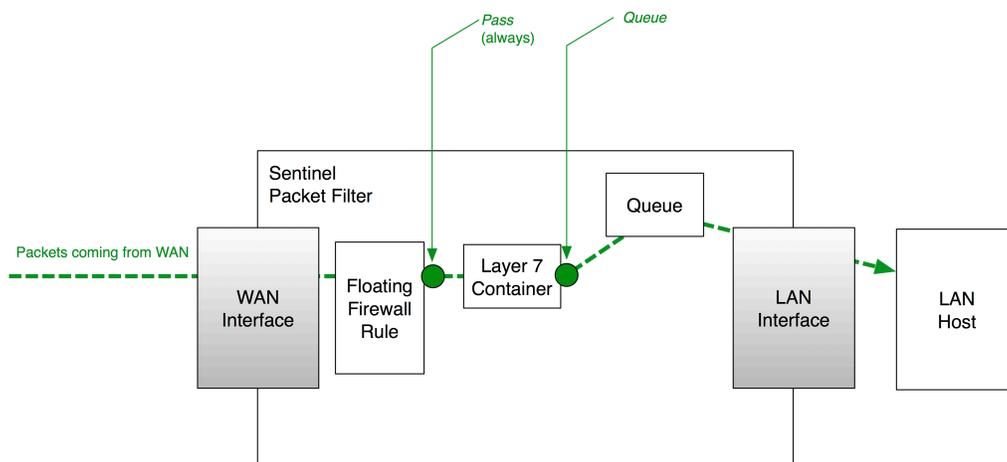


FIGURE 7-6 Sentinel PF queuing inbound traffic via Layer 7 packet inspection. The firewall rule is always configured to *pass* traffic to the Layer 7 container. The *queue* decision is taken by the container.

▼ To Shape Traffic via Layer 7

- 1 In Firewall / Traffic Shaper, select the «Layer7» tab and click on the  button in order to create a new Layer 7 container.
- 2 Click on «Enable layer7 Container».
- 3 Enter optional container description, e.g. **Business Applications**.
- 5 Click on the  button in the Rule(s) section to add protocols to the container.
- 6 Select protocol.
- 7 Choose structure.
Structure can be *action* to block traffic, *queue* to assign traffic into a shaper queue, and *limiter* to send traffic through a limiter.
- 8 Choose behavior.
For the action structure, the only option available is to *block* traffic. For queue or limiter structures, select a queue or a limiter to assign the matched traffic to.
- 9 Repeat steps 5 to 8 for any additional protocols you need in this container. Click «Save» when finished to save the container.
- 10 Create a floating firewall rule with the Layer 7 parameter set to the container. This rule must always be set to pass traffic to the container.
Please refer to the *Rules* section of this Chapter on how to create rules.

To sum up all the techniques available, traffic shaping in Sentinel PF is always initiated by a firewall rule that matches and directs traffic to a scheduler queue or a limiter. Queues can enforce QoS for different traffic types by assigning different prioritization levels, MIR limits and CIR guarantees to different traffic types. You can build scheduling disciplines using a variety of scheduling algorithms, such as HFSC, CBQ and PRIQ, each having its own advantages and limitations. Limiters can only enforce MIR limits and are usually used to define hard limits on a per-IP basis, however other creative applications are possible.

There are two ways to match traffic before sending it to a queue or limiting it. The most straightforward way is to match the traffic by parameters available in common firewall rules. This is usually one or the range of ports, protocols (TCP, UDP, etc), IP addresses or subnets. In this case, a firewall rule is created that queues traffic via *queue* action, or applies a limiter via *In/Out* firewall rule parameter.

The other way is to match traffic using Layer 7 (application layer) packet inspection. In this case each packet will be inspected and analyzed against a library of known traffic “fingerprints”. The firewall rule directing traffic to Layer 7 container should always be configured with the *pass* action, so it will pass all the traffic to the container for inspection. The actual decision to block, queue or apply limiters is then taken inside the Layer 7 container itself. It is not possible to re-route traffic via alternate path because Layer 7 inspection occurs only after the connection has been established. The Layer 7 approach is also very CPU intensive approach that may limit maximum throughput of your Sentinel server.

The option of matching traffic at the application layer is attractive, however protocol volatility is a known vulnerability of this approach. Application protocols that are IETF standards are highly unlikely to change their behavior and change their “fingerprints” suddenly, although if programs misimplement them, anything can happen. It is rather safe to use Layer 7 to match long-term IETF standards. Open source non-standardized protocols are somewhat more likely to change abruptly, but changes are likely to be publically documented and, of course, the source code can be read to learn about them as a last resort. BusinessCom issues Layer 7 fingerprint library updates on periodic basis to adapt Sentinel PF to protocol updates, and this is included in standard Sentinel support plans. Proprietary protocols like Skype can change at any time without warning. The nature of the changes may be a closely kept secret. With proprietary protocols, the Layer 7 approach is more likely to fail.

One of the most effective yet conservative bandwidth management strategy would be to assign all the traffic to default, low-priority queue, and prioritize only known applications that are important for end users. In this case, volatile protocols for non-critical applications, such as P2P and Skype, can be supported at “best effort” level, while business traffic can be assured via CBQ or PRIQ prioritization, or even CIR assignments via HFSC.

Virtual IPs

The Virtual IPs menu is used to configure Virtual IPs (VIPs) in the Sentinel PF system. VIPs can be used in a number of circumstances, such as assigning additional IP addresses to an interface, CARP redundant deployments and NAT configuration.

There are four types of VIPs:

IP Alias

Used for adding more than one IP address to a network interface. Typical uses of IP aliasing are virtual hosting and reorganizing servers without having to update any other machines. Think of an IP alias as a little kid piggy-backed on their parent interface (which is has the primary address of the physical interface). Sentinel will respond to ping on an IP Alias, and services on the firewall that bind to all interfaces will also respond on the IP Alias VIP, unless the VIP is used to forward those ports in to another device. IP Aliases on their own do not sync to XML-RPC Configuration Sync peers because that would cause an IP conflict. One exception to this is IP Alias VIPs using a CARP VIP "interface" for their interface. Those do not result in a conflict, so they do synchronize.

Proxy ARP

Proxy ARP (Address Resolution Protocol) is a technique by which a device on a given network answers the ARP queries for a network address that is not on that network. The ARP Proxy is aware of the location of the traffic's destination, and offers its own MAC address in reply, effectively saying, "send it to me, and I'll get it to where it needs to go." Serving as an ARP Proxy for another host effectively directs LAN traffic to the Proxy. The "captured" traffic is then typically routed by the Proxy to the intended destination via another interface or via a tunnel. Proxy ARP VIPs do not sync to XML-RPC Configuration Sync peers because that would cause an IP conflict.

CARP

CARP VIPs are mostly used with redundant deployments utilizing CARP. For information on using CARP VIPs, see *Redundancy* chapter.

Other

The only function of adding an Other VIP is to making that address available in the NAT configuration. This is useful when you have a public IP block routed to your WAN IP address, IP Alias, or a CARP VIP.

Note: Sentinel PF will not respond to pings destined to *Proxy ARP* and *Other* VIPs regardless of your firewall rule configuration. With *Proxy ARP* and *Other* VIPs, you must configure NAT to an internal host for ping to function. See the *NAT* section of this chapter for more information.

VIPs can be configured with a number of options:

| Option | Description |
|-----------------------|--|
| Type | Type of VIP, described above. |
| Interface | Interface the VIP is assigned to. |
| IP Address(es) | The IP address or a subnet. |
| Virtual IP Password | Password required by CARP. Otherwise can be left blank. See <i>Redundancy</i> chapter. |
| VHID Group | In a real environment, CARP interfaces will need unique identification numbers known as a VHID. This VHID or Virtual Host Identification will be used to distinguish the host on the network. See <i>Redundancy</i> chapter. |
| Advertising Frequency | CARP advertising frequency. See <i>Redundancy</i> chapter. |
| Description | Optional VIP description. |

Web GUI: Services

Services menu items are used to configure Sentinel PF network services: Network Monitor, captive portal, DHCP server and relay, DNS forwarder, Dynamic DNS services, IGMP proxy, load balancer, NTP, OLSR, PPPoE server, Proxy server, Proxy server (TP), RIP and SNMP daemons, Snort NIDS (Network Intrusion Detection System), UPnP and NAT-PMP and Wake on LAN.

Network Monitor ◆

Network Monitor is an external package developed by BusinessCom Networks. Sentinel PF relies on Network Monitor to track bandwidth usage of individual IP addresses and subnets in your network. This is one of the most used tools to troubleshoot network performance issues, as it allows you to see the top IPs of who uses how much bandwidth in your network. Network Monitor outputs chart showing bandwidth utilization over 2 day, 8 day, 40 day and 400 day periods. Additionally, each IP address utilization can be exported at intervals of 3.3 minutes, 10 minutes, 1 hour or 12 hours in a *cdf* format file, or to a backend database server.

In addition to bandwidth charts, Network Monitor provides a simple protocol breakdown, identifying FTP, HTTP, P2P, TCP, UDP and ICMP traffic.

Network Monitor can be configured with a number of options:

| Option | Description |
|-----------|---|
| Interface | The interface that Network Monitor will bind to. This is usually LAN if you want to track bandwidth utilization for hosts in your local network. |
| Subnet | Leave blank if default interface subnet is used. Otherwise, you can specify subnet(s) on which Network Monitor will report. (separate with ';' for multiple subnets, e.g. 192.168.1.0/24;10.0.0.0/24) |

(Table continued)

| | |
|----------------|---|
| Skip Intervals | Number of intervals (2.5 minute) to skip between graphing. Default 0. |
| Graph Cutoff | Graph cutoff is how many KB (kilobytes) must be transferred by an IP before it is graphed. Default 1024. |
| Promiscuous | Put interface in promiscuous mode to score to traffic that may not be routing through the host machine. This will chart traffic that is not routed via Sentinel. |
| output_cdf | Log data to cdf files. |
| recover_cdf | Read back the cdf file on startup. You will want output_cdf and recover_cdf enabled if you would like to preserve your charts after Sentinel reboots. |
| Filter | Libpcap format filter string used to control what Network Monitor sees. Please always include ip in the string to avoid strange problems. Leave blank for default settings. |
| Draw Graphs | This defaults to <i>true</i> to graph the traffic Network Monitor is recording. Set this to <i>false</i> if you only want cdf output or you are using the database output option. Network Monitor will use very little RAM and CPU if this is set to false. |
| Meta Refresh | Set META REFRESH seconds (default 150, use 0 to disable). This specifies how often your web browser will be asked to reload the Network Monitor chart to update the display. |

▼ To Show Bandwidth Chart For Individual IP

- 1 **Open Status / Network Monitor menu.**

Note: You may need to force-reload the page in your browser to see the latest Network Monitor information. On Chrome, Mozilla Firefox and Internet Explorer browsers, use **Ctrl+F5**. On Apple Safari, use **⌘+r**.

- 3 **Select Daily, Weekly, Monthly or Yearly chart by clicking on the links at the top.**
- 4 **Click on the bandwidth usage bar next to the IP address in the table to show its bandwidth utilization charts.**

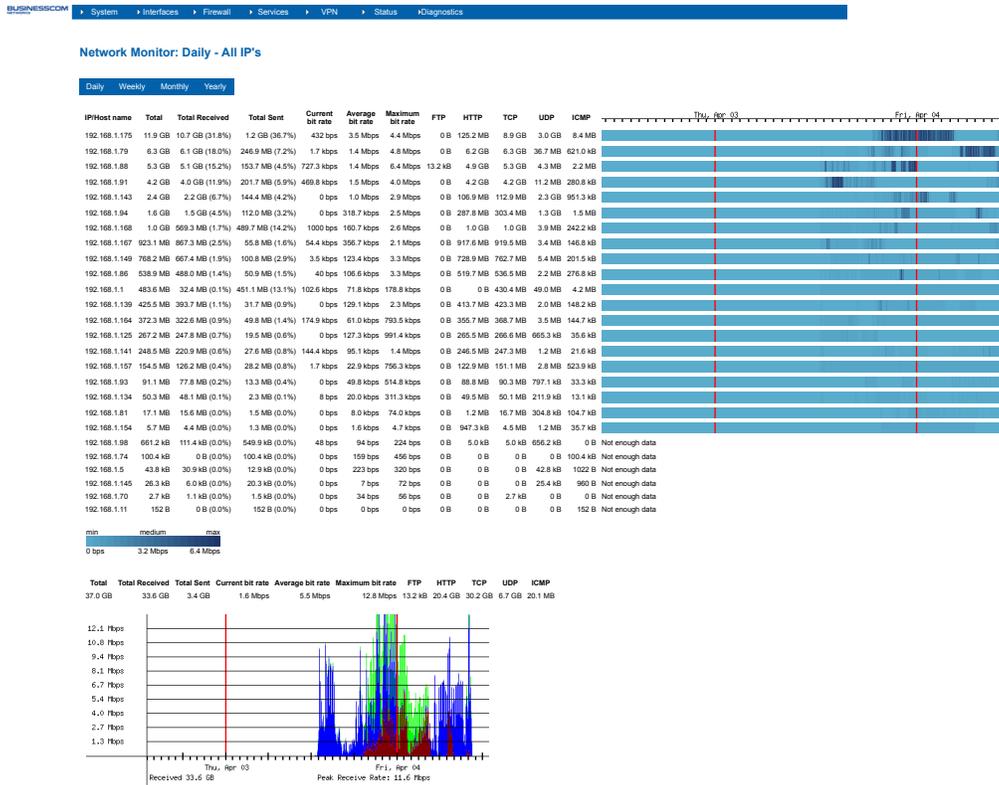


FIGURE 8–1 Example Network Monitor report.

Note: Network Monitor charts are based on long sampling rates (daily, weekly, monthly or yearly) that may average out peaks.

Captive Portal

The captive portal technique forces a host on a network to undergo the authentication process via web log in before being able to use Internet normally. A captive portal turns client's Web browser into an authentication device. This is done by intercepting all packets, regardless of address or port, until the user opens a browser and tries to access the Internet. At that time the browser is redirected to a web page which may require authentication. Captive portals are used at many Wi-Fi Hot Spots, and can be used to control wired access (e.g. apartment houses, hotel rooms, business centers, "open" Ethernet jacks) as well.

A Sentinel server can be turned into a captive portal by enabling this feature in the *Services / Captive Portal* menu.

Captive portal can be configured with a number of options:

| Option | Description |
|--|--|
| Enable captive portal | Enables captive portal on an interface. |
| Interfaces | Select the interface(s) to enable for captive portal. You can select multiple interfaces by pressing Ctrl and clicking on the interface. Usually captive portals are enabled on the LAN interface. |
| Maximum concurrent connections | This setting limits the number of concurrent connections to the captive portal HTTP(S) server, per client IP address. This does not set how many users can be logged in to the captive portal, but rather how many users can load the portal page or authenticate at the same time. Possible setting allowed is: minimum 4 connections per client IP address, with a total maximum of 100 connections. |
| Idle timeout | Clients will be disconnected after this amount of inactivity, specified in minutes. They may log in again immediately, though. Leave this field blank for no idle timeout. |
| Hard timeout | Clients will be disconnected after this amount of time, regardless of activity. They may log in again immediately, though. Leave this field blank for no hard timeout (not recommended unless an idle timeout is set). |
| Pass-through credits allowed per MAC address | This setting allows passing through the captive portal without authentication a limited number of times per MAC address. Once used up, the client can only log in with valid credentials until the waiting period specified below has expired. Recommended to set a hard timeout and/or idle timeout when using this for it to be effective. |

(Table continued)

| | |
|--|---|
| Waiting period to restore pass-through credits | Specified in hours. Clients will have their available pass-through credits restored to the original count after this amount of time since using the first one. This must be above 0 hours if pass-through credits are enabled. |
| Logout popup window | Clients will be disconnected after this amount of time, regardless of activity. They may log in again immediately, though. Leave this field blank for no hard timeout (not recommended unless an idle timeout is set). |
| Pre-authentication redirect URL | Use this field to set <code>\$PORTAL_REDIRURL\$</code> variable which can be accessed using your custom captive portal <code>index.php</code> page or error pages. |
| After authentication Redirection URL | If you provide a URL here, clients will be redirected to that URL instead of the one they initially tried to access after they've authenticated. |
| Concurrent user logins | If this option is set (i.e. disabled), only the most recent login per username will be active. Subsequent logins will cause machines previously logged in with the same username to be disconnected. |
| MAC filtering | If this option is set (i.e. disabled), no attempts will be made to ensure that the MAC address of clients stays the same while they're logged in. This is required when the MAC address of the client cannot be determined (usually because there are routers between Sentinel PF and the clients). If this is enabled, RADIUS MAC authentication cannot be used. |
| Pass-through MAC Auto Entry | <p>If this option is set, a MAC passthrough entry is automatically added after the user has successfully authenticated. Users of that MAC address will never have to authenticate again. To remove the passthrough MAC entry you either have to log in and remove it manually from the Pass-through MAC tab or send a POST from another system to remove it. If this is enabled, RADIUS MAC authentication cannot be used. Also, the logout window will not be shown.</p> <p>You can also set <i>Enable Pass-through MAC automatic addition with username</i>. If this option is set, with the automatically MAC passthrough entry created the username, used during authentication, will be saved. To remove the passthrough MAC entry you either have to log in and remove it manually from the Pass-through MAC tab or send a POST from another system to remove it.</p> |
| Per-user bandwidth restriction | If this option is set, the captive portal will restrict each user who logs in to the specified default bandwidth. RADIUS can override the default settings. Leave empty or set to 0 for no limit. |

(Table continued)

| | |
|--------------------|---|
| Authentication | <p>Defines how users are authenticated. There are three options:</p> <p><i>No authentication</i>: all authentication is disabled.</p> <p><i>Local User Manager / Vouchers</i>: authentication via users created in Sentinel PF User Manager or via Captive Portal Vouchers.</p> <p><i>RADIUS Authentication</i>: authentication via external RADIUS server.</p> <p>If you specify <i>RADIUS Authentication</i>, then you will need to enter IP address, Port and Shared Secret of a primary and, optionally, secondary RADIUS servers in your network. You can also enable <i>send RADIUS accounting packets</i> so RADIUS accounting packets will be sent to the primary RADIUS server.</p> <p>Additionally, RADIUS authentication will support re-authentication feature. If reauthentication is enabled, Access-Requests will be sent to the RADIUS server for each user that is logged in every minute. If an Access-Reject is received for a user, that user is disconnected from the captive portal immediately.</p> <p>Instead of authenticating via username and password, you can use <i>RADIUS MAC Authentication</i>. If this option is enabled, the captive portal will try to authenticate users by sending their MAC address as the username and the password entered below to the RADIUS server.</p> <p>Full RADIUS functionality will not be entered in this book. Please refer to additional documentation that is available from BusinessCom.</p> |
| MAC Address Format | <p>This option changes the MAC address format used in the whole RADIUS system. Change this if you also need to change the username format for RADIUS MAC authentication.</p> <p><i>default</i>: 00:11:22:33:44:55 <i>singledash</i>: 001122-334455 <i>ietf</i>: 00-11-22-33-44-55 <i>cisco</i>: 0011.2233.4455 <i>unformatted</i>: 001122334455</p> |
| | <p>HTTPS login, server name, certificate, private key and intermediary certificate.</p> <p>If enabled, the username and password will be transmitted over an HTTPS connection to protect against eavesdroppers. A server name, certificate and matching private key must also be specified. The <i>server name</i> will be used in the form action for the HTTPS POST and should match the Common Name (CN) in your certificate (otherwise, the client browser will most likely display a security warning). Make sure captive portal clients can resolve this name in DNS and verify on the client that the IP resolves to the correct interface IP on Sentinel PF. Certificates are to be passed in X.509 PEM format, and the RSA private key should be in PEM format.</p> |

You can also specify *portal page*, *authentication error page* and *logout page* contents to customize your captive portal. This can be used to create custom branded portal pages. You can upload an HTML/PHP file for your custom pages while configuring the captive portal, or you can leave these fields blank to keep the default Sentinel PF content. For the *portal page*, make sure to include a form (POST to "") with a submit button (name="accept") and a hidden field with name="redirurl" and value="". Include the "auth_user" and "auth_pass" and/or "auth_voucher" input fields if authentication is enabled, otherwise it will always fail. Example code for the form:

```
<form method="post" action="$PORTAL_ACTION$">
  <input name="auth_user" type="text">
  <input name="auth_pass" type="password">
  <input name="auth_voucher" type="text">
  <input name="redirurl" type="hidden"
value="$PORTAL_REDIRURL$">
  <input name="accept" type="submit" value="Continue">
</form>
```

The contents of the *authentication error page* are displayed when an authentication error occurs. You may include `$PORTAL_MESSAGE$`, which will be replaced by the error or reply messages from the RADIUS server, if any. The contents of the *logout page* displayed on authentication success when the logout popup is enabled.

You can use the File Manager tab to upload files to your captive portal HTTP(S) server. Any files that you upload here with the filename prefix of *captiveportal-* will be made available in the root directory of the captive portal HTTP(S) server. You may reference them directly from your portal page HTML code using relative paths. Example: you've uploaded an image with the name *captiveportal-test.jpg* using the file manager. Then you can include it in your portal page like this:

```

```

In addition, you can also upload *.php* files for execution. You can pass the filename to your custom page from the initial page by using text similar to:

```
<a href="/captiveportal-
aup.php?redirurl=$PORTAL_REDIRURL$">Acceptable usage policy</a>
```

The total size limit for all files is 1 MB.



Caution – Changing captive portal configuration will disconnect all clients. Don't forget to enable the DHCP server on your captive portal interface. Make sure that the default/maximum DHCP lease time is higher than the timeout entered on this page. Also, the DNS forwarder needs to be enabled for DNS lookups by unauthenticated clients to work.

Additional tabs of the captive portal menu are discussed below.

Pass-through MAC, Allowed IP addresses and Allowed Hostnames tabs are used to manage MAC and IP addresses and hostnames defined as *pass-through*. This allows those users to access Internet through the captive portal automatically without being taken to the portal page. Specifying a pass-through IP and hostname can be used for a web server serving images for the portal page, or a DNS server on another network, for example. By specifying *from* addresses, it may be used to always allow pass-through access from a client behind the captive portal.

Sentinel PF captive portal provides a tool to generate and manage *vouchers*. Vouchers are special codes that can be used to gain Internet access through the captive portal. Each roll of vouchers is cryptographically generated and includes a set time limit. Vouchers are commonly implemented in places authenticated, but time-limited, Internet access is desired without needing to provide a username and password to users. For example, in coffee shops, hotels, and airports. Users simply enter their voucher code in the portal login form and are granted access for as long as the voucher is valid. *Voucher rolls* can be exported as a *CSV* file, and some companies have even integrated the exported voucher lists into their point of sale applications to print a voucher on customer receipts.

Please note that the voucher time limit does not stop if the user log outs. The time counter begins when the user logins first and ends at the set time limit. Only one concurrent connection can be established with a voucher. A voucher can be used from multiple computers, however only one concurrent connection will be allowed. Vouchers are encrypted using a pair of *public key* and *private key*.



Caution – Do not change keys after vouchers have been generated. This will render them useless.

The database with stored vouchers can be synchronized to a slave Sentinel server. This works similarly to the XML-RPC configuration synchronization. When configured, the synchronization will copy the voucher rolls to the target Sentinel server, and also push information about active vouchers to the target unit as the vouchers are used.

Vouchers can be managed from the Services / Captive Portal menu, Vouchers tab. They can be configured with a number of options:

| Option | Description |
|-----------------|---|
| Enable Vouchers | Enables vouchers. |
| Voucher Rolls | Create, generate and activate Rolls with Vouchers that allow access through the captive portal for the configured time. Click on the <i>add (+)</i> button to generate a new roll. You can download the voucher roll in the <i>CSV</i> format file by clicking on the “information” button. |

(Table continued)

| | |
|---------------------------------|---|
| Voucher public key | Paste an RSA public key (64-bit or smaller) in PEM format here. This key is used to decrypt vouchers. Click on <i>Generate new key</i> to automatically generate a key. |
| Voucher private key | Paste an RSA private key (64-bit or smaller) in PEM format here. This key is only used to generate encrypted vouchers and doesn't need to be available if the vouchers have been generated offline. Click on <i>Generate new key</i> to automatically generate a key. |
| Character Set | Voucher codes are generated with the specified character set. It should contain printable characters (numbers, lower case and upper case letters) that are hard to confuse with others. Avoid e.g. 0/O and l/1. |
| # of Roll Bits | Reserves a range in each voucher to store the Roll # it belongs to. Allowed range: 1..31. Sum of Roll+Ticket+Checksum bits must be one bit less than the RSA key size. |
| # of Ticket Bits | Reserves a range in each voucher to store the Ticket# it belongs to. Allowed range: 1..16. Using 16 bits allows a roll to have up to 65535 vouchers. A bit array, stored in RAM and in the config, is used to mark if a voucher has been used. A bit array for 65535 vouchers requires 8 KB of storage. |
| # of Checksum Bits | Reserves a range in each voucher to store a simple checksum over Roll # and Ticket#. Allowed range is 0..31. |
| Magic Number | Magic number stored in every voucher. Verified during voucher check. Size depends on how many bits are left by Roll+Ticket+Checksum bits. If all bits are used, no magic number will be used and checked. |
| Invalid Voucher Message | Error message displayed for invalid vouchers on captive portal error page (<i>\$PORTAL_MESSAGE\$</i>). |
| Expired Voucher Message | Error message displayed for expired vouchers on captive portal error page (<i>\$PORTAL_MESSAGE\$</i>). |
| Synchronize Voucher Database IP | IP address of master nodes to synchronize voucher database and used vouchers from. <i>Note: Synchronize Voucher Database IP should be setup on the slave nodes and not the master node!</i> |
| Voucher Sync Port | This is the port of the master voucher nodes. Example: 443 |
| Voucher sync username | Username of the master voucher nodes. |
| Voucher sync password | Password of the master voucher nodes. |

Note: Specifying the Voucher Database Synchronization options will not record any other value from the other options. They will be retrieved/synced from the master.

▼ To Start Using Vouchers In Your LAN

- 1 In Firewall / Captive Portal, select the «Vouchers» tab and click «Enable Vouchers».
- 2 In the «Voucher Rolls», click on the  button to add a new voucher roll.
- 3 Enter Roll number, minutes per ticket (voucher) and count (number of vouchers you need).
- 4 Enter optional voucher roll comment, e.g. `Vouchers for Hotel Guests`.
- 5 The voucher roll should appear in the «Voucher Rolls» table. Click on the  button to download the voucher roll in the *CSV* format to your computer.
The file will contain voucher codes that you can distribute to your customers.
- 6 In Firewall / Captive Portal, select the «Captive portal» tab.
- 7 Change Authentication to «Local User Manager / Vouchers».
- 8 Click «Save».
If DHCP server is already enabled on the LAN interface then stop here, otherwise continue with steps below.
- 9 In Services / DHCP server, select the «LAN» tab.
- 10 Click «Enable DHCP server on LAN interface».
- 11 Enter the IP address range for the DHCP pool. E.g. `192 . 168 . 1 . 100 to 192 . 168 . 1 . 200`.
- 12 Click «Save».
- 13 Test your configuration. Once accessed from the LAN, the client web browser must be presented with the captive portal and «Enter Voucher Code» prompt before Internet access is allowed.
If you haven't been using DHCP connection during this procedure, you may need to set your computer's LAN IP address to be obtained automatically via DHCP in order to be routed via captive portal.

Note: The procedure above describes the most basic voucher usage setup. You may need to adjust various settings to configure this service to match your requirements. Information on current vouchers usage in your captive portal is available in the Status / Captive Portal menu.

Depending on your networking environment and how firewall rules are set up, you may also want to block *pass* rules from LAN to Sentinel in order make sure local hosts will not bypass your captive portal.

DHCP Relay and DHCPv6 Relay

In case you need to provide DHCP service to part of the network outside of the common segment, you can relay DHCP broadcasts on that network to a DHCP server from another network.

▼ To Enable DHCP Relay

- 1 **Open Services / DHCP Relay menu.**
- 2 **Select interface or multiple interfaces for DHCP Relay to listen on.**

Note: You can not enable DHCP Relay on an interface that has DHCP Server running on it.

- 3 **Select whether to append circuit ID and agent ID to requests.**
If this is checked, the DHCP relay will append the circuit ID (Sentinel PF interface number) and the agent ID to the DHCP request.
- 4 **Specify desination DHCP server.**
This is the IP address of the server to which DHCP requests are relayed. You can enter multiple server IP addresses, separated by commas.
- 5 **Click «Save»**

DHCP Server and DHCPv6 Server/RA

Sentinel PF can provide DHCP service to any network connected to it. The Services / DHCP Server menu has multiple tabs, per each interface. Those tabs are identical, so only one instance will be described below.

DHCP Server can be configured with a number of options:

| Option | Description |
|----------------------|---|
| Enable DHCP Server | Enables DHCP Server on the selected interface. |
| Deny unknown clients | If this is checked, only the clients defined in the static mapping table will get DHCP leases from this server. |

Note: This option should not be used as a master security measure. It is possible for an attacker to spoof MAC address and hard-code IP address configuration to obtain a lease. Additionally security measures, such as captive portal authentication and static ARP entries should be used to secure the system from unauthroized access.

(Table continued)

| | |
|--------------------|---|
| Range | The DHCP IP pool. Must be within Available Range indicated above this setting. |
| WINS Servers | WINS Servers that will be announced to DHCP leases. |
| DNS Servers | DNS Servers that will be announced to DHCP leases. Leave blank to use default DNS servers configured in the Sentinel PF operating system in <i>General Setup</i> section. |
| Gateway | Default gateway that will be announced to DHCP leases. Leave blank to use Sentinel as the default gateway, which is the most common use. When using CARP redundancy, fill in the CARP IP on this interface here. |
| Domain Name | Domain name that will be announced to DHCP leases. Leave blank to use default domain configured in the Sentinel PF operating system in <i>General Setup</i> section. |
| Domain Search List | Domain Search List controls the DNS search domains that are provided to the client via DHCP. If you have multiple domains and you use short hostnames on them, provide a list of domain names here, separated by a semicolon. Clients will attempt to resolve hostnames by adding the domains, in turn, from this list before trying to find them externally. If left blank, the <i>Domain Name</i> option is used. |
| Default Lease Time | This is used for clients that do not ask for a specific expiration time. The default is 7200 seconds. |
| Maximum Lease Time | This is the maximum lease time for clients that ask for a specific expiration time. The default is 86400 seconds. |
| Failover Peer IP | If this system is part of a failover setup such as a CARP cluster, enter the Failover peer IP address. This should be the real IP address of the other system in this subnet, not a shared CARP address. |
| Static ARP | This enables static ARP entries. <u>Note: If this option is enabled, only the machines listed in the static ARP table will be able to communicate with the firewall on this interface.</u> |
| Dynamic DNS | Enter the dynamic DNS domain which will be used to register client names in the DNS server. |
| NTP Servers | Default NTP (Network Time Protocol) servers that will be announced to DHCP leases. |
| TFTP Servers | Default TFTP server that will be announced to DHCP leases. This is usually referred as <i>DHCP option 66</i> . |
| LDAP URI | LDAP URI will send an LDAP server URI to the client if requested. This is usually referred as <i>DHCP option 95</i> . It should take the form of a fully qualified LDAP URI, such as ldap://ldap.example.com/dc=example,dc=com . |

(Table continued)

| | |
|-------------------------------|---|
| Enable network booting | Enter an IP address from which boot images are available, and a file name for the boot image. Both of these fields must be configured for network booting to work properly. You may also optionally specify a Root Path String to target a specific device as the client's root filesystem device, such as iscsi:(servername):(protocol):(port):(LUN):targetname. |
| Additional BOOTP/DHCP Options | Enter the DHCP option number and the value for each item you would like to include in the DHCP lease information. For additional information on BOOTP/DHCP options, please refer to <i>IANA - Dynamic Host Configuration Protocol (DHCP) and Bootstrap Protocol (BOOTP) Parameters</i> here: http://www.iana.org/assignments/bootp-dhcp-parameters/bootp-dhcp-parameters.xml |

Static DHCP mappings allow you to control which PCs can obtain DHCP leases. A static mapping table allows to bind a combination of IP and MAC addresses to a known user, and allow that user to obtain a lease. The static mapping can also be configured without IP, by specifying just the MAC address so the IP address will be allocated from the default DHCP pool.

▼ To Add Static Mapping

- 1 **In Firewall / DHCP Server, select the interface of the DHCP server, scroll to the bottom of the page and click on the  button in order to create a new static mapping.**
- 2 **Enter MAC address of a known client.**
- 3 **Optionally, enter IP address and a Hostname of a known client.**
If IP address is specified then this will specify the preferred IP address for this known client in the pool. This is *not* a reservation. If the preferred IP address is already taken by someone else then Sentinel DHCP Server will assign another IP from the default DHCP pool.
- 3 **Enter optional description of a client, e.g. Pete 's workstation.**
- 4 **Click «Save»**

DNS Forwarder

The DNS Forwarder is a caching DNS resolver. With DNS Forwarder enabled, Sentinel becomes a DNS server that forwards all requests to the DNS servers configured in the System / General Setup menu.

Note: If the DNS forwarder is enabled, the DHCP Server (if enabled) will automatically serve the LAN IP address as a DNS server to DHCP clients so they will use the forwarder. The DNS forwarder will use the DNS servers entered in System / General setup or those obtained via DHCP or PPP on WAN if the *Allow DNS server list to be overridden by DHCP/PPP on WAN* is checked. If you don't use that option (or if you use a static IP address on WAN), you must manually specify at least one DNS server on the System / General setup page.

▼ To Enable DNS Forwarder

- 1 **Open Services / DNS Forwarder menu.**
- 2 **Click «Enable DNS Forwarder»**
- 3 **Optional: select «Register DHCP leases in DNS forwarder».**
If this option is set, then machines that specify their hostname when requesting a DHCP lease will be registered in the DNS forwarder, so that their name can be resolved. You should also set the domain in System / General setup to the proper value.
- 4 **Optional: select «Register DHCP static mappings in DNS forwarder».**
If this option is set, then DHCP static mappings will be registered in the DNS forwarder, so that their name can be resolved. You should also set the domain in System / General setup to the proper value.
- 5 **Optional: select «Resolve DHCP mappings first».**
If this option is set, then DHCP mappings will be resolved before the manual list of names below. This only affects the name given for a reverse lookup (PTR).
- 6 **Enter advanced options if required.**
- 7 **Click «Save».**

You can also specify DNS host overrides for the DNS forwarder. This is used to force Sentinel PF to resolve a specific host differently than it otherwise would via default DNS servers. A split DNS infrastructure is a solution to the problem of using the same domain name for internally and externally accessible resources. It allows to resolve the difference in how internal and external clients should access resources using the same hostname.

Similarly, you can specify overrides on the domain level, by specifying authoritative DNS server to be used for this domain. One example of where this is commonly deployed is in small business networks with a single internal server with Active Directory, usually Microsoft Small Business Server. The DNS requests for the Active Directory domain name must be resolved by the internal Windows Server for Active Directory to function properly. Adding an override for the Active Directory domain pointing to the internal Windows server's IP address ensures these records are resolved properly whether clients are using Sentinel PF as a DNS server or the Windows Server itself.

▼ To Add DNS Host Override

- 1 In Services / DNS Forwarder, scroll to the «Host Overrides» table.
- 2 Click on the  button to add a new override.
- 3 Specify Host, Domain and IP address for the override.
- 4 Enter optional override description, e.g. `Access For Local Clients`.
- 5 Click «Save»

▼ To Add DNS Domain Override

- 1 In Services / DNS Forwarder, scroll to the «Domain Overrides» table.
- 2 Click on the  button to add a new override.
- 3 Specify Domain and IP address of the authoritative DNS server for this domain.
- 4 Enter optional override description, e.g. `Resolve for AD`.
- 5 Click «Save»

Dynamic DNS

Dynamic DNS is a method of maintaining a Domain Name System (DNS) record in a server, pointing to a changing IP address of a host. Dynamic DNS services are commonly used to assign static hostname to a server with dynamic IP address, such as servers obtaining their IP address via DHCP on the WAN connection.

Sentinel PF Dynamic DNS service supports proprietary Dynamic DNS updates, such as those found on services like DynDNS, freeDNS, OpenDNS, HE.net, DHS and others – via the DynDNS tab. Sentinel also supports *RFC 2136* style updates to servers that support this RFC – via the RFC 2136 tab.

▼ To Use Dynamic DNS

- 1 In **Services / Dynamic DNS**, select the corresponding tab (**DynDNS** or **RFC 2136**).
- 2 Click on the  button to add a new Dynamic DNS client.
- 3 **Select the network interface to monitor.**
This is the interface that will provide the IP address to Sentinel Dynamic DNS service.
- 4 **Enter desired hostname.**
With RFC 2136 style updates, enter TTL, key name, key type, HMAC-HD5 key and target DNS server IP. You can also specify to use TCP instead of UDP.
- 5 **With DynDNS servers, enter desired MX record.**
- 6 **With DynDNS servers, enable wildcards, if necessary.**
Enabling wildcards will resolve all sub-domains to the IP address of your host name. For example, if your host name is *example.dyndns.org*, enabling wildcard will make **.example.dyndns.org* resolve the same as *example.dyndns.org*.
- 7 **With DynDNS servers, enter username and password to your Dynamic DNS service provider.**
- 8 **Enter optional Dynamic DNS service description, e.g. `Resolve for example.dyndns.org`.**
- 9 **Click «Save».**

Note: You must configure a DNS server in System / General setup, or allow the DNS server list to be overridden by DHCP/PPP on WAN for dynamic DNS updates to work.

IGMP Proxy

The Internet Group Management Protocol (IGMP) is a communications protocol used by hosts and adjacent routers on IP networks to establish multicast group memberships. IGMP is an integral part of IP multicast. Sentinel PF IGMP Proxy service is forwarding IGMP frames and commonly is used when there is no need for more advanced protocol like PIM.

▼ To Enable IGMP Proxy

- 1 In Services / IGMP Proxy, click on the  button to add a new proxy.
- 2 Select network interface to be used by the proxy.
- 3 Enter optional proxy description, e.g. **IGMP Proxy 1**.
- 4 **Select Interface Type.**

The *upstream network interface* is the outgoing interface which is responsible for communicating to available multicast data sources. There can only be one upstream interface.

Downstream network interfaces are the distribution interfaces to the destination networks, where multicast clients can join groups and receive multicast data. One or more downstream interfaces must be configured.
- 5 **Enter TTL Threshold.**

Defines the TTL threshold for the network interface. Packets with a lower TTL than the threshold value will be ignored. This setting is optional, and by default the threshold is 1.
- 6 **Add CIDR-masked Network entries to control what subnets are allowed to have their multicast data proxied.**

In the *Network(s)* section, click on the  button to add a network with CIDR mask.
- 7 Click «Save».

Load Balancer

The Sentinel PF Load Balancer service is used to load balance servers for the purpose of distributing traffic between multiple internal services for load balancing and redundancy purposes. Please note this is *not* load balancing for multiple WAN connections. To make a distinction between WAN load balancing and the Load Balancer service, we will introduce the *server load balancing* term, and this technique is described in this chapter.

Server load balancing is most commonly used to distribute load among multiple Web servers running in the local network. It can also be used to load balance any other TCP based service, such as SMTP servers, or for DNS. The Sentinel PF Load Balancer service is able to monitor groups of hosts for availability, which is determined by checking for a specific service common to a host group. When availability is confirmed, Layer 3 and/or Layer 7 forwarding services are set up to allow the connection.

Sentinel PF Load Balancer configuration consists of multiple sections:

Pools define the list of internal servers to be load balanced, describe which port the services are provided and the monitoring method to be used by the Load Balancer service.

Virtual Servers define the IP and port for the load balancing virtual server to listen on, and the appropriate pool to direct the incoming traffic to.

Monitors define custom monitoring methods to confirm internal server availability.

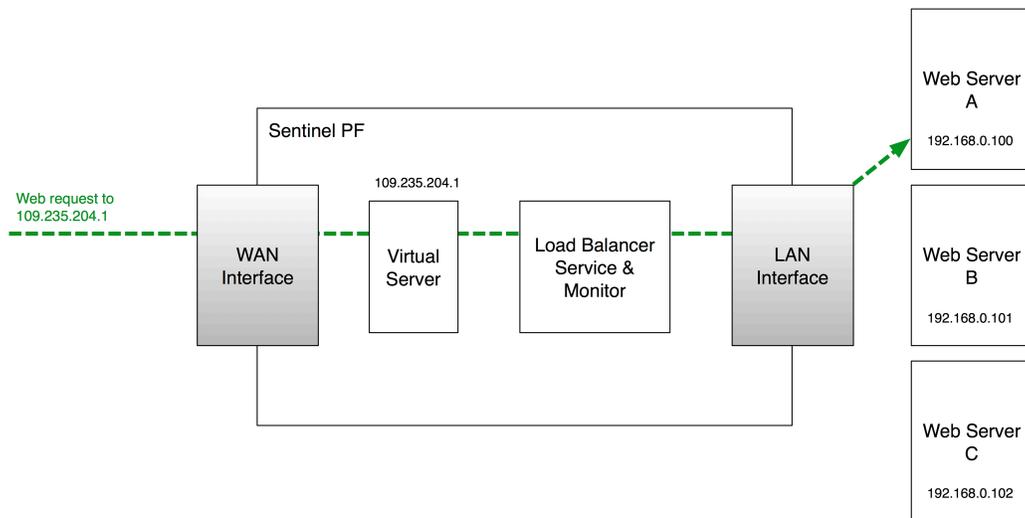


FIGURE 8-2 Sentinel PF Load Balancer Service Operation Example.

▼ To Create Internal Server Pool

- 1 In Services / Load Balancer, select the «Pools» tab and click on the  button to add a new pool.
- 2 Enter pool name, e.g. `webservers`.
- 3 Enter optional pool description, e.g. `web Servers Pool`.
- 4 Enter port pool servers are listening on.
- 5 In the «Add Item to Pool» section, specify the monitor type and server IP address to be added into the pool.
Default monitor types defined are *ICMP*, *TCP*, *HTTP*, *HTTPS* and *SMTP*. You can define custom monitors in the *Monitors* tab.
- 6 **Review pool members.**
You can enable and disable individual servers in the pool by moving them to the *Enabled* and *Pool Disabled* sections with > and < buttons in the *Current Pool Members* section.
- 7 Click «Save».

▼ To Create Virtual Server

- 1 In Services / Load Balancer, select the «Virtual Servers» tab and click on the  button to add a new virtual server.
- 2 Enter virtual server name, e.g. `webserver`.
- 3 Enter optional server description, e.g. `Load Balanced Virtual Web Server`.
- 4 Enter IP address and port for the virtual server to listen on.
- 5 Specify virtual server pool incoming traffic to be directed to.
- 6 **Optionally, specify Fail Back Pool.**
This is the alternate pool that clients are directed to if all the servers in your primary pool are down. This could be a pool with one server with a maintenance announcement page, for example.
- 7 **Select relay protocol.**
By default this is the *TCP* mode which acts as the port forward. Your servers will see source IP of the client. In the *DNS* mode, the load balancer acts as a DNS proxy. In the *DNS* mode, your clients will see Sentinel as the source.
- 8 Click «Save».

NTP

The Sentinel PF NTP (Network Time Protocol) service enables Sentinel to act as the NTP server on selected interfaces. The clocks of NTP clients will be synchronized to Sentinel's clock. There is virtually no configuration required for the NTP service except for the selection of interface(s) the service has to be provided on.

Note: Sentinel NTP service is available on all interfaces by default. Selecting no interfaces will listen on all interfaces with a wildcard. Selecting all interfaces will explicitly listen on only the interfaces/IPs specified.

OLSR

OLSR (Optimized Link State Routing) is a routing protocol for mobile ad-hoc networks, also known as MANETs. The protocol is pro-active, table driven and utilizes a technique called multipoint relaying for optimized message flooding. The Sentinel PF OLSR service also implements a popular optional *link quality* extension. Sentinel PF OLSR implementation is *RFC 3626* compliant with respect to both core and auxiliary functioning.

Sentinel PF OLSR service can be configured with a number of options:

| Option | Description |
|----------------------------------|--|
| Enable OLSR | Enables the dynamic mesh linking daemon. |
| Link Quality Level | Specifies <i>link quality</i> OLSR extension level. |
| Interfaces | Select the interfaces that OLSR will bind to. |
| Enable HTTPInfo Plugin | Enables the OLSR stats web server. |
| HTTPInfo Port | Port that HTTPInfo will listen on. |
| Allowed host(s) | Hosts that are allowed to access the HTTPInfo web service. |
| Allowed host(s) subnet | Enter the subnet mask in form 255.255.255.0. |
| Enable Dynamic Gateway | Enables the OLSR Dynamic Gateways feature. |
| Announce self as Dynamic Gateway | Enables the OLSR Dynamic Gateways Announcing feature. |
| Announce Dynamic local route | IP/netmask of Dynamic local route. |
| Ping and Poll | Pings this host to ensure connectivity. <i>Poll</i> specifies how often to look for <i>inet gw</i> , in seconds. |
| Enable Secure Mode and Key | Enables the secure mode. The <i>Key</i> section holds the secure OLSR key information. |

PEP

Sentinel PF provides an external package for client access to the BusinessCom PEP (Performance Enhancing Proxy), which is a WAN optimization solution. The current protocol version implemented is 4.2.0. PEP client establishes a connection with BusinessCom PEP gateway and opens a local HTTP proxy or TCP socks port to route traffic through.

PEP can be configured with a number of options:

| Option | Description |
|-------------------------------|---|
| Enable PEP Daemon | Enables the PEP client daemon. |
| Username | Your PEP gateway username, assigned by BusinessCom. |
| Password | Your PEP gateway password, assigned by BusinessCom. |
| PEP Gateway IP and UDP Port | This is the IP address or hostname and UDP port of the PEP gateway, separated by a colon. For example: gw-london-1.bcsate11ife.net:200. |
| PEP Client Outbound UDP Port | Outbound port for the PEP UDP traffic, default is 2002. |
| Download Bandwidth (min, max) | Specify bandwidth (in kbps) to use for the download direction, from PEP gateway to PEP client. Specify minimum and maximum in case of variable bitrate (shared) connection. |
| Upload Bandwidth (min, max) | Specify bandwidth (in kbps) to use for the upload direction, from PEP client to PEP gateway. Specify minimum and maximum in case of variable bitrate (shared) connection. |
| MTU | Maximum Transfer Unit in bytes, default is 500. |
| MRU | Maximum Receiving Unit in bytes, default is 1400. |
| Flush Time | Buffer flush time, specified in ms, default is 50. Long flush times will improve outbound traffic compression, however they are not recommended for burstable or low traffic environments, as the responsiveness of real-time applications may be impaired. 500 ms may be a good option for large DVB or SCPC trunks. |
| Local IP Address and Port | This is the local IP address and port of the PEP service, separated by a colon. For example: 192.168.1.1:3000. Note: Default port is 3000. Port must be in the 1035-65535 range. We do <i>not</i> recommend using port 3128, as this may conflict with the default Squid configuration. |
| PEP Service Type | Select the service type PEP shall provide via the <i>Local IP Address</i> : <i>HTTP</i> , <i>SOCKS</i> or <i>Compressed HTTP</i> (HTTP with image re-compression and Web content optimization). |
| Enable Logging | Enable PEP daemon logging. Note: Sentinel D2, D2W and other NanoBSD-based Sentinel models should keep this option turned off due to limited storage. |

Note: Compressed HTTP service type may not be available on all PEP gateways. Please consult with BusinessCom representative before enabling this option. This corresponds to internal service #3.

▼ For Transparent PEP HTTP Redirect

1 In Services / PEP, configure PEP client daemon.

Use *HTTP* or *Compressed HTTP* as PEP service type. Make sure PEP client daemon is running. Please refer to the above section of this Chapter for additional information on the PEP client configuration.

2 In Firewall / NAT, select «Port Forward» tab and create a port forward rule.

Interface: LAN

Protocol: TCP

Source Address: any

Source Ports: any

Destination Address: any

Destination Ports: 80

NAT IP: 127.0.0.1 (should equal *Local IP Address and Port* of PEP daemon)

NAT Ports: 3000 (should equal *Local IP Address and Port* of PEP daemon)

3 Create similar port forward rules for other interfaces facing users.

4 Check if your HTTP sessions are routed via PEP.

A simple way would be to point your web browser to *http://www.whatismyip.com* and check the IP detected. A valid response with connection initiated via PEP gateway would be as following:

Your IP Address is:

PEP Gateway IP Address

Proxy Detected:

PEP Gateway IP Address

1.1 localhost:3128 (squid/2.7.STABLE9), 1.0 gw-falkenstein:12345 (squid/2.7.STABLE9)

(*gw-falkenstein* is the name of BusinessCom PEP gateway in Germany used in illustrative purposes only, your actual PEP gateway may be different.)

PPPoE Server

Sentinel PF PPPoE (Point-to-Point Protocol over Ethernet) service enables Sentinel to implement a network protocol for encapsulating PPP frames inside Ethernet frames. This service can be used to make Sentinel to act as an access concentrator for the local network by forcing users to authenticate before gaining network access.

The PPPoE server can be enabled in the Services / PPPoE Server menu, by clicking the  button and adding a new server. By default, it uses Sentinel PF PPPoE server user database to authenticate PPPoE clients. PPPoE server can optionally use an external RADIUS server for authentication.

PPPoE server can be configured with a number of options:

| Option | Description |
|--------------------------|---|
| Enable PPPoE server | Enables server on the selected interface. |
| Interface | Interface for the server to bind to. |
| Subnet netmask | Enter subnet mask as bit counts (as in CIDR notation) to be assigned to PPPoE clients. E.g. entering 24 will result in 255.255.255.0 netmask selected. See <i>Appendix B – Netmask/CIDR Translation Table</i> . |
| No. PPPoE users | Amount of PPPoE clients to allow. |
| Server address | Enter the IP address the PPPoE server should give to clients for use as their <i>gateway</i> . Typically this is set to an unused IP just outside of the client range. Note: This should NOT be set to any IP address currently in use on this firewall. |
| Remote address range | Specify the starting address for the client IP address subnet. |
| Description | Optional description for this PPPoE service. |
| DNS servers | If entered these DNS servers will be announced to all PPPoE clients, otherwise LAN DNS and one WAN DNS will be announced to all clients. |
| RADIUS | When <i>Use a RADIUS server for authentication</i> is set, all users will be authenticated using the RADIUS server specified below. The local PPPoE server user database will not be used. You can optionally enable <i>RADIUS accounting</i> , and use a <i>Backup RADIUS server</i> . |
| NAS IP Address | RADIUS server NAS IP Address |
| RADIUS Accounting Update | RADIUS accounting update period in seconds. |

(Table continued)

| | |
|--------------------------------|---|
| RADIUS issued IPs | Issue IP Addresses via RADIUS server. |
| RADIUS server Primary | Enter the IP address, authentication port and accounting port (optional) of the RADIUS server. Standard ports are 1812 and 1813 for RADIUS accounting. |
| RADIUS primary shared secret | Enter the shared secret that will be used to authenticate to the RADIUS server. |
| RADIUS server Secondary | Enter the IP address, authentication port and accounting port (optional) of the backup RADIUS server. Standard ports are 1812 and 1813 for RADIUS accounting. |
| RADIUS secondary shared secret | Enter the shared secret that will be used to authenticate to the backup RADIUS server. |
| User(s) | This is local PPPoE server users database that will be used to authenticate clients when RADIUS is <i>not</i> used. |

Proxy Server ◆

Sentinel PF features *Squid* which is an external package that provides caching proxy functionality for Web traffic. Squid is originally derived from the DARPA (Defense Advanced Research Projects Agency) funded *Harvest* project. It supports HTTP, HTTPS, FTP and *gopher* protocols. The primary function of a proxy server is to reduce bandwidth and improve response times by caching and reusing frequently-requested web pages. Squid keeps meta data and especially hot objects cached in RAM, caches DNS lookups, supports non-blocking DNS lookups, and implements negative caching of failed requests. It supports SSL, extensive access controls, and full request logging. By using the lightweight Internet Cache Protocol, Squid caches can be arranged in a hierarchy or mesh for additional bandwidth savings.

The proxy server can be configured with a number of options, separated into multiple tabs:

General Settings

| Option | Description |
|--------------------------|--|
| Proxy Interface | The interface(s) the proxy server will bind to. This is usually set to <i>LAN</i> interface if you want to provide proxy services to your local network. |
| Allow users on interface | If this field is checked, the users connected to the interface selected in the <i>Proxy interface</i> field will be allowed to use the proxy, i.e., there will be no need to add the interface's subnet to the list of allowed subnets. This is just a shortcut. |
| Transparent proxy | If transparent mode is enabled, all requests for destination port 80 will be forwarded to the proxy server without any additional configuration necessary. |

(Table continued)

| | |
|---|--|
| Bypass proxy for Private Address Space (RFC 1918) destination | If enabled, proxy is bypassed for Private Address Space (<i>RFC 1918</i>), and it is routed directly through the firewall. |
| Bypass proxy for these source IPs | Do not forward traffic from these source IPs, CIDR nets, hostnames, or aliases through the proxy server but directly through the firewall. Separate multiple entries by semi-colons. (This applies only to transparent mode) |
| Enable logging | This will enable the access log. Don't switch this on if you don't have much disk space left. Logging is required for proxy status reporting to work. <u>Note: Sentinel D2, D2W and other NanoBSD-based Sentinel models should keep this option turned off due to limited storage.</u> |
| Log store directory | The directory where the log file will be stored. The default is <i>/var/squid/log</i> . Do not change from default, otherwise proxy status reporting may not work. <u>Note: Do not end the directory entry with a slash.</u> |
| Log rotate | Defines how many days of logfiles will be kept. Rotation is disabled if left empty. |
| Proxy port | This is the port the proxy server will listen on. The default is 3128. |
| ICP port | This is the port the proxy server will send and receive ICP queries to and from neighbor caches. Leave this blank if you don't want the proxy server to communicate with neighbor caches through ICP. <u>Note: ICP is a Web caching protocol used to exchange hints about the existence of URLs in neighbor caches. Caches exchange ICP queries and replies to gather information to use in selecting the most appropriate location from which to retrieve an object.</u> |
| Visible hostname | This is the URL to be displayed in proxy server error messages. |
| Administrator email | This is the email address displayed in error messages to the users. |
| Language | Select the language in which the proxy server will display error messages to users. |
| Disable X-Forward | If not set, the proxy server will include your system's IP address or name in the HTTP requests it forwards. |
| Disable Via | If not set, the proxy server will include a <i>Via</i> header in requests and replies as required by <i>RFC 2616</i> . |

(Table continued)

What to do with requests that have whitespace characters in the URI

Defines behavior for whitespace URIs:

strip: The whitespace characters are stripped out of the URL. This is the behavior recommended by *RFC 2396*.

deny: The request is denied. The user receives an "Invalid Request" message. (Not recommended)

allow: The request is allowed and the URI is not changed. The whitespace characters remain in the URI.

encode: The request is allowed and the whitespace characters are encoded according to *RFC 1738*.

chop: The request is allowed and the URI is chopped at the first whitespace. (Not recommended)

Use alternative DNS servers for the proxy server

If you want to use other DNS servers than the Sentinel DNS Forwarder, enter DNS server IP addresses here, separated by semi-colons.

Suppress Squid Version

If set, suppress Squid version string info in HTTP headers and HTML error pages. We recommend to keep this enabled for security purposes.

Custom Options

You can put your own custom options here, separated by semi-colons. They will be added to the configuration. They need to be *squid.conf* compatible options, otherwise the proxy server will *not* work.

Upstream Proxy

Sentinel PF Proxy Server can direct all requests to another proxy server which is called *upstream proxy*. For example, BusinessCom PEP service can be used as an upstream proxy for the HTTP traffic. See *PEP* section of this chapter.

| Option | Description |
|-------------------|---|
| Enable forwarding | This option enables the proxy server to forward requests to an upstream server. |
| Hostname | Enter here the IP address or host name of the upstream proxy. |
| TCP port | Enter the port to use to connect to the upstream proxy. |
| ICP port | Enter the port to connect to the upstream proxy for the ICP protocol. Use port number 7 to disable ICP communication between the proxies. |
| Username | If the upstream proxy requires a username, specify it here. |
| Password | If the upstream proxy requires a password, specify it here. |

Cache Mgmt

This is the cache management section that defines how cache objects are stored and retrieved by the Proxy Server. Sentinel PF Proxy Server operates with the *cache directory* that can be sized according to the disk space available on your Sentinel server. The cache directory holds *objects* which represent cached content. With the *ufs* hard disk cache system, the cache directory is organized in two levels: *L1* and *L2*:

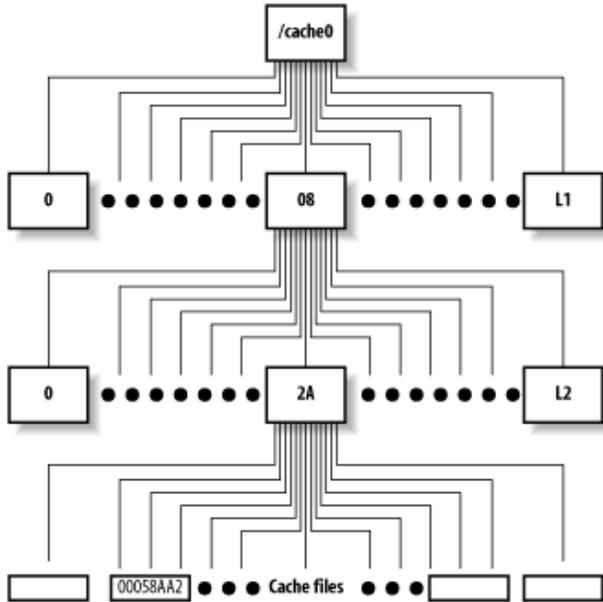


FIGURE 8-3 Sentinel PF Proxy Server *ufs* Cache Directory Structure.

The amount of L1 subdirectories is a number defined in advance. Each L1 subdirectory contains up to 256 L2 subdirectories. The default configuration is 256 L1 subdirectories, which results in $16 \times 256 = 4,096$ L2 subdirectories where cached objects can be stored.

It is possible to specify the amount of L1 subdirectories to regulate how cached objects will be organized on the Sentinel's hard drive. The optimal amount of L1 subdirectories will depend on the number of factors, such as the total amount of cache disk space allocated and average size of each cached object. The Proxy Server needs to be able to find cached content fast once it is requested by the client, and since it is easier to search smaller directories, the directory structure has to be created the way that only a few hundred files will end up stored at each L2 subdirectory.

This is an optional fine-tuning, and in most of the cases, due to hardware optimizations, all Sentinel servers will provide excellent performance with the default 16 L1 subdirectories setting.

| Option | Description |
|----------------------------|--|
| Hard disk cache size | <p>This is the amount of disk space (in megabytes) to use for cached objects.</p> <hr/> <p>Note: It is safe to use up to 70% of hard drive space available on your Sentinel server. Do not set this to very high amount to avoid filling up the entire disk.</p> |
| Hard disk cache system | <p>Specifies how cache is stored on disk:</p> <p><i>ufs</i> is the traditional Squid storage format and a default setting. Any disk I/O blocks the Squid process rendering it unable to process the connections. This is the slowest, however most safe mode.</p> <p><i>aufs</i> uses POSIX-threads to avoid blocking the main Squid process on disk I/O. Formerly known as <i>async-io</i>.</p> <p><i>diskd</i> uses a separate process to avoid blocking the main Squid process on disk I/O. This is similar to <i>aufs</i>, however slower in operation.</p> <p><i>null</i> Does not use any storage. Recommended for Sentinel systems with no disk storage, such as Sentinel D2.</p> |
| Hard disk cache location | <p>The directory where cache files will be stored. The default is <i>/var/squid/cache</i>. Do not change from default, otherwise proxy status reporting may not work.</p> <hr/> <p>Note: Do not end the directory entry with a slash.</p> |
| Memory cache size | <p>This is the amount of physical RAM (in Megabytes) to be used for negative cache and in-transit objects. This value should not exceed more than 50% of the installed RAM. The minimum value is 1 MB. The default value is 500 MB for Sentinel Blackbird, Cirrus and Sierra and 2000 MB for Westwind.</p> |
| Minimum object size | <p>Objects smaller than the size specified (in kilobytes) will not be saved on disk. The default value is 0, meaning there is no minimum.</p> |
| Maximum object size | <p>Objects larger than the size specified (in kilobytes) will not be saved on disk. The default value is 2000 kilobytes.</p> |
| Maximum object size in RAM | <p>Objects smaller than the size specified (in kilobytes) will be saved in RAM. The default value is 32.</p> |
| Level 1 subdirectories | <p>The number of first-level cache subdirectories which will be created under the cache directory. Default value is 16.</p> |
| Memory replacement policy | <p>The memory replacement policy determines which objects are purged from RAM when space is needed. The default policy for memory replacement is <i>Heap GDSF</i>. Available options are:</p> <p><i>LRU: Last Recently Used Policy</i> - The LRU policies keep recently referenced objects. i.e., it replaces the object that has not been accessed for the longest time.</p> |

(Table continued)

Heap GDSF: Greedy-Dual Size Frequency - The Heap GDSF policy optimizes object-hit rate by keeping smaller, popular objects in cache. It achieves a lower byte hit rate than LFUDA though, since it evicts larger (possibly popular) objects.

Heap LFUDA: Least Frequently Used with Dynamic Aging - The Heap LFUDA policy keeps popular objects in cache regardless of their size and thus optimizes byte hit rate at the expense of hit rate since one large, popular object will prevent many smaller, slightly less popular objects from being cached.

Heap LRU: Last Recently Used - Works like LRU, but uses a heap instead.

Note: If using the *LFUDA* replacement policy, the value of Maximum Object Size should be increased above its default of 12 KB to maximize the potential byte hit rate improvement of LFUDA.

| | |
|--------------------------|---|
| Cache replacement policy | The cache replacement policy decides which objects will remain in disk cache and which objects are replaced to create space for the new objects. The default policy for cache replacement is <i>Heap LFUDA</i> . Please see the type descriptions specified in the memory replacement policy for additional detail. |
| Low-water-mark in % | Cache replacement begins when the swap usage is above the <i>low-low-water mark</i> and attempts to maintain utilization near the <i>low-water-mark</i> . |
| High-water-mark in % | As swap utilisation gets close to the high-water-mark object eviction becomes more aggressive. |
| Do not cache | Enter each domain or IP address on a new line that should never be cached. |
| Enable offline mode | Enable this option and the proxy server will never try to validate cached objects. The offline mode gives access to more cached information than the proposed feature would allow (stale cached versions, where the origin server should have been contacted). |

Access Control

This tab specifies who can access the Proxy Server and for what destination domains and ports. If *Allow users on interface* is unchecked in the General settings tab, you will need to specify how users shall authenticate.

| Option | Description |
|-----------------|--|
| Allowed subnets | Enter each subnet on a new line that is allowed to use the proxy. The subnets must be expressed as CIDR ranges (e.g.: 192.168.1.0/24). |
| | Note: Proxy interface subnet is already an allowed subnet. All the other subnets won't be able to use the proxy. |

(Table continued)

| | |
|-------------------------|---|
| Unrestricted IPs | Enter each unrestricted IP address on a new line that is not to be filtered out by the other access control directives set in this page. |
| Banned host addresses | Enter each IP address on a new line that is not to be allowed to use the proxy. |
| Whitelist | Enter each destination domain on a new line that will be accessible to the users that are allowed to use the proxy. You also can use regular expressions. |
| Blacklist | Enter each destination domain on a new line that will be blocked to the users that are allowed to use the proxy. You also can use regular expressions. |
| External Cache-Managers | Enter the IPs for the external Cache Managers to be allowed here, separated by semi-colons. |
| acl safeports | This is a space-separated list of <i>safe ports</i> in addition to the already defined list: <i>21 70 80 210 280 443 488 563 591 631 777 901 1025-65535</i> |
| acl sslports | This is a space-separated list of ports to allow SSL "CONNECT" in addition to the already defined list: <i>443 563</i> |

Note: Authentication cannot be enabled while transparent proxy mode is enabled.



Caution: After changing the amount of Level 1 subdirectories, it is necessary to rebuild the cache directory. To rebuild directory, stop the Proxy Server Service in the Sentinel PF Dashboard by clicking on the  button. Then, issue the following commands in the CLI (see *CLI* section of *Chapter 3 – Command Line Interface*):

```
#rm -rf /var/squid/cache
#squid -z
2013/03/18 21:33:19| Creating Swap Directories
#exit
```

After this, restart the Proxy Server Service by clicking on the  button. After about 30 seconds, make sure the Proxy Server is running – the  sign should appear next to it.

Traffic Mgmt

This is the traffic management section that defines some advanced bandwidth management policies for traffic passing through Proxy Server, mostly based on the *throttling* technique. This can be an effective bandwidth management approach for controlling HTTP based traffic flows, especially when coupled with *transparent proxy* setting. For example, you can throttle multimedia files, so large file downloads will make a slow start with gradual throughput increase thereafter, allowing some “elbow room” for critical Web applications.

| Option | Description |
|-------------------------------------|---|
| Maximum download size | Limit the maximum total download size to the size specified here (in kilobytes). Set to 0 to disable. |
| Maximum upload size | Limit the maximum total upload size to the size specified here (in kilobytes). Set to 0 to disable. |
| Overall bandwidth throttling | This value specifies (in <i>kilobytes</i> per second) the bandwidth throttle for downloads. Users will gradually have their download speed increased according to this value. Set to 0 to disable bandwidth throttling. |
| Per-host throttling | This value specifies the download throttling per host. Set to 0 to disable this. |
| Throttle only specific extensions | Leave this checked to be able to choose the extensions that throttling will be applied to. Otherwise, all files will be throttled. |
| Throttle binary files | Check this to apply bandwidth throttle to binary files. This includes compressed archives and executables. |
| Throttle CD images | Check this to apply bandwidth throttle to CD image files. |
| Throttle multimedia files | Check this to apply bandwidth throttle to multimedia files, such as movies or songs. <u>Note: This has no effect on streaming services like YouTube.</u> |
| Throttle CD images | Check this to apply bandwidth throttle to CD image files. |
| Throttle other extensions | Comma-separated list of extensions to apply bandwidth throttle to. |
| Finish transfer if less than x KB | If the transfer has less than x KB remaining, it will finish the retrieval. Set to 0 to abort the transfer immediately. |
| Abort transfer if more than x KB | If the transfer has more than x KB remaining, it will abort the retrieval. Set to 0 to abort the transfer immediately. |
| Finish transfer if more than x % | If more than x % of the transfer has completed, it will finish the retrieval. |

Auth Settings

This is the authentication section that defines if and how Proxy Server clients shall authenticate before being allowed to pass traffic through it. If authentication is enabled then clients will be presented with *Authentication prompt* to log in with their credentials. Sentinel PF Proxy Server can work with LDAP, RADIUS and NT domain authentication servers, or it can use an internal user database.

| Option | Description |
|--|---|
| Authentication method | Select an authentication method. This will allow users to be authenticated by local or external services. |
| LDAP version | Enter LDAP protocol version (2 or 3). |
| Authentication server & port | Enter here the IP or hostname of the server that will perform the authentication. In the port section, enter the port to use to connect to the authentication server. Leave this field blank to use the authentication method's default port. |
| NT domain | Enter NT domain. |
| LDAP server user DN, password, base domain, username DN attribute, search filter | LDAP server authentication settings. |
| RADIUS secret | The RADIUS secret for RADIUS authentication. |
| Secondary NT servers | Comma-separated list of secondary servers to be used for NT domain authentication. |
| Authentication prompt | This string will be displayed at the top of the authentication request window. |
| Authentication processes | The number of authenticator processes to spawn. If many authentications are expected within a short timeframe, increase this number accordingly. |
| Authentication TTL | his specifies for how long (in minutes) the proxy server assumes an externally validated username and password combination is valid (Time To Live). When the TTL expires, the user will be prompted for credentials again. |
| Require authentication for unrestricted hosts | If this option is enabled, even users tagged as <i>unrestricted</i> through <i>Access Control</i> are required to authenticate to use the proxy. |
| Subnets that don't need authentication | Enter each subnet or IP address on a new line (in CIDR format, e.g.: 10.5.0.0/16 , 192.168.1.50/32) that should not be asked for authentication to access the proxy. |

Local Users

This is the Sentinel PF Proxy Server local user database that will be used for authentication purposes, if the latter is enabled.

▼ For PEP Acceleration with HTTP Proxy Cache

- 1 In Services / PEP, configure PEP client daemon.**
Use *HTTP* or *Compressed HTTP* as PEP service type. Make sure PEP client daemon is running. Please refer to the *PEP* section of this Chapter for additional information on the PEP client configuration.
- 2 In Services / Proxy Server, configure Proxy server.**
- 3 Enable «Transparent Proxy» in «General» tab.**
Save settings at this stage.
- 4 In «Upstream Proxy» tab, configure upstream proxy.**
Use *PEP Local Address and Port* as the upstream proxy. Use ICP port 7. Username and password is not required, as HTTP PEP service does not support authentication.
- 5 Check if your HTTP sessions are routed via PEP.**
A simple way would be to point your web browser to *http://www.whatismyip.com* and check the IP detected. A valid response with connection initiated via PEP gateway would be as following:

Your IP Address is:
PEP Gateway IP Address

Proxy Detected:
PEP Gateway IP Address

```
1.1 localhost:3128 (squid/2.7.STABLE9), 1.0 gw-falkenstein:12345 (squid/2.7.STABLE9)
```

(*gw-falkenstein* is the name of BusinessCom PEP gateway in Germany used in illustrative purposes only, your actual PEP gateway may be different.)

Note: If you are using multiple Sentinel servers configured with redundant CARP then each server will require a unique PEP account for gateway in order to avoid conflicts with concurrent sessions.

Proxy Server (TP) ◆

Sentinel PF features *Tinyproxy* which is an external package that provides non-caching proxy functionality for Web traffic. It appears as *Proxy Server (TP)* in the Services menu. Tinyproxy is a light-weight HTTP/HTTPS proxy daemon. Designed from the ground up to be fast and yet small, it is an ideal solution for use cases such as embedded deployments such as Sentinel D2 and D2W, where a full featured HTTP proxy is required, but the system resources for a larger proxy are unavailable.

The proxy server can be configured with a number of options:

| Option | Description |
|---|--|
| Enable proxy server | Enables the proxy server. |
| Proxy Interface | The interface(s) the proxy server will bind to. This is usually set to <i>LAN</i> interface if you want to provide proxy services to your local network. |
| Transparent mode | If the following option is enabled, the proxy server will work in transparent mode. |
| Use another «via» proxy name | If the following option is enabled, the string supplied will be used as the host name in the Via header. |
| «via» proxy name | Specifies the “via” string. |
| Listen port | Specify the port which proxy will listen on, default is 8888. |
| Timeout | The maximum time (in seconds) a connection is allowed to stay inactive before it is closed by proxy, default is 600. |
| Maximum clients | Number of clients that can connect at the same time. |
| Use upstream proxy | Turns on upstream proxy support. If enabled, all the traffic passing through the proxy server will be routed via the upstream proxy specified. |
| Upstream proxy socket | IP address and port for upstream proxy server, separated by colon. Example: 127.0.0.1:3000 to forward traffic to the PEP service. |
| Bypass proxy for these destination URIs/IPs | Do not proxy traffic going to the specified host names, networks, or IPs. Separated by semicolons. |
| Minimum spare servers | Tinyproxy always keeps a certain number of idle child processes so that it can handle new incoming client requests quickly. Minimum and maximum spare servers control the lower and upper limits for the number of spare processes. I.e. when the number of spare servers drops below the minimum value then Tinyproxy will start forking new spare processes in the background and when the number of spare processes exceeds maximum then Tinyproxy will kill off extra processes. These settings set the lower limit for the number of spare servers which should be available. Default is 5. |

(Table continued)

| Option | Description |
|----------------------------|---|
| Maximum spare servers | These settings set the upper limit for the number of spare servers which should be available. Default is 20. |
| Start servers | The number of servers to start initially, default is 10. This should usually be set to a value between minimum and maximum spare servers. |
| Maximum requests per child | This limits the number of connections that a child process will handle before it is killed. The default value is 0 which disables this feature. This option is meant as an emergency measure in the case of problems with memory leakage. In that case, setting this to a value of e.g. 1000, or 10000 can be useful. |
| Enable URLs filter | Turns on URL filter. |
| Default filter policy | Change the policy for the URLs list. The default filtering policy is to allow everything that is not matched by a filtering rule. Setting policy to <i>Deny</i> changes the policy to deny everything but the domains or URLs matched by the filtering rules. |
| URL list | Put URLs separated by semicolon here for filtering purposes. |
| Enable log | Enable logging. <u>Note: Sentinel D2, D2W and other NanoBSD-based Sentinel models should keep this option turned off due to limited storage.</u> |
| Log level | Sets the log level. Messages from the set level and above are logged. For example, if the Log level was set to <i>Warning</i> , then all log messages from <i>Warning</i> to <i>Critical</i> would be output, but <i>Notice</i> and below would be suppressed. Allowed values are: <ul style="list-style-type: none"> • Critical (least verbose) • Error • Warning • Notice • Connect (log connections without Info's noise) • Info (most verbose) |

Note: In some very high load environments with PEP acceleration running on embedded Sentinel servers like D2 and D2W, even Tinyproxy is too demanding of the RAM resources (approx. 3 MB per client). In these situations, we recommend using transparent non-caching redirect technique described in *For Transparent PEP HTTP Redirect how-to* of the *PEP* section in this chapter.

RIP

Sentinel PF features RIP (Routing Information Protocol) daemon to manage the network routing tables. It uses Routing Information Protocol, RIPv1 (*RFC 1058*), RIPv2 (*RFC 1723*), and Internet Router Discovery Protocol (*RFC 1256*) to maintain the Sentinel PF kernel routing table. The RIPv1 protocol is based on the reference 4.3BSD daemon.

When RIP service is enabled, Sentinel PF RIP daemon listens on the udp socket for for Routing Information Protocol packets. It also sends and receives multicast Router Discovery ICMP messages. If the host is a router, routed periodically supplies copies of its routing tables to any directly connected hosts and networks. It also advertises or solicits default routes using *Router Discovery* ICMP messages.

Sentinel PF RIP daemon can be configured with a number of options:

| Option | Description |
|----------------|---|
| Enable RIP | Enables the Routing Information Protocol daemon. |
| Interfaces | Select the interface(s) that RIP will bind to. |
| RIP Version | Select which RIP version the daemon will listen/advertise using. |
| RIPv2 Password | Specify a RIPv2 password. This password will be sent in the clear on all RIPv2 responses received and sent. |

SNMP

Sentinel PF features SNMP (Simple Network Management Protocol) daemon that allows to monitor some Sentinel PF parameters remotely, such as network traffic, network flows, packet filter queues and general system information such as CPU, memory and disk usage. To list available MIBs (Management Information Bases), you can use the *snmpwalk* command from a host that has *Net-SNMP* or equivalent software installed, e.g.:

```
$ snmpwalk -c public 192.168.1.1
SNMPv2-MIB::sysDescr.0 = STRING: sentinel.localdomain 974156609 FreeBSD
8.1-RELEASE-p13
...
(list of MIBs truncated)
```

A MIB database is organized as a tree. The upper structure is defined in *RFC 1155* and *RFC 1213*. The internal nodes of the tree are organized by a particular function. MIB variables are stored as the leaves of this tree. The full list of MIBs available will not enter into the discussion of this book, however you can refer to SNMP related RFCs, such as *RFC 2790 - Host Resources MIB* for additional information.

In addition, Sentinel PF SNMP daemon provides loadable modules:

TABLE 8-1 Sentinel PF SNMP Daemon Loadable Modules

| Module | Description |
|----------------|---|
| MibII | <p>One of the standard MIBs defined in <i>RFC 1213</i>. It has the following management groups:</p> <p><i>system</i> Objects that pertain to system operation, such as uptime.</p> <p><i>interfaces</i> Network interface related objects: status, traffic, errors, etc.</p> <p><i>at</i> Address translation group (deprecated).</p> <p><i>ip, icmp, udp, edp</i> IP routing; ICMP errors, discards, UDP and EGP statistics, EGP neighbor tables.</p> <p><i>transmission</i> and <i>snmp</i> Media-specific MIBs, SNMP implementation performance.</p> |
| Netgraph | <p>Access to <i>graph</i> based kernel networking subsystem of FreeBSD kernel that Sentinel PF is built upon. For additional information please refer to: http://www.freebsd.org/cgi/man.cgi?query=netgraph&sektion=4</p> |
| PF | <p>The PF MIB allows for the querying of objects relating to Sentinel PF Packet Filter. This includes general information, packet, state table, log interface and source track counters; memory limits, protocol timeouts, interface statistics (rules, states, bytes, packets), table statistics, filter rule label counters and others.</p> |
| Host Resources | <p>The Host Resources MIB defines a uniform set of objects useful for the management of host computers. Host computers are independent of the operating system, network services, or any software application. The Host Resources MIB defines objects which are common across many computer system architectures. Please refer to <i>RFC 2790 - Host Resources MIB</i> for additional information.</p> |

Sentinel PF SNMP daemon can send SNMP *traps* which are basically notification messages that can be received by an SNMP trap receiver, such as *Net-SNMP's snmptrapd*, or 3rd party commercial software like Paessler PRTG, What's Up Gold, Solar Winds SNMP Trap Receiver, and others.

▼ To Enable SNMP Daemon

- 1 **In Services / SNMP, click on «Enable» checkbox next to «SNMP Daemon» caption.**
- 2 **Enter optional system location and contact.**
- 3 **Enter optional Community String.**
The community string is like a password, restricting access to querying SNMP to hosts knowing the community string. Use a strong value here to protect from unauthorized information disclosure.
- 4 **Optionally, enable SNMP Trap sending and configure SNMP Trap server (receiver), port and string.**
SNMP trap string will be sent to the trap receiver along with any trap information generated.
- 5 **Specify virtual server pool incoming traffic to be directed to.**
- 6 **Optionally, enable SNMP modules.**
- 7 **Click «Save».**

Snort (NIDS) ◆

Snort® is an external open source package developed by Sourcefire, Inc. available for Sentinel PF. It provides NIDS (Network Intrusion Detection System) and IPS (Intrusion Prevention System) functionality and enables real-time traffic analysis of IP networks for malicious activity.

Snort can be used to detect a variety of attacks and probes, such as buffer overflows, stealth port scans, CGI attacks, SMB probes, OS fingerprinting attempts, and much more. It is a de facto standard for NIDS and IPS worldwide. The basic threat prevention components of Snort are packet classifier, IP defragmenter, TCP reassembler, portscan processor and detection engine. Snort inspects incoming and outgoing traffic and matches traffic against a known rule set, and reports (NIDS) and/or takes preventive action (IPS) for malicious activity in your network.

Sentinel PF comes with Snort activated in the NIDS mode (*without* IPS) on the LAN interface. This allows network administrators to react on network incidents, however this is rather conservative. Snort can be enabled on the WAN interface as an IPS system in order to prevent malware to enter the local network in the first place, however setting up an effective IPS requires an in-depth knowledge of the ruleset, your networking environment and applications in use.

We recommend operating Snort in the NIDS mode to refine the ruleset and minimize the amount of false alarms first before switching to the IPS mode. This avoids blocking hosts for legitimate traffic. An example of a false alarm could be access to your SNMP daemon from a public IP host which may be classified as an information leak in some network and perfectly normal activity in others.



Caution: Without manual NIDS ruleset refinement, the default configuration for the IPS mode may be overly aggressive for many common environments. This may result in legitimate traffic and hosts being blocked.

The *Snort Interfaces* tab lists network interfaces Snort daemon is binded to. There can be multiple instances of the daemon running for different purposes and interfaces. Use this tab to add, edit and delete interfaces from/to the Snort system.

The *General Settings* tab allows to configure Snort settings that are applicable to all instances of the daemon, on all interfaces:

| Option | Description |
|-------------------------------|--|
| Install Snort.org rules | Allows to install NIDS ruleset provided by Snort.org community or Sourcefire VRT® (Vulnerability Research Team). This requires a subscription with these resources, defined by the <i>Oinkmaster code</i> . In case you would like to use these rulesets, select <i>Install Basic Rules or Premium Rules</i> and enter your code. Snort.org subscription is free of charge, however requires registration at http://www.snort.org/snort-rules/ . |
| Install Emergingthreats rules | <i>Emerging Threats</i> is an open source community that produces fastest moving and diverse Snort Rules. The project is funded by National Science Foundation and U.S. Army Research Office, and the ruleset is available free of charge. This is the default NIDS ruleset used on Sentinel PF systems. |
| Update rules automatically | Sentinel PF will automatically update NIDS rules from the source(s) selected. This defines the update interval. |
| Log Directory Size Limit | Maximum size of the Snort log directory. Default is 20% of the disk space available on your Sentinel system. |
| | <u>Note: The log directory is used to store log files, such as NIDS alerts. It can quickly grow very large, if limits are not set.</u> |

(Table continued)

| | |
|--|---|
| Remove blocked hosts every | Please select the amount of time you would like hosts to be blocked for. This settings refers to IPS (Intrusion Prevention System.) |
| Keep <i>snort</i> settings after deinstall | Prevents your NIDS settings to be removed in case you decide to deinstall the Snort package. |

The *Updates* tab allows to manually update NIDS rules and view update logs. Currently installed NIDS rulesets, also known as signature rulesets, are shown in the *Installed Signature Ruleset* section. These updates and installed rulesets apply to all Snort instances.

The *Alerts* tab allows to view latest NIDS alerts. You can select Snort instance to inspect (e.g. default NIDS on LAN interface or others, if you have configured additional instances). You can also save log files for further inspection with 3rd party tools, or remove (clear) logs.

The *NIDS Alert* table has the following fields:

TABLE 8-2 Sentinel PF NIDS Alerts Table Fields

| Column | Description |
|-------------|--|
| Date | Time and date the alert has been raised. |
| PRI | Alert priority. Priority of 1 (High) is the most severe incident, and priority 3 (Low) is the least severe. Priority 4 (Very Low) most probably does not indicate a security incident. |
| Class | Incident classification. (See <i>Appendix F – NIDS Incidents Classification</i>). |
| SRC | Source of the traffic that resulted in alert. |
| SRCPORT | Port at the source. |
| DST | Destination. |
| DSTPORT | Port at the destination. |
| SID | Snort Rule ID for the alert. |
| DESCRIPTION | Alert type and description. |

Note: You can suppress alerts for legitimate traffic by adding a particular Snort Rule ID (SID) into the *Suppress* list by clicking on the  button under the SID. NIDS Alerts with this SID will not be displayed in the Alerts table anymore.

The *Blocked* tab allows to view latest hosts blocked by the IPS. You can download list of hosts by pressing the *Download* button. You can also remove (clear) all or individual hosts from the table.

The *Whitelists* tab allows to create *whitelist* files for your Snort package rules that would be excluded from IPS block decisions. Please add all the IP addresses or networks you want to exclude from the IPS. Remember that the default whitelist only includes local networks.

The *Suppress* tab allows to control what NIDS alerts are suppressed. Sentinel PF ships with default *lansuppress* list. If a NIDS rule is suppressed from the NIDS Alert table, it will be added to this list. Network incidents matching suppress rules will not be reported. This function is used to prevent NIDS alerting on known legitimate traffic.

▼ To Create New NIDS Suppress List

- 1 In Services / Snort, select the «Suppress» tab.
- 2 Click on the  button to create a new suppress list.
- 3 Enter the suppress list name, e.g. `mysuppress`.
- 4 Enter optional list description, e.g. `NIDS Suppress List For Known Web Traffic`.
- 5 Add or remove individual entries in the «Advanced Pass Through» section.
An example of the NIDS suppress entry is provided below. It starts with the # symbol, followed by *gen_id*, *sig_id* and optionally other Snort parameters.

```
#(http_inspect) SIMPLE REQUEST  
suppress gen_id 119, sig_id 32
```
- 6 Click «Save».

▼ To Edit NIDS Suppress List

- 1 In Services / Snort, select the «Suppress» tab.
- 2 Click on the  button next to the suppress list you want to edit.
- 3 Add or remove individual entries in the «Advanced Pass Through» section.
- 4 Click «Save».

Advanced Pass Through NIDS entries follow standard Snort event filtering syntax described in the table below. There are three event processing filters available to suit various tasks. *Rate filters* (*rate_filter*) are used to change a rule action when the number or rate of events indicates a possible attack. *Event filters* (*event_filter*) are used to reduce the number of logged events for noisy rules. This can be tuned to significantly reduce false alarms. *Event suppression* (*suppress*) is used to completely suppress the logging of uninteresting events.

An example of valid Advanced Pass Through entries is provided below:

```
suppress gen_id 1, sig_id 1852, track by_src, ip 10.1.1.54
event_filter gen_id 1, sig_id 1851, type limit, track by_src,
count 1, seconds 60
rate_filter gen_id 135, sig_id 1, track by_src, count 100,
seconds 1, new_action log, timeout 10
```

Additional information on Snort syntax is available in *Snort User's Manual*, *Event Processing* section.

In addition to settings discussed above, Sentinel allows to configure every Snort instance individually. The instance configuration is available from the *Snort Interfaces* tab, by pressing the  button next to the subject instance. The following configuration tabs are available:

The *IF Settings* tab allows to configure the following Snort interface settings:

| Option | Description |
|---------------------------------|--|
| Enable | Enable Snort instance on this interface. |
| Interface | Select network interface for this instance. |
| Description | Optional Snort instance description for your reference. |
| Send alerts to main System logs | If enabled, Snort will send Alerts to the firewall's system logs. |
| Block offenders | Checking this option will automatically block hosts that generate a Snort alert. This option puts this instance into a “dumb” IPS mode, whereby offenders are blocked by NIDS alerts, even if there is no matching IPS rule. |
| Kill states | Checking this option will kill firewall states for the blocked IP. |
| Which IP to block | Selects the IP to block. Can be <i>src</i> for source, <i>dst</i> for destination or <i>both</i> for both. |
| Memory Performance | <p>Specifies the performance model for Snort. There are different queued match traffic pattern search methods available. Traffic matches are queued until the fast pattern matcher is finished with the payload, then evaluated.</p> <p><i>ac</i> - Aho-Corasick Full (high memory, best performance) <i>ac-bnfa</i> - Binary NFA (low memory, high performance) <i>lowmem</i> Keyword Trie (low memory, moderate performance) <i>ac-split</i> - Aho-Corasick Full (low memory, high performance).</p> <p>No queue search methods - The <i>nq</i> option specifies that matches should not be queued and evaluated as they are found.</p> <p>Other search methods (the above are considered superior to these)</p> <p><i>ac-std</i> - Aho-Corasick Standard (high memory, high performance) <i>acs</i> - Aho-Corasick Sparse (high memory, moderate performance) <i>ac-banded</i> - Aho-Corasick Banded (high memory, moderate performance) <i>ac-sparsebands</i> - Aho-Corasick Sparse-Banded (high memory, moderate performance)</p> |

(Table continued)

| | |
|---|--|
| <hr/> <p>Note: <i>lowmem</i> and <i>ac-bnfa</i> are recommended for low end systems. Sentinel systems use <i>ac</i> by default.</p> <hr/> | |
| Checksum Check Disable | If ticked, checksum checking on Snort will be disabled to improve performance. Most of this is already done at the firewall/filter level. |
| Home net, External net and Whitelist | Define home net (local network to protect), external net (outside world) and whitelist for this Snort instance. These are usually set to <i>default</i> in most Sentinel environments. Whitelist option is used only when <i>Block offenders</i> is turned on. |
| Suppression and filtering | Choose suppression or filtering list to use for this instance. <i>Default</i> option disables suppression and filtering. |
| Advanced configuration pass through | Additional Snort configuration as per standard Snort configuration file syntax. These arguments will be automatically inserted into the Snort configuration. |

The *Categories* tab allows to configure which NIDS and IPS rules are used for the instance.

| Option | Description |
|--------------------|---|
| Resolve Flowbits | If ticked, Snort will examine the enabled rules in your chosen rule categories for checked <i>flowbits</i> . Any rules that set these dependent flowbits will be automatically enabled and added to the list of files in the interface rules directory. Flowbits are used to communicate protocol state information between different NIDS rules. For example, one NIDS rule may detect the login stage over a clear text protocol and trigger a flowbit for other NIDS rules to process further traffic in a different manner. With the default Emerging Threats ruleset, the flowbit usage is limited and, thus, disabled by default, in order to preserve CPU resources. |
| Auto Flowbit Rules | Click to view auto-enabled rules required to satisfy flowbit dependencies from the selected rule categories below. Auto-enabled rules generating unwanted alerts should have their <i>GID:SID</i> added to the Suppression List for the interface. |
| Use IPS Policy | If ticked, Snort will use rules from the pre-defined IPS policy selected below. You must be using the Snort VRT rules to use this option. Selecting this option disables manual selection of Snort VRT categories in the list below, although Emerging Threats categories may still be selected if enabled on the Global Settings tab. These will be added to the pre-defined Snort IPS policy rules from the Snort VRT. |

(Table continued)

| | |
|------------|--|
| IPS Policy | <p>Available Snort IPS policies are: <i>Connectivity</i>, <i>Balanced</i> or <i>Security</i>.</p> <p><i>Connectivity</i> blocks most major threats with few or no false positives.</p> <p><i>Balanced</i> is a good starter policy. It is speedy, has good base coverage level, and covers most threats of the day. It includes all rules in <i>Connectivity</i>.</p> <p><i>Security</i> is a stringent policy. It contains everything in the first two plus policy-type rules such as Flash in an Excel file.</p> |
|------------|--|

The *Rules* tab allows to view individual NIDS rules, as well as enable, disable individual rules in different categories by clicking on the first column buttons next to the rules. The *Rules* table has the following fields:

TABLE 8-3 Sentinel PF NIDS Rules Table Fields

| Column | Description |
|-------------|---|
| SID | Snort Rule ID. |
| Proto | Matching IP protocol. |
| Source | Matching source. E.g. <i>\$EXTERNAL_NET</i> means the rule will match any packets coming from external net (outside world). <i>\$HOME_NET</i> matches packets coming from home net (local network). |
| Port | Matching port number. <i>Any</i> means any port. |
| Destination | Matching destination. |
| Message | NIDS Alert Message. |

Each rule can have multiple states:

TABLE 8-4 Sentinel PF NIDS Rule States

| State | Meaning |
|---|-------------------------------------|
|  | Rule is enabled by default. |
|  | Rule is disabled by default. |
|  | Rule is enabled (changed by user). |
|  | Rule is disabled (changed by user). |

All user changes can be cleared by pressing *Remove Enable/Disable* buttons for the current category or all categories.

The *Variables* tab allows to configure Snort variables that will be used to help NIDS with incident detection. For example, Emerging Threats SQL NIDS rules (*emerging-sql.rules* category) are using the *\$SQL_SERVERS* variable to detect attacks on local SQL servers. By default, Sentinel PF assumes that all your servers are located in the local network (defined by the *\$HOME_NET*). Similarly, variables can be used to define ports.

You may want to check all the available variables and specify values for servers and ports that do not match the default *\$HOME_NET* assumption or standard ports commonly assigned to the applications in use.

Server IP addresses can be specified using the CIDR notation, e.g.:

64.12.24.0/23, 64.12.28.0/23

Ports can be specified separated by comma, e.g.:

6665, 6666, 6667

All variables are named accordingly to protocols or applications, e.g. *DNS_SERVERS* are DNS servers, *FTP_SERVERS* are FTP servers, and so on. In most of the network environments, these variables should remain empty (return *\$HOME_NET*), as set by default.

The *Preprocessors* tab allows to configure Snort preprocessors. Preprocessors help in identifying possible attack packets before NIDS rules are applied, after the preprocessing stage various rules are applied on the packets (raw data) for detecting attacks based on the pattern matches. A list of preprocessors available on the Sentinel PF system is provided in the table below.

TABLE 8-5 Sentinel PF Snort Preprocessors

| Preprocessor | Description |
|--------------|---|
| Flow | This preprocessor helps keep a state flow log of packets passing through the Snort engine. |
| Stream5 | This preprocessor reassembles TCP packets and inspects them to detect attempted IDS evasion attacks. This preprocessor also detects port scans, state problems with a session, and records session information. |
| HTTP Inspect | This preprocessor handles all HTTP traffic to help speed it through to the rules engine. This preprocessor serves several purposes such as: HTTP traffic normalization, HTTP traffic profiling and normalization, possibly for each web server in your organization; and the ability to detect proxy usage. |
| RPC Decode | Listens for RPC protocol packets on certain ports, and then decodes the traffic on those ports to ASCII to be passed back to the Snort rules engine for comparison. |

(Table continued)

| | |
|------------------------|--|
| Normalizers | There is a variety of normalizers available for different protocols that decode traffic and pass it via NIDS system. Normalizers available are for FTP, telnet, POP, IMAP, SMTP, DCE/RPC2, SIP, GTP (GPRS Tunneling Protocol), DNS and SSL data protocols. |
| Performance Statistics | This preprocessor generates statistical information on the load Snort is under, sensor load, and several network performance measurements. |
| Portscan | This preprocessor detects portscans. <hr/> Note: Portscan preprocessor can be exceedingly sensitive. <hr/> |
| SCADA | This preprocessor decodes Modbus and DNP3 protocols that are used in SCADA networks and passes it through the NIDS system. If your network does not contain DNP3 or Modbus enabled devices, you may want to keep this preprocessor disabled. |



Caution: NIDS rules may be dependent on preprocessors! Disabling preprocessors may result in dependent rules being automatically disabled.

The HTTP Inspect preprocessor can be configured with the following options:

| Option | Description |
|------------------------|--|
| HTTP Server Flow Depth | -1 to 65535 (-1 disables HTTP inspect, 0 enables <i>all</i> HTTP inspect) Amount of HTTP server response payload to inspect. Snort's performance may increase by adjusting this value. Setting this value too low may cause false negatives. Values above 0 are specified in bytes. Recommended setting is maximum (65535). Default value is 300. |
| HTTP Server Profile | Choose the profile type of the protected web server. The default is <i>All</i> . <i>IIS_4.0</i> and <i>IIS_5.0</i> are identical to <i>IIS</i> except they alert on the double decoding vulnerability present in those two versions. |
| HTTP Client Flow Depth | -1 to 1460 (-1 disables HTTP inspect, 0 enables <i>all</i> HTTP inspect) Amount of raw HTTP client request payload to inspect. Snort's performance may increase by adjusting this value. Setting this value too low may cause false negatives. Values above 0 are specified in bytes. Recommended setting is maximum (1460). Default value is 300. |
| Disable HTTP Alerts | Tick to turn off alerts from the HTTP Inspect preprocessor. This has no effect on HTTP rules in the rule set. |

The Stream5 preprocessor can be configured with the following options:

| Option | Description |
|------------------|---|
| Max Queued Bytes | Minimum is 1024, Maximum is 1073741824 (default value is 1048576, 0 means Maximum). The number of bytes to be queued for reassembly for TCP sessions in memory. Default value is 1048576. |
| Max Queued Segs | Minimum is 1024, Maximum is 1073741824 (default value is 1048576, 0 means Maximum). The number of segments to be queued for reassembly for TCP sessions in memory. Default value is 2621. |
| Memory Cap | Minimum is 32768, Maximum is 1073741824 (default value is 8388608). The memory cap in bytes for TCP packet storage in RAM. Default value is 8388608 (8 MB). |

The Portscan preprocessor can be configured with the following options:

| Option | Description |
|-----------------|--|
| Sensitivity | <p><i>Low</i>: alerts generated on error packets from the target host; this setting should see few false positives.</p> <p><i>Medium</i>: tracks connection counts, so will generate filtered alerts; may false positive on active hosts.</p> <p><i>High</i>: tracks hosts using a time window; will catch some slow scans, but is very sensitive to active hosts.</p> |
| Ignore Scanners | Ignores the specified entity as a source of scan alerts. Entity must be a defined alias. Default value: <i>\$HOME_NET</i> . Leave blank for default value. |

The *Barnyard2* tab allows to configure the Barnyard2 output system for Snort. With the Barnyard2 enabled, Snort creates a special binary output format called *unified2*. The Barnyard2 system reads this file, and then resends the data to a database backend. It decouples output overhead from NIDS and allows it to run at full speed. NIDS alerts and logs can be logged to a MySQL database, specified in the *MySQL Settings* section.

Example MySQL Settings input:

```
output database: alert, mysql, dbname=snort user=snort
host=10.0.0.138 password=xyz
```

```
output database: log, mysql, dbname=snort user=snort
host=10.0.0.138 password=xyz
```

UPnP & NAT-PMP

Sentinel PF supports UPnP (Universal Plug & Play) and NAT-PMP (NAT Port Mapping Protocol) to allow software and devices for auto-configuration. These protocols can be configured with the following options:

| Option | Description |
|---|---|
| Enable UPnP & NAT-PMP | Enable UPnP & NAT-PMP daemon |
| Allow UPnP Port Mapping | Allows UPnP Port Mapping. This protocol is often used by Microsoft-compatible systems. |
| Allow NAT-PMP Port Mapping | Allows NAT-PMP Port Mapping. This protocol is often used by Apple-compatible systems. |
| Interfaces | Select network interface(s) for the daemon to bind to. |
| Maximum Download Speed | Maximum download speed per each port opened by UPnP, specified in kbps. |
| Maximum Upload Speed | Maximum upload speed per each port opened by UPnP, specified in kbps. |
| Override WAN address | By default, the UPnP service will configure port forwards and firewall rules to the WAN address. This setting will let you enter an alternate IP address, such as a secondary WAN address or a shared CARP address. |
| Traffic Shaping Queue | By default, rules created by UPnP will not assign traffic into a shaper queue. By entering the name of a queue into this field, traffic that passes due to a UPnP-created rule will fall into this queue. |
| Log packets handled by UPnP & NAT-PMP rules? | If enabled, the daemon will log all packets handled by UPnP & NAT-PMP rules in the Firewall logs. |
| Use system uptime instead of UPnP & NAT-PMP service uptime? | If enabled, the daemon will report Sentinel PF system uptime instead of the UPnP & NAT-PMP service (daemon) uptime. |
| By default deny access to UPnP & NAT-PMP? | If enabled, access to the daemon will be denied by default. |

You can also specify permissions in the *User specified permissions* sections using the following format: `[allow or deny] [ext port or range] [int ipaddr or ipaddr/cidr] [int port or range]`. Here is an example of a valid user specified permissions entry:

```
allow 1024-65535 192.168.0.0/24 1024-65535
```

Wake on LAN

The Wake on LAN service can be used to wake up (power on) computers by sending special "Magic Packets". The NIC in the computer that is to be woken up must support Wake on LAN and has to be configured properly (WOL cable, BIOS settings).

To wake up a single machine, select the Sentinel network interface it is connected to and enter the machine's MAC address in the `xx:xx:xx:xx:xx:xx` format, and click on the "Send" button to send the wake up packet.

You can store multiple MAC entries in the Wake on LAN table. Add clients by clicking on the  button and specifying the network interface, MAC address and optional description.

You can wake up individual clients by clicking on their MAC addresses in the Wake on LAN table, or you can wake all clients at once by clicking on the  button.

Web GUI: VPN

VPN menu items are used to configure Sentinel PF VPN services: IPSec, L2TP, OpenVPN and PPTP.

IPSec

IPSec is a protocol which sits on top of the Internet Protocol (IP) layer. It allows two or more hosts to communicate in a secure manner (hence the name). The Sentinel PF IPSec “network stack” is based on the KAME implementation, which has support for both protocol families, IPv4 and IPv6.

IPsec consists of two sub-protocols:

Encapsulated Security Payload (ESP), protects the IP packet data from third party interference, by encrypting the contents using symmetric cryptography algorithms (like Blowfish, 3DES).

Authentication Header (AH), protects the IP packet header from third party interference and spoofing, by computing a cryptographic checksum and hashing the IP packet header fields with a secure hashing function. This is then followed by an additional header that contains the hash, to allow the information in the packet to be authenticated.

ESP and AH can either be used together or separately, depending on the environment. To create an IPSec VPN tunnel, click on the  button in the *Tunnels* tab of the VPN / IPSec menu. The IPSec tunnel can be configured with a number of options:

| Option | Description |
|---------------------------|---|
| Disable this phase1 entry | If set, the tunnel will be disabled. |
| Interface | Select network interface to use for the VPN tunnel. |
| Remote gateway | Remote gateway (peer), i.e. the VPN router on the other side of the link this tunnel will be established with. |
| Description | Optional VPN tunnel description. |
| Authentication Method | Authentication method to use, <i>Mutual PSK</i> or <i>Mutual RSA</i> . This must match settings on the remote side. |

(Table continued)

| | |
|--------------------|--|
| Negotiation method | <p>Negotiation method to use – <i>main</i> or <i>aggressive</i>.</p> <p><i>Aggressive</i> is generally the most compatible with 3rd party IPSec implementations. It is faster because it sends all identifying information in a single packet. This makes it less secure than <i>main</i>.</p> <p><i>Main</i> is the most secure mode, though it also requires more packets between the peers to accomplish a successful negotiation. It is also much more strict, the identifier must be the remote side's IP address and not a custom identifier.</p> |
| My identifier | <p>Identification method for your Sentinel to use:</p> <p><i>My IP address</i> will use Sentinel PF network interface IP address (or VIP). The <i>Remote gateway</i> must see packets arriving from this IP address.</p> <p><i>IP address</i> allows to specify alternate IP address to use as an ID, in case Sentinel is behind NAT. You may specify Public IP address here.</p> <p><i>Distinguished name</i> use FQDN (Fully Qualified Domain Name) as ID. For example, example.host.com.</p> <p><i>User distinguished name</i> use e-mail address as ID. For example, vpn@example.com.</p> <p><i>ASN.1 distinguished name</i> for using Mutual RSA authentication, this can be the subject of the certificate being used, or a similar string.</p> <p><i>KeyID tag</i> allows to specify arbitrary string to use as an ID.</p> <p><i>Dynamic DNS</i> uses a hostname to resolve and use as an ID.</p> |
| Peer identifier | <p>Identification method for <i>Remote Gateway</i> (peer) to use:</p> <p><i>Peer IP address</i> will use the IP address specified in the <i>Remote Gateway</i> setting.</p> <p><i>IP address</i> allows to specify alternate IP address to use as an ID,</p> <p><i>Distinguished name</i> use FQDN (Fully Qualified Domain Name) as ID. For example, example.host.com.</p> <p><i>User distinguished name</i> use e-mail address as ID. For example, vpn@example.com.</p> <p><i>ASN.1 distinguished name</i> for using Mutual RSA authentication, this can be the subject of the certificate being used, or a similar string.</p> <p><i>KeyID tag</i> allows to specify arbitrary string to use as an ID.</p> <p><i>Dynamic DNS</i> uses a hostname to resolve and use as an ID.</p> |
| Pre-Shared Key | <p>Pre-Shared Key (PSK) for this tunnel. At least 15 characters is recommended.</p> |

(Table continued)

| | |
|----------------------|---|
| Policy Generation | <p>When working as a responder (as with mobile clients), this controls how policies are generated based on SA proposals. SA stands for Security Association and means a one-way tunnel. A VPN tunnel will usually have a pair of SA's for each direction. A policy enables tunnel peers to negotiate how to encrypt and send data.</p> <p><i>Require</i> or <i>On</i> will accept the first policy sent by a remote peer.</p> <p><i>Unique</i> will make Sentinel PF create and track unique policies per each client.</p> <p><i>Off</i> prevents policies from being generated automatically, instead relying on manual configuration.</p> |
| Proposal Checking | <p>This controls how Sentinel PF will respond on parameters proposed by the remote peer, in particular the action of lifetime length, key length, and PFS (Perfect Forward Secrecy) of the phase 2 selection on the responder side, and the action of lifetime check in phase 1. Please refer to <i>RFC 2412</i> for additional information on PFS.</p> <p><i>Obey</i> will accept and use parameters proposed by a remote peer.</p> <p><i>Strict</i> will ensure that only key length and lifetime values are used, unless peer's values are more secure. If PFS is used, PFS values must match on both sides.</p> <p><i>Claim</i> works similarly to <i>Strict</i>, except that it will notify the peer about an adjusted phase 2 lifetime and will use its own, if it is longer.</p> <p><i>Exact</i> will reject anything except an exact match of the values.</p> <p>By default, Sentinel PF uses <i>Claim</i> for normal tunnels and <i>Obey</i> for mobile IPsec tunnels.</p> |
| Encryption Algorithm | <p>Specifies phase 1 and phase 2 data encryption algorithm.</p> <hr/> <p>Note: Sentinel D2 servers feature hardware 128-bit AES acceleration, so it is the optimal setting.</p> |
| Hash Algorithm | <p>Specifies hash algorithm to verify authenticity of packet data.</p> |
| DH key group | <p>Diffie-Hellman key group. <i>1</i> = 768 bit, <i>2</i> = 1024 bit, <i>5</i> = 1536 bit.</p> |
| Lifetime | <p>Specifies how often the connection must be rekeyed, specified in seconds. 28800 seconds on phase 1 is a pretty standard configuration and is appropriate for most scenarios.</p> |
| NAT Traversal | <p>Set this option to enable the use of NAT-T (i.e. the encapsulation of ESP in UDP packets) if needed, which can help with clients that are behind restrictive firewalls. <i>Force</i> will make sure clients always use NAT-T.</p> |

(Table continued)

| | |
|---------------------|--|
| Dead Peer Detection | Dead Peer Detection is a periodic check that the host on the other end of the Ipsec tunnel is still alive. If a check fails, the tunnel is torn down by removing its associated SAD entries and renegotiation is attempted. The <i>Number of consecutive failures allowed before disconnect</i> specifies how many of Dead Peer Detection checks must fail before a tunnel is considered to be a down state. |
|---------------------|--|

Note: For network connections that may experience prolonged downtime due to, for example, poor weather conditions, it is better to leave Dead Peer Detection disabled.

Sentinel PF IPsec implementation supports mobile clients that can be authenticated via user name and password instead of a fixed IP address. Mobile clients can be configured in the *Mobile Clients* tab with the following options:

| Option | Description |
|--|---|
| IKE Extensions (Enable IPsec Mobile Clients Support) | Enables Mobile IPsec. |
| User Authentication | Defines authentication source for users. |
| Group Authentication | Defines authentication source for groups. |
| Virtual Address Pool | If enabled, Sentinel PF will provide virtual IP addresses to clients from the pool specified in CIDR notation. |
| Network List | This option controls whether the client will attempt to send all of its traffic across the tunnel, or only traffic for specific networks. If this option is checked, then the networks defined in the Local Network options for the mobile phase 2 definitions will be sent to the client. If this option is unchecked, the clients will attempt to send all of their traffic, including Internet traffic, across the tunnel. |
| Save Xauth Password | Allow clients to save Xauth passwords (Cisco VPN client only). Note: With iPhone clients, this does not work when deployed via the iPhone configuration utility, only by manual entry. |
| DNS Default Domain | If enabled, the specified domain will be pushed to clients as their default domain suffix for DNS requests. For example if this is set to <i>example.com</i> and a client requests <i>host</i> , then the DNS request will be attempted for <i>host.example.com</i> . |
| DNS Servers | If enabled, the specified DNS server(s) will be pushed to clients. |
| WINS Servers | If enabled, the specified WINS server(s) will be pushed to clients. |

(Table continued)

| | |
|------------------|---|
| Phase2 PFS Group | If enabled, Sentinel will provide the Phase2 PFS group to clients. This overrides all mobile phase2 settings. |
| Login Banner | If enabled, Sentinel will provide a login banner specified to clients. |

IPSec Pre-shared keys are managed from the *Pre-shared keys* tab whereby you can add keys and optional descriptions.

L2TP

L2TP (Layer 2 Tunneling Protocol) is another protocol implemented by Sentinel PF to support VPNs. It was published in 1999 as proposed standard *RFC 2661*. The entire L2TP packet, including payload and L2TP header, is sent within a User Datagram Protocol (UDP) datagram. L2TP does not provide encryption or strong authentication by itself.

The L2TP service can be configured with the following options:

| Option | Description |
|----------------------|---|
| Enable L2TP Server | Enables L2TP service. |
| Interface | Network interface the L2TP service to be provided on. |
| Server address | Enter the IP address the L2TP server should give to clients for use as their <i>gateway</i> . Typically this is set to an unused IP just outside of the client range. Note: This should <i>not</i> be set to any IP address currently in use on this firewall. |
| Remote address range | Specify the starting address for the client IP address subnet. |
| Subnet netmask | Netmask specified as bit counts. E.g. entering 24 will result in 255.255.255.0 netmask selected. See <i>Appendix B – Netmask/CIDR Translation Table</i> . |
| Number of L2TP Users | Specify number of L2TP users (clients). |
| Secret | Specify optional <i>secret</i> shared between peers. Required on some devices/setups. |
| Encryption type | Specifies which protocol to use for authentication. |
| L2TP DNS Servers | Primary and secondary DNS servers assigned to L2TP clients. |
| WINS Server | WINS server assigned to L2TP clients. |

(Table continued)

Use a RADIUS server for authentication

When set, all users will be authenticated using the RADIUS server specified below. The local user database will not be used.

| | |
|--------------------------|--|
| Enable RADIUS accounting | Sends accounting packets to the RADIUS server. |
| RADIUS server | Enter the IP address of the RADIUS server. |
| RADIUS shared secret | Enter the <i>shared secret</i> that will be used to authenticate to the RADIUS server. |
| RADIUS issued IP's | Issue IP Addresses via RADIUS server. |

Unless RADIUS authentication is used, L2TP clients (users) should be configured in the Users tab, whereby you can specify username, password and optional IP address in case you want a user to be assigned a specific IP address.

OpenVPN

OpenVPN is an SSL-VPN solution that accommodates a wide range of configurations, including remote access and site-to-site VPNs. It supports clients on a wide range of operating systems including all the BSDs, Linux, Apple OS X, Solaris and Windows. The OpenVPN security model is based on SSL. It implements OSI Layer 2 or 3 secure network extension using the SSL/TLS (Secure Sockets Layer/Transport Layer Security) protocol. There are several types of authentication methods that can be used with OpenVPN: shared key, X.509 (also known as SSL/TLS or PKI), user authentication via local, LDAP, and RADIUS, or a combination of X.509 and user authentication. For shared key, a single key is generated that will be used on both sides. SSL/TLS involves using a trusted set of certificates and keys.

OpenVPN client software packages for various operating systems can be downloaded on the *OpenVPN Technologies, Inc. web site*: <http://openvpn.net>

Sentinel PF allows to run multiple OpenVPN server instances at the same time. It can also run multiple client instances, connecting to remote OpenVPN based servers. A server instance can be added on the *Server* tab of the OpenVPN menu. Sentinel PF OpenVPN server is configured with the following options:

| Option | Description |
|----------|---|
| Disabled | Set this option to disable this server without removing it from the list. |

(Table continued)

| | |
|--|---|
| Server Mode | <p>This is the role for the server, which specifies how routers or users will connect to this server instance.</p> <p><i>Peer to Peer (SSL/TLS)</i>: A connection between local and remote networks that is secured by SSL/TLS.</p> <p><i>Peer to Peer (Shared Key)</i>: A connection between local and remote networks that is secured by Shared Key.</p> <p><i>Remote Access (SSL/TLS)</i>: An SSL-VPN setup with mobile clients and per-user X.509 certificates, secured by SSL/TLS.</p> <p><i>Remote Access (User Auth)</i>: An SSL-VPN setup with mobile clients whereby clients must provide username and password to connect. This option is not recommended due to security weakness, unless authentication is done externally via RADIUS or LDAP.</p> <p><i>Remote Access (SSL/TLS + User Auth)</i>: The most secure SSL-VPN setup with mobile clients and per-user X.509 certificates, secured by SSL/TLS <i>and</i> requirement for clients to provide username and password to connect.</p> |
| Protocol | Protocol for the VPN connection to operate on. |
| Device Mode | <p>Select the mode for OpenVPN operation. <i>tun</i> works on OSI Layer 3, and <i>tap</i> works on OSI Layer 2 and is capable of routing and bridging.</p> <hr/> <p>Note: Not all clients support <i>tap</i>, so <i>tun</i> is the recommended default setting.</p> |
| Interface | Network interface for the OpenVPN server instance to run. |
| Local port | Local port for the OpenVPN connection. |
| Description | Optional OpenVPN server instance description. |
| Enable authentication of TLS packets | Enable asymmetric cryptography for authentication of key exchange as per <i>RFC 5246</i> and <i>RFC 6176</i> . |
| Automatically generate a shared TLS authentication key | Automatically generate TLS key for authentication purposes. If this is not enabled, you will need to provide your shared key in the TLS Authentication section. |
| Peer Certificate Authority | Certificate authority (CA) to sign certificate for this OpenVPN server instance. If none is available then use <i>Cert Manager</i> to create a new CA or specify existing CA on the <i>CAs</i> tab. |
| Peer Certificate Revocation List | A list of CA's that are no longer considered valid. |
| Server Certificate | Certificate for this server instance. If none is available then use <i>Cert Manager</i> to create a new certificate or specify existing one on the <i>Certificates</i> tab. |

(Table continued)

| | |
|-----------------------------|--|
| Server DH Parameters Length | The Diffie-Hellman (DH) key exchange parameters (in bits) are used for establishing a secure communications channel. |
| Encryption algorithm | Cryptographic cipher to be used for this connection. Note: Sentinel D2 servers feature hardware 128-bit AES acceleration, so it is the optimal setting. |
| Hardware Crypto | Hardware cryptographic engine to be used for this connection, if available. |
| Certificate Depth | When a certificate-based client logs in, do not accept certificates below this depth. Useful for denying certificates made with intermediate CA's generated from the same CA as the server. |
| Strict User/CN Matching | When authenticating users, enforce a match between the <i>common name</i> of the client certificate and the <i>username</i> given at login. |
| Tunnel Network | This is the virtual network used for private communications between this server and client hosts expressed using CIDR (eg. 10.0.8.0/24). The first network address will be assigned to the server virtual interface. The remaining network addresses can optionally be assigned to connecting clients. |
| Redirect Gateway | Force all client generated traffic through the tunnel. This will replace the default gateway at clients. |
| Local Network | This is the network that will be accessible from the remote endpoint. Expressed as a CIDR range. You may leave this blank if you don't want to add a route to the local network through this tunnel on the remote machine. This is generally set to your LAN network. |
| Concurrent connections | Specify the maximum number of clients allowed to concurrently connect to this server. |
| Compression | Compress tunnel packets using the LZO algorithm. |
| Type of Service | Set the TOS IP header value of tunnel packets to match the encapsulated packet value. |
| Inter-client communication | Allow communication between clients connected to this server. |
| Duplicate Connections | Allow multiple concurrent connections from clients using the same Common Name. Note: This is not generally recommended, but may be needed for some scenarios. |
| Dynamic IP | Allow connected clients to retain their connections if their IP address changes. |
| Address Pool | Provide a virtual adapter IP address to clients (see <i>Tunnel Network</i>). |

(Table continued)

| | |
|----------------------------|--|
| DNS Default Domain | Provide a default domain name to clients. |
| DNS Servers | Provide a DNS server list to clients. |
| NTP Servers | Provide an NTP server list to clients. |
| Enable NetBIOS over TCP/IP | If this option is <i>not</i> set, all NetBIOS-over-TCP/IP options (including WINS) will be disabled. |
| Node Type | Possible options: <i>b-node</i> (broadcasts), <i>p-node</i> (point-to-point name queries to a WINS server), <i>m-node</i> (broadcast then query name server), and <i>h-node</i> (query name server, then broadcast). |
| Scope ID | A NetBIOS Scope ID provides an extended naming service for NetBIOS over TCP/IP. The NetBIOS scope ID isolates NetBIOS traffic on a single network to only those nodes with the same NetBIOS scope ID. |
| WINS Servers | Provide a WINS server list to clients. |
| Advanced | Enter any additional options you would like to add to the OpenVPN server configuration here, separated by a semicolon, e.g.: push "route 10.0.0.0 255.255.255.0"; |

Client specific overrides can be created in the *Client Specific Overrides* tab of the OpenVPN menu. This allows to specify an alternate tunnel network or *Redirect Gateway* option, as well as specify different DNS default domain, DNS servers, NTP servers, NetBIOS options and advanced configured – all on a per client basis. You can also block specific clients using this tab, if this may be required.

There is a set of wizards on the *Wizards* tab that allow simple, step-by-step guided setup of OpenVPN server instances in the Sentinel PF environment.

A VPN client instance can be added on the *Client* tab of the OpenVPN menu. Sentinel PF OpenVPN client is configured with the following options:

| Option | Description |
|-------------|--|
| Disabled | Set this option to disable this client without removing it from the list. |
| Server Mode | This is the role for the client, which specifies how the client connect to a remote server instance. <i>Peer to Peer (SSL/TLS)</i> : A connection between local and remote networks that is secured by SSL/TLS. <i>Peer to Peer (Shared Key)</i> : A connection between local and remote networks that is secured by Shared Key. |

(Table continued)

| | |
|--|---|
| Protocol | Protocol for the VPN connection to operate on. |
| Device Mode | Select the mode for OpenVPN operation. <i>tun</i> works on OSI Layer 3, and <i>tap</i> works on OSI Layer 2 and is capable of routing and bridging. Note: Not all clients support <i>tap</i> , so <i>tun</i> is the recommended default setting. |
| Interface | Network interface for the OpenVPN server instance to run. |
| Local port | Local port for the OpenVPN connection. Set this option if you would like to bind to a specific port. Leave this blank or enter 0 for a random dynamic port. |
| Server host or address | Server host or address. |
| Server port | Server port. |
| Proxy host or address | Optional proxy host or address. |
| Proxy port | Optional proxy port. |
| Proxy authentication extra options | Specify authentication method for proxy: <i>none</i> , <i>basic</i> or <i>ntlm</i> . With authentication, specify username and password. |
| Infinitely resolve server | Continuously attempt to resolve the server host name. Useful when communicating with a server that is not permanently connected to the Internet. |
| Description | Optional OpenVPN client instance description. |
| Enable authentication of TLS packets | Enable asymmetric cryptography for authentication of key exchange as per <i>RFC 5246</i> and <i>RFC 6176</i> . |
| Automatically generate a shared TLS authentication key | Automatically generate TLS key for authentication purposes. If this is not enabled, you will need to provide your shared key in the TLS Authentication section. |
| Peer Certificate Authority | Certificate authority (CA) to sign certificate for this OpenVPN server instance. If none is available then use <i>Cert Manager</i> to create a new CA or specify existing CA on the <i>CAs</i> tab. |
| Client Certificate | Certificate for this client instance. If none is available then use <i>Cert Manager</i> to create a new certificate or specify existing one on the <i>Certificates</i> tab. |
| Encryption algorithm | Cryptographic cipher to be used for this connection. Note: Sentinel D2 servers feature hardware 128-bit AES acceleration, so it is the optimal setting. |
| Hardware Crypto | Hardware cryptographic engine to be used for this connection, if available. |

(Table continued)

| | |
|--------------------------|--|
| Tunnel Network | This is the virtual network used for private communications between this client and the server expressed using CIDR (eg. 10.0.8.0/24). The first network address is assumed to be the server address and the second network address will be assigned to the client virtual interface. |
| Remote Network | This is a network that will be routed through the tunnel, so that a site-to-site VPN can be established without manually changing the routing tables. Expressed as a CIDR range. If this is a site-to-site VPN, enter here the remote LAN here. You may leave this blank to only communicate with other clients. |
| Limit outgoing bandwidth | Maximum outgoing bandwidth for this tunnel. Leave empty for no limit. The input value has to be something between 100 bytes/sec and 100 Mbytes/sec (entered as bytes per second). |
| Compression | Compress tunnel packets using the LZO algorithm. |
| Type of Service | Set the TOS IP header value of tunnel packets to match the encapsulated packet value. |
| Advanced | Enter any additional options you would like to add to the OpenVPN server configuration here, separated by a semicolon, e.g.: push "route 10.0.0.0 255.255.255.0"; |

PPTP

PPTP (Point-to-Point Tunneling Protocol) is another VPN solution. It uses a control channel over TCP and a GRE tunnel operating to encapsulate PPP packets. The protocol was developed by a consortium formed by Microsoft, Ascend Communications, 3Com and others. A specification for PPTP is available in *RFC 2637*.



Caution: PPTP is no longer considered a secure VPN technology because it relies upon MS-CHAPv2 which has been compromised. If you continue to use PPTP be aware that intercepted traffic can be decrypted by a third party, so it should be considered unencrypted. We advise migrating to another VPN type such as OpenVPN or IPsec.

A PPTP server instance can be configured with the following options:

| Option | Description |
|------------------|--|
| PPTP redirection | Enter the IP address of a host which will accept incoming PPTP connections. (This requires <i>Redirect incoming PPTP connections</i> to be enabled.) |

(Table continued)

| | |
|---|--|
| Enable PPTP server | Enables PPTP server instance. |
| No. PPTP users | Number of PPTP users (clients) to support. |
| Server address | Enter the IP address the PPTP server should give to clients for use as their <i>gateway</i> . Typically this is set to an unused IP just outside of the client range. <u>Note: This should <i>not</i> be set to any IP address currently in use on this firewall.</u> |
| Remote address range | Specify the starting address for the client IP subnet. |
| PPTP DNS Servers | Primary and secondary DNS servers assigned to PPTP clients. |
| WINS Server | WINS server assigned to PPTP clients. |
| Use a RADIUS server for authentication | When set, all users will be authenticated using the RADIUS server specified below. The local user database will not be used. |
| Enable RADIUS accounting | Sends accounting packets to the RADIUS server. |
| Secondary RADIUS server for failover authentication | When set, all requests will go to the secondary server when primary fails. |
| [Secondary] RADIUS server | Enter the IP address of the RADIUS server. |
| [Secondary] RADIUS shared secret | Enter the <i>shared secret</i> that will be used to authenticate to the RADIUS server. |
| RADIUS issued IP's | Issue IP Addresses via RADIUS server. |
| RADIUS NAS IP | IP address for RADIUS NAS. |
| Require 128-bit encryption | When set, only 128-bit encryption will be accepted. Otherwise 40-bit and 56-bit encryption will be accepted as well. Note that encryption will always be forced on PPTP connections (i.e. unencrypted connections will not be accepted). |

Unless RADIUS authentication is used, PPTP clients (users) should be configured in the Users tab, whereby you can specify username, password and optional IP address in case you want a user to be assigned a specific IP address.

◆ ◆ ◆ CHAPTER 10

Web GUI: Status

Status menu items are used to monitor Sentinel PF: CARP (failover), Dashboard, DHCP Leases, Filter Reload, Gateways, Interfaces, IPSec, Load Balancer, NTP, OpenVPN, Package Logs, Proxy report, Queues, RRD Graphs, Services, System Logs, Traffic Graph, UPnP & NAT-PMP.

CARP (failover)

Provides a status overview of CARP interfaces and their Virtual IPs.

Dashboard

Returns user to the Sentinel PF dashboard overview.

DHCP Leases and DHCPv6 Leases

Provides a status overview of DHCP leases, their IP and MAC addresses, hostnames, lease start and end and lease type.

Filter Reload

Used to reload Sentinel PF packet filter. Click *Reload Filter* button to force immediate reload.

Gateways

Provides a status overview of gateways, their names, monitoring IP address, RTT, packet loss and description.

iDirect Monitor ◆

The iDirect Monitor is implemented via *Huginn* external package developed by BusinessCom Networks. It provides a status overview of the iDirect VSAT modems. The Huginn monitoring daemon can be configured with the following options on the *Daemon Configuration* tab:

| Option | Description |
|----------------------------------|---|
| Enable Daemon | Enables Huginn monitoring daemon. |
| Use WAN Gateway | Assume Sentinel WAN gateway's IP address points to the iDirect satellite modem. |
| iDirect satellite router address | IP address of the iDirect satellite router. |
| User Name | iDirect satellite router user name, typically <i>admin</i> . |
| User Password | iDirect satellite router password. |
| Query Mode | iDirect satellite router query mode. <i>Normal</i> assumes standard telnet query via port 23. <i>Secure</i> uses SSH via port 22. |
| Query Timeout | Connection timeout in seconds. Default is 10. |
| Update Period and Time Format | Update period in hours or minutes. |
| Enable daemon log | Enable Huginn error reporting. |

The iDirect modem information is accessible via *System Info*, *Remote Status* and *RX Info* tabs. You can click on *Refresh Now* button to force Huginn an immediate status update.

Interfaces

Provides a status overview of network interfaces, their MAC and IP addresses, configuration (subnet mask, gateway, DNS servers), media status. This menu also provides a breakdown of IP packets traveling in and out of the interface as total amount of packets, packets passed and packets blocked, as well as I/O errors and collisions.

IPSec

Provides a status overview of IPSec connections, local, remote and local/remote networks IP addresses and description on the *Overview* tab. Tabs *SAD* and *SPD* provide overview of IPSec Security Associations and Security Policies. The *Logs* tab provides access to IPSec log files.

Load Balancer

Provides a status overview of load balancer policies: name, mode, servers and monitor status and description. The *Virtual Servers* tab provides an overview of defined virtual servers.

Network Monitor

Provides per-IP bandwidth usage breakdown.

NTP

Provides a status overview of the NTP daemon.

OpenVPN

Provides a status overview of OpenVPN client and server connections.

Package Logs

Logs associated with Sentinel PF external software packages.

Proxy Report ◆

The Proxy Report service is implemented via LightSquid/SQStat external package and it provides a status overview of Sentinel PF's Squid proxy server. The Proxy Report service can be configured with the following options on the *Settings* tab:

| Option | Description |
|-------------------|--|
| Language | Select language for the report. |
| Bar color | Report color scheme. |
| Report scheme | Report design scheme. |
| IP resolve method | Select IP to Name resolve method (take effect only on new data): <i>IP</i> - return IP address <i>Demo</i> - return AUTHNAME, else DNSNAME, else IP <i>DNS</i> - return DNSNAME <i>Simple</i> - return AUTHNAME else IP <i>SMB</i> - return SMB name of PC <i>Squidauth</i> - return AUTHNAME else IP, allow cyrillic name |
| Refresh scheduler | Select data refresh period. |
| Skip URL | Exclude URLs from the report. |

The Proxy Report is accessible via *Lightsquid Report* tab. It provides an overview of cache hit rate, local users, shown as specified by the *IP resolve method*, and the corresponding proxy *Connect* events, bytes transferred and percentage.

Note: You may need to force-reload the page in your browser to see the latest proxy report. On Chrome, Mozilla Firefox and Internet Explorer browsers, use Ctrl+F5. On Apple Safari, use ⌘+r.

The *Proxy State* tab allows to view real-time proxy server connections and throughput.

Queues

Provides a status overview of traffic shaper queues.

RRD Graphs

Provides charts related to various Sentinel PF components:

System Tab: Packet filter throughput, number of states (connections), Sentinel server CPU utilization and RAM usage.

Traffic Tab: Network interface traffic charts.

Packets Tab: Packet forwarding rates (in pps, packets per second).

Quality Tab: Gateway QoS related information, such as RTT (Round Trip Time), packet loss and delay average.

Customer Tab: Generates custom charts as per user input.

Services

Provides a status overview of Sentinel PF services (daemons). You can stop any service by pressing the  button next to the service name, or start (or restart) a service by pressing .

System Logs

Provides access to various Sentinel PF logs via the *syslog* service, represented by multiple tabs in this menu. This includes logs for System, Firewall, DHCP, Portal Auth (Captive Portal), IPSec, PPP, VPN, Load Balancer, OpenVPN, NTP and Wireless connections. Sentinel PF logger can be configured using the *Settings* tab.

Note: If logging to a remote server, syslog sends UDP datagrams to port 514 on the specified remote syslog server. Be sure to set *syslogd* on the remote server to accept syslog messages from Sentinel PF.

Traffic Graph

Provides real-time traffic graph for network interface throughput.

Note: The Adobe SVG Viewer, Firefox 1.5 or later or other browser supporting SVG is required to view the graph.

UPnP & NAT-PMP Status

Provides a status overview of Sentinel UPnP and NAT-PMP services.

◆ ◆ ◆ CHAPTER 11

Web GUI: Diagnostics

Status menu items are used to access diagnostic features of Sentinel PF: ARP Table, Authentication, Backup/Restore, Command Prompt, DNS Lookup, Edit File, Factory Defaults, Halt System, Limiter Info, Packet Capture, pfInfo, pfTop, Ping, Reboot, Routes, SMART Status, States, States Summary, System Activity Tables and Traceroute.

ARP Table

Provides a status overview of Sentinel PF ARP Table with a list of MAC addresses and their associated IP addresses, hostnames and network interface.

Authentication

The *Authentication* menu allows to test credentials with Authentication Servers configured in Sentinel PF.

Backup/Restore

Provides backup/restore functionality for Sentinel PF configuration.

▼ To Backup Current Configuration

- 1 **In Diagnostics / Backup/Restore, select the «Backup/Restore» tab.**
- 2 **In «Backup configuration» section, select backup area(s).**
Select *ALL* if you would like to backup all configuration.
- 3 **Choose if you would like to backup configuration for external packages like Snort and Squid. This is controlled by the «Do not backup package information» checkbox.**
We recommend not to enable this box, as external Sentinel PF packages often play crucial roles in many Sentinel deployments.
- 4 **Choose if you would like to encrypt the configuration file.**
If encryption is chosen, provide username and password.
- 5 **Choose if you would like to backup RRD data.**
RRD data holds historic throughput statistics. In most of the cases, you do not need to backup RRD.
- 6 **Click «Download configuration» button.**

▼ To Restore Configuration From File

- 1 **In Diagnostics / Backup/Restore, select the «Backup/Restore» tab.**
- 2 **In «Restore configuration» section, select backup area(s).**
Select *ALL* if you would like to restore all configuration.
- 3 **Choose a local file to restore configuration from.**
- 4 **Choose if the file is encrypted.**
If encryption is chosen, provide username and password.
- 5 **Click «Restore configuration» button.**

Note: Sentinel PF will reboot after the configuration is restored.

The *Config History* tab allows to see a log of Sentinel PF configuration changes. You can revert to any previous configuration by pressing the  button next to the configuration point in the history. You can also download the configuration by pressing the  button, or remove the configuration from history using the  button.

▼ To View Changes Between Configurations

- 1 In *Diagnostics / Backup/Restore*, select the «*Config History*» tab.
- 2 Select the first configuration point by pressing the radio button next to it.
- 3 Select the second configuration point by pressing the radio button next to it.
- 4 Click «*Diff*» button.
Configuration changes will be presented in the *diff* format. Old values are shown with red background and (-) sign, and new values are shown with green background and (+) sign.

Command Prompt

This menu option provides low-level shell access to Sentinel PF operating system based on the FreeBSD kernel via Web GUI.



Caution – Low-level shell access is not required and shall not be used for any bandwidth management purposes. Sentinel provides low-level access for remote troubleshooting by BusinessCom staff only.

DNS Lookup

Enables to resolve hostnames or IP addresses via Sentinel PF DNS service.

Edit File

Enables to edit files on Sentinel PF storage.



Caution – This is for remote troubleshooting by BusinessCom staff only.

Factory Defaults

This menu option resets Sentinel PF configuration to factory defaults. Please consider *Appendix C – Factory Default Settings* for reference.

▼ To Reset Sentinel

- 1 Enter the **Diagnostics / Factory Defaults** menu.
- 2 Confirm reset by pressing **«Yes»** button.



Caution – With Reset to factory defaults, all your configuration will be lost.

Halt System

The *Halt System* menu allows to do a soft shutdown (halt) of a Sentinel server. This is the preferred method of shutting down Sentinel.

Limiter Info

Provides an overview of configured bandwidth *limiters*.

Packet Capture

The Sentinel PF *Packet Capture* feature allows to capture IP packets for detailed traffic analysis. The output packet capture may contain link level headers, ARP/RARP packets, TCP and UDP packets, name server requests and responses, decoded SMB/CIFS packets, NFS and AFS requests and replies, RIP AppleTalk, NBP and ATP packets and other information. A packet capture session is configured with the following options:

| Option | Description |
|----------------|---|
| Interface | Select the interface on which to capture traffic. |
| Address Family | Select the type of traffic to be captured, either <i>Any</i> , <i>IPv4 only</i> or <i>IPv6 only</i> . |

(Table continued)

| | |
|--------------------|---|
| Host Address | This value is either the Source or Destination IP address or subnet in CIDR notation. The packet capture will look for this address in either field. This value can be a domain name or IP address, or subnet in CIDR notation. If you leave this field blank, all packets on the specified interface will be captured. |
| Port | The port can be either the source or destination port. The packet capture will look for this port in either field. Leave blank if you do not want to filter by port. |
| Packet Length | The Packet length is the number of bytes of each packet that will be captured. Default value is 0, which will capture the entire frame regardless of its size. |
| Count | This is the number of packets the packet capture will grab. Default value is 100. Enter 0 (zero) for no count limit. |
| Level of Detail | This is the level of detail that will be displayed after hitting <i>Stop</i> when the packets have been captured. <u>Note: This option does not affect the level of detail when downloading the packet capture.</u> |
| Reverse DNS Lookup | This check box will cause the packet capture to perform a reverse DNS lookup associated with all IP addresses. <u>Note: This option can cause delays for large packet captures.</u> |

After setting up the capture, click on *Start* button to start capture. The capture will stop if the packet *Count* is reached, or you can stop the capture at any time by pressing the *Stop* button. Packet capture results will appear in the *Packets Captured* section on the same page. You can download the capture as a *cap* file by clicking the *Download Capture* button.

pflInfo

The *pfInfo* menu provides a comprehensive overview of Sentinel PF packet filter statistics. The information shown on the page contains items such as bytes in/out, packets in/out and passed/blocked; state table entry count, search rate, insertion rate, removal rate, source tracking entry count, search rate, insertion rate, removal rate; counter statistics for various types of special packets, counters for packets dropped due to exceeding limits such as max states per IP, state table max size, source node table size, frag table size, number of allowed tables, and maximum number of table entries; state timers for TCP, UDP, and other connections and per-interface packet counters.

pfTop

This menu option provides text based overview of network connections generated by the *pfTop* utility. pfTop displays the active packet filter states and rules, and periodically updates this information. The output is identical to the CLI version of pfTop, thus refer to the *Chapter 3 - Command Line Interface, View Network Connections* section.

Ping

The ping utility uses the ICMP protocol's mandatory ECHO_REQUEST datagram to elicit an ICMP ECHO_RESPONSE from a host or gateway. The usage is identical to the CLI version of Ping host, thus refer to the *Chapter 3 - Command Line Interface, Ping host* section.

Reboot

The *Reboot* menu allows to do a soft reboot of a Sentinel server. This is the preferred method of rebooting Sentinel.

Routes

The *Routes* menu shows Sentinel PF routing tables for IPv4 and IPv6. You can enable name resolution in the tables for your convenience. The following information is displayed per each route:

TABLE 11-1 Sentinel PF Route Table Data

| Field | Meaning |
|-------------|---|
| Destination | Destination host or subnet in CIDR format. |
| Gateway | Gateway used to reach <i>Destination</i> . |
| Flags | Route flags (see <i>Table 11-2, Sentinel PF Route Flags</i>) |
| Refs | Number of references to this route. |
| Use | Count of lookups for this route. |
| MTU | Maximum Transmission Unit for this route. |
| Netif | Network interface used for this route. |
| Expire | Expiration time for this route. |

TABLE 11-2 Sentinel PF Route Flags

| Flag | Meaning |
|------|---|
| 1 | Protocol specific routing flag #1 |
| 2 | Protocol specific routing flag #2 |
| 3 | Protocol specific routing flag #3 |
| B | Blackhole – Sentinel PF will discard packets on this route. |
| b | This route represents a broadcast address. |
| C | Generates a new route based upon this route for machines Sentinel PF connects to. This type of route is normally used for local networks. |
| c | Protocol-specific version of <i>C</i> flag. |
| D | Created dynamically. |
| G | Gateway: send anything for this destination on to this remote system, which will figure out from there where to send it. |
| H | Host entry (network otherwise). |
| L | Valid protocol to link address translation. |
| M | Modified dynamically. |
| R | Host or net unreachable. |
| S | Static route, manually added. |
| U | The route is <i>up</i> , i.e. active and usable. |
| W | Route was generated as a result of cloning. |
| X | External daemon translates proto to link address. |

SMART Status

The *SMART Status* menu allows to diagnose Sentinel's built in SSD or HDD storage via SMART. (Self-Monitoring, Analysis and Reporting Technology) system. The purpose of SMART is to warn a user of impending drive failure while there is still time to take action, such as copying the data to a replacement device. The following SMART tests are available:

Info: Display general information about the storage device, such as model, number, device ID, firmware version, user capacity, sector size, ATA version and standard and SMART support status.

Health: Perform overall SMART health self-assessment test.

SMART Capabilities: Display SMART capabilities of a storage device, such as error logging capability, offline surface scan, self-test and other.

Attributes: Display raw SMART attributes of a storage device.

All: Display all information and perform SMART health self-assessment test.

The same SMART Status menu can be used to perform storage device tests. There are multiple test modes available:

Offline: This immediately starts the test in foreground mode. This command can be given during normal system operation. The effects of this test are visible only in that it updates the SMART Attribute values, and if errors are found they will appear in the SMART error log.

Short: runs SMART Short Self Test (usually under ten minutes). This command can be given during normal system operation. The Short test checks the electrical and mechanical performance as well as the read performance of the disk. Results appear in logs.

Long: This is a longer and more thorough version of the Short Self Test described above. Results appear in logs.

Conveyance: This test's primary purpose is to test the drive after it has been physically relocated to determine if any components have been damaged by the move. It should only take a few minutes to complete.

SMART error and self test logs can be viewed in the *View Logs* section. Background tests can be aborted in the *Abort tests* section.

States

The *States* menu provides an overview of states (connections) established from/to Sentinel server. The information provided is protocol, source, router and destination and connection state. It is possible to remove a state entry manually (break the connection) from Sentinel PF by clicking on the  button. States can be filtered using an expression. A list of possible states is provided in the *Table 11-3, Sentinel PF TCP Connection States*.

It is possible to reset *all* Sentinel PF states using the *Reset States* tab. Resetting the state tables will remove all entries from the corresponding tables. This means that all open connections will be broken and will have to be re-established. This may be necessary after making substantial changes to the firewall and/or NAT rules, especially if there are IP protocol mappings (e.g. for PPTP or IPv6) with open connections. The firewall will normally leave the state tables intact when changing rules.

Note: If you reset the firewall state table, the browser session may appear to be hung after clicking *Reset*. Simply refresh the page to continue.

The TCP layer in the host at each end of a TCP connection keeps its own variable containing the state of the connection, using the connection states defined in *RFC 793*.

TABLE 11-3 Sentinel PF TCP Connection States

| State | Typical Situation |
|--------------|---|
| LISTEN | Waiting for a connection request. |
| SYN_SENT | Application has initiated connection and is waiting for a matching connection request after having sent a connection request. |
| SYN_RECEIVED | Waiting for a confirming connection request acknowledgement after having both received and sent a connection request. |
| ESTABLISHED | Open, established connection. Data transfer phase. |
| FIN_WAIT_1 | Waiting for a connection termination request, or an acknowledgment of a previous termination request. |
| FIN_WAIT_2 | Waiting for a connection termination request. |
| CLOSE_WAIT | Waiting for a connection termination request from the local user. |
| CLOSING | Waiting for a connection termination request. |
| LAST_ACK | Waiting for an acknowledgment of the connection termination request previously sent. |
| TIME_WAIT | Waiting for enough time to pass to be sure the remote side received the acknowledgment of its connection termination request. |
| CLOSED | No connection. |

Similarly, states for UDP connections are also reported:

TABLE 11-4 Sentinel PF UDP Connection States

| State | Typical Situation |
|------------|--|
| SINGLE | Connection has sent a single packet. |
| MULTIPLE | Connection has sent multiple packets. |
| NO_TRAFFIC | No traffic received from the remote end. |

For example, a state table entry *SINGLE:NO_TRAFFIC* indicates that a connection has sent a single packet to the remote end, however has not yet received any traffic back.

States Summary

The *States Summary* menu provides a summary overview of states (connections) established from/to Sentinel server, with breakdown by source IP, destination IP, total per IP and by IP pair. The information provided is protocol, number of states, and the amount of source and destination ports utilized.

System Activity

The *System Activity* menu provides an overview of running Sentinel PF processes, CPU and memory usage.

Tables

The *Tables* menu provides access to diagnostic tables for various Sentinel PF components.

Traceroute

This menu option provides *traceroute* functionality. Traceroute utilizes the IP protocol TTL (Time To Live) field and attempts to elicit an ICMP TIME_EXCEEDED response from each gateway along the path to some host, in order to trace the route from Sentinel to the host. By default UDP datagrams are used, however user can force using ICMP ECHO instead.

Note: Traceroute may take a while to complete. You may hit the *Stop* button on your browser at any time to see the progress of failed traceroutes.

NanoBSD

This menu allows to control various aspects of the BSD operating system that is the foundation of the Sentinel PF. We do not recommend end users to change any values in this menu without supervision of BusinessCom representative.

NDP Table

Provides a status overview of Sentinel PF IPv6 NDP (Network Discovery Protocol) Table with a list of MAC addresses and their associated IP addresses, hostnames and network interface. This is roughly analogous to the ARP table for IPv4.

Sockets

Provides an overview of all socket connections and listening sockets for both IPv4 and IPv6 and associated information, as specified in the table below.

TABLE 11-5 Sentinel PF Socket Information

| Column | Description |
|-----------------|---|
| USER | The user who owns the socket. |
| COMMAND | The command which holds the socket. |
| PID | The process ID of the command which holds the socket. |
| FD | The file descriptor number of the socket. |
| PROTO | The transport protocol associated with the socket for Internet sockets, or the type of socket (stream or datagram) for UNIX sockets. |
| ADDRESS | (UNIX sockets only) For bound sockets, this is the file-name of the socket. For other sockets, it is the name, PID and file descriptor number of the peer, or (<i>none</i>) if the socket is neither bound nor connected. |
| LOCAL ADDRESS | (Internet sockets only) The address the local end of the socket is bound to (see FreeBSD manual <i>getsockname(2)</i>). |
| FOREIGN ADDRESS | (Internet sockets only) The address the foreign end of the socket is bound to (see FreeBSD manual <i>getpeername(2)</i>). |

Test Port

This page allows you to perform a simple TCP connection test to determine if a host is up and accepting connections on a given port. This test does not function for UDP since there is no way to reliably determine if a UDP port accepts connections in this manner. No data is transmitted to the remote host during this test, it will only attempt to open a connection and optionally display the data sent back from the server.

| Option | Description |
|------------------|--|
| Host | IP address of destination host to test a port. |
| Port | Destination port number. |
| Source Port | Source port number. This should typically be left blank. |
| Show Remote Text | Shows the text given by the server when connecting to the port. Will take 10+ seconds to display if checked. |
| Source Address | IP address where the connection should be initiated from. |
| IP Protocol | IP protocol to use for the test: <i>IPv4</i> or <i>IPv6</i> |

Note: If you force IPv4 or IPv6 and use a hostname that does not contain a result using that protocol, it will result in an error. For example if you force IPv4 and use a hostname that only returns an AAAA IPv6 IP address, it will not work.

◆ ◆ ◆ CHAPTER 12

Redundancy

This chapter discusses redundant Sentinel server setup for high availability environments.

Sentinel Redundancy Overview

The Sentinel PF operating system offers enterprise-class high availability capabilities with stateful failover, allowing the elimination of Sentinel as the single point of failure in your network. This is provided by a combination of CARP (Common Address Redundancy Protocol), *pfSync* and XML-RPC configuration synchronization. Often this is simply referred to as CARP, though CARP is just a part of the redundant solution.

CARP is a protocol similar to Cisco VRRP (Virtual Router Redundancy Protocol) that features better licensing freedom. Each Sentinel server in a CARP group has its own unique IP address assigned, and also has the shared CARP VIPs (Virtual IP Addresses) assigned as well. CARP multicasts «heartbeat» packets to other Sentinel servers in the group. If any of the servers fails to respond to the heartbeat packet, the server is considered as failed. If the heartbeat response arrives not as fast as expected, the responding server is assigned a lower priority in the group.

The *pfSync* technology enables the synchronization of the firewall state table from the master Sentinel to secondary Sentinel servers in a group. This allows client connections to survive during the fail-over event. Please note that *pfSync* must be enabled on all servers in a group to work.

The Sentinel PF XML-RPC technology allows to maintain synchronized system configuration over multiple Sentinel servers in a group. The XML-RPC supports synchronization of users and groups, certificates, route tables and gateways, firewall rules and schedules, aliases, settings for NAT, IPSec, OpenVPN, DHCP, Wake on LAN, Load Balancer, Virtual IPs, Traffic Shaper (queues, limiters and Layer 7 rules), DNS forwarder and captive portal. Other packages like Snort and Squid will require manual configuration synchronization. XML-RPC must be enabled only on the master server, and all other servers should have XML-RPC *disabled*.

Example CARP Redundant Configuration

This section will describe an example redundant Sentinel configuration. Three network interfaces will be used per each Sentinel server: LAN, WAN and an optional interface OPT1 for pfSync traffic.

The optional interface is used solely to synchronize configuration and firewall states between primary and secondary Sentinel servers in the group. This has been done due to security considerations, as pfSync does not use any authentication method and its traffic is not encrypted.

We begin with IP address assignments, provided in tables below. Please note that shared CARP IP address for the optional interface is not required.

TABLE 12-1 WAN IP Address Assignment: Example Redundant Configuration

| IP Address | Usage |
|------------|-----------------------------------|
| 10.0.66.10 | CARP shared IP Address |
| 10.0.66.11 | Primary Sentinel server, WAN IP |
| 10.0.66.12 | Secondary Sentinel server, WAN IP |

TABLE 12-2 LAN IP Address Assignment: Example Redundant Configuration

| IP Address | Usage |
|-------------|-----------------------------------|
| 192.168.1.1 | CARP shared IP Address |
| 192.168.1.2 | Primary Sentinel server, LAN IP |
| 192.168.1.3 | Secondary Sentinel server, LAN IP |

TABLE 12-3 OPT1 (Optional) IP Address Assignment: Example Redundant Configuration

| IP Address | Usage |
|------------|--------------------------------------|
| 172.16.1.2 | Primary Sentinel server, pfSync IP |
| 172.16.1.3 | Secondary Sentinel server, pfSync IP |

Figure 12-1 provides a network diagram for the example CARP network described above.

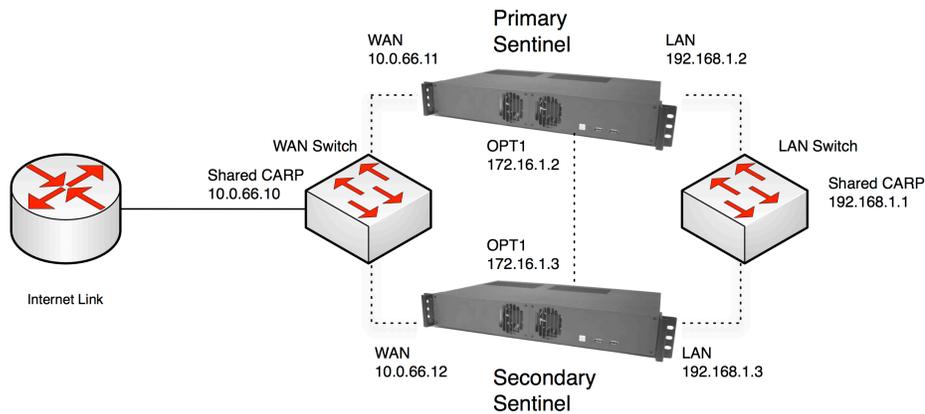


FIGURE 12-1 Example Redundant Configuration Diagram.

The next step is to configure the primary Sentinel server:

▼ Redundancy Configuration Example

1 Add Shared CARP Virtual IP address on WAN Interface.

In Firewall menu, Virtual IPs menu, click on the  button to add a new Virtual IP. Specify Type: CARP, Interface: WAN, IP Address: 10.0.66.10/24, arbitrary Virtual IP password (any string). VHID Group can be 1, unless there are other CARP or VRRP networks connected, otherwise it should be unique for this setup. Advertising Frequency Skew is 0 for this machine (Master). Description: WAN CARP VIP.

2 Add Shared CARP Virtual IP address on LAN Interface.

In Firewall menu, Virtual IPs menu, click on the  button to add a new Virtual IP. Specify Type: CARP, Interface: LAN, IP Address: 192.168.1.1/24, arbitrary Virtual IP password (any string – does not have to match the WAN VIP). VHID Group can be 2, unless there are other CARP or VRRP networks connected and using that number. The VHID Group should be different from the WAN VHID Group (1) assigned earlier. Advertising Frequency Skew is 0 for this machine (Master). Description: LAN CARP VIP.

3 Configure outbound NAT for CARP.

In Firewall menu, NAT menu, select *Outbound* tab, enable *Manual Outbound NAT* and click *Save*. Edit the Auto created rule for LAN to WAN by clicking on the  button. In the *Translation* section, specify WAN CARP VIP as the Address. Click on *Save* to apply settings. Now outbound traffic is routed via shared WAN CARP IP address.

(Table continued)

- 4 Configure OPT1 interface for pfSync.**
Configure OPT1 interface in the *Interfaces* menu. The configuration is no different from any other interface. If OPT1 is not present in the interfaces menu then you will need to assign it first. In the *Interfaces / (assign)* menu, *Interface assignments* tab, click on the  button to add OPT1.
- 5 Add OPT1 interface firewall rule.**
In the *Firewall / Rules* menu, create firewall rules to allow all traffic from/to the OPT1 interface. This is secure, assuming optional interfaces on primary and secondary Sentinel servers will be connected with a cross-over cable.
- 6 Set up DHCP server for CARP.**
In case Sentinel will act as DHCP server for the network, the DHCP server has to be adjusted to assign CARP IP as the gateway IP to clients. In the *Services / DHCP Server* menu, change the *Gateway* field to Shared LAN CARP VIP (192.168.1.1). Set the *Failover Peer IP* to the actual LAN IP of the secondary Sentinel server (192.168.1.3). This will allow DHCP service on both Sentinel servers to maintain a common set of leases. Click *Save*.
- 7 Assign interfaces and configure IP addresses on secondary Sentinel server.**
Repeat steps 1 to 7, adjusted for the perspective of the secondary server in this network. Please note the Advertising Frequency Skew should be 1 for this machine (Backup).
- 8 Set up configuration synchronization on secondary Sentinel server.**
On the secondary Sentinel server, go to *System / High Avail. Sync* menu. Check *Synchronize States*, pick OPT1 as the *Synchronize Interface*, and for the *pfSync Synchronize Peer IP*, enter the IP address for the primary system's OPT1 interface (172.16.1.2). Click *Save* when finished. Do not set any other values on this page.
- 9 Set up configuration synchronization on primary Sentinel server.**
On the primary Sentinel server, go to *System / High Avail. Sync* menu. Check *Synchronize States*, pick OPT1 as the *Synchronize Interface*, and for the *pfSync Synchronize Peer IP*, enter the IP address for the secondary system's OPT1 interface (172.16.1.3). Enter the backup system's OPT1 IP again in *Synchronize Config to IP*. Enter Web GUI admin password in the *Remote System Password* section. Click *Save* when finished.

(Table continued)

10 Initial Sync.

When the synchronization settings are saved on the primary, it will automatically copy the settings from the primary to the secondary Sentinel for each selected option on the *High Avail. Sync* page. This includes the proper outbound NAT settings for CARP, the firewall rules for the OPT1 interface, and even the CARP VIPs. Within 30 seconds, the initial configuration sync should have finished. The DHCP server settings are synchronized, and the system is smart enough to adjust the failover IP as needed to make sure that DHCP failover works.

Note: You should *not* setup *Synchronize Config to IP* from the secondary Sentinel to the primary. There are protections that should prevent this synchronization loop from causing harm, but it will clutter your logs with error messages and should never be configured this way. You should setup *Synchronize States* on both the primary and the secondary, as described above.

The procedure above assumes both Sentinel servers connect to the same WAN and LAN switch which could be a single point of failure. To avoid this, the deployment can be adapted to use a pair of switches on WAN and LAN interfaces:

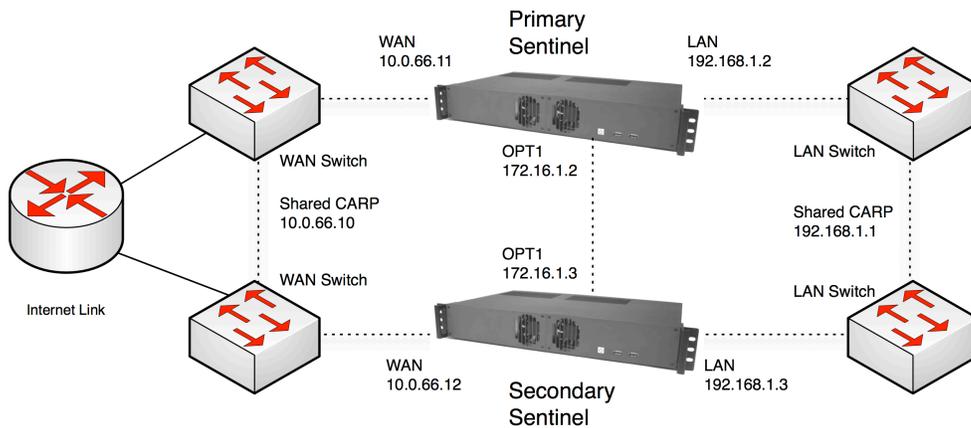


FIGURE 12-2 Example Redundant Configuration Diagram With Multiple Switches.

As long as switches are interconnected and do not bridge on either of the firewalls, this is safe with any type of switch. Where using bridging, or where multiple interconnections exist between the switches, care must be taken to avoid Layer 2 loops.

Multiple WAN Connections

This chapter discusses Sentinel setup for high availability environments with multiple WAN gateways. Topics include failover and load balancing between multiple WAN connections.

Mutli-WAN Environment

In addition to CARP redundancy, the Sentinel PF operating system supports multiple WAN connections for failover purposes. Multiple WAN gateways can also be load balanced to assure the most effective bandwidth utilization in a network.

Consider the following expansion of redundant CARP setup provided earlier:

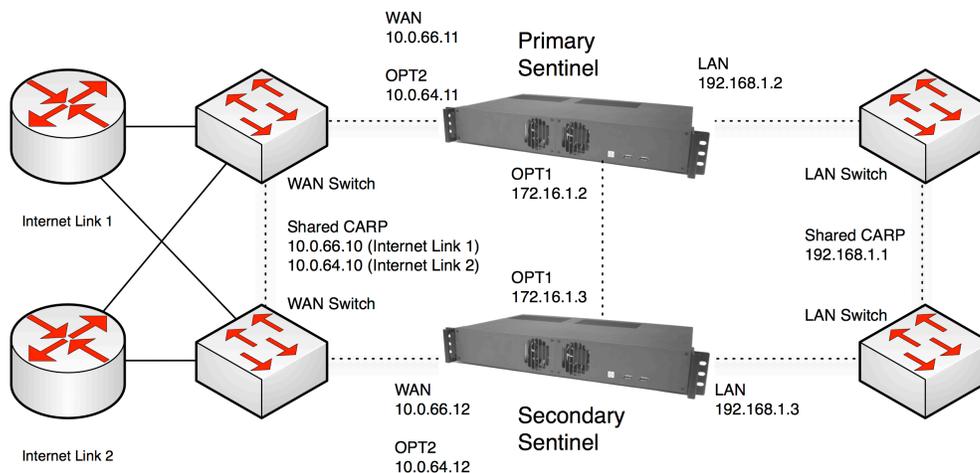


FIGURE 13-1 Example Redundant Configuration Diagram With Multiple Switches and Internet connections.

In this diagram, the network has two Internet links. There are two Shared CARP IPs, one per each WAN connection, as well as each Sentinel server has 2nd optional interfaces (OPT2) configured for the Internet Link 2.

Example Redundant Multi-WAN Configuration

The CARP/pfSync part of the redundancy is configured similarly as described in *Chapter 12 – Redundancy*, thus it will not be repeated in this chapter. The Multi-WAN load balancing part is configured as following:

▼ Multi-WAN Example: Configuration

1 Create OPT2 Gateway.

In System, Routing menu, select the *Gateways* tab click on the  button to add a new gateway for the OPT2 interface. Do not make it a default gateway.

2 Create Gateway Group.

In System, Routing menu, select the *Groups* tab click on the  button to add a new gateway group. Assign tiers in *Gateway Priorities* to both WAN and OPT2 gateways in this group. Tiers available are from 1 to 5, with Tier 1 being the highest priority gateway.

For load balancing scenario: assign Tier 1 to both WAN and OPT2 gateways. Traffic will be load balanced between both Internet connections, and failover will work automatically.

For weighted balancing scenario: you can assign different weights to each gateway in the group in case the amounts of bandwidth available on these interfaces are different. Weights are configured in the *System / Routing* menu, *Gateways* tab, *Advanced* section. For example, if you have 10 Mbps available on the WAN connection and 5 Mbps on the OPT2 connection, you can assign weight 2 to WAN and weight 1 to OPT2. In this example, for three connections that will go out, two will be routed to WAN, and one will be routed to OPT2.

For failover scenario: assign different Tiers to WAN and OPT2 gateways.

Notes:

Remember to use Gateway Groups in firewall rules in order to enable load balancing, failover, or policy-based routing. Without rules directing traffic into the Gateway Groups, they will not be used.

Due to the limited number of Ethernet interfaces available, the example configuration discussed above, with OPT1 interface used for pfSync traffic and OPT2 for backup Internet connection, can be implemented with Sentinel Westwind servers only.

State Filter Expressions

The expression filter selects which states will be displayed. It is based on the *tcpdump filtering language*. The following is based on the *tcpdump manual page*, modified for state filtering. The filter expression consists of one or more primitives. Primitives usually consist of an *id* (name or number) preceded by one or more qualifiers. There are three different kinds of qualifiers:

TABLE A-1 pfTop Qualifiers

| Qualifier | Meaning |
|-----------|--|
| type | Specify which kind of address component the <i>id</i> name or number refers to. Possible types are <i>host</i> , <i>net</i> and <i>port</i> . If there is no type qualifier, <i>host</i> is assumed. |
| dir | Specify the address component (<i>src</i> , <i>dest</i> , <i>gateway</i>) that <i>id</i> applies. Possible directions are <i>src</i> , <i>dst</i> , <i>gw</i> , <i>src</i> or <i>dst</i> , <i>src</i> and <i>dst</i> . If there is no dir qualifier, <i>src</i> or <i>dst</i> or <i>gw</i> is assumed. |
| proto | Restrict the match to a particular protocol. Possible protocols are: <i>ah</i> , <i>carp</i> , <i>esp</i> , <i>icmp</i> , <i>ip</i> , <i>ip6</i> , <i>pfsync</i> , <i>tcp</i> , and <i>udp</i> . If there is no protocol qualifier, all protocols consistent with the type are assumed. |

In addition to the above, there are some special primitive keywords that don't follow the pattern and arithmetic expressions. All of these are described below. More complex filter expressions are built up by using the words *and*, *or*, and *not* to combine primitives. Allowable primitives are:

TABLE A-2 pfTop Primitives

| Primitive | Meaning |
|----------------------|---|
| dst host <i>host</i> | True if the IP destination field of the state is <i>host</i> , which may be either an address or a name. |
| gw host <i>host</i> | True if the IP gateway field of the state is <i>host</i> . |
| src host <i>host</i> | True if the IP source field of the state is <i>host</i> . |
| host <i>host</i> | True if either the IP source or destination or gateway of the state is <i>host</i> . If <i>host</i> is a name with multiple IP addresses, each address will be checked for a match. |
| dst net <i>net</i> | True if the IP destination address of the state has a network number of <i>net</i> . |
| gw net <i>net</i> | True if the IP gateway address of the state has a network number of <i>net</i> . |
| src net <i>net</i> | True if the IP source address of the state has a network number of <i>net</i> . |

TABLE A-2 pfTop Primitives (Continued)

| Primitive | Meaning |
|--|---|
| net <i>net</i> | True if either the IP source, destination or gateway address of the state has a network number of <i>net</i> . Any of the above host or net expressions can be prepended with the keywords, <i>ip</i> , or <i>ip6</i> . |
| dst port <i>port</i> | True if the packet is IP/TCP or IP/UDP and has a destination port value of <i>port</i> . The <i>port</i> can be a number or name number or ambiguous name is used, only the port number is checked; |
| port <i>port</i> | True if either the source, destination or gateway port of the state is <i>port</i> . Any of the above port expressions can be prepended with the keywords <i>tcp</i> or <i>udp</i> , as in: <i>tcp src port port</i> which matches only TCP states whose source port is <i>port</i> . |
| inbound, in | True if the state has an inbound direction. |
| outbound, out | True if the state has an outbound direction. |
| proto <i>proto</i> | True if the IP protocol type of the state is <i>proto</i> . <i>proto</i> can be a number or name, such as <i>icmp</i> , <i>udp</i> , or <i>tcp</i> . |
| rnr <i>num</i> | True if the state was generated with the rule number <i>num</i> in the main ruleset. |
| ah, carp esp, icmp, pfsync, tcp, udp <i>expr relop expr</i> | Abbreviations for: proto <i>p</i> where <i>p</i> is one of the above protocols. True if the relation holds, where <i>relop</i> is one of '>', '<', '>=', '<=', '=', '!=', and <i>expr</i> is an arithmetic expression composed of integer constants (expressed in standard C syntax), the normal binary operators ('+', '-', '*', '/', '&', ' '), a length operator, and special state data accessors. The following expressions can be used to access numerical fields inside a state: <i>inp</i> , and <i>outp</i> return input and output packet counts. <i>inb</i> , and <i>outb</i> is for input and output bytes transferred through the state. <i>age</i> is the seconds since the state is created, and <i>exp</i> is the number of seconds left before the state expires. |

Primitives may be combined using a parenthesized group of primitives and operators. Allowable primitives and operators are:

TABLE A-3 pfTop Primitives and Operators for Combining

| Primitives and Operators |
|-------------------------------|
| Negation ('!' or 'not') |
| Concatenation ('&&' or 'and') |
| Alternation (' ' or 'or') |

Negation has highest precedence. Alternation and concatenation have equal precedence and associate left to right. Expression arguments must be passed to pftop as a single argument. Since the expression usually contains shell metacharacters, it should be placed in quotes.

Netmask/CIDR Translation Table

| | | | |
|-----------------|-------------------------------------|-----|-----------------------|
| 255.255.255.255 | 11111111.11111111.11111111.11111111 | /32 | Host (single address) |
| 255.255.255.254 | 11111111.11111111.11111111.11111110 | /31 | Unuseable |
| 255.255.255.252 | 11111111.11111111.11111111.11111100 | /30 | 2 useable |
| 255.255.255.248 | 11111111.11111111.11111111.11111000 | /29 | 6 useable |
| 255.255.255.240 | 11111111.11111111.11111111.11110000 | /28 | 14 useable |
| 255.255.255.224 | 11111111.11111111.11111111.11100000 | /27 | 30 useable |
| 255.255.255.192 | 11111111.11111111.11111111.11000000 | /26 | 62 useable |
| 255.255.255.128 | 11111111.11111111.11111111.10000000 | /25 | 126 useable |
| 255.255.255.0 | 11111111.11111111.11111111.00000000 | /24 | "Class C" 254 useable |
| | | | |
| 255.255.254.0 | 11111111.11111111.11111110.00000000 | /23 | 2 Class C's |
| 255.255.252.0 | 11111111.11111111.11111100.00000000 | /22 | 4 Class C's |
| 255.255.248.0 | 11111111.11111111.11111000.00000000 | /21 | 8 Class C's |
| 255.255.240.0 | 11111111.11111111.11110000.00000000 | /20 | 16 Class C's |
| 255.255.224.0 | 11111111.11111111.11100000.00000000 | /19 | 32 Class C's |
| 255.255.192.0 | 11111111.11111111.11000000.00000000 | /18 | 64 Class C's |
| 255.255.128.0 | 11111111.11111111.10000000.00000000 | /17 | 128 Class C's |
| 255.255.0.0 | 11111111.11111111.00000000.00000000 | /16 | "Class B" |
| | | | |
| 255.254.0.0 | 11111111.11111110.00000000.00000000 | /15 | 2 Class B's |
| 255.252.0.0 | 11111111.11111100.00000000.00000000 | /14 | 4 Class B's |
| 255.248.0.0 | 11111111.11111000.00000000.00000000 | /13 | 8 Class B's |
| 255.240.0.0 | 11111111.11110000.00000000.00000000 | /12 | 16 Class B's |
| 255.224.0.0 | 11111111.11100000.00000000.00000000 | /11 | 32 Class B's |
| 255.192.0.0 | 11111111.11000000.00000000.00000000 | /10 | 64 Class B's |
| 255.128.0.0 | 11111111.10000000.00000000.00000000 | /9 | 128 Class B's |
| 255.0.0.0 | 11111111.00000000.00000000.00000000 | /8 | "Class A" |
| | | | |
| 254.0.0.0 | 11111110.00000000.00000000.00000000 | /7 | |
| 252.0.0.0 | 11111100.00000000.00000000.00000000 | /6 | |
| 248.0.0.0 | 11111000.00000000.00000000.00000000 | /5 | |
| 240.0.0.0 | 11110000.00000000.00000000.00000000 | /4 | |
| 224.0.0.0 | 11100000.00000000.00000000.00000000 | /3 | |
| 192.0.0.0 | 11000000.00000000.00000000.00000000 | /2 | |
| 128.0.0.0 | 10000000.00000000.00000000.00000000 | /1 | |
| 0.0.0.0 | 00000000.00000000.00000000.00000000 | /0 | IPv4 space |

Factory Default Settings

The table below provides factory default settings for Sentinel PF operating system.

TABLE C-1 Factory Default Settings

| Setting | Value |
|-----------------------------|--|
| WAN IP Address | Automatic via DHCP |
| LAN IP Address and Subnet | 192.168.2.1/24 (netmask 255.255.255.0) |
| Login | admin |
| Password | sentinel pf |
| Web GUI Default Port | 8765 |
| SSH CLI Default Port | 22 |
| Package Repository Location | sentinelpf.bcsatellite.net |

Please note: default password has space character.

Notes:

Default password has a space character.

Sentinel PF has DHCP service enabled on LAN interface by default.

NOTES:

Additional Support

The following table provides an overview of additional interfaces and their features supported on various Sentinel hardware servers.

Optional support means none of the interfaces are included by default, however can be installed as an option. *Optional Ethernet Interfaces* show OPT Ethernet interfaces available in addition to the standard WAN and LAN interfaces. *Sentinel PF Version* indicates the first Sentinel PF operating system software version that announced full support for the server hardware.

TABLE D-1 Additional Interfaces and Features Supported

| Server | Sentinel PF Version | Optional Ethernet Interfaces | 802.1Q VLAN | 802.1 QinQ | Wireless Interfaces | Serial PPP | 802.3ad LACP |
|--------------------|---------------------|------------------------------|-------------|------------|---------------------|------------|--------------|
| Sentinel Blackbird | 2.0.2 | 1 | Yes | Yes | Optional | Optional | Yes |
| Sentinel Westwind | 2.0.2 | 3 | Yes | Yes | Optional | Optional | Yes |
| Sentinel Sierra | 2.1 | 0 | Yes | No | Optional | Optional | No |
| Sentinel Cirrus | 2.1 | 0 | Yes | No | Optional | Optional | No |
| Sentinel D2 | 2.0.2 Nano | 1 | Yes | No | Optional | No | No |
| Sentinel D2W | 2.0.2 Nano | 1 | Yes | No | Yes | No | No |
| Sentinel 3 | 2.1.5 Nano | 1 | Yes | No | Yes | No | No |

Note: At least 2 optional Ethernet interfaces are required for redundant CARP/Multi-WAN environment.

DiffServ Classification & Marking

Network traffic entering a DiffServ domain is subjected to classification and conditioning. Traffic may be classified by many different parameters, such as source address, destination address or traffic type and assigned to a specific traffic class. Traffic classifiers may honor any DiffServ markings in received packets or may elect to ignore or override those markings.

The Per-Hop Behavior is determined by the DS field of the IP header. The DS field contains a 6-bit Differentiated Services Code Point (DSCP) value. Most networks use the following commonly-defined Per-Hop Behaviors:

TABLE E-1 DSCP PHBs

| DSCP | Behavior |
|---------------------------|---|
| Default | Best-effort traffic |
| EF (Expedited Forwarding) | Low-loss, low-latency traffic |
| AF (Assured Forwarding) | Delivery assurance under prescribed conditions |
| Class Sector PHBs | Maintain backward compatibility with IP Precedence field. |

The IETF defines Expedited Forwarding behavior in *RFC 3246*. The EF PHB has the characteristics of low delay, low loss and low jitter. These characteristics are suitable for voice, video and other realtime services. EF traffic is often given strict priority queuing above all other traffic classes. Because an overload of EF traffic will cause queuing delays and affect the jitter and delay tolerances within the class, EF traffic is often strictly controlled through admission control, policing and other mechanisms. The recommended DSCP for expedited forwarding is 101110_B (46 or $2E_H$). The EF traffic can be matched in Sentinel PF operating system by selecting *EF* as the DiffServ Code Point.

The IETF defines the Assured Forwarding behavior in *RFC 2597* and *RFC 3260*. Assured forwarding allows the operator to provide assurance of delivery as long as the traffic does not exceed some subscribed rate. Traffic that exceeds the subscription rate faces a higher probability of being dropped if congestion occurs. The AF behavior group defines four separate AF classes with Class 4 having the highest priority. Within each class, packets are given a drop precedence (high, medium or low).

The combination of classes and drop precedence yields twelve separate DSCP encodings from AF11 through AF43 (see table below).

TABLE E-2 Assured Forwarding DSCP Encodings

| Precedence | Class 1 (Lowest) | Class 2 | Class 3 | Class 4 (Highest) |
|------------------|---------------------|---------|---------|----------------------|
| Low Drop Rate | AF11 | AF21 | AF31 | AF41 |
| Medium Drop Rate | AF12 | AF22 | AF32 | AF42 |
| High Drop Rate | AF13 | AF23 | AF33 | AF43 |

Some measure of priority and proportional fairness is defined between traffic in different classes. Should congestion occur between classes, the idea is that traffic in the higher class is given priority. Rather than using strict priority queuing, more balanced queue servicing algorithms such as fair queuing are likely to be used. If congestion occurs within a class, the packets with the higher drop precedence are supposed to be discarded first. To prevent issues associated with tail drop, more sophisticated drop selection algorithms such as random early detection (RED) are often used. The AF traffic can be matched in Sentinel PF operating system by selecting *AFxx* as the DiffServ Code Point where *xx* is the class and drop rate precedence.

For additional information on DiffServ, please consider *DiffServ RFCs* as per the table below:

TABLE E-3 DiffServ RFCs

| RFP Number and Title |
|--|
| RFC 2474—Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers |
| RFC 2475—An Architecture for Differentiated Services |
| RFC 2597—Assured Forwarding PHB Group |
| RFC 2983—Differentiated Services and Tunnels |
| RFC 3086—Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification |
| RFC 3140—Per Hop Behavior Identification Codes (Obsoletes RFC 2836) |
| RFC 3246—An Expedited Forwarding PHB (Obsoletes RFC 2598) |
| RFC 3247—Supplemental Information for the New Definition of the EF PHB (Expedited Forwarding Per-Hop Behavior) |
| RFC 3260—New Terminology and Clarifications for Diffserv (Updates RFC 2474, RFC 2475 and RFC 2597) |
| RFC 4594—Configuration Guidelines for DiffServ Service Classes |
| RFC 5865—A Differentiated Services Code Point (DSCP) for Capacity-Admitted Traffic (Updates RFC 4542 and RFC 4594) |
| RFC 3289—Management Information Base for the Differentiated Services Architecture |
| RFC 3290—An Informal Management Model for Diffserv Routers |
| RFC 3317—Differentiated Services Quality of Service Policy Information Base |

NIDS Incidents Classification

The following table describes Sentinel PF NIDS network incidents and malicious activity classification, and default priorities assigned to these events. Typical recommended course of action is provided, however this does not cover all of the system administrator's actions required in case of such an event.

TABLE F-1 Sentinel PF NIDS Incidents Classification

| Class | Default Priority | Description | Typical Recommended Course of Action |
|--|------------------|---|--|
| Attempted Administrator Privilege Gain | 1 (High) | Attempts to obtain administrator, superuser or <i>root</i> privileges for software (or hardware) in your network. These alarms do not confirm successful privilege escalation, rather just signal the attempt took place. An access to the <i>ADMIN\$</i> share on a Windows server is an example of the privilege gain attempt. | In case traffic comes from an unknown source, use firewall to restrict access to the software (or hardware) to authorized sources only. Optionally, move sensitive daemons like <i>sshd</i> to non-standard ports. |
| Attempted User Privilege Gain | 1 (High) | Attempts to obtain user level access to software (or hardware) in your network that do not fit common usage. These alarms do not confirm successful privilege escalation, rather just signal the attempt took place. A modification of the <i>.rhosts</i> file to bypass control access is an example of the privilege gain attempt. | In case traffic comes from an unknown source, use firewall to restrict access to the software (or hardware) to authorized sources only. Optionally, move sensitive daemons like <i>sshd</i> to non-standard ports. |
| Inappropriate Content was Detected | 1 (High) | A content is passing through your network that may be inappropriate in some environments or jurisdictions. This may include sexually explicit material, web sites related to dangerous and illegal activities, hate crimes, hacking and so on. | Depending on your corporate policy, this may range from disciplinary action to installing an external content filter to prevent access to inappropriate resources. |
| Potential Corporate Policy Violation | 1 (High) | Applications and/or content is passing through your network that may violate your corporate policies. This class includes FTP and SSH login attempts, VNC servers and similar remote administration tools, illegal software downloads, Peer to Peer (P2P) and TOR traffic, and similar. Many of the traffic with this classification may be legitimate. | Use firewall to block/reject traffic for applications that violate your corporate policy. |

TABLE F-1 Sentinel PF NIDS Incidents Classification (Continued)

| Class | Default Priority | Description | General Recommended Course of Action |
|---|------------------|--|---|
| Executable code was detected | 1 (High) | Attempts to inject executable code (<i>shellcode</i>) into hosts on your network. Shellcodes are used as one of the final attack stages to gain administrator or user level access to protected systems. This classification does not confirm the access has been obtained. | Treat the target host in your network as potentially compromised. Immediately use firewall to restrict access to authorized sources only. Check the target host for potential rootkits or other signs of exploited vulnerabilities. Optionally, move sensitive daemons like <i>sshd</i> to non-standard ports. |
| Successful Administrator Privilege Gain | 1 (High) | Alert after an attacker has successfully obtained administrator, superuser or <i>root</i> privileges for software (or hardware) in your network. The rules in this classification detect signatures of completed attacks and suspicious file transfers that confirm the event. For example, rules may include traffic patterns produced by a known <i>telnet</i> daemon exploit. | Treat the target host in your network as most probably compromised. Shut down the target host immediately. Check for potential rootkits or other signs of exploited vulnerabilities. Optionally, move sensitive daemons like <i>sshd</i> to non-standard ports. Use firewall to restrict access to authorized sources only. |
| A Network Trojan was detected | 1 (High) | A network trojan traffic has been detected in your network. | Shut down the target (or source) host in your local network. Use antivirus software to inspect operating system for trojans. |
| Unsuccessful User Privilege Gain | 1 (High) | Attempts to obtain user level access to software (or hardware) in your network that do not fit common usage, without success. This may be attributed to <i>brute force</i> attacks. | Use firewall to restrict access to the software (or hardware) to authorized sources only. Optionally, move sensitive daemons like <i>sshd</i> to non-standard ports. |
| Web Application Attack | 1 (High) | Attempts of attacks on Web applications in your network. | Use firewall or VPN to restrict access to the software (or hardware) to authorized sources only. Check your Web application for potential vulnerabilities. Optionally, move corporate Web applications to non-standard ports. |
| Attempted Denial of Service | 2 (Medium) | Attempts of Denial of Service attacks on software (or hardware) in your network. | Use firewall to restrict access to the software (or hardware) to authorized sources only. Contact your ISP to block the attack upstream. |

TABLE F-1 Sentinel PF NIDS Incidents Classification (Continued)

| Class | Default Priority | Description | General Recommended Course of Action |
|---|------------------|---|---|
| Attempted Information Leak | 2 (Medium) | Attempts of information gathering on your network that can potentially be used to plan attacks. This does not confirm information has been released. This class includes unusually fast Web application scans, SNMP access, security vulnerability scans, automatic injection tools, username scans and similar recon activities. | Use firewall to restrict access to the software (or hardware) to authorized sources only. |
| Potentially Bad Traffic | 2 (Medium) | This class encompasses traffic that may potentially indicate a compromised system in your network. An example event in this class could be raw local storage access via Web application. | If attack comes from unknown source, treat the target host in your network as potentially compromised. Immediately use firewall to restrict access to authorized sources only. Check the target host for potential rootkits or other signs of exploited vulnerabilities. |
| Attempt to login by a default username and password | 2 (Medium) | Attempts to login to hardware (or software) in your network using default credentials. This may include DSL or WiFi routers that are often left with factory default access settings. | If all default credentials are disabled everywhere in your LAN, this may be safely ignored, however recommend using firewall to restrict access to the software (or hardware) to authorized sources only. |
| Detection of a Denial of Service Attack | 2 (Medium) | Denial of Service attacks on software (or hardware) in your network. | Use firewall to restrict access to the software (or hardware) to authorized sources only. Contact your ISP to block the attack upstream. |
| Misc Attack | 2 (Medium) | Miscellaneous attacks on software (or hardware) in your network. This includes uncategorized «exotic» attacks, such as TTL exceeded in transit attacks, NNTP overflow attempts, etc. | Use firewall to restrict access to the software (or hardware) to authorized sources only. Contact your ISP to block the attack upstream in case of high traffic volumes. |
| Detection of a non-standard protocol or event | 2 (Medium) | Detection of odd traffic in your network that does not conform to protocol standards. This may include backdoors. | If attack comes from unknown source, treat the target host in your network as potentially compromised. Use firewall to restrict access to authorized sources only. Check the target host for potential rootkits or other signs of exploited vulnerabilities at earliest convenience, especially with persistent alerts. |

TABLE F-1 Sentinel PF NIDS Incidents Classification (Continued)

| Class | Default Priority | Description | General Recommended Course of Action |
|---|------------------|---|--|
| Decode of an RPC Query | 2 (Medium) | Possible attempts to exploit RPC (Sun Remote Procedure Call) vulnerabilities. | If attack comes from unknown source, treat the target host in your network as potentially compromised. Use firewall to restrict access to authorized sources only. Check the target host for potential rootkits or other signs of exploited vulnerabilities. |
| Denial of Service | 2 (Medium) | Denial of Service attacks on software (or hardware) in your network. | Use firewall to restrict access to the software (or hardware) to authorized sources only. Contact your ISP to block the attack upstream. |
| Large Scale Information Leak | 2 (Medium) | Successful information gathering activity on your network that can potentially be used to plan attacks. This class includes unusually fast Web application scans, SNMP access, security vulnerability scans, automatic injection tools, username scans and similar recon activities. | Use firewall to restrict access to the software (or hardware) to authorized sources only. Depending on the type of the leak, you may want to change credentials on all critical systems. |
| Information Leak | 2 (Medium) | See above. | See above. |
| A suspicious filename was detected | 2 (Medium) | This class alerts suspicious file transfer activity in your network that may be attributed to compromised systems. As an example, this could be an attempt to use FTP to retrieve <i>passwd</i> and <i>shadow</i> files that hold superuser and root level details to access hosts on your network. This classification does not confirm any systems have been compromised or information was leaked. | Use firewall to restrict access to authorized sources only. Check the target host for potential rootkits or other signs of exploited vulnerabilities. Optionally, change credentials and move sensitive daemons like <i>sshd</i> to non-standard ports. |
| An attempted login using suspicious username was detected | 2 (Medium) | This class alerts suspicious login activity in your network that may be attributed to compromised systems. As an example, this could be an attempt to log in as root via telnet session. This classification does not confirm any systems have been compromised or information was leaked. | Use firewall to restrict access to authorized sources only. Check the target host for potential rootkits or other signs of exploited vulnerabilities. Change credentials. Optionally, move sensitive daemons like <i>sshd</i> to non-standard ports. |

TABLE F-1 Sentinel PF NIDS Incidents Classification (Continued)

| Class | Default Priority | Description | General Recommended Course of Action |
|--|------------------|---|--|
| A system call was detected | 2 (Medium) | Attempts to use system calls to OS kernels that may be used to inject executable code (shellcode) into hosts on your network. Shellcodes are used as one of the final attack stages to gain administrator or user level access to protected systems. This classification does not confirm the access has been obtained. | Use firewall to restrict access to authorized sources only. Check the target host for potential rootkits or other signs of exploited vulnerabilities. Update/patch operating system software on a target system. |
| A client was using an unusual port | 2 (Medium) | Detection of standard traffic on non-standard ports in your network. This may include backdoors or possible corporate policy violation. | Use firewall to restrict access to authorized sources only. Check the target host for potential rootkits or other signs of exploited vulnerabilities. Run anti-virus check on the target host. |
| Access to a potentially vulnerable Web application | 2 (Medium) | A potentially vulnerable Web applications is accessed in your network. | Use firewall or VPN to restrict access to Web application to authorized sources only. Check your application for potential vulnerabilities. Update/patch software on a target system. |
| Generic ICMP Event | 3 (Low) | Generic ICMP activity. | Ignore unless refers to critical service in your network. |
| Misc activity | 3 (Low) | Miscellaneous low level protocol activity without any particular category. (See alert description for additional information.) | Ignore, unless persistent. Otherwise, use firewall to block source. |
| Network Scan | 3 (Low) | Potential network scans for open ports and active IPs in your network. | Ignore, unless persistent. False positives are common. Use firewall to block source in case of frequent alerts. Run anti-virus check on the source host. |
| Not Suspicious Traffic <i>or</i> Unknown Traffic | 3 (Low) | In most environments, this class refers to traffic that is generally not suspicious, such as telnet access or authentication failures. Some restricted environments may require this traffic to be handled differently. | Ignore, unless your corporate policy requires otherwise. |
| Generic Protocol Command Decode | 3 (Low) | Odd activity with generic protocols. (See alert description for additional information.) | Ignore, unless refers to critical service in your network. |
| A suspicious string was detected | 3 (Low) | Potential malicious activity involving suspicious strings – e.g. domain names, user agents, and similar. | Ignore, unless persistent. Otherwise, use firewall to block this traffic. |

Sentinel D2 RS-232 Console Cable

The following diagram provides wiring for RS-232 Console Cable used to connect to Serial D2 console port. This assumes DB9 male connector is available on the client PC.

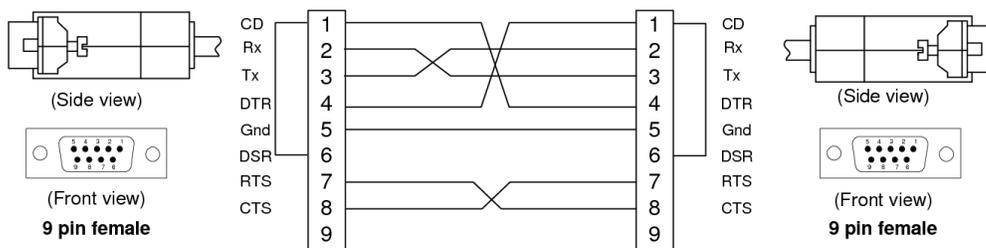


FIGURE G-1 Sentinel D2 RS-232 Console Cable.

Notes:

Although only Rx, Tx and Gnd connections are used by Sentinel D2 (no flow control), standard null modem cable is described above for future compatibility purposes.

In case a client PC does not have serial interface, you can use RS-232 to USB adapters. Depending on the configuration, some RS-232 to USB adapters may require an additional null modem adapter.

Troubleshooting

The following table describes some Sentinel PF incidents. Recommended course of action is provided.

Boot

TABLE H-1 Sentinel PF Boot Troubleshooting

| Symptom | Applicable Platforms | Recommended Course of Action |
|---|----------------------|--|
| Sentinel unresponsive to connections. Each LED on the front panel is flashing twice in an infinite loop after power on. | D2 and D2W | Sentinel PF operating system is waiting for user input. This could be due to interface assignment request, or a hardware error. Use RS-232 cable to connect via console to provide the input required. In case of unknown error, reboot Sentinel by powering off and on and provide console output to BusinessCom Customer Service Department. |

Run-Time

TABLE H-2 Sentinel PF Run-Time Troubleshooting

| Symptom | Applicable Platforms | Recommended Course of Action |
|---|----------------------|---|
| Intermittent web browsing performance. Sentinel is close to 100% of RAM usage after some time. | All | If proxy server is enabled, check if the service RAM setting is configured correctly. Please refer to <i>Chapter 8 - Web GUI: Services, Proxy Server</i> section, <i>Memory cache size</i> setting in cache management. Check PEP and Proxy services logging is disabled on Flash based Sentinel models such as D2 and D2W. |
| <i>Status: System logs: kernel: ath0: stuck beacon; resetting (bmiss count 4) messages appear</i> | Sentinel 3 | This can be safely ignored. Usually appears in noisy wireless environments. |

Services

TABLE H-3 Sentinel PF Services Troubleshooting

| Symptom | Applicable Platforms | Recommended Course of Action |
|--|----------------------|---|
| Windows clients can't join wireless Access Point on Sentinel. <i>Status: System logs: Wireless show WPA: invalid MIC in msg 2/4 of 4-way handshake error messages.</i> | All | <p>Wireless handshake procedures in different operating systems are not always compatible with each other. Here are recommended settings for the Sentinel wireless interfaces running in the Access Point mode, and Windows clients:</p> <p>WPA Mode: <i>WPA</i> WPA Key Management Mode: <i>Pre Shared Key</i> Authentication: <i>Open System Authentication</i> WPA Pairwise: <i>Both</i>.</p> |

BUSINESSCOM[®]
NETWORKS

BusinessCom Networks Limited
Glacis Road, Portland House, Suite 2
Gibraltar, GX11 1AA