

Firewall Capabilities

- Filtering by source and destination IP, IP protocol, source and destination port for TCP and UDP traffic
- Limit simultaneous connections on a per-rule basis
- pfSense software utilizes pf, [an advanced passive OS/network fingerprinting utility](#) to allow you to filter by the Operating System initiating the connection. Want to allow FreeBSD and Linux machines to the Internet, but block Windows machines? pfSense software allows for that (amongst many other possibilities) by passively detecting the Operating System in use.
- Option to log or not log traffic matching each rule.
- Highly flexible policy routing possible by selecting gateway on a per-rule basis (for load balancing, failover, multiple WAN, etc.)
- Aliases allow grouping and naming of IPs, networks and ports. This helps keep your firewall ruleset clean and easy to understand, especially in environments with multiple public IPs and numerous servers.
- Transparent layer 2 firewalling capable - can bridge interfaces and filter traffic between them, even allowing for an IP-less firewall (though you probably want an IP for management purposes).
- Packet normalization - Description from the pf scrub documentation - "'Scrubbing' is the normalization of packets so there are no ambiguities in interpretation by the ultimate destination of the packet. The scrub directive also reassembles fragmented packets, protecting some operating systems from some forms of attack, and drops TCP packets that have invalid flag combinations."
 - Enabled in the pfSense software by default
 - Can disable if necessary. This option causes problems for some NFS implementations, but is safe and should be left enabled on most installations.
- Disable filter - you can turn off the firewall filter entirely if you wish to turn your pfSense software into a pure router.

State Table

- The firewall's state table maintains information on your open network connections. The pfSense software is a stateful firewall, by default all rules are stateful.
- Most firewalls lack the ability to finely control your state table. The pfSense software has numerous features allowing granular control of your state table, thanks to the abilities of FreeBSD's ported version of pf.
- Adjustable state table size - there are multiple production pfSense installations using several hundred thousand states. The default state table size varies according to the RAM installed in the system, but it can be increased on the fly to your desired size. Each state takes approximately 1 KB of RAM, so keep in mind memory usage when sizing your state table. Do not set it arbitrarily high.
- On a per-rule basis:
 - Limit simultaneous client connections
 - Limit states per host
 - Limit new connections per second
 - Define state timeout
 - Define state type
- State types - the pfSense software offers multiple options for state handling.
- Keep state - Works with all protocols. Default for all rules.
- Sloppy state - Works with all protocols. Less strict state tracking, useful in cases of asymmetric routing.
- Synproxy state - Proxies incoming TCP connections to help protect servers from spoofed TCP SYN floods. This option includes the functionality of keep state and modulate state combined.

- None - Do not keep any state entries for this traffic. This is very rarely desirable, but is available because it can be useful under some limited circumstances.
- State table optimization options - pf offers four options for state table optimization.
- Normal - the default algorithm
- High latency - Useful for high latency links, such as satellite connections. Expires idle connections later than normal.
- Aggressive - Expires idle connections more quickly. More efficient use of hardware resources, but can drop legitimate connections.
- Conservative - Tries to avoid dropping legitimate connections at the expense of increased memory usage and CPU utilization.

Network Address Translation (NAT)

- Port forwards including ranges and the use of multiple public IPs
- 1:1 NAT for individual IPs or entire subnets.
- Outbound NAT
- Default settings NAT all outbound traffic to the WAN IP. In multiple WAN scenarios, the default settings NAT outbound traffic to the IP of the WAN interface being used.
- Advanced Outbound NAT allows this default behavior to be disabled, and enables the creation of very flexible NAT (or no NAT) rules.
- NAT Reflection - NAT reflection is possible so services can be accessed by public IP from internal networks.
-
- **Limitations: PPTP / GRE Limitation** - The state tracking code in pf for the GRE protocol can only track a single session per public IP per external server. This means if you use PPTP VPN connections, only one internal machine can connect simultaneously to a PPTP server on the Internet. A thousand machines can connect simultaneously to a thousand different PPTP servers, but only one simultaneously to a single server. The only available work around is to use multiple public IPs on your firewall, one per client, or to use multiple public IPs on the external PPTP server. This is not a problem with other types of VPN connections. PPTP is insecure and should no longer be used.

High Availability

- The combination of CARP, pfsync, and our configuration synchronization provides high availability functionality. Two or more firewalls can be configured as a failover group. If one interface fails on the primary or the primary goes offline entirely, the secondary becomes active. The pfSense software also includes configuration synchronization capabilities, so you make your configuration changes on the primary and they automatically synchronize to the secondary firewall.
- The firewall's state table is replicated to all failover configured firewalls. This means your existing connections will be maintained in the case of failure, which is important to prevent network disruptions.

Limitations: Only works with static public IPs, does not work with stateful failover using DHCP, PPPoE, or PPTP type WANs.

Multi-WAN functionality enables the use of multiple Internet connections, with load balancing and/or failover, Server Load Balancing

Server load balancing is used to distribute load between multiple servers. This is commonly used with web servers, mail servers, and others. Servers that fail to respond to ping requests or TCP port connections are removed from the pool.

- for improved Internet availability and bandwidth usage distribution.
-

Virtual Private Network (VPN)

The pfSense software offers three options for VPN connectivity, IPsec and OpenVPN.

IPsec

IPsec allows connectivity with any device supporting standard IPsec. This is most commonly used for site to site connectivity to other pfSense installations and most all other firewall solutions (Cisco, Juniper, etc.). It can also be used for mobile client connectivity.

OpenVPN

OpenVPN is a flexible, powerful SSL VPN solution supporting a wide range of client operating systems.

PPPoE Server

The pfSense software offers a PPPoE server. A local user database can be used for authentication, and RADIUS authentication with optional accounting is also supported.

Reporting and Monitoring

RRD Graphs

The RRD graphs in the pfSense software maintain historical information on the following.

- CPU utilization
- Total throughput
- Firewall states
- Individual throughput for all interfaces
- Packets per second rates for all interfaces
- WAN interface gateway(s) ping response times
- Traffic shaper queues on systems with traffic shaping enabled

Real Time Information

- Historical information is important, but sometimes it's more important to see real time information.
- SVG graphs are available that show real time throughput for each interface.
- For traffic shaper users, the Status -> Queues screen provides a real time display of queue usage using AJAX updated gauges.
- The front page includes AJAX gauges for display of real time CPU, memory, swap and disk usage, and state table size.

Dynamic DNS

- A Dynamic DNS client is included to allow you to register your public IP with a number of dynamic DNS service providers.
- Custom - allowing defining update method for providers not specifically listed here.
- DNS-O-Matic
- DynDNS
- DHS
- DNSexit
- DyNS
- easyDNS
- freeDNS
- HE.net
- Loopia
- Namecheap
- No-IP
- ODS.org
- OpenDNS
- Route 53
- SelfHost
- ZoneEdit
- A client is also available for RFC 2136 dynamic DNS updates, for use with DNS servers like BIND which support this means of updating.

Captive Portal

- Captive portal allows you to force authentication, or redirection to a click through page for network access. This is commonly used on hot spot networks, but is also widely used in corporate networks for an additional layer of security on wireless or Internet access. For more information on captive portal technology in general. The following is a list of features in the pfSense Captive Portal:
- Maximum concurrent connections - Limit the number of connections to the portal itself per client IP. This feature prevents a denial of service from client PCs sending network traffic repeatedly without authenticating or clicking through the splash page.
- Idle timeout - Disconnect clients who are idle for more than the defined number of minutes.
- Hard timeout - Force a disconnect of all clients after the defined number of minutes.
- Logon pop up window - Option to pop up a window with a log off button.
- URL Redirection - after authenticating or clicking through the captive portal, users can be forcefully redirected to the defined URL.
- MAC filtering - by default, pfSense filters using MAC addresses. If you have a subnet behind a router on a captive portal enabled interface, every machine behind the router will be authorized after one user is authorized. MAC filtering can be disabled for these scenarios.

- Authentication options - There are three authentication options available.
- No authentication - This means the user just clicks through your portal page without entering credentials.
- Local user manager - A local user database can be configured and used for authentication.
- RADIUS authentication - This is the preferred authentication method for corporate environments and ISPs. It can be used to authenticate from Microsoft Active Directory and numerous other RADIUS servers.
- RADIUS capabilities
- Forced re-authentication
- Able to send Accounting updates
- RADIUS MAC authentication allows captive portal to authenticate to a RADIUS server using the client's MAC address as the user name and password.
- Allows configuration of redundant RADIUS servers.
- HTTP or HTTPS - The portal page can be configured to use either HTTP or HTTPS.
- Pass-through MAC and IP addresses - MAC and IP addresses can be white listed to bypass the portal. Any machines with NAT port forwards will need to be bypassed so the reply traffic does not hit the portal. You may wish to exclude some machines for other reasons.
- File Manager - This allows you to upload images for use in your portal pages.
- Limitations: "Reverse" portal, i.e. capturing traffic originating from the Internet and entering your network, is not possible.
- Only entire IP and MAC addresses can be excluded from the portal, not individual protocols and ports.