



CIPA Filter

10/31/2011 - 7.0

1. Introduction

CIPAFilter is a powerful routing platform capable of delivering an evolving tool set to protect your enterprise. CIPAFilter's philosophy is to provide a cutting edge, well rounded, and aggressive network control solution to meet your current and future needs, commit to keeping your platform up to date, and to provide support far above and beyond simple problem solving. We believe our philosophy is what sets our product firmly apart from our competitors. Our technicians will talk you through installation every step of the way. We will advise you as to which features are a good fit for your organization and help you implement them. We will protect your return on investment by assuring that you can always easily upgrade to our latest platforms, often at no cost. And of course, we will help you solve any problems you encounter with our product, even when it's caused by a third party agent, software package, or network problem.

The CIPAFilter is a powerful solution for many of your networking problems.

- Our anti-virus systems will protect your network by scanning Mail (SMTP, POP3 and IMAP) and Web (HTTP and FTP) protocols.
- The content filter can block pornographic e-mail and web surfing to provide a more professional environment for your students or employees.
- IPSEC VPN tunneling can secure your data as it is transmitted between branch offices.
- Download redirection can save your workstations from spyware.
- Web Usage Reporting, Authentication, Bandwidth Reporting and Bandwidth Control helps you keep your employees on task and stop problems before they become serious.
- Anti-Spam systems can save you hours a day.

And of course, CIPAFilter provides all the back end tools you need to keep your network running smoothly such as DHCP, dynamic and static routing, Firewall, SMTP, Network Access Control and Whitelisting/Blacklisting.

This reference will provide the information necessary for an experienced network engineer to configure and operate these services. However, this manual is not intended to replace our excellent phone support. We will help you use our product every step of the way. Many customers never even log on to our router. Please feel free to call us at 1-800-24DERBY if you have any questions or need advice during the lifetime of your product. There is no need to try to install the device without our help if you don't feel comfortable with these instructions.

2. Table of Contents

Table of Contents

1. Introduction	2	6. Routing	16
2. Table of Contents	4	7. Port Forwarding	17
3. Interface Conventions	6	8. Stateful Firewall and Layer 7 Filtering	19
4. Installation	7	TCP/UDP Firewall	19
As a replacement for your router	8	Example Firewall Configurations	21
As a Bridge	8	ICMP Firewall	23
Between your router and your network	9	9. DHCP Server	24
As a server	10	10. VLAN	25
5. IP Settings	11	11. Clustering	26
Dot Notation to CIDR Notation Translation Table	12	12. User Manager	27
Field Descriptions	13	13. Content Filtering	28
		Authentication	29
		Field Descriptions	29

14. Group Permissions	34	Router-Router VPN	47
Groups	34	21. Network Access Control	49
Technologies	35	22. Anti-Spam	51
Automatic Blacklists	37	23. E-mail Archive	56
Temporary Whitelist Entries	38	24. Client Software	57
Manual Blacklists	38	25. Database Maintenance	58
Blacklist Examples	40	Backup Push	59
Super Whitelist	41	26. Firmware Updates	60
Global Lists	41	Automatic Updates	60
15. Permissions Scheduling	42	Manual Update	60
16. Internet Reports	43	27. Override Console	61
17. Notifications	44	28. Config Save/Restore	62
18. Bandwidth Control	45	29. Troubleshooting	63
19. Bandwidth Reports	46	30. Appendix I: Dot Notation to CIDR Notation Translation Table	64
20. VPN	47	31. Appendix II: POSIX Regular Expressions	65
Client-Server VPN	47		

3. Interface Conventions

If an option name in the web based user interface is highlighted with a blue link, that indicates the link will bring up a help document describing the option.

Clicking a **Save Changes** button will cause the configuration options on any particular page to be saved but not activated. **Save and Apply** will cause the configuration options to be saved and activated. If there is no **Save and Apply** option on any given page, that means the router must be restarted to activate the changes.

4. Installation

In most cases you will want to consult with CIPAFilter to decide what way the router can best be installed to meet your needs. A full over the phone consultation during installation is included in the standard one year maintenance and service agreement that comes with your router.

CIPAFilter is usually installed one of four ways, replacing your existing router, as a second router connected to your existing router with a crossover cable, a bridge spliced into the link between your existing router and your switch bank, or as a server on your network.

Of the four ways that CIPAFilter is usually installed, installation as a bridge is the easiest. Installation replacing your existing router is most recommended.

The CIPAFilter ships with 10.1.2.3/8 assigned to its first Ethernet interface and 192.168.0.1/24 assigned to its second Ethernet interface. If your network does not use 10.0.0.0/8 IP addresses you can simply assign an IP like 10.1.2.1/8 (netmask 255.0.0.0) to your workstation and plug the CIPAFilter into your hub or switch.

If your network does use these IP addresses plugging your CIPAFilter in without taking the proper precautions could cause loss of Internet access or other problems. If you're not sure call us for assistance.

Go to <https://10.1.2.3> in your web browser to begin configuration. Default user name is root and default password is derby. The password can be changed on the **User Manager** page.

As a replacement for your router

Often you will find the best way to install CIPAFilter is as a replacement for your current router. CIPAFilter is a top of the line router and implements virtually all the functionality of other routing systems. In this case you would configure CIPAFilter with the IP addresses, routes and NAT settings currently on your router, and then replace your current router.

As a Bridge

The easiest way to install CIPAFilter is as a bridge. Simply unplug your existing router from your switch bank, connect CIPAFilter to that port on your router via a cross over cable, then connect the cable that was plugged into your router to the other port on CIPAFilter. Check **Interface Bridging** on the **IP Settings** page on the web based interface. Your network should now operate normally, and you can activate CIPAFilter features one by one as you're ready to implement them.

Between your router and your network

If you do not wish to replace your current router, you can connect CIPAFilter and your router with a crossover cable, and then the CIPAFilter to the rest of your network. This essentially "splices" the CIPAFilter into the cable running from your router to your network. To do this you must first pick a small subnet from a private range you aren't using. If you're using 10.0.0.0/8 you might pick 192.168.0.0/24, for example. Assuming the default gateway you use for your clients is 10.0.0.1, you would:

1. Replace the 10.0.0.1/8 IP address on your main router with 192.168.0.1/24.
2. Configure your CIPAFilter with 192.168.0.2/24 on the outside interface and 10.0.0.1/8 on the inside interface.
3. Disconnect your main router from the hub, and run a crossover cable from that Ethernet port on your main router to the outside Ethernet port on the CIPAFilter.
4. Connect the inside Ethernet port on your CIPAFilter to the port on your hub or switch where your router was connected.
5. Create a route for 10.0.0.1/8 with a gateway of 192.168.0.2 in your main router, alternatively you can choose to activate NAT on the 10.0.0.0/8 subnet on the CIPAFilter and not create a route. If you choose the NAT option be certain that the 192.168.0.2/24 interface is selected as the **Primary Internet Connection**.

As a server

Many features of the CIPAFilter can be used simply by assigning the CIPAFilter an IP address on your network and connecting it like any other server.

Disadvantages:

- No transparent proxy.
- No POP3 Virus scanning.
- Firewall ineffective.
- You will often need to add firewall rules to your existing router to force users to use the proxy server in order to make CIPAFilter effective at protecting your network.
- Bandwidth control only effective for Internet traffic working through the proxy server.
- Bandwidth usage may not be able to monitor all traffic on your network.

5. IP Settings

Each interface has a primary IP address as well as optional additional secondary IP addresses. An IP address of 0.0.0.0/0 indicates that the interface is to be left unconfigured. Secondary IP addresses can be assigned to an interface by clicking **Add IP** near the bottom of each interface section.

IP addresses are specified using CIDR Notation, an Address/Subnet bits format. The subnet bits refer to the number of bits that are set in binary in the subnet mask from the left hand side. For example, a subnet mask of 255.0.0.0 in binary has 8 bits set from the left hand side. So, 10.0.0.25 with a subnet mask of 255.0.0.0 translates to 10.0.0.25/8 in CIDR Notation. You can use the following table to translate some common subnet masks (dot notation) into CIDR Notation.

Dot Notation to CIDR Notation Translation Table

Subnet Mask	Subnet Bits	IP Addresses
255.0.0.0	/8	16,581,375
255.255.0.0	/16	65,025
255.255.255.0	/24	256
255.255.255.128	/25	128
255.255.255.192	/26	64
255.255.255.224	/27	32
255.255.255.240	/28	16
255.255.255.248	/29	8
255.255.255.252	/30	4
255.255.255.254	/31	2
255.255.255.255	/32	1

NAT (Network Address Translation) is a system whereby private IP addresses (Such as 10.0.0.0/8 and 192.168.0.1/24) are mapped onto a public Internet IP from your ISP. To NAT a subnet of private IPs click the **NAT This Subnet** check box next to the private IP subnet and click the **Primary Internet Connection** radio button on the interface using your ISP assigned public IP. All traffic from the NATed subnet will now be mapped to the primary IP address of the **Primary Internet Connection** device.

Settings like host name and timezone can also be set on this page. If you are using the anti-spam features of CIPAFILTER make sure the host name and domain are correct. Make sure the hostname.domain combination resolves to an IP address on the CIPAFILTER and that the IP address reverses properly back to hostname.domain. This is important because many mail servers will refuse to trade e-mail with your CIPAFILTER if this is not done.

Field Descriptions

Interface Bridging

Clicking this checkbox will cause CIPAFILTER to become a bridge, bridging all it's Ethernet interfaces together into one logical interface. In this configuration CIPAFILTER acts like a switch. Installing CIPAFILTER this way may require the restarting of any CISCO routers on the subnet to clear their ARP cache. Any IP addresses you wish to configure on CIPAFILTER for management can be applied to any interface if this box is checked.

Windows Domain

The Windows Domain is the Windows NT Domain name served by your PDC.

Primary Internet Connection

The interface closest to the Internet should be chosen as your primary Internet connection. This is primarily used if you turn on NAT (Network Address Translation) on any of your other subnets. If a subnet is being NATed, all packets coming from it through the router are translated to appear as if they are coming from the IP you designate as the primary Internet connection.

Network Address Translation

Network address translation (NAT) is used to allow a group of machines with private IP addresses like 10.0.0.0/8 or 192.168.0.0/24 to access the Internet. Packets flowing from any subnet with **NAT This Subnet** checked will have their source address modified to reflect the IP address of your primary Internet connection. Packets coming back to the primary Internet connection will then be sorted and passed to their real destinations. You only have to check **NAT This Subnet** on a single IP from each subnet. For example, if you have five IPs on the 10.0.0.0/8 subnet, you only need to turn NAT on on any one of them to NAT all traffic from the 10.0.0.0/8 subnet.

DB Connect String

CIPAFILTER can use a remote CIPAFILTER as a database server for the purposes of scaling performance. This connect string is the PostgreSQL connect string that instructs CIPAFILTER to access a remote database server and should be configured by tech support.

Enterprise Connect String

This PostgreSQL connect string describes how a CIPAFILTER connects to the CIPAFILTER Enterprise monitoring product if it has been purchased by your organization and should be configured by tech support.

Primary and Secondary DNS

Enter external DNS servers here. We recommend leaving this blank to use CIPAFILTER's internal DNS server.

6. Routing

The CIPAFilter is a fully featured router capable of replacing the functionality of your existing routers.

Most customers will not need to reconfigure any features on this page. CIPAFilter will automatically route between any subnets specified on the **IP Settings** page.

7. Port Forwarding

Port Forwarding is a system by which connections to ports on the router can be forwarded on to servers inside your network that only have private IPs. Often times this feature will be used, for example, to forward RDP connections from your router to a server inside to enable remote management. When a client connects to a particular port on an IP address belonging to your router, the router instead causes it to connect to the **Forward To** address on the second port number entered.

In the example below, if your router has the IP 208.16.191.10, and you would like to enable RDP to a server with the address 10.0.0.5 and a server with the address 10.0.0.6, you would create two port forwards like the first two. This would allow you to connect to 208.16.191.10 with the default port to access the first server, and specify the port 4900 when you want to access the second server. In addition, line 3 demonstrates the use of port ranges to forward the 1000 ports starting at 64000 to 10.0.0.4. The **ALL** keyword in the ports field is used to forward all ports of a particular protocol to the same ports on another machine.

Note: Using the **ALL** keyword in the protocol field activates a 1-to-1 NAT between the **Connections To** IP and the **Forward To** IP. In addition to standard port forwarding (all traffic to the **Connections To** IP will be directed to the **Forward To** IP), with 1-to-1 NAT all outbound traffic from the **Forward To** address will also appear to come from the **Connections To** IP.

Forwarding port 443 or port 22 on your router's only outside IP will make the web interface or CIPAFilter's remote management system unreachable. Attempt to obtain a second IP for your router. If this is not feasible, be sure there is at least a second private IP address on the router, so you can manage the CIPAFilter yourself.

Protocol	Connections To	On Port	Forward To	On Port	Comment
TCP	208.16.191.10	4899	10.0.0.5	4899	RDP to 10.0.0.5
TCP	208.16.191.10	4900	10.0.0.6	4899	RDP to 10.0.0.6
UDP	208.16.191.10	64000:65000	10.0.0.4	64000:65000	1000 UDP ports to 10.0.0.4
TCP	208.16.191.11	ALL	10.0.0.7	ALL	ALL TCP ports to 10.0.0.7
ALL	208.16.191.12	ALL	10.0.0.8	ALL	1 to 1 Nat of 10.0.0.8 to 208.16.191.12

8. Stateful Firewall and Layer 7 Filtering

TCP/UDP Firewall

A stateful firewall is a secure, easy-to-use firewall that tracks all open connections and the state of the connection. A regular firewall only inspects packets.

With a regular firewall, if you wanted your client workstations to be able to access web servers on the Internet you would have to allow any machine on the Internet to access all the high ports on any client. This is a vulnerability many worms and trojans have been written to exploit.

With a stateful firewall, you specify how connections are to be made. For example, you can allow any client inside your network to make a connection to port 80 on any server outside. When that happens, only that server, and only from port 80, will be able to send packets back to the high port on your client that originated the connection.

In addition, CIPAFilter's firewall allows you to match rules on layer 7 (Application layer) protocols like chat, peer-to-peer, or external proxy servers. Eliminating these applications from your network is as easy as selecting **Chat**, **Peer-to-Peer**, or **Proxy Bypass** from the **Protocol** drop down and choosing **Reject** from the **Action** column.

The **Default Policy** of a firewall determines whether it normally drops or accepts connections. CIPAFilter ships with the **Default Policy** set to **ACCEPT**. This mode is acceptable if the CIPAFilter is in bridging mode behind another firewall. However, if the CIPAFilter is the firewall in your network, we recommend setting the **Default Policy** to **DROP**. The drop policy, while giving you specific control of the traffic passing through the firewall, will require the creation of rules for any traffic that you wish to allow through the firewall. Tech support will be happy to assist you with the creation of appropriate rules for your network.

Note: Connections are only matched against the firewall when they are first opened. If you change the firewall, any established connections will remain open even if the new firewall rules prohibit them. Also note that the firewall does not apply to CIPAFilter itself. CIPAFilter automatically adjusts it's own firewall based on the configuration of the system to assure proper operation.

Some Tips:

- ! can be used as "not" on subnets or ports.
- : can be used to denote port ranges, ex: 80:110 for 80-110.
- Accept allows the connection.
- Reject rejects the connection and sends an error message to the client.
- Drop loses all the packets and sends no error message back to the client.
- Do not use any spaces in the rules.
- The firewall is fully integrated with the port forwarding system. Make firewall rules using the actual private IP addresses of your internal servers, not the public IP address from which traffic is being port forwarded.

Example Firewall Configurations

<i>Accept/Reject</i>	<i>Protocol</i>	<i>Connections From</i>	<i>To</i>	<i>On Port</i>	<i>Result</i>
REJECT	TCP	10.0.0.14/32	0.0.0.0/0	80	Workstation with IP address 10.0.0.14 cannot use the web.
REJECT	TCP	64.215.40.0/24	10.0.0.15/32	80	The firewall is fully integrated with the port forwarding system. This rule would reject anybody that tries to connect to port 80 on the internal web server 10.0.0.15 unless the client was on the 64.215.40.0/24 subnet.
REJECT	ALL	10.0.0.25/32	0.0.0.0/0	ALL	Keep 10.0.0.25 from making connections to the outside world.
REJECT	CHAT	10.1.1.0/24	0.0.0.0/0	ALL	Computers in the 10.1.1.0/24 subnet can not use Internet chat programs.

The firewall rules are interpreted top down, so the first rule that matches a connection will determine its fate. For example:

Accept/Reject	Protocol	Connections From	To	On Port	Result
ACCEPT	TCP	208.20.191.0/24	10.0.0.15/32	80	Allows the 208.20.191.0/24 subnet to access the internal web server.
ACCEPT	TCP	64.215.40.0/24	10.0.0.15/32	80	Allows the 64.215.40.0/24 subnet to access the internal web server.
DROP	ALL	0.0.0.0/0	10.0.0.15/32	ALL	Drop all connections to the web server that have not already been accepted.
ACCEPT	Audio/Video	10.0.0.25/32	0.0.0.0/0	ALL	Allow 10.0.0.25 to receive audio and video traffic.
REJECT	Audio/Video	0.0.0.0/0	0.0.0.0/0	ALL	Rejects all other audio/video traffic.

The first two rules will allow machines on the 208.20.191.0/24 subnet and the 64.215.40.0/24 subnet to access the internal web server at 10.0.0.15. However, they will only be able to access the web service on that server. All connections going to all other services or from other subnets will be dropped on the third rule.

The fourth rule allows one IP address to download audio/video data from the Internet. The fifth rule blocks all other IP addresses from doing so.

ICMP Firewall

If you are experiencing a problem with ICMP traffic simply activate the ICMP firewall and check the types of ICMP packets you wish to let through the firewall.

Note: Always allow ICMP fragmentation-needed packets. Those packets are required by MTU Path Discovery. If fragmentation-needed packets are blocked, you may experience problems where you can transmit small amounts of data over a connection but large amounts cause the connection to hang.

9. DHCP Server

The **DHCP Server** is very easy to use. Simply check which interfaces you wish to serve DHCP on and enter a range of IP addresses to serve to your dynamic clients in the **Start IP** and **End IP** boxes. Machines can be assigned static IPs through DHCP by clicking **Add Mapping**. The name field is arbitrary and simply there to help you keep track of entries.

The **DHCP Leases** section displays information about the devices that have received IP addresses since the last time the router was restarted.

10. VLAN

The CIPAFilter can participate in VLANs on your network. Select the interface on the CIPAFilter that participates in the VLAN, enter the **VLAN Tag** and press **Add VLAN**. Any interface references throughout the CIPAFilter interface will reflect the VLAN tags created on this page.

11. Clustering

CIPAFilter is capable of being scaled up to hundreds of megabits of seconds of throughput through the use of our clustering technology. Discussion of this technology however is very complex and beyond the scope of this manual. Please consult with a CIPAFilter sales engineer to discuss clustering options for your environment.

12. User Manager

The **User Manager** allows you to add and remove users for web authentication and VPN services. The UID is the user id of each user. The CIPAFILTER will generate a unique UID for each new user. No two users should have the same user id and no user should have user id 0.

Three special accounts exist: "root", "admin", and "guest". The root user controls the password for the router's web interface. The guest user can access **Bandwidth Reporting** and the online manual. The admin user can access everything guest can, and in addition access **Internet Reporting, E-mail Archive, User Manager, and Group Permissions**.

Note: The admin and guest users cannot access the main page of the router. Consequently, they will require the exact URL link to the page(s) they will access on the CIPAFILTER.

13. Content Filtering

The first thing you must decide is whether or not to run individual subnets in proxy server mode or in transparent mode. With regular proxy services, each client is configured to use the CIPAFilter as a proxy server. For example, in Internet Explorer, the proxy settings can be found under the Internet Options menu. Alternatively, transparent proxy does not require each client to be configured, it simply intercepts all traffic on port 80 as it moves through the router. Transparent mode also supports the ability to implement blacklisting by domain name for HTTPS traffic.

Disadvantages of transparent proxy:

- No FTP support
- No FTP Antivirus
- No FTP Download Redirection
- No DNS lookup caching

In non-transparent mode, the CIPAFilter proxy server runs on port 8080 and supports HTTP, HTTPS, and FTP proxy. Please be advised, that even though CIPAFilter can properly proxy the HTTPS protocol, it is not possible to virus scan or content filter that protocol because it is encrypted.

Authentication

CIPAFILTER supports internal authentication via the **User Manager** and external authentication via LDAP.

CIPAFILTER external authentication systems support Windows Active Directory, Apple Open Directory, and Novell eDirectory. To utilize external authentication, configure your directory server with a user account that has read permissions for all your users and provide the authentication credentials here. In addition to that, a search base and the IP of your LDAP server must be provided.

Groups are defined on the **Group Permissions** page. Each user who logs in will be checked for membership in a group of the same name on the LDAP server. If a membership is found, the user will be configured according to the permissions of that group. If no membership is found, the permissions for the "default" group will be used.

CIPAFILTER checks these group memberships upon each user's first access and cache's the information for up to one hour, depending on protocol. Click **Save and Apply** on the **User Manager** page in the web interface to clear this cache.

Field Descriptions

Authorized Subnets

Only machines with IPs in subnets listed under **Authorized Subnets** will be allowed to use the proxy server. Subnets are in CIDR Notation. If 2 subnets overlap the smallest (most specific) subnet's configuration applies just like in the routing configuration.

Note: The exception to this rule is in the case of transparent proxy. If a less specific subnet has **Transparent Proxy** set to **Yes**, the router will still catch all traffic from the smaller included subnets as well and force them into the proxy server. This can be over-ridden by setting **Transparent Proxy** to **Disable**.

Transparent Proxy (Trans Proxy)

- **Yes+SSL** - Connections from this subnet, including HTTPS connections, will be transparently intercepted by the CIPAFilter.
- **Yes** - Connections from this subnet will be transparently intercepted by the CIPAFilter.
- **No** - Connections from this subnet will not be transparently intercepted by the CIPAFilter unless another subnet in the list intercepts them.
- **Disable** - A specific rule is created to prevent transparent interception of this subnet.

Require Authentication (Require Auth)

Users from this subnet must authenticate via user name and password or by using a transparent authentication method. If the user isn't using a transparent authentication method they will be prompted for a user name and password when they open their web browser. The user name and password is authenticated against the users on the **User Manager** page or the external LDAP authentication method. Each user's web traffic is logged by user name, if web monitoring is active, and can have individual "filtered" or "not filtered" settings.

Clients Use Subnet Group

Checking this box forces this subnet to use the group selected as the **Subnet Group** for the subnet.

Subnet Group

This drop down lists all the groups configured on the **Group Permissions** page. The group selected here will be used if **Require Authentication** is not enabled or **Clients Use Subnet Group** is enabled.

Note: If **Require Authentication** is not selected and an authentication method (see below) is selected, workstations with the CIPAFilter authentication client installed will be authenticated and the users filtered based upon their group membership. Workstations without the CIPAFilter authentication client will be filtered based upon the **Subnet Group** setting for the subnet. This method is referred to as **Optional Authentication**.

Manage Exceptions

When in transparent mode, subnets listed in Transparent Proxy Exceptions are not intercepted by the proxy server.

Authentication Method

This drop down lists all the authentication methods currently supported by CIPAFilter. These authentication methods come in transparent (Client) and non-transparent flavors. Transparent authentication methods do not need to prompt the user for a user name and password, whereas non-transparent do.

Internet Reporting

User activity and filter trips are collected for the **Internet Reporting** system to analyze. If this feature is not checked, CIPAFilter does not record user activity other than by e-mailing the administrator when the content filter is tripped.

Content Filter Reject HTML

When a user trips the blacklist or the content filter they are presented with an error page. This error page makes use of JavaScript variables that are prepended to the output to display details on the trip. You can examine this HTML by causing the content filter to trip, edit it to implement whatever customizations you desire, and then re-upload it here. When you do, do not include the JavaScript variables at the top of the page. Speak with tech support for details.

14. Group Permissions

This page is where you manage the permissions for groups of users. Each group has individual settings for the different filtering technologies available and a separate whitelist/blacklist. On this page, you can also edit the global whitelist/blacklist that applies to all groups.

Groups

Filtering and blacklist configuration is performed by group. The default group is called “default”. Additional groups can be created in the **Group Configuration** area.

When a user is authenticated, the LDAP tree is queried to determine if the user is part of any group matching the names of the groups configured here. If so, the first matching group applies to the user. If not, the default group applies.

There are also a global blacklist and a global whitelist that apply to all groups.

Group Configuration provides the following capabilities: add a new group, change a group name, change group order, and delete the current group. When adding a new group, you can select another similar group as a template to be used when the new group is created. To change or delete a group name, select the group to be changed prior to selecting **Group Configuration**. To change group order, simply drag-and-drop the group name to its new position.

Technologies

Multiple filtering technologies are available that can be applied to the users in any particular group. These technologies will not be applied to websites on the whitelist.

Web Based Proxies and Anonymizer Detection

This technology is designed to detect and block web based proxy services. This system uses fingerprints of popular websites along with a built in pre-detected blacklist of known proxy servers to provide a high success rate in eliminating this problem. Proxies detected by this system are shared by the rest of the CIAPFilters in our cloud, which provides a very high rate of discovery for new proxy servers that are made available on the Internet.

Safe Search Enforcement

This technology detects when a user accesses the popular search engines and enables the search engines built in "Safe Search" feature. This technology will reduce the amount of pornography and other objectionable content returned by search engines. Safe Search is also applied to YouTube.

Pornography Detection

Activates CIPAFILTER's pornography blocking system. Each time a website is blocked the administrator is e-mailed and the block can be recorded in the web monitoring system (If web monitoring is active).

Extreme Language

The Extreme Language technology blocks any page with offensive language on it. Pornography Detection must be selected for this technology to function.

Block Downloads

This option blocks users from being able to download restricted files. Files that are installable, executable or may otherwise contain viruses, spyware or trojans and are not required by the average user on a daily basis are considered restricted. The following extensions are considered restricted as of this writing: ".sit", ".dmg", ".ocx", ".cab", ".exe", ".com", ".bat", ".zip", ".tar", ".tgz", ".rar", and ".iso".

Block Flash

If this feature is activated the users will not be able to access Flash content.

NAC

This technology allows users in the defined group to temporarily authenticate unknown devices for network access via the Network Access Control feature of the Network Security suite.

Automatic Blacklists

Automatic Blacklists are lists of websites compiled by CIPAFilter and are updated twice daily from our corporate office. The Automatic Blacklists on the CIPAFilter include:

- **Gambling** – Sites which encourage gambling such as betting sites and online casinos.
- **Games** – Sites related to computer and electronic games, particularly sites where the user can play games. This blacklist also includes game download sites, game review websites, and websites devoted to the discussion of games. This category does not include games sites devoted to educational use only.
- **Games (DMOZ Enhanced)** – This category blocks additional game sites found at the Open Directory Project.
- **Hate, Drugs and Violence** – Sites related to racial hate, drugs, alcohol, and violent crime.
- **Image / Video** – This category blocks major image / video search sites such as YouTube.com, Flurl.com, etc. Also, sites for watching television such as Hulu are blocked. This category does not include sites that are primarily news websites that may include some video. This category also does not include image sites on which CIPAFilter enforces safe search such as Google images.
- **Online / Streaming Radio** - Major sites that offer online music/radio streaming, such as Pandora, Last.fm, SHOUTcast and Live365, as well as blocking for terrestrial radio broadcasting sites. Also included are services such as iTunes, Napster and Rhapsody.
- **Shopping and Travel** – Sites related to online auctions, travel and shopping.

- **Social Networking** – Sites offering an interactive network of photos, blogs, user profiles, etc.
- **Web-based E-mail** – Sites which offer free online web-based e-mail services, such as hotmail.com.

Temporary Whitelist Entries

This section displays any temporary whitelist entries created using the **Add to Whitelist** selection on the CIPAFilter block page. The **Manage Temporary Whitelist Entries** link allows an administrative user of the CIPAFilter to view the history of temporary whitelist entries and remove currently active items.

Note: **Add to Whitelist** can only be authenticated by the root and admin users of the CIPAFilter.

Manual Blacklists

The whitelist/blacklist system allows you to control web access by using a basic domain oriented syntax or a sophisticated regular expression URL matching technology to either allow or reject a websites based on their URL.

To block an entire website, simply enter it's domain. For example, "google.com" will block everything at , test.google.com, or any files therein.

Parts of a website like <http://games.yahoo.com> can be blocked as well using a regular expression syntax. These entries must be preceded by "REGEX:" and the syntax is complicated. Full support for industry standard POSIX regular expressions is provided. Please see Appendix II for regular expression syntax, or simply use the following examples as guidelines.

- Group characters with "("
- "*" matches any number of the preceding character or group
- "?" matches 0 or 1 of the preceding character or group
- "+" matches 1 or more of the preceding character or group
- "." matches any single character
- "^" matches the beginning of a URL
- "[^/]" matches anything not a "/"
- Precede special characters with a "\" to use them as literals

Blacklist Examples

<i>Example</i>	<i>Description</i>
REGEX: ^https?://[^\.]*\.edu/	"^" Matches the beginning of the URL. "http:/" matches that text in the URL. "[^\.]*" matches any number of characters that are not a ".". "\.edu/" matches ".edu/" at the end of the URL. The period must be preceded with a slash because otherwise it would be the special "Match any on character" symbol as opposed to the period literal. The effect of this rule is to match all .edu websites.
REGEX: ^https?://([^\.]*\.)?google\.com/	Matches google.com and all Google subdomains. The "[^\.]*\." is grouped and made optional so the rule won't match "hgoogle.com", but still matches "google.com" and "www.google.com".
REGEX: ^https?://games\.yahoo\.com/	Matches only the games.yahoo.com subdomain of yahoo. If a blacklist entry, users will still be able to use all other yahoo services.

Super Whitelist

The **Super Whitelist** is a special, automatic whitelist that applies to all groups. The **Super Whitelist** whitelists all operations including virus protection and authentication.

Global Lists

The **Global Whitelist** and **Global Blacklist** apply to defined groups if the **Global Lists** checkbox is selected.

15. Permissions Scheduling

Permissions Scheduling allows you to configure certain groups or subnets to be mapped to a different group at specific times of day.

If a subnet is mapped to a group all users on that subnet will be treated as members of the group during the times of day specified in the rule.

If a group is mapped to another group all users in the first group will be treated as members of the second group during the times of day specified by the rule.

Selecting **Invert** allows you to invert the times of day the rule is active with the times of day it is inactive. For example, you can select the hours school is in session with a rule and then select **Invert** in order to make a rule that will be in effect when school is out of session.

16. Internet Reports

The web usage reporting system is activated on the **Content Filtering** page. After it is activated, it collects information on every URL visited by your Internet users. It will track each user by user name and password if web authentication is activated, otherwise it can only track users by IP address.

17. Notifications

Notification Address

Informational notices from the CIPAFilter are sent here. This includes content filter trips and mail errors. This default to administrator@domain, where domain is the domain entered on the **IP Settings** page.

Daily, Weekly and Custom Reports

The CIPAFilter can send daily, weekly and custom reports to the email address specified.

18. Bandwidth Control

Bandwidth Control allows you to keep heavy bandwidth users from consuming your entire network bandwidth to the detriment of all other activities on your network. Bandwidth can be controlled by IP or subnet and is allocated in kilobits per second. Please be advised, that if a user is using their entire bandwidth allotment new packets to them will be delayed or dropped in order to limit them to their prescribed cap.

Within each IP or Subnet's bandwidth allocation ICMP, TOS 0x10 and ACK packets are prioritized to facilitate performance of real time protocols and pings and to eliminate wasteful packet re-transmissions due to delayed acknowledgments.

IP addresses or subnets not listed on this page are all lumped together in a single queue which is treated the same way as above, but is prioritized as a whole behind everything else. Therefore, IP addresses and subnets listed on this page will have their traffic served first, before any traffic from IP addresses or subnets not listed here.

When bandwidth is tight, all machines are first allocated their Committed Information Rate (CIR). After all CIRs have been met, the remaining bandwidth is divided between IP addresses in ratio with their CIRs, but no machine is allowed to exceed its maximum allocation.

19. Bandwidth Reports

Clicking on **Bandwidth Reports** brings up the "Bandwidthd" application. This software tracks the Internet usage of all clients on your network. Select an interface under **Select A Sensor** and click go to request a report. Custom reports and graphs can be requested by calling special URLs on the CIPAFILTER. Speak with tech support if you wish to embed these reports into other web consoles.

20. VPN

Client-Server VPN

CIPAFILTER supports PPTP VPNs for client VPN. PPTP is a very common standard and implementations are built into the operating system for Mac, Windows, Linux and even common hand-held appliances such as smart phones.

Simply configure a range of IPs in the **PPTP Range** box and add an appropriate user name and password to the **User Manager** page.

Warning: PPTP VPNs are only as secure as the passwords you use. We recommend that if you are using the link to transmit confidential information that you use very strong passwords.

Router-Router VPN

CIPAFILTER uses the more secure IPSEC VPN technology for router-to-router VPNs. CIPAFILTER's implementation is industry standard, based on open source technology, and will probably work with any IPSEC enabled device, but is designed primarily to interconnect two CIPAFILTER units. If you want to interconnect a CIPAFILTER with another device we will try to assist you, but IPSEC is a very complicated technology and we cannot guarantee the results.

To create a IPSEC tunnel simply enter the public IP address of the remote end point and the remote networks subnet. Click **Save and Apply** and a new key will be generated for you. Copy this key to the remote end and configure the remote end with the local public IP and subnet. Click **Save and Apply** on

the remote router and send a ping across the link to bring it up and test that it is functional. If you have any trouble establishing one of these links, please don't hesitate to call tech support so that we can troubleshoot any problems you may be having.

21. Network Access Control

Network Access Control (NAC) prevents the use of the internal network by unknown (“foreign”) devices. NAC operates in two modes: learning mode and normal mode. In learning mode, NAC simply collects information about the devices that it sees on the internal network. In normal mode, NAC prevents any unknown devices from accessing network resources.

NAC maintains three lists of devices: allowed, temporarily allowed and denied. Devices on the allowed list can access internal network resources and the Internet. Devices on the denied list cannot. Devices on the temporarily allowed list have full network access for a 12-hour period.

Devices can be added to each list either one-at-a-time or via a CSV file upload using the CIPAFILTER web interface. New, foreign devices can also gain temporary network access via a web browser authentication screen similar to the network logon screen used at hotels. Via the browser authentication screen, the user will enter their network LDAP credentials. If the user is in an LDAP group with the NAC technology selected on the **Group Permissions** page of the CIPAFILTER, the new device will be placed in the temporarily allowed list.

Network Access Control emits targeted ARP jamming pulses to segregate foreign devices from the known devices in the local broadcast domain. If the network is designed with multiple broadcast domains (a routed network), CIPAFILTER has small NAC units that can be deployed at remote facilities to communicate NAC information to the main CIPAFILTER unit.

Notes:

1. Network Access Control is only effective for IPV4 networks. Please ensure that any and all protocols other than IPV4 are disabled on internal devices to prevent foreign machines from accessing them via another protocol.
2. ARP is heavily cached by network devices including workstations. The ARP cache on each device may take several minutes to update with new device information. To expedite troubleshooting, you may need to clear the ARP cache on devices. A reboot of a device will clear the ARP cache.

22. Anti-Spam

CIPAFILTER's anti-spam system is designed to act as a proxy and a firewall to protect your mail server from the Internet. CIPAFILTER allows you to take full advantage of your mail server's advanced features without the risk of having it accessible to hackers and worms. We recommend placing your existing server behind the CIPAFILTER and configuring it with only a private IP address. Change your MX records to point to the CIPAFILTER, and configure your CIPAFILTER to route e-mail for your domains to the private IP address of your actual e-mail server. This will cause all e-mail to be delivered to CIPAFILTER, where it will be virus scanned, spam filtered, content filtered, then forwarded on to your actual mail server for processing.

Because of the problems with spam in recent years, many companies have become more cautious about the mail servers with which they communicate. Many unofficial rules have been adopted in a piecemeal fashion across the Internet. If your mail server is configured incorrectly, you may find that most people will receive your messages but for some they arrive marked as spam or not at all. Please allow us to help you through setting up your mail server to be compliant with all official and unofficial guidelines for message processing.

However, if you are an expert, and wish to configure the server yourself, please be certain to comply with the following guidelines. If you don't understand the reasoning behind any of the following, our tech support representatives will be happy to go over it with you.

- Always make sure the IP address pointed to by your MX record reverses to the same name that is contained in your MX record.
- If you are using CIPAFILTER for anti-spam do not use secondary MXs.
- Make sure the CIPAFILTER host and domain name match the name from your MX record.

To configure CIPAFilter to proxy e-mail for an existing server simply **Add Route** for your domain in the **Mail Routes** section to route e-mail for your domain to your mail server's private IP address.

To use CIPAFilter as an outgoing SMTP server, you must add the IP address of your internal mail server to the **SMTP Relay Authorized for the Following Subnets** list. If individual workstations will send mail directly to the Internet bypassing the mail server, their IP subnets must be added to the **SMTP Relay Authorized for the Following Subnets** list.

E-mail Content Filtering

E-mail Content Filtering applies the web content filtering technology to incoming and outgoing mail messages. If a message is inappropriate, a mail message is sent to the person in your organization who was either sending it or was the intended recipient and the message is blocked. This is great for filtering out pornographic spam.

Email Archive

Email-Archive can be set to one of three settings. The settings are:

- **Rcpt: cipafilter_email_archive** - is designed to work with email journaling support in products like Microsoft Exchange. Many mail server products support journaling, but some, like GroupWise, require third party software. If your mail server supports journaling, set the journaling address to cipafilter_email_archive@10.0.0.1 where 10.0.0.1 represents the CIPAFilter's internal IP address.
- **All Messages Passing Through** - If your mail server does not support journaling, you can still archive all messages that pass through the CIPAFilter using this selection.

- **None** - turns off the Email Archive functionality.

The archived messages can be browsed under **E-mail Archive**. If you are archiving email for retention purposes, be sure to speak with tech support about creating a suitable backup and recovery plan.

Email Footers

The CIPAFilter can automatically append a message footer to each email passing through the CIPAFilter. Simply compose your desired footer and upload it to the CIPAFilter. Footers can be selected for text-formatted e-mail messages and HTML formatted e-mail messages.

Mail Routes

These entries dictate where CIPAFilter will route e-mail after it has finished scanning it. Enter your domains here to route them to the IP address of your true mail server.

SMTP Relay Authorization

Only e-mail clients using IP addresses from subnets listed here can use CIPAFilter as an outgoing mail server. The **IP Ranges** entered here are Sendmail style IP ranges, which is just a text match to the text of the IP address.

Subnet	IP Range
10.0.0.0/8	10.
192.168.0.1/24	192.168.0.
64.215.40.0/25	64.215.40.
12.109.195.200/32	12.109.195.200

Anti-Spam Whitelist

E-mail to or from servers on this list are not affected by the anti-spam system. Enter the server IP, domain names, or subnet here if you have difficulty receiving e-mail from another party. Our anti-spam system works with all standards compliant mail servers, however some older and custom systems may have problems. See **SMTP Relay Authorization** for IP range syntax.

Note: The Anti-Spam Whitelist does not bypass virus scanning for email items.

Spam Forwarding

CIPAFilter can redirect all spam with a score above a certain threshold to a spam mailbox instead of delivering this mail to the intended recipient.

Anti-Spam Blacklist

All email from these domains, subnets or email addresses will be rejected with a 550 error.

23. E-mail Archive

This selection opens the window used to view the e-mail archive contained on the CIPAFilter. From the **E-mail Archive** interface, messages can be selected by mailbox, To address, and From address. You can search for specific text in the subject or body of the email messages. Messages can be selected within a specified date range. Selected messages can be downloaded to a file using EML format.

24. Client Software

The client software page contains the latest version of our CIPAFilter authentication client as well as a copy of the EICAR anti-virus test file.

The authentication client is supplied for both Microsoft Windows (32-bit and 64-bit) and Mac OS/X. The Windows client is in the form of a silent msi installer and is designed to be automatically deployed. If you are considering deploying this software, please consult with tech support for advice on deployment scenarios.

The EICAR anti-virus test file is a standard way of testing the effectiveness of your anti-virus system's configuration. It is a harmless text file that the various anti-virus vendors have all agreed to detect as a virus for the purpose of testing.

Please be certain to attempt a test virus transmission through each protocol you intend to protect to insure proper setup after installation.

Notes:

HTTP virus scanning requires web proxy to be activated and functioning in order to operate. HTTP anti-virus will work in both transparent and proxy server modes.

SMTP anti-virus requires that your MX record points to the CIPAFilter. Your mail can then be forwarded on to your internal mail server with the **Mail Routes** options on the **Anti-Spam** page.

FTP anti-virus only works with non-transparent web proxy. Your clients must be configured to use CIPAFilter as an FTP proxy.

25. Database Maintenance

The **Database Maintenance** page is where you manage maintenance and backup of your database. Here you configure your purge windows, examine the latest backup log, download a backup, or configure them to be pushed to a remote server.

Database backup and purge runs automatically every Saturday at 2:30AM. It can also be run manually at any time by clicking **Start Run**. Progress can be monitored by refreshing the page and viewing the log. Backup and purge applies only to the database stored locally on the router. The remote database is not affected by these settings in any way. If the remote database server is not a CIPAFILTER, the customer is responsible for implementing backup and purge.

If restoration of the database after failure is important for your implementation, the following responsibilities are incurred by the customer. Speak with tech support to help you determine a sensible test and backup procedure that meets your requirements.

- Check this page weekly for anomalies in the backup/purge log
- Periodically download database backups for use in recovery in the case of mirror failure
- Periodically test the backup by restoring a copy to a PostgreSQL server

Warning: Increasing purge windows will increase load, database size and backup size.

Backup Push

Backup push is where you configure the unit to upload a copy of the backup it performs each week. The backup is transferred to a Windows file share. Use of this feature is highly recommended, especially for **E-mail Archive** customers, in order to prevent data loss in the case of catastrophic failure of the unit.

26. Firmware Updates

Automatic Updates

The CIPAFilter can be configured to automatically receive firmware fixes to the current firmware version. To do so, select **Enabled** in the **Automatic Updates** section and press **Save and Apply**.

Manual Update

The CIPAFilter firmware can be updated by simply selecting the **Firmware Upgrade** button in the **Manual Update** section. The firmware upgrade process typically takes 2-3 minutes including the time required to restart the CIPAFilter.

27. Override Console

The Override Console is used by tech support to enter custom configuration changes. This page is used to work around problems specific to a single customer, or to implement customizations to the product for a single customer.

28. Config Save/Restore

This page allows you to easily download the configuration of your router after you make changes. Keeping a copy of your configuration on-hand will allow you to easily configure a new router in the event of a failure or an upgrade. From here you can also revert your router back to factory settings.

29. Troubleshooting

The **Troubleshooting** page runs diagnostics on the router. There should be no entries on this page marked as "FAIL". If there are, please contact tech support.

Network Diagnostics tests can also be run from this page to ping devices, perform a traceroute and perform DNS lookups.

30. Appendix I: Dot Notation to CIDR Notation Translation Table

Subnet Mask	Subnet Bits	IP Addresses
255.0.0.0	/8	16,581,375
255.255.0.0	/16	65,025
255.255.255.0	/24	256
255.255.255.128	/25	128
255.255.255.192	/26	64
255.255.255.224	/27	32
255.255.255.240	/28	16
255.255.255.248	/29	8
255.255.255.252	/30	4
255.255.255.254	/31	2
255.255.255.255	/32	1

31. Appendix II: POSIX Regular Expressions

regex - POSIX 1003.2 regular expressions

Regular expressions (``RE"s`), as defined in POSIX 1003.2, come in two forms: modern REs (roughly those of *egrep*; 1003.2 calls these ``extended" REs) and obsolete REs (roughly those of ed\(1\); 1003.2 `basic" REs). Obsolete REs mostly exist for backward compatibility in some old programs; they will be discussed at the end. 1003.2 leaves some aspects of RE syntax and semantics open; `(!) marks decisions on these aspects that may not be fully portable to other 1003.2 implementations.`

A (modern) RE is one(!) or more non-empty(!) *branches*, separated by ``|`. It matches anything that matches one of the branches.

A branch is one(!) or more *pieces*, concatenated. It matches a match for the first, followed by a match for the second, etc.

A piece is an *atom* possibly followed by a single(!) ``*'`, ``+''`, ``?'`, or *bound*. An atom followed by ``*'` matches a sequence of 0 or more matches of the atom. An atom followed by ``+''` matches a sequence of 1 or more matches of the atom. An atom followed by ``?'` matches a sequence of 0 or 1 matches of the atom.

A *bound* is ``{'` followed by an unsigned decimal integer, possibly followed by ``,'` possibly followed by another unsigned decimal integer, always followed by ``}'`. The integers must lie between 0 and `RE_DUP_MAX` (255(!)) inclusive, and if there are two of them, the first may not exceed the second. An atom followed by a bound containing one integer *i* and no comma matches a sequence of exactly *i* matches of the atom. An atom followed by a bound containing one integer *i* and a comma matches a sequence of *i* or more matches of the atom. An atom followed by a bound containing two integers *i* and *j* matches a sequence of *i* through *j* (inclusive) matches of the atom.

An atom is a regular expression enclosed in ``()'` (matching a match for the regular expression), an empty set of ``()'` (matching the null string)(!), a *bracket expression* (see below), ``.'` (matching any single character), ``^'` (matching the null string at the beginning of a line), ``$'` (matching the null string at the end of a line), a ``\'` followed by one of the characters ``^.[${}] *+?{\ \'` (matching that character taken as an ordinary character), a ``\'` followed by any other character(!) (matching that character taken as an ordinary character, as if the ``\'` had not been present(!)), or a single character with no other significance (matching that character). A ``{'` followed by a character other than a digit is an ordinary character, not the beginning of a bound(!). It is illegal to end an RE with ``\'`.

A *bracket expression* is a list of characters enclosed in `[]`. It normally matches any single character from the list (but see below). If the list begins with `^`, it matches any single character (but see below) *not* from the rest of the list. If two characters in the list are separated by `-`, this is shorthand for the full *range* of characters between those two (inclusive) in the collating sequence, e.g. `[0-9]` in ASCII matches any decimal digit. It is illegal(!) for two ranges to share an endpoint, e.g. `[a-c-e]`. Ranges are very collating-sequence-dependent, and portable programs should avoid relying on them.

To include a literal `\` in the list, make it the first character (following a possible `^`). To include a literal `-`, make it the first or last character, or the second endpoint of a range. To use a literal `-` as the first endpoint of a range, enclose it in `[.` and `.]` to make it a collating element (see below). With the exception of these and some combinations using `[` (see next paragraphs), all other special characters, including `\`, lose their special significance within a bracket expression.

Within a bracket expression, a collating element (a character, a multi-character sequence that collates as if it were a single character, or a collating-sequence name for either) enclosed in `[.` and `.]` stands for the sequence of characters of that collating element. The sequence is a single element of the bracket expression's list. A bracket expression containing a multi-character collating element can thus match more than one character, e.g. if the collating sequence includes a `ch` collating element, then the RE `[[.ch.]]*c` matches the first five characters of `chchcc`.

Within a bracket expression, a collating element enclosed in `[=` and `=]` is an equivalence class, standing for the sequences of characters of all collating elements equivalent to that one, including itself. (If there are no other equivalent collating elements, the treatment is as if the enclosing delimiters were `[.` and `.]`.) For example, if `o` and `are` are the members of an equivalence class, then `[[=o=]]`, `[[=are=]]`, and `[o]` are all synonymous. An equivalence class may not(!) be an endpoint of a range.

Within a bracket expression, the name of a *character class* enclosed in `[:` and `:]` stands for the list of all characters belonging to that class. Standard character class names are:

- alnum digitpunct
- alpha graphspace
- blank lowerupper
- cntrl printxdigit

These stand for the character classes defined in `wctype(3)`. A locale may provide others. A character class may not be used as an endpoint of a range.

There are two special cases(!) of bracket expressions: the bracket expressions ``[:<:]'` and ``[:>:]'` match the null string at the beginning and end of a word respectively. A word is defined as a sequence of word characters which is neither preceded nor followed by word characters. A word character is an *alnum* character (as defined by [wctype\(3\)](#)) or an underscore. This is an extension, compatible with but not specified by POSIX 1003.2, and should be used with caution in software intended to be portable to other systems.

In the event that an RE could match more than one substring of a given string, the RE matches the one starting earliest in the string. If the RE could match more than one substring starting at that point, it matches the longest. Subexpressions also match the longest possible substrings, subject to the constraint that the whole match be as long as possible, with subexpressions starting earlier in the RE taking priority over ones starting later. Note that higher-level subexpressions thus take priority over their lower-level component subexpressions.

Match lengths are measured in characters, not collating elements. A null string is considered longer than no match at all. For example, ``bb*'` matches the three middle characters of ``abbbc'`, ``(wee|week)(knights|nights)'` matches all ten characters of ``weeknights'`, when ``(.*)*'` is matched against ``abc'` the parenthesized subexpression matches all three characters, and when ``(a)**'` is matched against ``bc'` both the whole RE and the parenthesized subexpression match the null string.

If case-independent matching is specified, the effect is much as if all case distinctions had vanished from the alphabet. When an alphabetic that exists in multiple cases appears as an ordinary character outside a bracket expression, it is effectively transformed into a bracket expression containing both cases, e.g. ``x'` becomes ``[xX]'`. When it appears inside a bracket expression, all case counterparts of it are added to the bracket expression, so that (e.g.) ``[x]'` becomes ``[xX]'` and ``[^x]'` becomes ``[^xX]'`.

No particular limit is imposed on the length of REs(!). Programs intended to be portable should not employ REs longer than 256 bytes, as an implementation can refuse to accept such REs and remain POSIX-compliant.

Obsolete ("basic") regular expressions differ in several respects. `[]`, `+`, and `?` are ordinary characters and there is no equivalent for their functionality. The delimiters for bounds are `{` and `}`, with `{` and `}` by themselves ordinary characters. The parentheses for nested subexpressions are `(` and `)`, with `(` and `)` by themselves ordinary characters. `^` is an ordinary character except at the beginning of the RE or (!) the beginning of a parenthesized subexpression, `$` is an ordinary character except at the end of the RE or (!) the end of a parenthesized subexpression, and `*` is an ordinary character if it appears at the beginning of the RE or the beginning of a parenthesized subexpression (after a possible leading `^`). Finally, there is one new type of atom, a *back reference*: `\d` followed by a non-zero decimal digit *d* matches the same sequence of characters matched by the *d*th parenthesized subexpression (numbering subexpressions by the positions of their opening parentheses, left to right), so that (e.g.) `\([bc]\)1` matches `bb` or `cc` but not `bc`.

SEE ALSO

[regex\(3\)](#)

POSIX 1003.2, section 2.8 (Regular Expression Notation).

BUGS

Having two kinds of REs is a botch.

The current 1003.2 spec says that `)` is an ordinary character in the absence of an unmatched `(`; this was an unintentional result of a wording error, and change is likely. Avoid relying on it.

Back references are a dreadful botch, posing major problems for efficient implementations. They are also somewhat vaguely defined (does `a((b)*2)*d` match `abbbd`?). Avoid using them.

1003.2's specification of case-independent matching is vague. The "one case implies all cases" definition given above is current consensus among implementors as to the right interpretation.

The syntax for word boundaries is incredibly ugly.

This page was taken from Henry Spencer's regex package.