

# InfraNotes2020

- Infrastructure
  - DNS Hosting
  - Standards
  - Legacy Servers
  - Servers
  - Backups
  - sa-vm.apache.org install/migration notes
- Builds
  - mkupdates
  - nitemc
  - ruleqa

This page was created in July of 2020 during sa-vm1 sa-vm migration ( [INFRA-20449](#) - Getting issue details...  )

## Infrastructure

### DNS Hosting

PowerDNS web interface for easy management of spamassassin.org DNS records:

1. Open an SSH tunnel: `ssh -f sa-vm.apache.org -L 8090:localhost:8090 -N`
2. Open web interface: <http://localhost:8090>
3. Login with admin. (Password is encrypted in [sysadmins/accounts](#).)
  - a. Root can also create new users (`sqlite3 /var/www/nsedit/includes/pdns.users.sqlite3 - insert new into users table`)

Zone	Server	Contact	Notes
spamassassin.org	ns2.pccc.com	Kevin <a href="#">McGrail</a> kevin.mcgrail@mcgrail.com, kmcgrail@apache.org	Instant updates via NOTIFY
	ns2.ena.com	Dave Jones djones@ena.com, davej@apache.org	Instant updates via NOTIFY
	dns-master.sonic.net	ops-req@sonic.net, joe.muller@sonic.net	Hidden slave, 0-10 min delay of public slaves after NOTIFY
	<del>ns.hyperreal.org</del>	<del>Brian Behlendorf</del>	<del>Currently not used since DJBDNS doesn't support NOTIFY or EDNS over TCP</del>

### Standards

\*Ubuntu 20.04 LTS

\*Cron entries should be in new standard locations `/etc/cron.d`, `/etc/cron.daily`, etc. and avoid using user's crontab

\*Custom scripts should reside in `/usr/local/bin` if they are not directly related to [SpamAssassin](#) processing that should be in `/usr/local/spamassassin`

\*Symlink scripts from `/usr/local/bin` to `/etc/cron.d`, `/etc/cron.daily`, or `/etc/cron.weekly`. This provides easy discovery and future management by others on the sysadmins team.

\*Scripts and cron entries should mail output to the sysadmins mailing list

### Legacy Servers

\*minotaur.apache.org - handled various build and devel related tasks

\*hyperion.apache.org - likely a Solaris box that had backup data of next server

\*spamassassin.zones.apache.org - DIED - was replaced with spamassassin-vm

\*spamassassin.zones2.apache.org - deprecated by Infra, replaced by sa-vm1.apache.org

\*spamassassin-vm.apache.org - deprecated by Infra, replaced by sa-vm1.apache.org

\*buildbot, ruleqa, etc. are aliases of above deprecated servers

### Servers

Hostname	Function	Software	Configs/Location	Resource/URL
apachesf.sonic.net	Donated by Sonic	CentOS 7		sa-update.spamassassin.org (64.142.56.146)
sa-vm.apache.org	DNS Hidden Master	PowerDNS	/etc/powerdns/pdns.d/pdns.local.conf	spamassassin.org (DNS)

	Rsync Mirrors	rsyncd	/etc/rsyncd.conf	rsync.spamassassin.org
	RuleQA	apache2	/etc/apache2/sites-available/apache2-le-ssl.conf	ruleqa.spamassassin.org
	Nightly Masscheck	cron/scripts	/etc/cron.d/automc, /usr/local/spamassassin	

## Backups

An old backup of sa-vm1 etc exists in sa-vm.apache.org:/usr/local/spamassassin/backups.

We need to setup offsite backups that at least two of the SA sysadmins members can access.

sa-vm.apache.org OS is backedup by ASF infra.

Crashplan will be installed by KAM to backup everything including /usr/local/spamassassin.

## sa-vm.apache.org install/migration notes

```
# note that server uses an internal 10.x IP, and sa-vm.apache.org is an external NAT IP.
# /etc/hosts has some redirected names to localhost for ruleqa.spamassassin.org etc

apt install chrony
systemctl start chrony
systemctl enable chrony

apt install apache2 libapache2-mod-geoip libapache2-mod-php7.4 php7.4-sqlite3 php7.4-curl
a2enmod cgid
a2enmod cgi
a2enmod rewrite
a2enmod ssl
a2disconf serve-cgi-bin
a2dissite 000-default
a2dissite default-ssl

dpkg --purge geoip-database
mkdir -m 755 /usr/share/GeoIP
curl -o /etc/cron.weekly/geoip_update https://mailfud.org/geoip-legacy/geoip_update.sh
chmod 755 /etc/cron.weekly/geoip_update
## edit geoip_update, FILES="GeoIP GeoIPv6 GeoIPCity GeoIPCityv6 GeoIPASNum GeoIPASNumv6 GeoIPOrg GeoIPISP"
/etc/cron.weekly/geoip_update

groupadd -g 60000 automc
groupadd -g 60001 rsync
groupadd -g 60002 release
groupadd -g 60003 bbmass

useradd -u 60003 -g bbmass -d /usr/local/spamassassin/bbmass -s /bin/bash bbmass
useradd -u 60002 -g release -d /usr/local/spamassassin/release -s /bin/bash release
useradd -u 60001 -g rsync -G www-data,release -d /usr/local/spamassassin/rsync -s /bin/bash rsync
useradd -u 60000 -g automc -G www-data,rsync,release -d /usr/local/spamassassin/automc -s /bin/bash automc

rsync -vaH root@sa-vm1.apache.org:/usr/local/spamassassin/. /usr/local/spamassassin/.
rsync -vaH root@sa-vm1.apache.org:/var/www/. /var/www/.

systemctl stop systemd-resolved
# edit /etc/systemd/resolved.conf -> DNSStubListener=no
systemctl start systemd-resolved

apt install pdns-server pdns-backend-sqlite3 sqlite3 jq
systemctl stop pdns

apt install sysstat libalgorithm-diff-perl libalgorithm-diff-xs-perl \
libalgorithm-merge-perl libapparmor-perl libapt-pkg-perl libauthen-sasl-perl \
libbb-hooks-op-check-perl libbareword-filehandles-perl libcgi-fast-perl \
libcgi-pm-perl libclass-accessor-perl libclass-data-inheritable-perl \
libclass-dbi-abstractsearch-perl libclass-dbi-mysql-perl libclass-dbi-perl \
libclass-method-modifiers-perl libclass-singleton-perl libclass-trigger-perl \
libclass-xsaccessor-perl libclone-perl libconfig-file-perl \
libcrypt-openssl-bignum-perl libcrypt-openssl-rsa-perl libdate-manip-perl \
libdatetime-locale-perl libdatetime-perl libdatetime-timezone-perl \
libdbd-mysql-perl libdbi-perl libdbix-contextualfetch-perl \
libdevel-globaldestruction-perl libdigest-hmac-perl libdigest-sha-perl \
libdpkg-perl libencode-detect-perl libencode-locale-perl liberror-perl \
libexporter-tiny-perl libfcgi-perl libfile-fcntllock-perl \
libfile-listing-perl libfont-afm-perl libgd-perl libgeo-ip-perl \
libgeo-ipfree-perl libhash-merge-perl libhtml-form-perl libhtml-format-perl \
libhtml-parser-perl libhtml-tagset-perl libhtml-tree-perl \
```

```

libhttp-cookies-perl libhttp-daemon-perl libhttp-date-perl \
libhttp-message-perl libhttp-negotiate-perl libima-dbi-perl \
libimport-into-perl libindirect-perl libio-html-perl libio-socket-inet6-perl \
libio-socket-ssl-perl libio-stringy-perl liblexical-sealrequirehints-perl \
liblingua-en-inflect-perl liblist-moreutils-perl liblocale-gettext-perl \
liblwp-mediatypes-perl liblwp-protocol-https-perl libmail-dkim-perl \
libmail-spf-perl libmailtools-perl libmodule-implementation-perl \
libmodule-runtime-perl libmoo-perl libmultidimensional-perl \
libnet-cidr-lite-perl libnet-dns-perl libnet-http-perl libnet-ip-perl \
libnet-libidn-perl libnet-patricia-perl libnet-smtp-ssl-perl \
libnet-snmp-perl libnet-ssleay-perl libnet-xwhois-perl libnetaddr-ip-perl \
libparams-classify-perl libparams-validate-perl libregexp-assemble-perl \
librole-tiny-perl libsnmp-perl libsocket6-perl libsql-abstract-limit-perl \
libsql-abstract-perl libstrictures-perl libsub-exporter-progressive-perl \
libsub-name-perl libsvn-perl libterm-readkey-perl libtext-charwidth-perl \
libtext-iconv-perl libtext-wrapil8n-perl libtime-parsedate-perl \
libtime-piece-mysql-perl libtimedate-perl libtry-tiny-perl \
libuniversal-moniker-perl liburi-perl libwww-perl libwww-robotrules-perl \
libxml-libxml-perl libxml-namespacesupport-perl libxml-parser-perl \
libxml-sax-base-perl libxml-sax-expat-perl libxml-sax-perl \
libxml-simple-perl libyaml-libyaml-perl libyaml-perl libcompress-lz4-perl \
libxml-sax-expatxs-perl libbsd-resource-perl libarchive-zip-perl \
libio-string-perl libmath-int64-perl

apt install gnupg2 gnupg-agent pigz gnuplot git-svn dnsutils zip zsh tcsh \
gsfonts gsfonts-x11 pyzor razor lzop makedev mutt rename
systemctl stop gdm
systemctl disable gdm

wget https://cpan.metacpan.org/authors/id/J/JH/JHI/Statistics-DEA-0.04.tar.gz; tar xvfz Statistics-DEA-0.04.tar.
gz; cd Statistics-DEA-0.04; perl Makefile.PL; make install
wget https://cpan.metacpan.org/authors/id/J/JM/JMASON/IPC-DirQueue-1.0.tar.gz; ...
wget https://cpan.metacpan.org/authors/id/G/GA/GAAS/Digest-SHA1-2.13.tar.gz; ...
wget https://cpan.metacpan.org/authors/id/N/NW/NWELLNHOF/IP-Country-DB_File-3.03.tar.gz; ...

rsync -va root@sa-vml.apache.org:/usr/local/bin/*.sh /usr/local/bin/
rsync -va root@sa-vml.apache.org:/usr/local/bin/dns_compare /usr/local/bin/
apt install python python-dnspython

rsync -va root@sa-vml.apache.org:/etc/letsencrypt /etc/
apt install certbot python3-requests
# change to python3 --> /etc/letsencrypt/acme-dns-auth.py #! python3

rsync -va root@sa-vml.apache.org:/usr/local/spamassassin/automc/svn/automc/apache2-le-ssl.conf /etc/apache2/sites-
available/
rsync -va root@sa-vml.apache.org:/etc/apache2/sites-available/nsedit.conf /etc/apache2/sites-available/
a2ensite apache2-le-ssl
a2ensite nsedit
systemctl enable apache2
systemctl restart apache2

rsync -va root@sa-vml.apache.org:/etc/rsyncd.conf /etc/
systemctl enable rsync
systemctl start rsync

##
## final syncs after shutting down sa-vml services, crons commented out
##

rsync -vaHz --delete root@sa-vml.apache.org:/usr/local/spamassassin/. /usr/local/spamassassin/.
rsync -vaH --delete root@sa-vml.apache.org:/var/www/. /var/www/.
rsync -vaH root@sa-vml.apache.org:/etc/cron.d/automc :/etc/cron.d/svn /etc/cron.d
rsync -vaH root@sa-vml.apache.org:/etc/cron.hourly/setperms /etc/cron.hourly
rsync -vaH root@sa-vml.apache.org:/etc/cron.daily/checkDNShosting /etc/cron.daily

systemctl stop pdns
rm -f /var/lib/powerdns/pdns.sqlite3*
rsync -va root@sa-vml.apache.org:/var/lib/powerdns/pdns.sqlite3* /var/lib/powerdns/
sqlite3 /var/lib/powerdns/pdns.sqlite3
### UPDATE domainmetadata SET content='DEFAULT' WHERE kind='SOA-EDIT-API' AND content='INCEPTION-INCREMENT';
# also replace /var/www/nsedit/*/* INCEPTION-INCREMENT -> DEFAULT
systemctl start pdns
systemctl enable pdns

systemctl start apache2
systemctl enable apache2

rsync -va root@sa-vml.apache.org:/etc/letsencrypt /etc/

# check
# /etc/cron.d/* MAILTO=
# /usr/local/bin/* NOTIFY=
# uncomment cron

```

```
# fixes to masscheck, revisions r1880323, r1880320, r1880318, r1880316, r1880312, r1880309
```

## Builds

The sa-vm1 server TZ is UTC so cron entries will be in UTC.

## mkupdates

This section of scripts publishes new ruleset updates to the mirrors. There are currently two different rule daily updates. Both do lint tests against the latest version of [SpamAssassin](#) but the first one updates the 72\_scores.cf based on the masscheck contributions while the second one is a "blind" rule promotion and [tagged build of SVN rules](#) for the masscheck area setup later.

```
25 2 * * * automc ~/svn/trunk/build/mkupdates/do-stable-update-with-scores
*~/svn/masses/rule-update-score-gen/do-nightly-rescore-example.sh
*~/svn/masses/rule-update-score-gen/generate-new-scores.sh
*uses ~/tmp/generate-new-scores for SVN work area
*sorts out the usable corpus from the latest 'SVN revision' at the top of the submitter's log file which should match the latest tagged build of SVN rules
*${REVISION} LINE 123 NEEDS IMPROVEMENT!!! THIS SVN REVISION NEEDS TO BE CLOSELY TIED TO THE REVISION THAT WAS STAGED IN THE MASSCHECK RSYNC DIR.
*checks the sorted corpus for a minimum number of valid contributors and ham/spam
*~/svn/trunk/build/mkupdates/mkupdate-with-scores
*uses ~/tmp/sa-mkupdate for SVN working area
*gets latest SVN ${REVISION} from rulesrc/scores/score-set*
*masses -> perl Makefile.PL && make (complete build of SA and test)
*perl hit-frequencies
*garescorer - compiles and runs it, requires build/pga
*sends email if not enough masscheck submitters or usable ham/spam for the latest SVN revision
*creates ${REVISION}.tar.gz ${REVISION}.tar.gz.sha1 and ${REVISION}.tar.gz.asc in /var/www/automc.spamassassin.org/updates for mirrors to pull
*updates DNS TXT entries [0-3].3.3.updates.spamassassin.org and 0.4.3.updates.spamassassin.org – versions >= 3.4.1 have a CNAME to 3.3.3.updates.spamassassin.org
*Script rewrite notes:
*Make each primary step modular since these steps are common in other scripts
*Should check for minimum contributors of ham/spam up front and not waste resources if requirements not met
*These 3 scripts above all share the same temp working dir. This should be determined from config file or relative path of user's home dir for flexibility.
*Should be able to run the ham/spam processing in parallel and merge the results together to cut this time in half
*Temp working dir for the corpus should be persistent so the rsync copy will be faster.
*Usuable corpus symlink setup could be improved. Invalid stale corpus should be removed into an archive/excluded dir.
```

```
30 8 * * * automc ~/svn/trunk/build/mkupdates/run_nightly > /var/www/automc.spamassassin.org/mkupdates/mkupdates.txt
*Currently ${SA_VERSION} = "3.4.2"
*${REVISION} = latest SVN revision THIS NEEDS TO BE ADDRESSED!!! NEED TO PREVENT REVISION FROM MESSING UP THE MASSCHECK PROCESSING.
*creates new rules/active.list
*commits new rules/active.list
*runs spamassassin lint against the updated rules and checks in a tagged version of 'sa-update_${SA_VERSION}_${TSTAMP}'
*commits "promotions validated" and emails dev@spamassassin.apache.org
*if the earlier daily update did not successfully produce the ${REVISION}.tar.gz* files
**creates ${REVISION}.tar.gz ${REVISION}.tar.gz.sha1 and ${REVISION}.tar.gz.asc in /var/www/automc.spamassassin.org/updates for mirrors to pull
**updates DNS TXT entries [0-3].3.3.updates.spamassassin.org and 0.4.3.updates.spamassassin.org – versions >= 3.4.1 have a CNAME to 3.3.3.updates.spamassassin.org
*Script rewrite notes:
*Uses many of the same primary steps previous section so reuse the code and not have to maintain multiple versions
*Should be turned into generic script that can be run on demand via SVN trigger/polling
```

## nitemc

These run shortly after the build/mkupdates/run\_nightly to setup the masscheck download area based on the latest [tagged build of SVN rules](#).

```
34 8 * * 0-5 automc *~/svn/nitemc/corpora_runs >> ~/rsync/corpus/nightly-versions.txt
36 8 * * 0-5 automc *~/svn/nitemc/extract_to_rsync_dir nightly ~/rsync/corpus/nightly-versions.txt
34 8 * * 6 automc *~/svn/nitemc/corpora_runs >> ~/rsync/corpus/weekly-versions.txt
36 8 * * 6 automc *~/svn/nitemc/extract_to_rsync_dir weekly ~/rsync/corpus/weekly-versions.txt
```

## ruleqa

This updates the web interface for <http://ruleqa.spamassassin.org>.

```
5 2-20 * * * automc . /etc/profile; */usr/local/bin/runRuleQArefresh.sh
*${HOME}/svn/masses/rule-qa/corpus-hourly --dir=${HOME}/rsync/corpus
*${HOME}/svn/masses/rule-qa/automc/gen_info_xml
*${HOME}/svn/masses/rule-qa/automc/ruleqa.cgi -refresh
```