

Container init process

Problem

To enable running multiple processes, containers require process/service management. This is normally provided by some kind of init task (e.g. from sysvinit).

There are Docker-compatible replacements for full-fledged init's. Unfortunately they require either custom init scripts or service configurations (https://wiki.gentoo.org/wiki/Comparison_of_init_systems). The process of migration from OS-provided OpenRC init scripts is time consuming and error prone.

Usage of system's default sysvinit is hampered by following shortcomings:

- it mostly does not respond to Unix signals, which are used by Docker to manage containers (most importantly: signal termination),
- it does not stop properly on when Docker requests container to stop
 - attempt to stop container with init as PID 1 ends with error code 137:

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS	P
ORTS	NAMES				
b755c0f1b1d8	gentoo-base	"/sbin/init"	About a minute ago	Exited (137) 9 seconds ago	
	gentoo-base				

- when invoking shutdown from within container, init process remains running afterwards, keeping container in running state:

```
# docker-compose top
gentoo-base
UID      PID      PPID     C    STIME   TTY      TIME      CMD
-----
root    3510    3489     0    17:40   ?        00:00:00  init [0]
```

Solution

Nevertheless it is possible to use sysvinit inside Docker container as an init process. Required steps are following:

1. Change sysvinit to exit init process on hard shutdown (runlevel 0) with following patch

```
--- sysvinit-3.01/src/init.c      2021-12-13 20:21:26.000000000 +0100
+++ sysvinit-3.01/src/init.c      2022-04-18 01:21:47.966751774 +0200
@@ -2367,6 +2367,11 @@
     read_inittab();
     fail_cancel();
     setproctitle("init [%c]", (int)runlevel);
+
+    /*
+     * Exit on halt - causes Docker container to stop.
+     */
+    if (runlevel == '0') exit(0);
+
+    Write_Runlevel_Log(runlevel);
```

On Gentoo it's enough to put this patch inside `/etc/portage/patches/sys-apps/sysvinit/exit-on-halt.patch` and reemerge sysvinit.

2. Change Docker signal for container termination to SIGINT and set appropriate action in inittab

Container's `docker-compose.yml`:

```
services:  
  gentoo-base:  
    ...  
    stop_signal: SIGINT
```

/etc/inittab inside container - replace reboot action with shutdown:

```
# What to do at the "Three Finger Salute".  
ca:12345:ctrlaltdel:/sbin/shutdown -h now
```