

1 Alpine Linux v3.8 ShRx

2
3 This is an image of Alpine Linux v3.8, which is configured as a Shell-
4 receive (ShRx) appliance. When this image boots, it mounts a FAT
5 filesystem that is shared between the host's `payload/build/` and the
6 guest's `/mnt/vdb1`, and auto-executes `/mnt/vdb1/main.sh`.

7
8
9 Creating this media

10
11 This media was created using a manual process. Therefore, the following
12 steps should be taken to re-create this image.

- 13
14 1. Boot the default ISO -- this comes directly from the upstream
15 source. In this case, it comes from Alpine Linux. This is
16 with the repository, so we can always rebuild.
17 2. Install the media to a tiny boot disk. Do not set up any
18 network interfaces.
19 3. Insert a cron entry to, upon boot, mount the FAT filesystem.
20 This entry will auto-delete itself after mounting that FS, and
21 execute `/mnt/build/main.sh`.
22 4. Poweroff the machine, and copy the image. This is the 'release'.

23
24
25 Running QEMU to create this media

26
27 Create the boot media:
28 `make os/boot.qcow2`

29
30 Attaching a payload

31
32 This project requires a `payload/` directory, to contain a set of files
33 for running upon the machine's first, and only first, boot. This
34 directory must contain another, called `build/`, which shares its files
35 with the virtual guest's filesystem.

36
37 `payload/` should also contain a Makefile, which has `clean` and `all`
38 targets defined. shrx-alpine makes use of these during its own build
39 process; one may call `make all` from shrx-alpine, and `payload/Makefile`
40 will also get called. Once all is said and done, `payload/build/main.sh`
41 must exist. This will be the entrypoint to the virtual machine.

42
43
44 Running the payload installer

45
46 The first time this appliance boots, it expects there to be a /dev/vdb1
47 partition. This must be supplied to QEMU like a file, using a VirtIO
48 interface. The following example assumes a `payload/build/` directory,
49 which contains a `build/main.sh` executable file.

50
51 To build the "application-installed" image, run `make`. This will build
52 the `release/boot.qcow2` target, which will copy `os/boot.qcow2` to
53 `release/boot.qcow2`. Then, the VM will be started, and the cron line
54 to run `main.sh` will go, then the cron line will be removed, then the
55 machine will poweroff.

56
57 At this time, the files in `release/` are ready to use for production.

58
59
60 Building the release images

61
62 `make release`, which is also aliased by `make all`.

63
64
65 Running the release images

66
67 The images found in `release/` can be run with QEMU. There is a makefile
68 target called `run` that will run them. However, it should be noted that
69 once the `run` target activates, the images are no longer considered
70 release-worthy, as some bits may be twiddled.

71
72 Production images should be exported as soon as a `make release` target
73 is run.

74
75 Running `make run-release` will take it one step further, and run the
76 image locally in the terminal. Running `make irun-release` will do the

```
77 same thing, but with QEMU's '-serial mon:stdio' set, which redirects
78 C-c to the guest.
79
80
81 `main.sh` notes
82
83 This file, to be in the `payload/build/` directory, should perform the
84 following tasks to configure an application server.
85
86 x. Configure network, DNS, logging destination
87 x. Add application service to /etc/init.d and start at boot
88 x. Produce boot config file (for virtio, qemu, or proxmox)
89
90 This script will only run once, on the first boot. It does so via the
91 cron mechanism, which self-deletes after running successfully.
92
93 This should return a non-zero exit code if anything went wrong.
94
95
96 Why a manual process?
97
98 A manual process is used to create the release image, because it's very
99 hard to script commands inside the OS. There are many ways to skin a cat,
100 and this provides a very quick turnaround (debugging `expect` scripts is
101 time-consuming) and a minimal, solid platform for letting our application
102 code configure the guest.
103
104 It would be great if we had an automated procedure for making this kind
105 of image. However, once we have a master boot image, it doesn't matter
106 any more how we got it.
107
108 It could be envisioned a similar process taking place for other guest
109 OSes. It may not be possible to script all of them.
110
111
112 References
113
114 1. https://en.wikibooks.org/wiki/QEMU/Devices/Storage
115
```