

# SkiffOS: Minimal Cross-compiled Linux for Embedded Containers

Christian Stewart  
Aperture Robotics LLC.  
christian@aperturerobotics.com

**Abstract**—Embedded Linux processors are increasingly used for real-time computing tasks such as robotics and Internet of Things (IoT). These applications require robust and reproducible behavior from the host OS, commonly achieved through immutable firmware stored in read-only memory. SkiffOS addresses these requirements with a minimal cross-compiled GNU/Linux system optimized for hosting containerized distributions and applications, and a configuration layering system for the Buildroot embedded cross-compiler tool which automatically re-targets system configurations to any platform or device. This approach cleanly separates the hardware support from the applications. The host system and containers are independently upgraded and backed-up Over-the-Air (OTA).

**Index Terms**—Container; Embedded Linux; Robotics; Single Board Computer

## I. INTRODUCTION

Single Board Computer (SBC) vendors distribute "board support packages" (BSPs) containing a mutable root filesystem with a package manager. Each platform requires a unique combination of kernel, bootloader, and firmware to operate, as a result BSPs typically lack portability.

Most Linux distributions use binary package managers to install and upgrade applications and system files, managed by distribution maintainers. Application dependencies co-exist with system dependencies, often creating version conflicts resulting in out-of-date and potentially vulnerable system software [1]. Package availability and versioning can vary across processor architectures.

Package managers operate on a mutable filesystem with imperative management commands (i.e. install, uninstall, upgrade) [2]. Each modification to the system configuration diverges from the known initial state, and can result in unpredictable changes in system behavior such as broken boot, inaccessible ssh, or malfunctioning drivers.

Declarative configuration models allow for reproducible behavior from the OS. Embedded Linux systems such as IoT sensors, industrial automation, TV set-top boxes, webcams, kiosks, and other products often bundle a GNU/Linux system into read-only firmware based on declarative configurations. This guarantees reproducible behavior, but comes with significant disadvantages in usability: software typically cannot be installed or upgraded without a system re-build and re-flash.

## II. SKIFFOS: FIRMWARE FOR HOSTING CONTAINERS

Buildroot is a project providing a Makefile and KConfig-based tool for automated cross-compilation of GNU/Linux systems, including toolchains, kernels, drivers, userspaces, and bootloaders [3]. It includes over 2500 software packages providing support for all major system utilities, and provides an easy to understand structure and configuration language allowing for rapid development across a diverse set of hardware configurations [4].

SkiffOS builds a minimal cross-compiled system for hosting containers across a diverse set of compute platforms, with minimal variance in the user experience. It emulates traditional firmware approaches with an immutable operating system image which produces reproducible behavior, while providing the utility of package-manager based GNU/Linux distributions with containerized environments attached to persistent storage.

Configuration layers bundle configuration files, Buildroot extension packages, system files, build hooks, and installation/utility scripts into named packages. Configurations are supplied for numerous target systems and virtualization environments. Hardware-specific performance optimizations are enabled at build time, producing a system image tuned for the target computer. Makefile targets are available for formatting and installing the system to boot media. Multiple system configurations can be compiled in parallel with "workspaces."

The SkiffOS Git repository can be embedded in extension projects as a sub-module and extended with out-of-tree configuration packages and overrides. Buildroot is referenced as a sub-module within the SkiffOS Git tree. Checksumming, package version pinning, and reproducible offline builds are used to ensure that a given SkiffOS commit will always produce identical output.

Existing system userspaces can be imported and used directly as container images. Package managers can then be used to install and manage software independently from the host system. Multiple distributions and containerized applications can be run in parallel. Workloads can be defined as container images to enhance portability and reproducibility [5]. Container management platforms such as Kubernetes and Docker Swarm can be used to remotely deploy and monitor workloads [6].

### III. IMPLEMENTATION

SkiffOS is available under the MIT license, and references Buildroot as a sub-module under the GPLv2 license, with a patch series providing additional features and bug fixes. Changes are frequently submitted upstream to the Buildroot mailing list.

As of Release 2020.11.7 the target support table is:

System	Config Package	Kernel
Apple Macbook (Intel)	apple/macbook	5.11.2
BananaPi M1	bananapi/m1	5.11.2
BananaPi M1+/Pro	bananapi/m1plus	5.11.2
BananaPi M2+	bananapi/m2plus	5.11.2
BananaPi M3	bananapi/m3	5.11.2
Docker Container	virt/docker	N/A
Intel x86/64	intel/x64	5.11.2
Msft Windows (WSL)	virt/wsl	N/A
NVIDIA Jetson Nano	jetson/nano	4.9.140
NVIDIA Jetson TX2	jetson/tx2	4.9.140
Odroid C2	odroid/c2	tb-5.9.16
Odroid C4	odroid/c4	tb-5.9.16
Odroid HC1/2, XU3/4	odroid/xu	tb-5.9.16
Odroid U	odroid/u	tb-5.9.16
OrangePi Lite	orangepi/lite	5.11.2
OrangePi Zero	orangepi/zero	5.11.2
PcDuino 3	pcduino/3	5.11.2
PcEngines APU2	pcengines/apu2	5.11.2
Pi 0	pi/0	5.4.72
Pi 1	pi/1	5.4.72
Pi 3 (and 1/2)	pi/3	5.4.72
Pi 4	pi/4	5.4.72
Pine64 H64	pine64/h64	5.8.0
PinePhone	pine64/phone	megi-5.9.11
Pinebook Pro	pine64/book	5.11.2
Qemu (VM)	virt/qemu	5.11.2
Virtualbox (VM)	virt/qemu	5.11.2
Rockpro64	pine64/rockpro64	5.9.0

Legal information and licenses for all dependencies can be bundled together with the `make legal-info` command. Some board support packages include proprietary binary blobs (typically firmware) and are denoted as "Proprietary" in the produced licensing information bundle.

#### A. Configuration Layers

The SkiffOS base configuration, builds a minimal host operating system with hardware support, OpenSSH server, and system management tools. Enabling additional configuration layers adds use-case specific functionality. Configurations are organized into logical units called "layers" with the following structure:

- **cflags**: additional target compiler flags
- **buildroot**: buildroot configuration fragments
- **buildroot\_ext**: buildroot extensions
- **buildroot\_patches**: buildroot package patches
- **extensions**: utility commands
- **hooks**: scripts hooking pre/post build steps
- **kernel**: kernel configuration fragments
- **kernel\_patches**: kernel .patch files
- **root\_overlay**: root overlay files
- **metadata**: metadata files
  - **commands**: targets in "extensions" makefile
  - **dependencies**: comma-separated layers
  - **description**: single-line description
  - **unlisted**: if exists, hidden from "help"
- **resources**: support files
- **scripts**: used by extensions and/or hooks
- **uboot**: u-boot configuration fragments
- **uboot\_patches**: u-boot .patch files

Configuration layers can override options set by previous layers, making it simple to re-target configurations to various compute platforms by merging with the desired hardware configuration layer. The set of desired configuration layers is defined as an ordered comma-separated list, for example:

```
SKIFF_CONFIG=pi/4,core/gentoo
```

SkiffOS can be extended [7] by adding it as a sub-module of a project repo. Project configurations can then be specified in additional configuration layers.

#### B. Persistent data partition

The persist partition contains a `skiff` directory with:

- **connections**: network-manager connections
- **core**: skiff-core configuration and state
- **docker**: docker state and storage
- **etc**: configuration tree overlay
- **hostname**: system hostname
- **journal**: systemd-journald system logs
- **keys**: public keys for access to "root" user
- **ssh**: ssh server keys and persistent configuration

System startup scripts mount the persist partition and create the file structure at boot time. Most configurations will create a memory swapfile to avoid failure due to an out-of-memory condition. OpenSSH is configured for public-key access with the persist "keys" directory or in the OS image.

The persist partition is automatically resized to fit the remainder of the available storage space on first boot. Bootloaders such as u-boot are sometimes copied to the beginning of the storage media. Some systems use more complex partition layouts mandated by the hardware vendor. Skiff and Buildroot's flexible configuration language supports this variance between target systems.

### C. Over-the-air (OTA) upgrades

The SkiffOS system typically consists of fewer than five files, including the root filesystem squashfs/initramfs, kernel image, and kernel modules squashfs. The kernel and immutable boot system can be atomically upgraded by replacing these files at run-time. The `push_image.sh` script is provided, using `rsync` and `ssh` to upload the updated files to a running system. In some cases, the firmware and/or bootloader is also upgraded by the script.

```
./scripts/push_image.sh root@my-device-ip
```

### D. Kconfig configuration fragments

Buildroot organizes software components into "packages." Packages can be enabled and configured using the Kconfig language. The available options can be explored using the `make menuconfig` or `make xconfig` configuration menus. Configuration options are specified in "fragments" that are merged together into a single ".config" file and provided to the Buildroot build system. The Buildroot build system manages merging together kconfig and u-boot configuration fragments and applying patches.

Example buildroot configuration fragment:

```
BR2_LINUX_KERNEL_DEFCONFIG="versatile"
BR2_LINUX_KERNEL_CUSTOM_VERSION_VALUE="5.11.2"
```

Buildroot and the Linux kernel use the Kconfig configuration language. The available options can be explored with `make br/menuconfig` and `make br/kernel-menuconfig` configuration menus. Configuration fragments are merged together in `SKIFF_CONFIG` order followed by lexicographic filename order.

This example kernel configuration fragment enables the ext3/ext4 filesystem:

```
CONFIG_EXT3_FS=m
CONFIG_EXT3_FS_SECURITY=y
CONFIG_EXT3_FS_XATTR=y
CONFIG_EXT3_POSIX_ACL=y
CONFIG_EXT4_FS=y
CONFIG_EXT4_FS_SECURITY=y
CONFIG_EXT4_POSIX_ACL=y
```

The `m` option denotes a feature built as a kernel module instead of "built-in."

### E. Buildroot extensions

Configuration layers can extend Buildroot with packages in the `buildroot_ext` directory. The `buildroot_patches`

directory contains patches for Buildroot packages, i.e. fixes for platform errata. For convenience, kernel and uboot patches can also be specified in the `kernel_patches` and `uboot_patches` directories.

### F. SkiffOS extensions

The `extensions` configuration layer directory contains a Makefile which implements custom commands made available to the user in the help screen:

```
cmd/pi/common/format: Format a SD card.
cmd/pi/common/install: Install to a SD card.
```

Commands are prefixed by their layer name and can be executed as Makefile targets, for example, `make cmd/pi/common/install`, and are declared in the `metadata/commands` text file:

```
format Format a SD card and install bootloader.
install Installs to a formatted SD card.
```

### G. Root filesystem overlay

The `root_overlay` trees are copied to the target filesystem image at the end of the build in the order that the layers were specified. This is used to add additional configuration files or scripts to the root SkiffOS system image. For example, the `pi/common` layer adds configuration under the `etc` directory, and firmware configuration under the `usr` directory.

### H. Temporary local overrides

Configuration overrides can be specified in the `overrides` directory tree, which contains additional configuration layers implicitly added to the build. For example, temporary local buildroot overrides for all workspaces can be declared in `overrides/buildroot`, and kernel configuration fragments affecting the `pi4` workspace alone can be declared in `overrides/workspaces/pi4/kernel`.

## IV. SKIFF CORE: CONTAINERIZED ENVIRONMENTS

Skiff Core includes the Docker containerization runtime and a Go program for automating the creation and initial setup of containerized environments. It can be enabled with the `skiff/core` configuration layer. User sessions are routed to the container assigned to their account. Multiple OS distributions can be installed simultaneously on a single machine. Existing userspaces including vendor-provided software images can be imported as container images.

The `skiff-core` binary is configured as the user shell, and intercepts incoming SSH sessions to redirect them to the corresponding container. The container setup process is displayed if the container(s) are not yet ready. The usual userspace init daemon, typically `systemd`, is run as PID 1, and standard approaches for defining `systemd` services and the `systemctl` CLI tool function identically to when running without containerization.

The "core" system can be updated or rolled-back independently from the "root" operating system. Mountpoints are used to mount user home directories and other temporary paths so that "docker export" and "docker save" include system files only. Containers are portable between machines of similar architecture without target-specific configuration.

The container system is configured with a YAML file:

```
containers:
  core:
    image: skiffos/skiff-core-gentoo:latest
    mounts:
      - /dev:/dev
      - /etc/resolv.conf:/etc/resolv.conf:ro
      - /mnt/persist/data:/home
    [...]
users:
  core:
    container: core
    containerUser: core
    auth: {copyRootKeys: true}
images:
  skiffos/skiff-core-gentoo:latest:
    pull:
      policy: ifnotexists
      registry: quay.io
    build:
      source: /opt/skiff/coreenv/base
```

Several OS distribution configurations are available as configuration layers. If a pre-built image is unavailable, or the pull section of `skiff-core.yaml` is empty, the system will instead build the included `Dockerfile`.

OS Distribution	Config Package	Website
Gentoo	core/gentoo	gentoo.org
Manjaro	core/manjaro	manjaro.org
NixOS	core/nixos	nixos.org
Ubuntu	skiff/core	ubuntu.org

Additional images optimized for specific use cases are available:

Image	Description
gentoo-lto-exwm	LTO, Apps, Emacs X11 Workflow
gentoo-kde	KDE Desktop w/ apps
gentoo-lto	O3, Graphite, Link-time optimize
nasa-cfs	Flight systems framework
nasa-fprime	Flight systems framework
pinephone-neon	KDE Neon for PinePhone
-manjaro-kde	Manjaro KDE for PinePhone

The NASA Fprime[8] and NASA cFS[9] images are

currently based on Ubuntu. To pull, for example, the `gentoo-lto` image, which includes the Gentoo core image with `gentooLTO` optimizations, the Docker CLI command is:

```
docker pull skiffos/skiff-core-gentoo-lto:latest
```

When using "Skiff Core," the Docker container engine is used, which leverages Linux namespaces and not virtualization or emulation, implementing "OS-level virtualization as opposed to hardware virtualization" [10]. The PID namespace is used to allow running the usual system initialization and management daemons as the privileged PID 1 within the container.

The results of [11] show "an almost negligible impact of the [Docker container] layer in terms of performance, if compared to native execution." While throughput is not significantly affected, the results of [6] show that namespacing can cause a measurable signal processing delay. To mitigate potential impacts of processing latency, most container isolation is disabled.

## V. COMPARISON WITH EXISTING APPROACHES

This section discusses several of the current approaches in use today, and their disadvantages when applied to embedded Linux and/or robotics development:

1) *Binary Package Distributions*: Traditional board-support packages use binary package distribution systems such as Debian's Advanced Package Tool (APT). Several key disadvantages of this approach are lack of hardware-specific optimizations, reliance on third-party infrastructure for builds and maintenance, inability to reproduce a system from source code, and difficulties with portability. The rolling nature of package upgrades introduces unpredictable behavior when upgrading a system, particularly when a long time has passed since the previous upgrade [1].

Current widely-used package management tools install system files, firmware, and user applications together into a single mutable "root" filesystem. Imperative install, remove, and upgrade commands instruct the package manager to perform operations on the system. Interrupted package manager operations might leave some files in a partially written state. If the affected (now corrupted) files are critical to system boot and/or reachability, the only recourse to fix the system may be to remove the boot media from the machine, connect it to a different device, and fix the issue manually.

SkiffOS addresses the concern of always having reproducible boot-up and reachability behavior in embedded systems by booting to an immutable root filesystem image, mounting persistent storage, and running user applications and operating systems inside lightweight Linux containers. This mitigates the risks of mid-upgrade power brownout: the immutable portion of the system is always reachable, allowing the user can connect to the container and fix the problem without physical intervention.

Optimization of compute performance and reduction of energy usage are important considerations for resource-constrained embedded devices (the "power budget"). For most robotics applications, the "power budget" is a significant factor in determining maximum range/endurance. Compilation of the operating system and applications from source allows fine-tuning of build output to the hardware to make maximum use of energy saving optimizations.

Backing up the system by creating a full bit-for-bit copy is a time-consuming process which often leads to infrequent backups and, as a result, occasional data loss due to storage card failure. SkiffOS provides an alternate approach, in which existing OS distributions can be imported as containers. Containers are easy to back up, restore, roll-back, and are portable between machines of similar architecture. Any vendor-provided BSP can be imported as a container and used without sacrifices to workflow or compatibility.

2) *NixOS*: NixOS [2] argues that imperative package managers destructively update the state of the system, leading to "inability to roll back changes easily, to run multiple versions of a package side-by-side, to reproduce a configuration deterministically on another machine, or to reliably upgrade a system." NixOS includes a declarative package manager which performs modifications according to a specification of target system state. Nix can perform atomic upgrades with the ability to roll-back.

SkiffOS is designed around the same principles of immutability and declarative configuration, but provides these guarantees through compilation of the system in advance, loading an immutable and ephemeral root system at boot-up. It also addresses the other issues described by the NixOS paper, including running multiple systems side-by-side in Linux containers, and reproducing a system configuration deterministically at a later time.

NixOS breaks the Filesystem Hierarchy Standard (FHS) [12] declared by the Linux foundation and followed by most distributions, instead using a flat symbolic link-based structure. This causes incompatibility with software not compiled by Nix. On the other hand, Buildroot [3] (and therefore SkiffOS) follow the FHS for compatibility with existing glibc-based binaries.

SkiffOS focuses on providing a minimal "shim" to abstract away the differences between hardware (particularly Single-Board Computers) such that the containers are portable to new platforms. It does not distribute complex packages like web browsers, graphical interfaces, and other user applications. This responsibility is left to the distributions running in Linux containers attached to persistent storage and/or customization of the cross-compiled system.

NixOS has been integrated with SkiffOS as a "core" container configuration. This approach uses Buildroot to manage the hardware specifics, and NixOS to handle declarative system configuration and rolling application

upgrades in containers. Software not compiled by Nix can still be run in a concurrent container.

3) *Buildroot*: The Buildroot [3] project provides a automated system cross-compiler. Users typically download a Buildroot release, create a configuration by hand, and store this along with their other project files. It focuses on providing a toolset for Embedded Linux developers to produce system images for use with a single target platform and/or product.

SkiffOS extends Buildroot to simplify this process with an easy to understand configuration layering architecture, which merges together the platform support configs with the selected components to configure the build automatically. The configuration layers can be stored externally to the SkiffOS tree. Existing Buildroot projects can be adapted as configuration layers and ported to any of the available target configurations.

4) *Yocto Project*: The Yocto Project [13] is closely comparable to Buildroot, and is primarily focused on the python-based "bitbake" tool, which is derived from Gentoo's "portage," with a focus on cross-compilation of complex embedded Linux systems. This is contrasted with Buildroot's focus on simplicity, using the Makefile and Kconfig architecture [14]. Yocto system configurations are fragmented into "overlay" repositories from various sources, while Buildroot focuses on a consolidated and curated package tree with strict and opinionated code style. Yocto has many more packages than Buildroot.

Buildroot [3] forms a linear commit history with release checkpoints, and Skiff releases pin the version of the Buildroot sub-module along with the versions of the device firmware, kernels, and board support files. A given Git checkout is reproducible forever [4]. Yocto setups may suffer from loss of availability of overlay repositories, or de-synchronization between changes to the overlay repositories and changes to the core system configurations.

## VI. CONCLUSION

SkiffOS offers reproducible system behavior, offline compilation, immutable host system for containers attached to persistent storage, board-specific build re-targeting via config layers, and Over-The-Air upgrade/downgrade. The host system and/or containers can be easily customized for specific use cases with cross-platform configuration layers. The "Skiff Core" workflow provides an easy way to import any existing Linux userspace such that users will likely not realize they are now working in a container. Background workloads can run in parallel Linux containers, independently isolated from the user workspace(s) including resource management and quality-of-service (QoS).

SkiffOS is available under the MIT license at:

<https://github.com/skiffos/skiffos>

Buildroot is available under the GPLv2 license at:

<https://buildroot.org>

## REFERENCES

- [1] R. D. Cosmo, S. Zacchiroli, and P. Trezentos, “Package upgrades in FOSS distributions: Details and challenges,” *CoRR*, vol. abs/0902.1610, 2009. arXiv: 0902.1610. [Online]. Available: <http://arxiv.org/abs/0902.1610>.
- [2] E. Dolstra and A. Löh, “Nixos: A purely functional linux distribution,” in *Proceedings of the 13th ACM SIGPLAN International Conference on Functional Programming*, ser. ICFP ’08, Victoria, BC, Canada: Association for Computing Machinery, 2008, pp. 367–378, ISBN: 9781595939197. DOI: 10.1145/1411204.1411255. [Online]. Available: <https://doi.org/10.1145/1411204.1411255>.
- [3] *Buildroot*. [Online]. Available: <https://buildroot.org/> (visited on 03/17/2021).
- [4] J. Diamond and K. Martin, “Managing a real-time embedded linux platform with buildroot,” *Fermi National Accelerator Lab*, 2015. [Online]. Available: <https://www.osti.gov/biblio/1250794>.
- [5] P. Z. Vaillancourt, J. E. Coulter, R. Knepper, and B. Barker, *Self-scaling clusters and reproducible containers to enable scientific computing*, 2020. arXiv: 2006.14784 [cs.DC].
- [6] Y. Wang and Q. Bao, “Adapting a container infrastructure for autonomous vehicle development,” in *2020 10th Annual Computing and Communication Workshop and Conference (CCWC)*, 2020, pp. 0182–0187. DOI: 10.1109/CCWC47524.2020.9031129.
- [7] *Skiffos extension template*. [Online]. Available: <https://github.com/skiffos/skiff-ext-example> (visited on 03/17/2021).
- [8] R. Bocchino, T. Canham, G. Watney, L. Reder, and J. Levison, “F prime: An open-source framework for small-scale flight software systems,” 2018.
- [9] D. McComas, “Nasa/gsfcs’s flight software core flight system,” in *Flight Software Workshop*, vol. 11, 2012.
- [10] J. M. Mbongue, D. T. Kwadjo, and C. Bobda, *Performance exploration of virtualization systems*, 2021. arXiv: 2103.07092 [cs.DC].
- [11] R. Morabito, “A performance evaluation of container technologies on internet of things devices,” in *2016 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, 2016, pp. 999–1000. DOI: 10.1109/INFOCOMW.2016.7562228.
- [12] *Filesystem hierarchy standard*. [Online]. Available: [https://refspecs.linuxfoundation.org/FHS\\_3.0/fhs-3.0.pdf](https://refspecs.linuxfoundation.org/FHS_3.0/fhs-3.0.pdf) (visited on 03/17/2021).
- [13] *Yocto project*. [Online]. Available: <https://www.yoctoproject.org/> (visited on 03/17/2021).
- [14] T. P. Alexandre Belloni, “Buildroot vs. openembedded/yocto project: A four hands discussion,” in *Embedded Linux Conference*, 2016. [Online]. Available: [https://events.static.linuxfound.org/sites/events/files/slides/belloni-petazzoni-buildroot-oe\\_0.pdf](https://events.static.linuxfound.org/sites/events/files/slides/belloni-petazzoni-buildroot-oe_0.pdf) (visited on 03/17/2021).