FireSim

Running a FireSim Simulation: Booting Linux and Running Accelerated ResNet50

https://fires.im

@firesimproject

MICRO Tutorial 2024

Speaker: Vikram Jain

Berkeley Architecture Research

Tutorial Roadmap

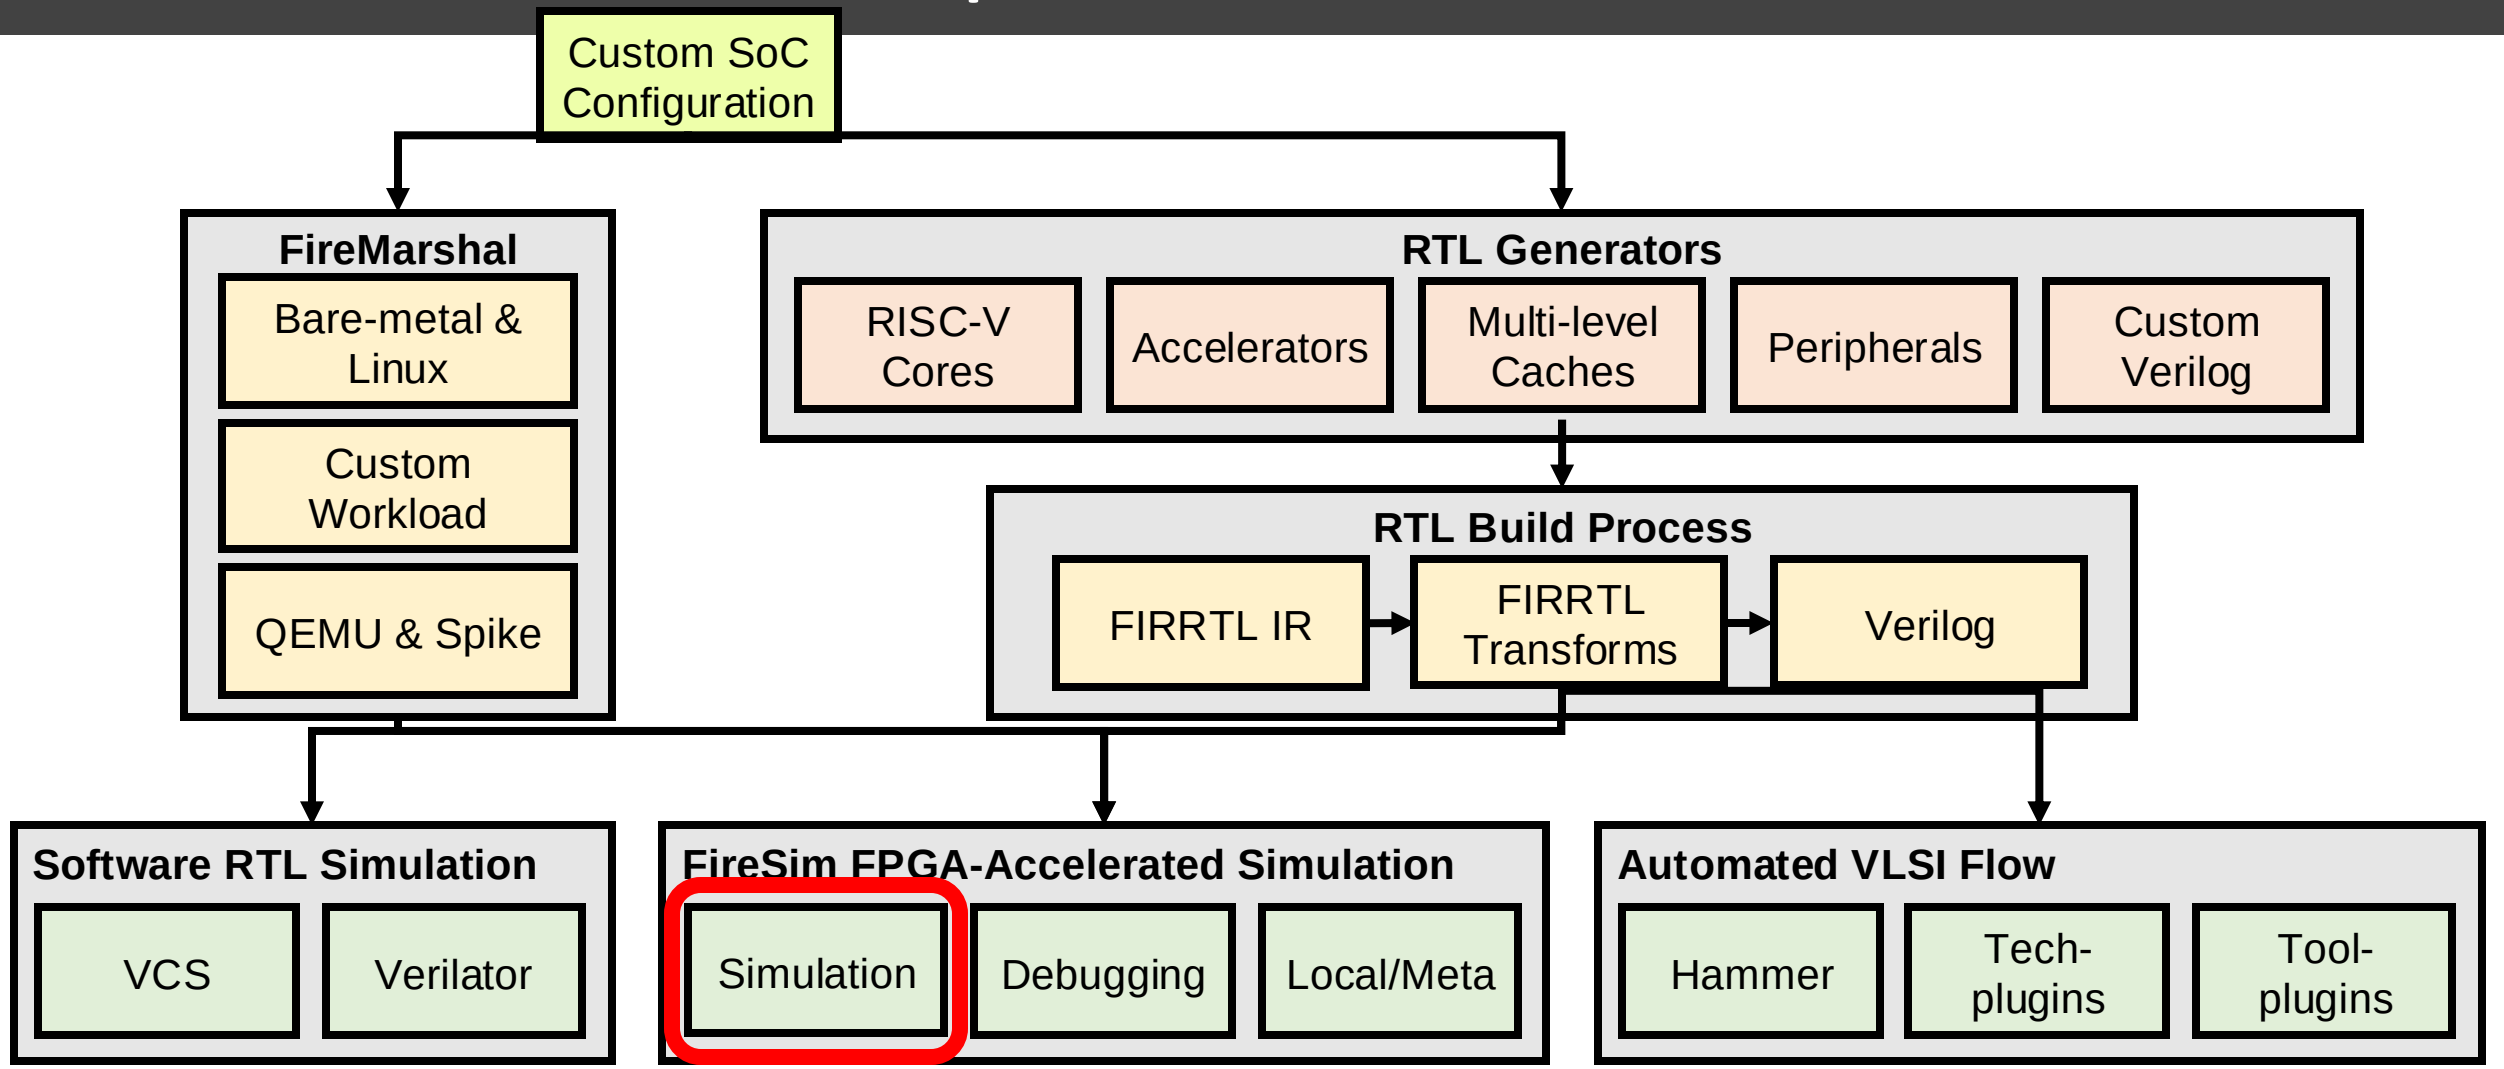

Agenda

- Configure and launch a simulation runfarm
- Boot Linux interactively on the target hardware
- Deploy new automated workloads
- Deploy a ResNet-50 image classification workload

Prefetching

- We will later be setting up and launching simulations
- To hide setup latency, edit `$FDIR/deploy/config_runtime.yaml` to match the following settings:

```
run_farm:
  recipe_arg_overrides:
    run_farm_hosts_to_use:
      - f1.4xlarge: 0
      - f1.2xlarge: 1

target_config:
    topology: no_net_config
    no_net_num_nodes: 1
    default_hw_config: firesim_rocket_singlecore_no_nic_l2_lbp

workload:
    workload_name: br-base-uniform.json
```

Prefetching

- We will later be setting up and launching simulations
- To hide setup latency:
  - Append the following entry to `config_hwdb.yaml`:

Make sure there are no duplicate entries

```
$ cd $FDIR/deploy
$ cat built-hwdb-entries/firesim_rocket_singlecore_no_nic_l2_lbp >>
    config_hwdb.yaml
```

  - Verify the format (with a *potentially different* unique AGFI ID):

```
firesim_rocket_singlecore_no_nic_l2_lbp:
    agfi: agfi-0e27eb94672e2f5a9
    deploy_triplet_override: null
    custom_runtime_config: null
```

In case `firesim buildbitstream` did not finish in time, a pre-populated entry is provided for you to use

Prefetching

- We will later be setting up and launching simulations
- To hide setup latency, run the following commands:

- First verify that you aren't inside another `tmux` session.
- If so, detach from the existing `tmux` session using `Ctrl+b` then `d`.

```
$ tmux new -s sim
$ firesim launchrunfarm && firesim infrasetup
```

What did we just do?

Runtime Configuration

What to simulate and *what infrastructure is required* is controlled by `$FDIR/deploy/config_runtime.yaml`

- Target-level: Assemble a simulated system from components
  - FPGA images of SoC hardware designs
  - Network topology
  - Workload definition
- Host-level: Specify which EC2 instances to use

config_runtime.yaml

The `run_farm` section specifies the number, type, and other launch parameters of instances to be managed

```
run_farm:
  base_recipe: run-farm-recipes/aws_ec2.yaml
  recipe_arg_overrides:
    run_farm_tag: mainrunfarm
    always_expand_run_farm: true
    launch_instances_timeout_minutes: 60
    run_instance_market: ondemand
    spot_interruption_behavior: terminate
    spot_max_price: ondemand
    default_simulation_dir: /home/centos
```

config_runtime.yaml

The `run_farm` section specifies the number, type, and other launch parameters of instances to be managed

```
run_farm:
  base_recipe: run-farm-recipes/aws_ec2.yaml
  recipe_arg_overrides:
    # ...
    run_farm_hosts_to_use:
      - f1.16xlarge: 0
      - f1.4xlarge: 0
      - f1.2xlarge: 0
      - m4.16xlarge: 0
      - z1d.3xlarge: 0
      - z1d.6xlarge: 0
      - z1d.12xlarge: 0
```

config_runtime.yaml

The `target_config` section specifies the high-level configuration of the system to simulate

```
target_config:
    topology: no_net_config
    no_net_num_nodes: 1
    link_latency: 6405
    switching_latency: 10
    net_bandwidth: 200
    profile_interval: -1
    default_hw_config: firesim_rocket_singlecore_no_nic_l2_lbp
    plusarg_passthrough: ""
```

`default_hw_config` references an entry from `config_hwdb.yaml`

config_runtime.yaml

The `workload` section specifies the software to be executed on the simulated nodes

```
workload:
    workload_name: br-base-uniform.json
    terminate_on_completion: no
    suffix_tag: null
```

Workload definitions live in `$FDIR/deploy/workloads/*.json`

Single-Node Simulation

What we modified in `config_runtime.yaml` earlier:

```
run_farm:
  recipe_arg_overrides:
    run_farm_hosts_to_use:
      - f1.2xlarge: 1

target_config:
    topology: no_net_config
    no_net_num_nodes: 1
    default_hw_config: firesim_rocket_singlecore_no_nic_l2_lbp

workload:
    workload_name: br-base-uniform.json
```

- Use a smaller f1.2xlarge instance (1 FPGA)
- Simulate one non-networked node without a switch model
- Load the single-core Rocket design without a NIC
- Boot Buildroot Linux workload

Launching Simulation Instances

```
$ firesim launchrunfarm
```

Already running in a `tmux` session; re-attach with `tmux attach -t sim`

```
FireSim Manager. Docs: https://docs.fires.im
Running: launchrunfarm

Waiting for instance boots: 0 f1.16xlarge
Waiting for instance boots: 1 f1.2xlarge
i-0c5c6894d0ac788af booted!
Waiting for instance boots: 0 f1.4xlarge
Waiting for instance boots: 0 m4.16xlarge
Waiting for instance boots: 0 z1d.12xlarge
Waiting for instance boots: 0 z1d.3xlarge
Waiting for instance boots: 0 z1d.6xlarge
The full log of this run is:
/home/centos/chipyard/sims/firesim/deploy/logs/2022-06-17--23-52-57-launchrunfarm-R50MKTLJ42O36MZZ.log
```

Launching Simulation Instances

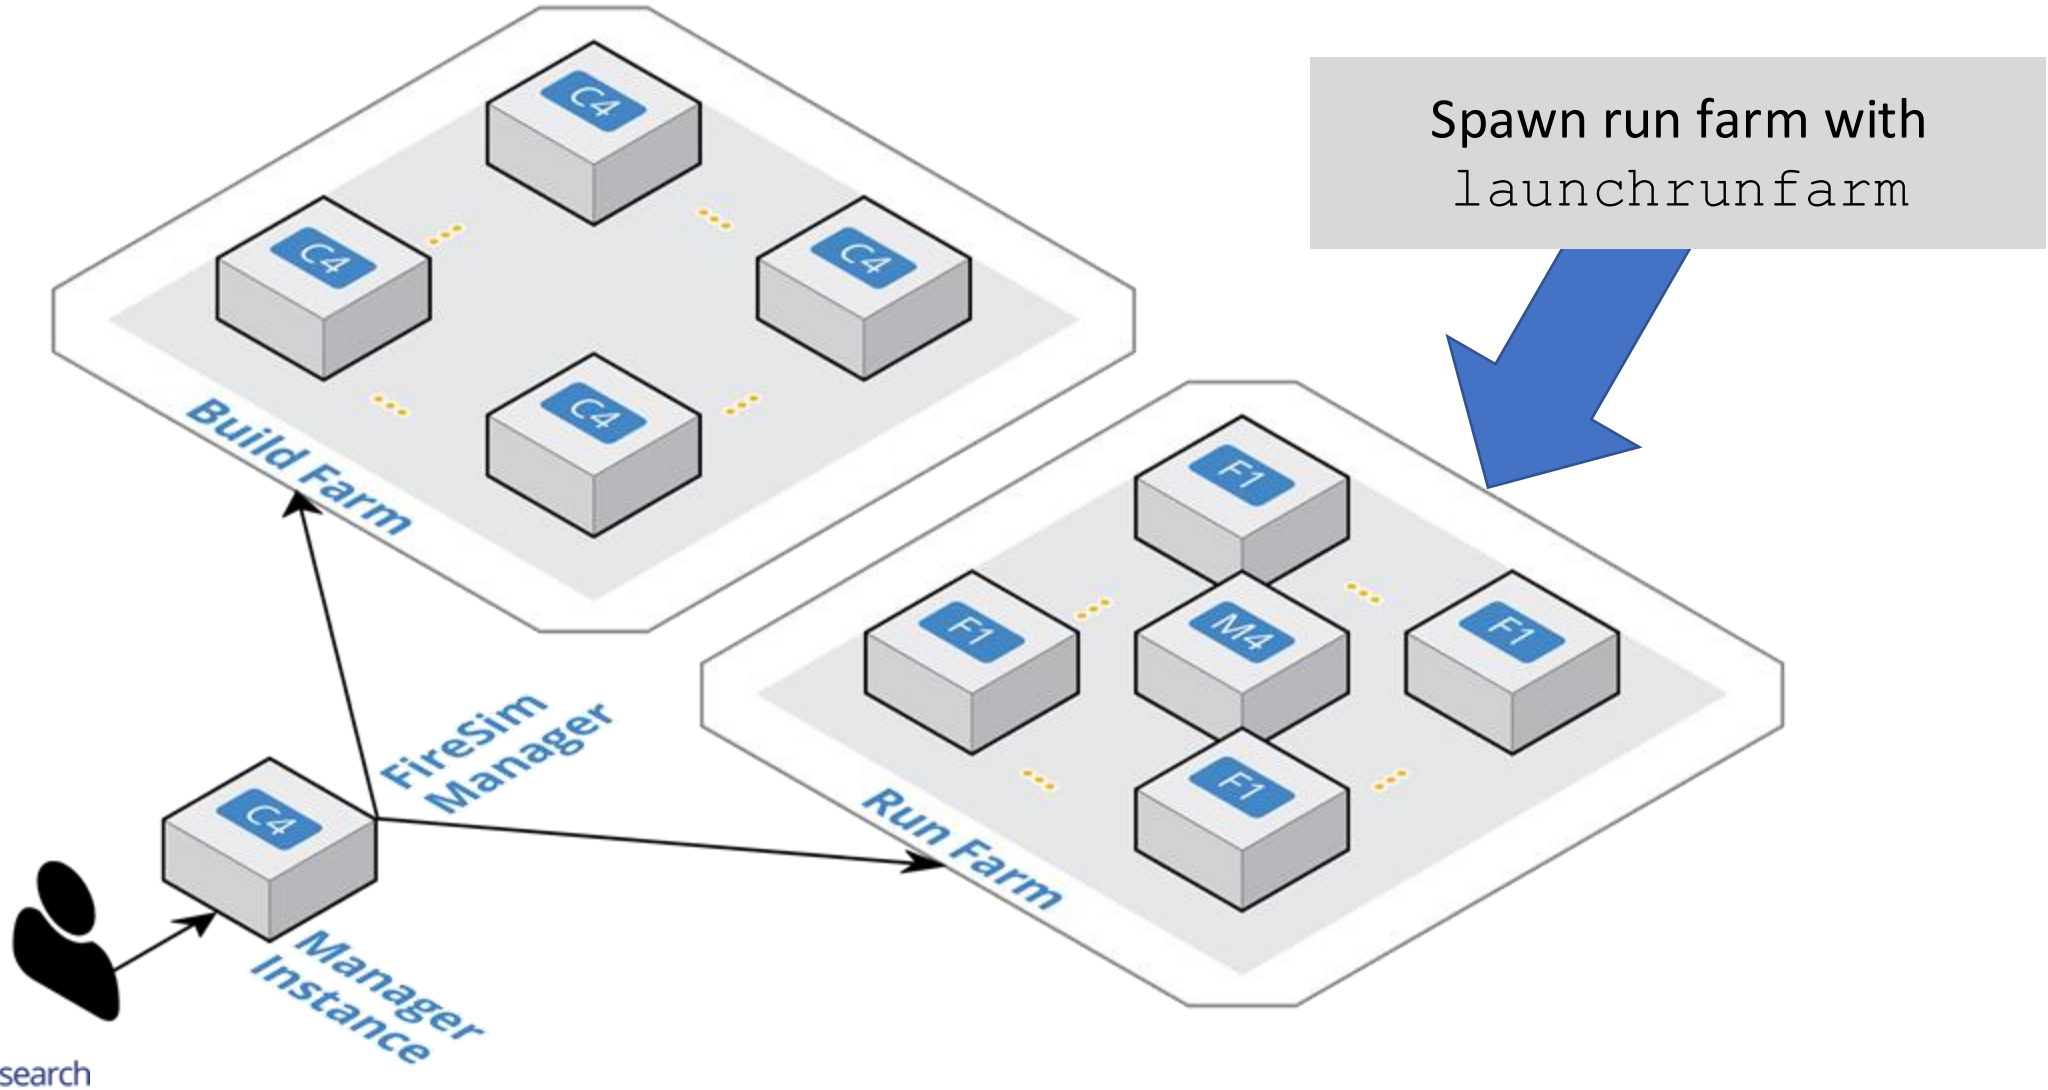

Deploying Simulation Infrastructure

```
$ firesim infrasetup
```

Already running!

This deploys various software prerequisites:

- Builds host-side simulation drivers for the specific build triplet
- Builds the switch model executable (if enabled)
- Collects information about simulation instances and transfers files
- Programs the FPGAs with the desired AGFIs

Deploying Simulation Infrastructure

```
$ firesim infrasetup
```

Already running!

```
FireSim Manager. Docs: https://docs.fires.im
Running: infrasetup

Building FPGA software driver for FireSim-WithDefaultFireSimBridges_WithFireSimHighPerfConfigTweaks_chipyard.RocketConfig-F90MHz_BaseF1Config
[192.168.3.52] Executing task 'instance_liveness'
[192.168.3.52] Checking if host instance is up...
[192.168.3.52] Executing task 'infrasetup_node_wrapper'
[192.168.3.52] Copying FPGA simulation infrastructure for slot: 0.
[192.168.3.52] Installing AWS FPGA SDK on remote nodes. Upstream hash: 1.12.0-72-gfed0aa6
[192.168.3.52] Unloading XRT-related Kernel Modules.
[192.168.3.52] Copying AWS FPGA XDMA driver to remote node.
[192.168.3.52] Unloading XDMA Driver Kernel Module.
[192.168.3.52] Loading XDMA Driver Kernel Module.
[192.168.3.52] Setting up remote node for qcow2 disk images.
[192.168.3.52] Loading NBD Kernel Module.
[192.168.3.52] Unloading NBD Kernel Module.
[192.168.3.52] Disconnecting all NBDs.
[192.168.3.52] Clearing FPGA Slot 0.
[192.168.3.52] Checking for Cleared FPGA Slot 0.
[192.168.3.52] Flashing FPGA Slot: 0 with agfi: agfi-0e27eb94672e2f5a9.
[192.168.3.52] Checking for Flashed FPGA Slot: 0 with agfi: agfi-0e27eb94672e2f5a9.
[192.168.3.52] Unloading XDMA Driver Kernel Module.
[192.168.3.52] Loading XDMA Driver Kernel Module.
[192.168.3.52] Starting Vivado hw_server.
[192.168.3.52] Starting Vivado virtual JTAG.
The full log of this run is:
/home/centos/chipyard/sims/firesim/deploy/logs/2022-06-18--00-13-05-infrasetup-SJJBIKPWYO20THF4.log
```

Deploying Simulation Infrastructure

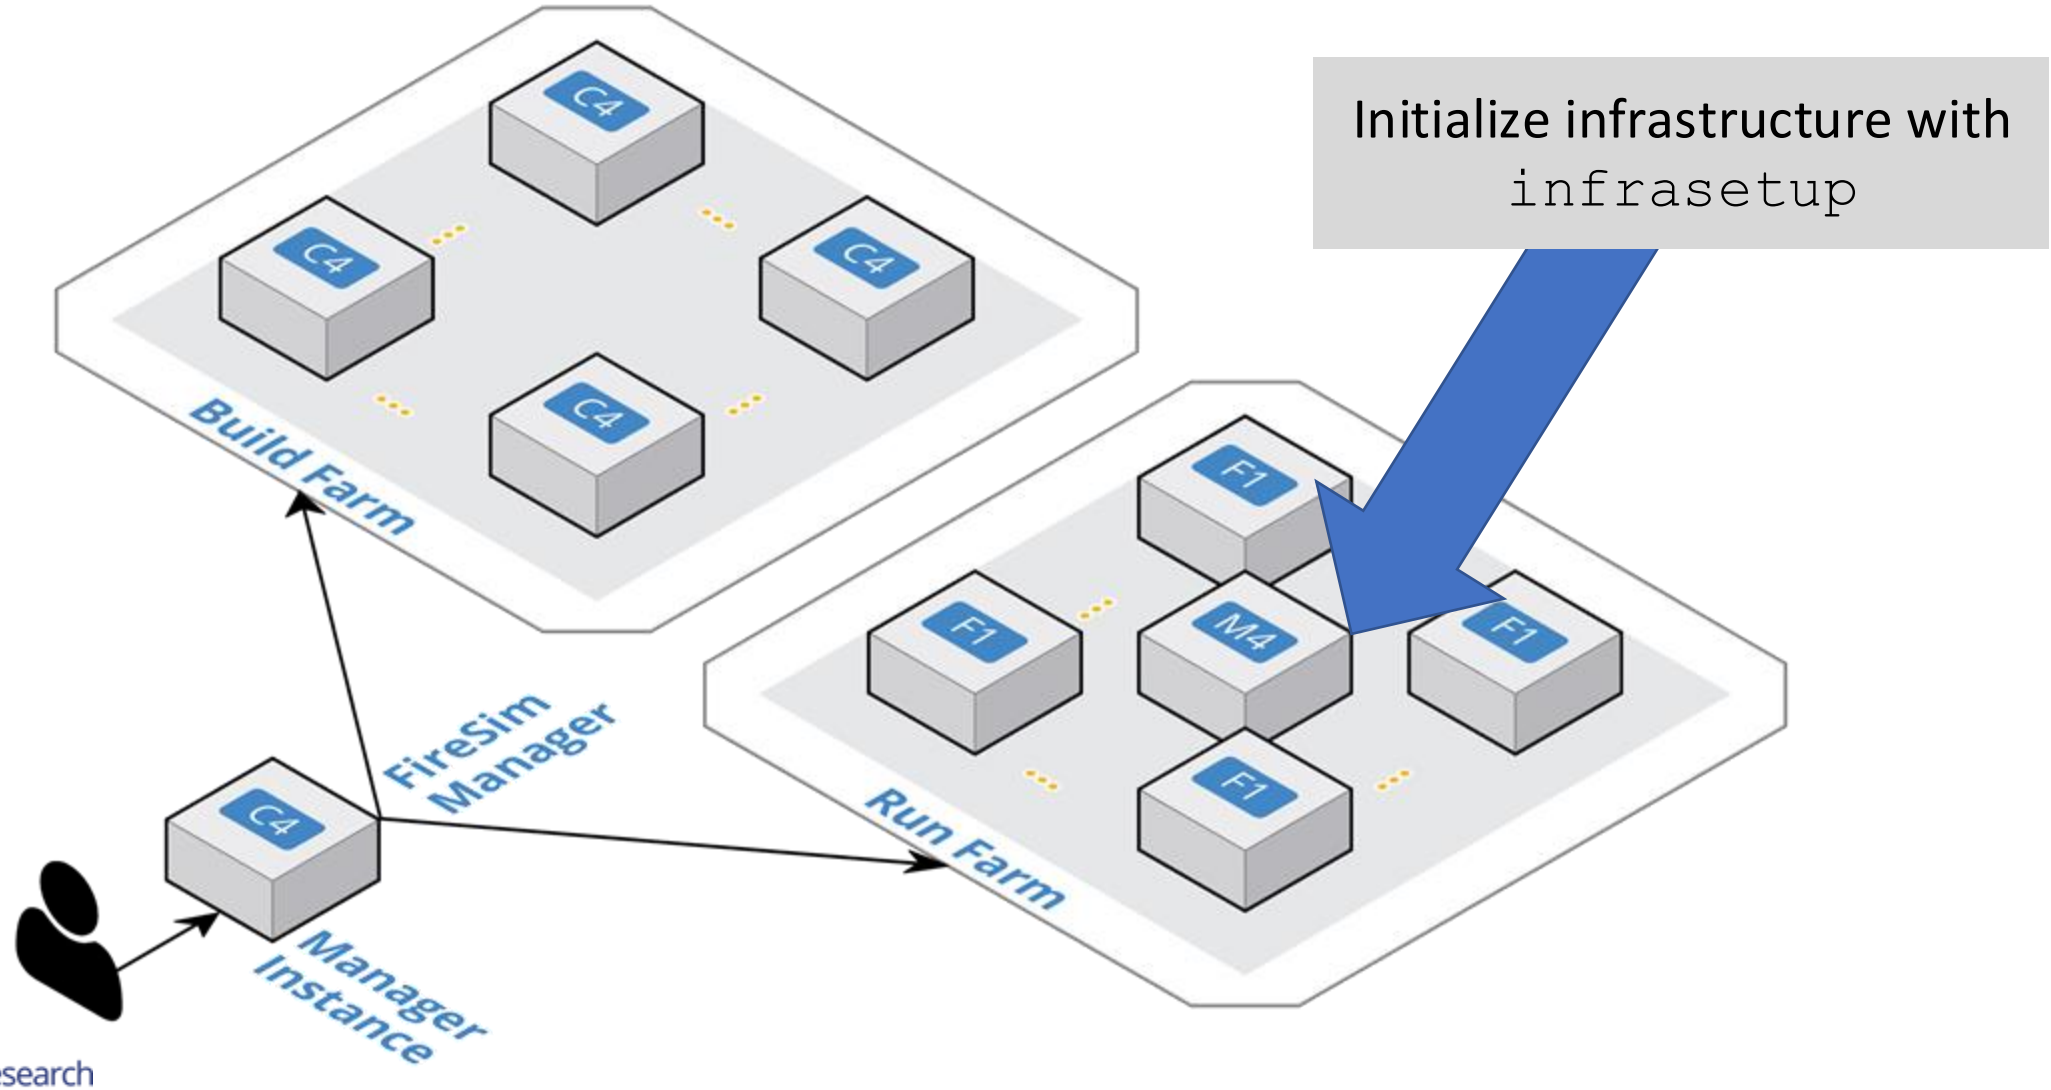

Running the Simulation

```
$ firesim runworkload
```

```
FireSim Manager. Docs: http://docs.fires.im
Running: runworkload

Creating the directory: /home/centos/chipyard/sims/firesim/deploy/results-
workload/2022-06-18--00-16-00-linux-uniform/
[192.168.3.52] Executing task 'instance_liveness'
[192.168.3.52] Checking if host instance is up...
[192.168.3.52] Executing task 'boot_switch_wrapper'
[192.168.3.52] Executing task 'boot_simulation_wrapper'
[192.168.3.52] Starting FPGA simulation for slot: 0.
[192.168.3.52] Executing task 'monitor_jobs_wrapper'
```

Running the Simulation

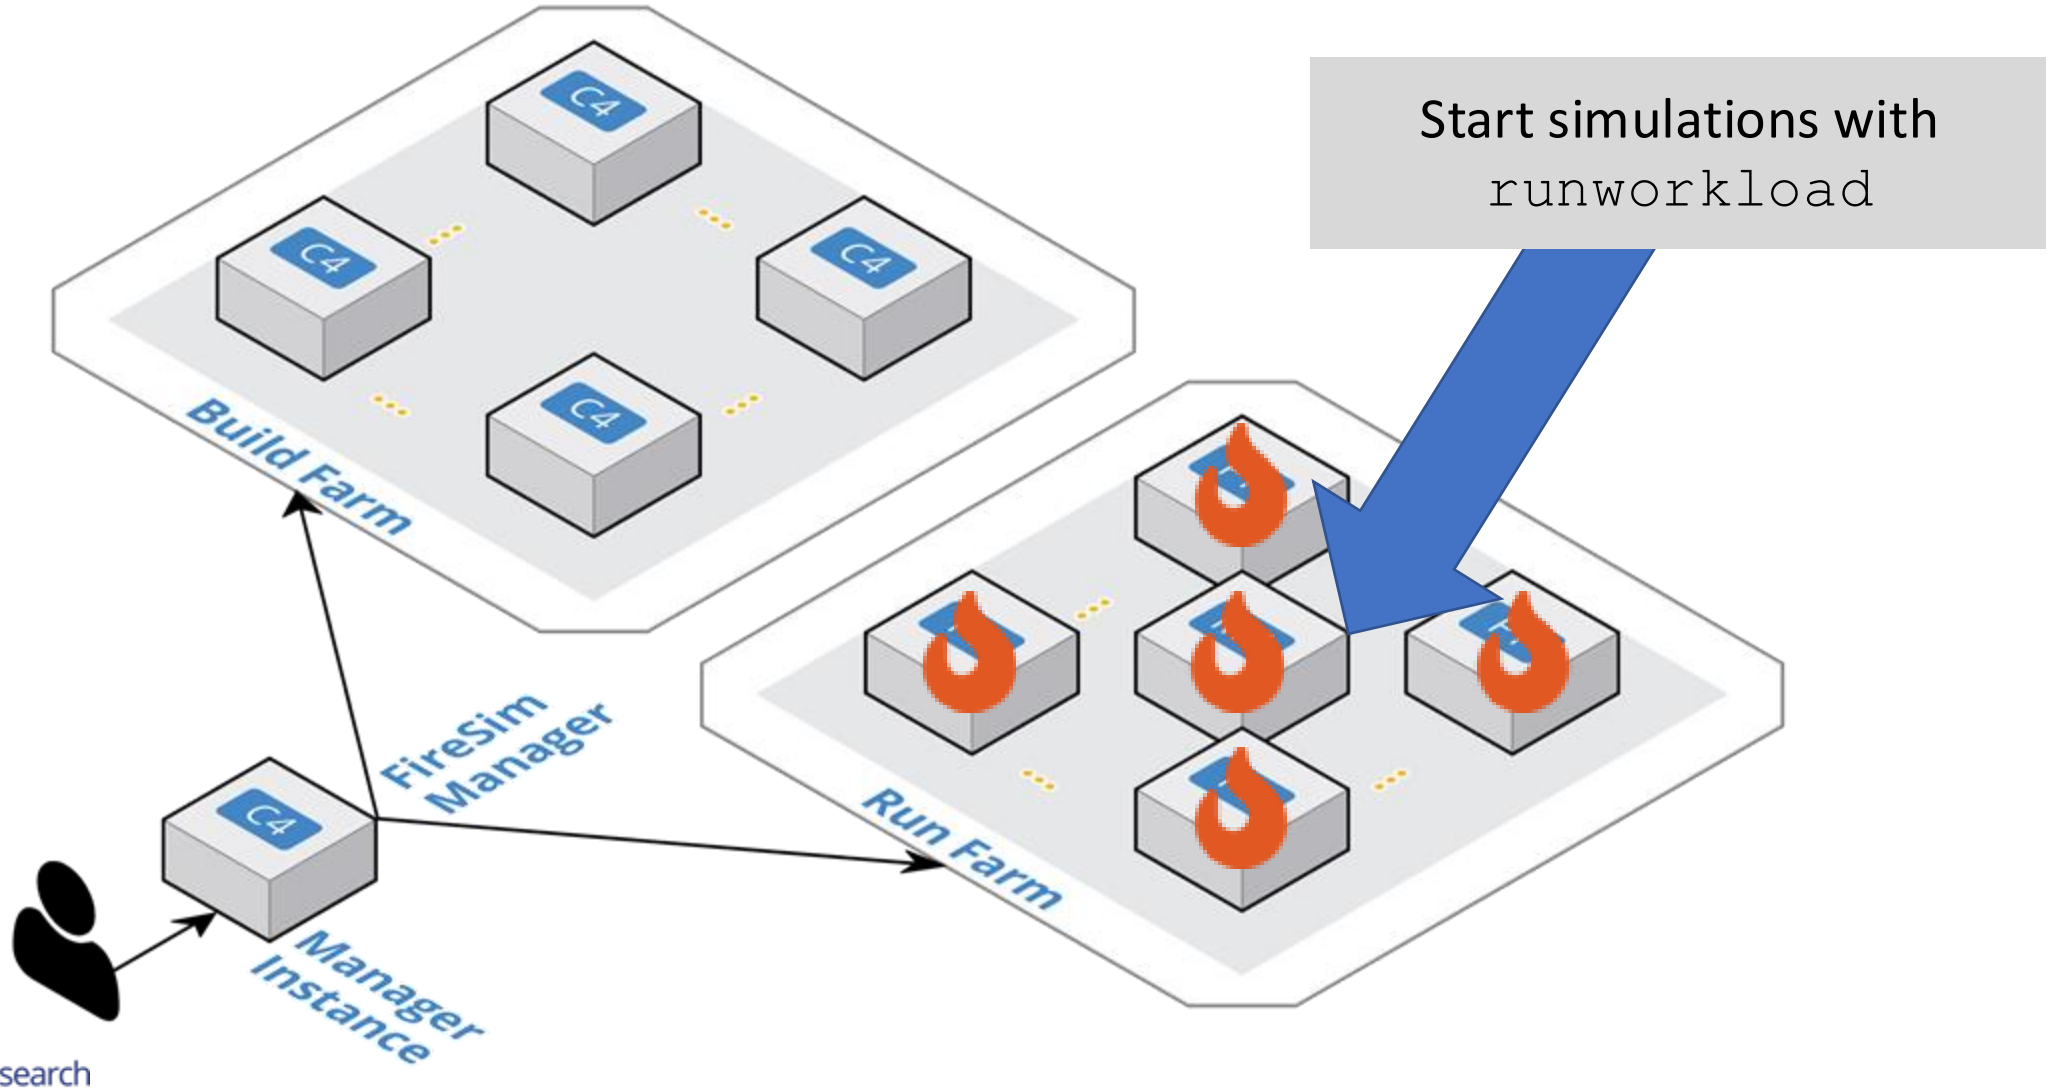

Monitoring the Simulation

You should see a live status report that refreshes periodically:

```
FireSim Simulation Status @ 2022-06-18 00:17:10.188191
--------------------------------------------------------------------------------
This workload's output is located in:
/home/centos/chipyard/sims/firesim/deploy/results-workload/2022-06-18--00-16-00-
linux-uniform/
This run's log is located in:
/home/centos/chipyard/sims/firesim/deploy/logs/2022-06-18--00-16-00-runworkload-
NEZCRUKBA2M44B9M.log
This status will update every 10s.
--------------------------------------------------------------------------------
Instances
--------------------------------------------------------------------------------
Hostname/IP: 192.168.3.52 | Terminated: False
--------------------------------------------------------------------------------
Simulated Switches
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
Simulated Nodes/Jobs
--------------------------------------------------------------------------------
Hostname/IP: 192.168.3.52 | Job: linux-uniform0 | Sim running: True
--------------------------------------------------------------------------------
Summary
--------------------------------------------------------------------------------
1/1 instances are still running.
1/1 simulations are still running.
--------------------------------------------------------------------------------
```

Interacting with the Simulation

Look for the run instance's IP address in the status:

```
FireSim Simulation Status @ 2022-06-18 00:17:10.188191
--------------------------------------------------------------------------------
This workload's output is located in:
/home/centos/chipyard/sims/firesim/deploy/results-workload/2022-06-18--00-16-00-
linux-uniform/
This run's log is located in:
/home/centos/chipyard/sims/firesim/deploy/logs/2022-06-18--00-16-00-runworkload-
NEZCRUKBA2M44B9M.log
This status will update every 10s.
--------------------------------------------------------------------------------
Instances
--------------------------------------------------------------------------------
Hostname/IP: 192.168.3.52 | Terminated: False
--------------------------------------------------------------------------------
Simulated Switches
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
Simulated Nodes/Jobs
--------------------------------------------------------------------------------
Hostname/IP: 192.168.3.52 | Job: linux-uniform0 | Sim running: True
--------------------------------------------------------------------------------
Summary
--------------------------------------------------------------------------------
1/1 instances are still running.
1/1 simulations are still running.
--------------------------------------------------------------------------------
```

Interacting with the Simulation

- On the *manager* instance, `ssh` into the run farm instance:

```
$ cd $FDIR
$ source sourceme-manager.sh
$ ssh 192.168.3.52
```

```
 ___ ___ ___   _     ___ _____  __    _  __ __ ___
| __| _ \/ __| /_\   |   \ __\ \ / /   /_\ |  \/  |_ _|
| _||  _/ (_ |/ _ \  | |) | _| \ V /   / _ \| |\/| || |
|_| |_|  \___/_/ \_\ |___/|___| \_/   /_/ \_\_|  |_|___|
AMI Version:          1.11.4
Xilinx Version:       2021.1
Readme:               /home/centos/src/README.md
AMI Release Notes:    /home/centos/src/RELEASE_NOTES.md
GUI/Cluster setup:    https://github.com/aws/aws-fpga/blob/master/developer_resources
```

- Then attach to the console of the simulated node:

```
$ screen -r fsim0
```

Logging Into the Simulated System

- Once Linux boots, the login prompt should appear over the console
- Log in as `root`

```
[    0.085714] EXT4-fs (iceblk): re-mounted. Opts: (null)
Starting syslogd: OK
Starting klogd: OK
Starting mdev... done.
Starting dropbear sshd: OK

Welcome to Buildroot
buildroot login: root
Password:
#
```

Logging Into the Simulated System

- Feel free to experiment with shell commands

```
# uname -a
# cat /proc/cpuinfo
# free -m
# vim
```

- When done, shut down the system

```
# poweroff -f
```

- This will also end the simulation

Finally, exit the `ssh` session with `Ctrl-d` to return to the manager instance

Running a ResNet50 simulation

Running ResNet50 on a Gemmini SoC

- Update `$FDIR/deploy/config_runtime.yaml` to point to our baremetal simulation of ResNet that we ran before

```
default_hw_config: firesim_gemmini_rocket_singlecore_no_nic
workload_name: resnet50-baremetal.json
```

- Then start another simulation:

```
$ firesim infrasetup && firesim runworkload
```

Make sure you do not launchfarm again

ResNet50

- ResNet50 is a CNN for embedded computer vision tasks

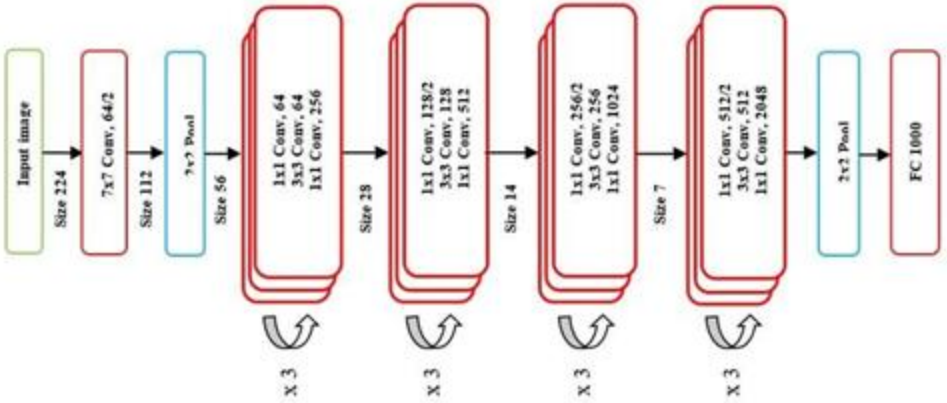

- Workload is predicting hardcoded images like…

Viewing the Simulation

Look for the run instance's IP address in the status:

```
FireSim Simulation Status @ 2022-06-18 00:17:10.188191
--------------------------------------------------------------------------------
This workload's output is located in:
/home/centos/chipyard/sims/firesim/deploy/results-workload/2022-06-18--00-16-00-
linux-uniform/
This run's log is located in:
/home/centos/chipyard/sims/firesim/deploy/logs/2022-06-18--00-16-00-runworkload-
NEZCRUKBA2M44B9M.log
This status will update every 10s.
--------------------------------------------------------------------------------
Instances
--------------------------------------------------------------------------------
Hostname/IP: 192.168.3.52 | Terminated: False
--------------------------------------------------------------------------------
Simulated Switches
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
Simulated Nodes/Jobs
--------------------------------------------------------------------------------
Hostname/IP: 192.168.3.52 | Job: resnet50-baremetal0 | Sim running: True
--------------------------------------------------------------------------------
Summary
--------------------------------------------------------------------------------
1/1 instances are still running.
1/1 simulations are still running.
--------------------------------------------------------------------------------
```

Viewing the Simulation

- On the *manager* instance, `ssh` into the run farm instance:

```
$ cd $FDIR
$ source sourceme-manager.sh
$ ssh 192.168.3.52
```

```
 ___ ___ ___   _     ___ _____   __    _   __ __ ___
| __| _ \/ __| /_\   |   \| __\ \ / /   /_\ |  \/  |_ _|
| _||  _/ (_ |/ _ \  | |) | _| \ V /   / _ \| |\/| || |
|_| |_|  \___/_/ \_\ |___/|___| \_/   /_/ \_\_|  |_|___|
AMI Version:          1.11.4
Xilinx Version:       2021.1
Readme:               /home/centos/src/README.md
AMI Release Notes:    /home/centos/src/RELEASE_NOTES.md
GUI/Cluster setup:    https://github.com/aws/aws-fpga/blob/master/developer_resources
```

- Then attach to the console of the simulated node:

```
$ screen -r fsim0
```

What you should see

- You predicted the images!

```
. . .
matmul_51: 299923
matmul_52: 342640
matmul_53: 235853
Prediction: 75 (score: 127)
Prediction: 900 (score: 127)
Prediction: 641 (score: 104)
Prediction: 897 (score: 99)

Total cycles: 88615903 (100%)
Matmul cycles: 7667548 (8%)
Im2col cycles: 0 (0%)
Conv cycles: 1778884 (2%)
Pooling cycles: 0 (0%)
Depthwise convolution cycles: 76301614 (86%)
Res add cycles: 758772 (0%)
Other cycles: 2109085 (2%)
PASS
```

What you should see

- You predicted the images!

```
. . .
matmul_51: 299923
matmul_52: 342640
matmul_53: 235853
Prediction: 75 (score: 127)
Prediction: 900 (score: 127)
```

- Detach from the screen by doing `Ctrl-a+d`
- Then exit out of the simulation instance with `Ctrl-d` to return to the manager instance

```
Pooling cycles: 0 (0%)
Depthwise convolution cycles: 76301614 (86%)
Res add cycles: 758772 (0%)
Other cycles: 2109085 (2%)
PASS
```

Capturing Results

- Once the workload terminates automatically, the results are copied to the manager instance:

```
FireSim Simulation Exited Successfully. See results in:
/home/centos/chipyard/sims/firesim/deploy/results-workload/2022-06-17--00-38-00-resnet50-baremetal/
```

- The exact directory path will contain a different timestamp
- Console output recorded in `resnet50-baremetal0/uartlog`
- HW configuration in `resnet50-baremetal0/HW_CFG_SUMMARY`