



# Using LAVA V2 for advanced KVM testing

Riku Voipio <[riku.voipio@linaro.org](mailto:riku.voipio@linaro.org)>





**Linaro  
connect**

Las Vegas 2016

ENGINEERS  
AND DEVICES  
WORKING  
TOGETHER

# Background

- LAVA supports KVM testing using vmgroups feature
- Easy to use, but rather restricted
  - Hardcoded qemu command line
  - Can't use for xen/libvirt/kvmtool for virtualization testing
  - Images need to be tailored for testing
- LAVA V2, aka pipeline to rescue
  - Explicit control of everything
  - After some bugs fixing, finally there!
- LAVA v2 Training session LAS16-TR05 wednesday 3pm



**Linaro  
connect**

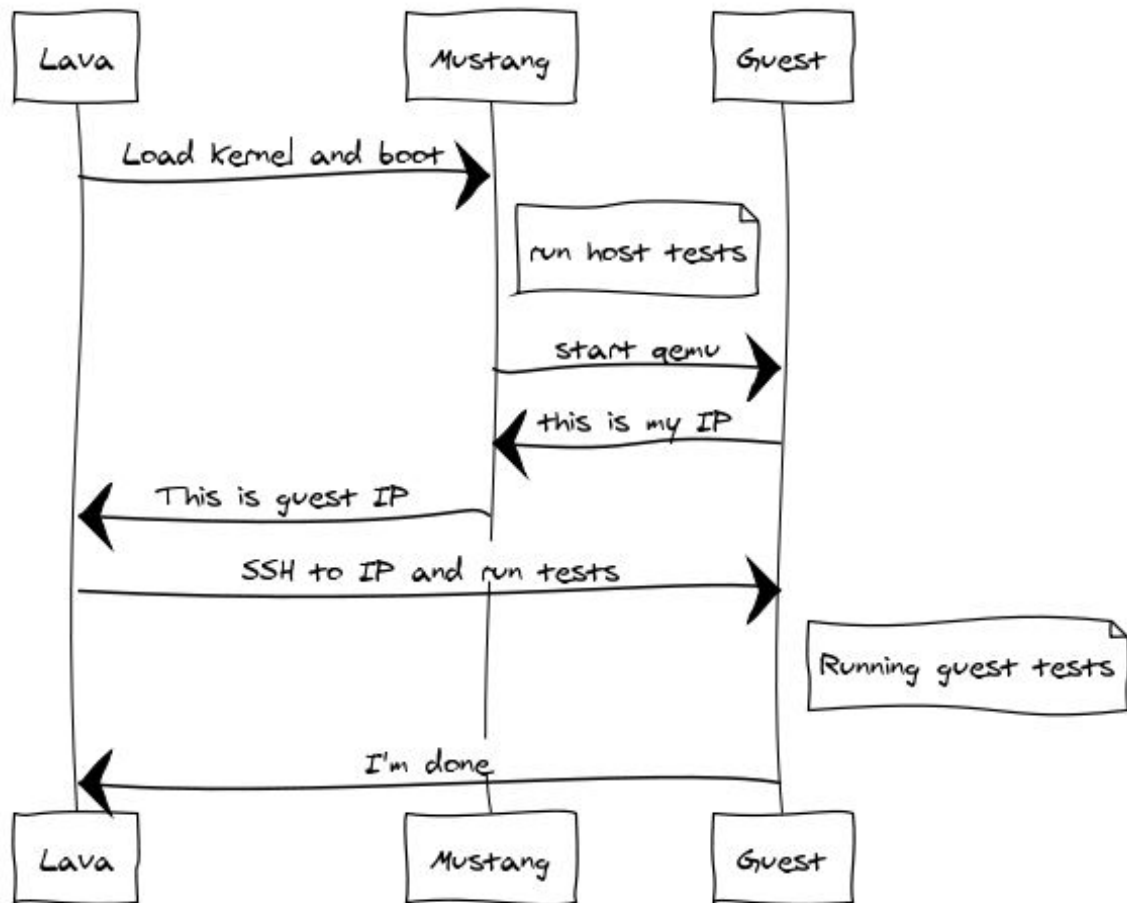
Las Vegas 2016

ENGINEERS  
AND DEVICES  
WORKING  
TOGETHER

# Basic pipeline KVM test

- Test that armv8 host boots an armv7 guest
- <http://people.linaro.org/~riku.voipio/kvm-pipeline.yaml>
- Interesting sections
  - Deploy: host
    - The kernel and nfsroot to use for test
  - Boot: host
    - Boot host using u-boot and wait for prompt
  - Boot: guest
    - Wait for guest IP to connect via SSH
  - Test: host
    - Start kvm guest with lava-test-shell definition
    - Send lava the IP of guest once guest is up
  - Test: guest
    - The tests run on guest

# KVM pipeline job





**Linaro  
connect**

Las Vegas 2016

ENGINEERS  
AND DEVICES  
WORKING  
TOGETHER

# Guest deployment steps

1. Download Guest OS using parameters
  - OS, UEFI firmware, possibly kernel
2. Configure Guest OS
  - <https://git.linaro.org/qa/test-definitions.git/blob/HEAD:/common/scripts/kvm-cloud/cloudinit.txt>
  - Install SSH keys for LAVA to log in
  - Phone home when boot is finished
3. Start QEMU
  - <https://git.linaro.org/qa/test-definitions.git/blob/HEAD:/common/scripts/kvm-cloud/start-kvm.sh#l105>
  - Uses daemonize option and saves boot log to file



**Linaro  
connect**

Las Vegas 2016

ENGINEERS  
AND DEVICES  
WORKING  
TOGETHER

# Download guest OS

- Any cloud image is good to go
  - Defined in job template

```
parameters:
```

```
    GUEST_ARCH: armv7l
```

```
    GUEST_IMAGE:
```

```
https://cloud-images.ubuntu.com/.../xenial-server-cloudimg-  
armhf-disk1.img
```

```
    GUEST_KERNEL:
```

```
https://cloud-images.ubuntu.com/.../xenial-server-cloudimg-  
armhf-vmlinuz-lpae
```

# Configure guest OS

- Supports any linux image with cloud-init support
- <https://cloudinit.readthedocs.io/en/latest/>
- Ubuntu, Debian, Fedora, CentOS tested

```
#cloud-config
```

```
users:
```

```
- name: linaro
  ssh-authorized-keys:
    - LAVA_KEY
  sudo: ['ALL=(ALL) NOPASSWD:ALL']
  groups: sudo
  shell: /bin/bash
```

```
runcmd:
```

```
- cp /home/linaro/.ssh/authorized_keys
  /root/.ssh/authorized_keys
- chown root:root /root/.ssh/authorized_keys
```

```
phone_home:
```

```
url: http://LOCALIP:8080/
```





**Linaro  
connect**

Las Vegas 2016

ENGINEERS  
AND DEVICES  
WORKING  
TOGETHER

# Start Qemu

- Start-kvm.yaml
  - <https://git.linaro.org/qa/test-definitions.git/blob/HEAD:/ubuntu/start-kvm.yaml>
  - Install and start qemu
  - Wait for guest start and signal lava with IP
- Stop-guest.yaml
  - Signal shutdown of guest
- Against the LAVA documentation advice to inline lava-send calls
  - See the link for the rationale:  
<https://validation.linaro.org/static/docs/v2/writing-multinode.html#controlling-synchronisation-from-the-test-shell>





# Bits from Multinode protocol



- Guest deploy uses lava-wait to stand-by until read
- Host will send the IP using lava-send
- The messages obviously need to match
- <https://validation.linaro.org/static/docs/v2/writing-multinode.html>

```
- deploy:  
  role:  
    - guest  
  connection: ssh  
  os: ubuntu  
  protocols:  
    lava-multinode:  
      - action: prepare-scp-overlay  
        request: lava-wait  
        messageID: ipv4  
        message:  
          ipaddr: $ipaddr
```



# Viewing Results

← → ↻ [https://validation.linaro.org/results/1113661/0\\_smoketest-guest-armv7l](https://validation.linaro.org/results/1113661/0_smoketest-guest-armv7l) ☆  

 Home Dashboard ▾ Results ▾ Scheduler ▾ API ▾ Help ▾ Instance: production  Riku Voipio ▾








LAVA / Results / Test job 1113661 / Suite 0\_smoketest-guest-armv7l

## Results for test suite 0\_smoketest-guest-armv7l - Test Job 1113660.1

Exports 

Test suite export:  CSV or  YAML

Show  entries

Name 	Testset 	Result 	Measurement 	Units 	logged 	Bug Links 
linux-linaro-ubuntu-pwd	—	✓ pass	—	—	Sept. 12, 2016, 1:52 p.m.	[0]
linux-linaro-ubuntu-uname	—	✓ pass	—	—	Sept. 12, 2016, 1:52 p.m.	[0]
linux-linaro-ubuntu-vmstat	—	✓ pass	—	—	Sept. 12, 2016, 1:52 p.m.	[0]
linux-linaro-ubuntu-ifconfig	—	✓ pass	—	—	Sept. 12, 2016, 1:52 p.m.	[0]
linux-linaro-ubuntu-lscpu	—	✓ pass	—	—	Sept. 12, 2016, 1:52 p.m.	[0]
linux-linaro-ubuntu-lsb_release	—	✓ pass	—	—	Sept. 12, 2016, 1:52 p.m.	[0]





**Linaro**  
**connect**  
Las Vegas 2016

## Tips and bits for LAVA V2

- Use unique values for each timeout
  - Helps you find out which timeout triggered!
- Use inline test-definitions while developing
- Keep lava-dispatcher codebase available...
- Try and find out!

ENGINEERS  
AND DEVICES  
WORKING  
TOGETHER





**Linaro**  
**connect**

Las Vegas 2016

ENGINEERS  
AND DEVICES  
WORKING  
TOGETHER

# Future directions

- Replicate this with Xen, libvirt, kvmtool
- Web frontend for results
- Selecting what tests to run
  - Previously we used only hackbench
  - With the pipeline job structure, we can choose any of the existing QA tests
- Migration tests (more than 1 node)
- Support cloud images / cloud-init with LAVA directly?



# Thank You

#LAS16

For further information: [www.linaro.org](http://www.linaro.org)

LAS16 keynotes and videos on: [connect.linaro.org](http://connect.linaro.org)

