

# OpenStack on ARMv8

Date Huang, NCU  
Gene Kuo, NCTU

# About Us

## 黃宇強 Date Huang

- **NCU 轉職 NCTU 中**
- **不務正業的資訊安全準研究生**
- **バカ baka**

## 郭靖 Gene Kuo

- **NCTU**
- **不務正業的電子工程系學生**
- **OpenStack App Hackathon Mentor (X)**

# Outline

**Intro**

**ARMv8 Pros and Cons**

**Possible Use Case**

**OpenStack on ARM Issue**

**Future Plans**

# Intro

- **Why DozenCloud**
  - **Native ARM VPS**
  - **Native ARMv8 vs QEMU Emulator**
- **Why ARMv8**
  - **Pros**
  - **Cons**

# Native vs Emulate

- **Lightweight**
- **Massive Speed Improvement**
  - **Cache**
  - **Pipeline**
  - **Out-of-Order Execution**
- **QEMU has lots of bugs**
  - **JIT compiler has some problems**



# Supported devices

## Android

- Android 4.4 to Android 6.0.1  
*(Android N will not supported until the official Android release)*
- Preferred resolution of 720x1280 pixels *(Not optimized for tablet)*
- Strong internet connection (Wi-Fi, 3G, or 4G)
- GPS and Location Services
- *Intel CPUs are not supported*

## Supported devices

Android

x86 不支援哭哭喔

x86 is not supported QQ

- Android 4.4 to Android 6.0.1  
*(Android N will not supported until the official Android release)*
- Preferred resolution of 720x1280 pixels *(Not optimized for tablet)*
- Strong internet connection (Wi-Fi, 3G, or 4G)
- GPS and Location Services
- *Intel CPUs are not supported*



ARMv8

# Pros

- **Native ARM Development Environment**
  - Compare with QEMU
- **Better Concurrency Performance**
  - 96 cores
- **Lower Power Consumption**

```

 1 [||||| 9.6%] 25 [|| 2.4%] 49 [|||| 4.8%] 73 [|| 3.0%]
 2 [||||| 10.3%] 26 [ 0.0%] 50 [|| 6.6%] 74 [|| 4.3%]
 3 [|| 1.8%] 27 [ 0.0%] 51 [||||| 9.0%] 75 [|| 4.2%]
 4 [ 0.0%] 28 [ 0.0%] 52 [|||| 8.4%] 76 [|| 2.4%]
 5 [|| 0.6%] 29 [ 0.0%] 53 [|||| 5.5%] 77 [||||||| 15.1%]
 6 [|| 6.1%] 30 [ 0.0%] 54 [|||| 7.8%] 78 [|||| 7.2%]
 7 [|| 2.4%] 31 [|| 3.0%] 55 [|||| 4.2%] 79 [|| 4.8%]
 8 [|| 0.6%] 32 [|| 2.4%] 56 [||||| 9.6%] 80 [||||||| 15.2%]
 9 [|| 3.6%] 33 [|| 0.6%] 57 [||||| 6.7%] 81 [|||| 7.9%]
10 [ 0.0%] 34 [ 0.0%] 58 [|| 3.7%] 82 [|||| 6.6%]
11 [ 0.0%] 35 [|| 3.0%] 59 [||||| 12.2%] 83 [|||| 8.4%]
12 [ 0.0%] 36 [ 0.0%] 60 [||||| 9.0%] 84 [|||| 6.0%]
13 [|| 1.8%] 37 [|||| 4.8%] 61 [|| 3.0%] 85 [||||| 5.4%]
14 [|| 1.8%] 38 [ 0.0%] 62 [|| 3.0%] 86 [|||| 7.4%]
15 [ 0.0%] 39 [|| 1.8%] 63 [|| 4.2%] 87 [ 0.0%]
16 [|| 2.4%] 40 [ 0.0%] 64 [||||| 12.5%] 88 [|| 1.8%]
17 [ 0.0%] 41 [|| 1.8%] 65 [|||| 9.1%] 89 [|||| 4.2%]
18 [|| 1.2%] 42 [|| 3.0%] 66 [||||||| 15.5%] 90 [|||| 9.1%]
19 [ 0.0%] 43 [|| 2.4%] 67 [||||||| 17.0%] 91 [|| 4.2%]
20 [ 0.0%] 44 [ 0.0%] 68 [||||| 7.9%] 92 [|| 3.0%]
21 [ 0.0%] 45 [ 0.0%] 69 [||||||| 16.0%] 93 [|||| 9.1%]
22 [ 0.0%] 46 [ 0.0%] 70 [|||| 8.0%] 94 [|||| 3.7%]
23 [|| 2.4%] 47 [ 0.0%] 71 [|||| 8.4%] 95 [||||| 9.6%]
24 [ 0.0%] 48 [ 0.0%] 72 [|||| 10.3%] 96 [|| 6.0%]
Mem[||||||||||||||||||||| 41.0G/126G] Tasks: 705, 702 thr; 3 running
Swp[ 0K/0K] Load average: 3.07 3.79 4.07
Uptime: 5 days, 13:52:36

```

PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
40698	date	20	0	8912	2240	0	S	0.0	0.0	0:00.00	(sd-pam)
25705	keystone	20	0	232M	128M	9764	S	0.0	0.1	0:00.00	(wsgi:keystone-ad -k start
25706	keystone	20	0	232M	128M	9764	S	0.0	0.1	0:00.88	(wsgi:keystone-ad -k start
25708	keystone	20	0	232M	128M	9764	S	0.0	0.1	4:29.52	(wsgi:keystone-ad -k start
25610	keystone	20	0	232M	128M	9764	S	0.0	0.1	4:30.47	(wsgi:keystone-ad -k start
25763	keystone	20	0	232M	128M	9764	S	0.0	0.1	0:00.00	(wsgi:keystone-ad -k start
25764	keystone	20	0	232M	128M	9764	S	0.0	0.1	0:00.88	(wsgi:keystone-ad -k start
25765	keystone	20	0	232M	128M	9764	S	0.0	0.1	4:31.60	(wsgi:keystone-ad -k start
25611	keystone	20	0	232M	128M	9764	S	0.0	0.1	4:32.55	(wsgi:keystone-ad -k start
25663	keystone	20	0	232M	128M	9764	S	0.0	0.1	0:00.00	(wsgi:keystone-ad -k start
25665	keystone	20	0	232M	128M	9764	S	0.0	0.1	0:00.88	(wsgi:keystone-ad -k start

F1help F2Setup F3Search F4Filter F5Tree F6SortBy F7Nice -F8Nice +F9Kill F10Quit

# Cons

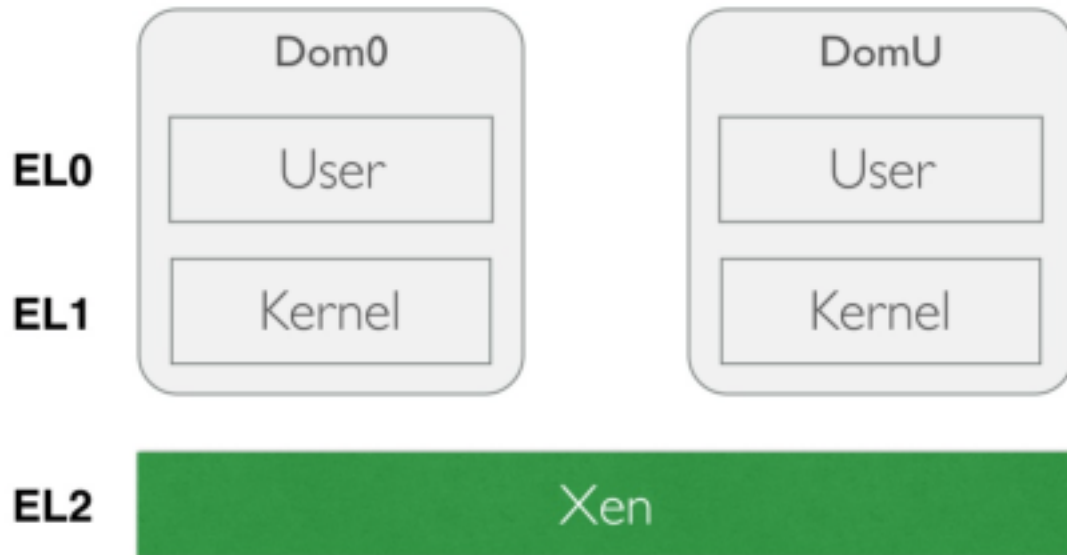
- **Single Core Performance Worse than Other Architecture**
- **Software isn't Ready**
  - **Compiling from source or patch it.**
- **Need More Optimization**
  - **KVM on ARM is worse than Xen on ARM**
  - **KVM on x86 is better than Xen on x86 (about 3x)**

THE #1 PROGRAMMER EXCUSE  
FOR LEGITIMATELY SLACKING OFF:

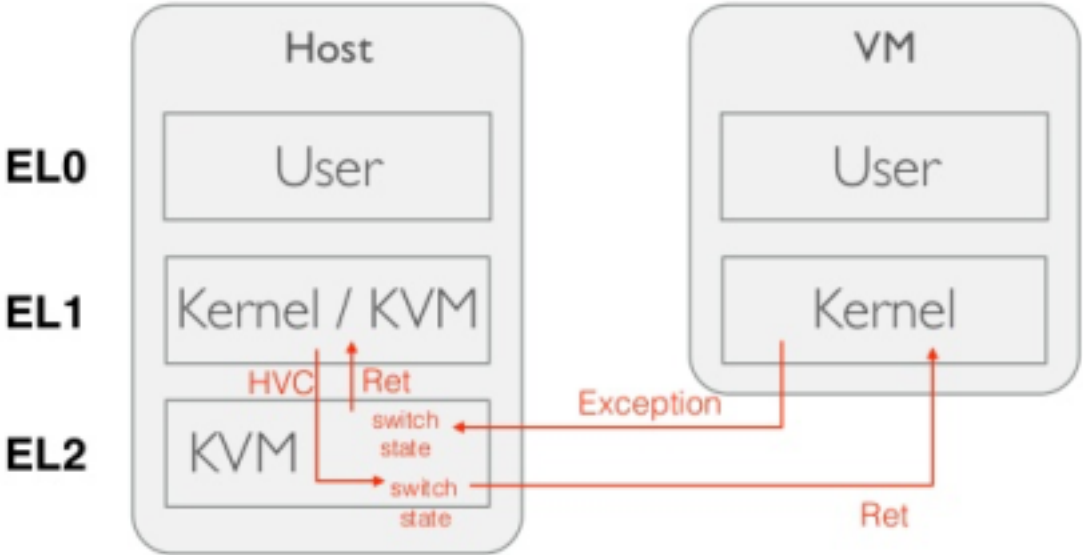
"MY CODE'S COMPILING."



# Type-1 Hypervisor



# KVM/ARM VM Exits



# Possible Use Case

- **ARM Virtual Private Server**
  - **OpenStack**
- **Native ARM CI server**
  - **Gitlab-CI**
- **Storage Server**
  - **Ceph**
- **CDN**



# Native CI

main.s 235 Bytes

```
1  .global _start
2
3  _start:
4      // fd = stdout
5      mov x0, #0
6      // buf
7      ldr x1, =msg
8      // len
9      mov x2, len
10     // syscall num
11     mov x8, #64
12     // syscall
13     svc #0
14
15     mov x0, #0
16     mov x8, #93
17     svc #0
18
19     .equ len, (end - msg)
20     msg:
21         .ascii "Hello World\n"
22     end:
```

```
gitlab-ci-multi-runner 1.3.0-beta.14.ga5bbe6e (a5bbe6e)
Using Docker executor with image gitlab-runner-prebuilt-arm64:a5bbe6e ...
Pulling docker image gitlab-runner-prebuilt-arm64:a5bbe6e ...
WARNING: Cannot pull the latest version of image gitlab-runner-prebuilt-arm64:a5bbe6e : Error:
WARNING: Locally found image will be used instead.
Running on runner-f8476e7d-project-1-concurrent-0 via gitlab...
Cloning repository...
Cloning into '/builds/date/armv8-hello'...
Checking out 368deb41 as master...
$ apt-get update
Hit:1 http://free.nchc.org.tw/ubuntu-ports xenial InRelease
Get:2 http://free.nchc.org.tw/ubuntu-ports xenial-updates InRelease [94.5 kB]
Hit:3 http://free.nchc.org.tw/ubuntu-ports xenial-security InRelease
Get:4 http://free.nchc.org.tw/ubuntu-ports xenial-updates/main Sources [185 kB]
Get:5 http://free.nchc.org.tw/ubuntu-ports xenial-updates/universe Sources [56.3 kB]
Get:6 http://free.nchc.org.tw/ubuntu-ports xenial-updates/main arm64 Packages [278 kB]
Get:7 http://free.nchc.org.tw/ubuntu-ports xenial-updates/universe arm64 Packages [151 kB]
Fetched 677 kB in 1s (648 kB/s)
Reading package lists...
$ apt-get install -y build-essential
Reading package lists...
Building dependency tree...
Reading state information...
build-essential is already the newest version (12.1ubuntu2).
0 upgraded, 0 newly installed, 0 to remove and 0 not upgraded.
$ make clean
rm -rf main.o a.out
$ make
gcc -c -nostdlib -o main.o main.s
gcc -nostdlib main.o
$ ./a.out

Build succeeded
```

```
_start:
```

```
    // fd = stdout  
    mov x0, #0  
    // buf  
    ldr x1, =msg  
    // len  
    mov x2, len  
    // syscall num  
    mov x8, #64  
    // syscall  
    svc #0  
  
    mov x0, #0  
    mov x8, #93  
    svc #0
```

```
$ apt-get install -y build-essential  
Reading package lists...  
Building dependency tree...  
Reading state information...  
build-essential is already the newest  
0 upgraded, 0 newly installed, 0 to r  
$ make clean  
rm -rf main.o a.out  
$ make  
gcc -c -nostdlib -o main.o main.s  
gcc -nostdlib main.o  
$ ./a.out  
  
Build succeeded
```

# Issue

- **Kernel Panic**
  - **Software Isn't Ready**
- **OpenStack QEMU & Libvirt Driver**
  - **32 on 64 KVM Problem**
- **Mitaka Experimental UEFI Support Bug**

32 on 64 KVM

# ARMv7 vs ARMv8

- **ARMv7**

- **32 bits**
- **ISA: ARM, Thumb**
- **CPU Mode: SVC, USR, FIQ, .....**

- **ARMv8**

- **32 and 64 bits**
- **ISA: AArch64, AArch32**
  - Only ARMv8-A has AArch64 ISA support

# AArch32 vs AArch64

- **AArch32**

- **Partially compatible with ARMv7**
- **Only Reserve co-processor CP10, CP10, CP11, CP14, and CP15 etc.**
- **Including A32, T32, which correspond to ARM ISA and Thumb ISA respectively**

- **AArch64**

- **No more co-processor**
- **Remove conditional execution**

Non-UEFI

# Non-UEFI VMs

**Need to assign kernel image and initramfs**

**Kernel patches will be problematic**

**Security issue**



UEFI

```
25 <vcpu placement='static'>16</vcpu>
26 <cputune>
27   <shares>16384</shares>
28 </cputune>
29 <resource>
30   <partition>/machine</partition>
31 </resource>
32 <os>
33   <type arch='aarch64' machine='virt'>hvm</type>
34   <loader readonly='yes' type='pflash'>/usr/share/AAVMF/AAVMF_CODE.fd</loader>
35   <nvram template='/usr/share/AAVMF/AAVMF_CODE.fd'>/var/lib/libvirt/qemu/nvram/inst
ance-0000000a_VARS.fd</nvram>
36   <boot dev='hd' />
37 </os>
38 <features>
39   <acpi/>
40   <apic/>
41   <gic version='3' />
42 </features>
43 <cpu mode='host-passthrough'>
44   <topology sockets='16' cores='1' threads='1' />
45 </cpu>
46 <clock offset='utc'>
47   <timer name='pit' tickpolicy='delay' />
48   <timer name='rtc' tickpolicy='catchup' />
```

# Patch

```
218     try:
219         self._domain.undefineFlags(
220             libvirt.VIR_DOMAIN_UNDEFINE_MANAGED_SAVE)
221     except libvirt.libvirtError:
222         LOG.debug("Error from libvirt during undefineFlags. %d"
223                 "Retrying with undefine", self.id)
224         self._domain.undefine()

218     try:
219         self._domain.undefineFlags(
220             libvirt.VIR_DOMAIN_UNDEFINE_MANAGED_SAVE |
221             libvirt.VIR_DOMAIN_UNDEFINE_NVRAM)
222     except libvirt.libvirtError:
223         LOG.debug("Error from libvirt during undefineFlags. %d"
224                 "Retrying with undefine", self.id)
225         self._domain.undefine()
```



VIA 9GAG.COM



9GAG is your best source of fun.

# Future Plans

- **Closed Beta at August**
  - **ARMv8 Development Course**
  - **Computer Security CTF Contest**
- **Open Beta**
  - **No idea**
  - **Financial Problem**
- **Native Android**

# 特別感謝

技嘉科技

鈺登科技

國立中央大學資工系先進防禦實驗室

國立交通大學

**GIGABYTE™**

**Edge-core**

NETWORKS

國立交通大學資工系



國立交通大學  
National Chiao Tung University



國立中央大學  
National Central University

# Demo

- **GIGABYTE R150-T60 ARMv8 Server**
  - Cavium ThunderX ARMv8 processors
  - DDR4 128GB RAM
- **Ubuntu 16.04**
- **OpenStack Mitaka**



# Summary

ARMv8 software ecosystem isn't complete yet

Use with CAUTION!!!

OpenStack on ARMv8 is usable but not complete

VNC, UEFI issue, and etc.....

x86\_64 still better for general purpose usage



# Ref

<http://wiki.csie.ncku.edu.tw/embedded/ARMv8>

[https://en.wikipedia.org/wiki/ARM\\_architecture](https://en.wikipedia.org/wiki/ARM_architecture)

<https://support.pokemongo.nianticlabs.com/hc/en-us/articles/221958248-Supported-devices>

<https://i.ytimg.com/vi/2sj2iQyBTQs/maxresdefault.jpg>

<https://review.openstack.org/#/c/335512/>

<http://www.slideshare.net/linaroorg/bkk16504-running-linux-in-el2-virtualization>