THE LINUX FOUNDATION
OPEN SOURCE SUMMIT

An Empirical Study of an Advanced Kernel Tailoring Framework

Junghwan Kang / ultract@nsr.re.kr

@ultractt

THE LINUX FOUNDATION

Contents

- Introduction
- Review
  - My Previous Work @ OSSummit NA 2017
- Advanced Features
- Demo
- Evaluation
- Discussion
- Conclusion

Introduction

THE LINUX FOUNDATION

Introduction

- Motivation of My Work
  - Minimize the Attack Surface of the Linux Kernel
  - Automate the Kernel Configuration
  - Produce a Stable Tailored Linux Kernel

More than
12,000 Options
(Has Prompts)

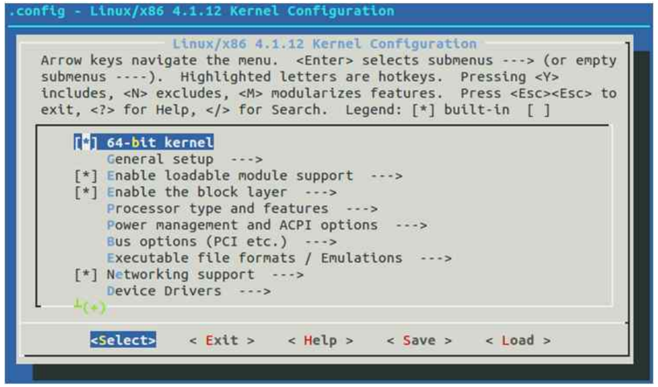

```
.config - Linux/x86 4.1.12 Kernel Configuration
                 Linux/x86 4.1.12 Kernel Configuration
 Arrow keys navigate the menu.  <Enter> selects submenus ---> (or empty
 submenus ----).  Highlighted letters are hotkeys.  Pressing <Y>
 includes, <N> excludes, <M> modularizes features.  Press <Esc><Esc> to
 exit, <?> for Help, </> for Search.  Legend: [*] built-in  [ ]

        [*] 64-bit kernel
            General setup  --->
        [*] Enable loadable module support  --->
        [*] Enable the block layer  --->
            Processor type and features  --->
            Power management and ACPI options  --->
            Bus options (PCI etc.)  --->
            Executable file formats / Emulations  --->
        [*] Networking support  --->
            Device Drivers  --->
        ┴(+)

      <Select>    < Exit >    < Help >    < Save >    < Load >
```

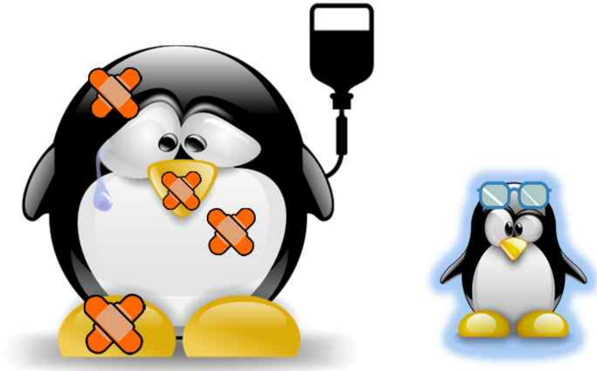

Introduction

- 1st Approach – Undertaker-Tailor
  - Uses Ftrace(Kernel Function Tracer)
  - Formulates Dependency Relationships of Kernel Configuration Options
  - Uses SAT Solver

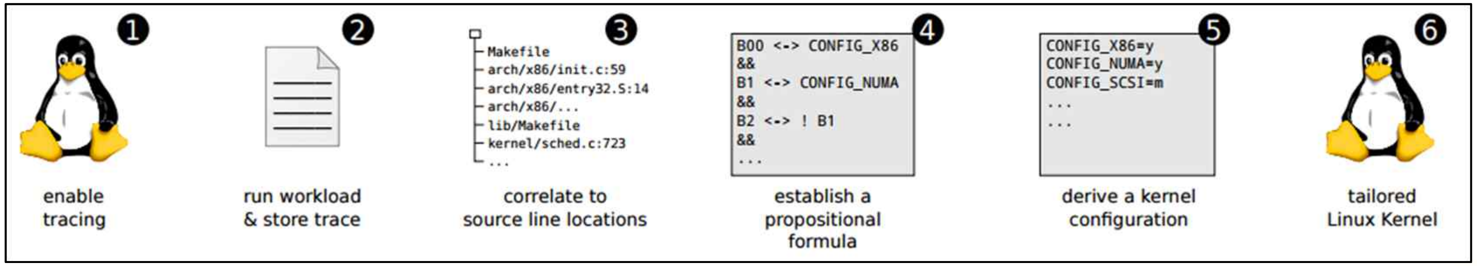

1
enable tracing
2
run workload & store trace
3
correlate to source line locations
4
establish a propositional formula
5
derive a kernel configuration
6
tailored Linux Kernel

```
Makefile
arch/x86/init.c:59
arch/x86/entry32.S:14
arch/x86/...
lib/Makefile
kernel/sched.c:723
...
```

```
B00 <-> CONFIG_X86
&&
B1 <-> CONFIG_NUMA
&&
B2 <-> ! B1
&&
...
```

```
CONFIG_X86=y
CONFIG_NUMA=y
CONFIG_SCSI=m
...
...
```

Workflow of Undertaker-tailor

Introduction

- 1st Approach – Undertaker-Tailor
  - Great! However, tailored kernels often fail to run
    - Failed to Boot Up ☹
    - Found Some Bugs & Fixed them

Introduction

- 2nd Approach – Localmodconfig
  - Command For Configuring the Kernel
  - Very Useful to reduce the # of Kernel Modules
    - Mostly Drivers Removed
    - Still Unnecessary Configuration Options…

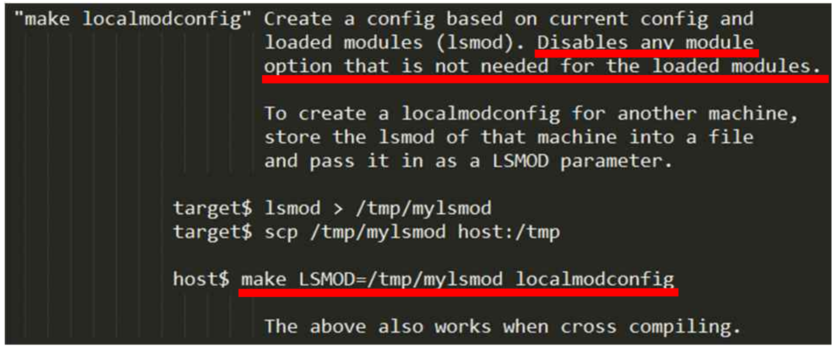

```
"make localmodconfig" Create a config based on current config and
                      loaded modules (lsmod). Disables any module
                      option that is not needed for the loaded modules.

                      To create a localmodconfig for another machine,
                      store the lsmod of that machine into a file
                      and pass it in as a LSMOD parameter.

              target$ lsmod > /tmp/mylsmod
              target$ scp /tmp/mylsmod host:/tmp

              host$ make LSMOD=/tmp/mylsmod localmodconfig

                      The above also works when cross compiling.
```

Introduction

- 3rd Approach – Kernel Tailoring Framework
  - Uses the Undertaker-Tailor with Some Fixes
  - Automates Kernel Tailoring Workflow
  - Checks Tailored Kernels if it includes essential configurations, by looking into
    - Boot State
    - System Logs, Kernel Modules
    - Peripherals(Keyboard, Mouse, Network, etc.)
  - **Got a Working Tailored Kernel!**
    - But, Not Boot Up Sometimes…
    - *I Needed Next Approaches for an Advanced Kernel Tailoring…*

Introduction

- 4th Approach – **Advanced Kernel Tailoring Framework**
  - Improves a Stability
    - Enables tailoring with fine-grained configuration options (Not Grouping)
    - Includes Various Conditions to Verify Tailored Kernels
      - Shows Relationships between Configuration Options & the Conditions
  - Supports for Other Linux Distributions
    - Debian
    - Ubuntu
    - …
  - Measures Performance of between a Tailored & Original Kernel
    - Lmbench (Micro-benchmark for Linux/UNIX/POSIX)
    - Phoronix-Test-Suite (Benchmark for Linux & Other Operating Systems)

Review - My Previous Work

※ Details of My Previous Work are
in a Presentation File at OSSummit NA 2017 ☺
(http://sched.co/BCsG)

Review – My Previous Work

- Design
  - Architecture

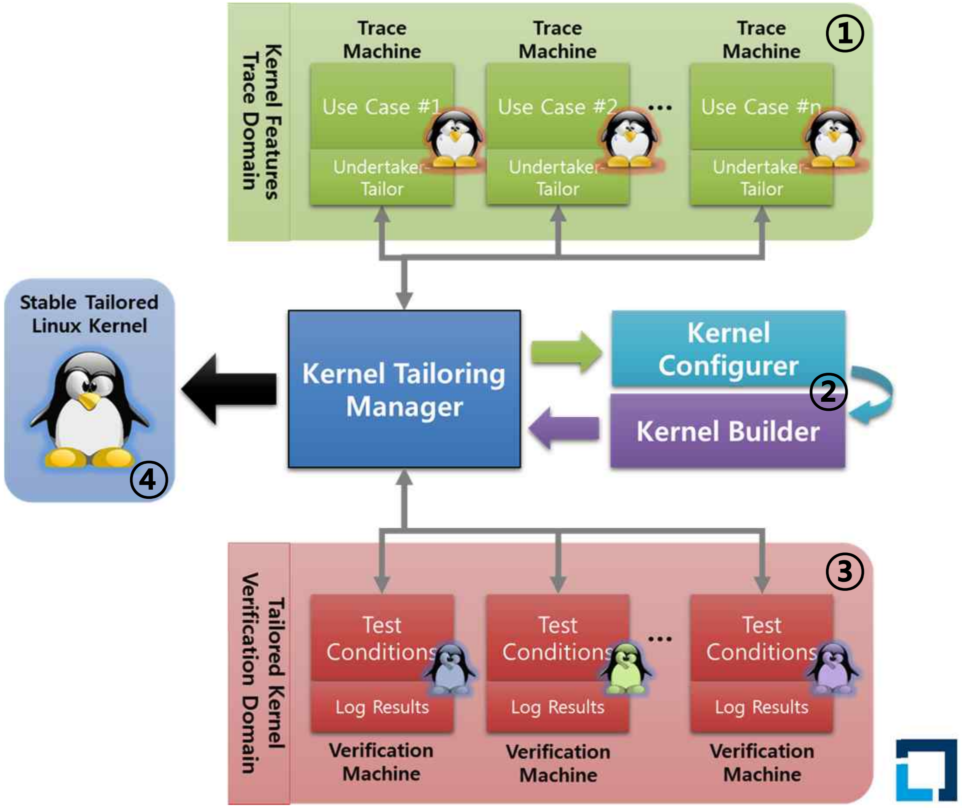

Review – My Previous Work

- Design
  - Workflow

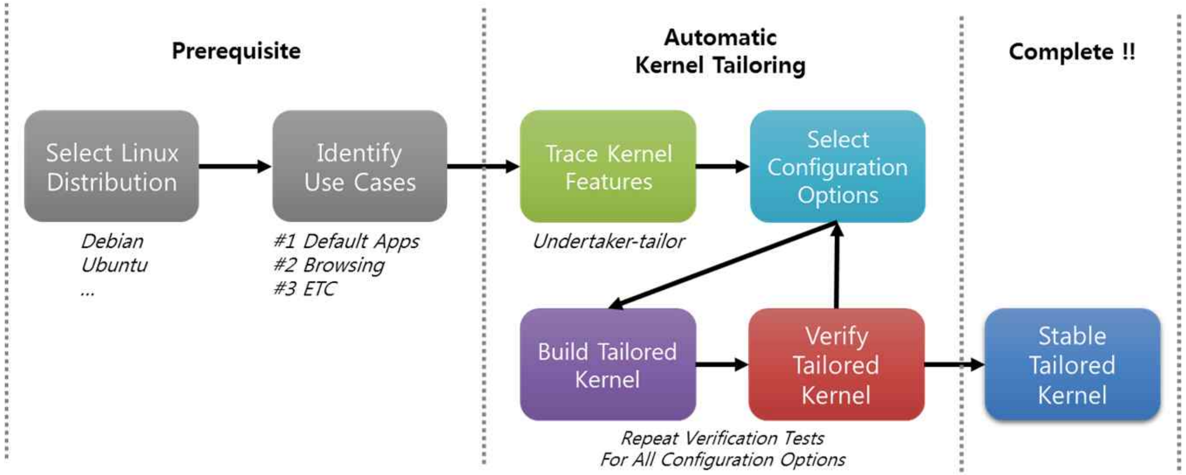

Review – My Previous Work

- Design
  - Kernel Configurer
    - Selects Configuration Options
      - Replenishes a Shortage of the Kernel Configuration by the Undertaker-Tailor
    - Groups Configuration Options For Tests
      - Reduces the number of Tests for Tailored Kernels (Configure & Build & Verify a Tailored Kernel)

Review – My Previous Work

- Design
  - Kernel Configurer

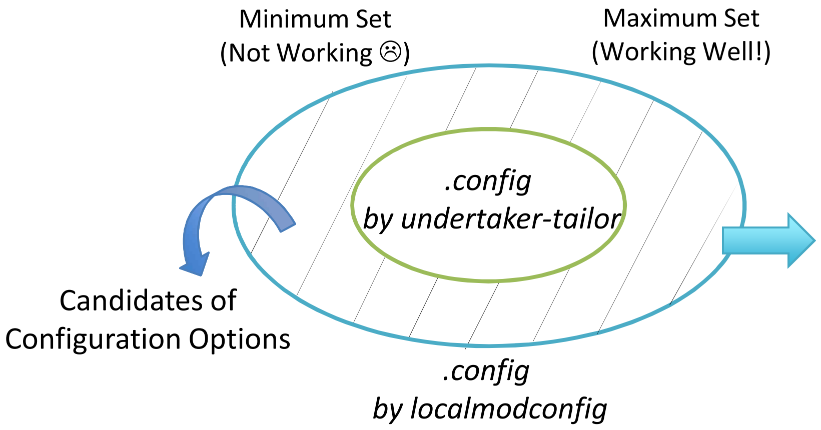

Select & Group the Configuration Options

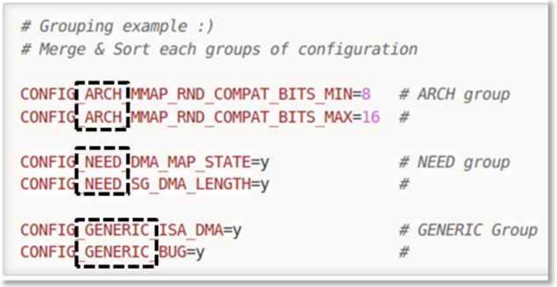

```
# Grouping example :)
# Merge & Sort each groups of configuration

CONFIG_ARCH_MMAP_RND_COMPAT_BITS_MIN=8    # ARCH group
CONFIG_ARCH_MMAP_RND_COMPAT_BITS_MAX=16   #

CONFIG_NEED_DMA_MAP_STATE=y               # NEED group
CONFIG_NEED_SG_DMA_LENGTH=y               #

CONFIG_GENERIC_ISA_DMA=y                  # GENERIC Group
CONFIG_GENERIC_BUG=y                      #
```

Review – My Previous Work

- Design
  - Kernel Configurer

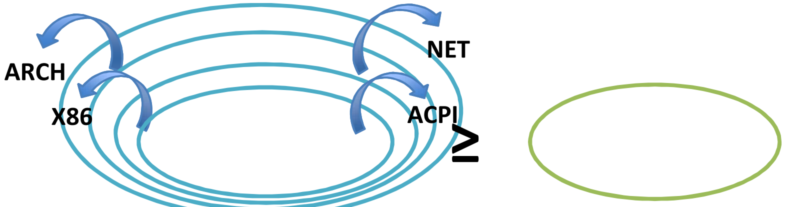

.config
by Kernel Tailoring Framework
by localmodconfig

.config
by undertaker-tailor

Review – My Previous Work

- Implementation

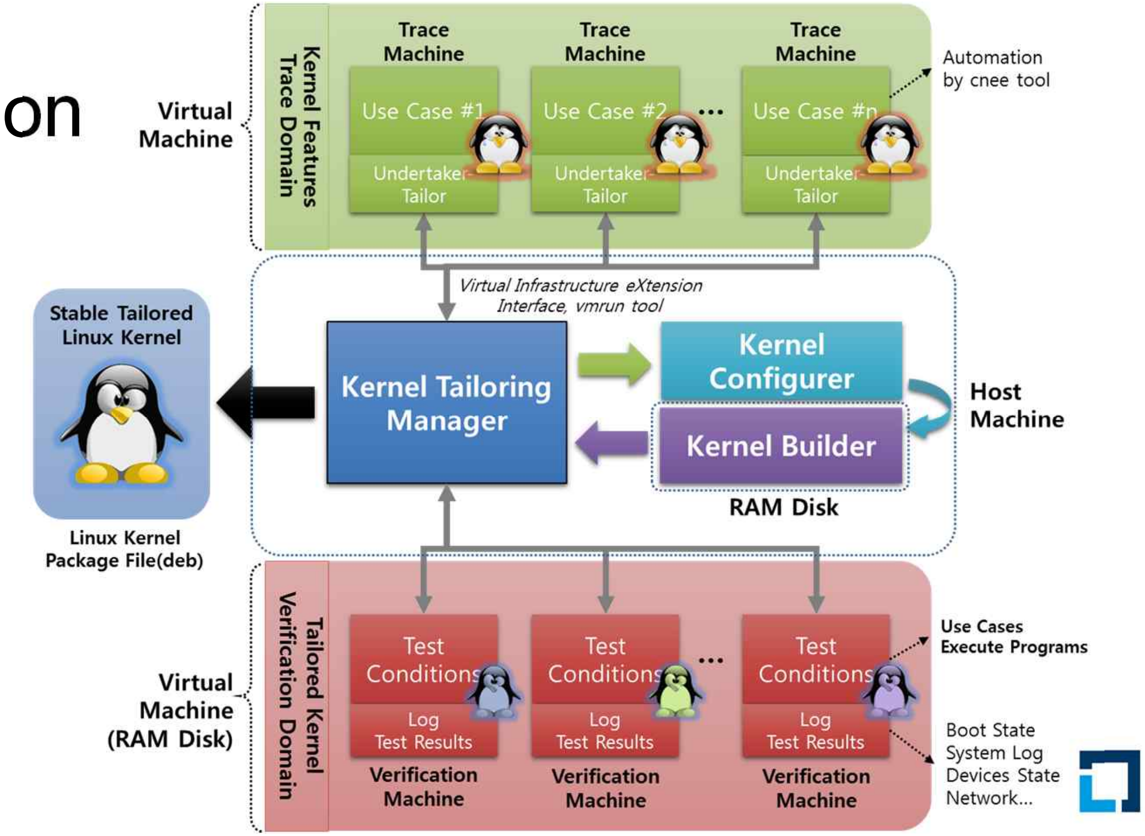

Review – My Previous Work

- Implementation
  - Multi-VMs for a Verification
    - # of Maximum VMs: 5

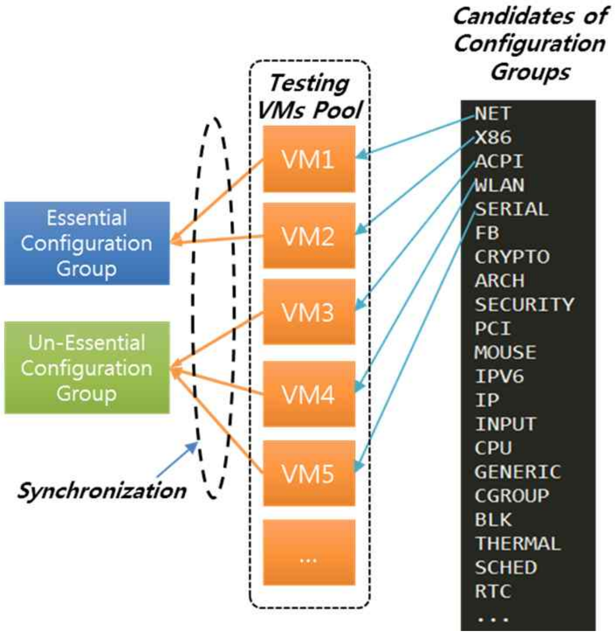

Review – My Previous Work

- Evaluation
  - Elapsed time: About 5 Hours(# of Verification VMs: 5)
  - Kernel Image Size: About ½ ↓
  - # of Kernel Modules: 110/3269 ≒ 3.4 %
  - **Got a Working Tailored Linux Kernel!!**
    - But, I found out that the Kernel doesn’t boot up sometimes ☹

Advanced Features

Advanced Features

- Fine-grained Kernel Tailoring
  - Not Grouping
    - Tailoring Each Kernel Configuration Option
    - Relationship with Conditions for a Verification

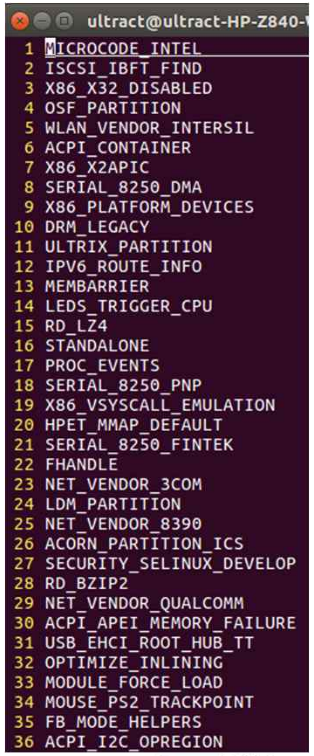

```
MICROCODE_INTEL
ISCSI_IBFT_FIND
X86_X32_DISABLED
OSF_PARTITION
WLAN_VENDOR_INTERSIL
ACPI_CONTAINER
X86_X2APIC
SERIAL_8250_DMA
X86_PLATFORM_DEVICES
DRM_LEGACY
ULTRIX_PARTITION
IPV6_ROUTE_INFO
MEMBARRIER
LEDS_TRIGGER_CPU
RD_LZ4
STANDALONE
PROC_EVENTS
SERIAL_8250_PNP
X86_VSYSCALL_EMULATION
HPET_MMAP_DEFAULT
SERIAL_8250_FINTEK
FHANDLE
NET_VENDOR_3COM
LDM_PARTITION
NET_VENDOR_8390
ACORN_PARTITION_ICS
SECURITY_SELINUX_DEVELOP
RD_BZIP2
NET_VENDOR_QUALCOMM
ACPI_APEI_MEMORY_FAILURE
USB_EHCI_ROOT_HUB_TT
OPTIMIZE_INLINING
MODULE_FORCE_LOAD
MOUSE_PS2_TRACKPOINT
FB_MODE_HELPERS
ACPI_I2C_OPREGION
```

Candidates of Configuration Options
(# of Candidates: 650 For Gooroom)

Advanced Features

- Fine-grained Kernel Tailoring
  - Only Selectable Configuration Options
    - Uses a Model File by the *undertaker-kconfigdump*
      - “HasPrompts”

Showing Selectable Configuration Options

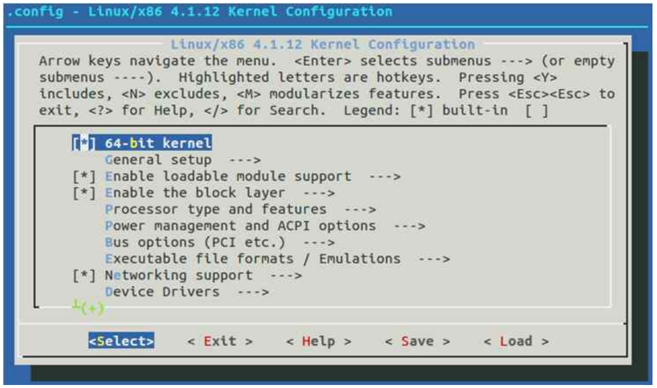

```
.config - Linux/x86 4.1.12 Kernel Configuration
                Linux/x86 4.1.12 Kernel Configuration
 Arrow keys navigate the menu.  <Enter> selects submenus ---> (or empty
 submenus ----).  Highlighted letters are hotkeys.  Pressing <Y>
 includes, <N> excludes, <M> modularizes features.  Press <Esc><Esc> to
 exit, <?> for Help, </> for Search.  Legend: [*] built-in  [ ]

     [*] 64-bit kernel
         General setup  --->
     [*] Enable loadable module support  --->
     [*] Enable the block layer  --->
         Processor type and features  --->
         Power management and ACPI options  --->
         Bus options (PCI etc.)  --->
         Executable file formats / Emulations  --->
     [*] Networking support  --->
         Device Drivers  --->
     └(+)

      <Select>    < Exit >    < Help >    < Save >    < Load >
```

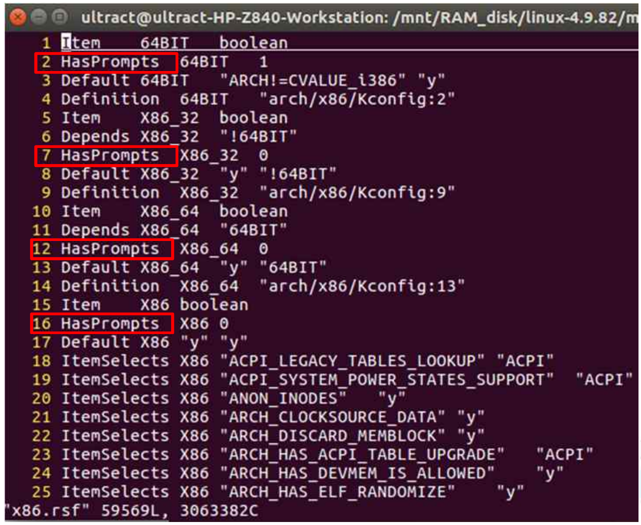

```
ultract@ultract-HP-Z840-Workstation: /mnt/RAM_disk/linux-4.9.82/m
Item     64BIT   boolean
HasPrompts  64BIT   1
Default 64BIT   "ARCH!=CVALUE_i386" "y"
Definition  64BIT   "arch/x86/Kconfig:2"
Item     X86_32  boolean
Depends X86_32  "!64BIT"
HasPrompts  X86_32  0
Default X86_32  "y" "!64BIT"
Definition  X86_32  "arch/x86/Kconfig:9"
Item     X86_64  boolean
Depends X86_64  "64BIT"
HasPrompts  X86_64  0
Default X86_64  "y" "64BIT"
Definition  X86_64  "arch/x86/Kconfig:13"
Item     X86 boolean
HasPrompts  X86 0
Default X86 "y" "y"
ItemSelects X86 "ACPI_LEGACY_TABLES_LOOKUP" "ACPI"
ItemSelects X86 "ACPI_SYSTEM_POWER_STATES_SUPPORT"  "ACPI"
ItemSelects X86 "ANON_INODES"   "y"
ItemSelects X86 "ARCH_CLOCKSOURCE_DATA" "y"
ItemSelects X86 "ARCH_DISCARD_MEMBLOCK" "y"
ItemSelects X86 "ARCH_HAS_ACPI_TABLE_UPGRADE"   "ACPI"
ItemSelects X86 "ARCH_HAS_DEVMEM_IS_ALLOWED"    "y"
ItemSelects X86 "ARCH_HAS_ELF_RANDOMIZE"    "y"
"x86.rsf" 59569L, 3063382C
```

“x86.rsf” File by undertaker-kconfigdump

Advanced Features

- Fine-grained Kernel Tailoring
  - Dependency between Configuration Options
    - Counts how other configuration options “Depend on” a particular configuration option (reverse dependency)
    - Tailoring in the order of degree of the dependency from lowest to highest

# of the Reverse Dependency

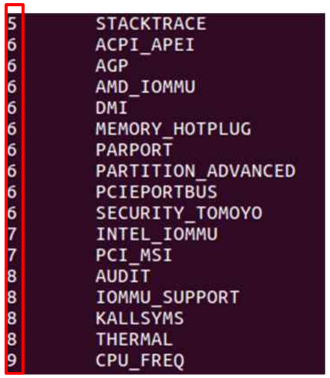

```
5        STACKTRACE
6        ACPI_APEI
6        AGP
6        AMD_IOMMU
6        DMI
6        MEMORY_HOTPLUG
6        PARPORT
6        PARTITION_ADVANCED
6        PCIEPORTBUS
6        SECURITY_TOMOYO
7        INTEL_IOMMU
7        PCI_MSI
8        AUDIT
8        IOMMU_SUPPORT
8        KALLSYMS
8        THERMAL
9        CPU_FREQ
```

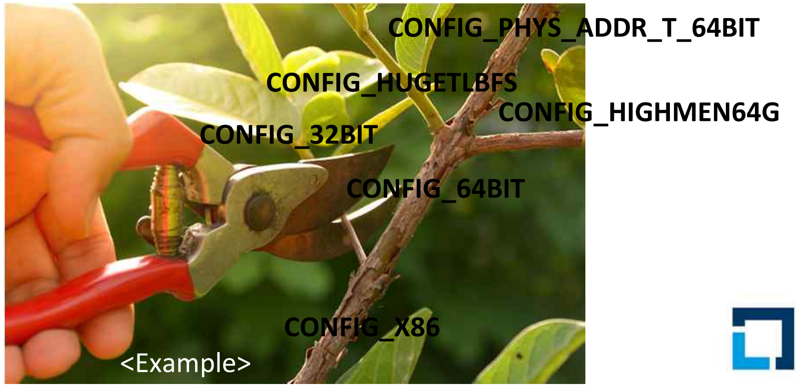

<Example>

Advanced Features

- Fine-grained Kernel Tailoring
  - Randomize Configuration Options
    - Minimize Dependency between Candidates of Configuration Options

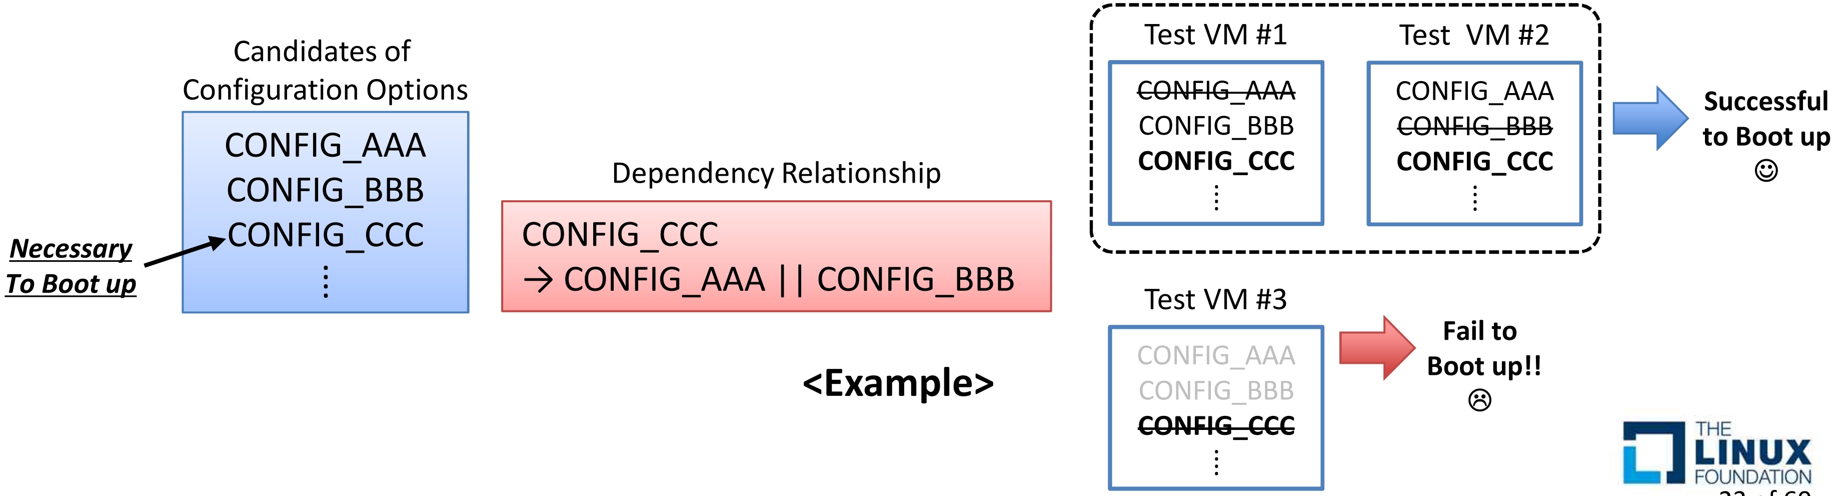

Candidates of Configuration Options

CONFIG_AAA
CONFIG_BBB
CONFIG_CCC
⋮

Necessary To Boot up

Dependency Relationship

CONFIG_CCC
→ CONFIG_AAA || CONFIG_BBB

<Example>

Test VM #1

~~CONFIG_AAA~~
CONFIG_BBB
CONFIG_CCC
⋮

Test VM #2

CONFIG_AAA
~~CONFIG_BBB~~
CONFIG_CCC
⋮

Successful to Boot up
☺

Test VM #3

CONFIG_AAA
CONFIG_BBB
~~**CONFIG_CCC**~~
⋮

Fail to Boot up!!
☹

Advanced Features

- Various conditions for a verification
  - Display
    - Resolution and Dimension
  - Network
  - Peripherals
    - Keyboard and Mouse
  - Security
    - Protection Mechanisms for the Linux Kernel
  - File Systems
  - Etc
    - Power State
    - System Logs (Journalctl)
    - Running Applications

Advanced Features

- Various conditions for a verification - Display
  - Resolution & Dimension
    - phoronix-test-suite system-info → Compare the Before and After
    - xdpyinfo or xrandr
      → Compare the Before and After

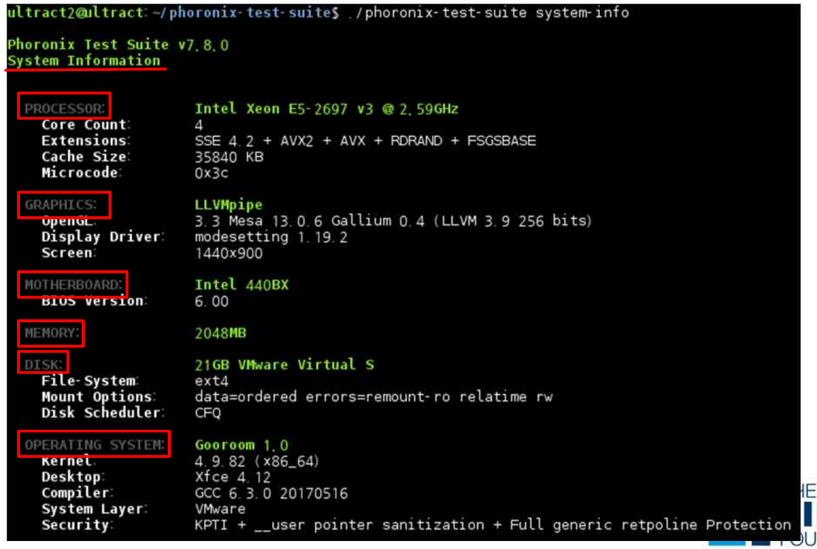

```
ultract2@ultract:~/phoronix-test-suite$ ./phoronix-test-suite system-info

Phoronix Test Suite v7.8.0
System Information


  PROCESSOR:              Intel Xeon E5-2697 v3 @ 2.59GHz
    Core Count:           4
    Extensions:           SSE 4.2 + AVX2 + AVX + RDRAND + FSGSBASE
    Cache Size:           35840 KB
    Microcode:            0x3c

  GRAPHICS:               LLVMpipe
    OpenGL:               3.3 Mesa 13.0.6 Gallium 0.4 (LLVM 3.9 256 bits)
    Display Driver:       modesetting 1.19.2
    Screen:               1440x900

  MOTHERBOARD:            Intel 440BX
    BIOS Version:         6.00

  MEMORY:                 2048MB

  DISK:                   21GB VMware Virtual S
    File-System:          ext4
    Mount Options:        data=ordered errors=remount-ro relatime rw
    Disk Scheduler:       CFQ

  OPERATING SYSTEM:       Gooroom 1.0
    Kernel:               4.9.82 (x86_64)
    Desktop:              Xfce 4.12
    Compiler:             GCC 6.3.0 20170516
    System Layer:         VMware
    Security:             KPTI + __user pointer sanitization + Full generic retpoline Protection
```

Advanced Features

- Various Conditions for a Verification
  - Network
    - IPv4
      - /bin/ip a | grep "192.168."
    - IPv6
      - /bin/ip a | grep "inet6 [a-z0-9]\+::[a-z0-9:]\+"
      - dmesg or journalctl | grep "Failed to insert module 'ipv6'"
    - Ping the Gateway

Advanced Features

- Various Conditions for a Verification - Peripherals
  - Keyboard & Mouse Device
    - /dev/input & udevadm(udev management tool) info
      - ID_INPUT_KEYBOARD, ID_INPUT_MOUSE
    - lsmod | grep 'psmouse‘

Advanced Features

- Various Conditions for a Verification
  - Security Mechanisms for the Linux Kernel
    - checksec → Compare the Before and After
      - Check Kernel Protection mechanisms.
        E.g. Restrict /dev/mem, ASLR, GCC stack protector support…
        (https://github.com/slimm609/checksec.sh)
    - phoronix-test-suite info → Compare the Before and After

Advanced Features

- Various Conditions for a Verification
  - File Systems
    - mount → Compare the Before and After
      - Filters Plugable(Dynamic) File Systems
        E.g. grep -v "binfmt_misc\|iso9660\|fusectl"

※ Verifiable by Other Conditions or Use-cases

Advanced Features

- Various Conditions for a Verification
  - Etc
    - Power State(Suspend & Hibernation)
      - grep "suspend" | /sys/power/disk
      - grep "disk" | /sys/power/state

        ※ https://www.kernel.org/doc/Documentation/power/
    - Journalctl → Compare the Before and After
    - phoronix-test-suite info → Compare the Before and After
    - Running Applications

Advanced Features

- Supports for Other Linux Distributions
  - Gooroom(Our Custom Desktop Linux ☺)
    - Beta 1.0 64bit, Kernel Ver 4.9
    - Xfce Desktop Environment, Lightdm
  - Debian
    - Stretch(9.4) 64bit Desktop, Kernel Ver 4.9
    - Gnome Desktop Environment, Lightdm
  - Ubuntu
    - Bionic Beaver(18.04) 64bit Desktop, Kernel Ver 4.15
    - Gnome Desktop Environment, Lightdm

Demo

Demo

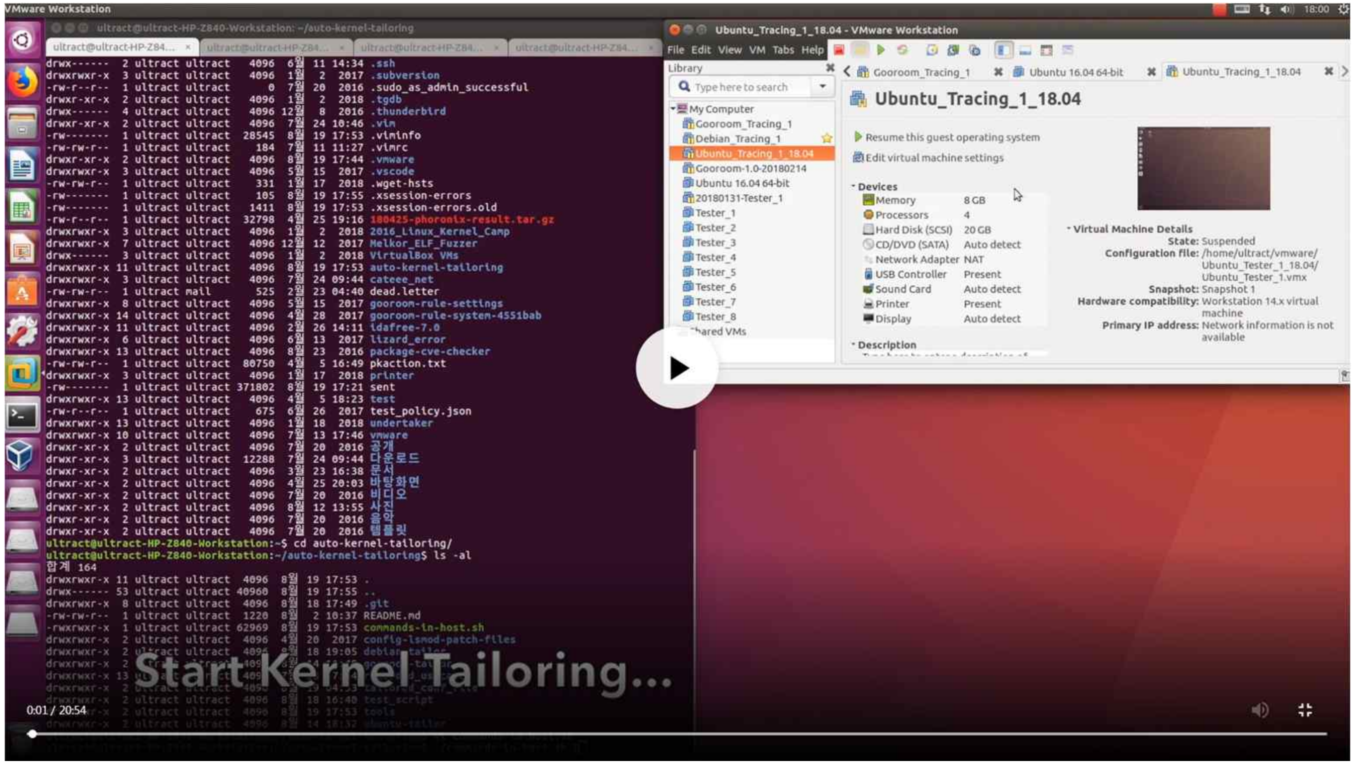

```
drwx------  2 ultract ultract   4096  6월 11 14:34 .ssh
drwxrwxr-x  3 ultract ultract   4096  1월  2  2017 .subversion
-rw-r--r--  1 ultract ultract      0  7월 20  2016 .sudo_as_admin_successful
drwxr-xr-x  2 ultract ultract   4096  1월  2  2018 .tgdb
drwx------  4 ultract ultract   4096 12월  8  2016 .thunderbird
drwxr-xr-x  2 ultract ultract   4096  7월 24 10:46 .vim
-rw-------  1 ultract ultract  28545  8월 19 17:53 .viminfo
-rw-rw-r--  1 ultract ultract    184  7월 11 11:27 .vimrc
drwxr-xr-x  2 ultract ultract   4096  8월 19 17:44 .vmware
drwxrwxr-x  3 ultract ultract   4096  5월 15  2017 .vscode
-rw-rw-r--  1 ultract ultract    331  1월 17  2018 .wget-hsts
-rw-------  1 ultract ultract    105  8월 19 17:55 .xsession-errors
-rw-------  1 ultract ultract   1411  8월 19 17:53 .xsession-errors.old
-rw-r--r--  1 ultract ultract  32798  4월 25 19:16 180425-phoronix-result.tar.gz
drwxrwxr-x  3 ultract ultract   4096  1월  2  2018 2016_Linux_Kernel_Camp
drwxrwxr-x  7 ultract ultract   4096 12월 12  2017 Melkor_ELF_Fuzzer
drwx------  3 ultract ultract   4096  1월  2  2018 VirtualBox VMs
drwxrwxr-x 11 ultract ultract   4096  8월 19 17:53 auto-kernel-tailoring
drwxrwxr-x  3 ultract ultract   4096  7월 24 09:44 cateee_net
-rw-rw-r--  1 ultract mail       525  2월 23 04:40 dead.letter
drwxrwxr-x  8 ultract ultract   4096  5월 15  2017 gooroom-rule-settings
drwxrwxr-x 14 ultract ultract   4096  4월 28  2017 gooroom-rule-system-4551bab
drwxrwxr-x 11 ultract ultract   4096  2월 26 14:11 idafree-7.0
drwxrwxr-x  6 ultract ultract   4096  6월 13  2017 lizard_error
drwxrwxr-x 13 ultract ultract   4096  8월 23  2016 package-cve-checker
-rw-rw-r--  1 ultract ultract  80750  4월  5 16:49 pkaction.txt
drwxrwxr-x  3 ultract ultract   4096  1월 17  2018 printer
-rw-------  1 ultract ultract 371802  8월 19 17:21 sent
drwxrwxr-x 13 ultract ultract   4096  4월  5 18:23 test
-rw-r--r--  1 ultract ultract    675  6월 26  2017 test_policy.json
drwxrwxr-x 13 ultract ultract   4096  1월 18  2018 undertaker
drwxrwxr-x 10 ultract ultract   4096  7월 13 17:46 vmware
drwxr-xr-x  2 ultract ultract   4096  7월 20  2016 공개
drwxr-xr-x  3 ultract ultract  12288  7월 24 09:44 다운로드
drwxr-xr-x  2 ultract ultract   4096  3월 23 16:38 문서
drwxr-xr-x  2 ultract ultract   4096  4월 25 20:03 바탕화면
drwxr-xr-x  2 ultract ultract   4096  7월 20  2016 비디오
drwxr-xr-x  2 ultract ultract   4096  8월 12 13:55 사진
drwxr-xr-x  2 ultract ultract   4096  7월 20  2016 음악
drwxr-xr-x  2 ultract ultract   4096  7월 20  2016 템플릿
ultract@ultract-HP-Z840-Workstation:~$ cd auto-kernel-tailoring/
ultract@ultract-HP-Z840-Workstation:~/auto-kernel-tailoring$ ls -al
합계 164
drwxrwxr-x 11 ultract ultract  4096  8월  19 17:53 .
drwx------ 53 ultract ultract 40960  8월  19 17:55 ..
drwxrwxr-x  8 ultract ultract  4096  8월  18 17:49 .git
-rw-rw-r--  1 ultract ultract  1220  8월   2 10:37 README.md
-rwxrwxr-x  1 ultract ultract 62969  8월  19 17:53 commands-in-host.sh
drwxrwxr-x  2 ultract ultract  4096  4월  20  2017 config-lsmod-patch-files
drwxrwxr-x  2 ultract ultract  4096  8월  18 19:05 debian_tailor
drwxrwxr-x  2 [illegible]
drwxrwxr-x 13 [illegible]
drwxrwxr-x  2 [illegible]
drwxrwxr-x  2 ultract ultract  4096  8월  18 16:40 test_script
drwxrwxr-x  2 ultract ultract  4096  8월  19 17:53 tools
drwxrwxr-x  2 ultract ultract  4096  8월  14 18:32 ubuntu-tailor
```

Ubuntu_Tracing_1_18.04

Devices

Memory	8 GB
Processors	4
Hard Disk (SCSI)	20 GB
CD/DVD (SATA)	Auto detect
Network Adapter	NAT
USB Controller	Present
Sound Card	Auto detect
Printer	Present
Display	Auto detect

Virtual Machine Details

State:	Suspended
Configuration file:	/home/ultract/vmware/Ubuntu_Tester_1_18.04/Ubuntu_Tester_1.vmx
Snapshot:	Snapshot 1
Hardware compatibility:	Workstation 14.x virtual machine
Primary IP address:	Network information is not available

※ This Video: https://youtu.be/fHceA4asiXU
Previous Work : https://youtu.be/fnnCn-Bxjnw

Evaluation

Evaluation

- Total Elapsed Time
  - Gooroom Beta 1.0
    - 7 Hours 55 Minutes
      - # of Verification VMs: 8
      - # of Candidates of Configuration Options: 650
  - Debian 9.4
    - 9 Hours 20 Minutes
      - # of Verification VMs: 8
      - # of Candidates of Configuration Options: 628
  - Ubuntu 18.04
    - 14 Hours 45 Minutes
      - # of Verification VMs: 8
      - # of Candidates of Configuration Options: 997

Evaluation

- Kernel Image & Initial Ramdisk & Kernel Modules
  - Gooroom Beta 1.0
    - Kernel Image Size
      - Tailored : 14,399,796 Bytes (≈ 72%)
      - Original : 20,090,752 Bytes, ※ Decompressed by extract-vmlinux
    - Initial Ramdisk Size
      - Tailored : 6,672,465 Bytes (≈ 20%)
      - Original : 34,078,719 Bytes
    - The Size of Kernel Modules
      - Tailored : 6,650,050 Bytes (≈ 0.04%), # of .ko : 91 (≈ 0.03%)
      - Original : 186,697,093 Bytes , # of .ko : 3,387

Evaluation

- Kernel Image & Initial Ramdisk & Kernel Modules
  - Debian 9.4
    - Kernel Image Size
      - Tailored : 12,289,612 Bytes (≈ 61%)
      - Original : 20,161,244 Bytes, ※ Decompressed by extract-vmlinux
    - Initial Ramdisk Size
      - Tailored : 5,910,123 Bytes (≈ 30%)
      - Original : 19,582,713 Bytes
    - The Size of Kernel Modules
      - Tailored : 5,026,255 Bytes (≈ 0.03%), # of .ko : 91 (≈ 0.03%)
      - Original : 189,458,941 Bytes , # of .ko : 3,387

Evaluation

- Kernel Image & Initial Ramdisk & Kernel Modules
  - Ubuntu 18.04
    - Kernel Image Size
      - Tailored : 20,951,272 Bytes (≈ 22%)
      - Original : 94,147,992 Bytes, ※ Decompressed by extract-vmlinux
    - Initial Ramdisk Size
      - Tailored : 12,377,995 Bytes (≈ 22%)
      - Original : 53,935,618 Bytes
    - The Size of Kernel Module
      - Tailored : 5,772,651 Bytes (≈ 0.02%), # of .ko : 64 (≈ 0.01%)
      - Original : 236,401,113 Bytes , # of .ko : 5,161

Evaluation

- Kernel Configuration File
  - Gooroom Beta 1.0

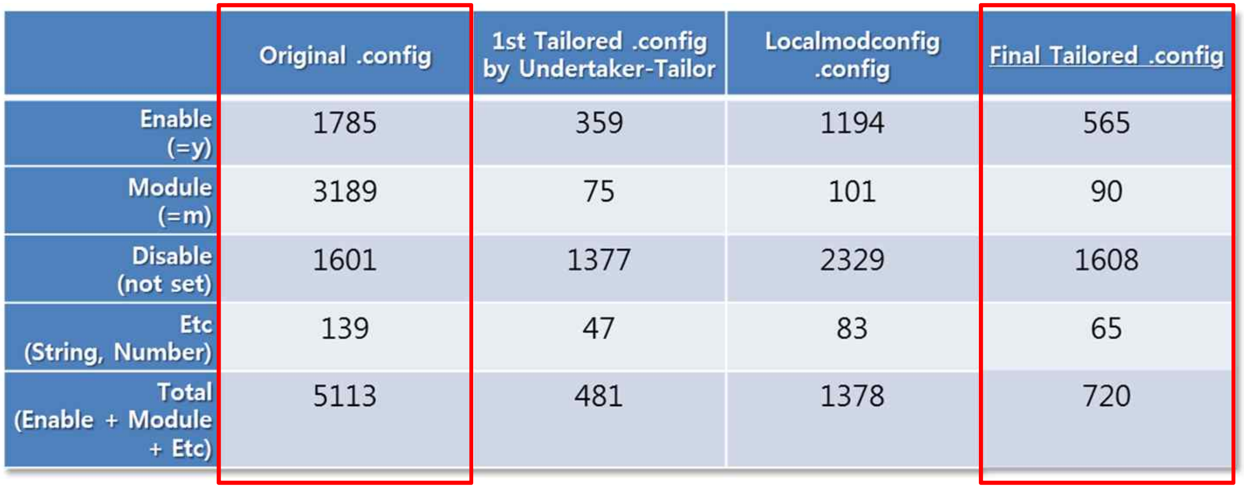

	Original .config	1st Tailored .config by Undertaker-Tailor	Localmodconfig .config	Final Tailored .config
Enable (=y)	1785	359	1194	565
Module (=m)	3189	75	101	90
Disable (not set)	1601	1377	2329	1608
Etc (String, Number)	139	47	83	65
Total (Enable + Module + Etc)	5113	481	1378	720

1785 → 565 (≈ 32%)

3189 → 90 (≈ 3%)

5113 → 720 (≈ 14%)

Evaluation

- Kernel Configuration File
  - Gooroom Beta 1.0

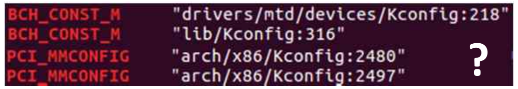

```
BCH_CONST_M      "drivers/mtd/devices/Kconfig:218"
BCH_CONST_M      "lib/Kconfig:316"
PCI_MMCONFIG     "arch/x86/Kconfig:2480"
PCI_MMCONFIG     "arch/x86/Kconfig:2497"
```

?

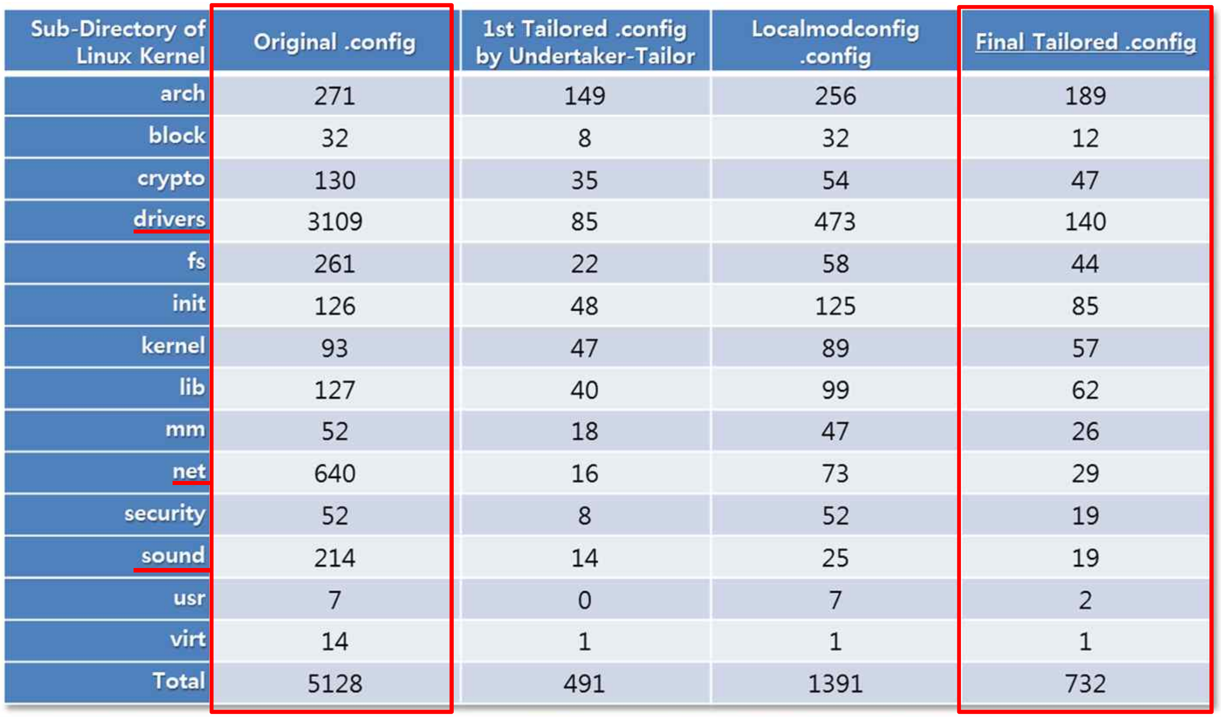

Sub-Directory of Linux Kernel	Original .config	1st Tailored .config by Undertaker-Tailor	Localmodconfig .config	Final Tailored .config
arch	271	149	256	189
block	32	8	32	12
crypto	130	35	54	47
drivers	3109	85	473	140
fs	261	22	58	44
init	126	48	125	85
kernel	93	47	89	57
lib	127	40	99	62
mm	52	18	47	26
net	640	16	73	29
security	52	8	52	19
sound	214	14	25	19
usr	7	0	7	2
virt	14	1	1	1
Total	5128	491	1391	732

3019 → 140 (≈ 5%)

640 → 29 (≈ 5%)

214 → 19 (≈ 9%)

Evaluation

- Kernel Configuration File
  - Debian 9.4

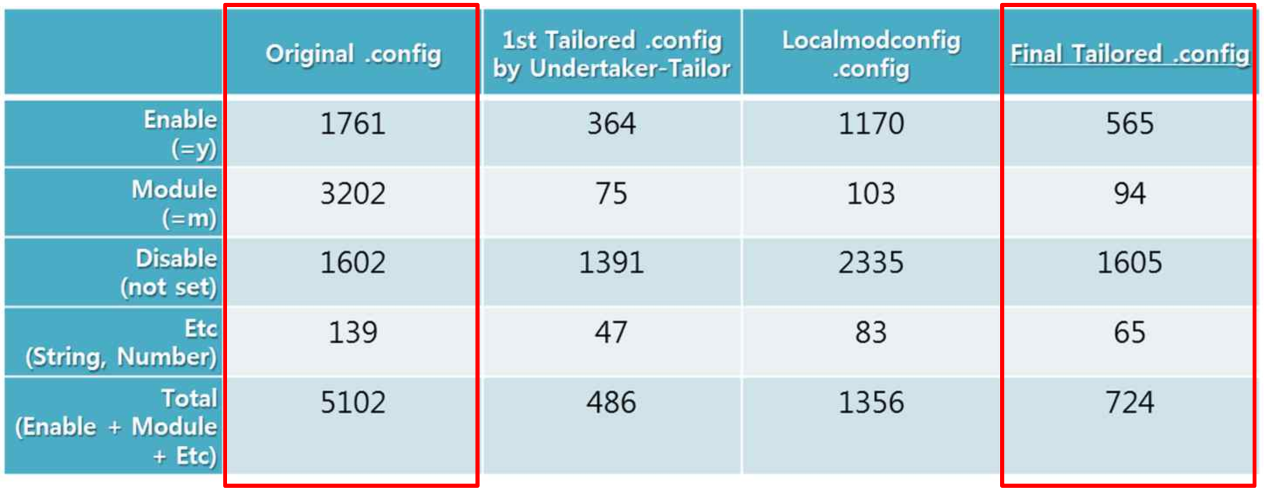

	Original .config	1st Tailored .config by Undertaker-Tailor	Localmodconfig .config	Final Tailored .config
Enable (=y)	1761	364	1170	565
Module (=m)	3202	75	103	94
Disable (not set)	1602	1391	2335	1605
Etc (String, Number)	139	47	83	65
Total (Enable + Module + Etc)	5102	486	1356	724

1761 → 565 (≈ 32%)

3202 → 94 (≈ 3%)

5102 → 724 (≈ 14%)

Evaluation

- Kernel Configuration File
  - Debian 9.4

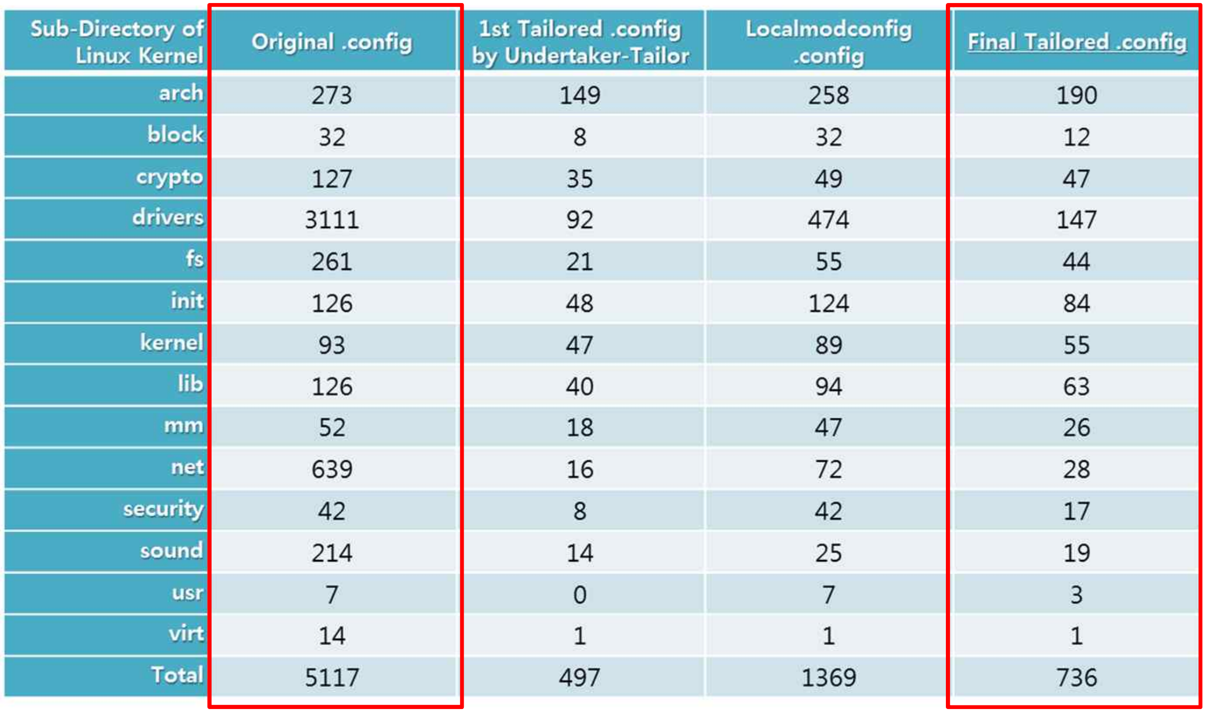

Sub-Directory of Linux Kernel	Original .config	1st Tailored .config by Undertaker-Tailor	Localmodconfig .config	Final Tailored .config
arch	273	149	258	190
block	32	8	32	12
crypto	127	35	49	47
drivers	3111	92	474	147
fs	261	21	55	44
init	126	48	124	84
kernel	93	47	89	55
lib	126	40	94	63
mm	52	18	47	26
net	639	16	72	28
security	42	8	42	17
sound	214	14	25	19
usr	7	0	7	3
virt	14	1	1	1
Total	5117	497	1369	736

3111 → 147 (≈ 5%)

639 → 28 (≈ 4%)

214 → 19 (≈ 9%)

Evaluation

- Kernel Configuration File
  - Ubuntu 18.04

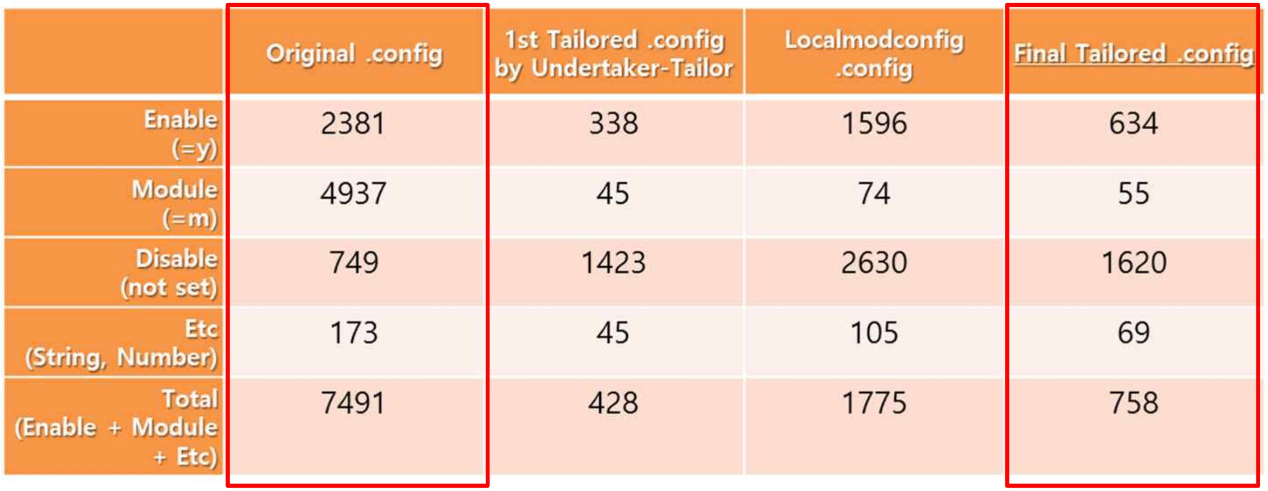

	Original .config	1st Tailored .config by Undertaker-Tailor	Localmodconfig .config	Final Tailored .config
Enable (=y)	2381	338	1596	634
Module (=m)	4937	45	74	55
Disable (not set)	749	1423	2630	1620
Etc (String, Number)	173	45	105	69
Total (Enable + Module + Etc)	7491	428	1775	758

2381 → 634 (≈ 23%)

4937 → 55 (≈ 1%)

7491 → 758 (≈ 10%)

Evaluation

- Kernel Configuration File
  - Ubuntu 18.04

New Directories

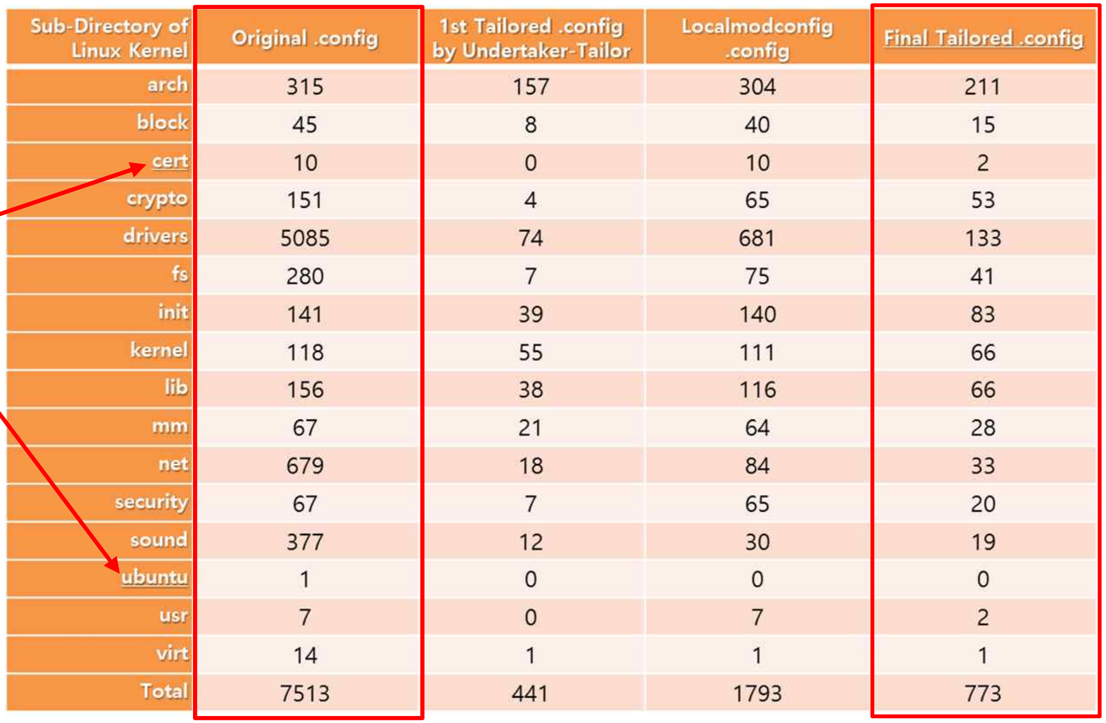

Sub-Directory of Linux Kernel	Original .config	1st Tailored .config by Undertaker-Tailor	Localmodconfig .config	Final Tailored .config
arch	315	157	304	211
block	45	8	40	15
cert	10	0	10	2
crypto	151	4	65	53
drivers	5085	74	681	133
fs	280	7	75	41
init	141	39	140	83
kernel	118	55	111	66
lib	156	38	116	66
mm	67	21	64	28
net	679	18	84	33
security	67	7	65	20
sound	377	12	30	19
ubuntu	1	0	0	0
usr	7	0	7	2
virt	14	1	1	1
Total	7513	441	1793	773

5085 → 133 (≈ 3%)

679 → 33 (≈ 5%)

377 → 19 (≈ 5%)

Evaluation

- Verification Log - Gooroom Beta 1.0

※ https://github.com/ultract/linux-kernel-tailoring-framework/tree/master/tailoring_log

[Boot Up]

BINFMT_SCRIPT
DEVTMPFS
EPOLL
FILE_LOCKING
FUTEX
INOTIFY_USER
MULTIUSER
RD_GZIP
SERIAL_8250
SHMEM
SIGNALFD
SYSFS
TIMERFD
TMPFS
TTY
UNIX
UNIX98_PTYS
VT

[Phoronix-test-suite]

DMI → Motherboard & BIOS Information
DMIID → Motherboard & BIOS Information
DRM_LEGACY → Graphics
IOSCHED_CFQ → Disk Scheduler - CFQ(Before), NOOP(After)
PACKET → No Internet Connectivity
PAGE_TABLE_ISOLATION → Security - KPTI
RETPOLINE → Security - Full generic retpoline Protection

[Journalctl Log]

ECRYPT_FS → Failed to find module 'ecryptfs'
IPV6 → device (enp2s1): addrconf6: failed to start neighbor discovery ...
NAMESPACES → Failed to start Hostname Service ...
PACKET → (Socket Filtering) are enabled in your kernel ...
PARPORT → Failed to find module 'lp', 'parport_pc', 'ppdev'
PRINTER → Failed to find module 'lp'
RETPOLINE → Spectre V2 : kernel not compiled with retpoline;
TMPFS_POSIX_ACL → Failed to apply ACL on /dev/dri/card0: Operation not supported ...

Evaluation

- Verification Log - Gooroom Beta 1.0

※ https://github.com/ultract/linux-kernel-tailoring-framework/tree/master/tailoring_log

[Checksec]

AUDIT → SELinux Enable
CC_STACKPROTECTOR_STRONG → GCC stack protector support
RANDOMIZE_BASE → Address space layout randomization
RELOCATABLE → Address space layout randomization
SECURITY → SELinux Enable
SECURITY_SELINUX → SELinux Enable
STRICT_DEVMEM → Restrict /dev/mem access

[File Systems]

DEFAULT_SECURITY_SMACK → smackfs
NAMESPACES → hugetlbfs
SECURITY → smackfs
SECURITY_SMACK → smackfs

[Peripherals]

INPUT_KEYBOARD
INPUT_MOUSE
KEYBOARD_ATKBD
MOUSE_PS2

[Network]

IPV6 → IPv6 Address Not Set
NAMESPACES → IPv4 Address Not Set
PACKET → IPv4 Address Not Set, Ping to Gateway Failed

[Power State]

HIBERNATION → /sys/power/disk, /sys/power/state
SUSPEND → /sys/power/disk
SWAP → /sys/power/disk, /sys/power/state

[Kernel Module]

MODULE_UNLOAD → Kernel Module Loading Failed

[Applications]

ADVISE_SYSCALLS → Browser Not Working - Fatal Error
NAMESPACES → Pulse Audio Not Working

Evaluation

- Verification Log - Debian 9.4

※ https://github.com/ultract/linux-kernel-tailoring-framework/tree/master/tailoring_log

[Boot Up]

BINFMT_SCRIPT
DEVTMPFS
EPOLL
EXT4_USE_FOR_EXT2
FILE_LOCKING
FUTEX
INOTIFY_USER
MULTIUSER
RD_GZIP
SHMEM
SIGNALFD
SYSFS
TIMERFD
TMPFS
TTY
UNIX
UNIX98_PTYS
VT

[Phoronix-test-suite]

DMI → Motherboard & BIOS Information
DMIID → Motherboard & BIOS Information
IOSCHED_CFQ → Disk Scheduler - CFQ(Before), NOOP(After)
NET_VENDOR_REALTEK → No Internet Connectivity
PACKET → No Internet Connectivity
PAGE_TABLE_ISOLATION → Security - KPTI
RD_LZ4 → No Internet Connectivity
RETPOLINE → Security - Full generic retpoline Protection

[Journalctl Log]

IPV6 → device (enp2s1): addrconf6: failed to start neighbor discovery ...
NAMESPACES → Failed to start Hostname Service ...
NET_VENDOR_REALTEK → setsockopt(udp, IP_ADD_MEMBERSHIP)(0.0.0.0): No such device
PACKET → (Socket Filtering) are enabled in your kernel ...
PARPORT → Failed to find module 'lp', 'parport_pc', 'ppdev'
PRINTER → Failed to find module 'lp'
RD_LZ4 → setsockopt(udp, IP_ADD_MEMBERSHIP)(0.0.0.0): No such device
RETPOLINE → Spectre V2 : kernel not compiled with retpoline; no mitigation available!
SERIAL_8250 → bad device "/dev/ttyS0" given
TMPFS_POSIX_ACL → Failed to apply ACL on /dev/dri/card0: Operation not supported ...
VT_CONSOLE → /dev/ttyS0: not a tty

Evaluation

- Verification Log - Debian 9.4

※ https://github.com/ultract/linux-kernel-tailoring-framework/tree/master/tailoring_log

[Checksec]

AUDIT → SELinux Enable
CC_STACKPROTECTOR_STRONG → GCC stack protector support
RANDOMIZE_BASE → Address space layout randomization
RELOCATABLE → Address space layout randomization
SECURITY → SELinux Enable
SECURITY_SELINUX → SELinux Enable
SLAB_FREELIST_RANDOM SLAB freelist randomization
STRICT_DEVMEM → Restrict /dev/mem access
VMAP_STACK Virtually-mapped kernel stack

[File Systems]

NAMESPACES → hugetlbfs

[Peripherals]

INPUT_KEYBOARD
INPUT_MOUSE
KEYBOARD_ATKBD
MOUSE_PS2

[Network]

IPV6 → IPv6 Address Not Set
NAMESPACES → IPv4 Address Not Set
PACKET → IPv4 Address Not Set, Ping to Gateway Failed

[Power State]

HIBERNATION → /sys/power/disk, /sys/power/state
SWAP → /sys/power/disk, /sys/power/state

[Kernel Module]

MODULE_UNLOAD → Kernel Module Loading Failed

[Applications]

NAMESPACES → Pulse Audio Not Working

Evaluation

- Verification Log - Ubuntu 18.04 ※ https://github.com/ultract/linux-kernel-tailoring-framework/tree/master/tailoring_log

[Boot Up]

BINFMT_SCRIPT

DEVTMPFS

EPOLL

EXT4_FS

FUTEX

INOTIFY_USER

MULTIUSER

RD_GZIP

SERIAL_8250

SERIAL_8250_CONSOLE

SHMEM

SIGNALFD

TIMERFD

TMPFS

UNIX

UNIX98_PTYS

VT

[Phoronix-test-suite]

DMI → Motherboard & BIOS Information

DMIID → Motherboard & BIOS Information

IOSCHED_CFQ → Disk Scheduler - CFQ(Before), NOOP(After)

PACKETT → No Internet Connectivity

PAGE_TABLE_ISOLATION → Security - KPTI

RETPOLINEE → Security - Full generic retpoline Protection

VIRTIO_BALLOON → No Internet Connectivity

[Journalctl Log]

FILE_LOCKING → [autospawn] core-util.c: lock: Permission denied ...

FUSE_FS → Failed to find module 'fuse'

INPUT_EVDEV → cannot open input layer

IPV6 → device (enp2s1): addrconf6: failed to start neighbor discovery ...

OSF_PARTITION → Failed to mount Mount unit for core, revision 5145

PACKET → (Socket Filtering) are enabled in your kernel ...

PARPORT → Failed to find module 'lp', 'parport_pc', 'ppdev'

PARPORT_PC Failed to find module 'parport_pc'

POSIX_TIMERS Failed to call clock_adjtime(): Function not implemented

PRINTER → Failed to find module 'lp'

PRINTK → activation of module imklog failed

RETPOLINE → Spectre V2 : kernel not compiled with retpoline; no mitigation available!

SQUASHFS_XZ → squashfs: SQUASHFS error: Filesystem uses "xz" compression

TMPFS_POSIX_ACL → Failed to apply ACL on /dev/dri/card0: Operation not supported ...

Evaluation

- Verification Log - Ubuntu 18.04

※ https://github.com/ultract/linux-kernel-tailoring-framework/tree/master/tailoring_log

[Checksec]

VMAP_STACK → Virtually-mapped kernel stack
HARDENED_USERCOPY → Hardened Usercopy
SLAB_FREELIST_RANDOM → SLAB freelist randomization
CC_STACKPROTECTOR_STRONG → GCC stack protector support
RANDOMIZE_BASE → Address space layout randomization
AUDIT → SELinux Enable
SECURITY_SELINUX → SELinux Enable
SECURITY → SELinux Enable

[File Systems]

SQUASHFS → squashfs
SQUASHFS_XZ → squashfs
CONFIGFS_FS → configfs
FUSE_FS → fuse.gvfsd-fuse
MISC_FILESYSTEMS → pstore

[Peripherals]

INPUT_KEYBOARD
INPUT_MOUSE
KEYBOARD_ATKBD
MOUSE_PS2

[Network]

PACKET → IPv4 Address Not Set, Ping to Gateway Failed
IPV6 → IPv6 Address Not Set

[Power State]

HIBERNATION → /sys/power/disk, /sys/power/state
SUSPEND → /sys/power/disk
SWAP → /sys/power/disk, /sys/power/state

[Kernel Module]

MODULE_UNLOAD → Kernel Module Loading Failed

[Applications]

FILE_LOCKING → Pulse Audio Not Working

Evaluation

- Boot Up Time - Gooroom Beta 1.0
  - Tailored Kernel Image

    ※ system-analyze

    - Startup finished in **1.577s** (kernel) + 2.930s (userspace) = **4.507s**
    - Startup finished in **1.410s** (kernel) + 2.928s (userspace) = **4.338s**
    - Startup finished in **1.523s** (kernel) + 3.241s (userspace) = **4.764s**
  - Original Kernel Image
    - Startup finished in **2.695s** (kernel) + 3.324s (userspace) = **6.020s**
    - Startup finished in **2.839s** (kernel) + 3.502s (userspace) = **6.341s**
    - Startup finished in **2.836s** (kernel) + 3.082s (userspace) = **5.918s**

Evaluation

- Boot Up Time - Debian 9.4
  - Tailored Kernel Image

    ※ system-analyze

    - Startup finished in **1.416s** (kernel) + 6.751s (userspace) = **8.167s**
    - Startup finished in **1.450s** (kernel) + 6.649s (userspace) = **8.100s**
    - Startup finished in **1.442s** (kernel) + 6.598s (userspace) = **8.041s**
  - Original Kernel Image
    - Startup finished in **1.845s** (kernel) + 7.243s (userspace) = **9.089s**
    - Startup finished in **1.800s** (kernel) + 7.228s (userspace) = **9.029s**
    - Startup finished in **2.053s** (kernel) + 6.992s (userspace) = **9.046s**

Evaluation

- Boot Up Time - Ubuntu 18.04
  - Tailored Kernel Image

    ※ system-analyze

    - Startup finished in **1.724s** (kernel) + 5.912s (userspace) = **7.636s**
    - Startup finished in **1.662s** (kernel) + 4.319s (userspace) = **5.982s**
    - Startup finished in **1.737s** (kernel) + 5.660s (userspace) = **7.397s**
  - Original Kernel Image
    - Startup finished in **3.931s** (kernel) + 5.752s (userspace) = **9.683s**
    - Startup finished in **3.980s** (kernel) + 4.162s (userspace) = **8.143s**
    - Startup finished in **3.894s** (kernel) + 3.793s (userspace) = **7.688s**

Evaluation

- Performance – Lmbench on the Gooroom
  - Most of the Test Results are Similar, except Some Test Items

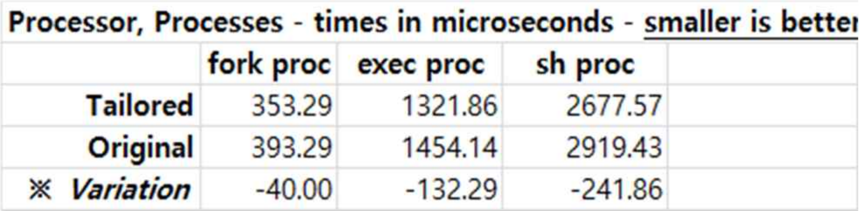

Processor, Processes - times in microseconds - smaller is better

	fork proc	exec proc	sh proc
Tailored	353.29	1321.86	2677.57
Original	393.29	1454.14	2919.43
※ *Variation*	-40.00	-132.29	-241.86

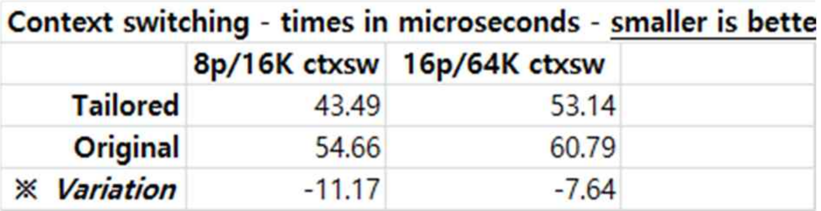

Context switching - times in microseconds - smaller is bette

	8p/16K ctxsw	16p/64K ctxsw
Tailored	43.49	53.14
Original	54.66	60.79
※ *Variation*	-11.17	-7.64

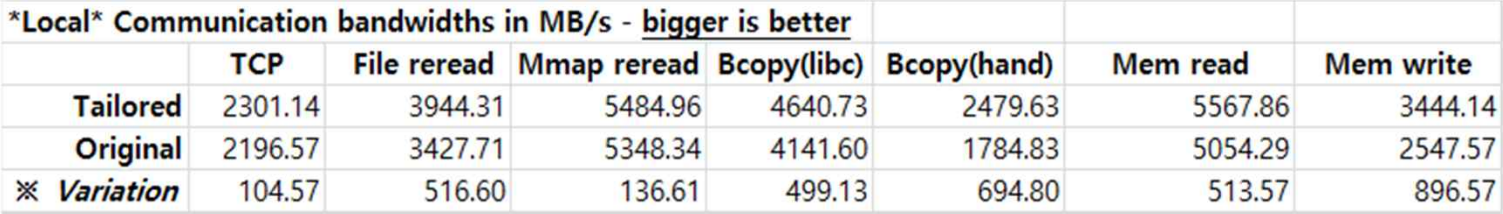

Local Communication bandwidths in MB/s - bigger is better

	TCP	File reread	Mmap reread	Bcopy(libc)	Bcopy(hand)	Mem read	Mem write
Tailored	2301.14	3944.31	5484.96	4640.73	2479.63	5567.86	3444.14
Original	2196.57	3427.71	5348.34	4141.60	1784.83	5054.29	2547.57
※ *Variation*	104.57	516.60	136.61	499.13	694.80	513.57	896.57

Evaluation

- Performance - Phoronix-test-suite on the Gooroom
  - I'll show you the original results
    - ※ https://github.com/ultract/linux-kernel-tailoring-framework/tree/master/performance_test

Discussion

Discussion

- Fine-grained Kernel Tailoring
  - Considering the dependency & Randomizing the Configuration Options
    - Reduced a failure rate of the kernel tailoring empirically
      - The tailored Kernel is always working well ☺
  - The Relationship between conditions for a verification and the Configuration Options
    - Useful to make whitelist for the kernel tailoring

Discussion

- Fine-grained Kernel Tailoring
  - Takes longer than the previous method
    - More than 2 hours at the Gooroom
  - Reduces candidates of configuration options by selectable options(“HasPrompts”) thankfully

Discussion

- The Performance of the Tailored Kernel
  - A little better performance
  - To understand the reason, I need an analysis about the results more…

Discussion

- The Performance of Tailored Kernel
  - It is difficult to collect configuration options about the performance by undertaker-tailor & tailoring framework
  - The Configuration Options need to be added by hand
    - I refer to the linux performance and tuning guidelines
    - I added what the configuration options are in the original .config already

Discussion

- Conditions for the Verification
  - The conditions are found out heuristically
    - Trial and error
    - Comparing the before and after
  - H/W Spec, Drivers & Modules, Applications, Etc
  - It need to be formalize and organize
  - *The more conditions are added, the more configuration options are gathered…*

Discussion

- Desktop Manager Issues for the Verification
  - Xfce or Lightdm is better than Gnome or Gdm
    - A vm using Gnome is slow to be revert and play
    - Gdm service can’t be restarted properly for the use-cases and the verification
  - xfce4-terminal and gnome-terminal
    - They have different options to execute use-cases and the Verification scripts

Discussion

- I have troubles to make Kconfig model files on the Ubuntu
  - undertaker-kconfigdump can’t handle “imply” attribute of the kconfig
    - “imply”(weak select) → “select”

      ※ https://www.kernel.org/doc/Documentation/kbuild/kconfig-language.txt

Discussion

- The limitation of the Localmodconfig
  - It only includes configuration options of inserted modules via the insmod command
- The kernel tailoring is only for a virtual machine yet
  - I need another new approach for the physical machine
    - How to automate to trace kernel features and verify tailored kernels like the virtual machines ??

Conclusion

Conclusion

- We looked into the several approaches for the kernel tailoring
  - Undertaker-tailor
  - Localmodconfig
  - My Kernel tailoring framework
- Advanced features of the kernel tailoring framework
  - Fine-grained kernel tailoring
    - Enhanced Stability of a Tailored Kernel
    - Relation between Configuration Options & Various Verification Conditions
  - Supported for other linux distributions
    - Debian, Ubuntu
  - A little performance benefit
- Future work
  - Formalizing or organizing the Conditions for a Verification
  - Kernel tailoring toward the physical machine ☺

Questions?

(https://github.com/ultract/linux-kernel-tailoring-framework)

THE LINUX FOUNDATION

OPEN SOURCE SUMMIT

THE LINUX FOUNDATION