Devicetree BOF

ELCE 2017
Prague, Czech Republic

Frank Rowand, Sony

October 23, 2017
171023_0006

Agenda

- Devicetree Workshop at Kernel Summit 2017
- questions, comments, issues, concerns from the crowd
- Overlays
- Plumbers 2017
- Plumbers 2016
- Devicetree Specification
- documentation
- commit statistics
- dtc compiler
- Status of debug tools

My Goal

Do NOT show all of the slides

Devicetree Workshop 2017

Prague, Czech Republic
October 26 -- Thursday

Time: 9:00am-5:30pm (Lunch from 12:30-2:30)
Location: Athens room - Hilton Prague

If you plan to attend, make sure you update
your OSSunmitE/ELCE registration to include
the DT Workshop

Devicetree Workshop 2017

Thursday 10/26

9:30 Welcome and Schedule bashing

===Tooling & Schema===	9:40 - 11:10
===Runtime usage===	11:50 - 12:30
===DTS maintenance issues===	14:30 - 16:00
=== More stuff===	16:00 - 17:20+

Devicetree Workshop 2017

```
 9:30 (10min) Welcome and Schedule bashing
 9:40 ( 5min) Encoding and Schema checking: Framing the problem
 9:45 (15min) DT YAML encoding overview
10:00 (20min) YAML encoding discussion
10:20 (15min) DT Schema format - option 1
10:35 (15min) DT Schema format - option 2
10:50 (20min) DT Schema discussion - what should go in the spec?
11:50 (20min) Code Generation from DT
12:10 (20min) Runtime memory consumption
14:30 (15min) Overlay maintenance plan
14:45 (15min) Avoiding duplicate descriptions
15:00 (15min) Criteria for accepting board files
15:15 (15min) Location for maintaining bindings - how to handle foreign bindings
15:30 (15min) Sharing Generic bindings
15:45 (15min) ABI Stability
16:00 (30min) [break and overflow discussion]
16:30 (20min) DT health check
16:50 (15min) devicetree.org update
17:05 (15min) EBBR Discussion
17:20 Closing and feedback
```

Devicetree Workshop 2017

```
 9:30 Welcome and Schedule bashing
 9:40 Encoding and Schema checking: Framing the problem
 9:45 DT YAML encoding overview
10:00 YAML encoding discussion
10:20 DT Schema format - option 1
10:35 DT Schema format - option 2
10:50 DT Schema discussion - what should go in the spec?
11:50 Code Generation from DT
12:10 Runtime memory consumption
14:30 Overlay maintenance plan
14:45 Avoiding duplicate descriptions
15:00 Criteria for accepting board files
15:15 Location for maintaining bindings - how to handle
         foreign bindings
15:30 Sharing Generic bindings
15:45 ABI Stability
16:00 [break and overflow discussion]
16:30 DT health check
16:50 devicetree.org update
17:05 EBBR Discussion
17:20 Closing and feedback
```

What do you want to talk about?

questions

comments

issues

concerns

Overlays - a gating factor

“a gating factor” is not meant to imply that these issues are a comprehensive list

Overlays - a gating factor

there needs to be some restrictions around what the overlays can touch.

We can't have it be wide open and then lock things down later and break users.

Overlays - a gating factor

On 10/18/17 14:46, Frank Rowand wrote:

> On Wed, 2017-10-18 at 10:44 -0500, Rob Herring wrote:

>> The issue remains that the kernel is not really setup to deal with any
>> random property or node to be changed at any point in run-time. I
>> think there needs to be some restrictions around what the overlays can
>> touch. We can't have it be wide open and then lock things down later
>> and break users.

> That paragraph is key to any discussion of accepting code to apply overlays.
> Solving that issue has been stated to be a gating factor for such code from
> the beginning of overlay development.

Overlays - a gating factor

I do not want to enable overlays when there is fundamental breakage in the implementation

Overlays - a gating factor

I do not want to enable overlays when there is fundamental breakage in the implementation

Simple real world overlay usage exists out of mainline

This does not prove lack of fundamental breakage

Overlays - a gating factor

Devicetree source files with hand coded overlay internal information are not acceptable

Overlay internal information is not a stable API, in the sense that the format can change when the DTB format version changes

Overlays - gating factor - STATUS

- restrict what overlays can touch
  ==> discussed, no momentum
  ==> connectors appear to be the way forward
  ==> apply overlay(s) early boot or pre-boot
  may be another approach

- fundamental breakage in the implementation
  ==> slow progress

- overlay internal information in source form
  ==> in the pipeline, maybe 4.15-rc1

dtc - overlays - Linux v4.15-rc1 ??

dtc creates the .dtb **OVERLAY INTERNAL DATA**

Do not hand code overlay internal data nodes
in DTS source:

fragment@
__overlay__
__fixup__
__local_fixup__
__symbols__

Currently in Rob's dt/next branch

dtc - overlays - example

```
$ diff -b -u old.dts new.dts
--- old.dts
+++ new.dts
@@ -1,13 +1,7 @@
 /dts-v1/;
 /plugin/;

-/ {
-        fragment@0 {
-                target-path = "/soc/base_fpga_region";
-                #address-cells = <1>;
-                #size-cells = <1>;
-
-                __overlay__ {
+&fpga_region {
                         ranges = <0x00000000 0x00000000 0xc0000000 0x00040000>,
                                  <0x00000001 0x00000000 0xff200000 0x00001000>;

@@ -28,6 +22,4 @@
                                 interrupt-parent = <&intc>;
                                 interrupts = <0 21 4>;
                         };
-                };
-        };
 };
```

dtc - overlays - example - old.dts

```
/dts-v1/;
/plugin/;

/ {
        fragment@0 {
                target-path = "/soc/base_fpga_region";
                #address-cells = <1>;
                #size-cells = <1>;

                __overlay__ {
                        ranges = <0x00000000 0x00000000 0xc0000000 0x00040000>,
                                 <0x00000001 0x00000000 0xff200000 0x00001000>;

                        external-fpga-config;

                        #address-cells = <2>;
                        #size-cells = <1>;

                        fpga_pr_region0 {
                                compatible = "fpga-region";
                                fpga-bridges = <&freeze_controller_0>;
                                ranges;
                        };

                        freeze_controller_0: freeze_controller@100000450 {
                                compatible = "altr,freeze-bridge-controller";
                                reg = <0x00000001 0x00000450 0x00000010>;
                                interrupt-parent = <&intc>;
                                interrupts = <0 21 4>;
                        };
                };
        };
};
```

dtc - overlays - example - new.dts

```
/dts-v1/;
/plugin/;

&fpga_region {
        ranges = <0x00000000 0x00000000 0xc0000000 0x00040000>,
                 <0x00000001 0x00000000 0xff200000 0x00001000>;

        external-fpga-config;

        #address-cells = <2>;
        #size-cells = <1>;

        fpga_pr_region0 {
                compatible = "fpga-region";
                fpga-bridges = <&freeze_controller_0>;
                ranges;
        };

        freeze_controller_0: freeze_controller@100000450 {
                compatible = "altr,freeze-bridge-controller";
                reg = <0x00000001 0x00000450 0x00000010>;
                interrupt-parent = <&intc>;
                interrupts = <0 21 4>;
        };
};
```

.dtsi source vs overlay .dtsi

With the new dtc --

Overlay .dts file contains directives:

```
/dts-v1/;
/plugin/;
```

.dtsi include file does not

Use include as .dtsi or overlay

---------- base tree ----------

```
$ expand fpga_tree.dts
/dts-v1/;

/ {
        soc {
                intc: interrupt_ctrl {
                };
                fpga_region: base_fpga_region {
                };
        };
};

/include/ "fpga_plugin_or_dtsi.dts"
```

---------- overlay ----------

```
$ expand fpga_overlay.dts
/dts-v1/;
/plugin/;

/include/ "fpga_plugin_or_dtsi.dts"
```

Use include as .dtsi or overlay

```
$ expand fpga_plugin_or_dtsi.dts
&fpga_region {
        ranges = <0x00000000 0x00000000 0xc0000000 0x00040000>,
                 <0x00000001 0x00000000 0xff200000 0x00001000>;

        external-fpga-config;

        #address-cells = <2>;
        #size-cells = <1>;

        fpga_pr_region0 {
                compatible = "fpga-region";
                fpga-bridges = <&freeze_controller_0>;
                ranges;
        };

        freeze_controller_0: freeze_controller@100000450 {
                compatible = "altr,freeze-bridge-controller";
                reg = <0x00000001 0x00000450 0x00000010>;
                interrupt-parent = <&intc>;
                interrupts = <0 21 4>;
        };
};
```

Fundamental Breakage

Locking

Memory Leaks
- drivers can NOT be expected to correctly have direct access to devicetree internal data

Apply / Remove dependencies
- frameworks
- devices
- drivers (static and modules)
- other overlays

Subsystem support

Restrict what overlays can touch

Restrict what overlays can touch

Related to uses cases

We do not need to enable every use case at the same time, but we must be aware of other use cases any time we choose how to implement a given use case

Some Use Case Categories

- add-on card exposes small set of signals
  (eg Grove connector)

- add-on card exposes most or all of SOC's pads
  (eg Beaglebone)

- fpga

Orthogonal to Use Case

- connector is stackable / daisy chain vs single target

- single connector of a given type on the board vs multiple connectors of the same type on the board

  Multiple connectors lead to wanting to use a single relocatable overlay dtb instead of hard-coding a dtb to be tied to a specific connector on the board

Use Cases -- tool

If the only tool you have is a hammer, then every problem you have looks like a nail.

Use Cases -- tool

> If the only tool you have is a hammer, then every problem you have looks like a nail.

If your problem is a bolt, then you either need to extend the capabilities of your hammer tool, or create a new tool.

Use Cases -- that are not nails

- mezzanine cpu cards
- device tree fixups
- system configuration
- more...

Overlays

- U-Boot overlay support
- Connectors (sockets and plugs)
- Overlay Managers
  - capemgr

Overlays

- Examples of use cases
  - beaglebone
  - raspberry pi
  - minnowboard
  - C.H.I.P.
  - Arduino
  - seeedstudios Grove 4 pin connectors
  - others?

Overlays

- Combinatorial explosion of .dts / .dtb files

  example:

  Devicetree Hardware Autoconfiguration
  Hans de Goede
  ELC Europe 2016

Plumbers 2017 Summary

Was not scheduled -- not enough interest / commitment

Plumbers 2016 Summary

Device Tree Schema Verification

Grant Likely (slides and etherpad)

Hardware Description vs Configuration vs Policy

(slides and etherpad)

Overlays

(etherpad)

Slides:

http://elinux.org/Device_tree_future#presentation_material_2

Etherpad Notes:

http://elinux.org/Device_tree_plumbers_2016_etherpad

Devicetree Specification

Devicetree Specification 0.1 supersedes ePAPR for the Linux kernel, continues to evolve

https://www.devicetree.org/specifications/

Mail list, Build Instructions, etc

https://www.devicetree.org/collaborate/

Repository

https://github.com/devicetree-org/devicetree-specification

Devicetree Documentation

elinux.org/Device_Tree_Reference

- becoming more complete
- contributions and comments welcome

arch/*/boot/dts/ commits

v4.2	638
v4.3	592
v4.4	666
v4.5	725
v4.6	682
v4.7	722
v4.8	674
v4.9	719
v4.10	768
v4.11	632
v4.12	658
v4.13	687

v4.2.. arch/*/boot/dts/ commits

```
  405  arc
41261  arm
 8820  arm64
   32  cris
   74  h8300
   22  metag
  793  mips
   21  nios2
  617  powerpc
   15  sh
   86  xtensa
```

arch/*/boot/dts/ commits

	arm	arm64
v4.2	568	36
v4.3	535	46
v4.4	553	67
v4.5	571	137
v4.6	511	149
v4.7	559	117
v4.8	498	156
v4.9	489	195
v4.10	519	235
v4.11	456	151
v4.12	434	217
v4.13	381	283

.../bindings/ commits

```
v4.2        488
v4.3        411
v4.4        418
v4.5        467
v4.6        428
v4.7        445    2016 07 24
v4.8        418    2016 10 02
v4.9        453
v4.10       491
v4.11       421
v4.12       468
v4.13       460
```

scripts/dtc/ commits

```
v4.2          2
v4.3          0
v4.4          0
v4.5          0
v4.6          3
v4.7          1       2016 07 24
v4.8          1       2016 10 02
v4.9          0
v4.10         0
v4.11         1
v4.12         6
v4.13         6
v4.14-rc1     0
              2  Rob's dt/next 171020
```

drivers/of/ commits

```
v4.2        29
v4.3        16
v4.4        38
v4.5        22
v4.6        29
v4.7        40      2016 07 24
v4.8        30      2016 10 02
v4.9        24
v4.10       28
v4.11       19
v4.12       32
v4.13       36
v4.14-rc1   32
            26  Rob's dt/next 171019
```

v4.7.. drivers/of/ commits

(July 24, 2016 ..)

What have patch topics been?

Very imprecise topic count (useless, but interesting)

commit short description

ignore leading “of: “

strip trailing “:.*”

sort unique and count

v4.7.. drivers/of/ commits

(July 24, 2016 ..)

```
  17 overlay
  11 of_mdio
   8 device property
   7 of/platform
   6 of/fdt
   5 of/irq
   5 irq
   4 of/numa
   4 base
   3 of_graph
   3 of/device
   3 device
   2 unittest
   2 reserved_mem
   2 overlay.c
   2 of/unittest
   2 of/resolver
   2 of_pci
   2 of/pci
   2 of_numa
   2 mdio_bus
   2 fdt
   2 drivers/of
  12 fix ...
with Rob's dt/next 171019
```

dtc compiler

dtc - Devicetree Build Warnings

Rob has been enhancing dtc error checks

Enabled for “W=1” builds

$ make V=0 W=1 qcom-apq8074-dragonboard.dtb

```
make[1]: Entering directory `/local/frowand_nobackup/src/git_linus/build/dragon_linus_4.10'
  DTC     arch/arm/boot/dts/qcom-apq8074-dragonboard.dtb
Warning (unit_address_vs_reg): Node /memory has a reg or ranges property, but no unit name
Warning (unit_address_vs_reg): Node /soc/spmi@fc4cf000/pm8941@0/vadc@3100/die_temp
Warning (unit_address_vs_reg): Node /soc/spmi@fc4cf000/pm8941@0/vadc@3100/ref_625m
Warning (unit_address_vs_reg): Node /soc/spmi@fc4cf000/pm8941@0/vadc@3100/ref_1250
Warning (unit_address_vs_reg): Node /soc/spmi@fc4cf000/pm8941@0/vadc@3100/ref_gnd h
Warning (unit_address_vs_reg): Node /soc/spmi@fc4cf000/pm8941@0/vadc@3100/ref_vdd h
make[1]: Leaving directory `/local/frowand_nobackup/src/git_linus/build/dragon_linus_4.10'
```

Kernel Configuration Info -- OLD

```
dt_to_config                                                    \
    arch/arm/boot/dts/qcom-apq8074-dragonboard.dts              \
    --short-name                                                \
    --config ${KBUILD_OUTPUT}/.config                           \
    --config-format                                             \
    > dragon_config_info
```

```
$ grep -i coincell dragon_config_info
```

```
# -d-c-----n--F : coincell@2800 : qcom,pm8941-coincell : drivers/misc/qcom-coincell.c : CONFIG_QCOM_COINCELL : n
# CONFIG_QCOM_COINCELL is not set
# CONFIG_QCOM_COINCELL=y
```


-d-c-----n--F : coincell@2800 : qcom,pm8941-coincell : .....
CONFIG_QCOM_COINCELL is not set
CONFIG_QCOM_COINCELL=y

Debug Tools -- semi-OLD

scripts/dtc/dt_prop

- Compare properties accessed on target system vs a device tree (dtX)
- available on elinux.org
- Plan to submit to mail list “any day now”

  ==> Stalled, awaiting some of Frank's bandwidth

dt_prop example snippets

```
$ dt_prop --td dmesg_4.5-rc5_160307_2100 qcom-apq8074-dragonboard.dts

# --- dmesg_4.5-rc5_160307_2100
# +++ qcom-apq8074-dragonboard.dts
 /dts-v1/;


                // *****   i2c@f9964000 disabled  *****
                i2c@f9964000 {
+                       #address-cells;
+                       #size-cells;
+                       clock-names;
+                       clocks;
                        compatible;
+                       interrupts;
+                       reg;
                        status;
                };
```

Debug Tools - OLD

dt_node_info, dt_stat

- Provide info about device tree nodes from /proc/device-tree files
- proof of concept on elinux.org
- Stalled, awaiting some of Frank's bandwidth

dt_node_info example 1

```
$ dt_node_info coincell
===== devices

===== nodes
/soc/spmi@fc4cf000/pm8941@0/qcom,coincell@2800 qcom,

===== nodes bound to a driver

===== nodes with a device

===== nodes not bound to a driver
/soc/spmi@fc4cf000/pm8941@0/qcom,coincell@2800 qcom,

===== nodes without a device
/soc/spmi@fc4cf000/pm8941@0/qcom,coincell@2800 qcom,
```

dt_node_info example 2

```
$ dt_node_info coincell
===== devices
/sys/devices/platform/soc/fc4cf000.spmi/spmi-0/0-00/

===== nodes
/soc/spmi@fc4cf000/pm8941@0/qcom,coincell@2800 qcom,

===== nodes bound to a driver

===== nodes with a device
/soc/spmi@fc4cf000/pm8941@0/qcom,coincell@2800 qcom,

===== nodes not bound to a driver
/soc/spmi@fc4cf000/pm8941@0/qcom,coincell@2800 qcom,

===== nodes without a device
```

Debug Tools - OLD

dtc: dts source location annotation

- Provide source locations from .dts & .dtsi
- Several proof of concept versions on devicetree-compiler list
- Stalled, awaiting some of Frank's bandwidth

source location annotation

```
----- short format  -----

sdhci@f9824900 { /* qcom-apq8074-dragonboard.dts:14 */
      compatible = "qcom,sdhci-msm-v4"; /* qcom-msm8974.dtsi:240 */
      reg = <0xf9824900 0x11c 0xf9824000 0x800>; /* qcom-msm8974.dtsi:241 */
      reg-names = "hc_mem", "core_mem"; /* qcom-msm8974.dtsi:242 */
      interrupts = <0x0 0x7b 0x0 0x0 0x8a 0x0>; /* qcom-msm8974.dtsi:243 */
      interrupt-names = "hc_irq", "pwr_irq"; /* qcom-msm8974.dtsi:244 */
      clocks = <0xd 0xd8 0xd 0xd7>; /* qcom-msm8974.dtsi:245 */
      clock-names = "core", "iface"; /* qcom-msm8974.dtsi:246 */
      status = "ok"; /* qcom-apq8074-dragonboard.dts:17 */
      bus-width = <0x8>; /* qcom-apq8074-dragonboard.dts:15 */
      non-removable; /* qcom-apq8074-dragonboard.dts:16 */
}; /* qcom-apq8074-dragonboard.dts:18 */
```

THE END

Thank you for your participation...

Questions?

Comments?

Resources

http://elinux.org/Device_Tree_presentations_papers_articles
http://elinux.org/Device_Tree_presentations_papers_articles#debug

http://elinux.org/Device_Tree_Reference

Resources

dtx_diff
dtc --annotate
dt_node_info

Solving Device Tree Issues:
Frank Rowand, elce 2015
http://elinux.org/images/0/04/Dt_debugging_elce_2015_151006_0421.pdf
(In this presentation, dtx_diff was named dtdiff.)

Supporting material for: Solving Device Tree Issues:
http://elinux.org/Device_Tree_frowand
section: Embedded Linux Conference Europe (ELCE) - October 6, 2015

dt_to_config

Solving Device Tree Issues - Part 2:
Frank Rowand, LinuxCon Japan 2016
http://elinux.org/images/5/50/Dt_debugging_part_2.pdf

Resources

dt_prop

Solving Device Tree Issues - Part 3:
Frank Rowand, elce 2016
http://elinux.org/images/e/e5/Dt_debugging_part_3.pdf

Supporting material for: Solving Device Tree Issues - Part 3:

kernel patches
scripts/dtc/dts_diff
scripts/dtc/dt_prop
http://elinux.org/Device_Tree_frowand
section: Resources for "Solving Device Tree Issues - Part 3" talk

How to get a copy of the slides

1) frank.rowand@sony.com

2) http://elinux.org/Device_Tree

3) http://events.linuxfoundation.org