

# Best practices for long-term support and security of the device-tree

Alison Chaiken

Mentor Embedded Software

October 24, 2013

[alison\\_chaiken@mentor.com](mailto:alison_chaiken@mentor.com)



**mentor**  
embedded

[mentor.com/embedded](http://mentor.com/embedded)

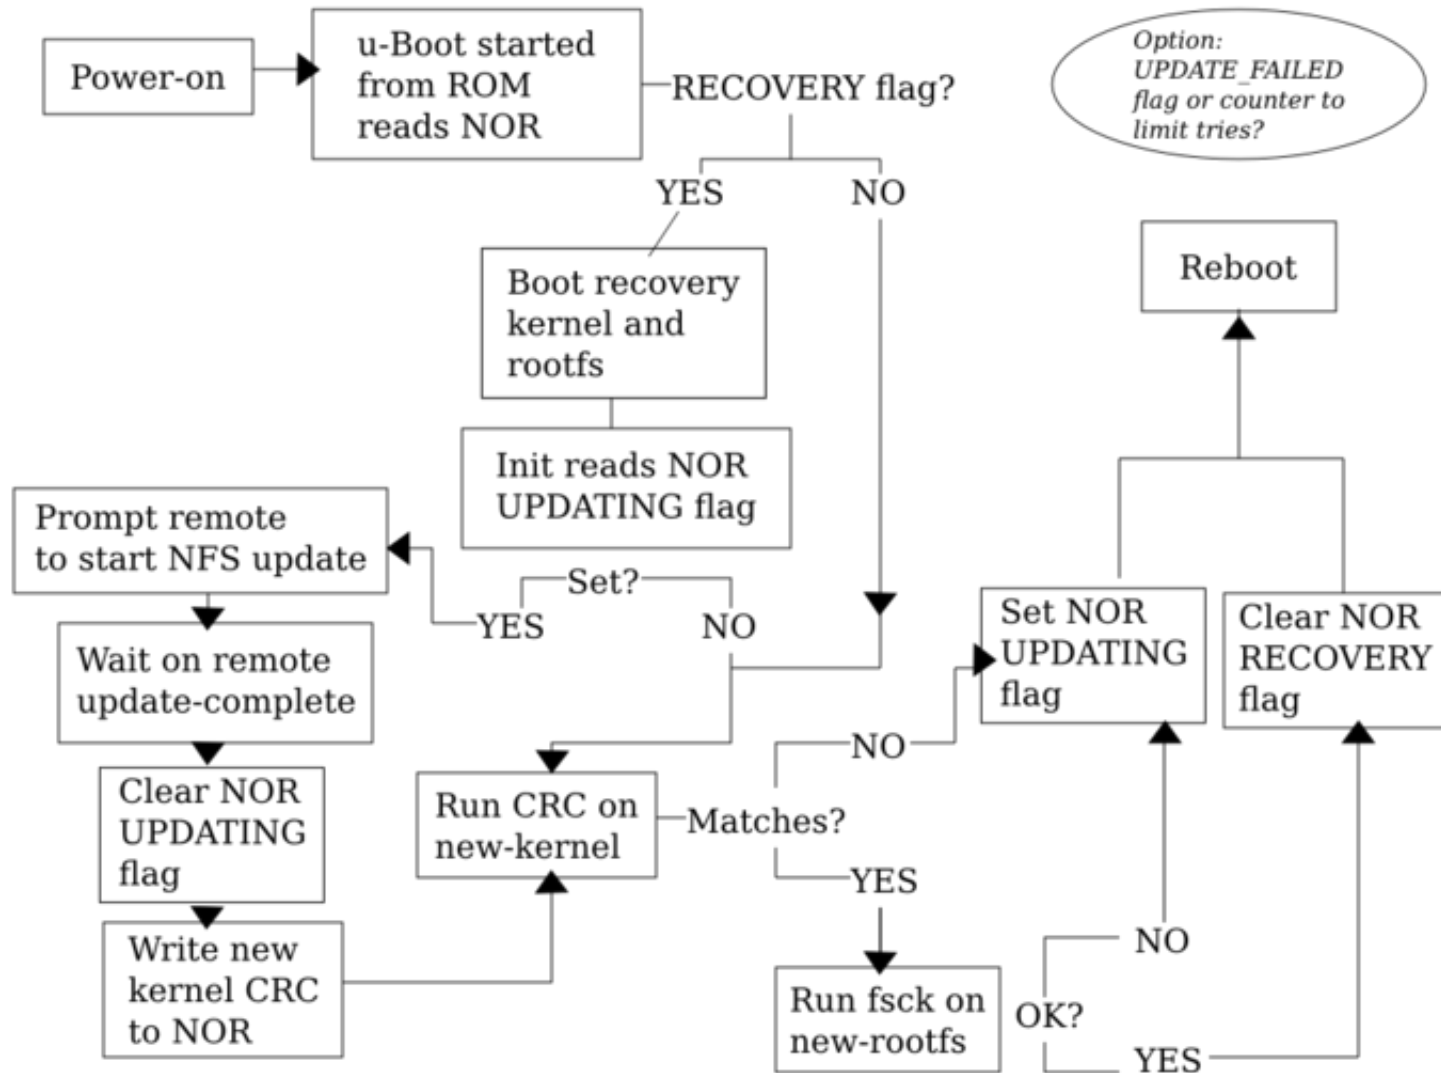
# Agenda

- How DT complicates updates
- Strategies for updates including DTBs
- Best practices in creating DTS
- Tools to make DT comprehensible to non-experts
- Best-practice summary



[mentor.com/embedded](http://mentor.com/embedded)

# Updates are already tricky without DTB



# Consider security patch in 2025 for 2015MY car



New kernel needed; will device-tree update be required?  
How to match and manage the files?

# Advocate changing DTB? No, but . . .

- Consider battery upgrades: many MCUs.
- 2013MY: 5-year battery warranty.



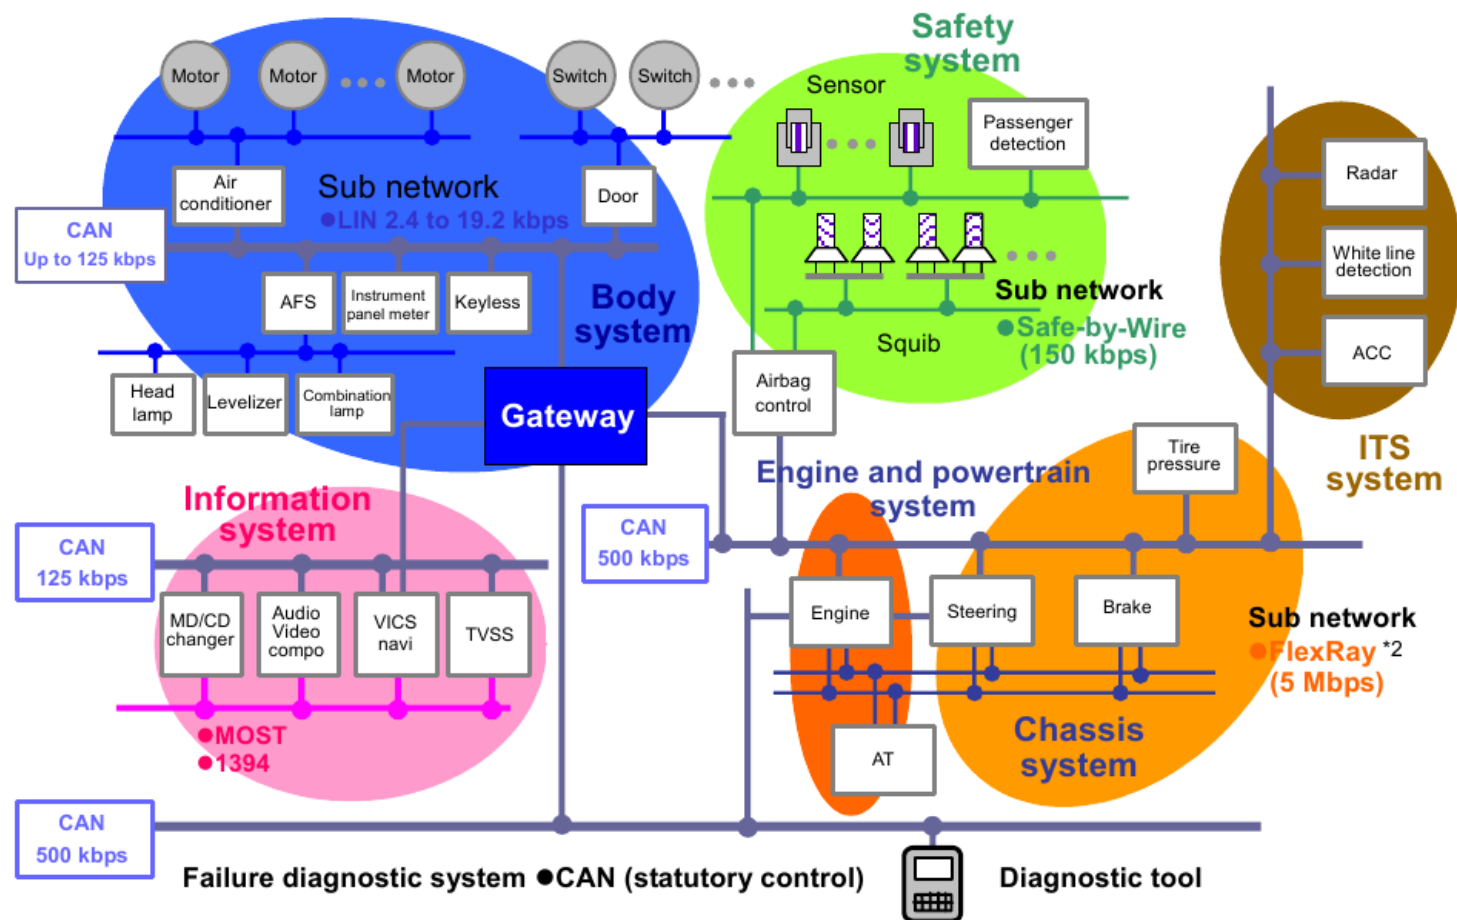
- DT's are supposed to be HW *description*.
- Lots of *configuration* already in DT's:
  - MTD partition tables; boot device selection; pinmux.
- Newer kernels have features old DTs won't support.



# Why DTB updates are hard



# Multiple processors in a LAN may need update



Copyright Renesas, "Introduction to CAN", with permission.

2014 Mercedes S-Class: 200 MCUs on Ethernet, CAN-FD and LIN.

# Careful: DTS and Kconfig must be compatible

From: Jason Gunthorpe <jgunthorpe <at> obsidianresearch.com>  
Subject: **Re: DT bindings as ABI [was: Do we have people interested in device tree janitoring / cleanup?]**  
Newsgroups: [gmane.linux.kernel](#), [gmane.linux.ports.arm.kernel](#)  
Date: 2013-07-25 21:37:53 GMT (7 weeks, 3 days, 4 hours and 54 minutes ago)

We use DT has a kernel configuration input. Our environment is designed to guarantee 100% that the kernel and DT match exactly. DT very deliberately isn't an ABI boundary in our systems.

Drivers move and CONFIG options change name over time.



# *Unintentionally* changing drivers

- *Obvious*: driver behavior is modified when its DT node is.
- *Less obvious*: Driver behavior changes if clock, pinmux or PS voltage are modified by *another driver's* DT node.
- Koen Kooi:
  - “IMPORTANT: booting the existing am335x-bone.dts will blow up the HDMI transceiver after a dozen boots with an uSD card inserted because LDO will be at 3.3V instead of 1.8. MMC support for AM335x still isn't in . . .”



# Three possible approaches



# Flattened Image Tree addresses LTS problems

- .its files record versions of DTS, kernel, misc. files and their CRCs.
- Enables OTA with one transferred archive.
- ChromeOS “Verified Boot” provides extensions for signed images.
- Glass added signing of .itb's to U-boot.
- Introduced by [Fernandes at ELC2013](#).
- Alternative: record selections in Bitbake recipe or Android-like manifest.

# DTS runtime configurability via Overlays

- Implements hot-plug for “capes”.
- *Proposal:* use overlays as a DTB update method.
  - Not suitable for security improvements or boot devices.
  - Similar in spirit to unionfs in Knoppix.
- See **Pantelis' talk** and **paper**.

# Hypervisors in device-tree

## **[PROPOSAL] ARM/FDT: passing multiple binaries to a kernel**

*From: Andre Przywara Date: Tue, 03 Sep 2013*

Example:

```
/chosen {
    #size-cells = <0x1>;
    #address-cells = <0x1>;
    module@0 {
compatible = "xen,linux-zimage", "xen,multiboot-module", "boot,module";
        reg = <0x80000000 0x003dcff8>;
        bootargs = "console=hvc0 earlyprintk ro root=/dev/sda1 nosmp";
    };
    module@1 {
compatible = "xen,linux-initrd", "xen,multiboot-module", "boot,module";
        reg = <0x08000000 0x00123456>;
    };
};
```

## Another path to safe updates?



# Making DTS more maintainable



# [RFC 00/15] Device Tree schemas and validation

- *From:* Benoit Cousson <[bcousson@xxxxxxxxxxxxx](mailto:bcousson@xxxxxxxxxxxxx)>
- *Date:* Tue, 24 Sep 2013 18:52:06 +0200

=== What is a schema? ===

A schema is a file describing the constraints that one or several nodes of a dts must conform to. Schemas must use the file extension ".schema". A schema can check that:

- A node has a given property
- An array has a valid size
- A node contains the required children.
- A property has the correct type.
- A property has the correct value.

# ***Proposal: explicit inheritance in design***

- Create schemata matching device classes: board, cpu, platdev, <arch>
  - Each CPU node would validate against cpu.schema,
  - each ARM core node against arm.schema,
  - each V4L2 camera against v4l2\_capture.schema.
- Validated DTSI compiled into top-level DTS.
- New validator enforces inheritance.



# Consider XML for DT schema?

- Suggested on LKML by Arend van Spriel
- Lots of existing manipulation tools for XML.
- Reaction basically favorable:

From: James Bottomley <James.Bottomley@HansenPartnership.com>  
Subject: **Re: [Ksummit-2013-discuss] DT bindings as ABI [was: Do we have people interested in device tree**  
Newsgroups: [gmane.linux.ports.arm.kernel](#), [gmane.linux.kernel](#)  
Date: 2013-07-28 05:11:16 GMT (7 weeks, 10 hours and 56 minutes ago)

XSLT is a transform language ... you'd use it say to transform xml to dtc, so it would be an integral component of an xml/xslt based schema.

If you want actually to describe and have validated the xml schema itself, then you'd use xsd (XML schema description language) and its associated tools.


I'm not saying you **\*should\*** do this, just that it's possible (plus I've just blown my kernel cred by knowing about xml, sigh).

James

# JSON: The Fat-Free Alternative to XML

Presented at XML 2006 in Boston, December 6 by Douglas Crockford

“JSON is a natural representation of data for the C family of programming languages.”



This is an example of a JSON object.

```
{
  "name": "Jack (\\"Bee\\" ) Nimble",
  "format": {
    "type": "rect",
    "width": 1920,
    "height": 1080,
    "interlace": false,
    "frame rate": 24
  }
}
```

"translated into Chinese, French, German, Italian, Japanese, Korean, and Spanish. . . . formalized in RFC 4627. . . . support for . . . C, C++, C#, . . . Erlang, Haskell, Java, . . . Lua, Perl, Objective-C, . . . PHP, Python,. . . Ruby, ."

# Level separation

- Explicit inheritance in nesting of device-trees:

```
$ grep okay linux/arch/arm/boot/dts/*dt si | wc
```

```
181    732  11530
```

```
$ grep chosen linux/arch/arm/boot/dts/*.dt si | wc
```

```
22     80   1103
```

- Decisions about shared resources in low-level include files = premature optimization!
- Caution with status “okay”, “chosen,” “config-on-boot” and configuration of (power-supply, pinmux and clocks) at “leaf” level.

# Tools: existing and desirable

- DTC extensions for schema validation.
- Verbose mode for DTC.
- Tools to answer the questions:
  - “Do my Kconfig and DTS match?”
  - “Which of this huge set of DTS files can produce the DTB in the shipping product?”
- lshw-like tool to present /proc DT state.
- Graphviz/dotty extensions to visualize FDT.
- Support for signing DTBs which are not appended.
- Put dt.gz in /boot along with other config.gz?

# What DTS output does preprocessor produce? *makedts* for ARM and x86

Advantage over fdt dump: outputs strings as characters, not ASCII codes.

Invoke with 'makedts <full-path-to-dts-file>' :

```
#!/bin/bash
DTC_CPP_FLAGS="-E -Wp,-MD,$BASE.pre.tmp -nostdinc /
-larch/$SRCARCH/boot/dts -larch/$SRCARCH/boot/dts/include /
-undef -D__DTS__ -x assembler-with-cpp"

#run C preprocessor
$CC $DTC_CPP_FLAGS -o $BASE.tmp $1

#run DTC
scripts/dtc/dtc -O dts -o $BASE.out.dts -b 0 -i arch/arm/boot/dts /
-d $BASE.dtc.tmp $BASE.tmp
```

For x86, script first compiles dtc itself.

# Summary: Best practice candidates

- Preserving *sets* of {bootloader configs, DTS sets, Kconfig and kernel sources}
  - FIT? BB recipes? Android-style “repo” manifests?
- Validation of device-trees:
  - Via new schemata and Warren's [binding checklist](#).
  - CRC-based via bootloader.
- Design of *maintainable* device-trees via
  - Level separation.
  - Overlays as update mechanism to failsafe DTB.
  - Update into hypervisor domains.

# Conclusions



[Weekly edition](#)  
[Archives](#)

[Kernel](#)  
[Calendar](#)

[Security](#)  
[Subscribe](#)

[Distributions](#)  
[Write for LWN](#)

[Contact Us](#)  
[LWN.net FAQ](#)

[Adzerk](#)  
[Google AdSense](#)

## Object-oriented design patterns in the kernel, part 1

Logged in as  
**alison**

[My Account](#)  
[Unread comments](#)  
[Log out](#)

**Weekly Edition**  
[Return to the](#)

Despite the fact that the Linux Kernel is mostly written in C, it makes broad use of some techniques from the field of object-oriented programming. Developers wanting to use these object-oriented techniques receive little support or guidance from the language and so are left to fend for themselves. As is

**June 1, 2011**

This article was contributed  
by Neil Brown

- DT = application of agreed best practices to kernel.
- Much criticized, but change is hard!


# Following slides are extras



# x86 users will want DTS, too

## *Announcing the* **Quark Family of SoCs**

Open Architecture  
Standard Industry Software Support  
Fully Synthesizable



The key word for most analysts in that statement is synthesizable, which means that customers can add their own IP around the core. ARM for example let's companies license its CPU core and then add their own co-processors, or other components to create chips optimized for a wide variety of projects and industries. How they would do this in practice is unclear as Mangano says that

GigaOm, 9/10/2013

# Implications of “Device trees as ABI”

- Export all bindings to header file?
- Put dt.gz in /boot along with other config files?
- Warn about unstable bindings with dtc or require explicit dtc switch to include unstable bindings?




**Logged in as**  
**alison**  
[My Account](#)  
[Unread comments](#)

By **Jonathan Corbet**  
July 30, 2013

## Device trees as ABI

Last week's [device tree article](#) introduced the ongoing discussion on the status of device tree maintainership in the kernel and how things needed

# A book about design patterns for embedded



# Uncle Bob Martin's 5 Agile Design Principles

1. Dependency Inversion Principle
2. Interface Segregation Principle
3. Liskov Substitution Principle
4. Open-Closed Principle
5. Single Responsibility Principle

From 1998, C++ Report, available at [objectmentor.com](http://objectmentor.com)