Tools and Techniques for Finding Hidden Performance Bottlenecks

Presented by: Bill Roman
Software Architect, Datalight

Santa Clara, CA USA
August 2008

Agenda

- Define Performance
- Analyzing Performance Data
  - Raw access
  - Driver level
  - File system / application level
- Use Case Analysis

Datalight

Performance:How/Where?

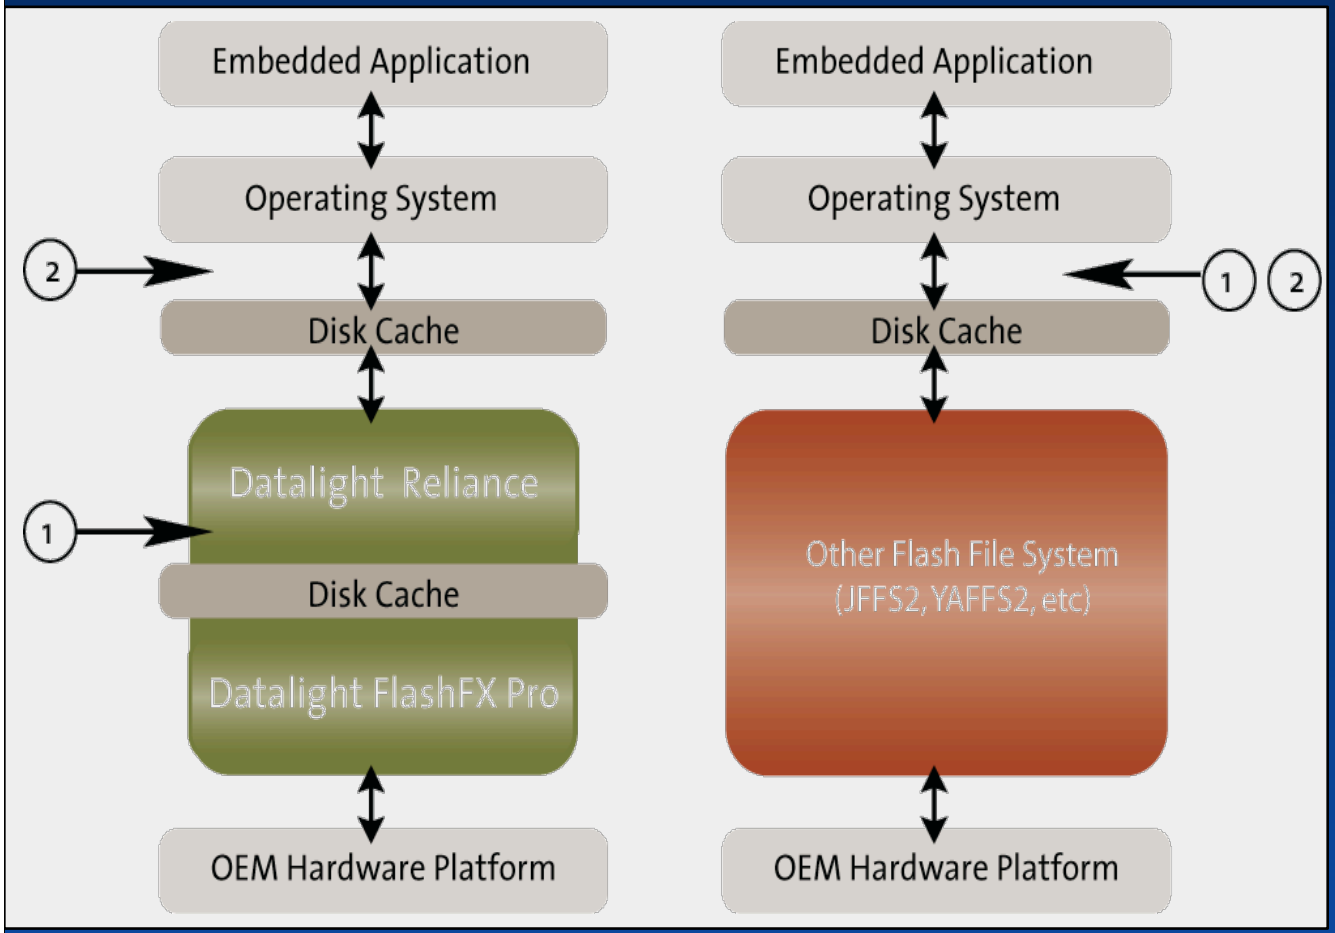

1. When measured here, the Datalight flash file system platform delivers 400% faster sequential write performance than JFFS2

2. When measured here, the Datalight flash file system platform delivers 2000% faster sequential write performance than JFFS2

Note: Write performance does not include erase time. Our experience suggests that accommodations for erase time impact would reduce write speed to approximately 2 MB/s (from 2.5 MB/s) for the Datalight platform and reduce the performance of JFFS2 by a similar percentage.

Performance Measurements

- Hardware
  - Throughput Rates
  - Bus & CPU Speeds
  - MIPS & Power Efficiency
- Software
  - Algorithm Efficiency
  - Software Layers
- User Experience
  - Boot Speed
  - Device Responsiveness

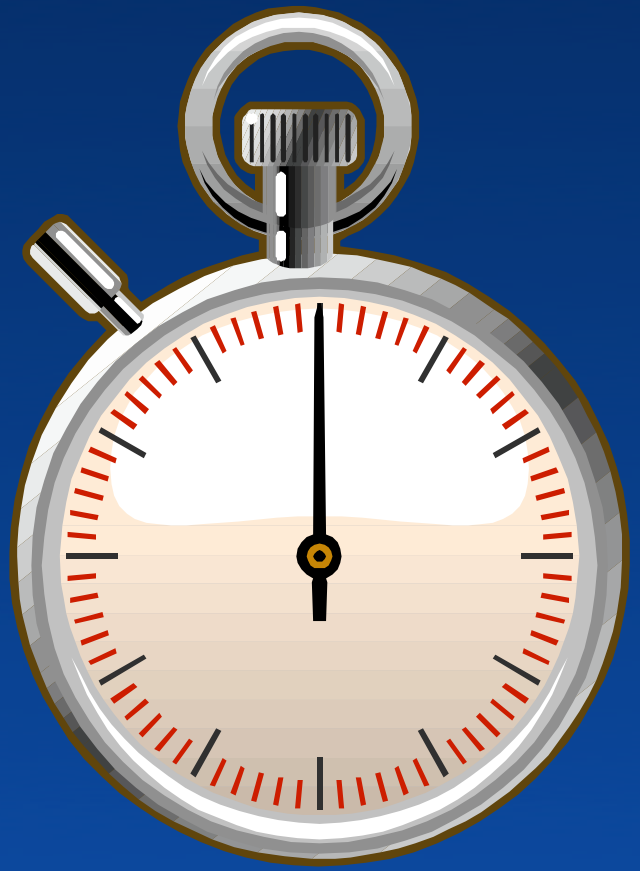

Datalight

Ordering performance

- Hardware Interface level
  - Raw access
  - Mount Speed
- Driver level
- Application level
  - CPU Utilization
  - Boot Speed
  - Buffers & Cache
  - Use case

Datalight

Mount Speeds

- The rate at which a device is able to initiate a file system and make it available for use
- In case of boot scenarios, mount speed is very important, especially if the boot image is on the file system
- Faster mount speeds hence enable quicker boot times

Hardware Interface

- Today's example:
  - FMSLTest
    - RAW flash access
    - Source code available in <install dir>/common/tests
      - Fmsltst.c
      - Fmsltst.h
      - Dlperflog.h

Datalight

FMSL Test Output Example

```
FlashFX FMSL Unit Test Datalight FlashFX
Pro v3.3.0
Linux Edition for 386
Copyright (c) 1993-2007 Datalight, Inc.
Patents: US#5860082, US#6260156.
Random Seed = 21010
Testing FMSL performance
     Write:      80807 KB per second
     Read:       213960 KB per second
     Erase:      1690 KB per second
```

What does raw level testing reveal?

- Difference between data sheet numbers and in system maximums.
- Analyze the system bus.
- Specific flash routines can be implemented to improve specific operations:
  - Multi page plane support
  - Cache read modes
  - Burst write modes
  - Interleaving

Datalight

Driver Level Performance

- Today's Example
  - VBF Test
    - Sector emulation
    - Compaction routines
    - Wear leveling
    - Source code available in <install dir>/common/tests
      - Vbftst.c
      - Vbftst.h
      - Dlperflog.h

VBF Test Output Example

```
FlashFX VBF Unit Test
Datalight FlashFX Pro v3.3.0
Linux Edition for 386
Copyright (c) 1993-2007 Datalight, Inc.
Patents: US#5860082, US#6260156
Testing VBF performance
     Write:              42862 KB per second
     Read:              200928 KB per second
     Discard:             9100 KB per second
     Partition mount:        5 ms per mount
```

What does driver level testing reveal?

- Disk efficiency
- Improvements can be made through:
  - Cushion size
  - Cache use
  - Background Compactions

Datalight

Application Level Performance

- Script for use case
- Provides sustained scenarios
  - Invokes garbage collections
  - Invokes block erases
  - Injects read and write errors / corrections
- Defines customer expectations

Datalight

What does application level testing reveal?

- Relationships between the application, file system interface, and flash driver.
- Optimizations can be made in:
  - Discard interface
  - File system efficiency
    - Relationship between the amount of data requested for a write and the actual amount of data written.
- Application use case analysis

Datalight

Summary / Q & A

- Test performance across the platform
- Targeted testing guides development efforts
- Starting with application level testing is inefficient

Datalight