

Logical Block Protection

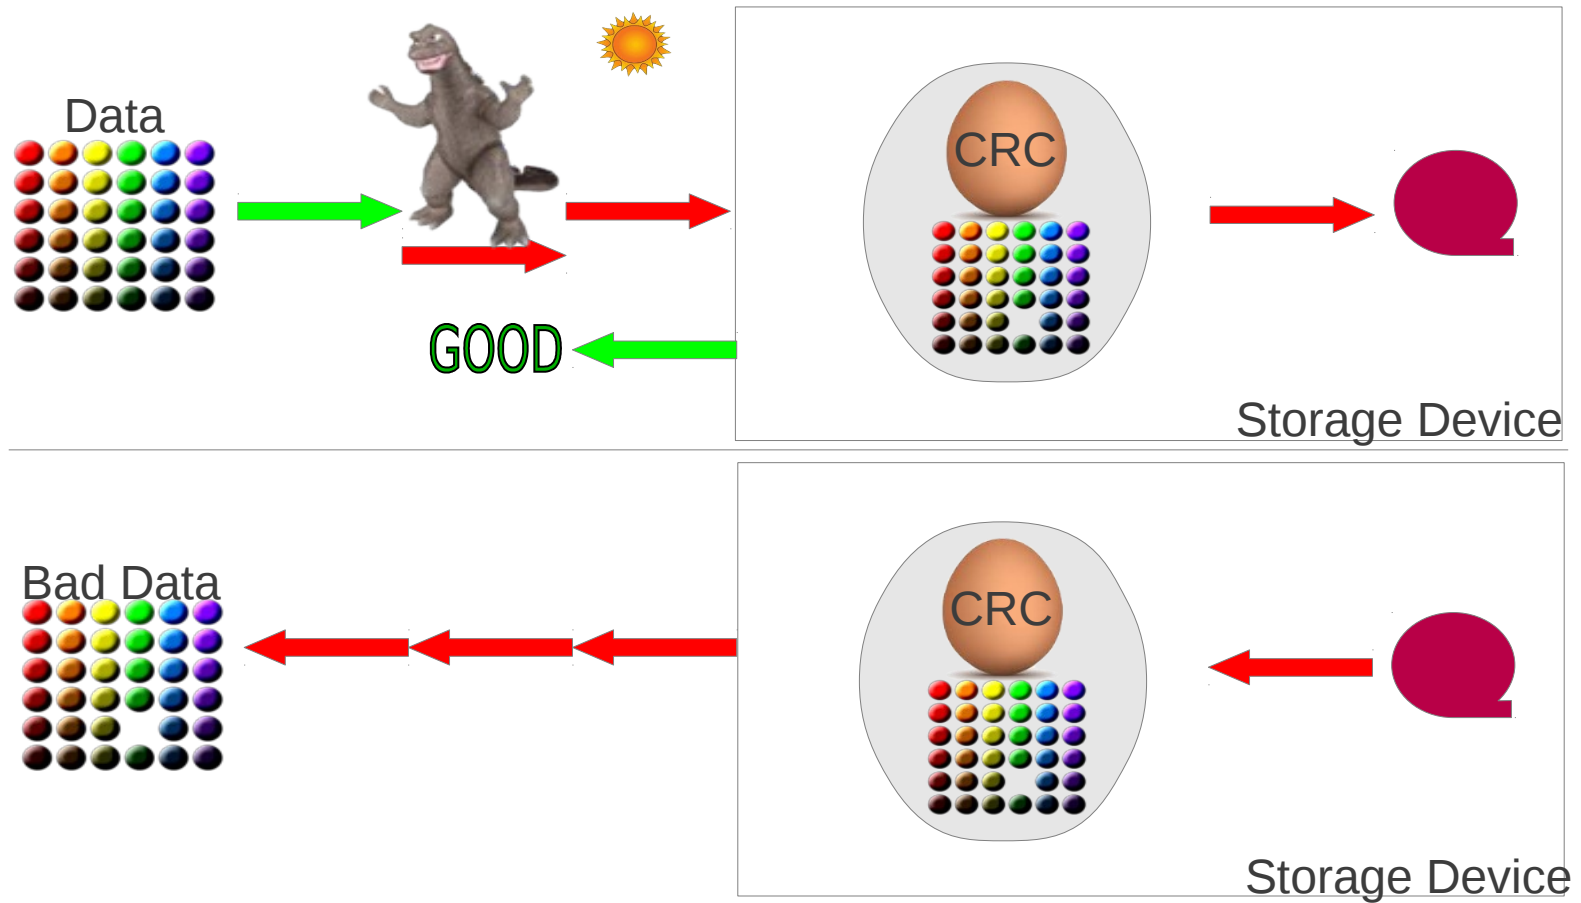
Joseph Nemeth

Tech Lead, Storage Software


Judith Schmitz

MTS, Storage Software



Device CRC...with line errors




Application Generated CRC



Device Generated CRC





Logical Block Protection

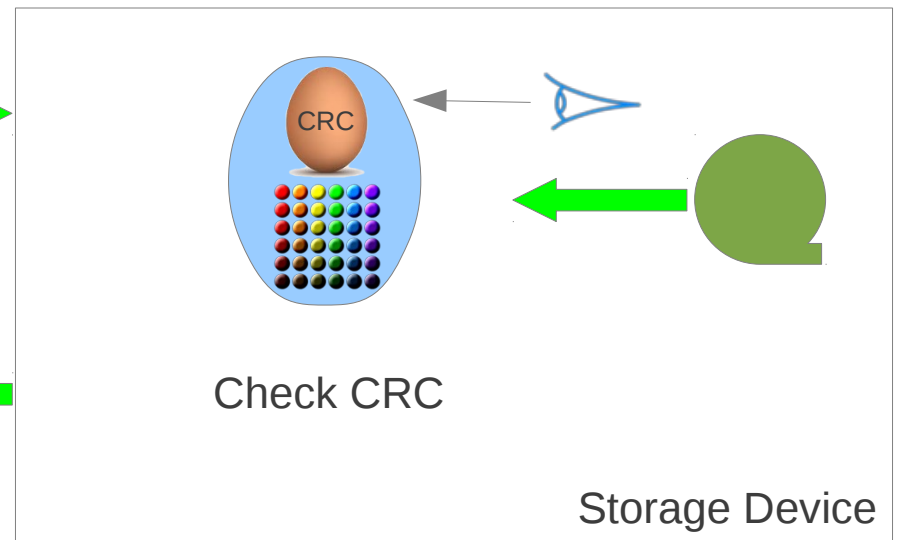



Tape Verification with LBP

Application
CRC



Logical
Block
Protection



Benefits of LBP

- Detects data transmission corruption
 - Immediately
 - Without writing to media
- Allows record-by-record data verification by storage device
 - Without data transmission
 - Does not burden application host with data
 - Can run at full media speedext

Availability in Drives

- Some tape drives now support LBP
 - LTO5 (with proper firmware) & later
 - T10KC (with proper firmware) & later
 - Some vendors call it Data Integrity Validation
- Drive vendors specify the CRC algorithms that are supported with LBP
 - LTO : Reed-Solomon
 - T10K: Reed-Solomon or CRC32C

DMF Support

- DMF 6.1 (ISSP3.1) will support LBP on T10K drives, using CRC32C
- Use of LBP by DMF is optional
 - Specified in the dmf.conf file
 - **LOGICAL_BLOCK_PROTECTION on** in **volumegroup** stanza
 - Off by default

Read Checksum Error: 2 ways

- Software: when DMF does verification
- SCSI: when drive reads the tape block
- Handled in dmatrc
 - Tries to read the next block from the tape
 - If that is good, writes that to the file
 - Tries to recover the bad blocks of the file from a 2nd copy.

Write Checksum Error

- SCSI: when drive receives the tape block for write
- Handled like any other I/O error we'd get when writing to the tape.
 - The DMF process stops writing to the tape.
 - If errors continue with subsequent processes, DMF will either set the HLOCK or HVFY flag on the volume.

New Media Format

- DMF has historically checksummed its tape blocks, but
 - Did not use either of the above algorithms
 - The checksum was not placed at the end of the tape block
- So DMF uses a new format for its tape blocks when LBP is enabled

Enabling the New Media Format

- You can enable the new format in `dmf.conf` without LBP
- **CHECKSUM_TYPE `crc32c`** in the **`volumegroup`** stanza
- Default is **legacy**

More on the New Media Format

- DMF won't switch formats in the middle of a tape
 - The tape is written either in the old format (default), or it is written in the new format.
- Tape labels are not checksummed.
- No new data in the label
 - When reading existing tapes, DMF detects checksum type from the first block.

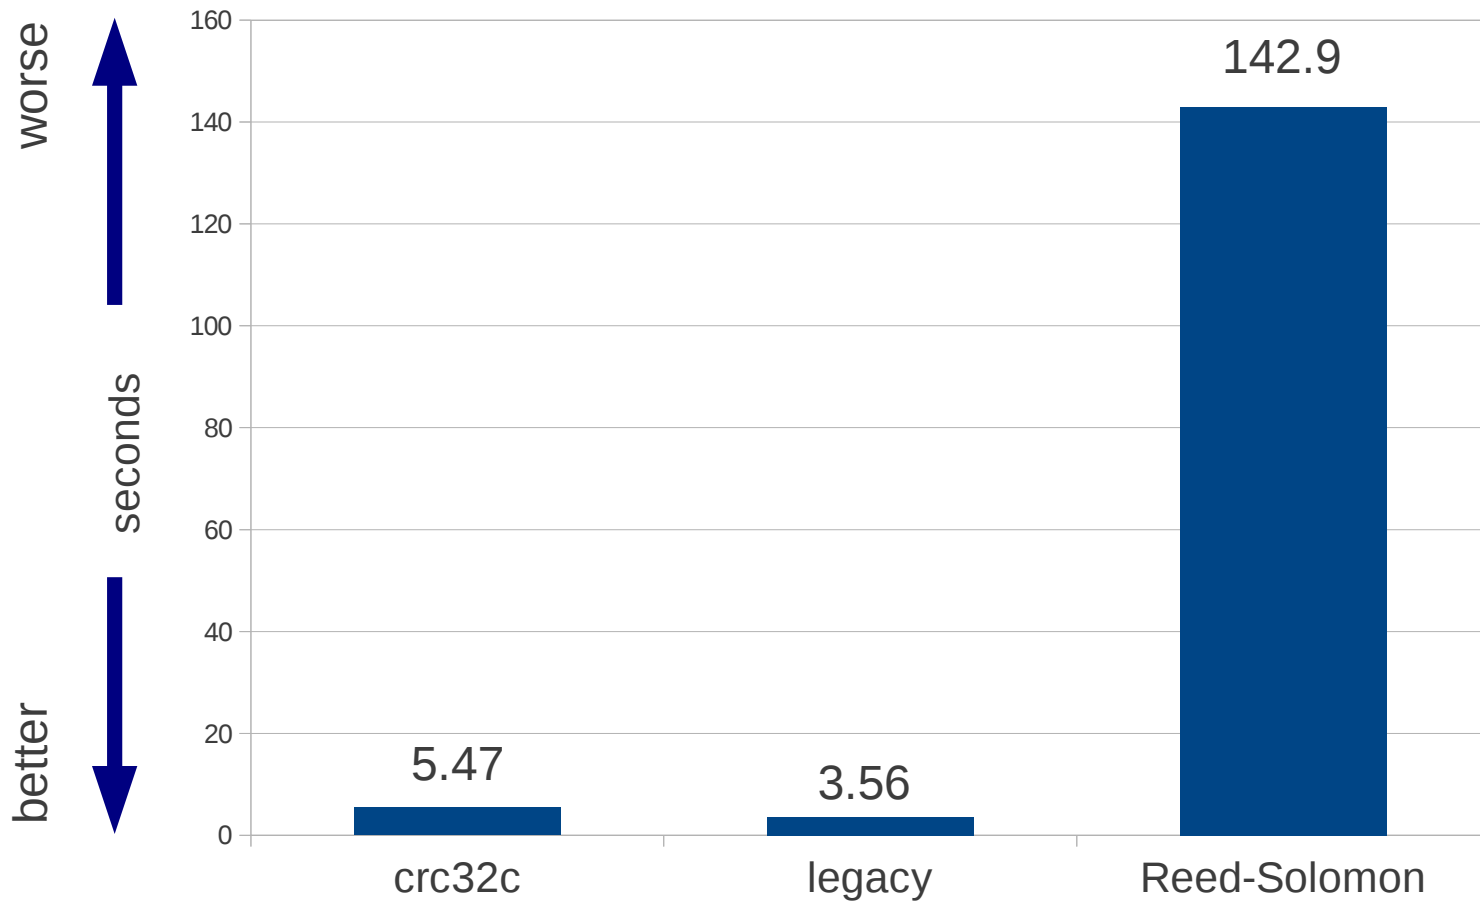
Hardware Support

- If the CPU hardware supports it, the CRC32C algorithm can be performed by specialized hardware instructions.
 - Intel's CRC32 instruction in SSE4.2

Checksum Speed Comparison

- The checksum for a 524288-byte block was computed 100000 times
- `crc32c`
 - with Intel 64 bit CRC32 instruction
- `legacy`
 - DMF's 16 bit checksum, software
- Reed-Solomon code, software

Checksum Speed Comparison (2)



sggi

