



I/O Fault Diagnosis in Software Storage Systems

Eric Schrock
Staff Engineer
Sun Microsystems

About This Presentation

- Based on experience developing Sun Storage 7000 series of appliances
- Leverages Solaris technologies
 - > Fault Management Architecture
 - > ZFS
- Principles are generic to all systems

What is Fault Diagnosis?

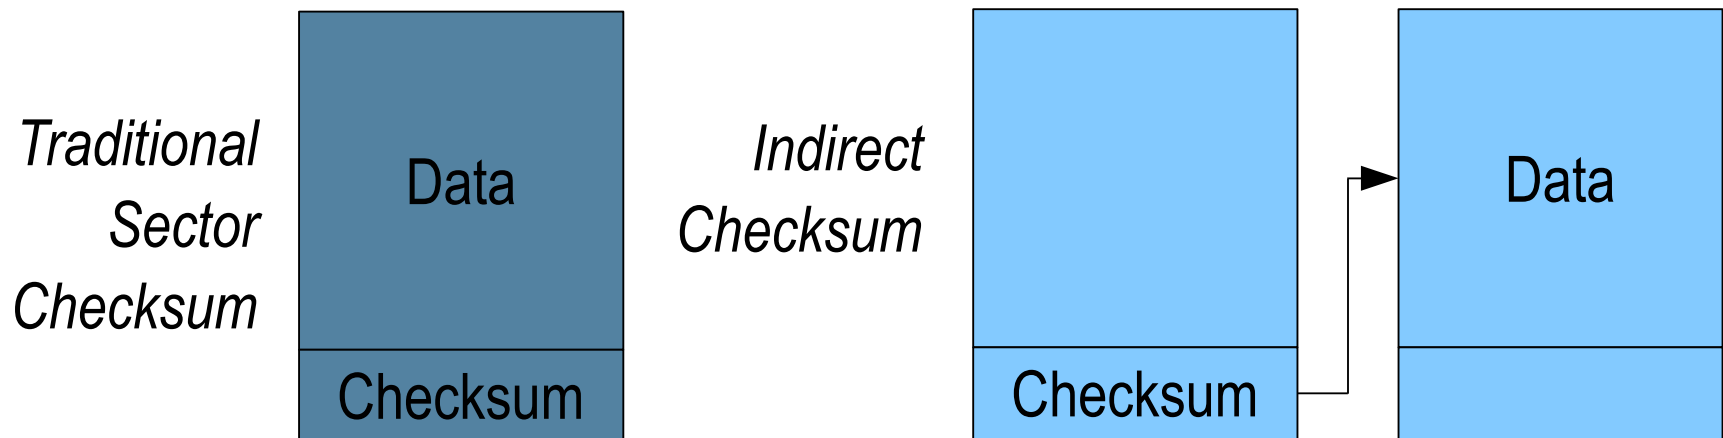
- Distinguish between *errors* and *faults*
 - > Traditional systems present a series of errors
 - > Root cause requires detailed knowledge of the system
 - > Often results in incorrect diagnosis
- Provide rich structured telemetry
- Clear indication of affected component
 - > FRU – Field Replaceable Unit
 - > ASRU – Automatic System Recovery Unit
- Common diagnosis and reporting infrastructure

I/O Subsystem Diagnosis

- Core system diagnosis
 - > CPUs, DIMMs, PCI devices
- Transport diagnosis
 - > Fatal SCSI errors
- Health monitoring
 - > SMART, overtemp, self-test failure
- Fabric diagnosis
 - > Bad cable in multipathed config
 - > Use of SMP to understand SAS topology

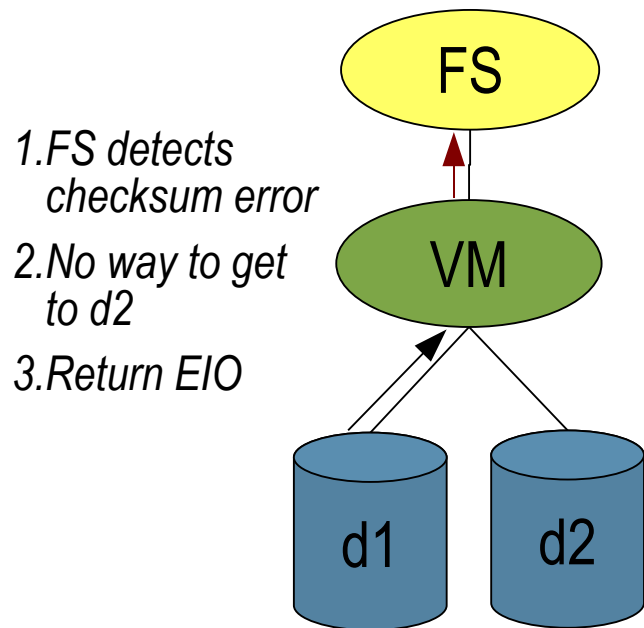
Filesystem Diagnosis

- Everything must be checksummed
 - > Not all faults manifest as transport errors
- Checksum must be stored in separate LBA
 - > Or location-specific identifier included in checksum
 - > Protects against phantom reads, misdirected writes

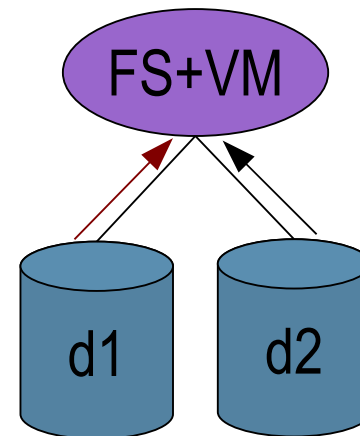


Integrated Volume Manager

- FS knows where disjoint checksum is stored
 - > Needs to request other copies on bad checksum
 - > Requires integration with volume manager



1. FS detects checksum error
2. Gets good data from d2
3. Corrects bad data on d1
4. Return good data



Diagnosing a fault

- Feed errors to *diagnosis engine*
- Correlate telemetry to create list of suspects
 - > SERD engines commonly used
- Generate fault against physical component
 - > Independent of error source
- Proactive repair
 - > Activate hot spare, retire pages, offline CPU core
- Present human-readable description of fault
 - > Details, impact, recommended action, response

Physical FRUs

- OS device abstractions have no physical meaning
 - > /dev/hdd0, /dev/dsk/c1t50000fea45620bd0
 - > Identification impossible with large numbers of disks
- Unified topology for describing disks
 - > Internal disks enumerated using XML maps
 - > External enclosures enumerated using SES
- Human-readable labels
 - > *“The number of I/O errors associated with disk 'JBOD40/HDD 6' has exceeded acceptable levels.”*

Fault Management

- Awareness of faults
 - > Email, SNMP notifications
 - > Phone home to vendor or central location
- Provide means of identifying FRUs
 - > Human readable labels
 - > Fault LEDs
 - > Highlight within system images
- N-way management
 - > Aggregate faults across machines
 - > Single “dashboard” for viewing state of many machines

One Vendor Plug

“What is your biggest problem with storage?”

- Find and harass me during break



I/O Fault Diagnosis in Software Storage Systems

Eric Schrock

Eric.Schrock@sun.com

<https://blogs.sun.com/eschrock>