Solaris Volume Manager : RAID Overview

SVM RAID Overview

- RAID Advisory Board
- RAID Levels

`

RAID Advisory Board

- Industry Body
- Defines official RAID levels
- Unofficial levels also exist
- RAID vs EDAP
  - > RAID – defines on-disk layout of data
  - > EDAP – defines broad criteria for whole storage system

RAID Levels – RAID 0

- Striping with no parity
- No data protection
- Data written in chunks
  - Chunk size is defined by the interlace value
- Performance improved by spreading I/O across drives

RAID Levels – RAID 0

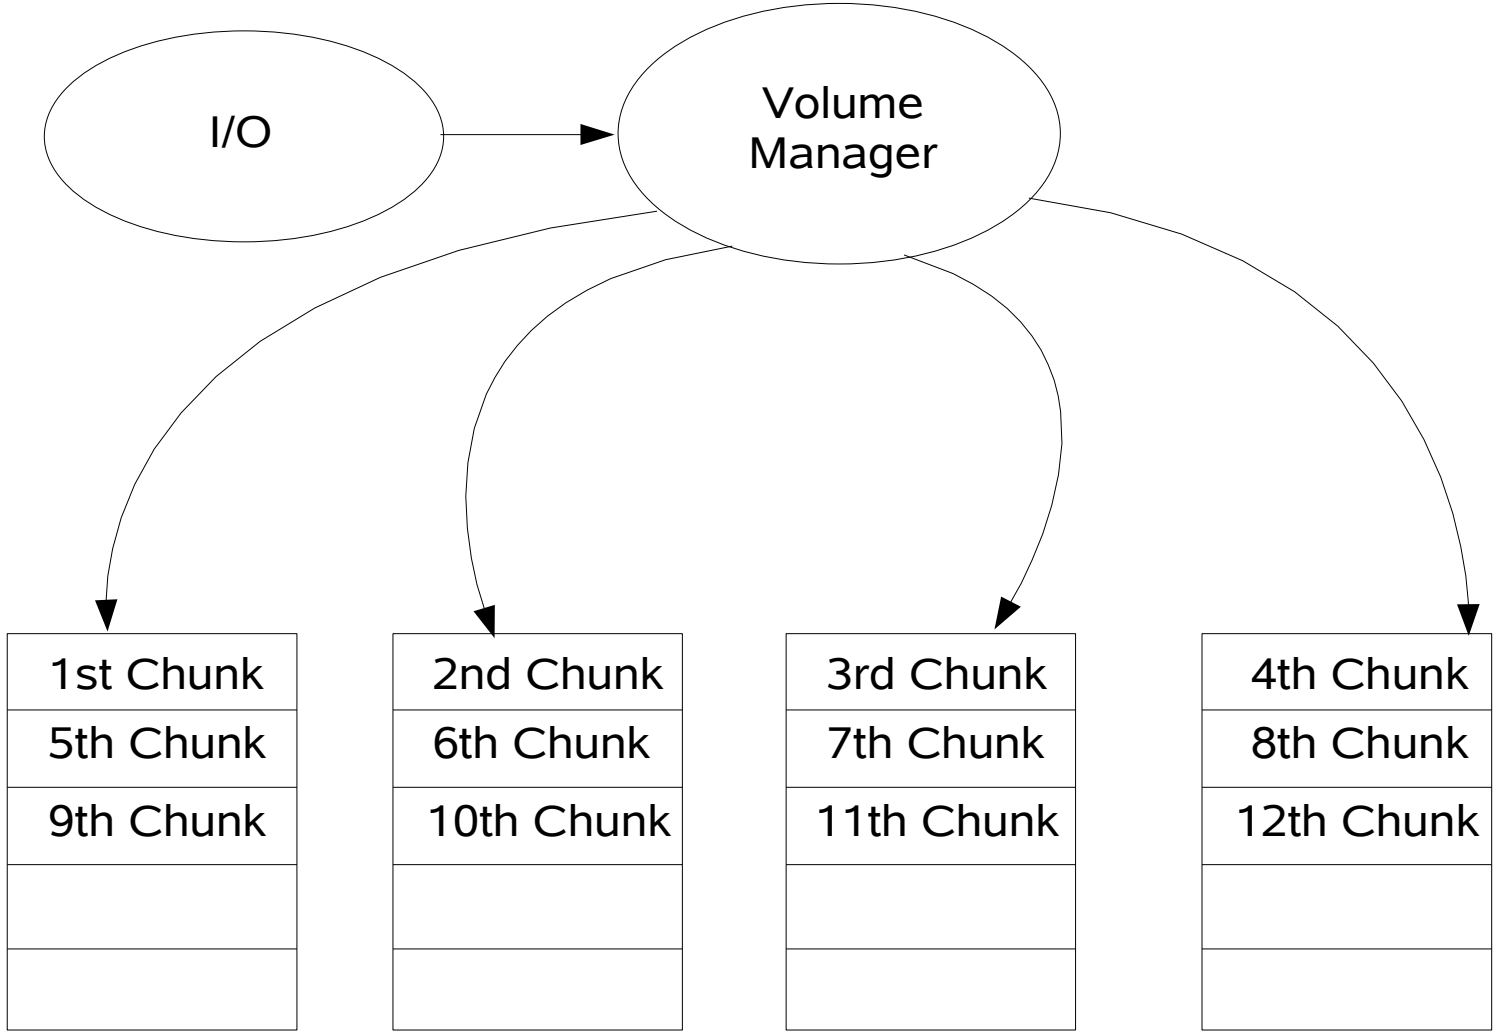

RAID Levels – RAID 1

- Disk Mirroring
- Complete data protection
- Data written to multiple copies
- Write performance degrades through multiple I/O's
- Read performance can improve

RAID Levels – RAID 1

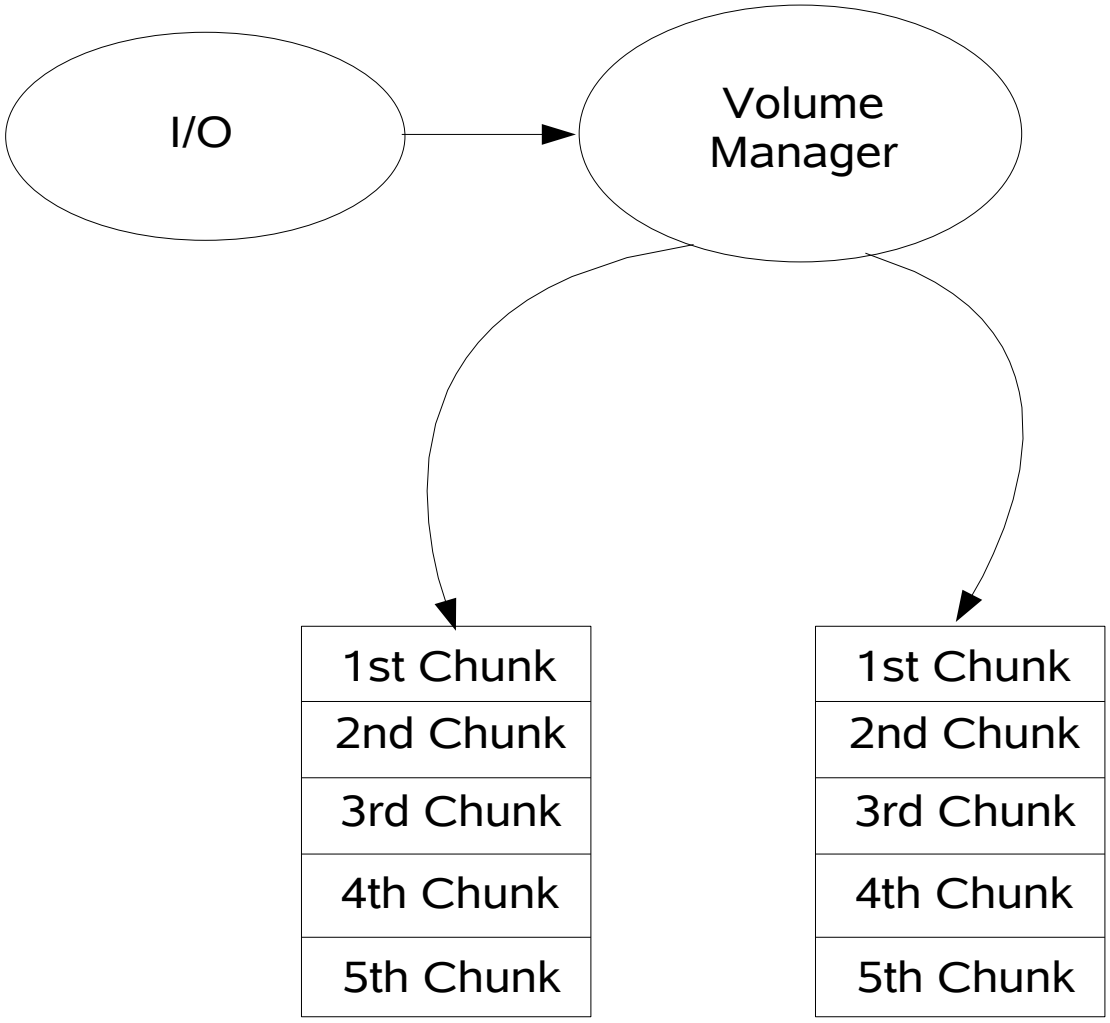

RAID Levels – RAID 2

- Striping with dedicated ECC data disk
- Data protection for single disk failure
- Not used in practice

RAID Levels – RAID 3

- Striping with dedicated parity disk
- Data protection for single disk failure
- Data written in chunks – as a stripe
  - Must write a complete stripe, with parity each time
- Read performance improved – as a stripe
- Write performance can be poor
  - Only good for full-stripe writes
  - Partial stripe writes must read old data and re-generate parity
- Rarely Used

RAID Levels – RAID 4

- Striping with dedicated parity disk
- Data protection for single disk failure
- Much larger interlace values than RAID-3
  - > Each read can be drawn from a single spindle
  - > Allows other reads to be serviced from other spindles
- Not used in practice

RAID Levels – RAID 5

- Striping with distributed parity data
- Data protection for single disk failure
- Data written in chunks – as a stripe
  - > Must write a complete stripe, with parity each time
  - > Parity data spread through all drives to avoid hot-spots
- Read performance improved – as a stripe
- Write performance can be poor
  - > Only good for full-stripe writes
  - > Partial stripe writes must read old data and re-generate parity

RAID Levels – RAID 5

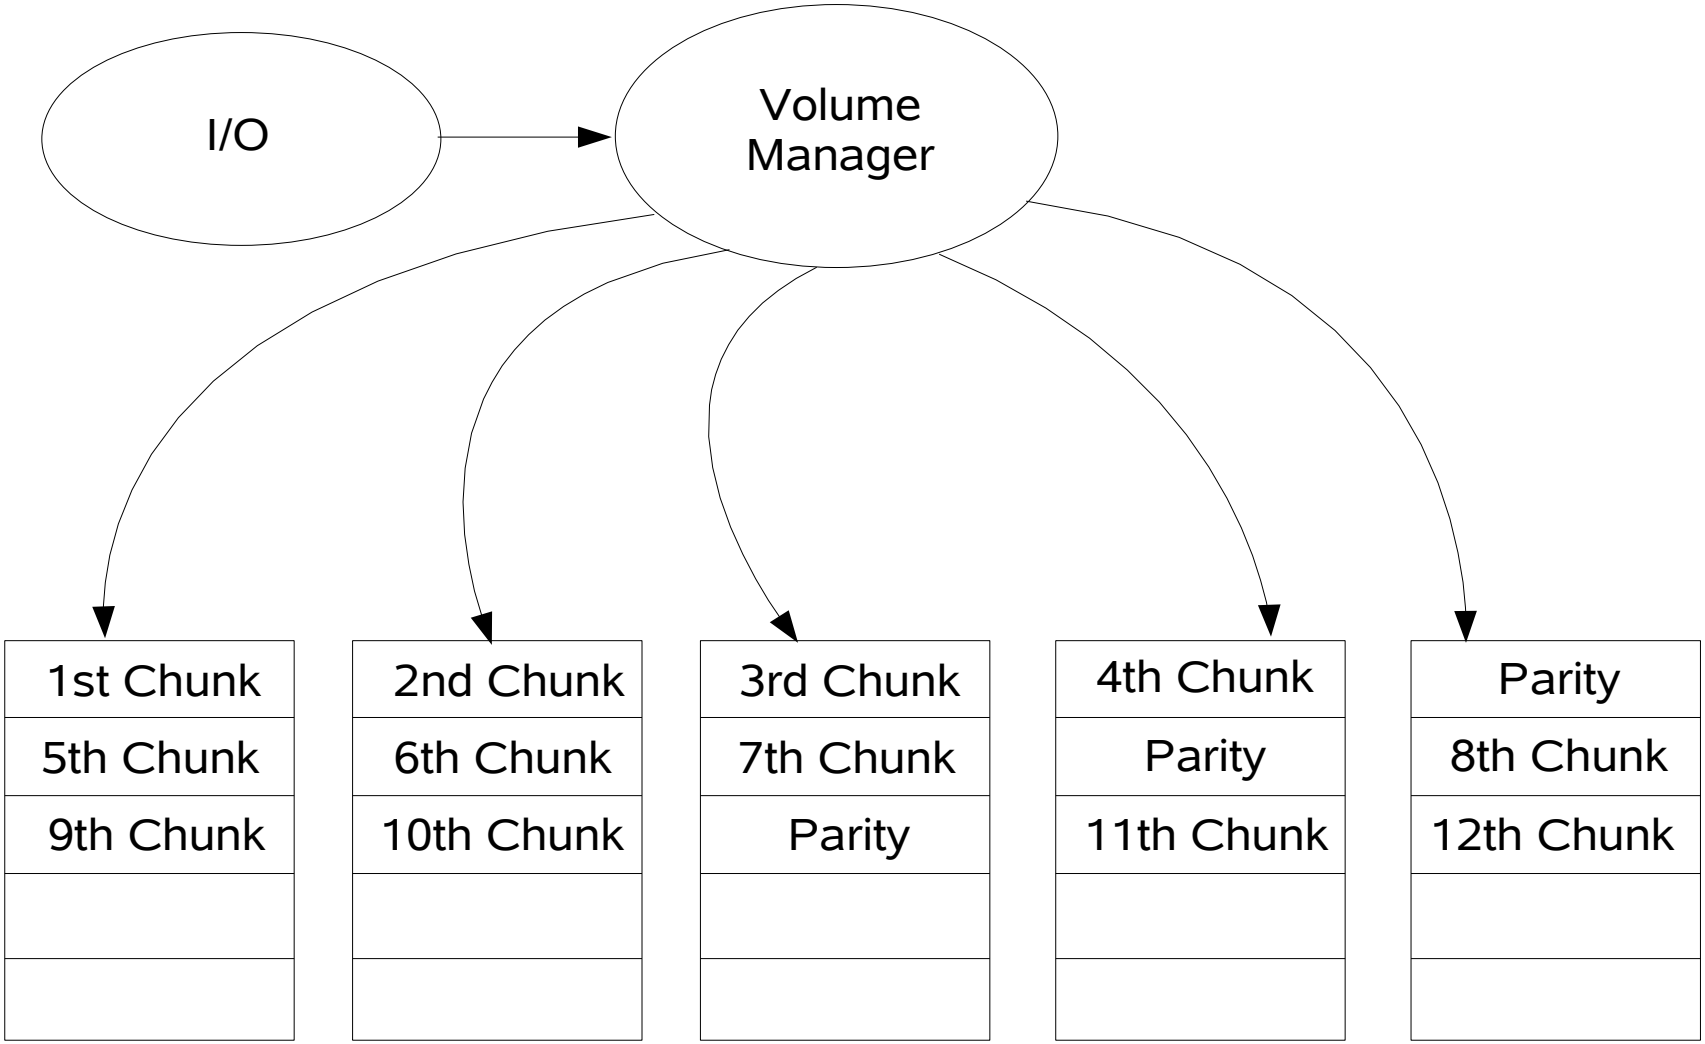

RAID Levels – RAID 6

- Striping with multiple, distributed parity data blocks
- Data protection for double disk failure
- Data written in chunks – as a stripe
  - > Must write a complete stripe, with parity each time
  - > Parity data spread through all drives to avoid hot-spots
- Read performance improved – as a stripe
- Write performance can be poor
  - > Only good for full-stripe writes
  - > Partial stripe writes must read old data and re-generate parity

RAID Levels – RAID 6

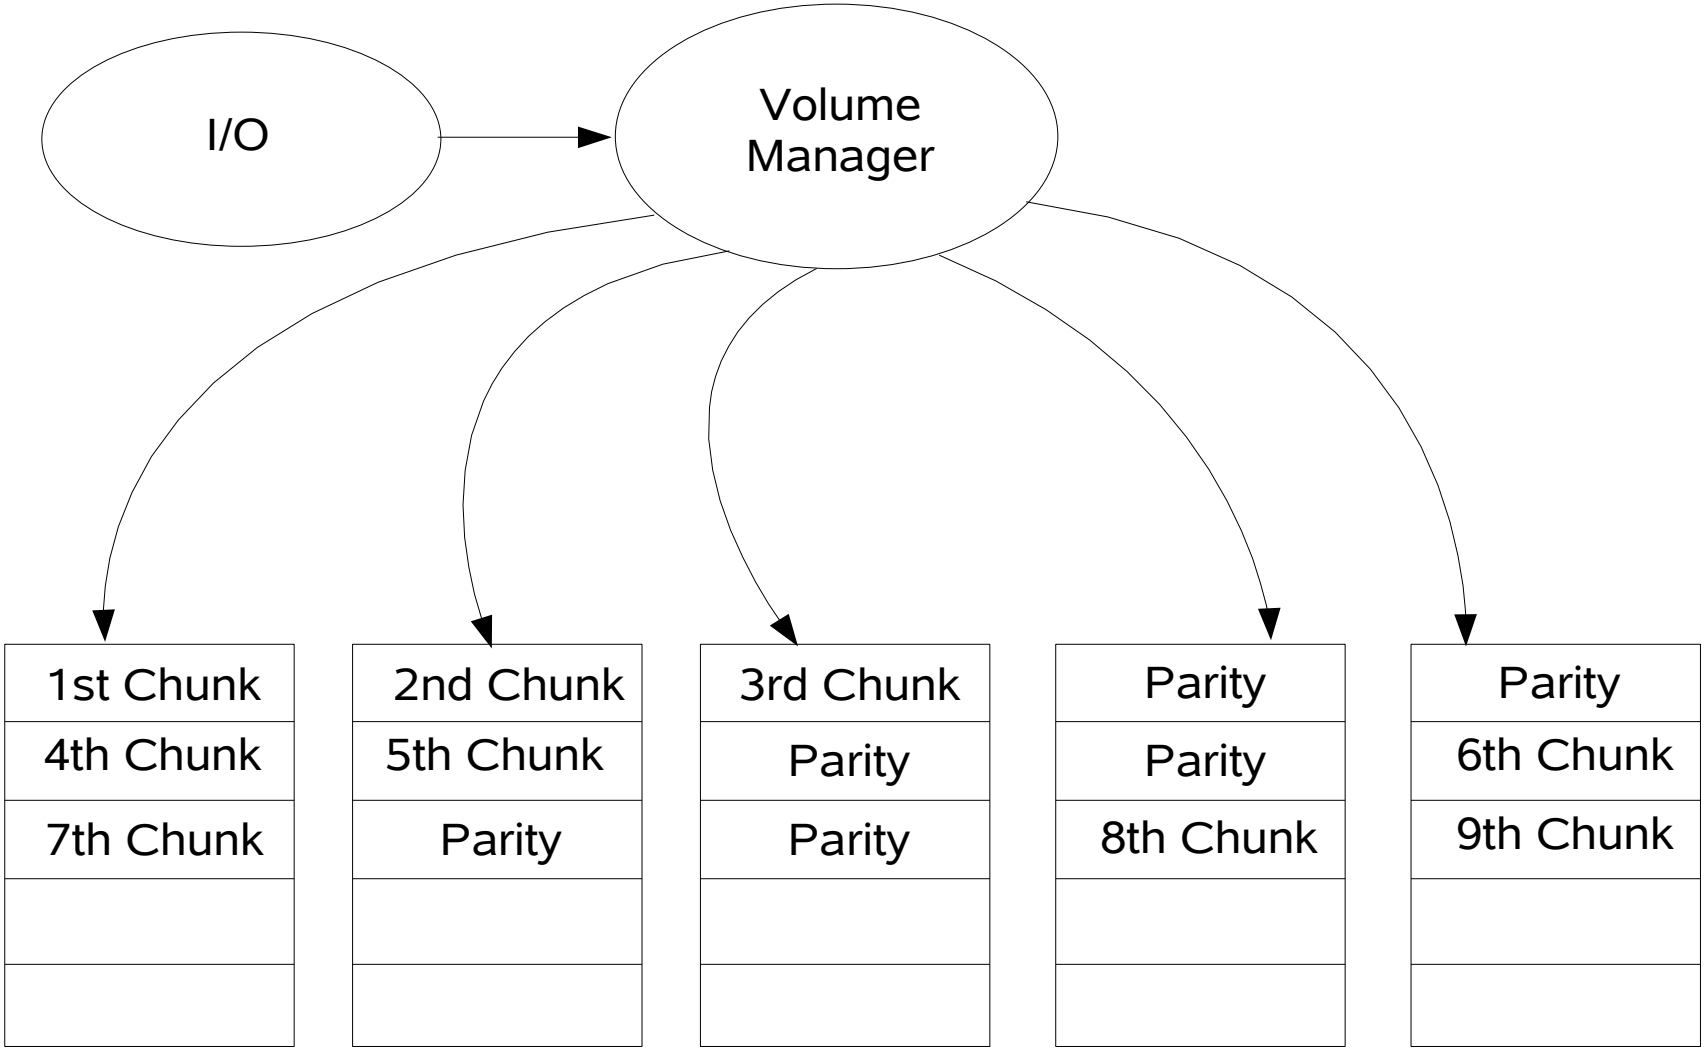

SVM RAID Overview