


AdaptiveZ: intuitive idea

- Divide the address space in zones
 - Created dynamically each time new disks are added
 - Each zone has its own heterogeneous striping "policy"
- When new disks are added, only one zone is restriped




Initial setup

- Let's assume a first array
 - 3 disks 1Gbytes each
 - Striping units \rightarrow 128KB
 - 8192 per disk
 - 24576 in total
 - Initially we have one zone




Adding 2 disks: naïve way

- Adding 2 disks
 - 2 disks 2Gbytes each
 - Striping units \rightarrow 128KB
 - 16384 per disk
 - 32768 of new storage
 - 57344 final capacity
- Create a new zone with only two disks
 - No data movement
- Problem
 - NOT balanced
 - Only new data in new disks
 - No increase in parallelism




Adding 2 disks: load balancing

- Create two new zones
 - Zone1: old data
 - Theoretical minimum movement to balance load
 - Assume new disks should have 2 times more SU
 - $C_{old} * (1 - C_{old} / (C_{new} + C_{old}))$
 - 14043 SU
 - Zone2: new data
 - Uses all disks
 - Only for new data
- Problem
 - Old data loses parallelism



Adding 2 disks: more parallelism

- Create two new zones
 - Zone1: old data
 - Increase 15% the theoretical minimum
 - $1.15 * C_{old} * (1 - C_{old} / (C_{new} + C_{old}))$
 - 16149 SU
 - Restripe these blocks on all disks
 - Find the adequate UF for this zone to guaranty global disk UF
 - Small compared to full restripe
 - Zone2: new data



Adding 2 more disks

- Create two new zones
 - Zone1: old data
 - Size = $1.15 * C_{old} * (1 - C_{old} / (C_{new} + C_{old}))$
 - Depending on size
 - We create 2 zones
 - We merge zones
 - Restripe these blocks on all disks
 - Zone2: new data




AdaptiveZ parameters

- AdaptRaid parameters are used
 - One set per zone
- % increase on the size to restripe could be changed
 - Our evaluation shows that 15% is a good tradeoff
 - and suggest no modifications

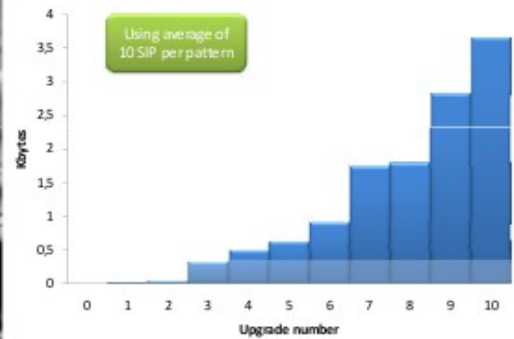
Computing the location of a block

- Formulas
 - Zone = Search in zone tree
 - Then use AdaptRaid mechanisms and metadata
- Metadata
 - Zone // Tree of AdaptRaid patterns
 - // At most 2 new zones per upgrade
 - // If upgrades every 6 months:
 - // 40 zones in 10 years
 - // 6 levels in the tree

Load balance after 10 upgrades



Metadata size after 10 upgrades



Performance results

