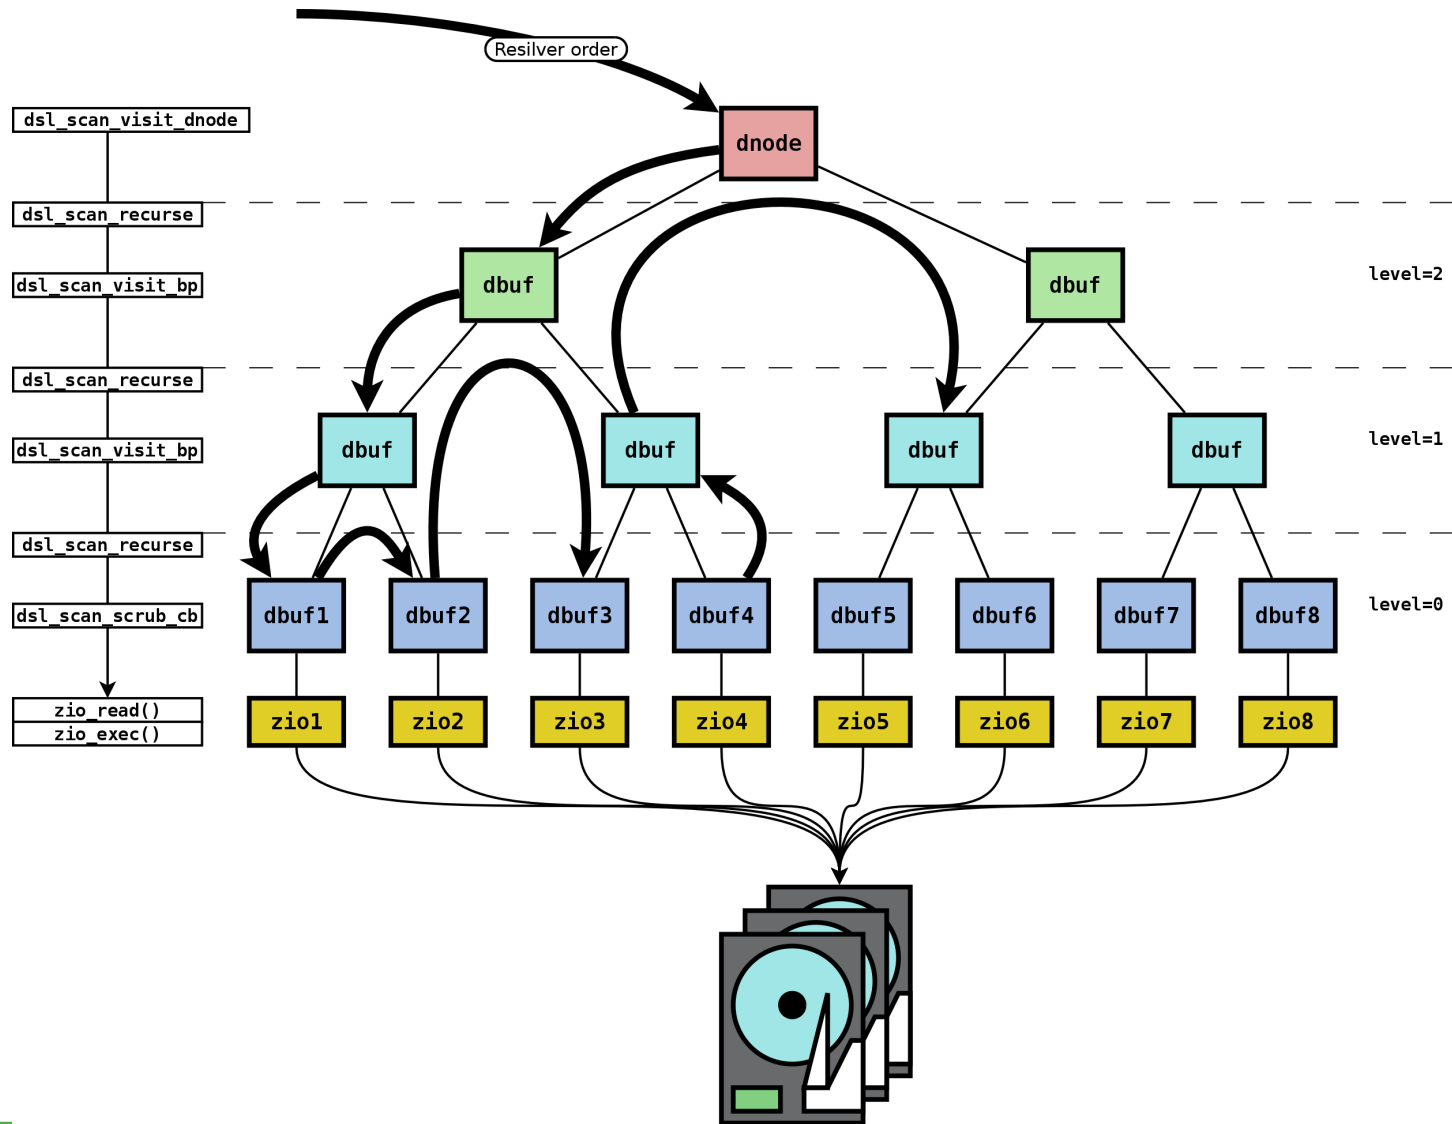


Better resilver – progress 1 yr on

SAŠO KISELKOV | kernel engineer

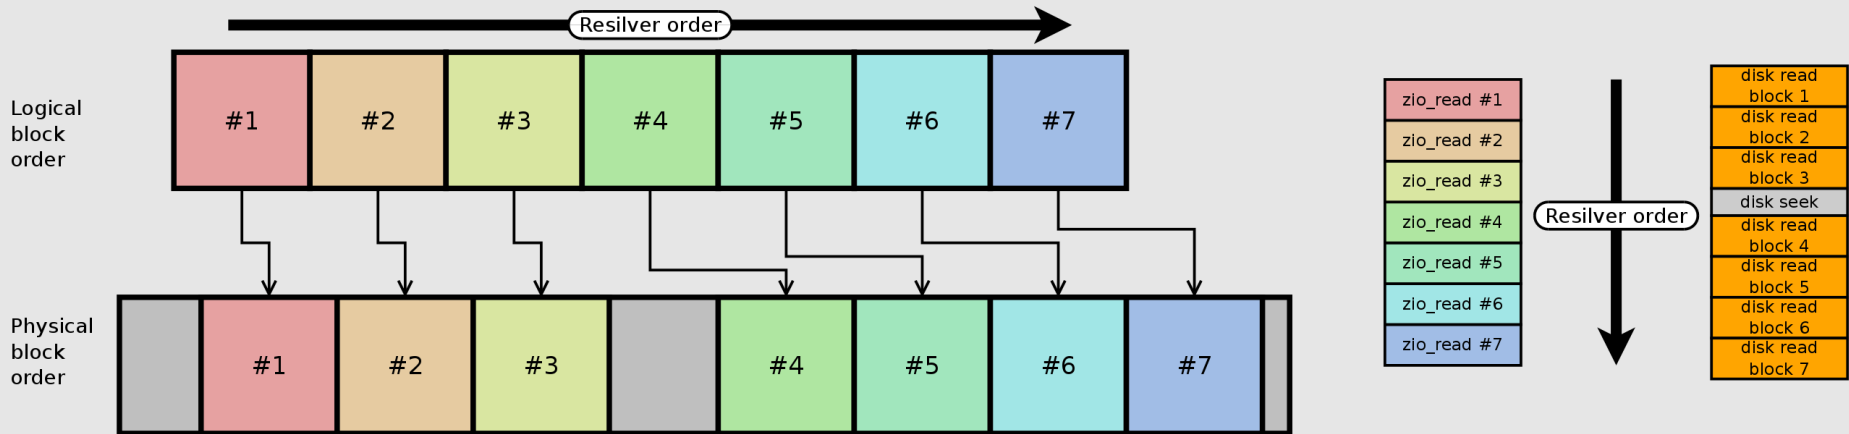
OCTOBER 2017

Better resilver

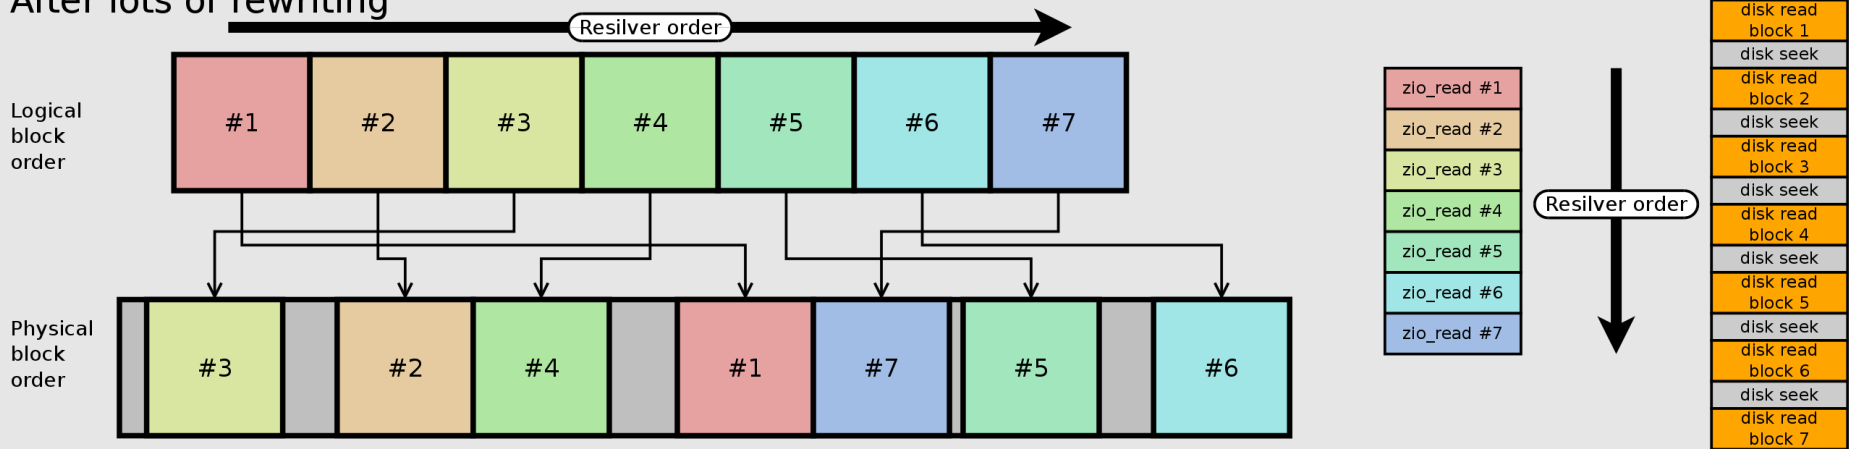


Better resilver

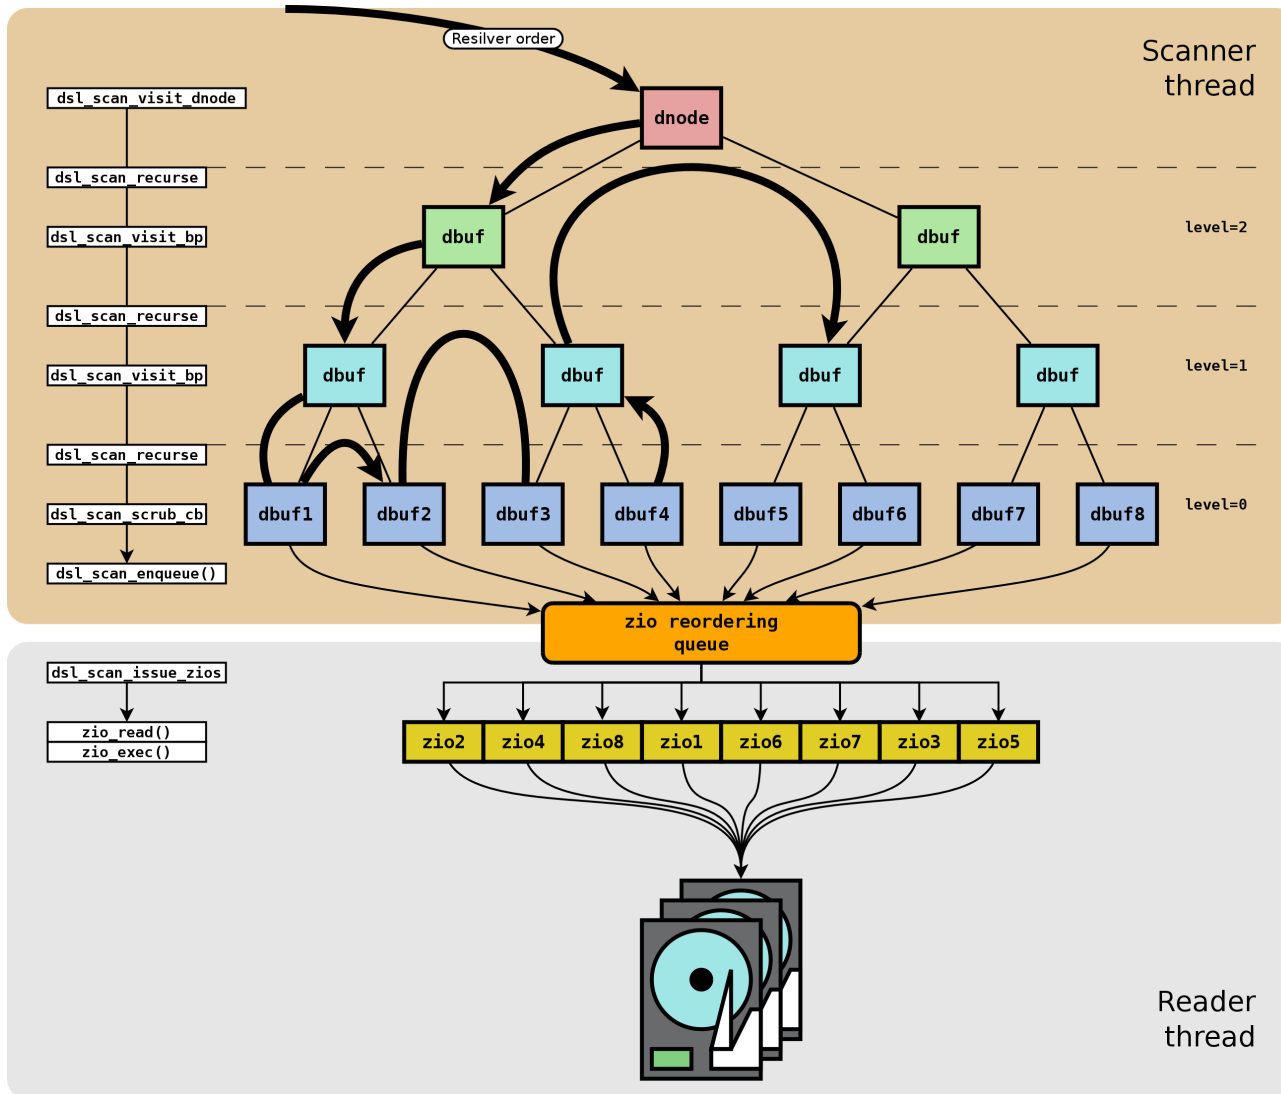
On initial writing



After lots of rewriting



Better resilver



Better resilver


Further improvements

- Shipped in February
- New default algorithm on NexentaStor 5.0.3 and newer
- No disk format change
- Involved Matt & George in fine-tuning the design
- Tom Caputi from Datto improved resilver prefetcher (ZoL)

Better resilver

Performance

Promises, promises...



Better resilver

Performance

raw data size: 21.1 TiB (12% full)

recordsize: 128k

pool geometry: 12 x 8-drive raidz2, 2x slog (200GB SAS SSD)

main pool devices: SEAGATE ST2000NM0023 1.82TB

ALGO	TEST	TIME (hh:mm)	RATE (MB/s)	DELTA	CACHESZ (MB)	GAPSZ (MB)
improved	scrub	10h44m	573	3.1x	26121	2
stock	scrub	32h45m	188	----	----	----

Better resilver

Performance

raw data size: 26.2 TiB (pool 36% full)

recordsize: 16k

pool geometry: 2 x 5-drive raidz2


devices: 10x HGST 8TB 7k2 3.5" SAS (HUH728080AL5204)

KERNEL	TEST	TIME (hh:mm)	RATE (MB/s)	DELTA	CACHESZ (MB)	GAPSZ (MB)
stock	scrub	49h16m	154.9	----	----	----
improved	scrub	14h40m	520	3.4x	1309	2
improved	scrub	13h49m	552	3.6x	6548	2

Better resilver

Performance

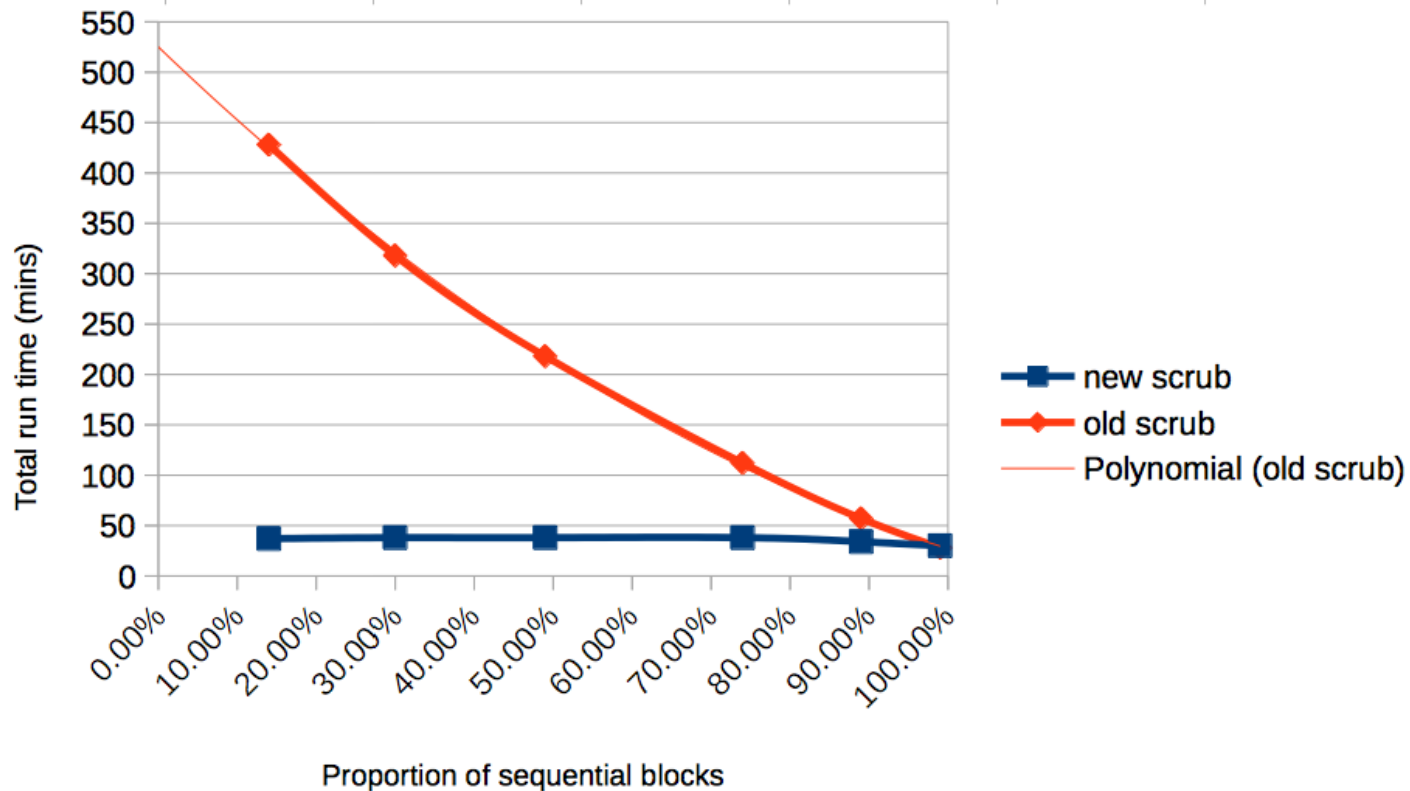
Customer boxes (X speedup from old algo)



Better resilver

Performance


- Effects of randomization



Better resilver

Performance

- Effects of randomization



www.nexenta.com