



STORAGE DEVELOPER CONFERENCE

SNIA ■ SANTA CLARA, 2016

Iometer for Zoned Block Devices

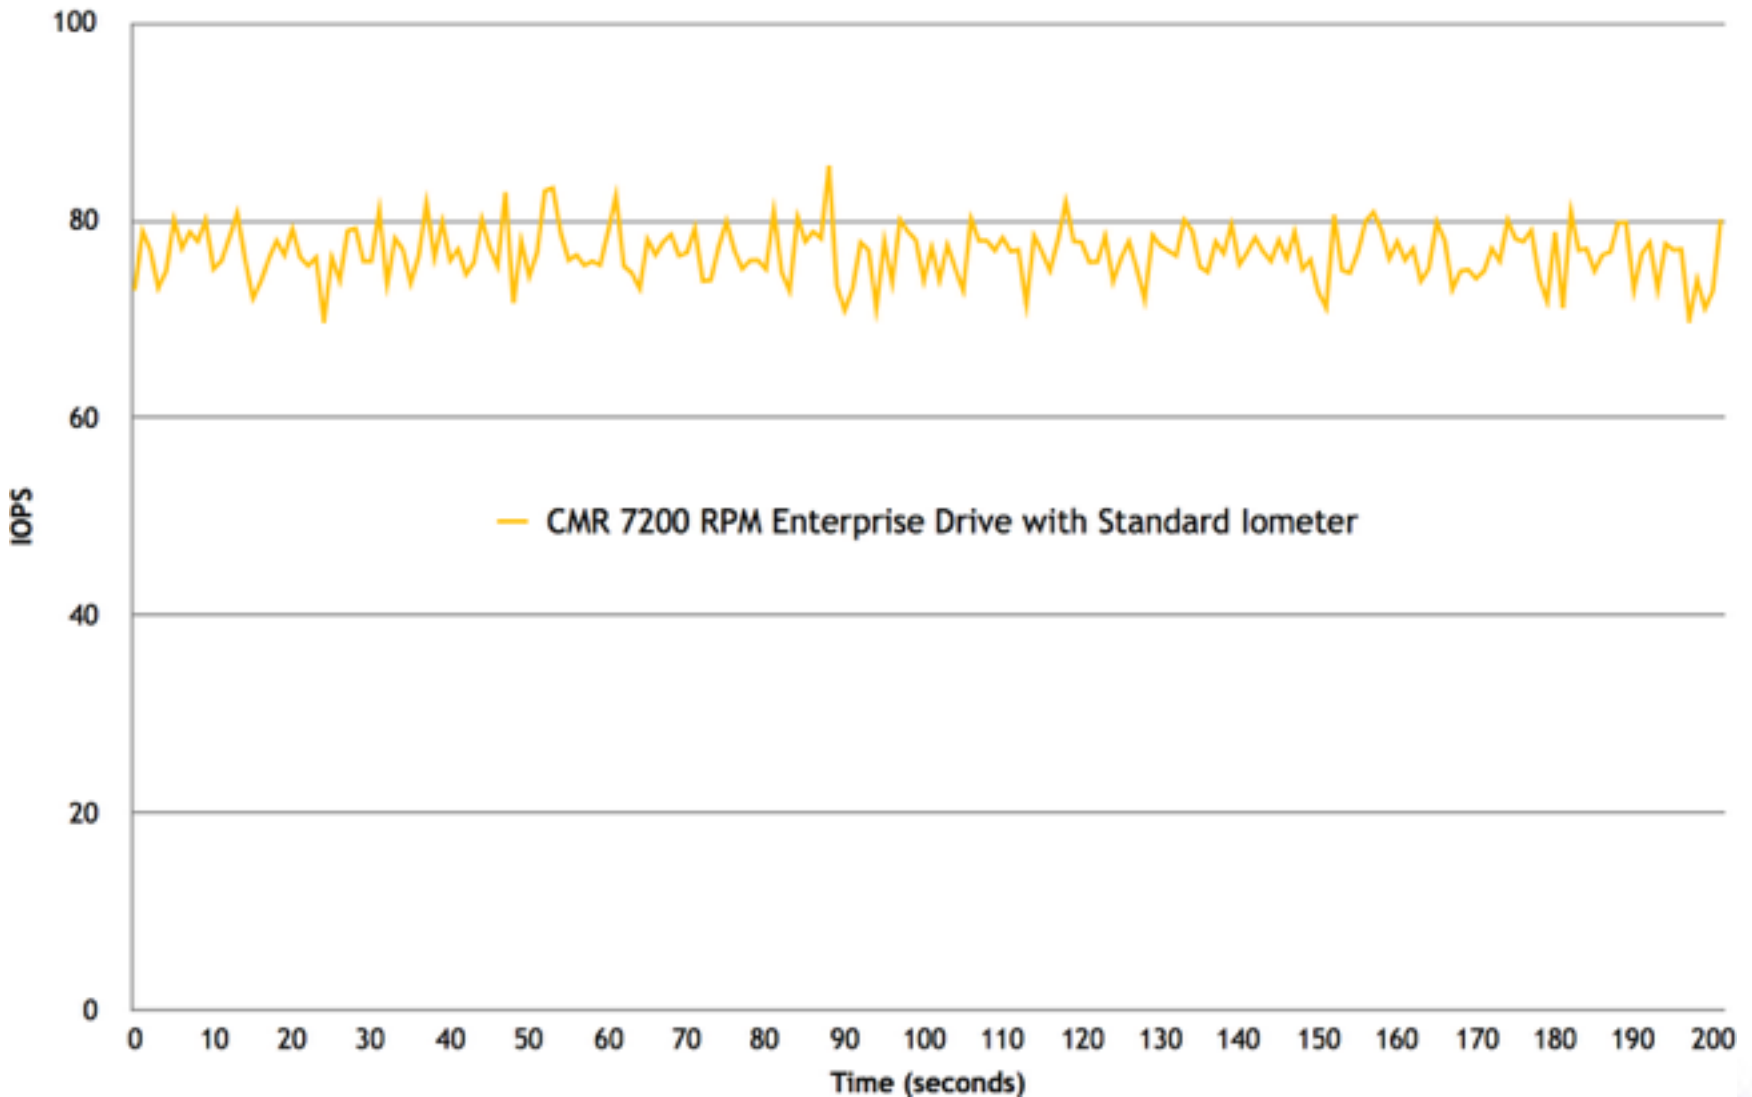
Microbenchmarking SMR HDDs

Muhammad Ahmad


Gregory Frazier

Seagate Technology

CMR 8K Block Random Writes - WCD QD=1



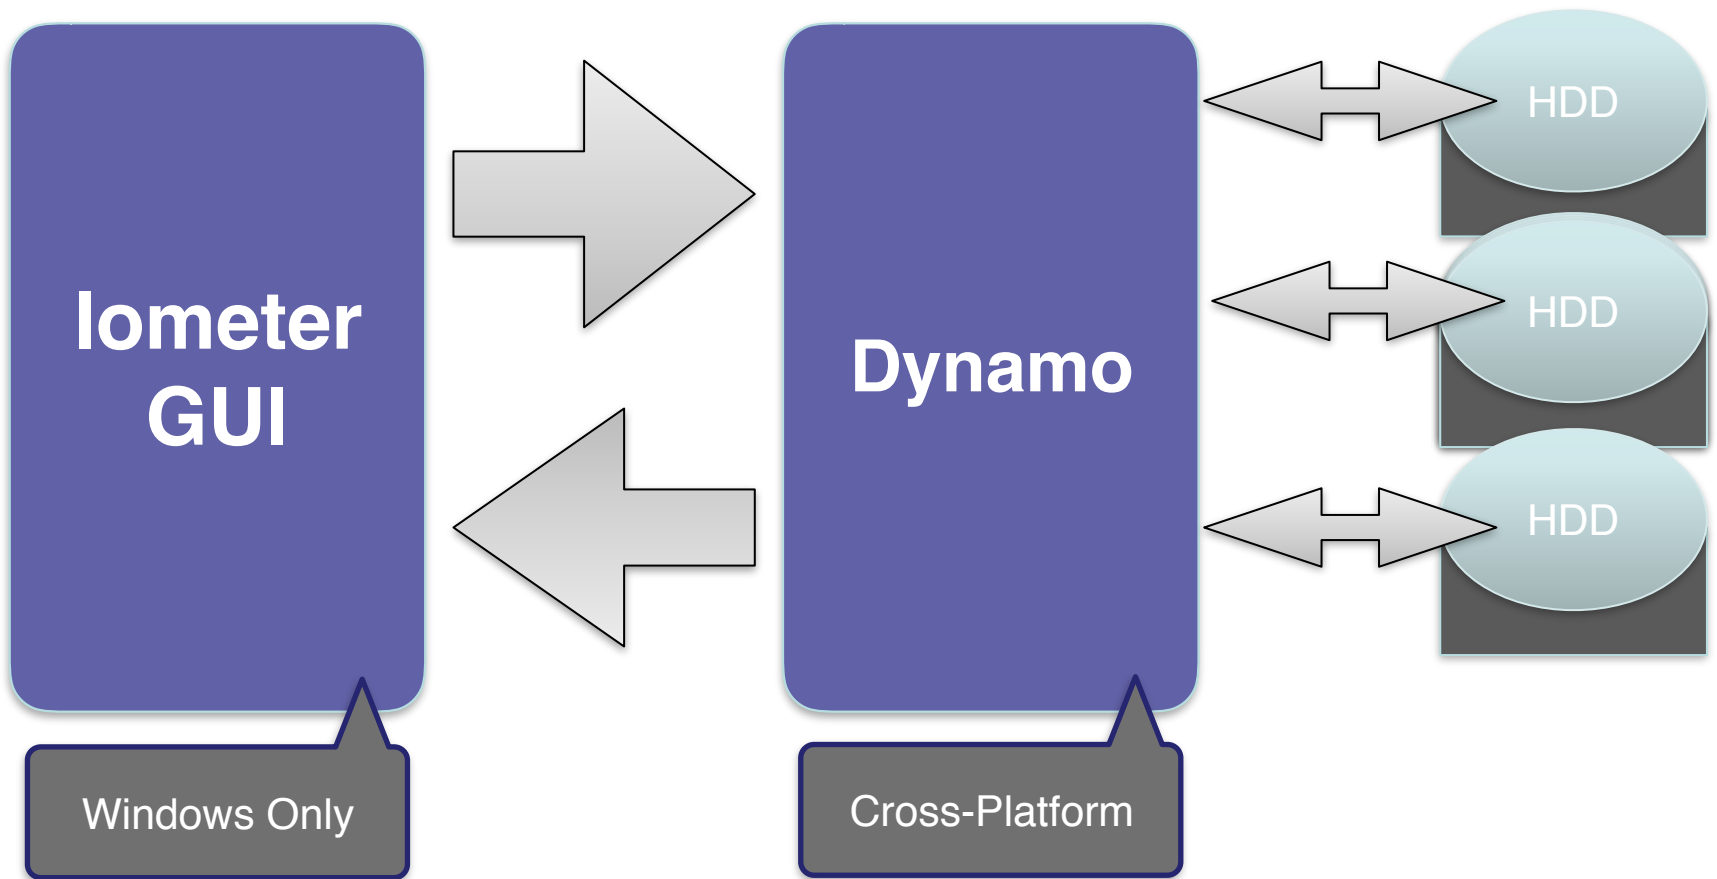
ZBD 8K Block Random Writes - WCD QD=1



What's wrong with Iometer?

- ❑ Industry Standard Benchmark tool with no support/intelligence for Zoned Block Devices (ZBD)
- ❑ Customers need a way to microbenchmark ZBDs to validate primitive performance capabilities that are a function of seek speed, spin speed and media data rate

Iometer Big Picture



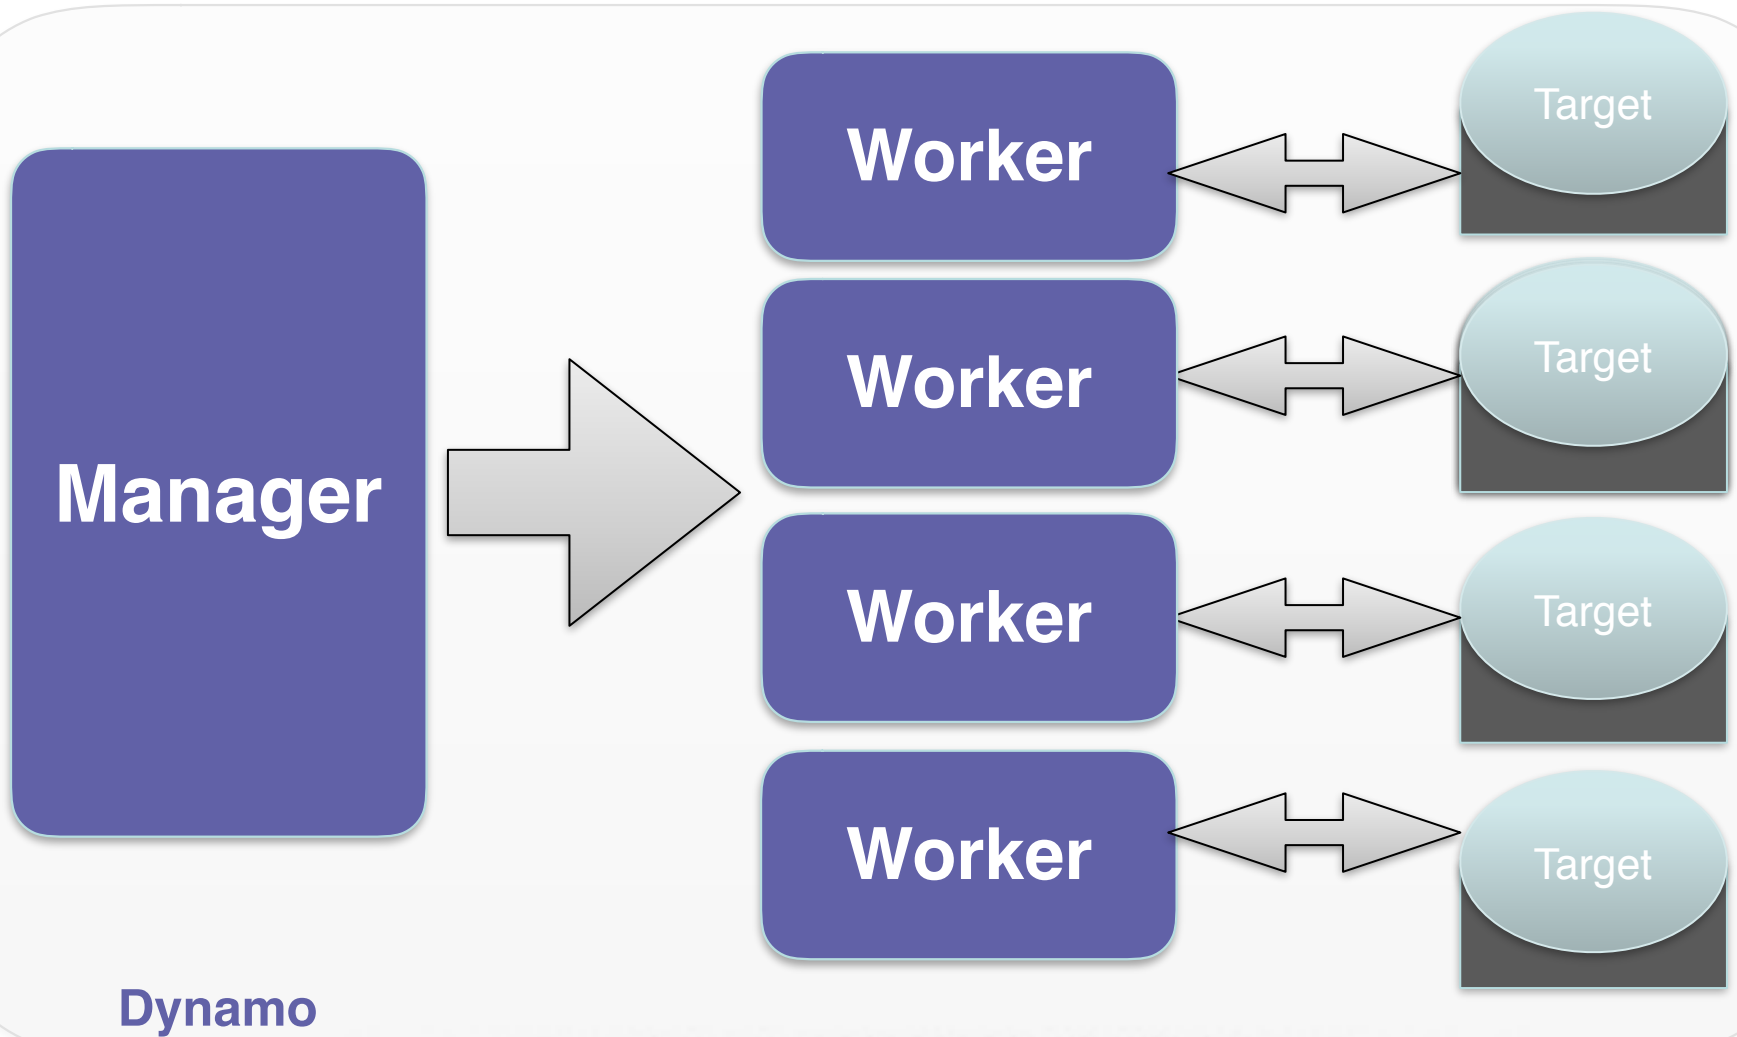
Iometer Design Challenges

- ❑ No support for issuing individual ATA/SCSI commands to enumerated raw devices
- ❑ Device Identification (e.g. Model Number, size) are determined through OS APIs

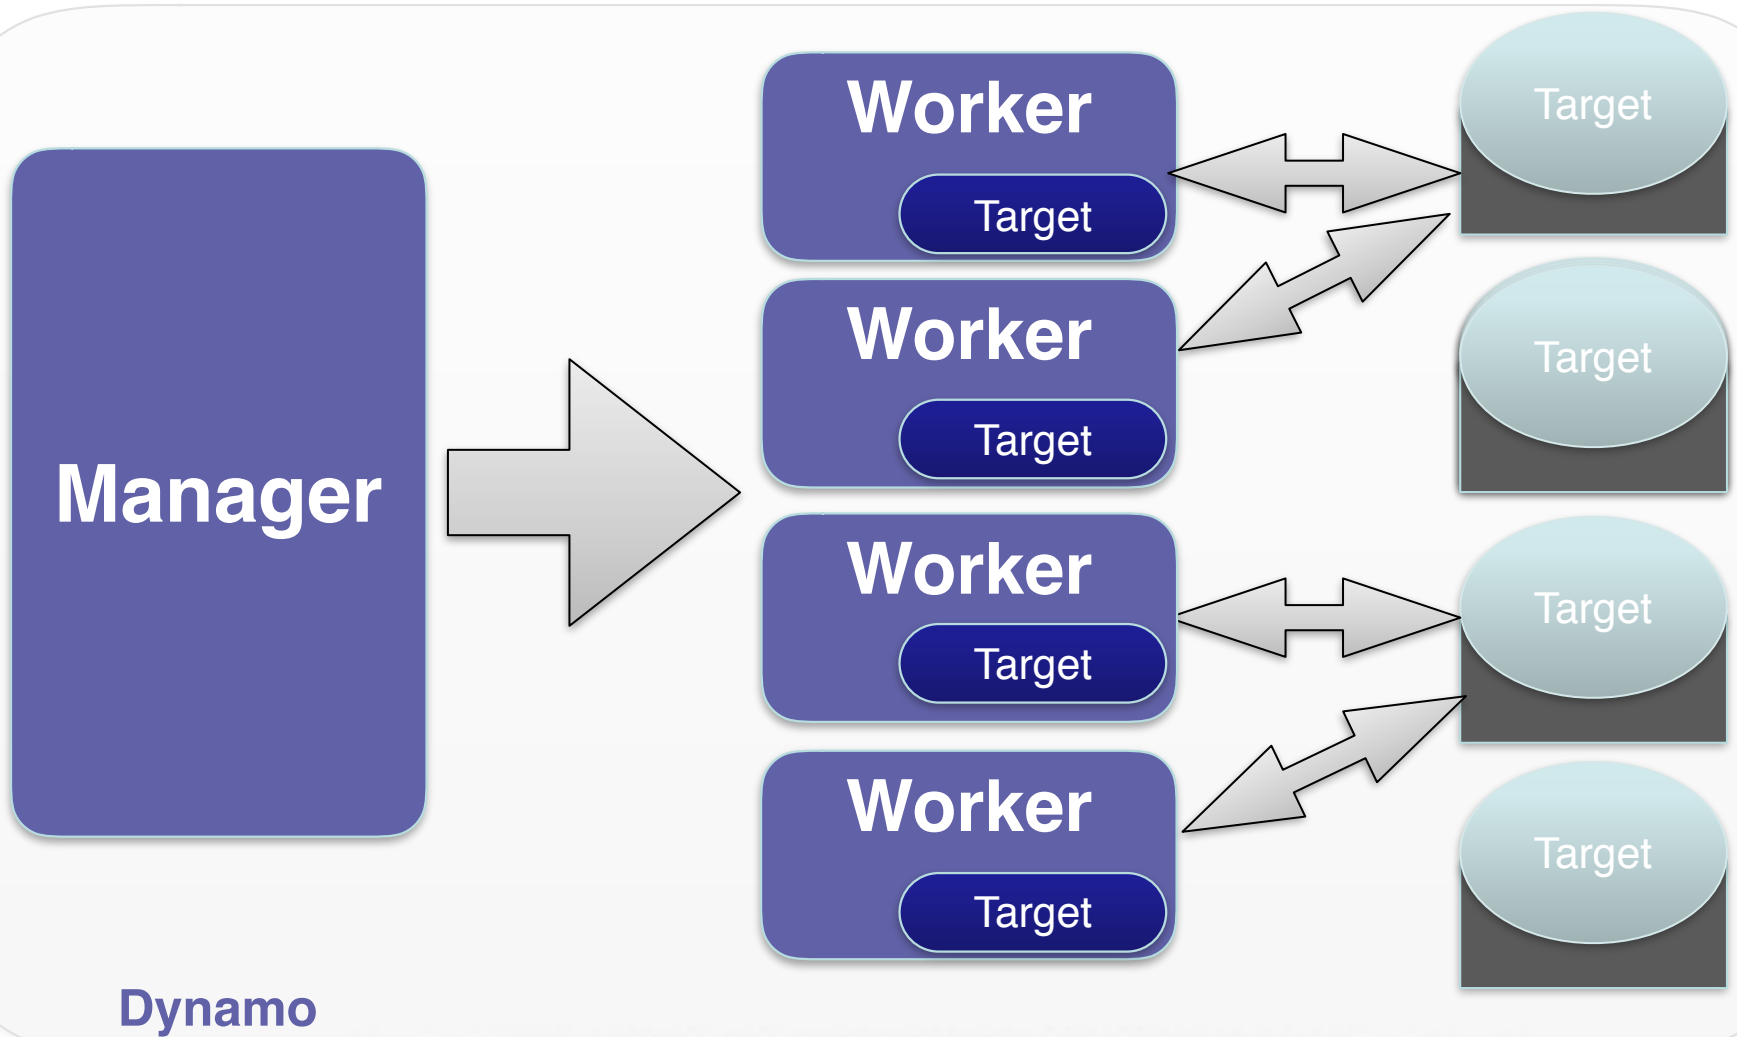
Iometer Design Challenges

- ❑ IOs (Read & Writes) are issued to volumes & raw devices based on an offset (no concept of Zones or even LBAs)
- ❑ IO Manager has a hands off approach towards worker threads who are unaware of each other
- ❑ Comparing CMR to ZBDs in a ZBD-friendly workload

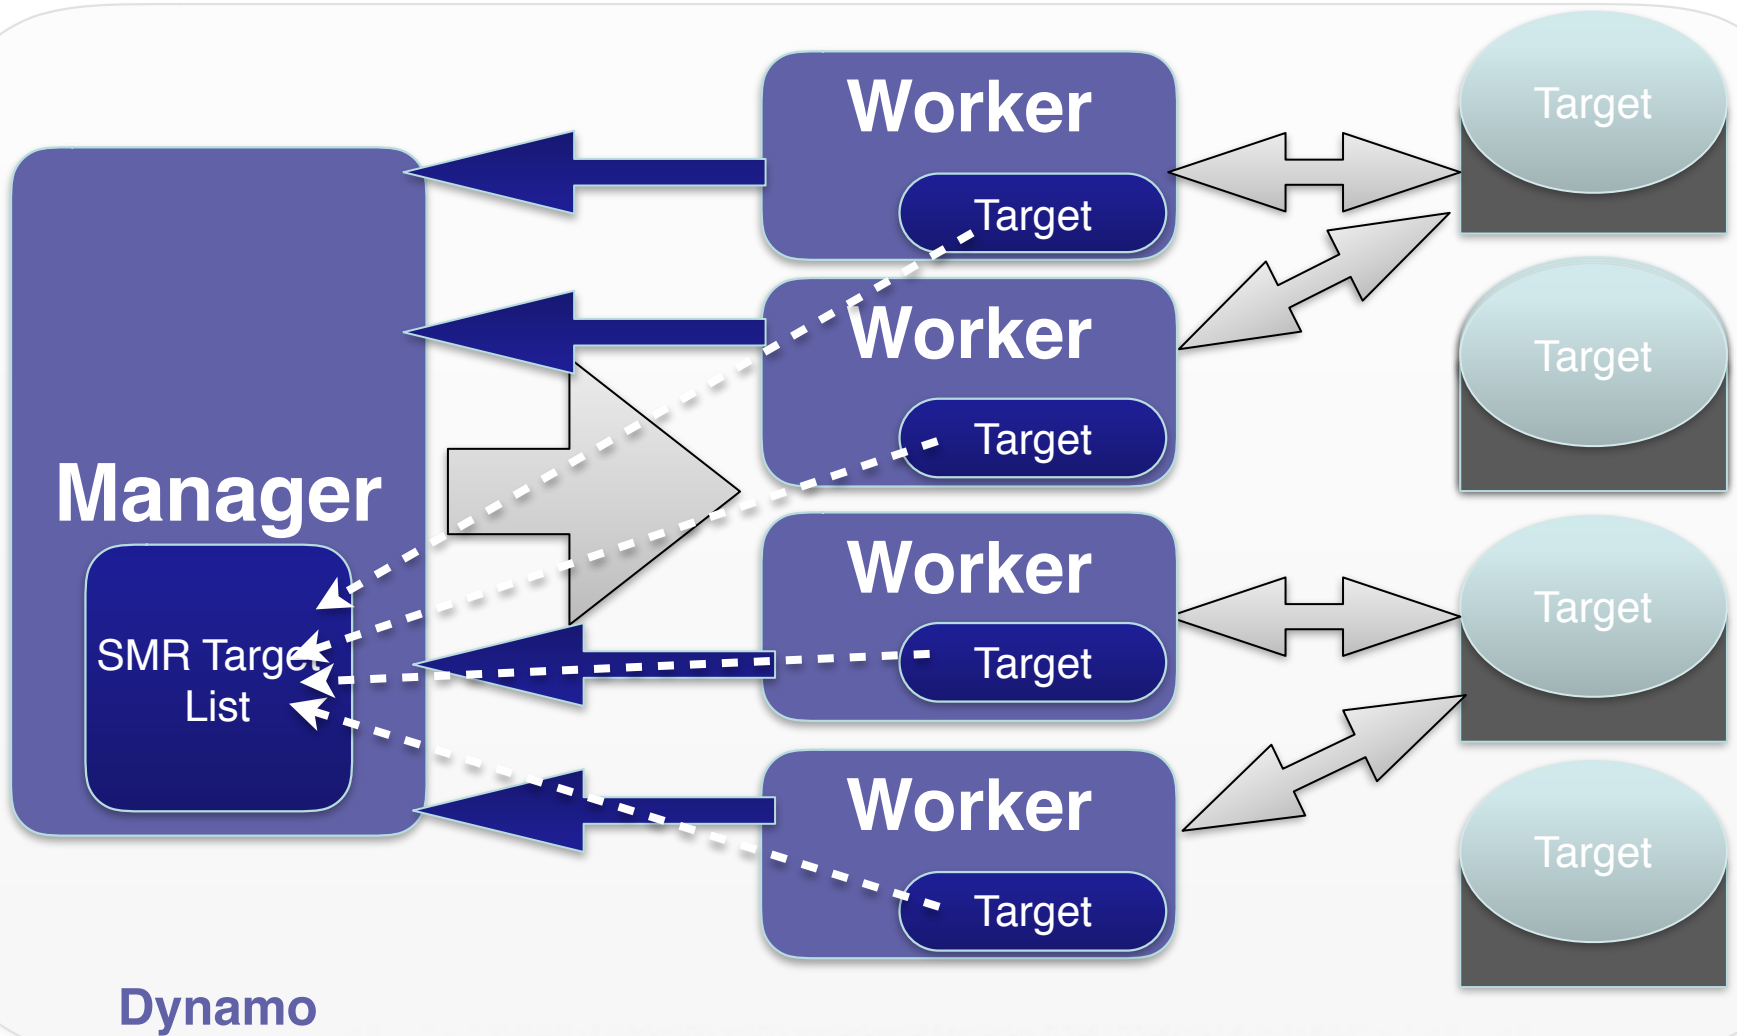
Dynamo Big Picture



Dynamo Challenge



Dynamo Changes



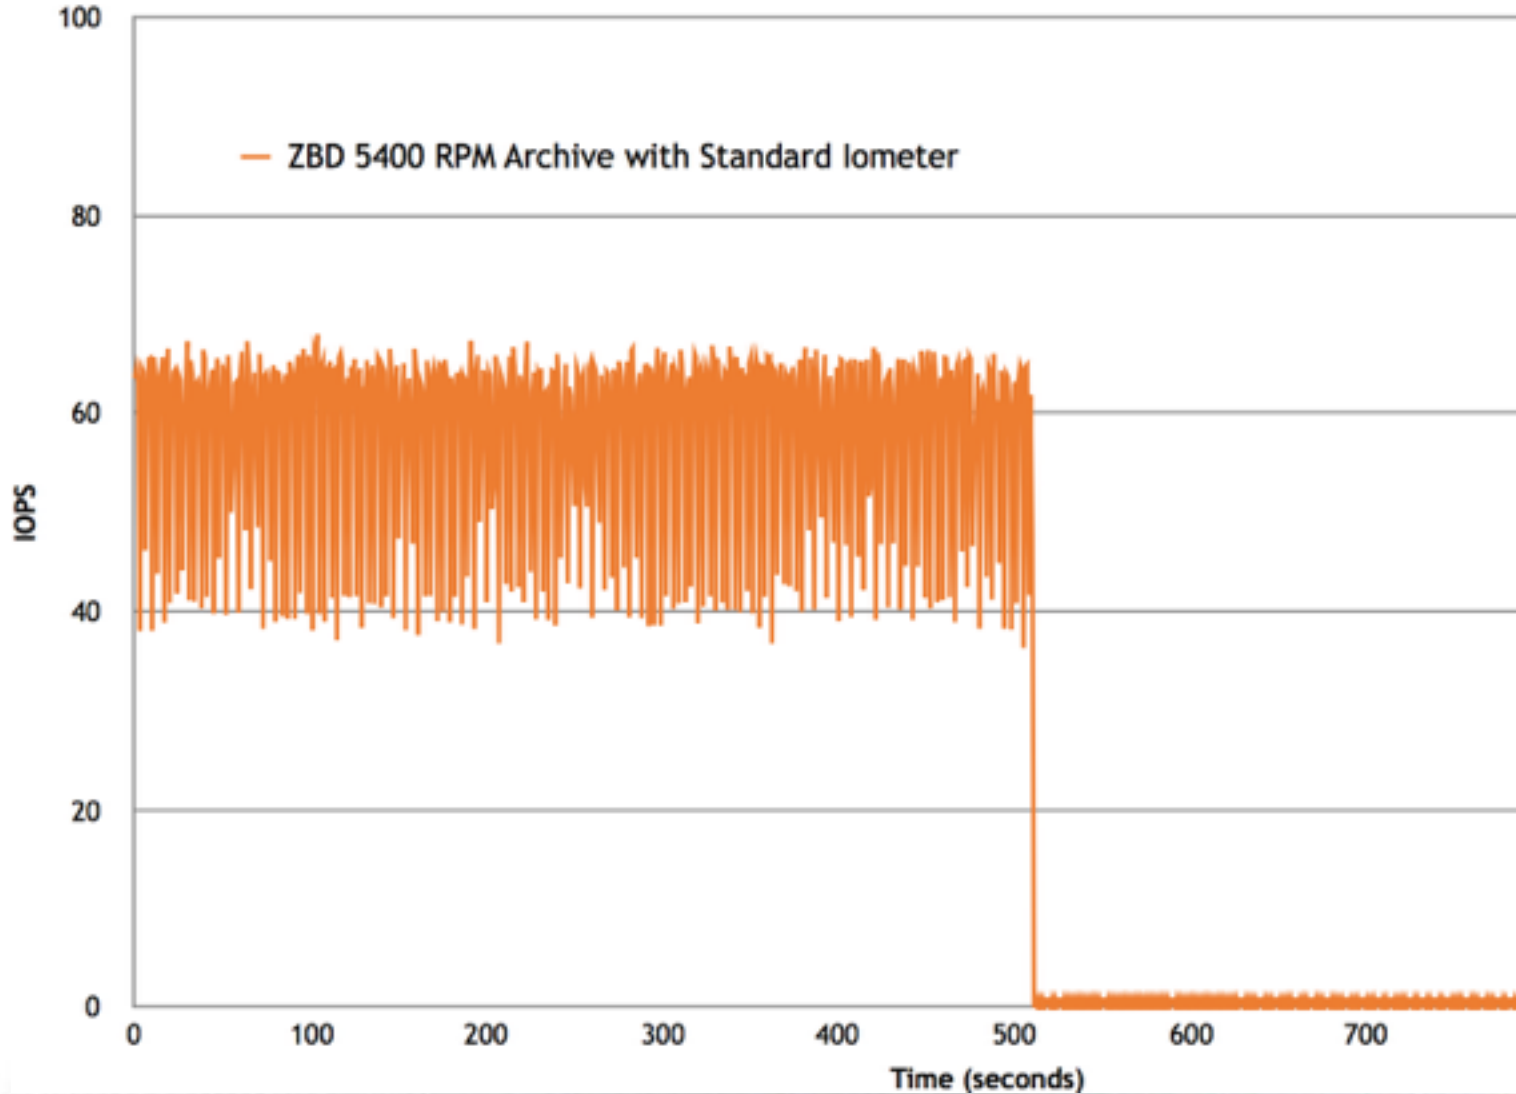
SMR Iometer Change List Highlights

- ❑ New SMRTarget Class
 - ❑ ZBD Type
 - ❑ Number of Zones
 - ❑ Max Open Zones
 - ❑ Open Zone List (index)
 - ❑ Zone List (Descriptor)
 - ❑ Mutex (Threadsafe)
 - ❑ Function that supports these data structures


SMR Iometer Change List Highlights

- ❑ TargetDisk Class
 - ❑ Added ZBD identifiers
 - ❑ Command Passthrough transport (ZAC/Windows)
 - ❑ Modified **Seek** / Write / Read
- ❑ Manager Class
 - ❑ m_SMRTgts <vector>
 - ❑ force_ZBD_Flag
 - ❑ max_open_zones

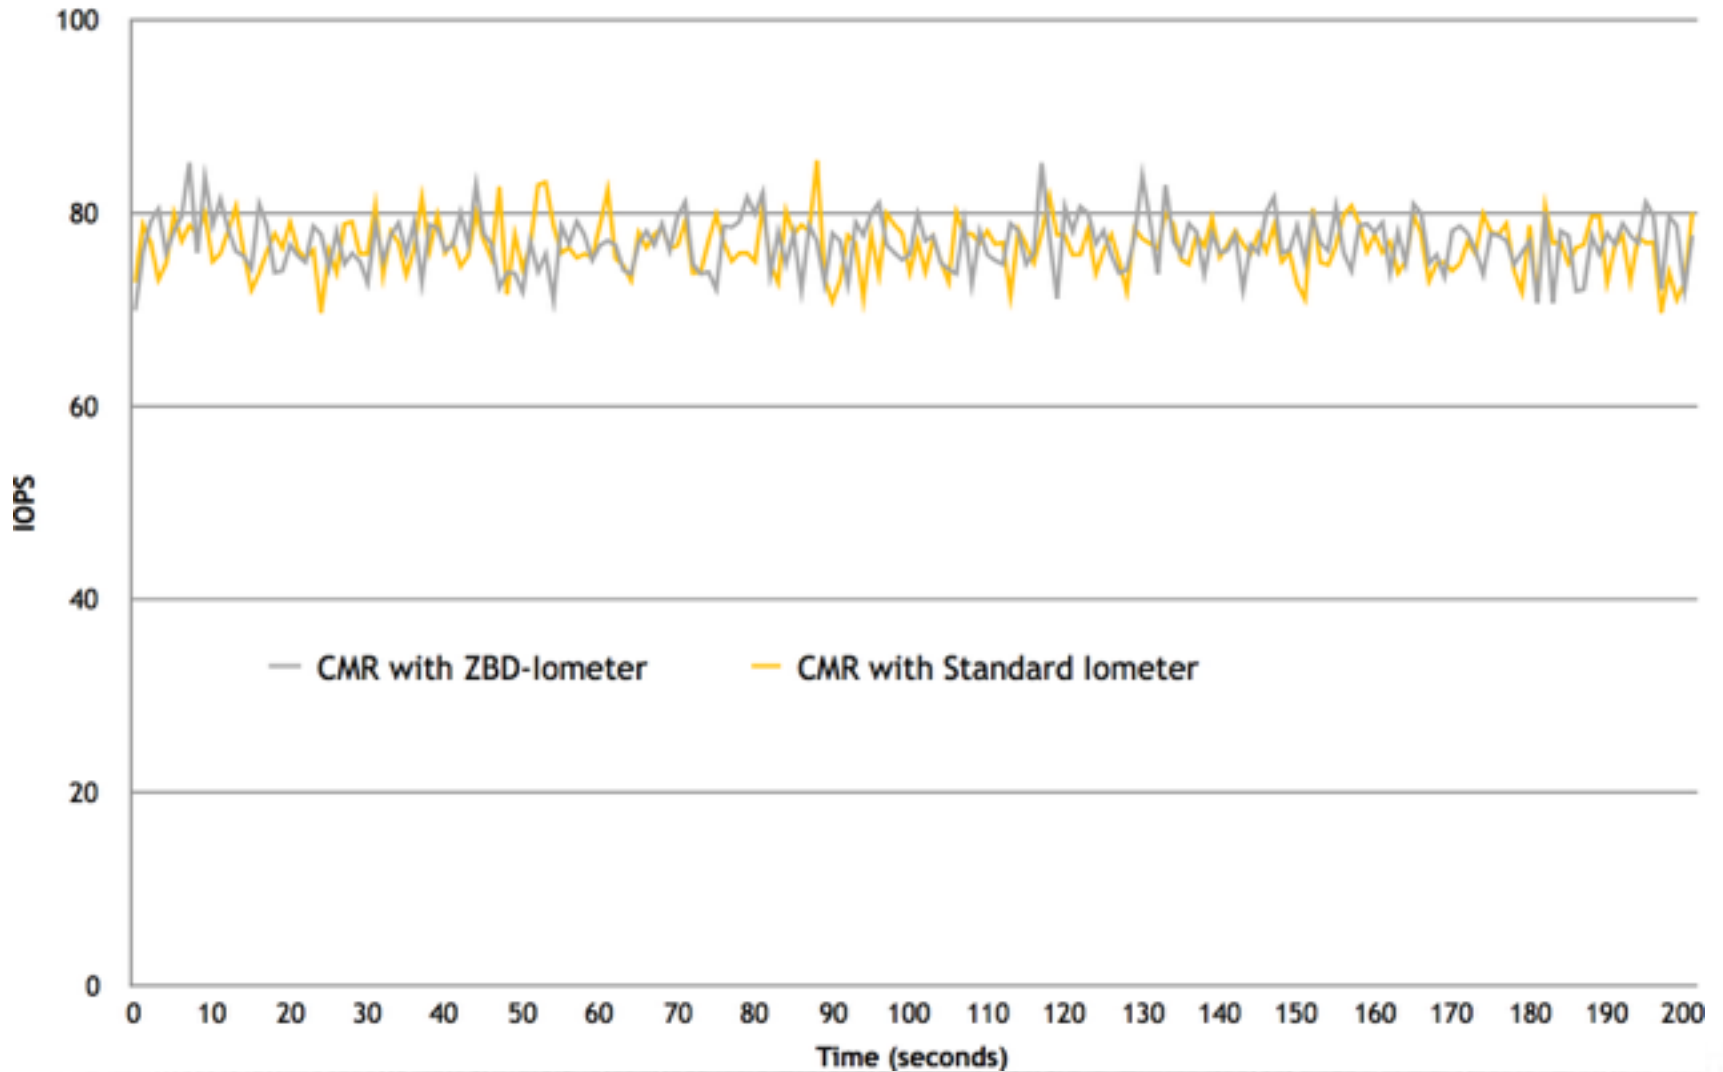
ZBD 8K Block Random Writes - (Recap)



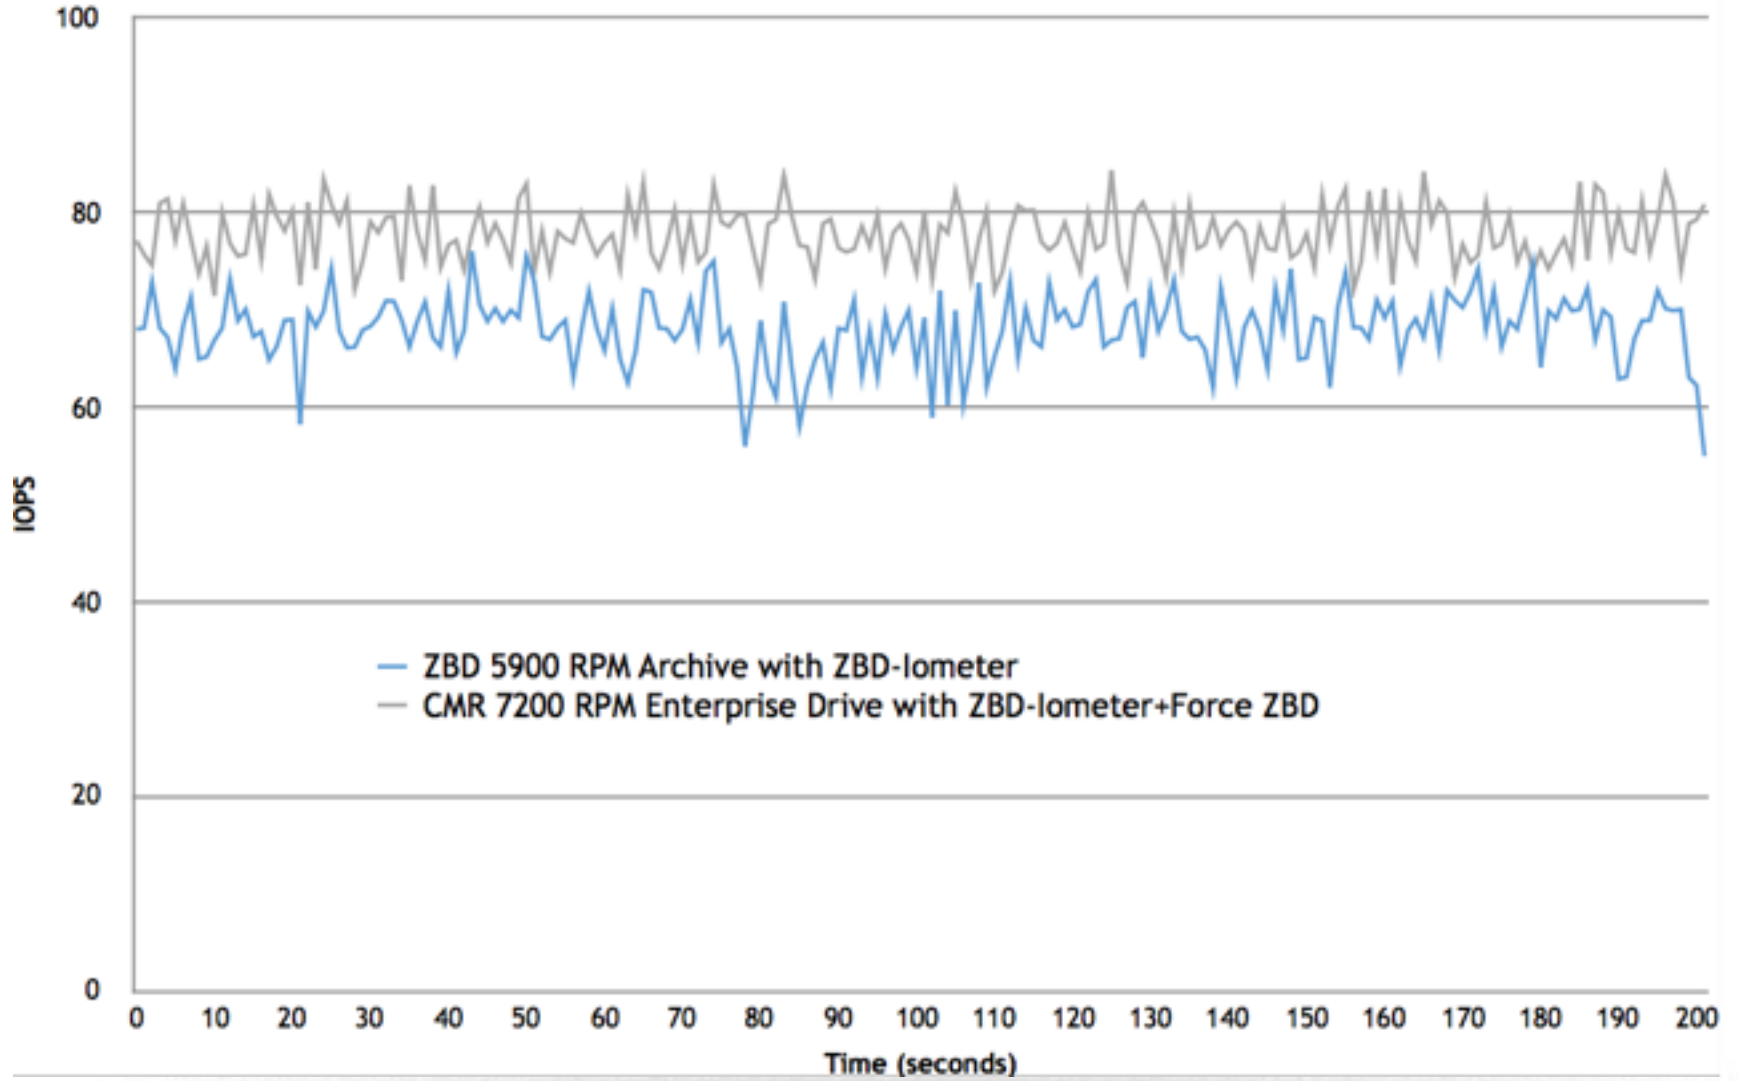
ZBD 8K Zone Random Writes - ZBD-Iometer



CMR Drive Iometer vs ZBD-Iometer



CMR vs ZBD - 8K Zone Random Writes



SMR Iometer ToDo:

- ❑ Adding Dynamo support for other OSes (Linux)
- ❑ Updating ZAC/ZBC commands
- ❑ Logic support for Host Managed Devices
- ❑ Testing Error Paths
- ❑ Visual Studio 2013/15

References

- ❑ Repository

<https://github.com/Seagate>

- ❑ Original Iometer Project

<http://iometer.org>