Storage Performance Benchmarking: Introduction and Fundamentals

Ken Cantrell / NetApp

Mark Rogov / EMC

July 30, 2015

About The Speakers

Ken Cantrell

NetApp

Manager Perf Engineering

@kencantrelljr

Mark Rogov

EMC

Systems Engineer

@rogovmark

Dr. J Metz

Cisco

R&D Engineer

@drjmetz

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Storage Performance Benchmarking

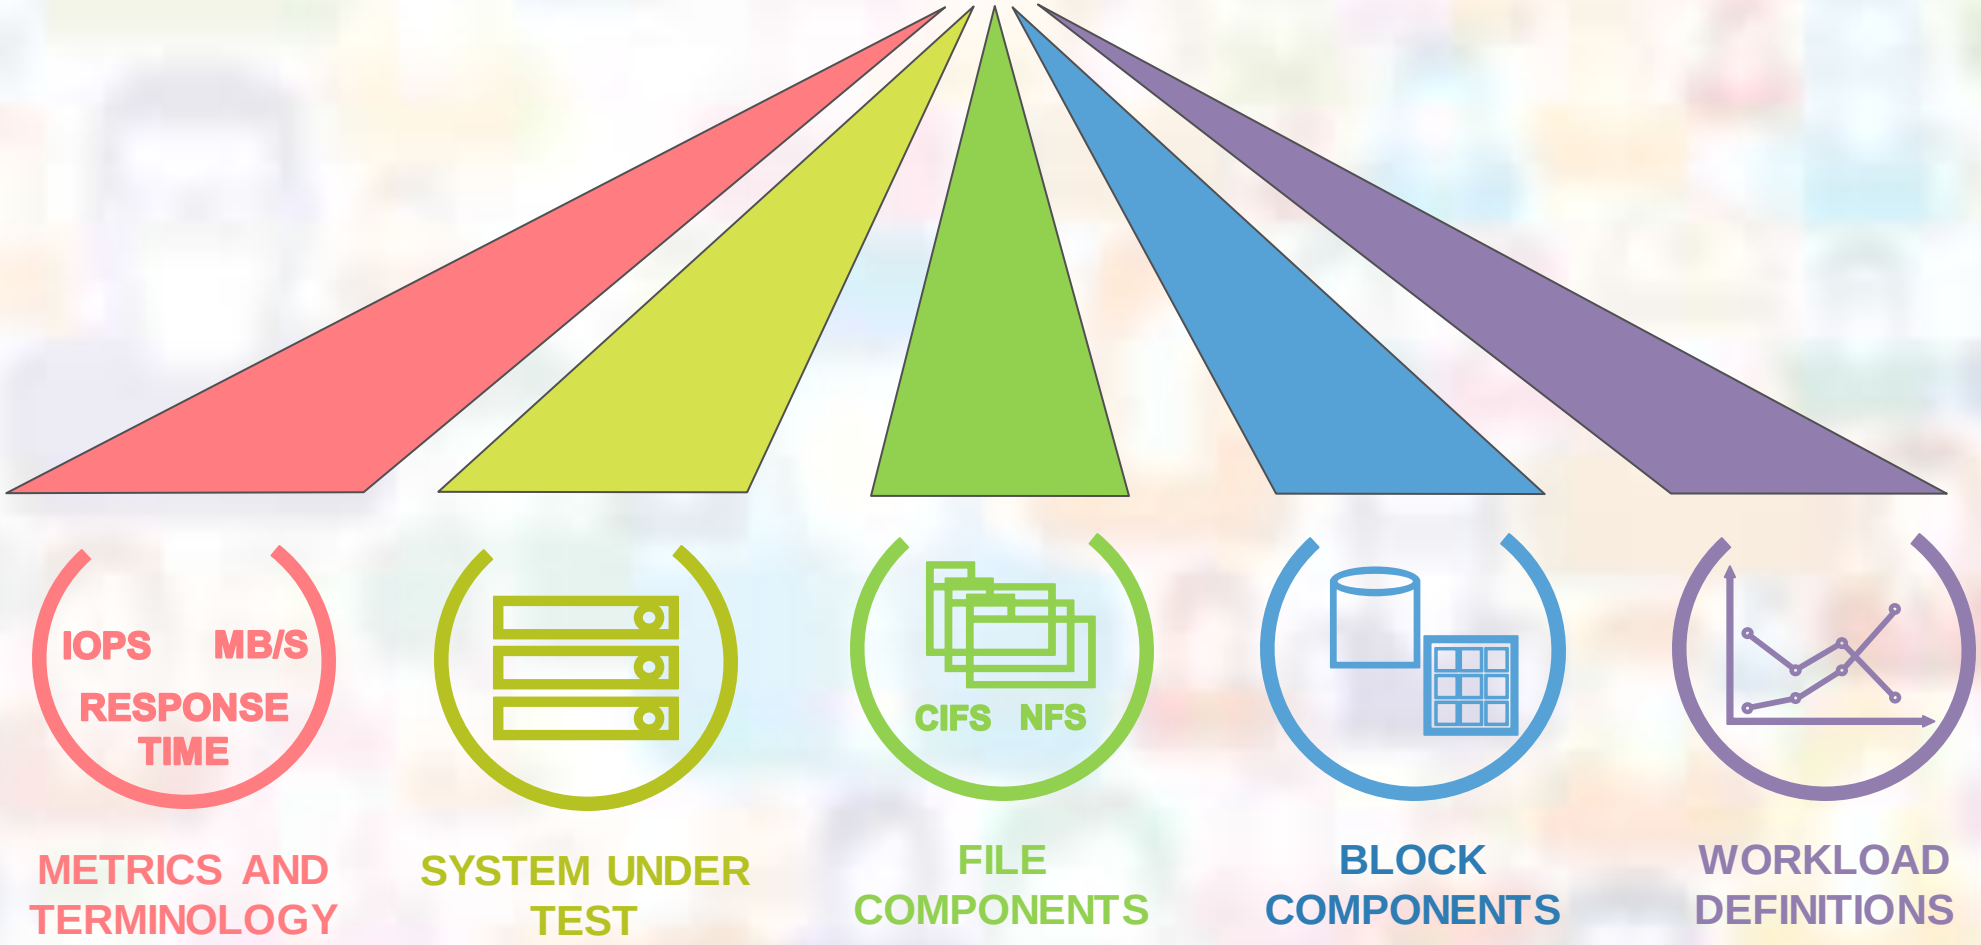

Storage Performance Benchmarking

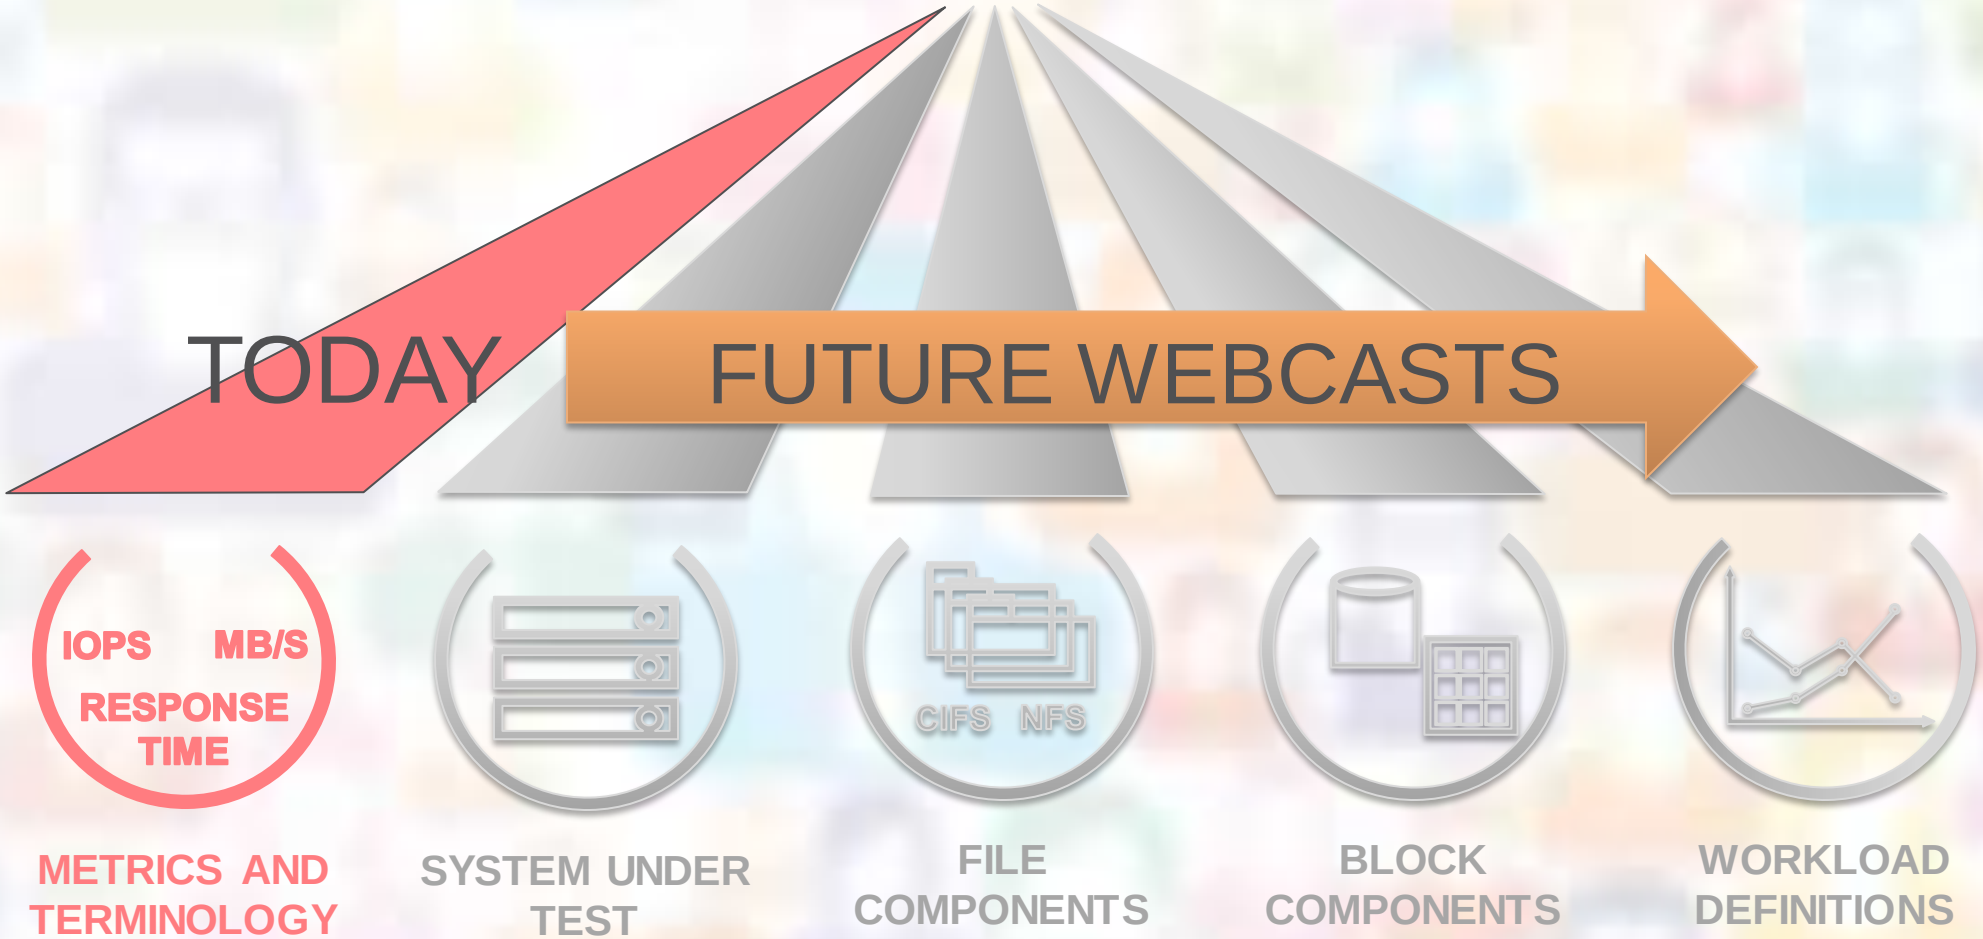

INTRO BLOCK FILE GRAPH

Three Great Buckets Of Testing

BAKE-OFFS

SINGLE OBSCURE METRIC

UN*X DD
IOMETER
IOZONE
VDBENCH

BENCHMARKS

BALANCED PROFESSIONAL APPROACH

SPEC SFS
SPC1

LOGIN VSI
VMMARK

BUSINESS

RUN PRODUCTION APPLICATIONS

DATABASE
COMPUTE FARM
VDI

Block / SAN Terminology

Block / SAN Terminology

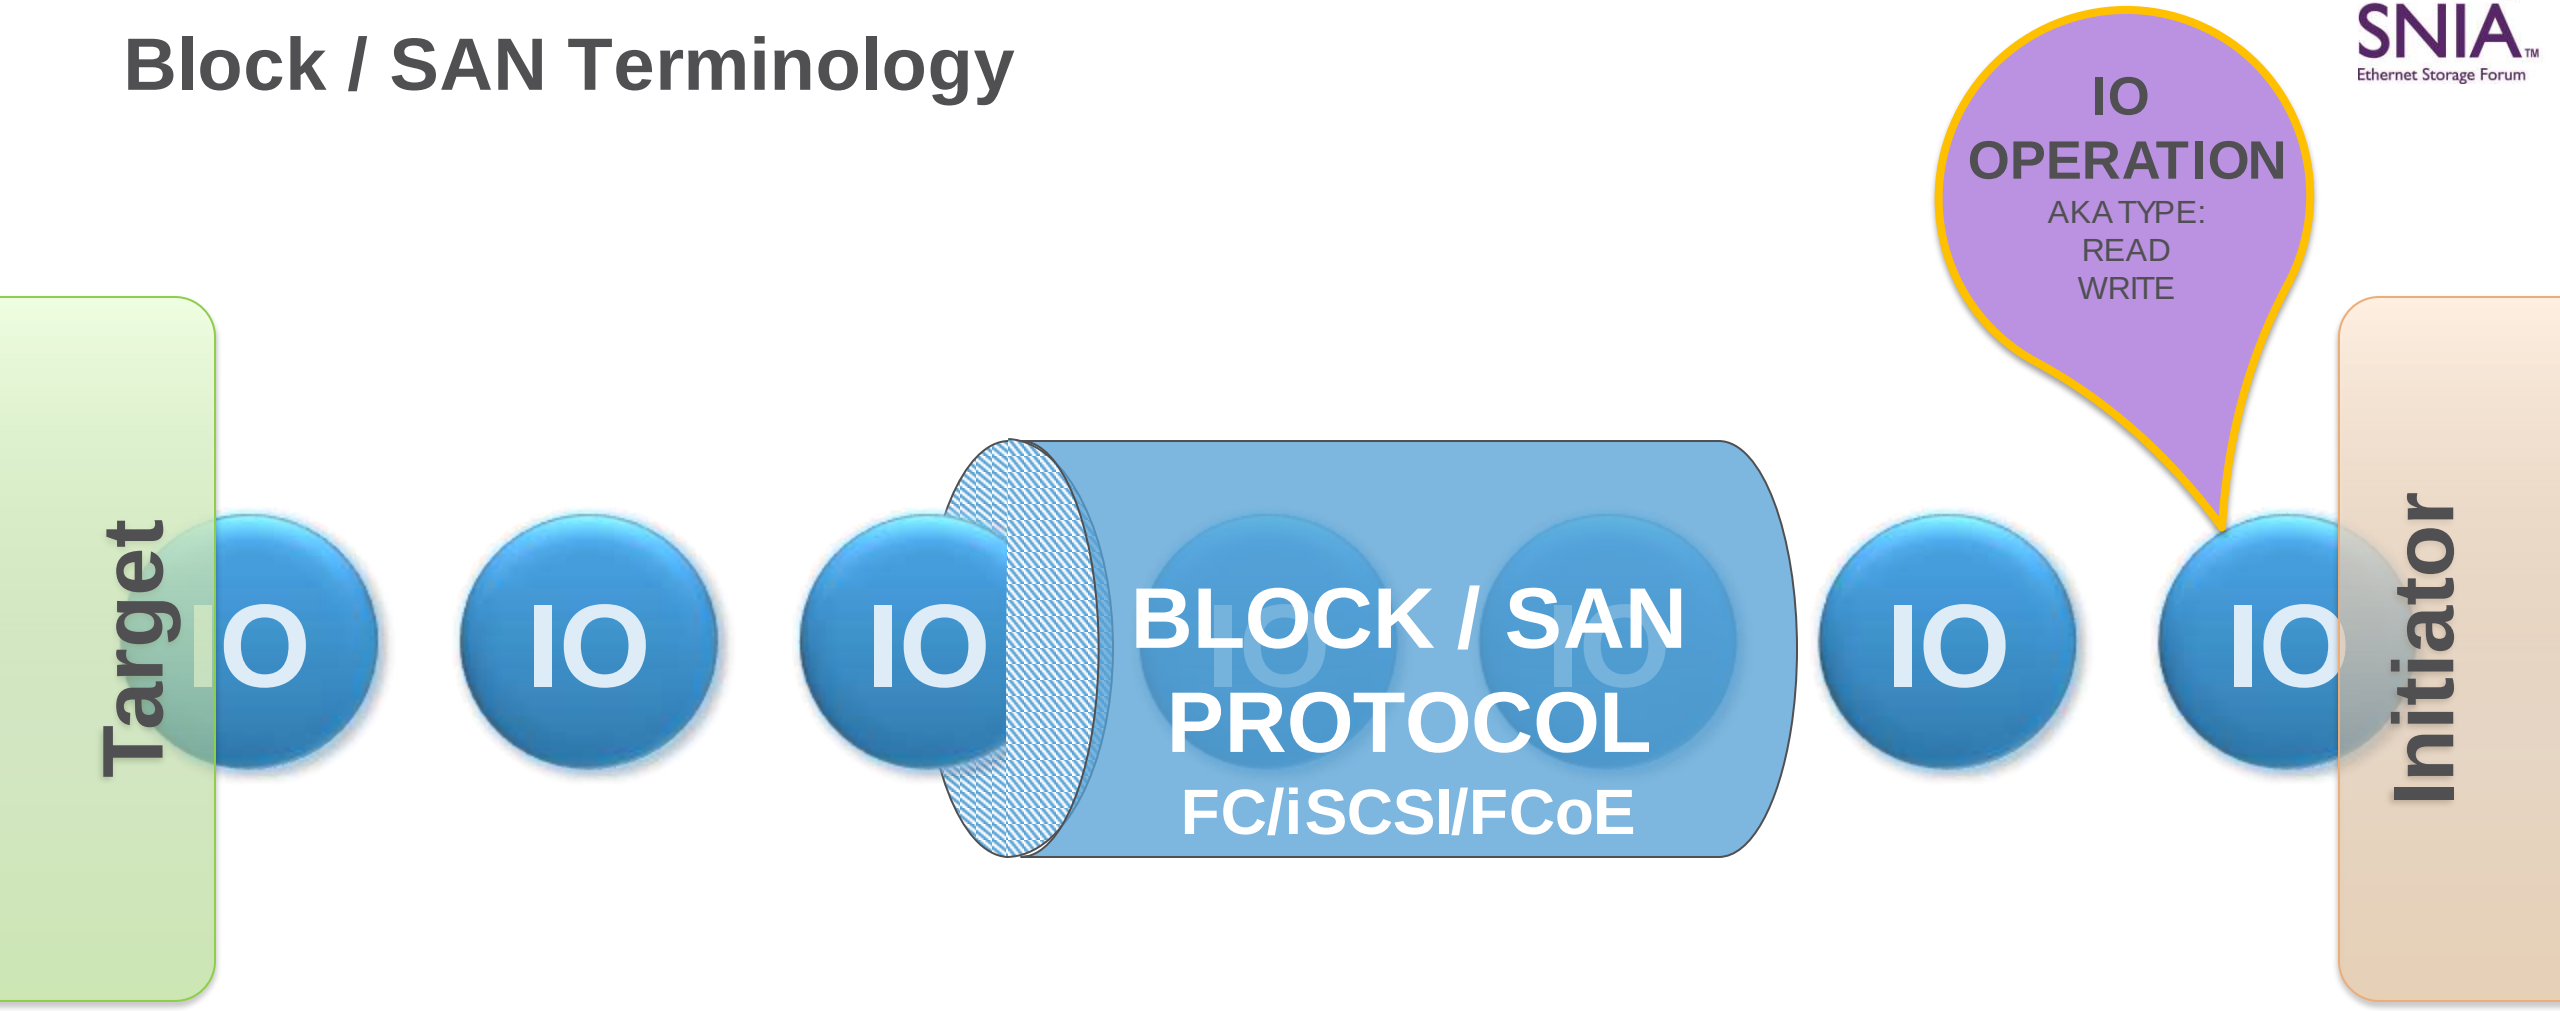

Block / SAN Operations

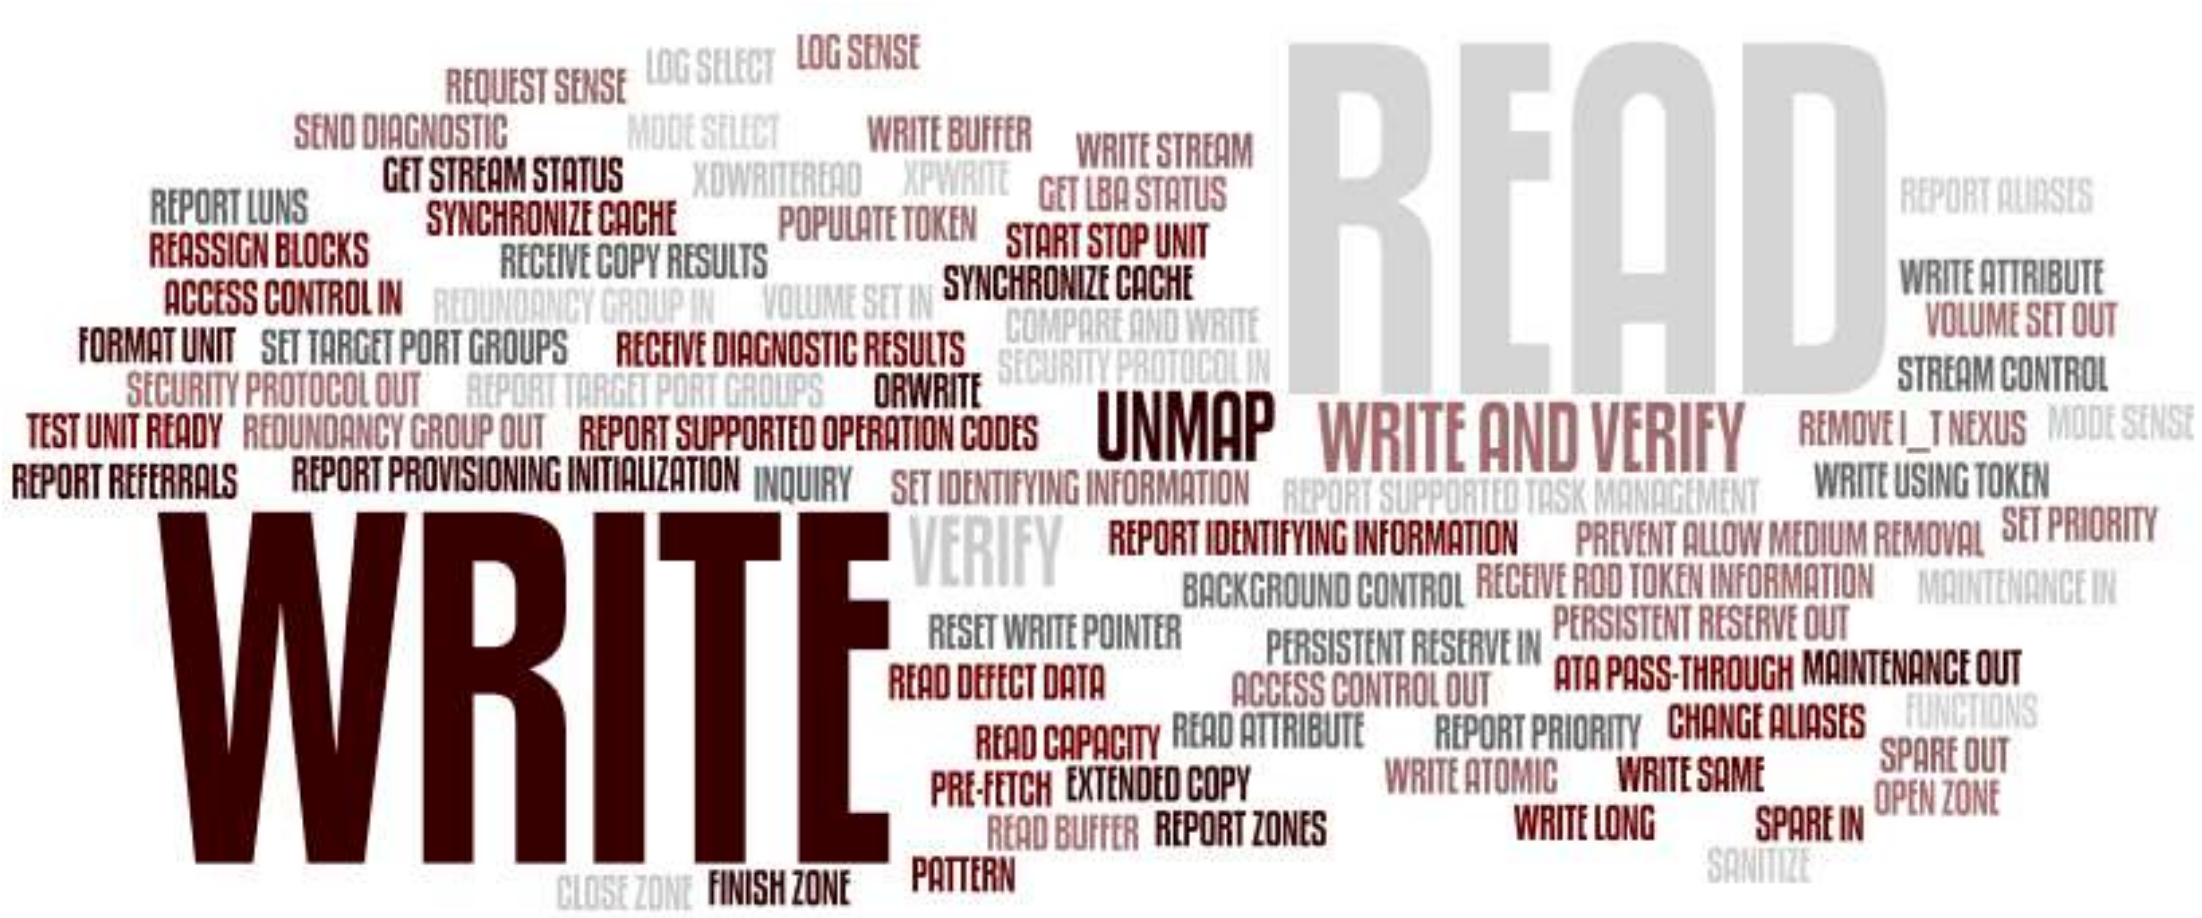

Source: http://www.t10.org/cgi-bin/ac.pl?t=f&f=sbc4r07c.pdf
Graphic: http://www.wordle.net/create

Block / SAN Terminology

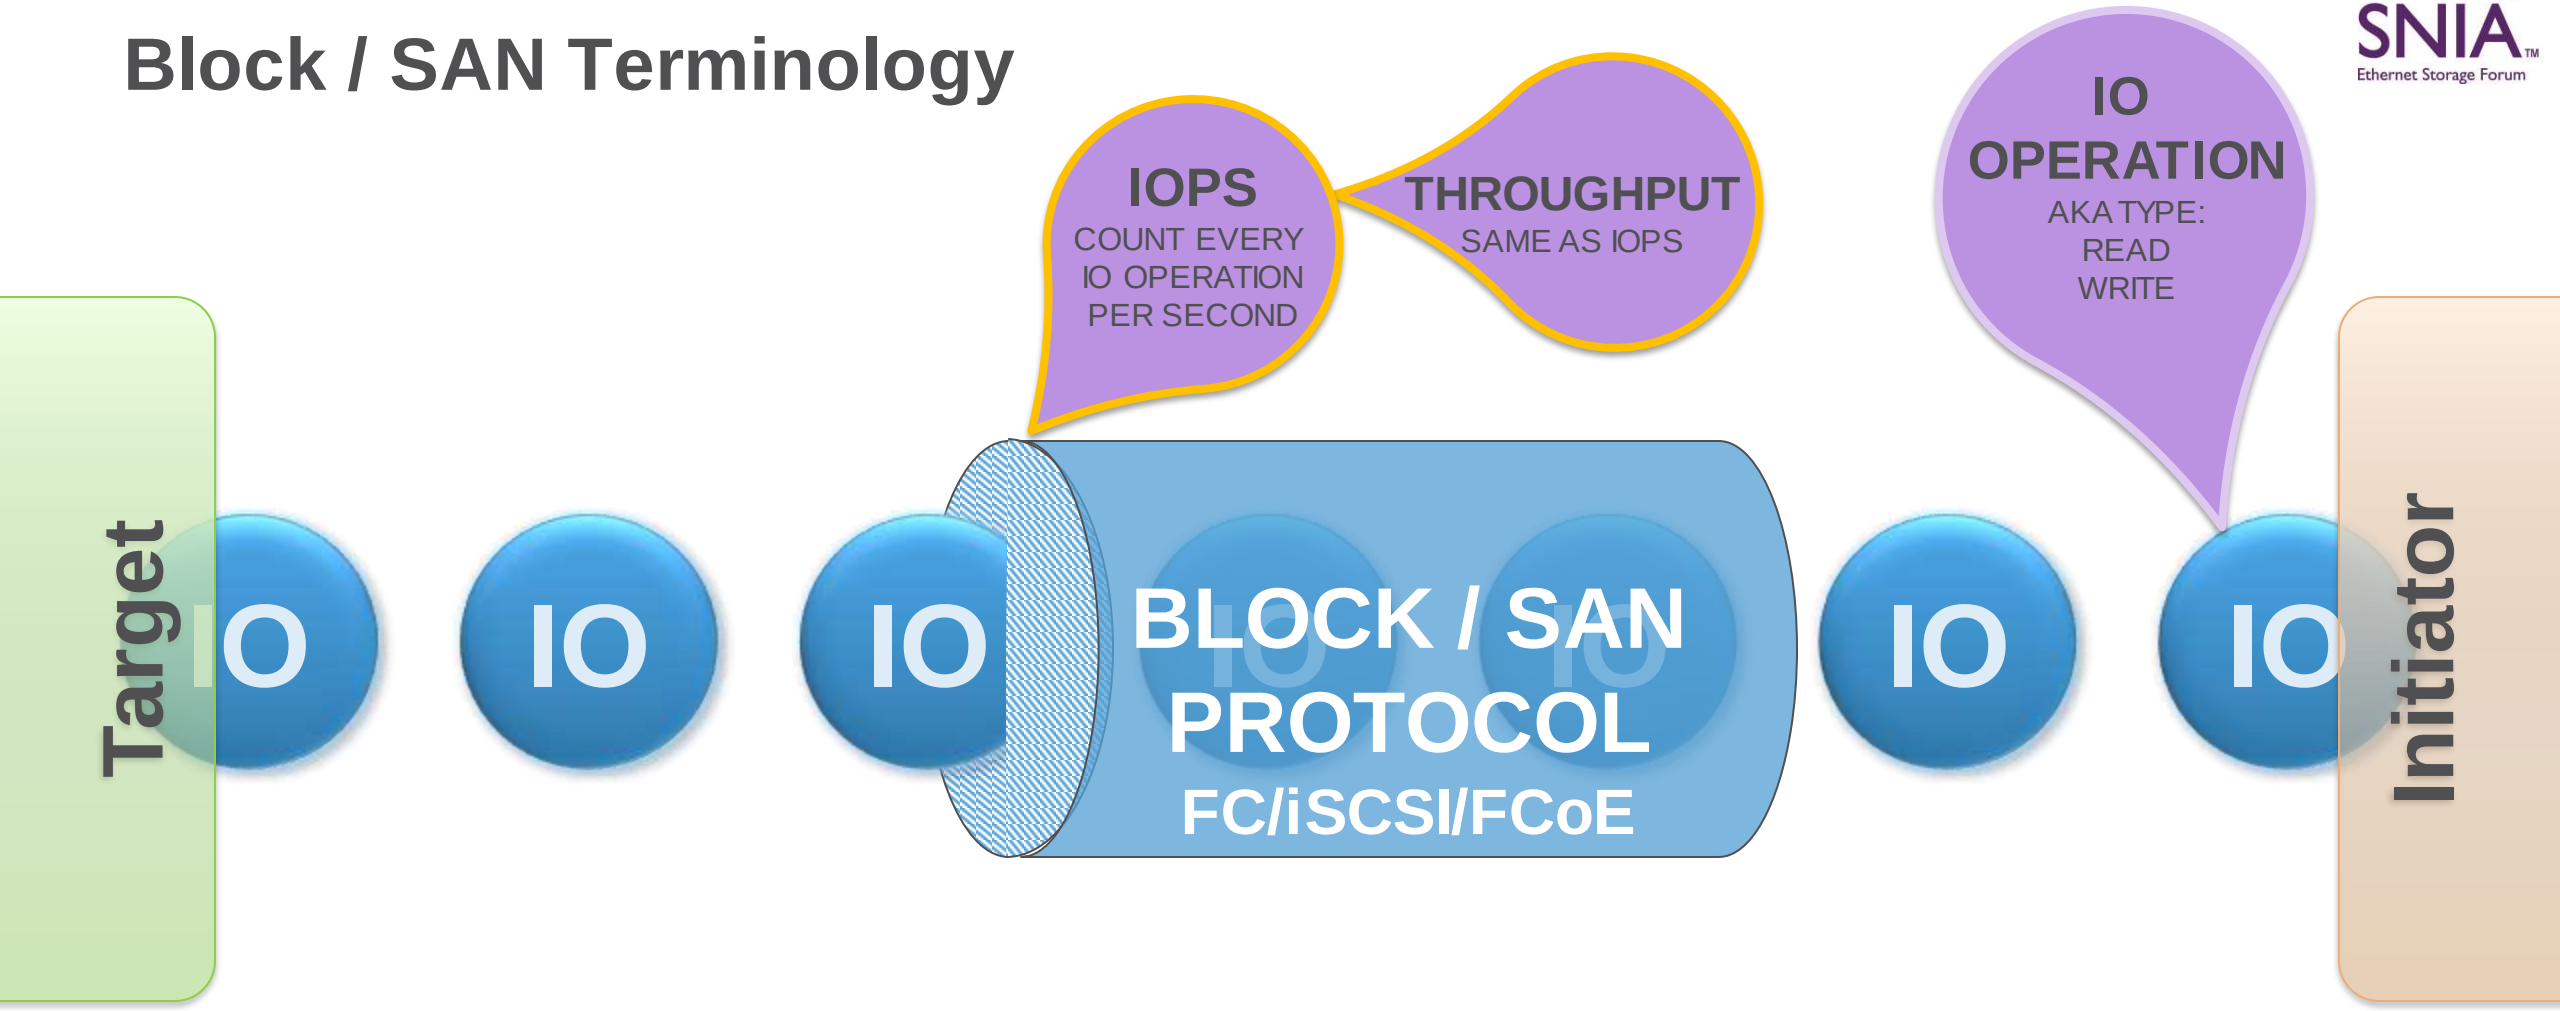

Where To Measure IOPS?

DISK DRIVES

STORAGE CONTROLLER

HOSTS

Where To Measure IOPS?

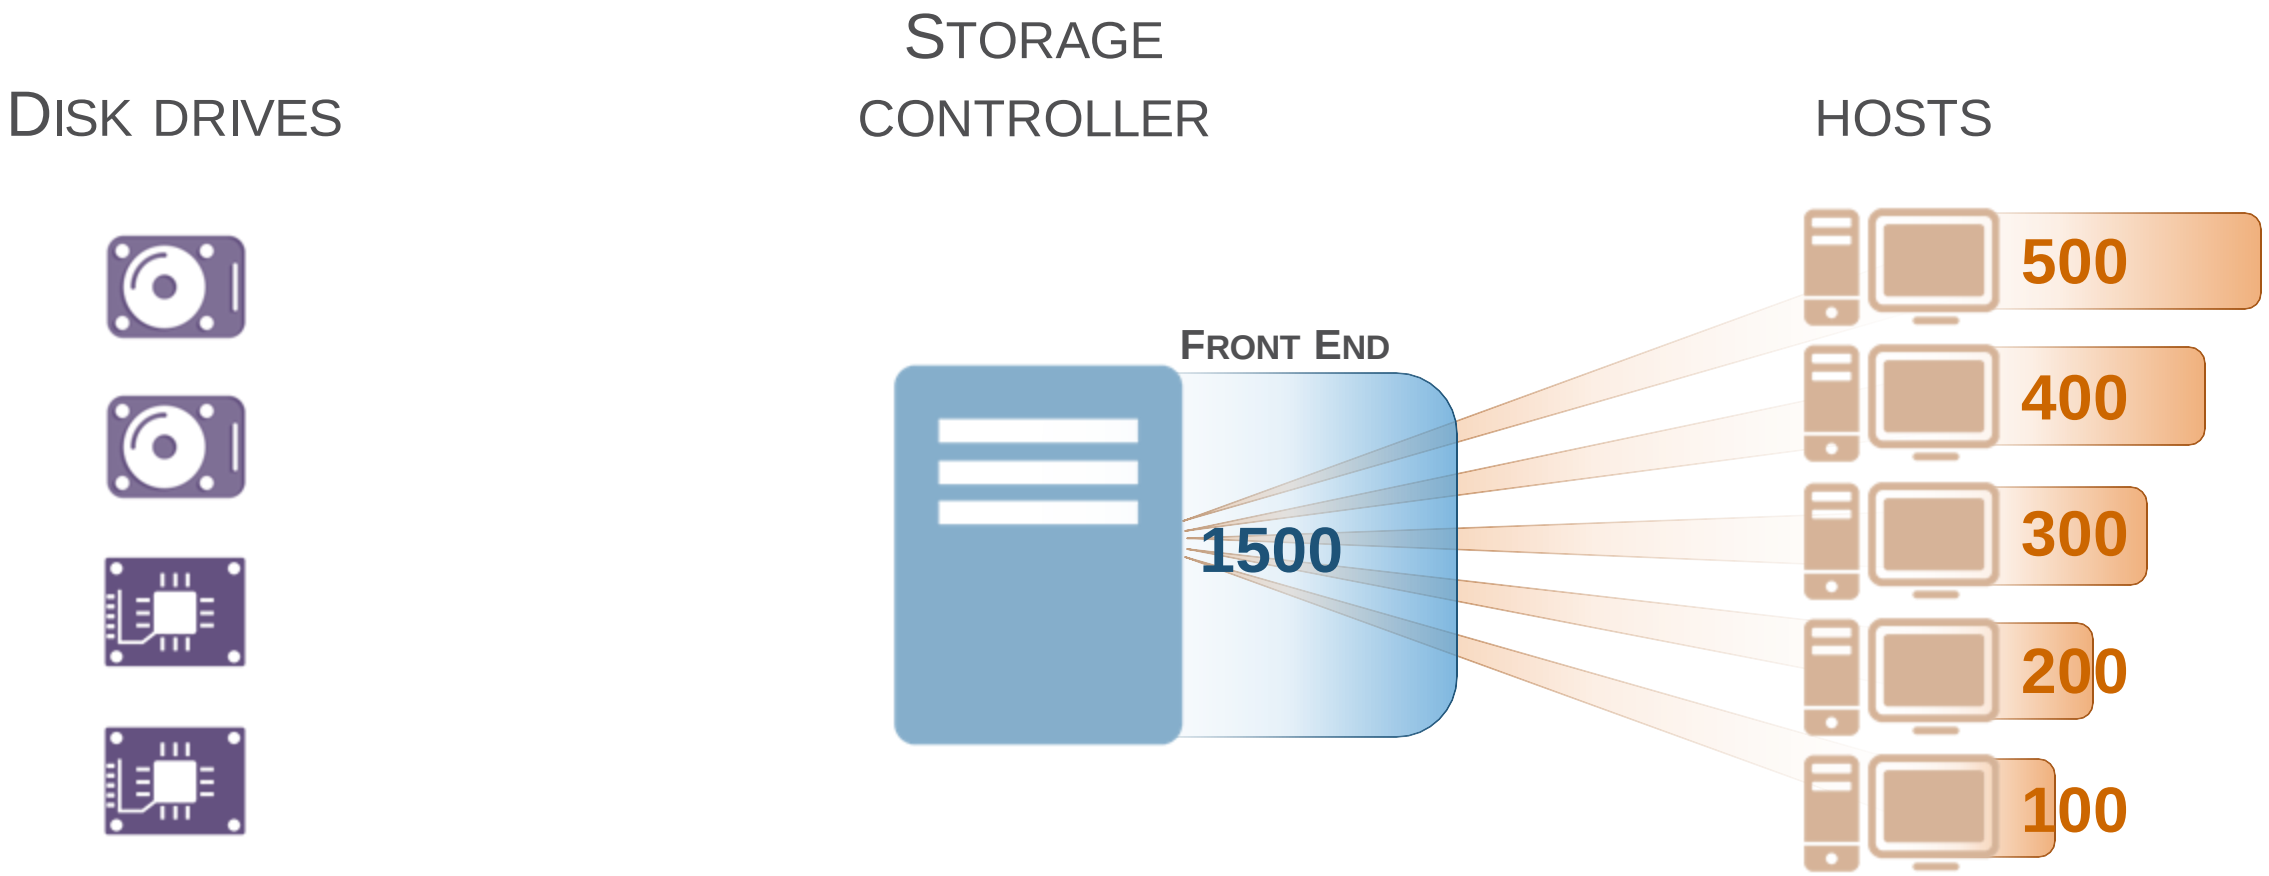

Where To Measure IOPS?

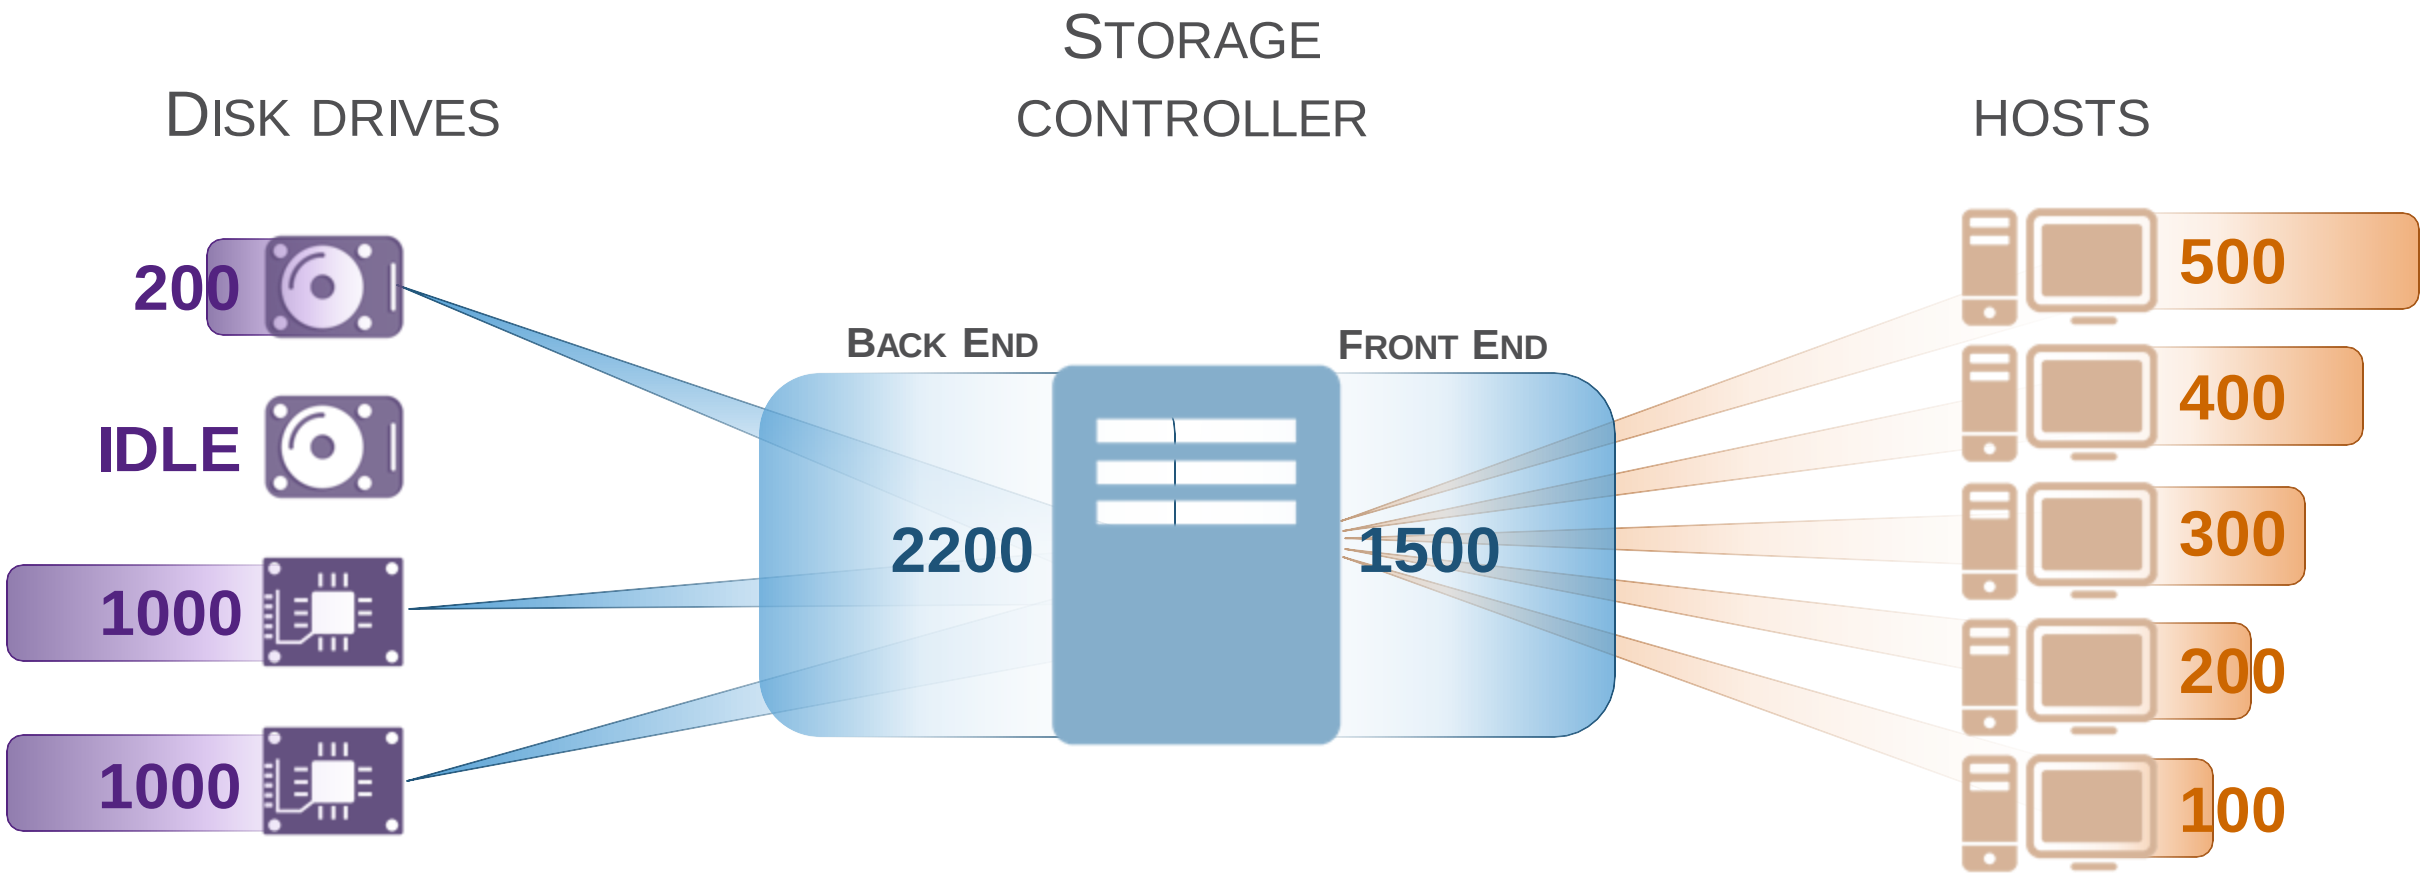

Block / SAN Terminology

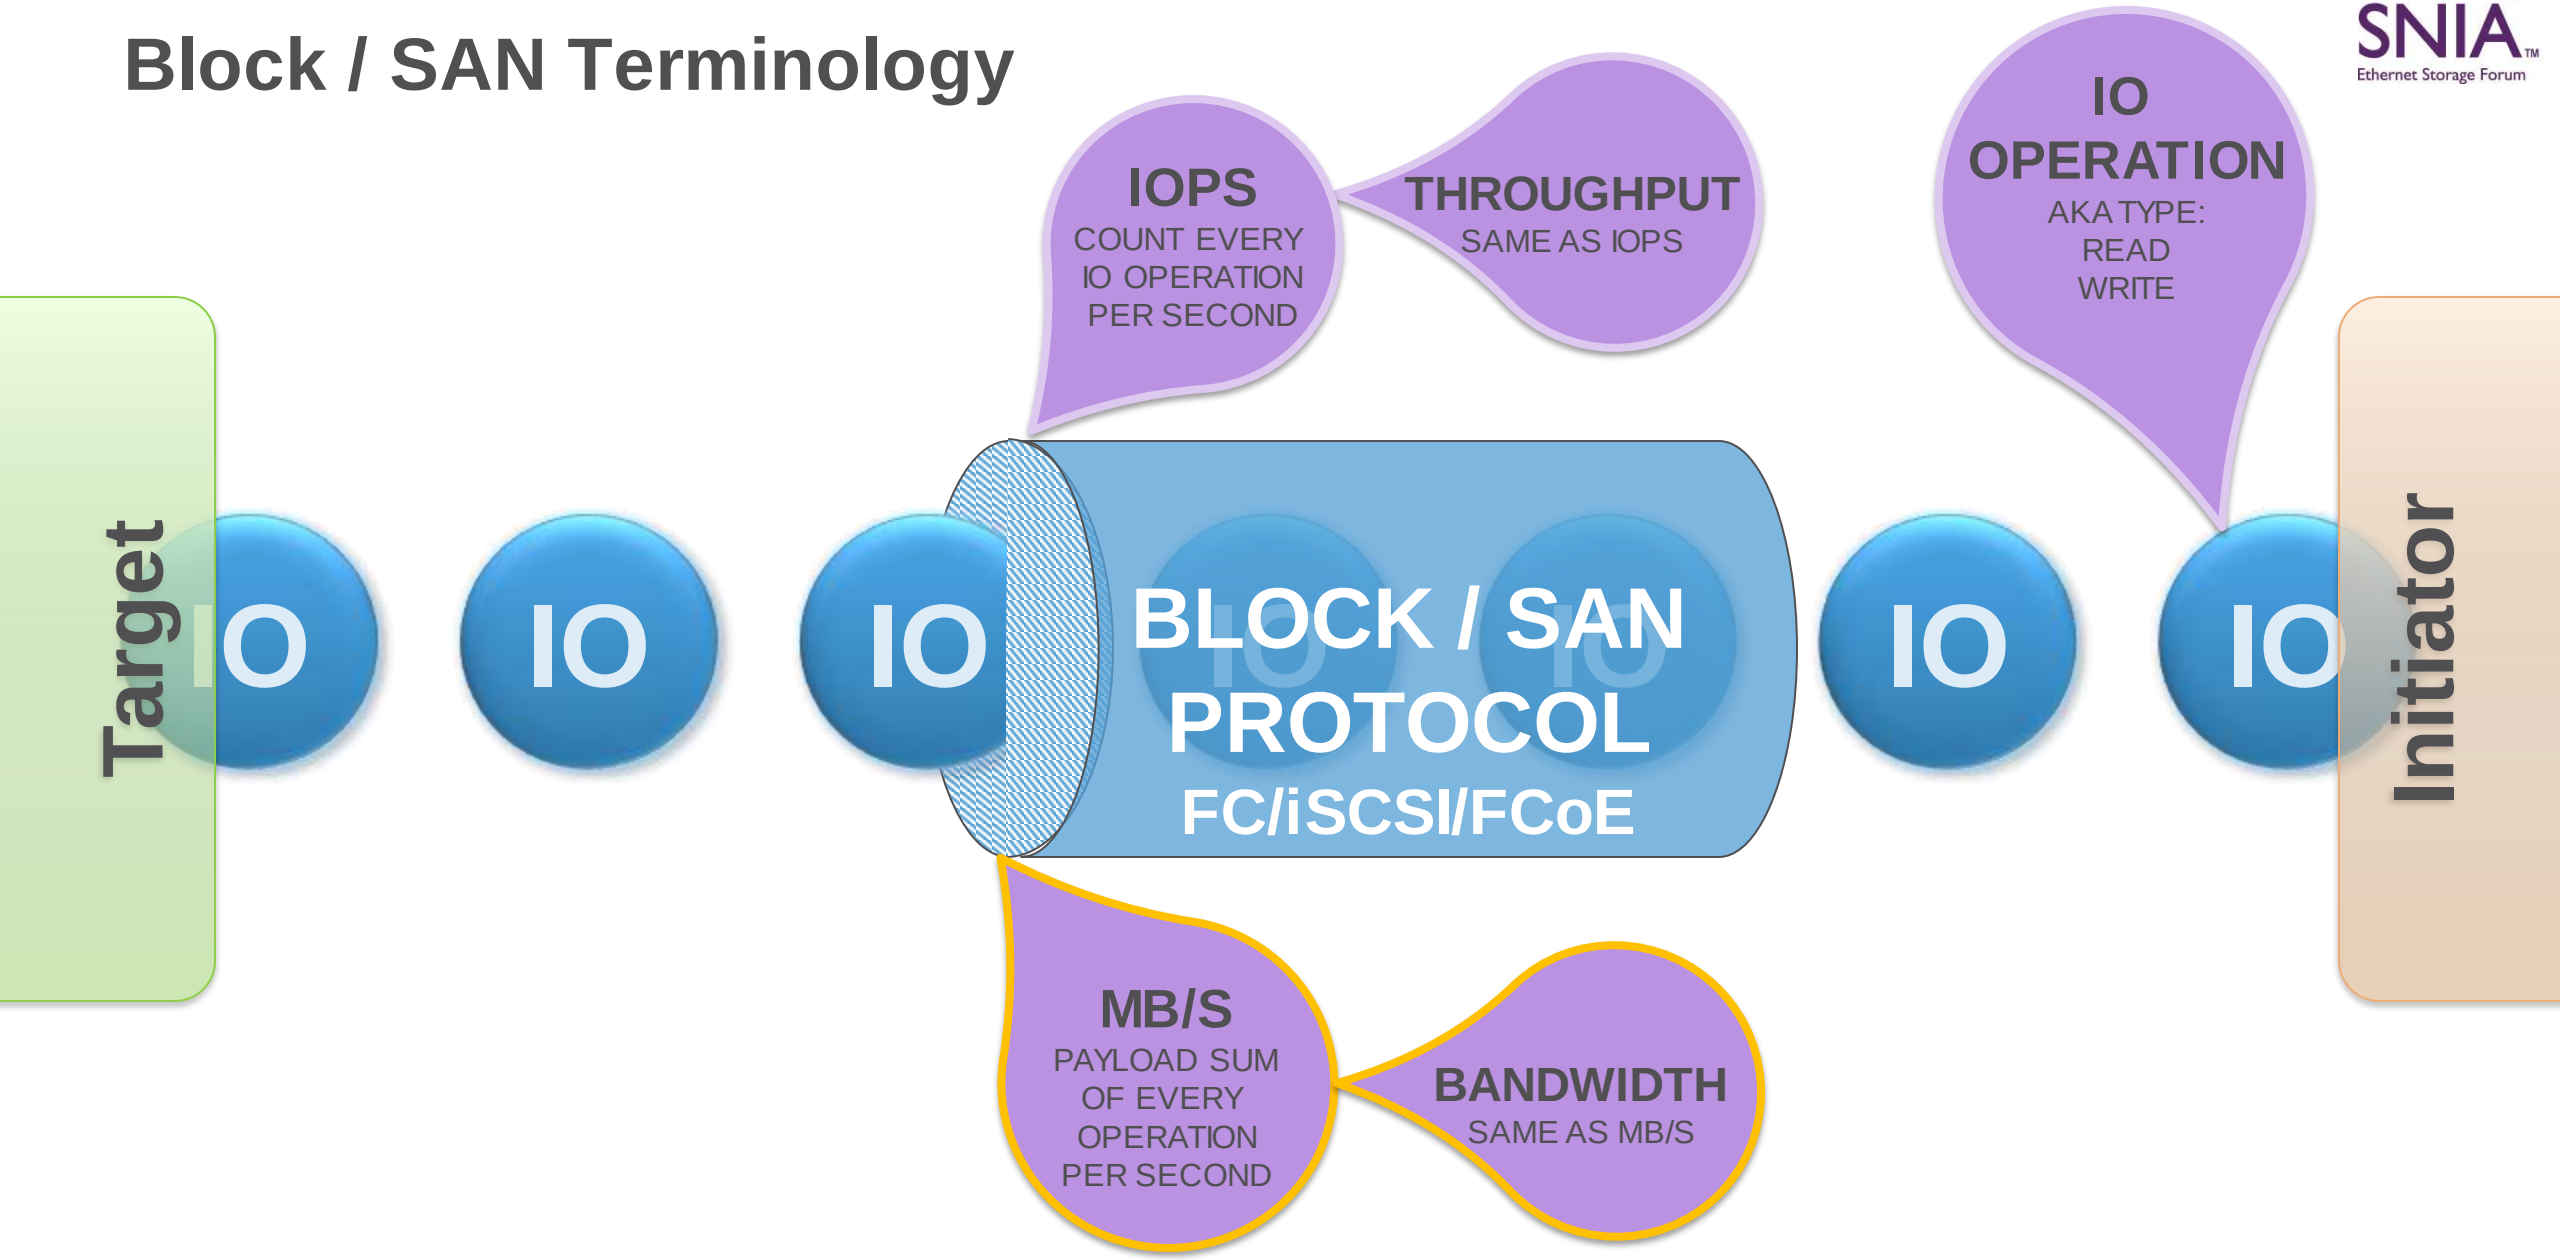

MB/s Vs. MiB/s

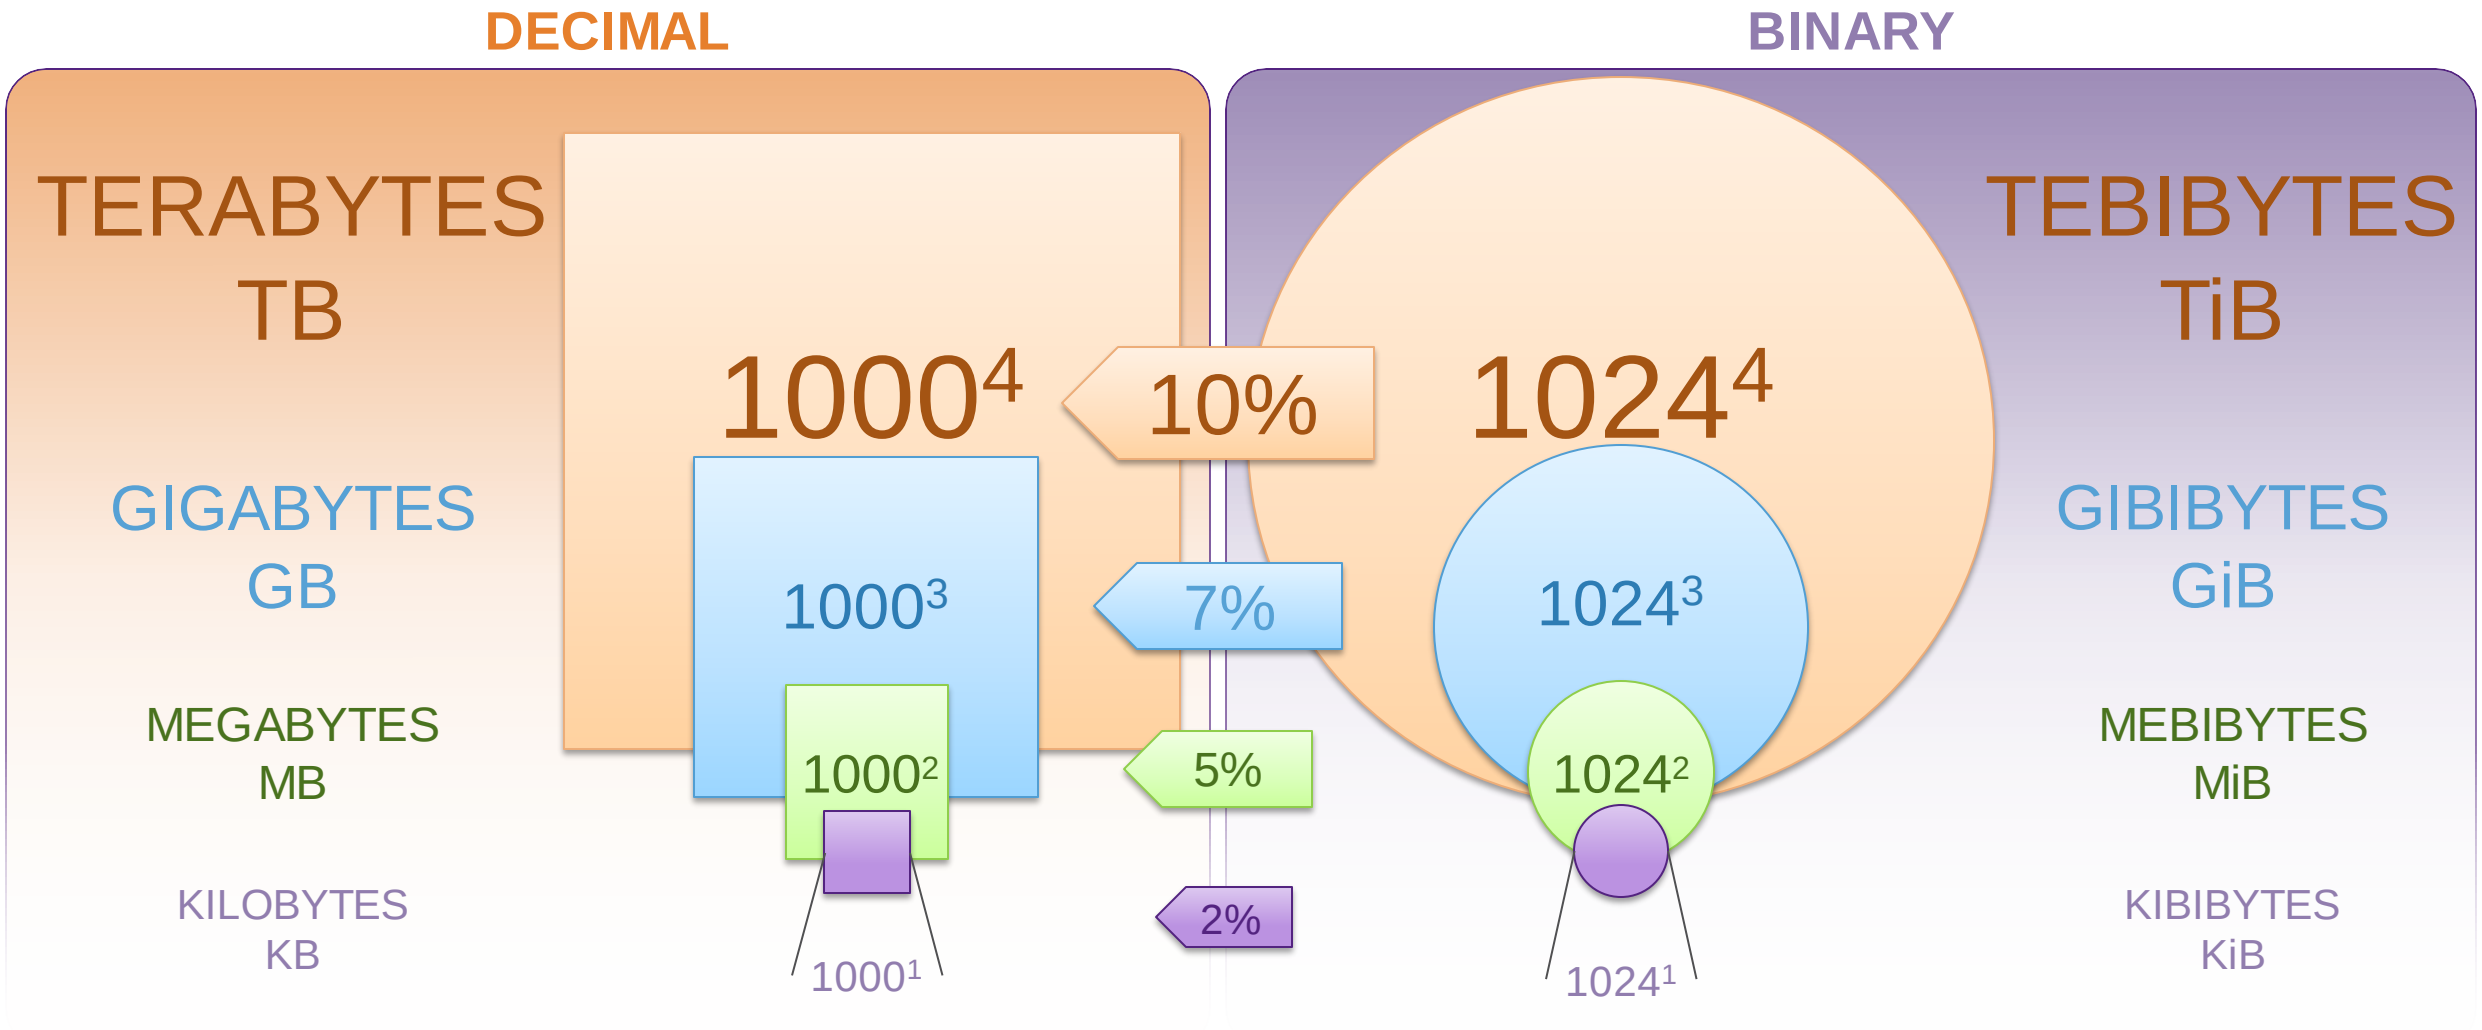

Block / SAN Terminology

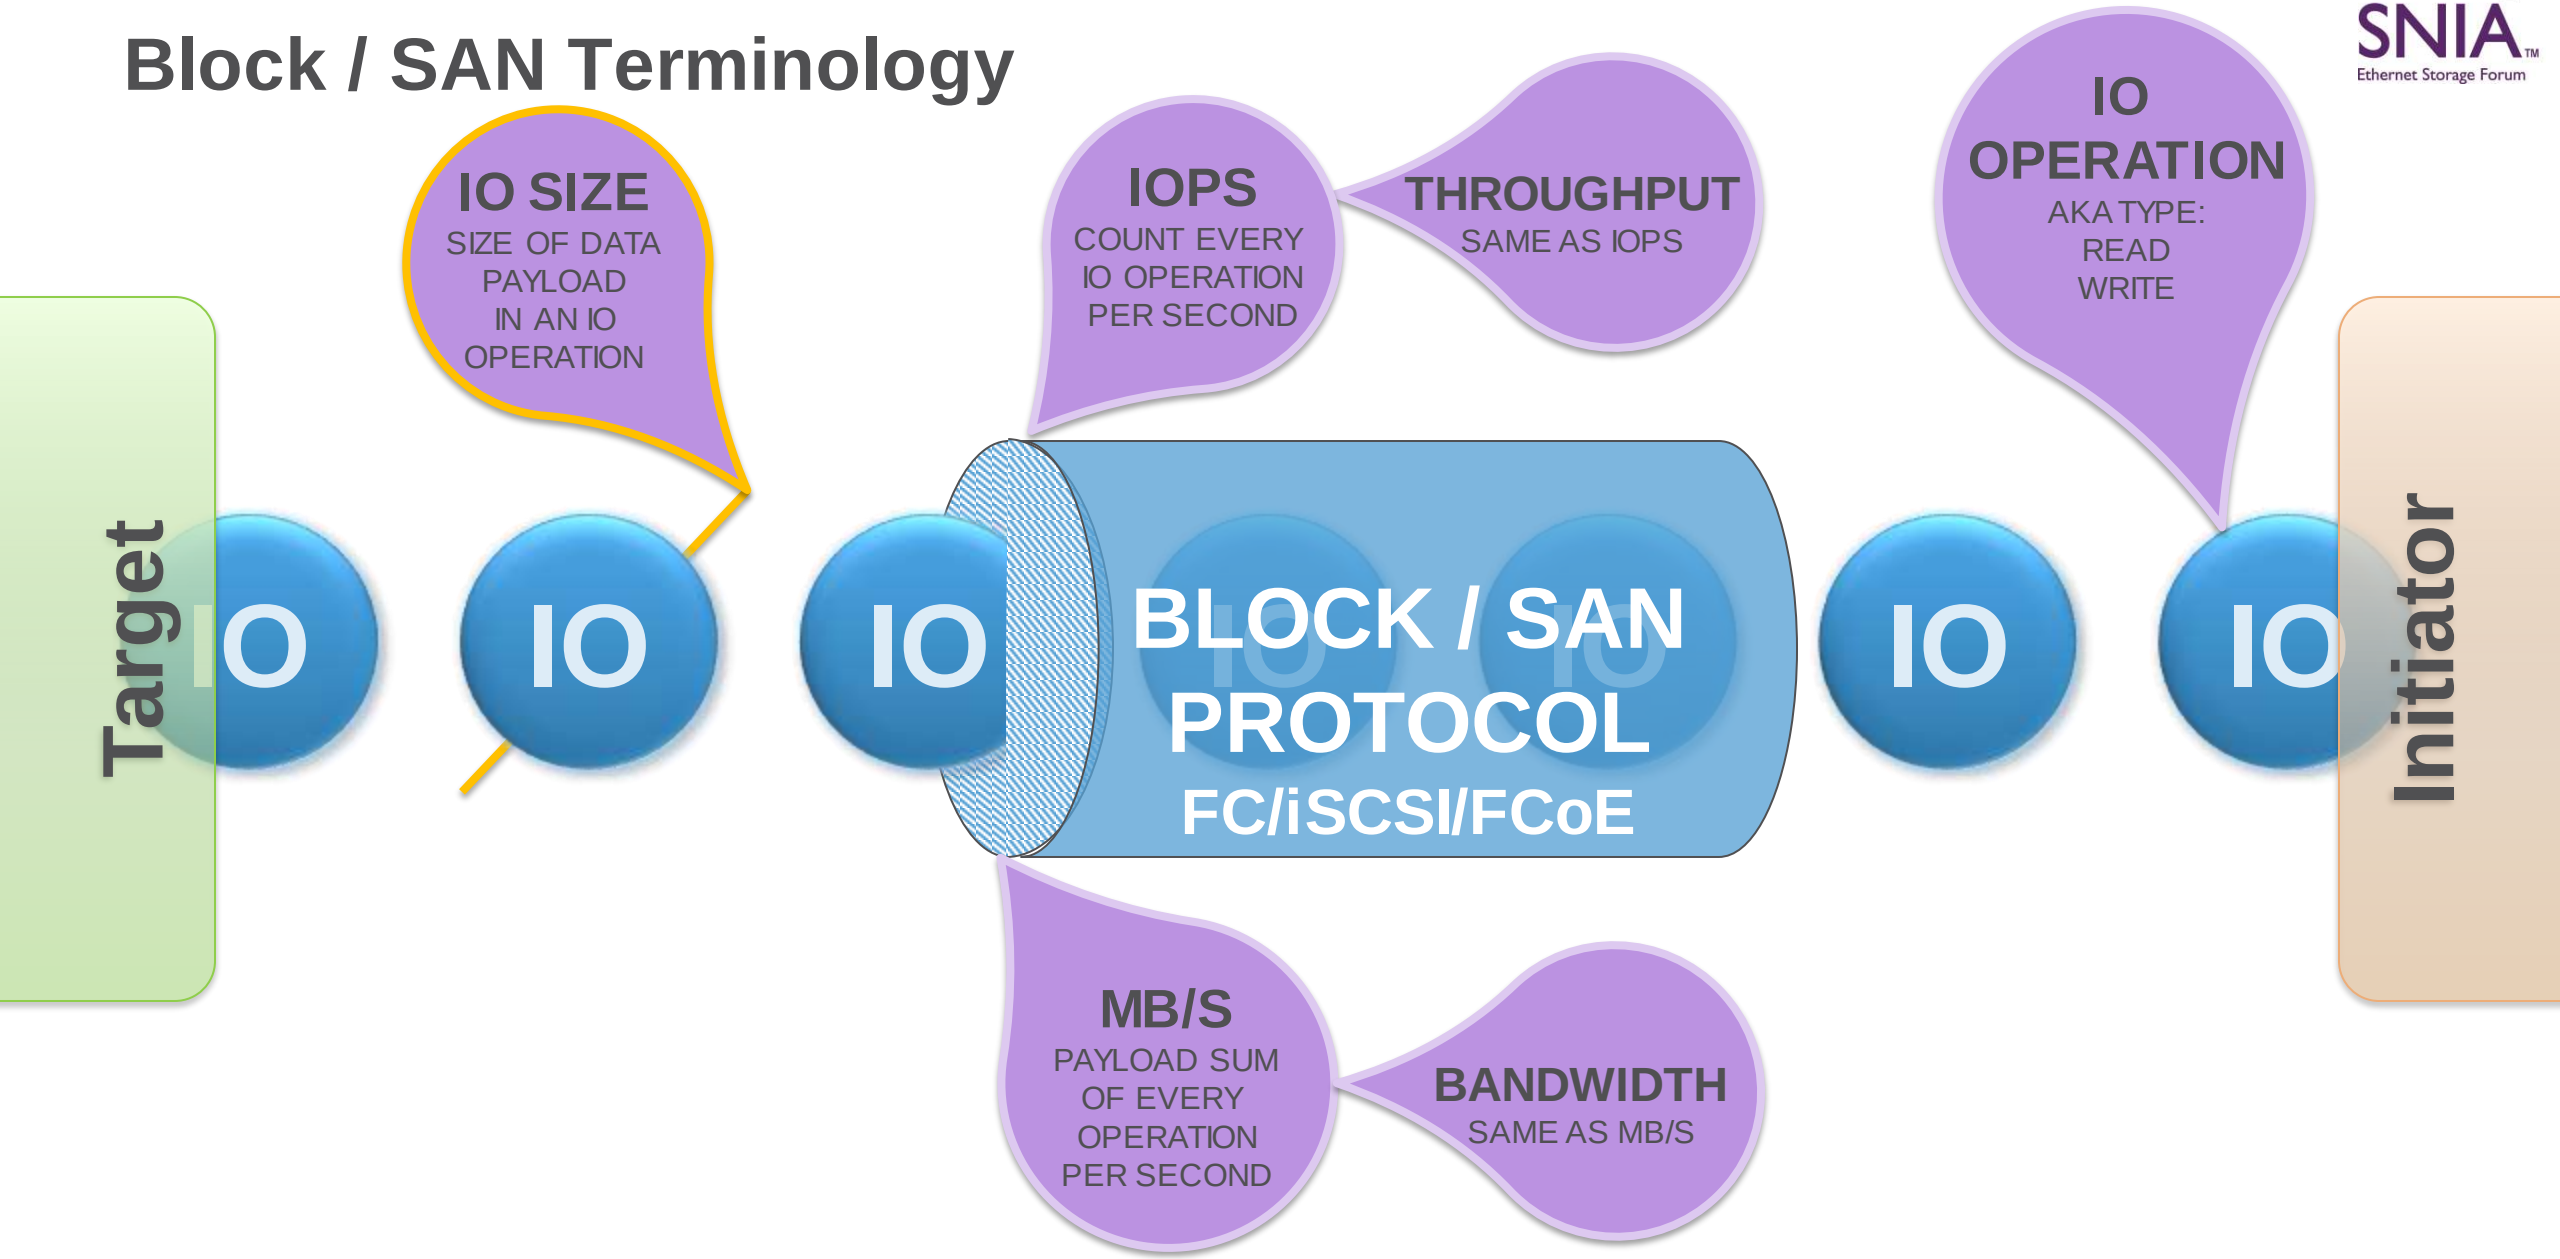

MB/s (or MiB/s), IO Size, And IOPS

$$MiB/s = IO\ Size * IOPS$$

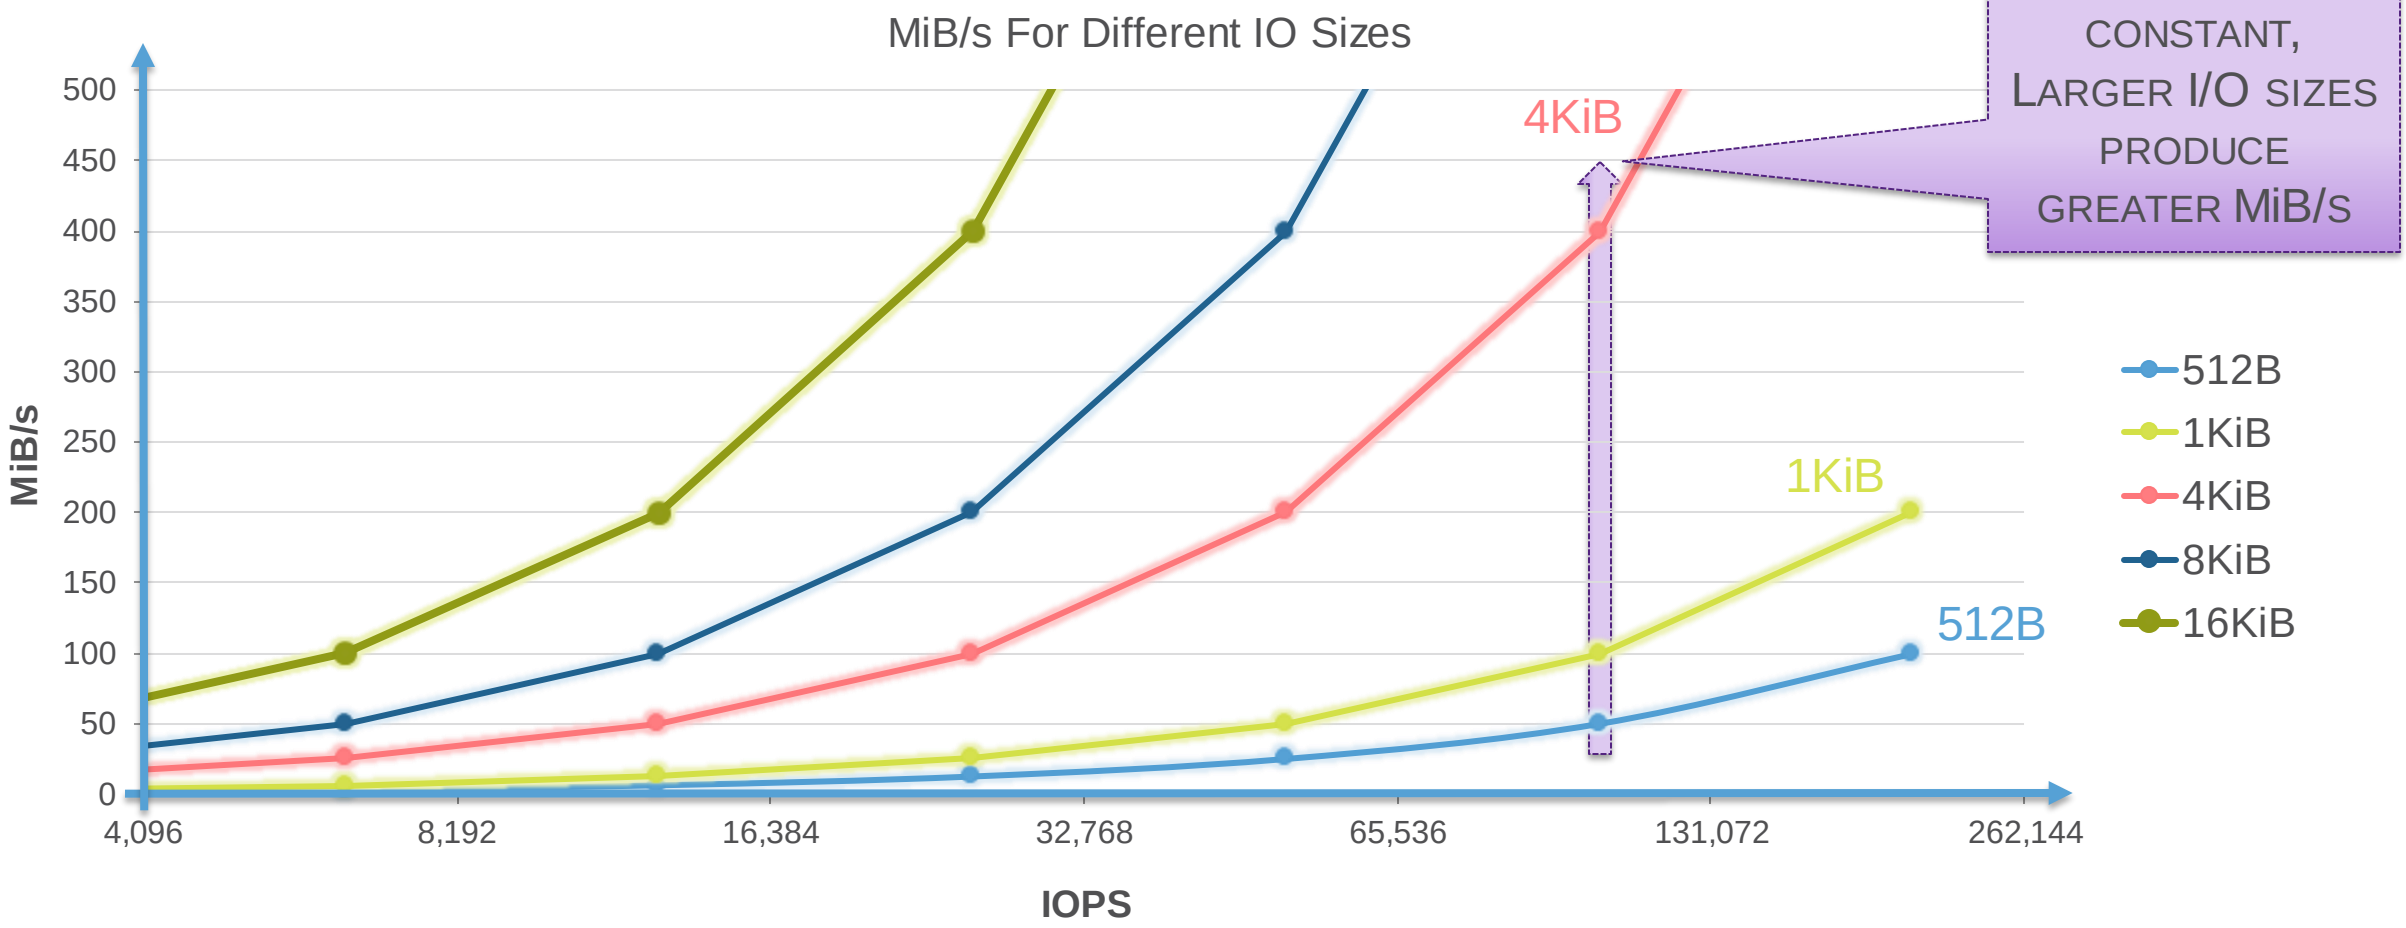

MB/s (or MiB/s), IO Size, And IOPS

$$MiB/s = IO\ Size * IOPS$$

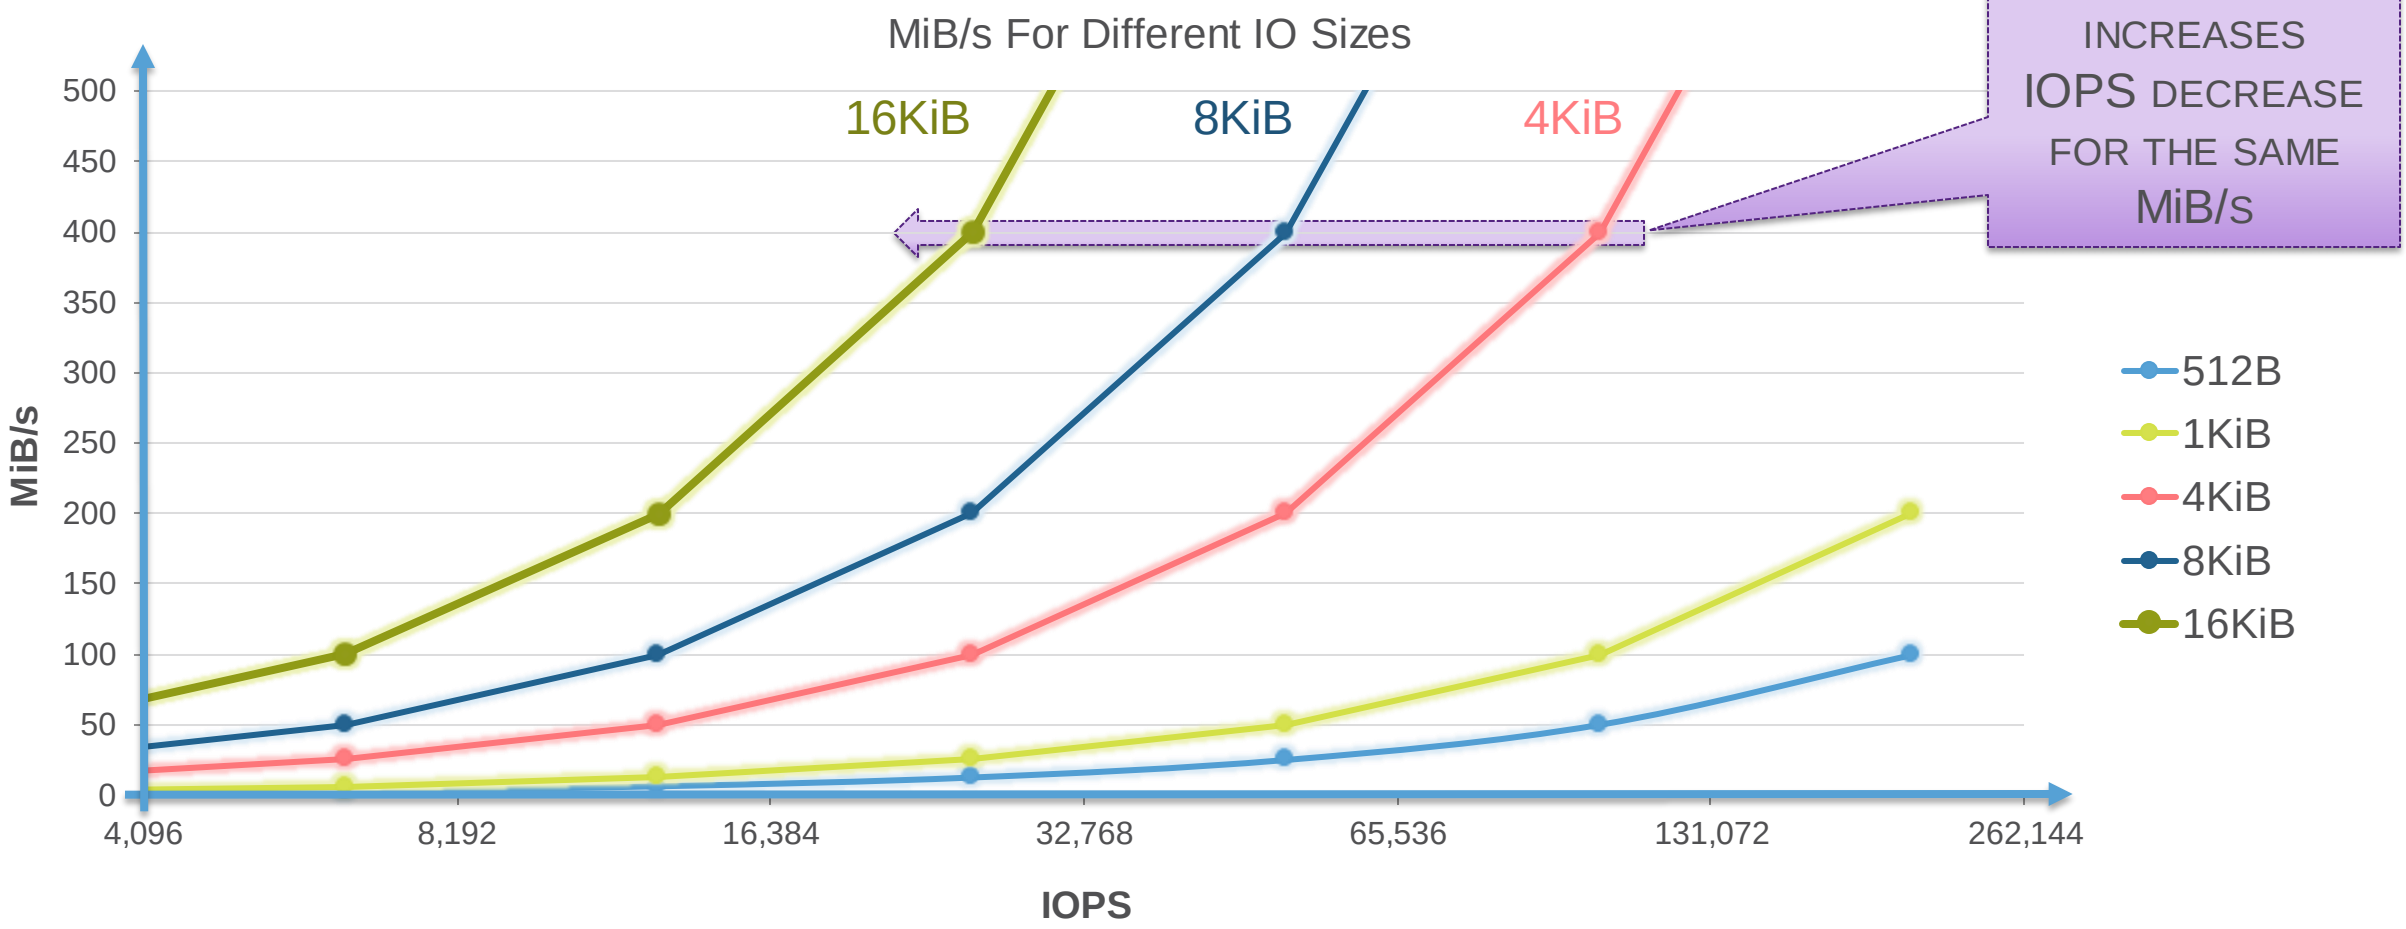

Block / SAN Terminology

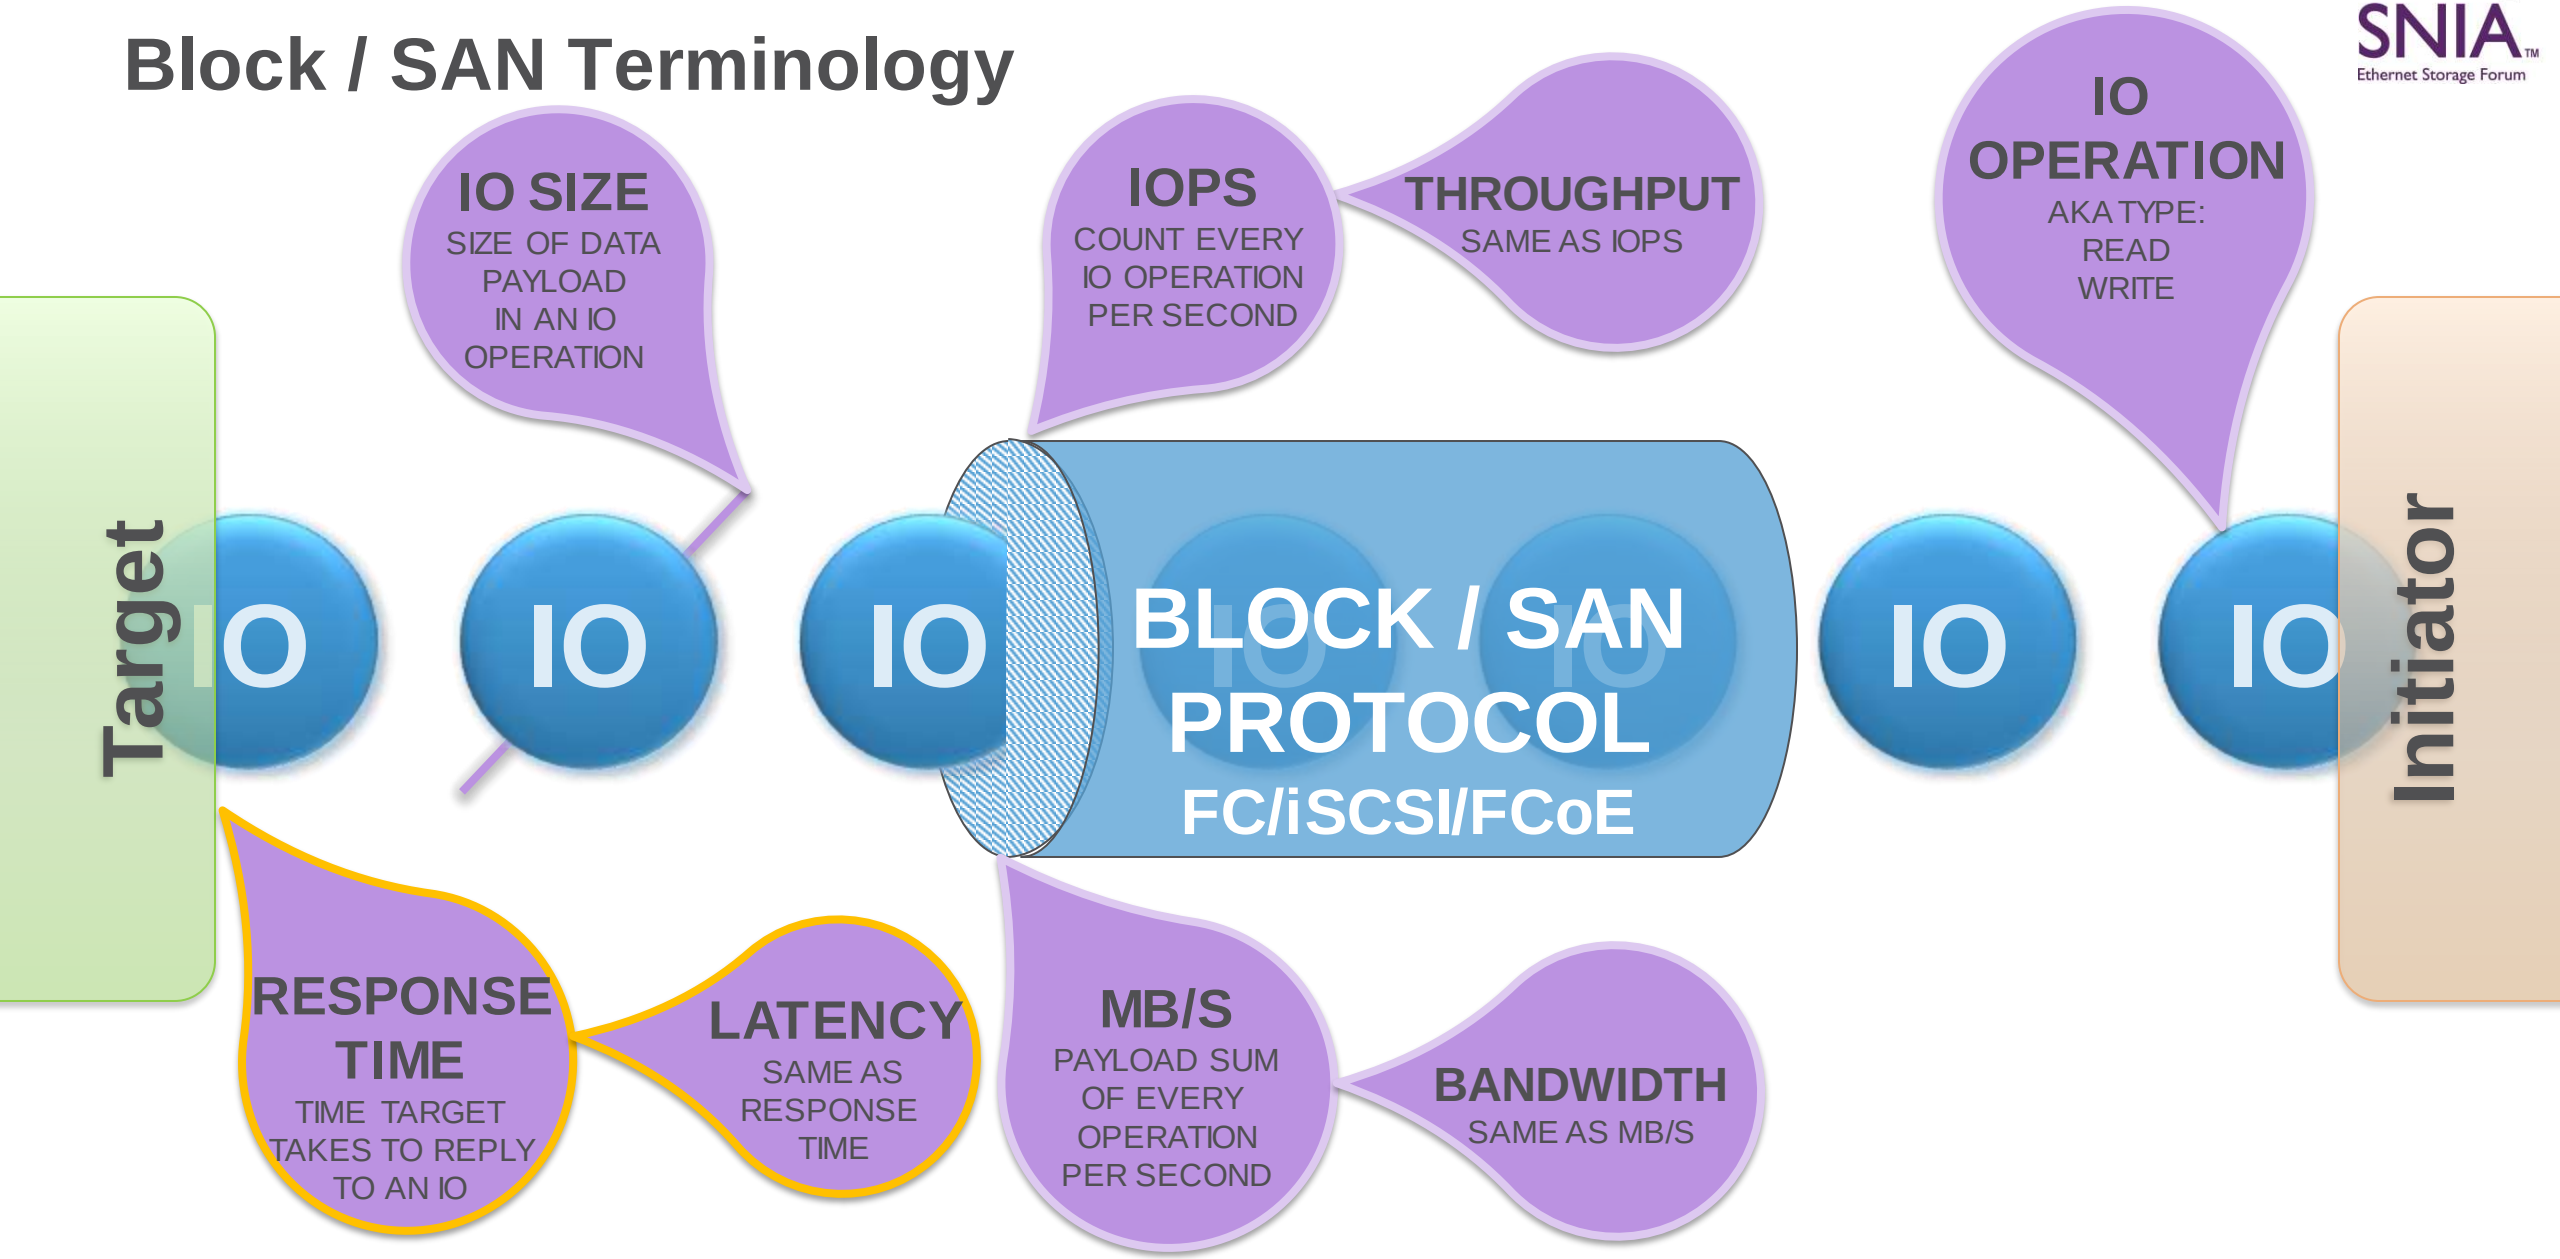

Aggregate Property Of Response Time

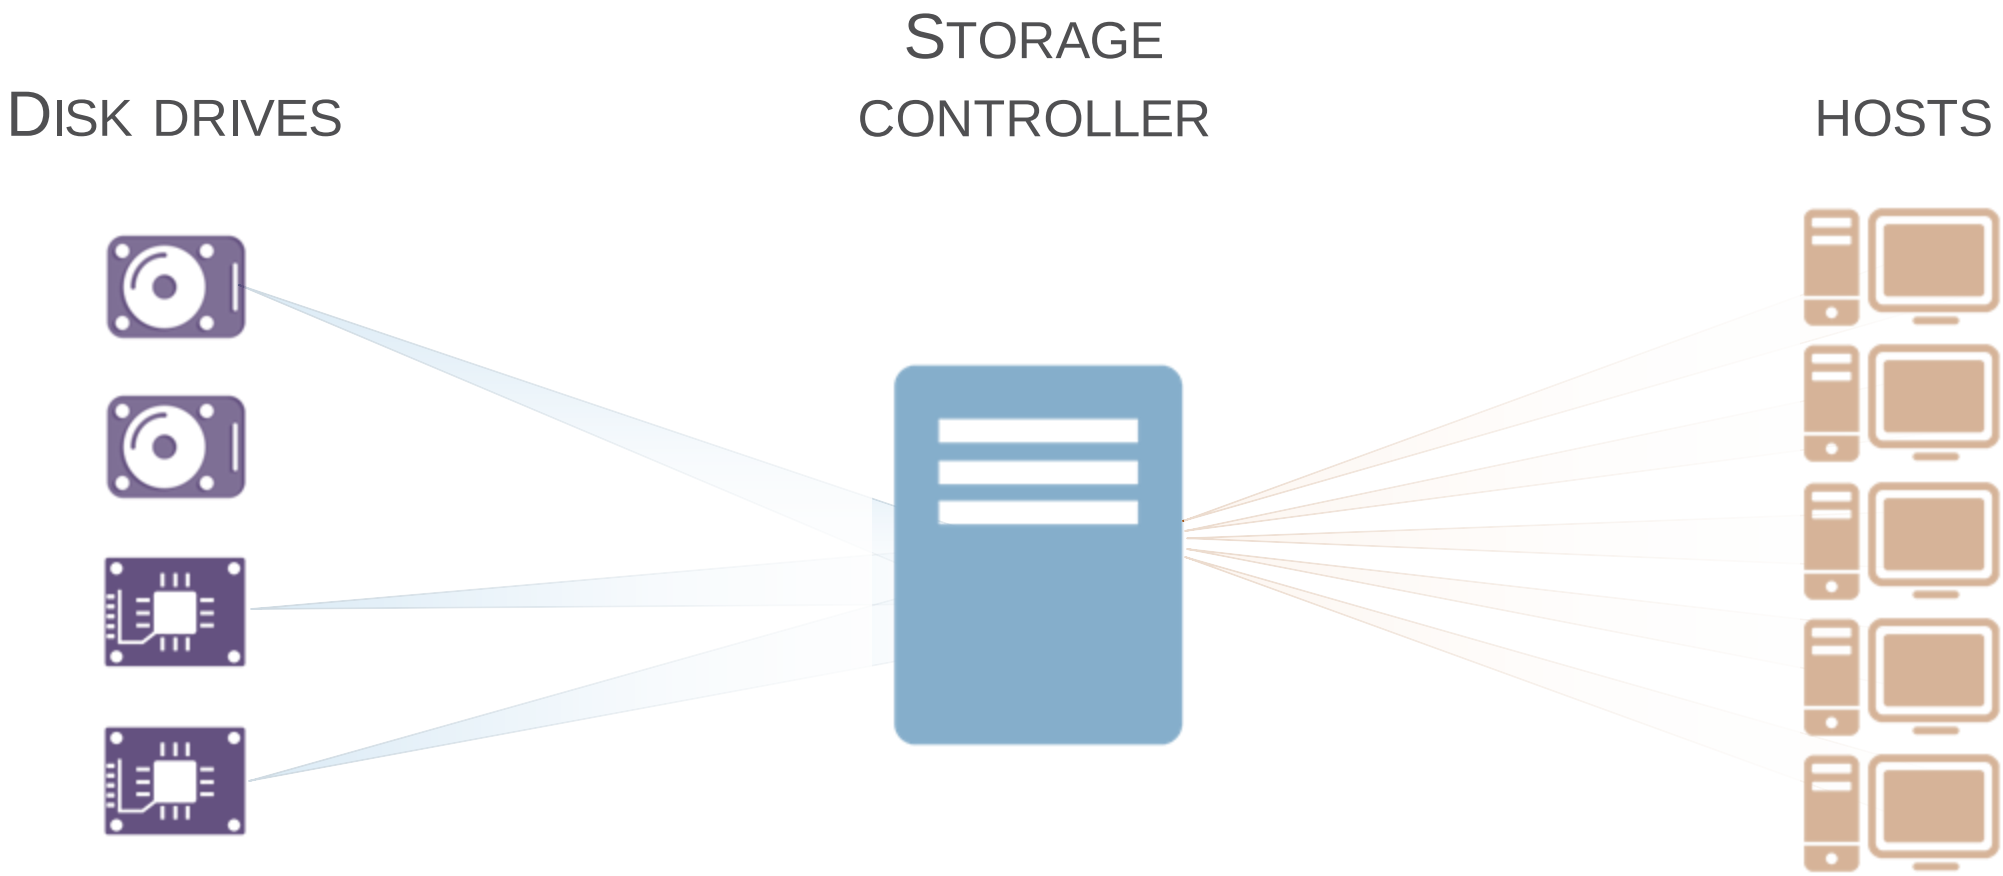

Aggregate Property Of Response Time

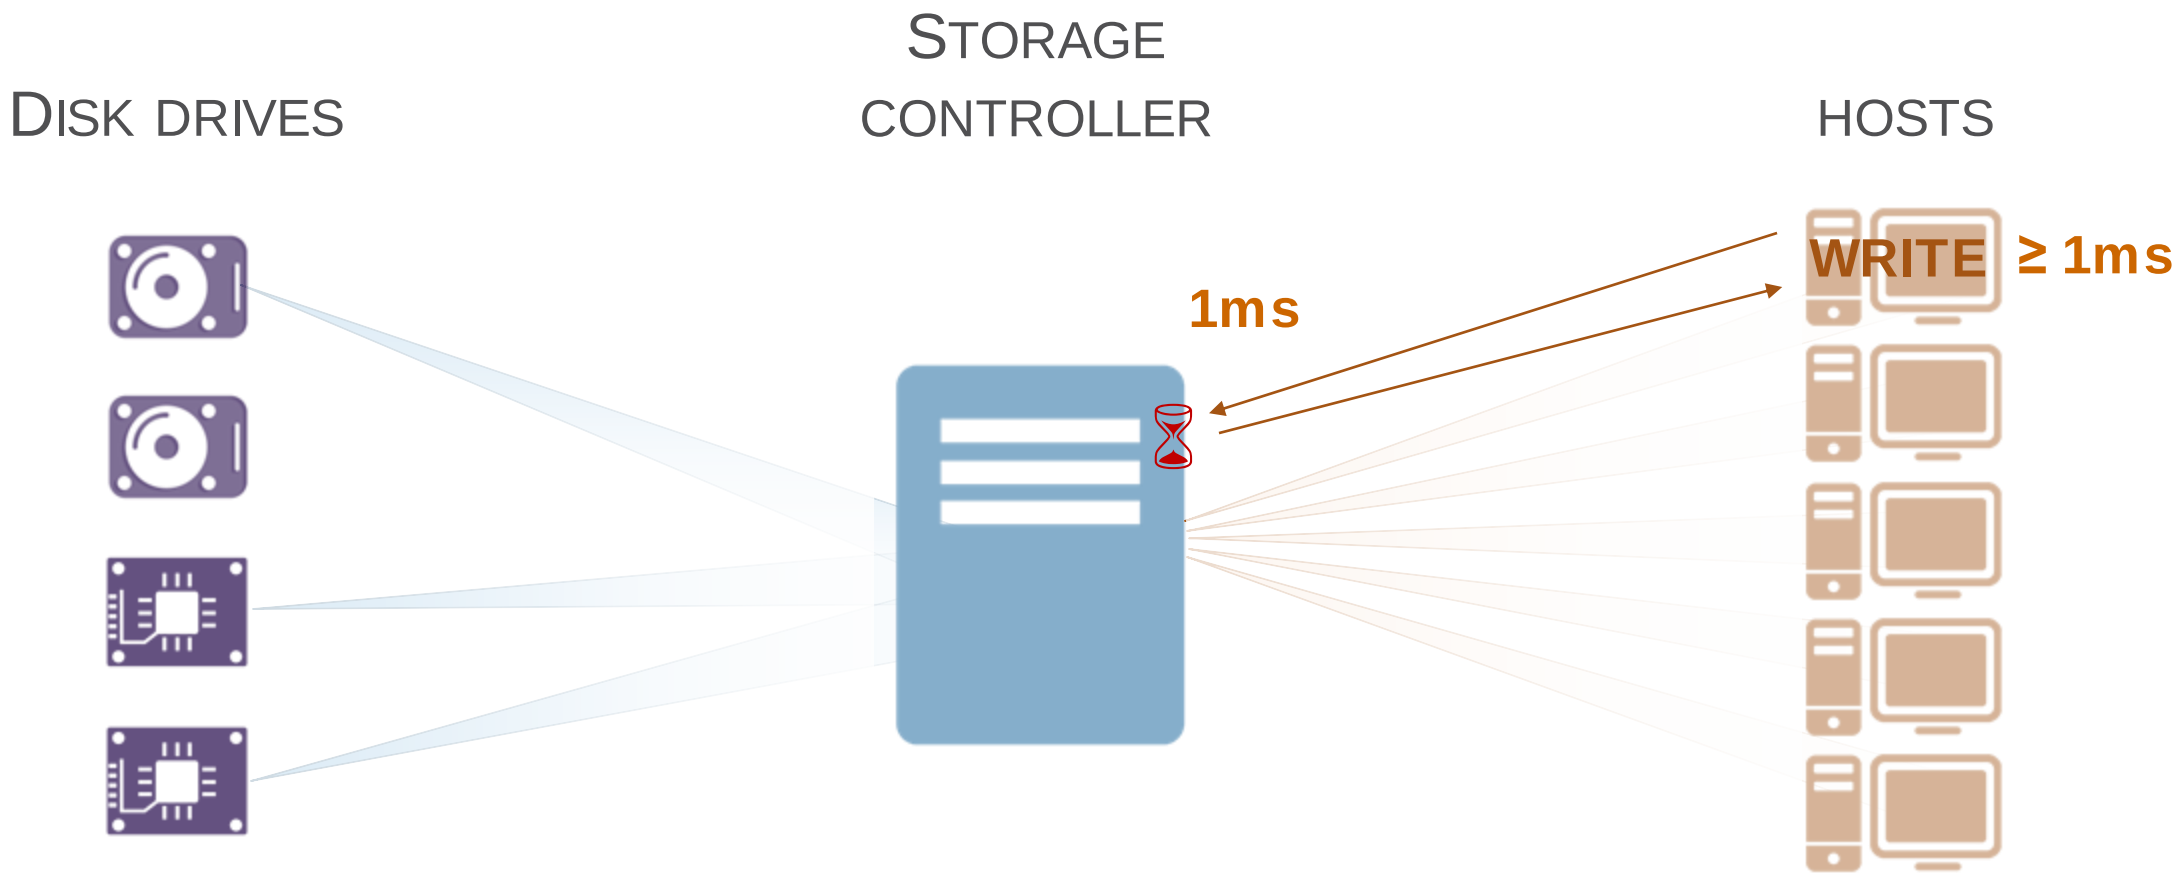

Aggregate Property Of Response Time

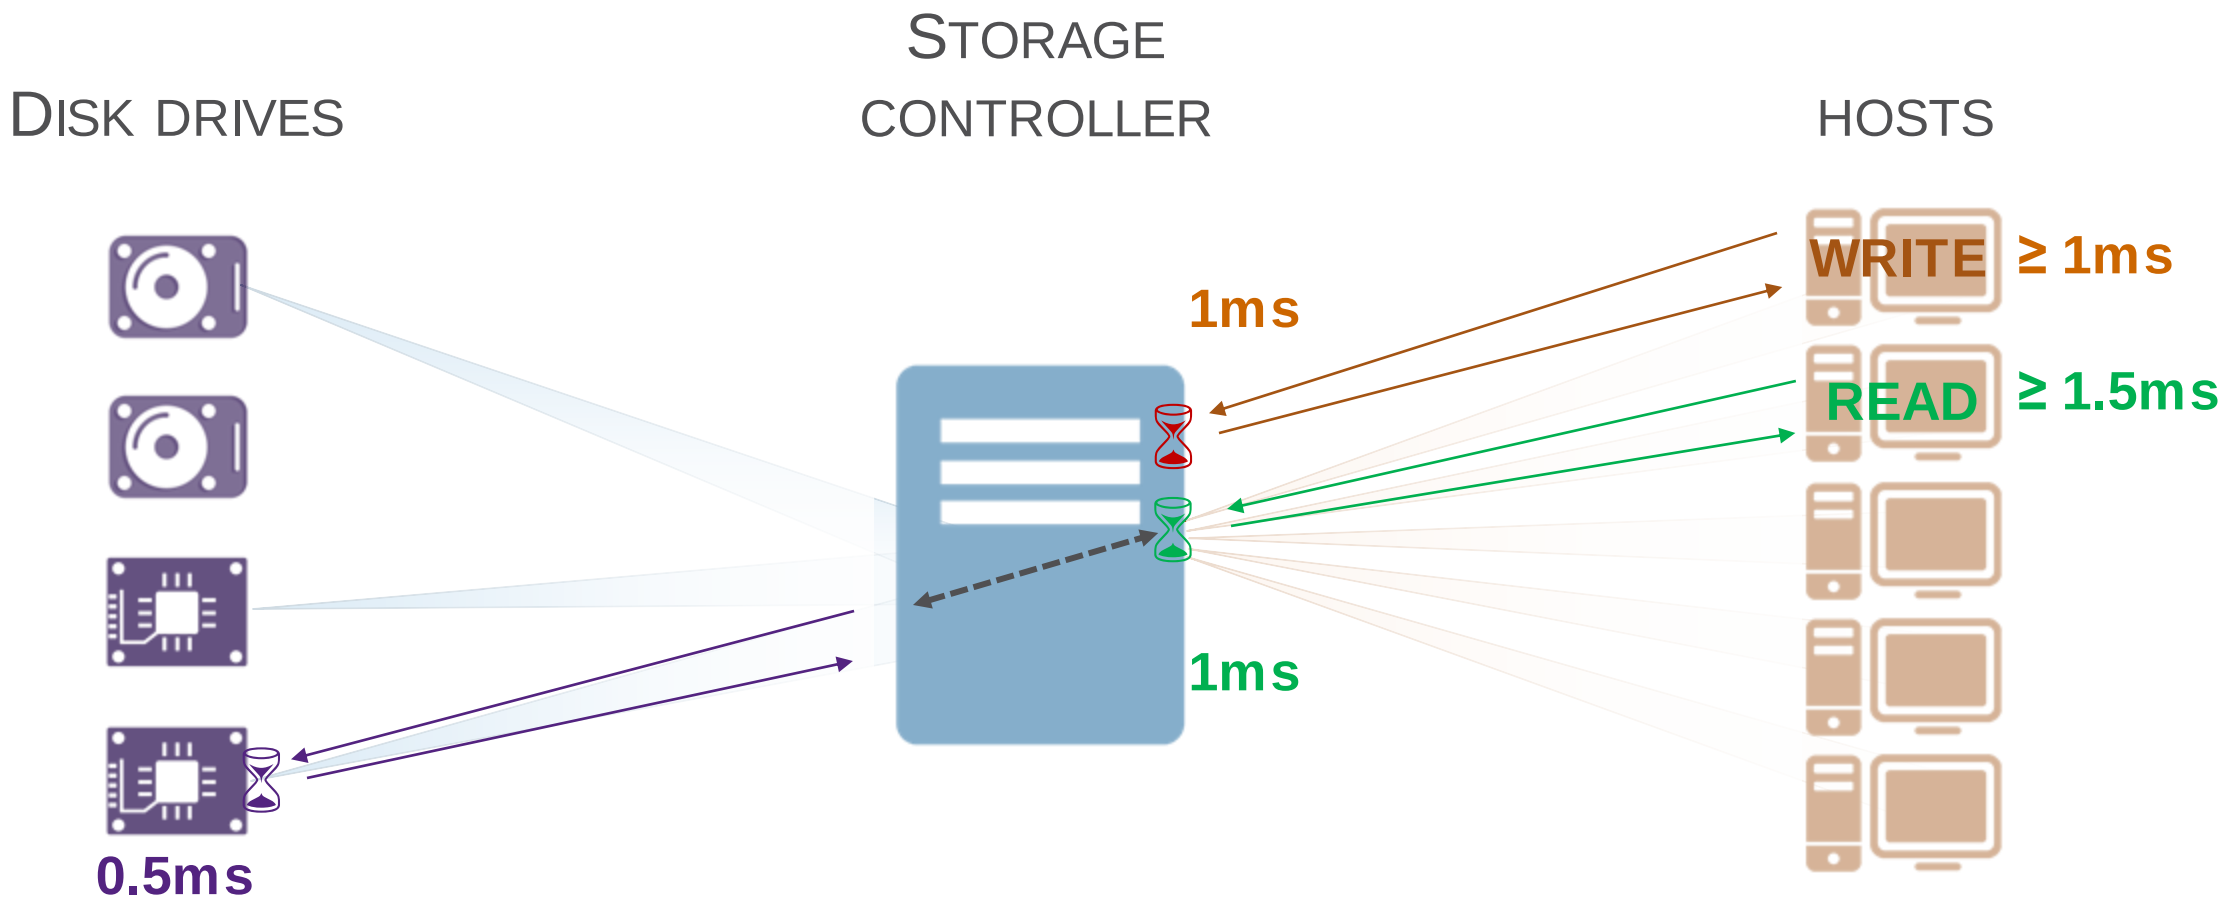

Aggregate Property Of Response Time

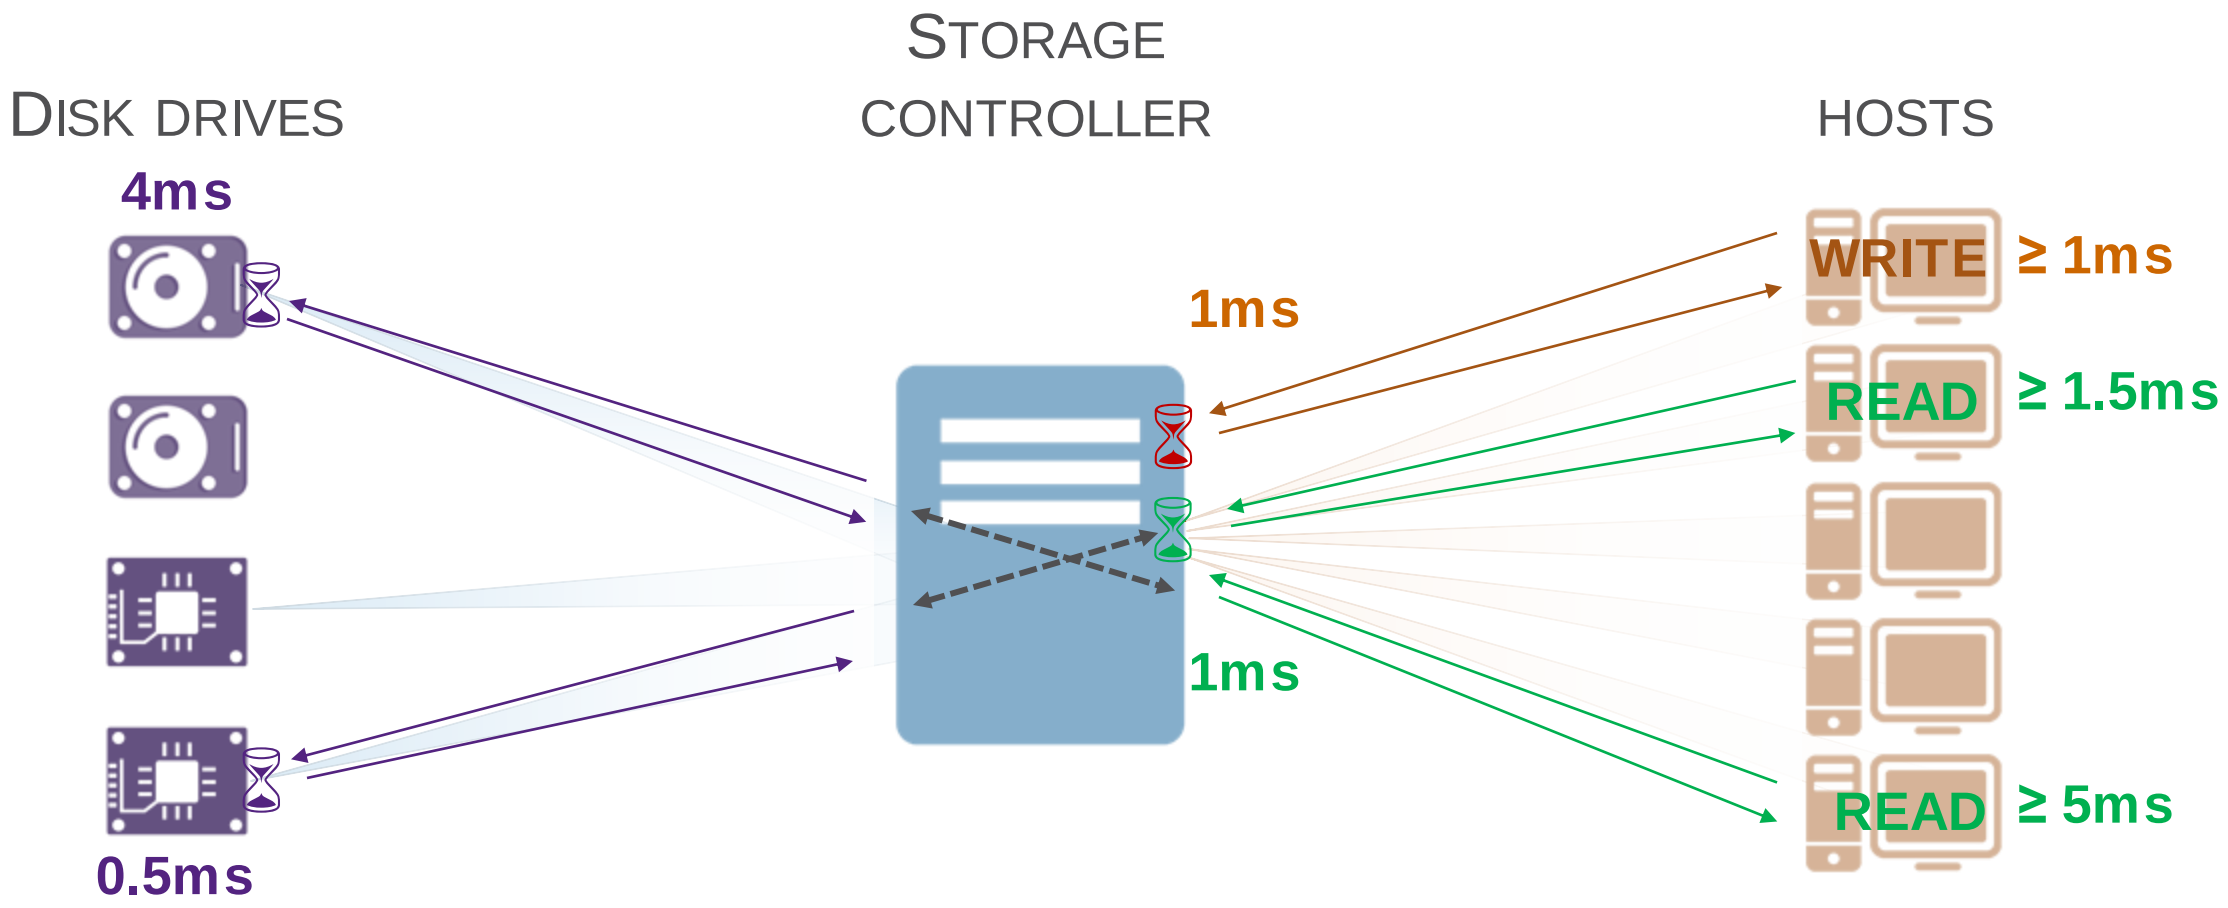

File / NAS Terminology

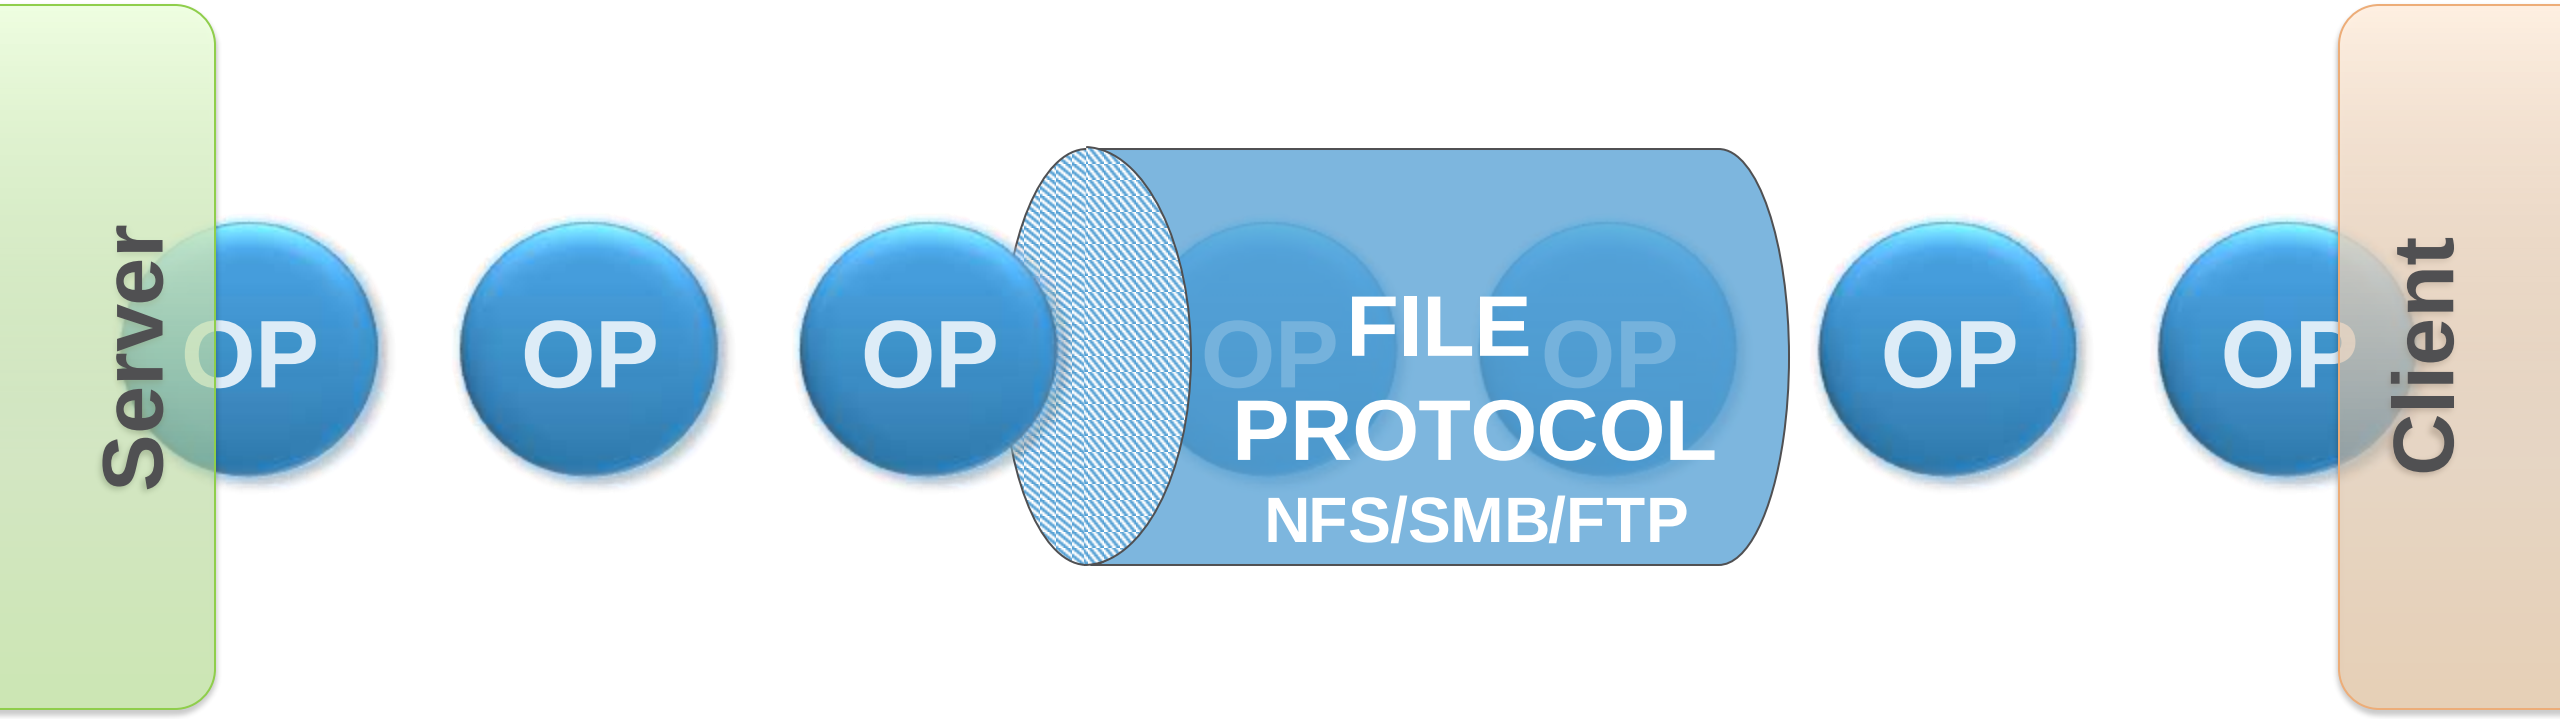

File / NAS Terminology

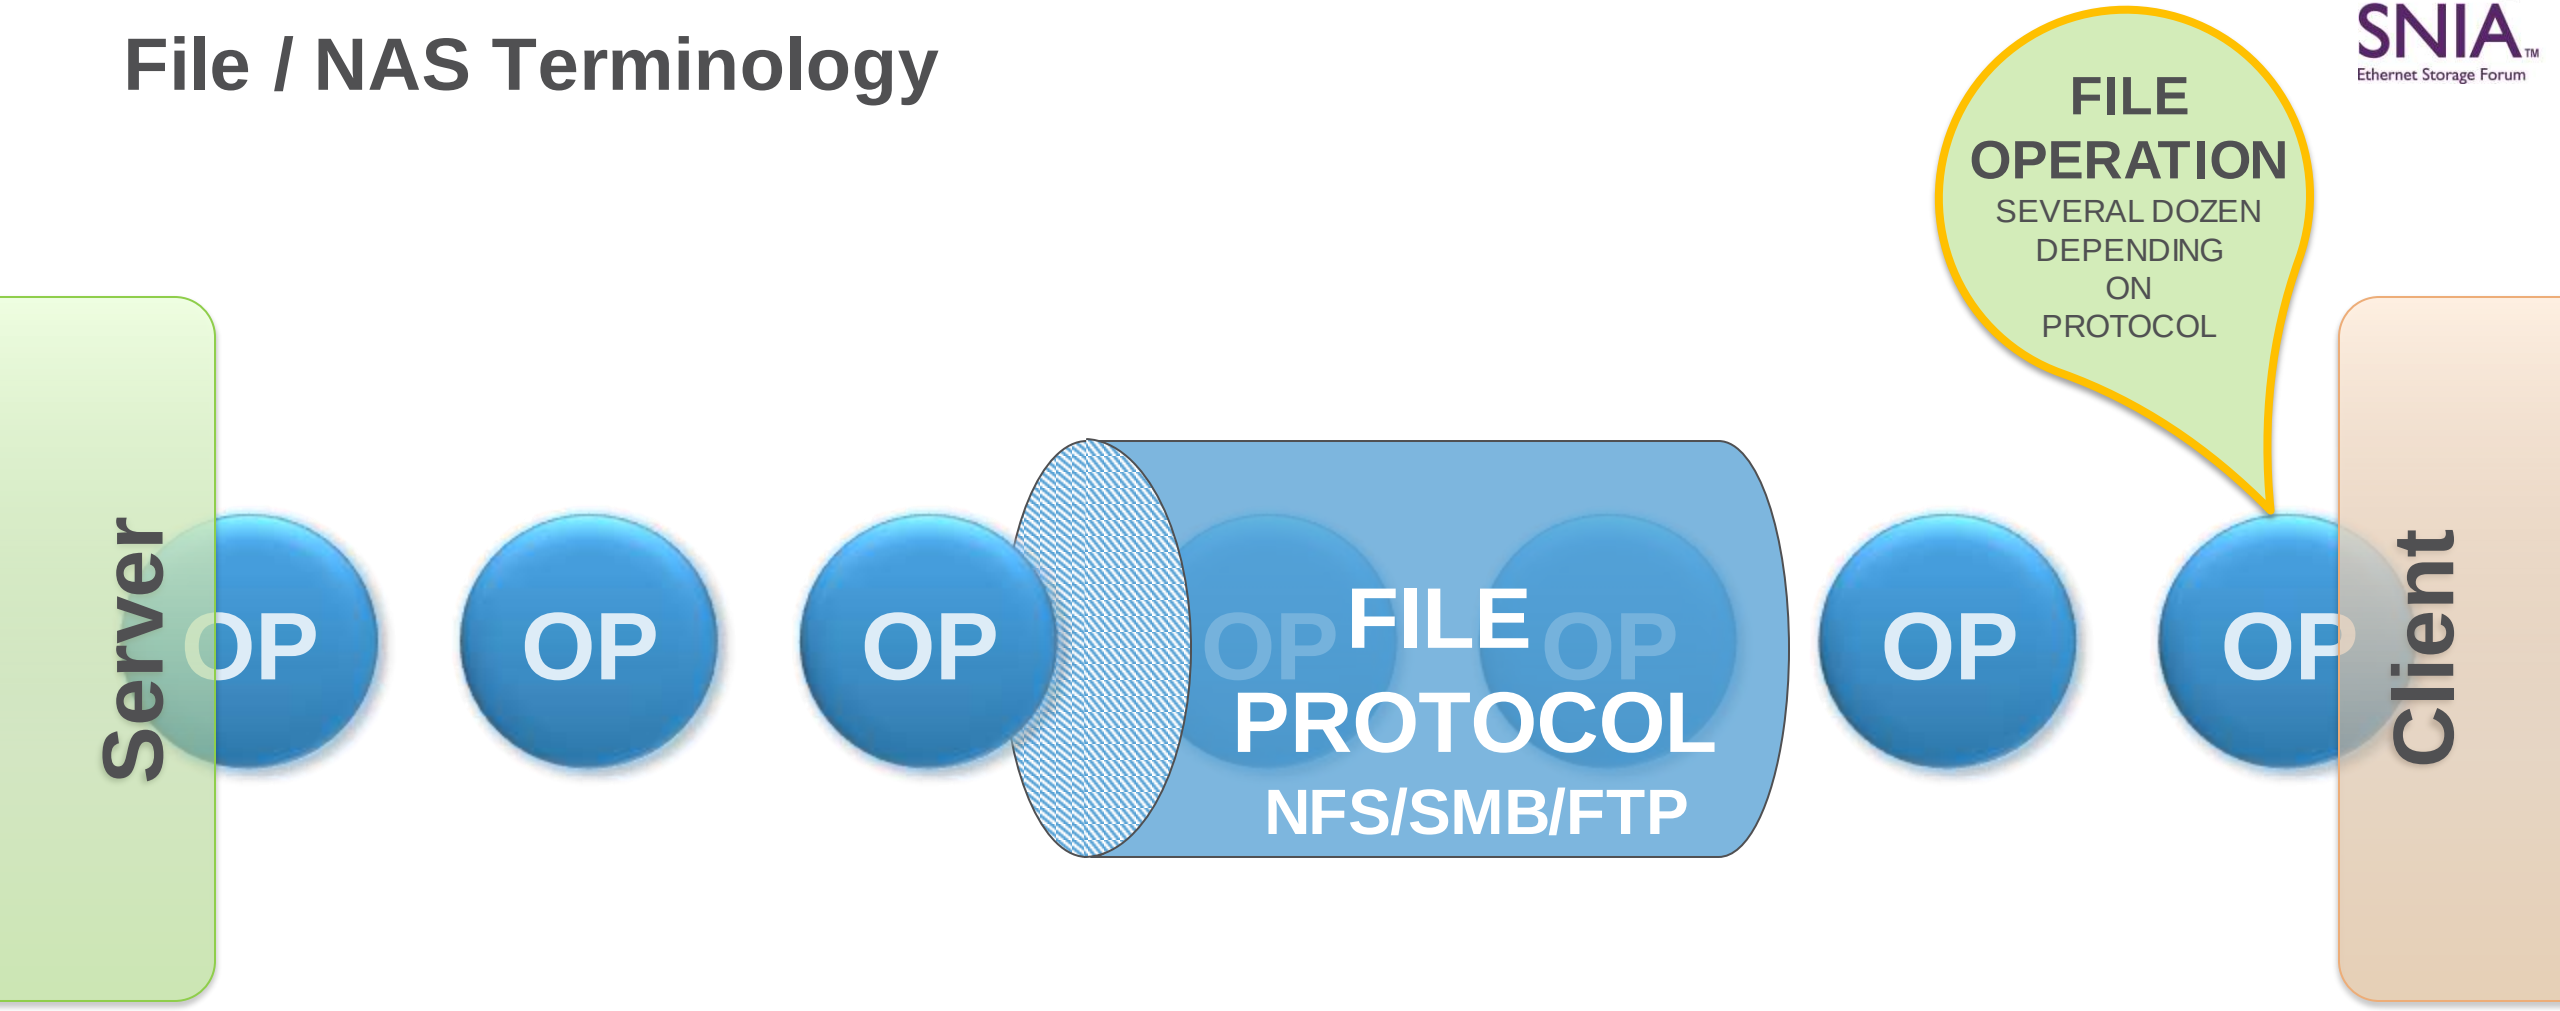

NFSv4 Operations

Source: https://www.ietf.org/rfc/rfc3530.txt
Graphic: http://www.wordle.net/create

SMB3 Operations

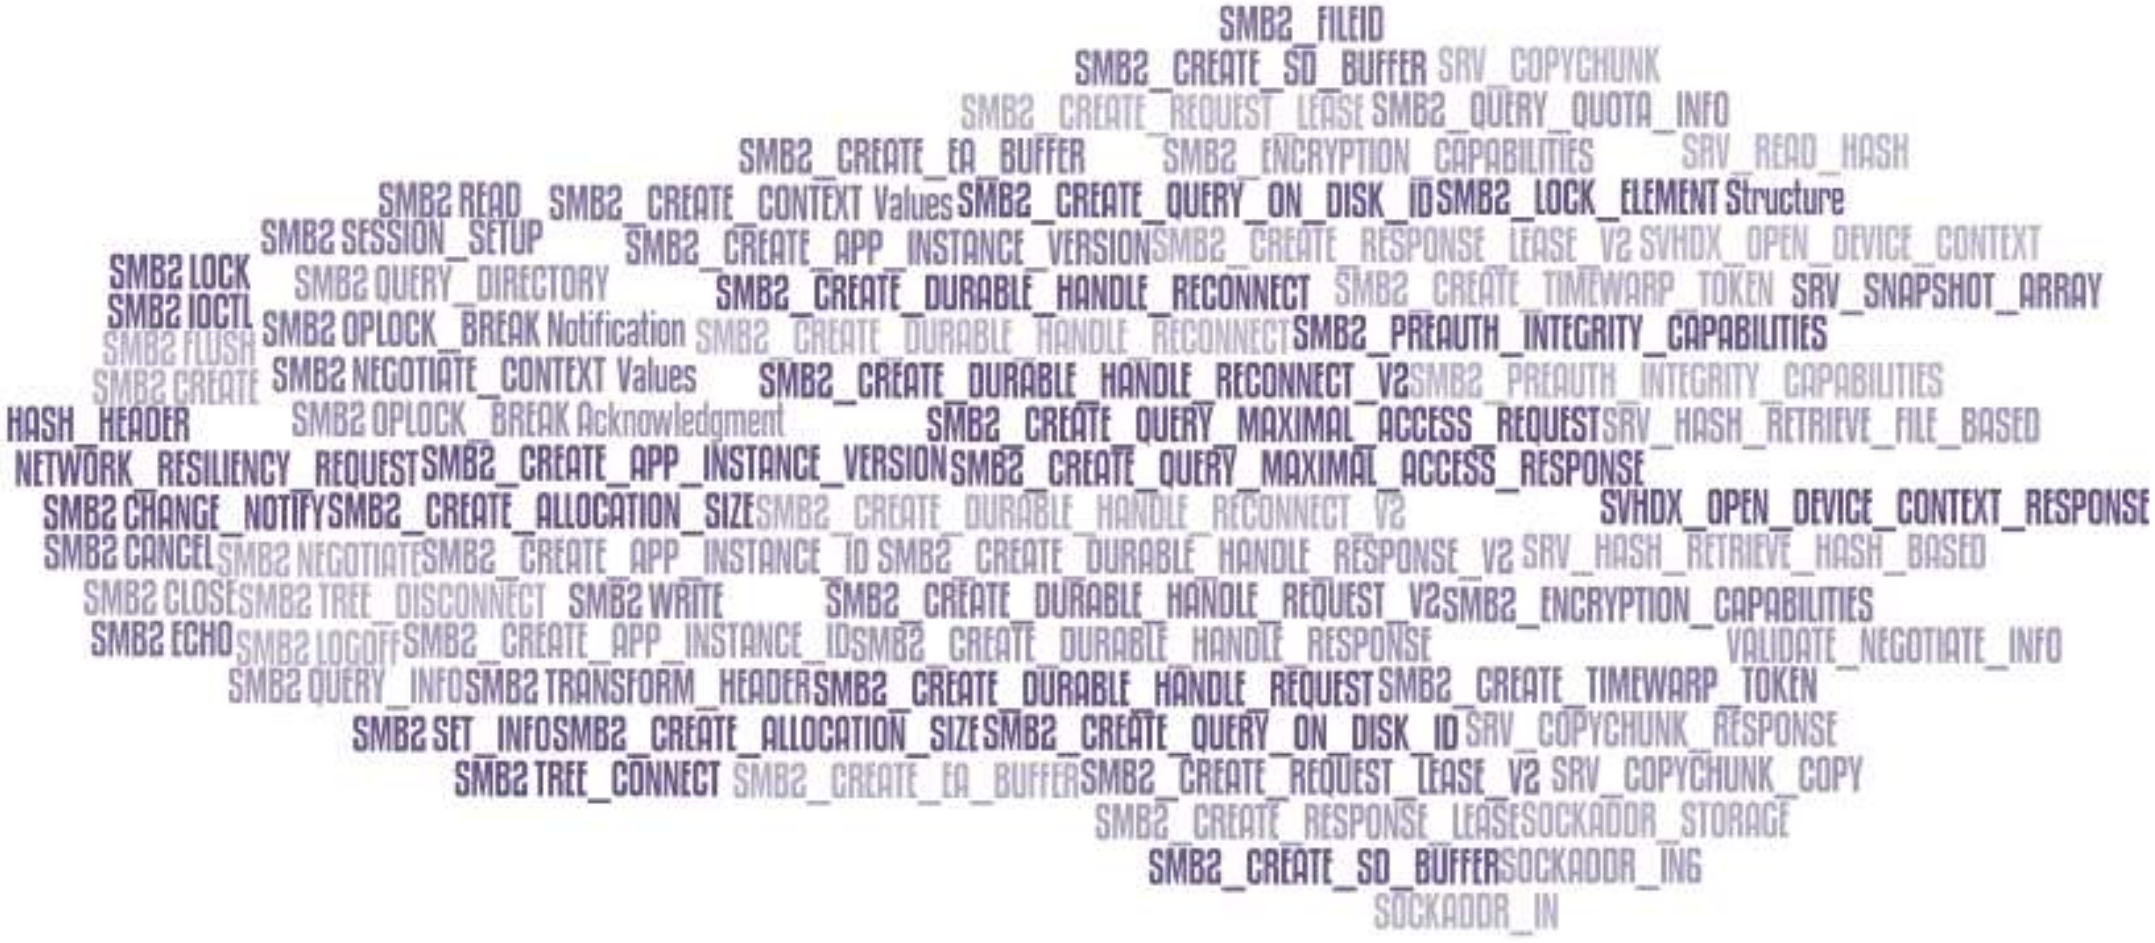

Source: https://msdn.microsoft.com/en-us/library/cc246482.aspx (use [MS-SMB2].pdf)
Graphic: http://www.wordle.net/create

File / NAS Terminology

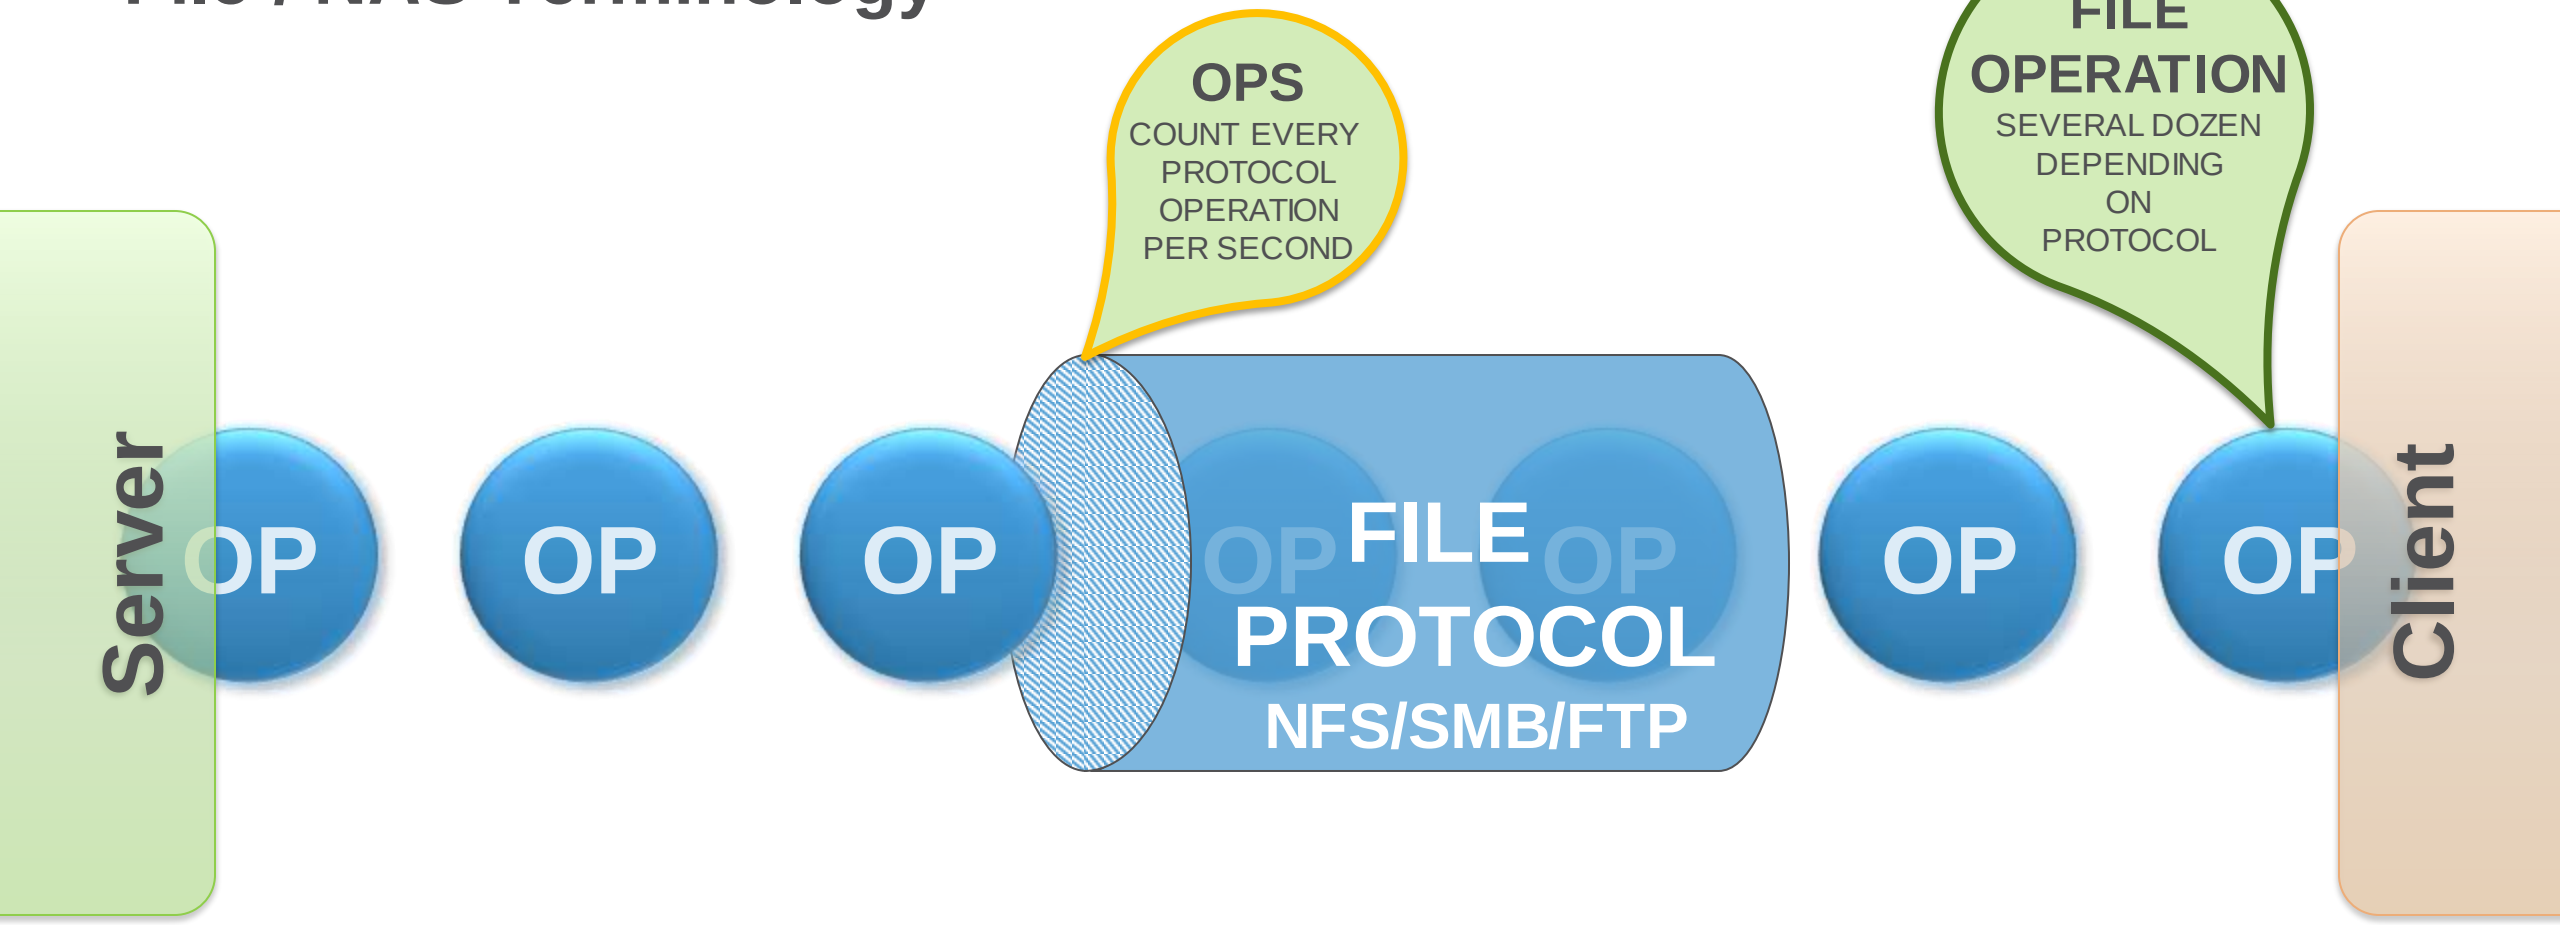

File / NAS Terminology

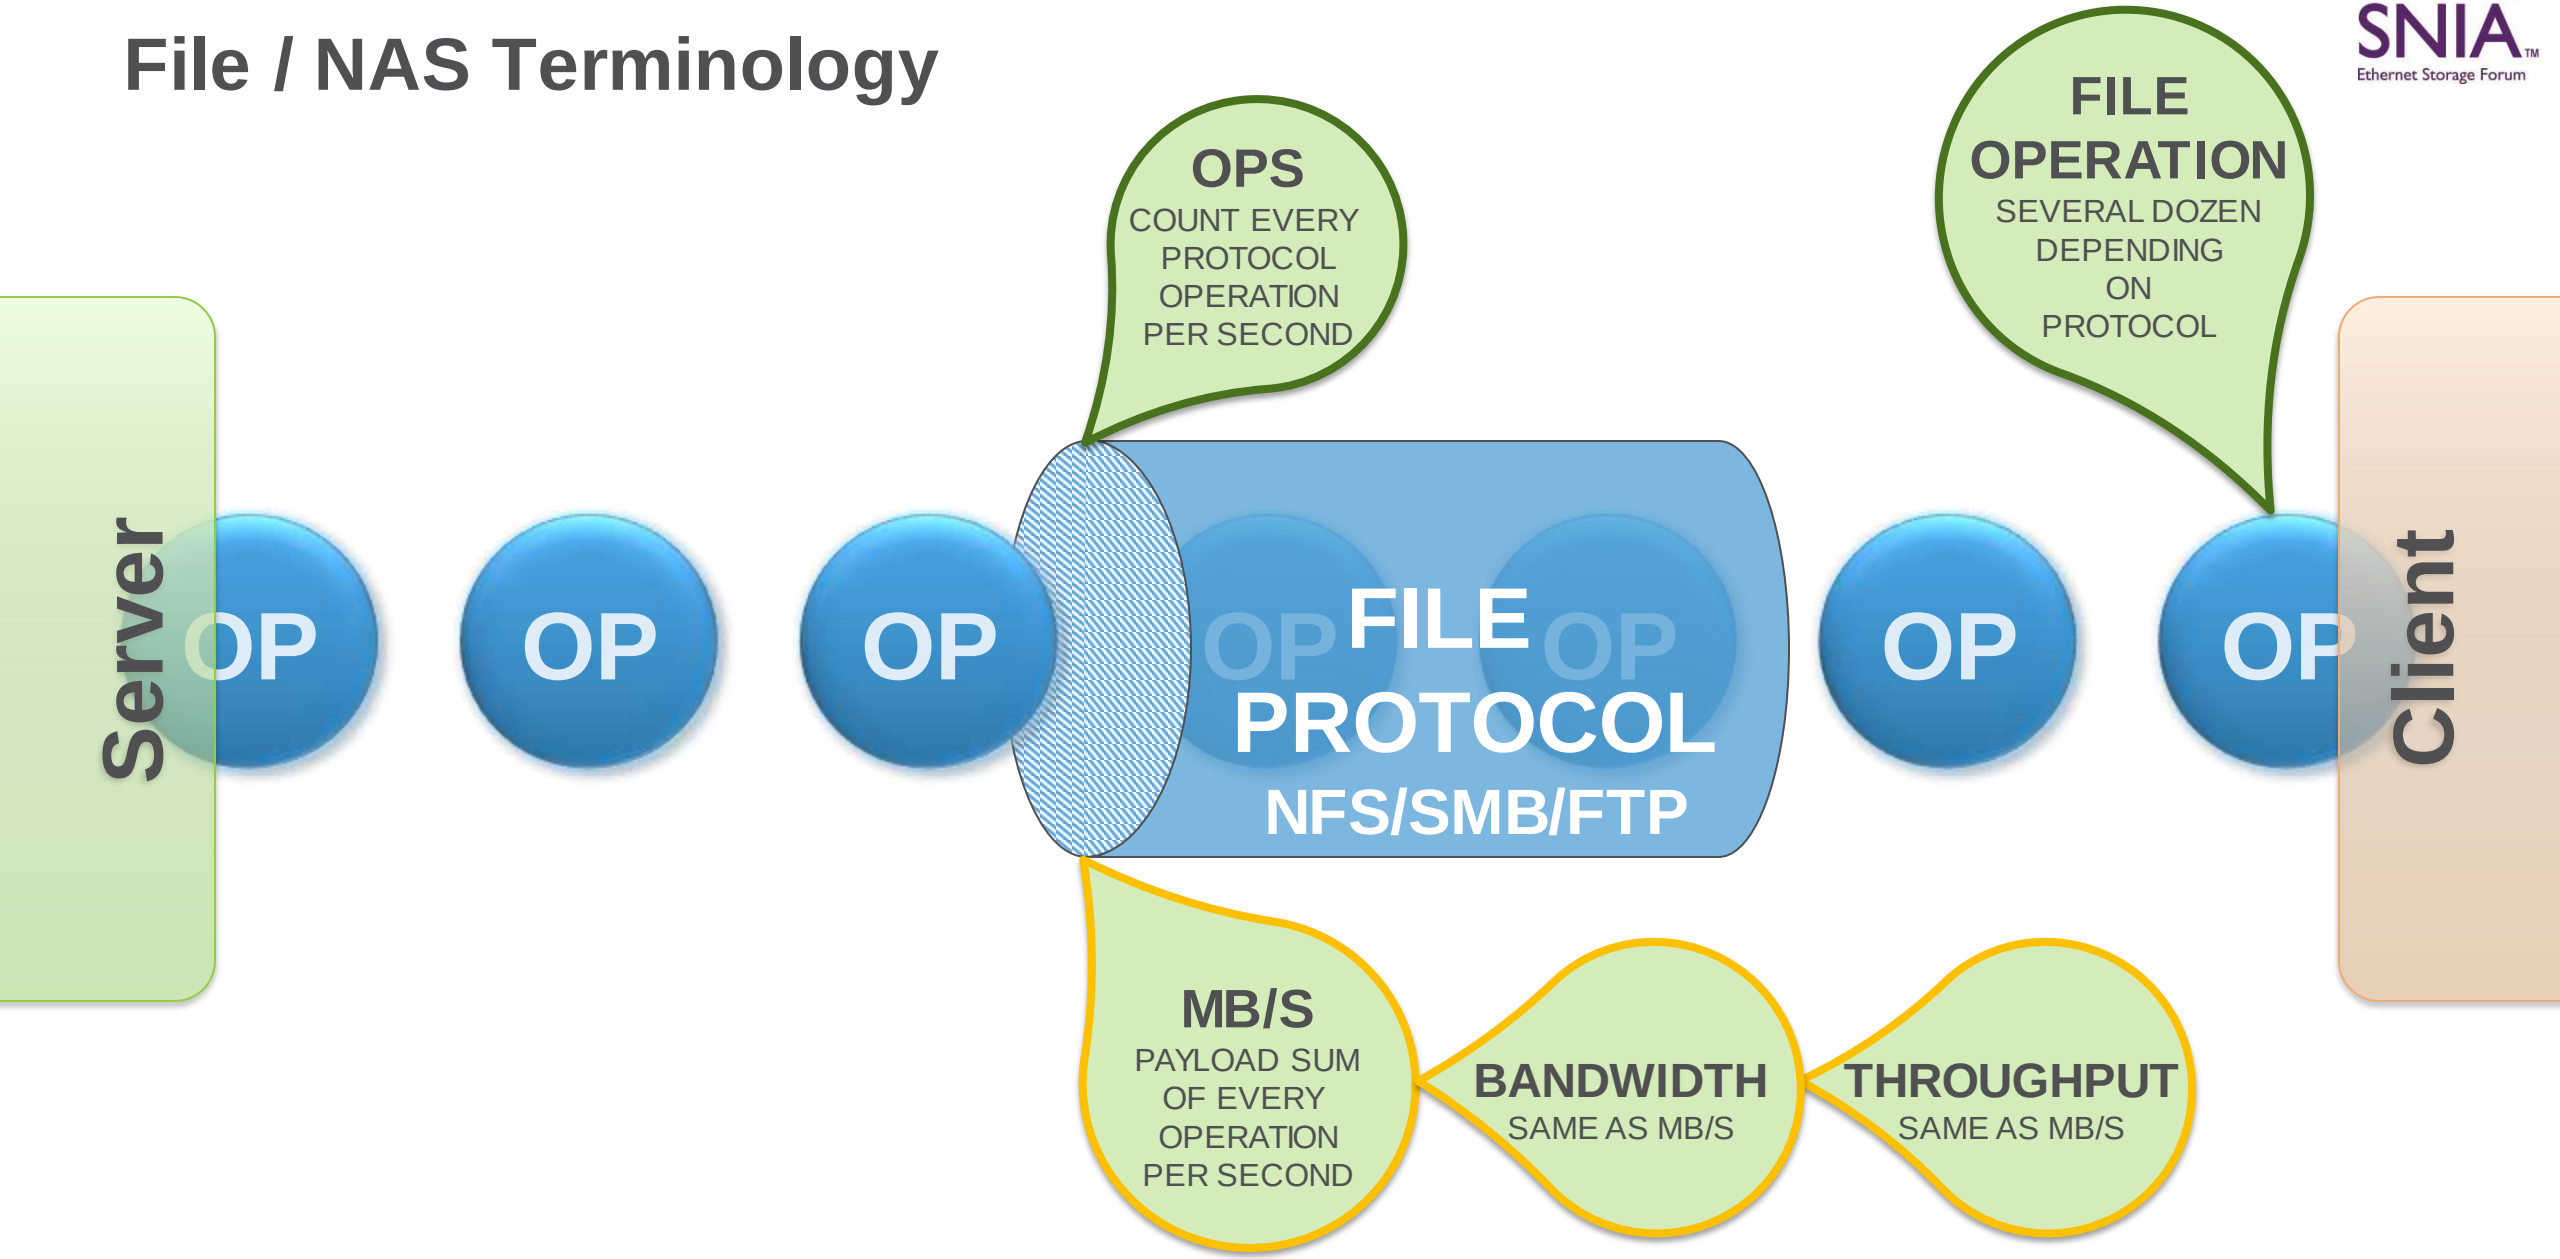

File / NAS Terminology

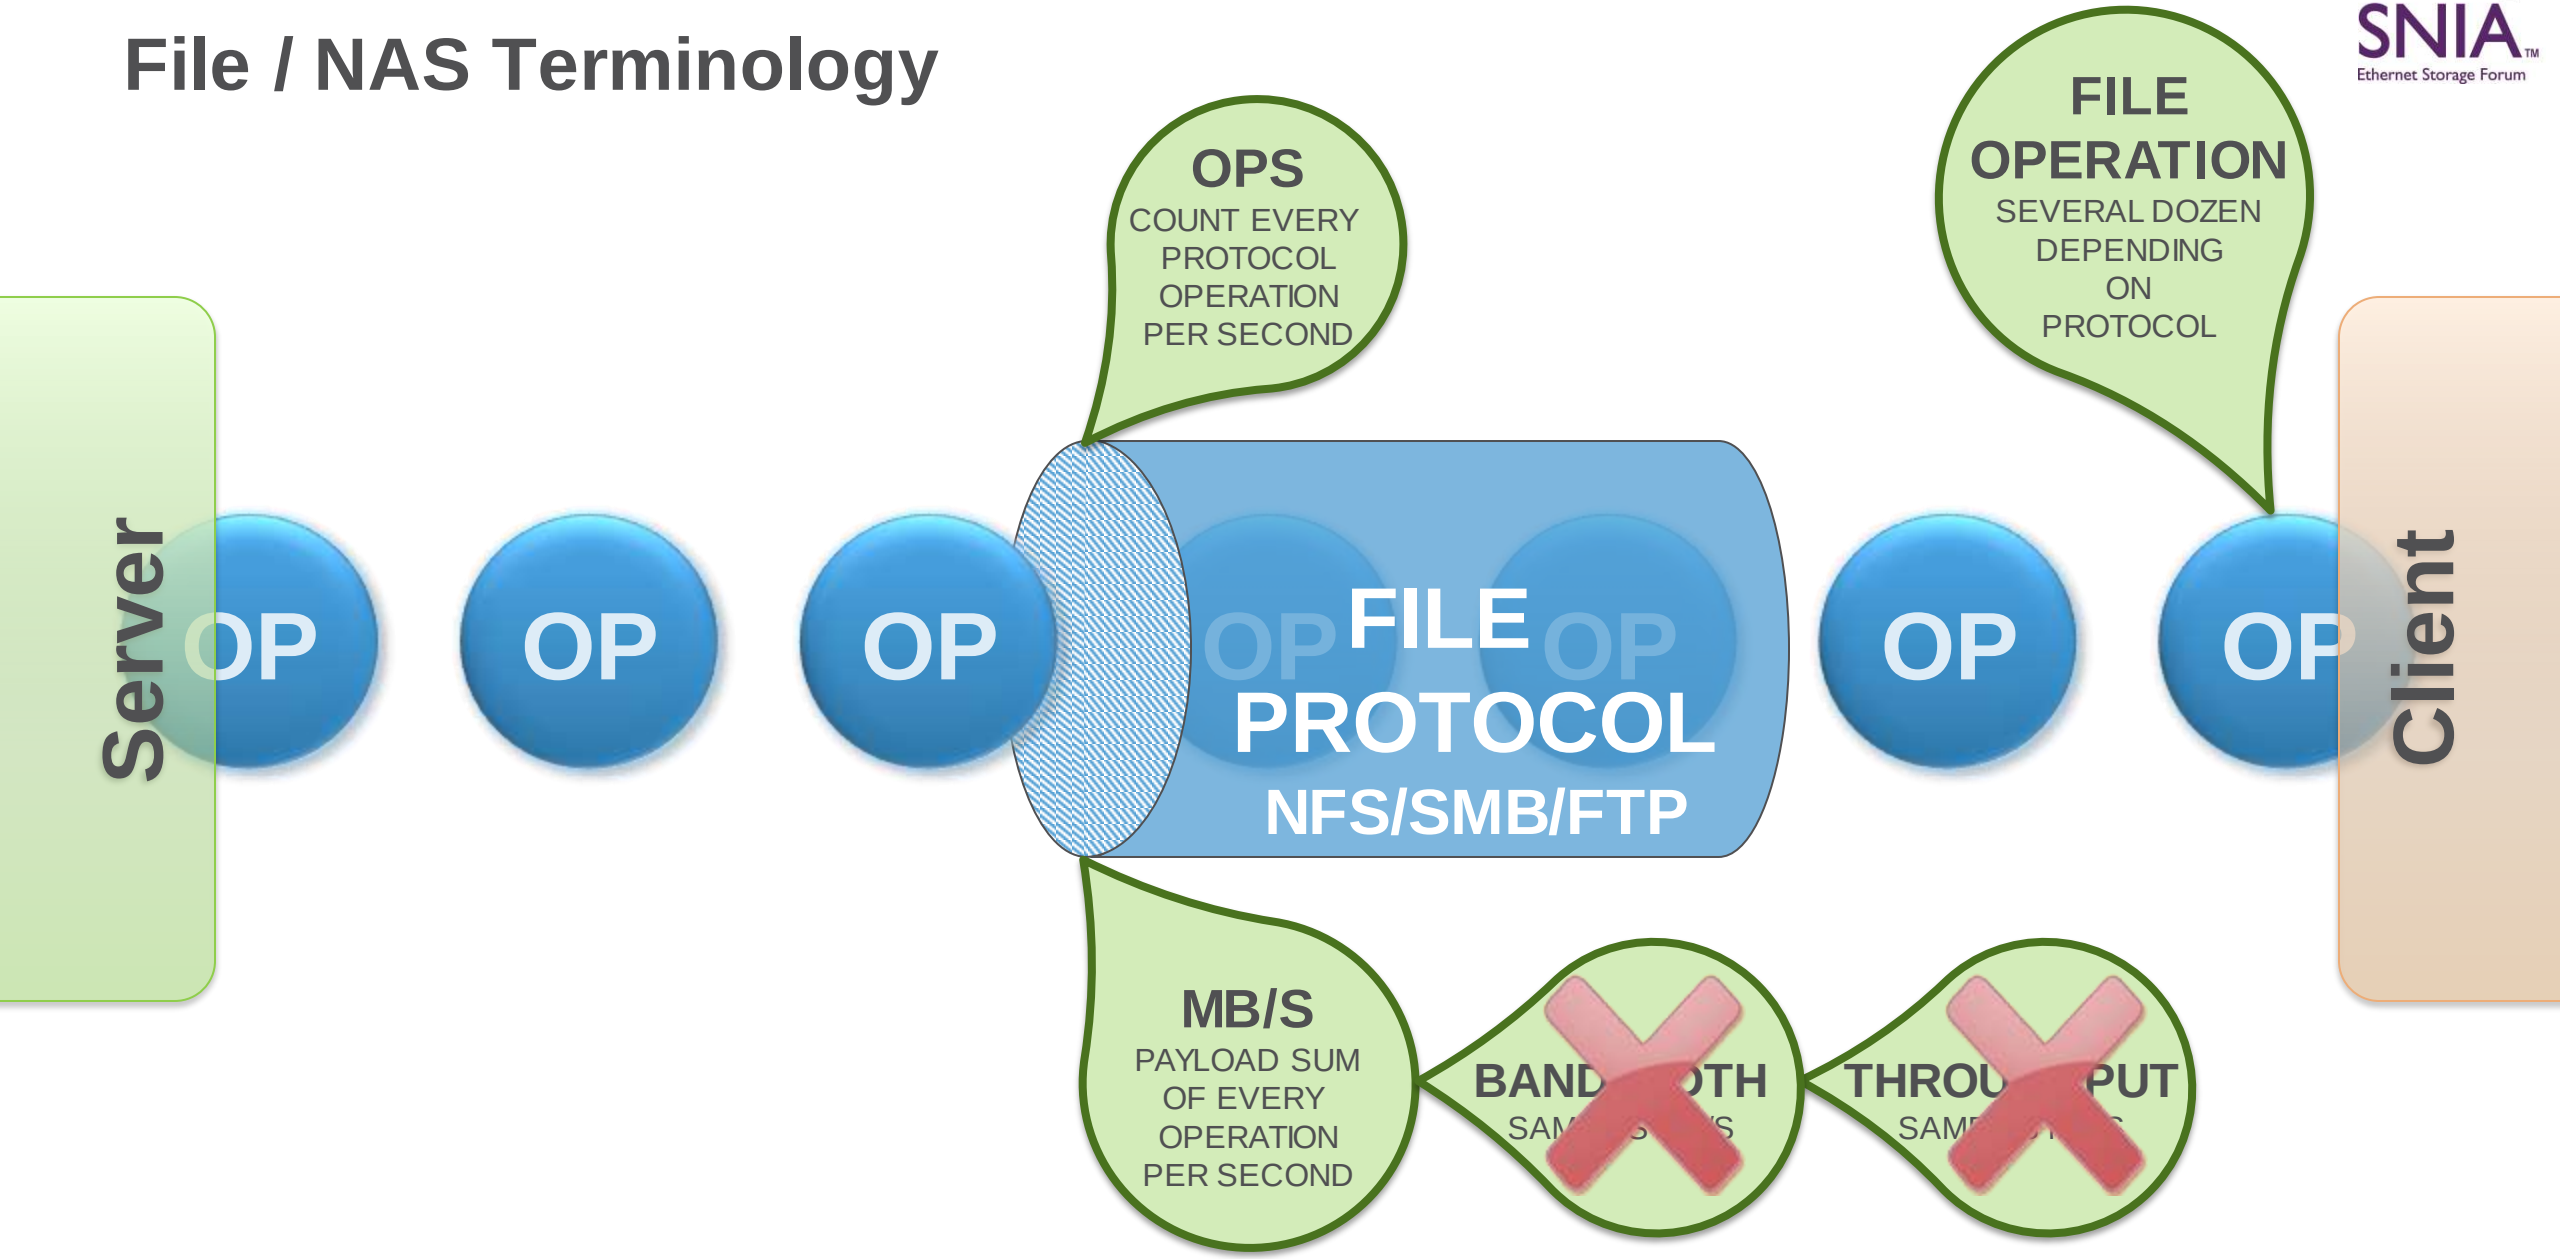

File / NAS Terminology

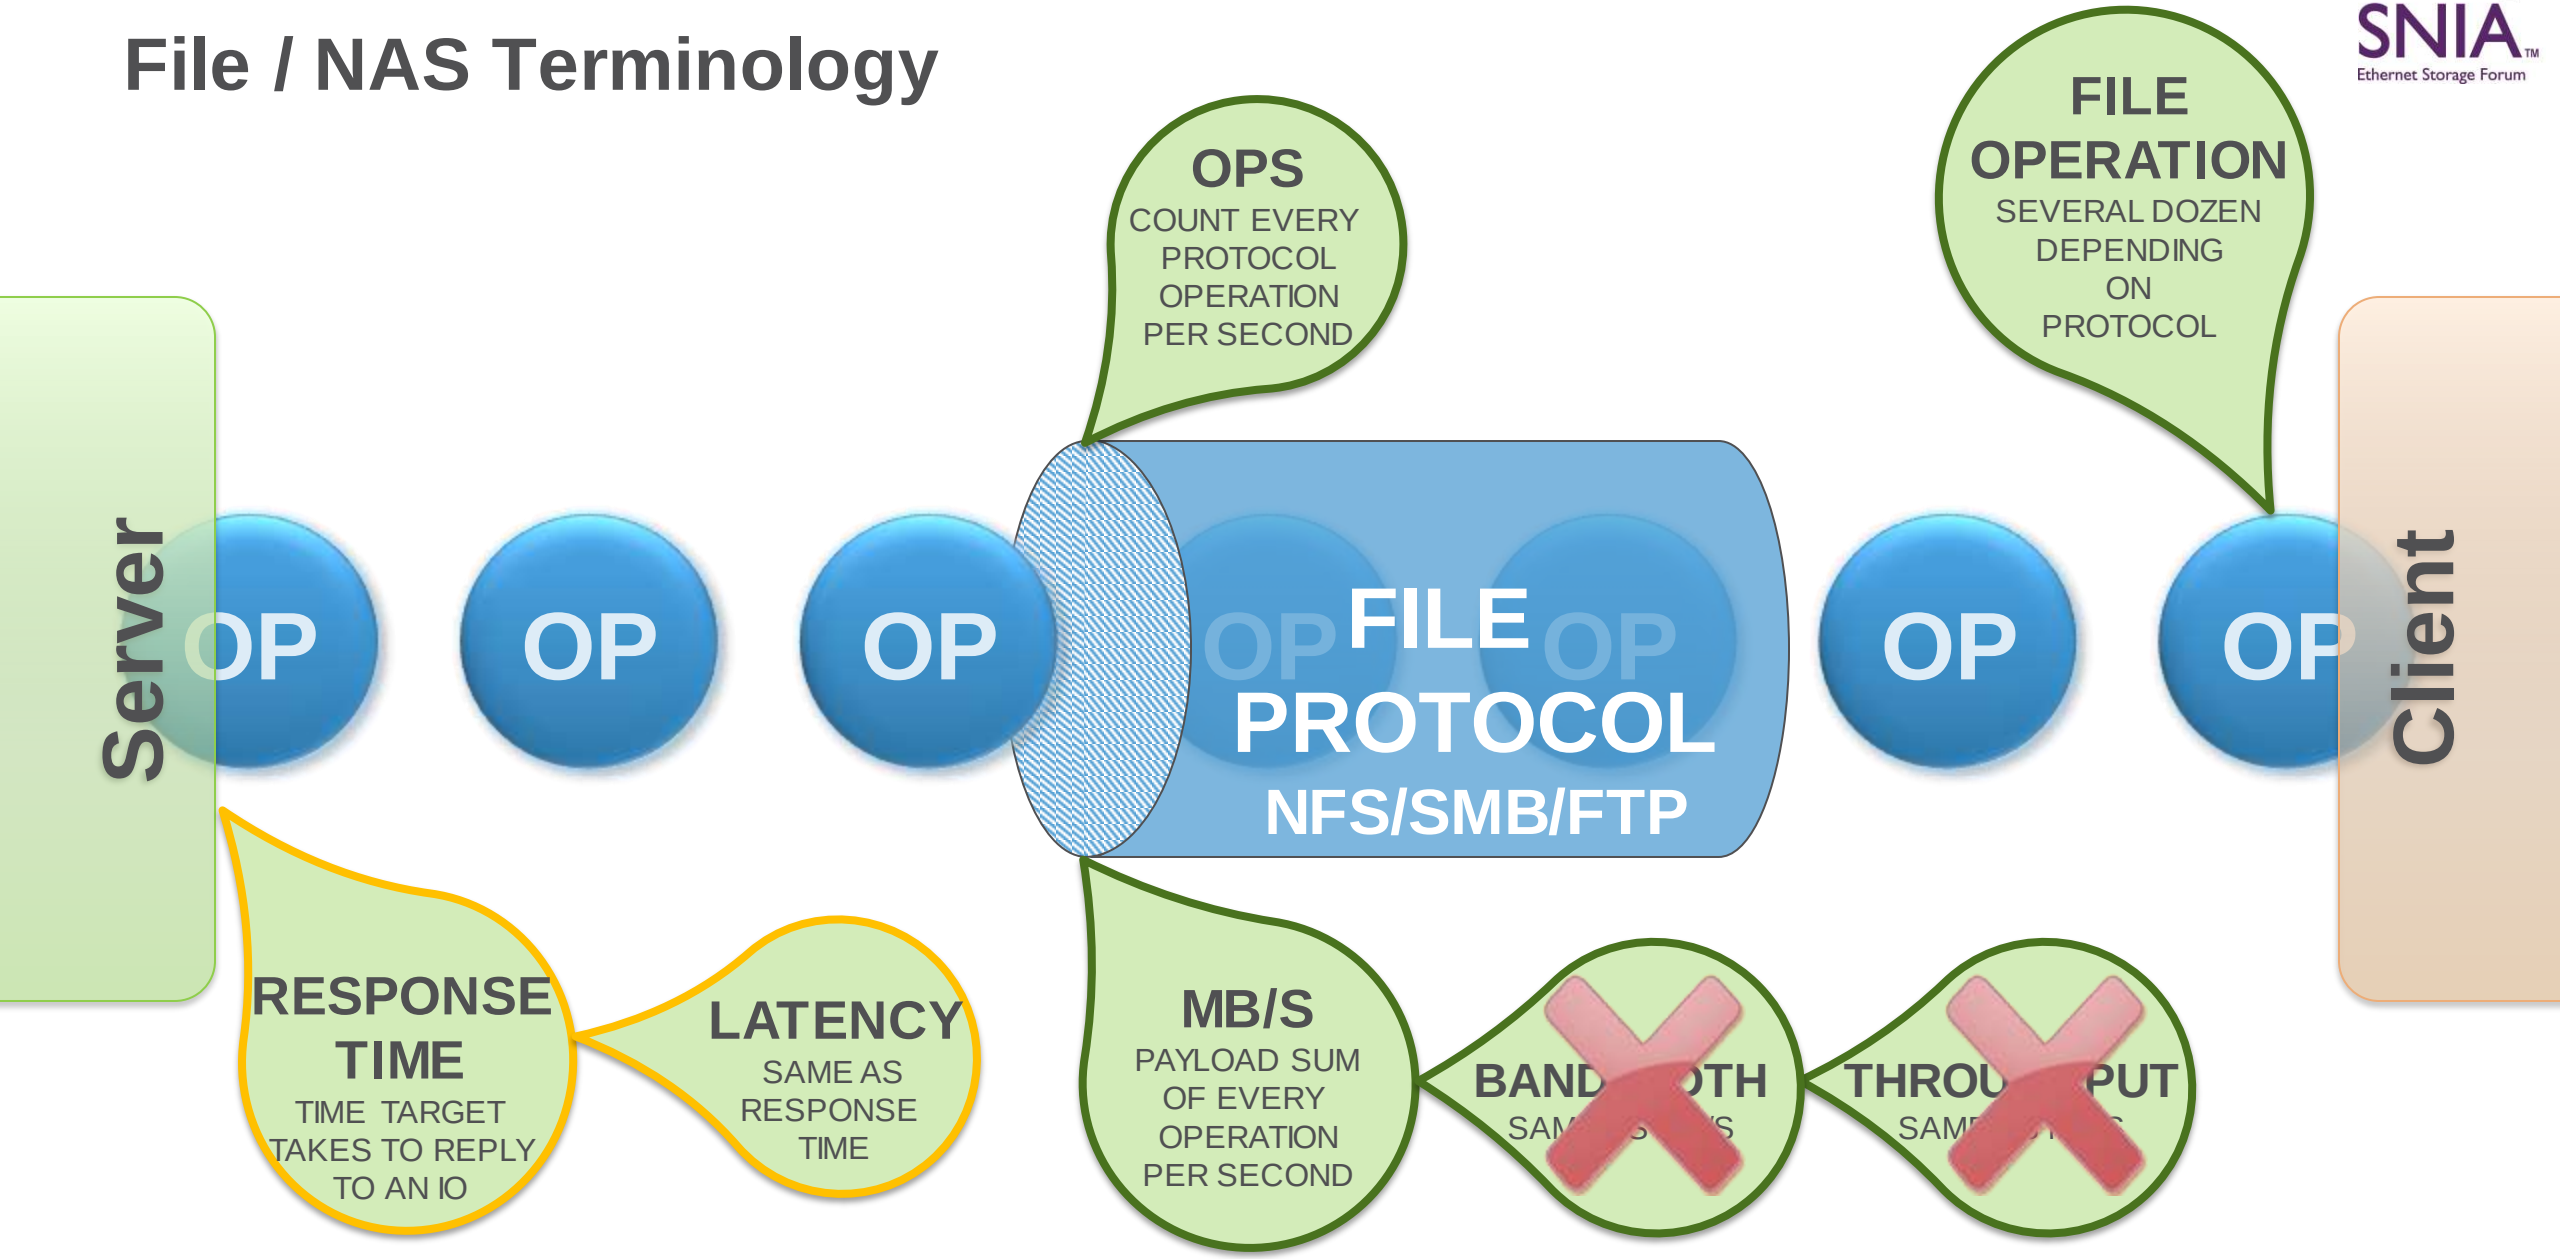

Difference Between OPS and MB/s (MiB/s)

- Two Otherwise Identical Runs With Different NFS Operations Mix (64 KiB IO Size)

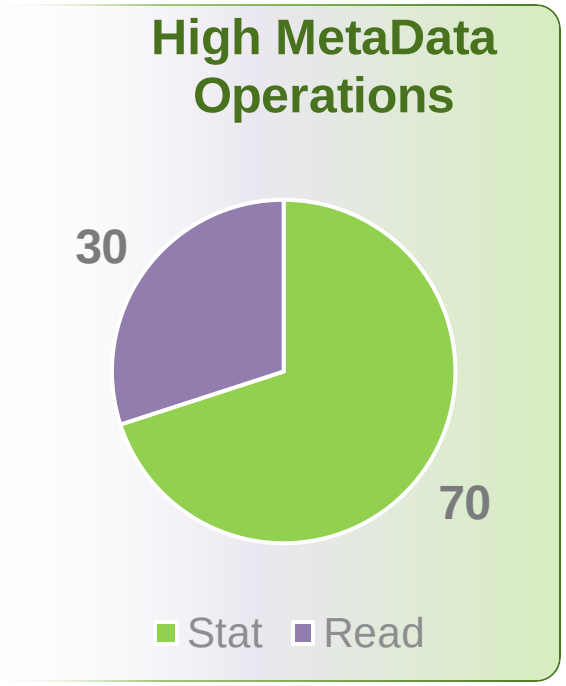

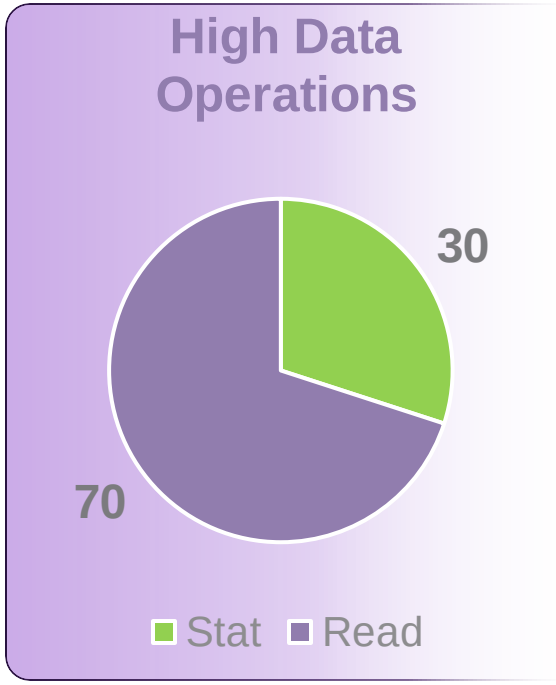

Graph Fun Or Why “It Depends” Is Relevant?

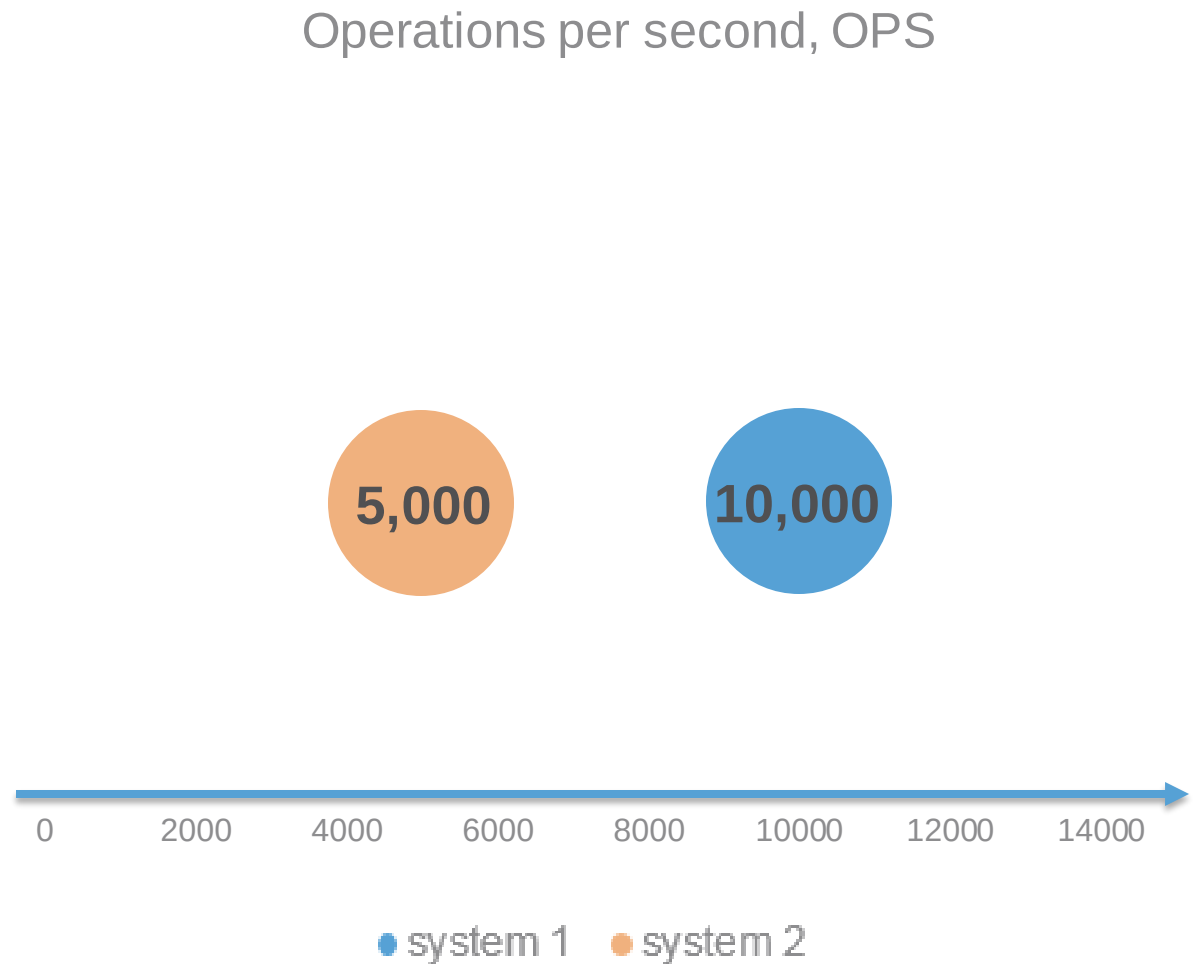

Graph Fun Or Why “It Depends” Is Relevant?

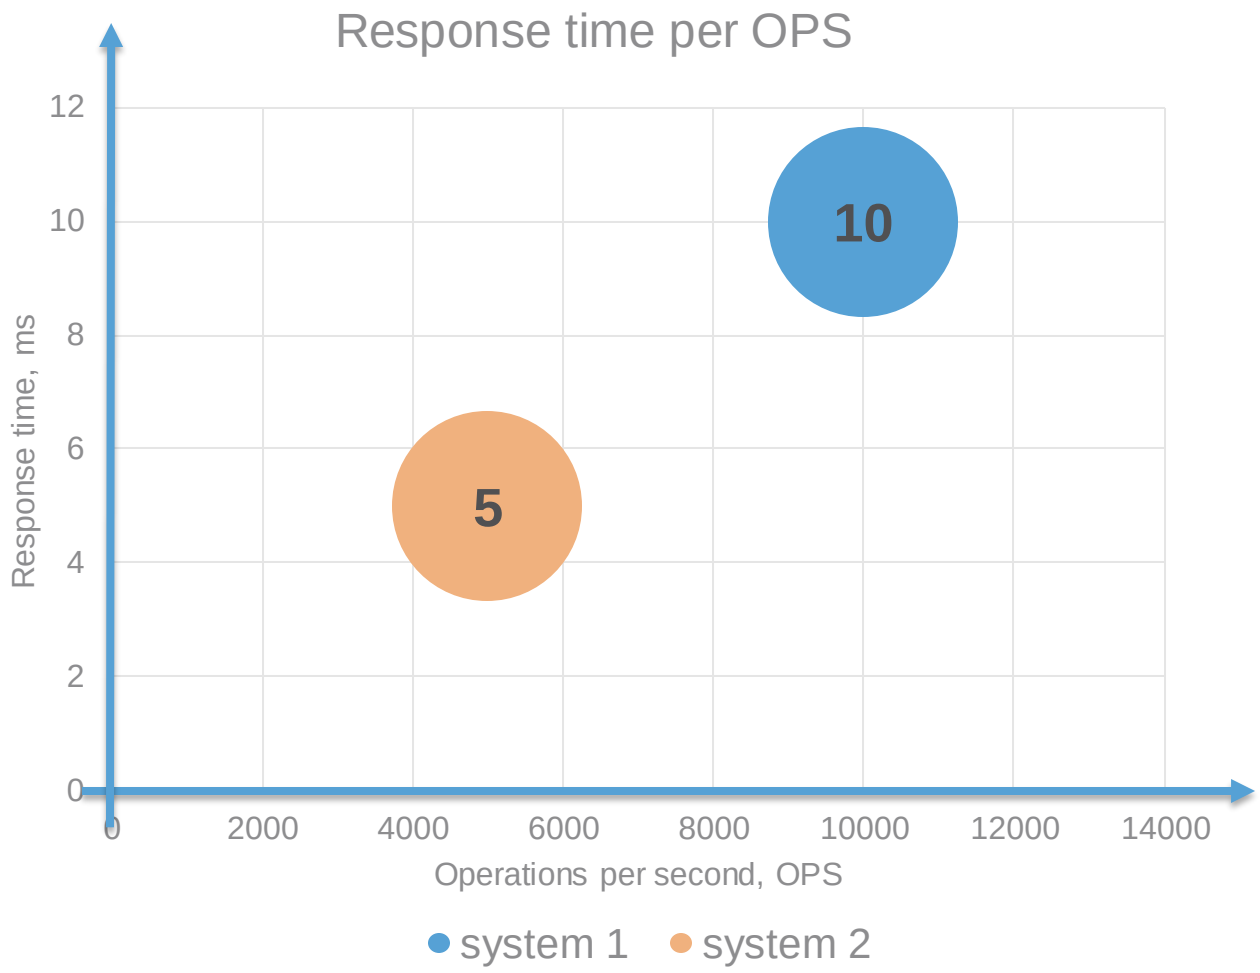

Graph Fun Or Why “It Depends” Is Relevant?

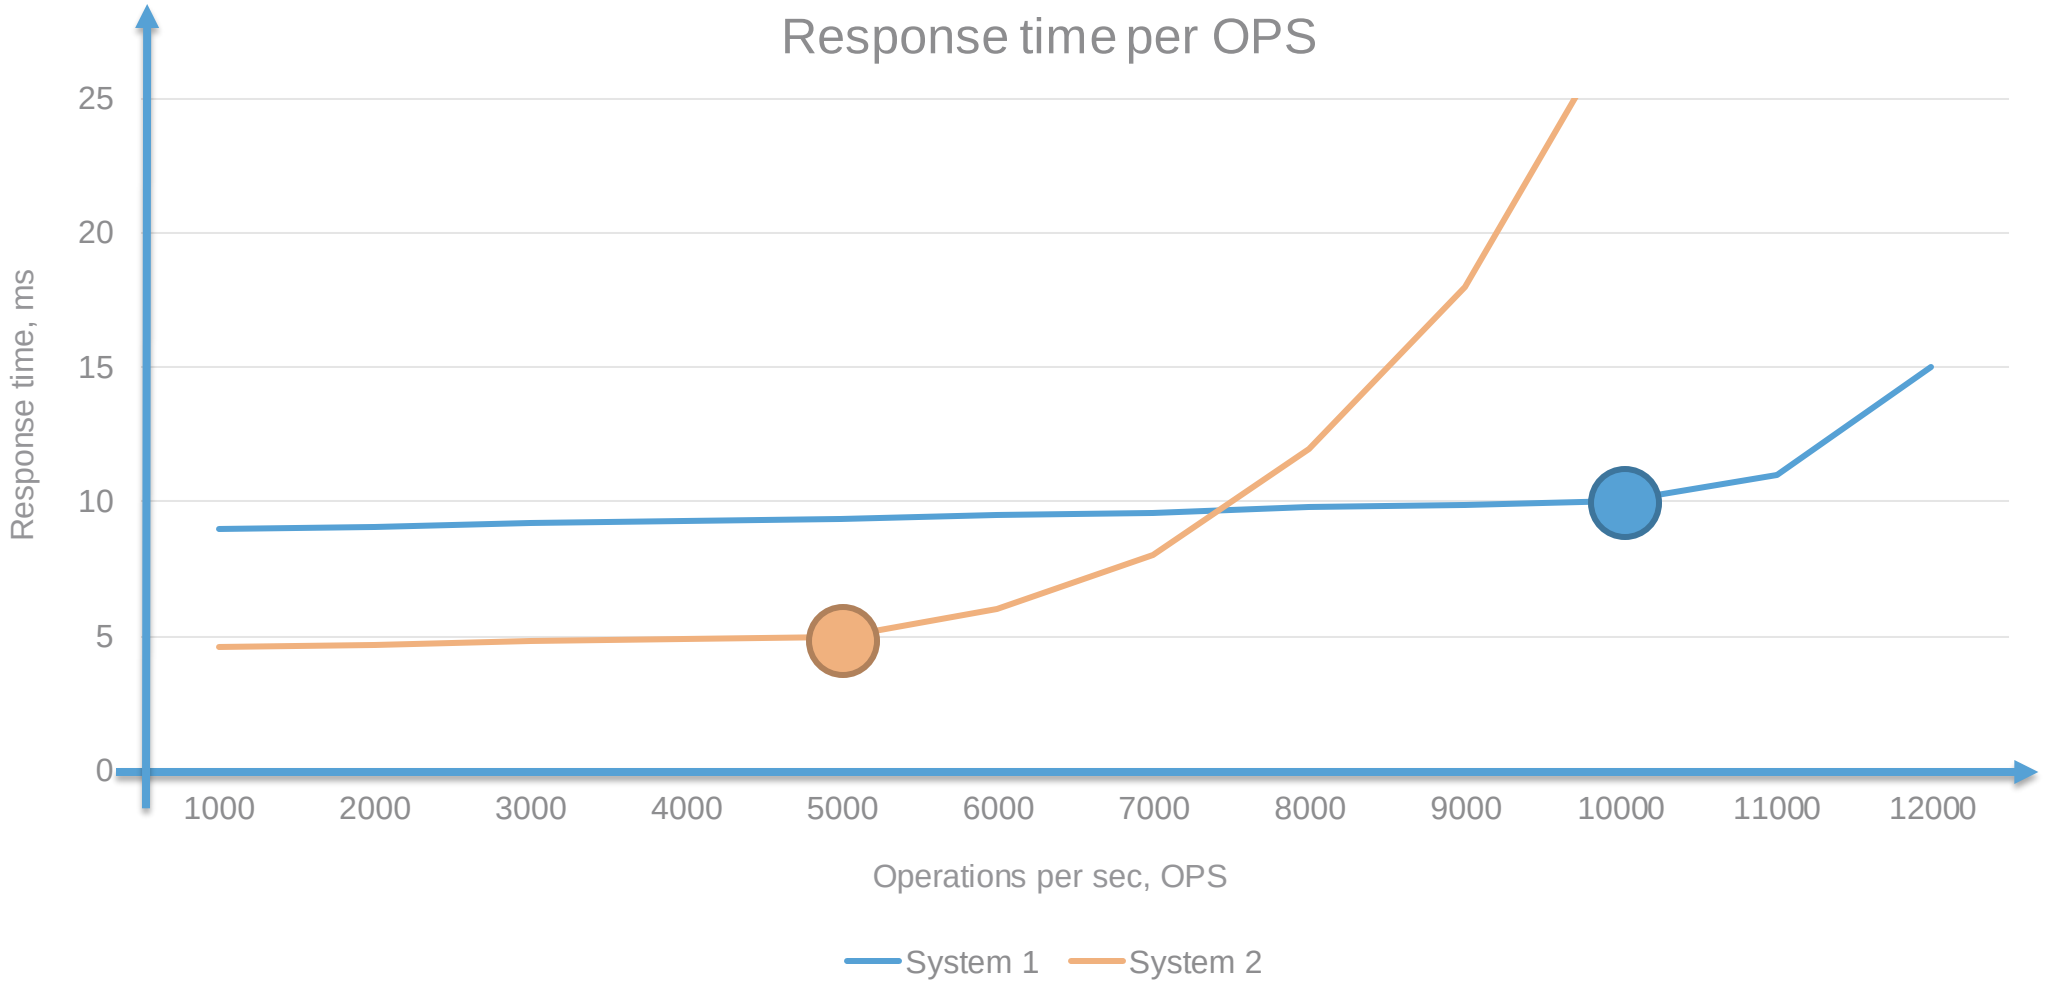

Graph Fun Or Why “It Depends” Is Relevant?

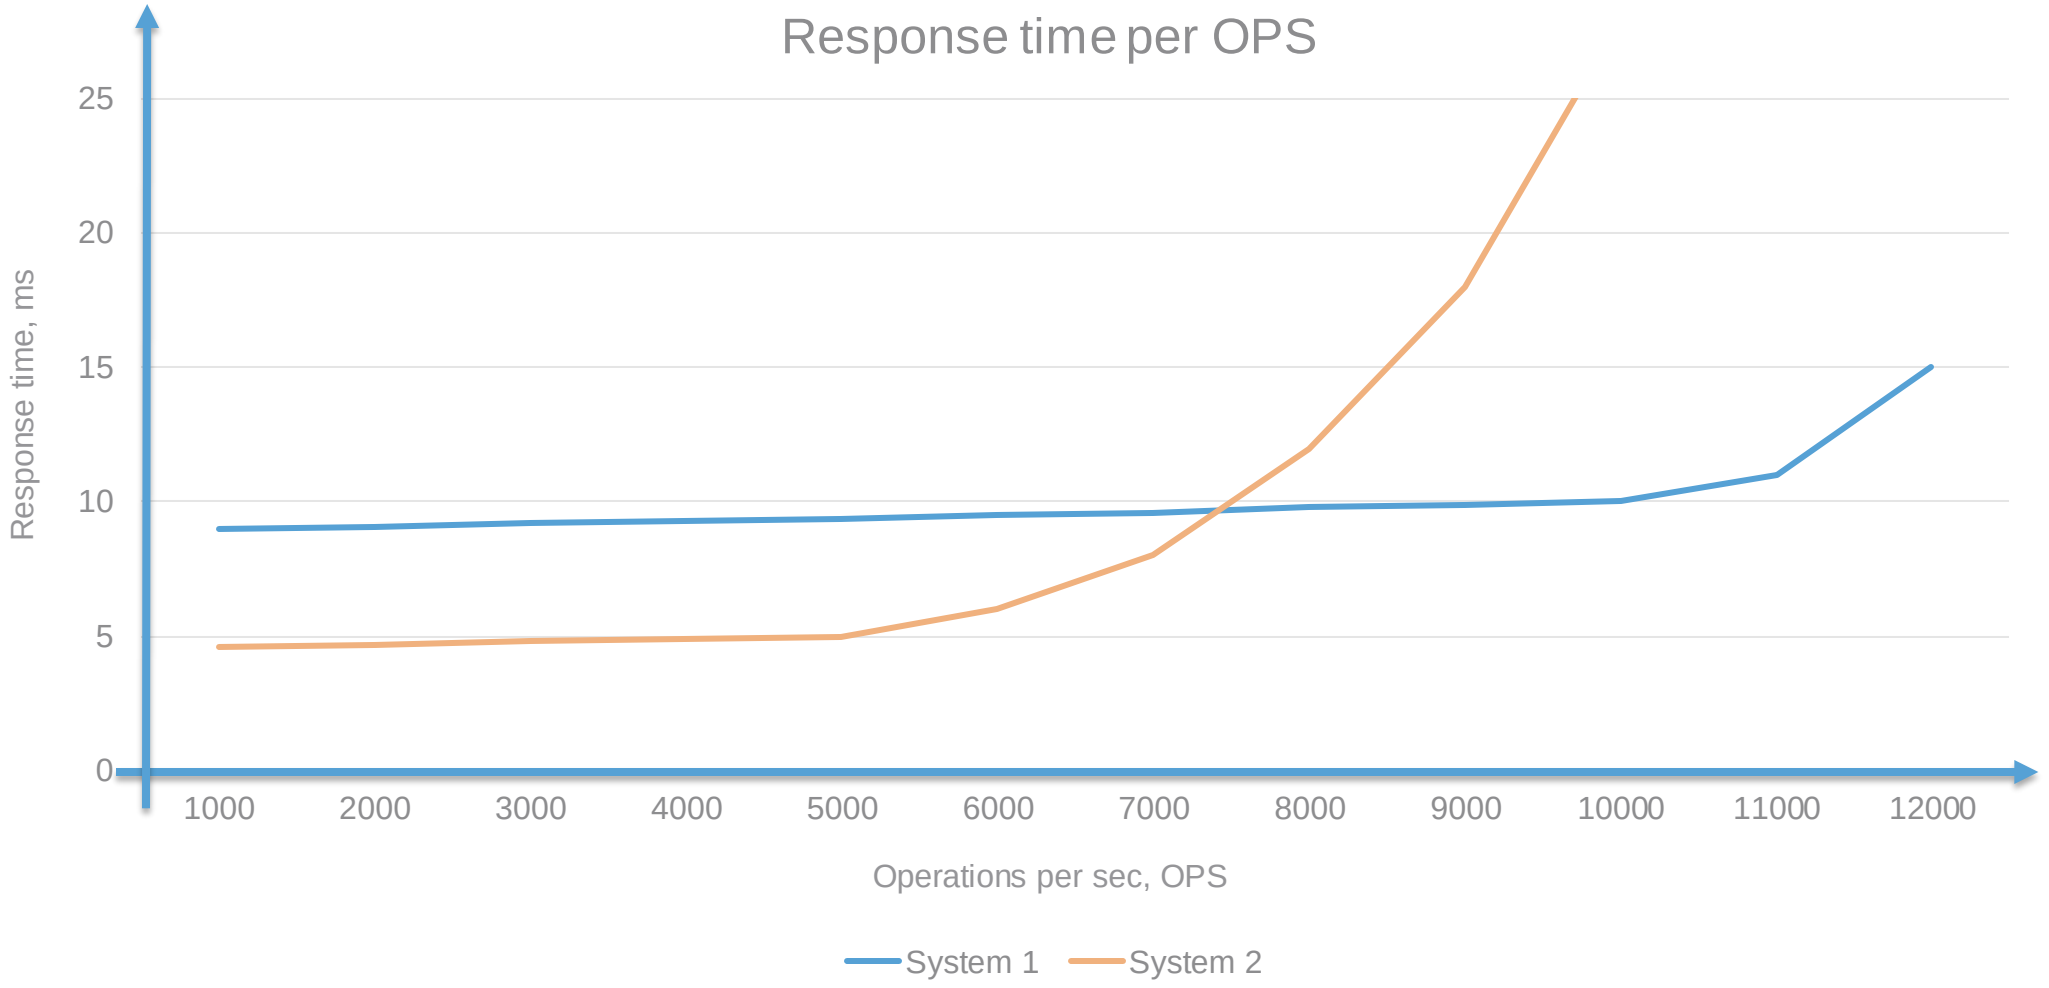

Graph Fun Or Why “It Depends” Is Relevant?

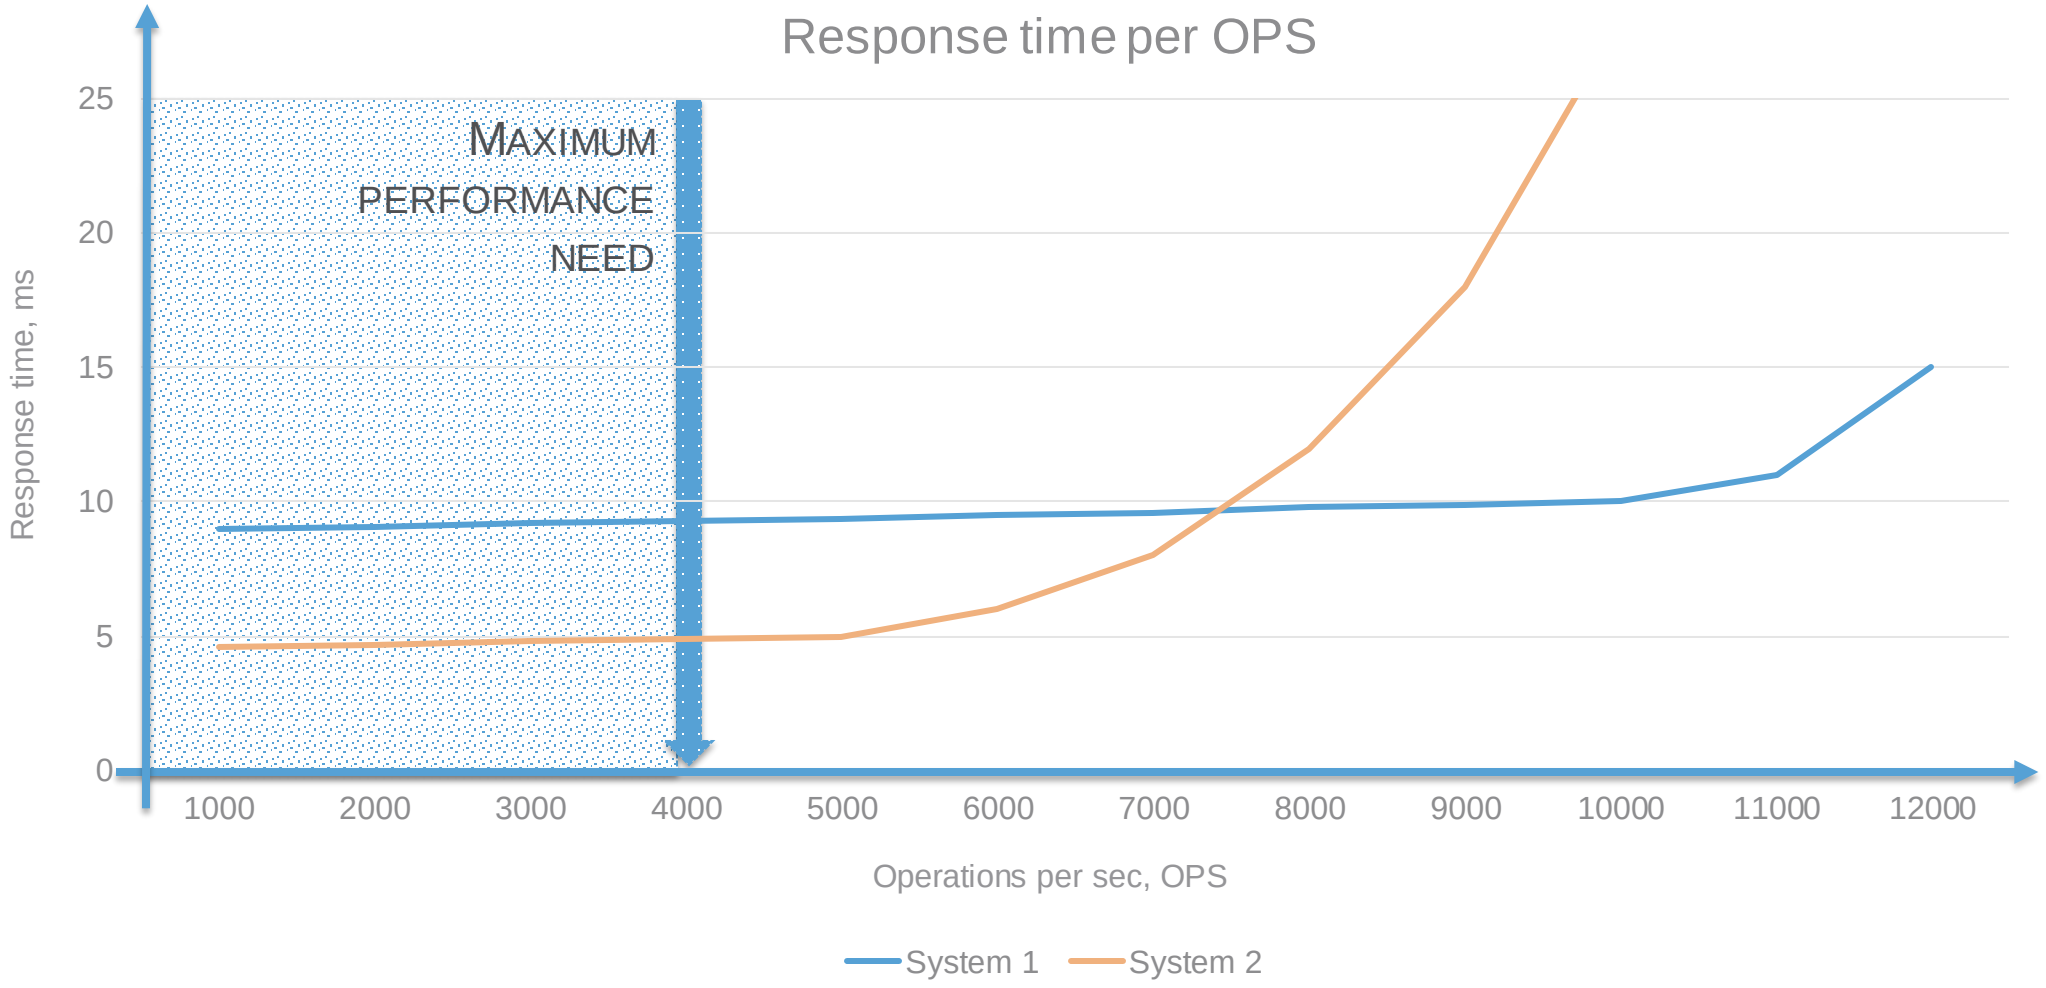

Graph Fun Or Why “It Depends” Is Relevant?

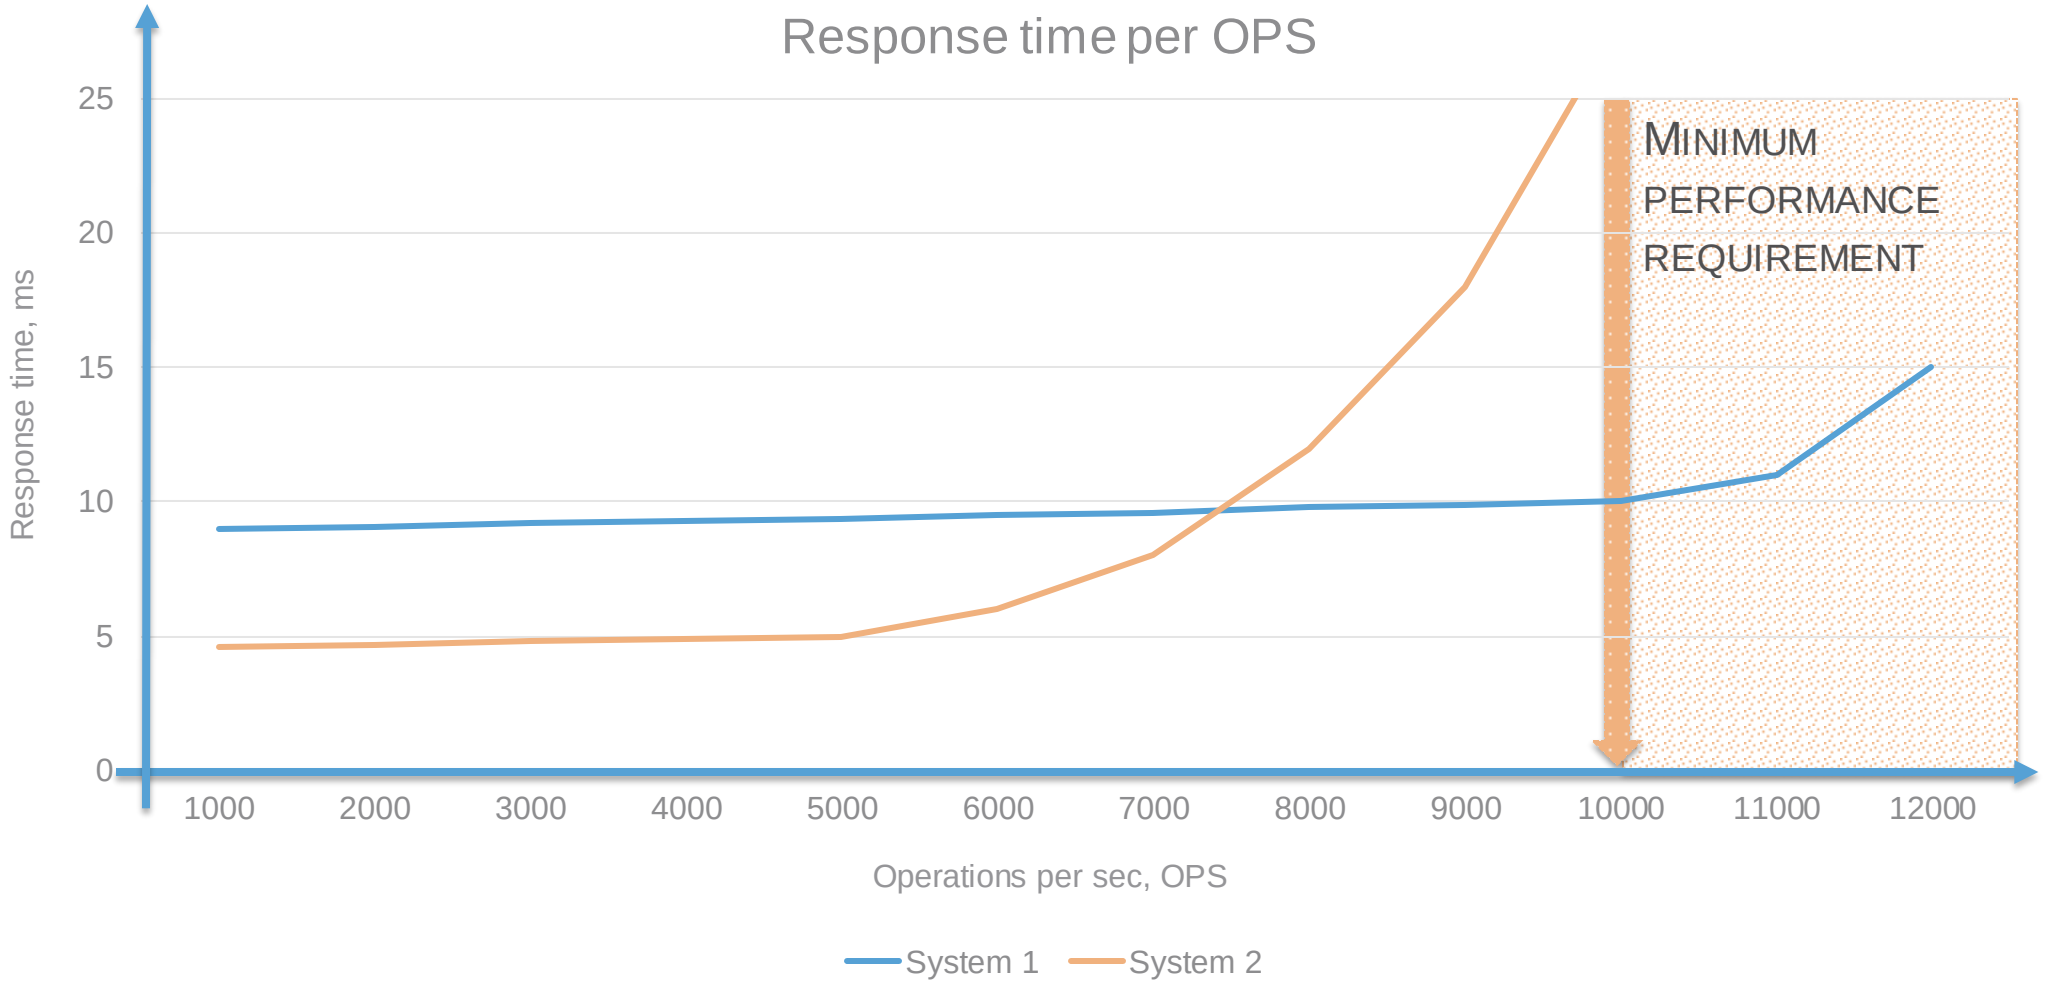

Graph Fun Or Why “It Depends” Is Relevant?

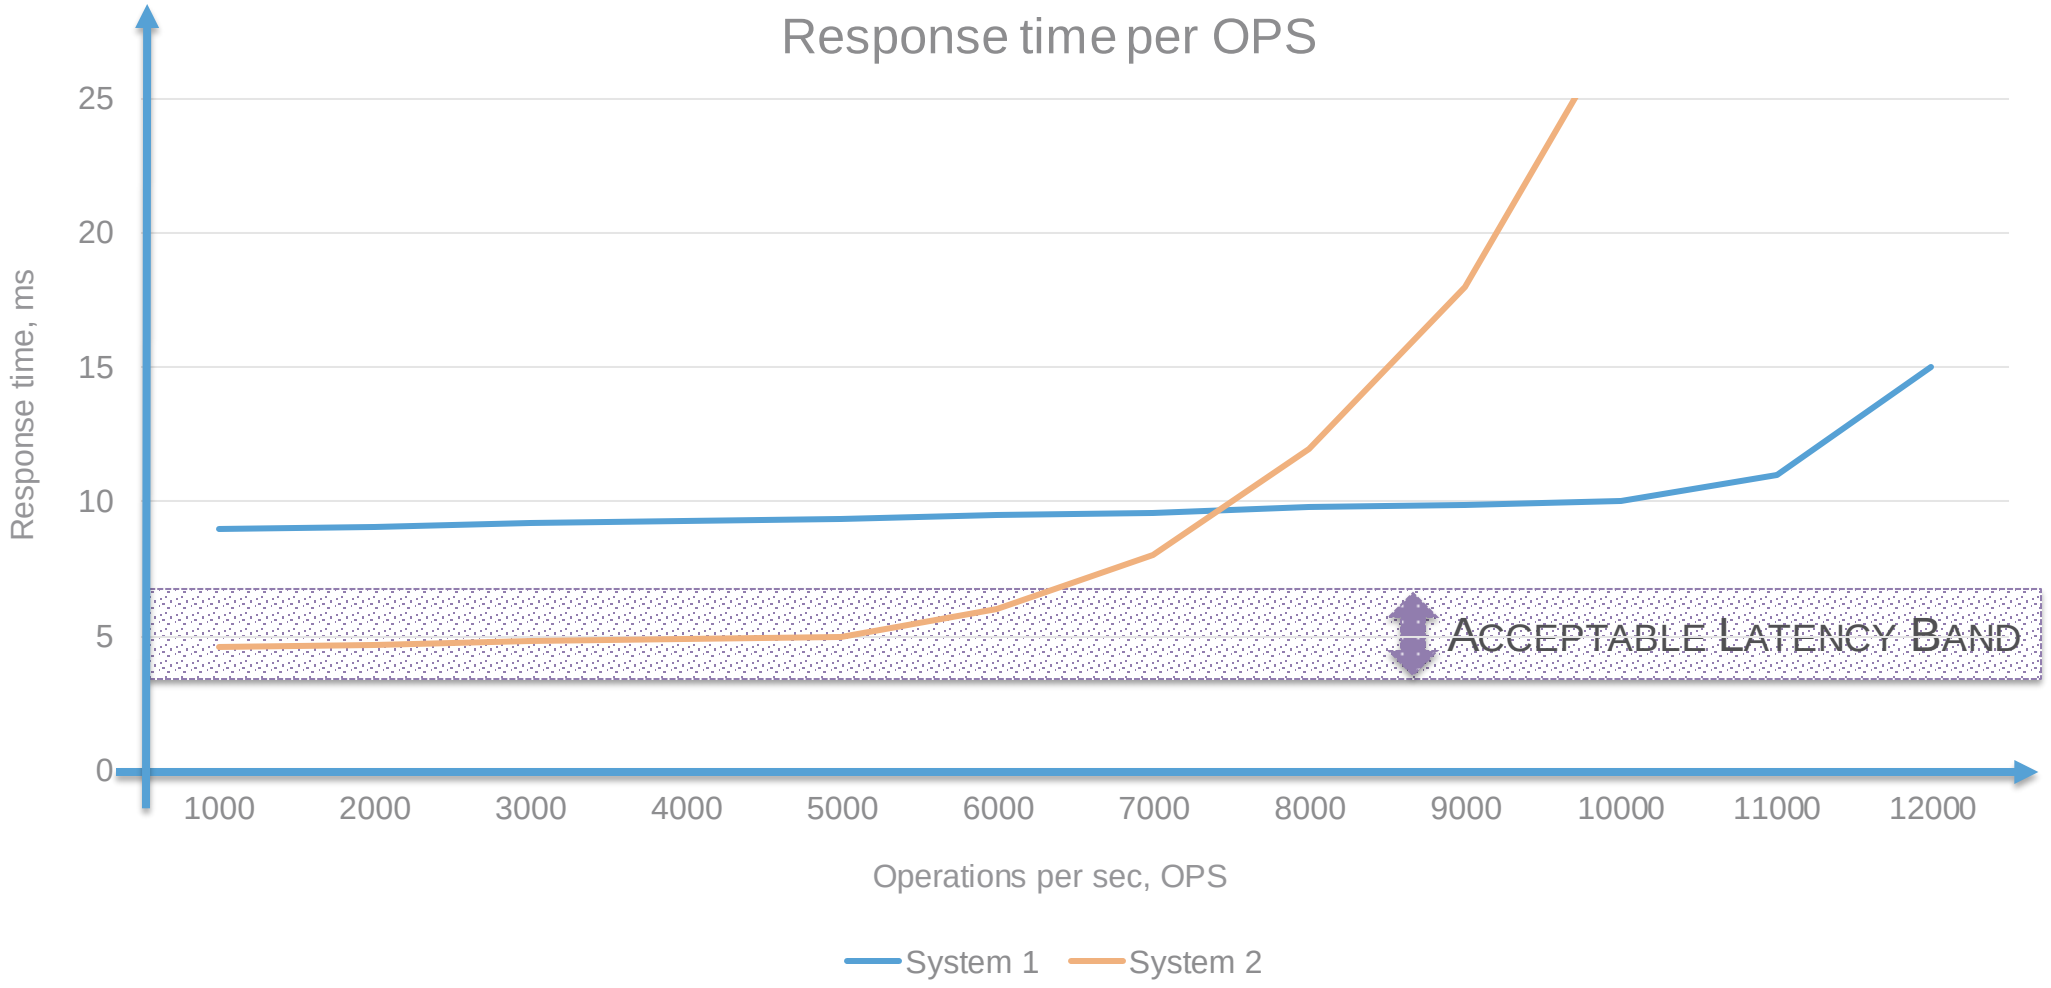

Variance

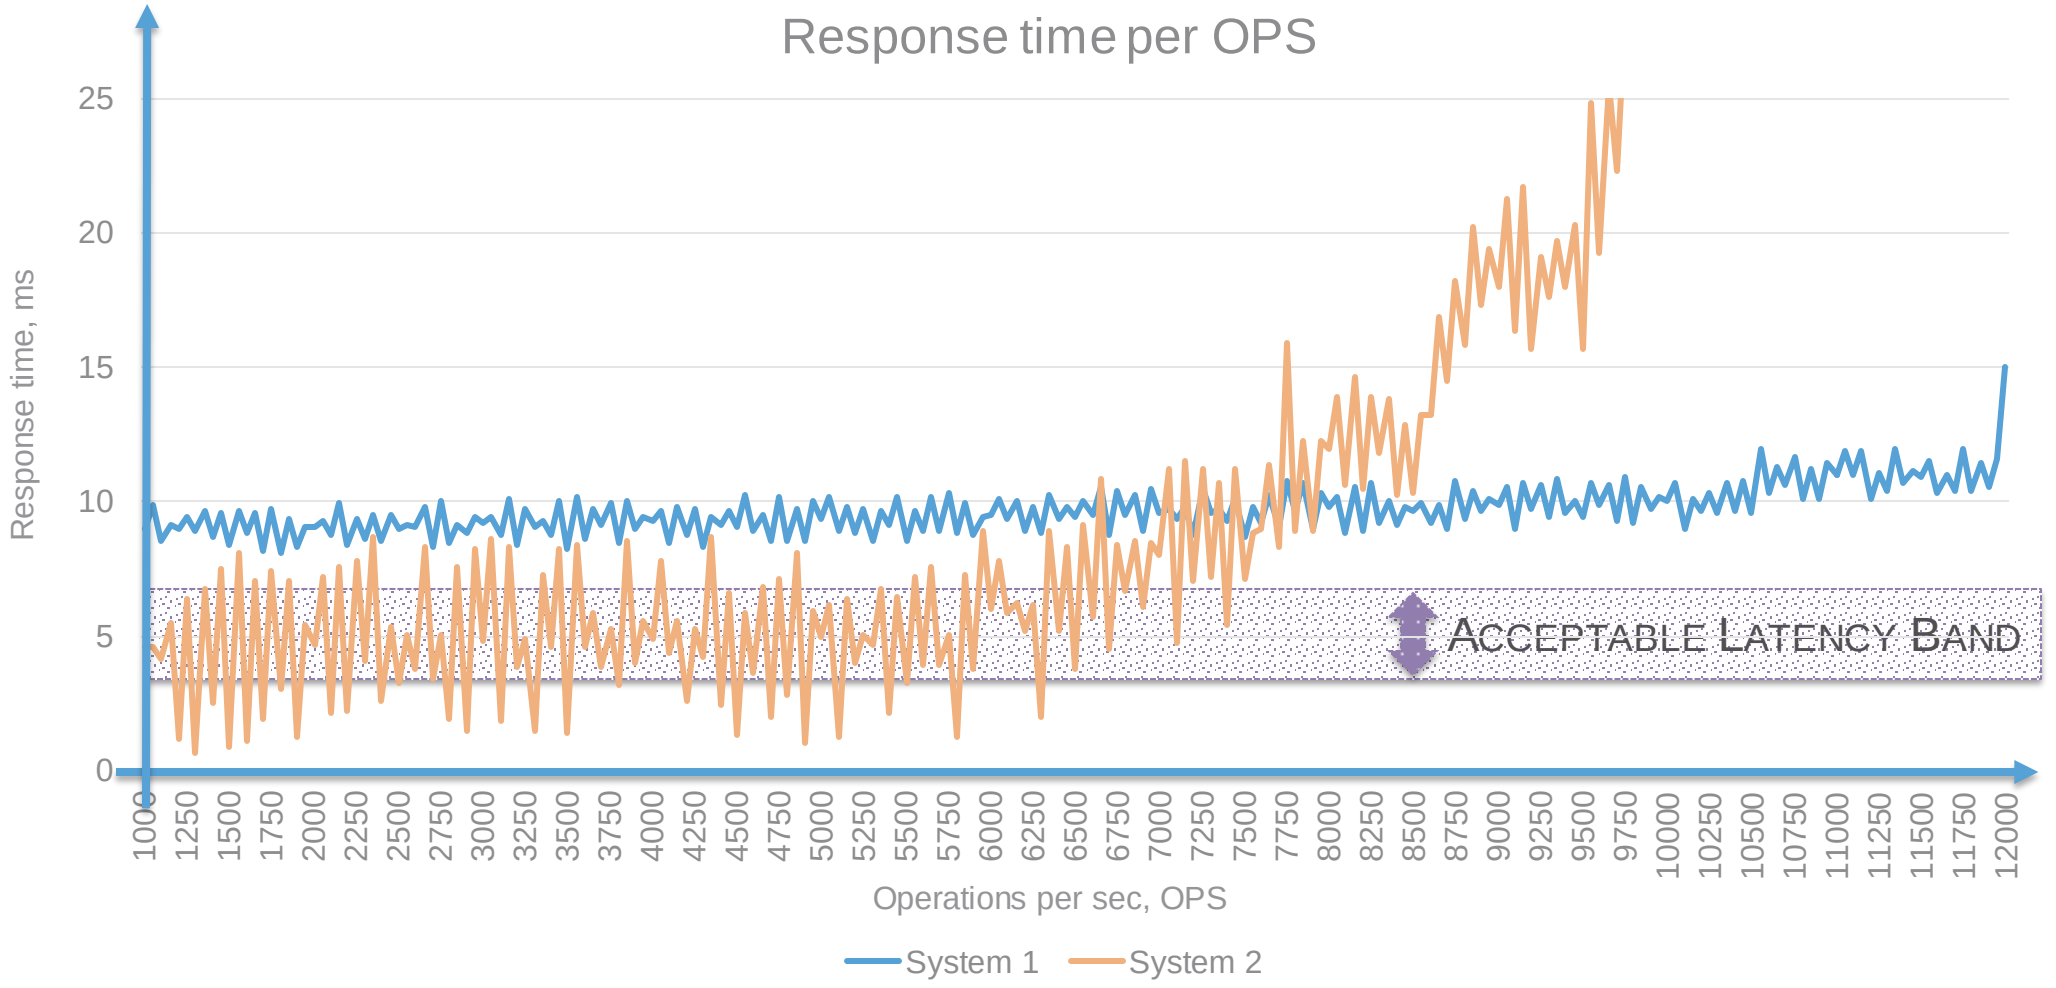

Other Metrics To Consider

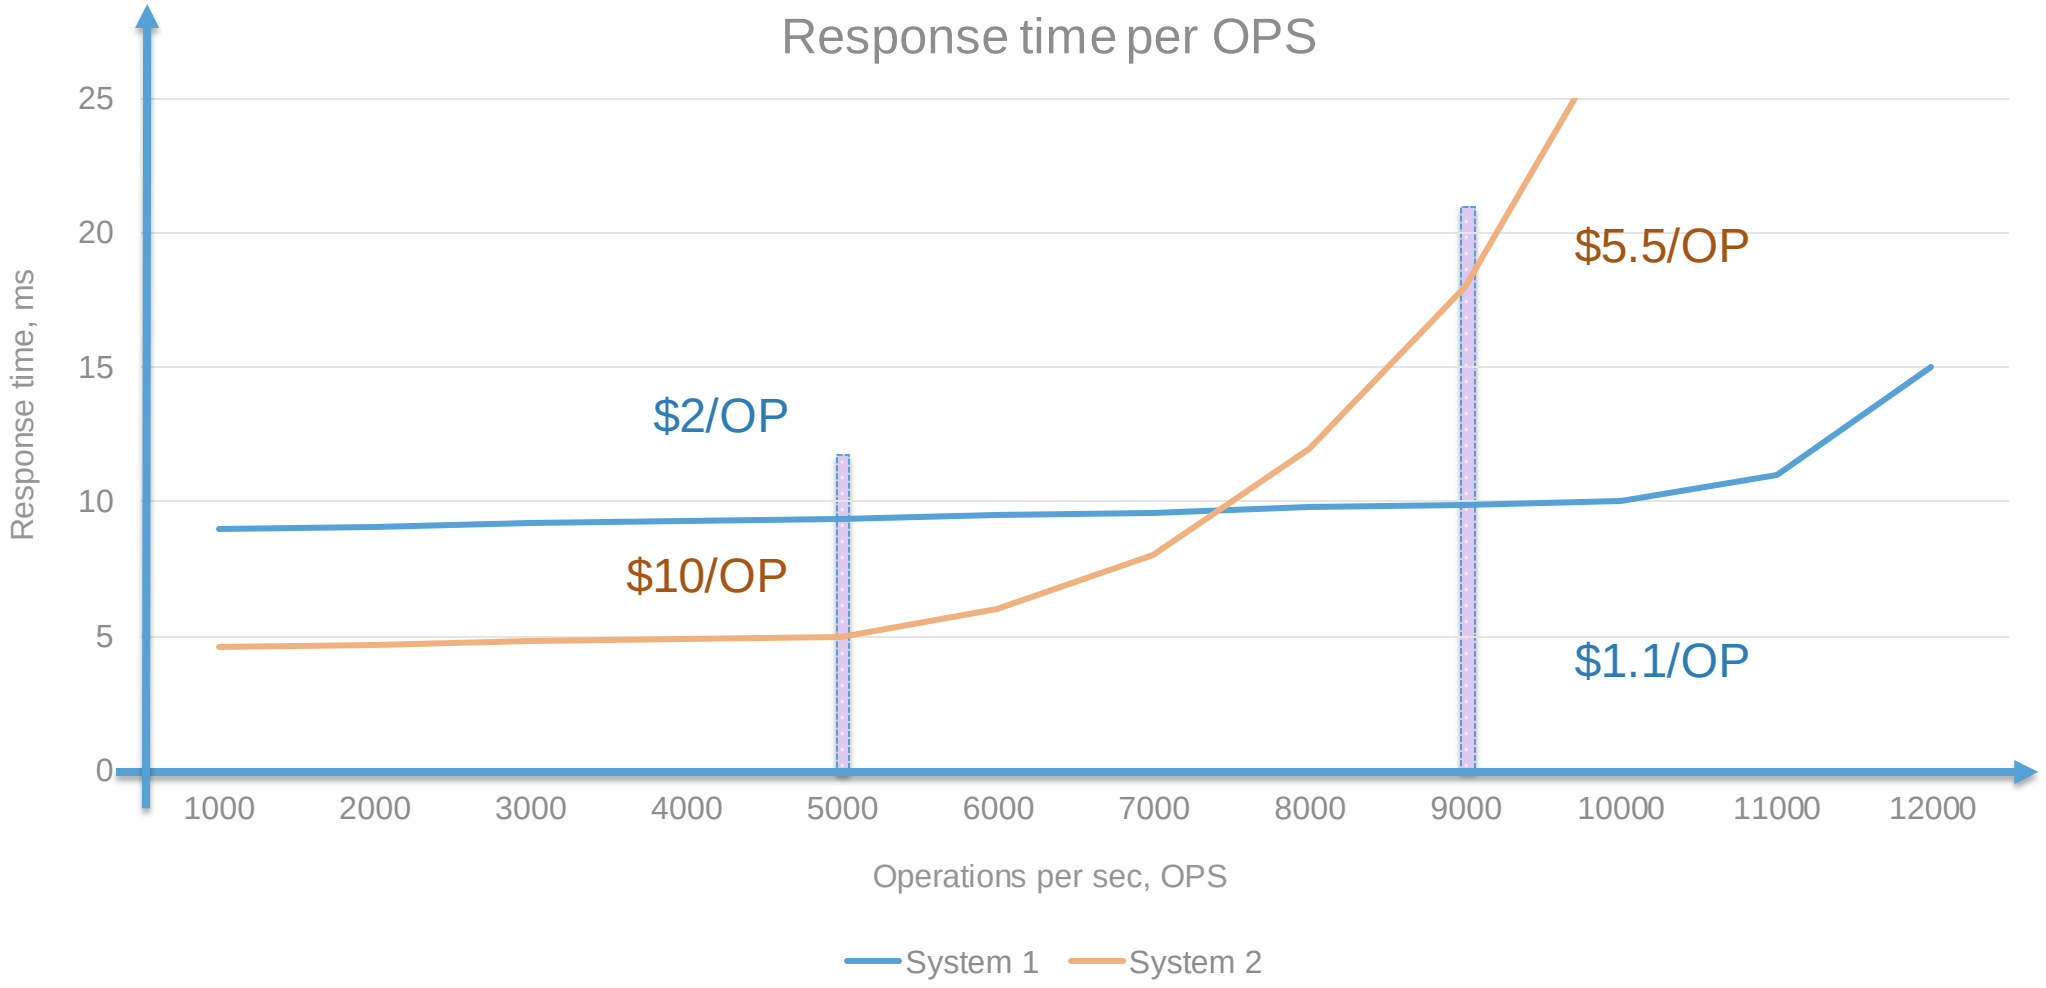

Other Metrics To Consider

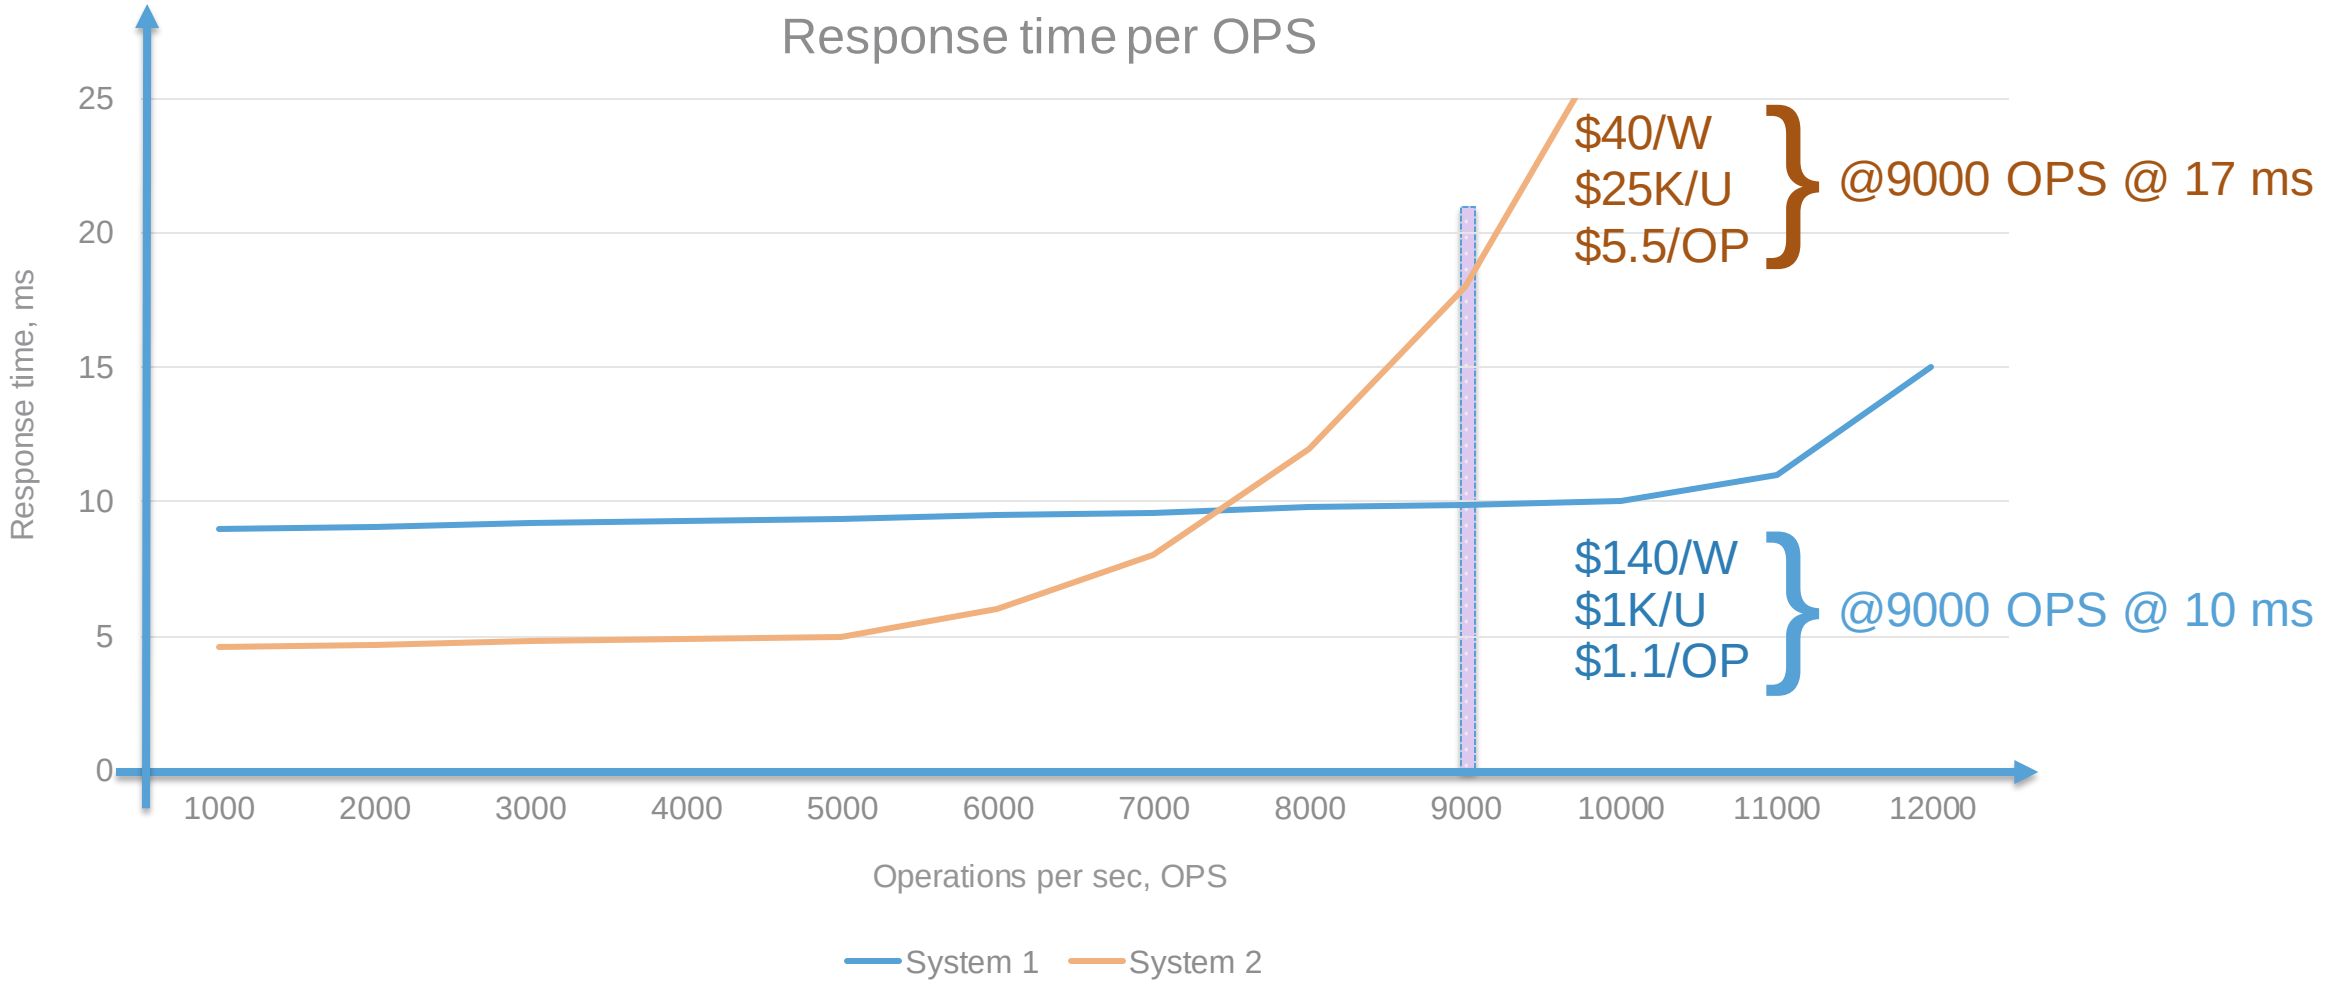

Other Metrics To Consider

CONTEXT MAKES METRICS MATTER

Summary

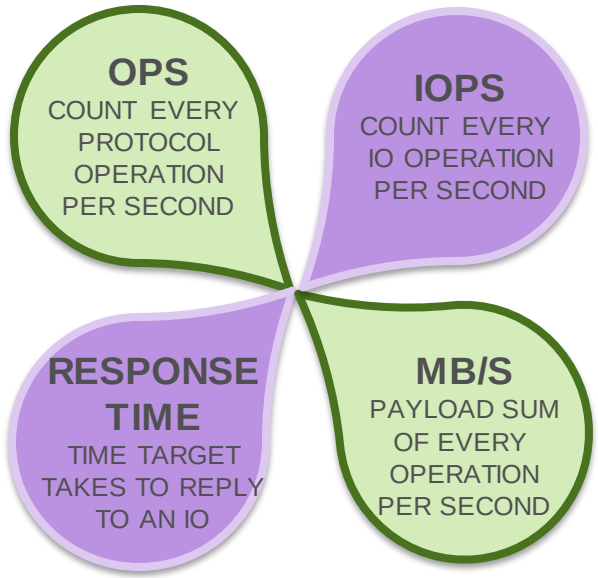

TERMINOLOGY

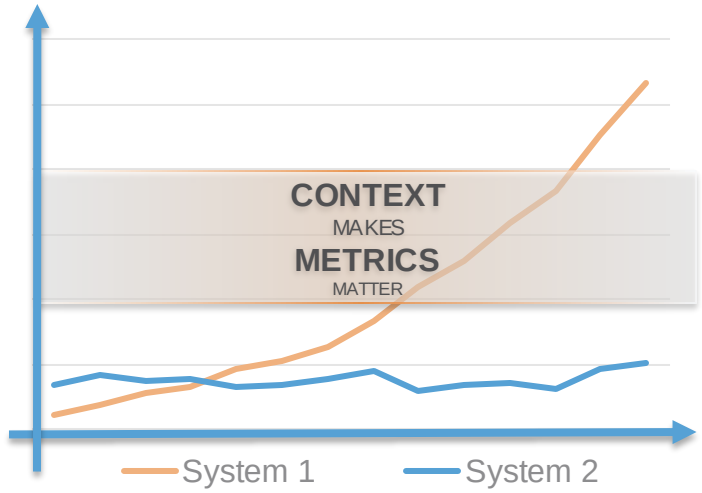

GRAPH FUN

Storage Performance Benchmarking

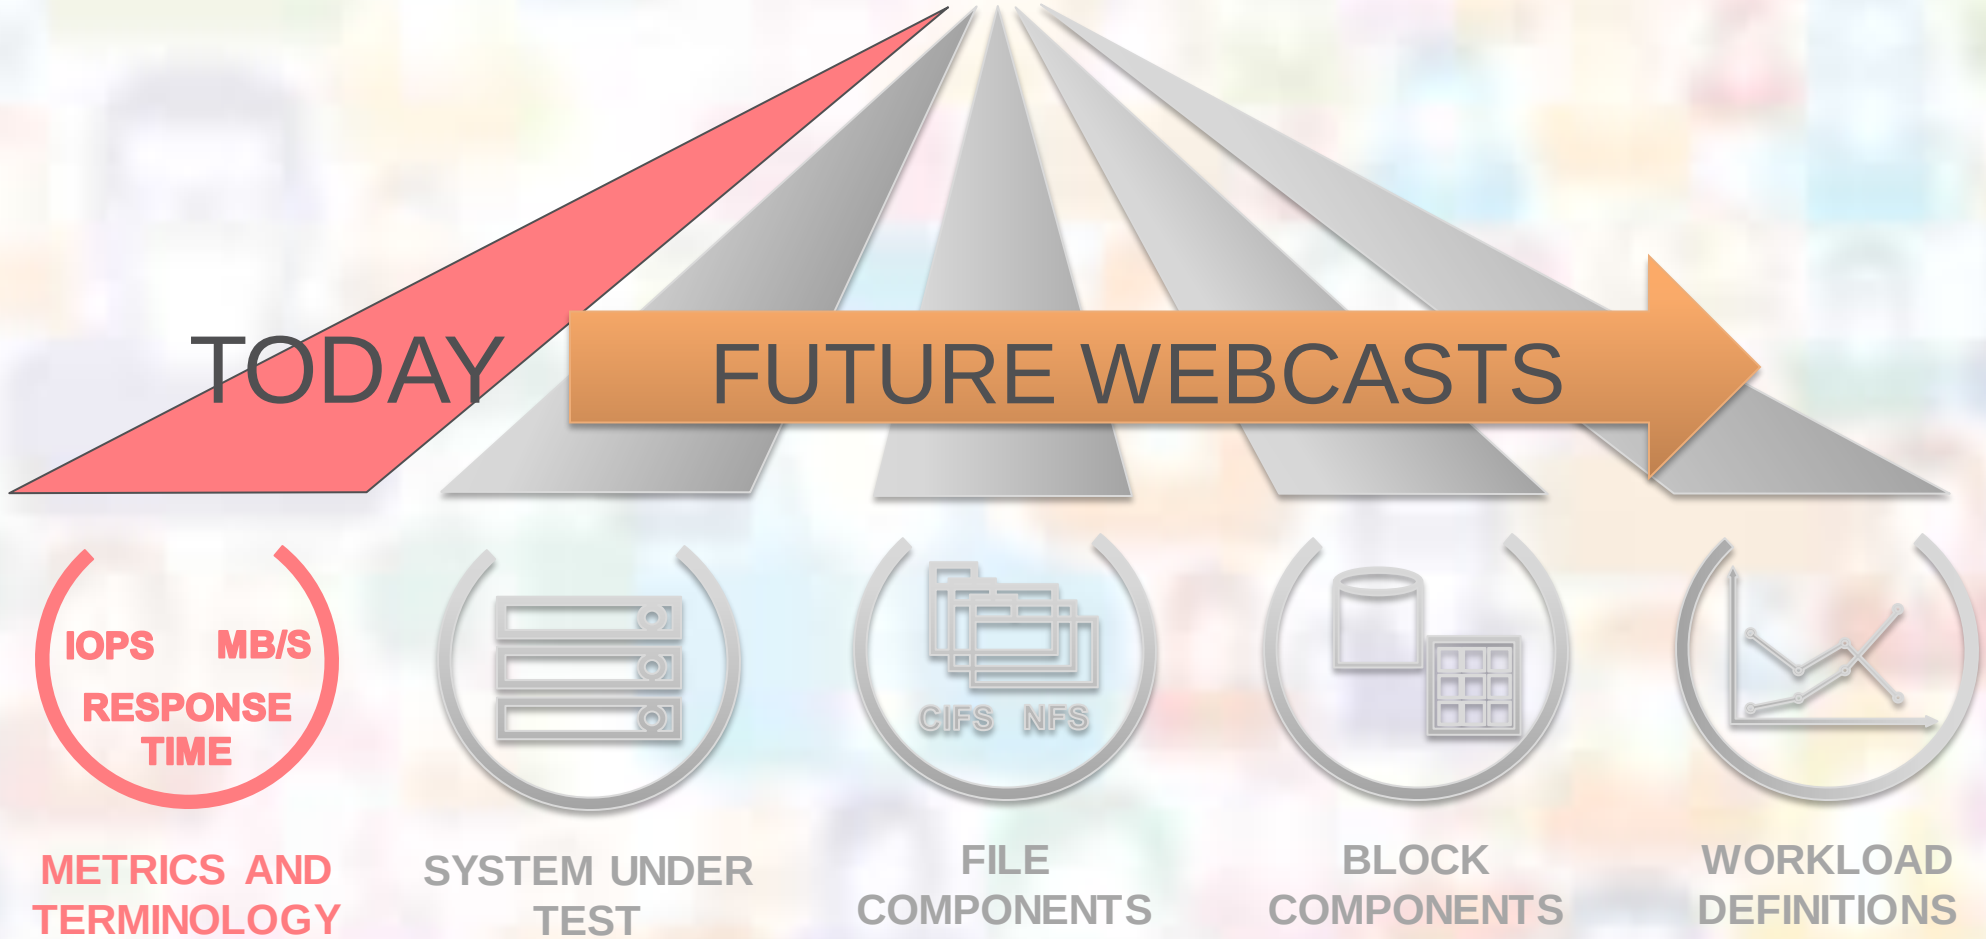

After This Webcast

- This Webcast and a PDF of the slides will be posted to the SNIA Ethernet Storage Forum (ESF) website and available on-demand
  - http://www.snia.org/forums/esf/knowledge/webcasts
- A full Q&A from this webcast, including answers to questions we couldn't get to today, will be posted to the SNIA-ESF blog
  - http://sniaesfblog.org/
- Follow us on Twitter @SNIAESF
- Next Webcast - October 21st
  - "Storage Performance Benchmarking: Part 2"

SNIA™
Ethernet Storage Forum

THANK YOU!