



STORAGE DEVELOPER CONFERENCE

SNIA ■ SANTA CLARA, 2016

# Introducing the EDA Workload for the SPEC SFS® 2014 Benchmark

**Jig Bhadaliya – Dell EMC**

**Nick Principe – Dell EMC**

# Agenda

- ❑ SPEC SFS 2014 Intro
- ❑ What is EDA (Electronic Design Automation)?
- ❑ Characteristics of EDA workloads
  - ❑ Defining an EDA workload
  - ❑ Analysis of trace data from EDA environments
- ❑ EDA Workload Definition for SPEC SFS 2014
- ❑ Example data
- ❑ Other SFS 2014 load generator enhancements

# SPEC

## Standard Performance Evaluation Corporation

- ❑ **The Standard Performance Evaluation Corporation (SPEC)** is a non-profit corporation formed to establish, maintain and endorse a standardized set of relevant benchmarks that can be applied to the newest generation of high-performance computers. SPEC develops benchmark suites and also reviews and publishes submitted results from member organizations and other benchmark licensees
- ❑ [www.spec.org](http://www.spec.org)
- ❑ SPEC and SPEC SFS are registered trademarks of the Standard Performance Evaluation Corporation.
- ❑ SPEC, the SPEC logo and SPEC SFS are registered trademarks of the Standard Performance Evaluation Corporation, reprint with permission.

# Disclaimer

- ❑ The SPEC SFS 2014 SP2 release, and the EDA workload definition, as represented in this presentation, is *pre-release* software
- ❑ The benchmark framework, workload, and features are still under internal SPEC review and may change before final release of SPEC SFS 2014 SP2 and/or the EDA workload.

# SPEC SFS 2014 Overview

- ❑ SPEC SFS 2014 is an industry-standard storage solution benchmark
  - ❑ Realistic, Solution-based workloads
    - ❑ DATABASE, SWBUILD, VDA, VDI, plus EDA (in SP2)
  - ❑ Measures application-level performance
    - ❑ Engages whole storage stack: application to “disk”
  - ❑ Ability to measure a broad-range of products and configurations
    - ❑ Traditional, Multi-Tier, All-Flash, Cloud, and more
  - ❑ Not just NAS!
    - ❑ Ability to test any fully-featured file system

5

# SPEC SFS 2014 Further Reading

- ❑ For a more extensive introduction to SFS 2014
  - ❑ SDC 2014: SPEC SFS 2014 – The Workloads and Metrics, an Under-the-Hood Review
    - ❑ Spencer Shepler, Nick Principe, Ken Cantrell
    - ❑ [http://www.snia.org/sites/default/files/SpencerShepler\\_SPEC\\_Under-the-Hood\\_Review\\_Final.pdf](http://www.snia.org/sites/default/files/SpencerShepler_SPEC_Under-the-Hood_Review_Final.pdf)
    - ❑ <http://spec.org/sfs2014/presentations/benchmarking.html>
  - ❑ SDC 2015: Application-Level Benchmarking with SPEC SFS 2014
    - ❑ Nick Principe, Vernon Miller
    - ❑ [http://www.snia.org/sites/default/files/SDC15\\_presentations/performance/Principe\\_MillerApplication\\_Level\\_Benchmarking\\_v1.6.pdf](http://www.snia.org/sites/default/files/SDC15_presentations/performance/Principe_MillerApplication_Level_Benchmarking_v1.6.pdf)
    - ❑ <https://www.youtube.com/watch?v=4wfeM1q0zHA>

# What is EDA?

- ❑ Electronic Design Automation
  - ❑ Represents software tools and workflows for designing semiconductor chips
  - ❑ Dozens of software tools used to design a chip from specification to fabrication
  - ❑ Very compute-heavy process – high concurrency
  - ❑ Storage is often the performance bottleneck

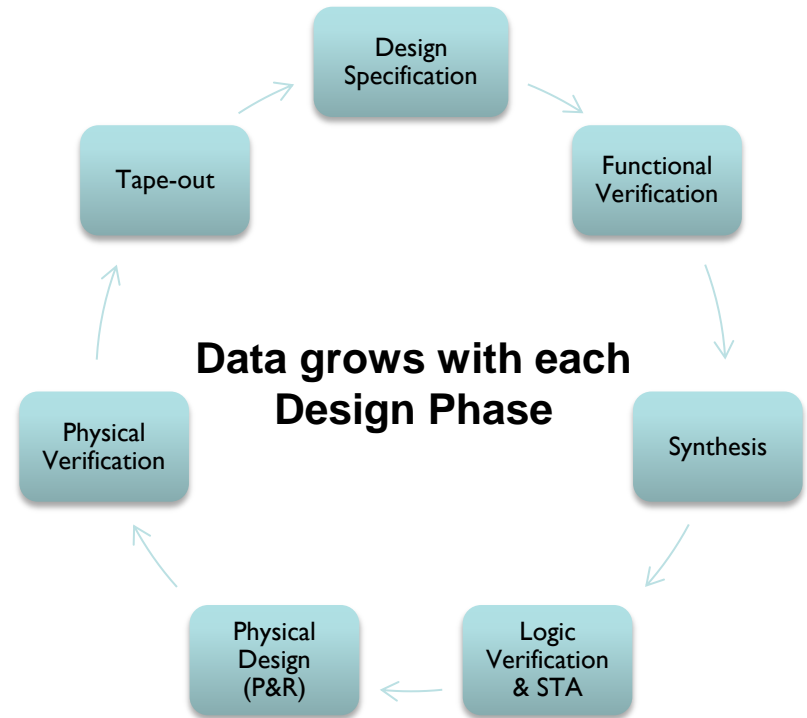
# Characteristics of EDA Workloads

- ❑ Millions to billions of small files
  - ❑ With a small % of large files
- ❑ Mixed random and sequential I/O
  - ❑ Metadata intensive
- ❑ Two high level design phases
  - ❑ Frontend Design
    - ❑ Millions of small files
    - ❑ Transient output of terabytes that requires fast storage
  - ❑ Backend Design (aka Physical Design & Verification)
    - ❑ Larger file sizes with more sequential access patterns
- ❑ Many EDA jobs running concurrently
  - ❑ Storage sees a combined and random mix of workloads from both multiple design phases

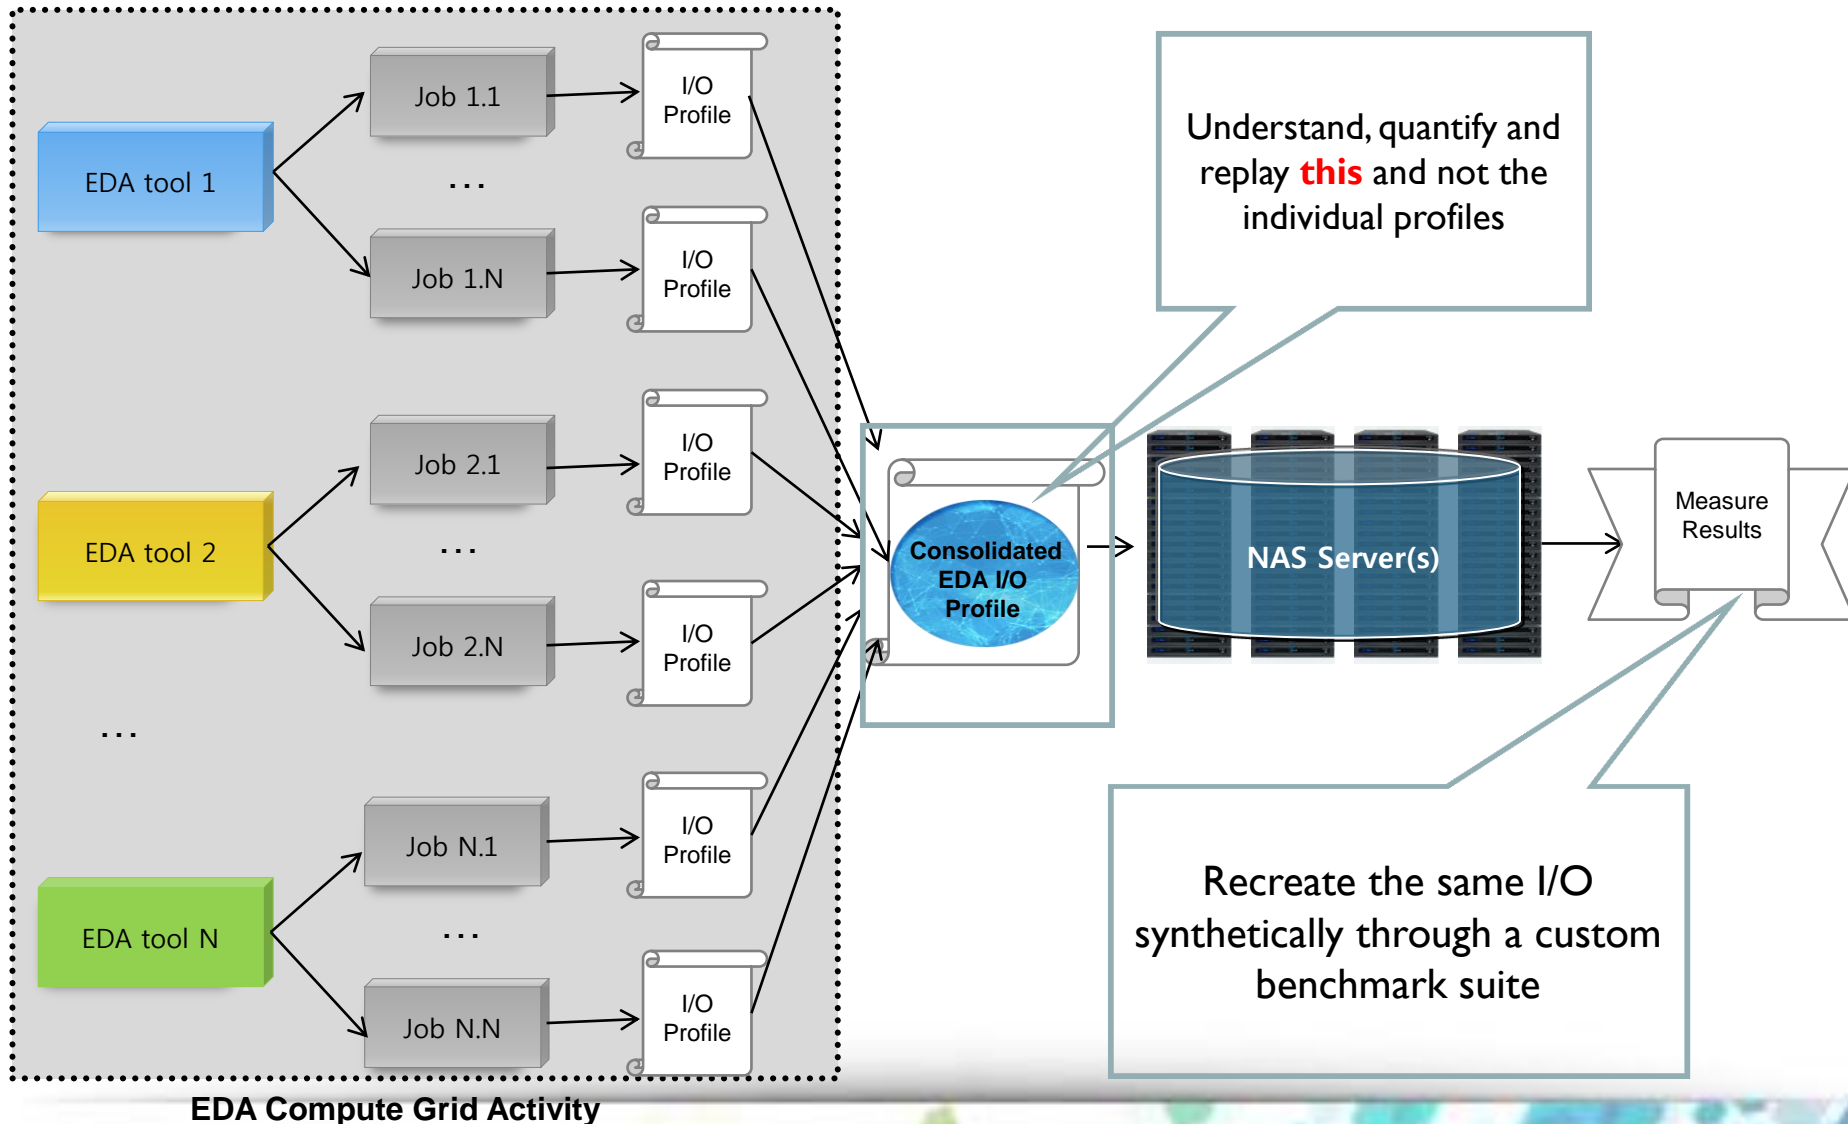


# Characteristics of EDA Workloads

- ❑ RTL Description (chip in HDL) – tiny in size in terms of storage
- ❑ Frontend Design Phases
  - ❑ Millions of small files
  - ❑ Requires high levels of concurrency, not throughput
  - ❑ Transient output of simulations results in TBs of data
  - ❑ Scratch space yet requires Tier1 storage in terms of performance
- ❑ Physical Design and Verification
  - ❑ File sizes increase
  - ❑ Data access pattern becomes more sequential than random
  - ❑ Single top level GDS-II file could be in tens of GBs




# Understanding EDA Workload



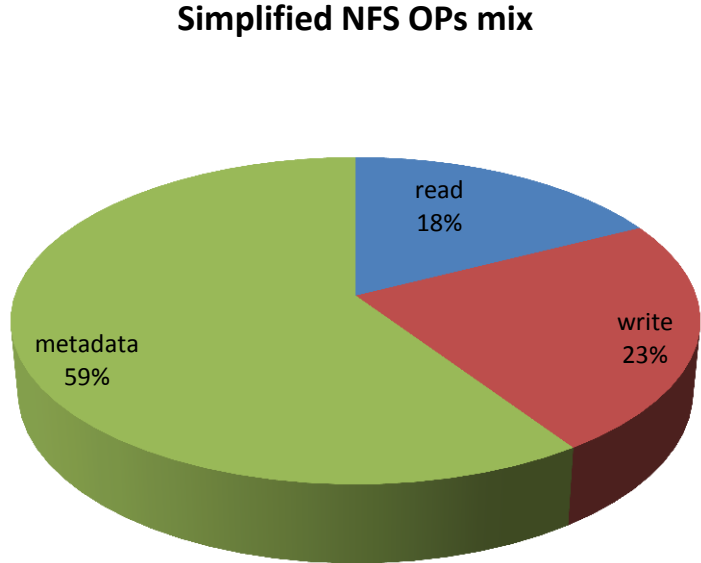

EDA Compute Grid Activity

# EDA Workload - profile over time



- ❑ Notice the sustained pattern and predictability to the workload.
- ❑ Activity is fairly proportionate, with periodic spikes of metadata and read activity
- ❑ Notice the sustained pattern and predictability to the workload.
- ❑ Activity is fairly proportionate, with periodic spikes of metadata and read activity

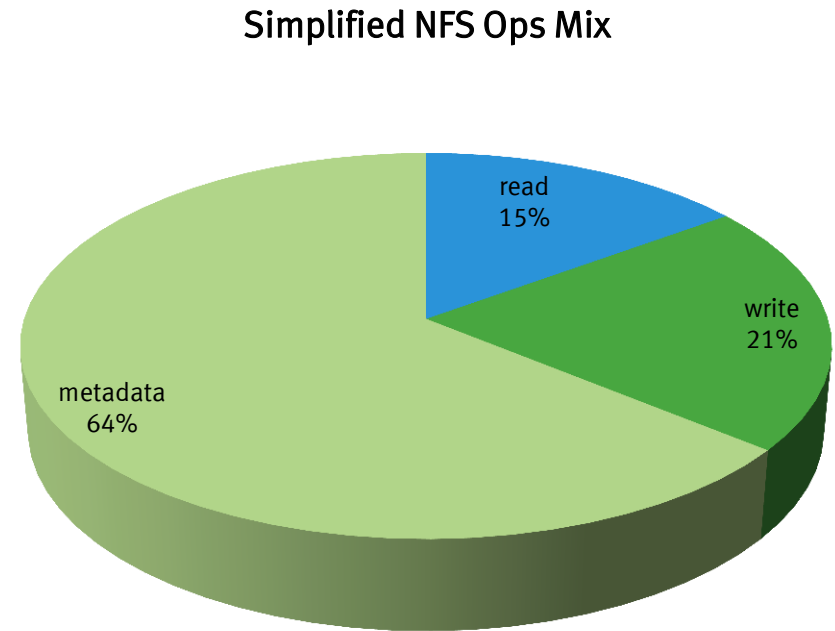

# EDA Workload Profiles – Peer 1



- ❑ Large sample of 1.84 trillion NFS OPs, collected from 20 NAS systems over 300+ days
- ❑ Profile mix is: 59/23/18 (metadata/write/read)
- ❑ Workload is attribute intensive due to a large number of files, plus deep and wide directory structure.

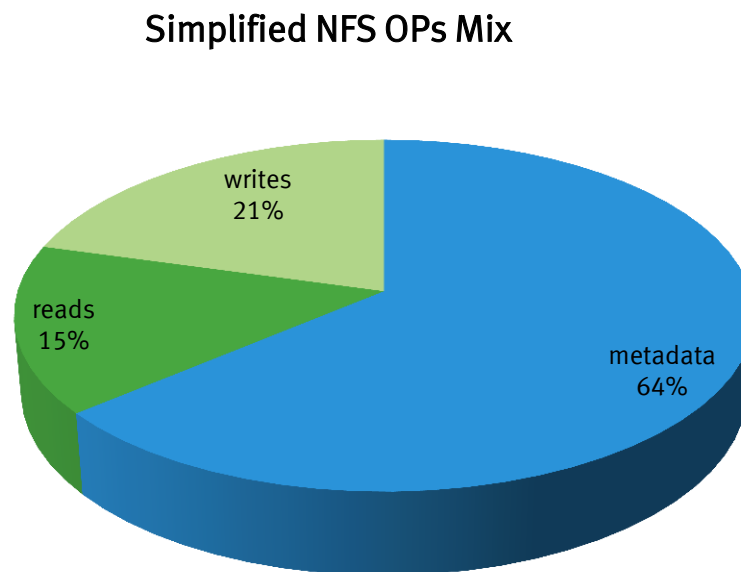
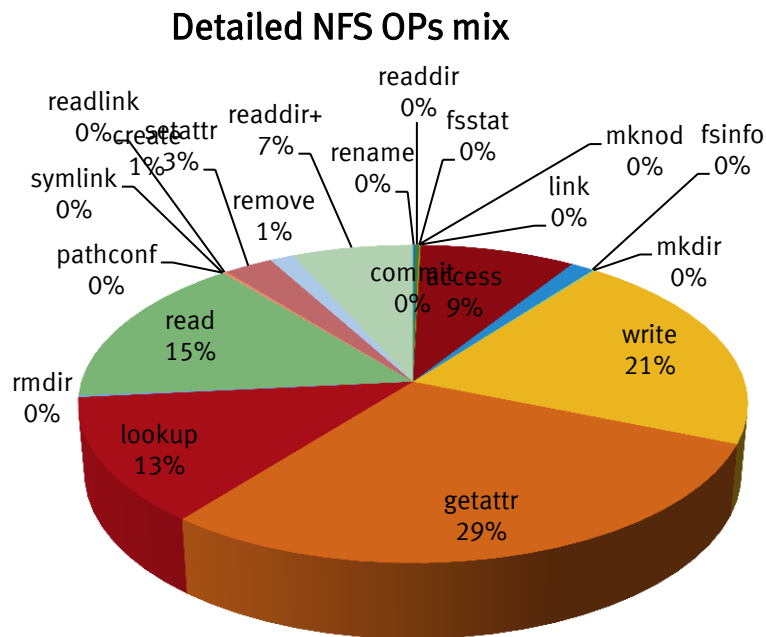


# EDA Workload Profiles – Peer 2



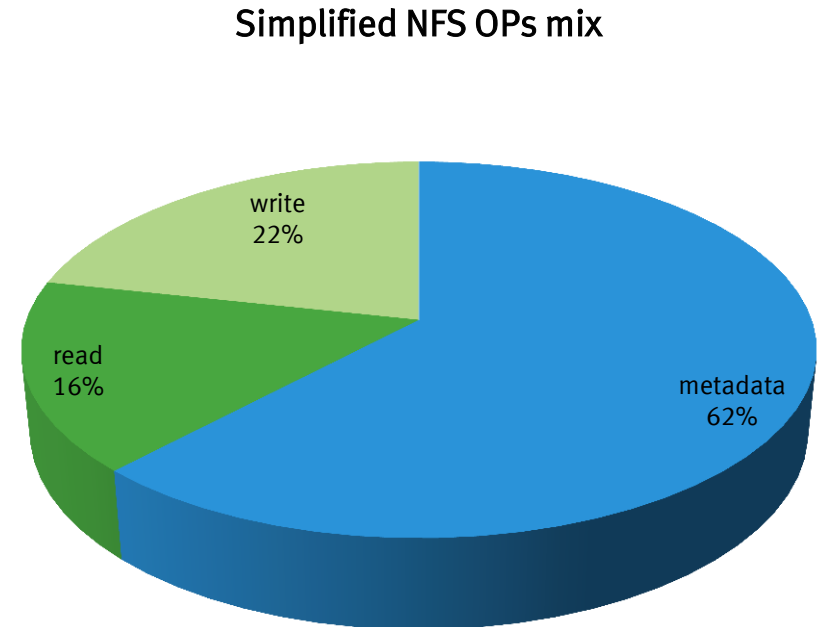
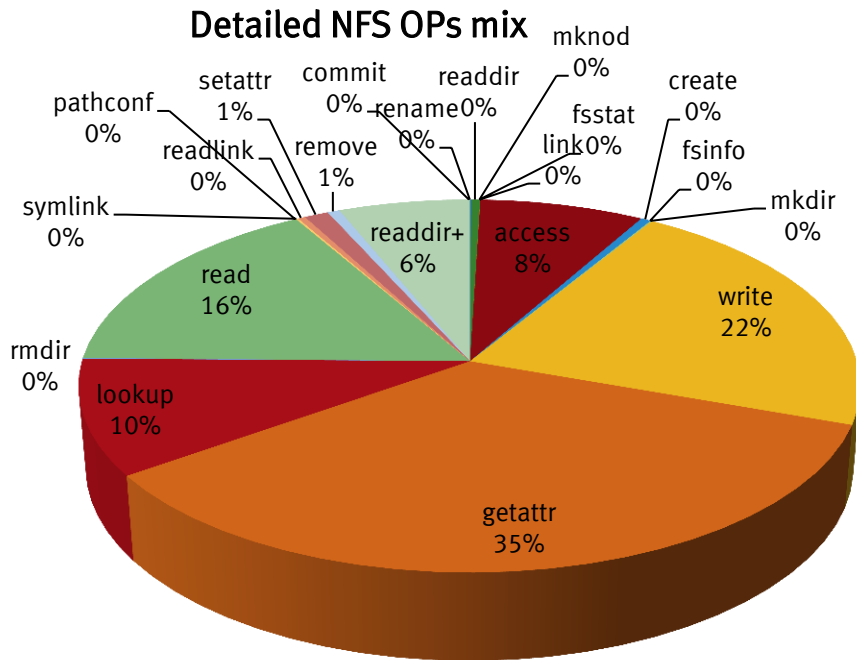
- ❑ This sample is based on 3.35 billion NFS operations from another EDA company
- ❑ Profile has similar mix as the previous peer

# EDA Workload Profiles – Peer 3



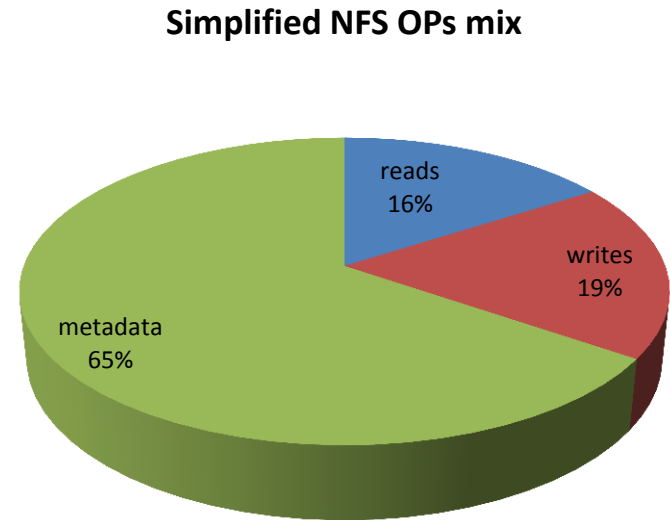
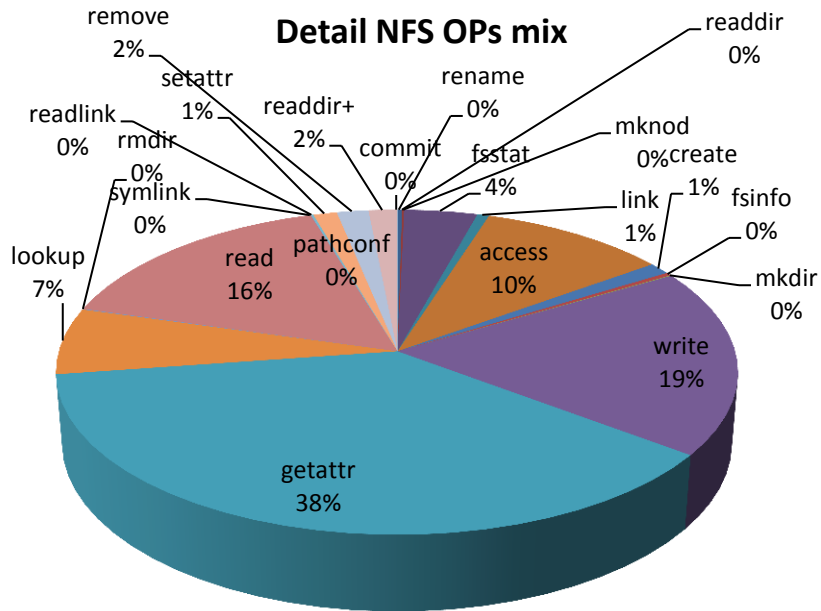
- ❑ This sample based on 7 \* NAS Systems
- ❑ Sample size of 188+ billion NFS OPs total over and uptime of 120+ days
- ❑ 60+% metadata OPs similar to other EDA peers
- ❑ Slightly more writes than reads, similar to previous profiles

# EDA Workload Profiles – Peer 4



- ❑ Analysis is based on 589 billion NFS operations
- ❑ Workload is metadata intensive
- ❑ Mix is similar to other EDA peers


# EDA Workload Profiles – Peer 5



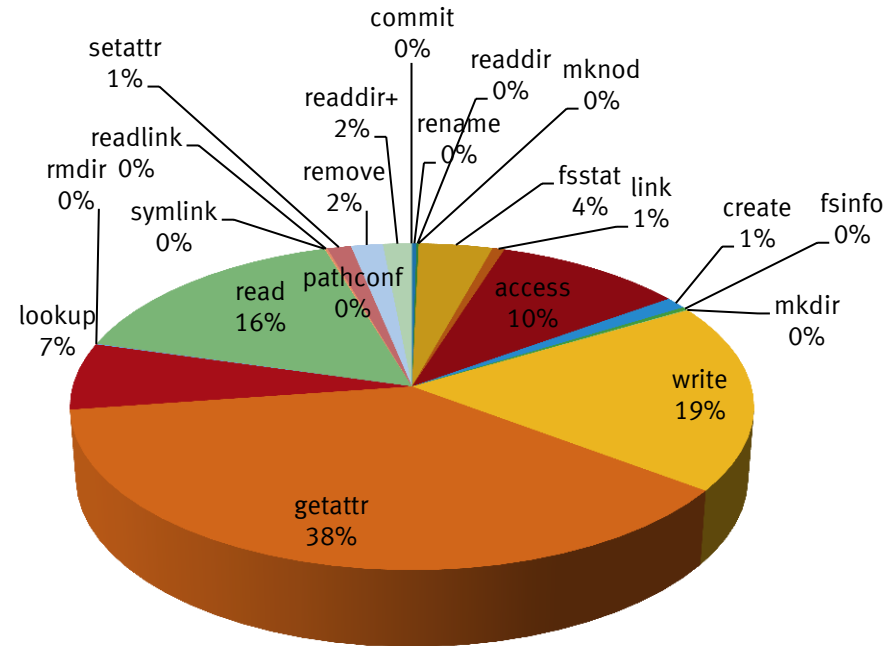
- ❑ Four NAS systems with uptime of 110 days
- ❑ Sample size: 110 billion NFS OPs
- ❑ Consistently similar to other peers



# EDA Workload Profiles – Peer Comparison



Peer 1 : Detailed NFS Ops Mix



Peer 5 : Detailed NFS Ops Mix

□ EDA Peer-1 profile is strikingly similar to EDA Peer-5

# Studies of EDA workloads globally

Customer	# of File Servers	Sample size (NFS Ops)	Reads (%)	Writes (%)	Meta-Reads (%)	Meta-Writes (%)	Creates (%)	Deletes (%)	Others (%)
	32	10,363,761,458,104	13.86%	10.41%	71.73%	1.52%	0.61%	0.99%	0.89%
	6	2,768,210,838,022	4.12%	3.42%	91.03%	0.40%	0.45%	0.46%	0.11%
	20	1,840,771,676,833	17.52%	23.07%	57.06%	0.90%	0.64%	0.33%	0.47%
	42	1,816,825,667,071	9.40%	23.89%	63.14%	1.05%	0.67%	1.24%	0.59%
	32	953,879,205,595	26.23%	7.12%	64.18%	1.24%	0.43%	0.41%	0.40%
	4	758,777,007,509	7.42%	7.74%	83.96%	0.33%	0.23%	0.19%	0.12%
	1	589,580,191,282	16.36%	21.61%	59.53%	1.14%	0.48%	0.57%	0.32%
	13	563,845,422,292	16.93%	22.01%	57.78%	1.34%	0.72%	0.73%	0.49%
	2	527,817,134,890	7.40%	9.15%	79.37%	1.04%	1.27%	1.13%	0.62%
	4	295,284,923,196	21.32%	18.66%	58.82%	0.47%	0.32%	0.28%	0.13%
	7	188,352,047,947	15.45%	20.70%	58.07%	2.71%	1.30%	1.30%	0.48%
	4	105,589,268,942	15.91%	18.63%	59.99%	1.27%	1.17%	1.70%	1.27%
	4	92,387,333,892	5.31%	13.84%	66.43%	9.06%	1.99%	2.09%	1.27%
	4	71,137,375,389	26.08%	30.89%	39.22%	2.67%	0.47%	0.29%	0.37%
	8	37,971,948,841	6.71%	19.58%	59.02%	0.44%	0.41%	0.45%	13.38%
	1	3,350,836,259	15.00%	20.90%	61.28%	0.71%	0.97%	0.90%	0.23%

# EDA Workload Definition for SPEC SFS 2014 – Motivation

- Why produce a standard EDA workload within SPEC SFS 2014?
  - Commonly used benchmarking tools are often not representative of production EDA workloads
  - Various different EDA applications produce different storage workloads
  - EDA tool-specific testing is often inconclusive due to lack of scale; not considering aggregate workload
  - In aggregate, there is sufficient pattern to produce a representative combined workload

# EDA Workload Definition for SPEC SFS 2014

- ❑ Significant differences in I/O pattern between EDA Frontend vs. Backend design flows
- ❑ Two component workloads:
  - ❑ EDA\_FRONTEND
  - ❑ EDA\_BACKEND

# EDA Workload Definition for SPEC SFS 2014

## EDA\_FRONTEND

Business Metric	Category		Value	
	Workload name	EDA_FRONTEND		
	Metric	JOBS		
Capacity/Bus. Met.	TBD			

File Operation Distribution	Operation	%	Operation	%
	read	0	read file	7
	mmap read	0	rand read	0
	write	0	write file	10
	mmap write	0	rand write	15
	rmw	0	append	0
	mkdir	1	readdir	0
	create	2	unlink	1
	unlink2	1	stat	39
	rename	0	access	15
	locking	0	copyfile	0
	statfs	0	chmod	1
	pathconf	0		

Read Transfer Size Distribution	Slot	Start	End	%
	0	1	511	4
	1	512	1023	0
	2	1024	2047	0
	3	2048	4095	2
	4	4096	8191	43
	5	8192	16383	30
	6	16384	32767	21
	7	32768	65535	0
	8	65536	65536	0
	9	131072	131072	0
	10	1	1	0
	11	1	1	0
	12	1	1	0
	13	1	1	0
	14	1	1	0
15	1	1	1	

Write Transfer Size Distribution	Slot	Start	End	%
	0	1	511	25
	1	512	1023	10
	2	1024	2047	15
	3	2048	4095	18
	4	4096	8191	27
	5	8292	16383	3
	6	16384	32767	2
	7	32768	65535	0
	8	65536	65536	0
	9	131072	131072	0
	10	1	1	0
	11	1	1	0
	12	1	1	0
	13	1	1	0
	14	1	1	0
15	1	1	0	

Miscellaneous	Option	%	Option	Value
	write commit	15	background	0
	direct	0	sharemode	0
	osync	0	uniform size dist	0
	geometric	50	rand dist behavior	0
	compress	50	align	0

Thresholds	Threshold	Value
	proc oprate	75%
	global oprate	95%
	workload variance	5%
	proc latency	n/a ms
	global latency	n/a ms

Execution Parameters	Parameter	Value
	procs	3
	oprate per proc	100
	dirs per proc	10
	files per dir	10
	avg file size	8 KiB

# EDA Workload Definition for SPEC SFS 2014

## EDA\_BACKEND

Business Metric	Category		Value	
	Workload name	EDA_BACKEND		
	Metric	JOBS		
	Capacity/Bus. Met.	TBD		

File Operation Distribution	Operation	%	Operation	%
	read	50	read file	0
	mmap read	0	rand read	0
	write	50	write file	0
	mmap write	0	rand write	0
	rmw	0	append	0
	mkdir	0	readdir	0
	create	0	unlink	0
	unlink2	0	stat	0
	rename	0	access	0
	locking	0	copyfile	0
	statfs	0	chmod	0
	pathconf	0		

Read Transfer Size Distribution	Slot	Start	End	%
	0	1	511	0
	1	512	1023	0
	2	1024	2047	0
	3	2048	4095	0
	4	4096	8191	0
	5	8192	16383	0
	6	16384	32767	0
	7	32768	65535	49
	8	65536	65536	51
	9	131072	131072	0
	10	1	1	0
	11	1	1	0
	12	1	1	0
	13	1	1	0
	14	1	1	0
15	1	1	0	

Write Transfer Size Distribution	Slot	Start	End	%
	0	1	511	0
	1	512	1023	0
	2	1024	2047	0
	3	2048	4095	0
	4	4096	8191	0
	5	8292	16383	0
	6	16384	32767	0
	7	32768	65535	45
	8	65536	131072	55
	9	131072	131072	0
	10	1	1	0
	11	1	1	0
	12	1	1	0
	13	1	1	0
	14	1	1	0
15	1	1	0	

Miscellaneous	Option	%	Option	Value
	write commit	15	background	0
	direct	50	sharemode	0
	osync	5	uniform size dist	0
	geometric	50	rand dist behavior	0
	compress	50	align	0

Thresholds	Threshold	Value
	proc oprate	75%
	global oprate	95%
	workload variance	5%
	proc latency	n/a ms
	global latency	n/a ms

Execution Parameters	Parameter	Value
	procs	2
	oprate per proc	75
	dirs per proc	5
	files per dir	10
	avg file size	10 MiB

# Example Data

- [Expected by presentation time]

# SPEC SFS 2014 EDA Workload Addition

- ❑ The EDA workload is currently the anchor feature for our next slipstream release of the SFS 2014 benchmark
  - ❑ SPEC SFS 2014 SP2
- ❑ This will be a performance-neutral release for all existing workloads
  - ❑ All currently published SFS 2014 results remain valid and comparable to SP2 results
  - ❑ There will be no distinction on the results page between SP1 and SP2
- ❑ The upgrade to SP2 is expected to be free for all SPEC SFS 2014 license holders

24



# SPEC SFS 2014 SP2

## Additional Features

- There are a few new features that will be added to the SPEC SFS 2014 Benchmark suite in SP2, along with the EDA Workload:
  - unlink2 op type (remove non-empty files)
  - Dedupable dataset options
  - Dedupe and compression granule size options
  - Ability to have different file, directory, and file size distributions per component workload
  - Encrypted dataset option
  - Reduced memory consumption

25

# Key Takeaways

- ❑ EDA is a unique and interesting workload not adequately represented by other benchmarking tools
- ❑ EDA workload in SPEC SFS 2014 SP2 is based upon a large number of traces from multiple customers
- ❑ SPEC plans to include the EDA workload in the SPEC SFS 2014 SP2 and additional helpful features for custom workloads
- ❑ The SPEC SFS 2014 SP2 release is expected to be a free upgrade for existing license holders

26

# Q&A

- Thank you for attending! Please remember to submit feedback on our session!