PURESTORAGE®

Tech Talk

“Why Pure”

PURESTORAGE®

HPC/EDA & Pure Storage

Bikash Roy Choudhury

Director, HPC/EDA/DevOps Solutions- Pure Storage

Chip and Software Design Headlines & Opportunities

- Faster Time to Result
- Lower Costs
- Maximize Competitiveness
- Reduce Dev Cycle Times
- Scale & Grow as Business Dictates
- Improve Quality

Third Quarter Sales Hit By Delivery Delays

Micron Is First to Deliver 3D Flash Chips With More Than 200 Layers › 232-layer NAND makes tiny 2-terabyte products that deliver data 50 percent faster

HOME · COMPUTING · NEWS

Intel has 500 bugs to fix in its next supercomputer chips

Meta tells employees it will freeze hiring and restructure some teams in cost cutting effort

Micron Is First to Deliver 3D Flash Chips With More Than 200 Layers › 232-layer NAND makes tiny 2-terabyte products that deliver data 50 percent faster

3D-Stacked CMOS Takes Moore's Law to New Heights › When transistors can't get any smaller, the only direction is up

The Chip Shortage, Giant Chips, and the Future of Moore's Law › IEEE Spectrum's biggest semiconductor headlines of 2021

U.S. Passes Landmark Act to Fund Semiconductor Manufacturing › CHIPS and Science Act of 2022 provides billions for new fabs and other incentives

The First High-Yield, Sub-Penny Plastic Processor › It took a major redesign for cheap flexible chips to reach their promise

Nvidia's Next GPU Shows That Transformers Are Transforming AI › The neural network behind big language processors is creeping into other corners of AI

Autodesk to Launch Construction-Specific Cloud Collaboration Platform

Pure Storage Tech Talk "Why Pure"

EDA Design Trends

- Increased compute requirements on-demand
  - Faster job completion times
  - Reduce compute per Watt
- Massive data growth during chip design
  - Sub 10nm designs generate more data
  - Silicon sizes dropping to <20A˚ with multi-die
- Lots of consolidation in the industry
  - Increased mixed workloads and environments
  - Shift towards AI-designed commercial tape-outs
- Transition into Hybrid Cloud for design workloads
  - Compute, Compute and more Compute for complex designs
  - Scalable performance and data continuity on-premises/ Cloud

https://www.intel.com/content/www/us/en/it-management/intel-it-best-practices/data-center-strategy-paper.html
https://www.intel.com/content/dam/www/central-libraries/us/en/documents/semiconductors-and-intel-introduction.pdf

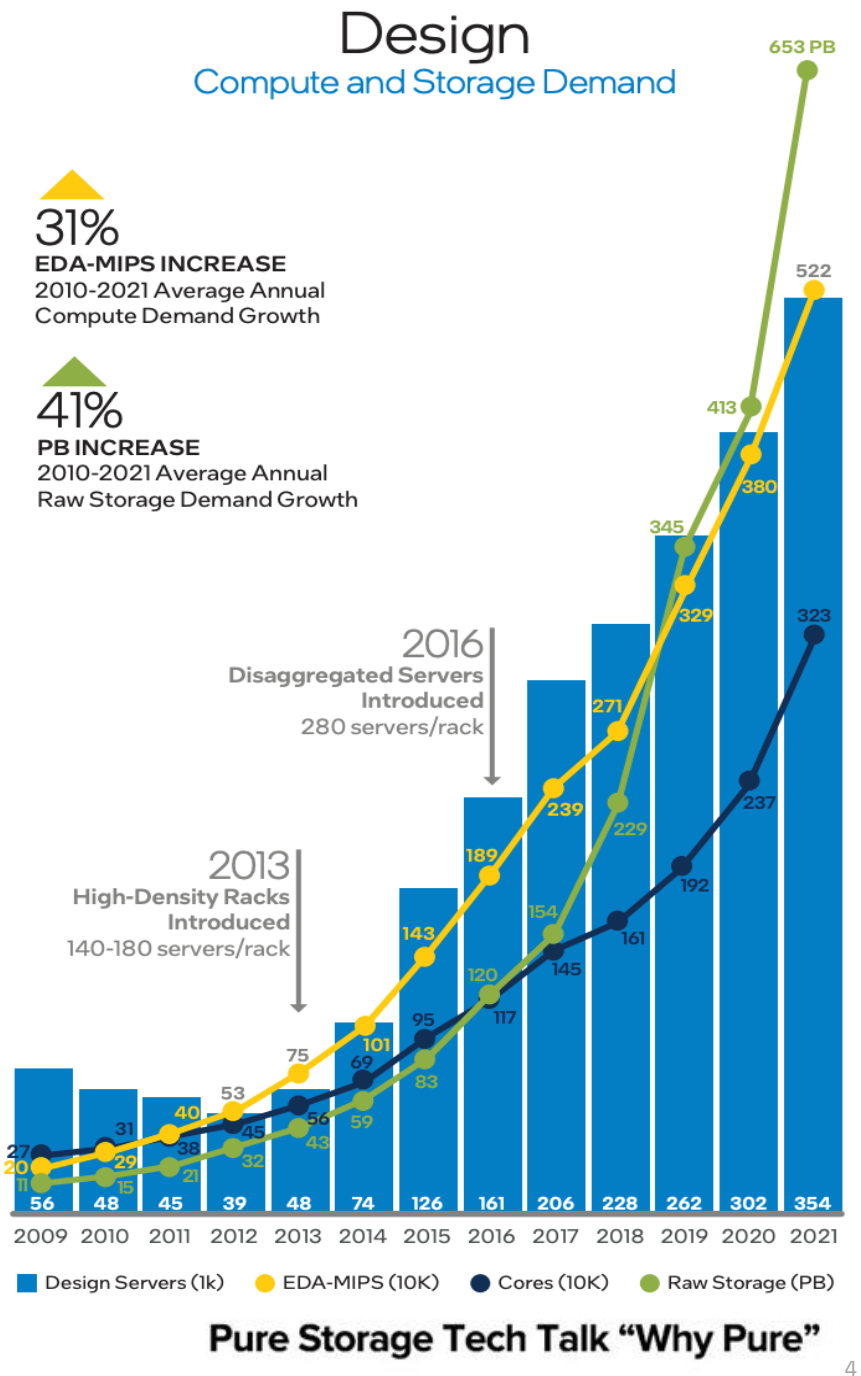

PURESTORAGE®

Factors driving Chip Design Costs

Challenges

- Design complexity w/higher EDA Tool license Costs
- Slow design turnaround time impacting TTM
- Long running jobs with limited scalability
- Slow infrastructure with limited scalability - compute, network, storage

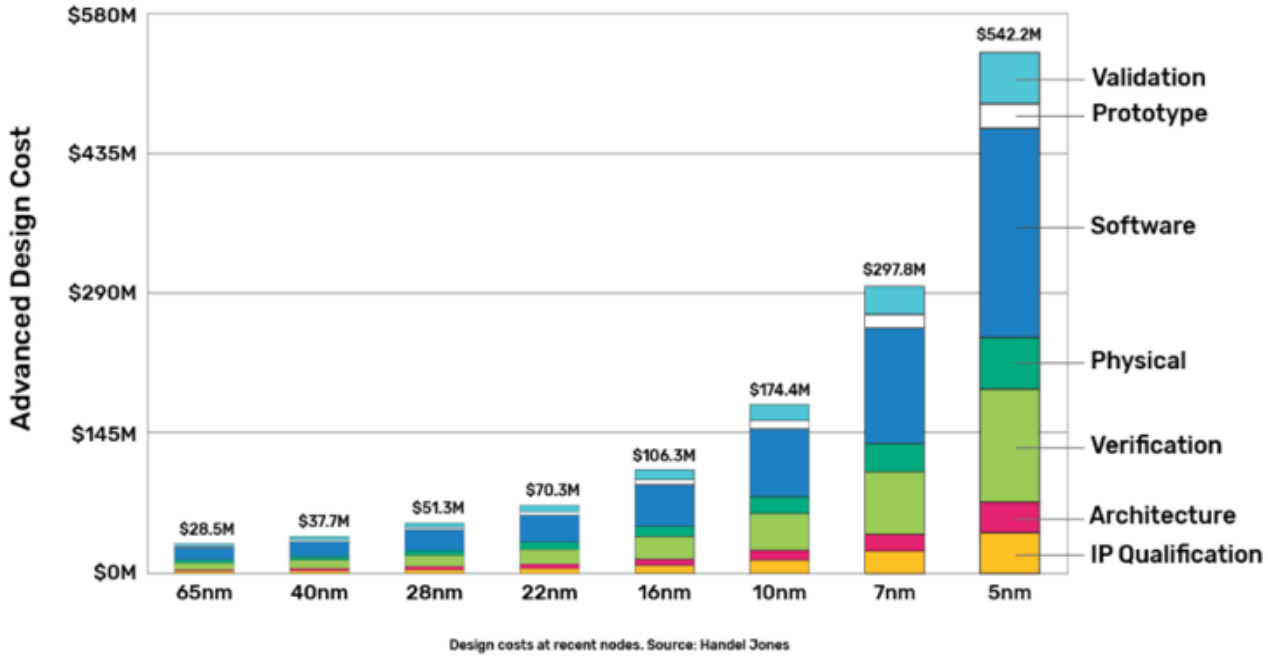

Design costs at recent nodes. Source: Handel Jones

Source: https://www.chipestimate.com/How-to-address-SiP-challenges-with-EDA-tools-and-IP/Cadence/Technical-Article/2020/09/08

EDA Businesses Needs

Scale more cores in compute farm for incremental growth

Enable growth with headroom and non-disruptive upgrades

Deliver designs faster with reduced job completion times.

Keep your pipeline of designs flowing, optimize (or lower) EDA Tool License Costs.

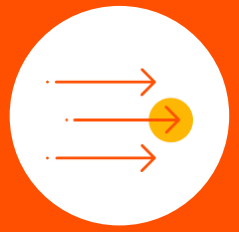

Accelerate SW Delivery Pipeline (Decrease TTM)

Improve developer productivity with increased software quality

Provide scalable and efficient performance

Scale storage and compute independently (disaggregated) and reduce data center footprint

Simplify Storage Management

Monitor and reporting of EDA Infra end-points

Mobilize data for hybrid cloud

Enable array/host based data mobility and build enterprise hybrid cloud solutions with no compromises

EDA Workflow: Potential Bottlenecks

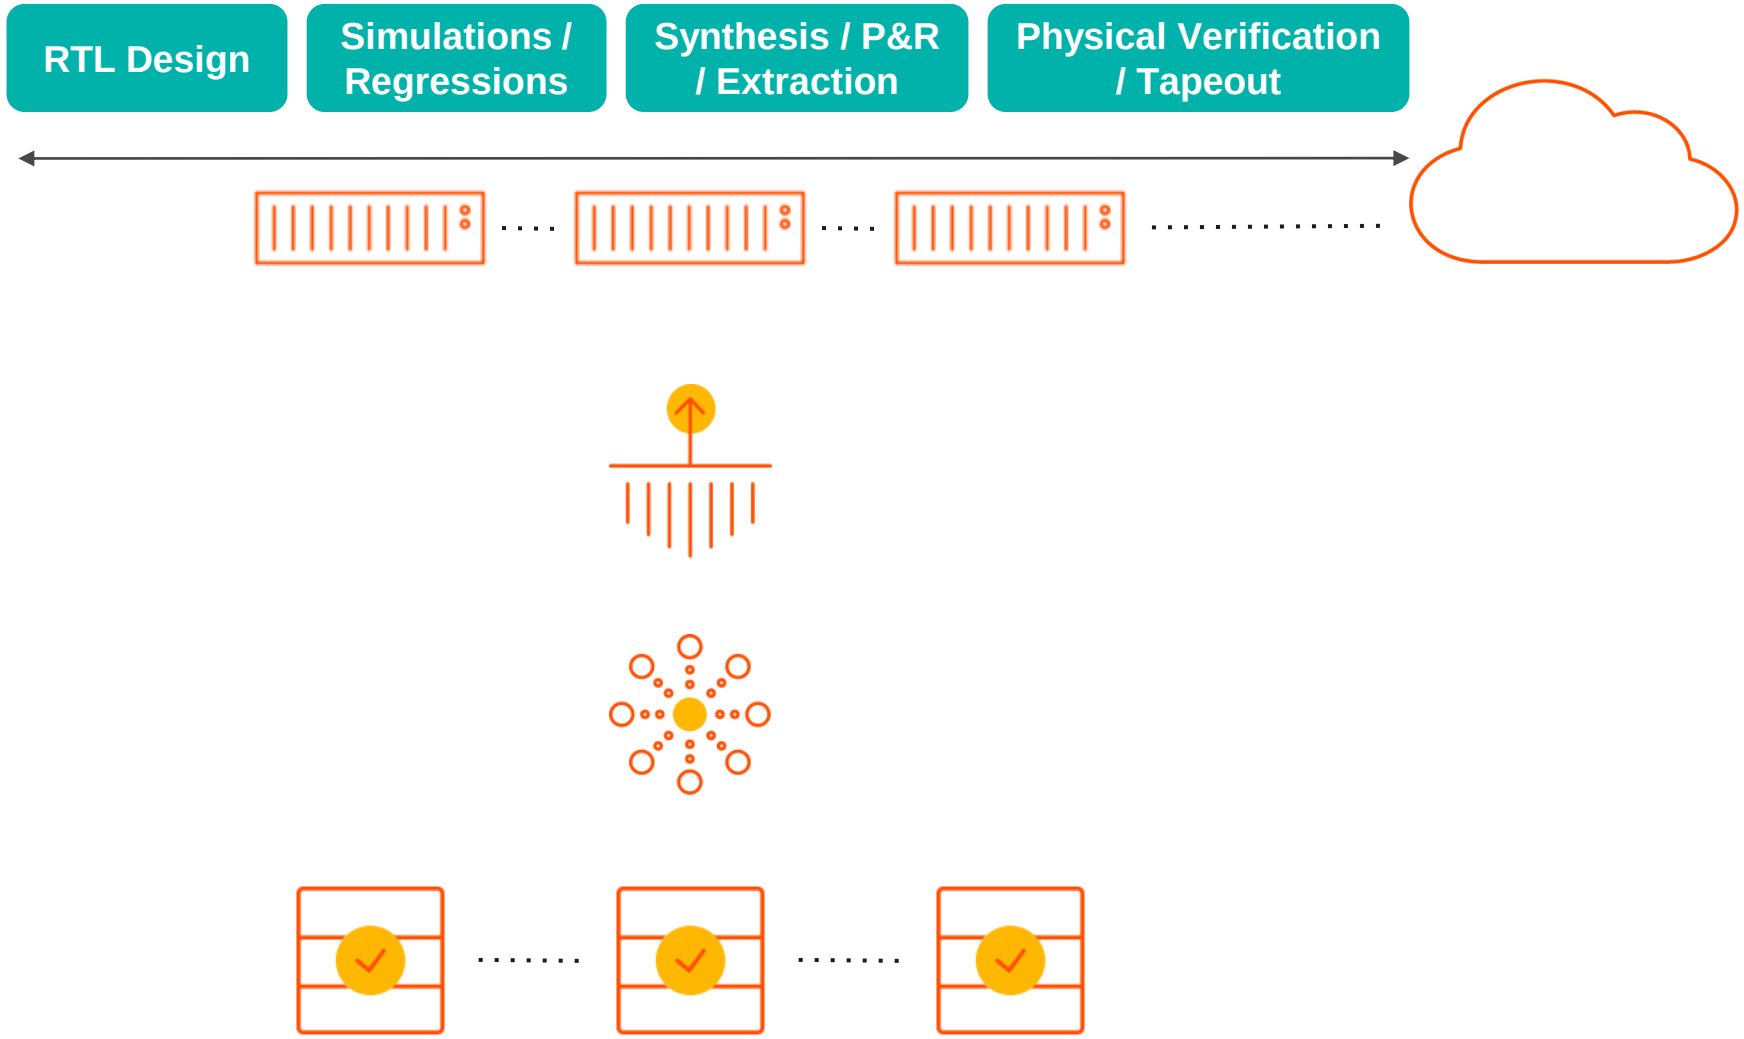

COMPUTE
(Scale on-demand - Flexibility to extend into cloud)

SCHEDULER
(24/7 Queue - Resubmit failed jobs - not storage aware)

NETWORK
(IO bottleneck/ Saturating network from compute hosts to Storage)

STORAGE
(Concurrent Data access at scale - variable/competing workloads)

PURESTORAGE®

Accelerate Chip Design Workflows and Tools

Shorten Chip Design Cycles in EDA

Major EDA Provider	>10X Build run-time Improvement
Top 10 Semiconductor	Enable > 200M more simulations every year
Social Media AI	~50 GB/s throughput for deep learning
Top 10 Semiconductor	Reached 70 GB/s and 2.3 million IOPS
Major EDA Provider	40% Questa regression runtime improvement
Top 10 Semiconductor	60%+ PrimeTime runtime improvement
Top EDA Provider	Achieved 1.4 million IOPs during simulation phase
Various EDA Companies	RapidFile Toolkit results in up to 60x speed up of millions of file operations

Consolidate Semiconductor Workloads to Accelerate Chip Design Workflows

CHIP DESIGN WORKFLOW

SW Dev Version Control	Build & RTL Regressions	Design/Analog Simulations	Synthesis/ P&R	Timing/ Power/Noise	Signoff/ Tapeout

COMPUTE LAYER – ON PREMISES / CLOUD

JOB SCHEDULER – PBS/LSF/UNIVA

CONSOLIDATED WORKLOADS ON FLASHBLADE

Verilog Source Code	Build Binary Packages	Scratch	Project Data	EDA Tools	IP Libraries	Design Data

NFS

- Parallel Access Tested in > 500K Core Environments
- Metadata >30 Million IOPS
- Bandwidth >300 GB/s
- Fast Delete >50K files/sec
- High File Count 2^56
- Capacity 19+ PB (2:1)

PURESTORAGE®

Reduce Logical Verification time in EDA Process by 50%

FlashBlade vs. FlashBlade//S over NFSv3

- Increase IOPs for metadata only workloads by > 2.5x
- Reduce overall latency by 50%
- Scale more compute cores by 50%
- Increase data reduction by > 30%

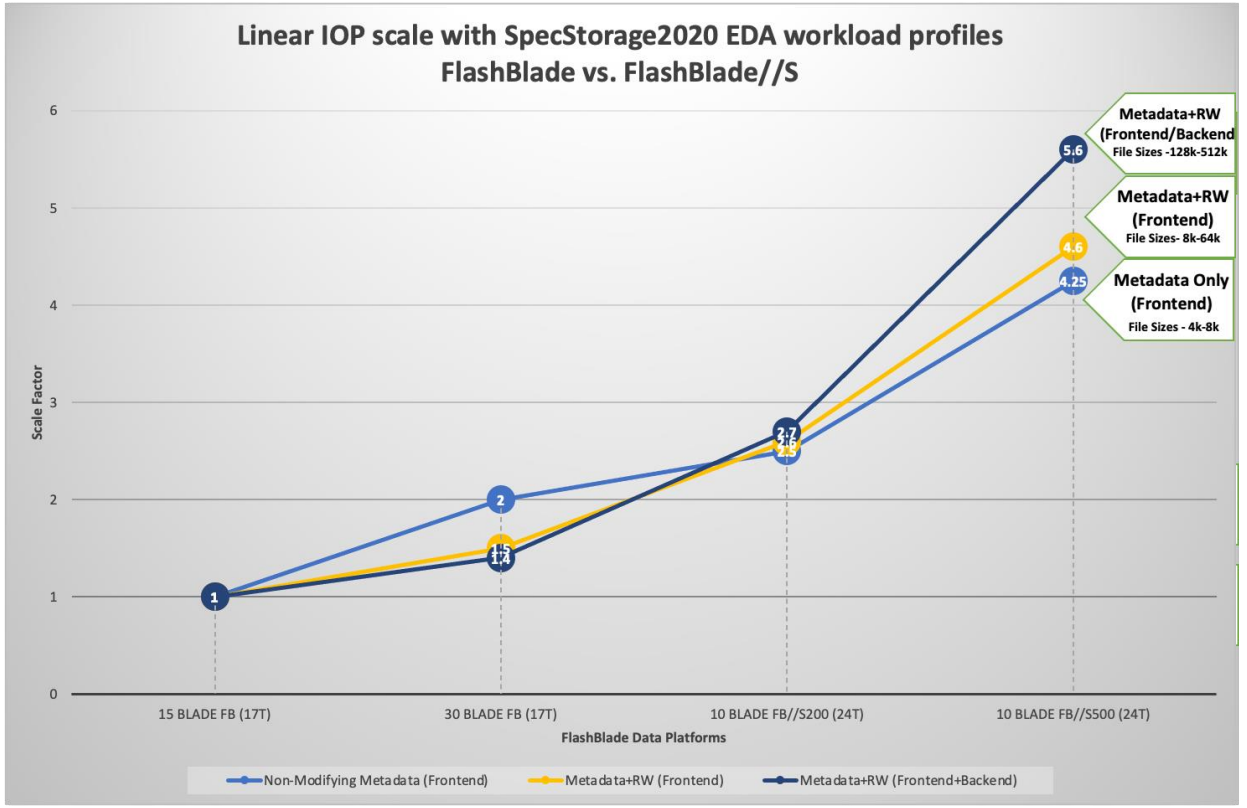

Reduce Regression Time During Tapeout by 50%

FlashBlade vs. FlashBlade//S over NFSv3

- Increase bandwidth by 75%
- Reduce write latency by 60%
- Scale compute cores up to 50% more

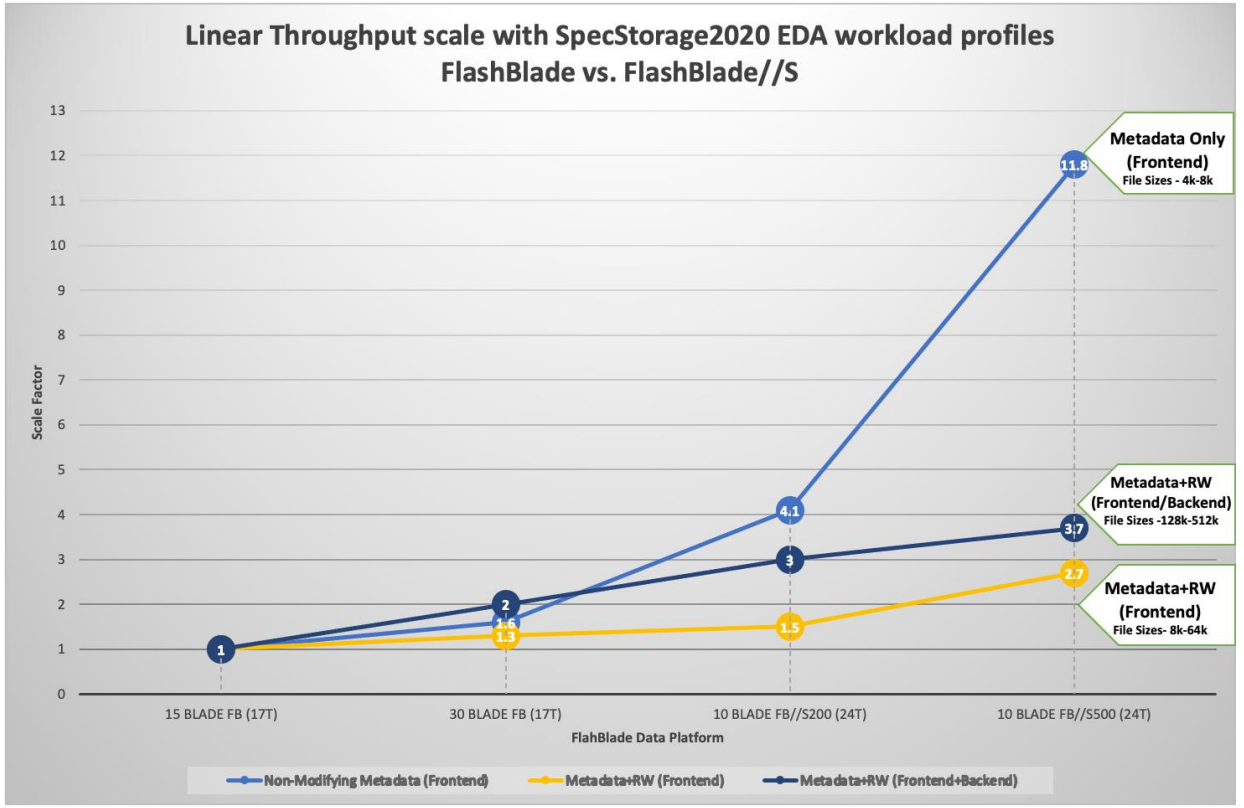

EDA Generational Performance Better by 4X

FlashBlade//S: IOPS vs Latency

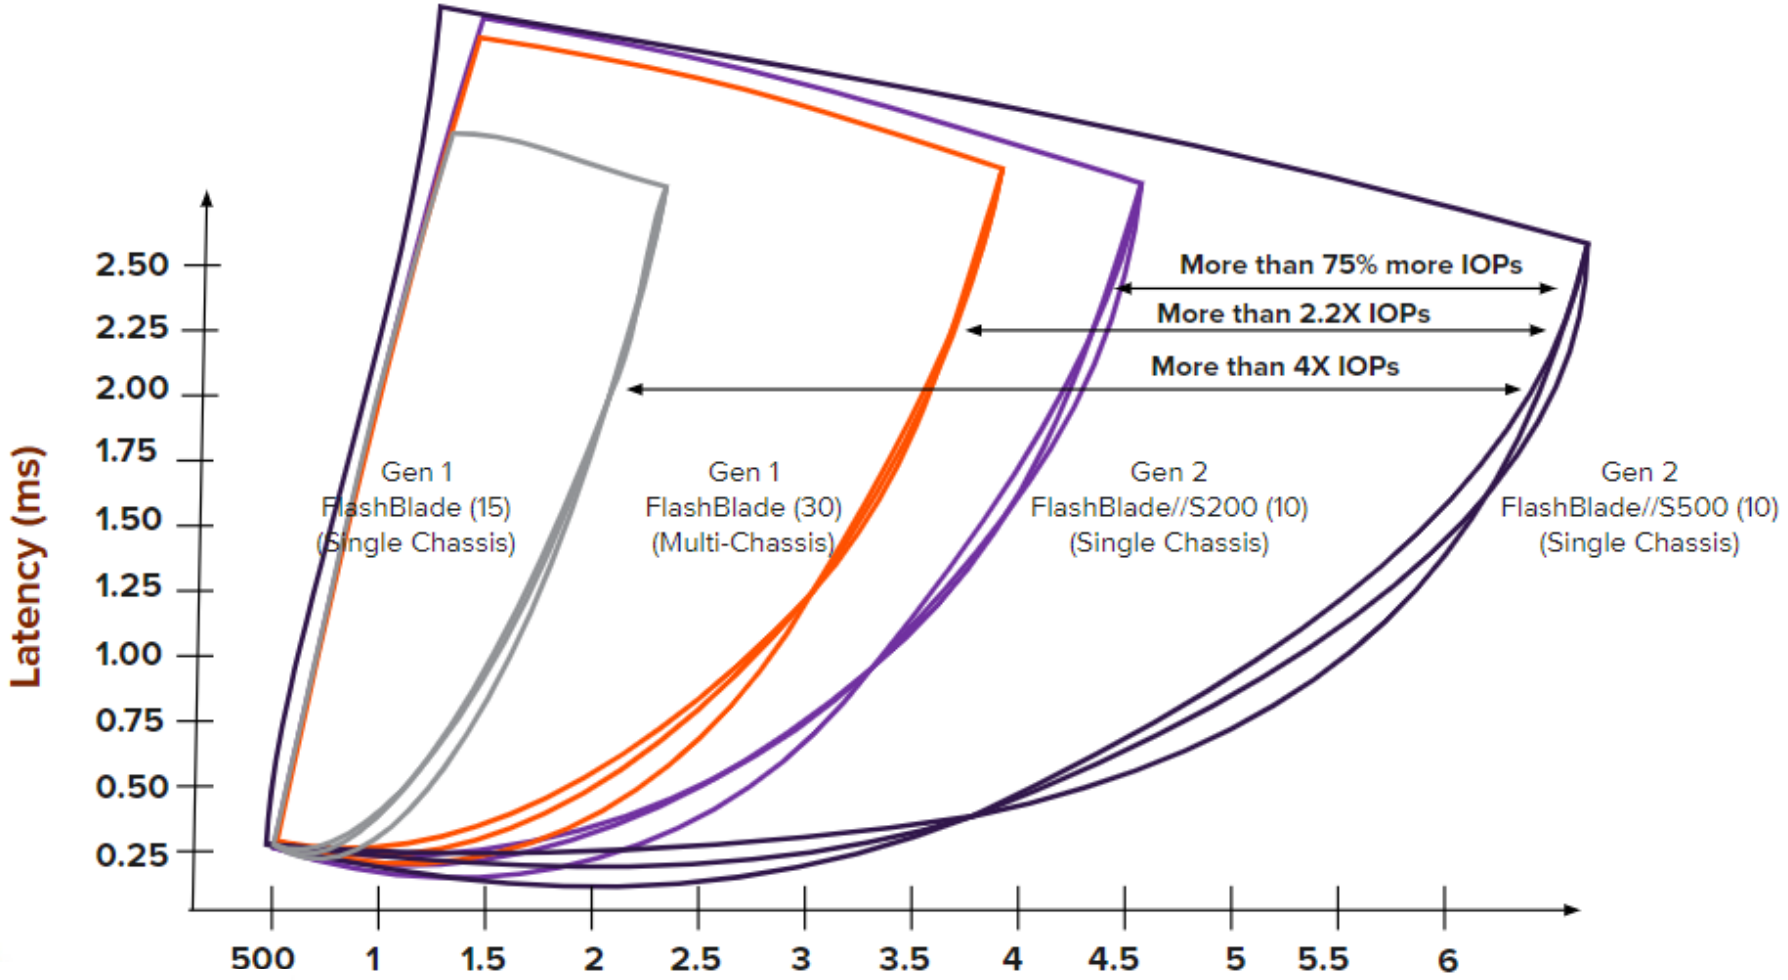

EDA Generational Performance Up to 3X Better

FlashBlade//S Bandwidth vs Latency

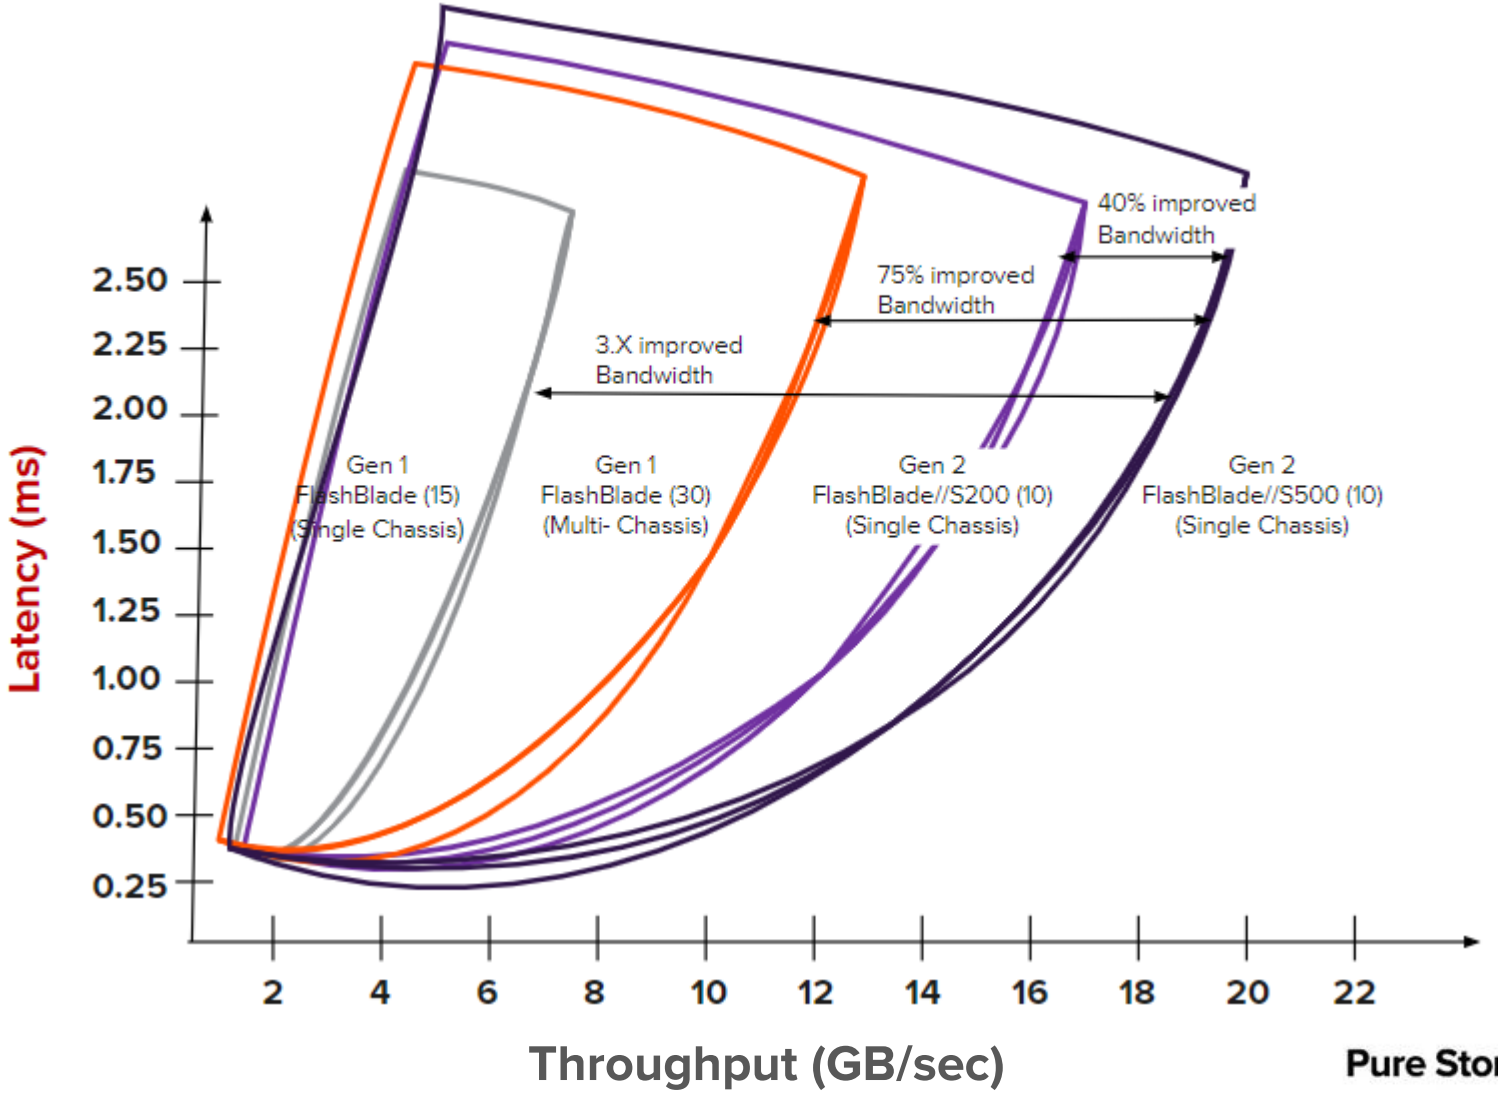

PURESTORAGE®

Pure Storage Tech Talk “Why Pure”

Consolidate Semiconductor Design Workflows with FlashBlade//S

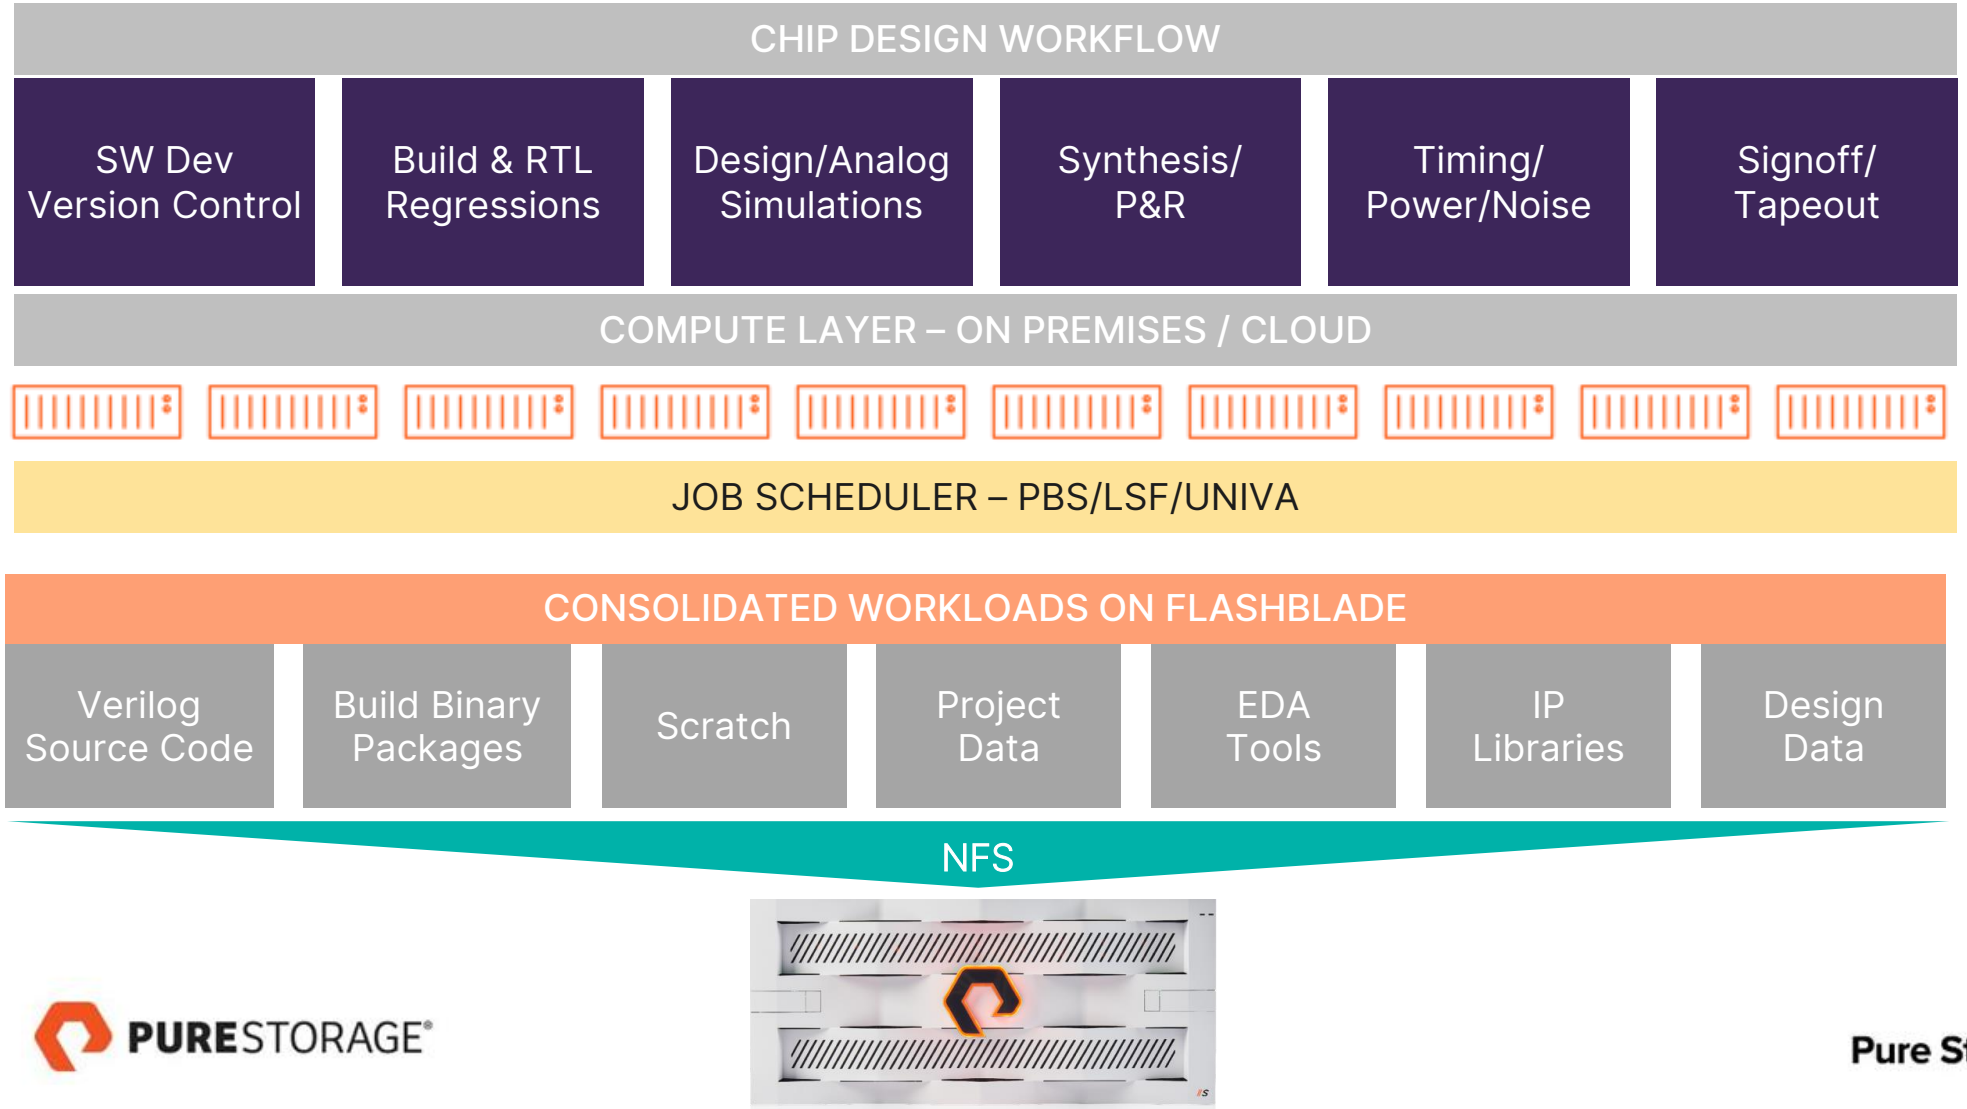

Post-production challenges

- Unused design data on Flash Storage
- Restricting onboarding new projects
- Inefficient capacity & energy usage
- Reduced ROI

PURESTORAGE®

Pure Storage Tech Talk “Why Pure”

Data Lifecycle with FB//S and FB//E

Array and host based data mobility

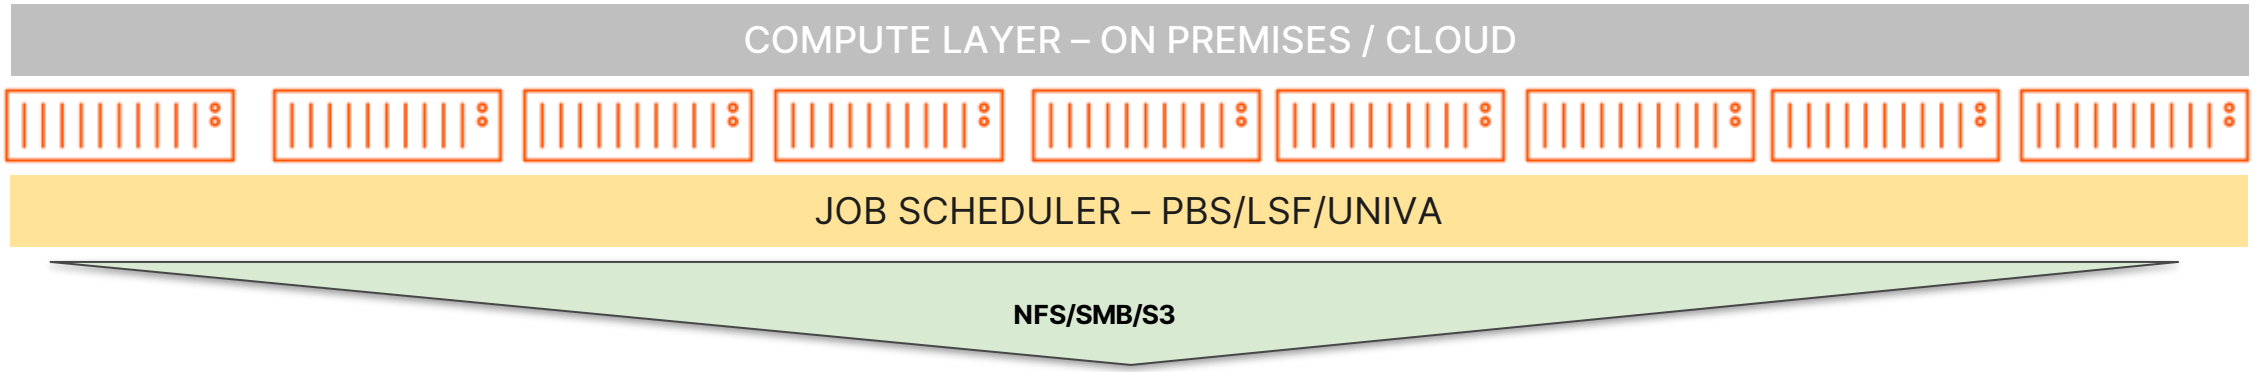

Active Design Projects

Scalable Performance

Content Repository

Scalable Capacity

Frozen data

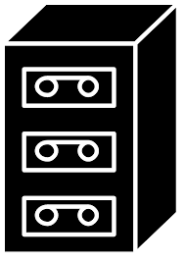

Tape Library

Native Array Level Data Mobility from FB//S to FB//E

Snapshot based Bi-directional file system Array level replication

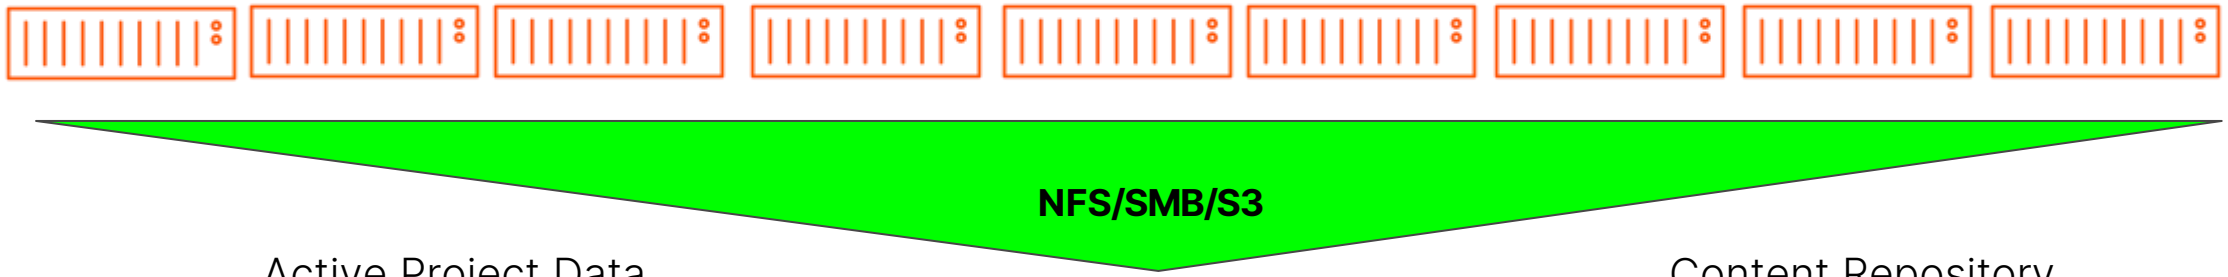

Host based select Data Mobility with Rapid File Toolkit (RFT)

Using RFT/ptar/pcopy to move select files/directories from //S to //E

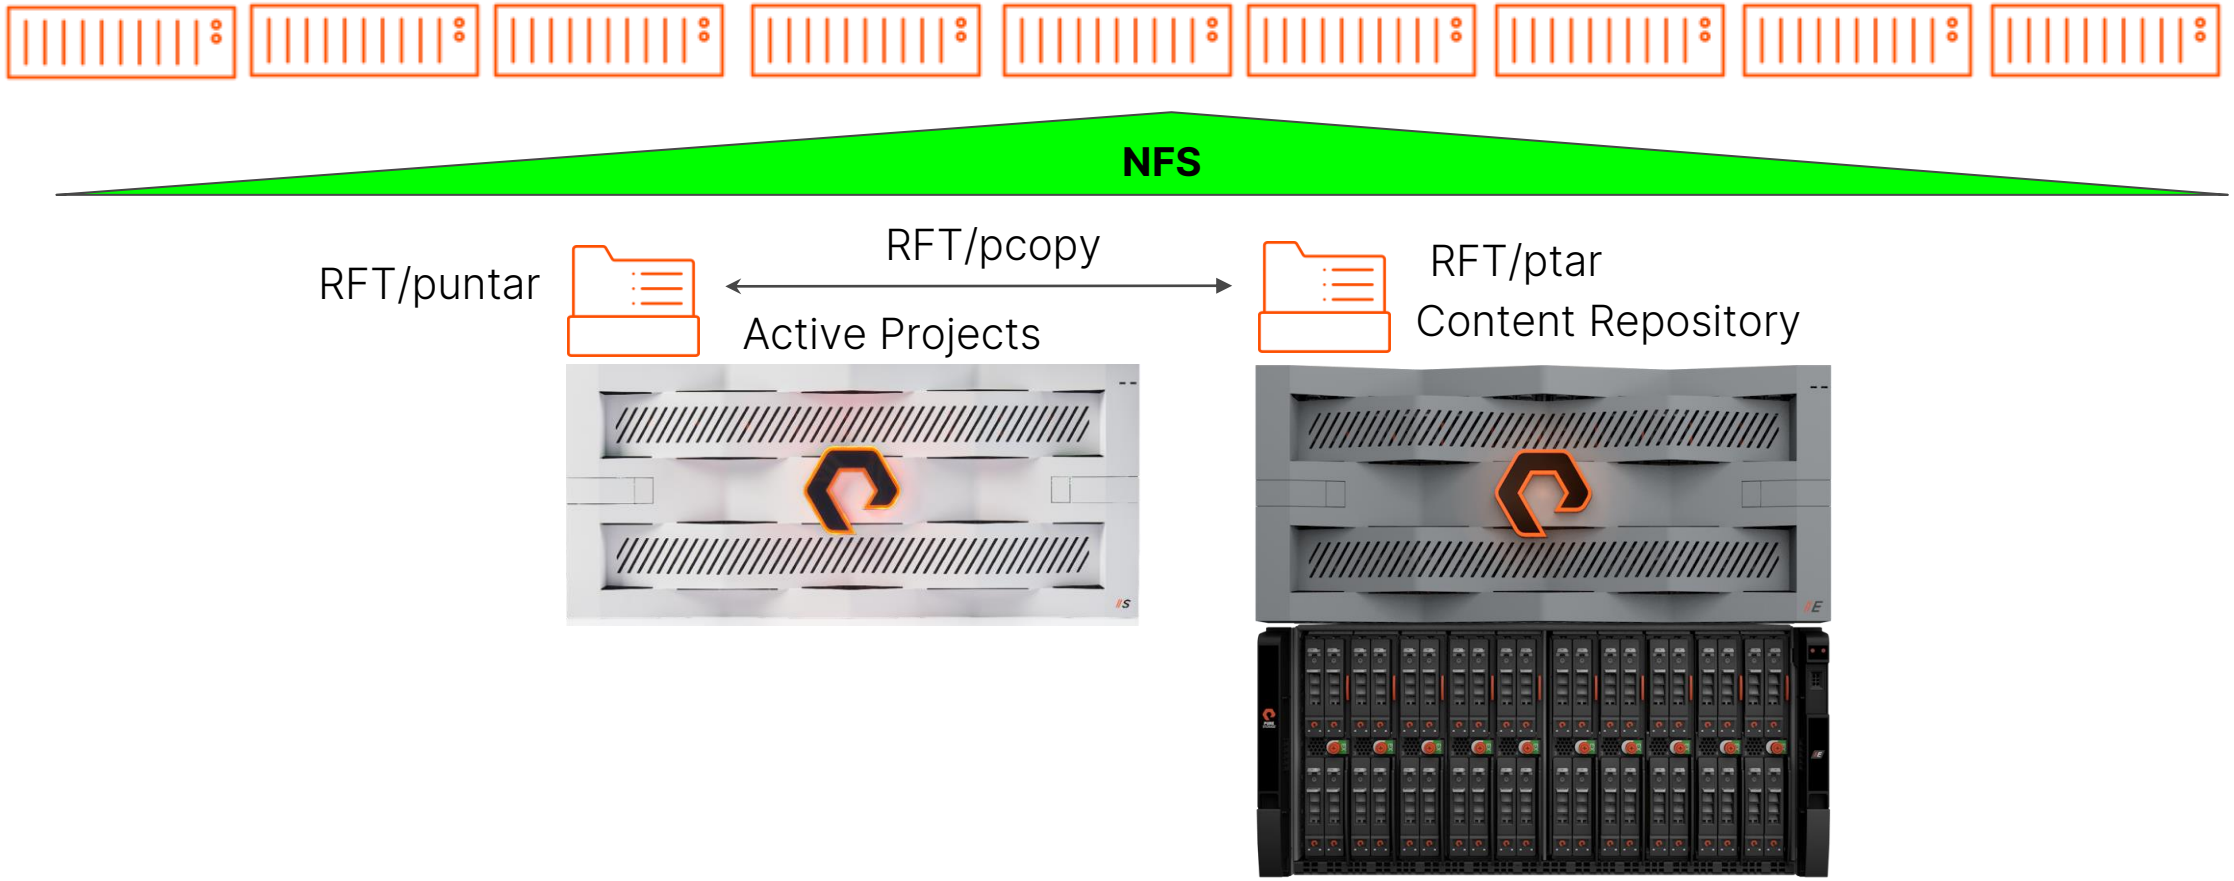

Data Lifecycle for Chip design workloads

Benefits with Pure Storage FB//S and FB//E

- Active project data can move to archive storage after the project is over.
  - New projects can start after old data is moved to FB//E and the source file systems are removed on FB//S.
  - Common tools and libraries continue to move the changes to FB//E until all the collaborative projects are over.

- Design data, tools and libraries can be moved to FB//E at the array and host level after project completion
  - Snapshot based bidirectional file system replication allows to move the entire filesystem to FB//E.
  - Host based tool like RFT provides the ability to move select files and directories to FB//E faster.

- Chip design projects and data can be archived on all flash arrays for better cost efficiency
  - Data continuity with FB//E allows data access over standard file and object protocols
  - The FB//E is not designed for scalable performance; performance is limited

Smart job scheduling with LSF/ELIM using FB/APIs

Benefits with Pure Storage FB//S and FB//E

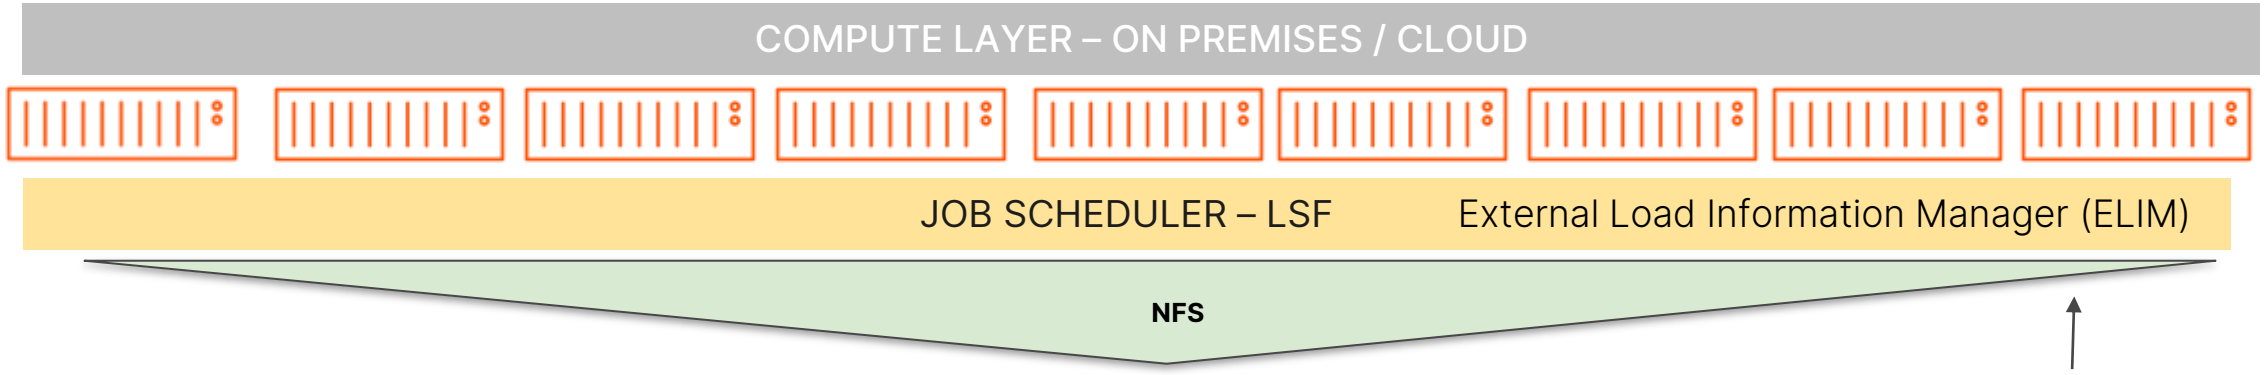

Active Design Projects

```
fb = PurityFb('xx.xx.xx.xx')
try:
   vault_pass = subprocess.check_output("vault kv get -tls-skip-verify -
field=apitoken_hyd teams/radops/API_token", shell=True)
   res = fb.login(vault_pass)
except rest.ApiException as e:
   print("Exception when logging in to the array: %s\n" % e)
if res:
  try:
   usec_per_read_op = fb.arrays.list_arrays_nfs_specific_performance()
   usec_per_write_op = fb.arrays.list_arrays_nfs_specific_performance()
   print("2 usec_per_read_op
{}".format(usec_per_read_op.items[0].usec_per_read_op),end=" ")
   print("usec_per_write_op {}".format(usec_per_write_op.items[0].usec_per_write_op))
  except rest.ApiException as e:
   print("Exception when listing array performance: %s\n" % e)
time.sleep(8)
```

- Control job submission in the queue when the storage is busy.
- Optimizes EDA tool license costs by limiting job submission to the queue beyond a threshold array latency value.
- Integrating FB/APIs with ELIM smoothens the job submission to the queue and preemptively identifies storage saturation.

EDA Workflow Solutions

Pure Storage Portfolio offerings for semiconductor design and manufacturing organizations

Performance and Capacity @ Scale	Cloud	Data Mobility	Observability

PERFORCE

SYNOPSYS®

JFrog Artifactory

git

Grafana Prometheus

PURESTORAGE®

Pure Storage Tech Talk “Why Pure”

Connected Cloud Software Dev & EDA Workflows

- Improve Engineering productivity
- Faster data recovery from failures
- Scalable performance and capacity
- Optimizing EDA tool license costs
- Cost efficiency with improved Data reduction ratio
- Data security and sovereignty with Hybrid Cloud

- Faster Software Delivery
- Accelerate design verification
- Managed EDA workflows
- Cloud bursting for EDA workflows

On Premises

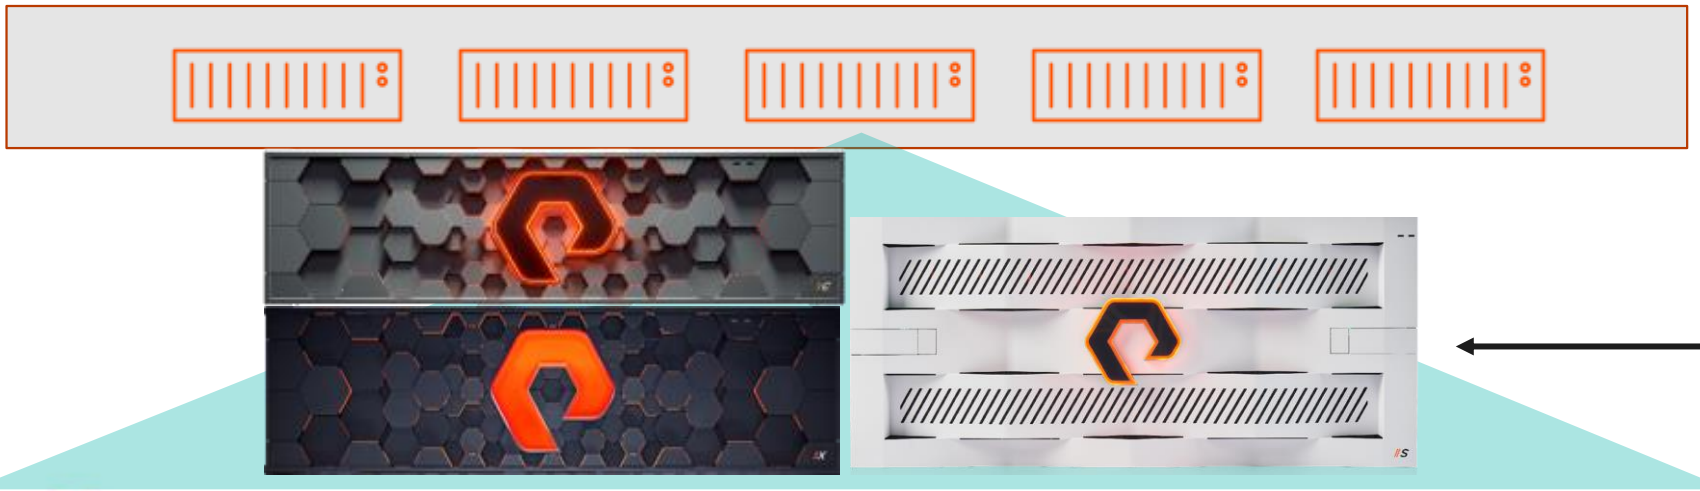

Hybrid Cloud

EQUINIX

PURESTORAGE®

Pure Storage Tech Talk “Why Pure”

PURESTORAGE®

Thank You!

홈페이지 : **www.purestorage.com/kr**

페이스북 : **www.facebook.com/purestoragekorea**

블로그 : **www.blog.naver.com/purestorage_korea**

유튜브 : **www.youtube.com/@PureStoragekr**

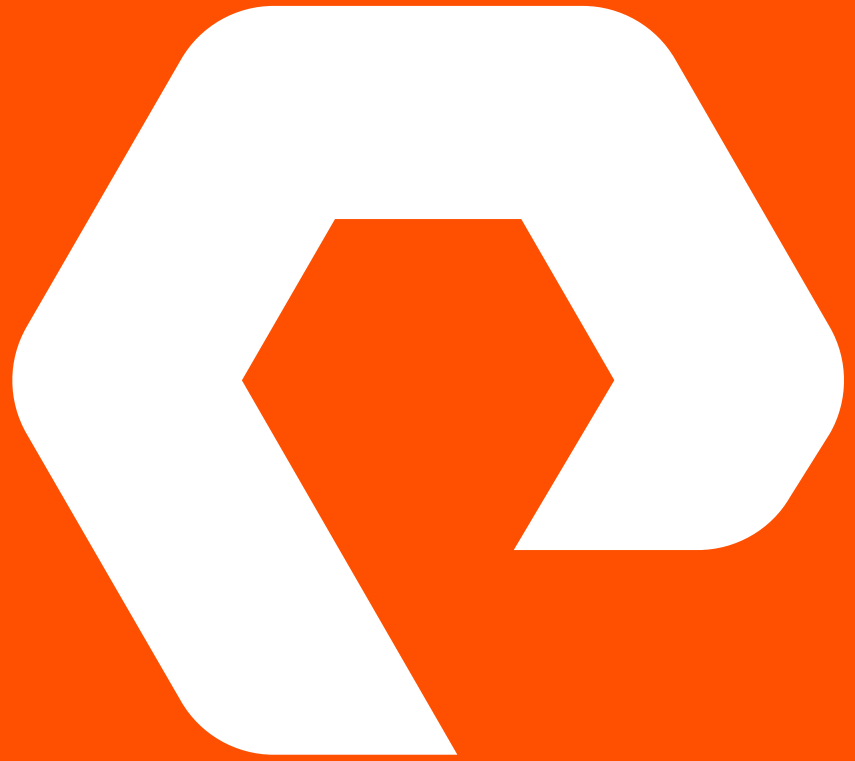

Uncomplicate Data Storage, Forever