



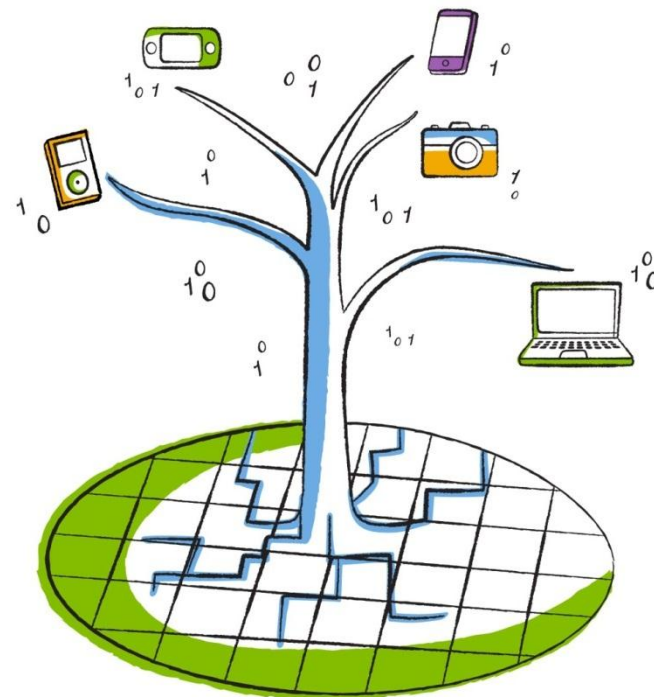
NetApp®

Go further, faster®

# Accelerating Software Releases With Continuous Delivery

Kumaraswamy Namburu


April 17, 2014





# Market-Leading Storage Solutions

## *Shared Storage Infrastructure*



**Clustered Data ONTAP®**  
for Shared Infrastructure

## *Dedicated Storage Solutions*

**Flash Arrays**  
for ultra-high performance

**E-Series Systems**  
for price/performance at scale

**StorageGRID®**  
for web-scale object storage

# Sound Familiar?

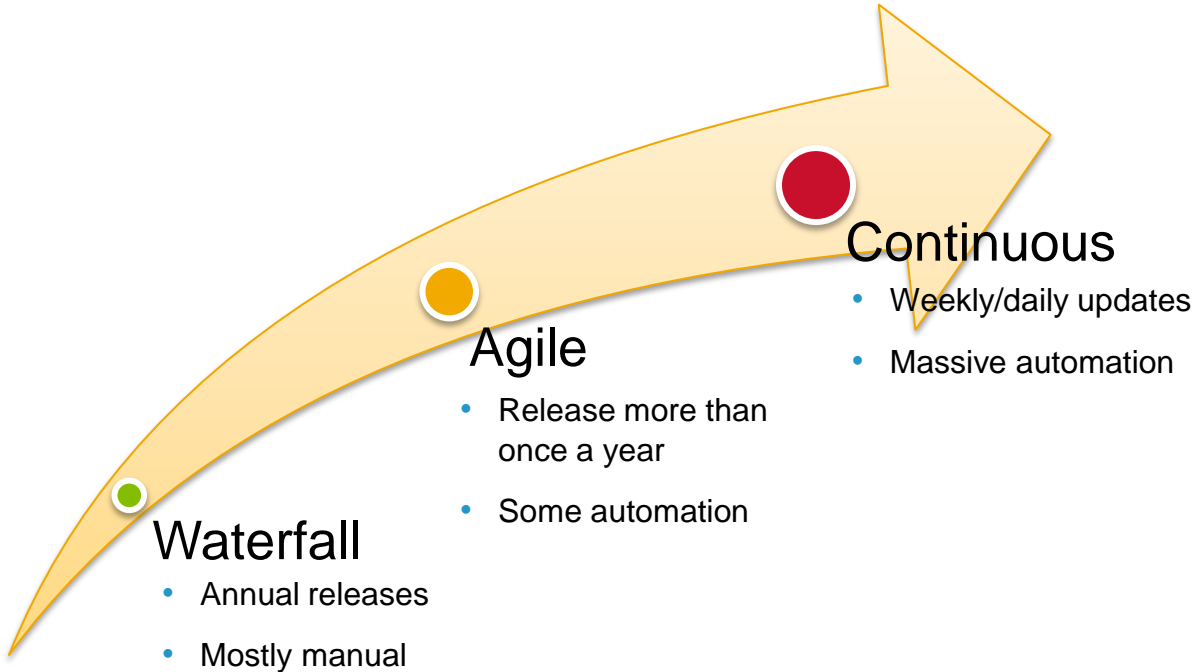




# Perspective & Environment

- 3,722 Developers
- 9 sites
- 25 Million Lines of Code
- Single Development Branch
- Continuous Delivery

# Accelerating Product Delivery is Critical



---

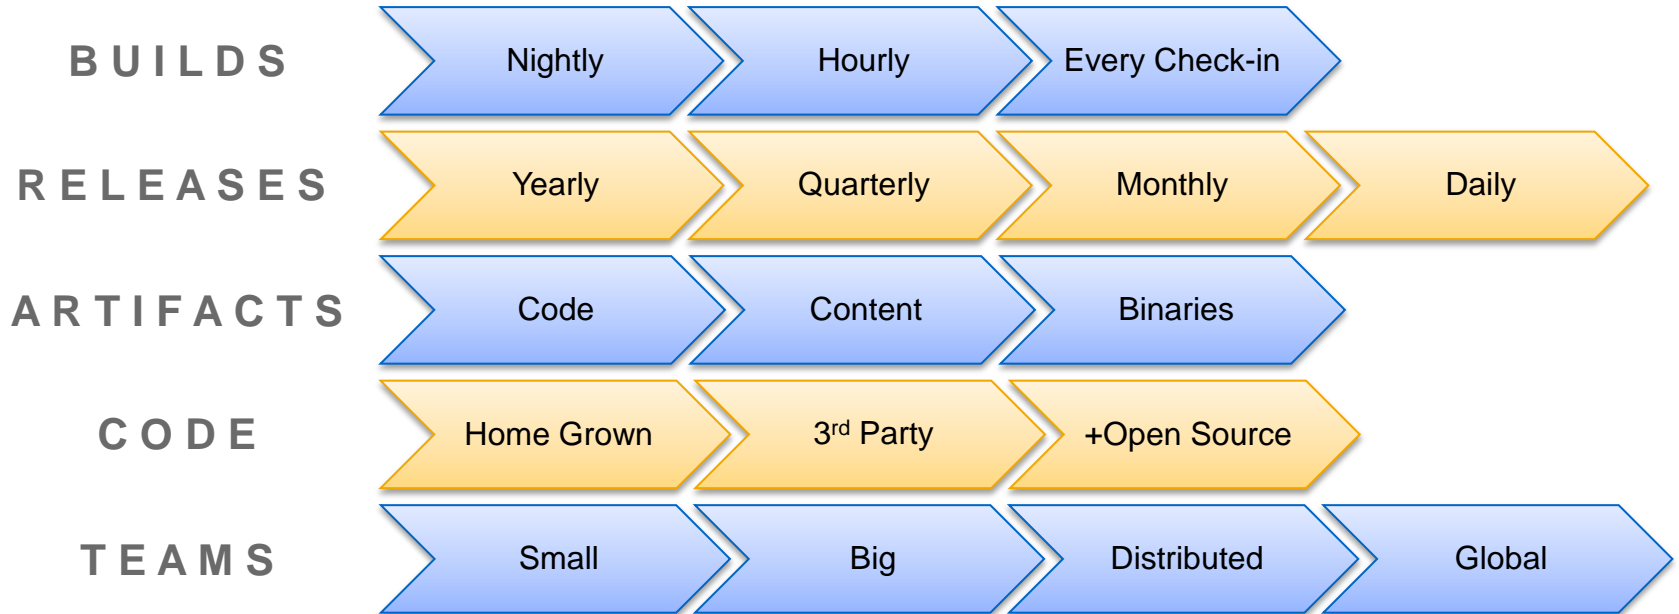
“The days when a successful organization could release software once every 12 to 18 months are over.”

---

“Continuous Delivery is Reshaping the Future of ALM,”  
Kurt Bittner, Forrester, July 2013



# Accelerated Delivery Means...





# 1000 foot view of NetApp Build

**PERFORCE** **SCM**

- Source Code in Perforce

**IBM** RedHat 6.4 **CISCO** **Compute**

- Diskless Clients ( NFSboot from filers)
- In-house developed distributed Engine
- NFSv3 mounts for build workspaces

**NetApp™** NetApp Storage (6280 filers) **cDOT (8.2.x)**

- Faster workspace Snapshots/ Flexclones
- SFO (fail over) & Junction paths - cDOT
- Diskspace footprint reduction with Flexclone

**NetApp™**  
Storage Tier

**File System – NFSv3**

**Storage Efficiency – Acceleration – Collaboration**

- FlexClone**
- Compression**
- FlexVol®**
- Deduplication**
- Flash Cache**
- 

**Data Protection**

- SnapMirror**

**High Availability Scale Up / Scale Out**

- cDOT, cDOT 8.2**

**Backup**

- SnapVault, Volume SnapMirror**



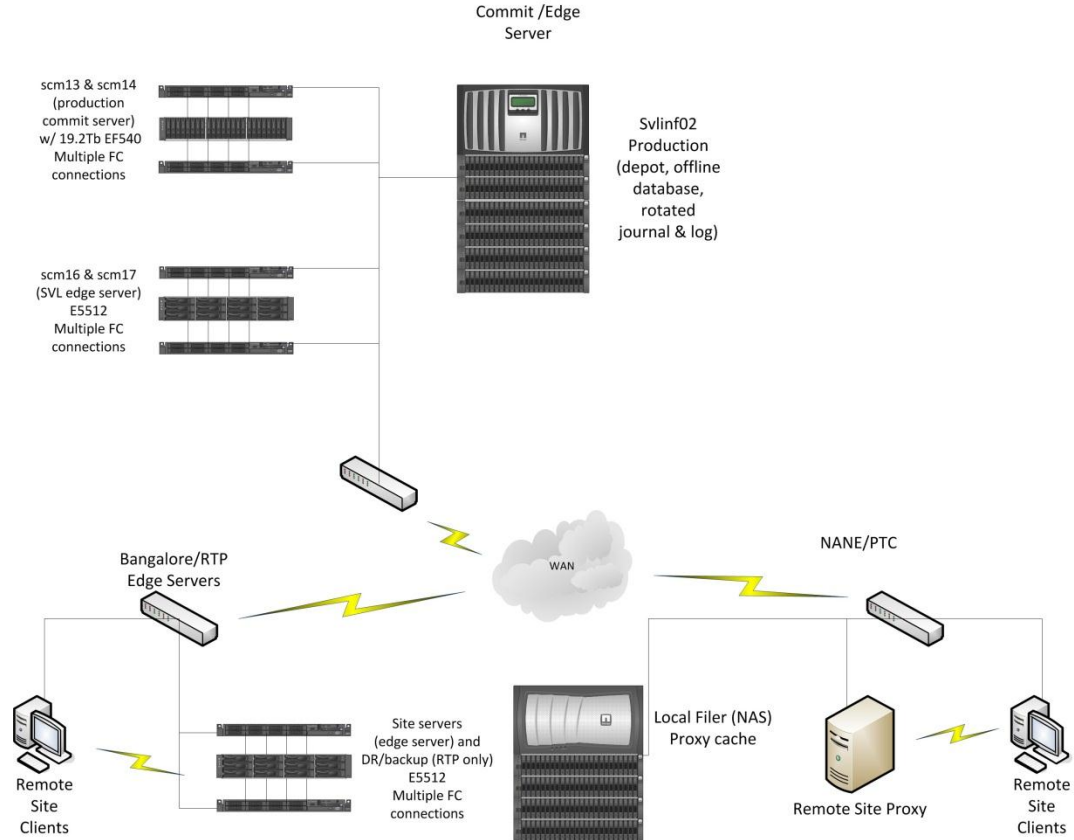
# Perforce details

- One central P4d instance
  - Database size: 980G
  - Daily journal size: 44G
  - Commands run daily: 1.7M
  - Users: 3,722
  - Depots: 570G
  - Clients: 246K
  - Four proxies distributed geographically





# Perforce Topology @NetApp





# Perforce database locking

## ■ Problem Statement


- Perforce metadata is stored in 40 Berkeley DB format files in P4ROOT
- The various processes invoked on the P4 server for each command locks these DB files (as needed) to ensure data integrity.
- This reduces Perforce **concurrency** and reduces overall performance
- For large Perforce sites, this is the [single biggest factor affecting performance](#).

## ■ Solution

- Deploying EF540 ( all-SSD enterprise storage) in this environment helps to solve this problem



# Performance Results (EF540 Vs local DAS SSD)



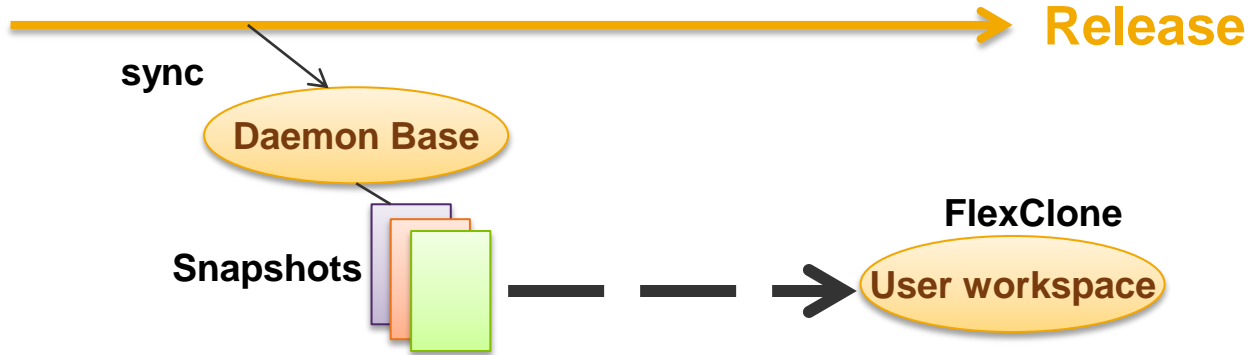


# Perforce Operations

		Before EF540 (local SSD) (seconds)	After EF540 (seconds)	Degree of Improvement
P4 flush	SVL	35.5	2.1	16x
	RTP	37.7	2.1	18x
	PTC	40.7	2.2	18x
	NB	35.3	3.1	11x
	Average*	37.3	2.375	15x
WS_merge	SVL	19.0	4.9	3x
	RTP	42.8	7.5	5x
	PTC	26.9	11.0	2x
	NB	63.6	39.5	1.5x
	Average*	38.075	15.725	2.5x
P4 resolve	SVL	424.1	122.7	3x
	RTP	271.9	2.5	100x
	PTC	388.5	5.0	75x
	NB	443.9	20.7	20x
	Average*	382.1	37.725	10x

**Note:** Perforce Server located in SVL  
Average \* = All sites are equally weighted

# Faster workspaces from SCM



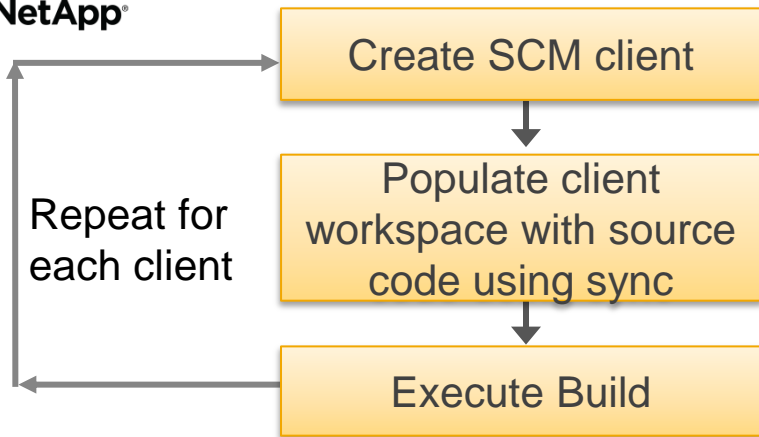
## ■ Daemon

- A SCM client for each codeline that is setup at the time of the codeline creation.
- Process involves syncing this client to pick up latest changes on the codeline, and run the needed builds.
- Once the build is complete, a snapshot of the client is preserved.
- By using these snapshots as their client base instead of creating individual workspaces, developers are able to get populated, prebuilt workspaces in matter of minutes

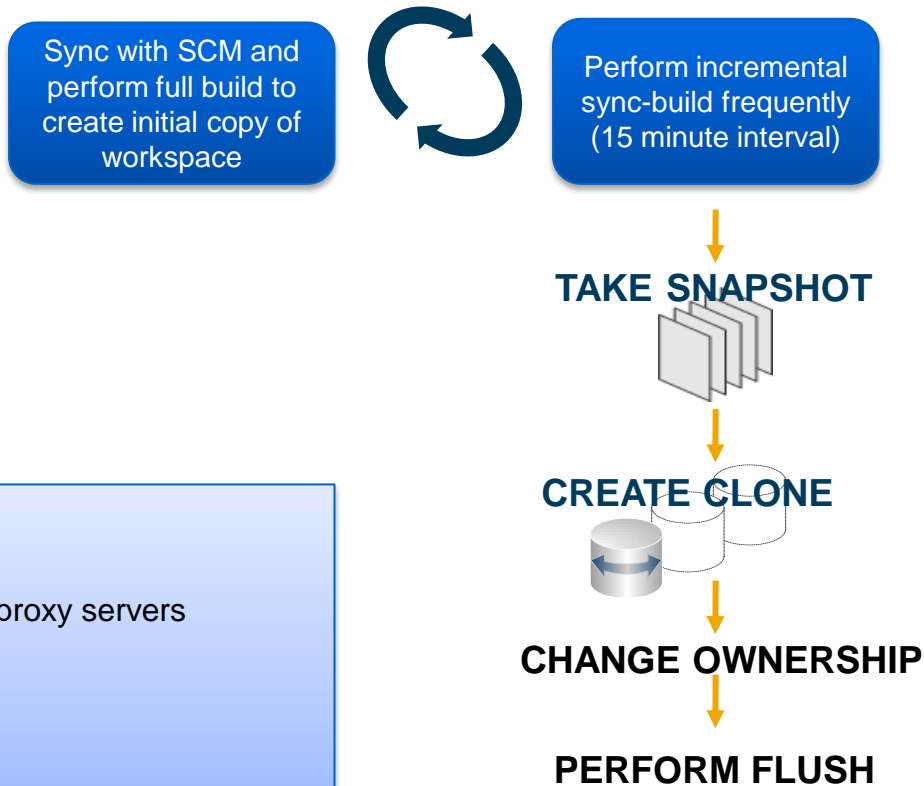


NetApp®

## Traditional



## Optimized





### Benefits of NetApp Technology

- Workspace availability in minutes (<5 mins) vs. hours
- P4 sync vs P4 flush reduce the workloads on Perforce & proxy servers
- Leverages NetApp FlexClone capabilities
- Debug capability reduced to minutes
- Reduced disk consumption
- Improved consistency and stability in process
- Measured against developer time, multi-million dollar benefit



# Before



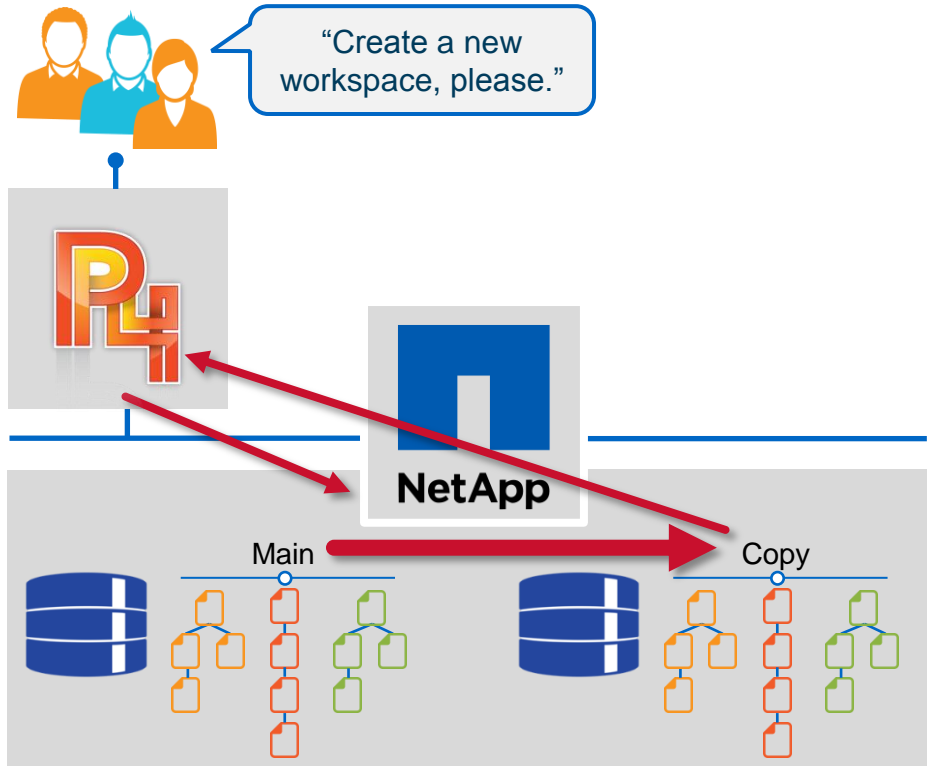
Stored in file system structure; not a "blob"



- Perforce reads/writes the files, as needed
- Various interfaces – network, fiber channel, etc.



# After

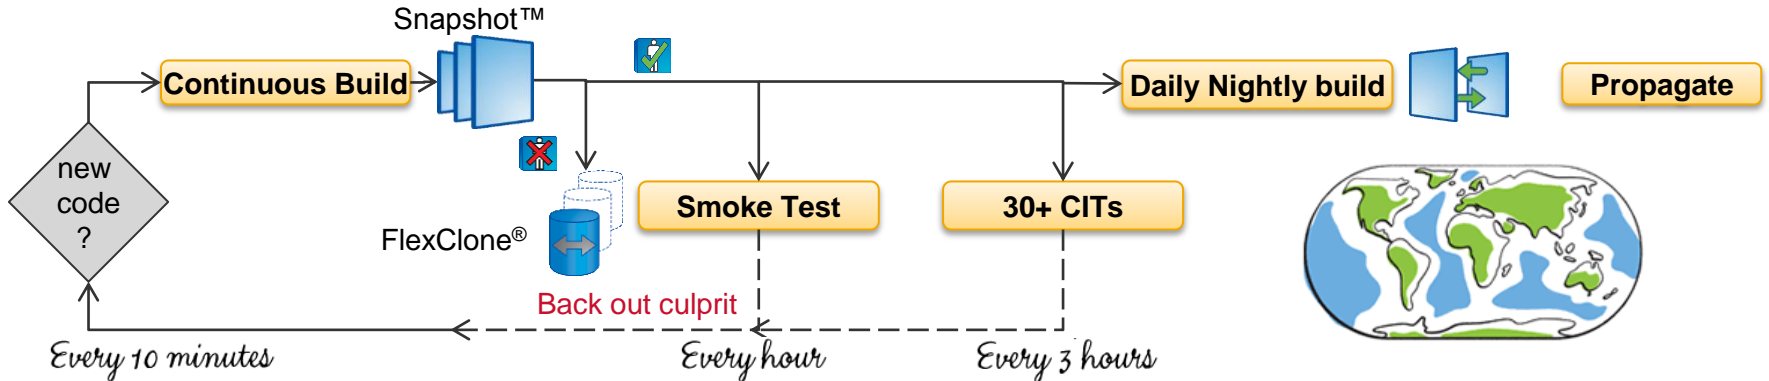



- NetApp executes the move
- Perforce is updated when it completes
  - Updates Perforce database
- Benefits from both architectures
  - Perforce's use of standard files
  - NetApp's file handling capabilities





# Continuous Integration





# Accelerating Innovation

## Standard workspace methods

80 minutes

## Hard links

40 minutes

## Workspace cloning


- 2-3 minutes
- No build necessary
- Minimal data transfer work
- **Saves NetApp 100 person years/year**





# Continuous Integration Benefits

- One development branch
- Continuous delivery of features
- Less overhead on branching
- Faster bug identification



# Q&A





# NetApp Collateral List




NetApp®

## White Papers:

- WP-7181 Optimizing EDA Workflows

## Best Practice TRs:

- TR-4143 Optimizing VCS Performance
- TR-4164 Perforce Deployment Guide
- TR-4142 Backup and Recovery for Perforce
- TR-4067 NFS Implementation Guide on cDOT
- TR-4063 pNFS Best Practices for cDOT
- TR-3183 Optimizing RedHat NFS Clients
- TR-4270 Optimizing Cadence Virtuoso Liberate
- TR-4238 Optimizing VCS Performance on cDOT
- TR-4239 Optimizing VCS Performance with pNFS
- TR-4237 Making LSF Storage-aware




**NetApp**

Collateral Kiosk

[www.netapp.com/hightech](http://www.netapp.com/hightech)




## Stay Connected

**Join us in the High-Tech Industry User Group:**  
<https://communities.netapp.com/groups/high-tech-industry-user-group>

**Follow us on Twitter:**  
 @NetApp  
 Hashtag: #NetAppHiTech

**'Like' us on Facebook:**  
[Facebook.com/NetApp](https://www.facebook.com/NetApp)

**Code Reader:**  
[app.scanlife.com/appdownload/dl](http://app.scanlife.com/appdownload/dl)





NetApp®

*Thank you*

