




Go further, faster®

ParaSwift: File I/O trace modeling for the future

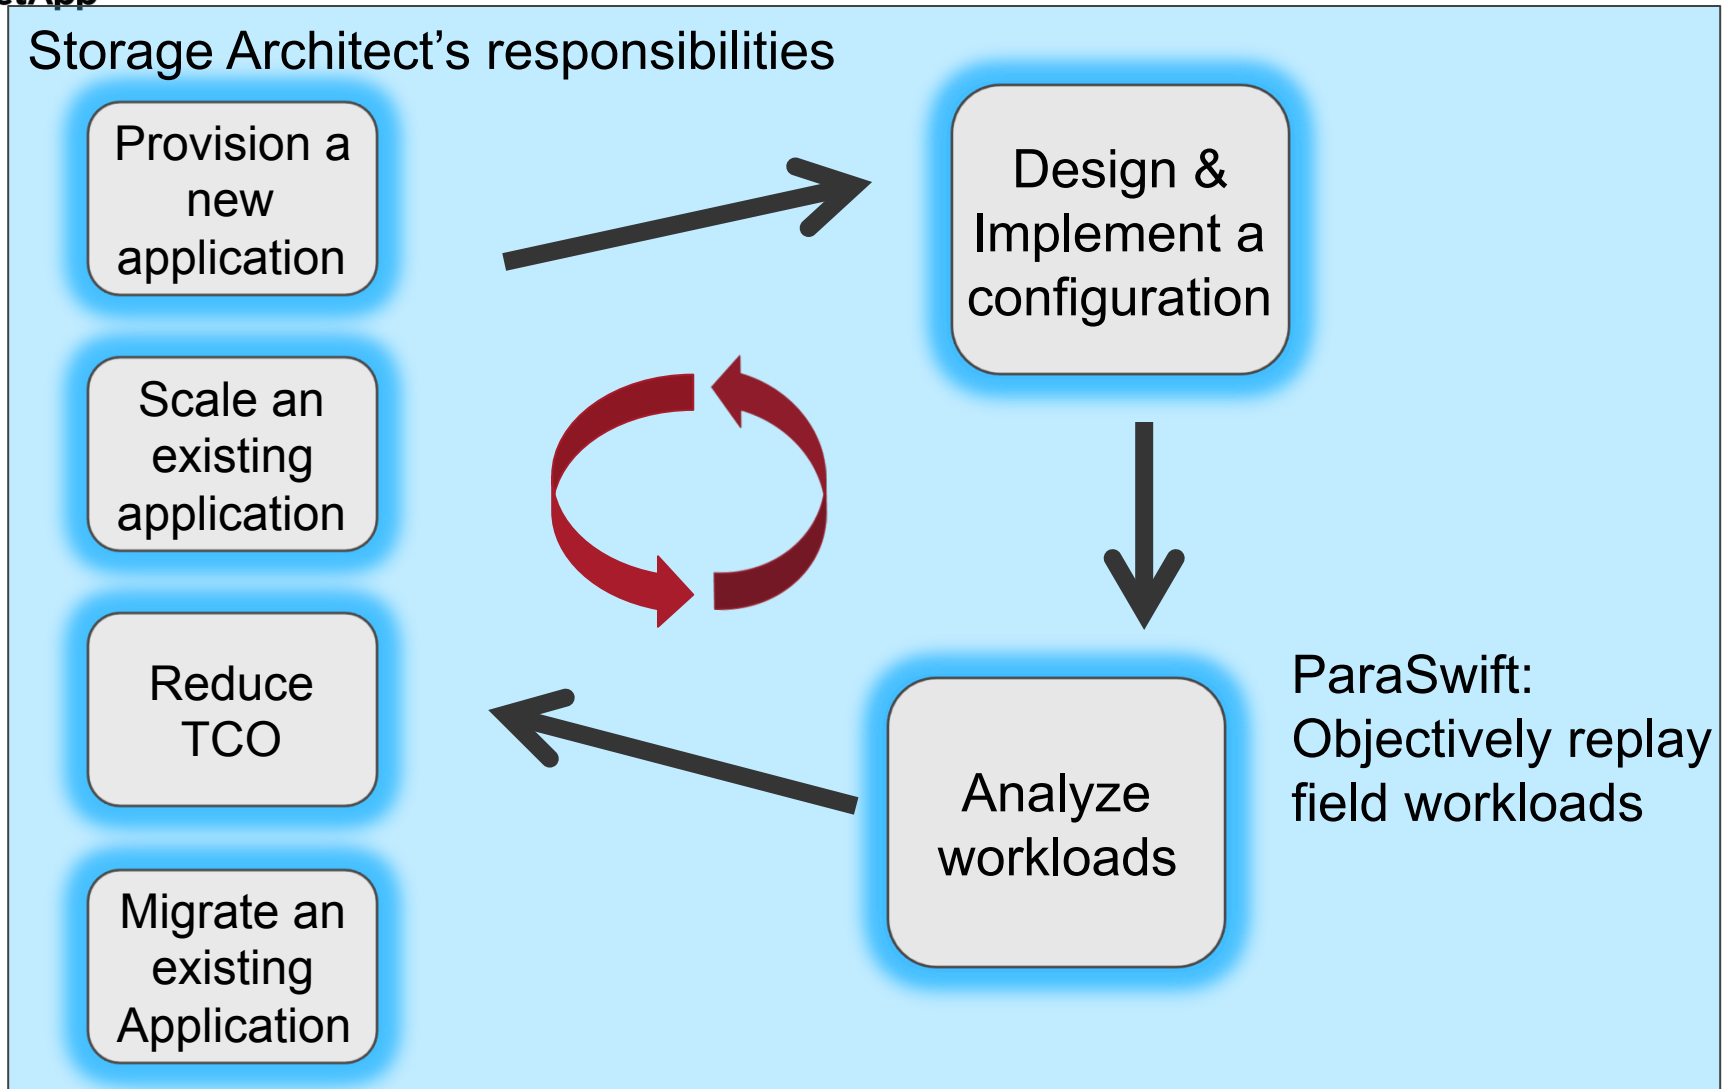
Rukma Talwadker, Kaladhar Voruganti
Advanced Technology Group
NetApp (I) Pvt. Ltd





NetApp®


Where ParaSwift Can Help?






Bird's Eye View of the ParaSwift Ecosystem

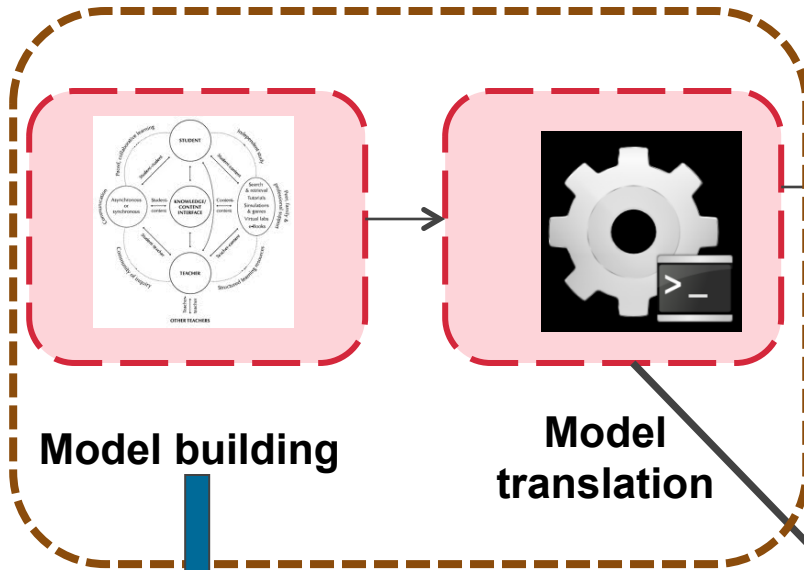
Trace customer applications



ParaSwift Ecosystem



Trace Capture Ecosystem



Model building


Model translation




Trace Library

Load DynamiX*

Workload



FS Hierarchy Creation



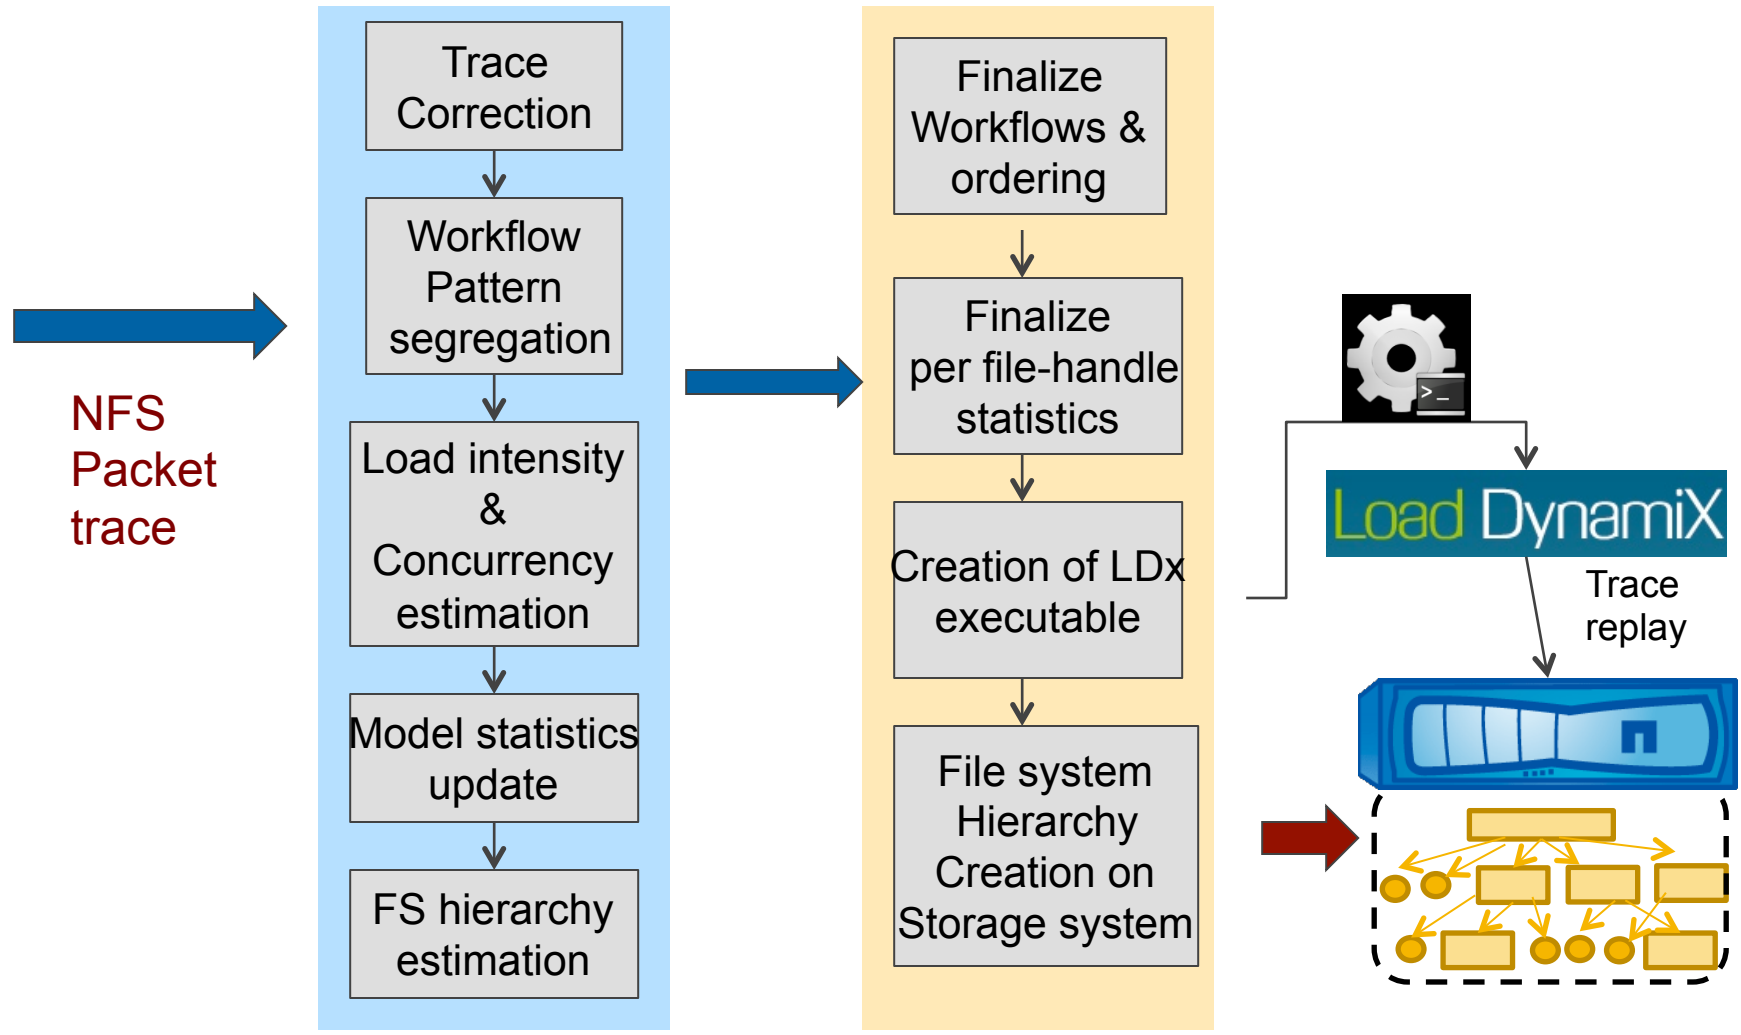
* Third party component



ParaSwift's Trace Modeling Flow

Inline components

Offline components





Validation Metrics

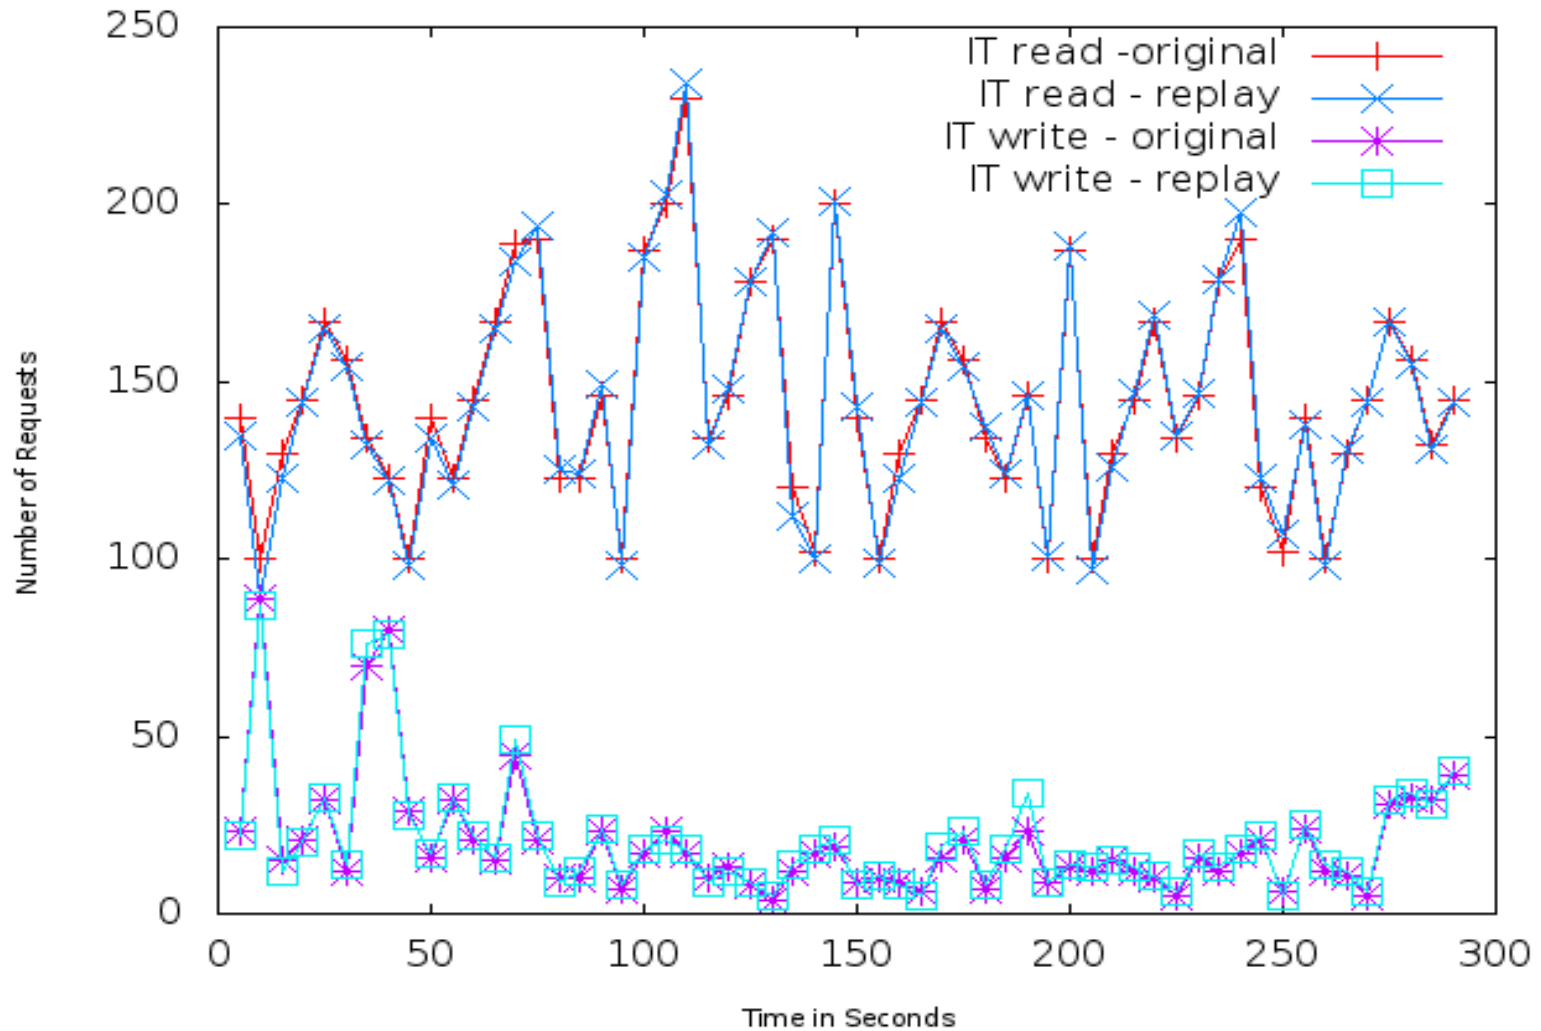
We validate the ability of **ParaSwift** to regenerate the workload by comparing the **original trace** with the **regenerated trace**

1. Appropriate operation mix
2. Appropriate read/write file a) offsets and b) I/O size distributions
3. Appropriate write properties; file a) partial b) full vs c) appends.
4. Misaligned writes
5. Number of unique files per minute
6. Read/Write payloads

We use Root Mean Square (RMS) error for comparisons. **~92% accuracy**



Results: Accuracy of Regeneration





Thank you

