

SMB 2.0 implementation in Data ONTAP

Ankit Mathur

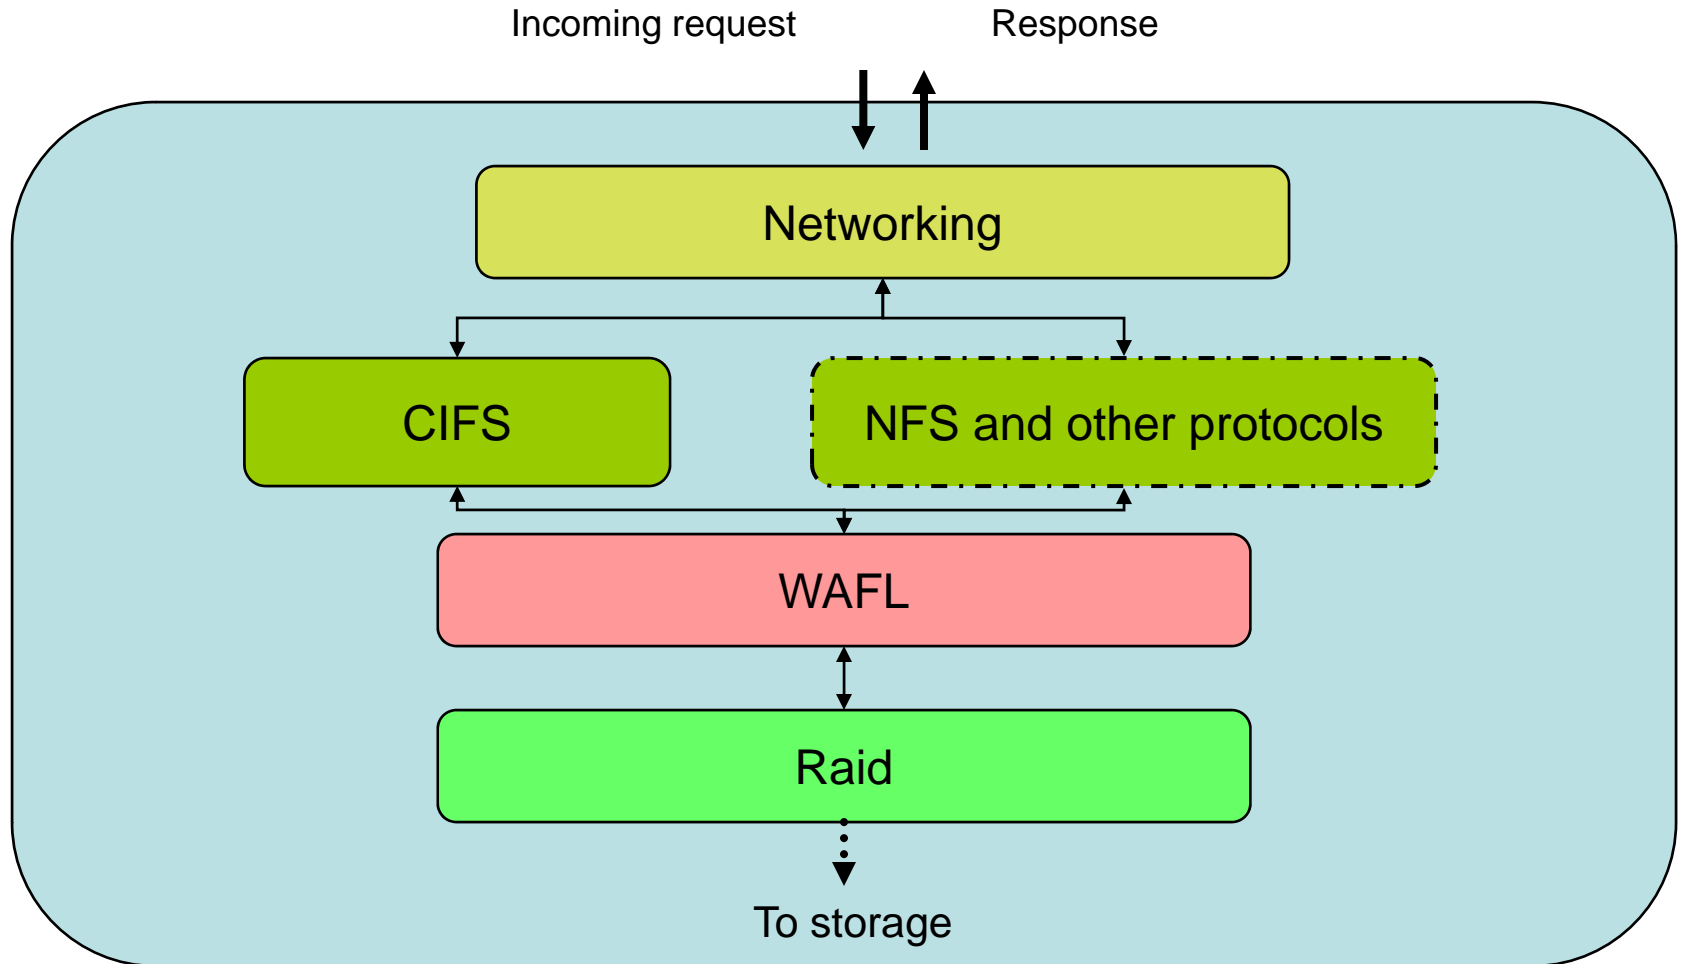
NetApp, Inc

- ❑ CIFS in Data ONTAP
- ❑ Adding SMB 2.0 support
- ❑ Testing and Inter-op
- ❑ Wish list
- ❑ Summary

CIFS in Data ONTAP

- ❑ Data ONTAP has logical protocol processing domains
- ❑ Data is zero copied across these domains
- ❑ Networking makes up call in to CIFS
- ❑ CIFS sends messages to WAFL for file system operations setting a Finish routine for WAFL to call back

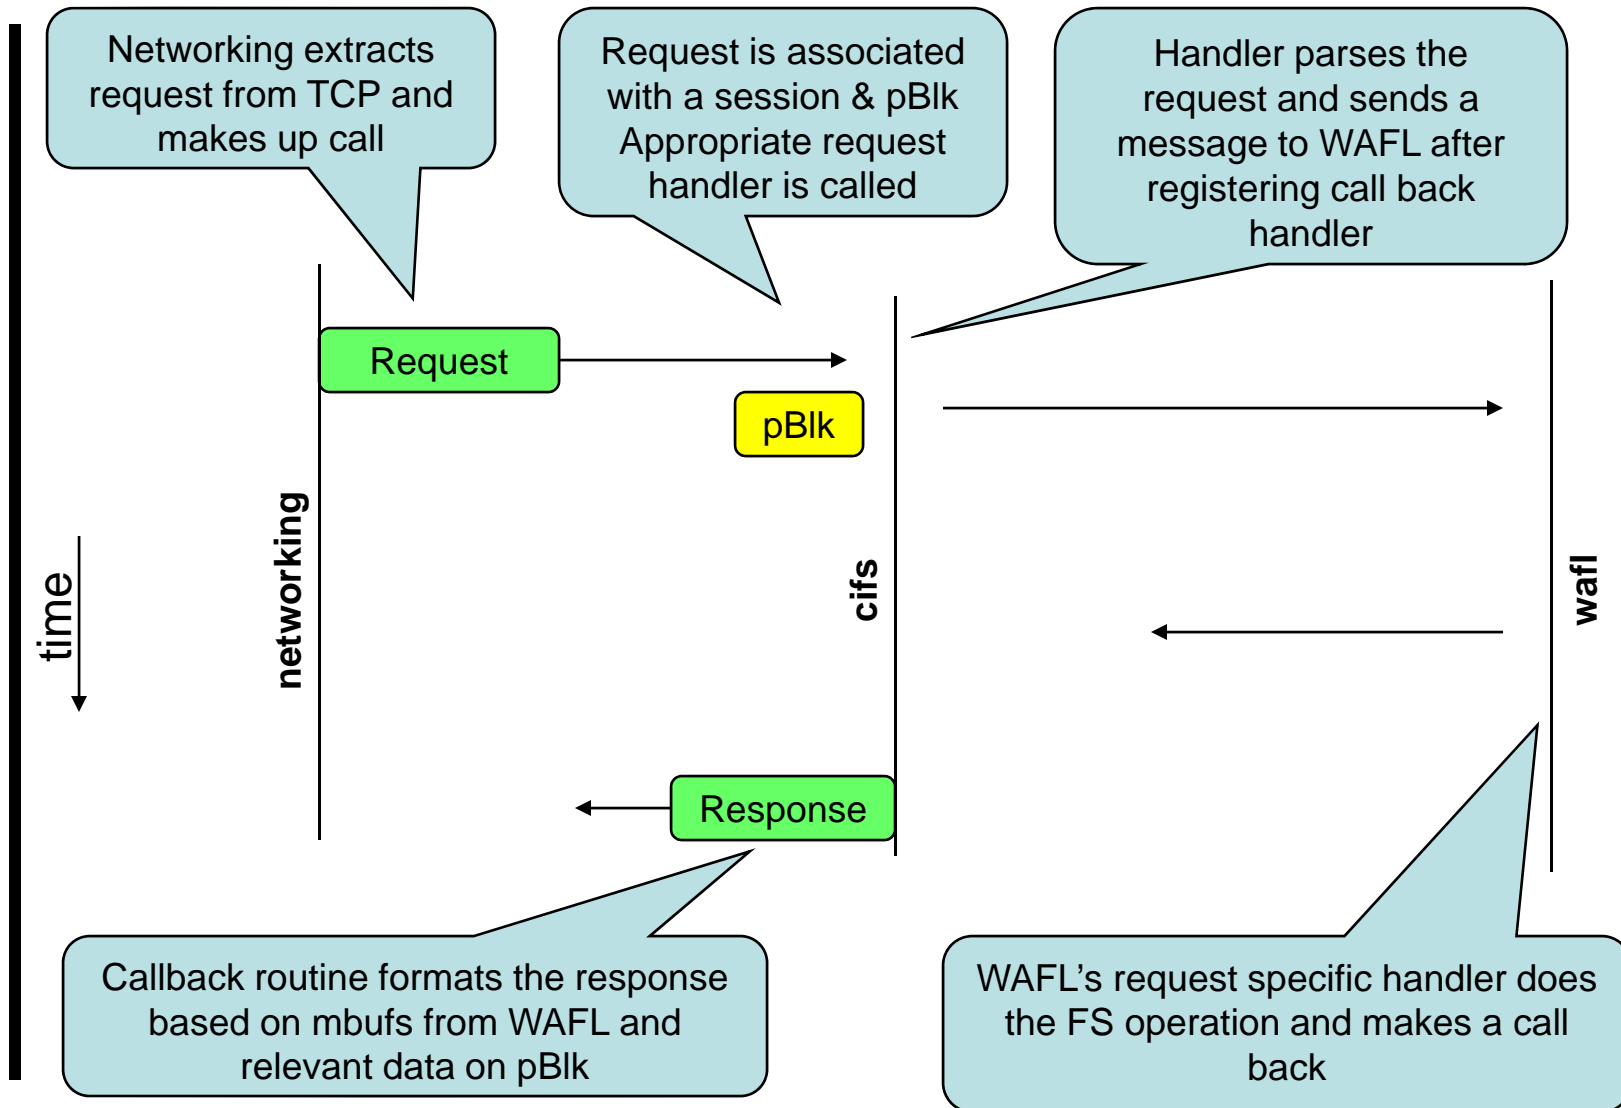
Framework in Data ONTAP (contd.)



Data ONTAP SK micro kernel

- ❑ Uses pre-allocated memory chunks for each request
 - pBlks
- ❑ pBlks have references to:
 - ❑ Received request
 - ❑ Relevant parsed data
 - ❑ WAFL specific data
- ❑ pBlk remains associated with request until response is sent

Request processing lifecycle



Adding SMB 2.0 support

Laying the SMB 2.0 foundation

- ❑ Test client capable of running against partially-implemented server is required
- ❑ Sort protocol commands on the basis of essentialness in setting up a connection
 - ❑ Implementing in this order allows using real client for tests sooner

Adding support in Data ONTAP

- ❑ Adding request-specific handlers and response-specific finish routines
- ❑ Compounded request processing support
 - ❑ Related – using pBlk as common carrier
 - ❑ Unrelated – forking out to multiple pBlks

- ❑ Message signing – much improved against its predecessor in SMB 1.0 and fairly to the book
 - ❑ State management across packets is simpler than SMB 1.0
- ❑ Async responses
 - ❑ requires synchronization – interim response should go before final
 - ❑ Credit granting required in the interim response
- ❑ Interaction between credit granting, interim responses and compounded requests can pose state management challenges

Testing and Inter-op

- ❑ Testing tool to validate partial server implementation is a big value add
- ❑ Test tools should be regularly verified against Windows servers
 - ❑ Documentation can mismatch server behavior
- ❑ Credit granting issues can be stealthy and appear only after exchanging lots of requests

- ❑ Found gaps in the available documentation and final expected server behavior
- ❑ Interaction helped nail down expected behaviors in unique scenarios
 - ❑ Multiple async responses to related compounded requests

Wish list

- ❑ A construct or response that could mandate retry-after 'x' duration
- ❑ More control over durable file handles retry mechanism
- ❑ Making session independent of underlying TCP connections
- ❑ Method to promote the granted oplock level

Summary

- ❑ Implementation targeting a changing documentation can get challenging
- ❑ Testing tools for partial implementation are indispensable
- ❑ Async responses and credit granting don't go hand-in-hand cleanly

Summary – cont'd

- ❑ Lack of path based requests in SMB 2.0 makes implementation cleaner but can bring down performance in certain load-mix
- ❑ Absence of Trans Secondary equivalent can limit the protocol

Questions