

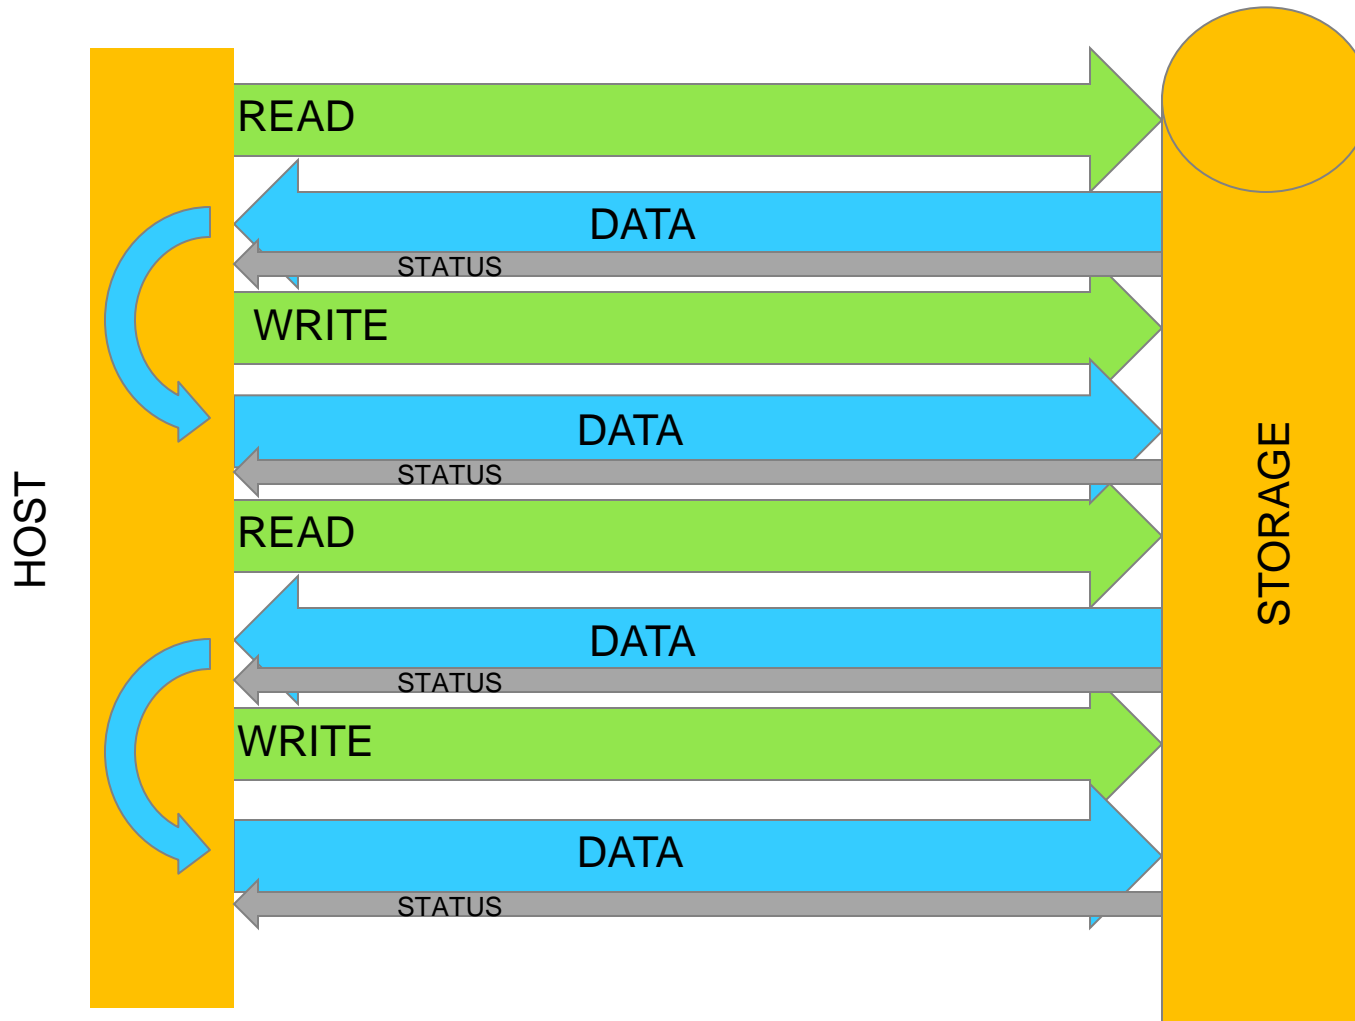
Storage Data Movement Offload

Frederick Knight
NetApp

Storage Data Movement Offload

- ❑ What is Storage Data Movement Offload?
- ❑ Why is it important?
- ❑ What are the uses?
- ❑ How does it work?

Moving Data the old way



What Is It?

- ❑ Storage Data Movement Offload
 - ❑ Also called - Copy Offload
 - ❑ Storage device copies the data
 - ❑ Data does not move through the host
- ❑ Two Techniques
 - ❑ Single command current data copy
 - ❑ Multi-command Read/Write model

Why Is It Important?

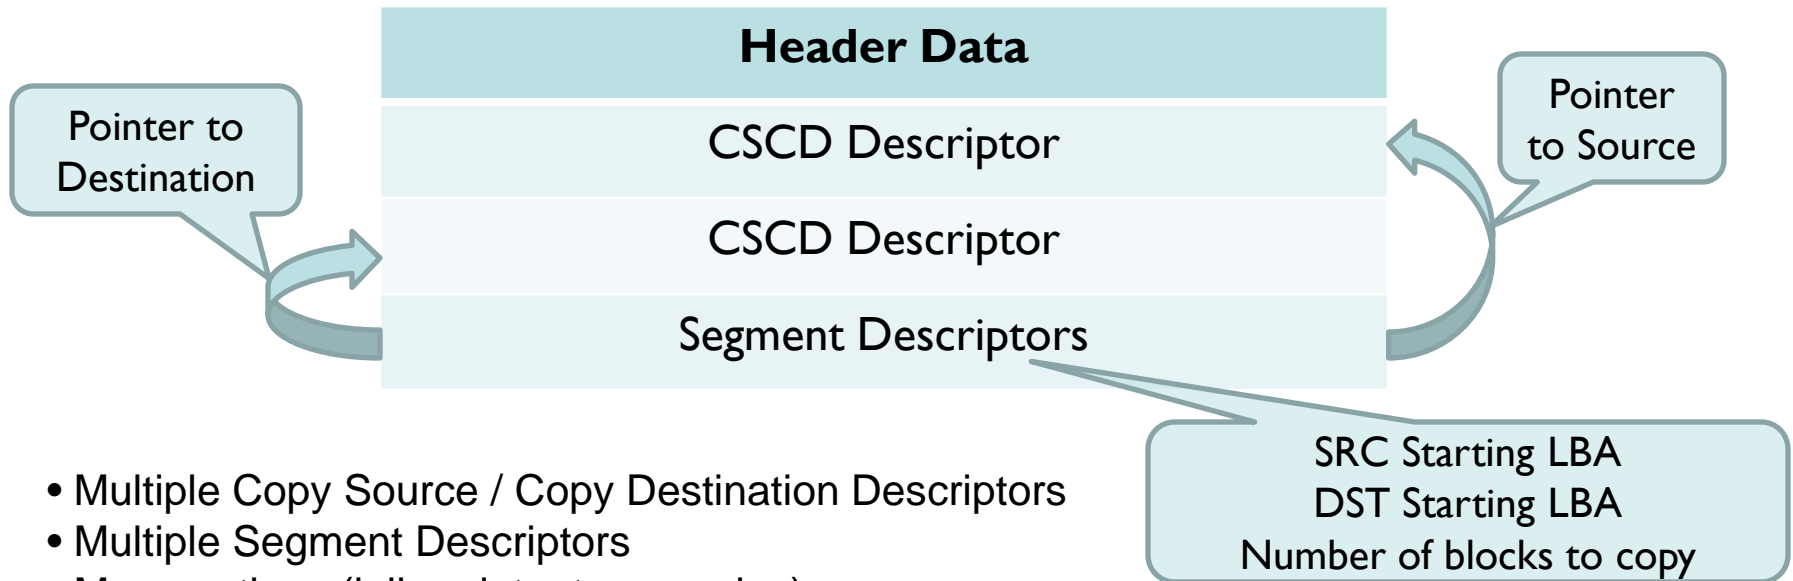
- ❑ Copies data faster
- ❑ Saves bandwidth for real work
 - ❑ Saves I/O bus (fabric/network) bandwidth
 - ❑ Saves host I/O adapter bandwidth
 - ❑ Saves host memory bandwidth
- ❑ Offloads work from the host to storage
- ❑ Gets you back to real work faster

How Copy Offload works

- ❑ Copy Manager in the device
 - ❑ Manages the movement of the data
- ❑ Two models of operation
 - ❑ EXTENDED COPY (XCOPY) operations
 - ❑ TOKEN based operations

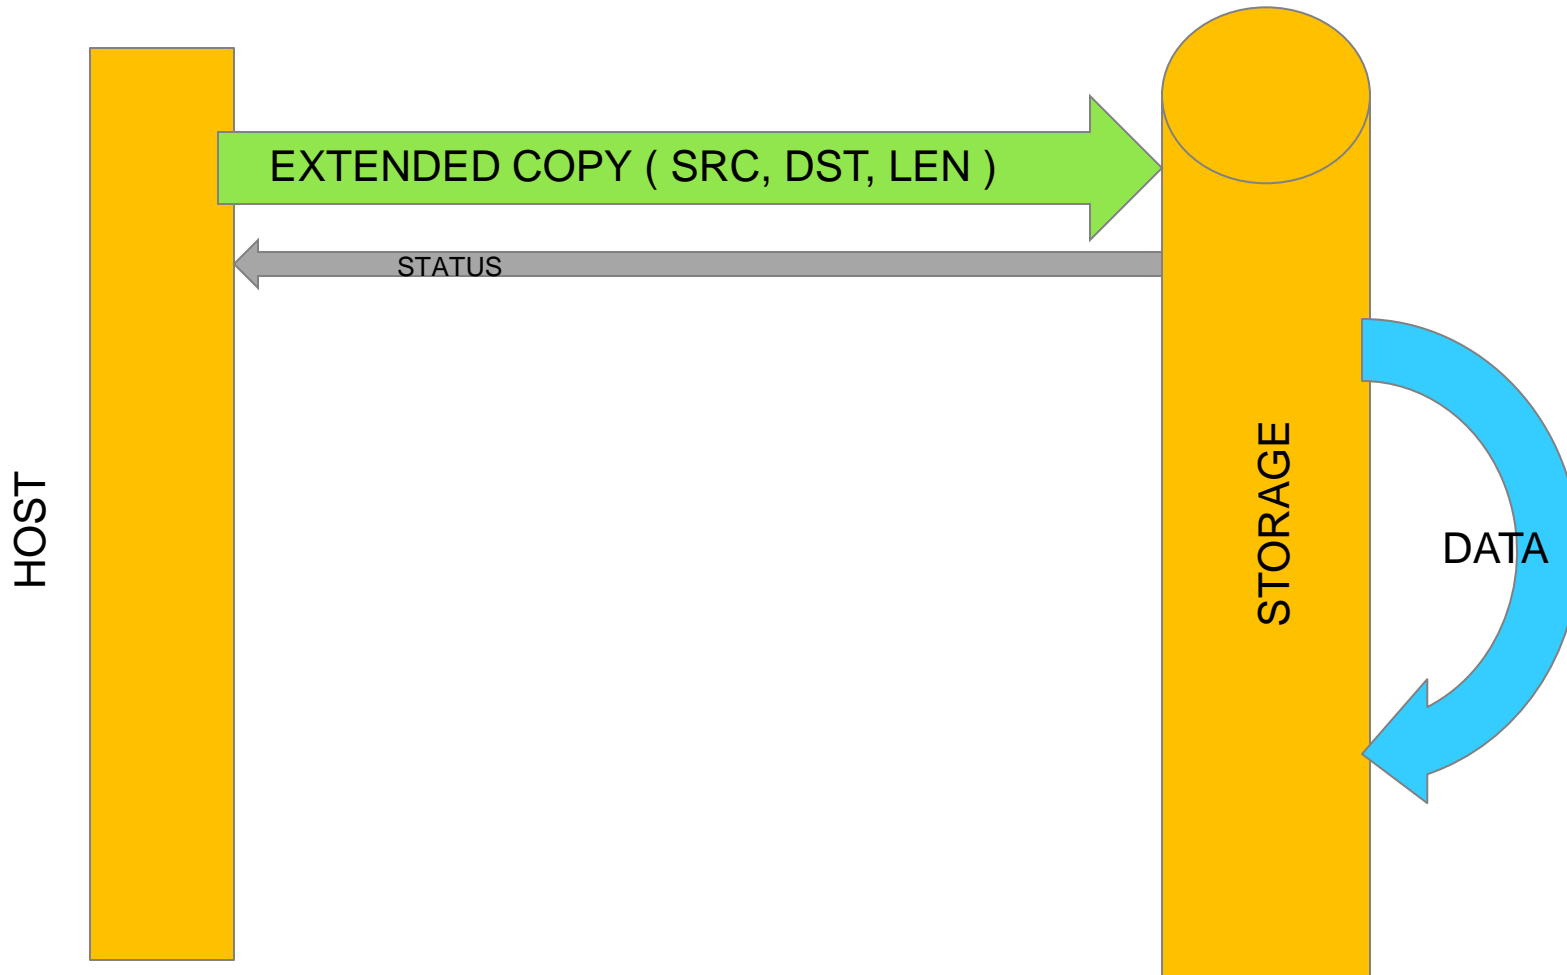
- ❑ Copy Manager in the device
 - ❑ Manages the movement of the data
 - ❑ Operates on data currently in the blocks
- ❑ EXTENDED COPY command provides:
 - ❑ Identifiers for the Logical Units
 - ❑ The SRC starting LBA
 - ❑ The DST starting LBA
 - ❑ The length
 - ❑ A whole lot more optional stuff (including tapes)

How XCOPY works

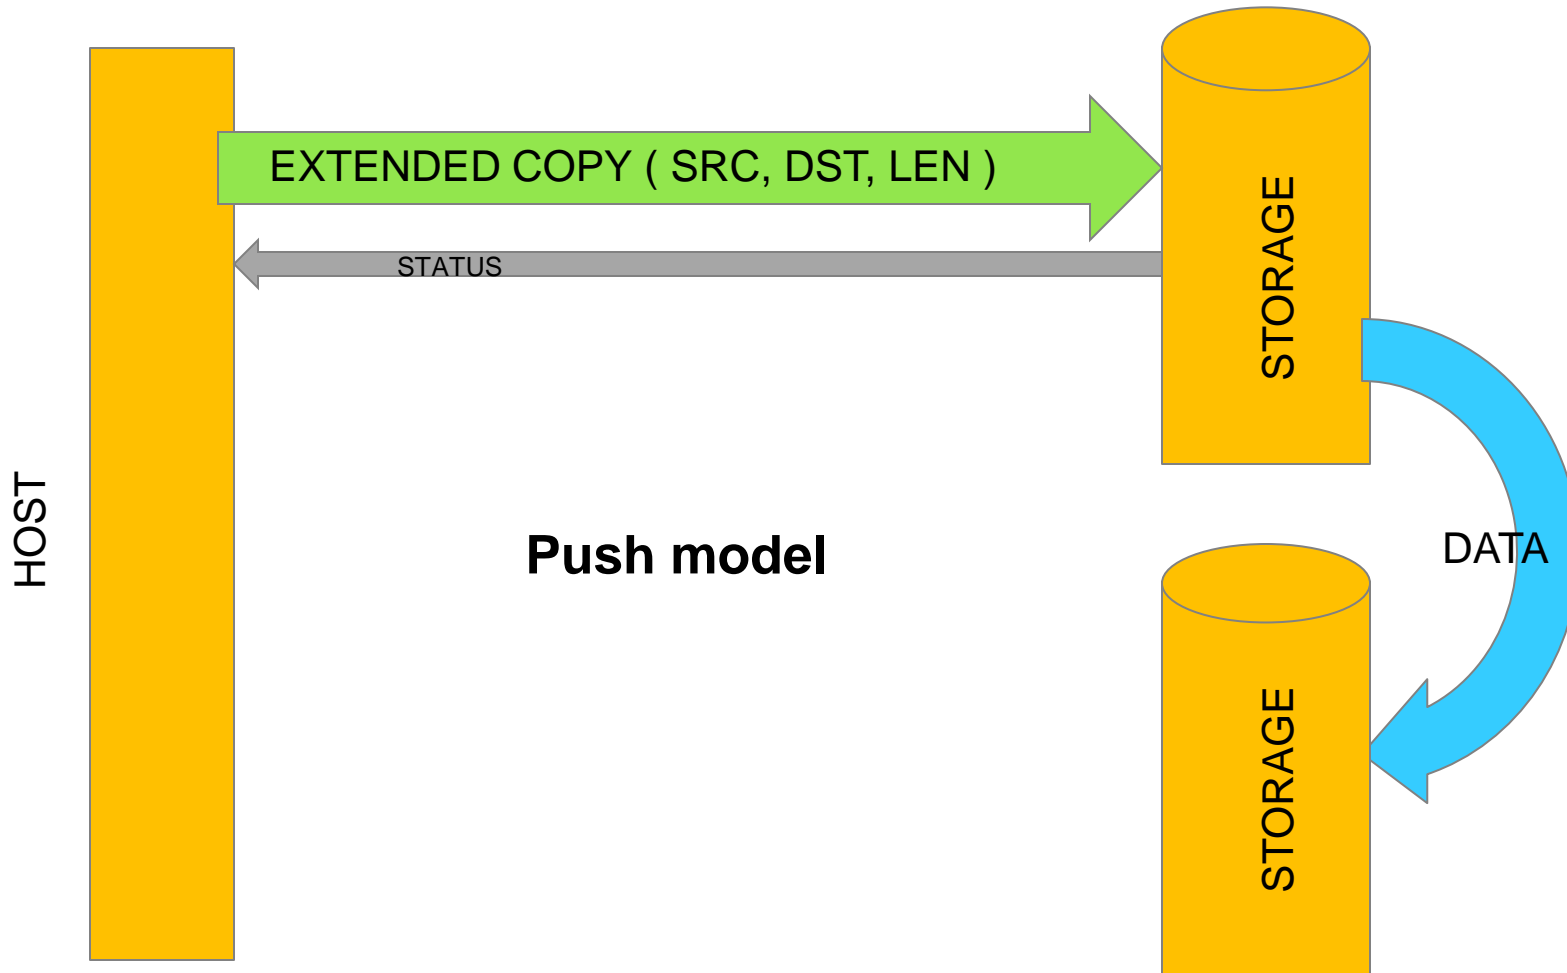


- Multiple Copy Source / Copy Destination Descriptors
- Multiple Segment Descriptors
- Many options (inline data, tape copies)
- CSCD descriptors provide access to:
 - Different LBAs on the same LUN
 - Different LUNs in the same Array
 - Different Arrays in a SAN
- Extremely complicated

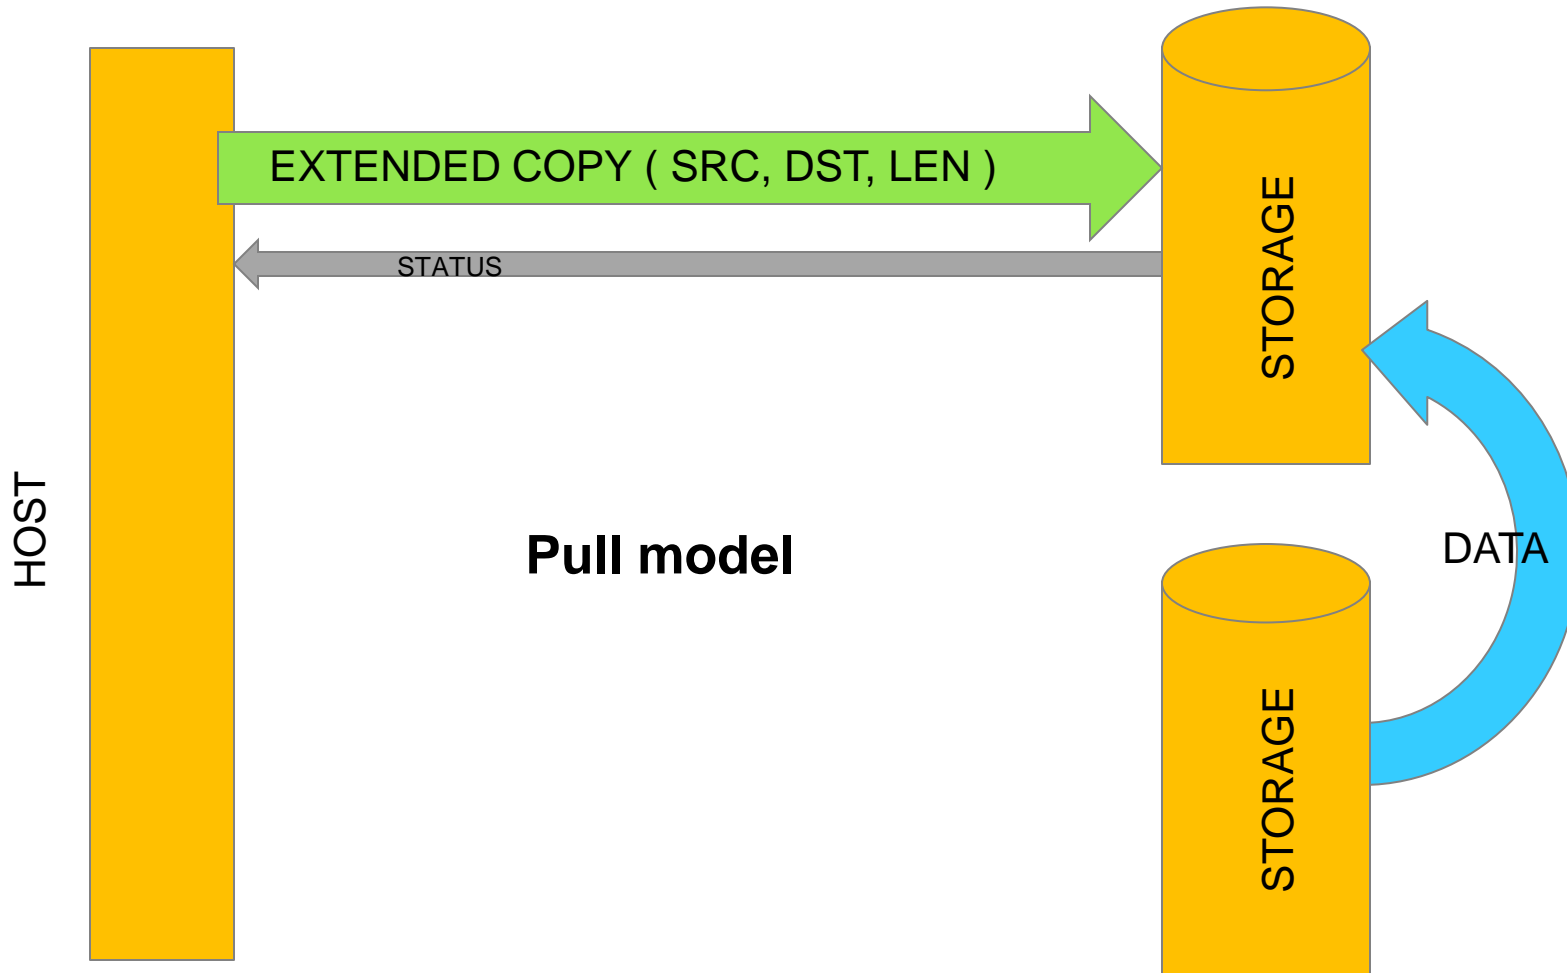
XCOPY Data Movement



XCOPY Data Movement



XCOPY Data Movement



How Token Operations work?

- ❑ Copy Manager in the device
 - ❑ Manages the movement of the data
 - ❑ Operates on frozen instance of data
- ❑ Creates a TOKEN that represents the data
 - ❑ If you have the TOKEN, then you have the data
 - ❑ Tokens are 512 bytes in size
 - ❑ Freezes the source data
 - ❑ Point in time snapshot/clone/copy
 - ❑ Write tracking of source LBAs

How Token Operations work?

- ❑ Token operations require multiple commands
- ❑ POPULATE TOKEN command
 - ❑ Send list of LBAs to source device (read)
- ❑ RECEIVE ROD TOKEN INFORMATION command
 - ❑ Gets the status & the token
- ❑ WRITE USING TOKEN command
 - ❑ Send token to destination with list of write LBAs
 - ❑ Writes data from the frozen instance

How Token Operations work

- POPULATE TOKEN sent to source LUN
 - Tagged with a List Identifier
 - With list of LBAs to READ
- Creates Point in time Instance of requested data
 - Frozen point in time copy
 - LBA list for write tracking

CDB (83h / 10h)

List Identifier

Parameter Data Header

Multiple LBA/length pairs

How Token Operations work

- RECEIVE ROD TOKEN INFORMATION sent to source LUN
 - Uses the same List Identifier (same as POPULATE TOKEN)
 - Obtains status of the matching POPULATE TOKEN command
 - Obtains the 512 byte TOKEN created by POPULATE TOKEN

CDB (84h + 07h)

Matching List Identifier

Parameter Data Header

Status

TOKEN value

How Token Operations work

- WRITE USING TOKEN sent to destination LUN
 - Supplies the 512 byte TOKEN value
 - Supplies a list of LBAs
 - Uses a new List Identifier

CDB (83h + 11h)

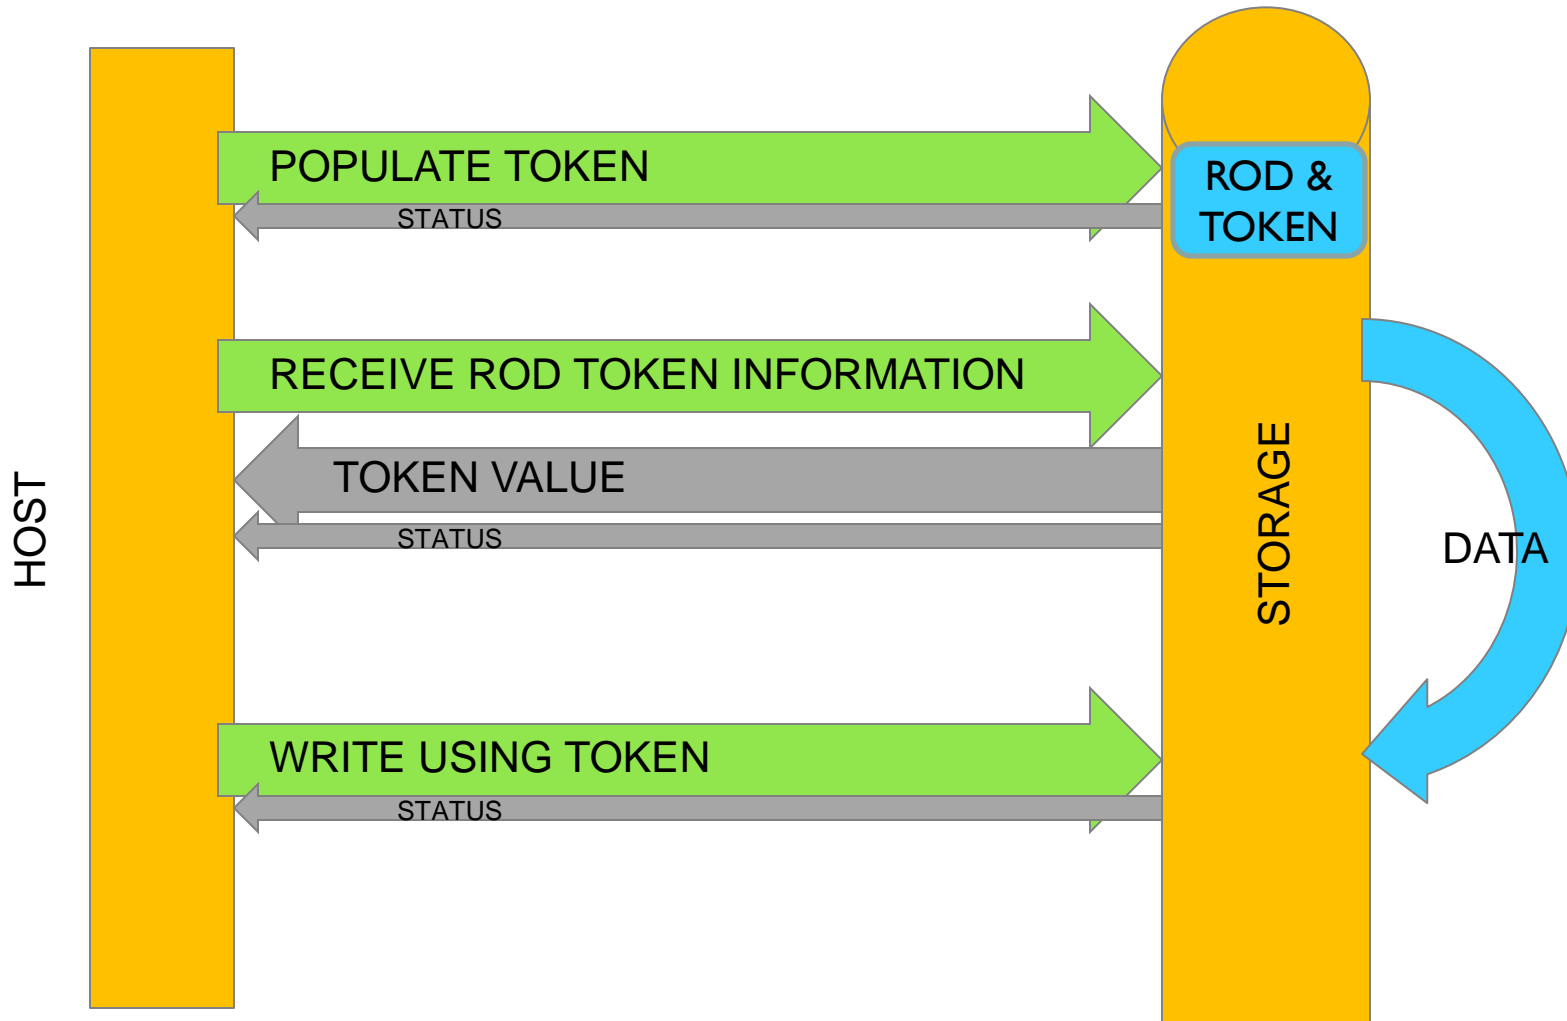
List Identifier

Parameter Data Header

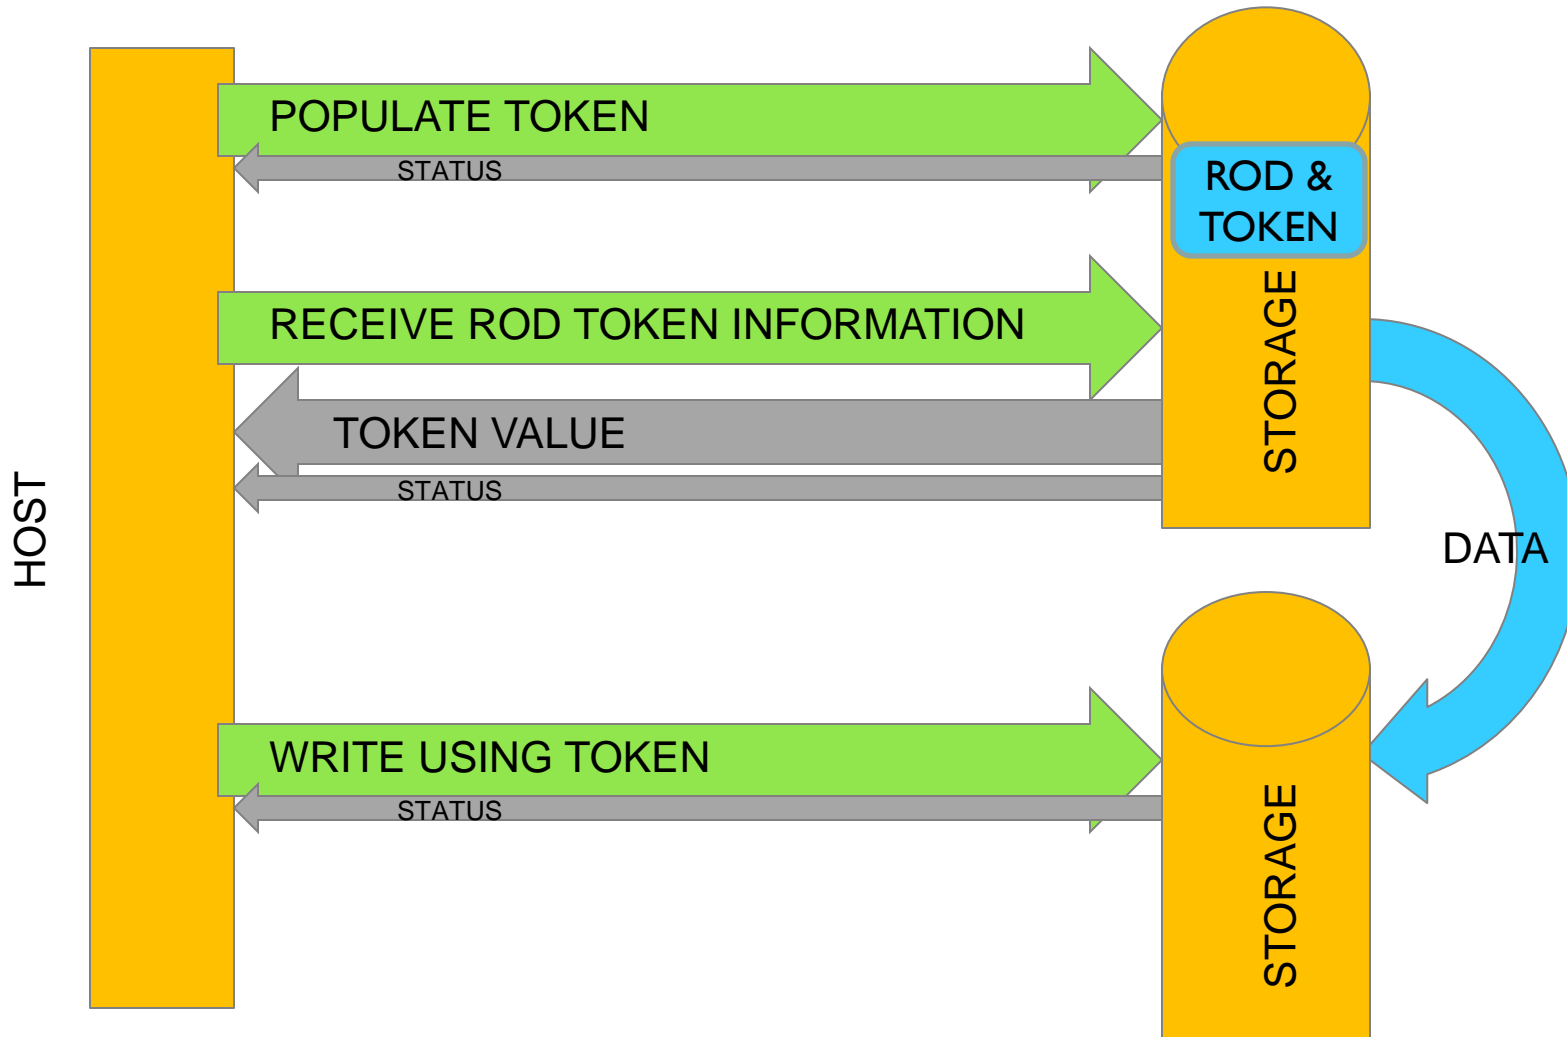
TOKEN value

Multiple LBA/length pairs

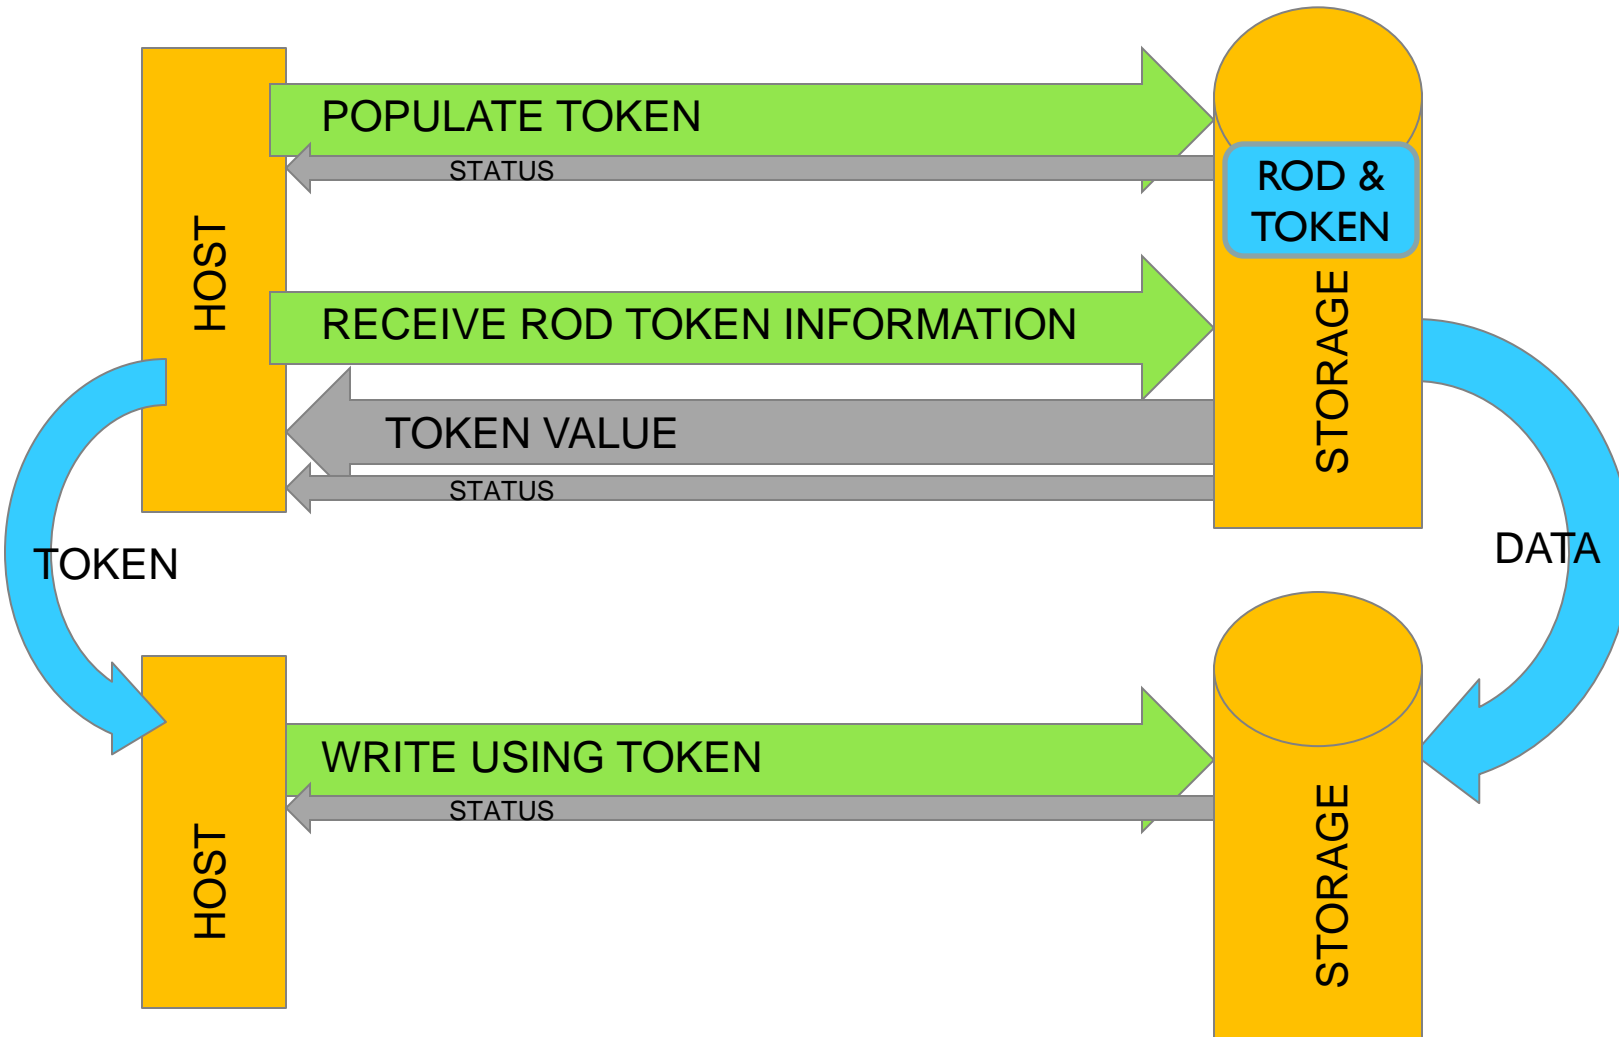
Token based Data Movement



Token based Data Movement



Token based Data Movement



What are the uses?

- ❑ Virtual Machine Provisioning
 - ❑ Multiple XCOPY commands that copy from same source to different destinations
 - ❑ Single XCOPY command with multiple segments descriptors referencing many CSCD destinations
 - ❑ Operates on data currently in the blocks

What are the uses?

- ❑ Virtual Machine Provisioning
 - ❑ POPULATE TOKEN command
 - ❑ Creates a point in time instance of the data
 - ❑ If you have the TOKEN, then you have the data
 - ❑ Lots of WRITE USING TOKEN commands (with the same TOKEN) to different destinations – the TOKEN is multi-use

What are the uses?

- Initialization
 - WRITE SAME command
 - Write the same data to lots of LBAs
 - WRITE USING TOKEN command
 - Special defined TOKEN to write zeros to lots of LBAs

What are the uses?

- ❑ User level COPY command
 - ❑ Hard using XCOPY
 - ❑ POPULATE TOKEN - READ
 - ❑ With all the file extents
 - ❑ WRITE USING TOKEN - WRITE
 - ❑ With the new list of file extents

What are the uses?

- ❑ Other possible uses
 - ❑ Defragmentation
 - ❑ POPULATE TOKEN (with fragment extents)
 - ❑ WRITE USING TOKEN (with single extent)
 - ❑ BACKUP operations

- ❑ Learn what APIs are available
- ❑ Ask Vendors to provide APIs