

EXTENDS *Naturals, Sequences*

Our “Test Data”. Each of these is a set of ids, relevant only in that they are distinct from each other.

CONSTANTS

USERIDS, A set of *userIds* to test with (one per user)
SERVERS, A set of *serverIds* (each one will “create” a new server)
METADATAS, A set of *metadata* versions.
IMAGES A set of image versions.

Our variables update each step and represent the state of our modeled system.

VARIABLES

These variables are relevant to the implementation.

databaseState, *databaseState[key]* = What is stored for this key
blobStoreState, *blobStoreState[key]* = What is stored for this key
serverStates, *serverStates[serverId]* = What the server is doing

This variable is used to observe the state of the system to check if it’s doing the right thing. Think of it like the test harness.

Represents all the write requests and read responses sent to/from the system.

operations

Represents every variable in this model.

$vars \triangleq \langle databaseState, blobStoreState, serverStates, operations \rangle$

You strongly type math with math. Here is where we say which types are allowed for all our variables. *TypeOk* is set as an Invariant, which means we expect it always to be true. It will fail if false, effectively giving us type checking. The declarations that follow are part of *TypeOk*, separated to make it clearer.

Allows all the values to also be *UNSET*, which is a distinct value not to be confused for the others.

$UserVal \triangleq USERIDS \cup \{“UNSET”\}$
 $MetadataVal \triangleq METADATAS \cup \{“UNSET”\}$
 $ImageVal \triangleq IMAGES \cup \{“UNSET”\}$

Describes all possible states a server can be in.

$ServerStateVal \triangleq$
 $[$
 $\quad state : \{$
 $\quad \quad current:$
 $\quad \quad “waiting”, \quad next: StartWrite \text{ or } StartRead$
 $\quad \quad after: StartWrite$

```

    "started_write", next: WriteMetadata or FailWrite
    after: WriteMetadata
    "wrote_metadata", next: WriteBlobAndReturn or FailWrite
    after: StartRead
    "started_read", next: ReadMetadata
    after: ReadMetadata
    "read_metadata" next: ReadBlobAndReturn
  },
  userId : UserIdVal,
  metadata : MetadataVal,
  image : ImageVal
]

```

Represents an action that occurred on the *API* boundary. Used for observability.

$OperationValue \triangleq [type : \{ "READ", "WRITE" \},$
 $userId : UserIdVal,$
 $metadata : MetadataVal,$
 $image : ImageVal]$

The full type specification for all variables in the system

$TypeOk \triangleq$

The database state contains a mapping of *userIds* to *metadatas*. It can also be "UNSET", representing a case where there is no *metadata*. Note: we make this specific to our problem. If this were a more general problem, it might look like:

$databaseState \in [KEYS \rightarrow RECORDS].$

$\wedge databaseState \in [USERIDS \rightarrow MetadataVal]$

The blob store state contains a mapping of *userIds* to images. Note: we make this specific to our problem. If this was a more general problem, it might look like:

$blobStoreState \in [KEYS \rightarrow BLOBS].$

$\wedge blobStoreState \in [USERIDS \rightarrow ImageVal]$

The *serverStates* store the current states for each server, allowing us to build a state machine describing our system. Implemented as a mapping between servers and all their possible states.

$\wedge serverStates \in [SERVERS \rightarrow ServerStateVal]$

$\wedge operations \in Seq(OperationValue)$

When the model starts, everything begins unset. Unlike standard testing every possible state will be explored, so we don't need to initialize for specific scenarios.

$Init \triangleq$

$\wedge databaseState = [u \in USERIDS \mapsto "UNSET"]$
 $\wedge blobStoreState = [u \in USERIDS \mapsto "UNSET"]$
 $\wedge serverStates = [s \in SERVERS \mapsto [state \mapsto "waiting",$
 $userId \mapsto "UNSET",$
 $metadata \mapsto "UNSET",$

$image \mapsto \text{"UNSET"}$

]]

$\wedge operations = \langle \rangle$

State Machine: All of the states are functions of s (server), because the only actively modeled actors in this system are our servers, but there can be multiple working simultaneously.

$StartWrite(s) \triangleq$

Writing only starts when a server is waiting

$\wedge serverStates[s].state = \text{"waiting"}$

This will try every combination of $userId$, $metadata$ and $image$ (one at a time). We store this throughout the state lifecycle. Next states will refer to this

$\wedge \exists u \in USERIDS, m \in METADATAS, i \in IMAGES :$

$serverStates'$ means the next state of $serverStates$

$\wedge serverStates' = [serverStates \text{ EXCEPT } \begin{array}{l} \text{update only server } s \\ ![s].state = \text{"started_write"}, \text{ update state} \\ \text{set values for the upcoming write} \\ ![s].userId = u, \\ ![s].metadata = m, \\ ![s].image = i \end{array}$

Record the write for observability

$\wedge operations' = Append(operations,$

This is created with "record" symantics, which is why \mapsto not $=$ is used

[

$type \mapsto \text{"WRITE"},$

$userId \mapsto u,$

$metadata \mapsto m,$

$image \mapsto i$

])

We need to list every unchanged variable.

Not changing is a behavior too.

$\wedge UNCHANGED \langle databaseState, blobStoreState \rangle$

$WriteMetadata(s) \triangleq$ Represents a successful database write

Established an alias to make code more compact

LET $currentState \triangleq serverStates[s]$

IN

Metadata writing happens directly after write is started

$\wedge currentState.state = \text{"started_write"}$

Database is transactional/consistent. We can therefore model this happening in one step

$\wedge databaseState' = [databaseState \text{ EXCEPT } \begin{array}{l} ![currentState.userId] = currentState.metadata \end{array}]$

$\wedge serverStates' = [serverStates \text{ EXCEPT } \text{This is how the state advances}]$
 $\quad \quad \quad ! [s].state = \text{"wrote_metadata"}$
 $\wedge \text{UNCHANGED } \langle blobStoreState, operations \rangle$
 $WriteBlobAndReturn(s) \triangleq \text{Represents a successful blob store write}$
 $\text{LET } currentState \triangleq serverStates[s]$
 IN
 $\quad \quad \quad \text{Metadata writing happens directly after write is started}$
 $\quad \quad \quad \wedge currentState.state = \text{"wrote_metadata"}$
 $\quad \quad \quad \text{Blob store has read after write consistency. We can therefore}$
 $\quad \quad \quad \text{model it happening in one step}$
 $\quad \quad \quad \wedge blobStoreState' = [blobStoreState \text{ EXCEPT }]$
 $\quad \quad \quad \quad \quad \quad ! [currentState.userId] = currentState.image]$
 $\quad \quad \quad \wedge serverStates' = [serverStates \text{ EXCEPT } \text{update only server } s]$
 $\quad \quad \quad \quad \quad \quad \text{Process done once blob is written}$
 $\quad \quad \quad \quad \quad \quad ! [s].state = \text{"waiting"}$
 $\quad \quad \quad \wedge \text{UNCHANGED } \langle databaseState, operations \rangle$

$FailWrite(s) \triangleq$
 $\quad \quad \quad \text{In our model, a server can only fail if it is writing. We don't need to do this, but it cuts down}$
 $\quad \quad \quad \text{state space we don't care about. We are worried about writes failing in a bad spot causing}$
 $\quad \quad \quad \text{errors in future reads and writes. We don't model spontaneous read failures.}$
 $\quad \quad \quad \text{Will only get to this state if writing}$
 $\quad \quad \quad \wedge serverStates[s].state \in \{ \text{"started_write"}, \text{"wrote_metadata"} \}$
 $\quad \quad \quad \wedge serverStates' = [serverStates \text{ EXCEPT }]$
 $\quad \quad \quad \quad \quad \quad ! [s].state = \text{"waiting"},$
 $\quad \quad \quad \quad \quad \quad ! [s].userId = \text{"UNSET"},$
 $\quad \quad \quad \quad \quad \quad ! [s].metadata = \text{"UNSET"},$
 $\quad \quad \quad \quad \quad \quad ! [s].image = \text{"UNSET"}]$
 $\quad \quad \quad \text{Nothing happens with the database and blob store. Everything this server did stays done,}$
 $\quad \quad \quad \text{anything left undone stays undone.}$
 $\quad \quad \quad \wedge \text{UNCHANGED } \langle databaseState, blobStoreState, operations \rangle$

We model reading in detail because reading and writing are occurring at the same time, and may interact with each other in unexpected ways.

$StartRead(s) \triangleq$
 $\quad \quad \quad \text{Reading only starts when a server is waiting}$
 $\quad \quad \quad \wedge serverStates[s].state = \text{"waiting"}$
 $\quad \quad \quad \wedge \exists u \in USERIDS : \text{When we start reading we pick a user id}$
 $\quad \quad \quad \quad \quad \quad serverStates' = [serverStates \text{ EXCEPT } \text{update only server } s]$
 $\quad \quad \quad \quad \quad \quad ! [s].state = \text{"started_read"},$
 $\quad \quad \quad \quad \quad \quad ! [s].userId = u]$
 $\quad \quad \quad \text{Reading doesn't changed stored state}$
 $\quad \quad \quad \wedge \text{UNCHANGED } \langle databaseState, blobStoreState \rangle$
 $\quad \quad \quad \wedge \text{UNCHANGED } operations$

$ReadMetadata(s) \triangleq$
 LET $currentState \triangleq serverStates[s]$
 IN
 Once the read has started, the first thing we do is Read *Metadata*
 $\wedge currentState.state = \text{"started_read"}$
 $\wedge serverStates' =$
 $[serverStates \text{ EXCEPT } \text{update only server } s$
 $![s].state = \text{"read_metadata"},$
 Assembles the read request from whatever is in the database for that user.
 $![s].metadata = databaseState[currentState.userId]]$

 Reading doesn't changed stored state
 $\wedge \text{UNCHANGED } \langle databaseState, blobStoreState \rangle$
 $\wedge \text{UNCHANGED } operations$

$ReadBlobAndReturn(s) \triangleq$
 LET $currentState \triangleq serverStates[s]$
 IN
 Blob is read after *metadata* is read
 $\wedge currentState.state = \text{"read_metadata"}$
 $\wedge serverStates' = [serverStates \text{ EXCEPT } \text{update only server } s$
 $![s].state = \text{"waiting"},$
 Assembles the read request from whatever is in the database for
 that user.
 $![s].image = blobStoreState[currentState.userId]]$

 $\wedge operations' = Append(operations,$
 Read returns the state it built up during the read process.
 [
 $type \mapsto \text{"READ"},$
 $userId \mapsto currentState.userId,$
 $metadata \mapsto currentState.metadata,$
 $image \mapsto blobStoreState[currentState.userId]$
])
 Reading doesn't changed stored state
 $\wedge \text{UNCHANGED } \langle databaseState, blobStoreState \rangle$

The *Next* section determines what states will be chosen on every step.

$Next \triangleq$
 For every step, pick a server and have it advance one state
 $\exists s \in SERVERS :$
 $\vee StartWrite(s)$
 $\vee WriteMetadata(s)$
 $\vee WriteBlobAndReturn(s)$

$\vee \text{FailWrite}(s)$
 $\vee \text{StartRead}(s)$
 $\vee \text{ReadMetadata}(s)$
 $\vee \text{ReadBlobAndReturn}(s)$

The *Spec* describes what the describe system DOES. First it starts in the *Init* state. Then for every step use *Next* state: represented as $\Box \text{Next}$. In temporal logic \Box means “for all states.” However let’s imagine this is part of a larger system; sometimes this system will do nothing. That is represented by $[\text{Next}]_{\text{vars}}$, meaning: $\text{Next} \vee \text{UNCHANGED vars}$. See learning material for a better explanation of temporal logic operators. Note: The spec in this case describes what the system DOES, not what it should do. Basically this is our system under test, and we describe Invariants (below) and properties (discussed later) to alert us if the system does something wrong/unexpected

$$\text{Spec} \triangleq \text{Init} \wedge \Box [\text{Next}]_{\text{vars}}$$

Invariants: These are things that should always be true about the system. If they become false during any step, an error will occur with a trace that shows you the series of steps that let it to be violated. This is very powerful. The first invariant we saw was *TypeOk*: the types are expected to always conform to the expected type system, and if not we want to know why.

$\text{ConsistentReads} \triangleq$
 If there are no operations, they are consistent
 $\vee \text{operations} = \langle \rangle$
 $\vee \forall i \in 1 \dots \text{Len}(\text{operations}) :$ For every read operation
 LET $\text{readOp} \triangleq \text{operations}[i]$ IN
 $\vee \wedge \text{readOp.type} = \text{“READ”}$
 There must exist a write operation
 $\wedge \vee \exists j \in 1 \dots i :$
 LET $\text{writeOp} \triangleq \text{operations}[j]$ IN
 $\wedge \text{writeOp.type} = \text{“WRITE”}$
 With the same data
 $\wedge \text{readOp.userId} = \text{writeOp.userId}$
 $\wedge \text{readOp.metadata} = \text{writeOp.metadata}$
 $\wedge \text{readOp.image} = \text{writeOp.image}$
 \vee Ignore unset reads
 $\wedge \text{readOp.metadata} = \text{“UNSET”}$
 $\wedge \text{readOp.image} = \text{“UNSET”}$
 $\vee \text{readOp.type} = \text{“WRITE”}$ Ignore writes

One of the best things about invariants is that if they were ever going to be tripped, you’ll hear about it. Unlike testing, where sometimes a confluence of events leads to a test passing when it shouldn’t, the model checker will try every possible state, so if it ever messes up, you’ll know.

This is used for model checker configuration so the simulation doesn’t go on forever.

$\text{StopAfter3Operations} \triangleq$
 $\text{Len}(\text{operations}) \leq 3$
