


# DATALUTION: A TOOL FOR CONTINUOUS SCHEMA EVOLUTION IN NOSQL-BACKED WEB APPLICATIONS

STEFANIE SCHERZINGER  
STEPHANIE SOMBACH  
KATHARINA WIECH  
MEIKE KLETTKE  
UTA STÖRL



OSTBAYERISCHE  
TECHNISCHE HOCHSCHULE  
REGENSBURG







development IDE

↕ *commit*




code repository

Development Environment



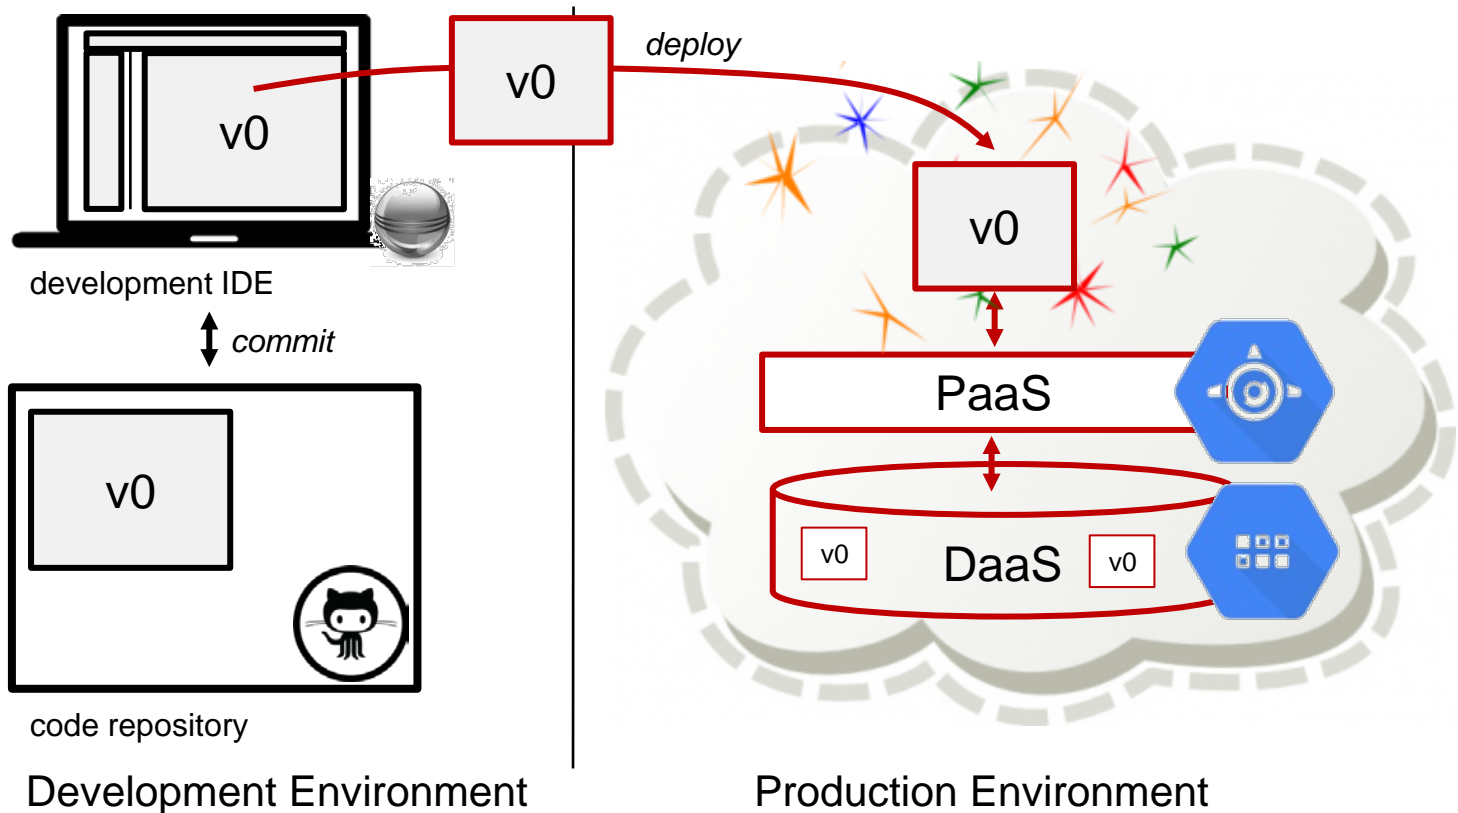

development IDE

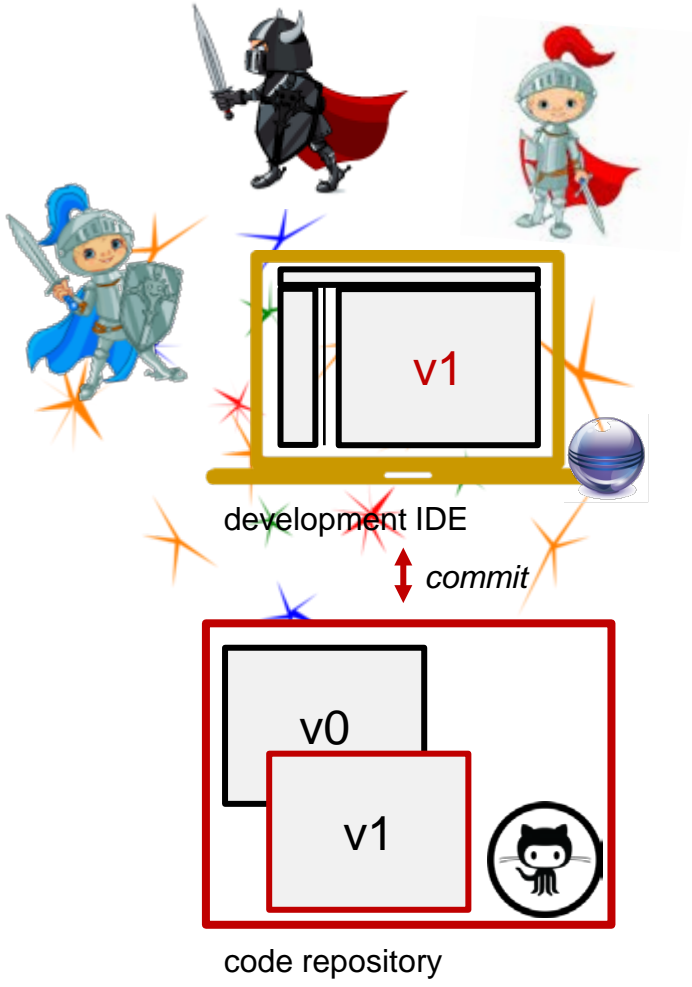

↕ *commit*



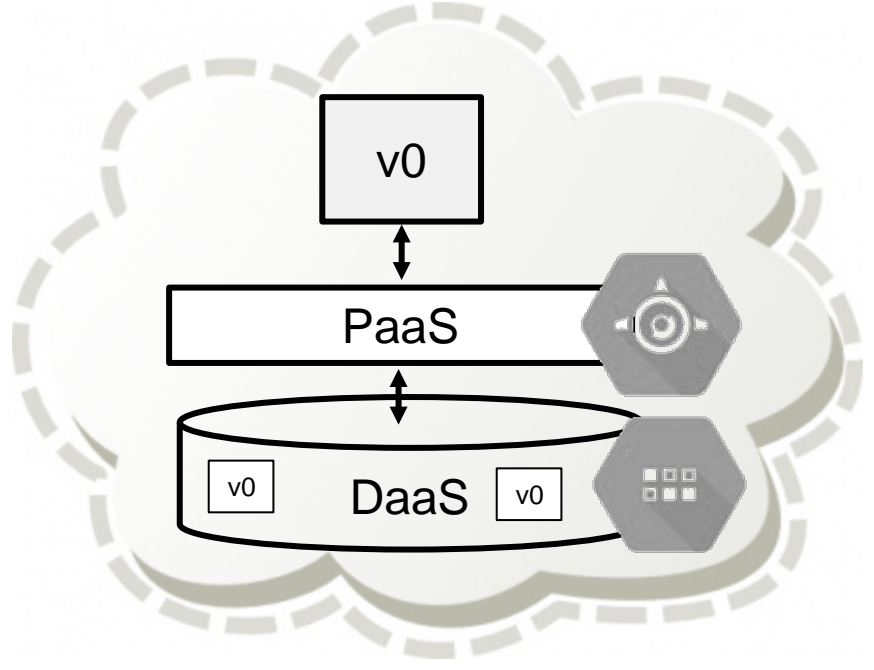
code repository

Development Environment


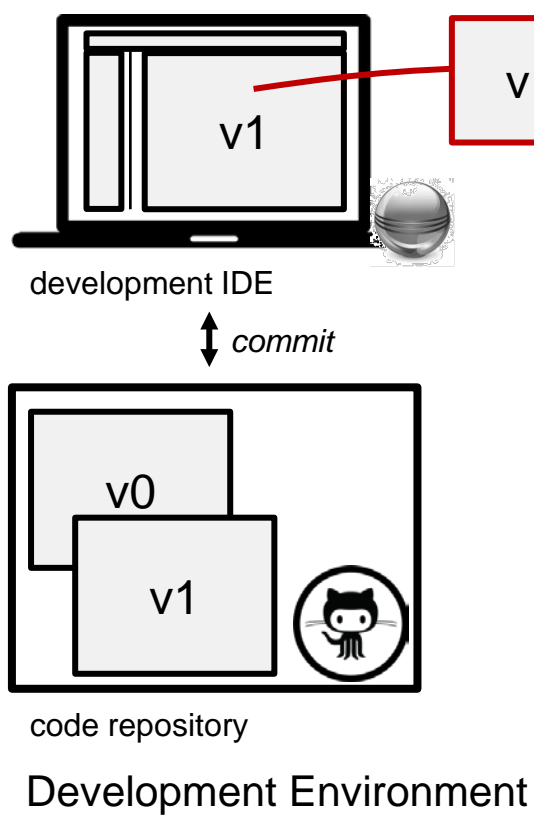





Development Environment



Production Environment



# DESIDERATUM

The application code declares a schema.

The application code evolves.

Thus, we need to address  
**schema evolution:**

- Eager
- Lazy with Object-NoSQL Mappers
- Lazy with Datalution



# EXAMPLE: GAMING APPLICATION

## Release 1


- `Player(ID, NAME)`
- `Mission(ID, TITLE, PID)`

## Release 2

- Players carry a property `SCORE`:  
add `Player.SCORE = 50`


## Release 3

- Missions carry their player's score  
copy `Player.SCORE` to `Mission`  
where `Player.ID = Mission.PID`

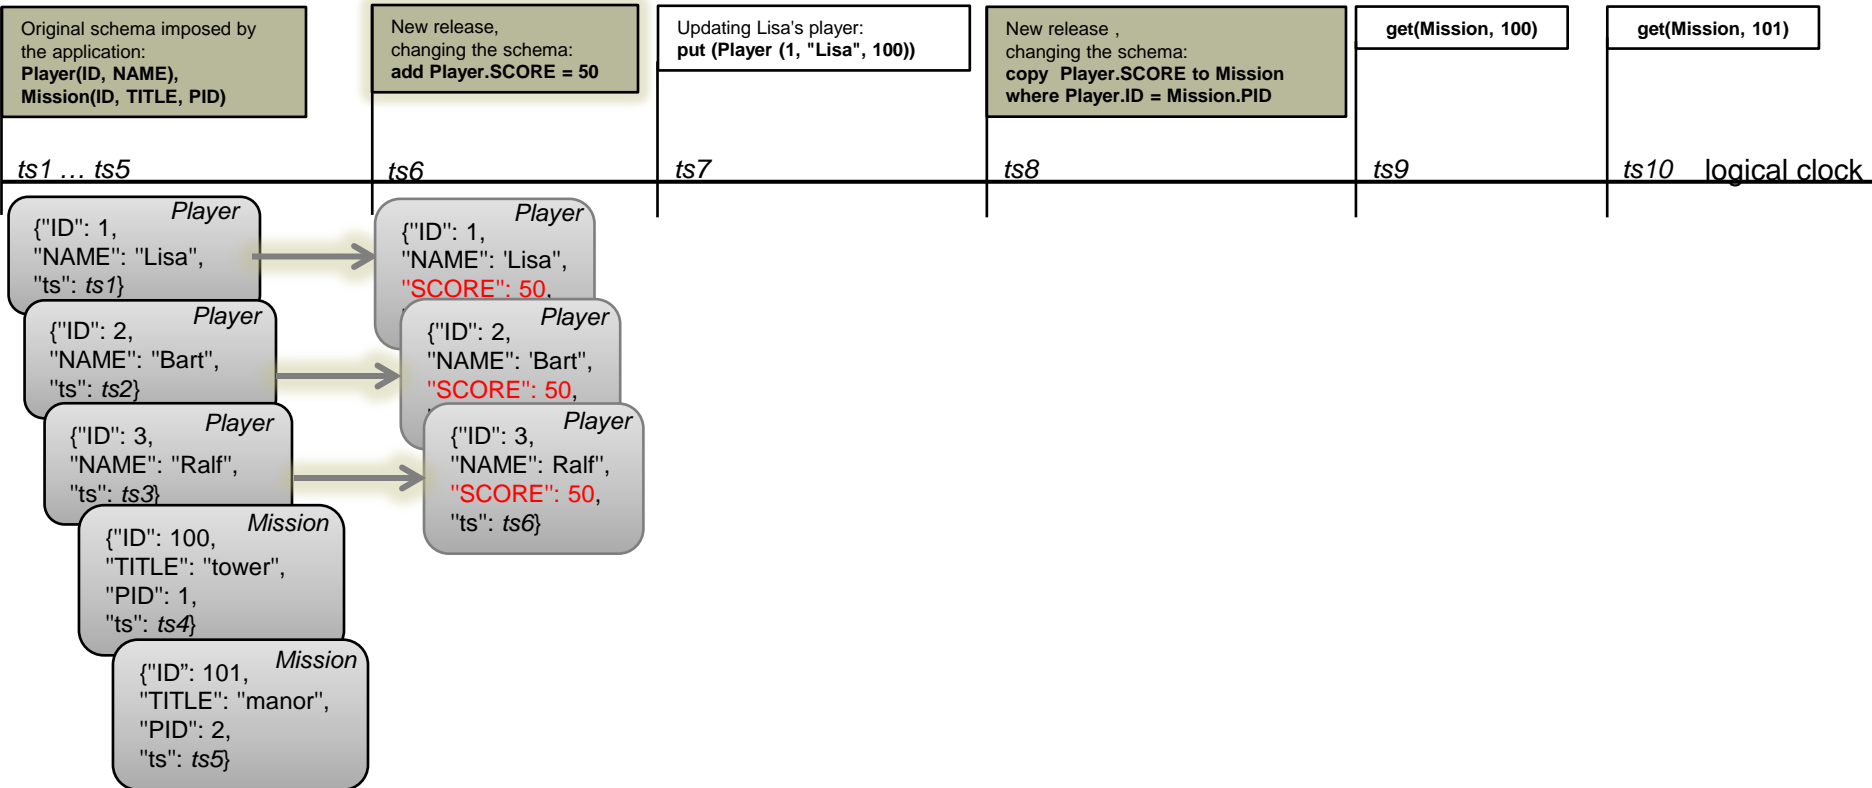




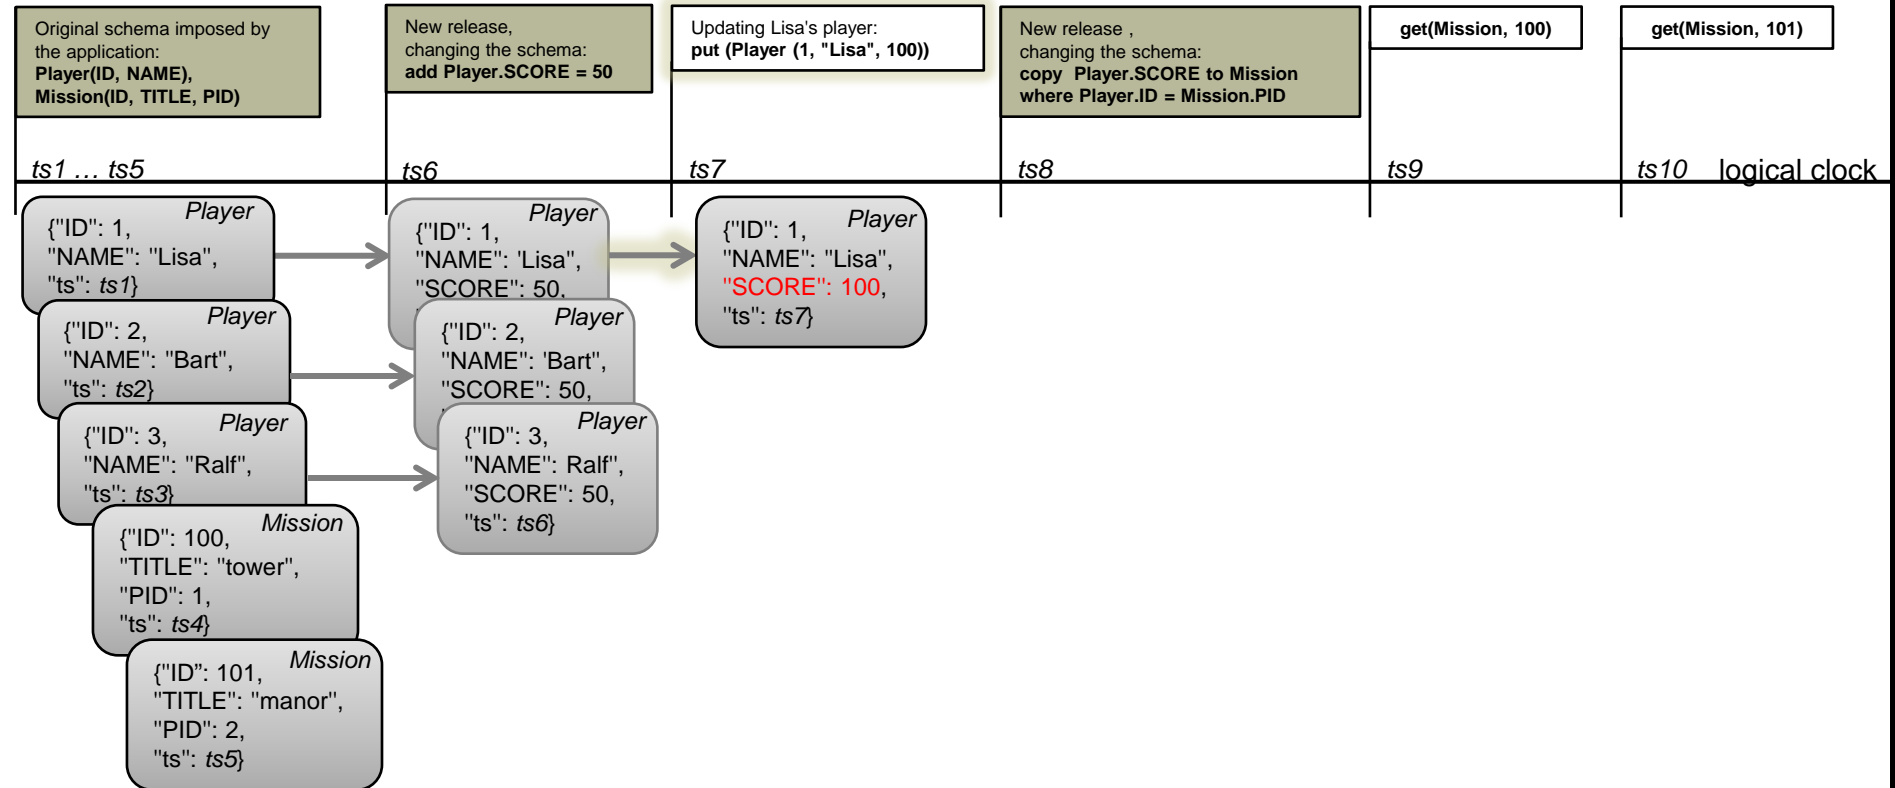
# EAGER MIGRATION



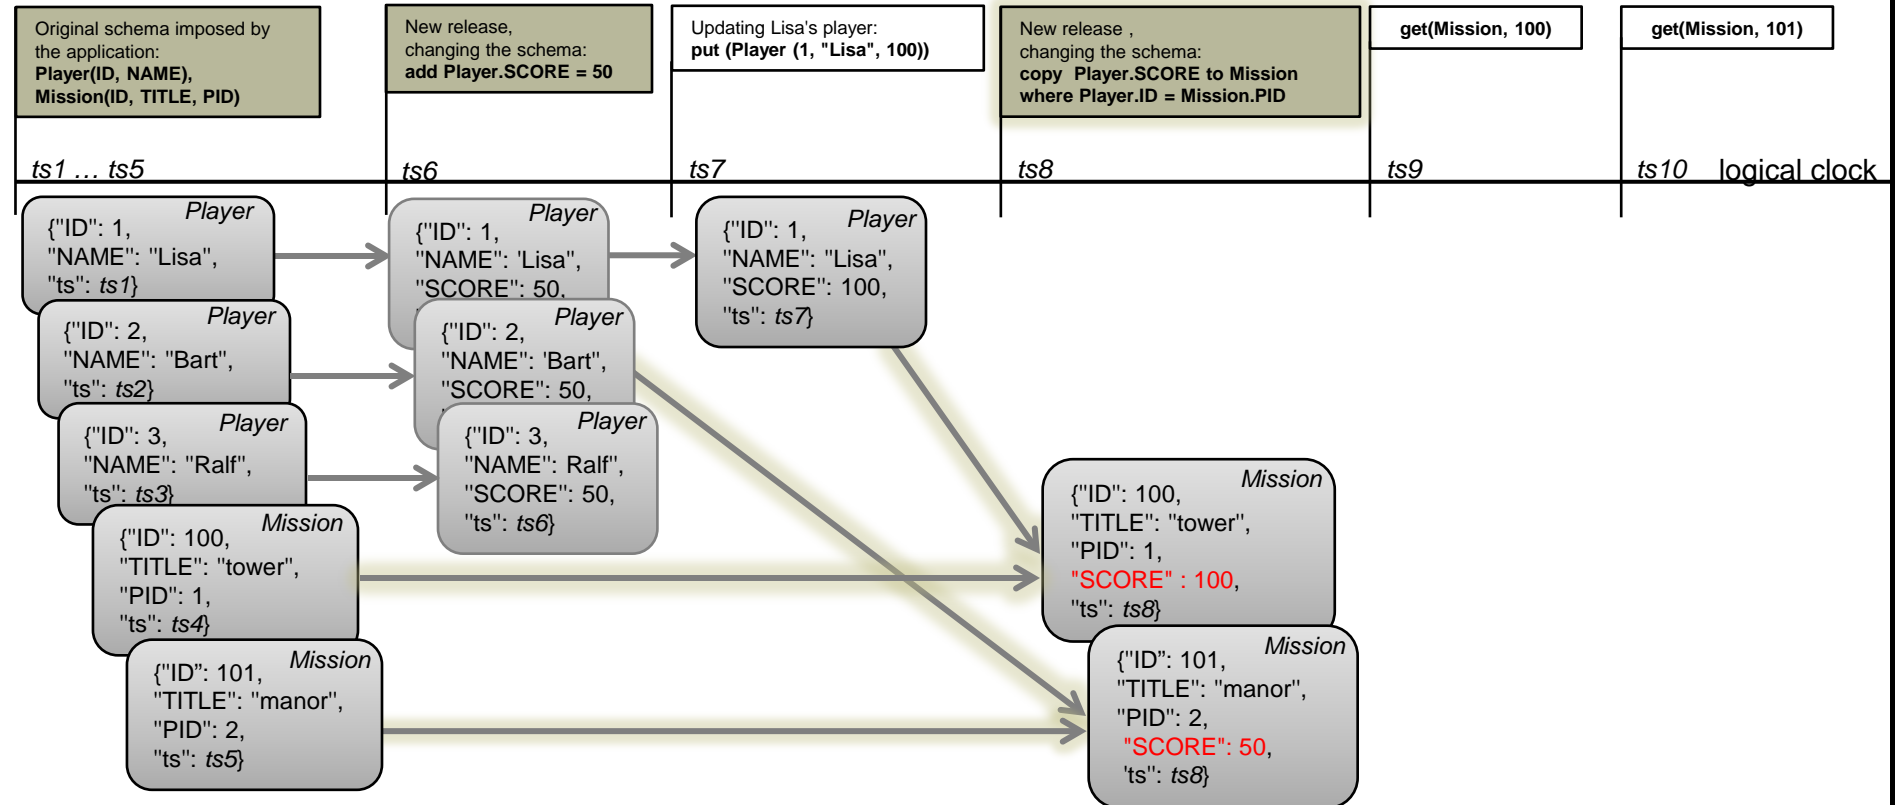
# EAGER MIGRATION



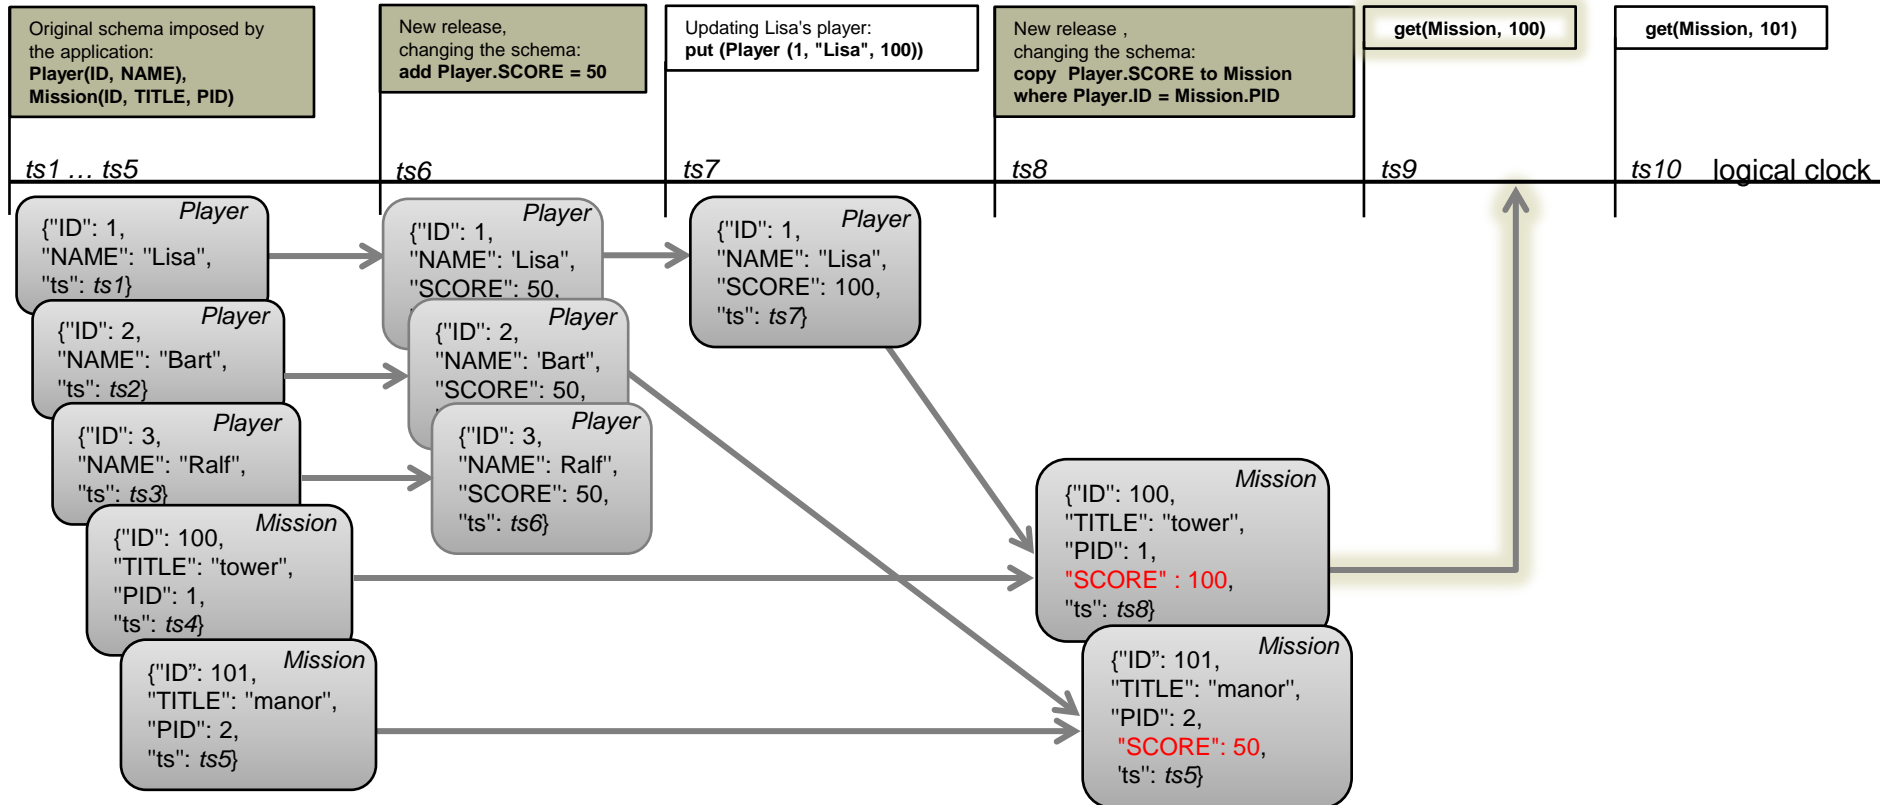
# EAGER MIGRATION



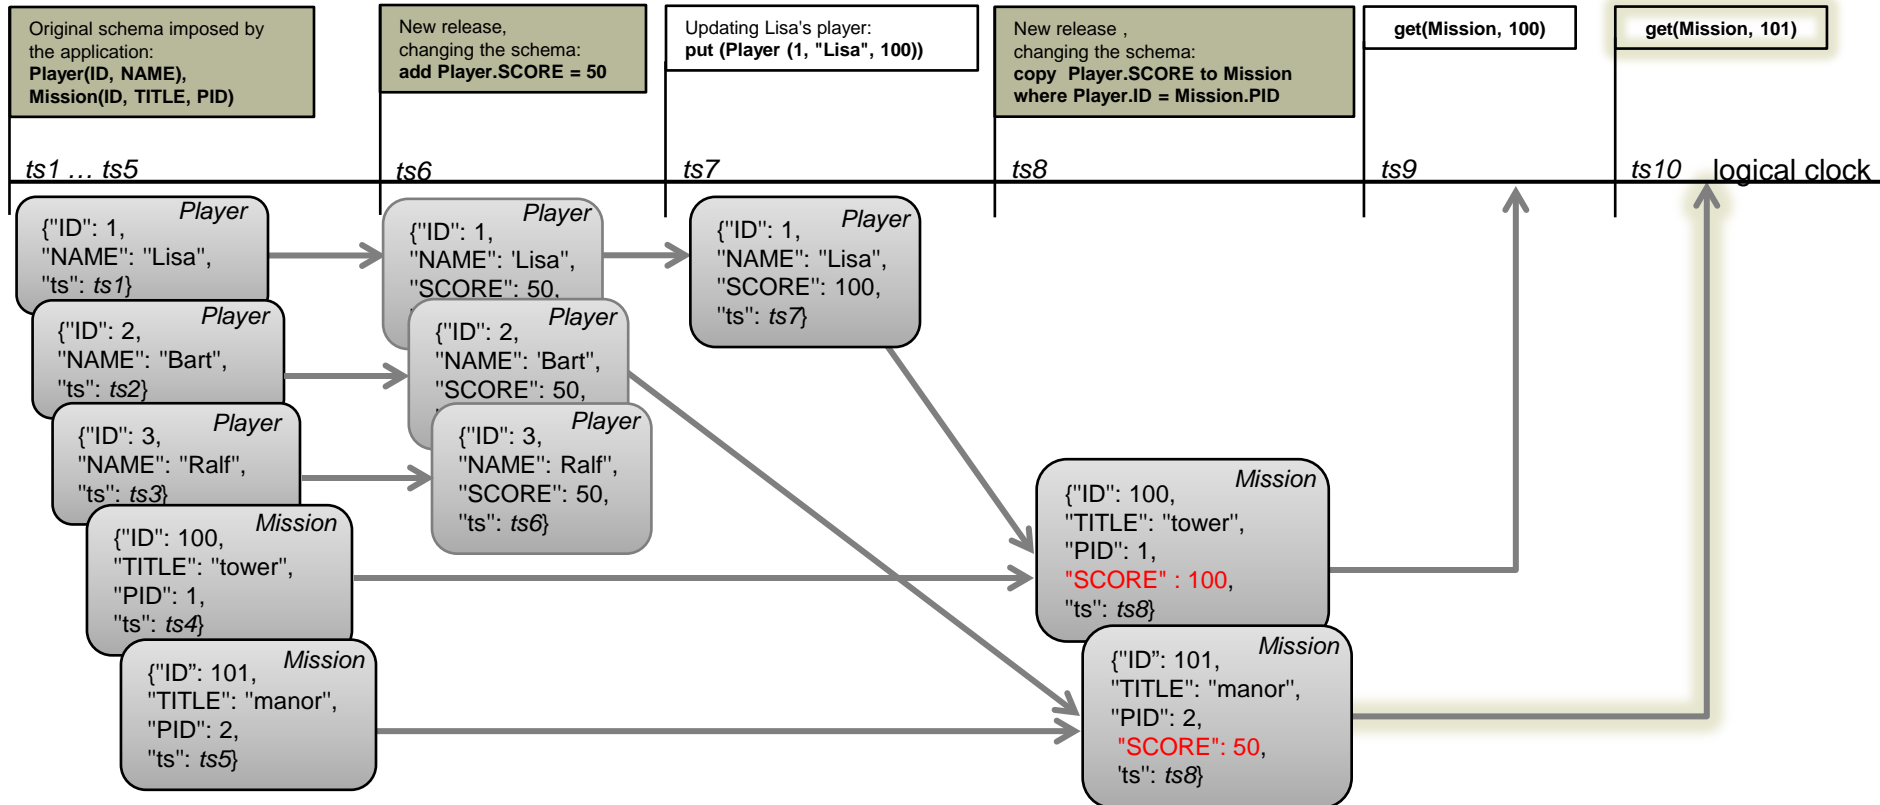
# EAGER MIGRATION



# EAGER MIGRATION



# EAGER MIGRATION



# LAZY EVOLUTION WITH OBJECT-NOSQL MAPPERS

Original schema imposed by the application:

```
@Entity
class Player{
  @Id
  Integer ID;

  String NAME;
  ...
}
```

New release, changing the schema:

```
@Entity
class Player{
  @Id
  Integer ID;

  String NAME;
  Integer SCORE = 50;
}
```

*get(Player, 1)*

*put(...)*

logical clock

Player  
{ "ID": 1, "NAME": "Lisa" }

Player  
{ "ID": 2, "NAME": "Bart" }

Player  
{ "ID": 3, "NAME": "Ralf" }


Player  
{ "ID": 1, "NAME": "Lisa", "SCORE": 50 }

Convenient "quick fix"  
for simple changes.

Long-term:  
Maintenance  
nightmare.




# LAZY MIGRATION IN DATA LUTION

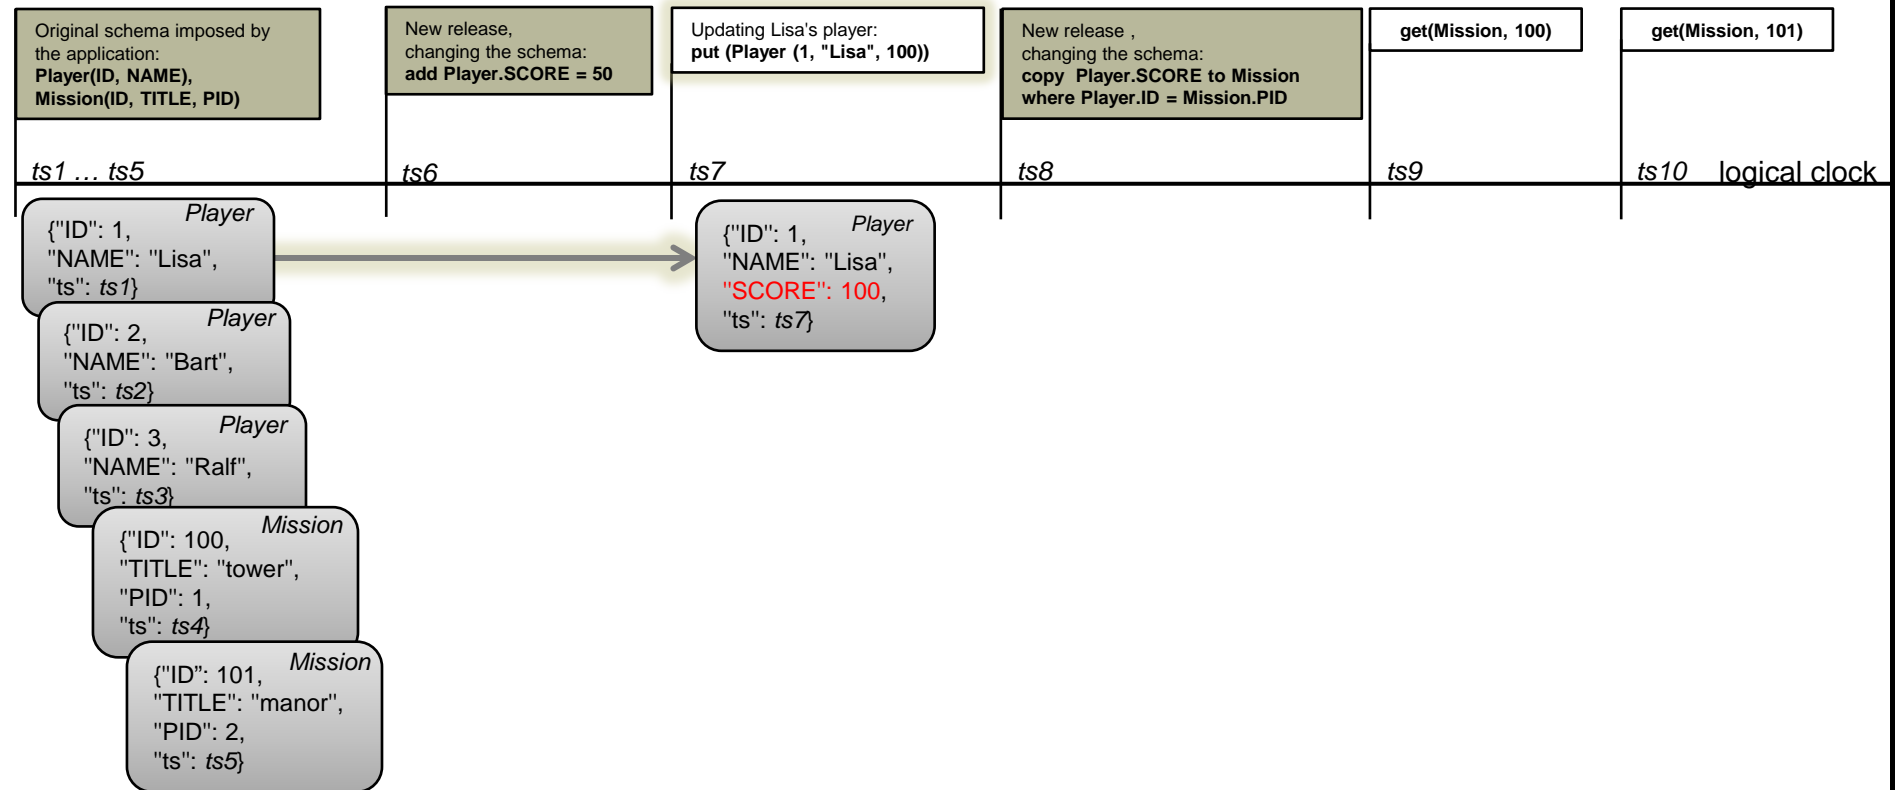




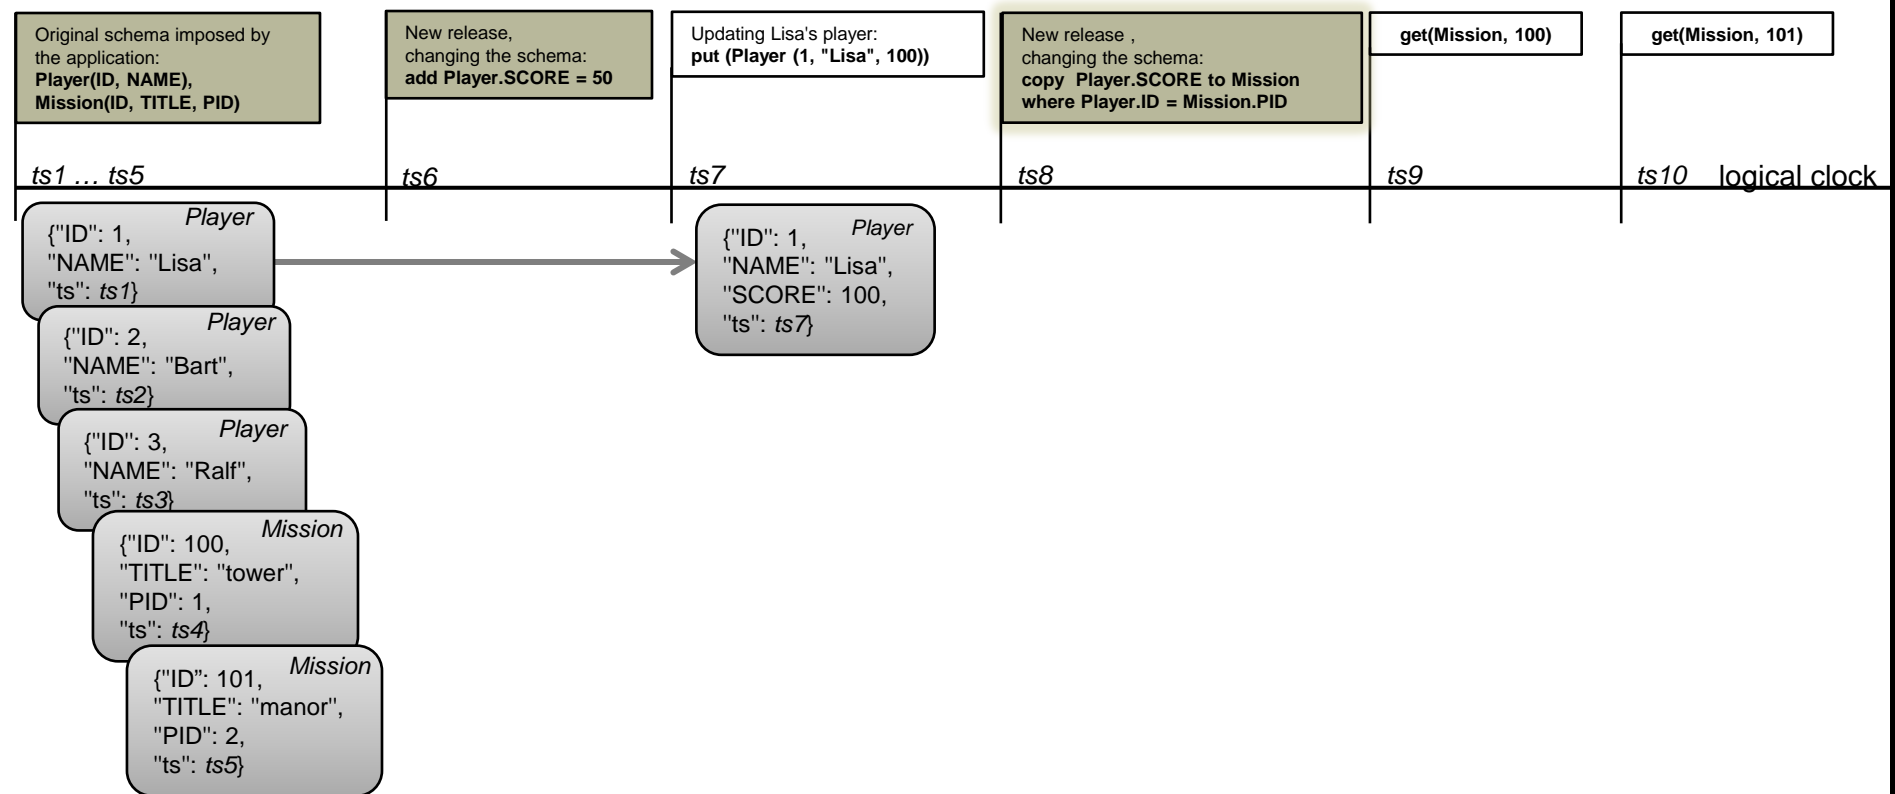
# LAZY MIGRATION IN DATALUTION



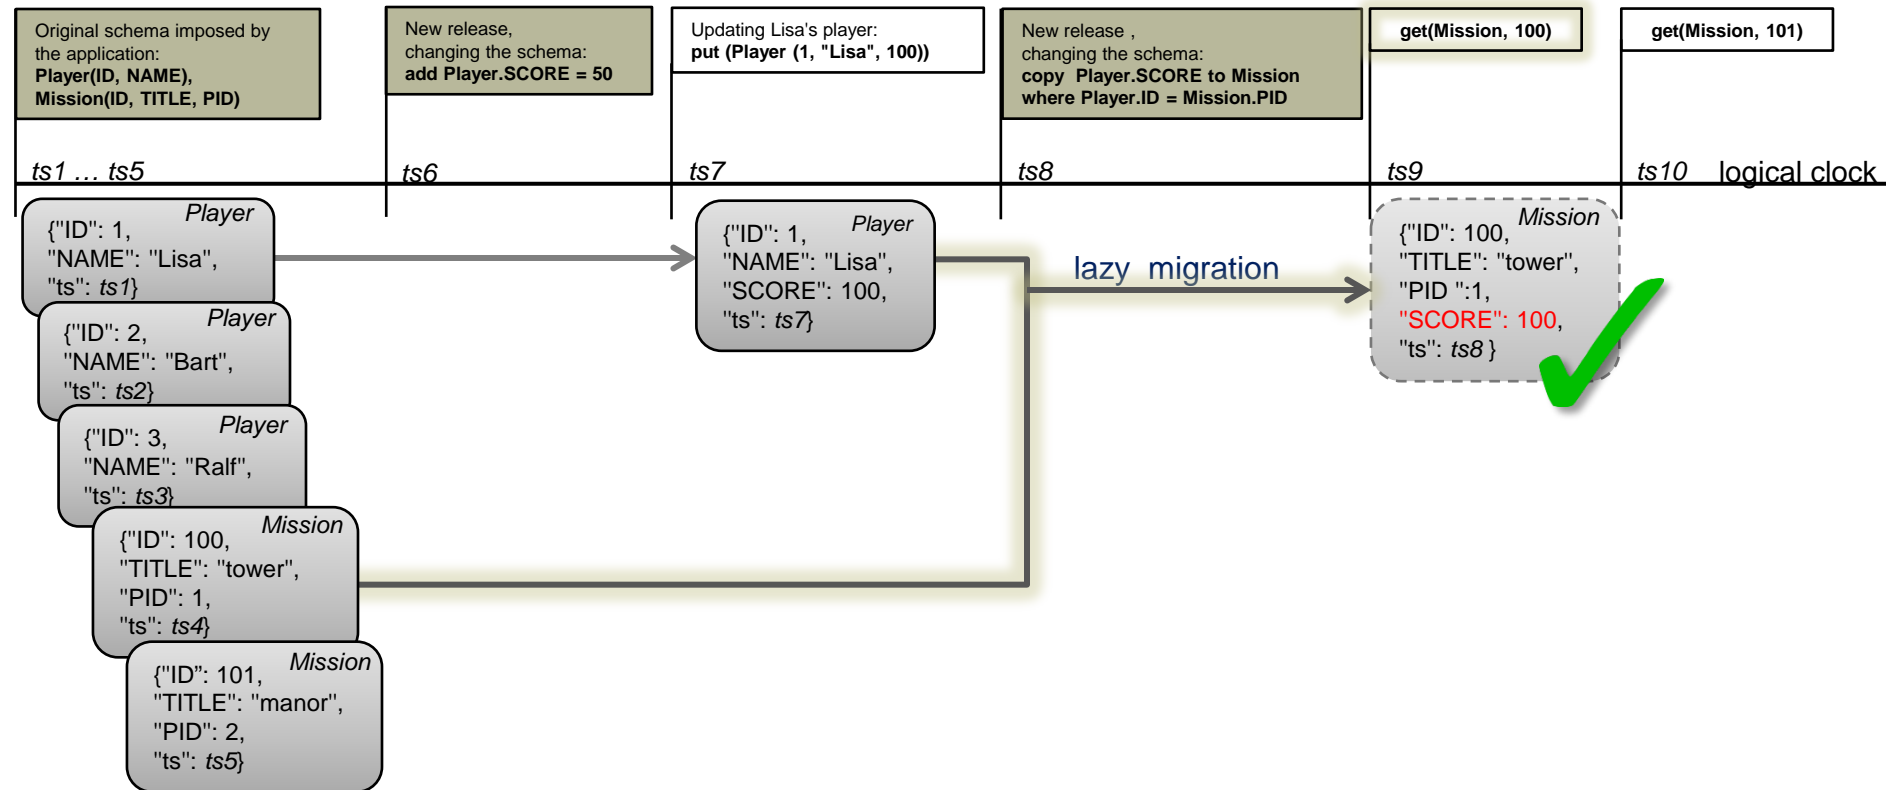
# LAZY MIGRATION IN DATALUTION



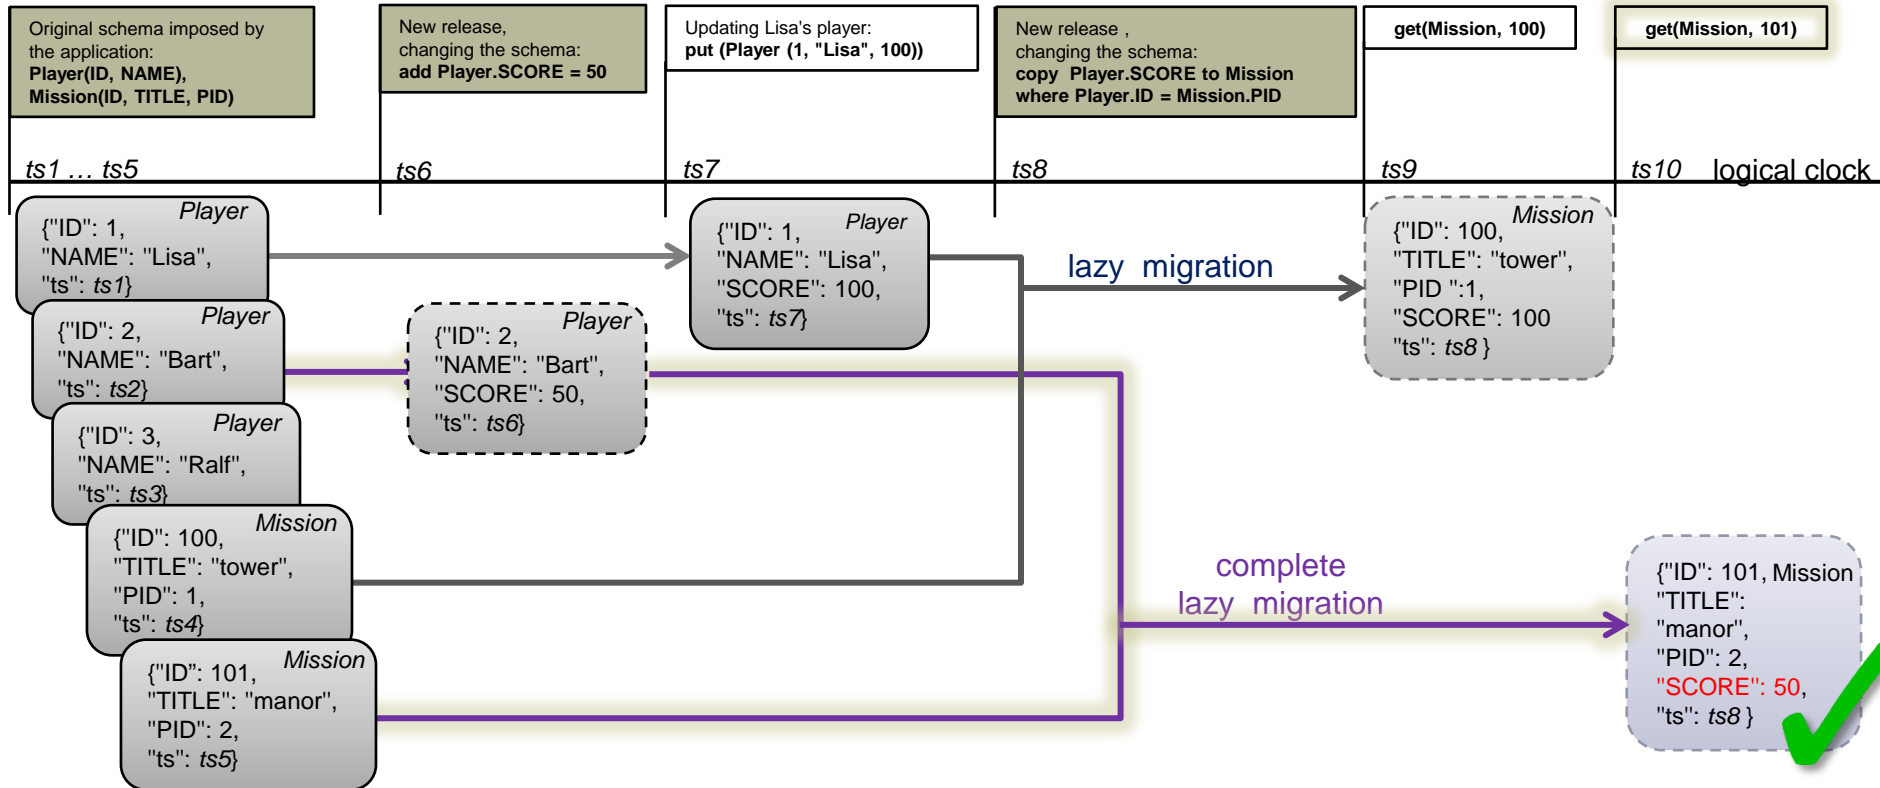
# LAZY MIGRATION IN DATALUTION



# LAZY MIGRATION IN DATALUTION



# LAZY MIGRATION IN DATALUTION



# DATALOG MODEL (NONRECURSIVE, STRATIFIED)

```
a1: put(Player(1, "Lisa"));
a2: put(Player(1, "Lisa S."));
```

```
    r1: Player(1, "Lisa", ts1).
    r2: Player(1, "Lisa S.", ts2).
```

```
a3: get("Player", 1);
```

```
    r3: legacyPlayer(ID, TS) :-
        Player(ID, _, TS), Player(ID, _, NTS), TS < NTS.
    r4: latestPlayer(ID, TS) :-
        Player(ID, _, TS), not legacyPlayer(ID, TS).
    r5: getPlayer(ID, NAME, TS) :-
        Player(ID, NAME, TS), latestPlayer(ID, TS).
```

transient rule – derived facts not kept around for incremental evaluation

Let  $\text{kind}[r](ID, P_1, \dots, P_n)$  be the schema imposed by the current application release.  $ts$  denotes a fresh timestamp associated with release  $r$ .

i) **add**  $\text{kind}.P_{n+1} = v$ , where  $P_{n+1}$  is a new property name and  $v$  is a default value (in the new version of the entity,  $P_{n+1}$  has value  $v$ ):  
 $\overline{\text{kind}[r+1]}(ID, P_1, \dots, P_n, v, ts) \quad :- \quad \text{kind}[r](ID, P_1, \dots, P_n, \text{OTS}), \text{latestkind}[r](ID, \text{OTS}) .$

ii) **delete**  $\text{kind}.P_i$   
 $\overline{\text{kind}[r+1]}(ID, P_1, \dots, P_{(i-1)}, P_{(i+1)}, \dots, P_n, ts) \quad :- \quad \text{kind}[r](ID, P_1, \dots, P_n, \text{OTS}), \text{latestkind}[r](ID, \text{OTS}) .$

Let  $\text{kindS}[r](ID, S_1, \dots, S_n)$  and  $\text{kindT}[r](ID, T_1, \dots, T_m)$  be the current source and target schema imposed by the application.

iii) **copy**  $\text{kindS}.S_i$  **to**  $\text{kindT}$  where  $\text{kindS}.ID = \text{kindT}.T_j$   
 $\overline{\text{kindT}[r+1]}(ID\_T, T_1, \dots, T_m, S_i, ts) \quad :- \quad \text{kindT}[r](ID\_T, T_1, \dots, T_m, \text{TS}_T), \text{latestkindT}[r](ID\_T, \text{TS}_T),$   
 $\quad \quad \quad \text{kindS}[r](ID\_S, S_1, \dots, S_n, \text{TS}_S), \text{latestkindS}[r](ID\_S, \text{TS}_S), ID\_S = T_j .$   
 $\text{kindT}[r+1](ID\_T, T_1, \dots, T_m, \text{null}, ts) \quad :- \quad \text{kindT}[r](ID\_T, T_1, \dots, T_m, \text{TS}_T), \text{latestkindT}[r](ID\_T, \text{TS}_T),$   
 $\quad \quad \quad \text{not kindS}[r](ID\_S, S_1, \dots, S_n, \text{TS}_S), ID\_S = T_j .$   
 $\text{kindS}[r+1](ID, S_1, \dots, S_n, ts) \quad \quad \quad :- \quad \text{kindS}[r](ID, S_1, \dots, S_n, \text{OTS}), \text{latestkind}[r](ID, \text{OTS}) .$

iv) **move**  $\text{kindS}.S_i$  **to**  $\text{kindT}$  where  $\text{kindS}.ID = \text{kindT}.T_j$ , with the same first two rules as for copy, as well as the following rule:  
 $\overline{\text{kindS}[r+1]}(ID, S_1, \dots, S_{(i-1)}, S_{(i+1)}, \dots, S_n, ts) \quad :- \quad \text{kindS}[r](ID, S_1, \dots, S_n, \text{OTS}), \text{latestkind}[r](ID, \text{OTS}) .$

# DATALUTION: DATALOG-BASED

- **Eager migration: Incremental bottom-up evaluation**
- **Lazy migration: Incremental top-down evaluation**
  - Employing sideways information passing strategies
  - Exploiting uniqueness of identifiers
- **Both strategies always yield the same result**
- **Progress:**
  - Theory in DBPL@SPLASH'15 paper
  - Demo of PoC Datalution at QUDOS'16
  - Ongoing: Integration with NoSQL data store

