

# **Sanity: The Less Server Architecture for Cloud functions**

Shripad J Nadgowda, Nilton Bila, Canturk Isci

IBM T J Watson Research Center

# Agenda


- Background ( check )
- Motivation ( for Change )
- Sanity Architecture ( What, Why and How ? )
- ( Validation by ) Evaluation
- ( Take away ) Conclusion
- ( Open for ) Discussion

# Getting on the same page

—————> Data Read/Write      - - - - -> Function trigger

Event/Data Sources

Compute Platform




# Getting on the same page

—————> Data Read/Write      - - - - -> Function trigger

Event/Data Sources

Compute Platform



# Eureka Moment!

*What if*

Data from the input set are **Equivalent** And

associated functions **Deterministic** (or idempotent)

*Then,*

can we **Avoid** execution of functions, and still

**De-duplicate** the output results ?



# Sincere tribute


## **Insanity:**

Doing same thing over and over again, and expecting different results



## **Sanity:**

De-duplicate execution of cloud functions for equivalent data events



# Validation: Equivalent data

- Common data sources for Cloud functions:
  - IoT/Sensor data (e.g. weather), social media (e.g. tweets), user-activity (e.g. click stream), system monitor data (e.g. Prometheus)
- Bounded range of values
  - E.g. temperature data to be (-20C to 50C)
- Temporal duplication
  - E.g. data from a fixed sensors, system monitors
- Spatial duplication
  - E.g. data from geo-distributed sensors
- Semantically equivalent data

# Validation: Deterministic functions

[AWS Documentation](#) » [Amazon EC2](#) » [API Reference](#) » [Making API Requests](#) » Ensuring Idempotency

## Ensuring Idempotency

Filter View:

An *idempotent* operation completes no more than one time.



IBM Cloud Functions Apache OpenWhisk

### Action semantics

The following sections describe details about Cloud Functions Actions.

### Statelessness

Action implementations are stateless, or *idempotent*. While the system does not enforce this property, it is not guaranteed that any state maintained by an Action is available across invocations.

- Delivery of function invocations is not currently guaranteed. As the Cloud Firestore and Cloud Functions integration improves, we plan to guarantee "at least once" delivery. However, this may not always be the case during beta. This may also result in multiple invocations for a single event, so for the highest quality functions ensure that the functions are written to be [idempotent](#).





Google Cloud Platform



# Sanity: Less-server Architecture


—————> Sanity Deduplication      —————> Data Read/Write      - - - - -> Function trigger

Event/Data Sources



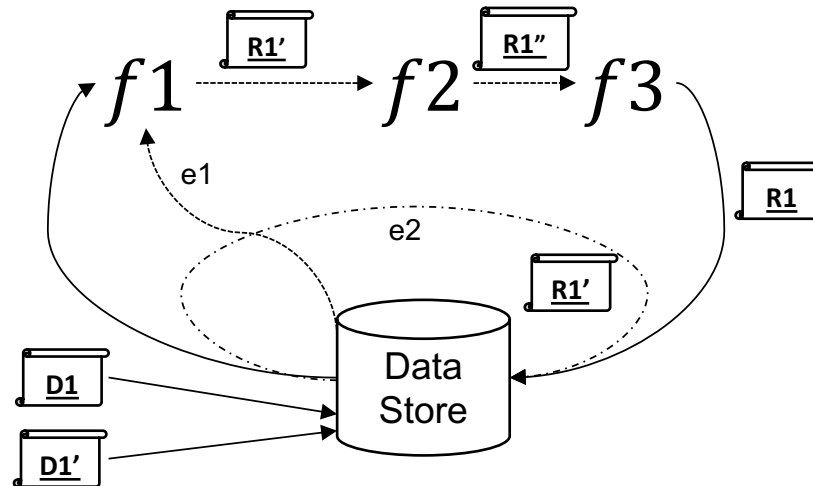
D1 ~ D1'  
D2 ~ D2'

Compute Platform



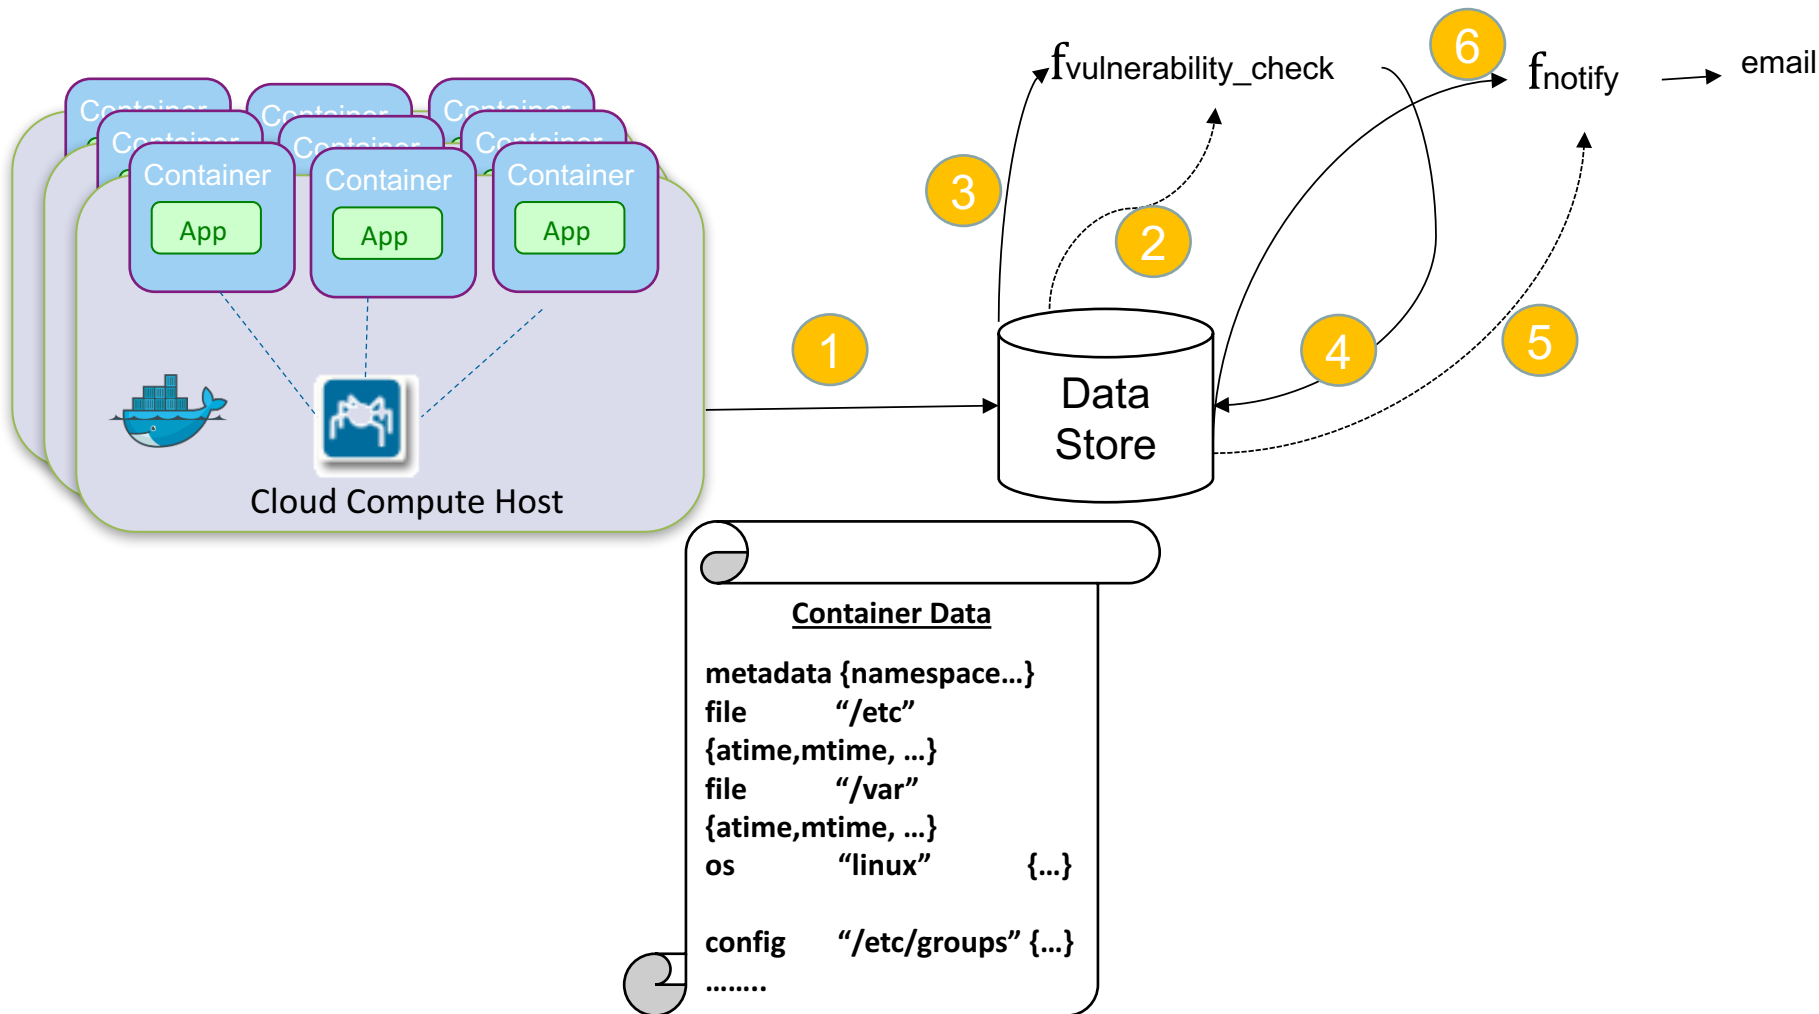
# Sanity: Less\_server Architecture

—————> Sanity Deduplication      —————> Data Read/Write      - - - - -> Function trigger



Extending Sanity to Sequence of Cloud functions

# Sanity Use case: Vulnerability Analysis

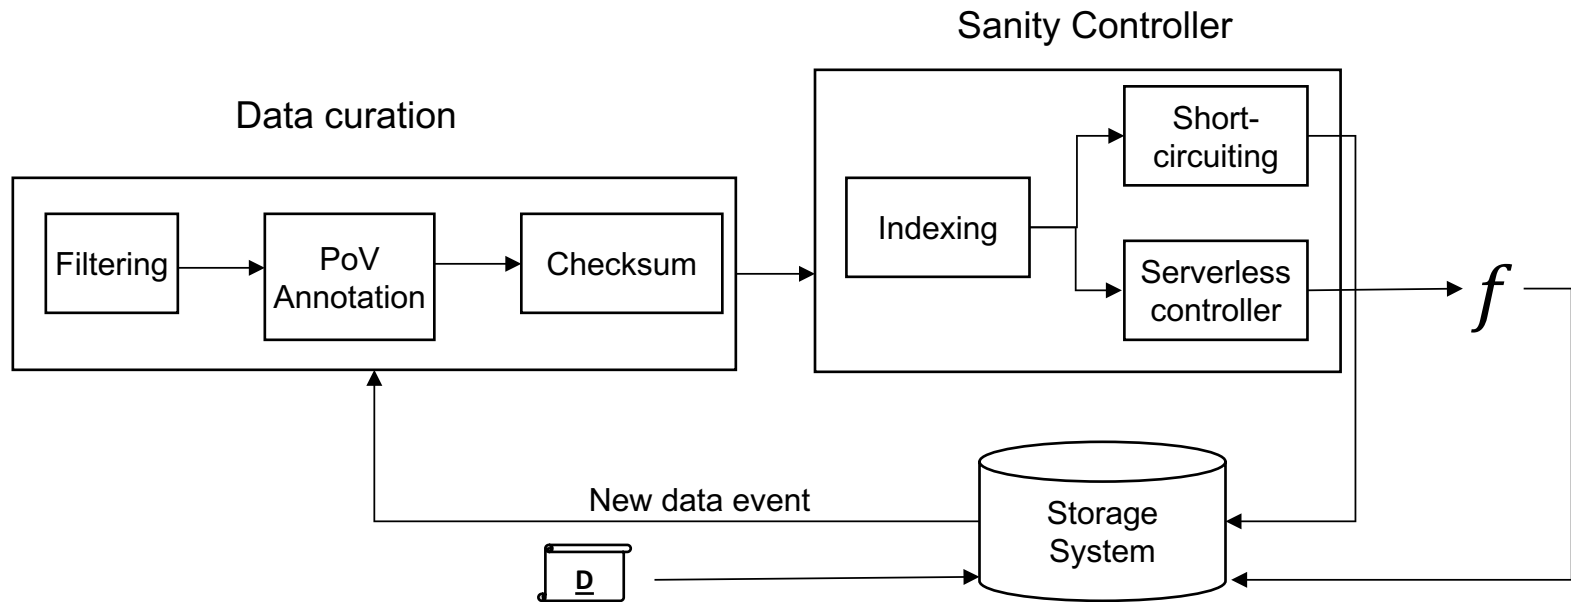


# Sanity: Mind the Gap...

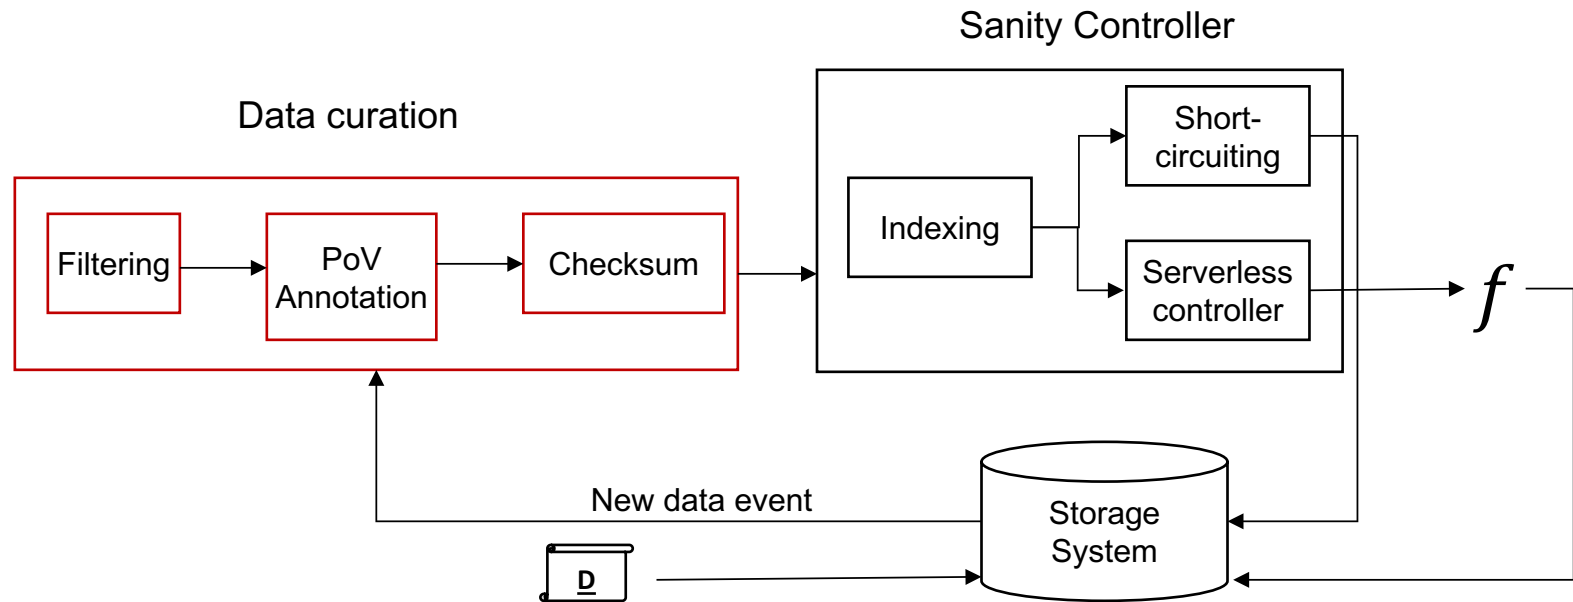


- Considering ONLY storage-closed loop functions
  - reads data from storage
  - writes result back to the storage
- External stimuli are avoided
  - stimulate external events like sending email, slack, SMS etc.

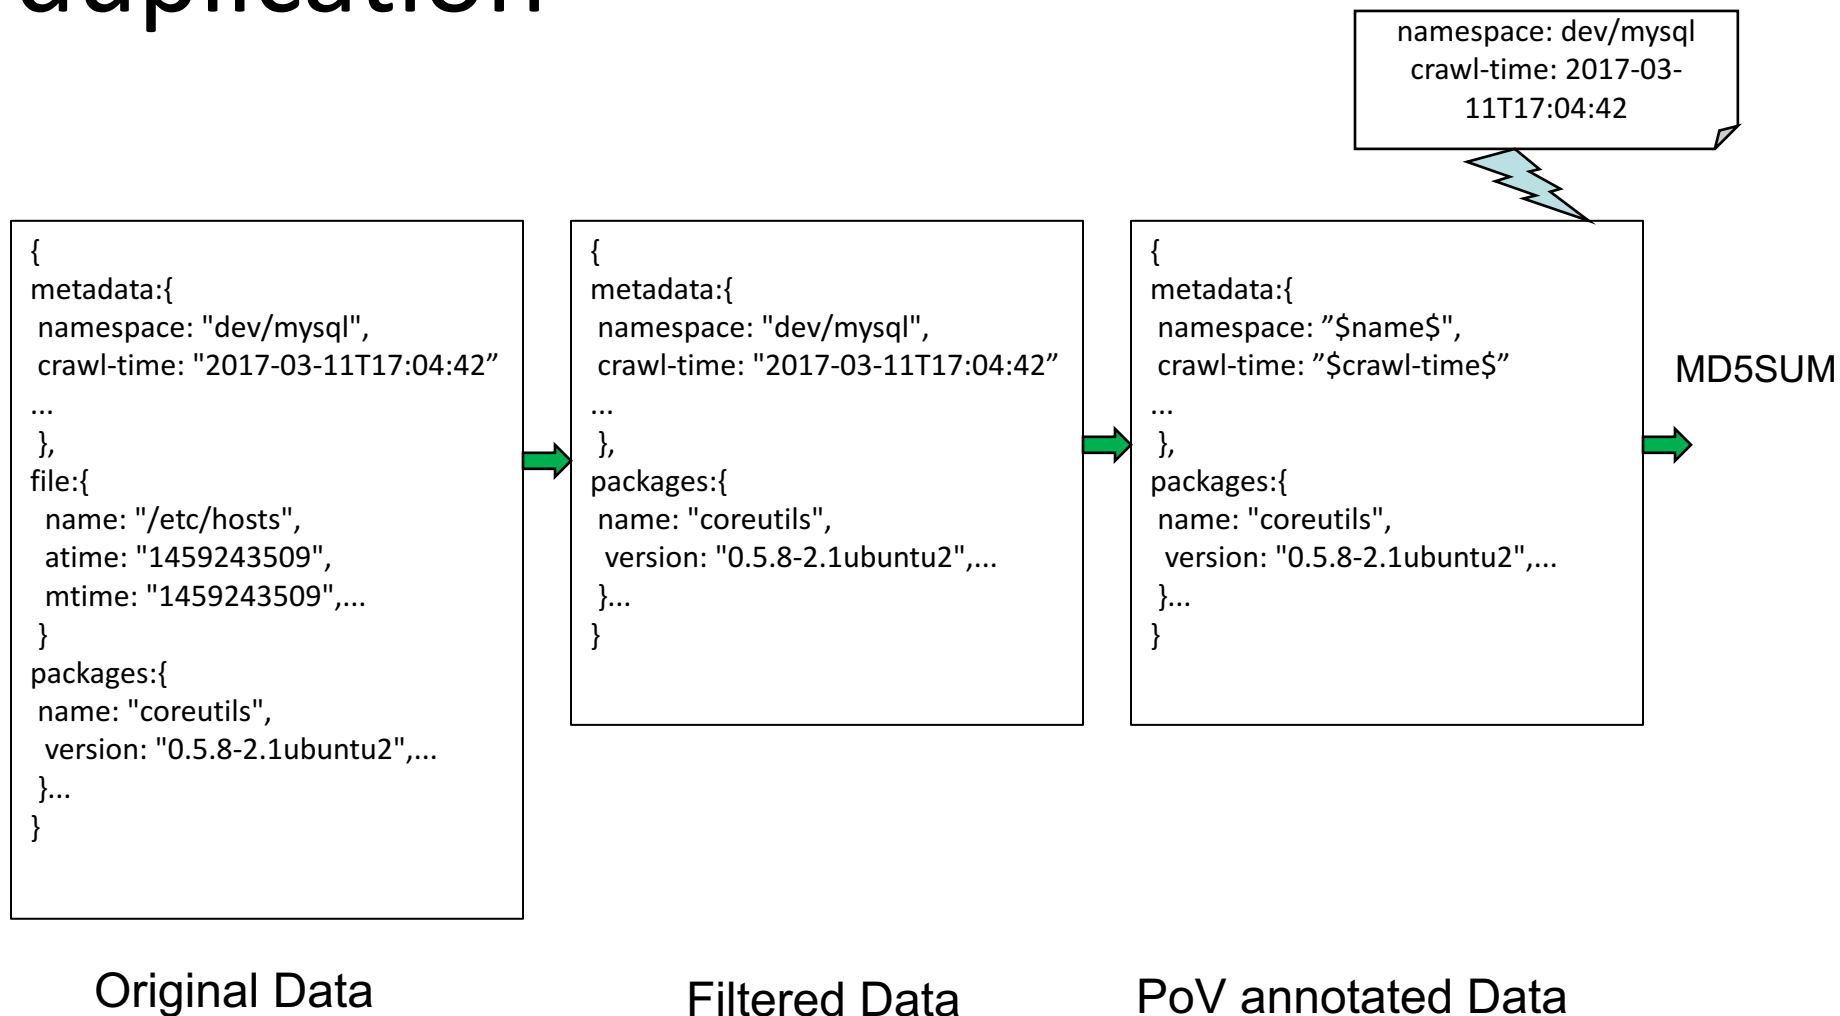
# Sanity Use case: Architecture



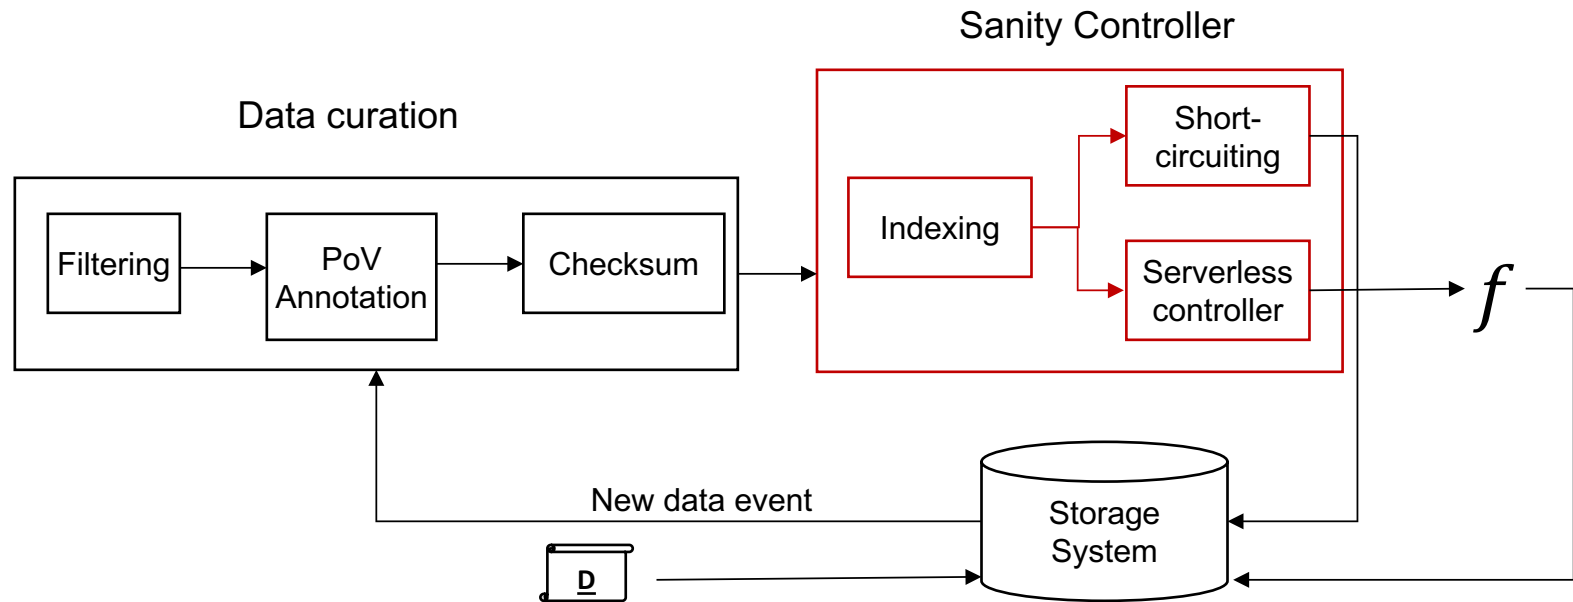
# Sanity Use case: PoV based de-duplication



# Sanity Use case: PoV based de-duplication

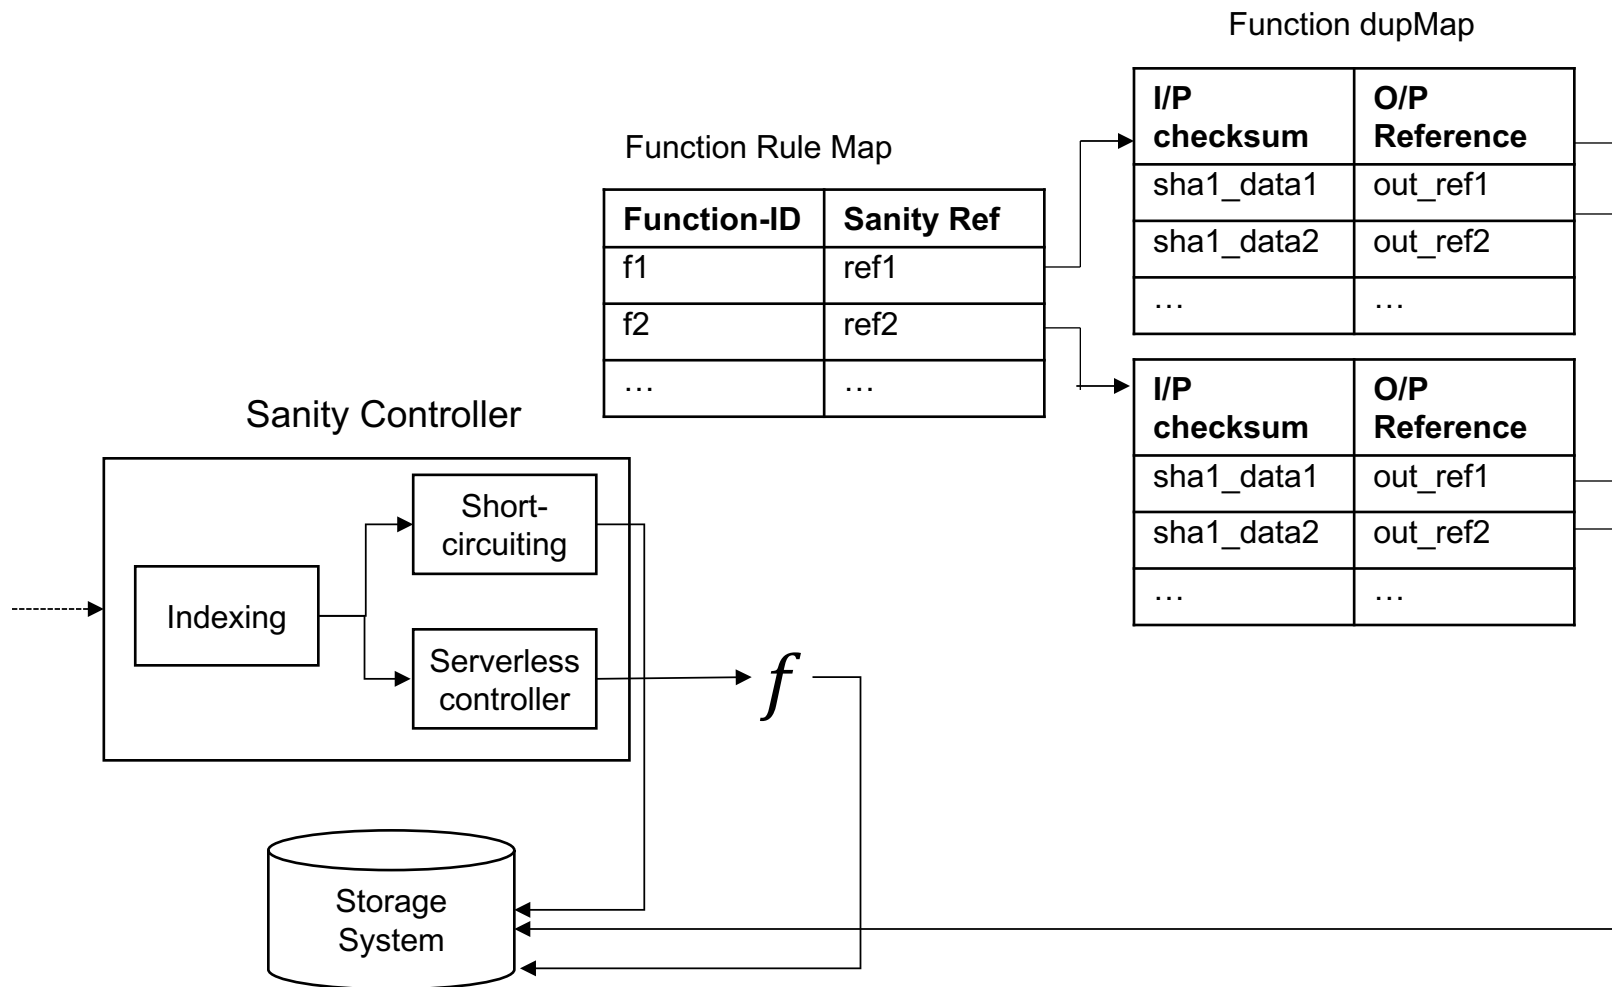


# Sanity Use case: Controller





# Sanity Use case: Controller



2 GB Memory for 40K unique data entries

# Sanity Use case: Evaluation

<b>Computational task</b>	<b>Time</b>
PoV annotation and filtering	4.7 ms
SHA1 computation	0.05 ms
Indexing and output de-duplication	0.02 ms
<b>Total</b>	<b>5 ms</b>

Sanity Deduplication Overhead

<b>Computational task</b>	<b>Time</b>
Input data retrieval	772 ms
Function computation	333 ms
Output data storage	59 ms
<b>Total</b>	<b>1165 ms</b>

Function execution stats

# Conclusion


## Disaggregation with Cloud functions

- Data and Compute are managed independently




## Data events are largely semantically equivalent

- Presenting an opportunity to de-duplicate data



## Cloud functions are commonly deterministic

- Presenting an opportunity to de-duplicate data



- Cloud functions can be efficiently de-duplicated avoiding their redundant execution
- Scale serverless platform by requiring less-server

# Thank You

Contact : [nadgowda@us.ibm.com](mailto:nadgowda@us.ibm.com)