STREAMS IN THE DIGITAL NATIVES

What does a data engineer do anyway?

Messaging: "sipping from the fire hose"

Big Data: many orders of magnitude bigger!

100 (560) 101 - 611



Data Engineering:
Channel the data river
Transform to a usable form

Failure Modes

4111

The system is complex We need a key simplifying model

The Unit Hydrograph (J.C.I Dooge, 1967)





-- confluent

Adoption of Event Streaming

60%

Fortune 100 Companies Using Apache Kafka



Linearly Scalable Architecture



Single topic:

- Many producers machines
- Many consumer machines
- Many Broker machines
 No Bottleneck!

If Messaging remembers, Databases don't have to!

BEFORE







CQRS with Kafka

Using events to build a view (DB, Cache, Stream Processor)

Query by customer Id Events

onfluent



- Event stream is source of truth
- View can be a DB, Cache or Stateful Stream Processor
 - View can be re-derived from the event stream

Contextual Event-Driven Apps

