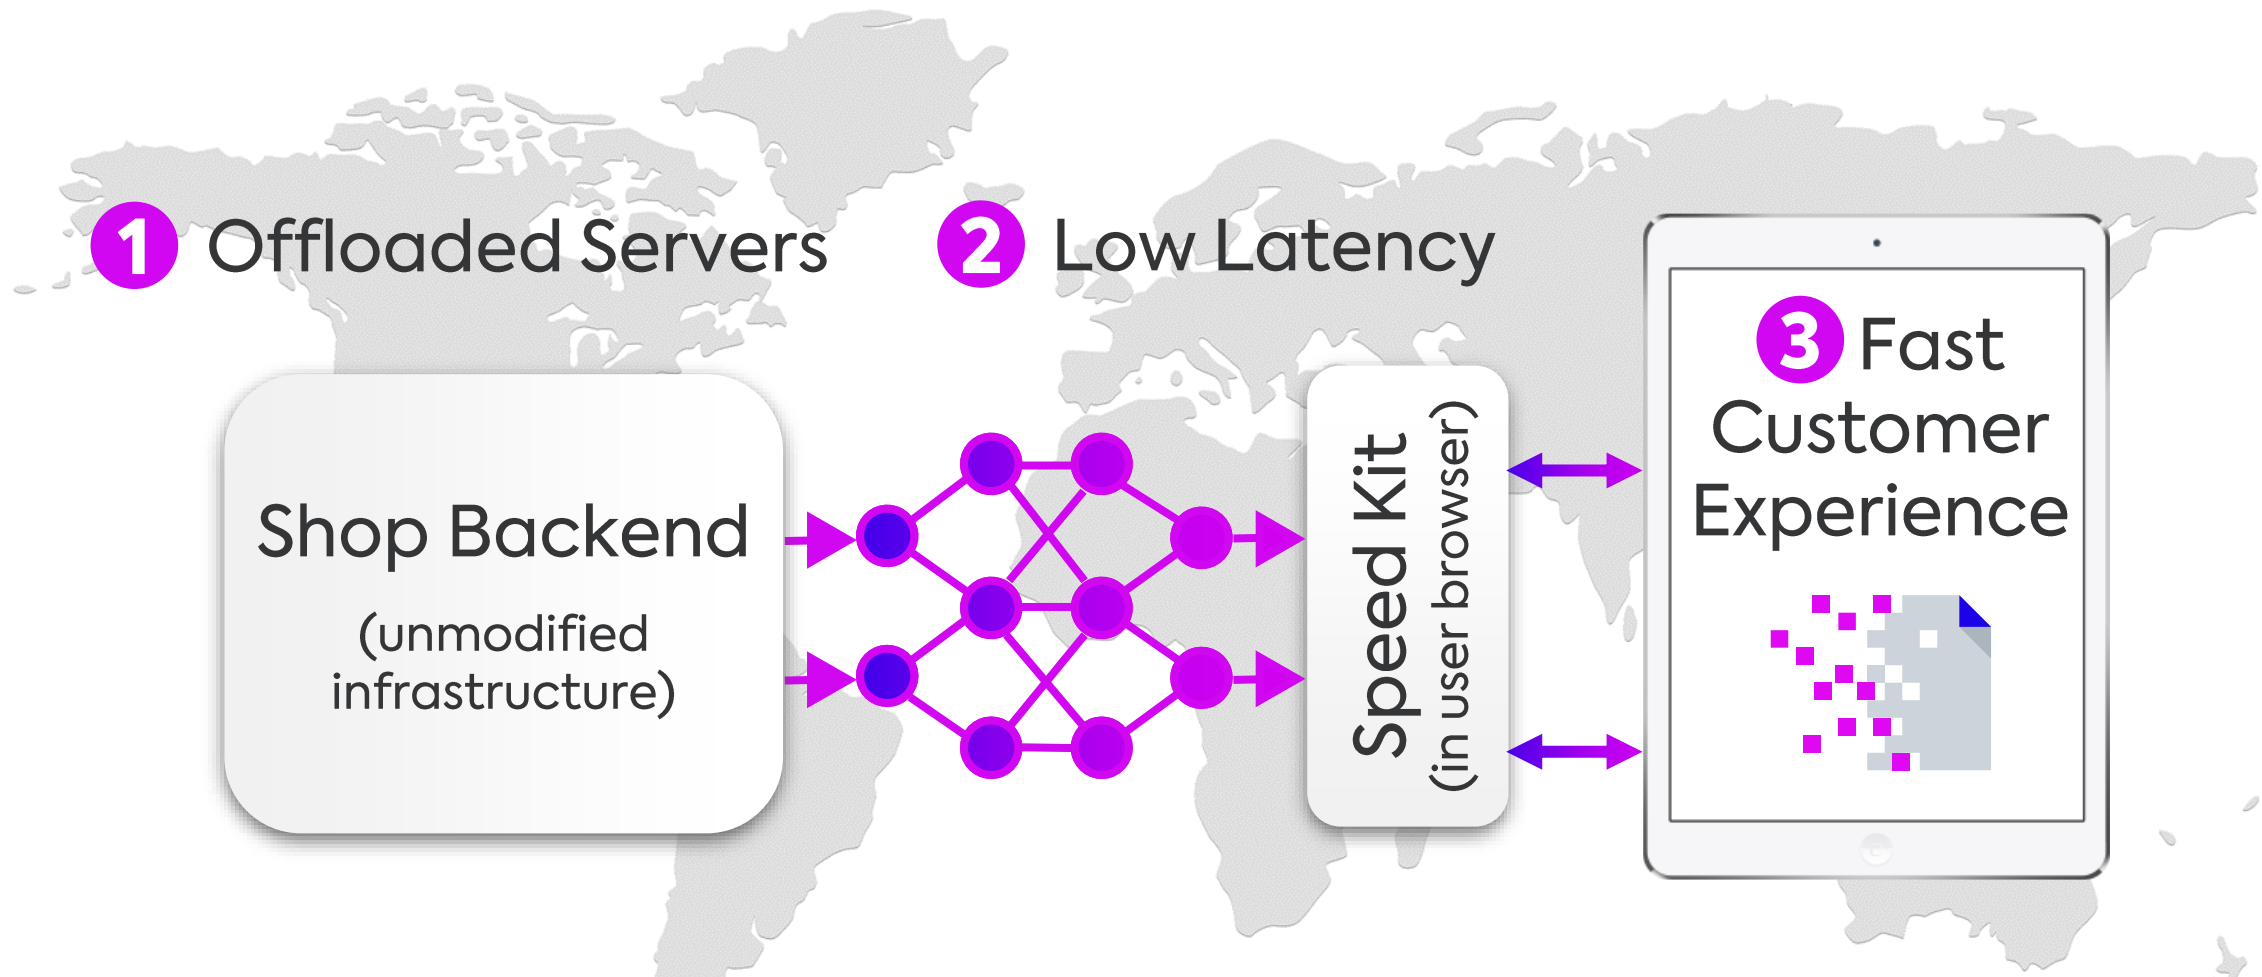



Beaconnect

Continuous Web Performance A/B Testing at Scale

Wolfram Wingerath, Benjamin Wollmer, Markus Bestehorn, Stephan Succo,
Sophie Ferrlein, Florian Bücklers, Jörn Domnik, Fabian Panse,
Erik Witt, Anil Sener, Felix Gessert, Norbert Ritter


Speed Kit: Web Performance Plugin



 Wolfram Wingerath, Felix Gessert, et al.: Speed Kit: A Polyglot & GDPR-Compliant Approach For Caching Personalized Content, ICDE (2020).

Split Testing for Web Performance

Speed Kit Users



Tracking
→

vs.



Tracking
←

Normal Users

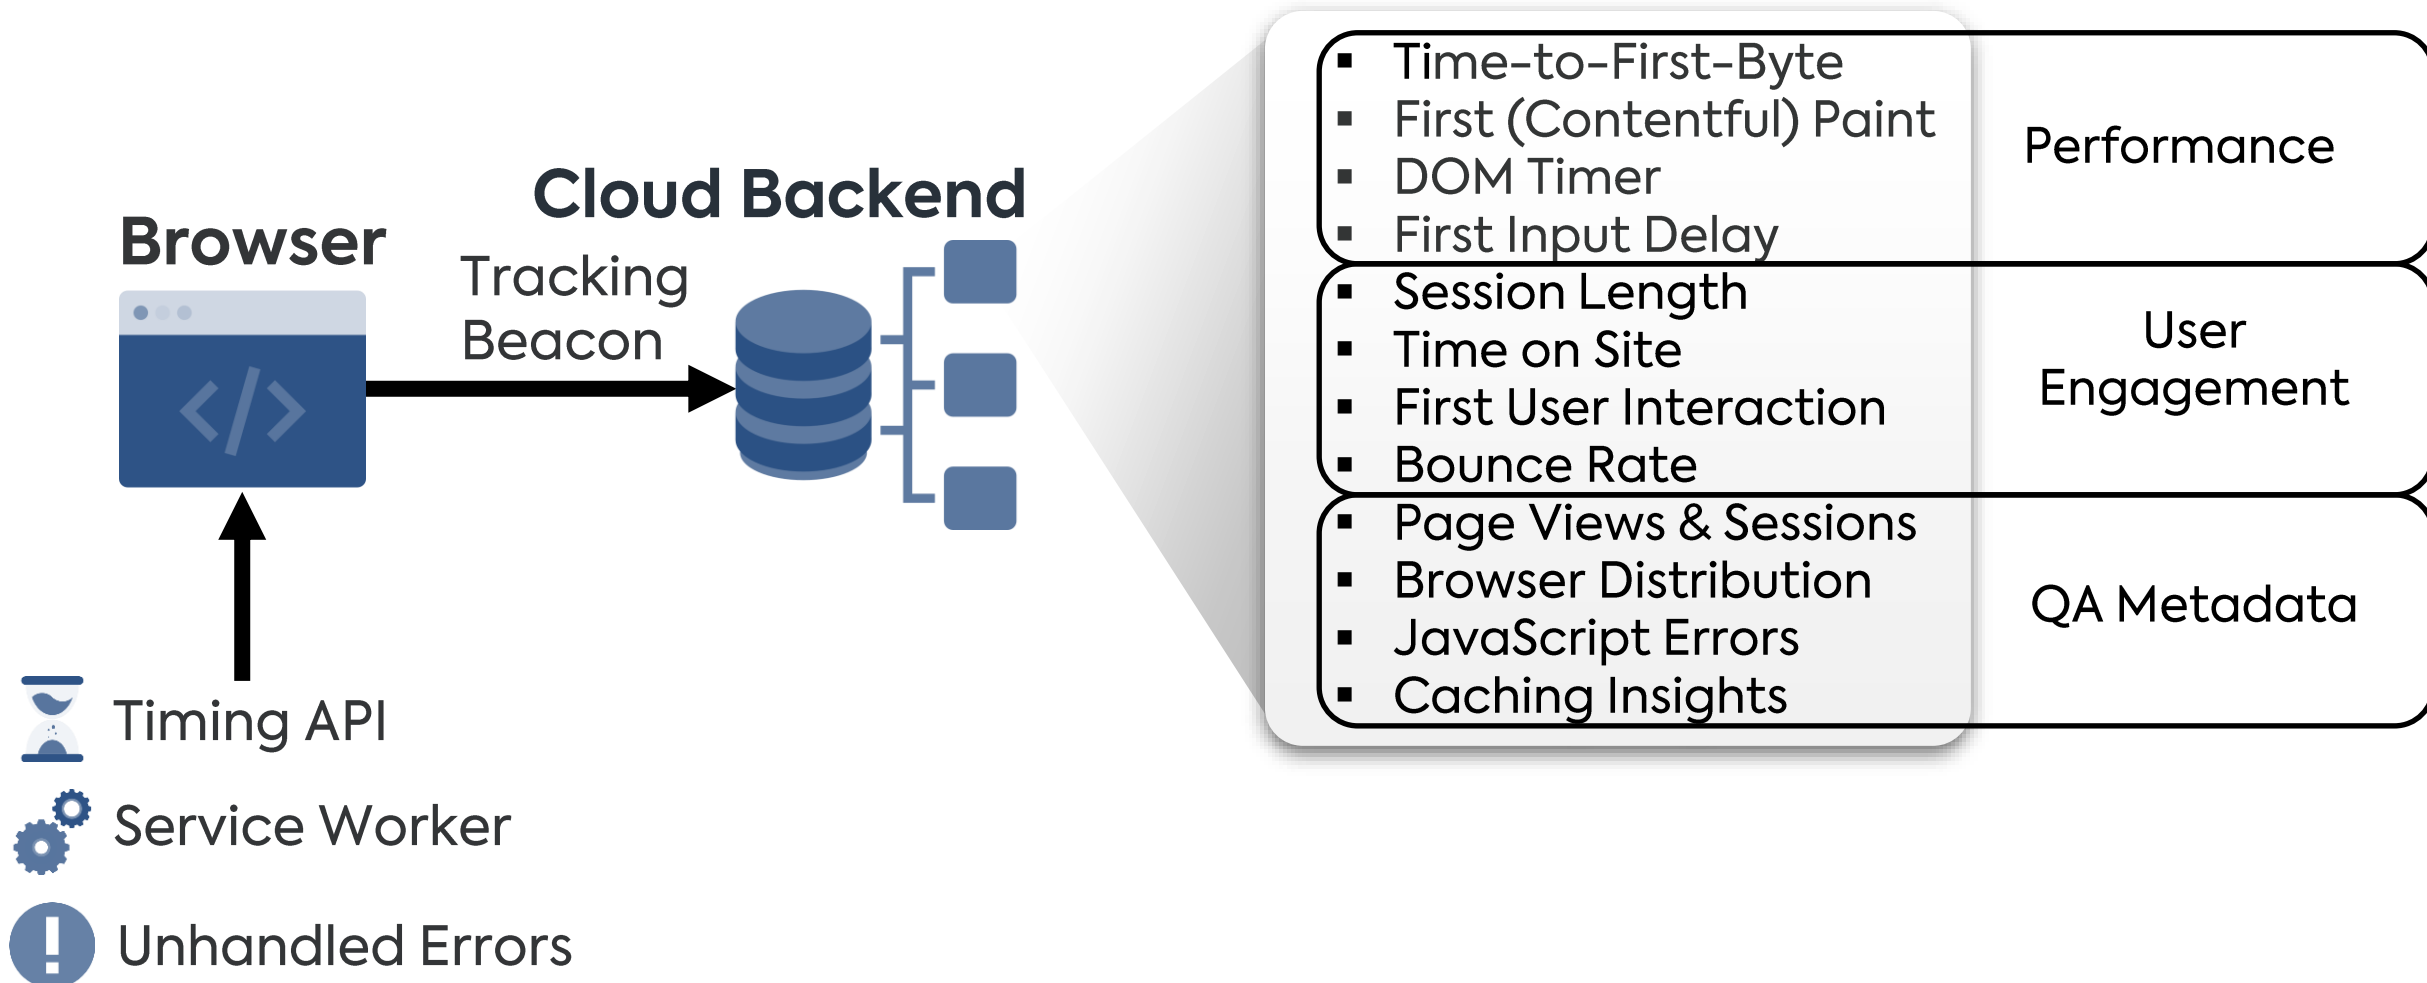


- Speed Kit enabled

- **Measurable uplift:**
 - + Performance
 - + User engagement
 - + ...

- Speed Kit disabled
(no acceleration)

Real-User Monitoring (RUM)



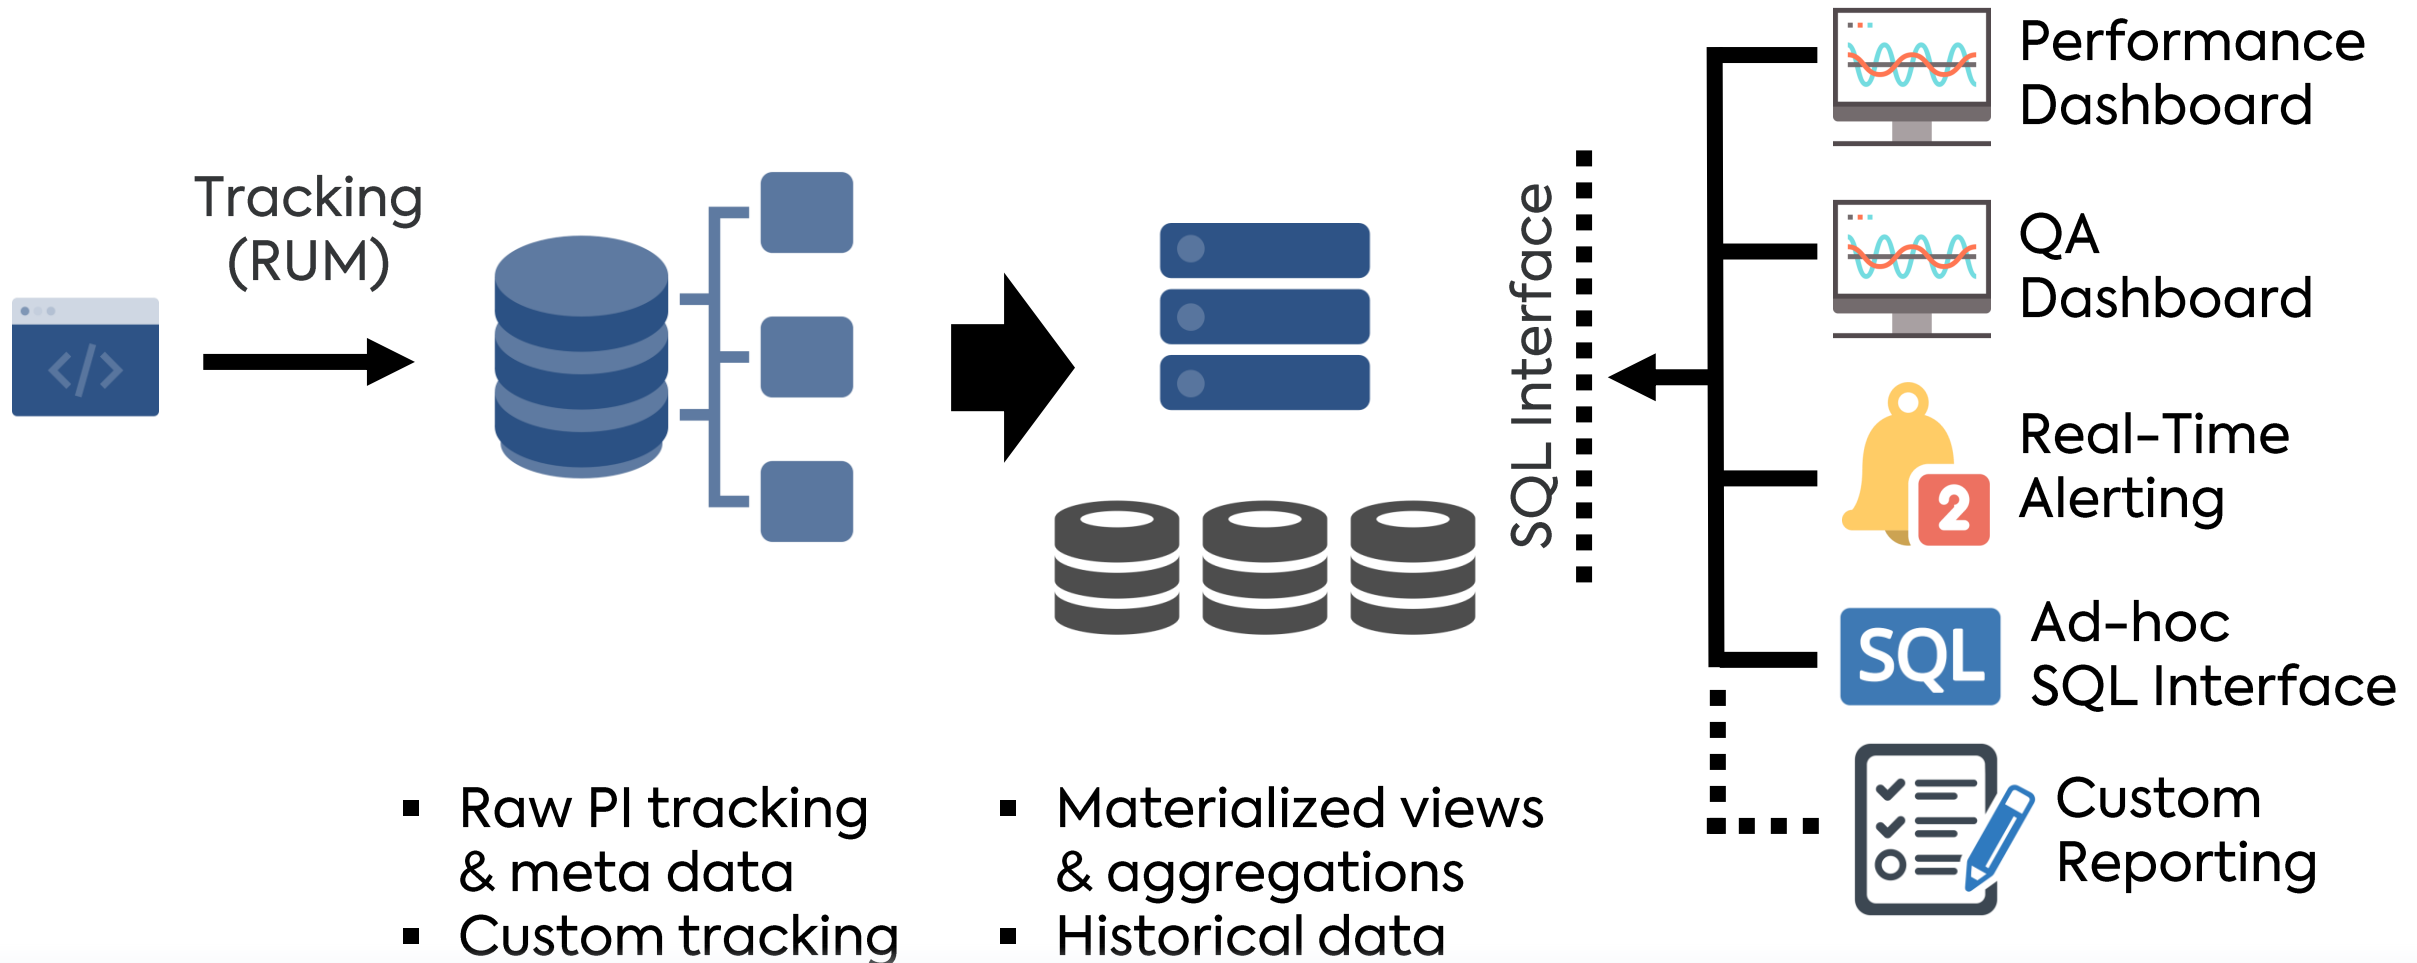
Real-User Monitoring (RUM)

Collection

Ingestion

Analytics

Reporting



Real-User Monitoring (RUM)

Collection



Tracking
(RUM)




Ingestion



- Raw PI tracking & meta data
- Custom tracking


Analytics



- Materialized views & aggregations
- Historical data

Reporting

SQL Interface



Performance Dashboard



QA Dashboard



Real-Time Alerting

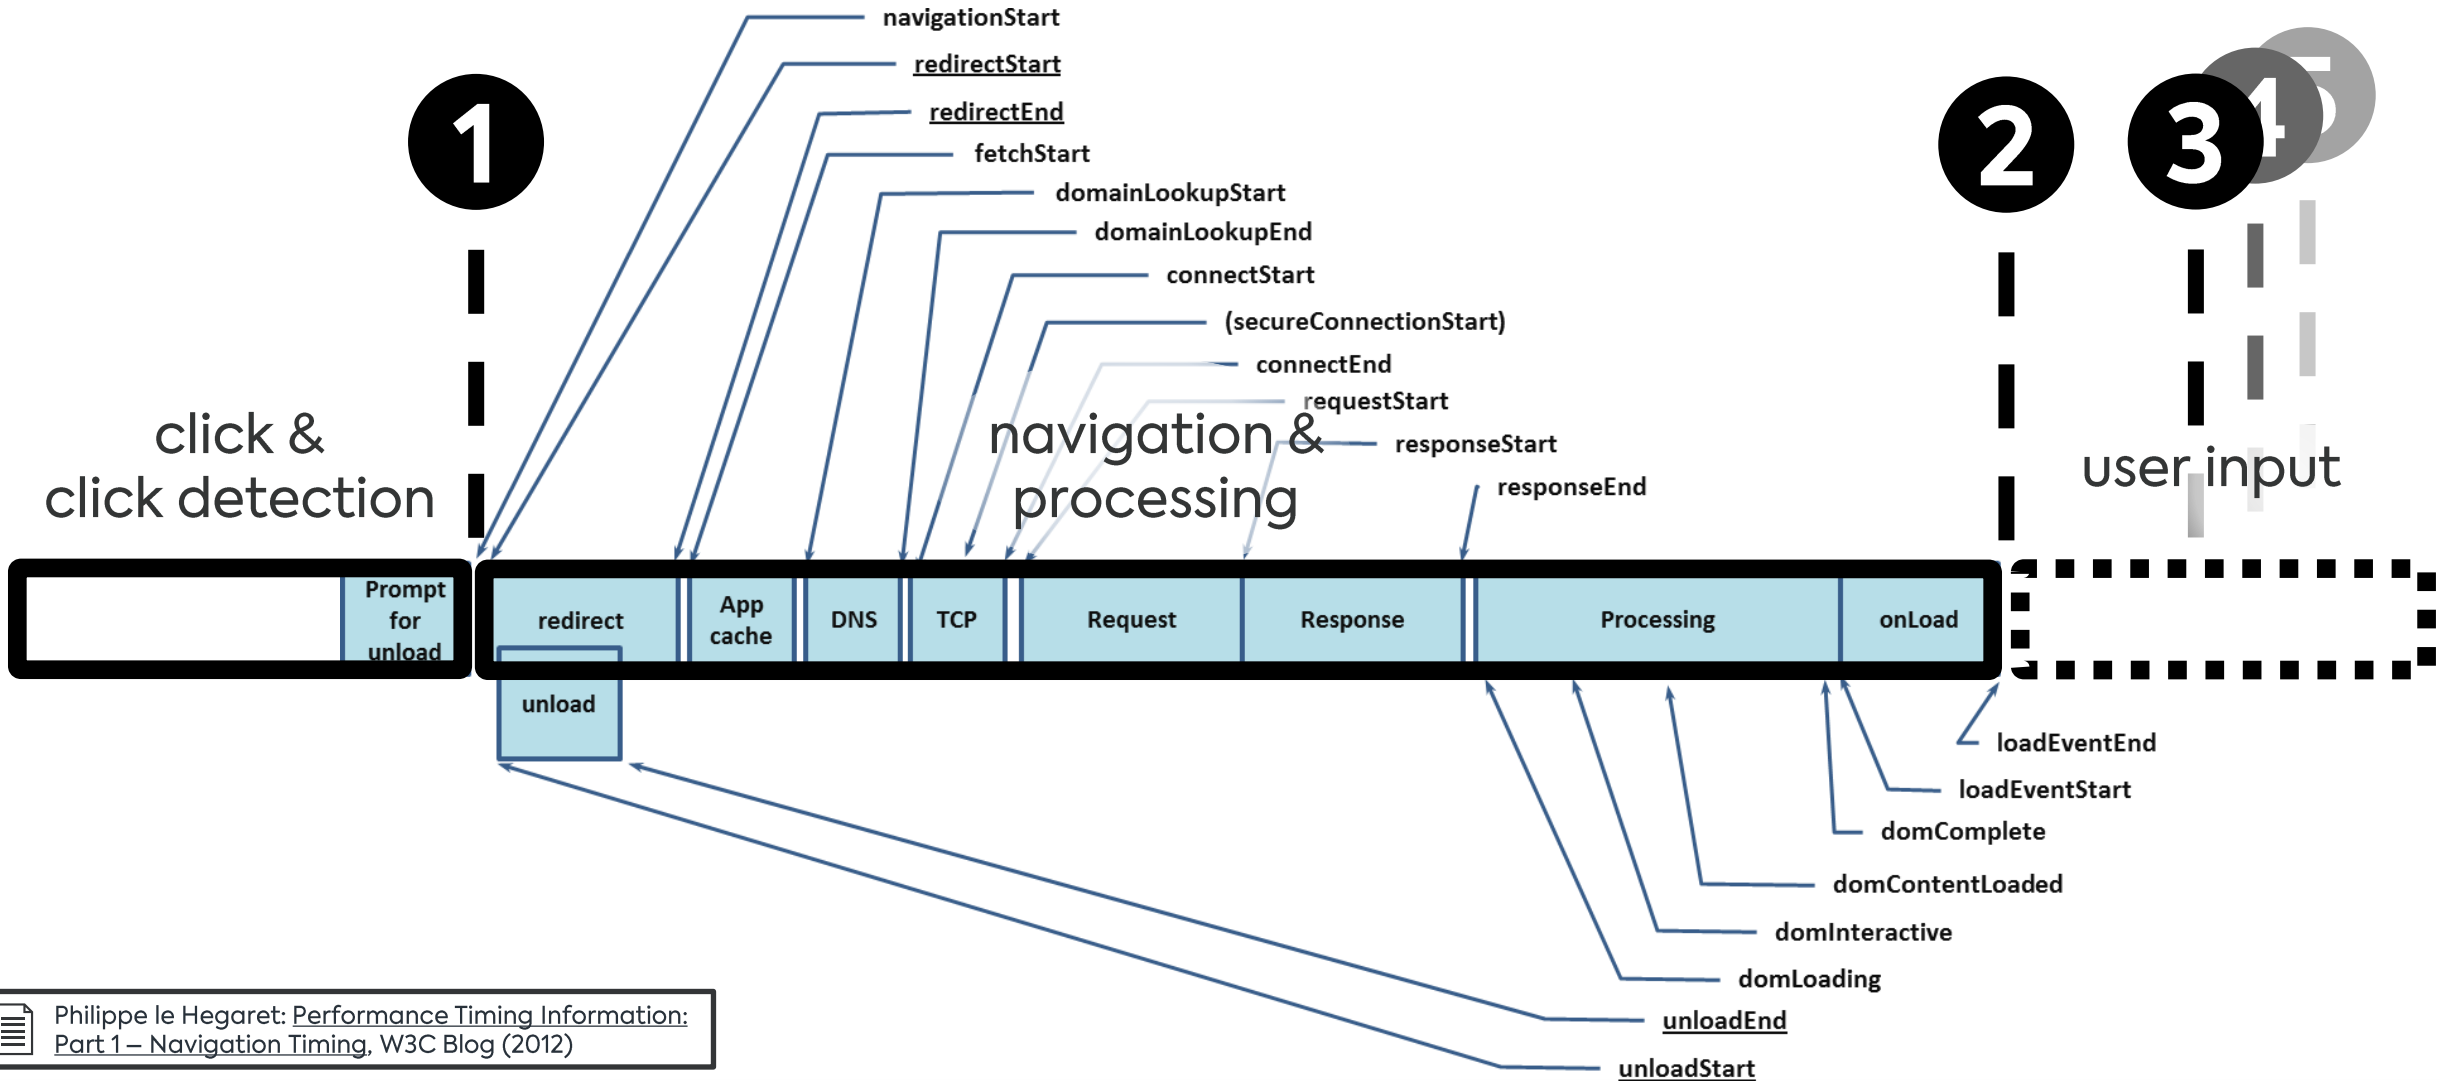


Ad-hoc SQL Interface



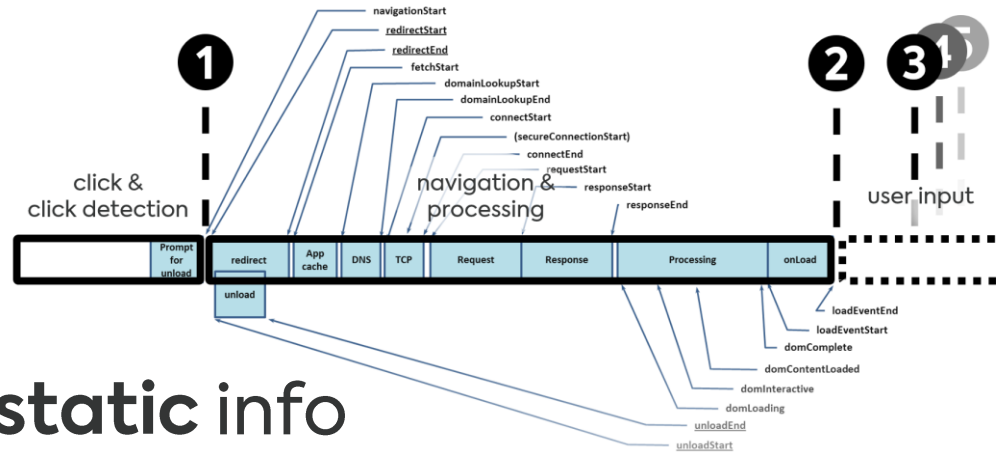
Custom Reporting

When to Send Data Beacons?



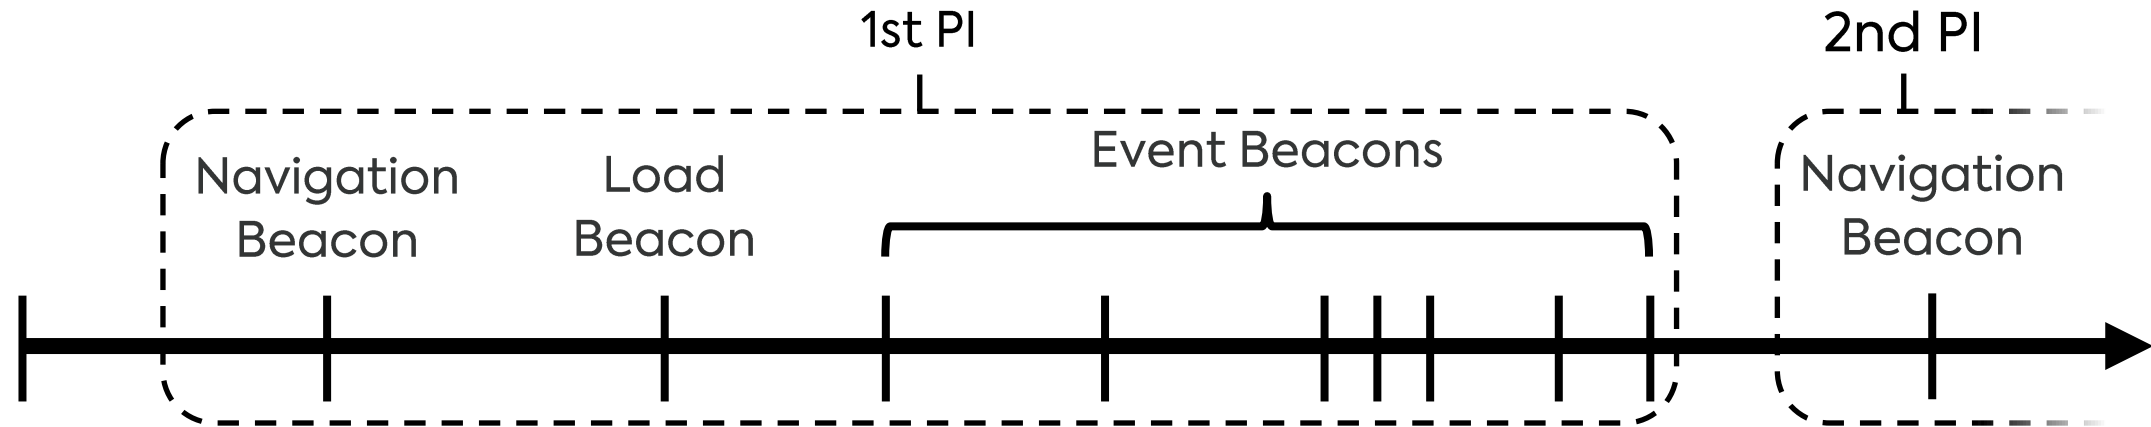
Philippe le Hegaret: Performance Timing Information: Part 1 – Navigation Timing, W3C Blog (2012)

Types of Data Beacons



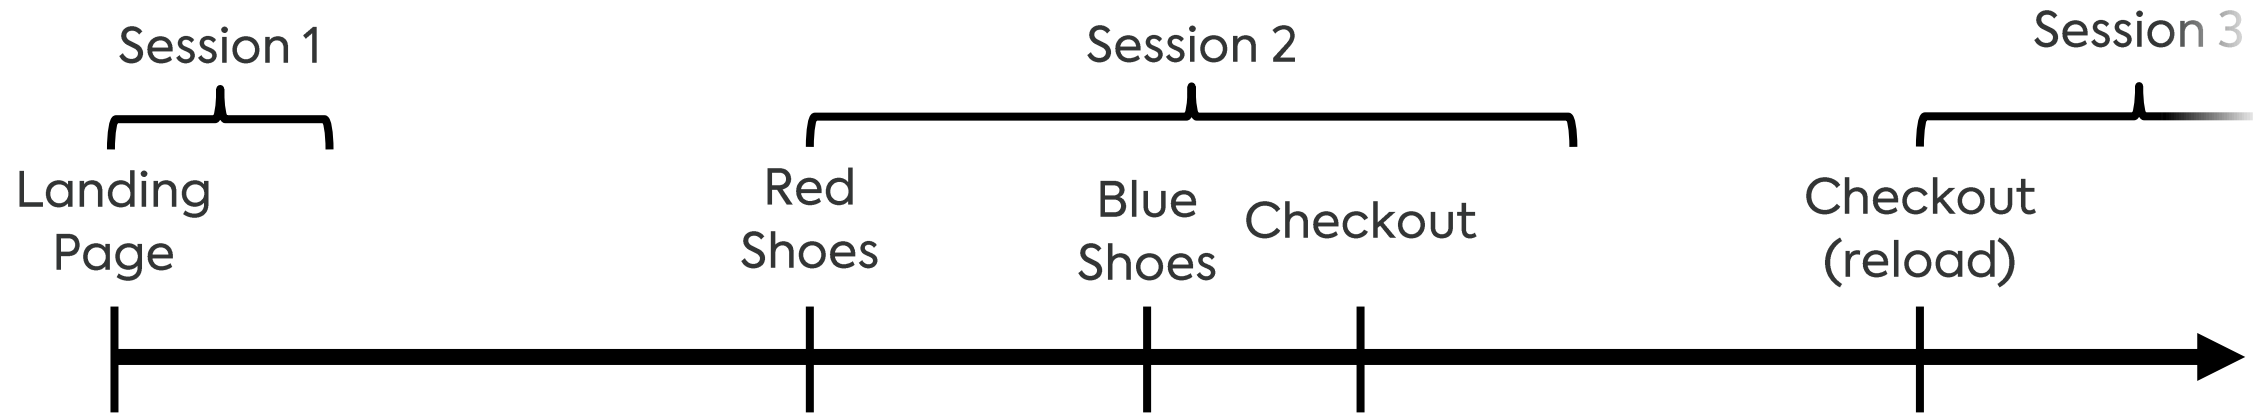
1. **1 for static info**
(URL, user agent, session ID, ...)
2. **1 for timings**
(TTFB, load time, FCP, ...)
3. **0-n for events**
(first input, add-to-cart, ...)

Aggregating Page Impressions



- **Beacon Join → PI:** How do we handle events that come late?
 - Simply wait 5 minutes?
 - Wait for next PI or session timeout?
 - ...?
- How to resolve **user agents**?


Aggregating Session Data





- **Bounces & Session End:** find out when and where people leave
- **Session timeout** after 30 minutes of *inactivity*

Batch Processing & Latency

Alerting



Processing Time




Trend Analysis



- Simple metrics / little context
 - Counters
 - Extreme values
 - Specific errors

- Complex aggregations / huge time windows
 - Average session length
 - Performance by month
 - Seasonal effects

Beaconnect: Continuous Processing

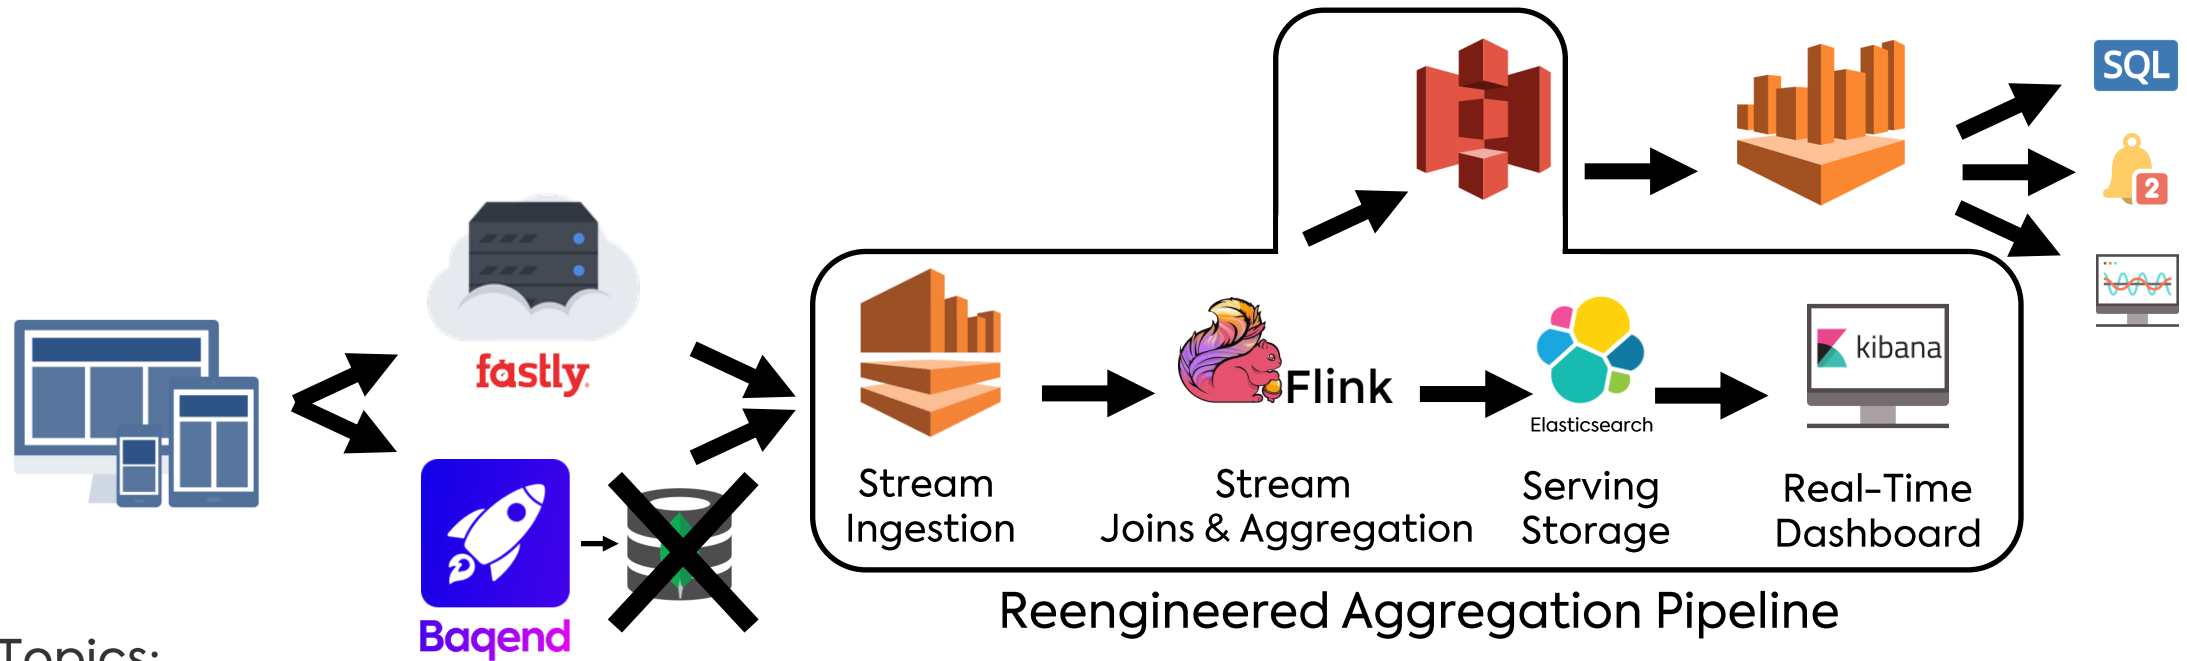


Key Topics:

- ✓ Continuous Aggregation
- ✓ Real-Time Reporting & Analytics
- ✓ Extreme Scalability

 F. Gessert, W. Wingerath. Batching Was Yesterday: Real-Time Tracking & Analysis For 100+ Million Visitors, Flink Forward (2021)

Beaconnect: Continuous Processing

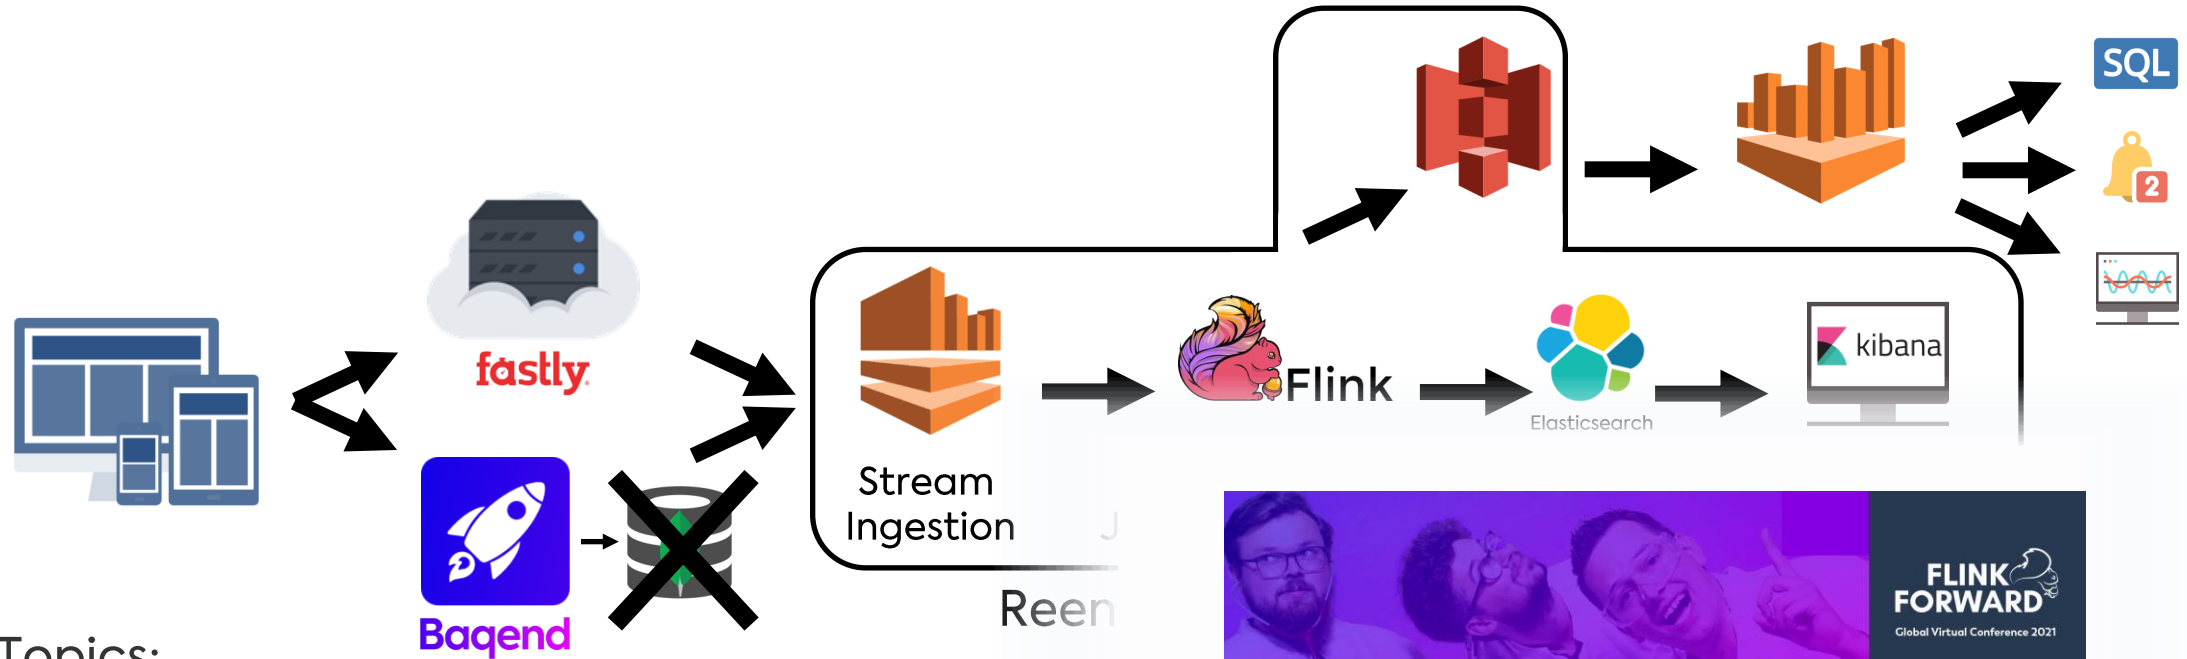


Key Topics:

- ✓ Continuous Aggregation
- ✓ Real-Time Reporting & Analytics
- ✓ Extreme Scalability


 F. Gessert, W. Wingerath. Batching Was Yesterday: Real-Time Tracking & Analysis For 100+ Million Visitors, Flink Forward (2021)

Beaconnect: Continuous Processing




Key Topics:

- ✓ Continuous Aggregation
- ✓ Real-Time Reporting & Analytics
- ✓ Extreme Scalability




Real-Time Dashboarding Example



Traffic* Across 3 Orders of Magnitude

Overall Monthly Traffic

- >100M users
- >200M sessions
- >650M PIs
- >3B data beacons



*Monthly Page Impressions (PIs)

3-Level Aggregation

Partial Page Impressions (PPIs) Enhanced Data Beacons

Time	Browser	Device	Test Group	First Contentful Paint (FCP)
11:05:04.578	Firefox	Mobile	Speed Kit	127ms
11:06:48.139	Chrome	Mobile	Original	958ms


1-Min. Time Windows Immediate Aggregates (Storage)

	Browser	Device	Test Group	First Contentful Paint (FCP)
11:05	Firefox	Mobile	Speed Kit	{200ms: 1, 500ms: 2}
	Firefox	Mobile	Original	{600ms: 2, 800ms: 5}
	Safari	Desktop	Original	{1100ms: 1}
11:06	Firefox	Mobile	Speed Kit	{200ms: 3}
	Chrome	Mobile	Speed Kit	{400ms: 2}
	Opera	Tablet	Original	{700ms: 1, 1300ms: 2}
	Safari	Desktop	Original	{600ms: 4, 900ms}

Arbitrary Time Windows Real-Time Reporting (Dashboard Queries)

11:05	Browser	Device	Test Group	First Contentful Paint (FCP)
–	Firefox	Mobile	Speed Kit	{200ms: 4, 500ms: 2}
11:06				

Dashboarding Query Efficiency



$$\text{Compaction Ratio: } CR = 1 - \frac{|\text{intermediate aggregates}|}{|PIs|}$$

Ongoing Research

- **Prefetching** & Click Prediction
- **Connection-Aware**
 - Compression & Push
 - Image Optimization
- **Vision-based**
 - Image Optimization
 - Regression Testing (QA)
 - Speed Metrics (SI/FMP)
- **Anomaly detection**
on Real-User Data
- Content-based
Staleness Minimization
- Workload-based
TTL Estimation
- Config Optimization
- Dynamic Block Inference
- Dual & Adaptive **Bloom Filter**

Thanks!

Questions?

Material available at:

Contact:

baqend.com/publications
research@baqend.com