

Benchmarking Fast-Data Platforms for the *Aadhaar* Biometric Database

Yogesh Simmhan^(✉), Anshu Shukla, and Arun Verma

Department of Computational and Data Sciences,
Indian Institute of Science, Bangalore, India

simmhan@cds.iisc.ac.in, shukla@ssl.secr.iisc.in, arun.verma100@gmail.com

Abstract. *Aadhaar* is the world’s largest biometric database with a billion records, being compiled as an identity platform to deliver social services to residents of India. *Aadhaar* processes streams of biometric data as residents are enrolled and updated. Besides ~1 million enrollments and updates per day, up to 100 million daily biometric authentications are expected during delivery of various public services. These form critical Big Data applications, with large volumes and high velocity of data. Here, we propose a stream processing workload, based on the *Aadhaar* enrollment and Authentication applications, as a Big Data benchmark for distributed stream processing systems. We describe the application composition, and characterize their task latencies and selectivity, and data rate and size distributions, based on real observations. We also validate this benchmark on Apache Storm using synthetic streams and simulated application logic. This paper offers a unique glimpse into an *operational* national identity infrastructure, and proposes a benchmark for “fast data” platforms to support such eGovernance applications.

1 Introduction

The Unique Identification Authority of India (UIDAI) manages the national identity infrastructure for India, and provides a biometrics-based unique 12-digit identifier for each resident of India, called *Aadhaar* (which means *foundation*, in Sanskrit). *Aadhaar* was conceived as a means to identify the social services and entitlements that each resident is eligible for, and ensures transparent and accountable public service delivery by various government agencies.

The scope of UIDAI is itself narrow. It maintains a database of unique residents in India, with uniqueness guaranteed by their 10 fingerprints and iris scan; assigns a 12 digit *Aadhaar* number to the resident; and as an authentication service, validates if a given biometric matches a given *Aadhaar* number by checking its database. Other authorized government and private agencies can use this UIDAI authentication service to ensure that a specific person requesting service is who they are, and use their *Aadhaar* number as a primary key for determining service entitlements and delivery. However, to guarantee the privacy of residents, UIDAI does not permit a general lookup of the database based on just the biometrics, to locate an arbitrary person of interest.

UIDAI holds the world's largest biometric repository. As of writing, it has enrolled 981 M of the 1,211 M residents of India in its database¹. It continues to voluntarily register pending and newly eligible ones, at an operational cost of about USD 1 per person. It also currently performs up to 2 M authentications *per day* to support a few social services, and this is set to grow to 100 M per day as the use of *Aadhaar* becomes pervasive. The *Aadhaar* database is 5× larger than the next publicly-known biometric repository, the US Homeland Security's OBIM (VISIT) program, which stores 176 M records on fingerprints, and processes 40× fewer transactions at 88.73 M authentications *each year* (or 245 K *per day*), as of latest data available from 2014 [10].

Clearly, the *Aadhaar* repository offers a unique Big Data challenge, both in general and specially from the *public sector*. The present software architecture of UIDAI is based on contemporary open source enterprise solutions that are designed to scale-out on commodity hardware [3]. Specifically, it uses a Staged Event-Driven Architecture (SEDA) [12] that uses a publish-subscribe mechanism to coordinate the execution flow of logical application stages over batches of incoming requests. Within each stage, Big Data technologies *optimized for volume*, such as HBase and Solr, and even traditional relational databases like MySQL, are used. As part of the constant evolution of the UIDAI architecture, one of the goals is to reduce the latency time for enrollment of new residents or updation of their details, that takes between 3–30 days now, to something that can be done interactively. Another is to ensure the scalability of the authentication transactions as the requests grow to 100's of millions per day, and evaluate the applicability of emerging Big Data platforms to achieve the same.

Distributed stream processing systems such as Apache Storm and Apache Spark Streaming have gained traction, off late, in managing the data velocity. Such “fast data” systems offer dataflow or declarative composition semantics and process data at high input rates, with low latency, on distributed commodity clusters. In this context, this paper proposes benchmark workloads, motivated by realistic national-scale eGovernance applications, to evaluate the *quality of service* and the host *efficiency* that are provided by *Fast Data* platforms to process high-velocity data streams.

More generally, this paper highlights the growing importance of eGovernance workloads rather than just enterprise or scientific workloads [2]. As emerging economies like China, India and Brazil with large populations start to digitize their governance platforms and citizen-service delivery, massive online applications can pose unique challenges to Big Data platforms. For e.g., technology challenges with the HealthCare.gov insurance exchange website in the US to support the Affordable Care Act are well known². There are inadequate benchmarks and workloads to help evaluate Big Data platforms for such public sector

¹ India 2011 Census, and live statistics from <https://portal.uidai.gov.in>.

² Healthcare.gov: CMS Has Taken Steps to Address Problems, but Needs to Further Implement Systems Development Best Practices, www.gao.gov/products/GAO-15-238.

applications. This work on characterizing applications for the identity infrastructure of the world’s largest democracy is a step in addressing this gap.

The rest of the paper is organized as follows: in Sect. 2, we provide context for the UIDAI data processing workloads; in Sects. 3 and 4, we describe the composition of the *Aadhaar* enrollment and authentication dataflows, respectively, including characteristics of their input stream data sizes and rates; in Sect. 5, we evaluate this benchmark for Apache Storm using synthetic streams and tasks based on the real distributions; in Sect. 6, we review related work on Big Data benchmarks; in Sect. 7, we discuss how this benchmark can be expanded and generalized, and its relevance on the field; and finally offer our conclusions in Sect. 8.

2 Background

UIDAI aims to provide a standard, verifiable, non-repudiable identity for residents of India. This biometric-based identity distinguishes itself from other traditional forms of physical and digital identities in several ways. It is *unique, universal and non-repudiable*, which cannot be said for passports and driving licenses (not universal) or birth certificates and utility bills (not guaranteed to be authentic and unique). *Aadhaar* guarantees uniqueness of the individual by using their 10 fingerprints and iris scans of both eyes. It is also *electronically verifiable*, since it is just a number and the associated biometric of a person – it is not a “physical” identity card that can be lost, stolen or duplicated. Further, unlike digital identities like OAuth or OpenID, it also proves *physical presence* since the biometric has to be provided by the individual for authentication. Lastly, the fact that it has covered over 80% of eligible residents of India (5 years and older) makes it as close to universal as possible for a voluntary national identity program at this scale.

UIDAI offers two categories of services, one to enroll residents’ demographic and biometric data into *Aadhaar*, and maintain them up to date; and another to authenticate users who provide their *Aadhaar* number and a biometric. We briefly offer context here, and in the next two sections, we drill down into the workloads themselves. Additional background can be found elsewhere [3].

2.1 Enrollment and Update

Residents who have not enrolled into *Aadhaar* register themselves with an authorized enrollment agency, and provide their *demographic details* (i.e., Name, Date of Birth, Address, EMail, Mobile Number) and their *biometrics* (photo, 10 fingerprints and both iris scans). This data is captured offline (presently), encrypted, digitally signed by the agency, and the *encrypted* enrollment packet for each resident uploaded daily to the UIDAI servers. As part of the enrollment pipeline, the resident’s data is validated using basic sanity checks and, importantly, their biometrics is compared against every other resident’s biometrics present in the database to ensure duplicate registrations are not performed. This de-duplication

provides the authoritative uniqueness to *Aadhaar*. Once these checks pass, a random and unique 12 digit number is assigned and mailed to the resident. The packets remain encrypted right after data collection, stay encrypted on the network and on disk, and are decrypted in-memory, just in time, during data insertions, queries or comparisons.

The update process is similar, and allows residents who already have an *Aadhaar* number to change their transient demographic details, such as address and phone number, or even update their biometrics that could degrade with age.

2.2 Authentication and KYC

The basic *Authentication service* verifies if a given *Aadhaar* number matches a biometric or a demographic information provided by a resident, and returns a True/False response. This allows an authorized agency to request verification on whether the individual providing the biometric indeed matches the *Aadhaar* number they have on record, or if the demographic data that has been provided matches the verified one stored as part of the *Aadhaar* enrollment.

The service accepts the *Aadhaar* number and a combination of the following to be provided – fingerprint, iris scan, demographic (gender, age, etc.), One Time Personal Identification Number (OTP) – in an encrypted and signed manner. It then lookups up that *Aadhaar* number, matches the given biometrics and/or demographics against the details stored for that specific *Aadhaar* record, and returns a boolean response. The optional use of an OTP ensures that an authentication is done only with the consent of the individual who requests an OTP from UIDAI and provides it to the agency they are authorizing to perform their authentication.

Another *Know Your Client (KYC)* service extends this authentication to also allow the retrieval of demographic data and photo (but not fingerprint or iris scans) by an authorized agency, upon informed consent by the resident. This eases the burden of proof that residents have to provide in order to sign up for public or banking services, and promotes financial inclusion.

3 Enrollment Workload

In this section, we describe the proposed fast-data benchmark based on the *Aadhaar* enrollment application, followed by the authentication application in the next section. As part of the workload, we characterize the dataflow compositions used to process streaming data, the event rate and size distributions that they process, and the required quality of service in terms of end-to-end processing latency.

3.1 Enrollment Dataflow

A high level streaming dataflow composition for the Enrollment Workload is shown in Fig. 1. It also shows the *latency* per task to process an input packet,

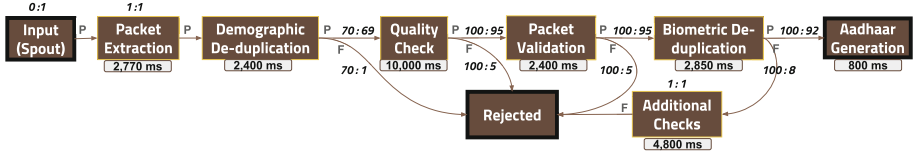


Fig. 1. *Enrollment dataflow.* Tasks are labeled with the average latency time in milliseconds. The selectivity is given for each outgoing edge. “P” edges are taken by events that pass the check at a task, while “F” edges are taken by events that fail a check.

based on observations of the existing application logic, and its *selectivity*, i.e., the ratio between the number of items generated on an output edge for a certain number input items consumed.

The dataflow starts when encrypted enrollment packets (files) are uploaded to UIDAI by the registering agencies. The input packets’ checksums have been verified at a DMZ to detect tampering, duplicate uploads, and malware/virus. The **Input** task emits such validated packets. Next, the **Packet Extraction** task decrypts the packets in-memory and inserts the demographics fields and the photo into a MySQL database and a Solr index, and the entire encrypted packet is archived into a distributed, geographically replicated store.

The **Demographic De-duplication** task locates existing residents whose demographics, such as name, age and pincode (zipcode), are similar to the incoming packet based on a Solr index search. Candidates that fuzzy-match the demographics have their biometrics preemptively checked against the input packet, and matching input packets, estimated at about 2%, are sent to the **Rejected** task. This avoids a full biometric de-duplication *across all residents* for these inputs.

Then, a **Quality Check** task does sanity checks on the demographics and the photo to ensure that the photo appears genuine (e.g., based on gender and age), the names and addresses appear valid, detects language transliteration errors during data entry, etc. An estimated 5% of enrollments are rejected by this task. Inputs that pass this QC arrive at the **Packet Validation** task where its digital signature is verified, and is checked to ensure it was generated by an operator who has been certified, is authorized to operate in that pincode, and matches their assigned supervisor and agency. Typically, 95% of packets pass this check.

Next, the **Biometric De-duplication** task performs a cross-check of each input packet’s iris and fingerprints with every other registered biometric stored in *Aadhaar*. Three independent Automated Biometric Identification Systems (ABIS) are used here. The input biometric is first inserted into all of them, whereupon they extract a biometric template that they index. Then one of the ABIS is chosen, based on a weighting function, to verify if there is a match for this input from among all indexed biometrics. About 8% of inputs typically find a match and are flagged for potential rejection. Actual rejection of packets depends on **Additional Checks** of biometric and demographic data, a small

number of which require manual checks, before being passed onto the **Rejected** task.

Packets that pass the prior tasks successfully are sent to the **Aadhaar-Generation** task which assigns a unique, random 12 digit number to this packet, creates a master record for the number, and forwards it for printing and mailing to the resident. Packets that are rejected in any prior step are retained by the **Rejected** task for auditing in a reject master database.

3.2 Enrollment Data Stream

The transaction rates for the enrollment dataflow change over time, and have to be captured to effectively evaluate the runtime performance of the application. Further, the input data sizes are larger than typical messages to stream processing systems, and hence their variation in sizes needs to be reflected too. Figure 2a shows the probability distribution function (PDF) for the encrypted input enrollment packet size that is uploaded by the enrollment agency and contains the demographics and biometrics. This distribution is from data collected on a single day in September 2015, and is representative of the overall trend. We see that half the packets have a 2–3 MB size, with the overall size ranging from 1–5 MB. We keep the input and output data sizes the same, for the purposes of the workload.

Next, Fig. 2b shows the average hourly input data rates within a recent 24 h period. This has a bi-modal distribution, with peaks in the late morning and late evening. The rate is low early in the day and by 11AM–12PM, ~65,000 packets/h are uploaded in batches by the field agencies after a morning session of registrations. Similarly, enrollment uploads from the afternoon session peaks at 6–7PM. Note that in 2013 when a bulk of the enrollments took place, the enrollment rate was much larger at 1.3M packets per day. The output rate from individual bolts and the entire dataflow is determined by the selectivity.

In future, we expect the enrollment agencies to be constantly connected to the Internet, and the enrollment packet uploaded immediately upon capture. So, based on the latency achieved by the streaming enrollment pipeline, an *Aadhaar*

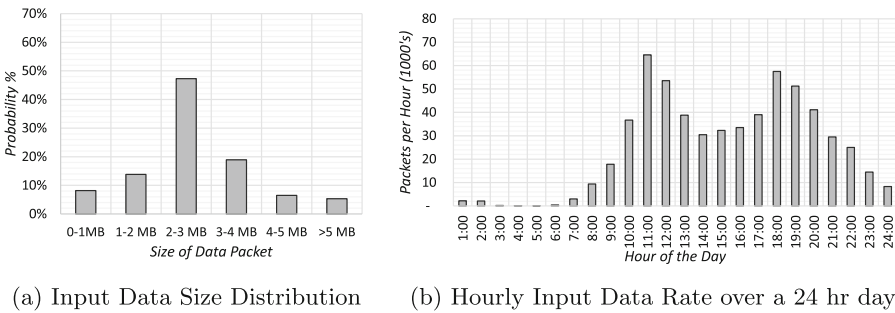


Fig. 2. Probability distribution of the input data stream to the *enrollment dataflow*.

number can be assigned and returned interactively. The input rate distribution is expected to smoothen with such a model.

3.3 Expected Quality of Service

Each stream processing workload has a service level agreement (SLA) defined for it based on the end-to-end latency to process a single request. For the enrollment pipeline, the nominal throughput requirement based on a batch-processing model is to complete processing all packets that are received in a 24-hour period, within that 24-hour period. However, given that the advantage of a stream processing approach is to provide lower latency, we define the quality of service expected for processing each packet to be 10 min. This would allow residents to eventually enroll for *Aadhaar* interactively at a service center and be assigned an ID in the same session.

4 Authentication Workload

4.1 Authentication Dataflow

The authentication dataflow composition is shown in Fig. 3. This dataflow is pre-dominantly a linear pipeline with a selectivity of 1:1. As before, the latency is shown for each task based on an observational snapshot of the dataflow from Sep 2015.

The **Packet Validation** stage operates over HTTPS requests that arrive and is responsible for validating the authenticating agency and user, parsing the XML request and verifying the digital signatures of the request. Subsequently, the **Packet Decryption** task decrypts and extracts the packet’s contents and verifies its integrity. The parameters of the request are then parsed by the **Verify Authorization** task, and a check performed on the type of authentication being done (demographic, biometric) and whether this request is a replay of a previous request. These tasks determine if the request is valid, and is being performed by an authorized entity.

After that, the *Aadhaar* number present in the request is used by the **Query Resident Data** task to lookup and retrieve the resident’s demographic data from the backend HBase storage. The **Biometric and Demographic Match** task performs one or more of the following operations based on the type of authentication requested. It checks if an OTP or a PIN number, if present in the request, is valid. It may verify if the retrieved demographic matches the one in the request, if provided. And if a biometric is passed within the request, it checks if this biometric data matches the one stored for that *Aadhaar* number in the backend ABIS biometric system. These checks determine if the authentication failed or was successful.

As a measure of security, the **Resident Notification** task asynchronously notifies the owner of the *Aadhaar* number using their registered email or mobile number that an authentication was performed. Then an XML response is created with the results of the authentication and digitally signed by the **Create**



Fig. 3. *Authentication dataflow.* Tasks are labeled with the average latency time in milliseconds. Selectivity for all tasks is 1:1.

Response task. Finally, an audit record for the request is created, statistics on the authentication rates and latencies updated for business analytics, and the HTTPS response transmitted to the authenticating agency by the **Audit Log & Send** task.

4.2 Authentication Data Stream

The input stream to the authentication dataflow is presently more uniform than the enrollment dataflow, even as the number of users performing authentication will increase with time. Input requests are about 4 KB in size, and we use this constant size per message for our workload. The hourly input data rate is shown in Fig. 4. This is over $20\times$ faster than the enrollment dataflow, and is also expected to grow in future. The current base rate for authentications is 150 requests/s, with two sharp peaks of about 500 requests/s in the morning and the evening – *Aadhaar* is used by federal employees to clock in and out of office each day, and these peaks reflect the requests by this attendance service³. In future, the number of authentications are expected to rise to 100 M requests during a working day, or an average of about 2,500 requests/s.

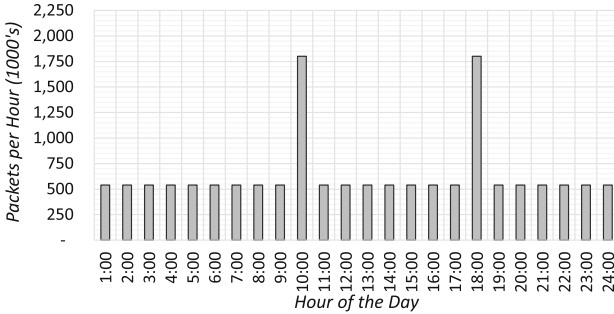


Fig. 4. Hourly input request rate over a 24 h day passed to the *authentication dataflow* with a base rate of 150 requests/s. The two peaks of 500 requests/s during the morning and evening periods reflect requests by the federal attendance service.

³ Biometric Attendance Service, <http://attendance.gov.in>.

4.3 Quality of Service Expected

Authentication requests are inherently used by interactive applications. Hence the latency requirements for processing them are tighter. The SLA for end-to-end latency to process a single authentication request by the dataflow is set to 1000 ms, with most of the requests expected to be completed within 500 ms. In addition, there are network latencies for transmitting the request and response between the data center and the authenticating agency that add to the round trip time, which we do not consider here.

5 Experimental Validation

The motivation for the proposed workload is to evaluate the scalability of stream processing systems that can orchestrate the dataflow for the given inputs data streams, and return results within the specified SLAs. Reducing the computational resources required to achieve these quality of service metrics is an additional goal. We implement the proposed workload⁴ and validate it on the Apache Storm distributed streaming platform.

5.1 Workload Generation

We compose the Enrollment and Authentication pipelines as two *topologies* in Apache Storm. Tasks are defined using a *synthetic bolt* logic that performs in-memory string operations for the given latency duration, and in the process consume CPU resources. The actual workload logic at UIDAI would have performed XML serialization or deserialization, encryption or decryption, remote NoSQL or a relational queries, and so on. Depending on the task, the load on the bolt itself would be limited to string parsing, integer arithmetic (both CPU bound) or waiting for query responses (idle CPU). So, in the absence of access to the actual logic themselves, the string processing performed by the synthetic bolt is a conservative approximation of the expected computational effort taken by each task in the dataflow.

The synthetic bolt is configured for each task to perform its computation for the given *latency* duration for that task. The bolt is also configured with the *selectivity* on each of the output edges for that task, and based on this, a certain fraction of input packets are routed to the downstream bolt(s) on the relevant edges. If the selectivity matches, the bolt logic passes the input packet to the relevant downstream bolt(s) without change.

We have developed an *event stream generator* tool that uses the given input rate and size distributions of the enrollment and authentication packets to pre-generate events with the appropriate relative timestamps and payload sizes. Synthetic events of these sizes are pre-fetched into memory by a *spout* logic within the Storm topology, and replayed at the appropriate relative time intervals to match the required input rate distribution. In particular, the spout uses a

⁴ Code and data generator at <https://github.com/dream-lab/bigdata-benchmarks>.

multi-threaded and distributed mechanism to ensure that the data rate can be maintained at even 10,000's of events/s, if necessary. As we show, the results confirm that the observed input rates and sizes for the experiments match the reference input distribution that is given. The spout can also be configured with different time scaling factors to speed-up or slow-down the rates while maintaining a proportional inter-arrival time between events.

5.2 Deployment

We run the topology on a 24-node commodity cluster, with each node having 8 AMD Opteron 2.6 GHz Cores and 32 GB RAM, connected by GigaBit Ethernet, and running Apache Storm v0.9.4 on OpenJDK v1.7 and CentOS 7. Storm *supervisors* that execute the tasks of the topology run on 23 nodes and 1 node is dedicated to management services such as Master, Zookeeper and Nimbus. There are eight resource *slots* in each supervisor (host), one per CPU core, on which the Storm *workers* can execute one or more task threads.

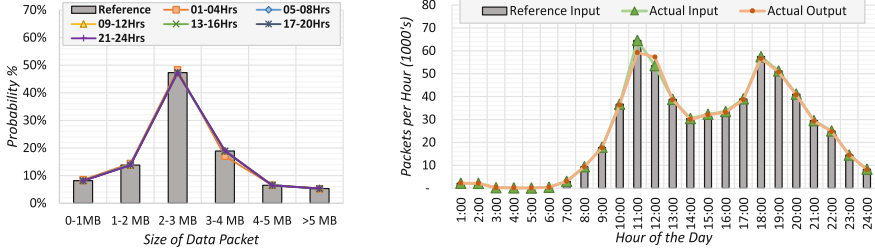
Storm topologies can be configured to use only a subset of the supervisor slots in the cluster and also the degrees of parallelism (number of threads) assigned to each task in the topology. By default, each task runs on a single thread on a single slot. Based on the SLA required for each pipeline and data rate distribution, we determine the *minimum number of slots* that should be allocated to the topology and the *degree of parallelism per task* required to meet the SLA for the benchmarks. These configurations are based on the latency time per task in the dataflow such that there are adequate threads to process data arriving at the *peak rate* for that workload, and ensure that there are adequate CPU cores to sustain the compute requirements of these threads, yet without punitive context-switching overheads.

For the enrollment topology, the total degree of parallelism required to meet these latencies and peak rate comes out to 475 threads, shared uniformly across all the tasks, and running on 9 nodes (72 cores) of the cluster. For the authentication topology, we arrive at a total degree of parallelism of 514 threads, distributed proportionally across all tasks based on their contribution to the overall latency of the topology, and using 19 nodes (152 cores). The spouts that generate the input streams in parallel for the enrollment and authentication dataflows are included within this count and they take up 1 thread and 10 threads, respectively, while the terminal sink tasks in the topologies take up 1 thread each.

5.3 Results

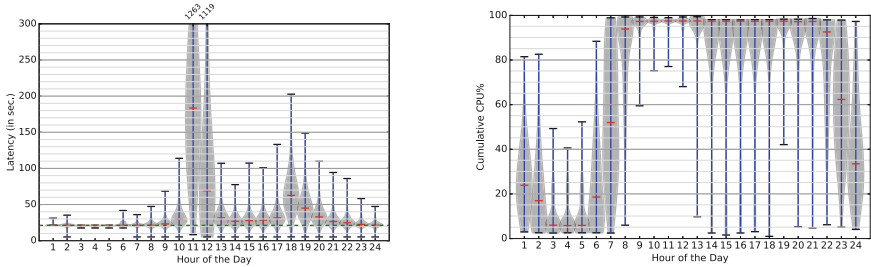
Enrollment Workload. We perform a 24-hour benchmark run where data streams that follow the given rate and size distributions are generated and passed as input to the Enrollment topology in Storm for a whole day. Overall, about 591,270 input packets were generated during this period.

Figures 5a and b show characteristics of the expected and actual input stream. We are able to validate that the event generator can generate input events with the same size distributions as the reference size distribution (Fig. 5a,



(a) Expected and actual input data sizes. Observed rates are grouped into 4-hour periods. We see that all periods have the same distribution. (b) Expected hourly input rate (bar), and the actual hourly input and output rates (lines) over a 24 hr period.

Fig. 5. Probability distribution of the expected and observed input data stream to the *enrollment workload* experiment. (Color figure online)



(a) Violin plot of latency distribution. The minimum theoretical latency for a successful enrollment is 21.22 secs, shown as a dashed green horizontal line. (b) Violin plot of CPU utilization% sampled every second for 9 nodes, for a $1/10^{th}$ time duration experiment that is scaled back to 24 hours for plotting.

Fig. 6. Violin plot on hourly distribution of *latency* per packet (s) and *CPU utilization* for the *enrollment workload* experiment.

gray bars), not just cumulatively for the entire 24-hour period but also for individual 4-hour periods (Fig. 5a, lines). Likewise, we see from Fig. 5b that the hourly input rates generated by the spouts (green line/triangle) match the reference input rate distribution (gray bars). The figure also shows that the output event rate (orange line/circle) from the topology closely matches the input rates (green line/triangle). The output rate falls behind by $\sim 5,000$ packets/h at 11 AM, where the input rate peaks, but it is able to compensate for this and catch up within the next two-hour period.

Drilling in further into the *event latencies*, Fig. 6a shows a violin plot⁵ during each hour for the end-to-end latency for every event packet in that hour.

⁵ The *violin plot* is a generalization of a box and whiskers plot. The minimum, median and maximum values are marked with a dash on the vertical line. The width of the horizontal shaded region around each vertical bar represents the relative frequency of packets having that latency value.

Here, we note that the average hourly *median* latency per enrollment packet is 35.74s, relative to the theoretical lower bound makespan of 22.22s per event that excludes the network time and other overheads. Except for the 11 AM and 12 PM hours, when we have a peak input rate and the system is catching up, the *maximum* observed latency time for packets during all other hours is less than 202s per packet – well below the 600s SLA.

For the 11 AM and 12 PM hours, the median latency is 183s and 69s, respectively, with a peak at 1,263s. The evening input rate spike at 6 PM does not have such a strong impact on the latencies, and they go up only marginally to a median value of 62s and a peak of 202s. In all, 17,700 events out of a total of 591,270 events (i.e., 2.99%) had an SLA violation, where the end-to-end latency per event was greater than 600s. In fact, *for 97% of the inputs, we are able to complete the enrollment pipeline within just 5 min.*

In order to evaluate the *resource efficiency* of Apache Storm in processing the workload, we sample the CPU utilization every second on each of the 9 active nodes for this topology. To keep the logs manageable, this is done for a shorter experiment that ran for $1/10^{th}$ the whole-day duration (i.e., 2h and 24min), while keeping the same size and rate distributions scaled down in time. Figure 6b shows an hourly violin plot of the CPU utilization% sampled every second per node, with each bar having $(9 \times 60 \times 60 \times \frac{1}{10})$ samples. As the pipeline warms up in the initial hour, we see that the CPU utilization is at about 20%, and beyond that, we see a correlation of the utilization with the input rate. Once we hit the morning peak corresponding to the 11 AM period, the CPU utilization is close to 100% for all nodes in the cluster during the entire hour, exhibiting high efficiency. The median utilization remains close to 100% for 13 straight periods, until the effects of the morning and evening peaks wear off by the 10 PM period. The average CPU utilization across all nodes for the entire period was high at 70%, despite the input rate variation.

Authentication Workload. The event generator for the authentication workload was also run for a 24-hour period and about 15,480,000 events are passed to the authentication topology during that time. Figure 5 shows that the reference and actual hourly input rates for the topology. The observed input rates are able to match precisely with the much higher, albeit smaller in size, expected input rates for the authentication workload compared to the enrollment workload. Since one output event is generated by each input event passed to the topology, the figure also shows that the output rate is able to sustain and match the input rate at the hourly granularity, even during the AM and PM spikes.

We examine the end-to-end event *latency* distributions for every hour using the violin plot shown in Fig. 8a. We see that except for the 10 AM and 6 PM hours that correspond to the input rate spikes, the bulk of the latencies for all other hours are tightly bounded, and fall predominantly in the range of 350–500ms. This is in relation to the theoretical lower bound latency of 250ms that does not take network costs into account. During the two spikes, the latency both increases and its distribution widens, with the most of the events falling in

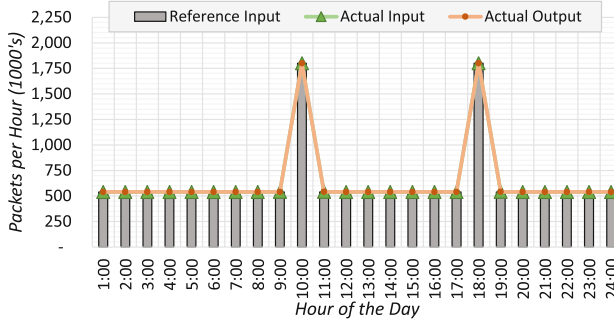
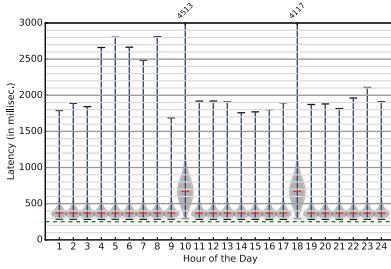
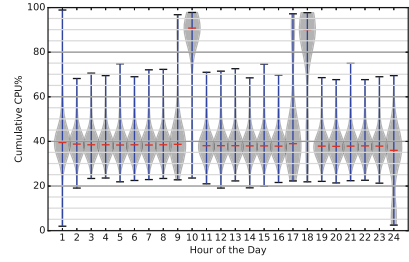


Fig. 7. Hourly reference and actual input data rates passed to the *authentication dataflow* during a 24h experiment run, and the actual hourly output rates from the dataflow.



(a) Violin plot of latency distribution. The minimum latency for a successful authentication is 250 ms, shown as a dashed horizontal line.



(b) Violin plot of CPU utilization sampled every second for 19 nodes, for a $1/10^{th}$ time duration experiment that is scaled back to 24 hours for plotting.

Fig. 8. Violin plot of hourly distribution of *latency* per request (ms) and *CPU utilization* for the *authentication workload* experiment.

the latency range of 550–1,000 ms, but within the SLA of 1 s. While there are outliers in all hours that sometimes reach 4,500 ms, overall, we see that *the SLA is met for 99.98% of the input events*, and violated for only 2,430 events out of over 15 M.

The violin plot of the *CPU utilization* across 19 nodes is shown in Fig. 8b, for a shorter experiment that ran for $1/10^{th}$ the period. Unlike the enrollment workload, we see that the median utilization stays low at about 40% for most of the hours except for the 10 AM and 6 PM input rate spikes when the median goes up to 93%. These utilization values however drop back in the next hour, indicating that the computation load stabilizes quickly after the input rate drops back to the base load of 150 requests/s. The average utilization across all hours and nodes is 43% (Fig. 7).

Relative to the enrollment workload, a higher fraction of input events fall well within their SLA limits for the authentication workload but we also observe a

lower CPU utilization. This indicates that we can potentially trade-off resource efficiency and latency in the case of the authentication dataflow. The resource allocation could potentially be reduced for better efficiency, though the latency distribution for the two peaking hours may come closer to or go beyond the SLA limits. Similarly, the enrollment workload could improve its SLA beyond 97% by increasing the resource allocation, but consequently have a resource utilization lower than 70%. This also motivates the need for elastic resource allocation over time.

6 Related Work

Benchmarking distributed stream processing systems (DSPS) using real world workloads can help identify which of their features are impacted by unique input data characteristics, composition capabilities, and latency requirements. Several benchmarks have been proposed in this context.

The *Linear Road Benchmark* [1] simulates a highway toll system for motor vehicles with variable tolling, and was meant to compare Data Stream Management Systems (DSMS) with Database Management Systems (DBMS). The input to the benchmark is from a traffic model that has variable number of vehicles, but each emitting tuples at a uniform rate and with the same type. So while the input rate is variable, each message is of a fixed, small size. The metrics for evaluation are the response time and accuracy of the query, and the maximum sustained rate within a specified response time, but it does not consider resource utilization. While not developed for distributed stream processing systems, which deal with opaque messages and user logic, it can be adopted to validate DSPS. Our proposed benchmark is particularly tuned for DSPS and simulates the behavior of a real-world eGovernance workload, with variable message sizes and resource efficiency as an additional metric.

StreamBench [8] proposes 7 micro-benchmarks on 4 different synthetic workload suites generated from real time web logs and network traffic. Different workload suites are classified keeping performance metrics in mind. Performance workload and Multi-recipient performance workload suites measure the throughput and latency by pushing up the input rate, using single and multi receiver respectively; Fault tolerance workload suite measures the throughput penalty factor by deliberately causing nodes to fail; and Durability workload suite measures the fraction of time for which the framework is available when running the experiments for long durations. While we do not consider durability or fault-tolerance directly, our benchmarks are run for 24-hour, and this can be extended to longer periods to test durability. In addition to message rate variability, we also consider larger message sizes and variations in message sizes that can impact the resource utilization, which we additionally consider as a metric for comparison between DSPS platforms.

The *IBM Streams benchmark* [9] uses a email spam detection application over the Enron email dataset, and does a relative comparison of IBM's Infosphere Streams and Apache Storm. While they are able to reproduce the data sizes

(KBs) and rates from real emails from over a decade ago, their dataflow itself does not capture any application logic or even sleep times. So this benchmark is more a measure of the network overheads when running the dataflow, rather than computational resources used by, or the consequent latency of, the dataflow. Hence it has limited value.

While such stream processing benchmarks are useful, they do not consider message payloads that are large (MBs) in size, provide strict latency SLA requirement based on real applications, or capture characteristics of eGovernance services. Our *Aadhaar* benchmarks complements these efforts.

Chronos [5] is a generic data generation framework for stream benchmarks. Its focus is not to actually benchmark or stream the input data, but to generate data that can then be streamed to perform benchmarking. Chronos can create large scale synthetic time-series input data that mimics the distributions and correlations that it identifies from a given sample data. It uses Latent Dirichlet Allocation (LDA) method for extracting temporal dependencies from the sample data and preserves correlation among columns. This avoids users having to mine the data for complex patterns and distributions themselves for generating benchmark input specifications. This complements our work, where the distributions are explicitly given based on historic observations, and the focus is on benchmarking streams with different size and rate characteristics over the given dataflow definitions.

Other comprehensive benchmarks have developed for Big Data processing, beyond just fast-data processing which is our emphasis. *Hibench* [6] includes 10 micro workloads covering SQL Queries, Machine learning, Graph Computation and micro-benchmarks like word count and sort. *SparkBench* [7] comprises of different application types like Graph computation, SQL queries along with streaming application for Spark Streaming. This is a mixed-workload for different dimensions of Big Data. It uses resource consumption and data processing rate as the metrics for evaluation, but processing latency is not evaluated.

Big Bench [4] is modeled on the TPC-DS benchmark [11] and uses a retail industry scenario. It offers both semi-structured and unstructured data for data variety; analytics queries over click logs and user reviews, contributing to data volume; and an extract-transform-load (ETL) pipeline for data velocity. This is evolving as a community Big Data benchmark and has been validated on Teradata Aster and Hive platforms, among others. Variations to TPC-DS have also been proposed [13]. Similar to these Enterprise Big Data workload suites, one can envision an eGovernance Big Data workload suite, with this paper on the velocity dimension being a starting point.

7 Discussion

7.1 Extending the Workloads

The fast-data workloads that we have proposed are based on initial observations on the structure and latencies of the two *Aadhaar* dataflows, and the distributions of their input packet rates and sizes. There are several opportunities to

expand upon these, both to increase the value derived from the workloads and to evaluate other Big Data platforms that support these applications.

The tasks in the two dataflows themselves are diverse in terms of the resources that they consume. Our use of a string processing synthetic task is a conservative proxy for the actual tasks. We can further *categorize the tasks in the dataflow based on their resource intensity* into CPU intensive, I/O intensive, memory intensive or idling tasks – the latter of which runs a remote query but does not consume much resources on the task’s local host. Based on this, we can have different categories of simulated tasks be used in the workload which will help to model their resource consumption better. This will offer more genuine estimates of the degrees of parallelism required for each task, service level agreements that can be met, and the resource efficiency possible.

Occasionally, new and better biometric identification algorithms may be considered to improve the robustness of authentication. Then, all prior biometric data collected during enrollment may be *reprocessed* to evaluate the new algorithms. This reprocessing can use a variation of the enrollment pipeline to operate in a bulk mode, and compare the quality of the old algorithm with the new. There are also one-off scenarios when *bulk authentications* are performed for demographic verification by specific agencies, often during off-peak hours at night. Such high-throughput variations to the proposed low-latency workloads are worth considering.

The data rate distributions themselves are bound to change over time, and with *Aadhaar*’s use of multiple regional data centers, the load on the backend may have *region specific trends* that can be captured. It may also be possible to identify rate distributions specific to different authenticating sources, and generate synthetic *cumulative distributions* that blend them in different phases for representative, multi-modal, longer-term simulation runs. These can also help identify extreme spikes that may be possible when different distributions coincidentally peak at the same time, and help understand the behavior of the system during exigencies.

The trade-off that we observe between resource efficiency and SLA violations also provides the opportunity for developing and evaluating *resource allocation strategies* for fast-data platforms. The rate variation combined with the different probabilities for paths taken through the dataflow means that the load on each task is not uniform across time. Thus, these workloads can be used to benchmark the agility of fast-data platforms to intelligently acquire and release elastic Cloud resources to achieve the SLA while also reducing the (actual or notional) monetary cost for using virtualized resources.

While the stream processing dataflows coordinate the execution of the business logic, there are other high *data volume* workloads that are performed on the backend platforms that actually host the *Aadhaar* data and respond to queries from the streaming pipelines. These span NoSQL and relational databases that host demographic data, biometric databases that index and query over fingerprint and iris data, distributed file-systems that archive raw enrollment data and results, and audit logs from billions of authentications that are useful for mining.

Each of these are a Big Data platform case study in itself, and deserve attention as part of a eGovernance benchmark that cuts across Big Data dimensions.

7.2 Practical Considerations

The proposed *Aadhaar* benchmarks are validated using the Apache Storm distributed stream processing system. Other platforms such as Apache Spark Streaming and InfoSphere Streams could similarly be validated. Our results from evaluating Storm show that such fast-data platforms can achieve the SLAs that are required for the enrollment and authentication dataflows, and have the potential to significantly improve the quality of service for the end user.

However, such a validation of the orchestration platform is just one piece of a complex architecture, and cannot be construed to making an immediate operational impact within *Aadhaar*. The current SEDA model, which a distributed stream processing system could conceivably replace, is one of many Big Data platforms that work together to sustain the operations within UIDAI. Switching from a batch to a streaming architecture can have far-reaching impact, both positive and negative (e.g., on robustness, throughput), that needs to be understood, and the consequent architectural changes to other parts of the software stack validated. Making any design change in a complex, operational system supporting a billion residents each day is not trivial.

There are also logistical considerations since the overall application workflow relies on agents on the field and technology availability outside the data center. Some of the tasks in the dataflow, such as **Quality Check** in the enrollment dataflow, has human agents in the loop who complement automated quality checks. These have to be modeled better to guarantee the notional SLAs we observe, or there should be sufficient confidence in complete automation. Mobile field offices that collect enrollment data do not have Internet connectivity round-the-clock, since they may be at remote villages, and currently upload the data in batches each day. Moving to an interactive enrollment process necessitates constant Internet connectivity, which may be possible in the near future but requires additional resources and planning.

So, in summary, our proposed benchmarks are able to quantitatively verify the intuition that stream processing platforms will offer better SLAs than the current batch design, and offer a meaningful, real-world workload to verify fast-data platforms. But the caveats mentioned above in operationalizing such an architecture limit their immediate impact within *Aadhaar*.

8 Conclusions

In this article, we have proposed a Big Data benchmark for high velocity, *fast-data* applications based on an *eGovernance* workload. This addresses a gap in existing benchmarks that are based on web or enterprise applications, and are often volume-oriented.

Some characteristics of the workload stand out. The enrollment dataflow's input messages have a larger size, atypical of event streams for fast-data platforms that tend to be in the order of KBs in size. The bi-modal distribution of data streams for both the workloads is seen in event streams from other domains too, and is consistent with human activity patterns that are intrinsic to eGovernance platforms. The enrollment dataflow also has control-flow built-in, when packets fail validation and take a different path, and this impacts the probability with which different paths are taken. This is captured by the outgoing edge's selectivity. Both of these variations impact resource utilization, and motivate the need for elastic resource provisioning. Better resource allocation models for determining the degree of parallelism per task will also be valuable, and help meet the SLAs while conserving resources.

The *Aadhaar* workload is unique in its scale, being the largest of its kind in the world, but the characteristics of this workload can be seen in other public sector services such as the Department of Motor Vehicles or the Passport Office issuing IDs, and hence generalizable. Offering a benchmark based on these eGovernance workloads allows us to validate Big Data platforms for these socially important services, and offers the research and practitioner community additional transparency into the internal working of such mission-critical operations. While these results can also help improve UIDAI's Big Data architecture and their SLAs, the practical limitations discussed earlier stand.

Besides the many extensions to this work that was discussed earlier, it is worth examining other such public sector workloads to understand features are intrinsic to them, and set them apart from enterprise workloads. This will help design more effective Big Data solutions and potentially open up research opportunities.

Acknowledgments. We are grateful for inputs provided by Dr. Vivek Raghavan from UIDAI, and UIDAI's public reports in preparing this article. The views and opinions of authors expressed herein do not necessarily state or reflect those of the Government of India or any agency thereof, the UIDAI, nor any of their employees.

References

1. Arasu, A., Cherniack, M., Galvez, E., Maier, D., Maskey, A.S., Ryvkina, E., Stonebraker, M., Tibbetts, R.: Linear road: a stream data management benchmark. In: VLDB (2004)
2. Baru, C., Marcus, R., Chang, W. (eds.): Use cases from NIST big data requirements working group V1.0. Technical report M0180 v15, NIST (2013). <http://bigdatawg.nist.gov>
3. Dalwai, A. (ed.): Aadhaar technology and architecture: principles, design. best practices and key lessons. Technical report, Unique Identification Authority of India (UIDAI) (2014)
4. Ghazal, A., Rabl, T., Hu, M., Raab, F., Poess, M., Crolotte, A., Jacobsen, H.A.: BigBench: towards an industry standard benchmark for big data analytics. In: ACM SIGMOD (2013)

5. Gu, L., Zhou, M., Zhang, Z., Shan, M.C., Zhou, A., Winslett, M.: Chronos: an elastic parallel framework for stream benchmark generation and simulation. In: IEEE ICDE (2015)
6. Huang, S., Huang, J., Dai, J., Xie, T., Huang, B.: The HiBench benchmark suite: characterization of the mapreduce-based data analysis. In: Agrawal, D., Candan, K.S., Li, W.-S. (eds.) *New Frontiers in Information and Software as Services*. LNBIP, vol. 74, pp. 209–228. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-19294-4_9](https://doi.org/10.1007/978-3-642-19294-4_9)
7. Li, M., Tan, J., Wang, Y., Zhang, L., Salapura, V.: Sparkbench: a comprehensive benchmarking suite for in memory data analytic platform spark. In: *ACM International Conference on Computing Frontiers* (2015)
8. Lu, R., Wu, G., Xie, B., Hu, J.: Stream bench: towards benchmarking modern distributed stream computing frameworks. In: *IEEE/ACM UCC, 2014* (2014)
9. Nabi, Z., Bouillet, E., Bainbridge, A., Thomas, C.: *Of Streams and Storms*. Technical report, IBM (2014). <https://github.com/IBMStreams/benchmarks>
10. Office of the Chief Financial Officer: Office of Biometric Identity Management Expenditure Plan: Fiscal Year 2015 Report to Congress. Technical report, Office of Biometric Identity Management, Homeland Security, United States (2015)
11. Poess, M., Smith, B., Kollar, L., Larson, P.: TPC-DS, taking decision support benchmarking to the next level. In: *ACM International Conference on Management of Data (SIGMOD)*, pp. 582–587. ACM (2002)
12. Welsh, M., Culler, D., Brewer, E.: SEDA: an architecture for well-conditioned, scalable internet services. In: *ACM SOSP* (2001)
13. Zhao, J.-M., Wang, W.-S., Liu, X., Chen, Y.-F.: Big data benchmark - big DS. In: Rabl, T., Jacobsen, H.-A., Raghunath, N., Poess, M., Bhandarkar, M., Baru, C. (eds.) *WBDB 2013*. LNCS, vol. 8585, pp. 49–57. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10596-3_5](https://doi.org/10.1007/978-3-319-10596-3_5)

Big Data Benchmarking

6th International Workshop, WBDB 2015, Toronto, ON,
Canada, June 16-17, 2015 and 7th International
Workshop, WBDB 2015, New Delhi, India, December
14-15, 2015, Revised Selected Papers

Rabl, T.; Nambiar, R.; Baru, C.; Bhandarkar, M.; Poess,
M.; Pyne, S. (Eds.)

2016, IX, 129 p. 60 illus., Softcover

ISBN: 978-3-319-49747-1