OpenSearchCon INDIA

ODD is the new ELK?

How ODD Powers Log Management at 400TB-Per-Day Scale

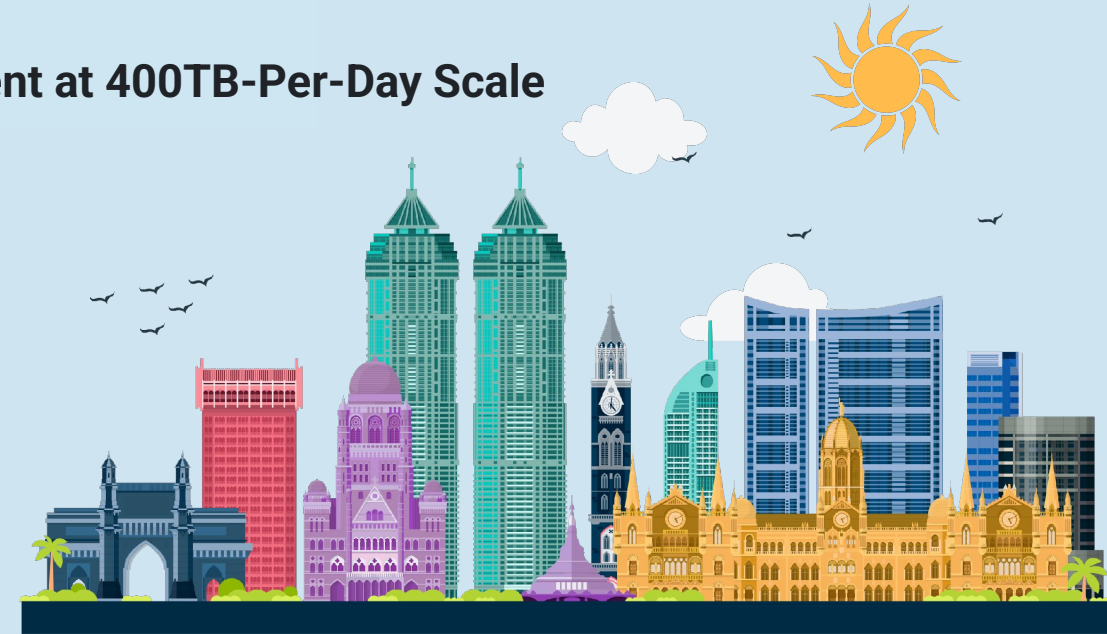

Speakers

OpenSearchCon INDIA

Rashmi Ramanathan

Principal Engineer - Systems, Observability

freshworks

Shaik Buden Saheb

Staff Engineer - Systems, Observability

freshworks

Freshworks at a glance

2010
Founded

FRSH
IPO September 2021

$838M
2025 Annual Revenue Guidance

~4,500+
Employees

75,000
Total Customers

Recognition
Winner of AI Excellence Awards, Best SaaS in AI and Winner of Customer Excellence Awards

freshworks

Freshworks Observability Platform

OpenSearchCon INDIA

Logs
Centralized log management for deep debugging and historical analysis.

Metrics
Real-time performance data and infrastructure health tracking.

Traces
Distributed tracing to visualize requests across microservices.

Frontend Observability
RUM monitoring for real user experience and performance.

Profiles
Continuous profiling to optimize code-level performance.

Synthetic Monitoring
Proactive testing of APIs and workflows from global locations.

Alerts: Alerting capabilities on all signals

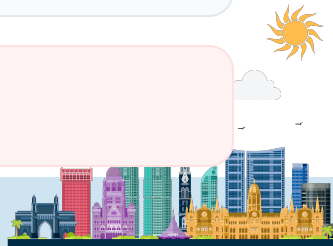

Observability Platform Mantra

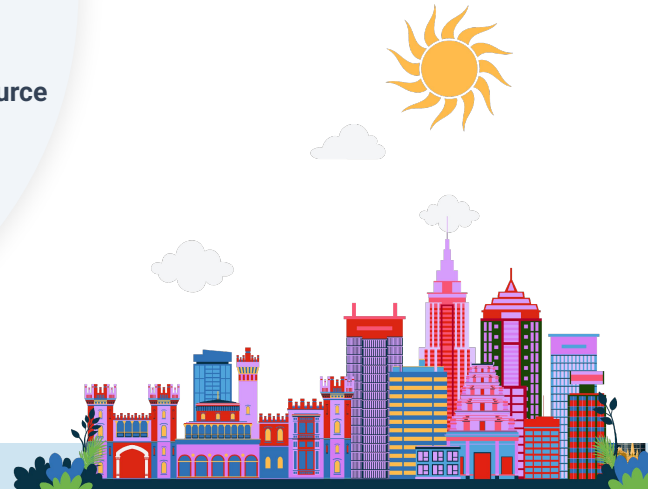

Who Are We & Why This Talk?

OpenSearchCon INDIA

Running one of the largest ODD deployments in production

Logs

~8 Mil. events / sec

~450 TB / day

Traces

~4 Mil. events / sec

~100 TB / day

Metrics

~20 Mil. samples / sec

Logs Infrastructure

Regions	Clusters	Retention	Nodes
5 AWS	7 ODD	21 Days	~1000+

6 GB/s

sustained ingestion

Goal

Share journey, patterns, and prove open source can compete at scale

The Problem - Log Management at Scale

The Challenge

Exponential Growth: Log volumes from microservices, containers, and distributed systems.

Vendor Lock-in: Rising costs and limited flexibility with proprietary solutions.

Extensibility: O11y under single pane of Glass. Need for RUM, synthetic monitors, traces, and profiles beyond just logs.

Key Pain Points

High Cost per GB
Unsustainable pricing models from commercial vendors.

Limited Customization
Rigid tools that don't fit domain-specific needs.

Data Sovereignty
Concerns over control and compliance of sensitive data.

Why ODD? The Migration Decision

OpenSearchCon INDIA

From ELK to ODD:

Open Source Commitment
No licensing surprises

AWS Integration
Native S3 support, serverless options

Active Community
OpenSearch Foundation backing

API Compatibility
Smoother migration path

Extensibility - Plugin architecture for custom needs

The Promise
Enterprise-grade observability without vendor lock-in

ODD Stack Overview

The Three Pillars of ODD

1. OpenSearch

Search & Analytics

Distributed engine for heavy lifting.

- Horizontal scaling
- Multi-tenancy
- Rich query DSL
- Aggregations

2. Data Prepper

Ingestion Pipeline

Server-side data processing.

- Transform & enrich
- Route logs
- Pluggable architecture

3. Dashboards

Visualization Layer

Unified UI for observability.

- Custom dashboards
- Extensible via plugins

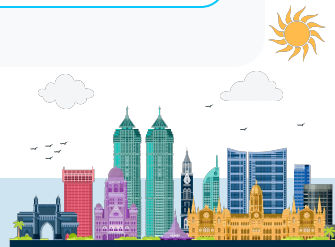

ODD Architecture

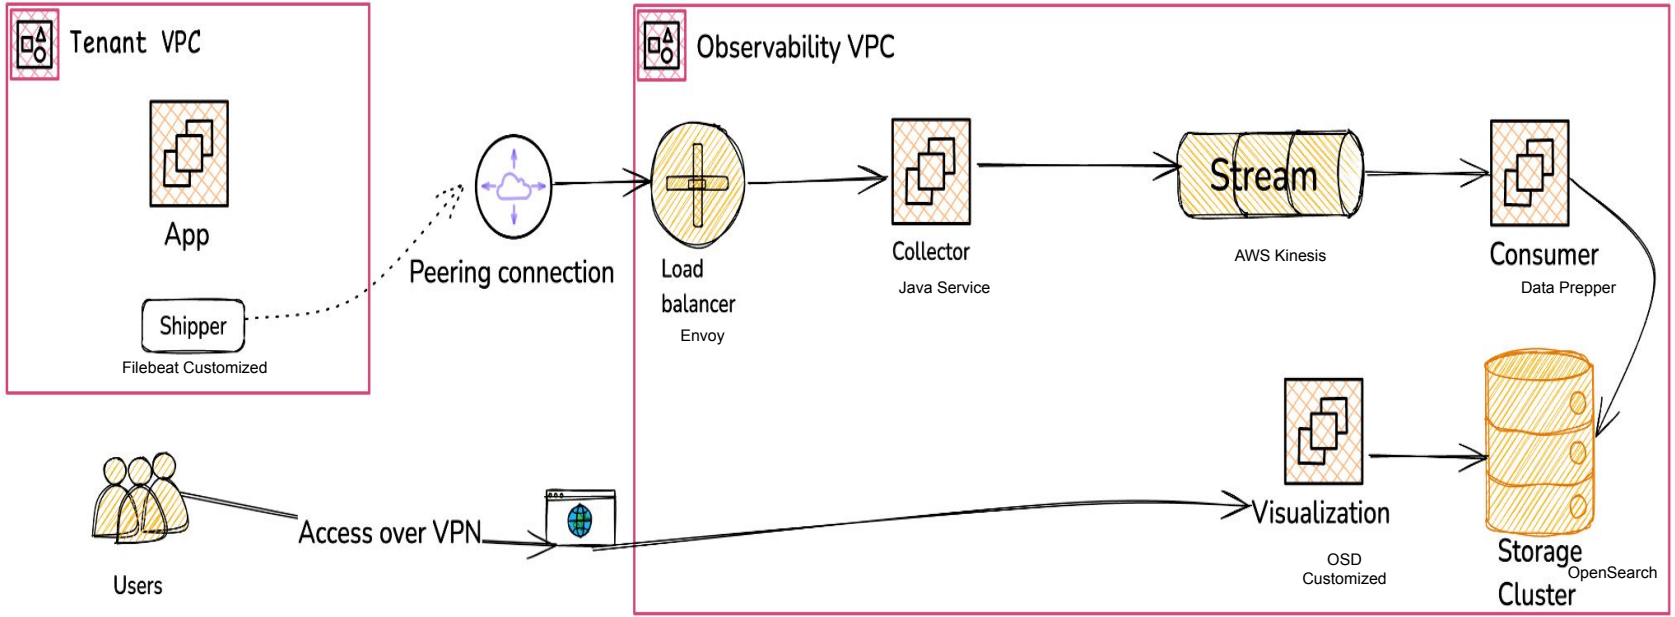

Data Prepper - The Transformation Engine

Why Data Prepper is Critical:

- **Performance:** Pre-processes before indexing → reduces OpenSearch load
- **Built-in Processors:** grok, JSON parsing, date conversion
- **Custom Plugins:** Our secret sauce

Benefits:

↘ **30% reduction** in cost compared to logstash

✅ **Consistent schema** enforcement

Multi-destination routing (logs → OpenSearch, Traces → Clickhouse)

Scaling OpenSearch - The Core

Cluster Design

Master Nodes
3 per cluster (cluster coordination)

Data Nodes
~1,000 nodes (hot)

Search Nodes
~ 200 nodes (query routing)

Snapshots
Tier to S3

Sharding Strategy

- **Index Management:** Rollover policy — time-based and size-based triggers
- **Isolation:** Hard-pinned tenancy — one tenant, one index, one subcluster
- **Distribution:** Shard-per-node layout — data spread across the cluster
- **Availability:** 1:1 replica ratio — one replica per primary shard

Opensearch Dashboards

OpenSearchCon INDIA

Logs Search & Analytics

Unified exploration of structured and unstructured log data.

Extensibility: Custom Plugins

1. Trace Visualization

Advanced analytics for distributed traces.

2. RUM Analytics

Real-user performance monitoring.

3. Profiling

Deep-code performance analysis.

4. Alerts

Proactive incident management.

5. Synthetic

API & UX uptime checks.

Architecture: Abstracts the underlying datasource for a unified control plane.

Opensearch Dashboards on RUM data

OpenSearchCon INDIA

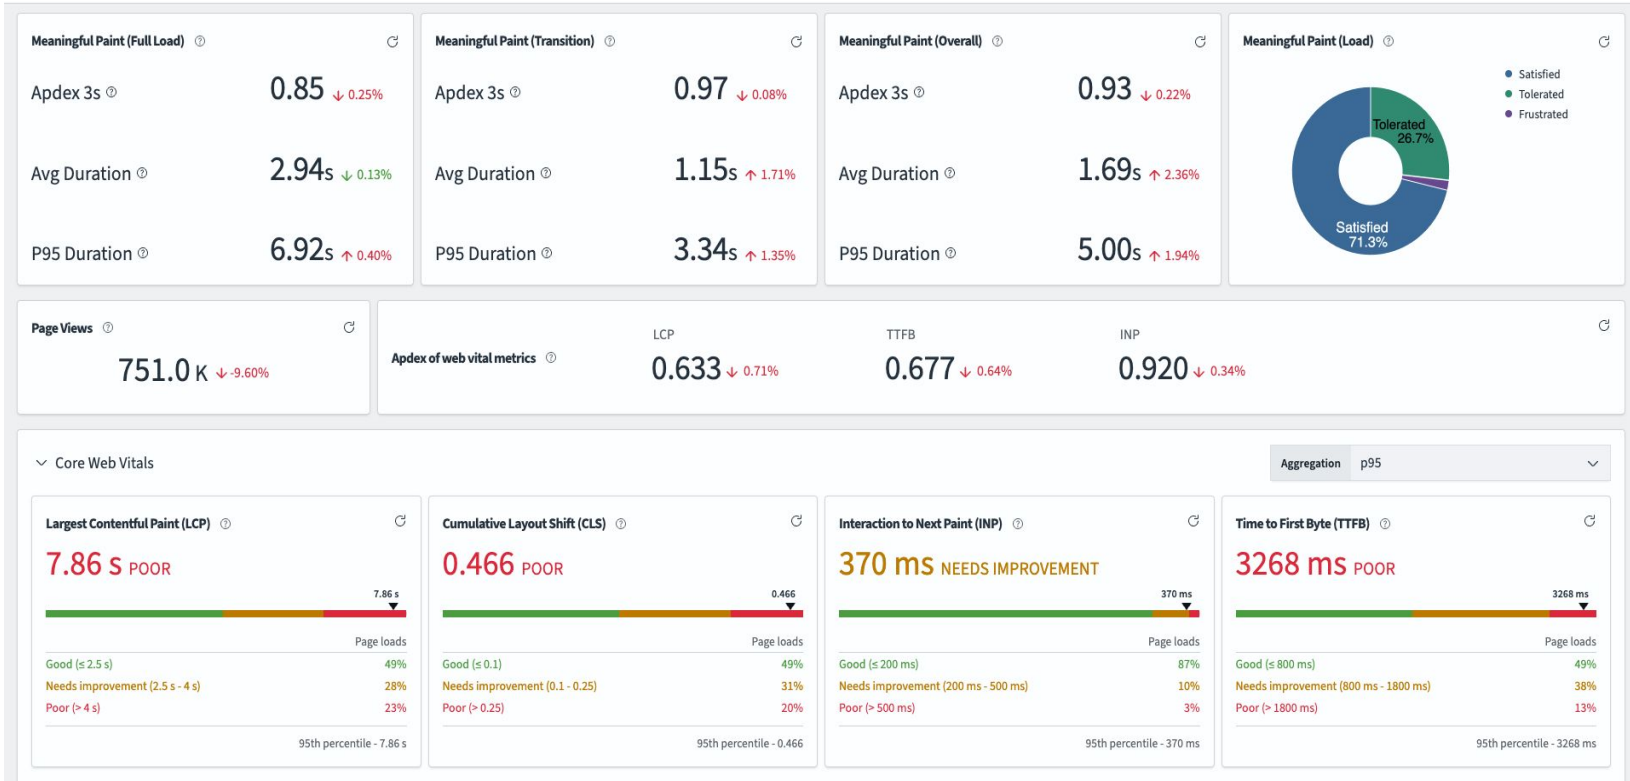

Extensibility - Custom Plugins

1. Data Prepper Processors

- Custom log parsers and transformers for performance
- HTTP output plugin for Clickhouse

2. OpenSearch Dashboards Plugins

📊 RUM (Real User Monitoring) dashboard
🔍 Synthetic monitoring integration
🔗 Distributed tracing viewer
🔥 Profiling/flame graph visualizations
🚨 Custom alerting workflows

Plugin Development

Backend / Core

Java/Kotlin for OpenSearch/Data Prepper

Frontend / UI

React/TypeScript for Dashboards

Results & Impact (Real numbers)

OpenSearchCon INDIA

Data Prepper

~12K ops/sec

Throughput per single data-prepper pod.

OpenSearch

200K requests/hour

Optimized serving performance for search requests.

OpenSearch Dashboards

~500 unique daily users

Handling ~4K daily requests with high engagement.

Impact Summary: Significant operational scale achieving high throughput and consistent user adoption.

Lessons Learned

What We'd Tell Our Past Selves:

📋
Start with ISM policies
Automate lifecycle from day 1

👀
Monitor ingestion lag obsessively
It's your canary

📏
Shard size matters
Too large = instability, too small = overhead

💰
Budget for Data Prepper
It's not "just a proxy"

🔌
Plugin development is powerful
Don't be afraid to extend

❄️
S3 remote snapshots are a game changer
Embrace cold tier

🤝
Community is your ally
File issues, contribute back

Contributing Back to ODD

Our Open Source Contributions:

🛠️ Data Prepper

- PR #5888
- PR #5881
- PR #5875
- PR #6425
- PR #6361
- PR #6423
- PR #6424
- PR #5272

📊 OpenSearch Dashboards

PR #5635: Enhancement and Bug Fixes

🤝 Join the Community

We believe in the power of open source. Check out our contributions and get involved with the OpenSearch and data-prepper project.

Thank you!

Let's discuss in Q&A