

# Stopping performance regression via changepoint detection

SREcon25 Americas  
March 27, 2025

Shanthini Velan, Software Engineering Team Lead  
Joseph Cirella, Software Engineer

[TechAtBloomberg.com](https://TechAtBloomberg.com)

© 2025 Bloomberg Finance L.P. All rights reserved.

Engineering

Bloomberg

# What is a changepoint?

*Moments of abrupt change in the behavior of a time series are often cause for alarm as they may signal a significant alteration to the data generating process.*

## - An Evaluation of Change Point Detection Algorithms




Figure 2 from [Bayesian Online Change Point Detection](#)

# What are we going to talk about?

Our experience building a **changepoint detection product** to pinpoint significant performance changes.

## Including

- **Why** changepoint detection?
- **When** are **probabilistic** methods right for you?
- **Challenges** we encountered
- **Thought-process** during decisions



A decorative particle trail in the top right corner, composed of many small, multi-colored dots (green, blue, red, purple) that form a curved, comet-like shape pointing towards the top right.

# Why?

**TechAtBloomberg.com**

© 2025 Bloomberg Finance L.P. All rights reserved.

**Bloomberg**

Engineering

# Ticker Plant

Ticker Plant feeds the vast majority of market data into Bloomberg's systems. If you use:

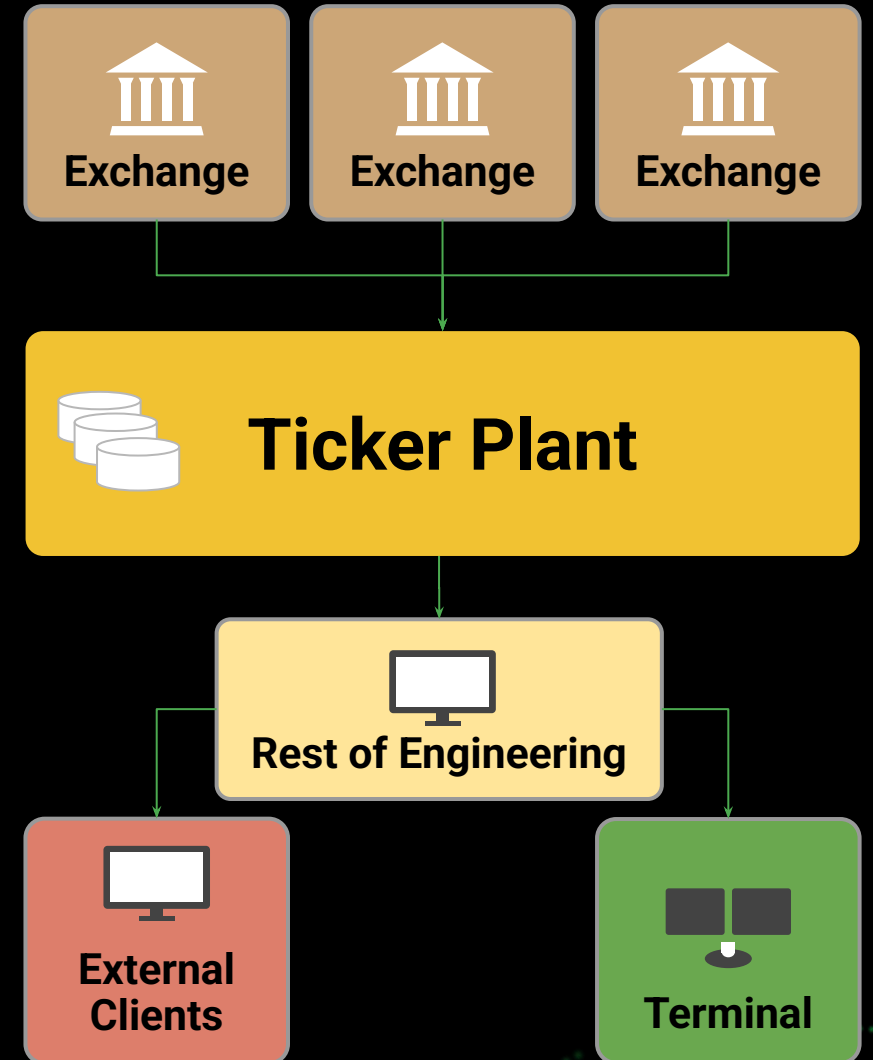
- The Bloomberg Terminal
- Data License
- B-PIPE

Then, you probably rely on Ticker Plant.

We service historical data requests and offer **real-time publishing**. Clients that are technically-advanced often care a **lot** about **latency**. **When our latency changes, so does much of Bloomberg's.**

**TechAtBloomberg.com**

© 2025 Bloomberg Finance L.P. All rights reserved.



# Low Latency – a feature of our product

Low latency is crucial to the success of market participants

We need to be able to know whether the efforts of our Ticker Plant teams are what's making our system's performance better or worse

We need to map actions to outcomes

**Bloomberg**

Engineering

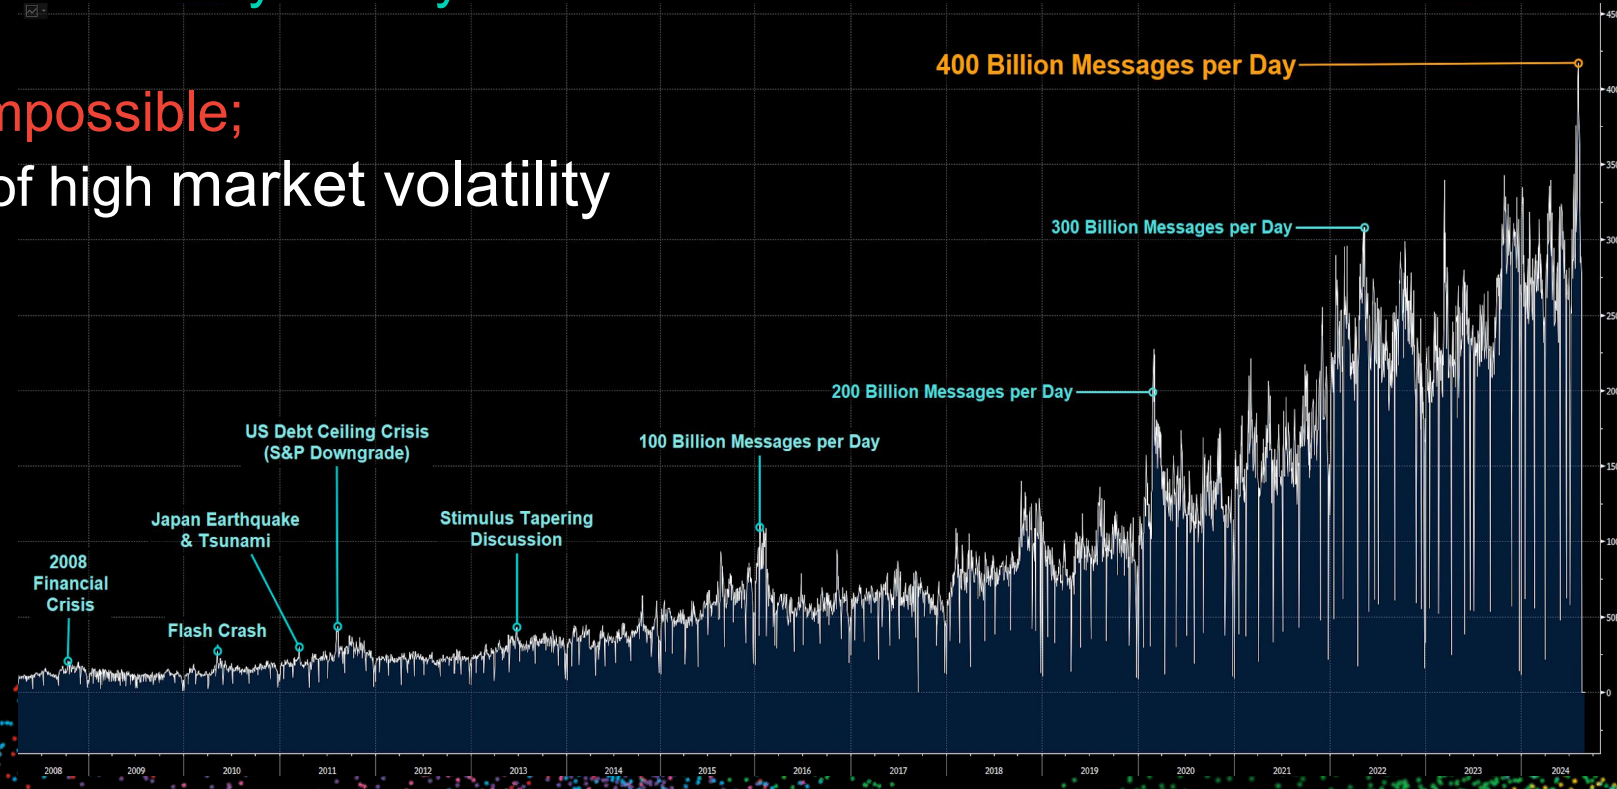
# What makes analyzing Ticker Plant latency hard?

- Latency is often influenced by factors beyond our control
  - News & politics
  - Black swans
  - Data feed bugs
- Expectations change by instrument and event
- Rarely-traveled code paths with bursty activity

Sometimes avoiding latency is impossible;

Latency will increase in periods of high market volatility

*Market messages received daily*




# Actions to outcomes

What is **measurable** and **actionable**?

- Code rollouts
- Bugs
- Feature flags
- Hardware changes
- OS changes
- Back-pressure from downstream systems

Bloomberg | INFR



7.X → 8.X →  
9.X

S0 → S1 → S2 → S3 → S4

TechAtBloomberg.com

© 2025 Bloomberg Finance L.P. All rights reserved.


Linux Penguin source: Larry Ewing ([lewing@isc.tamu.edu](mailto:lewing@isc.tamu.edu)) and [The GIMP](#)

Bloomberg

Engineering



# Motivations



Breadth & depth of data >>> current ability to extract insights

We have many clients who are interested in a **unique subset of the market**; this is something **we can't know about ahead of time**

Establishing **realistic performance tests** that cover the **universe of market scenarios** and **client access patterns** is **prohibitively expensive**

**Bloomberg**

Engineering



# Action

**TechAtBloomberg.com**

© 2025 Bloomberg Finance L.P. All rights reserved.

**Bloomberg**

Engineering

# A gentle push

## A Problem Appears

Decreases in performance sweeps across machines

- No clear pattern of stages or code rollout
- It's not even obvious when it started

## Solution

Using a rolling confidence interval, we were able to flag changepoints in performance metrics

With our changepoints known, we could pull all the relevant metadata regarding builds, machines, tags, and feature flags for programmatic analysis



Generated with ChatGPT / DALL-E


# The findings



*Worktime over time for a single machine ( $\mu s$ ); lower is better*




# What algorithm should we use?



*Three separate latency series; lower is better*

# What algorithm should we use?



*Lower is better; relatively consistent*


**TechAtBloomberg.com**

© 2025 Bloomberg Finance L.P. All rights reserved.

**Bloomberg**

Engineering

# What algorithm should we use?



*Lower is better; consistently elevated*


**TechAtBloomberg.com**

© 2025 Bloomberg Finance L.P. All rights reserved.

**Bloomberg**

Engineering

# What algorithm should we use?



*Lower is better; more ambiguous*

**TechAtBloomberg.com**

© 2025 Bloomberg Finance L.P. All rights reserved.

**Bloomberg**

Engineering



# Selection

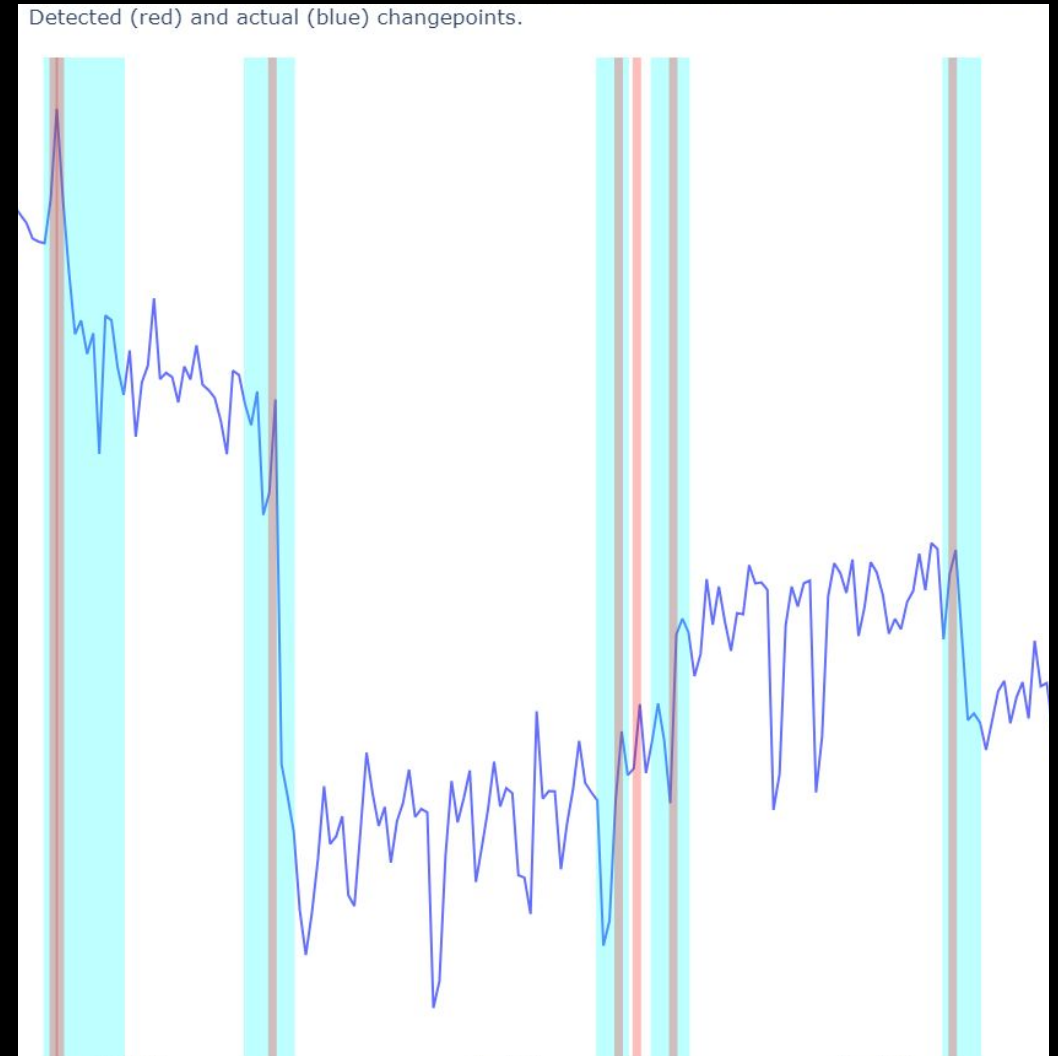
We decided to go with **Bayesian Online Change Point Detection**

**Traditional** techniques are often competitive with modern techniques

This method **fit our expectations for changepoints** (despite our data not obviously being in the **exponential-family**)

**TechAtBloomberg.com**

© 2025 Bloomberg Finance L.P. All rights reserved.



# Challenges

**TechAtBloomberg.com**

© 2025 Bloomberg Finance L.P. All rights reserved.

**Bloomberg**

Engineering

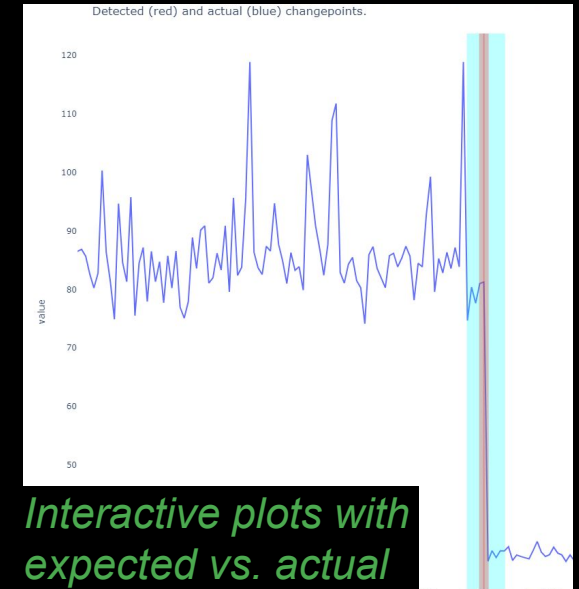
# Expected challenges

## Did well:

- Early tooling like **visualization and analysis**
- Investing in **synthetic data generation**

## Did not do well (mostly MDLC):

- Tracking **inter-annotator agreement**
- **Version** all of the data clearly
- **Document & organize** experiments
- **Narrowed the subset of data** we focused on



*Final\_final\_v2\_actual.pdf  
naming*

+	partial-results-with-interval	👤
+	results-with-daily-data	👤
+	results-with-interval-05-30	👤
+	with-merged-rolling-window	👤
+	with-rolling-window	👤

post\_bugfix\_part


Synthetic\_data\_r

# Unexpected challenges - Expected Changes

Dealing with **expected** changepoints

Our detector worked at consistently pointing out:

- Market **open/close**
- Machine **maintenance**
- Week**ends**
- Long Holidays



*Watch out! A market is opening and closing!*



# Unexpected challenges - Expected Changes

Upkeeping configuration files is surprisingly **complex & toil-intensive**

```
1 {
2   1 {
3     2 {
4       3 1 {
5         4 2 {
6           5 3 "periods": [
7             6 4 {
8               7 5 "data_source": "market-segment-123",
9                 8 6 "series_type": "some metric",
10                9 7 "description": "Cool exchange",
11               10 8 "initial_date": "2222-02-01",
12              11 9 "final_date": "*",
13             12 10 "start": ["15", "2", "*", "*", "1"]
14            13 11 "end": ["15", "12", "*", "*", "6"],
15           14 12 },
16          15 13 {
17            16 14 "data_source": "market-segment-123",
18              17 15 "series_type": "some metric",
19              18 16 "description": "Cool exchange",
20              19 17 "initial_date": "2222-02-01",
21              20 18 "final_date": "*",
22              21 19 "start": ["15", "4", "*", "*", "6"]
23             22 20 "end": ["15", "10", "*", "*", "0"]
24            23 21 },
25          24 22 ],
26          25 23 "one_offs": [
27            26 24 {
28              27 25 "data_source": "market-segment-123",
29                28 26 "series_type": "some metric",
30                29 27 "description": "Unexpected maintenance",
31               30 28 "start": "2222-02-22T02:22:22"
32              31 29 "duration": "00:08:00:00",
33             30 30 },
34            31 31 ],
35          32 32 ],
36        33 33 },
37      34 34 },
38    35 35 },
39  }
```



*Generated with ChatGPT / DALL-E*

# Unexpected challenges - Clutter & Noise

SETTINGS

Q Search data rows

run_id	Cluster	Start time	End time	Confidence%	Mean%	Mean+, ms	Node	Hostname	Group	Arch	DC	Node Stage	Pint	Links	Ticket
2	17	2024-01-21 17:00 -5	-	93	14	0.00457	6036		1			s4		LINKS	CREATE
3	2	2024-01-20 17:00 -5	-	75	29	0.01388	36						N/A		CREATE
12	0	2024-01-21 17:00 -5	-	92	25	0.02744	6875						N/A		CREATE
13	9	2024-01-21 17:00 -5	-	72	34	0.04334	2827						N/A		CREATE
14	9	2024-01-21 13:00 -5	-	92	44	0.02109	6921		2			s2		LINKS	CREATE

Rows per page: 5 1-5 of 12229 < >

Our initial UI  
(mock data)

# Unexpected challenges - Clutter & Noise

MARK ACCURACY

SETTINGS

UTC

Q e2e-test

	IsTP ↑	run_id	Target	Occurrence time	Confidence%	Mean+, ms	UserTags	CpUUID	Ticket
<input checked="" type="checkbox"/>	?	1367	Work Time	2024-09-06 00:05	56	116.45093999999999	e2e-test, test	2787780f-467b-4da9-a866-bc7d953855f9	<div>CREATE</div>
<input checked="" type="checkbox"/>	?	1367	Work Time	2024-09-16 00:05	49	45.784319999999994	e2e-test, test	0c3f882f-c126-47e4-b0ad-0d72f0bfb905	<div>CREATE</div>
<input type="checkbox"/>	?	1367	Work Time	2024-09-17 00:05	40	-107.69729799999999	e2e-test, test	298f95ea-5317-488c-9799-6a40aa037979	<div>CREATE</div>
<input type="checkbox"/>	✖	1367	Work Time	2024-09-05 00:05	65	9.385812999999999	e2e-test, test	9d5f2b5f-b2d0-4532-bb4a-eb942eb109eb	<div>CREATE</div>
<input type="checkbox"/>	✓	1367	Work Time	2024-09-10 00:05	77	-64.00916	e2e-test, test	d3c64cbc-1157-4532-a858-7ca9b0c85827	<div>CREATE</div>

Rows per page:

10

1-5 of 5

<

>

Our updated UI  
(mock data)

# Learning Summary

**TechAtBloomberg.com**

© 2025 Bloomberg Finance L.P. All rights reserved.

**Bloomberg**

Engineering



# Are Statistical Methods Right For Me?

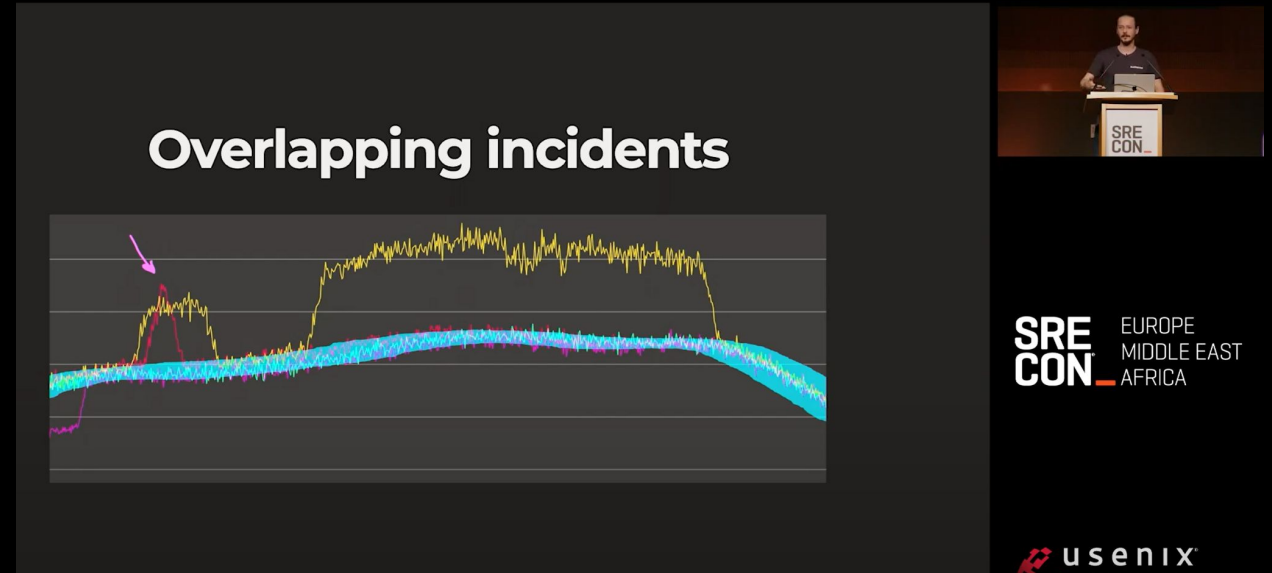
Consider if you have some of the following:

- Many **degrees of freedom** for system perturbation
- A **noisy** environment
- The behavior of your system changes over time
- Your system has **distributed responsibility**
- A **lot of data**

# What should I watch out for?

A **lot** more engineering work will be spent on **usability**, rather than on the core detection system

**Also check out:** [Anomaly Detection in Time Series from Scratch Using Statistical Analysis](#) - Ivan Shubin, SREcon24 EMEA



# Thank you!

# Questions?

<https://www.bloomberg.com/careers>

Contact us:

**Shanthini Velan** [svelan2@bloomberg.net](mailto:svelan2@bloomberg.net)

**Joseph Cirella** [jcirella@bloomberg.net](mailto:jcirella@bloomberg.net)



#MAKEITHAPPENHERE

Engineering

# Bloomberg

**TechAtBloomberg.com**

© 2025 Bloomberg Finance L.P. All rights reserved.