Production Engineering when trading Billions of Dollars a Day

Pedro Flemming

Jane Street

https://www.sec.gov/newsroom/press-releases/2013-222
https://www.sec.gov/files/litigation/admin/2013/34-70694.pdf

“Knight Day”

Trading is really scary

Robust redundant monitoring is crucial

Noisy alerts are worse than useless

Empower your support to take action

Jane Street

Production Engineering when trading Billions of Dollars a Day

Pedro Flemming

Jane Street

The Trading Environment

Jane Street

The Trading Environment

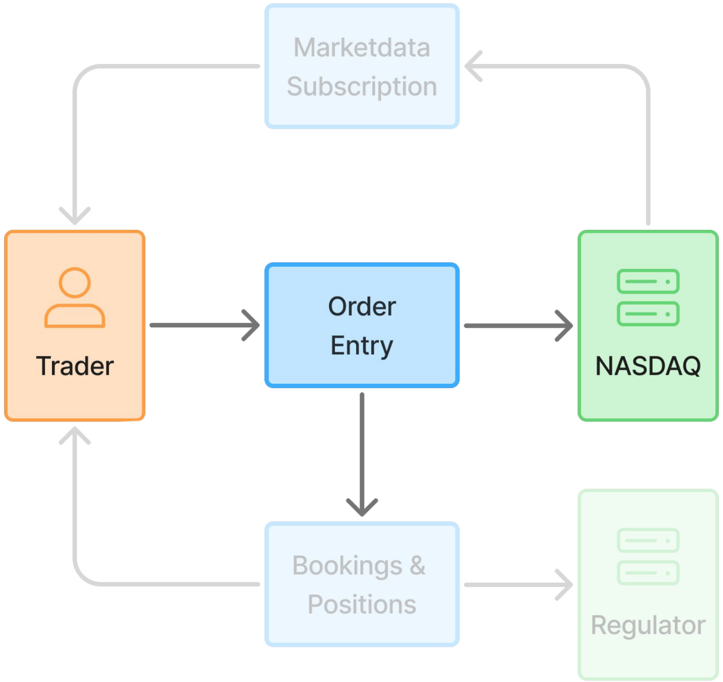

Jane Street

Anatomy of an Order

I want to **sell** **1 share** of **J-Com Co.** at a price of **610,000** **yen**.

Jane Street

Anatomy of an Order

I want to sell 1 share of J-Com Co. at a price of 610,000 yen.

I want to sell 610,000 shares of J-Com Co. at a price of 1 yen.

https://www.nytimes.com/2005/12/09/business/worldbusiness/japan-rebukes-exchange-for-costly-trading-error.html
https://www.jpx.co.jp/english/corporate/news-releases/0063/b5b4pj000000pr1b-att/20150904_E.pdf

Jane Street

Technical Issues in Trading

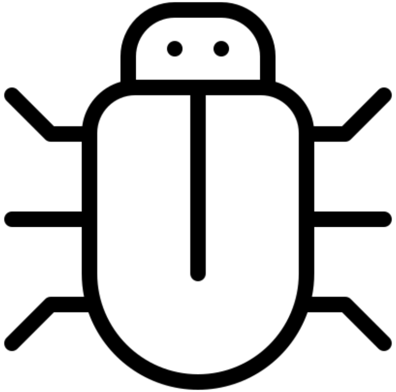

Bugs

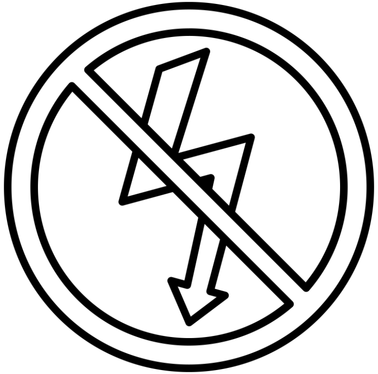

Outages

Jane Street

Nuances of the Trading Domain

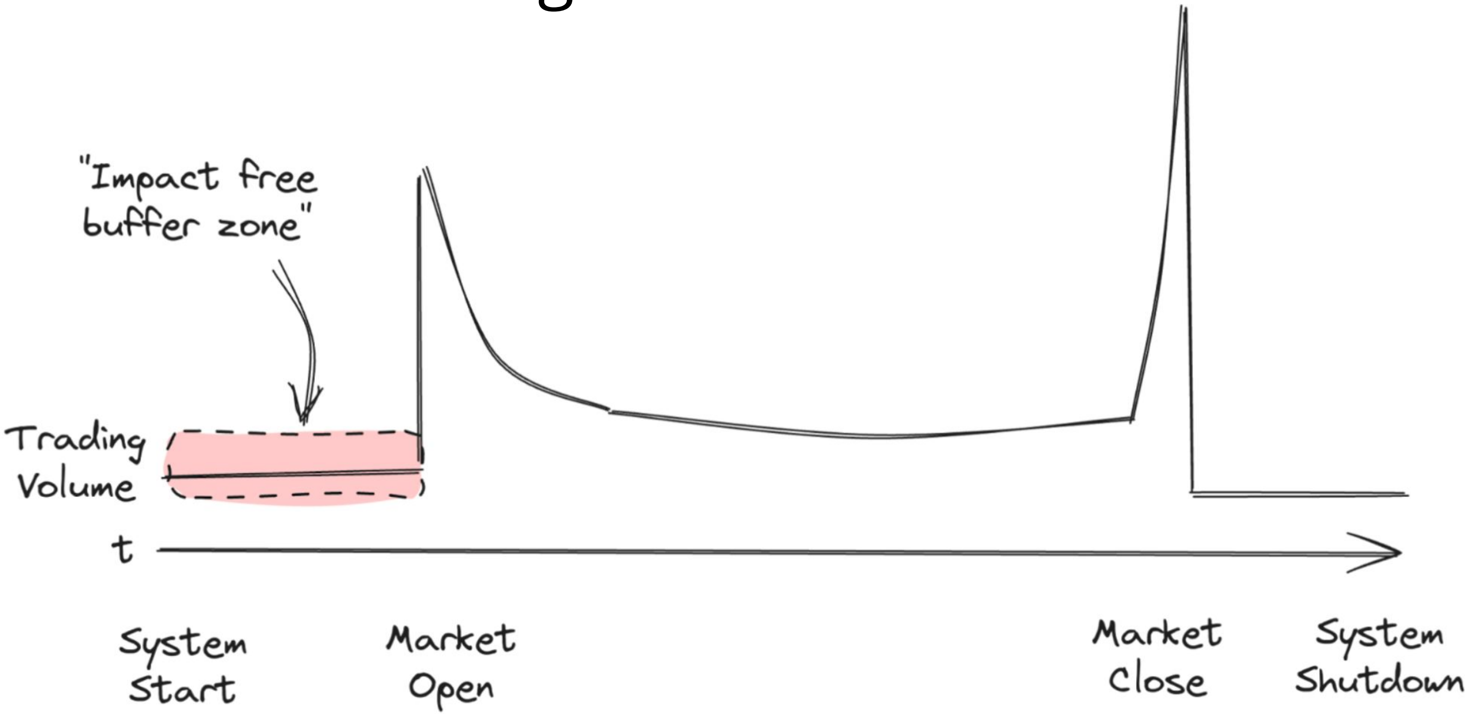

Time-dependent

Jane Street

Nuances of the Trading Domain

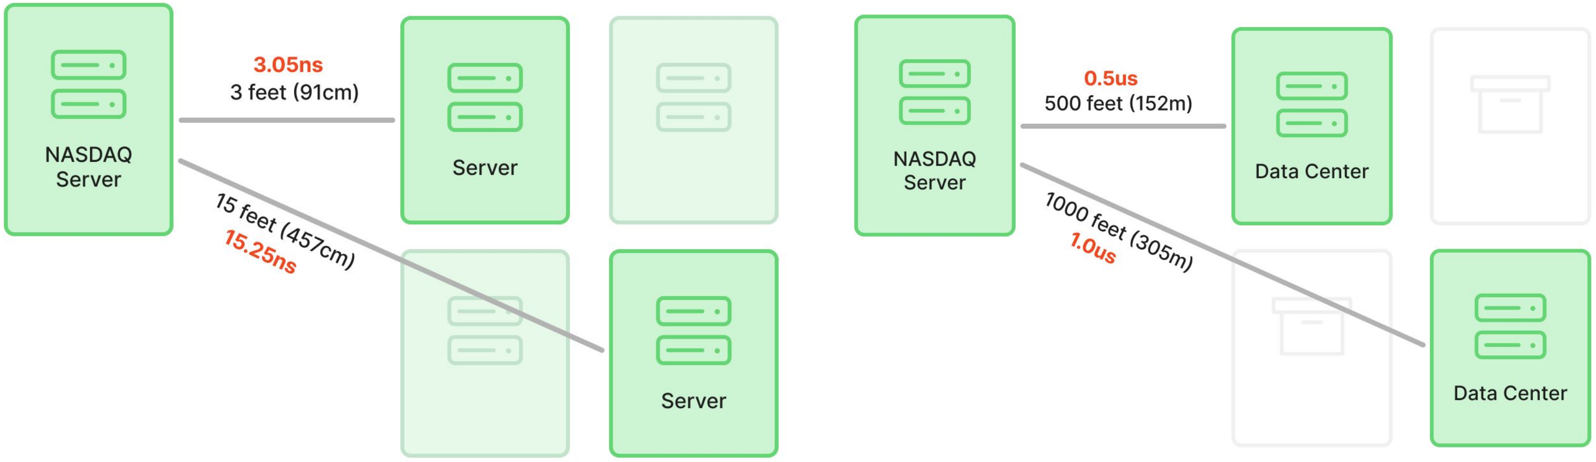

Fine-tuned technology

Jane Street

Nuances of the Trading Domain

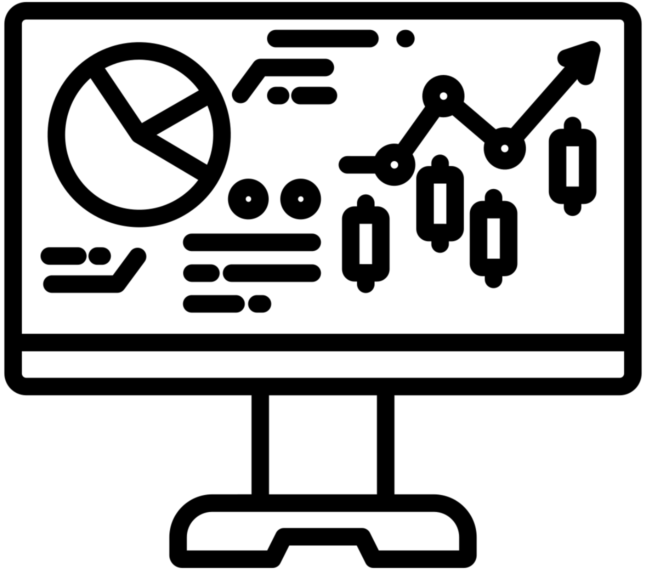

Trading is complicated

Nuances of the Trading Domain

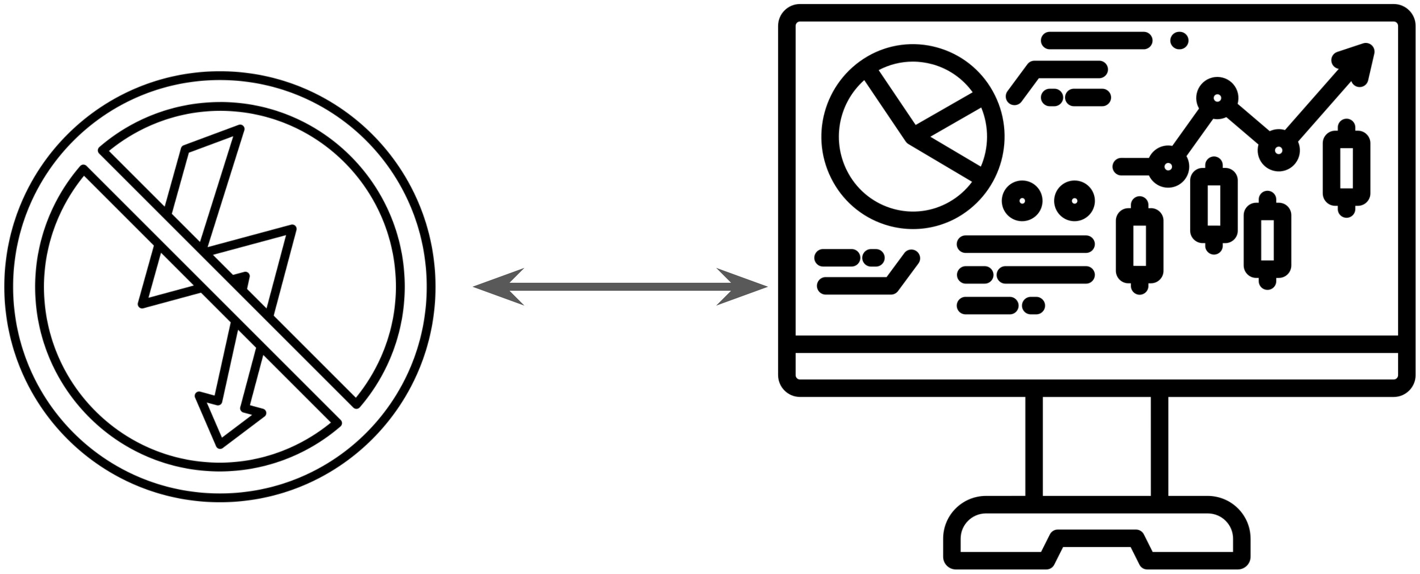

Your response to an incident depends on the business context

Monitoring

Monitoring is not optional

Jane Street

Monitoring

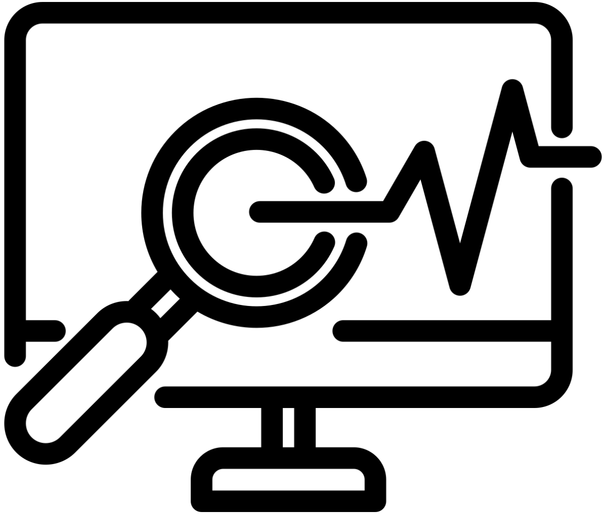

Technical Health

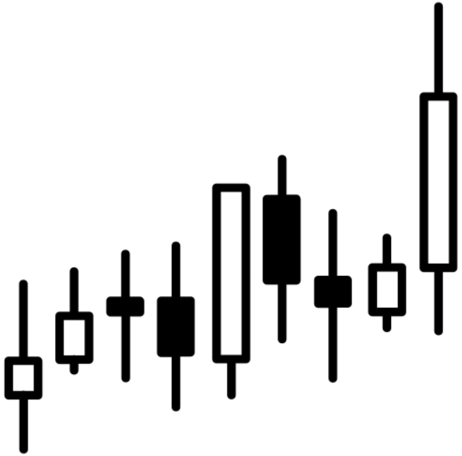

Trading Impact

Jane Street

The Trading Environment

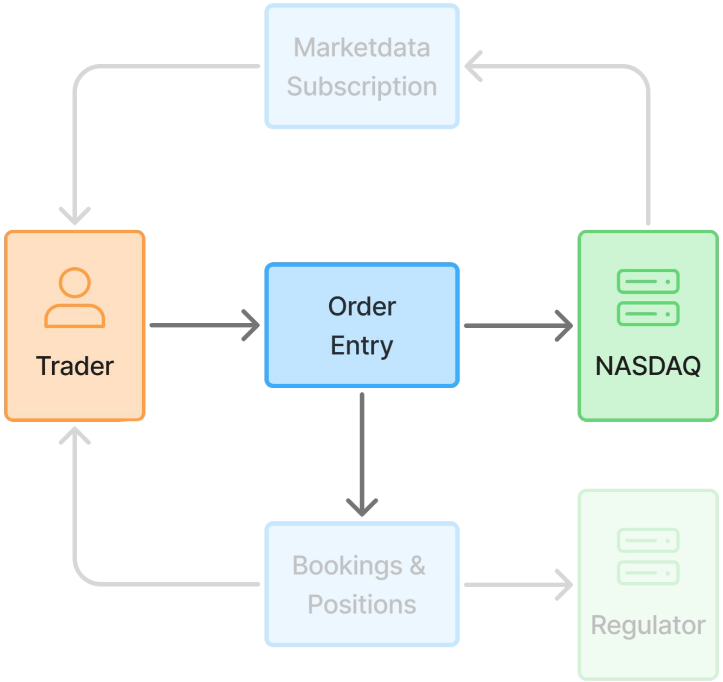

Basic message flows

Send New Orders

Cancel Orders

Receiving Acks or Rejects

Receive Executions

Jane Street

SLO-based Monitoring & Alerting

SLO: Order entry service uptime above 99.99%.

Problem: Any downtime requires immediate attention.

SLO-based Monitoring & Alerting

SLO: Error rate below 0.01%. i.e. 99.99% of all our requests should succeed.

Problem: Any given order can be critical to send or cancel.

SLO-based Monitoring & Alerting

SLO: External response latency. p99 latency should be below 1ms.

Yes! Increased average latency can be a leading indicator of issues.

However, tail latency also matters and can cause immediate losses.

Jane Street

Event-based Monitoring & Alerting

Event-based alerting for live trading systems

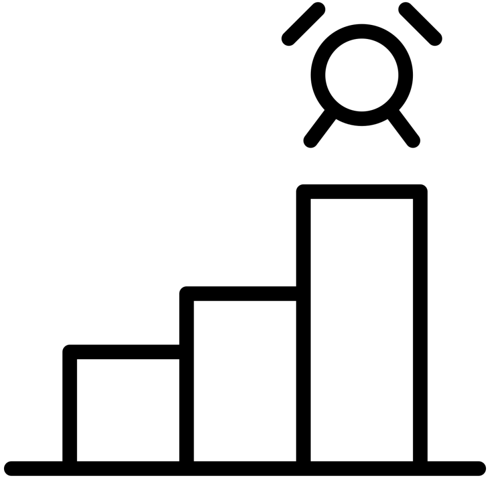

Metrics-based alerting for long-term health

Jane Street

What does support typically look like?

You’re a NY engineer supporting orderflow connections to our exchanges

US exchanges core trading sessions run from **9:30am to 4:00pm ET**

You and your secondary support engineer are on the morning shift and online at 9:00am to set up and catch up with London (“follow-the-sun” rotation)

The NY support day is split in two shifts: 9am-1pm, 1pm-5pm

People are usually on support roughly 1 week every 4-6 weeks

Incident Response

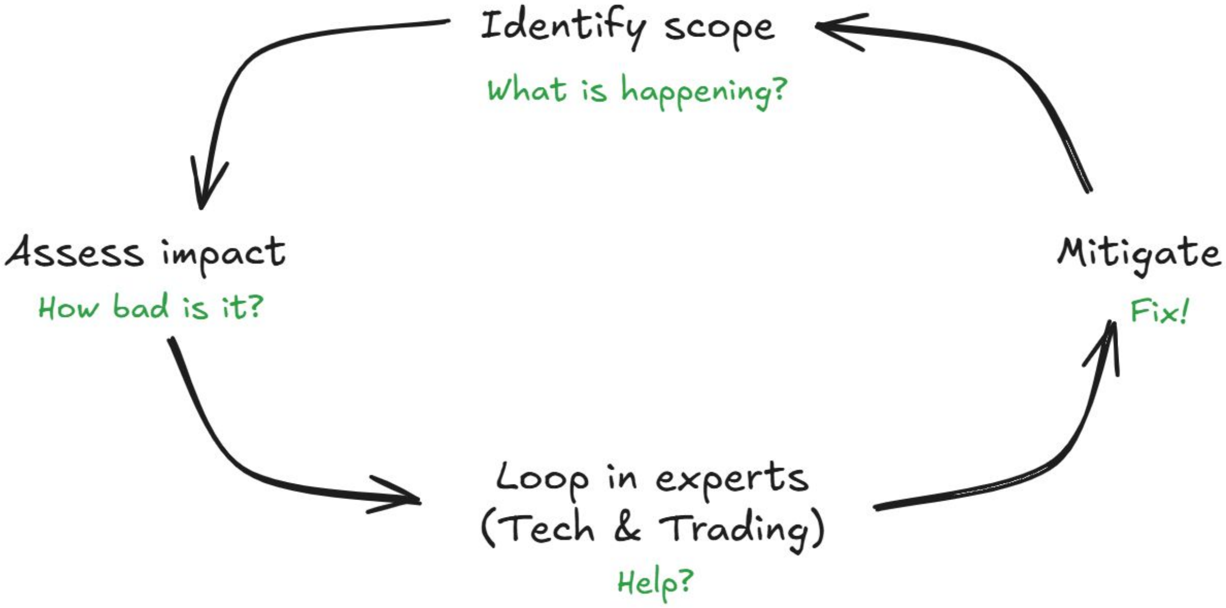

An incident log

[09:30] We're seeing rejects only on NASDAQ, so we conclude that the impact is limited to trading on that exchange

[09:31] When looking at the rejects, we see that they are all coming for orders from a trading system for trading in a subset of US ETFs

[09:33] Speaking to the ETFs desk, they confirm that they are seeing these issues and that the system isn't able to send out any orders

[09:34] We halt trading from this system entirely so that no further messages are going out. Since related versions of this system are used in other parts of the world, we are broadcasting a message to all trading desks

[09:38] We notice that the system in question was recently rolled to a new version that changed the message format we sent to the exchange. We confirm that these are invalid and are rolling back the system

[09:43] We are back online!

Post-mortem analysis

Impact: What was the overall cost? How much money did we lose?

Cause: How did this change go into production? Why didn't our tests catch it?

Discovery: Could we have noticed the issues sooner (e.g. with canaries)?

Mitigation: Why did it take 5 minutes to rollback the system?

Hiding mistakes is much worse than making them

Reliability Projects

Connectivity Monitor

Exchange	Status	Message Rate	Latency	Actions
NASDAQ	healthy	12 msg/s	**1.5ms**	see details
NYSE	healthy	8 msg/s	0.8ms	see details
ARCA	**disconnected**	0 msg/s	-	see details

see more

Jane Street

Automated Performance Benchmarking

Order Management System at **version 175.32** (run 2025-03-27, 11:08:33)

Name	Count	Mean	Min	Max	Δ Mean
order acceptance	21231	202.47ns	192.27ns	242.15ns	+10.53ns
handle cancel request	15546	413.62ns	273.43ns	533.81ns	+5.23ns
basket trade x10	8391	1.02us	0.82us	1.22us	-0.01us

Global Halt Service

Scope	Status	Actions	Reason
NYSE	**halted**	unhalt	Out of hours, no trading coverage
USD/GBP	**halted**	unhalt	Technical issues in FX reporting
system-A	**halted**	unhalt	Upstream metadata errors in Options symbology
> $100k Bonds	unhalted	halt	First production trades

Jane Street

Wrapping up

Production Engineering at this Trading Firm

In a business where high-value decisions are made in split-seconds, the ability for your support to react **correctly** and **quickly** is absolutely critical.

Actionable **event-based** alerts enable quick response

Monitoring systems are **critical** and need to be the most robust system

Your support needs **business knowledge** to make high-quality decisions

Empower your Production Engineers to **take action** in an emergency

For more details, check out our podcast!
https://signalsandthreads.com/solving-puzzles-in-production/

Jane Street

Thank you!

Jane Street