



# How ROSA enables MarketAxess to deliver our first AI driven product

Post Trade

Brian Yu, Rajeev Sachdeva, Arun Johny

## *High level*

- Reasons for cloud adoption
- Cloud Release Process - GitOps approach with quality gates
- SensAI and its journey
- Upcoming journey and challenges

# Hands off approach to SDLC through innovation

- **GitOps**

- ArgoCD syncs env with source of truth
- No more config drift, “source code not updated after EBF”

- **Automated release process per sprint**

- No human intervention

- **Scalability**

- HPA (Horizontal Pod Scaling) to scale pods on demand
- Autoscaling cluster add/remove worker nodes on demand
- No more manual requests for extra VMs, firewalls, certs, install Java, deploy code

- **Security**

- Shift left approach with Checkmarx and Twistlock scan integrated into secure CI/CD pipeline

- **Audit Trail**

- As per cloud SDLC process, SRE monitor tool is created for tracking quality gates in pipeline

- **Observability**

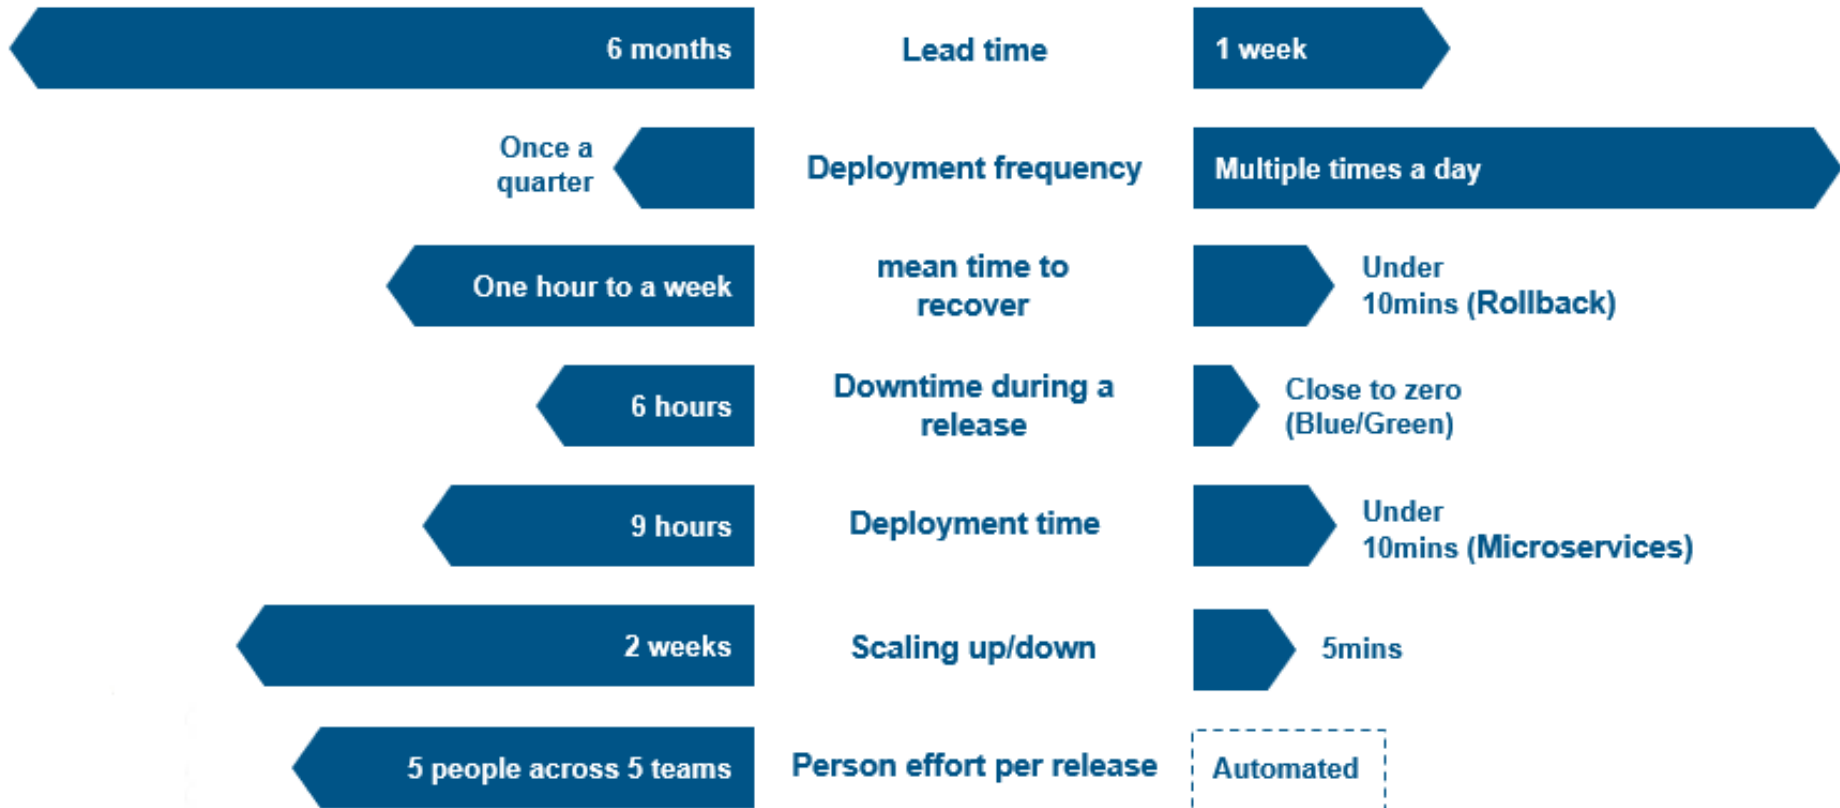
- Datadog alerts on config drift, worker nodes limit, etc



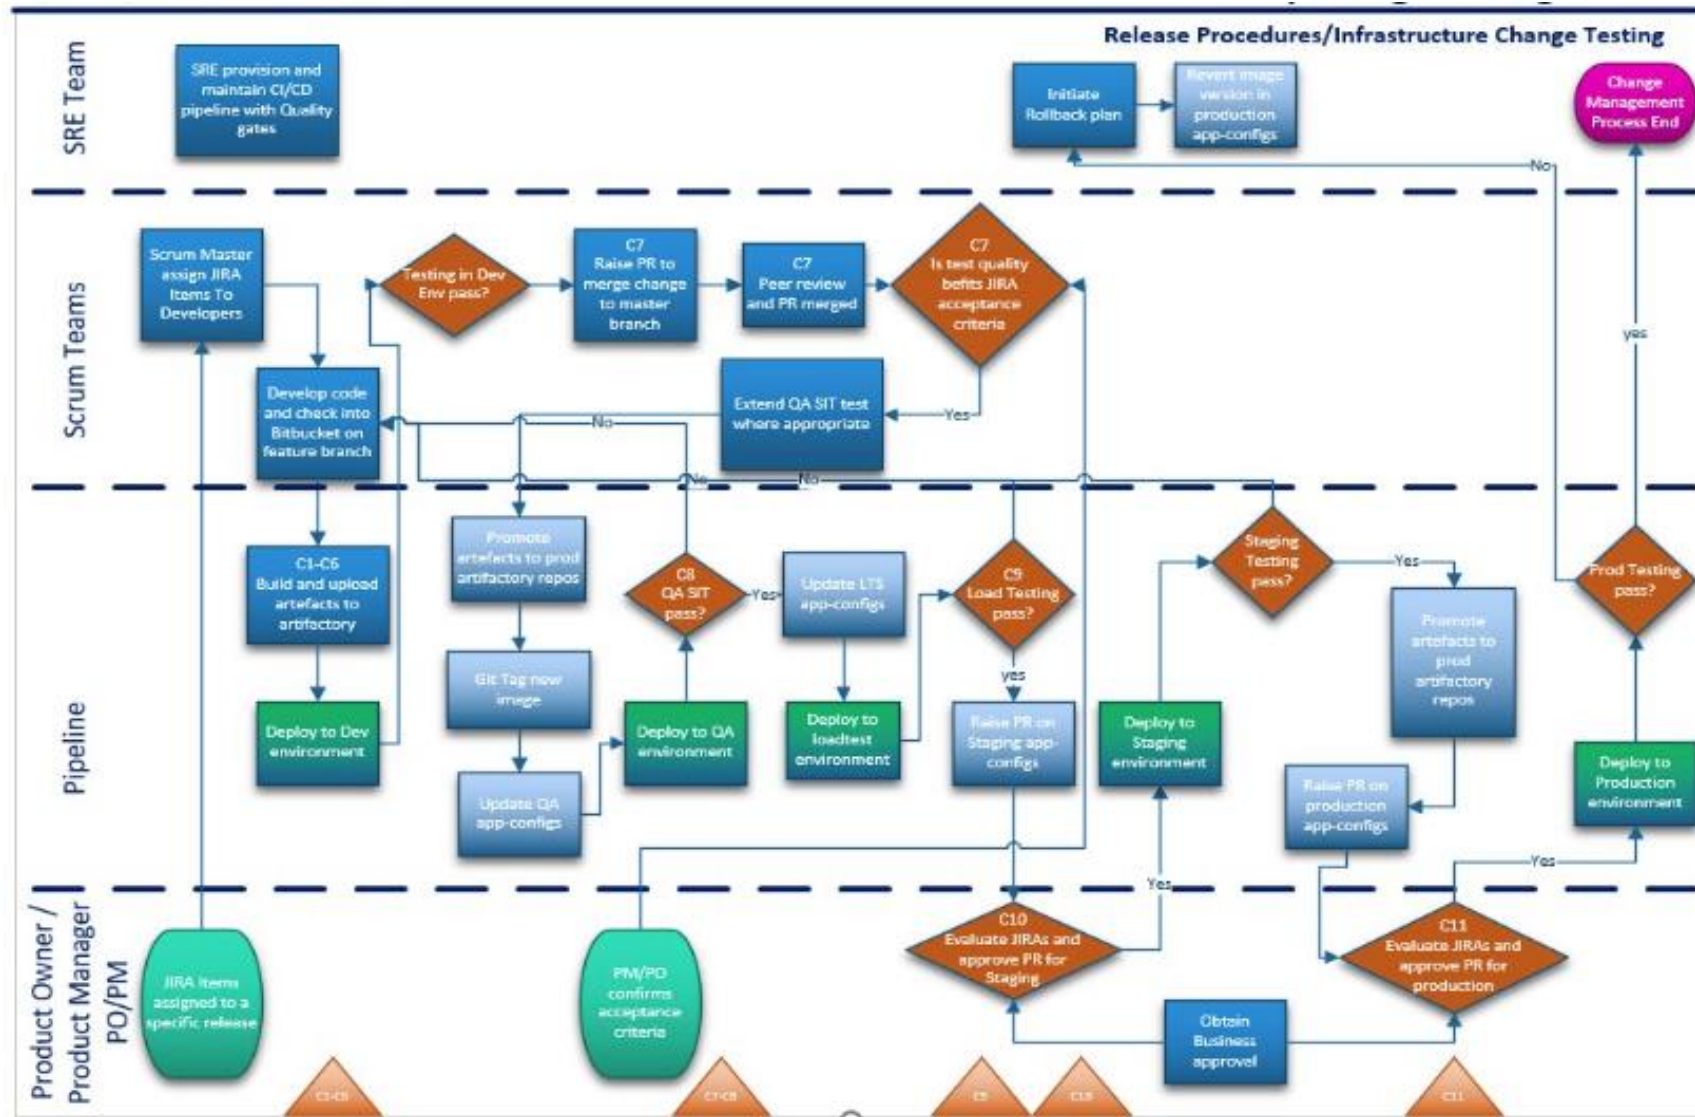
**Hands  
Off!**

# Statistics

# Baseline vs New World



# Cloud Release flow chart



# SRE Monitor - Release Audit tool

Audit trail on PR Approvals

Applications released by date

Release notes from JIRA

SRE Monitor Cluster Releases Applications
Summary Release Notes

## Cluster Releases

Here is a list of all the clusters with their releases.

LOADTEST

PRODUCTION

PTCLUS PTSFTR PTHAL PTSRE

Show  entries

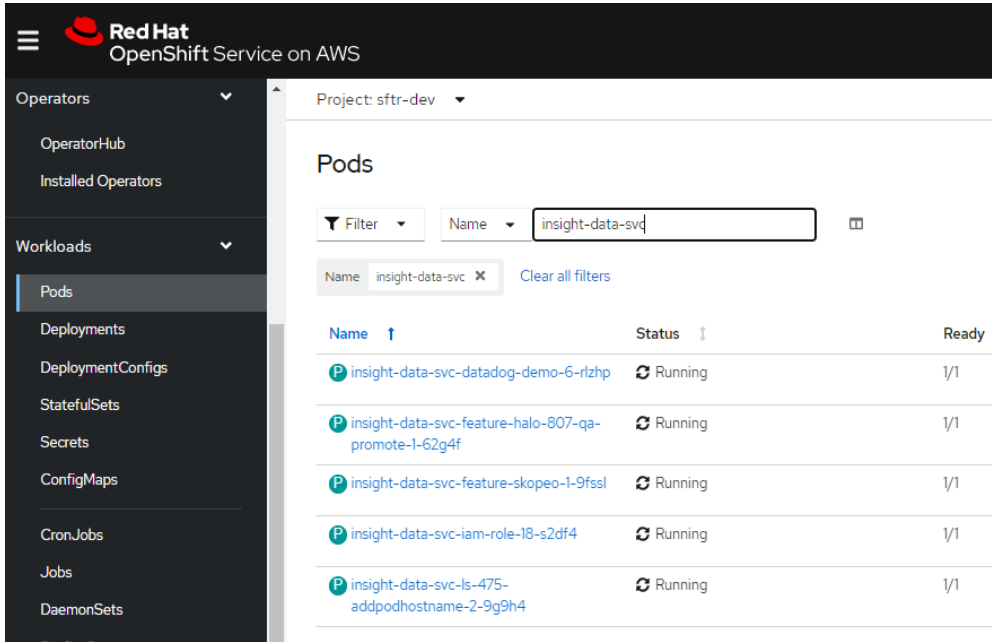
ID	JIRAs	Merged by	Reviewers	Release Date	Creation Date	Apps Released	More Info
<a href="#">213</a>		<span>ajohny</span>	<span>pandre</span> <span>rahsingh</span>	2022/06/20 10:45:20	2022/06/20 10:35:20		<a href="#">Details</a>
<a href="#">181</a>	<a href="#">HALO-1558</a> <a href="#">HALO-1562</a>	<span>ajohny</span>	<span>rahsingh</span>	2022/06/06 21:54:06	2022/06/06 16:41:06	<span>hal-db-extract:1.0.13</span> <span>hal-onetime-lei-load:1.0.7</span>	<a href="#">Details</a>
<a href="#">143</a>	<a href="#">HALO-1546</a> <a href="#">HALO-1538</a>	<span>rahsingh</span>	<span>rahsingh</span> <span>lurban</span> <span>nmos</span>	2022/05/23 09:27:23	2022/05/23 09:05:23	<span>hal-batch:1.4.1</span> <span>hal-service:1.10.0</span> <span>hal-ui:1.1.38</span>	<a href="#">Details</a>

**HALO-1546**  
No release notes available

**HALO-1538**  
No release notes available

**HALO-1270**  
This will enable user to see pair method having value (OTC and TVTIC) and message id fields on the transaction search screen. User can also filter on those fields.

# Build on Master branch



Developers checks on Pod status in ROSA

Developer generates PR to master branch and peer review

PO reviews test cases

Jenkins build on Master branch triggered via webhooks


Quality gates - Audit alert, Checkmarx, Unit testing, sonarqube, Twistlock

Image and artifacts pushed to artifactory (preprod repos)

✓ sfttr-poller-app < 37

Branch: master 17m 39s Changes by vdua, itcm-uk, RSachdeva

Commit: - 12 days ago Branch indexing



# Pipeline driven QA release

Applications / qa-app-of-sfr APPLICATION DE

APP HEALTH Healthy CURRENT SYNC STATUS Synced To master (fa669c0) LAST SYNC RESULT Sync OK To 9b3f85a

Author: Jenkins <jenkins-@marketaxess.com> Comment: AUTOMATED COMMIT - QA Deploy...  
 Author: Jenkins <jenkins-@marketaxess.com> Comment: AUTOMATED COMMIT - QA Deploy...

qa-app-of-sfr Healthy 2 months

- qa-sfr-billing processor 2 months
- qa-sfr-client-response-app 2 months
- qa-sfr-es-indexer a month
- qa-sfr-insight 2 months
- qa-sfr-liquibase 2 days
- qa-sfr-poller-app 2 months

ArgoCD monitors QA app-configs and update QA env

promote-from-qa < 112 > Pipeline Changes Tests Artifacts Logout

Branch: master 16m 49s Changes by RSachdeva  
 Commit: 2150850 15 days ago Started by user rsachdeva



Promote-from-qa Jenkins job invoked

Raise PR for LTS release - <1s

- > Shell Script <1s
- > Send Slack Message <1s




# Pipeline driven LTS release

LTS PR created with QA SIT test results and quality gates report

## LTS Merge Request

[Overview](#) [Diff](#) [Commits](#) [Builds](#)






 svc-rosa-jenkins-to-bitbucket created a pull request 24 Jun 2022

## SUMMARY

**Job name:** sftr-sit  
**Job number:** 192  
**Job branch:** master  
**Job result:** SUCCESS  
**Job test URL:** [link](#)  
**Job Cucumber URL:** [link](#)

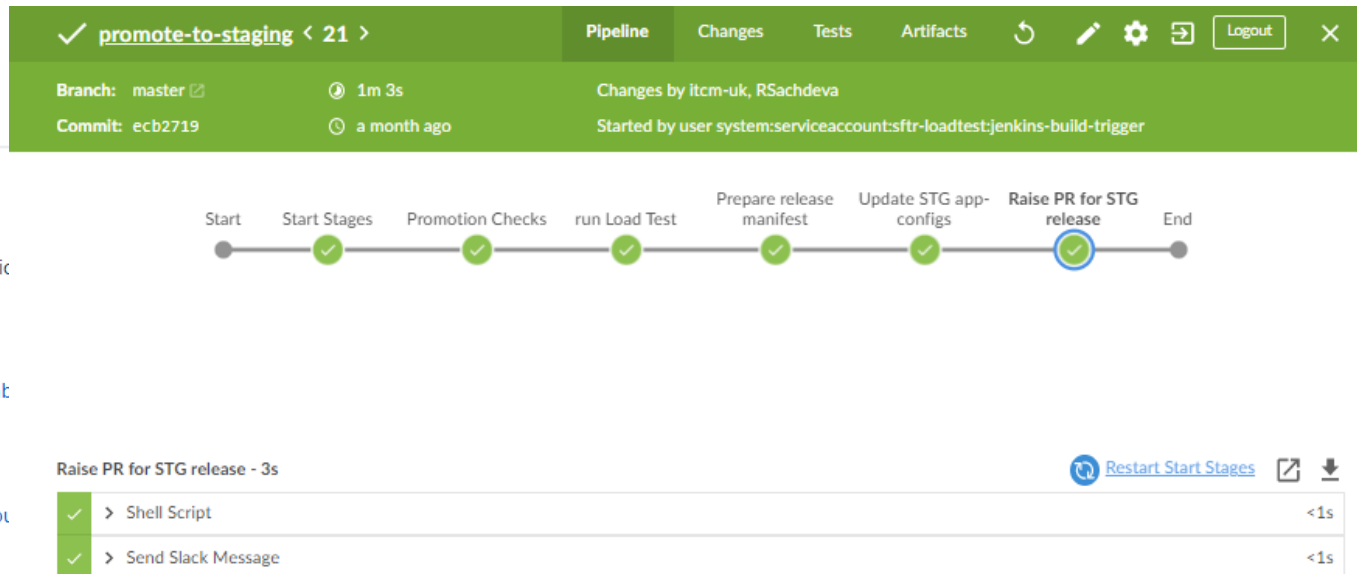
## CHANGELOG

### sftr-poller-app

- Tag Name:** sftr-poller-app1.21.0  
**Description:** MOD-125: Handled File Already Exist scenario
  - QUALITY GATES**
    - SONAR** (GATING ENABLED   
Result is OK   
Dashboard: <http://sonarqube.marketaxess.com/dash>
    - TWISTLOCK** (GATING ENABLED   
Vulnerability threshold check results: PASS   
Compliance threshold check results: PASS   
Link to the results in Console: <https://app3.prismacloud>
- Build Job URL:** [link](#)

ArgoCD monitors LTS app-configs and update LTS env

Promote-to-staging Jenkins job invoked



The screenshot shows the ArgoCD pipeline interface for the 'promote-to-staging' job. The pipeline is currently running on the 'master' branch with commit 'ecb2719'. The pipeline consists of several stages: Start, Start Stages, Promotion Checks, run Load Test, Prepare release manifest, Update STG app-configs, Raise PR for STG release, and End. The 'Raise PR for STG release' stage is currently active and highlighted with a blue circle. Below the pipeline overview, the details for the 'Raise PR for STG release' stage are shown, including a 'Restart Start Stages' button and a list of tasks: 'Shell Script' and 'Send Slack Message', both of which are completed.

✓ promote-to-staging < 21 > Pipeline Changes Tests Artifacts Logout

Branch: master 1m 3s Changes by itcm-uk, RSachdeva  
Commit: ecb2719 a month ago Started by user system:serviceaccount:sftr-loadtest:jenkins-build-trigger

Start Start Stages Promotion Checks run Load Test Prepare release manifest Update STG app-configs Raise PR for STG release End

Raise PR for STG release - 3s Restart Start Stages


- ✓ > Shell Script <1s
- ✓ > Send Slack Message <1s

# Pipeline driven STG release

STG PR created with LTS  
SIT test results and  
quality gates report

## STG Merge Request

[Overview](#) [Diff](#) [Commits](#) [Builds](#)

 [svc-rosa-jenkins-to-bitbucket](#) created a pull request 30 May 202

## SUMMARY

No Loadtest test have been run.

## CHANGELOG

### sftr-poller-app

- **Tag Name:** sftr-poller-app1.7.0  
**Description:** MOD-73: Loadtest NPE - updated app version ;  
**Build Job URL:** [link](#)
- **Tag Name:** sftr-poller-app1.9.0  
**Description:** updated openjdk version to 1.13 in Dockerfile

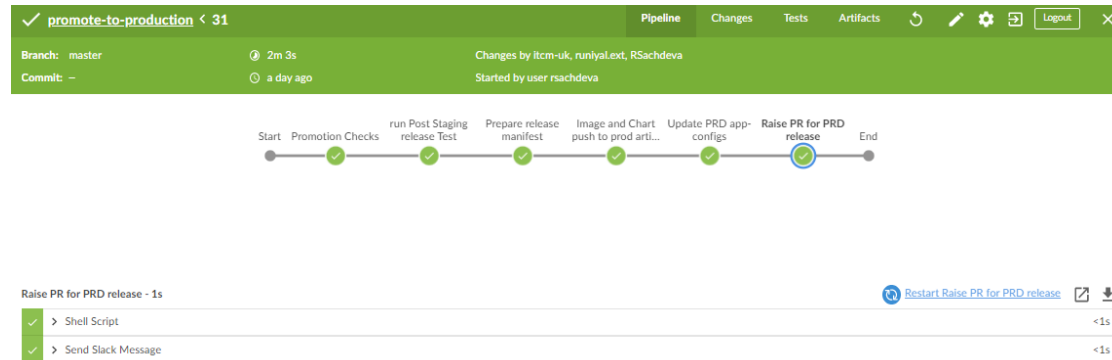
### sftr-processor-app

- **Tag Name:** sftr-processor-app1.7.0  
**Description:** MOD-83: Upgraded JDK Version to 1.13  
**Build Job URL:** [link](#)

ArgoCD monitors STG  
app-configs and update  
STG env

Promote-to-production  
Jenkins job invoked

Image push to  
Artifactory (Prod repos)



✓ promote-to-production < 31

Branch: master 2m 3s Changes by ltrcm-uk, runiyal.ext, RSachdeva  
Commlt: - a day ago Started by user rsachdeva

Start Promotion Checks run Post Staging release Test Prepare release manifest Image and Chart push to prod arti... Update PRD app-configs Raise PR for PRD release End

Raise PR for PRD release - 1s

- ✓ > Shell Script <1s
- ✓ > Send Slack Message <1s

# Pipeline driven PRD release

PRD PR created

ArgoCD monitors PRD app-configs and update PRD env

## PRD Merge Request

[Overview](#) [Diff](#) [Commits](#) [Builds](#)

 svc-rosa-jenkins-to-bitbucket created a pull r

## CHANGELOG

### sftr-poller-app

No Changes

### sftr-processor-app

- **Tag Name:** sftr-processor-app1.27.0  
**Description:** version updated  
**Build Job URL:** [link](#)
- **Tag Name:** sftr-processor-app1.26.0  
**Description:** Enabled SQL logs  
**Build Job URL:** [link](#)


# SensAI and its journey

- Bring your own model
  - Models created using ML workbench tools and other notebooks can be integrated ( SageMaker , H2O )
  - Scoring to be performed on the exported model from a research environment
  
- Scalability and Orchestration
  - Ability to load and perform scoring on daily and on demand
  - Cost Management
  - Use specialized hardware and custom configuration
  
- Speed to market
  - Agile release tools
  - Flexibility across hybrid cloud

# What is SensAI?




Service to provide reporting accuracy to the customers. Uses vast transaction data to analyse every trade.



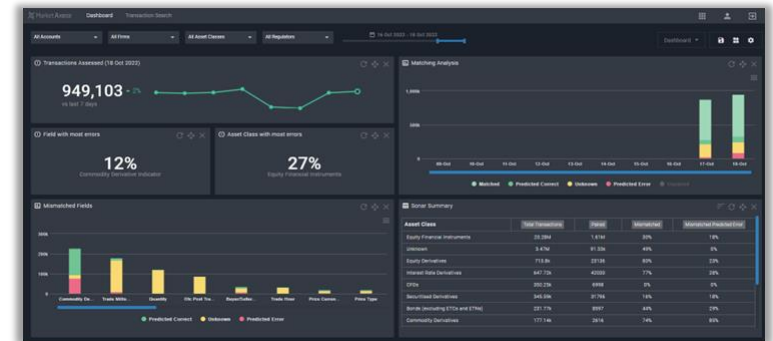
Rapid benchmarking of transactions of participants and counterparties. Using a mixture of Statistical analysis and machine learning based processes.



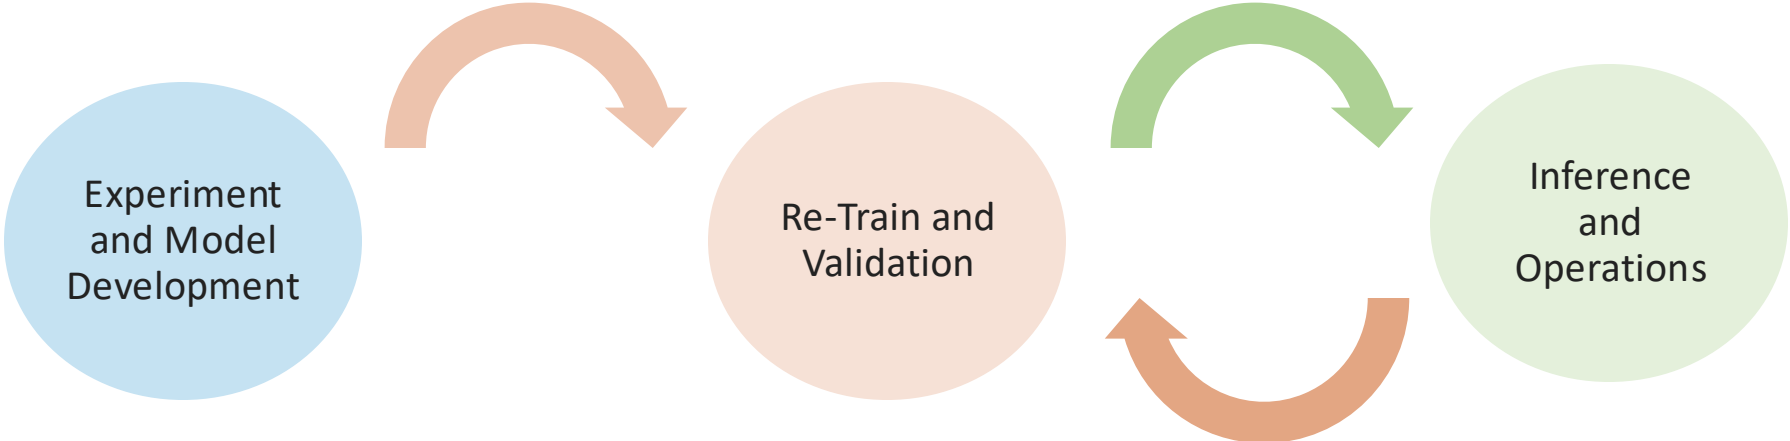
MarketAxess received ~5bn transactions in 2023. Data problems include counterparty identification, inconsistent use of price types and the various indicators and flags.



Using AI driven models – we have added multiple field-based prediction to detect and report the inaccuracy to the customers



# SensAI Workflow



- Create new models in a disconnected environment.
- Feature engineering to improve the accuracy scoring
- Embed with Red Hat AI – using Kubeflow and PyTorch

Workbench deployed to different Environments via MLOPs

Validation and deployment

Resource sizing/ pod configuration

Stat analysis need 30 days of paired data with the ability to deploy larger instances

Production promotion

Built-in security and operator integration

Custom machine pool and inference on demand

Feedback and data drift capture

Different operators ( Athena , MongoDB and Kafka )

# Requirements

- **Scalability**

- Autoscaling cluster add/remove worker nodes on demand

- **User story**

- On a given evening, when data becomes available, I want to run some scripts in loadtest environment to train and test my AI model. I'd need to throw a lot of data at it, so we need plenty of processing power which can be costly. Once the run is finished, I want to shut it all down to save cost.

- **Scheduling, pre-emption and eviction concept in Kubernetes**

- <https://kubernetes.io/docs/concepts/scheduling-eviction/taint-and-toleration/>
- Taints and tolerations are a flexible way to steer pods *away* from nodes or evict pods that shouldn't be running. A few of the use cases are
  - **Nodes with Special Hardware:** In a cluster where a small subset of nodes have specialized hardware (for example GPUs), it is desirable to keep pods that don't need the specialized hardware off of those nodes, thus leaving room for later-arriving pods that do need the specialized hardware.
  - This can be done by tainting the nodes that have the specialized hardware and adding a corresponding toleration to pods that use the special hardware.

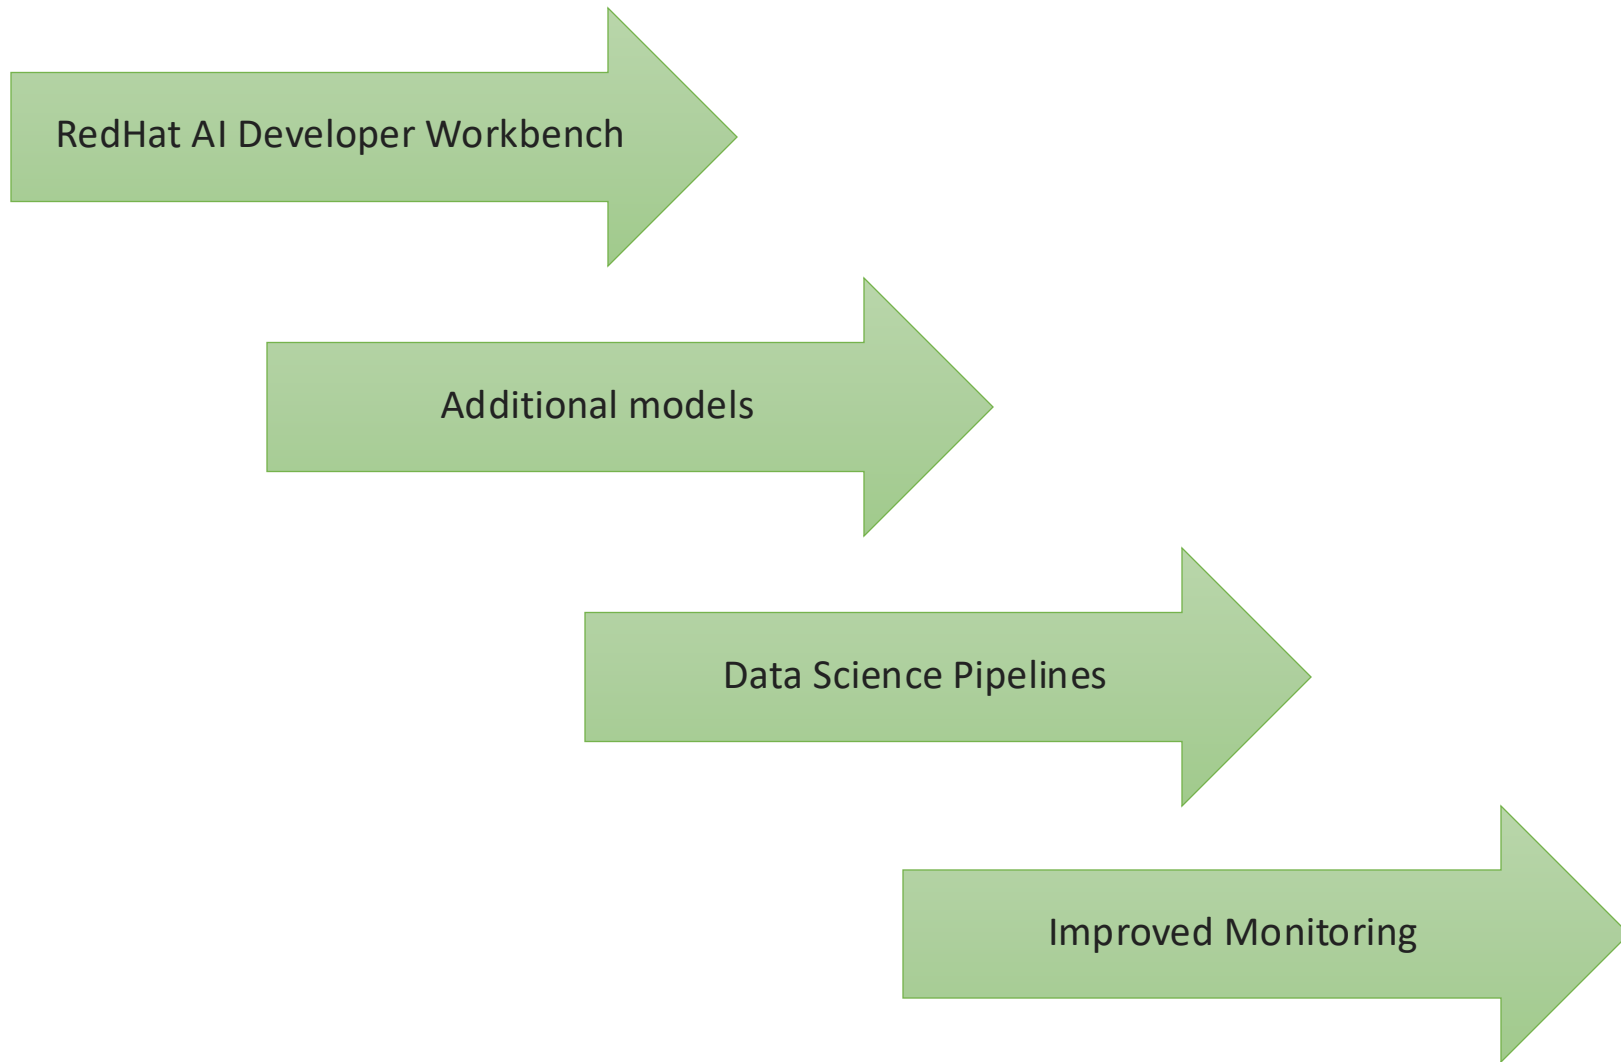
# Taints and toleration in ROSA

▼	high-mem-workers	r5.4xlarge	us-east-1b, us-east-1c, us-east-1a	Min: 0, Max: 3	Enabled	⋮
<b>Labels</b>						
workload = memoryIntensive						
<b>Taints</b>						
memorytype = intensive:NoSchedule						
<b>Autoscaling</b>						
Min nodes per zone	Max nodes per zone					
0	1					
▶	post-trade-workers	m5.2xlarge	us-east-1b, us-east-1c, us-east-1a	Min: 3, Max: 15	Enabled	⋮

```
spec:  
  template:  
    spec:  
      serviceAccountName: hal-dl-extract-service-account  
      {{- if .Values.sonarpool.enabled }}  
      nodeSelector:  
        workload: memoryIntensive  
      tolerations:  
        - key: memorytype  
          value: intensive  
          operator: Equal  
          effect: NoSchedule  
      {{- end }}
```



# Journey Ahead



# Customer experience

- ROSA cluster upgrade and maintenance
  - Change log and release notes
  - Simple to follow instructions
- Redhat support portal
  - SLA on four severity level and response time
  - One hour for severity 1
  - Post upgrade support available to remove any obstacles due to human error during an upgrade
  - Excellent know how available to investigate cluster health issue and propose timely remediations
- As an important customer
  - Dedicated account manager, customer executive and architect
  - Shape approach on delivery of key products on topics
  - Dedicate slack channel to establish direct communication with the account manager

# Upcoming Milestones

Developer / Sandbox Environment

KEDA

Load Test on Demand

Blue Green Deployment

Chaos Engineering / FIS

Advanced Cluster Management