

TODO1

Gaining and **EDGE** on Fighting Fraud

Our journey to MACHINE LEARNING with a purpose (and efficiently)

EDGAR OSUNA, Ph. D.

VP DATA & ANALYTICS

- About TODO1
- Use case summary
- Knime + Cloud = Results
- Knime + Cloud + **EDGE** = Better Results
- Challenges and next steps

- Mission to humanize the interaction between customers and their financial institutions through our Digital Banking Multichannel solution
- Born from DNA of leading Latam Banks and operating for 20+ years
- 500+ Employees in Latam and US
- Processing 7.5+ Billion transactions per year for 15+ million customers
- Relatively new to the Cloud and Machine Learning landscapes


- Transactional fraud mitigation in the digital interactions between [financial] institutions and their customers
- Typical digital interactions include Login, password change, personal information update, transfers and all types of payments, etc.
- Profiler leverages on a wide array of machine learning models using customer's behavior, device information, biometrics, fraud patterns and context to provide real-time decision capabilities to fight fraud
- Can run integrated with TODO1's other solutions or standalone
- Fully hosted and maintained by TODO1

- Significant transactional/behavioral/device/biometrical FRAUD data is hard to get
- Feature engineering is tough
- Real-time (< 100ms) response is also hard to manage
- Sheer volume makes it even more complex (i.e. > 350 txns per sec.)
- Down time is not an option
- Scalable to accommodate wide array of institutions

- Running Load balancer, Knime servers / Rabbit MQ / 2x Executors
- System in production in the Cloud and fully operational for the past 18 months managing 4.5+ billion transactions with NO down-time
- A Bank with 8+ million customers, peaks of 350+ transactions per second with overall improvement of +100% in detection rate at the same false positive rate vs market leader.
- Round-trip median time approx. 90ms
- ...but we want to do better (faster, more efficient, scalable, safer)

- Opportunity:

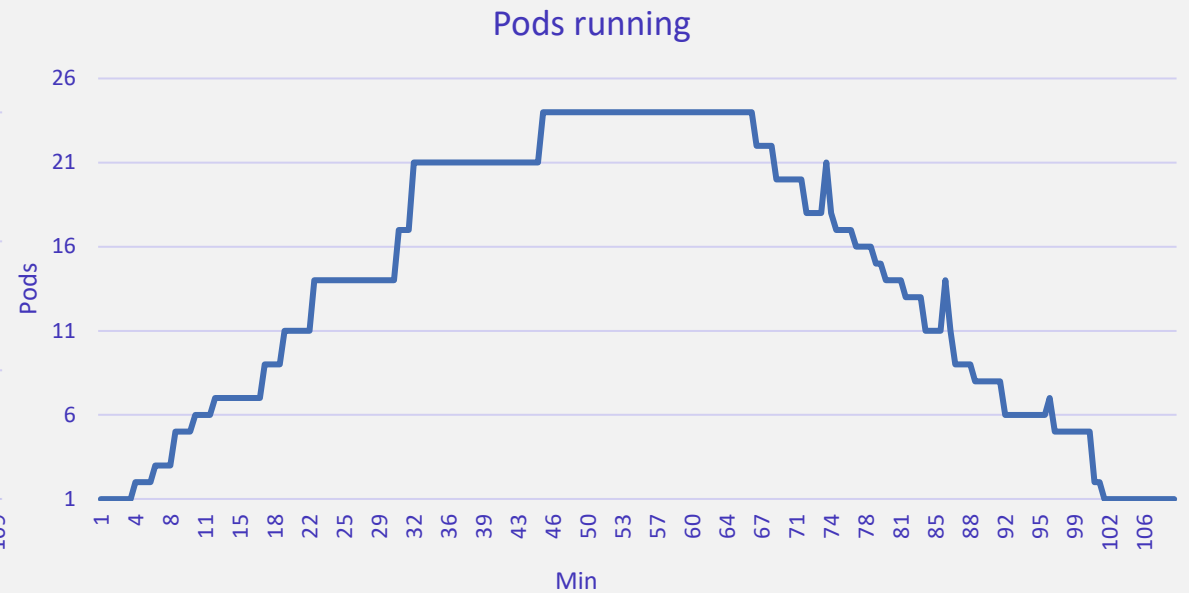
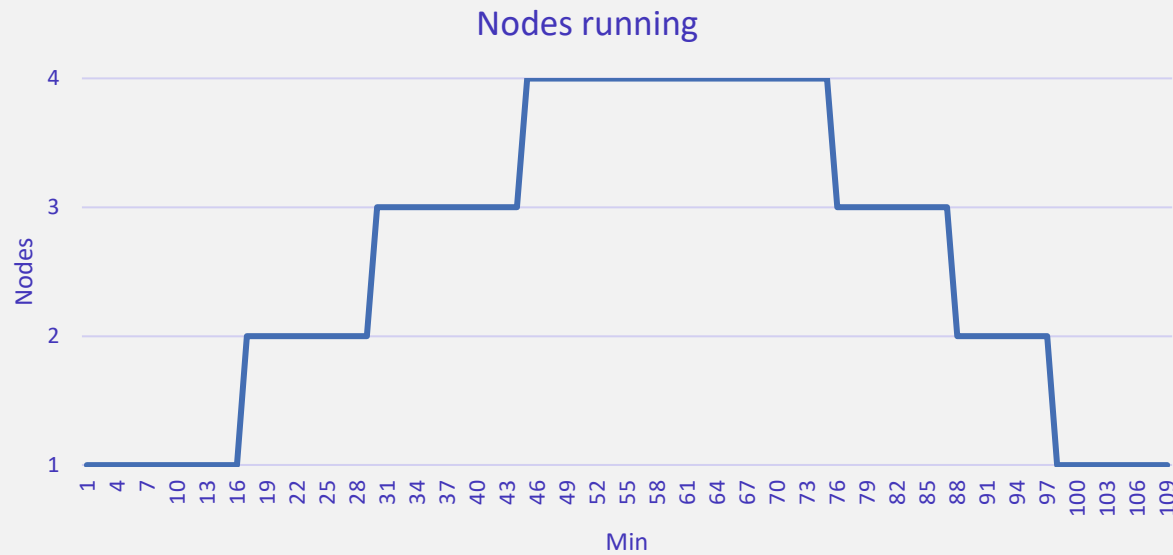
- Transactional demand is variable w/ smooth changes
- Current setup runs over static hardware assignment w/ excess capacity
- Real-time scoring is light and does not need the *full weight* of Knime Server



Illustrative numbers

- What we tested:


- Knime EDGE on an Azure Kubernetes Cluster (AKS) using same VM type but a fraction smaller than production
- Same workflows as production environment with real-life transactions
- Simulated ramp-up and ramp-down demand to benchmark performance



Illustrative numbers

- Results:

- Similar response time
- Same throughput with approx. 50% of the computational resources (vcpu)
- When combined with variable demand we achieve a 75% resource reduction



Illustrative numbers

- Continue benchmarking tests (i.e. how small can the VM be ..?)
- Roll out infrastructure cost optimization
- Accommodate challenger models in production
- Improve run-time performance (proprietary nodes ?)
- New use cases
- Evangelize others in the organization



Thank you!!!



Edgar Osuna

VP – Data & Analytics

Copyright © 2021 TODO1 Services, Inc. All rights reserved. TODO1 and the TODO1 Logo are trademarks or registered trademarks of TODO1 Services Inc. or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This document is provided for informational purposes only and is not intended as advertising. All warranties relating to the information in this document, either express or implied, are disclaimed to the maximum extent allowed by law. The information in this document is subject to change without notice.