BLUEHAT SHANGHAI 2019

User Profile Based Anomaly Detection

Fighting Malicious Insider Threat at Office365

Lei He

Principal Engineering Manager, Microsoft Corporation

2019-05-30

Office 365

The Context - Office365 Security Requirements

Access Control

- Cannot use Corp identity to depoly or access, 2FA must for access
- OCEs need to meet clearance expecation
- Isolation between capacity and control plane
- Zero standing access and least privelege access
- Customer lockbox & BYOK

Security Monitoring

- ETW-based detailed raw telemetry, with ability to instrument new telemetry in < 1 day
- Centralized processing raises alerts in < 15 minutes for security issues
- At least 1 attack/day simulated against service to validate monitoring/response
- Ability to trigger security response workflows in < 15 minutes, and instrument new workflow in < 1 day

Secrets Management

- All secrets must be stored in secure container
- Secrets should never leave boundary
- Secrets should be rolled immediately on exposure
- Data Trustee should have full control of secrets in soverign environment

Anti Virus Patching

- All prod end point scanned every day
- All medium and higher rated vulnerabilitis should be patched or exempted
- All web apps should be scanned once a month
- Central reporting and tracking

*Customer lockbox onboarding is must only for enterprise service accessing customer data.

Overview

Problem

- Compromised/Disgruntled Malicious Insiders Remain a Top Threat
- Supervised ML Detections Trained to Detect “Known” Patterns; Need Coverage for Unknown Patterns

Solution

- User Profile Based Detection to Flag Abnormal Activities
- For O365 Data Center Security, User == DevOps
- Can be Generalized to Other Entity Profiles

O365 Data Center Monitoring

Automated Near Real Time Multi Tier Detection and Response

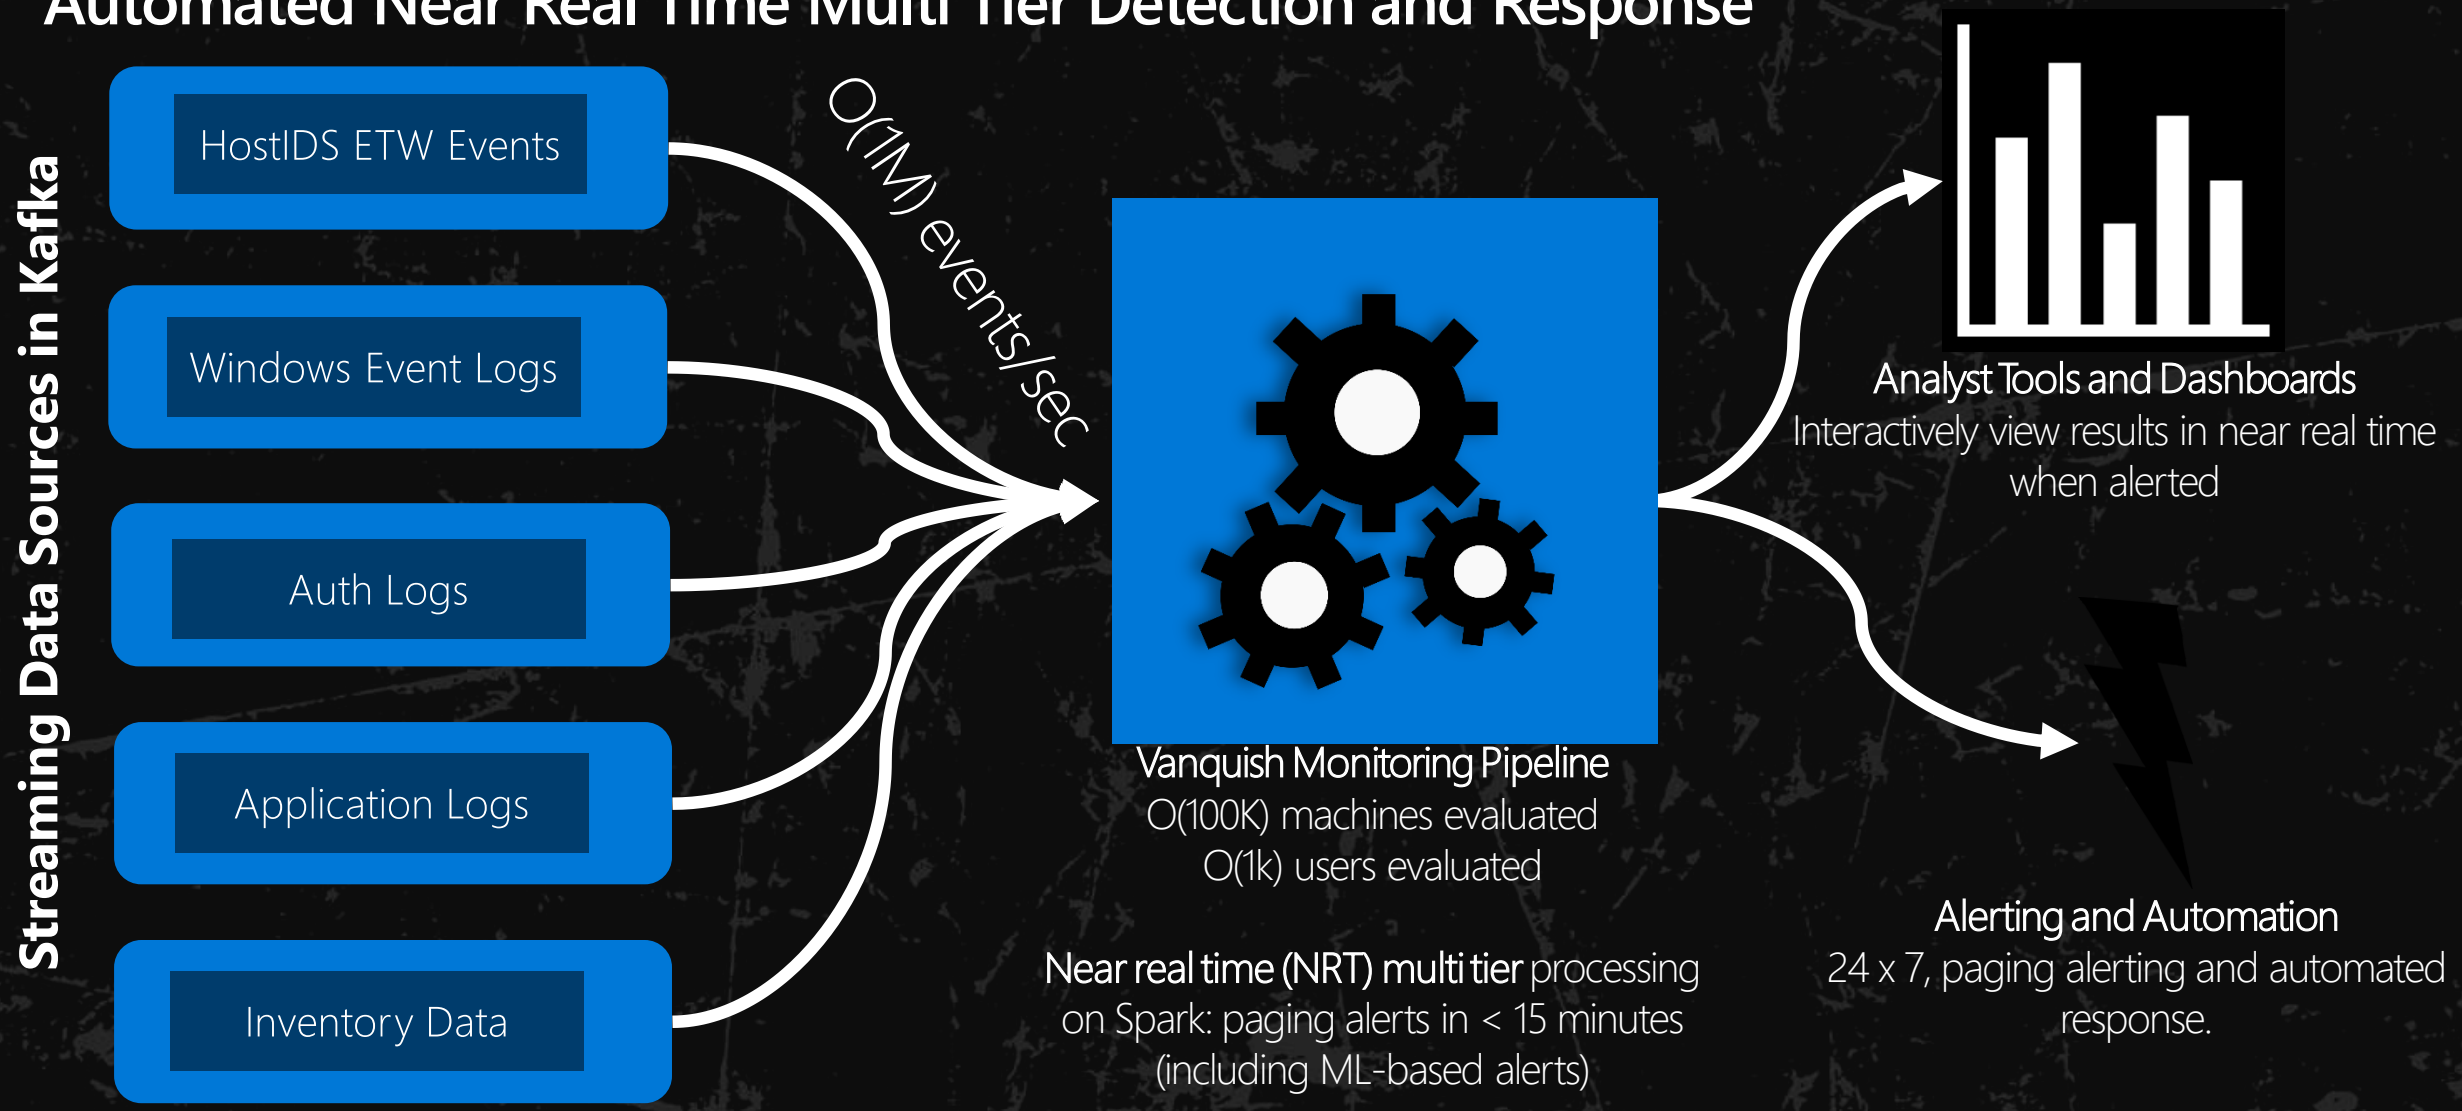

Model Feature Selection

Observation: DevOps Activities Exhibit Patterns

Logon Activities:

- Where did the user login from?
- When did the user login?
- What did the user login to?

Just In Time Elevations:

- What Elevation Role did the user request?
- Who approved the elevation request?

Actions:

- Which workflow did the user perform?
- Which processes did the user start?
- Which detections did the user trigger?

Features:

- Machines logged on to
- Machine Role
- Capacity Unit Scope
- IPs logged on from
- IPs connected to
- Other Atomic Detections Triggered
- Roles Elevated to
- Elevation Purpose
- Approver
- Time of Day
- Process and Workflow Ran

The Model

1. **Compute a complete Historical Profile for Every User**
   a) Each Feature Is a List of { Value: Occurrence }
2. **Every 2 minutes, activities of the past 6-hour long user sessions are evaluated**
3. **Each Current Session is compared with the user's Historical Profile**
   a) A feature **similarity score** is calculated for the value(s) in the session. **The higher the historical occurrence, the more similar/normal**.

$$Sim_i = MIN(Phis_0, Phis_1, Phis_j, \dots, Phis_k)$$

   b) A session **Anomaly Score** is generated from the similarity score of the n features.

$$\text{Anomaly Score} = 1 - \text{AVG}(Sim_0, Sim_1, Sim_2, \dots, Sim_n)$$

4. **A Test Dataset is curated from historical positive events to evaluate performance**
5. **A Session Score threshold is set for alerting by Alert Tolerance (5 per week)**

Model Tuning

Model Tuning Workflow

Malicious != Anomalous

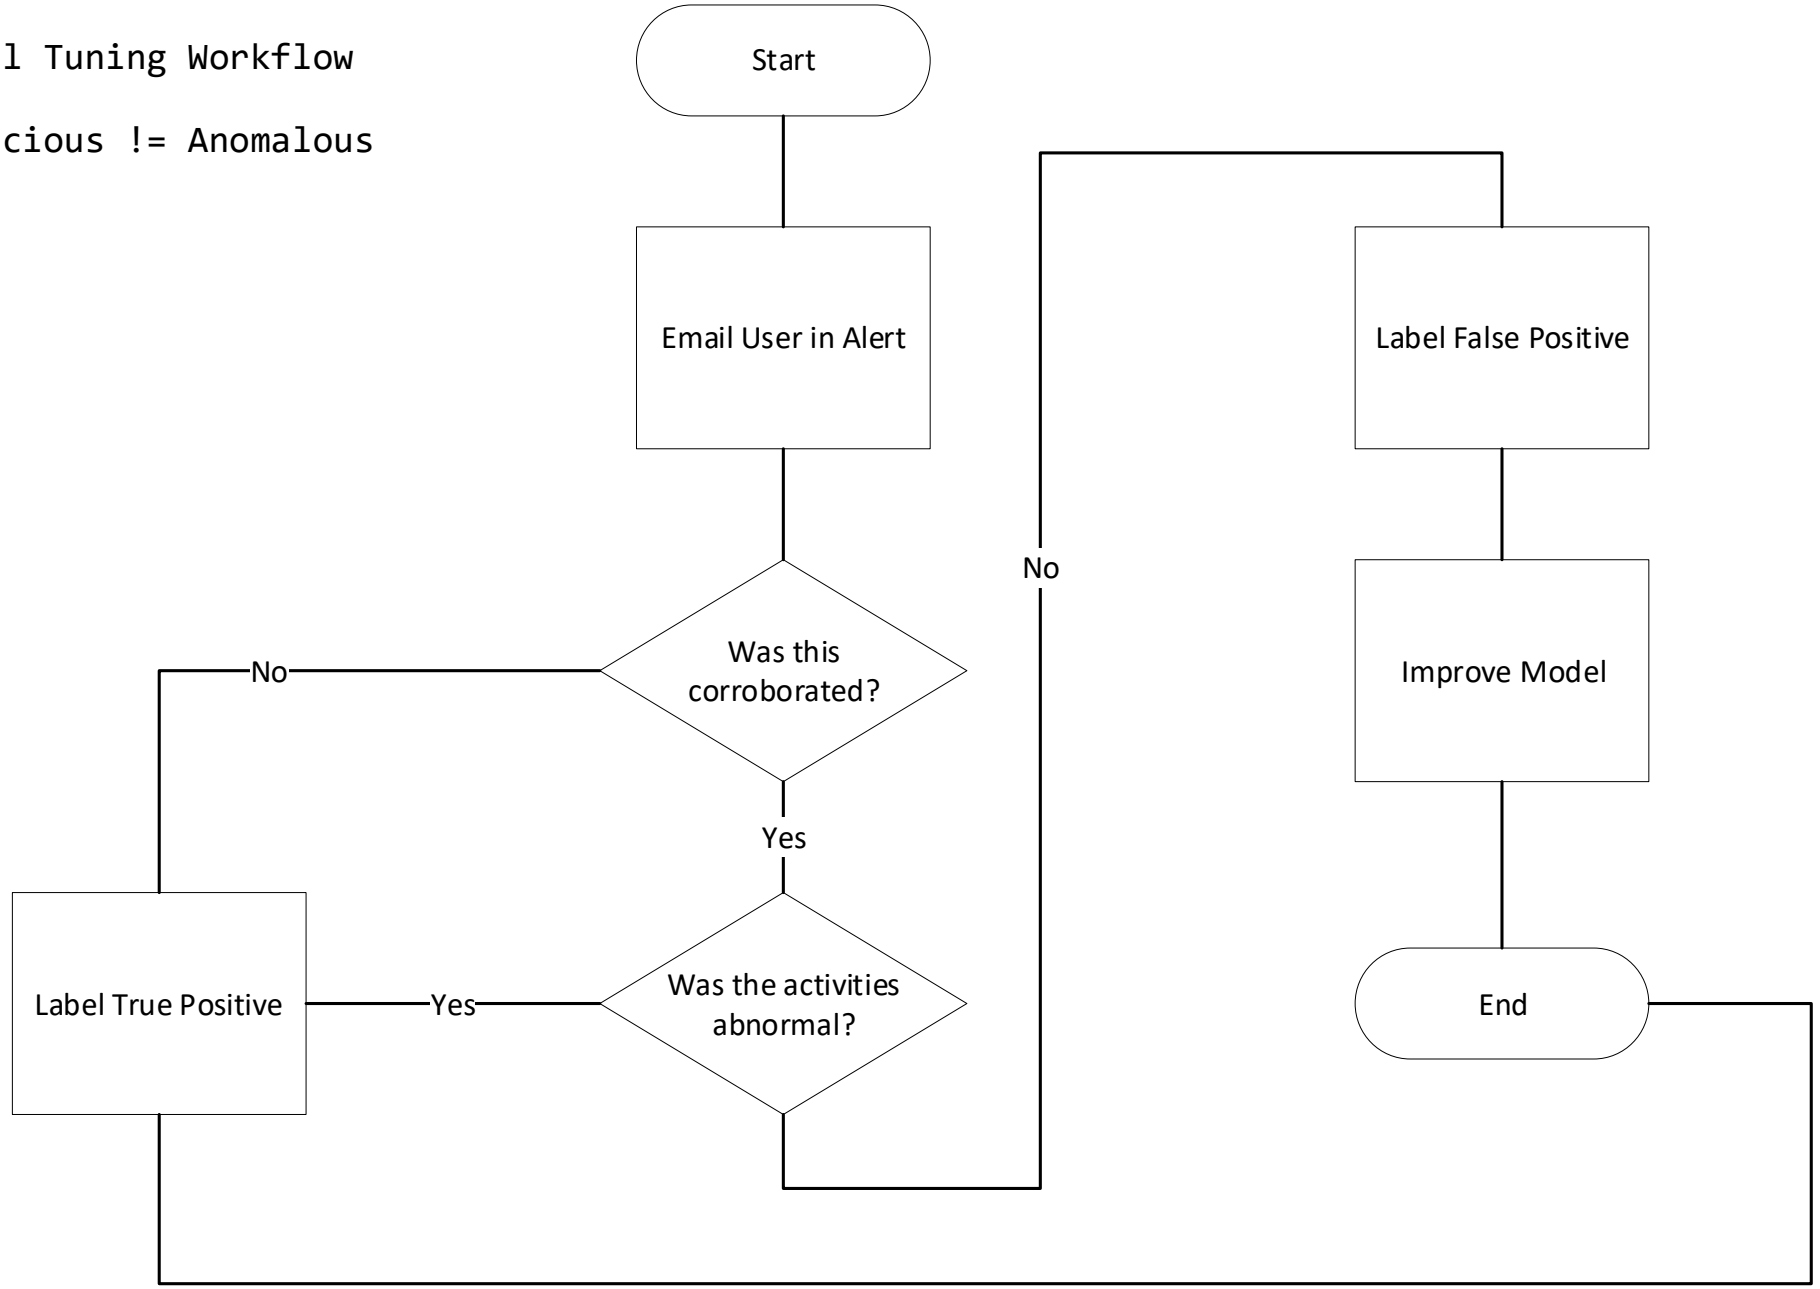

Model Tuning Patterns

The New Team Member:

Set minimum threshold for volume of activities

Sparse Features:

Set minimum threshold of valid features

Double Counting:

Remove auto generated activities

Changing Teams:

Use team profile

Renamed Commands:

Case insensitive comparison

High Entropy Features:

Replace target machine name or IP by {Capacity Unit, Capacity Type}

Downgrade the weight of user source IP

Downgrade free text features (elevation reason)

Model Performance

Success Metric: Paging Volume

- Paging alerts average ~5 per week since shipped late 2018

Success Metric: Precision

- ~90% true positives (Last Month 16/18)

Types of True Positives

- Penetration Testing
- One Off Incident (“Break Glass”) or Audit
- Operation Process Violation
- DevOps Poor Practice
- DevOps One off Testing or Experimenting

Call To Action

Apply Profile Based Anomaly Detection for Defense in Depth!

- Start with Quality Data
- Identify True Positives
- Offline Analysis before Online Implementation
- Continuous Tuning through Feedback Loop
- Explainable Model
- Actional Alerts

Q & A

BLUEHAT
SHANGHAI 2019

User Profile Based Anomaly Detection

Lei He

leih@Microsoft.com

www.linkedin.com/in/lei-he-security