https://dx.doi.org/10.14722/prism.2026.23008

NOD: Uncovering intense attackers' behavior through Nested Outlier Detection from SSH logs

The University of Utah

Authors:

Ghazal Abdollahi,
Hamid Asadi,
Robert Ricci

PRISM 2026,
NDSS

Intense attackers!

Photo credit: https://gcore.com/learning/how-to-use-ssh-secure-remote-access-ubuntu

- **Statistically extreme behaviors**
  - Volume
  - Reach
  - credential diversity

Why they matter?

- **Alert Fatigue Mitigation**
  - population-level reasoning
  - Strategic Prioritization
    - Coordinated Pattern Recognition
      - Upstream Defense Shifting

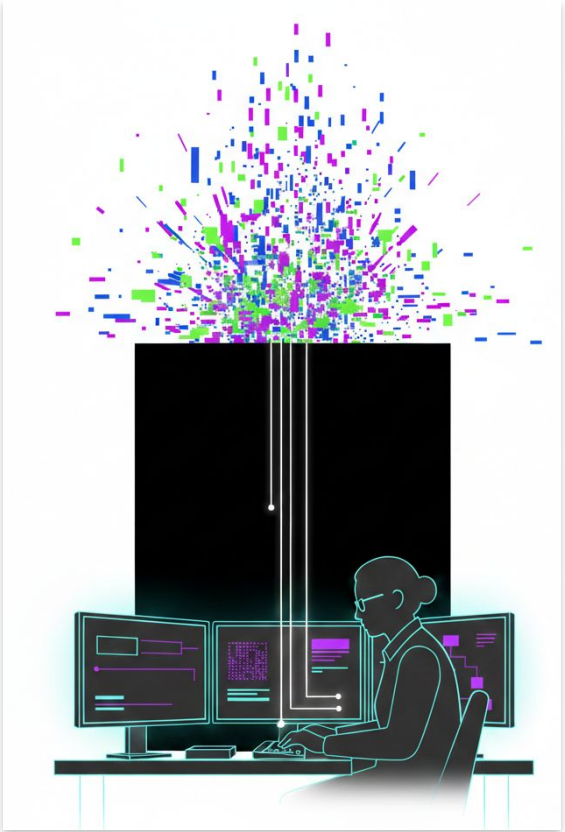

Photo credit: Google Gemini

Structural Imbalance in SSH Attacks

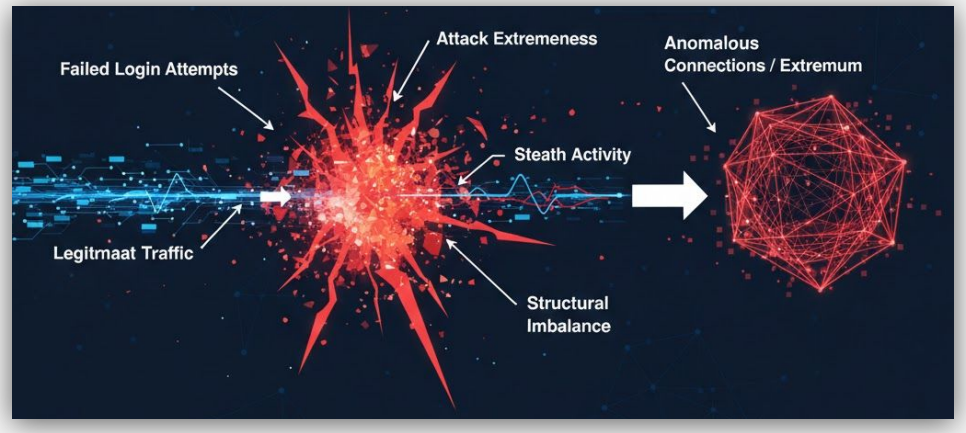

Photo credit: Google Gemini

- **Structural extremeness**
  - may persist or intensify over time?
- **Intensity shifts**
  - can reveal coordinated behavior?
- **Static snapshots**
  - cannot capture dynamic strategy?
- **Time exposes**
  - extremeness is episodic or organized?

Dataset

- 89,090,860 SSH authentication events
- 68 days of collection
- Aggregated into 77,886 source IP behavioral profiles
- Real-world brute-force telemetry

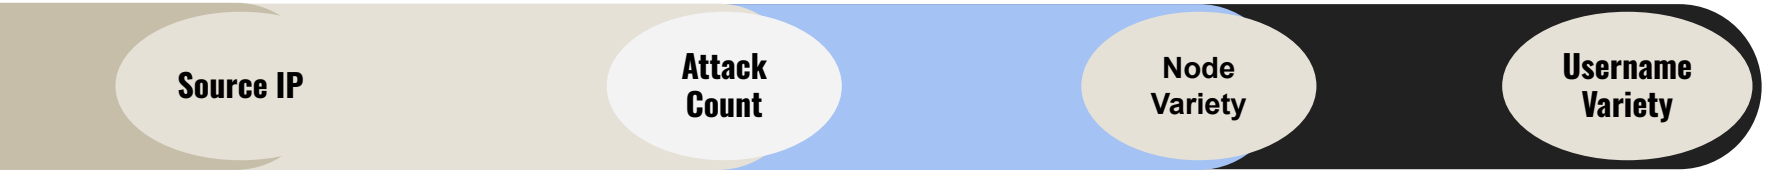

Source IP	Attack Count	Node Variety	Username Variety
attacker entity	Total login attempts per IP	Number of distinct destination nodes leveraged per IP	Number of distinct usernames leveraged per IP

NOD design summary

IP-aggregation
89 M logs → 77K

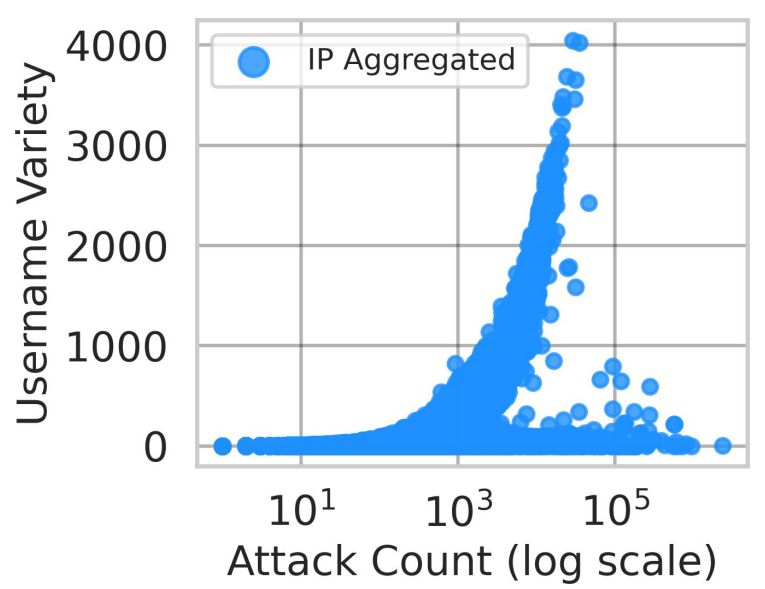

Tier 1

- Identifying Initial outliers

Tier 2

- Identifying Outliers of outliers

Tier 2 phenotypes

- Campaign identification

Mahalanobis Distance

- how far a point is from a distribution
- accounting for correlations between variables
- widely used for detecting multivariate outliers

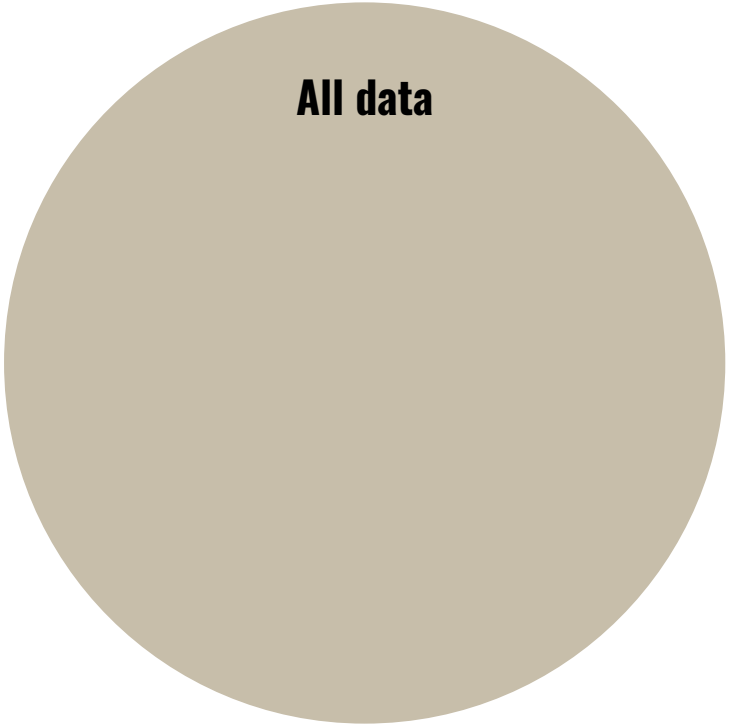

Mahalanobis Distance, T1

- how far a point is from a distribution
- accounting for correlations between variables
- widely used for detecting multivariate outliers

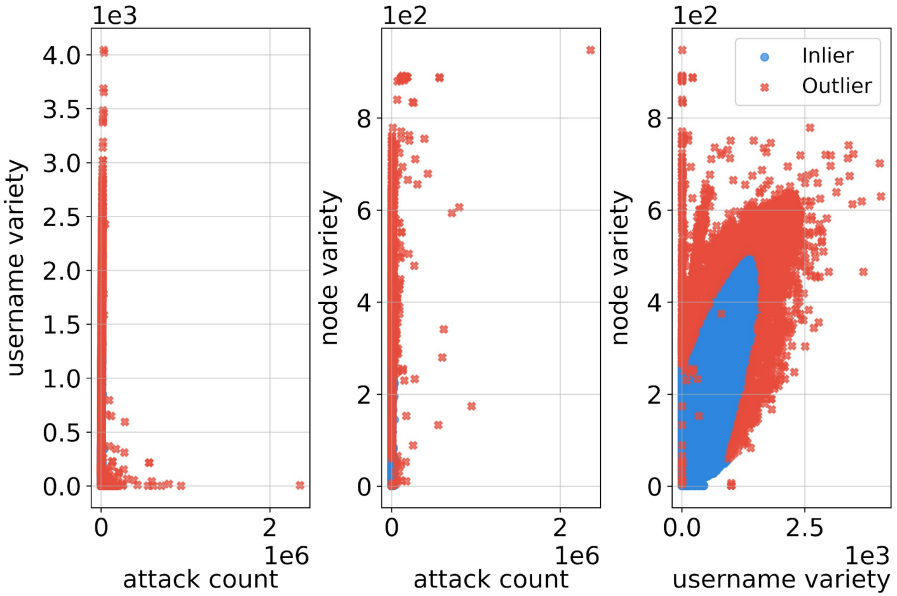

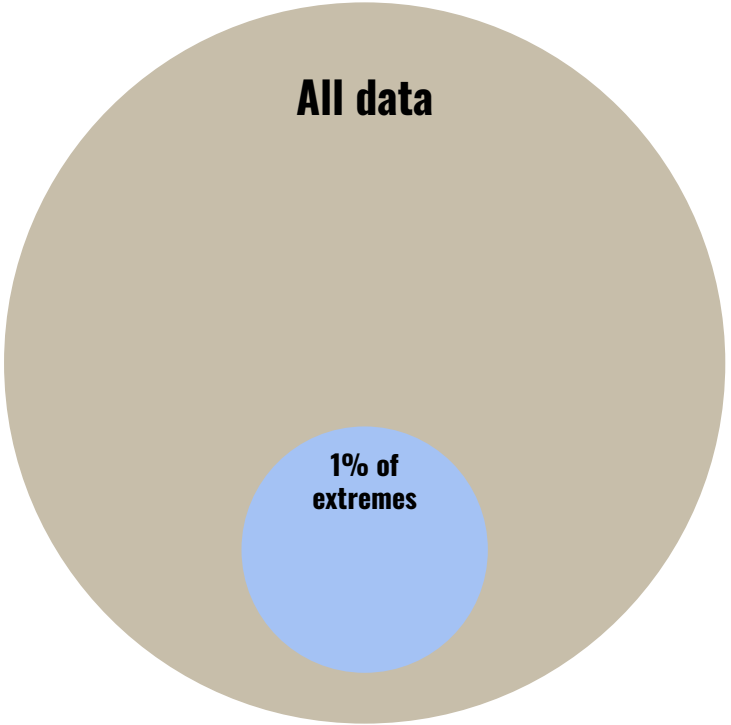

Mahalanobis Distance, T2

- how far a point is from a distribution
- accounting for correlations between variables
- widely used for detecting multivariate outliers

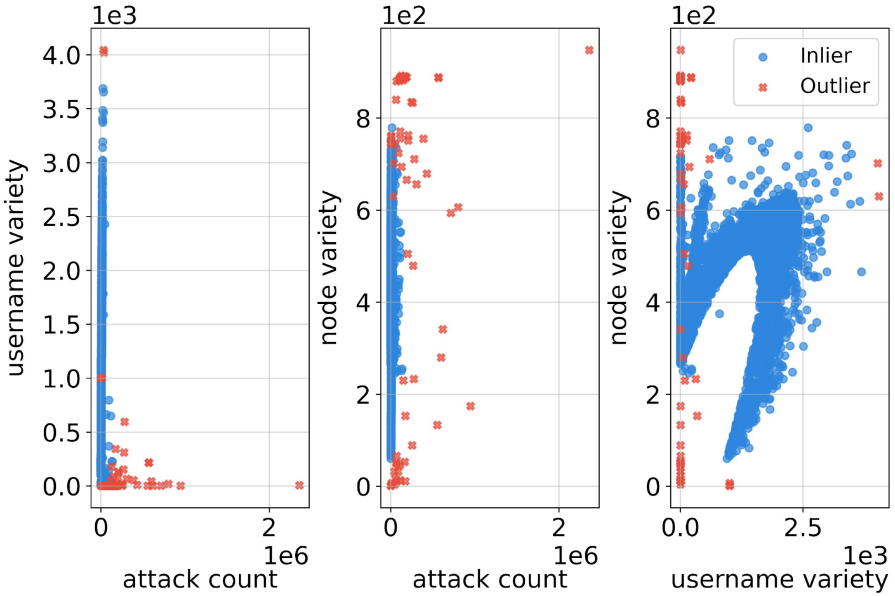

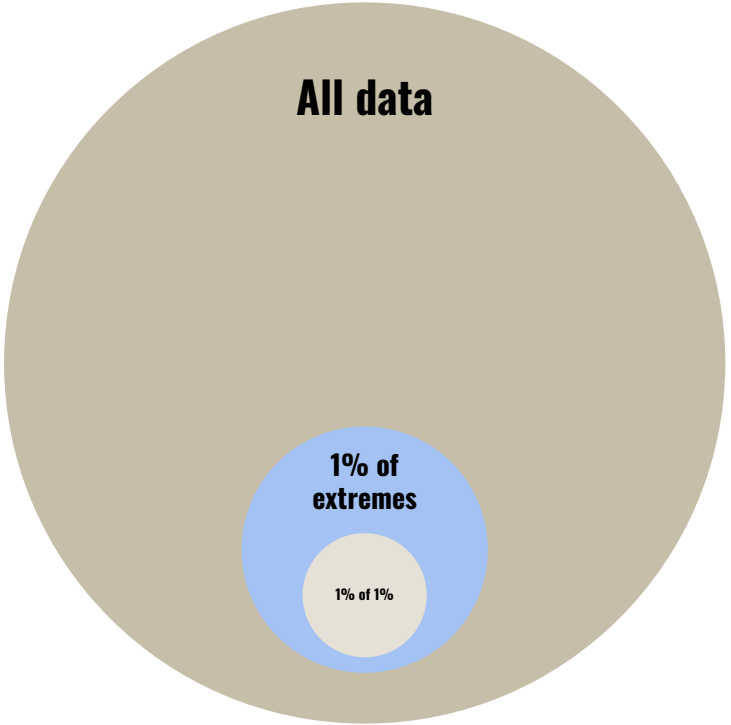

What we learnt from Tier 1 ?

Feature	Total Dataset	Tier-2	Tier-1 %
Attack count	89,090,860	15,974,515	**67**
Username Variety	12,386,995	14,846	**58.1**
Node Variety	5,043,482	40,059	**48.4**
Number of IPs	77,886	72	**6.7**

- Small minority, major impact!
- Broad operational reach!
- Drives credential entropy!
- Primary investigative layer!

What we learnt from Tier 2 ?

- Extremely small group!
- Disproportionate impact!
- Behaviorally extreme!
- Strategic signal layer!

Feature	Total Dataset	Tier-2	Tier-1 %	Tier-2 %
Attack count	89,090,860	15,974,515	**67**	**17.9**
Username Variety	12,386,995	14,846	**58.1**	**0.1**
Node Variety	5,043,482	40,059	**48.4**	**0.8**
Number of IPs	77,886	72	**6.7**	**0.09**

Our identified phenotypes (Tier 2)

- **High volume: 386 K**
- **high reach: 572**
- **and high credential diversity: 588**

1 Extreme Anomalies

Our identified phenotypes (Tier 2)

- High volume: 386 K
- high reach: 572
- and high credential diversity: 588

- **Very broad node reach: 752**
- **with relatively low credential diversity: 26**
- **sharing infrastructure context (ipinfo.io)**

1 Extreme Anomalies

2 Subnet & Infrastructure Clusters

Our identified phenotypes (Tier 2)

- High volume: 386 K
- high reach: 572
- and high credential diversity: 588

- Very broad node reach: 752
- with relatively low credential diversity: 26
- sharing infrastructure context (ipinfo.io)

- **One credential sprayed**
  - **across many target nodes: 751**

1. Extreme Anomalies
2. Subnet & Infrastructure Clusters
3. Single-Username Swarms

Our identified phenotypes (Tier 2)

- High volume: 386 K
- high reach: 572
- and high credential diversity: 588

- Very broad node reach: 752
- with relatively low credential diversity: 26
- Sharing infrastructure context (ipinfo.io)

- One credential sprayed
  - across many target nodes: 751

- **Attempt large username sets : 1001 (Jaccard similarity)**
  - **against very few targets: 3**

Correlation of outliers

- Warm for attack counts
- More mixed for node variety
- Most variable for username variety

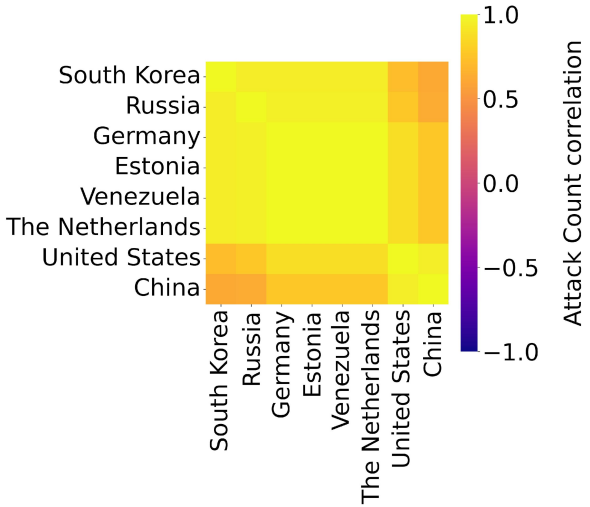

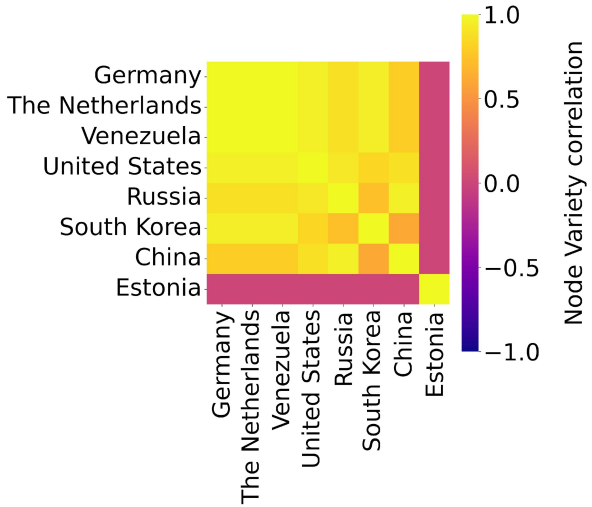

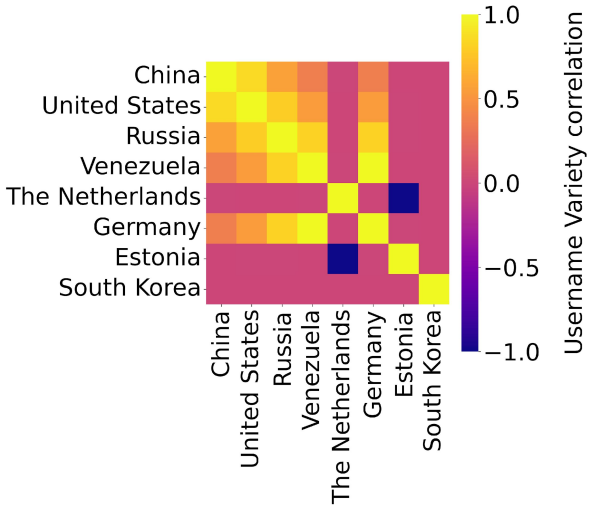

Pattern through time

- **Spikes ≠ new attackers**

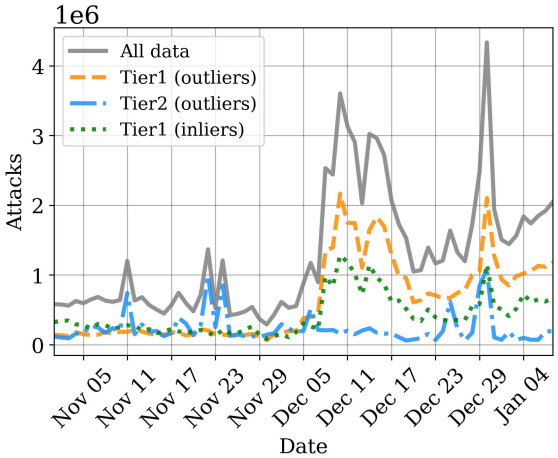

Pattern through time

- Spikes ≠ new attackers
- **Spikes = more usernames**

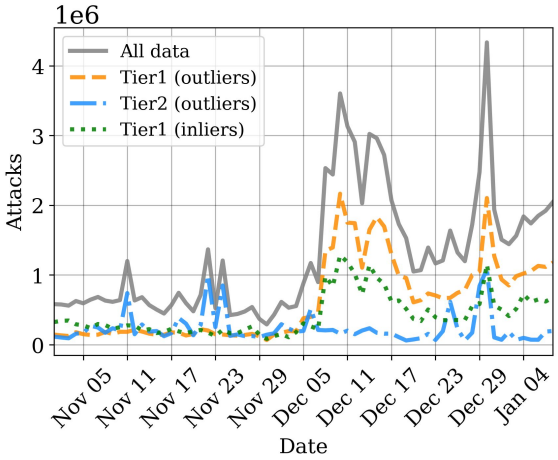

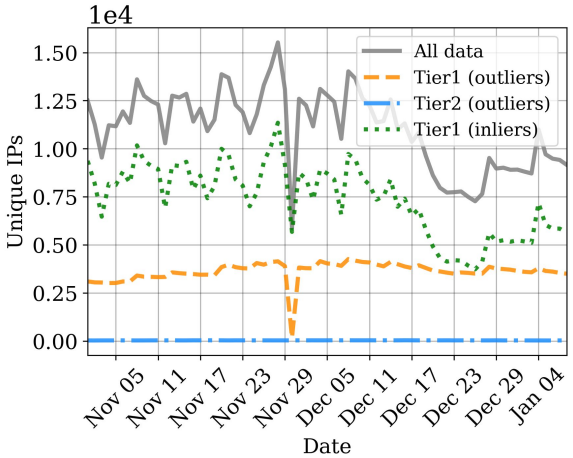

Pattern through time

- Spikes ≠ new attackers
- Spikes = more usernames
- **IP count more stable, usernames change!**

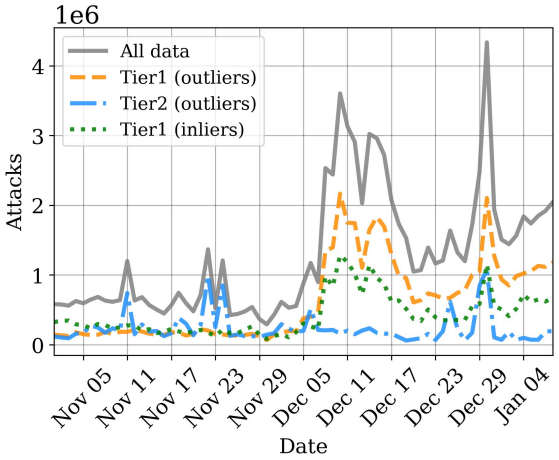

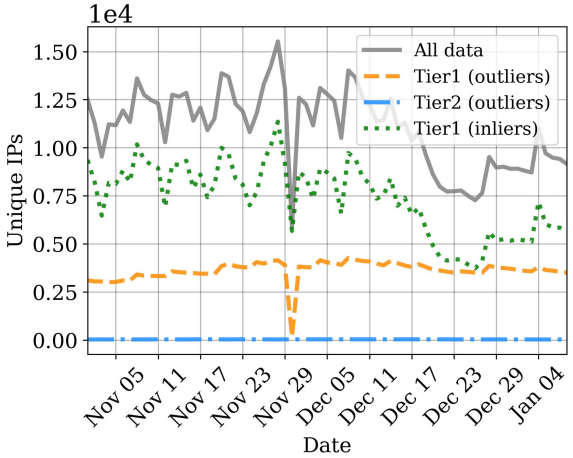

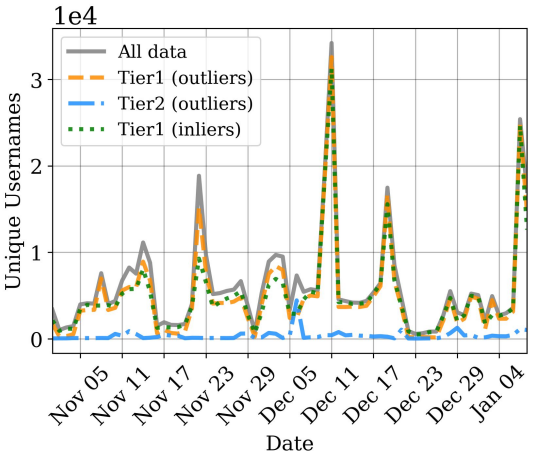

Pattern through time

- Spikes ≠ new attackers
- Spikes = more usernames
- **IP count more stable, usernames change!**

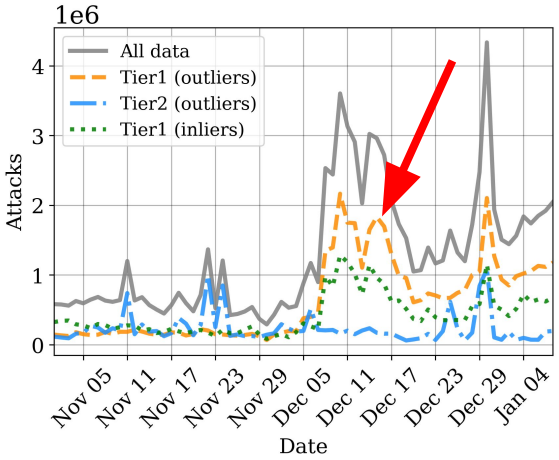

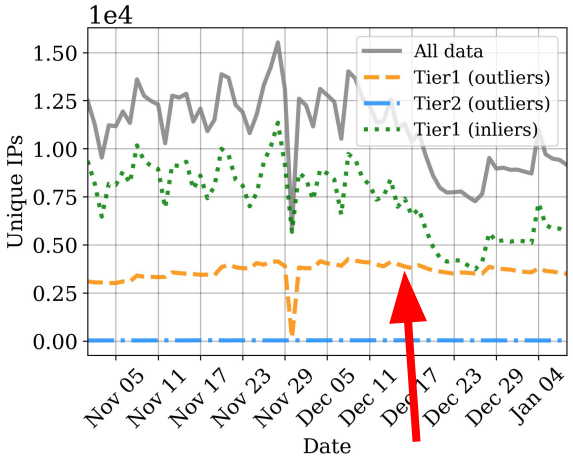

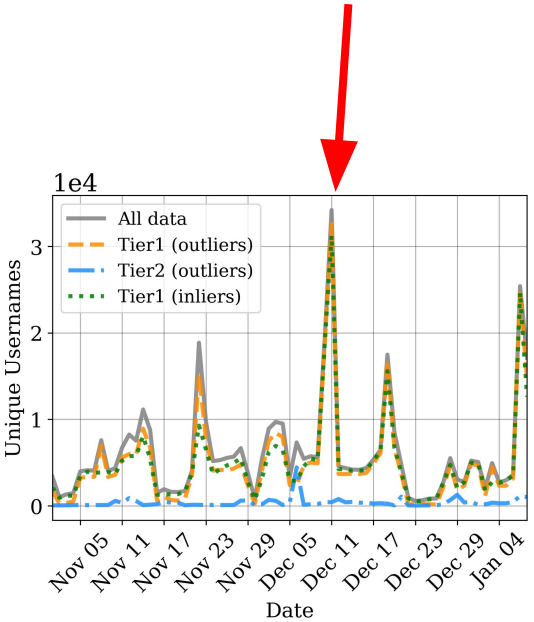

Conclusions

NOD transforms overwhelming SSH noise into a
strategic map of the most extreme and coordinated threat behaviors.

- **89 million events**
  - **→ attacker-level behavioral profiles**
  - **→ nested statistical extremeness**
  - **→ interpretable phenotypes**
- **Graded Triage**
- **Upstream Investigation**

Thank You!

Happy to talk more...!
My portfolio & contact :