Data-Driven Threat Intelligence: Metrics on Indicator Dissemination and Sharing (#ddti)

Alex Pinto
Chief Data Scientist
MLSec Project
@alexcpsec
@MLSecProject

Alexandre Sieira
CTO
Niddel
@AlexandreSieira
@NiddelCorp

Agenda

- Previously on #ddti
- Challenges at TI Sharing
- Measuring TI Sharing
- The Future of Sharing

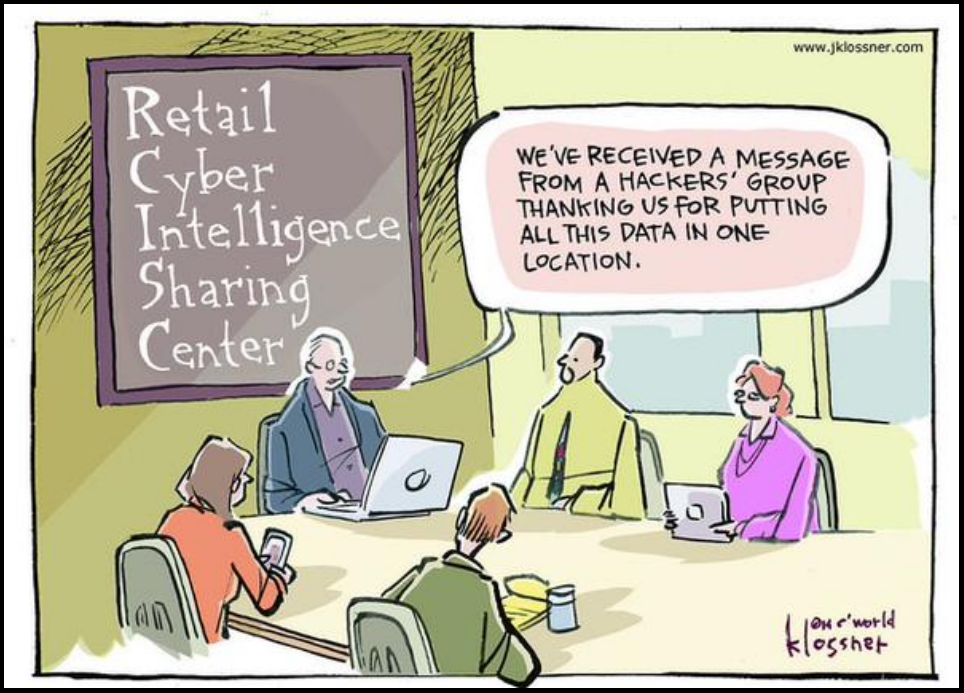

This is a data-driven webinar!

Please check your anecdotes at the door

Previously on #ddti

- Useful Methods and Measurements for Handling Indicators
  - Analysis of Threat Intelligence Feeds
  - Indirectly, a methodology for analyzing TI Providers

- Combine (https://github.com/mlsecproject/combine)
  - Gathers TI data (ip/host) from Internet and local files

- TIQ-Test (https://github.com/mlsecproject/tiq-test)
  - Runs statistical summaries and tests on TI feeds

TIQ-TEST - Tons of Threat-y Tests

Putting this threat intel data to work

- ~~NOVELTY – How often do the feeds update themselves?~~
- ~~AGING – How long does an indicator sit on a feed?~~
- ~~POPULATION – How does this population distribution compare to my data?~~
- OVERLAP – How do the indicators compare to the ones you got?
- UNIQUENESS – How many indicators are found only on one feed?

Overlap Test

More data is fine, but make sure it is different

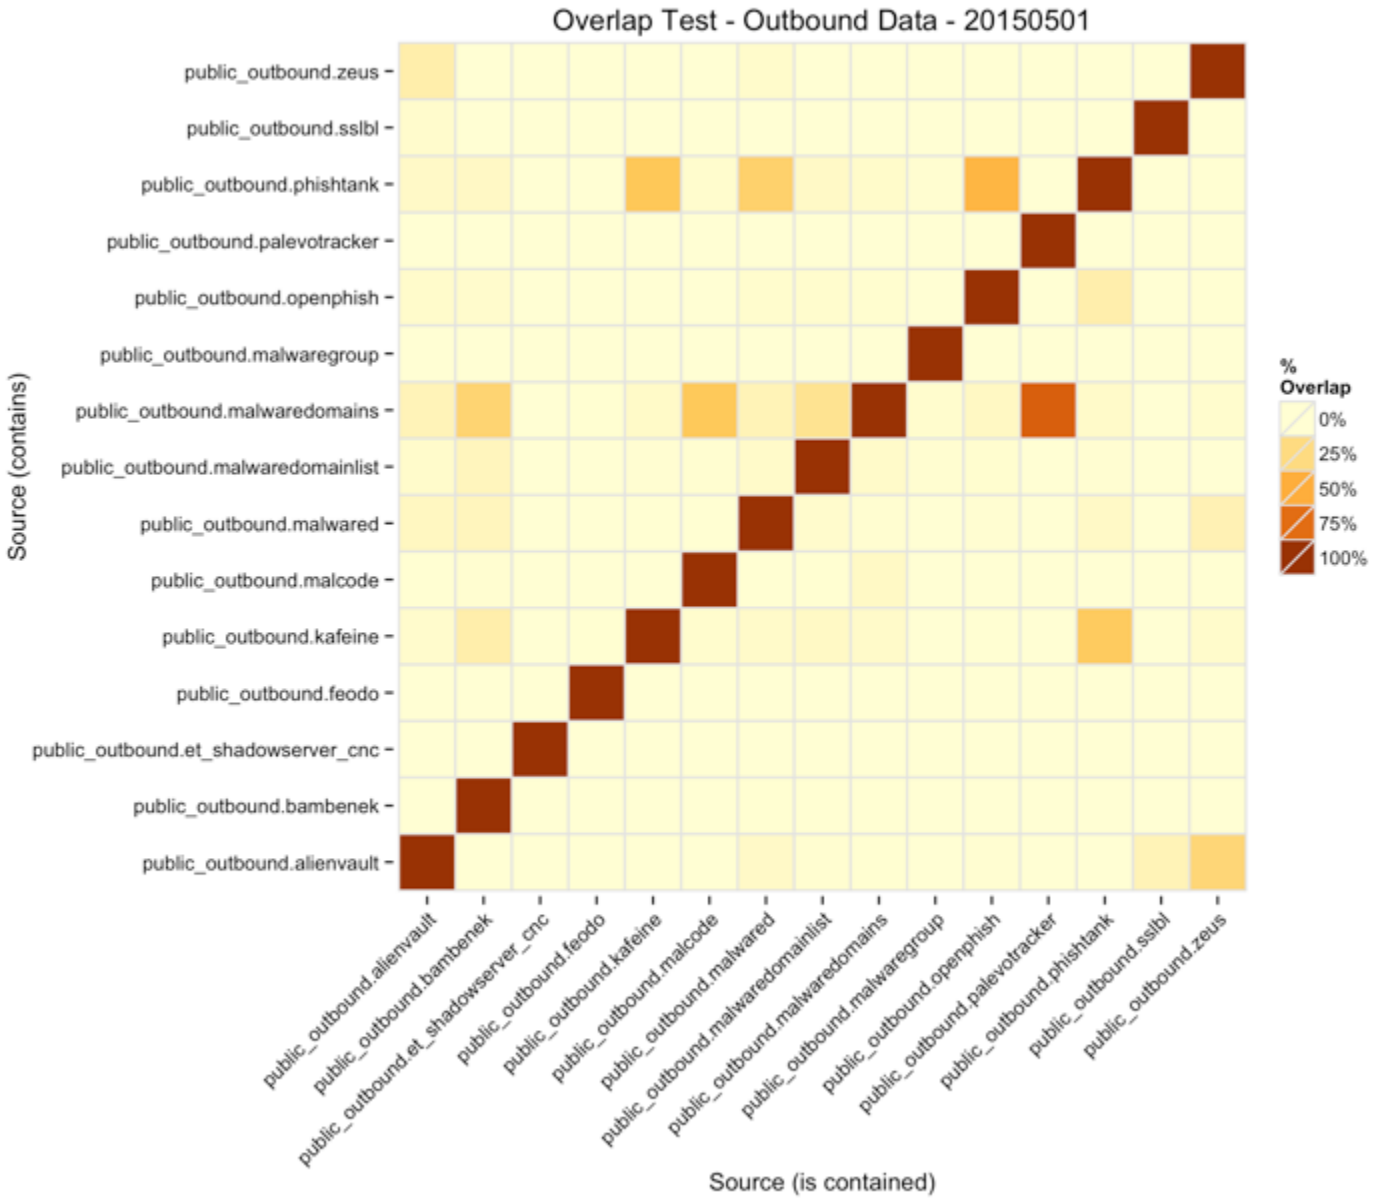
Overlap Test - Outbound Data - 20150501
Source (contains)
public_outbound.zeus
public_outbound.sslbl
public_outbound.phishtank
public_outbound.palevotracker
public_outbound.openphish
public_outbound.malwaregroup
public_outbound.malwaredomains
public_outbound.malwaredomainlist
public_outbound.malwared
public_outbound.malcode
public_outbound.kafeine
public_outbound.feodo
public_outbound.et_shadowserver_cnc
public_outbound.bambenek
public_outbound.alienvault
Source (is contained)
% Overlap
0%
25%
50%
75%
100%

Uniqueness Test

Can we tell if we are close to finding *all* the threats?

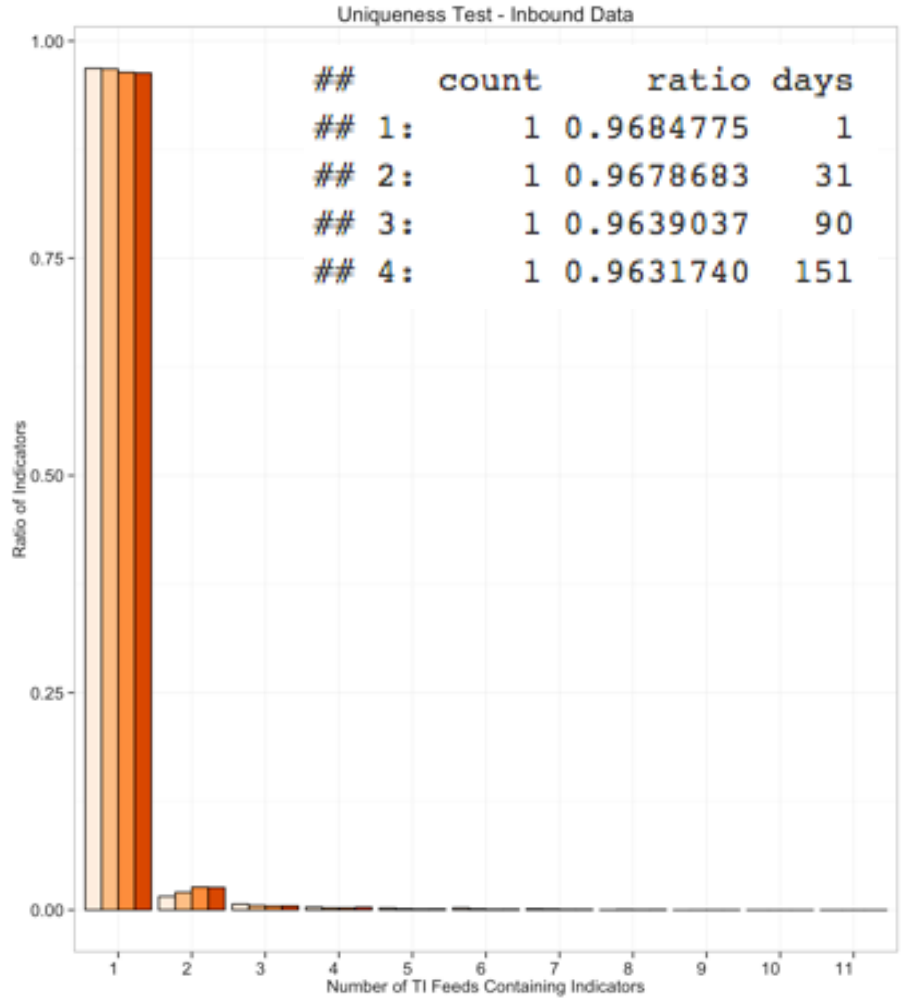
Uniqueness Test - Inbound Data
count ratio days
1: 1 0.9684775 1
2: 1 0.9678683 31
3: 1 0.9639037 90
4: 1 0.9631740 151
Ratio of Indicators
1.00
0.75
0.50
0.25
0.00
1 2 3 4 5 6 7 8 9 10 11
Number of TI Feeds Containing Indicators

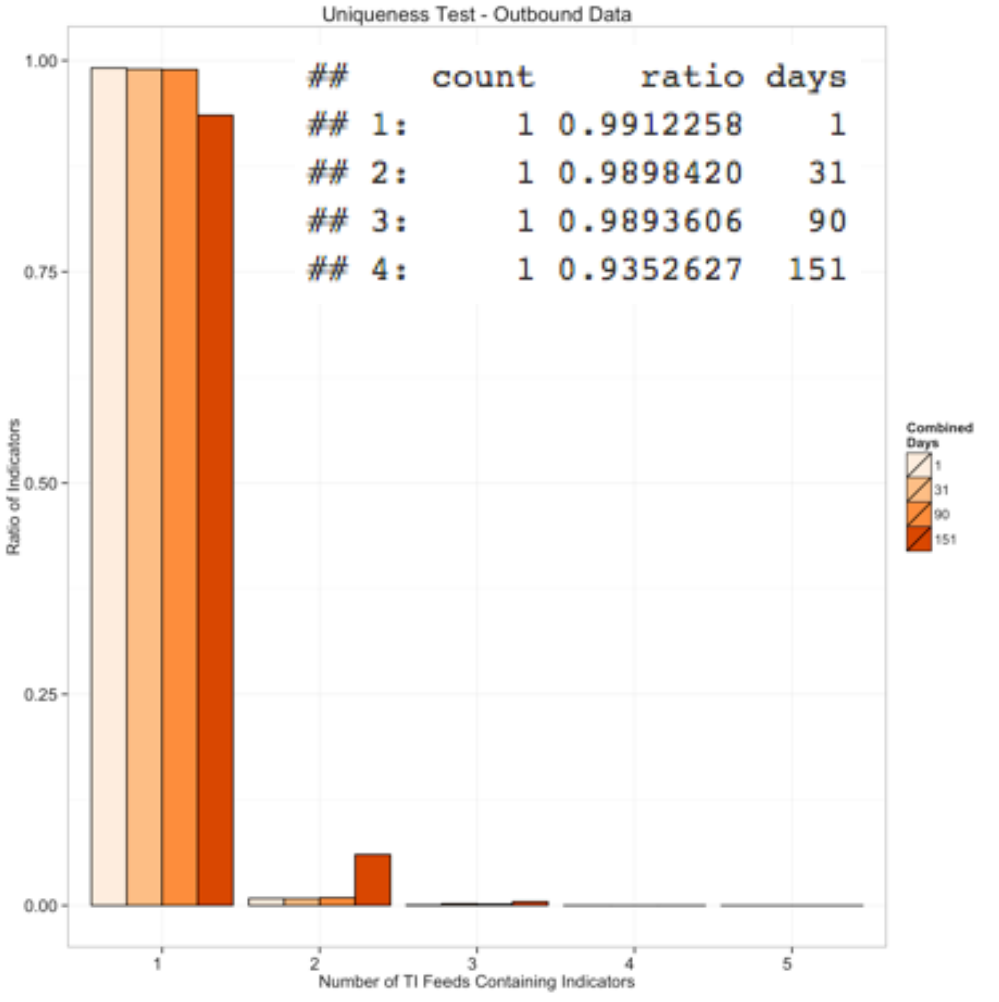
Uniqueness Test - Outbound Data
count ratio days
1: 1 0.9912258 1
2: 1 0.9898420 31
3: 1 0.9893606 90
4: 1 0.9352627 151
Ratio of Indicators
1.00
0.75
0.50
0.25
0.00
1 2 3 4 5
Number of TI Feeds Containing Indicators
Combined Days
1
31
90
151

I hate quoting myself, but…

2015 DATA BREACH INVESTIGATIONS REPORT

It is hard to draw a positive conclusion from these metrics, and it seems to suggest that if threat intelligence indicators were really able to help an enterprise defense strategy, one would need to have access to all of the feeds from all of the providers to be able to get the "best" possible coverage. This would be a Herculean task for any organization, and given the results of our analysis, the result would still be incomplete intelligence. There is a need for companies to be able to apply their threat intelligence to their environment in smarter ways so that even if we cannot see inside the whole lake, we can forecast which parts of it are more likely to have a lot of fish we still haven't caught.

MORE != BETTER

Threat Intelligence
Indicator Feeds

Threat Intelligence
Program

Constructive Feedback from the Internet:

“TI Sharing is TOTALLY going to solve this”

Right, folks? Right?

TI Sharing Solution Plan:

Or at least a rough straw man

1. The best Threat Intelligence is the one that you analyze from your own incidents (homegrown / organic intelligence)

2. There is strength in numbers – vertical herd immunity!

3. ????????

4. PROFIT!! (or at least SECURITY!!)

Issue 1 - BYOTI

451 Research

Spotlight

Threat intelligence: only for the 1%?

Analyst: Scott Crawford 1 Jul, 2015

Threat intelligence has become a booming area of information security, and with good reason. Attackers have the luxury of exploiting whichever weaknesses in a target best serve their intent. Defenders, on the other hand, must make the most of limited resources to defend all the most vulnerable aspects of critical information assets. Understanding the nature of current threats and adversary intent is essential to knowing how and where to place the most effective bets on defense.

If CONSUMING is for the 1%, what is the percentage of organizations able to PRODUCE?

Issue 2 - Herd Immunity

- We may be able to detect more ”virus strains” together but we are *terrible* at inoculation.
- The things we detect the most mutate too fast (Pyramid of Pain)
- Who didn’t get immunized, still gets sick (FOMO-TI)

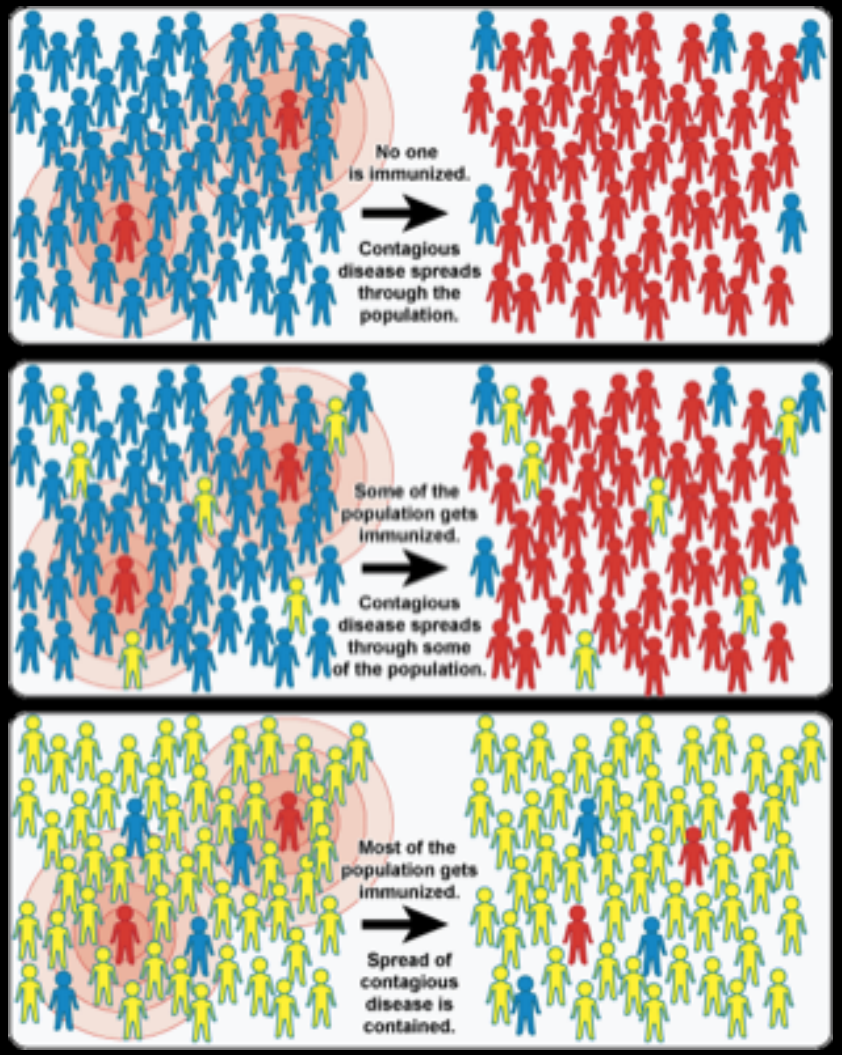

Source: www.vaccines.gov

The Cognitive Dissonances of TI Sharing

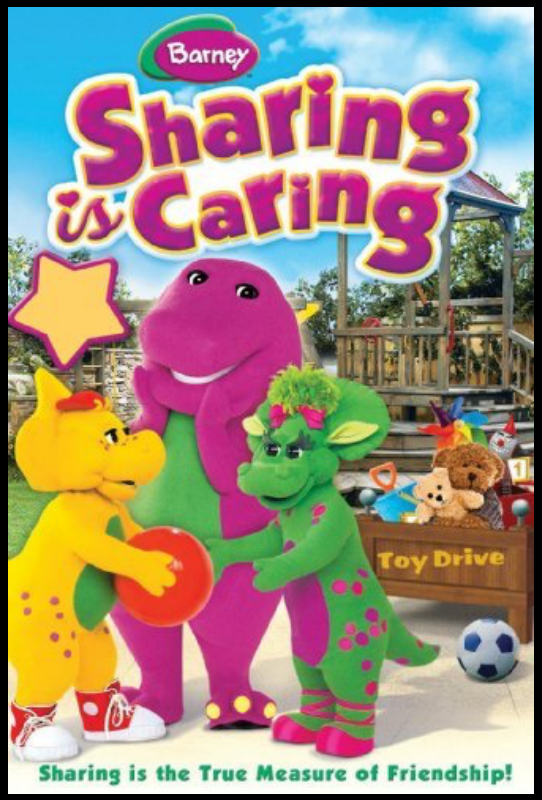

Everybody should share!

VS

The CIRCLE OF TRUST

The Two Sides of the Trust Coin

TRUST FALL

Do you trust the group enough to share?

Do you trust the group enough to consume?

How about measuring it?

We would like to thank the kind contribution of data from the fine folks at Facebook ThreatExchange and ThreatConnect

… and also the sharing communities that chose to remain anonymous. You know who you are, and we ❤ you too.

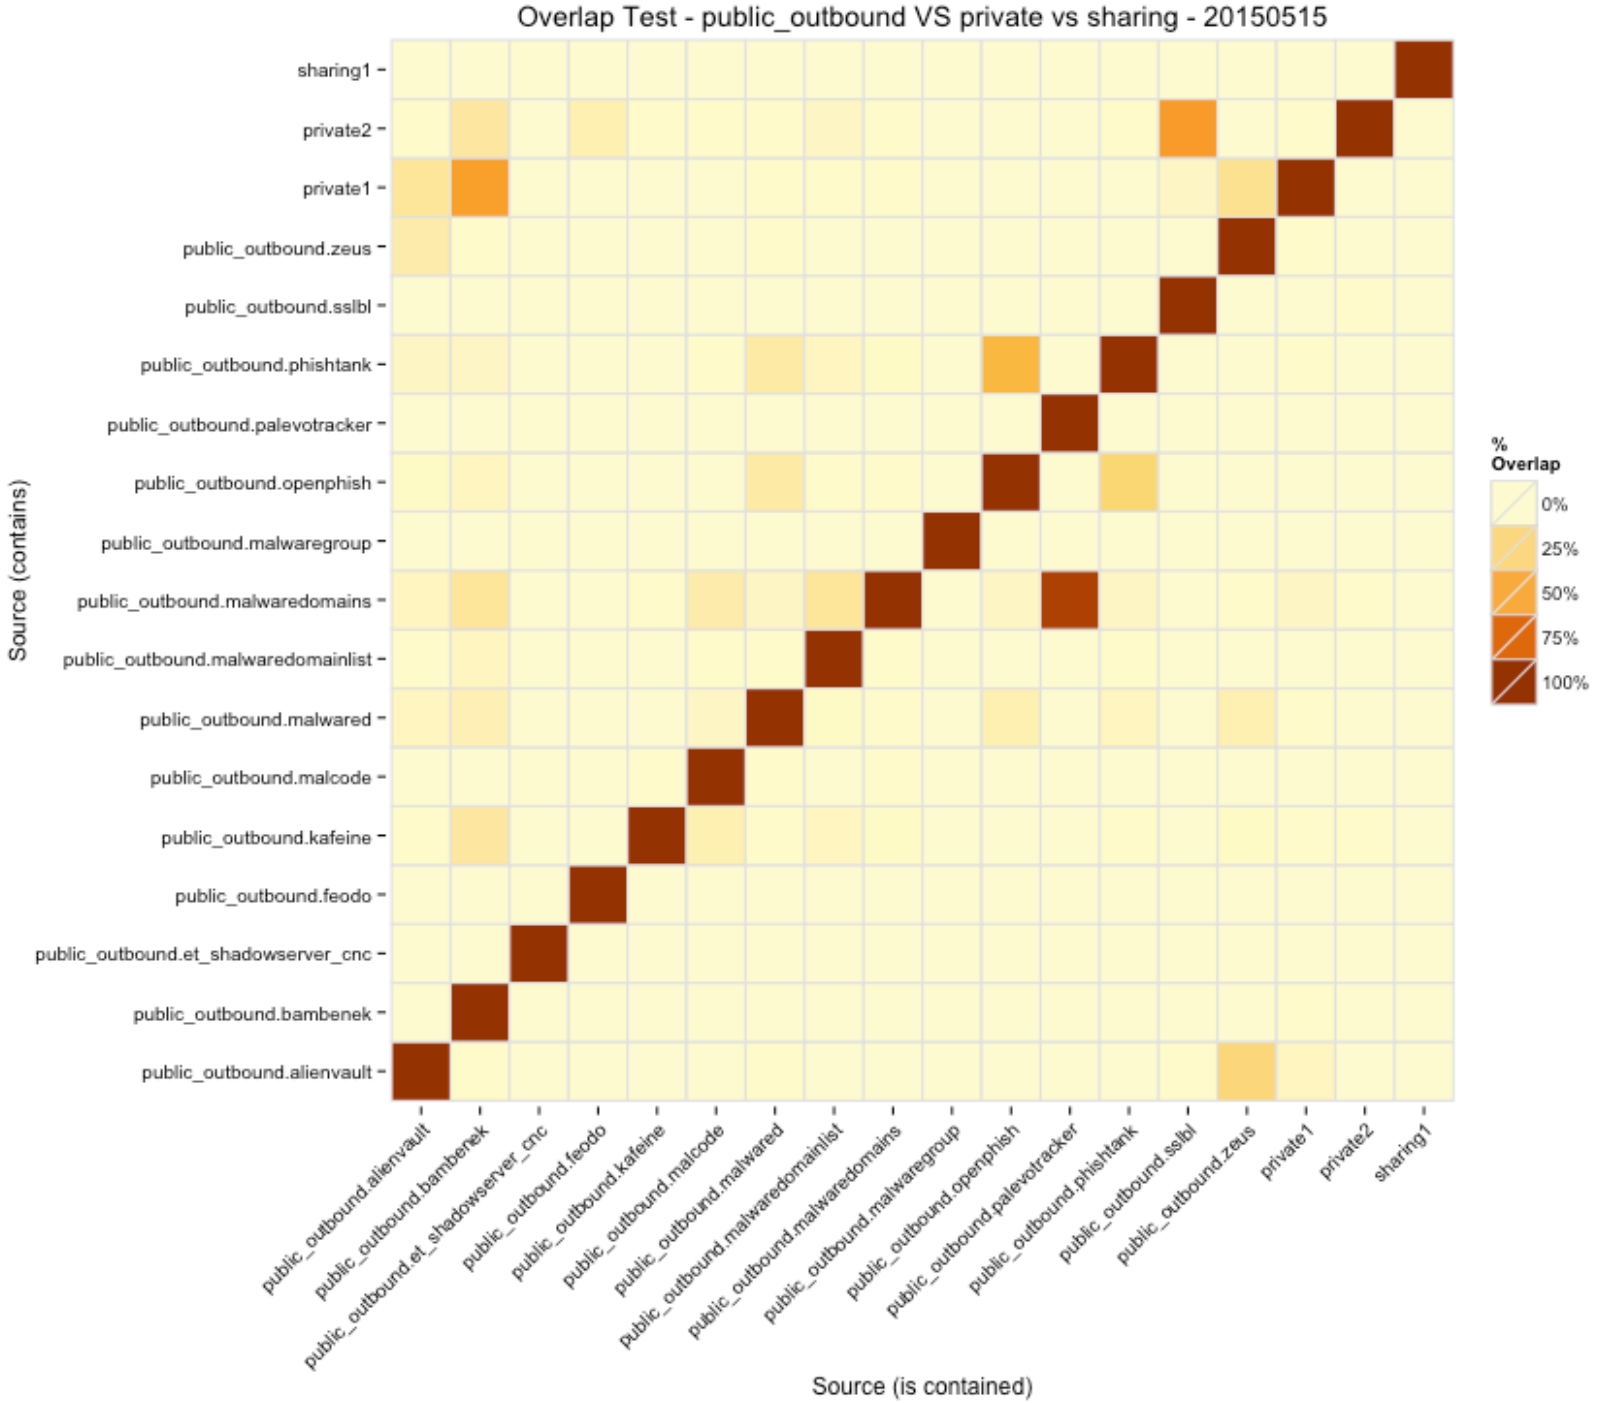
Overlap Test - public_outbound VS private vs sharing - 20150515
Source (contains)
sharing1
private2
private1
public_outbound.zeus
public_outbound.sslbl
public_outbound.phishtank
public_outbound.palevotracker
public_outbound.openphish
public_outbound.malwaregroup
public_outbound.malwaredomains
public_outbound.malwaredomainlist
public_outbound.malwared
public_outbound.malcode
public_outbound.kafeine
public_outbound.feodo
public_outbound.et_shadowserver_cnc
public_outbound.bambenek
public_outbound.alienvault
Source (is contained)
% Overlap
0%
25%
50%
75%
100%

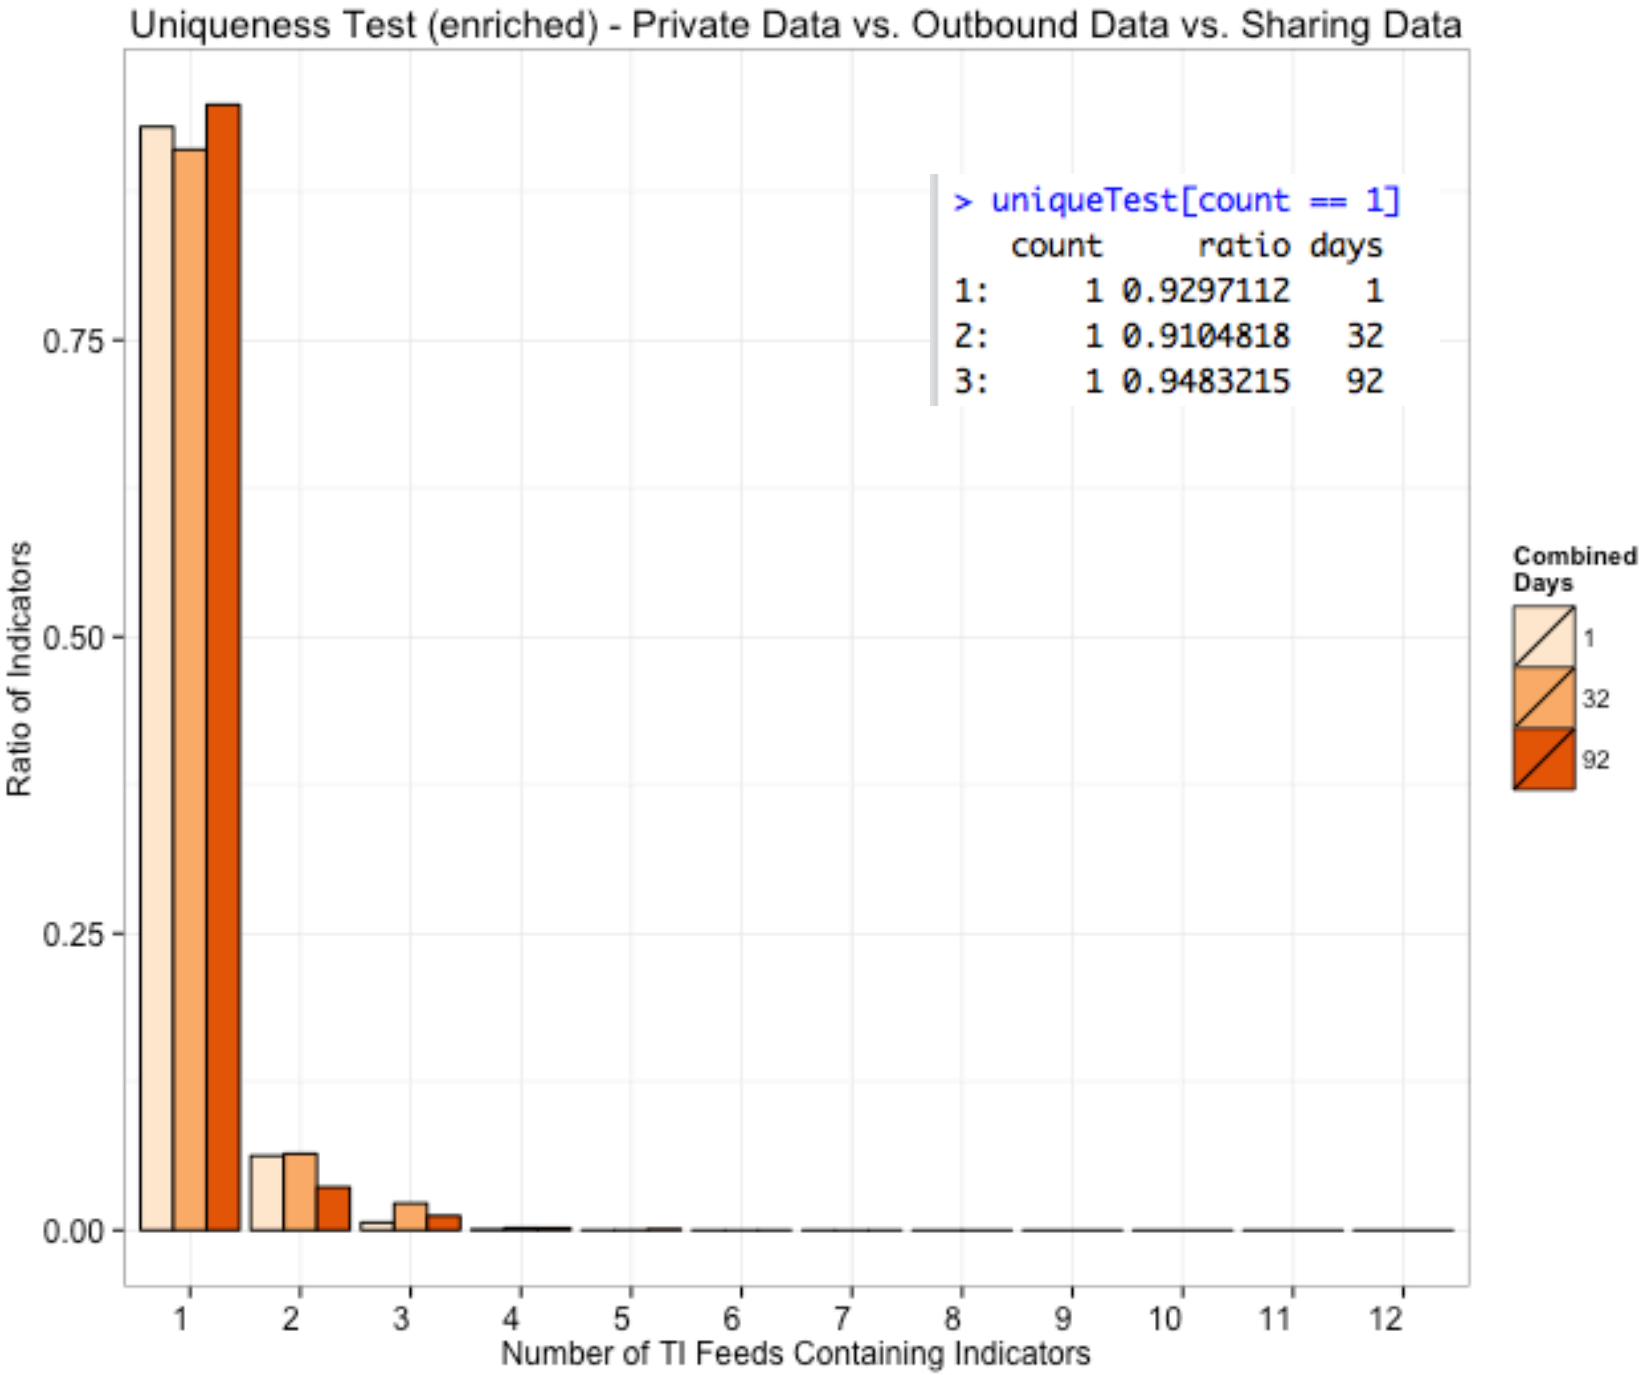
Uniqueness Test (enriched) - Private Data vs. Outbound Data vs. Sharing Data
> uniqueTest[count == 1]
count ratio days
1: 1 0.9297112 1
2: 1 0.9104818 32
3: 1 0.9483215 92
Ratio of Indicators
0.00
0.25
0.50
0.75
1
2
3
4
5
6
7
8
9
10
11
12
Number of TI Feeds Containing Indicators
Combined Days
1
32
92

Looks like we would get similar quality on a ”good” Threat Intelligence Sharing Platform as we would on a ”paid feed"

Activity Metric

Is there any actual sharing going on?

Large – 10.000s members

Small – High 10s members

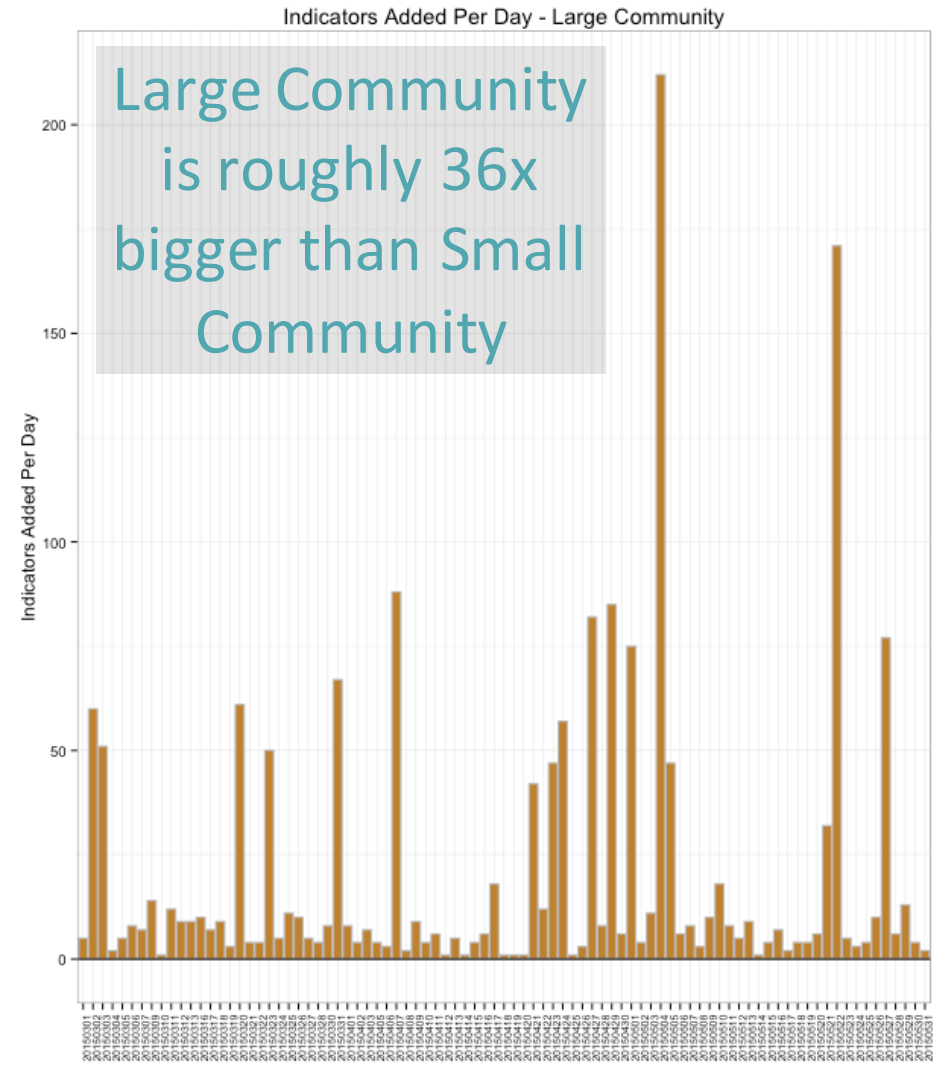

Large Community is roughly 36x bigger than Small Community

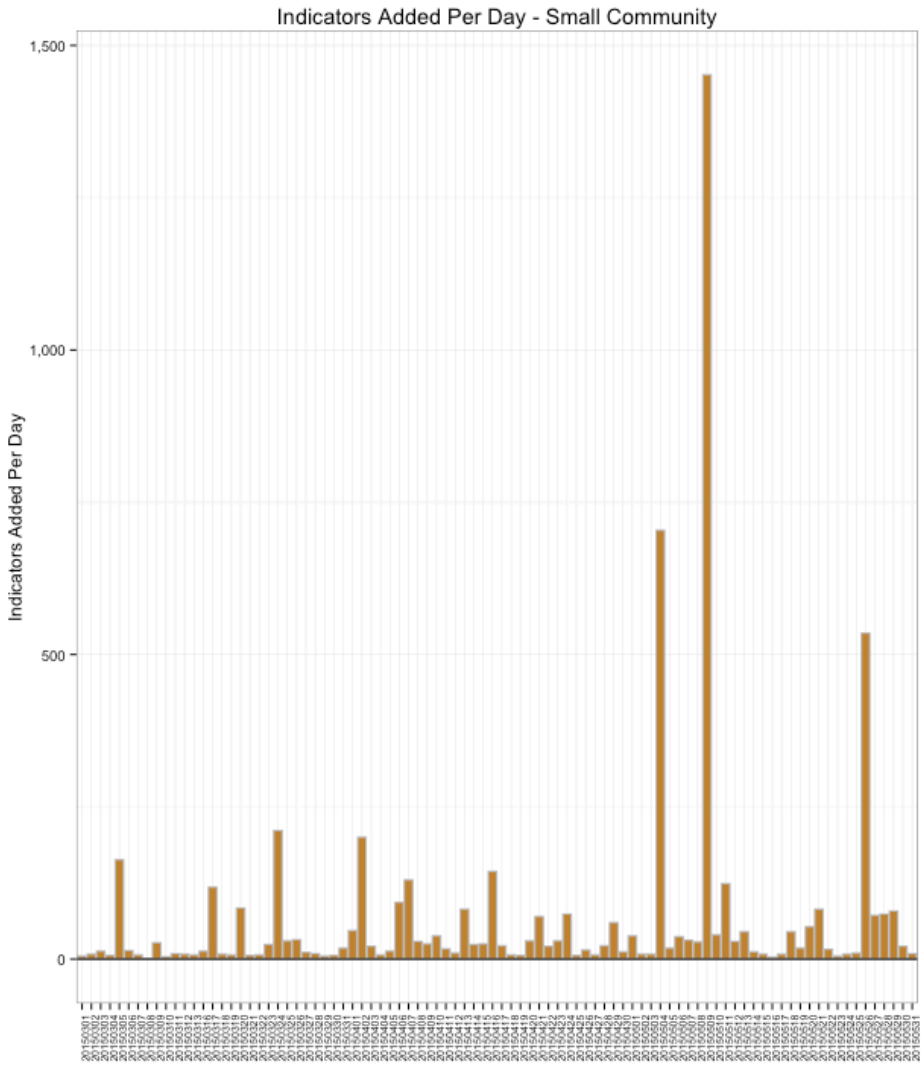

High 10s average

Low 100s average

Diversity Metric

Check your sharing privilege


Other Small Community (high tens)
Roughly 10% of the organizations share data into the community
Indicators Added Per User (Log10 Scale)
100,000
1,000
10
2
4
6
8

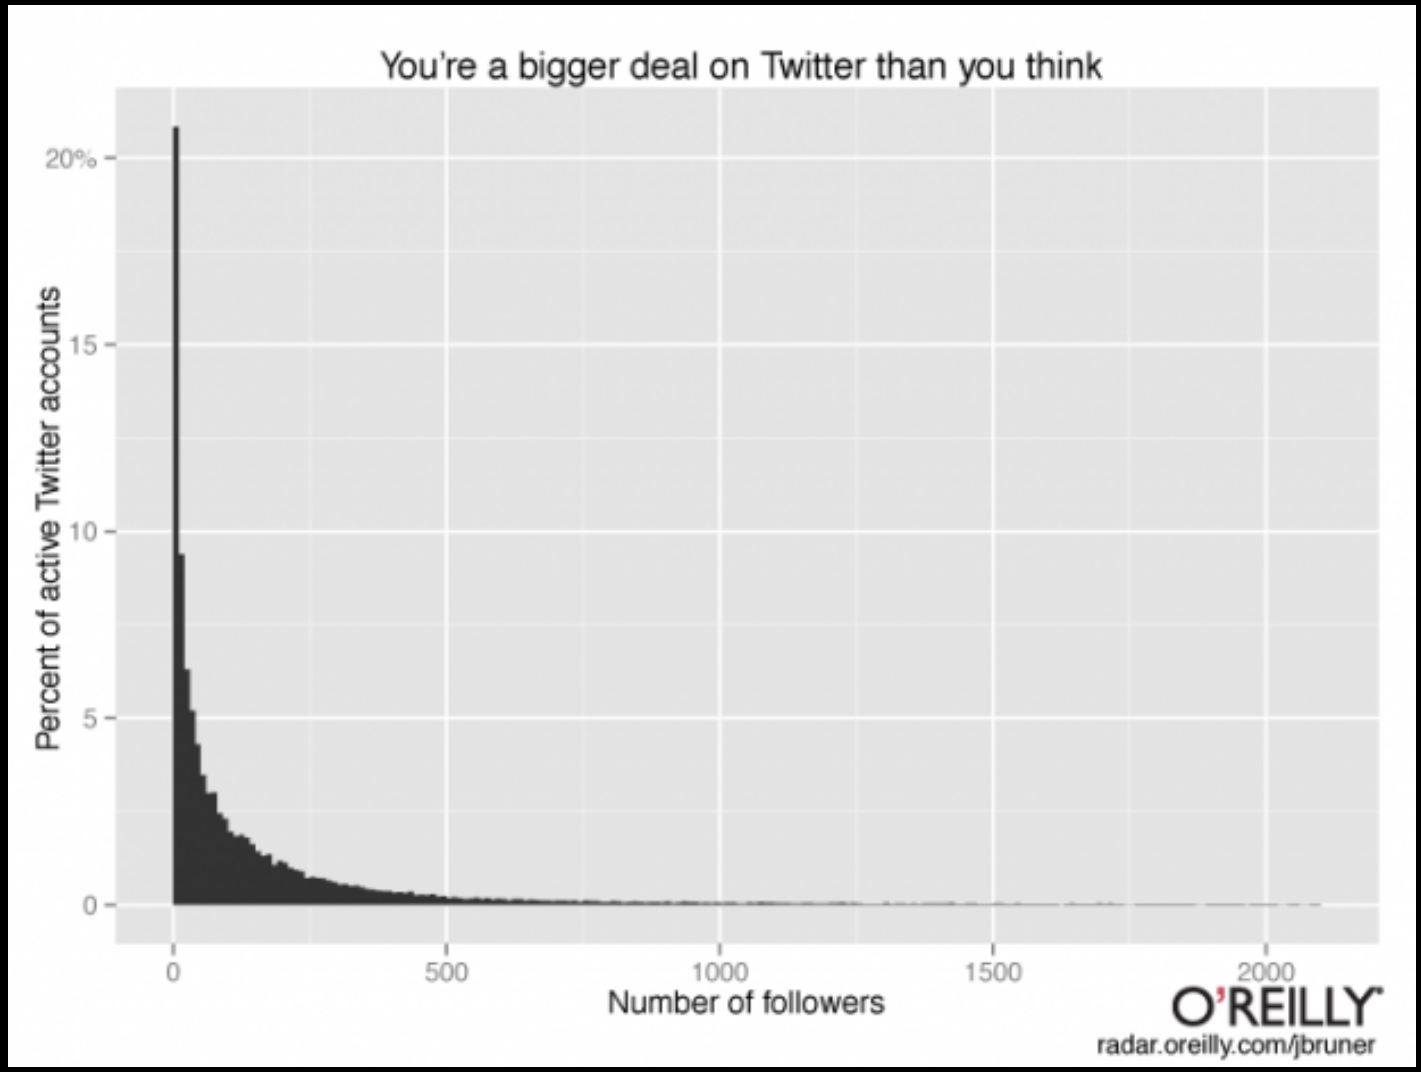

Some organizations are clearly in a better position operationally and legally to share. And that is expected due to our premises.

Feedback Metric

But is the data any good?


citation needed

Feedback Metric

- Almost no support on automation-driven platforms
- Some allow you to leave ”comments” or ”new descriptors” for the IOCs – even by counting those very low % in relation to new shared data

- Analyst-driven environments allow for collaboration on e-mails and forum posts to describe and refine strategies and best practices.

- How can we make this collaboration work on automation-driven platforms?

Trust Metric

Are we helping all the community or just a few orgs at a time?

citation needed

citation needed

Trust Metric

- The rough estimate seems to be that more than 50% of ”sharing” (IOCs, messages, etc) happens in ”private groups” inside the infrastructure of the sharing platform

- All communities have them:
  - Part of the DNA of the IC / cleared community
  - Offsets the trust equation, but defeats the ”herd immunity” argument
  - Usually MANDATORY on collaboration with LEA
- But then the ”good” data is not helping ”the community” is there any way we can reconcile?

The Future of Sharing

At the very least my humble opinion ☺

#squadgoals

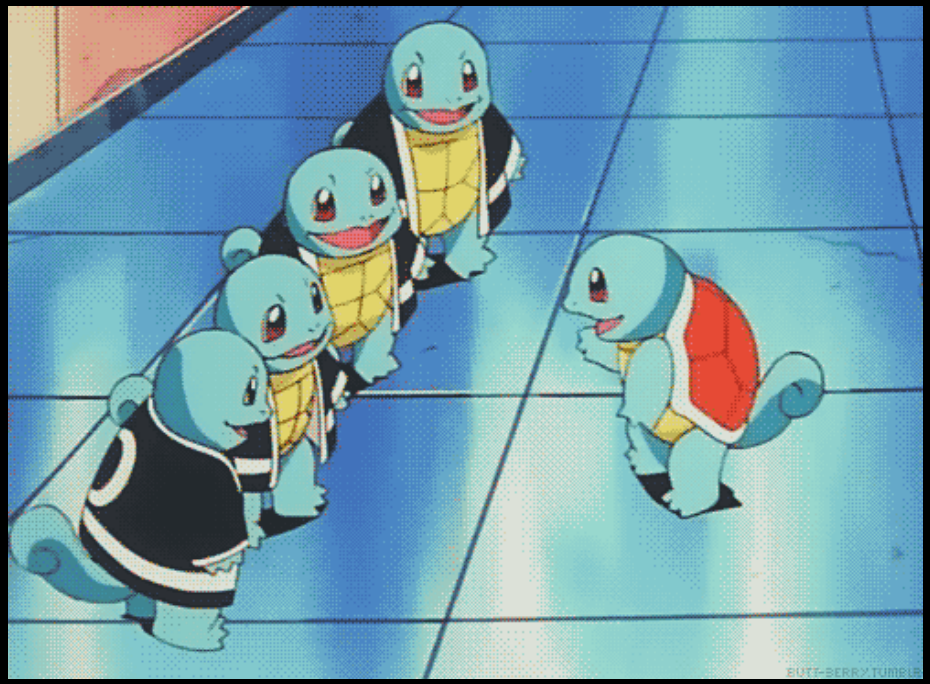

Increase the TRUST among peers

Reduce the TECHNICAL BARRIER for sharing useful information

TRUST: Anonymity + Good Curation

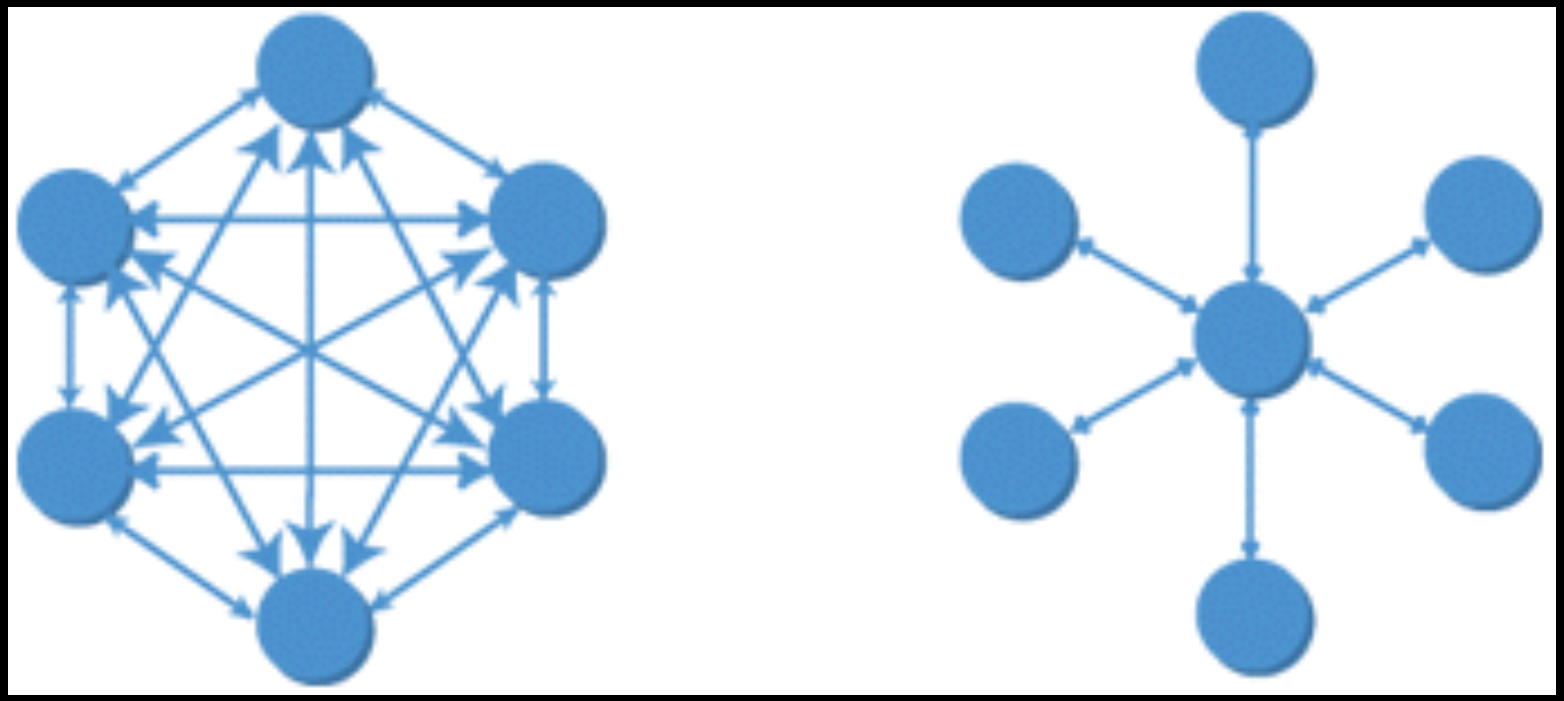

Some sharing communities accept anonymous submissions that they then curate and disseminate to all organizations

TECHNICAL BARRIER:
"Pyramid of Sharing"

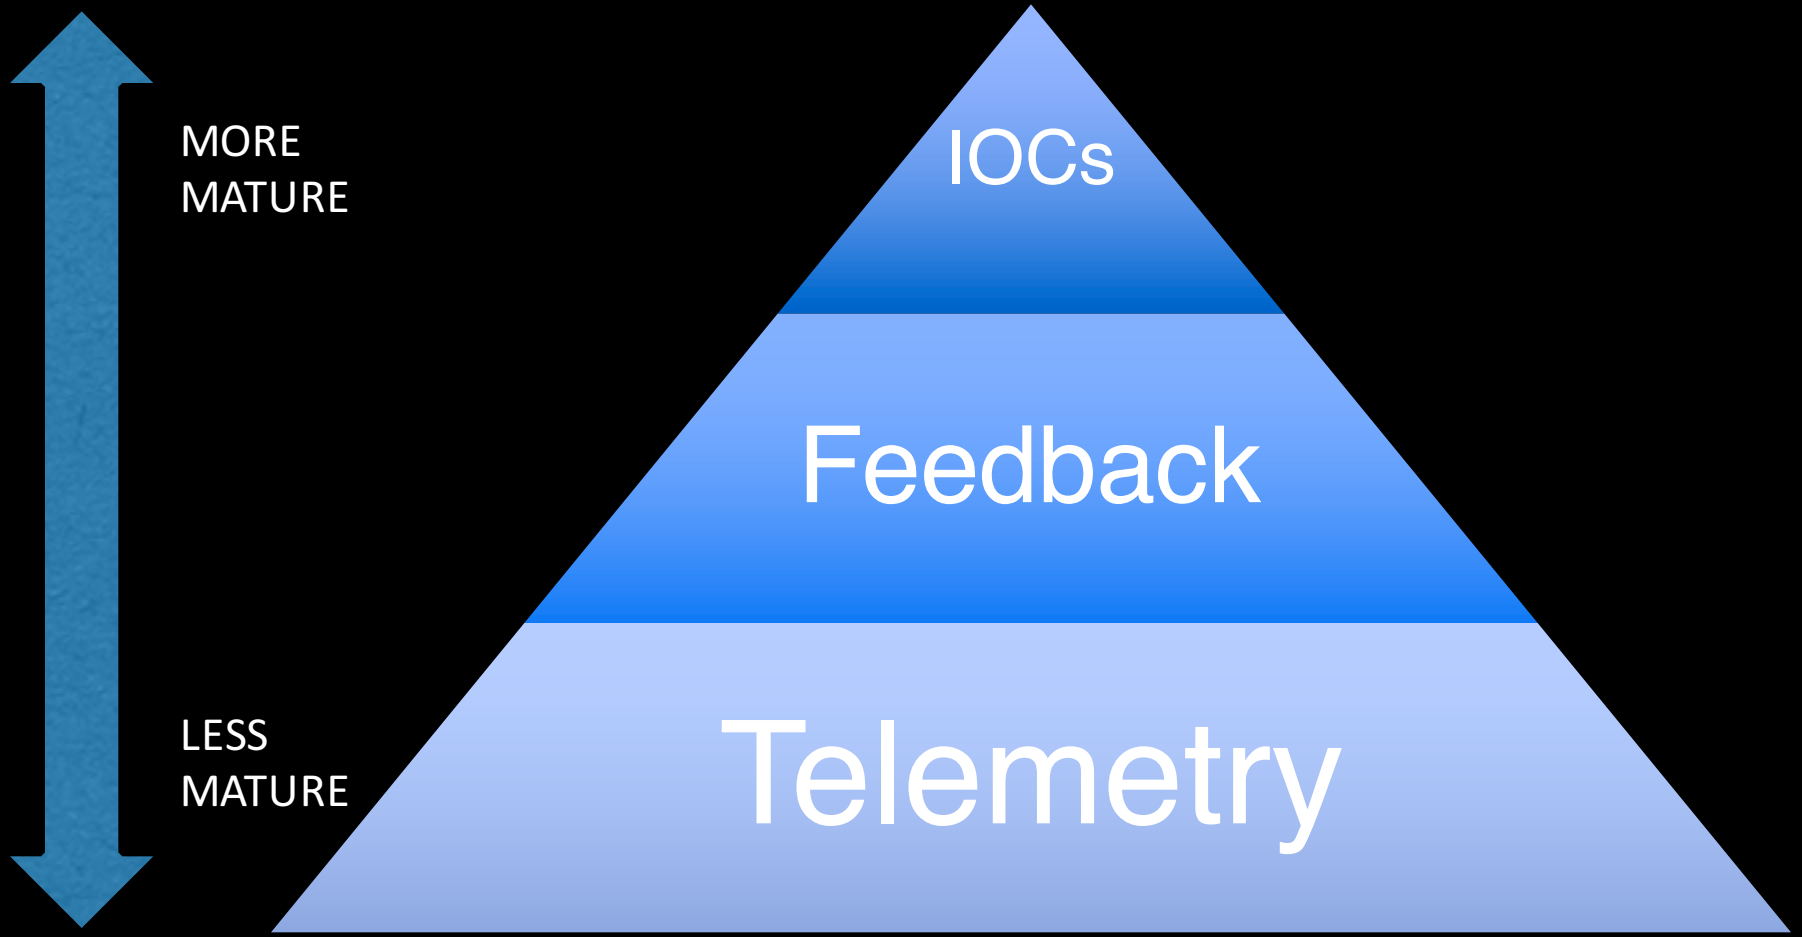

❤ and apologies to @DavidJBianco

Takeaways

- Intelligence Sharing is a very analyst-centric activity that we have been tasked with scaling out
- Data can be as good as a paid feed, but you have to be in the right circles of trust
- Does not solve analyst shortage and making the indicators / strategies operational into your environment

Thanks!

- Q&A?
- Feedback!

Alex Pinto
@alexcpsec
@MLSecProject

Alexandre Sieira
@AlexandreSieira
@NiddelCorp

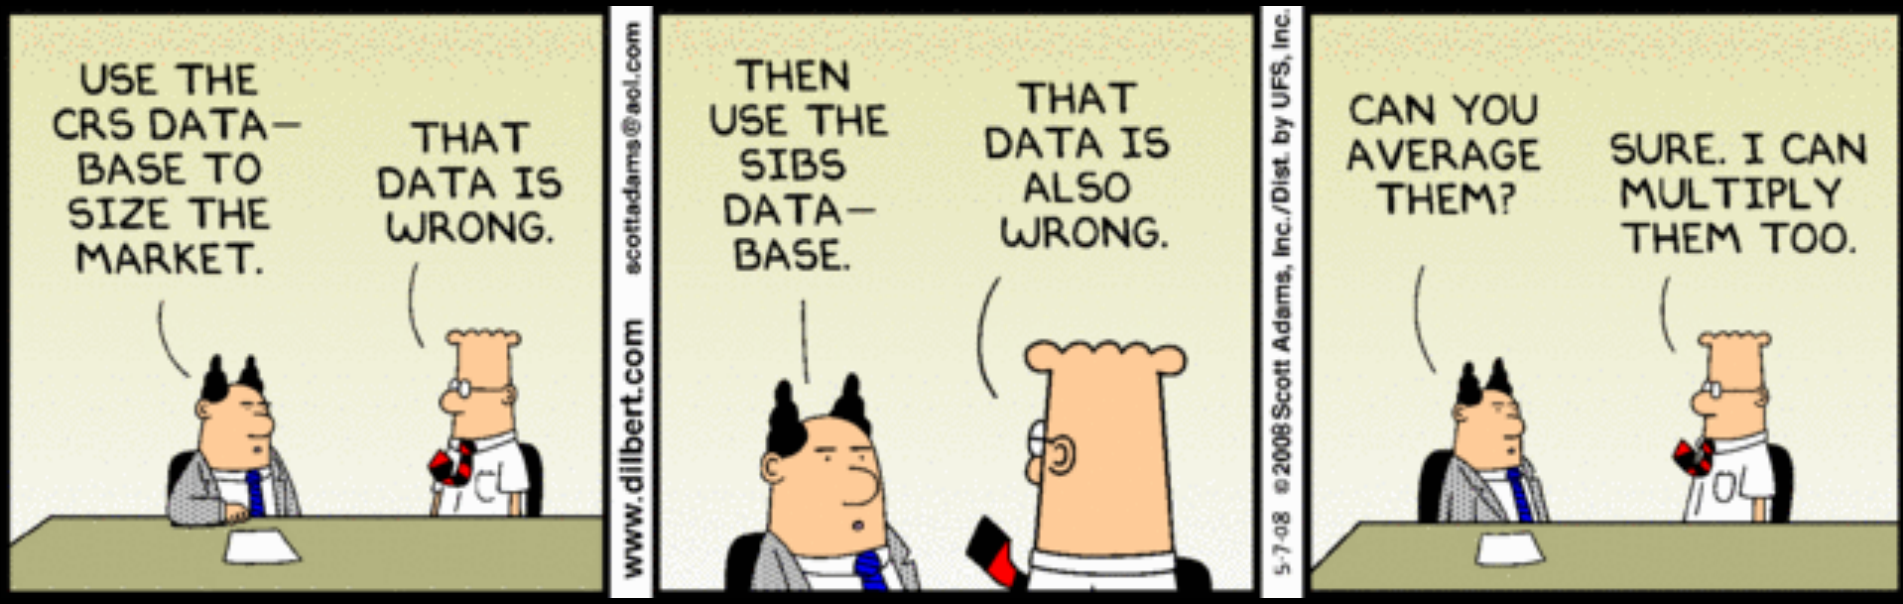

”The measure of intelligence is the ability to change."

- Albert Einstein