A7165: KISMET - Knowledge of Internet Structure: Measurement, Epistemology, and Technology

kc claffy
kc@caida.org

David Clark
ddc@csail.mit.edu

Overview

We will create a knowledge network that will improve the security and functioning of three key but inherently vulnerable systems that underpin all activity on the Internet: the addressing, routing, and naming systems. Lack of attention to security in early design decisions, persistent disagreement about the best path to improvement, and lowest-cost operational practices surrounding these three layers have allowed malicious actors to execute and scale harmful misbehavior including interception and disruption of traffic, denial of service and phishing attacks, and distribution and execution of malware.

Description

Abuse of the Domain Name System is rampant: the Internet Corporation for Assigned Names and Numbers (ICANN) reported in March that there were about 740K active domain names associated with abusive practices. The Internet routing system shares similar vulnerabilities but little reporting; only high-profile episodes make headlines, such as two uses of BGP in 2018 to steal cryptocurrency and create $29M of fraudulent advertising revenue.

The premise of our work is that we cannot improve security by responding to individual events. Rather we must shift the landscape by understanding the operation of these Internet systems and behavior patterns of malicious actors, with the goal of finding operational practices that will thwart attackers.

Our Open Knowledge Network is structured to yield two types of knowledge: technical knowledge (Stage 1), with significant scientific and engineering components (tool developments, data analysis), and actionable knowledge (Stage 2) of direct use to society. To enable Stage 1, we will gather data sources and expertise related to these systems, in order to support scientifically reproducible data analytics that can identify ongoing security threats, evaluate proposed mitigations, and track improvements. In Stage 2, we will cultivate an international operational community, drawing from disciplines including law and economics, to develop and socialize operational practices that can prevent or mitigate vulnerabilities. Our success metrics will include reduced “time to insight” regarding structural characteristics of these systems, methods to develop, evaluate, and instantiate enhanced practices in the Internet naming systems, and demonstrating the utility of the OKN in verifying compliance with enhanced practices. We will analyze models of sustainability for the OKN, learning from other critical industries, e.g., financial and energy sectors.

Our long-range goal is to use our OKN to identify, incentivize, and validate operational practices that improve Internet security.

Differentiators

Approaches to solving fundamental Internet vulnerabilities have been: (1) largely technical, failing to overcome non-technical barriers; (2) complex and expensive; (3) global in scope. These properties are not compatible with the pressure of a competitive ecosystem, or in some cases governments with counter-incentives to address infrastructure security issues. We have spent years studying the vulnerabilities and tracking the impact of proposed solutions. Our conclusion, re-affirmed by our Phase I work, is that the path to better security does not lie in proposals to make global changes to the Internet protocols, but in finding operational practices that regions of the Internet can implement to improve the security profile of those regions. We identify important changes in Internet traffic patterns that reduce the necessity for global agreement in order to improve security in regions of the Internet, increasing our likelihood of success.

Many organizations around the world generate data that can contribute to the OKN, some of whom are collaborators and partners in this effort. Navigating the diffuse nature of information about network infrastructure is a significant challenge. We will pursue a community-driven approach to sharing data using consistent formats and APIs, and identify metadata that enables discoveries that cut across data sets and domains. Our critical assets include partners with commercial network engineering expertise, data science techniques, and technology to manage data integrity, availability, and privacy.

Road Map

Y1: Establish OKN, including annotated maps of routing and naming systems

-- Demonstrate use of OKN to identify properties of network and DNS service providers associated with persistent misbehavior or large attack surfaces.

-- Develop tools to understand key challenges in secure network configuration.

-- Use OKN to verify compliance with operational practices required by the Mutually Agreed Norms for Routing Security initiative (MANRS).

-- Workshops with industry experts.

-- Develop extension to MANRS to prevent hijacks; validate benefit with OKN.

Y2: Generate and validate actionable knowledge.

-- Create and demonstrate utility of DNS+BGP knowledge graph to improve security properties

-- AI-driven assistant for security configuration.

-- Develop catalog of best practices for DNS ecosystem. Demonstrate use of OKN to verify utility. Socialize/refine with industry actors.

Partnerships

Securing the Internet requires national and international partnerships that combine previously disconnected data sources, user communities, operators, regulators, and other stakeholders to overcome conflicting incentives that hinder development and deployment of data-driven solutions. The leaders we have selected span government (ESnet, NIST, FCC), industry (Verisign, PIR, Comcast, Farsight, Interisle), non-profit (ISOC, DNS-OARC, SIDN Labs, M3AAWG), and academic (MIT, U Oregon, BYU, Virginia Tech, U Waikato) sectors. Our partners bring diverse data sources, but more importantly they bring deep understanding of network engineering, operations, cartography, curriculum development and training, research infrastructure, economics, policy development and other disciplines fundamental to the creation of this prototype OKN. This interdisciplinary collaboration is required to not only create and sustain the OKN, but also to use it to evaluate and apply technical knowledge in economic, legal, and social contexts.

Intellectual Property

All software source code brought into project will be BSD, GPL, or UCSD licensed. Data sharing will rely on CAIDA Acceptable Use Agreements as a basis, with modifications as needed by data owners.

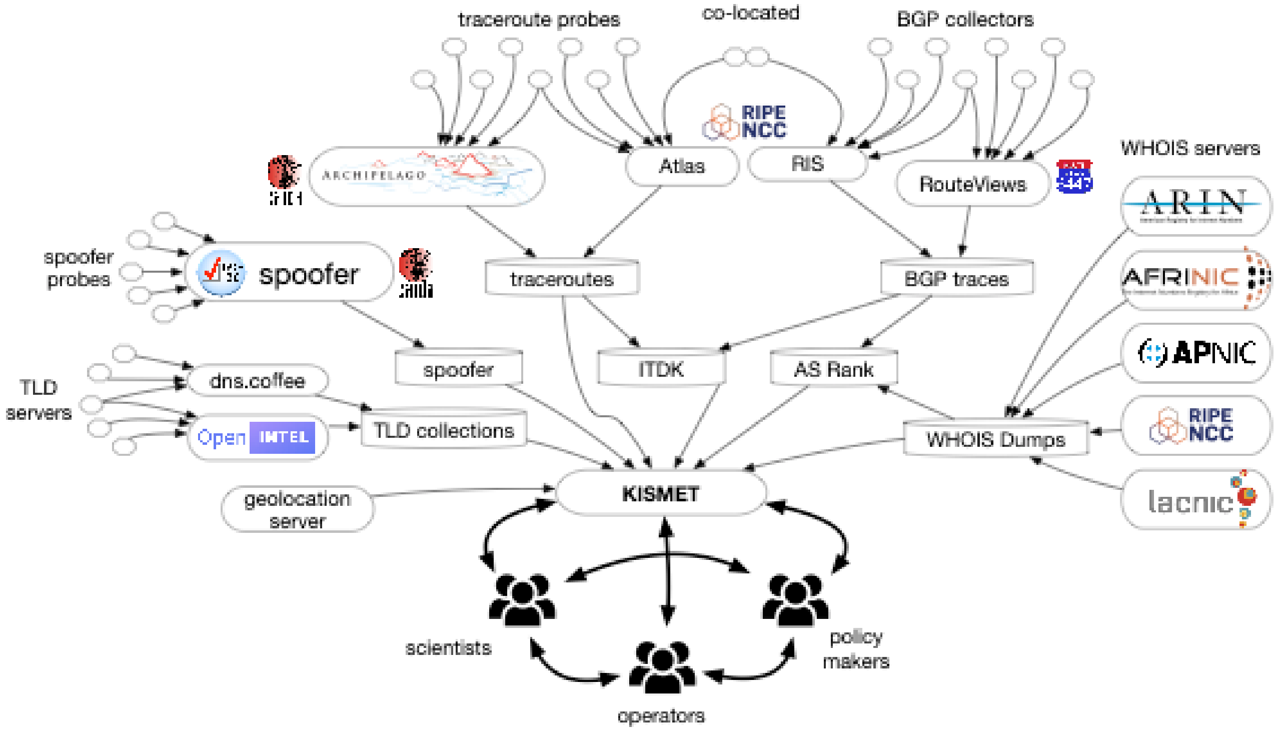
traceroute probes
co-located
BGP collectors
RIPE NCC
ARCHIPELAGO
Atlas
RIS
RouteViews
WHOIS servers
ARIN
AFRINIC
APNIC
RIPE NCC
lacnic
spoofer probes
spoofer
traceroutes
BGP traces
spoofer
ITDK
AS Rank
TLD servers
dns.coffee
Open INTEL
TLD collections
WHOIS Dumps
geolocation server
KISMET
scientists
policy makers
operators