Modeling a Mission-Aware Prioritization Scheme for Cyber Incidents

Lena Pons

Data Scientist

Cyber Security Foundations | CERT

Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213

Problem: Too little decision support for cyber incident priority

Current Situation: Too little support to prioritize alerts

How can we (1) prioritize alerts and (2) incorporate mission context?

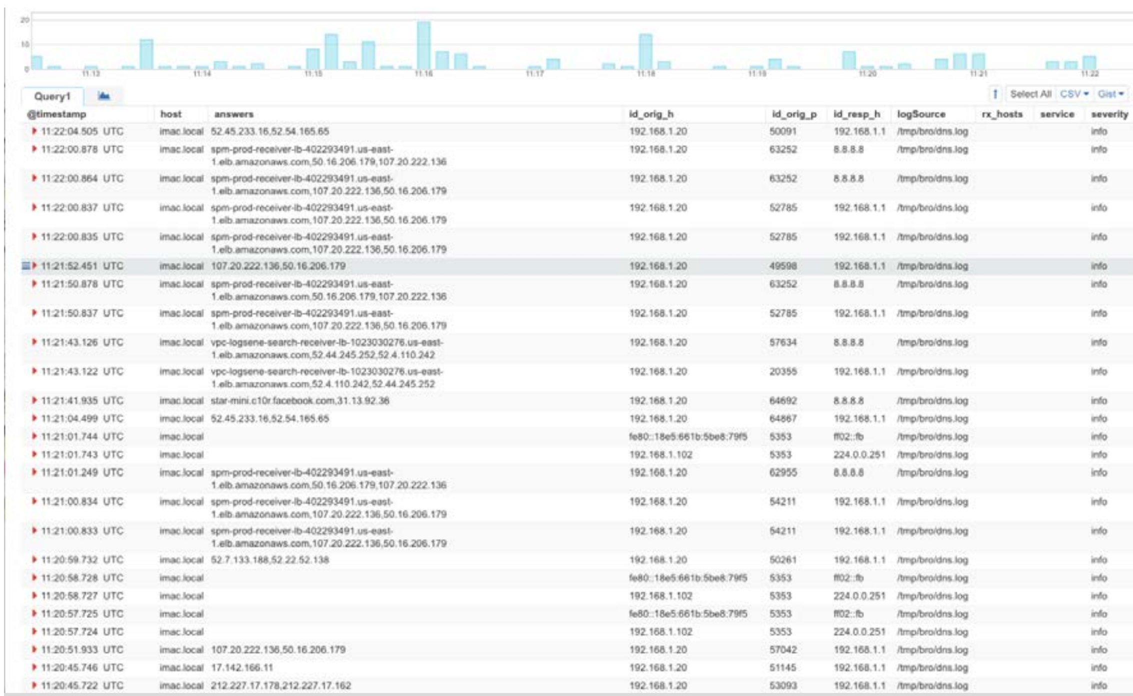

@timestamp	host	answers	id_orig_h	id_orig_p	id_resp_h	logSource	rx_hosts	service	severity
11:22:04.505 UTC	imac.local	52.45.233.16,52.54.165.65	192.168.1.20	50091	192.168.1.1	/tmp/bro/dns.log			info
11:22:00.878 UTC	imac.local	spm-prod-receiver-lb-402293491.us-east-1.elb.amazonaws.com,50.16.206.179,107.20.222.136	192.168.1.20	63252	8.8.8.8	/tmp/bro/dns.log			info
11:22:00.864 UTC	imac.local	spm-prod-receiver-lb-402293491.us-east-1.elb.amazonaws.com,107.20.222.136,50.16.206.179	192.168.1.20	63252	8.8.8.8	/tmp/bro/dns.log			info
11:22:00.837 UTC	imac.local	spm-prod-receiver-lb-402293491.us-east-1.elb.amazonaws.com,107.20.222.136,50.16.206.179	192.168.1.20	52785	192.168.1.1	/tmp/bro/dns.log			info
11:22:00.835 UTC	imac.local	spm-prod-receiver-lb-402293491.us-east-1.elb.amazonaws.com,107.20.222.136,50.16.206.179	192.168.1.20	52785	192.168.1.1	/tmp/bro/dns.log			info
11:21:52.451 UTC	imac.local	107.20.222.136,50.16.206.179	192.168.1.20	49598	192.168.1.1	/tmp/bro/dns.log			info
11:21:50.878 UTC	imac.local	spm-prod-receiver-lb-402293491.us-east-1.elb.amazonaws.com,50.16.206.179,107.20.222.136	192.168.1.20	63252	8.8.8.8	/tmp/bro/dns.log			info
11:21:50.837 UTC	imac.local	spm-prod-receiver-lb-402293491.us-east-1.elb.amazonaws.com,107.20.222.136,50.16.206.179	192.168.1.20	52785	192.168.1.1	/tmp/bro/dns.log			info
11:21:43.126 UTC	imac.local	vpc-logsene-search-receiver-lb-1023030276.us-east-1.elb.amazonaws.com,52.44.245.252,52.4.110.242	192.168.1.20	57634	8.8.8.8	/tmp/bro/dns.log			info
11:21:43.122 UTC	imac.local	vpc-logsene-search-receiver-lb-1023030276.us-east-1.elb.amazonaws.com,52.4.110.242,52.44.245.252	192.168.1.20	20355	192.168.1.1	/tmp/bro/dns.log			info
11:21:41.935 UTC	imac.local	star-mini.c10r.facebook.com,31.13.92.36	192.168.1.20	64692	8.8.8.8	/tmp/bro/dns.log			info
11:21:04.499 UTC	imac.local	52.45.233.16,52.54.165.65	192.168.1.20	64867	192.168.1.1	/tmp/bro/dns.log			info
11:21:01.744 UTC	imac.local		fe80::18e5:661b:5be8:79f5	5353	ff02::fb	/tmp/bro/dns.log			info
11:21:01.743 UTC	imac.local		192.168.1.102	5353	224.0.0.251	/tmp/bro/dns.log			info
11:21:01.249 UTC	imac.local	spm-prod-receiver-lb-402293491.us-east-1.elb.amazonaws.com,50.16.206.179,107.20.222.136	192.168.1.20	62955	8.8.8.8	/tmp/bro/dns.log			info
11:21:00.834 UTC	imac.local	spm-prod-receiver-lb-402293491.us-east-1.elb.amazonaws.com,107.20.222.136,50.16.206.179	192.168.1.20	54211	192.168.1.1	/tmp/bro/dns.log			info
11:21:00.833 UTC	imac.local	spm-prod-receiver-lb-402293491.us-east-1.elb.amazonaws.com,107.20.222.136,50.16.206.179	192.168.1.20	54211	192.168.1.1	/tmp/bro/dns.log			info
11:20:59.732 UTC	imac.local	52.7.133.188,52.22.52.138	192.168.1.20	50261	192.168.1.1	/tmp/bro/dns.log			info
11:20:58.728 UTC	imac.local		fe80::18e5:661b:5be8:79f5	5353	ff02::fb	/tmp/bro/dns.log			info
11:20:58.727 UTC	imac.local		192.168.1.102	5353	224.0.0.251	/tmp/bro/dns.log			info
11:20:57.725 UTC	imac.local		fe80::18e5:661b:5be8:79f5	5353	ff02::fb	/tmp/bro/dns.log			info
11:20:57.724 UTC	imac.local		192.168.1.102	5353	224.0.0.251	/tmp/bro/dns.log			info
11:20:51.933 UTC	imac.local	107.20.222.136,50.16.206.179	192.168.1.20	57042	192.168.1.1	/tmp/bro/dns.log			info
11:20:45.746 UTC	imac.local	17.142.166.11	192.168.1.20	51145	192.168.1.1	/tmp/bro/dns.log			info
11:20:45.722 UTC	imac.local	212.227.17.178,212.227.17.162	192.168.1.20	53093	192.168.1.1	/tmp/bro/dns.log			info

Operator perspective:

- Address cyber incidents of greatest concern.
- How do I know which are of greatest concern?
- How do I put the cyber alert in mission context?
- How do I divide my attention among multiple responsibilities?

Foundational Work: Mission-Specific Cyber Asset Criticality

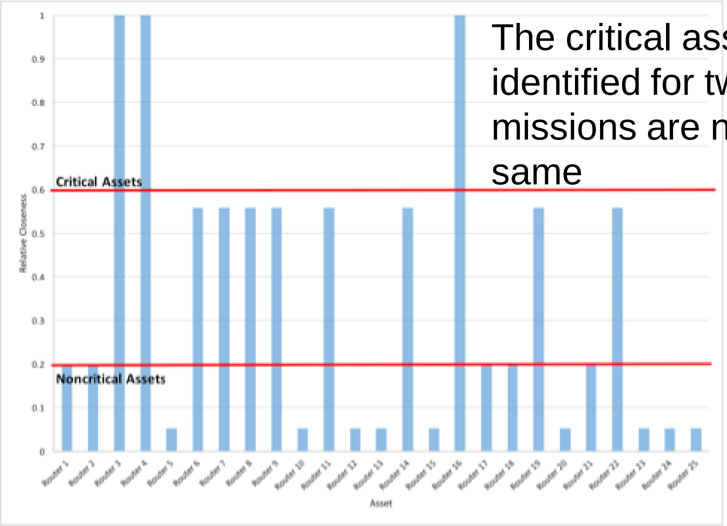

The critical assets identified for two missions are not the same

Figure 6. Asset criticality under the Sea Control Mission: average family weighting, transitive criteria, and uniform asset scoring.

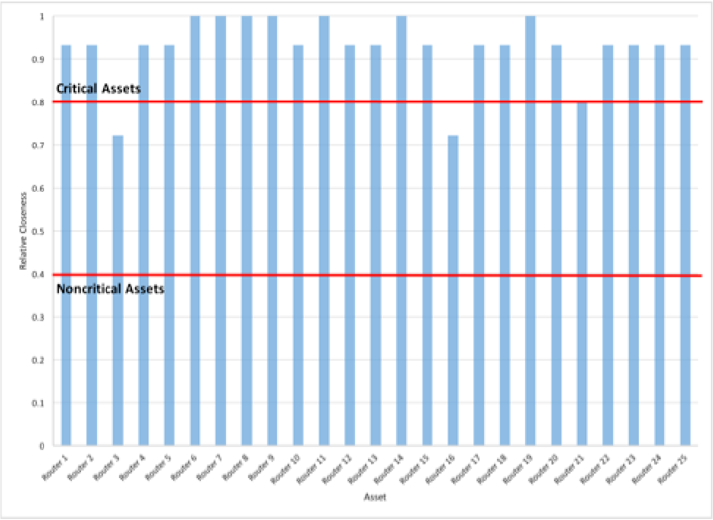

Figure 7. Asset criticality under the Power Projection Mission: average family weighting, transitive criteria, and uniform asset scoring.

Systems in the network are enumerated along the x-axis, and criticality to a particular mission is mapped along the y-axis

Asset Criticality Varies by Mission

Key conclusions:

- Cyber situational awareness should account for mission
- Situational awareness tools should account for changes in operational use of a system that are mission specific

Asset Criticality in Mission Reconfigurable Cyber Systems and its Contribution to Key Cyber Terrain

Peyton Price*, Nicholas Anthony Leyba*, Mark Gondree†, Zachary Staples*, Thomas Parker‡
*Naval Postgraduate School
nicholas.a.leyba.mil@mail.mil, zhstaple@nps.edu
†Sonoma State University
‡*thomas.c.parker@navy.mil*

Retrieval Optimization: Information vs Alert

Information Retrieval

*A **collection** of **documents** which contain **topics***

A **topic** is a collection of words which relate to an document relevance to an information need

Queue Retrieval

*A **queue** of **alerts** which contain **severity attributes***

Severity attributes are pieces of information that give clues to how to prioritize

An alert functions differently than a document, in that conventional methods of determining relevance are not appropriate

Ranked Retrieval & Unbounded Lists

Machine Learning (293M results) → ML + Natural Language Processing (181M results) → ML + NLP – Deep Learning (1.8M results) → ML + NLP - DL + Question Answering (75k results)

Ranking Results in Unbounded List

An unbounded list is just a collection of information where the number of items exceeds the number that will be read.

Ranking results in an unbounded list means we have to go beyond topic relevance to find what information is most useful

Ranking all the results in an unbounded list can be computationally very expensive, and not all that valuable to the user

Once you pass a threshold of returned results, anything below a certain value will not be read

Prioritize Alerts: Cascade Model

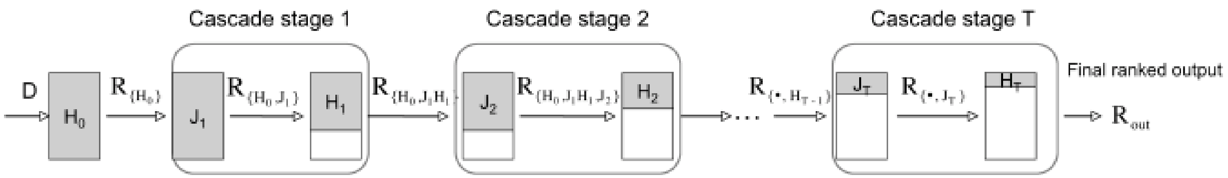

Figure 1: An example cascade. After an initial ranking function H_0, each stage consists of two sequential operations: J_t prunes the input ranked documents, then a local ranking function H_t refines the rank order of the retained documents. The new ranked list is passed to the next stage. The size of the shaded area denotes the size of the candidate documents. Subscripts for each ranked list denotes the sequence of actions applied.

- In an unbounded retrieval model, relevance cannot be the only determiner.
- More complex ranking schemes are computationally expensive, and degrade time performance.
- Reducing complexity by feature selection reduces relevance clarity.
- Apply increasingly complex ranking, while pruning less relevant results with each pass.

Lidan Wang, Jimmy Lin, and Donald Metzler. 2011. A cascade ranking model for efficient ranked retrieval. In Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval (SIGIR '11). ACM, New York, NY, USA, 105-114.

Candidate Cascade

Pass 1: Relevance

Alerts are presented in the order they are fired

1. INC-1234
2. INC-1235
3. INC-1236
4. INC-1237

…

Event meets the criteria for an alert

Pass 2: Timeliness / Persistence

Alerts are presented in order of event occurrence, oldest first

1. INC-1171
2. INC-1222
3. INC-1234
4. INC-1240

…

Older events are presented higher in the list, events that recur are also upweighted

Pass 3: Alert Severity

Alerts are prioritized with respect to scoring of incident severity

1. INC-1222
2. INC-1234
3. INC-1171
4. INC-1240

…

Severity scoring may include functional impact, observed activity, location of activity, actor characterization, information impact, recoverability, etc.

Pass 4: Mission context

Alerts are presented in order of most critical affected system

1. INC-1201
2. INC-1234
3. INC-1222
4. INC-1171

…

In the final pass, the ranking incorporates asset criticality scores that are mission aware

Cascade Model

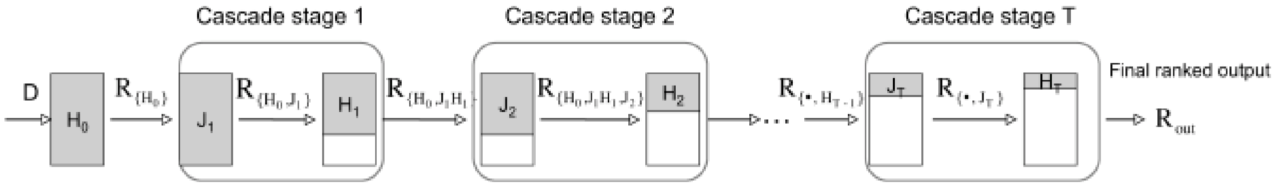

Applying Cascade to Queue

Cascade model allows for complex, multilayered event prioritization while reducing computational time.

Elements in the cascade can be modified for operational context.

Mission context is relevant to providing appropriate situational awareness.

Contact Information

Presenter

Lena Pons

Machine Learning Researcher

Telephone: +1 703.247.1374

Email: lepons@sei.cmu.edu

References

Mission-Aware Cyber Incident Prioritization

Backup Slides

Information Retrieval: Keyword-Based Relevance

Within the financial services sector, Anti-Money Laundering (AML) is a significant challenge for many institutions, often consuming large numbers of people and effort to manage the process and comply with the regulations. As a result, these same institutions are looking for new solutions to help them reduce the burden and increase the controls in this complex space. The combination of **artificial intelligence (AI)** and, more specifically, **machine learning (ML)**, are increasingly being considered as enablers of a better solution.
Despite its potential, however, adoption of **AI** and **ML** within Anti-Money Laundering has been relatively slow. This is due, in part, to the limited understanding of how **AI** and **ML** could be applied within compliance programs, and to the fact that regulators and compliance officers are often concerned that **AI** and **ML** are "black boxes" whose inner workings are not clearly
understood. Regulators typically require compl
understand and validate not just the outputs, bu
outcomes from AML **models** are derived. Des
concerns, we already see movement and applica
technologies.
Machine learning has been shown to be parti
conducting suspicious activity monitoring and t
monitoring, two key AML activities. A common
transaction monitoring, for example, is the gen
number of alerts, which in turn requires operati
and process the alerts. **ML** can **teach comput**
recognize suspicious behavior and to **classify** a
high, medium or lower risk. Applying **rules** to
classifications can facilitate the automatic closi
allowing humans to **supervise** the machines th
alerts rather than reviewing all of the alerts ma
better use of the time of these experts.

Google today announced the alpha launch of AutoML Vision, a new service that helps developers — including those with no **machine learning (ML)** expertise — build custom image recognition models. While Google plans to expand this custom **ML model** builder under the AutoML brand to other areas, the service for now only supports **computer vision models**, but you can expect the company to launch similar versions of AutoML for all the standard **ML** building blocks in its repertoire (think speech, translation, video, natural language recognition, etc.).
The basic idea here, Google says, is to allow virtually anybody to bring their images, upload them (and import their tags or create them in the app) and then have Google's systems **automatically** create a customer **machine learning model** for them. The company says that Disney, for example, has used this system to make the search feature in its online store more robust because it can now find all the products that feature a likeness of Lightning McQueen and not just those where your favorite talking race car was **tagged** in the text description.
The whole process, from importing **data** to **tagging** it and **training** the **model**, is done through a drag and drop interface. We're not talking about something akin to Microsoft's Azure ML studio here, though, where you can use a Yahoo Pipes-like interface to build, train and evaluate models. Instead, Google is opting for a system where it handles all of the hard work and trains and tunes your model for you.

The number of key terms that appear in a text determine whether it is relevant. Terms may be weighted, e.g. "ML" might be more relevant than "data"

Information retrieval systems traditionally order results based on the number and/or density of relevant terms that appear in the text, most systems will also factor in publication date and return more recent results first.

Prioritize Alerts: Query Construction

Machine Learning (293M results) → ML + Natural Language Processing (6.5M results) → ML + NLP - Deep Learning (1.8M results) → ML + NLP - DL + Question Answering (75,800 results)

Boolean Query Construction

More specificity reduces the number of returned results

Combining terms can also make searches more broad and return more results

Ranking Unbounded Results

Specificity alone does not provide information about relative relevance

Keyword searching does not account for timeliness, reliability of source, relevant terms not in query