IPOST REDUX:
A Continuous Monitoring Vision for NASA

Matt Linton

NASA Ames Research Center

X / II / MMXII

Congratulations – you're FISMA accredited! You're SECURE now!

See you in 3 years!

Our path to this point

- Began with work on NASA Nebula
- Experimental private cloud (now Open Stack)
- Contained intense, constant monitoring of hosts
  - Internal scan-on-launch
  - External scans constantly
  - Score threshold over X triggered remediation/containment actions
- Need for algorithms and scoring led to NASA/State department talks

RIP Nebula!

- Nebula went away (bye-bye!)
- Many lessons learned grew into their own projects
- NASA's IPOST vision arose from the success continuously monitoring under Nebula

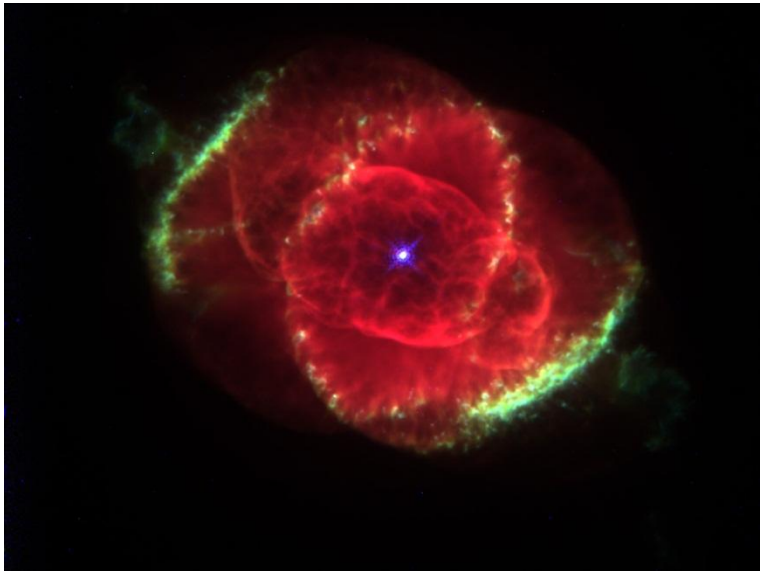

When old stars explode, new stars are born!

The stage is set and the deck stacked

- NASA's environment is challenging.
- It **IS**:
  - Very, very diverse.
  - Intensely focused on missions
  - Extremely geographically disparate
  - Simultaneously cutting-edge and obsolete
- It **LACKS**:
  - Strong chain of command / military discipline
  - Money

Primary Goals

- Our CM tool and data should be:
  - Automated *(duh)*
  - Rapid turnaround (7d internal / 30d maximum)
  - Tactically and Operationally relevant
  - Compelling and Relevant
    - *To the people who can actually fix things!*
  - Useful and Positive, **NOT punitive!**

Design Considerations

- Basic guidelines about how the interface operates:
  - Minimum number of navigational clicks to get to lowest-level of data granularity (shoot for 4)
  - Static URLs for any data view – allows email/IM conversations to include “click this to see what I’m talking about”
  - Pointers, wherever possible, to sources of authority for findings

Our Data Sources

- Entries are composed of multiple, correlated data sources:
  - Internal Data (Vulnerability information, detailed stuff)
  - External Data (Self-discovery of exposure to outside world)
  - Operational data (Self-awareness of responsibility, policy, etc)
  - Intel data (Information about threats, problems, etc)

Our Data Sources

- Key IDEAL Characteristics for sources of data:
  - Open data formats (Text, XML, etc)
  - Direct, API access
  - Open-source whenever possible

Internal Data

- Nessus/Foundstone* data
  - Produces detailed vulnerability information about hosts on the network
  - Basis for CVSS scores
    - This data is fairly sound
- KACE/Patchlink data
  - Produces detailed information about patch status for hosts on the network
  - Basis for “patch status” scores
    - (this data can be flaky)
- ** Foundstone is a secondary data source*

External Data

- Nessus external scanner
  - Dumps entire DNS tree,
  - Scans all ports on all hosts from external posture
- TCP/Netflow data
  - Data mine netflow across all borders to determine server/service relationships
    - (SYN/ACK, RST outbound)
- Google Search Results (Search API)
  - Search for site:youragency.gov,
  - Mine results into array

Operational Data

- Full dump of DNS trees (associate w/ hosts, build CNAME relationships)
  - Note: This discovers (*most*) cloud hosts too!
- Dump of asset database(s)
  - Associate sysadmin & security POC wherever possible (fallback to sysadmin)
- Dumps of DHCP logfiles + MAC association

Intel Data

- Datamine IR (Incident Response) tools and database to look for hostnames, sysadmins
  - Increase risk factor for repeated findings
- Datamine threat sources for specific vulnerable services
- Attach risk based on “best effort” adjudications for platforms/targets of frequent attack

You've got risk scores!... Now what?

- Building a risk scoring tool is only half the battle.
- Now how do you get people to use it?

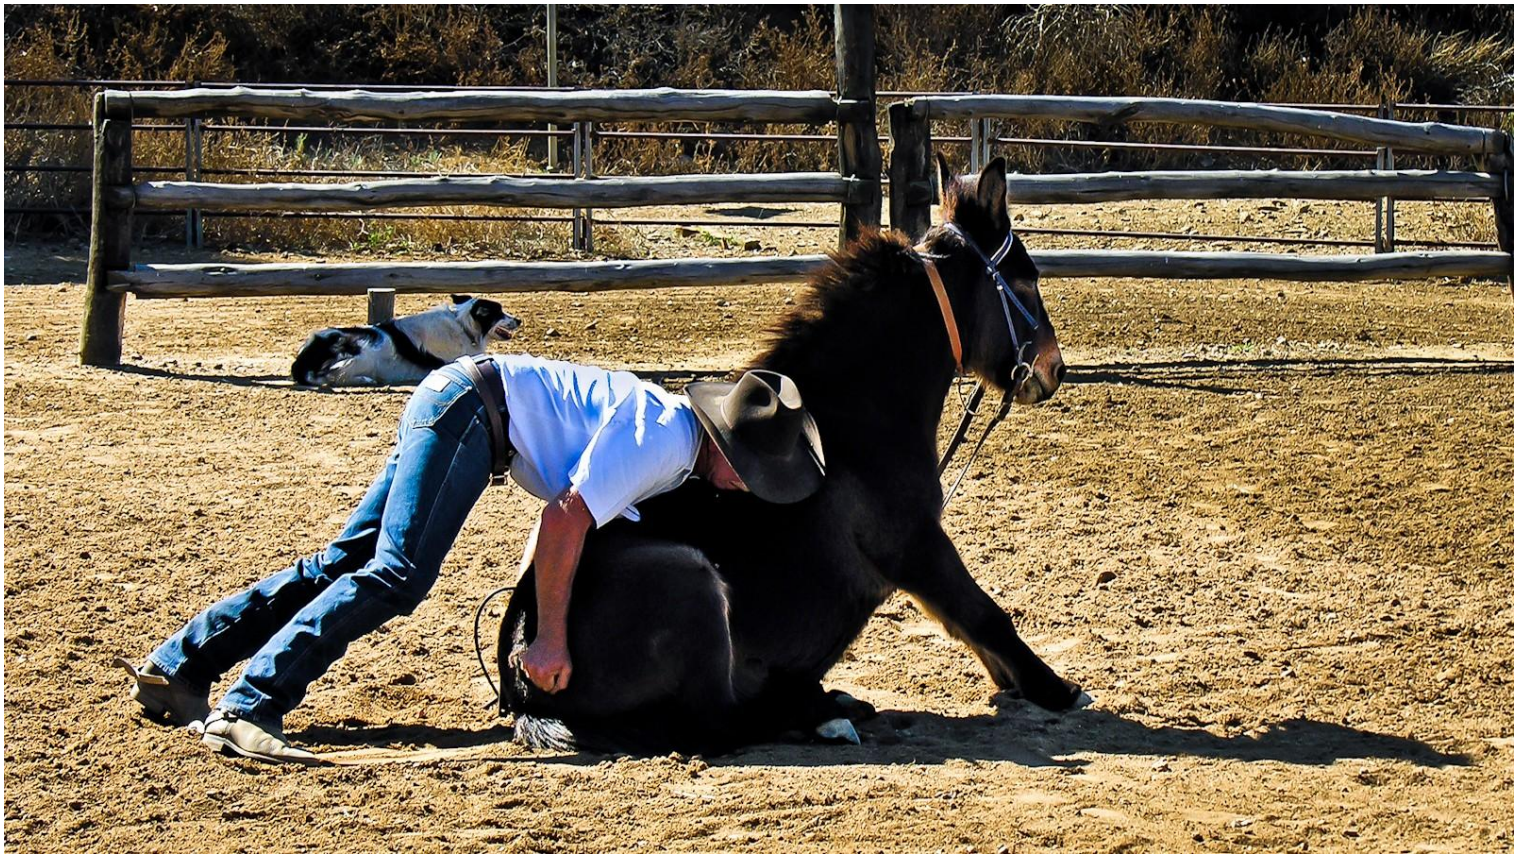

A NOTE ABOUT EGO:

- EGO and PRIDE are not deadly sins, they're mental API calls to your security and technical staff.
- The internet itself runs on ego.
- Geeks work because they're proud of being good at their work.
  - Root nameservers: Volunteer effort.
  - White-hats & Black-hats – EGO? Check.
  - Torvalds? Zimmerman? De Raadt? All positive examples of ego shaping the industry.

How to harness this power

- One key thing we're targeting is harnessing the power of pride, ego and social engagement.
  - Sysadmins see each others' scores
  - Presented in gamer-style interface
  - Status awards/trophies
    - "Superhero" award (positive)
    - "pwnie express" award (negative)

How to harness this power

IT Security Dashboard | Admins | Centers | Hosts | Reports | mlinton

System Administrators

Notes

This is a view of all sysadmins ordered by mean weighted CVSS score. Mean weighted CVSS scores is defined as total weighted CVSS score over number of hosts. Total weighted CVSS score is defined as the sum of the most recent weighted CVSS score for each of the sysadmin's hosts.

AUID is something we hope have access to in the future. For now, that column just contains a unique identifier.

Org information is also currently unavailable.

Administrators

AUID	First	Last	Center	Org	Total Score	Mean Score
3226	Je	B	GSFC		159.0	79.5
3204	M	M	GSFC		70.5	70.5
3010	M	M	GSFC		50.2	50.2
3216	Er	Cl	GSFC		126.6	42.2
3325	Al	C	GSFC		36.5	36.5
58	Sl	D	MSFC		91.5	30.5
2171	Da	Cl	ARC		322.6	29.3
2928	St	Fa	GSFC		142.9	28.6
2980	Ar	Fc	GSFC		26.8	26.8
3285	Ti	W	GSFC		21.7	21.7

How to harness this power

IT Security Dashboard | Admins | Centers | Hosts | Reports | mlinton

System Administrator

Chri

Center: ARC

Org:

Total Weighted CVSS Score: **173.4**

Number of Hosts: **31**

Mean Weighted CVSS Score: **5.6**

High Impact Vulnerabilities

Port	Service	Details	Instances	CVSS Score	Impact
443	www	+	11 – [redacted].gov [redacted]nasa.gov [redacted].nasa.gov [redacted]nasa.gov	6.4	31.46

Live Demo

Insert Browser Stuff Here

Initial Outcomes at NASA

- During Pilot phase at ARC and GSFC:
  - First week, a dozen sysadmins given access
  - 8/10 of the top “worst” hosts were immediately mitigated (had been thought decommissioned but weren’t)
  - Second week, few dozen SAs given access
  - Scores on particularly bad hosts dropped very quickly as tool use spread – mostly vulns / issues that weren’t known to the SAs.

Feedback from Sysadmins

"*I hate to admit that I saw my name flash up on your screen as a problem person. Most of my computers are extremely old and I inherited most of these problems I'm sure. If I need assistance in correcting some of the issues, is there a recommended source for assistance?*"

-- We referred this SA to a sysadmin user group where they got a lot of good advice.

Feedback from Sysadmins

Regarding access to the tool for a group with large amounts of systems:

"*This will be a really quick an easy way to monitor our systems, and CSA's. Great idea!*"

Regarding an exposed system with a critical flaw:

"*Sorry, I had NO IDEA that was the case. I will see about putting an abrupt stop to that nonsense. My boss gave his blessing.*"

Current Major Challenges

- Universal adoption – not everyone is open to the idea of “open”. Some stakeholders strongly object to the openness of the tool.
- Fairness – To succeed, stakeholders must be convinced there’s no bias in the scoring system. This can be difficult.
- False positives – need a robust system for dealing with these and keeping them from showing up.
- DHCP/NAT – Consistent attribution of hosts across their various IPs on various dates

Future Challenges

- IPv6 – No longer possible to do “discovery scans” across hosts. Your infrastructure MUST include netflow/network monitoring and aggregating the v6 auto-configuration logs.

Future Directions

- “Scan on Demand” – using APIs provided by vendors to initiate scans on-request for immediate “Did I fix it?” feedback to sysadmins
- “Scan on Connect” – tying DHCP/IPv6 auto-configuration logs to a scan initiation (pulled from Nebula) – will help attribution and potentially can kick hosts off if they fail
- More score-based and status-based situational gaming for the sysadmins, CSOs and users.

Obligatory last slide

- Email: matt@nasa.gov
- Twitter: mattatnasa

- BBS: Take time machine to 1991, connect to RenegadeBBS. *(remind me to buy Apple stock)*.