

Malware Analytics at Stream Rate: *Higher Analyst Productivity and Reduced Threat Exposure*

8/21/2012

Dr. Harold Jones
harold.jones@baesystems.com



Agenda

- Problem Statement
- Looking Under the Hood
 - Real-time Forensics
 - Stream Analytics
 - Analyst's Dashboard
- A Good Day in the Life of An Analyst
- Summary

PROBLEM STATEMENT

Three Mutually-Supportive Concepts

- **Real-time Network Forensics**

- Discover and reconstruct event history in seconds
- Implement an informed remedy before damage escalates



- **Malware Analytics at Stream Rate**

- Detect that you have been had (or, *maybe* just targeted 🤖) sooner
- Focus analyst attention on the biggest threat 🧑‍🚒 to the enterprise

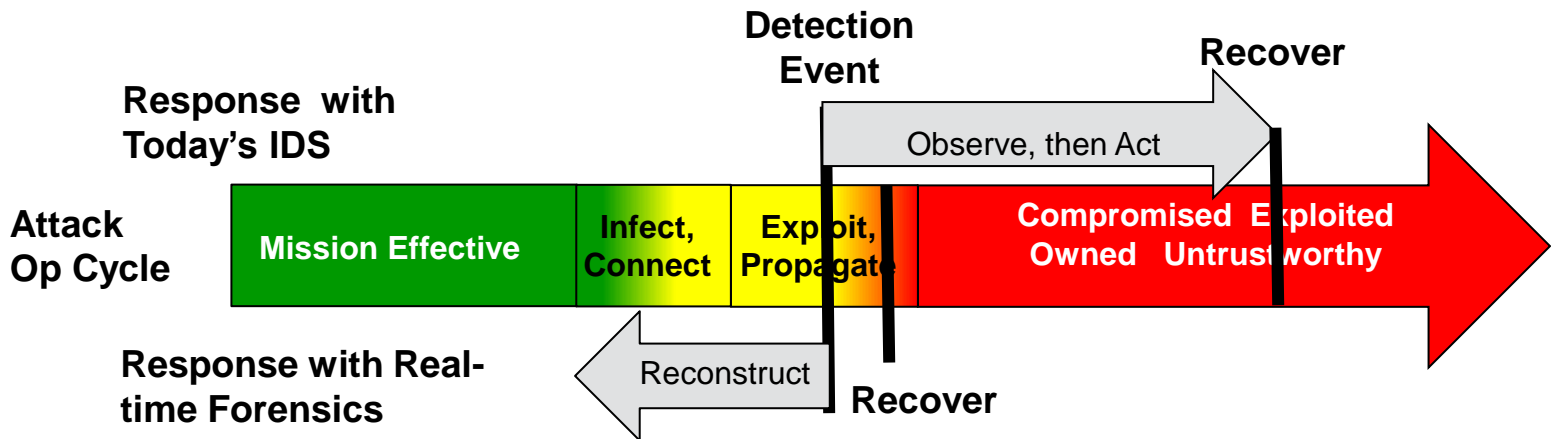
- **Give Analysts Productivity and Skill-Enhancing Tools**

- Open Source to reduce acquisition delays and budget pressures
- Rapid, guided search through historical data caches

Real-time Network Forensics

Dramatically Compressing the Analytic Timeline

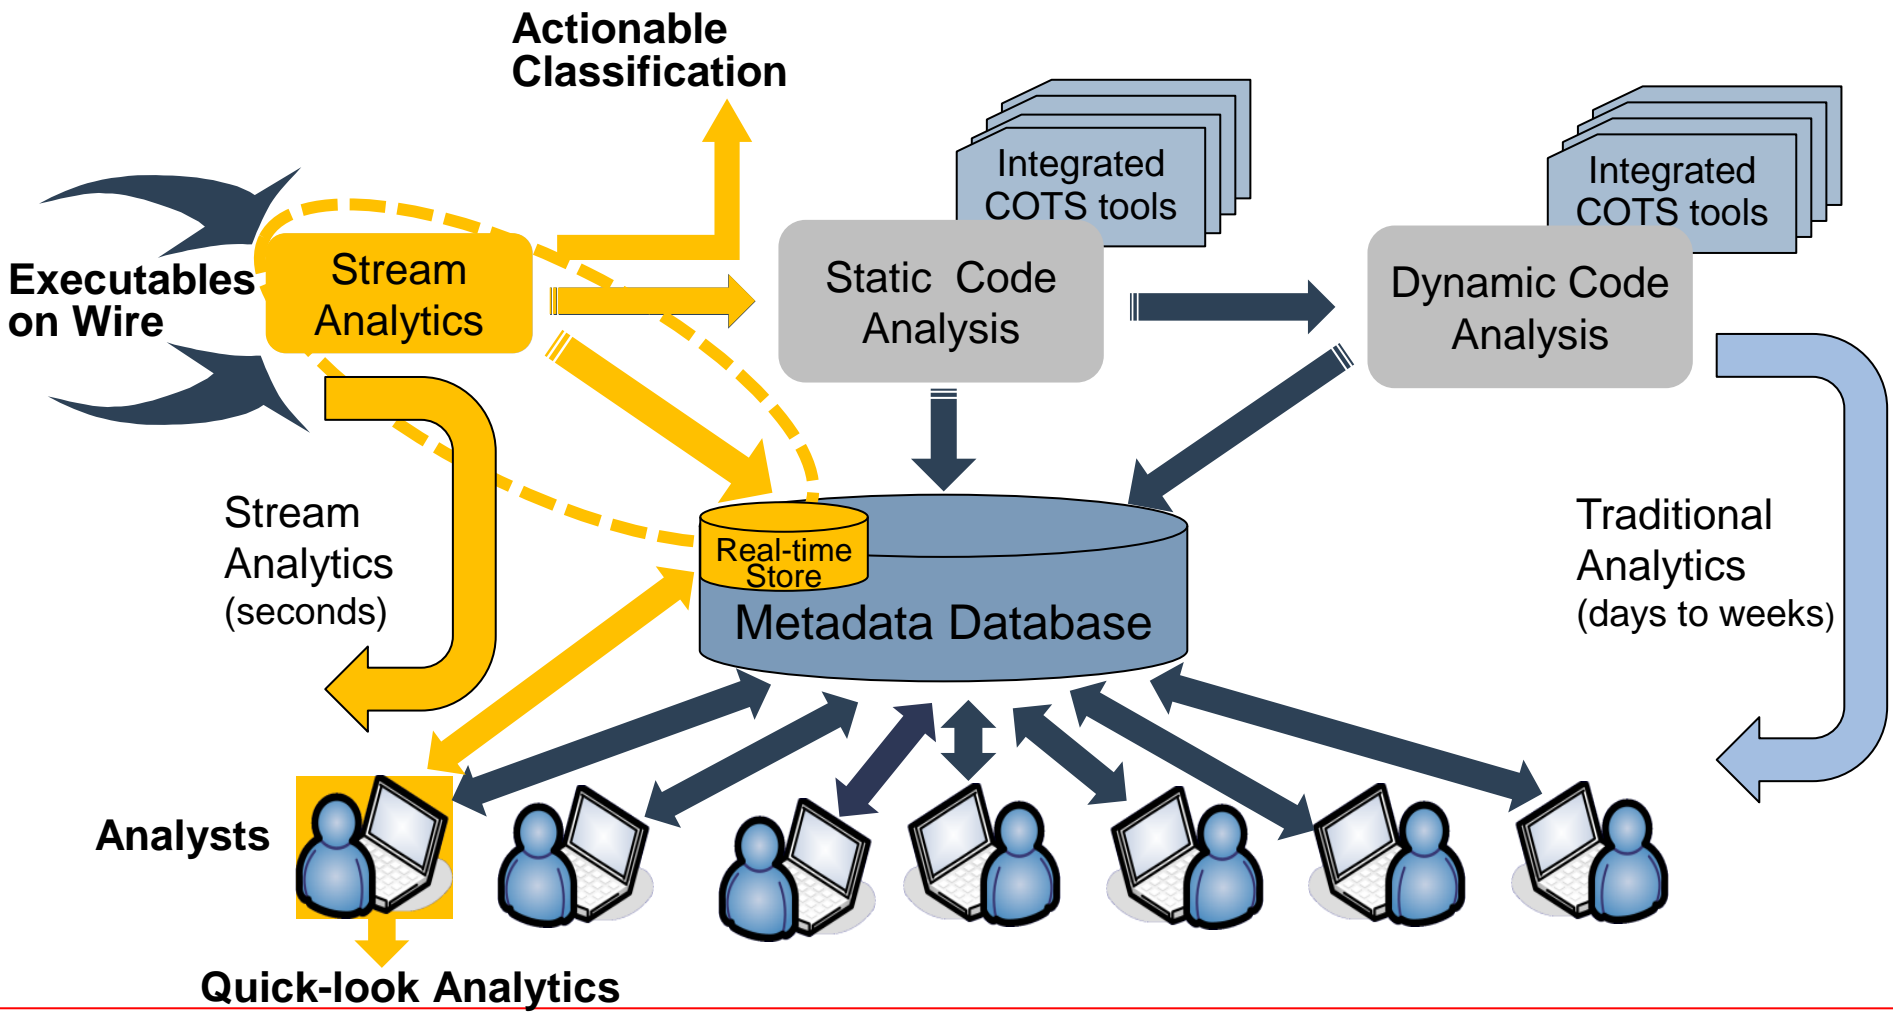
- A worrisome event is detected. What does the SOC team do?
- **Yesterday's Answer:** Observe adversary actions over weeks to map attack scope and ConOps. And sustain added damage.



- **The Network Forensics Answer:** Scroll back in time to reconstruct event history. *In seconds!!* Then act. Fast and decisively.

Stream Analytics As Precursor to Code Analysis

- Code analysis is slow and expensive, but essential for new attacks
- Stream Analytics is fast – actionable analytics for routine attacks in *seconds!*



Real-time, open source network forensics

- **Broad structural anomaly detection**
 - Analyzes every byte of every packet
 - Scales to 10Gps in multi-CPU card
- **Real-time analytics**
 - PEs, e-mail, web sessions...
 - Malware classification w/o unpacking
- **Scalable repository with real-time query via Splunk™ data mining engine**
 - Intuitive, extensible GUI
- **Operationally-validated**
- **Key tools migrating to open source via Suricata™ IDS (2012)**

™ Suricata is a trademark of the Open Information Security Foundation (OISF)



REAL-TIME NETWORK FORENSICS TO HALT INTRUSIONS BEFORE DAMAGE CAN BE DONE

BENEFITS

As the volume and sophistication of cyber attacks escalate, analytic timelines to deal with the onslaught have not kept pace in response. With current tools, analysts spend weeks or longer deciphering the scope of an incident, after which the ops staff has to clean up the destructive consequences. The net effect: analyst time is used inefficiently and the enterprise suffers unacceptable losses and needless damage. The NIDAR-FX suite of tools dramatically compresses the analytic timescale, enabling analysts to complete their assessments in minutes – not weeks. The benefit to the enterprise is dramatic and two-fold: **minimized operating costs and minimized risk of damage.**

Our Approach
NIDAR-FX is a real-time, software-based network forensics solution.

Behavioral Anomaly Detection
Our behavioral anomaly detectors extract powerful indicators of even the most sophisticated attack in real-time.

Stream Analytics
Stream analytics organize and classify potential system risks to generate high-confidence indicators of malicious activity in real-time.

Data Repository
Optimized with real-time search capabilities, our data repository places both current and historical data at the analyst's fingertips.

Intuitive Graphical Interface
Developed for analysis, by analysts, our NIDAR-FX Dashboard provides a clear understanding of the nature and scope of an incident within minutes.

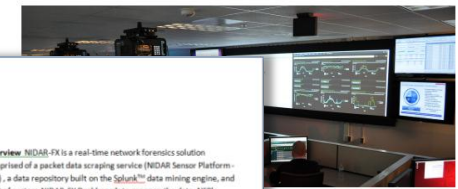
Broad protocol coverage makes it difficult for even sophisticated attackers to avoid detection

Statistical correlation yields high-confidence results

Faster analytic response time reduces enterprise losses

Higher analyst productivity lowers enterprise costs

Forensic search enables full discovery of attack scope in real-time



NIDAR-FX Dashboard Guide

1st Revision



Overview NIDAR-FX is a real-time network forensics solution comprised of a packet data scraping service (NIDAR Sensor Platform-NSP), a data repository built on the Splunk™ data mining engine, and a set of custom NIDAR-FX Dashboards to manage the data. NSP's mission is to collect session and behavioral anomaly data across a broad swath of protocols – its scope and design is covered in a companion document. This guide is intended to provide users with a visual guide to the pre-built dashboards that come with the standard NIDAR-FX deployment.

For each mission area, a summary Analyst's Dashboard provides a view of network traffic trends and the riskiest network exchanges in the selected time/address/content space. The full set of NIDAR-FX drilldown dashboards is intended to be used both as an alerting mechanism for likely network compromises and to facilitate decisive, analyst-driven investigations.

Section A - Site Traffic Analytics: contains screens to support high-level traffic analysis (pg. 3)

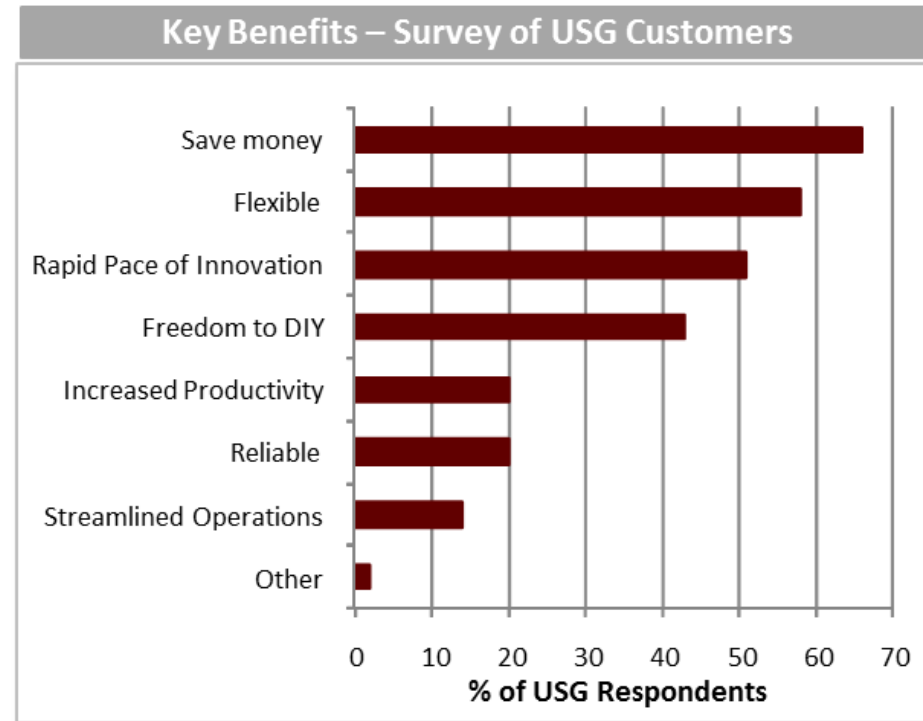
Section B - Portable Executables: specifically designed to support the rapid analysis of Portable Executables (PEs) as they transit a network (pg. 15)

Section C - Site SMTP Analytics: designed to support SMTP analytics as they might pertain to targeted phishing attacks (pg. 31)



Why Open Source? Why Suricata?

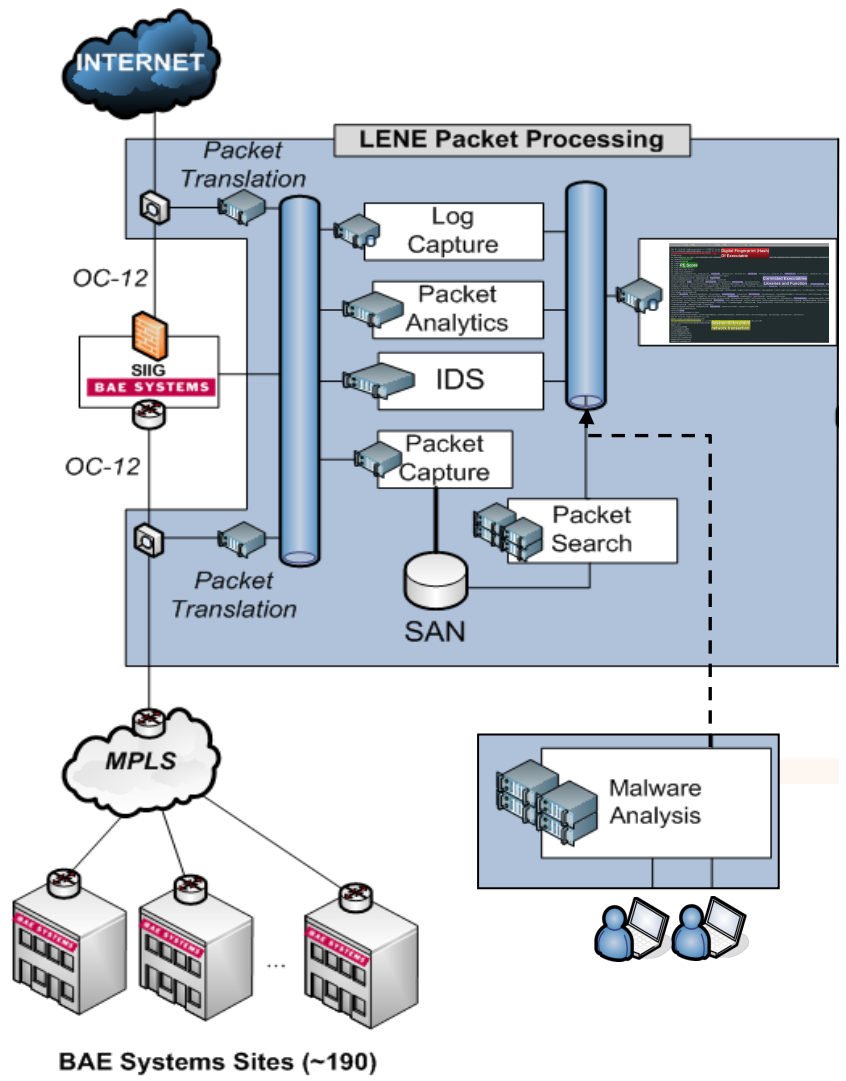
- **USG adaptation of OS is significant ...and growing**
- **Supportive policies beginning to emerge**
- **Why Suricata?**
 - A next-gen, multi-threaded IDS. Ideal for large pipes/data flows
 - Maintained by Open Information Security Foundation (OISF)
 - Sponsored by DHS/DS&T Homeland Open Security Technology (HOST)
 - Funding from DHS and SPAWAR
 - GPL license



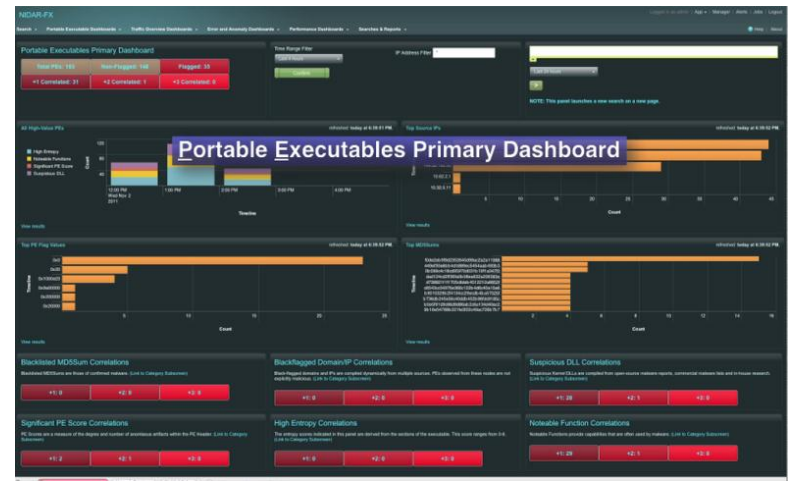
*Source: Market Connections and Lockheed Martin survey
"The Intersection of Open Source and Cloud" May 2011*

Validation at Operational Scale

NIDAR-FX and Suricata In BAE's Leading Edge NOSC Environment (LENE)



- **Free-range testing:** Ops “in the wild” on BAE’s 50K node network
- **Seeded testing:** Artifacts from Malware lab injected into flow
- **Real-time or batch testing**
- **Quick-look console**



LOOKING UNDER THE HOOD:

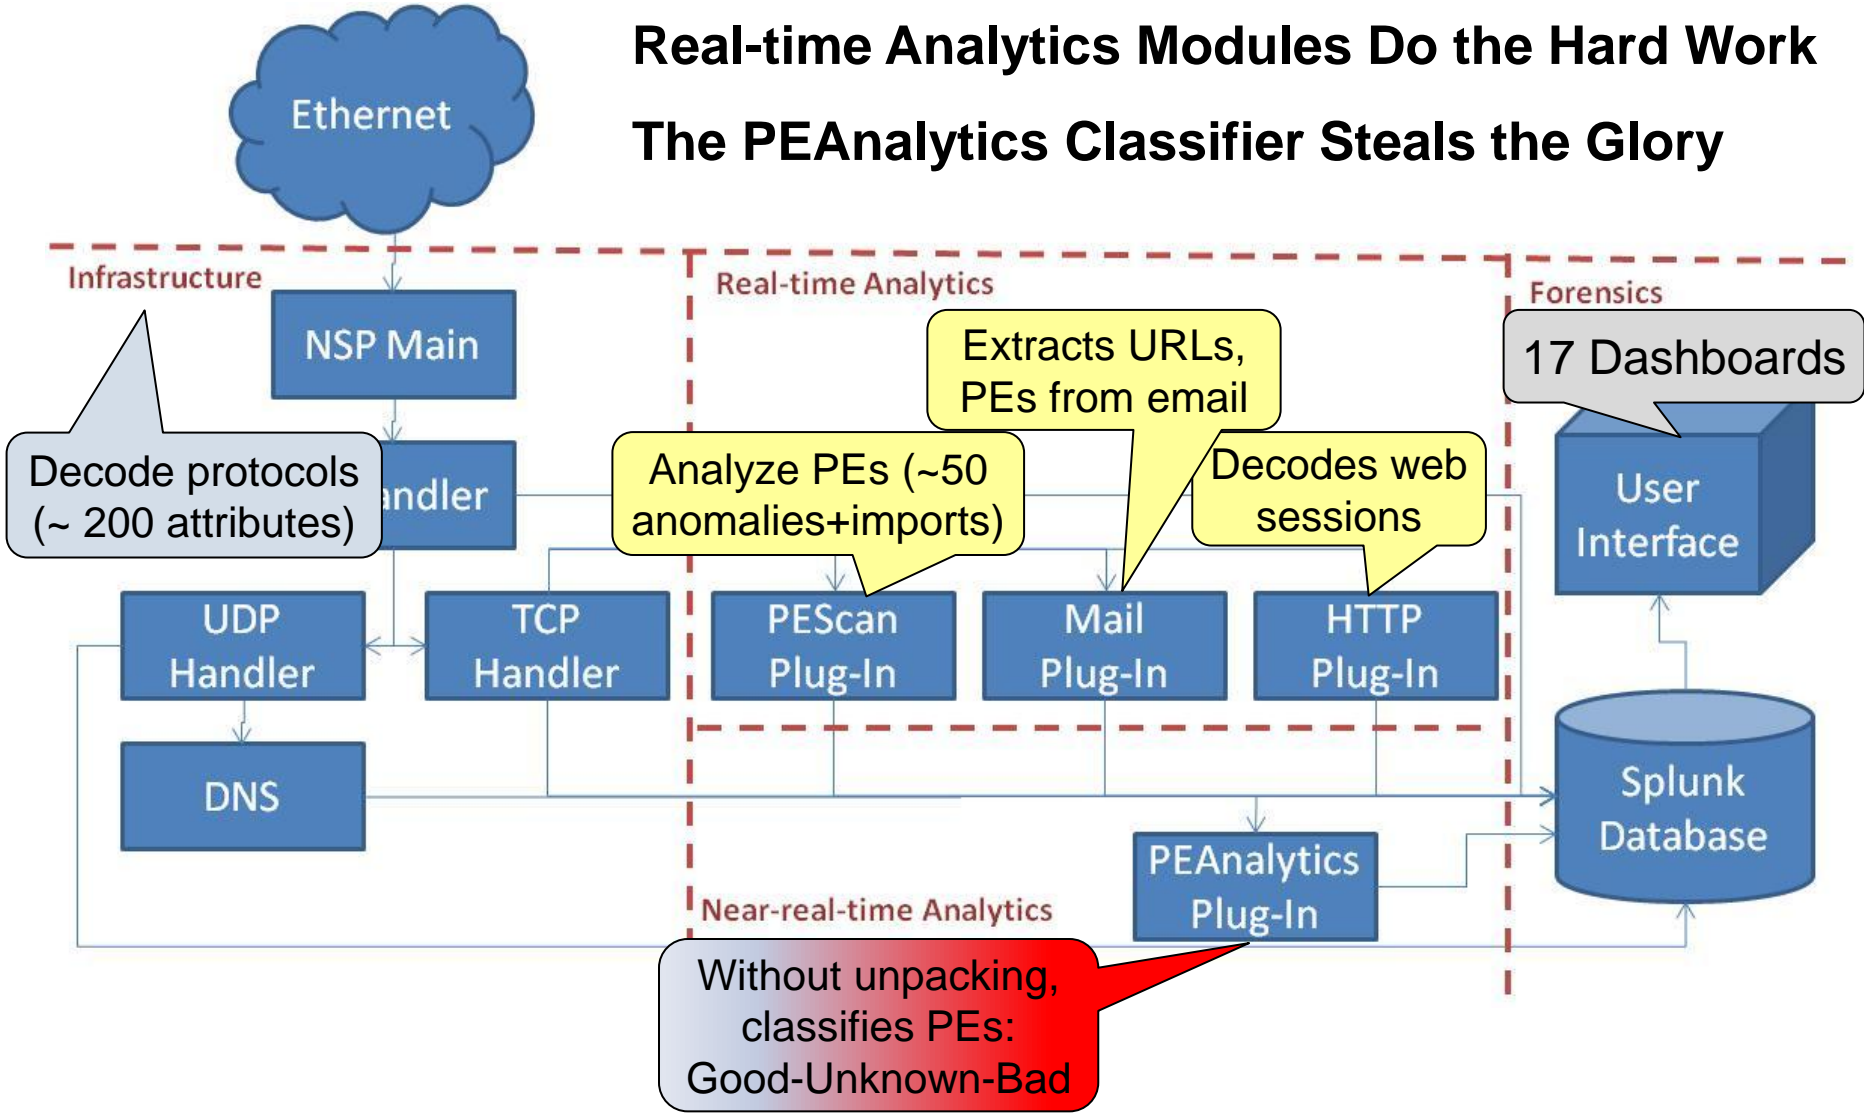
REAL-TIME FORENSICS

STREAM ANALYTICS

ANALYSTS' DASHBOARD

NIDAR-FX Architecture

Real-time Analytics Modules Do the Hard Work The PEAnalytics Classifier Steals the Glory

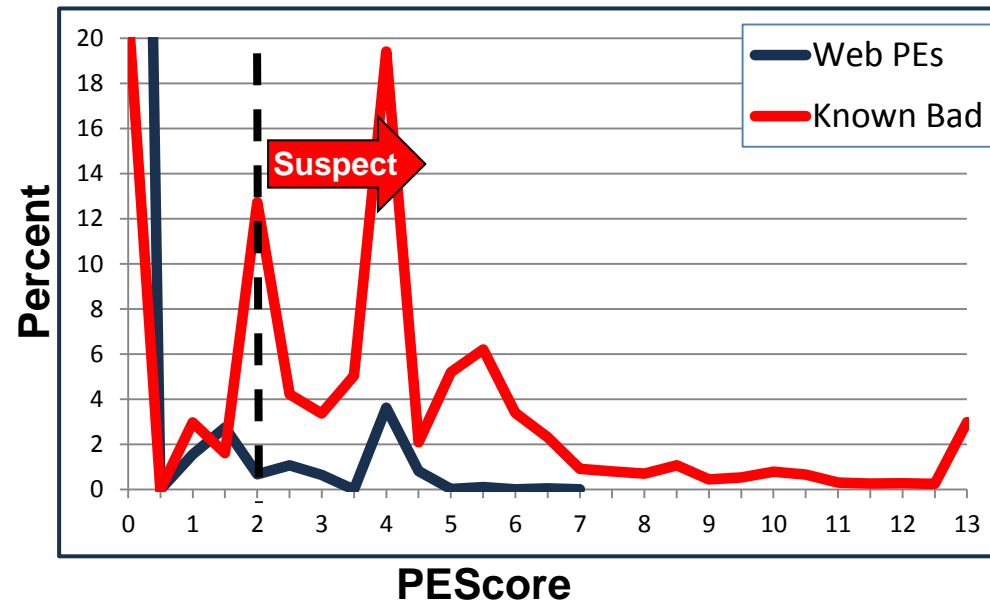


PE Attributes Separate Malware from Goodware

- **NIDAR-FX collects 3 attribute types:**
 - *Structural* – # of sections, privileges...
 - *Numerical* – Load point, offsets....
 - *Imports* – libraries, function calls...
- **Example: Structural anomalies in ~ 80% of malware (and some goodware)**

Executable sections with write privileges...
- **763K malware samples ⇒ 6000 unique anomaly patterns**

Anomaly Severity (PEScore) Separates Malware From “Goodware”



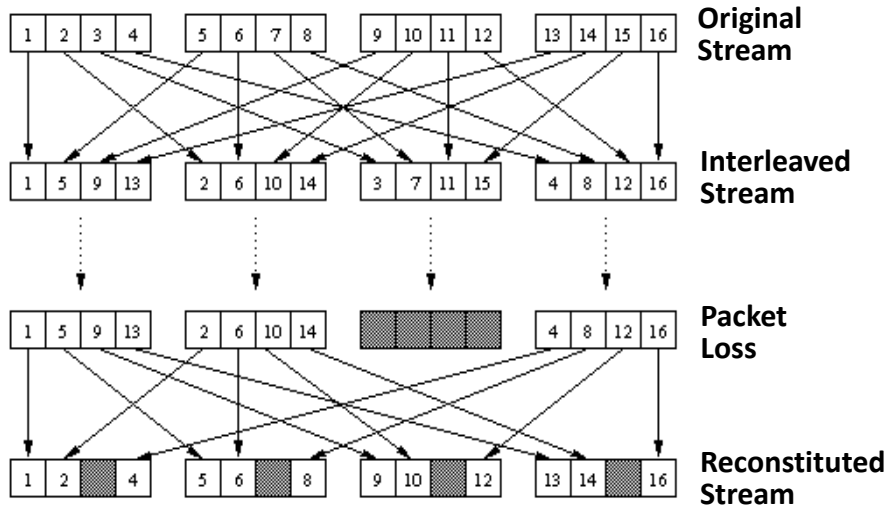
LOOKING UNDER THE HOOD:

REAL-TIME FORENSICS


STREAM ANALYTICS

ANALYSTS' DASHBOARD


Extracting Executables from Packet Flow



- Out-of-order packets
- Interlaced with other messages
- Transmission errors

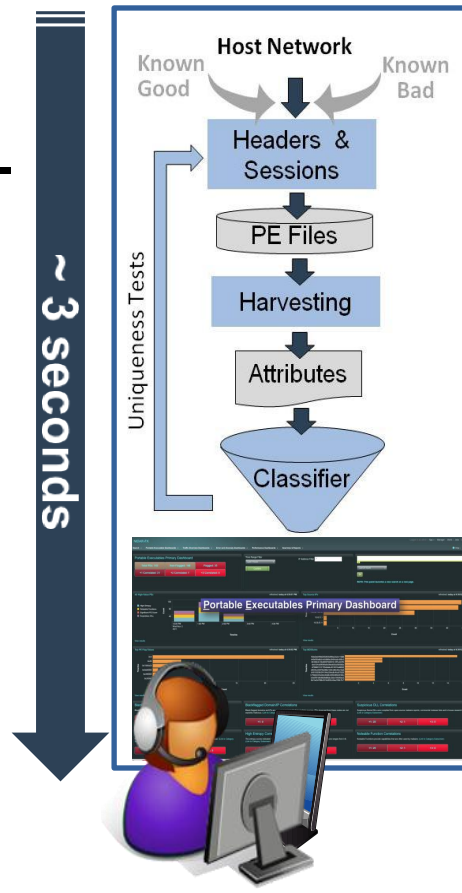


Windows Portable Executables (PEs):

- Microsoft PE 8.x format, but  take liberties
- Header exposes meta-data and structural anomalies
- Obfuscation is common
- Session “fragments” can be intentional

Extracting Attributes from Executables

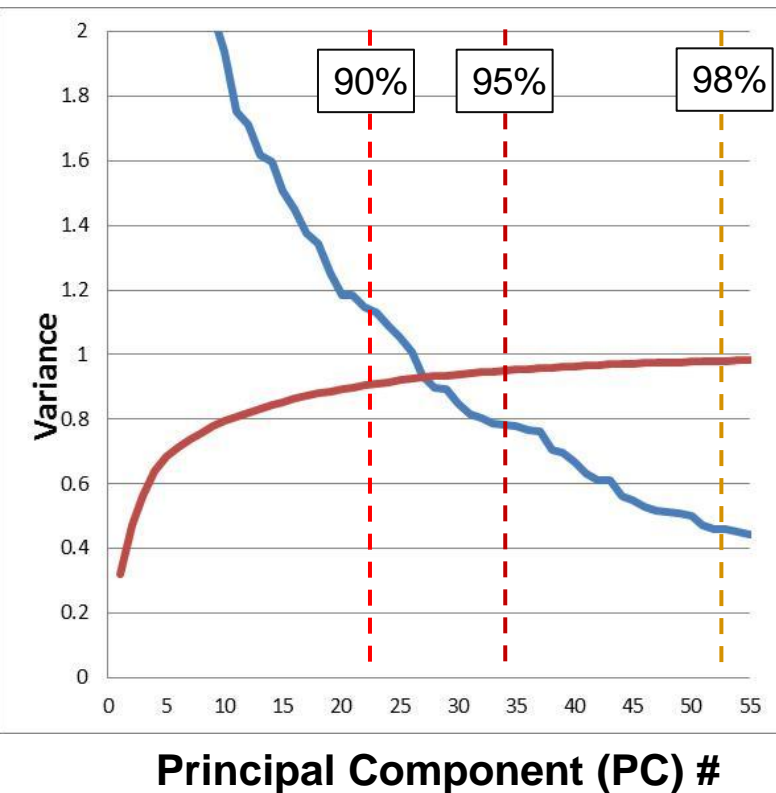
- **Attribute classes – structural, numerical, imports – define available features**
- **Classifiers need “right-sized” attribute set:**
 - *Too few:* loss of separability
 - *Too many:* over-training
 - *Too precise:* countermeasures
 - *Too hard:* encryption/unpacking
- **Without unpacking, PEAnalytics classifies PEs as:**
 - Benign ■ Malicious–known type ■ Unknown–high risk



Alternative classifiers (Naïve Bayes, random forest...)

Which Attributes Really Matter?

Principal Component Analysis (PCA) of 50K malware samples

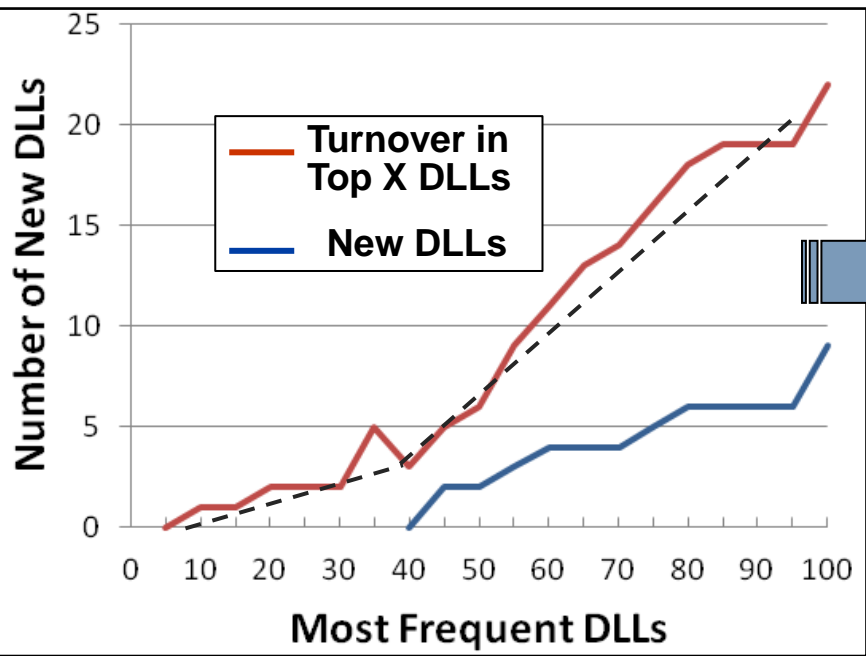


- 227 attributes (anomalies and DLLs)
- PCA suggests 22 (90%) to 53 (98%) independent (composite) variables
- But, malware families exhibit clusters of DLLs and anomalies:
 - 22 PCs → 135 “core” attributes
 - 53 PCs → 210 “core” attributes
- So, how many attributes do we need?

Population drift analysis is major factor in answering “how few”


Population Drift Impact on Classifiers

Attack Fingerprints Change (DLL Drift 3/11-3/12)



So Classifier Performance Erodes


Classifier Accuracy	Day 1	Day N
True Positive	97.6%	See below
False Positive	0%	↗




Current experiments are measuring classifier decay rate

Automating Attribute Selection and Training


All Attributes




Reduced Attribute Set



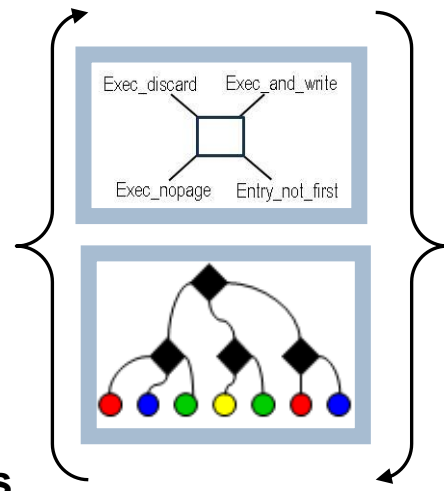
PCA*



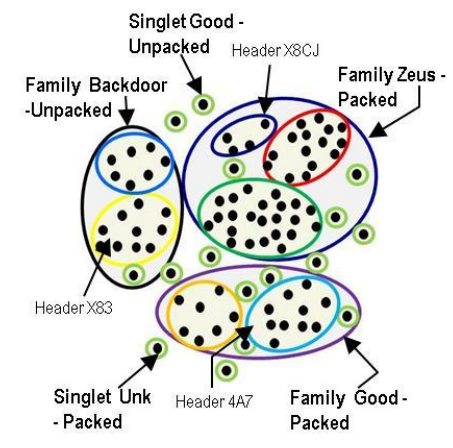
9x% variance



Classifiers



Class Separation

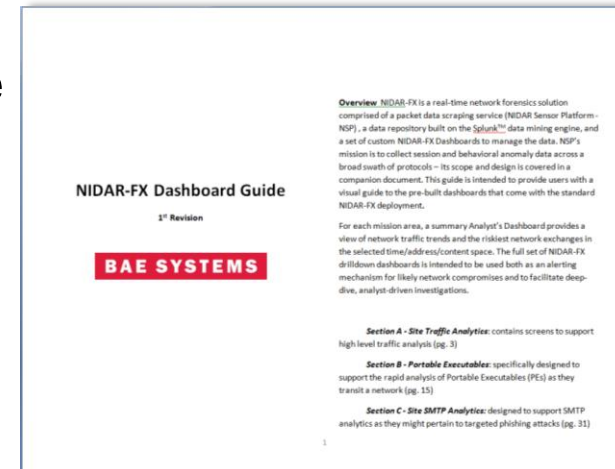


* Principal Components Analysis

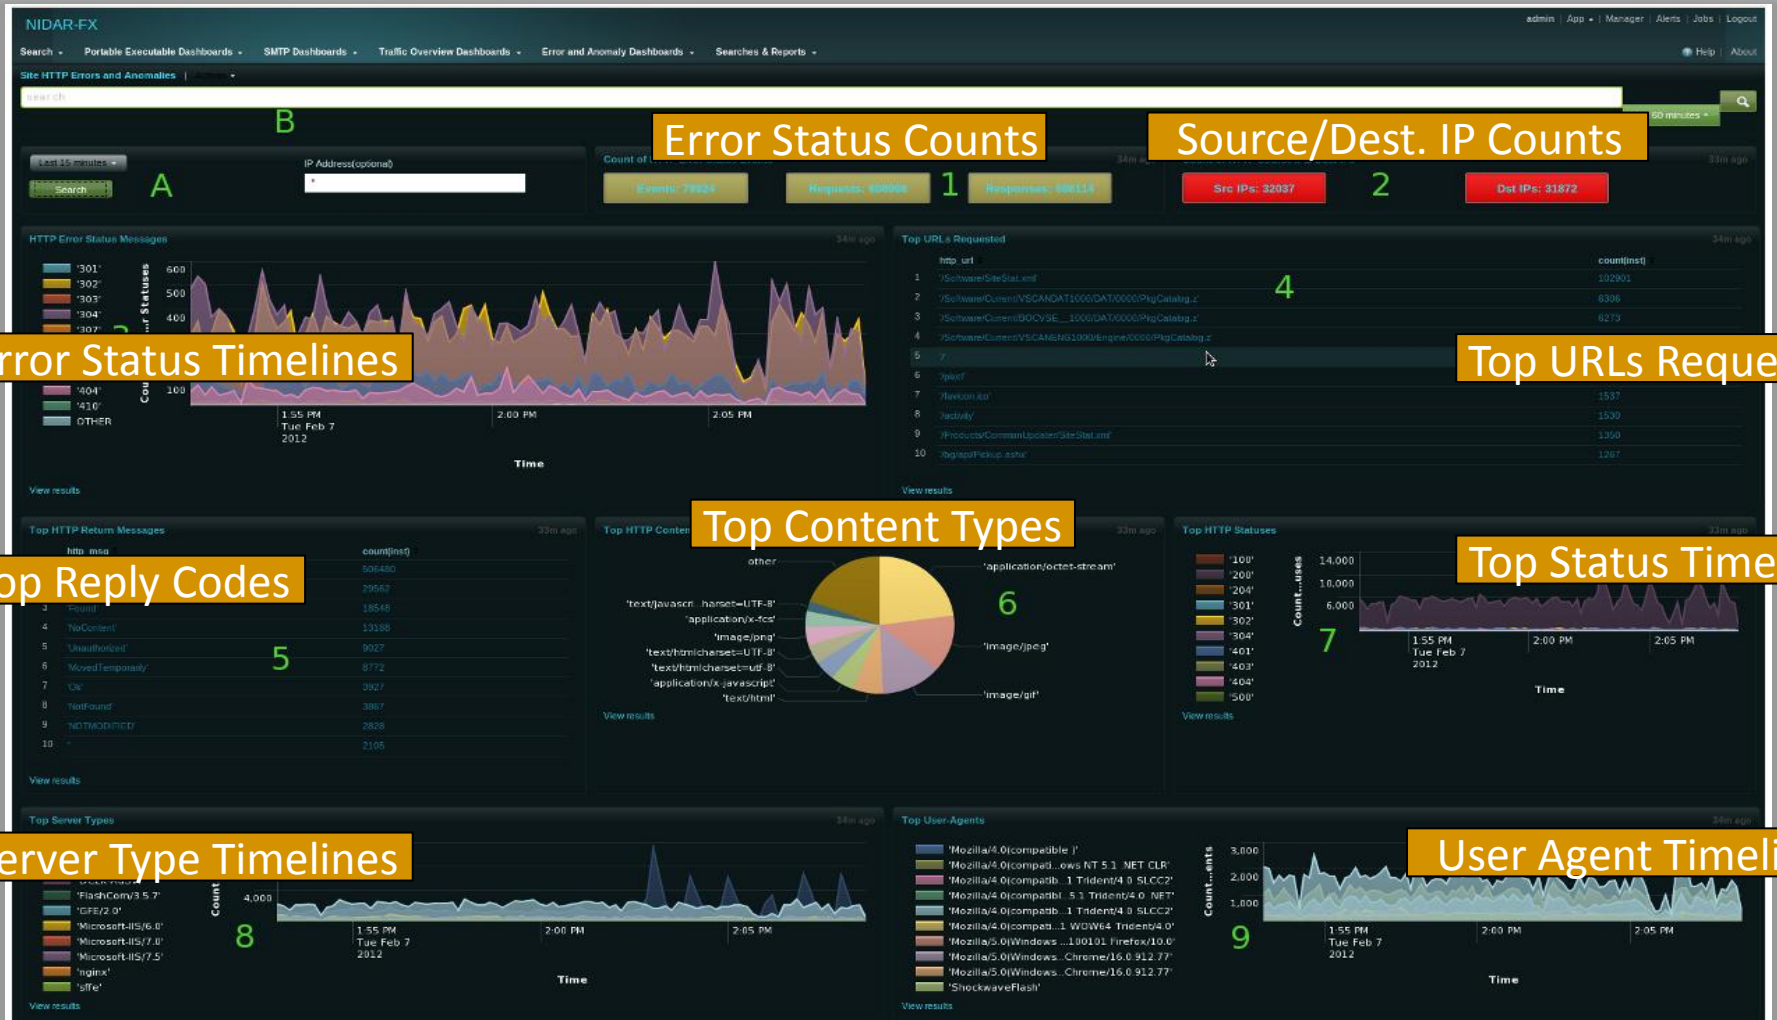
LOOKING UNDER THE HOOD:

REAL-TIME FORENSICS STREAM ANALYTICS ANALYSTS' DASHBOARD

- Built on Splunk™ data mining engine
- 18 standard screens, with drill-down
- Composable queries
- Toolkit for new screens



HTTP Error Summary Dashboard



Error Status Timelines

Error Status Counts

Source/Dest. IP Counts

Top URLs Requested

Top Reply Codes

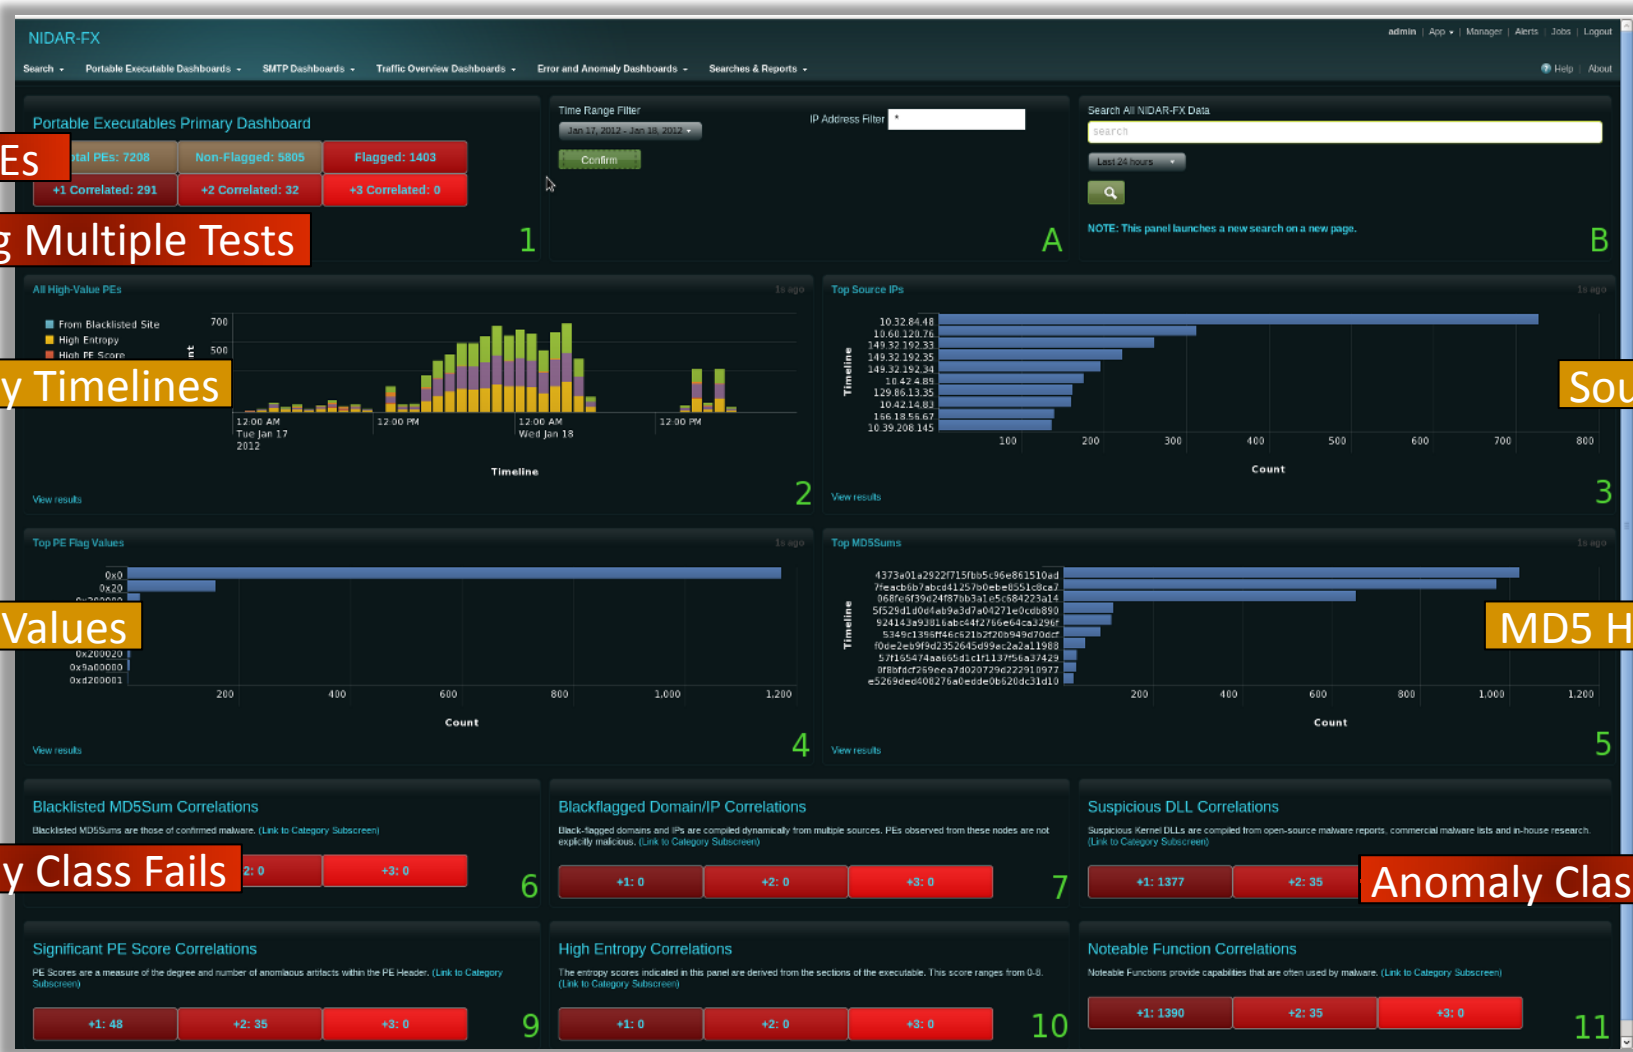
Top Content Types

Top Status Timelines

Server Type Timelines

User Agent Timelines

Malware Analytics Dashboard



Total PEs

Failing Multiple Tests

Anomaly Timelines

PE Flag Values

Anomaly Class Fails

Source IP

MD5 Hashes

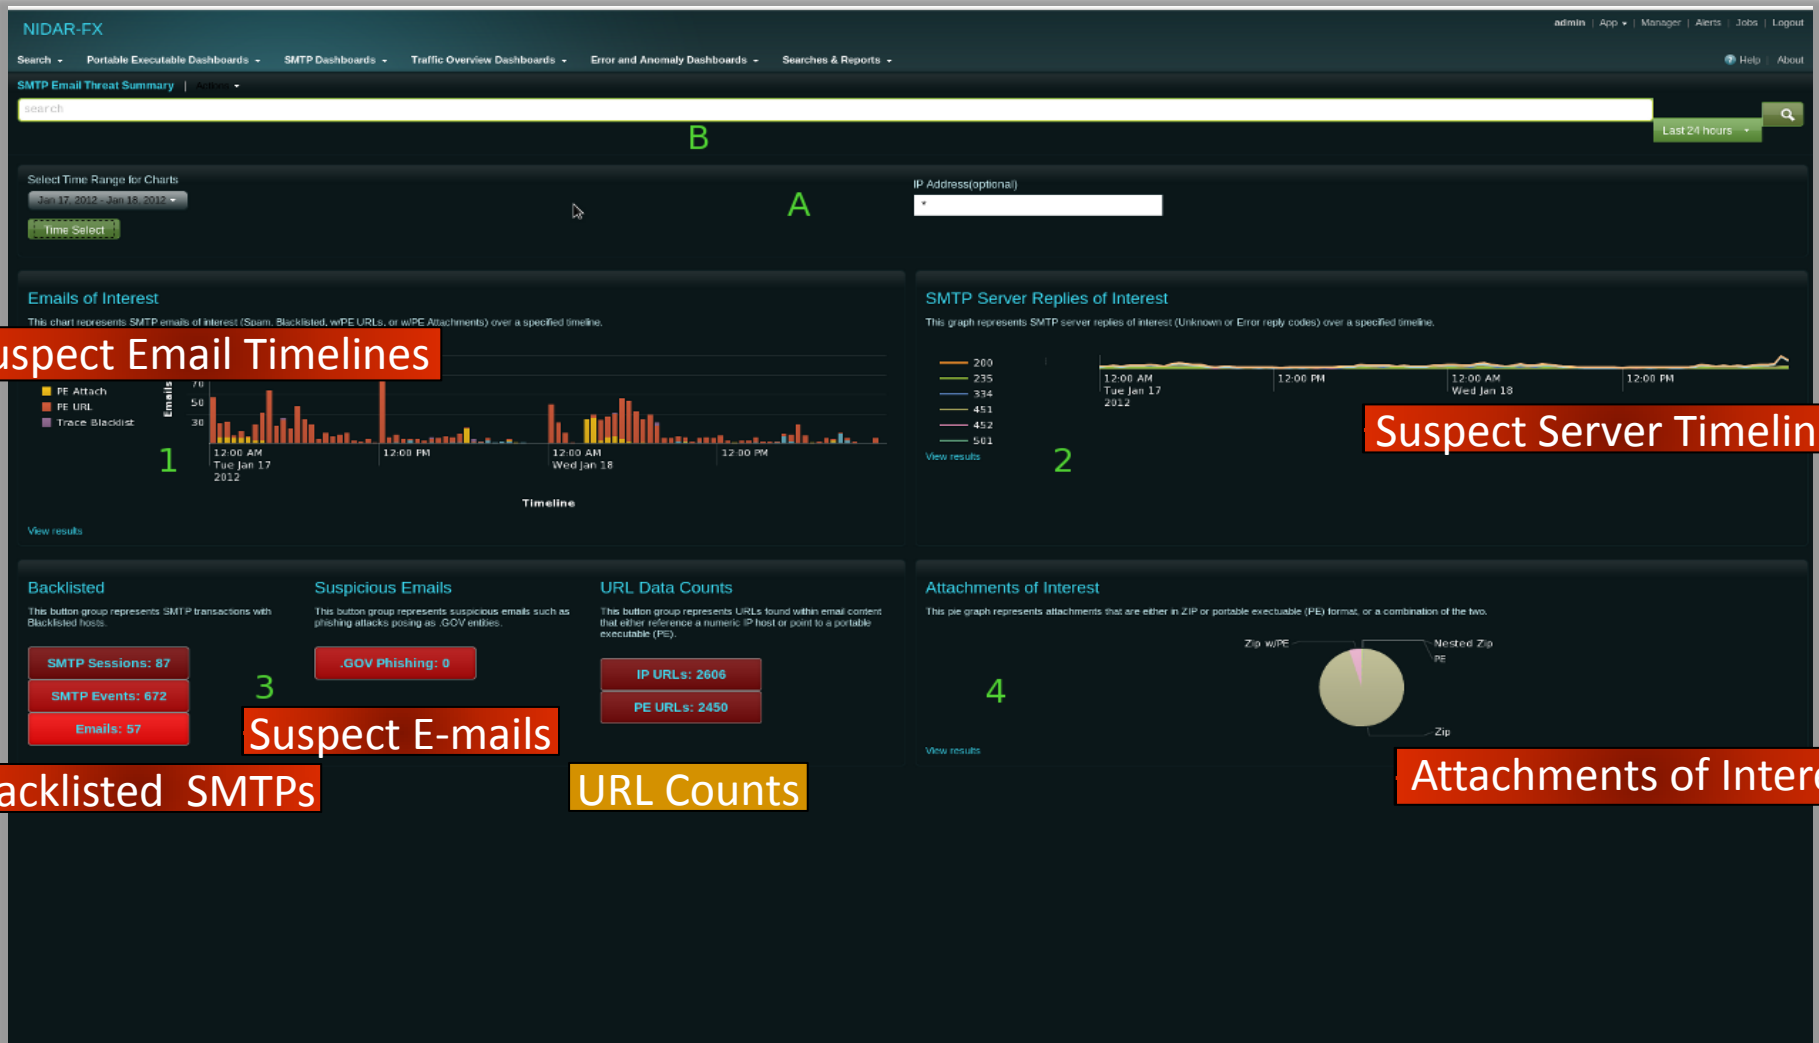
Anomaly Class Fails

Metadata Drill-down

Supporting evidence for the Tier-3 analyst

Nov 02 14:32:45 croda pe_printer-v1.1.2[20515] rec-pe
tcp_sid=9520c021f0084843061d0050 src_ip=149.32.192.33
pe_md5_sum=b736db245e56c40db452b86fdc8185c
blnd5=False
pe_hash=d69ce11a pe_family=010240001000142001440060000
pe_id=None pe_score=4
pe_flags=0x9a000000
pe_flagco
pe_flag_b
pe_flag_too_mully_sect
pe_flag_entry_not_first=2
pe_flag_vsize_expands=2
pe_import_dlls=oleaut32.dll, advapi32.dll, user32.dll, kernel32.dll, kernel32.dll, user32.dll, kernel32.dll, advapi32.dll, comctl32.dll, kernel32.dll, advapi32.dll, oleaut
oleaut32=SysFreeString, SysReAllocStringLen, SysAllocStringLen
advapi32=RegQueryValueExW, RegOpenKeyExW, RegCloseKey
user32=GetKeyboardType, LoadStringW, MessageBoxA, CharNextW
kernel32=GetACP, Sleep, VirtualFree, VirtualAlloc, GetSystemInfo, GetTickCount, QueryPerformance
Counter, GetVersion, GetCurrentThreadId, VirtualQuery, WideCharToMultiByte, MultiByteToWideChar, Is
GetModuleFileNameW, GetLocaleInfoW, GetCommandLineW, FreeLibrary, FindFirstFileW, FindClose, ExitProcess, UnhandledExceptionFilter, SetLastError, RaiseException, G
kernel32=TlsSetValue, TlsGetValue, LocalAlloc, GetModuleHandleW
user32=CreateWindowExW, TranslateMessage, SetWindowLongW, PeekMessageW, MsgWaitForMultipleObjects, MessageBoxW, LoadStringW, GetSystemMetrics, ExitWindowsEx, DispatchMessa
CharToOemW, CallWindowProcW
kernel32=WriteFile, WaitForSingleObject, VirtualQuery, VirtualProtect, VirtualFree, VirtualAlloc, SizeofResource, SignalObjectAndWait, SetLastError, SetFilePointer, SetEver
ResetEvent, RemoveDirectoryW, ReadFile, MultiByteToWideChar, LockResource, LoadResource, LoadLibraryW, LeaveCriticalSection,
InitializeCriticalSection, GetWindowsDirectoryW, GetVersionExW, GetUserDefaultLangID, GetThreadLocale, GetSystemInfo, GetStdHandle, GetProcAddress, GetModuleHandleW, GetM
GetLocalTime, GetLastError, GetFullPathNameW, GetFileSize, GetFileAttributesW, GetExitCodeProcess, GetEnvironmentVariableW, GetDiskFreeSpaceW, GetDateFormatW, GetCurrentPr
InterlockedExchange, InterlockedCompareExchange, FreeLibrary, FormatMessageW, FindResourceW, EnumCalendarInfoA, EnterCriticalSection, DeleteFileW, DeleteCriticalSection, C
CreateDirectoryW, CompareStringW, CloseHandle
advapi32=RegQueryValueExW, RegOpenKeyExW, RegCloseKey, OpenProcessToken, LookupPrivilegeValueW
comctl32=InitCommonControls
kernel32=Sleep
advapi32=AdjustTokenPrivileges
oleaut32=SafeArrayPtrOfIndex, SafeArrayGetUBound, SafeArrayGetLBound, SafeArrayCreate, VariantChangeType, VariantCopy, VariantClear, VariantInit
Nov 02 14:32:45 croda pe_printer-v1.1.2[
0515] rec-pe rectype=pe_section_info
tcp_sid=9520c021f0084843061d0050 src_ip=149.32.192.33 dst_port=80
pe_md5_sum=b736db245e56c40db452b86fdc8185c
Name=.text
VirtualSize=82900
VirtualAddress=4096
SizeOfRawData=82944
PointerToRawData=1024
PointerToRelocations=0
PointerToLinenumbers=0
NumberOfRelocations=0
NumberOfLinenumbers=0

E-mail Threat Summary Dashboard



A GOOD DAY IN THE LIFE ON AN ANALYST

NIDAR-FX Example:

Spearphishing Forensics

A fine July morning. The IT staff is kicking back



At noon, a .gov customer informs us of an e-mail hoax under his letterhead.

The hoax is directive in nature, and contains malicious links.

The reported attack was blocked at the BAE Systems gateway, but:

Who was targeted? Why?

Were we penetrated by an earlier/later phase?

If so, were we compromised?

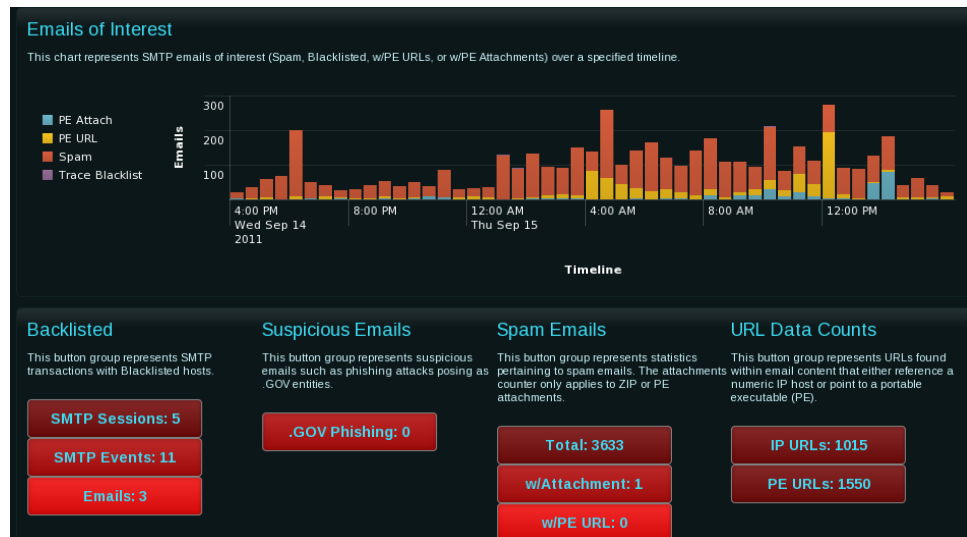

Was any sensitive info at risk?

First hr: Scrolling Back in Time To Discover “What”

NIDAR-FX Real-time Forensics

By 2 PM, our analysts had:

1. Found 18 suspect emails received between 4:20 and 5:53AM
2. Determined URL's are inactive
3. **Verified URL's weren't visited!**
4. Searched TCP session of 4:20 email. Found *[invalid address]* error code
5. Retrieved 4:20 e-mail TXT
6. Searched error codes from '.gov'. Found similar e-mails from other agencies!
7. Packaged data for follow-up



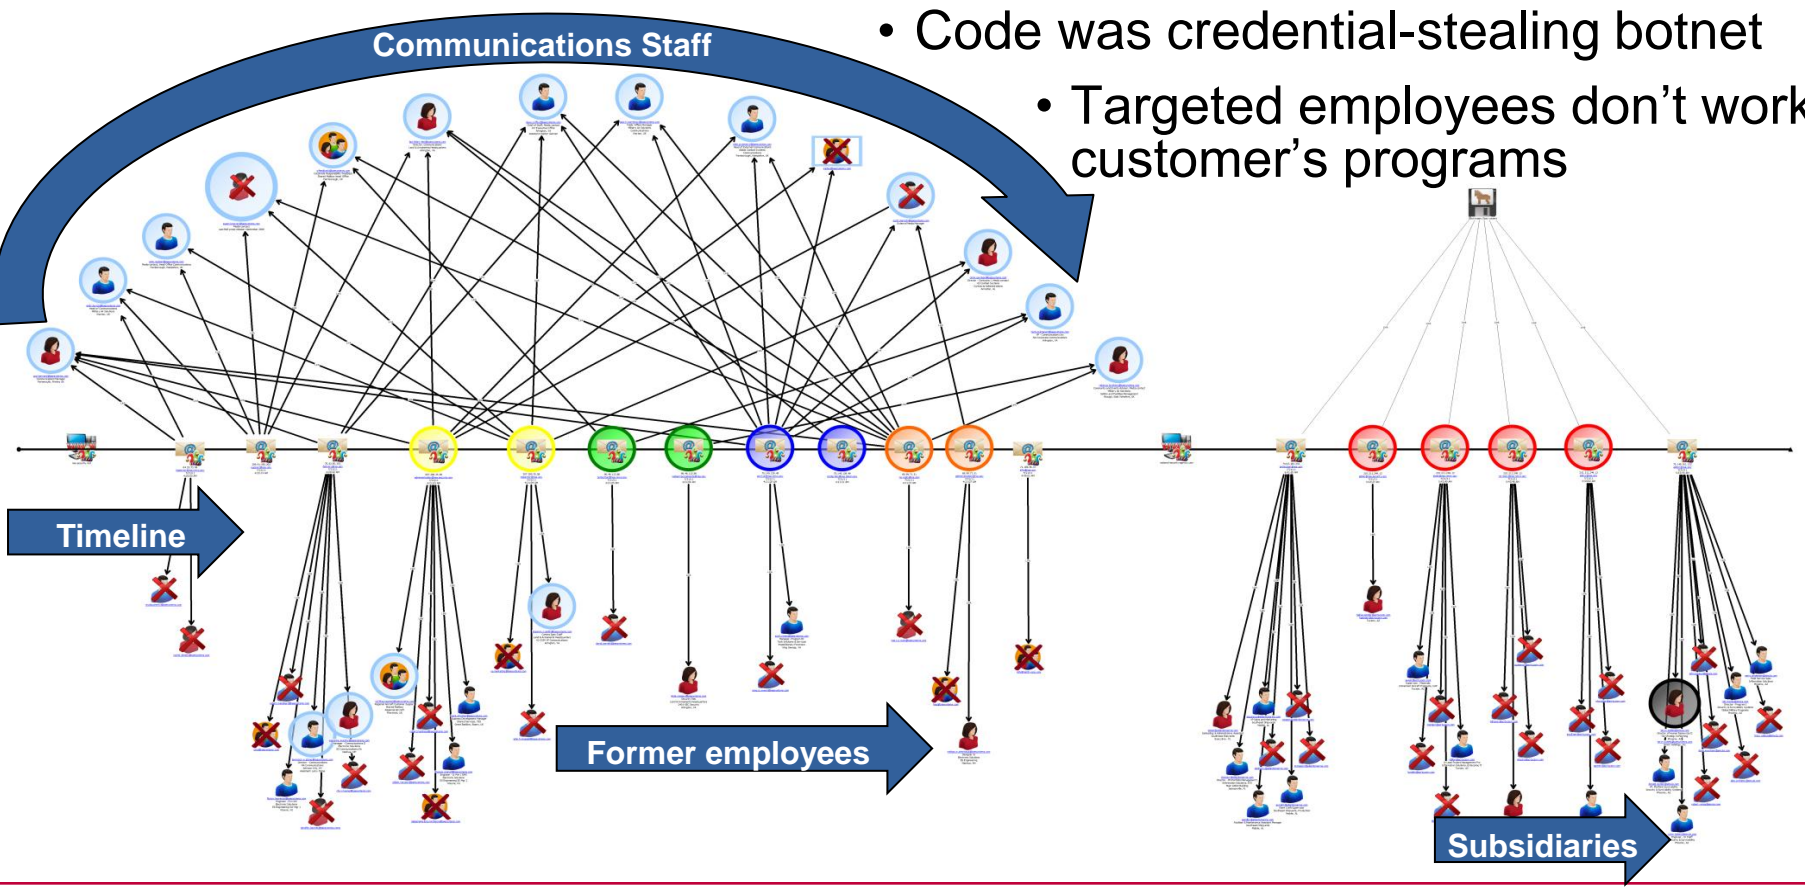
AND, begun deciphering the attack

Understanding “How” and “Why”

Adding Global Threat Analytics to Event Processing

Within 48 hrs, our analysts determined we were AOK:

- E-mail part of broad gov-theme spam campaign
- Code was credential-stealing botnet
- Targeted employees don't work customer's programs



SUMMARY

Summary

- **Real-time forensics is real... and dramatically accelerates malware analysis. BAE System's toolset available soon at <http://www.openinfosecfoundation.org>**

- **Accurate, real-time, sustainable classification is feasible.**

Population drift can be measured and managed


Automating the process is doable

- **Classifier design will ultimately matter. But first, we have to understand our data.**

More productive analysts

Faster Understanding of attacks

Less sustained damage



QUESTIONS?

Dr. Harold Jones
harold.jones@baesystems.com