tmforum
inform innovate accelerate optimize

Sharing Threat Intelligence to Mitigate Cyber Attacks

Architecture Proposal

GB976

Version 0.8

November, 2013

IPR Mode: RAND

Notice

Direct inquiries to the TM Forum office:

240 Headquarters Plaza,
East Tower – 10th Floor,
Morristown, NJ 07960 USA
Tel No. +1 973 944 5100
Fax No. +1 973 944 5110
TM Forum Web Page: www.tmforum.org

Table of Contents

List of Tables and Figures

Executive Summary

This document was created by expert members of the Threat Intelligence Sharing Catalyst project for Management World 2013. Its purpose is to both educate TM Forum members on the cyber security problem and explain why sharing threat intelligence is essential for getting ahead of today's sophisticated threats.

This paper examines every aspect of sharing threat intelligence – the winners and losers, problems associated with sharing, standards that provide support and it proposes architecture.

This document has gone through a second set of reviews by new members of the team.

Members of the team appreciate your interest in this topic, and welcome feedback on the contents of this paper.

1. Focus on the Problem

1.1. Threat Intelligence Sharing

SANS states that Threat Intelligence (TI) can enable defenders to establish a state of information superiority which decreases the adversary's likelihood of success with each subsequent intrusion attempt. Threat intelligence can be a force multiplier as organizations look to update their security programs and defenses to deal with increasingly sophisticated advanced persistent threats. Security managers need accurate, timely, actionable, and detailed information to continuously monitor new and evolving attacks with methods to exploit this information in furtherance of an improved defensive posture. Make no mistake about it: co-existing computer network defense contains a strong element of intelligence and counterintelligence that analysts and managers alike must understand and leverage[1].

The recent MNE7 project[2] has stated:

> "There is currently a gap in our ability to generate sufficient collaborative national and international situational awareness across the cyber domain. This generates a requirement for a generic and comprehensive framework that details the processes for generating sufficient collaborative national and international cyber situational awareness. Cyber situational awareness is required across the public and private sectors, and at a national and international level, if we are to truly provide a comprehensive and integrated response to sophisticated threats that can operate below the detection thresholds of many existing processes and systems."

In spheres of defensive action, military or commercial, collective action has proven to be both more operationally effective and cost effective for the individual entities in the collaboration. The imperative for Cyber Security is to scale the collaborations needed, and rapidly, to give the defenders the upper hand compared to the attackers. Whilst threat sharing is a starting point, and degrees of vulnerability, exploit and solution sharing occur in narrow communities today, broadening these to be the baseline for all businesses is essential to raising the threshold against attacks; ultimately to secure global commons[3] for safe and secure business in cyberspace for all.

Organizations across many industry sectors are discussing how to share Threat Intelligence; those with experience have found sharing to be very beneficial. For instance, the organizations in the financial industry in the United States have been able to assist each other by sharing information to enable protective defenses against detected threats through the Financial Sector Information Sharing and Analysis Center (FS-ISAC). The telecommunication sector have been sharing information to track down, mitigate, or stop threats including DDoS, mail abuse, and other domain specific threats for many years. In some cases, the sharing is ad-hoc through email or portals and in other cases specific data formats and protocols are used to enable the exchanges, such as the Messaging Malware

[1] http://computer-forensics.sans.org/blog/2013/02/11/sans-cyber-threat-intelligence-summit-22-mar-2013
[2] http://mne.oslo.mil.no:8080/Multinatio/MNE7produk/35CyberCon/file/3.5%20Concept%20of%20Employment.pdf
[3] http://www.act.nato.int/mainpages/globalcommons

Mobile Anti-Abuse Working Group (M3AAWG) who use the Abuse Reporting Format (ARF)[4] to mitigate mail abuse problems, the Anti-phishing Working Group (APWG) that uses the Incident Object Description Exchange Format (IODEF)[5] to mitigate phishing and eCrime attacks, and various communities such as the Financial Services Information Sharing and Analysis Center (FS-ISAC) that utilize the Structured Threat Intelligence eXpression (STIX)[6] for conveying full-spectrum threat intelligence information. Threat intelligence may also be provided by vendors or service providers using proprietary data formats, portals, or email.

Information sharing has proven successful within industry groups as well as in cross-industry groups within specific domain areas such as mail abuse and instant messaging, from which lessons can be learned and applied more broadly. Implementations of both IODEF and STIX are emerging, and will continue to grow at an accelerated pace due to industry pressures to consume intelligence through automated mechanisms. Improved standards based sharing models, collaborative analysis tools, and the ability to provide accurate, context rich, directed, and actionable intelligence delivered broadly in a timely manner is imperative to prevent today's sophisticated attacks[7].

We know that the adversary is determined, skilled, and fast. Recent estimates indicate that attackers are changing their TTPs (tactics, techniques, and procedures) every six months or less. To make things worse, there is a critical shortage of skilled resources capable of performing TI analysis. Furthermore, many companies fear consequences – political, financial, insurance, and other, associated with admitting they have experienced a cyber-attack, so they prefer to remain discrete.

The directed and actionable sharing of threat intelligence within an eco-system is necessary to leverage the skill sets of cyber experts broadly, while enabling exchanges through brokers for discrete sharing. With directed and actionable intelligence sharing, automation of mitigation controls becomes possible through vetted information directly updating devices or host monitoring services on an organizations network, eliminating the need for advanced resources at every organization..

1.2. How do we 'change the game' for sharing threat intelligence?

Responding quickly to cyber threats is of critical importance, and the only way to really change the game is to understand how the adversary works and predict where they might go next. We know that we can't win this battle by fighting alone, and while progress has been made to share Threat Intelligence within affected communities, improvements must be made to automate the exchanges and where possible the implementation of protective defenses. How do we change the game and become more effective at sharing than the adversary? We think the answer rests in having diverse analytic capabilities tied to an ecosystem that enables the direct and actionable sharing of intelligence information resulting in automated defensive monitoring and/or mitigation actions.

[4] RFC5965 An Extensible Format for Email Feedback Reports. Y. Shafranovich, et. al. Aug 2012.
[5] RFC5070 Incident Object Description Exchange Format. R. Danyliw, et. al 2007.
[6] http://stix.mitre.org
[7] "Provide timely, relevant situational awareness of potential threats, attacks, network status, and other critical information to support decision making for defense of DOD information networks including tailoring to support CC/S/A missions and operations." Chairman of the Joint Chiefs of Staff Instruction, Information Assurance and Support to Computer Network Defense (CND), CJCSI 6510.01F, 9 February 2011

What if we applied the crowdsourcing model as a type of analytic center? Applying a crowdsourcing approach to Threat Intelligence would involve being able to assemble an impromptu, virtual army of trusted cyber defenders to more quickly and comprehensively understand the threat and predict where they will go next. It would require the ability to create dynamic relationships with trusted sharing partners who have common threat interests, and being able to register and receive notifications when threats change. By transitioning today's more static "sharing" model to be a more dynamic "crowdsourcing" approach for Threat Intelligence, we could actually improve response times and predict attacks, leveraging the skill sets of analysts from many organizations. Furthermore, it is possible that an established, successful crowdsourcing Threat Intelligence solution could serve as a deterrent for cyber adversaries. The benefits of this model could result in improvements in shared data out to other aspects of the ecosystem to quickly protect against current threats.

How can the data be quickly disseminated and be immediately actionable through a broader ecosystem? Building from the crowdsourcing example, assume a dynamic group has come together using advanced analytic capabilities and identified a threat with high confidence and were assured of the accuracy of that data. In today's sharing models, each of the participants would leverage that data and implement monitoring or defensive protection measures mostly through manual tasks. If we progress our sharing models, we can leverage standards to consume and process that data for use within each of those organizations, improving the response times to implement mitigation measures. Although that helps the organizations with skilled resources that can participate in the information sharing communities, creating an ecosystem where actionable threat information can be disseminated more broadly to where it can be most effective is necessary to change the game. How do we create this ecosystem?

There are numerous types of analytic or sharing centers today, many with differing and complementary capabilities such as crowdsourcing analyzing data across threat vectors or more focused capabilities such as network, malware, or mobile threat analysis. Some of the analytic centers are sharing communities, while others are part of service provider offerings that further validate information and direct actionable vetted data to customers through threat intelligence feeds. Moving towards an architecture that supports standard data formats and protocols as inputs and outputs for data exchange between analysis centers or service providers and products will enable an effective ecosystem.

It has become clear that an ecosystem including advanced analytic capabilities and multiple sharing models is necessary to communicate as effectively as the attackers. Organizations within and across sectors will benefit greatly from the ability to share data that is critical for their cyber and other threat defenses. In addition to large organizations with many resources and access to centralized threat analytic centers, small and mid-sized organizations will greatly benefit from an agile ecosystem with sharing capabilities specific to their needs. The threat intelligence feeds from existing vendors could improve with the dissemination of shared knowledge, analysis, and vetting for more efficient (lower number of total resources to assess shared information) and effective malware detection, network threats, email abuse, or other types of intrusion prevention.

Another important aspect of this model is the ability to visually present meaningful, actionable data that is tailored to the needs of different users or groups within an organization or community. Vendor feeds can come from a variety of different tools and data sources creating a need to unify these disparate sources of information and in some cases take raw data and create new visualizations that aid in quickly understanding the volume or hostility level associated with any given threat. A Threat Intelligence Dashboard should provide role-based access to key security or threat indicators, regardless of the data sources, while at the

same time shielding the users from the potential complexities of the underlying systems and data structures. This type of dashboard should provide faster access to critical information, reduce analysis times, and help make the identification to remediation process more efficient.

1.3. Why it is needed?

Cyber espionage and advanced cyber-crime are no longer just a government problem; they are very real threats confronting anyone with information worth protecting. General Keith Alexander, Director of the National Security Agency and commander of the US Cyber Command, has said that threats are transitioning from "exploitation to disruption to destruction", and one simple, focused attack can have disastrous consequences[8]. Mitigating the threat is hard, time consuming, expensive, and usually done in isolation. This is not a fight that we can win alone. To change the game, we must leverage the force of a community.

1.4. Why does Threat Intelligence need to be shared?

There are several obvious wins that the community gains from sharing Threat Intelligence.

They include:

- Early warning of potential threats
- Lessen time commitment to understand the threat
- Reduce costs to obtain a larger understanding of the threat
- Obtain insights that would not be otherwise obvious
- Connect with other stake holders who are also experiencing the same problem
- Track / measurement of the threat, so that it's possible to explain and articulate the problem to decision makers
- Automate immediate detection and defense capabilities efficiently from trusted sources

1.5. What is being done today? Is it working?

With some exceptions cyber threats today are typically analyzed and tracked manually with various toolsets and multiple data structures, and reported on in various means, informally and formally through serialized reports, tippers, blogs, whitepapers, and emails. Often times, "circular reporting" of the same information results in costly redundant expenditures and poor response times. Sharing is most often done in a peer-to-peer manner and is often manual, squandering any potential to scale intelligence distribution and threat response efforts across large user bases.

There are currently numerous examples of intra-sector sharing where participant members share threat information using traditional methods such as portals, email, or personal

[8] http://www.wilsoncenter.org/event/cyber-gridlock-why-the-public-should-care

discussion. This form of sharing circle is often contact-driven and requires knowing an individual to vouch for formal inclusion into the sharing circle.

In the US, many of these intra-sector communities are currently evolving to leverage well-structured content shared through automated exchange mechanisms all within well-defined scopes as specified using explicit profiles of cyber threat representation languages such as STIX. The FS-ISAC is a pioneer in this area but numerous other sector-specific ISACs including the US National Health ISAC, the US Electricity Sector ISAC, the US Industrial Control System ISAC and others are actively following suit. ISACs and other groups are evaluating their use cases, determining what information is useful to exchange, and then moving to automation with the appropriate specifications and standards once they understand their requirements and sharing model in an effective eco-system.

NATO's Cyber Security Data Exchange and Collaboration Infrastructure (CDXI)[9] platform provides a knowledge management tool with the objective of facilitating automated information sharing. CDXI is following a use case driven approach and will support multiple interfaces to allow the market to drive adoption of standard data formats and protocols that best meet the participants needs.

There are also international operator-driven examples of inter-sector and government-industry sharing that has developed effective and efficient sharing models. The following operator-driven models have a broad impact while leveraging a small set of skilled experts.. The examples include efforts such as those mentioned above by M3AAWG and APWG, as well as the work by the Research and Education Information Sharing and Analysis Center (Ren-ISAC) using automated sharing through the Collective Intelligence Framework (CIF)[10], and the AbuseHelper network among European CERTs.

Government efforts are assisting the organizations within their country; examples include the US Department of Homeland Security NCCIC and CISCP programs and the UK Ministry of Defense's CISP program.

Some sectors are limited in what can be shared in terms of data classification, and sharing authorities are limited to one way sharing versus a preferred bi-directional or distributed sharing model. However, in most sharing circles, there is no structure for Threat Intelligence and no standard model for describing its data. The broad adoption of open and internationally accepted standards in an ecosystem with useful persistence of this data results in effective collaboration with quick response times to prepare for or prevent pending attacks.

Based on established current day environmental factors, we should have at our disposal the necessary ingredients to create a successful crowdsourcing environment for Threat Intelligence to complement the larger sharing eco-system. For instance, the enormous popularity of social media has turned strangers into virtual friends. Facebook's "like" feature virtually binds people into a community via a common interest. Furthermore, there is a growing awareness and acceptance of crowdsourcing problems. We can combine these concepts into an effective sharing ecosystem that supports crowdsourcing, and various solution implementations for Threat Intelligence that include not only the fundamental threat information but also rankings, statistics, and metrics to foster competition.

---

[9] http://ieeexplore.ieee.org/xpl/login.jsp?tp=&arnumber=6568369&url=http%3A%2F%2Fieeexplore.ieee.org%2Fxpls%2Fabs_all.jsp%3Farnumber%3D6568369

[10] Collective Intelligence Framework: http://code.google.com/p/collective-intelligence-framework/

2. Focus on the Solution:

2.1. Threat Intelligence Sharing

In order to counter sophisticated threats, and their targeted attacks, we must unite the efforts of security analysts and researchers across the industry. Together we could limit the threat's ability to operate, and predict where they would go next. Through 'crowdsourcing', our collective threat intelligence would increase over time and improve our ability to effectively cage our common cyber enemies.

2.2. What does good look like?

The MNE7 Project[11] defined high level requirements in the concept for sharing of information, including threat intelligence, to support common situational awareness. These requirements are shown in the following table and set a vision for what good could look like:

Ser	Requirement	Requirement Breakdown
1	Ability to determine dependency on cyberspace	Drives content of information monitored and shared
2	Framework/construct to enable the sharing of information	Facilitate anonymisation, work through communities of interest - Hub & Node
3	Standards for collaborative information sharing	Content, format, taxonomy, with whom
4	Confidence in shared cyber SA information	Trust, Assurance, standards, accreditation..
5	Ability to fuse/integrate information	Ability to fuse large quantities of data/information from diverse sources
6	Visualise cyber information/SA in context appropriate to level of decision maker	Ability to integrate Cyber SA with 'global' (other domain) SA. Display to show (potential) impact of cyber incident on primary business process/operation.
		Ability to determine (and display) what 'normal' activity looks like
		Ability to 'drill down' as required
		Ability to review historical activity
7	Automate processes where appropriate/possible	Ability to identify and analyse anomalies/ identify cyber threats to (components within) critical infrastructure /assets
8	Understand Legal issues around any proposed action	

[11] http://mne.oslo.mil.no:8080/Multinatio/MNE7produk/35CyberCon/file/3.5%20Concept%20of%20Employment.pdf

Table 1: The MNE7 Project[12] defined high level requirements

A potential Threat Intelligence Sharing solution could be made up of participants from one or more organizations, across a single or multiple communities, coming together to share what they know and work together to analyze their joint understanding of the threat. The solution could begin by enabling sharing across the organization, then incorporate peer to peer sharing, and finally allow threat sharing communities and intelligence feeds to be leveraged. The following sections detail the 5 major building blocks of such a potential threat intelligence sharing solution.

- Sharing Community
- Inputs
- Collaboration and Sharing
- Data Analytics
- Actionable Outputs

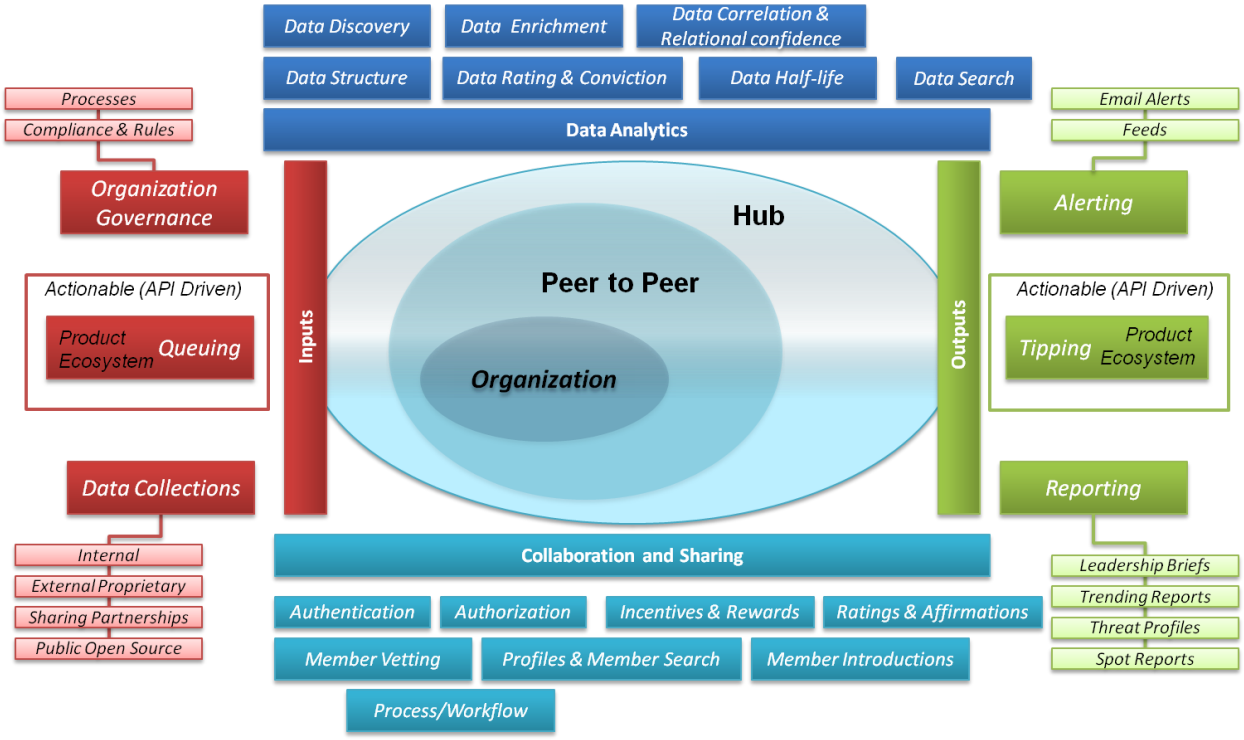

Figure 1: Threat Intelligence Sharing Components

Once we have confirmed our major building blocks, we can describe the underlying subcomponents.

[12] http://mne.oslo.mil.no:8080/Multinatio/MNE7produk/35CyberCon/file/3.5%20Concept%20of%20Employment.pdf

2.2.1. Sharing Community

Scope

- Internal – Sharing is limited to within an organization. No sharing with external entities.
- Peer-to-peer – One-on-one between organizations.
- Community Hub - Supports many-to-many sharing. The goal of the "hub" is to support rapid dissemination of information in the widest group possible given trust constraints. In other words, you can share immediately with the 5 companies whom are working together, verses having to share the same threat intelligence with all 5 companies independent of one another. This also allows the quickest community enrichment since all 5 companies are seeing changes to the data in real-time.
- Threat Intelligence Feeds – includes directed and actionable intelligence that may be focused to specific recipients or types of threats and may result in immediate actions taken to mitigate threats. Examples include fraud detection services to notify organizations about compromised accounts or credit cards, vetted intelligence provided to an intrusion prevention or security analytics system, malware detection products and service, or other host or network service to detect or prevent threats.

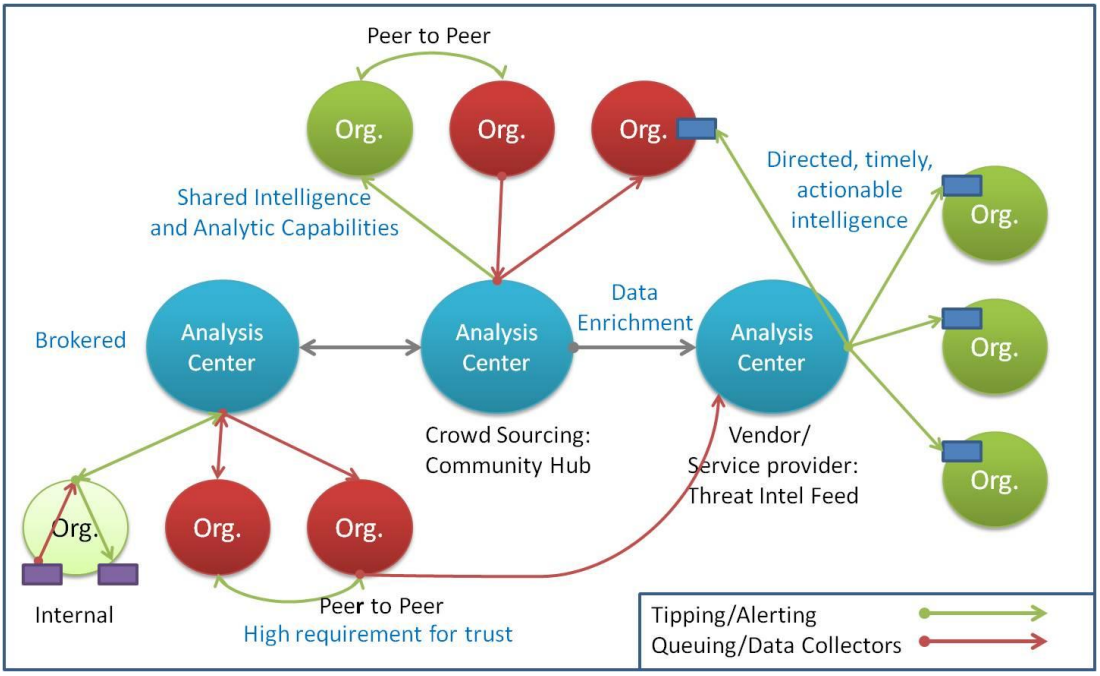

Figure 2: Information Sharing Ecosystem

The sharing communities depicted demonstrate the complexities of the larger ecosystem, where there are many types of analysis centers, service providers, and organizations that will exchange data to mitigate threats. The ecosystem must accommodate various types of inputs, outputs, analytic capabilities, and collaboration as described in the next four sub-sections to realize the full benefits of a connected ecosystem. The types of information or intelligence exchanged may vary between sharing partners, sharing communities, analysis centers, and service provider threat feeds to fully leverage the intelligence from detected threats, analytic skills, and automated capabilities of the ecosystem.

Analysis centers may include crowdsourcing community hubs with dynamic abilities to share data and utilize shared analytic resources, closed communities within an industry or region, a community of operators of a specific service, or threat intelligence service providers. Analysis centers may share information with directly connected organizations. However, to create an effective ecosystem with the ability to quickly respond to threats in an actionable way, sharing through common data formats and protocols with other analysis centers or threat feed providers is essential. Although intelligence threat feed providers may use proprietary formats and protocols to communicate with their own devices, the ability for these products to exchange data with a standards based interface provides organizations the ability to quickly respond to threats learned through other sources, eliminating the current need to copy and paste threat data within their environment.

Information sharing is a broad term encompassing many aspects of a well-connected ecosystem to attain the directed and actionable exchange of intelligence needed to combat current and future threats.

2.2.2. Inputs (Red Boxes)

There are three types of inputs into the Threat Sharing Platform. Governance includes all the human processes, management of risk, and usage of knowledge as it traverses the organization that go into information sharing. Data providers will come in many shapes and sizes and in general will be helpful to the community. An example of an intelligence provider would be Cyber Squared providing 1,000 indicators about Crimson OAK an APT Group to a telecommunications industry community. These 1,000 indicators must be injected into the community and fused into an existing understanding of risk to that industry and to those companies working together in the community.

For effective communication of the crowd sourced indicators, the inputs provided can be shared with other service providers who could then inject these indicators as inputs to products deployed by their customers, or a subset of their customers in a similar or connected sharing group. This step is essential to automate the necessary actions, enabling defenses from well-known and understood threats vetted by a vendor and disseminated broadly. (With respect to the Catalyst solution set, the inputs could range from the collective intelligence entered into Cyber Squared to perform the analysis to the next set of inputs out to vendors such as RSA's Security Analytics that can verify the inputs, perform additional data collection and analysis, then quickly operationalize the data (in the form of parsers, feeds, rules, etc.) out defenses to customers.) By leveraging automated inputs using strategic and directed methods, organizations can benefit from the analytic capabilities and verification from central locations, without having to track the details of threat actor patterns and other complex information sets exchanged in today's sharing circles. This model also reduces the number of cyber security experts needed in each organization, while enabling advanced capabilities to small organizations with few resources.

While a large number of intelligence providers are actively working toward representing their content in standardized formats such as IODEF or STIX, the reality is that most today do not leverage a standard to communicate what they know, rather they use a Comma Separated Value (CSV) file or proprietary means to share intelligence data with others. As this will likely be the case for some time, we must have a system in place by which we can import files/data of various types into the community in a useful way. Standards used within an ecosystem to direct intelligence data to maximize effectiveness are a very import building block and are described below.

Input Components

- Organization Governance – Policies (compliance and rules) and procedures (processes) that regulate sharing both internal to the originator and external to how the information can be used by the receivers. In some cases, policies can be imposed by government, internal/corporate, self-imposed amongst a community, or other sources.
    - Processes
    - Compliance & Rules
- Queuing - These are automatic feeds of inputs driven by automation. Processing of the incoming data will be completely automated, semi-automated, or completely manual.
- Data Collections – Threat information that is input for the purpose of analytics and/or sharing. This data can be input via a variety of mechanisms:
    - Internal – Examples of internal data sources includes: incident reports, malware discovered from an incident, router logs, firewall logs, intrusion detection system (IDS) logs, web server logs, anti-virus logs, host based event logs, network packet captures, email server, SIEM feeds, etc.
    - External (Service/Subscription based) – Often this is data available for purchase or subscription based. Examples of external data include: raw data feeds (malicious IP addresses, C2 domains, binaries/samples), indicator and malware sharing circles, and incident reports and threat intelligence reporting products (private reporting sources).
    - Sharing (Partnerships) – Organizations within the same industry may establish sharing partnerships. Examples of shared data include: incident reports, public or private passive DNS, formal Threat Intelligence reporting from external organizations, major organization events, and sharing indicators. Sharing may also include clients providing data to service providers to enable a broad and fast response to a new threat.
    - Public (Freely available) – Public sources of information that are often produced by analysts and private researchers that blog about threats.

2.2.3. Collaboration and Sharing (Light Blue Boxes)

We believe that sharing without collaboration will not be as strong a solution as sharing as part of a robust collaboration community. There are currently some examples of intra-sector sharing where participant members attempt to share threat information by sending structured data as attachments using traditional methods such as email. This simply does not scale as each organization and person involved has a different view of that data and email is not a medium that allows collaboration across the data, which is a must. Other models involve the use of a portal where members need to log into the portal to receive data, and then direct the data to where it is useful in their environment. This model is also flawed as it relies upon users logging in to view and share data, leaving the possibility of time gaps as to when information is received after it has been shared.

Collaboration and Sharing Components

- Authentication – The system must support authentication of individual users and machine to machine communications via the API. Authentication across a single system should support a commonly leveraged "identity" in the organization. For example, username and password may be all that is necessary or mutual certificate based authentication may be required.
- Authorization - The system should be capable of authorizing each user against roles/permissions inside the organization and across the various toolsets, peer to peer, and hub.
- Incentives & Rewards - People are more willing to share if there are incentives attached. Examples include: medals and points, unlocking features, etc.
- Ratings & Affirmations – Provides the ability to assign a value or confidence level to other users or organizations.
- Member Vetting - As trust within communities will be paramount, the provider of the system, or the owner of the community must have some way of vetting those organizations wishing to connect to it.
- Profiles & Member Search - Provides a mechanism for discoverable identity, so that other users can find those with whom they want to connect. Searching could be done across a variety of interests: geography, industry, size, same threat, same risk model, same legal, etc.
- Member Introductions - Facilitates two people from different organizations being introduced. Member introductions will require adherence to the organizations opting into specific interests. For example - I would like introductions from other financial institutions, larger than 500 people, in the North East United States.
- Process/Workflow - There is a process that is followed to keep knowledge of the threat moving. Workflow provides the ability for internal and community “action” requests and distribution of workloads across larger distributed teams.

2.2.4. Data Analytics (Dark Blue Boxes)

Given the sheer amount and complexity of data being inputted and shared, analytics is a critical capability for threat intelligence sharing. There are multiple specific requirements that need to be addressed as data is collected, shared, and leveraged. As the amount of data grows the need for automation grows as well.

Data Analytics Components

- Data Structure - Data and relationships within the ecosystem should have preference to structure for ease of sharing. This means that indicators, context data, and the relationships between data points within the system need to be structured wherever possible; however, the system should support unstructured data wherever required. Data structures may vary based on use cases and user groups to serve specific functions for that community. Differing structures may have relationships with each other to facilitate the needs of the larger ecosystem. For example, a spear phishing attack may be described with IODEF, but if the more full redacted email details were to be

exchanged, ARF could be embedded within IODEF. In addition, if the use case called for broader context around these details such as more refined and expressive indicators, structured characterization of attack behavior and resources (TTP), targeted vulnerabilities and weaknesses, attributed threat actors, related attack campaigns, or suggested courses of action, the STIX language could be utilized. Data structures will evolve over time, making it critical that they be maintained in a use case, user group driven model. The larger ecosystem may unite the various standards based data structures if appropriate for their sharing models. To continue from the example, ARF is currently maintained and used by numerous large-scale mail abuse operators ensuring as mail protocols evolve, so will their supported data format, ARF.

- Data Ratings - Are used to apply a "badness" value to an indicator, incident, adversary, or threat. Although these are done by each analyst or organization, they can additionally be averaged across communities, and within industries, so that more global trends can be determined. Ratings can be created manually or through automation.
- Data Correlation & Relational confidence: There should be automated and manual mechanisms for creating relationships between data in the system such that these relationships are retrievable and viewable within the system. The system should allow confidence in these relationships to be as much as possible assigned through automated analytics and but also allow manual user assignment.
- Data Half-Life - As the amount of data grows, the task of managing the data becomes more complicated. In addition, with automated data discovery and collection the task is simply too much for human beings to accomplish. Thus, the system needs to auto-convict indicators rating and rating confidence. Along that same line, many types of relations are not long lived and thus, you will want to depreciate the value of the relational confidence between certain types of associations over time. Each type of intelligence derived from various sources must have its own "half-life" scheme so that each type of data and relationship between data, created through various means, can depreciate over time as required.
- Data Search – A global mechanism for finding data within the system. This is especially important as the amount of data provided across many shares will be daunting to look through. Security will pose a problem in searching this data set and will require a non-traditional indexing and search capability be used.
- Data Discovery - A method of finding data external to the system that can be integrated into your existing data set. For example, you use your existing data to find relationships with external data providers and then create a relationship in the system to this external data, either by importing it or pointing to the external source.
- Data Enrichment - Ability to add context and validation to data within the system manually or through automated analytics.

2.2.5. Outputs: Actionable Intelligence (Green Boxes)

Sharing and automatic analysis accomplished at crowd-speed will give us a leg up on our perspective threats, however we will need to integrate that new found knowledge into our organizations in some way other than another lengthy whitepaper. We need

a way to consume actionable intelligence in real-time and leverage it in pursuit of defending our particular organizations.

Actionable Intelligence Components

- Reporting
    - Leadership Briefs - Formal presentations to executive leadership.
    - Trending Reports - Report on trends such as the number of incidents by a threat, key periods of time, or organizational events. They may provide historic context or metrics.
    - Threat Profiles - A comprehensive overview of activity that is being conducted against an organization by a hacker or group. Typically the threat profile details the activity along with capabilities and infrastructure, and may contain mitigation recommendations.
    - Spot Reports - Primarily focused on supporting incident response, but may contain Threat Intelligence elements such as targeting, attribution and historical activity.
    - Tipping - These are outputs driven by automation that merge tactical indicators with signatures, data feeds, data parsers or rules (product specific or standards based) which are tasked into specific products internal to the organization. Examples may include: intrusion detection, firewalls, email security, network monitoring, SIEM, Data Loss Protection (DLP).

  Note: Reports in this context can be offered in the form of a Threat Intelligence Dashboard with role-specific views for threat experts and leadership to visualize key threat indicators and attack profiles regardless of the data source.

- Alerting
    - Feeds - Automated outputs that provide tactical indicators based off analysis, and allow security personnel to identify current activity targeting the organization, or alert on activity that will target the organization in the future.
    - Email Alerts - Automated issuance of an email containing a tipper or feed.

In the example provided above where Cyber Squared shares out 1000 indicators to the telecom community, the intelligence is made actionable by providing the sets of indicators out to other relevant vendors in addition to the telecom sector participants. Other vendors can assist with the effective dissemination of the data, into the network or system management tools, where the intelligence can be applied effectively.

While the current model of sharing in industry groups has had a large impact to participants, the data is received in formats that are difficult to ingest and then deploy appropriate controls. This somewhat manual process has to be repeated by analysts within each participating organization. Organizations that are not part of sharing circles in their industries or regions may be left vulnerable to attacks that could be more effectively prevented. Automated threat feeds from intelligence providers as well as tools that operationalize this intelligence to

provide immediate defenses to organizations need to be considered as part of the larger sharing eco-system for effective use of shared intelligence.

2.3. How can standards help?

Standard data formats and protocols enable the interoperable exchange of information between different implementations of open source or vendor products, resulting in increased automation. For information sharing, we care about the ability to send or receive and process data with common interfaces to enable interoperability. Once the data is successfully exchanged, analytic capabilities and use of the shared intelligence may vary greatly between organizations, analysis centers, and intelligence providers. Hence standards are often focused primarily on exchange; however, holistically expressive representations such as STIX are seeing significant inroads in use not only for exchange but also within solutions and centers. Large analysis centers may track patterns or techniques of threat actors, analyze complex malware, correlate network or specific protocol activity across a diverse set of organizations, or other areas of research where each requires different types of data that may or may not be complementary. As such, meeting the use case and sharing model needs are key requirements when selecting a data format. Most data formats are flexible and extensible by design to enable to use of extensions as needed.

Understanding your use case and sharing model requirements are far more important than what standard or specification to format your data when embarking on a new intelligence sharing initiative. The sharing model must consider what data is useful to share with whom to create an effective eco-system. This may be accomplished through the use of comma separated value (CSV) files. Once the sharing model is at a mature state, evaluating data formats and protocol standards can be done to automate the flow of information to meet your requirements. Putting too much a focus on standards can detract from the real goals of eliminating threats, while leveraging the very small number of skilled resources that exist. Vendors may play key roles in effective sharing models to assist small and medium sized organizations that may never have enough resources to process and utilize threat intelligence. The APWG example provided below demonstrates an eco-system that considers these constraints, helping to provide protections to organizations and individuals alike.

Actionable intelligence is typically derived from the more complex analysis capabilities to be provided either directly from the analysis centers or through intelligence providers. Updates through intelligence provider feeds typically result in direct updates to the threat protections of an organization that can be broad, rapid, and impactful to the subscribers of the service. Standard data formats are needed to facilitate multiple types of information to be exchanged to support the range of activities of analysis centers, intelligence providers, and organizations. Analysis centers may require support for diverse data sets, whereas most organizations are more interested in directed and actionable intelligence with little or no interaction by their staff. In other words, ‘one size does not necessarily fit all’ for information sharing standards!

2.3.1. Early Leaders/Adopters

As described earlier, there are a few successful examples of communities using open and internationally accepted standards to facilitate their communications such as the M3AAWG operator community to mitigate mail abuse and the APWG mitigating anti-phishing and

eCrime behavior. Ren-ISAC through the open source Collective Intelligence Framework (CIF) performs higher level analytics and shares actionable intelligence directly with participants.

The APWG provides an excellent example of use case driven sharing with a highly efficient and effective sharing model. The APWG began its work focused on exchanges to thwart phishing attacks, leveraging very few skilled resources to have a global impact. Since then, they have added a second use case to combat eCrime. Members, including the financial sector and vendors participate in directed and actionable exchanges using the IODEF data format with explicit extensions (standard, such as RFC5901 for anti-phishing and private for eCrime) designed to meet their requirements. The APWG accepts IODEF formatted anti-phishing reports and collects them into a Cybercrime information warehouse. The APWG provides relevant feeds to various members who perform directed actions, having a global impact while leveraging a very small number of skilled resources.

This eco-system focused sharing model is highly efficient as it leverages an interconnected set of analysis centers that can have a broad impact using directed exchanges. The model eliminates the need for analysts at every organization to review threat intelligence, determine how to deploy controls, interact with law enforcement, and maintain updated block lists. This model benefits not only large and small organizations, but also individuals.

More recently, the FS-ISAC (inspiring emulation by numerous other ISACs) operationally shares threat intelligence today among its 4200+ financial institution members utilizing its Cyber Intelligence Repository (CIR) with capability for not only interactive query but also full machine-to-machine automation of exchange. The US Department of Homeland Security currently shares actionable operational threat intelligence with critical infrastructure and other industry players.

These efforts have had a tremendous effect assisting those who directly participate, leveraging the skilled analysts of members that include vendors to make use of any shared intelligence and distribute automated controls or mitigations more broadly. Until vendors can produce and consume this directed actionable intelligence into products in additional use case scenarios, information sharing is limited to those with analyst resources. The use of standard data formats and protocols are needed at product end points that interface with information sharing communities as well as to the service providers of those products deployed within the organization to implement recommended mitigation measures.

The ecosystem and architecture described in this paper and demonstrated in the TM Forum catalyst project provide a method to achieve the desired sharing of directed and actionable intelligence. How do we get there?

We need automated capabilities through standards in place to exchange data between products within an organization for the organizations that do have analyst resources so they don't have to consume their time with manual tasks. We also need connections in the ecosystem to provide automated actionable intelligence to both those who have analyst resources and can benefit from vetted intelligences as well as those who rely on vendor products with automated threat feeds to those products. The automation activities must be driven by use cases and what data is necessary to share with whom in order to have a broad impact. Examples of use case driven approaches are available in a paper published by RSA.[13]

[13] http://www.emc.com/collateral/emc-perspective/h12175-transf-expect-for-threat-intell-sharing.pdf

Standards Used by Catalyst

The catalyst demonstration primarily makes use of open and internationally accepted standards for data formats and protocols from two sources (the IETF Managed Incident Lightweight Exchange (MILE)[14] working group, and the DHS-sponsored/MITRE-led structured threat information standardization communities) to demonstrate actionable and directed sharing of information for the inputs and outputs of the architecture.

The IETF MILE working group standards[15] include the IODEF[16], in its current state of revision for the data format and Real-time Internetwork Defense (RID)[17]. IODEF provides a base data format to exchange information about incidents as well as commonly exchanged threat data with the goal of providing actionable intelligence can be context rich. In cases where it is desirable to exchange data that is not as commonly shared or is specific to a use case, the flexibility and extensibility common to IETF protocols are leveraged, in this case, extensions to IODEF are used. An example includes the use of the previously mentioned ARF format to exchange email information may be used when it is desirable to exchange the full email information (or redacted version of the full email) within the context of the IODEF message. This enables the ability to forward on the ARF portion of the message to mail operators in a format they expect. Leveraging existing community formats for specific extensions provides a long term support model. The capability to include ARF with IODEF is being discussed for support in the open source tool[18] used in the TM Forum catalyst demonstration. Another example of extending IODEF for a specific use case is the Jabber or XMPP operator community who currently use IODEF with an extension to support their abuse and incident related exchanges for instant messaging[19].

IODEF also supports a structured extension mechanism through the working group draft, Structured Cyber Security Information (SCI)[20] extensions to IODEF. This standard establishes expected locations to embed other complementary schemas over time through the use of an IANA registry. This approach assists developers as they may only choose to support a portion of the possible schema extensions based on the relevant use cases for their implementations or products and the registry can evolve over time to further structure the extension points of IODEF. IODEF's extensibility and flexibility helps to future proof the format with the use of a registry to maintain the current list of supported extensions to help developers narrow the scope of work for support.

The DHS-sponsored/MITRE-led structured threat information standardization community representations include STIX and some of its purpose-focused constituent representations including the Cyber Observable eXpression (CybOX) and the Malware Attribute Enumeration and Characterization (MAEC). CybOX provides an extremely expressive, flexible and extensible representation format for cyber observables whether observed instances or potentially observed patterns. It can express both dynamic events/actions and static object properties (with current structured representation for more than 80 different types of objects) as well as infinitely diverse and complex relational or logical compositions. MAEC provides an expressive, flexible and extensible representation format for characterizing malware structure

[14] IETF Managed Incident Lightweight Exchange (MILE) Charter: http://datatracker.ietf.org/wg/mile/charter/
[15] IETF MILE Working Group http://trac.tools.ietf.org/wg/mile/trac/wiki/WikiStart
[16] RFC5070-bis https://datatracker.ietf.org/doc/draft-ietf-mile-rfc5070-bis/
[17] RFC6545 Real-time Inter-network Defense (RID). K. Moriarty. 2012. and RFC6546 Transport of RID over HTTP/TLS. B. Trammell. 2012.
[18] https://github.com/RSAIntelShare/RID-Server/wiki/RSA-EMC-RID-Server
[19] IODEF in XMPP for incident reporting among XMPP operators: http://xmpp.org/extensions/xep-0268.html
[20] Draft-ietf-mile-sci-06.txt. T. Takahashi, et. al. Feb 2013.

and behavior. It leverages CybOX for expressing the actions and objects related to the malware and adds additional layers of abstraction for characterizing the malware-specific behaviors and mechanisms providing contextual meaning to those lower level actions. STIX provides an extremely expressive, flexible and extensible representation format for holistic threat intelligence. It provides expressivity for cyber observables, indicators, incidents, Tactics, Techniques and Procedures (TTPs), vulnerabilities or weaknesses targeted for exploit, threat actors, campaigns and courses of action. It supports expressing everything from the tiniest single piece of information to enormous bodies of interrelated intelligence.

The TM Forum catalyst demonstration and architecture continues to evolve and incorporate standards and best practices developed internally as well as those developed by industry partners and other SDOs. One such framework is the Collaborative Cyber Situational Awareness (CCSA) framework[21] which was applied to Phase 1 of the Catalyst demonstration, and STIX which was used during Phase 2 of the Catalyst.

2.3.2. Emerging Specifications

The authors of this paper would like to recognize other standards that are emerging in this space. While these standards were not leveraged by the Catalyst project or the participating vendors, it's important to note these for the sake of completion:

- Trusted Automated Exchange of Indicator Information (TAXii)[22]
- Traffic Light Protocol[23]
- Common Alerting Protocol OASIS[24]

2.3.3. TM Forum Frameworx and Threat Intelligence Sharing

Since TM Forum Frameworx have evolved over the past few years to address Security Management concerns, such as the incorporation of NIST's SCAP into the Information Frameworx[25], it is reasonable to expect that Frameworx will be extended to address best practice TI sharing standards in future releases.

---

[21] CCSA – http://mne.oslo.mil.no:8080/Multinatio/MNE7produk/35CyberFra/file/3.5%20Framework%20of%20Processes.pdf

[22] http://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=1&cad=rja&ved=0CC4QFjAA&url=http%3A%2F%2Ftaxii.mitre.org%2F&ei=yZyBUduBFMfV0gG61YHwAQ&usg=AFQjCNHdlpdKz8sKiEWjofJccb4rWvtX7A&bvm=bv.45921128,d.dmQ

[23] https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=1&cad=rja&ved=0CDAQFjAA&url=https%3A%2F%2Fwww.us-cert.gov%2Ftlp&ei=LJ2BUe3BC8Sx0AHSqoHwDw&usg=AFQjCNGtSqofXxdj9eSWEXpNRpqfKPqZXQ&bvm=bv.45921128,d.dmQ

[24] https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=1&cad=rja&ved=0CDEQFjAA&url=https%3A%2F%2Fwww.oasis-open.org%2Fcommittees%2Femergency%2F&ei=Vp2BUbOyBoe-0QH12IEw&usg=AFQjCNFEFY0jPHI9rF3l4lzEQKbTXWKRGA&bvm=bv.45921128,d.dmQ

[25] http://www.tmforum.org/Information/13654/home.html

2.4. Why the TM Forum?

The TM Forum is a sensible place to address the topic of Threat Intelligence sharing, for several reasons:

- Key TM Forum members are Communication, Internet, and Mobile Service Providers who are operate critical infrastructure
- The TM Forum offers an active collaboration program that brings together subject matter experts from across the globe to develop industry shaping best practices for Security Management issues
- The TM Forum has a liaison program which enables the exchange of information and collaboration around industry standards and best practices begin developed by other service delivery organizations
- The TM Forum runs a Catalyst program[26] which allows The short duration of these projects enables a winning formula for all involved:
    - Service providers accelerate trials and deployment of solutions
    - Suppliers rapidly progress their ability to easily integrate with partners
    - TM Forum best practices and standards gain enhancements in the newest and most pressing areas

While it is unlikely that the TM Forum will ever be the database of record for Threat Intelligence, it's an excellent forum for bringing together suppliers and consumers to

2.5. Who are the winners?

Applying a crowdsourcing approach to Threat Intelligence would allow assembly of an impromptu, virtual army of trusted cyber defenders. This would enable the community to more quickly and comprehensively understand the threat and predict where they will go next. It would also allow the ability to create dynamic relationships with trusted sharing partners who have common threat interests, and enable participants to register and receive notifications when threats change. By transitioning today's more static "sharing" model to be a more dynamic "crowdsourcing" approach for Threat Intelligence, we would actually improve response times and predict attacks. Furthermore, it is possible that through a successful crowdsourcing Threat Intelligence solution it could serve as a deterrent for cyber adversaries.

The larger eco-system benefits from the crowd-sourcing model in the ability to share intelligence information using methods that are immediately actionable. The targeted dissemination of threat intelligence for monitoring and defense capabilities will help us overcome the tremendous lack of skilled personnel with cyber security expertise.

[26] http://www.tmforum.org/CatalystProgram/786/home.html

2.6. What are the risks?

The Threat Intelligence Sharing architecture takes into account concerns of confidentiality and operational security to encourage sharing of relevant intelligence data. We also look at context and validation as key elements and recognize that we need to protect from the adversary using the solution against its members. Our team recognizes these issues, and has worked to mitigate, wherever possible, those barriers to the solution.

2.6.1. Litigation

Fear of embarrassment, litigation, and loss of confidence due to admission of compromise are some of the impediments to sharing between or amongst groups. Litigation is today a reason people 'say' they do not share. On one hand people are worried that they will be sued if they share externally, and on the other hand, there are beliefs that not sharing what you know, and that leading to a compromise of another network, could additionally lead to litigation. Fortunately, global efforts are underway to reduce the risk to the commercial companies in sharing incident information with the government or with other companies.

2.6.2. Accountability

Accountability concerns are a well-known constraint to threat sharing. A crowdsourcing solution must introduce a trust model where users have the ability to share anonymously, yet with data integrity intact and trust maintained. The Threat Intelligence Sharing architecture provides facilities to act as an "arbiter" of trust. This allows aliases and other means of remaining private when anonymity is desired, while at the same time keeping accountability intact. This capability may be critical to success of the solution, and is being treated as such.

2.6.3. Context

Also, existing solutions do little to nothing to address the issue of necessary "context" of threat information. Threat Intelligence is not actionable if its context has been stripped away or was never present in the first place. This creates a challenge as it is very difficult to leverage threat intel unless you know why it is bad, how it was derived, and what actions are recommended by those that have created it. Our solution takes context into account, and provides facilities to share context along with specific threat information.

2.6.4. Confidence

Additionally, if there hasn't been a vetting process and confidence associated with the intelligence, there will be no ability to act on it. There no such thing is a long list of addresses that has confidence blindly applied to the entire list. Confidence needs to be applied to each indicator and to each relationship across the dataset. Confidence also should be applied to those organizations creating the intelligence. That way, you can de-conflict multiple organizations and analysts confidence from one another. As the number of trusted analytic opinions grows, so does the ability to take action quickly on the data being crowdsourced.

2.6.5. Sneaky Adversaries

We also must admit that our solution could be used by an adversary as a counterintelligence tool. If an adversary was able to establish an account as a user they could attempt to know who and what organizations know about them, their infrastructure and techniques, and perhaps use the solution as a platform to conduct social engineering of other groups to share false flag information. The solution again as an arbiter must overcome this risk in vetted members, accountability, and crowd defined trust groups. Only with accountability for each member of the group will the trust of the crowd be kept intact.

3. Conclusion:

3.1. What is possible? What happens if we aren't successful?

More and more we see the power of the crowd, working together, combining strengths to overcome challenges. For too long, the cyber security community has tried to address security threats in point solutions. This approach is different, because we are coming together, sharing what we know, and jointly developing our responses.

Without a shared understanding of the threat, we will not be fully prepared for its advances. If we come together as a community and combine our strengths, we can collectively stop sophisticated cyber threats.

If we succeed we will make a significant contribution to securing the Global Commons in cyberspace for all businesses; the more that participate as a crowd as illustrated by this project the more effective the cyber security outcome will be.

4. Administrative Appendix

4.1. About this document

This is a TM Forum Guidebook. The guidebook format is used when:

- The document lays out a 'core' part of TM Forum's approach to automating business processes. Such guidebooks would include the Telecom Operations Map and the Technology Integration Map, but not the detailed specifications that are developed in support of the approach.
- Information about TM Forum policy, or goals or programs is provided, such as the Strategic Plan or Operating Plan.
- Information about the marketplace is provided, as in the report on the size of the OSS market.

4.2. Document History

4.2.1. Version History

Version Number	Date Modified	Modified by:	Description of changes
Version 0.1	26/Feb/2013	Adam Vincent	first draft of the document
Version 0.2	6/March/2013	Kathleen Moriarty	Edits, comments, and additional text
Version 0.3	21/March/2013	Adam Vincent, Christy Coffey	Updated Architecture drawing, added architecture definitions and risk section
Version 0.4	15/April/2013	Seth Geftic (RSA/EMC), Martin Huddleston (DSTL), Alex Hamerstone (TOA Tech)	Review and edits.
Version 0.5	25/April/2013	Christy Coffey (TM Forum)	Accepted edits, general cleanup, updated Why TM Forum & standards section.
Version 0.6	8/Aug/2013	Alicja Kawecki	Rebranding, minor cosmetic edits prior to web posting
Version 0.7	10/Sep/2013	Alicja Kawecki	Added IPR Mode to cover page
Version 0.8	24/October/2013	Christy Coffey (TM Forum), Sean Barnum	Updates to bring the paper current with

		(MITRE), & Kathleen Moriarty (EMC), Ward Cobleigh (Edge Technologies)	Catalyst Phase 2 work & the advancement of standards in the intel sharing space.

4.2.2. Release History

Release Number	Date Modified	Modified by:	Description of changes
<<Release Number >>	DD/MMM/YY	<<name>>	Description e.g. first issue of document

4.3. Company Contact Details

Company	Team Member Representative
AT&T	*Brian Rexroad*
Defence Science and Technology Laboratory	*Martin Huddleston*
Telstra Corporation	*Clive Reeves*
Bell Canada	*Blake Lindsay*
cVidya Networks Ltd.	*Gadi Solotorevsky* *Tal Eisner* *Amir Gefen*
EMC	*Kathleen Moriarty*
TOA Technologies, Inc	*Carpenter Mike* *Alex Hamerstone*
Birmingham City University	*Abdallah Ali*
MITRE	*Sean Barnum*
TM Forum	*Christy Coffey* *Jenny Rottinger*
Edge Technologies	*Ward Cobleigh*

4.4. Acknowledgments

This document was prepared by the members of the TM Forum Threat Intelligence Sharing Catalyst Phase 1 team, and updated by the Catalyst Phase 2 team to be included in the Frameworx 13.5 release.

