Cisco's Media Sharing Protection System White Paper

Introduction

In the business of online subscription-based media services such as NFL Game Pass, MLB.Com, Hulu, Netflix, or Amazon Prime Instant Video, one of the biggest threats is the sharing of credentials by multiple non-paying users (sharers), so they can all view indefinitely using one subscription account.

The following information describes the amounts of Media Sharing that is occurring in on line video subscription services as published by Parks Associates (PA) and The Diffusion Group (TDG):

- 500 million lost revenue because of credential sharing (Parks Associates)
- Lost revenue will grow to 550 million within the next 5 years (Parks Associates)
- 20% of adult broadband users access services using someone else's credentials (The Diffusion Group)
- More than 25% of 18-22 year old broadband users not living at home access services using someone else's credentials (Parks Associates)
- Incidence of illegal sharing of Netflix streams is 19.9% and HBO Go is 18% (The Diffusion Group)
- Most Service Providers have a limited grasp of illegal sharing based on credential sharing happening in their system (response from service providers we interviewed at CES)

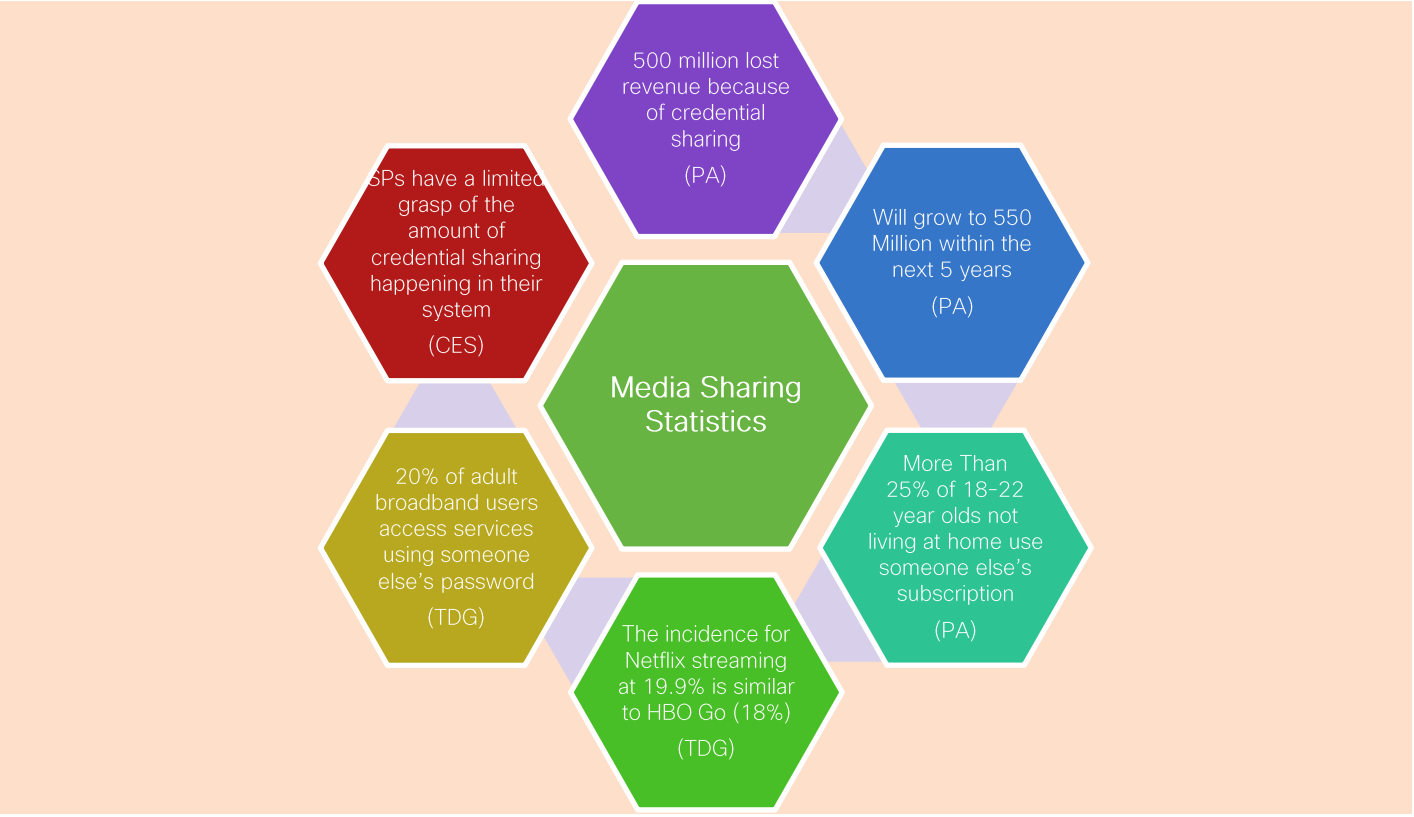

"Credentials" are the means by which clients prove their identities to the content provider or service provider system during a login or content request, such as username/password, secure cookies, OAUTH Tokens, or SAML assertions.

Credential sharing causes significant losses to service provider due to:

- Lost revenues from non-paying subscribers
- Added cost in providing wasted resources (CDN Bandwidth etc.) to non-paying sharers
- Poisoning of their brand based on proclamations of stolen account credentials from their system

However, sharing is much more complicated with subscription-based media than it is for standard banking or credit card sites, which are only worried about rogue activity based on stolen credentials. With subscription-based media, sharing may be performed by a multiplicity of users with a variety of motivations: some based on stolen accounts, some based on a form of personal or business collusion.

The primary users, sharing techniques, and motivations can be summarized in the following table:

Sharing Type	Users	Motivation	Knowledge of Sharers to Account Owner	Sharing Relationship
Casual Sharing (personal collusion / Known Sharing)				
Sharing with friends	account owner, sharer (friend)	benevolence	much knowledge	collusion
Business Sharing (business collusion / Known Sharing)				
Pooling Accounts	Many Poolers (one registers account)	Very cheap service	minimal	collusion
Swapping of accounts	account owner, sharer (business partner)	mutually beneficial	minimal	collusion
Selling valid or fraudulent accounts to many buyers	account owner, many sharers (buyers)	profit	minimal	collusion
Stolen Accounts (credentials purchased or extracted illegally / Unknown Sharing)				
Thief uses stolen credentials for his own viewing	account owner (innocent), sharer (thief)	free service	none	parasitical
Thief sells stolen credentials to many buyers	account owner (innocent), middleman (thief), sharers (buyers)	profit	none	parasitical

To combat this threat and react to each individual sharing type in the manner that best suits the business, the Service Provider needs to be confident in reporting of both the fact that sharing is taking place and the specific sharing type. In

addition, the Service Provider needs recommended proactive policies based on significant domain knowledge of each sharing type.

This white paper presents an innovative new product called Media Sharing Protection, which identifies both the sharing and sharing type with high confidence and also recommends proactive events to the Service Provider based on Cisco's strong domain knowledge of the media industry.

Detailed Media Sharing Protection Solution

Multiple Identities: Prevention vs Detection

As shown in the diagram below, the core tenet of secure communication is authentication that is based entirely on the maintenance and preservation of a secure identity. In a secure managed device one's identity is typically depicted by a private key/certificate or ID/symmetric key. These are usually maintained in either tamper-proof hardware, such as a SIM card, smart card or serialized chip, or tamper-proof software as typified by Digital Rights Management systems. However, in a standard Web environment, one's identity is not based on 'what one has', but 'what one knows'—namely a username and password or some activation code. As opposed to the aforementioned software initial identity, which is protected by secure storage,, in a web environment the credentials are more easily cloned because they may be easily stolen from non-encrypted personal files or non-encrypted databases on some Web server. Furthermore, as mentioned above, in certain environments such as online video subscriptions, the owner of the identity may be willing to share his identity with others for personal or business reasons. Hence, the key element of MSP is to rely on a third-type of identity–'how one behaves'–to detect whether the web identity/credentials, were shared or stolen. The basic idea here is that people are creatures of habit and typically act in consistent ways; hence, if multiple separate behaviors can be gleaned from account data, this usually signals multiple users sharing a common identity.

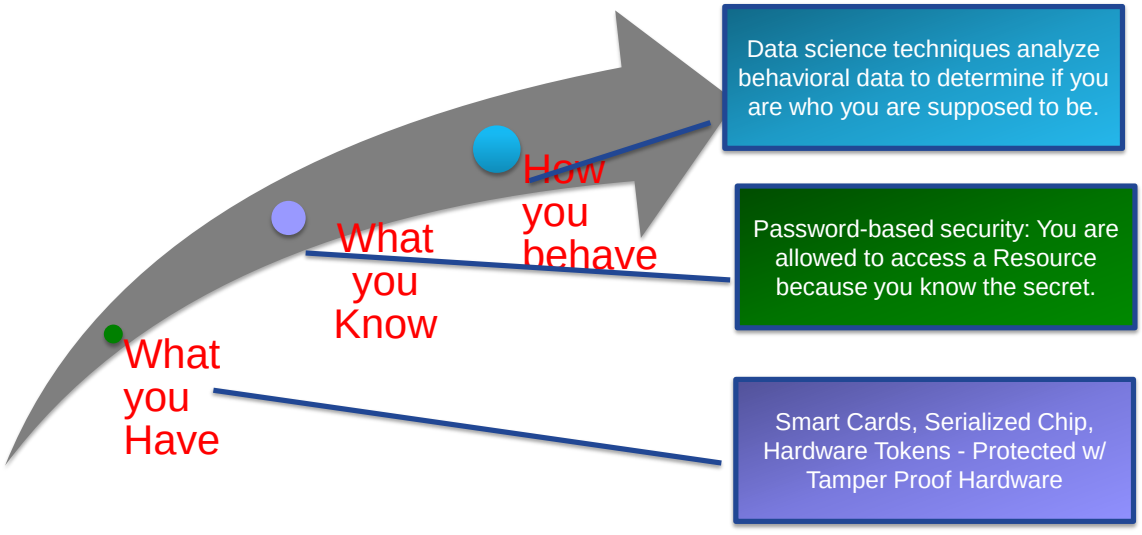

Evolution of the Secure Identity

Basic Mechanism

The basic process and mechanism and logical flow that is implemented in the MSP system to detect and respond to media credential sharing is defined by the activity/state diagram below:

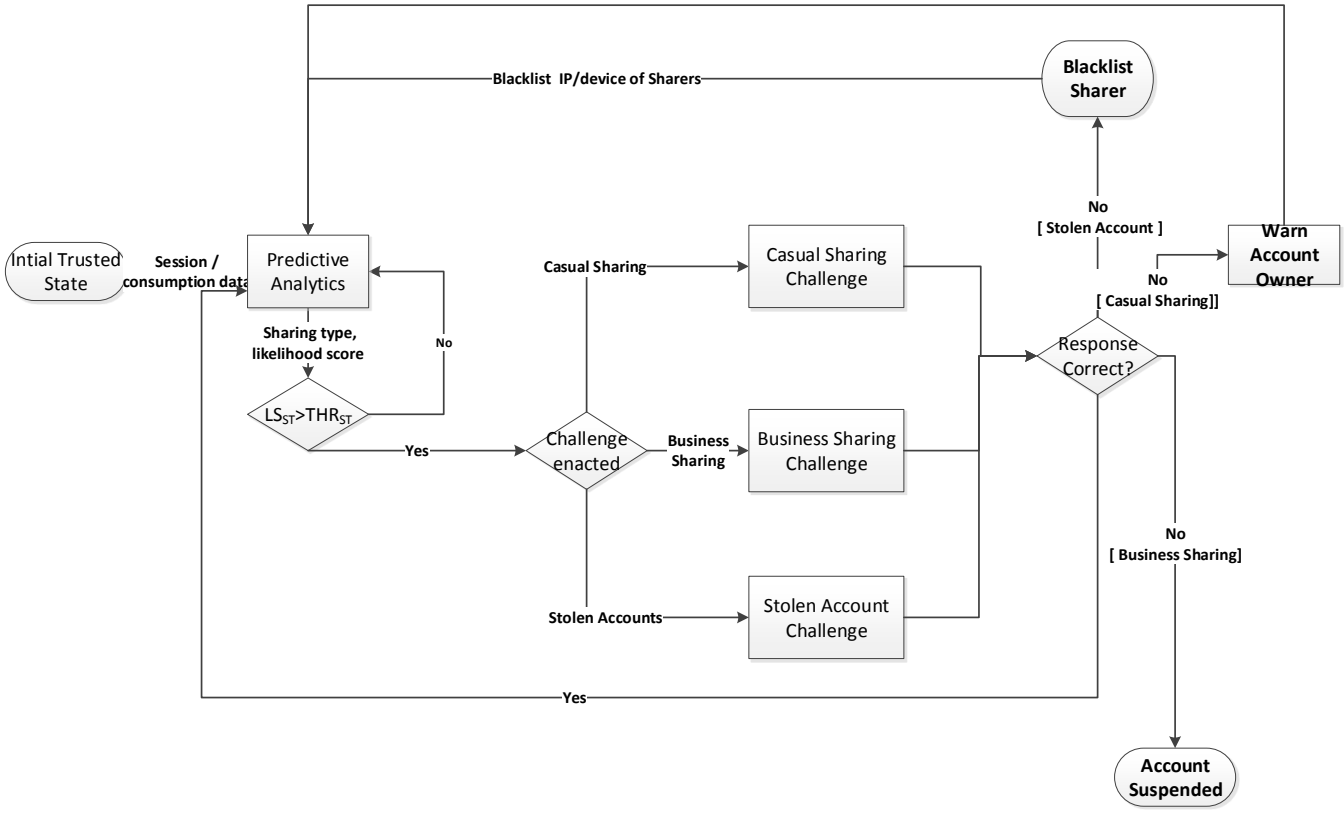

The core activity is the Predictive Analytics block where sharing is detected. Once sharing and the sharing type is detected using predictive analytics, we verify our conjecture with challenges and if the response by the user is incorrect, enforce a service provider anti-sharing policy per sharing type.

Cisco's solution uses data science routines to predict when sharing occurs. These routines not only predict media sharing, but also provide a "likelihood score" for each type of sharing described in the table above. The goal of this solution is to perform all of the following steps:

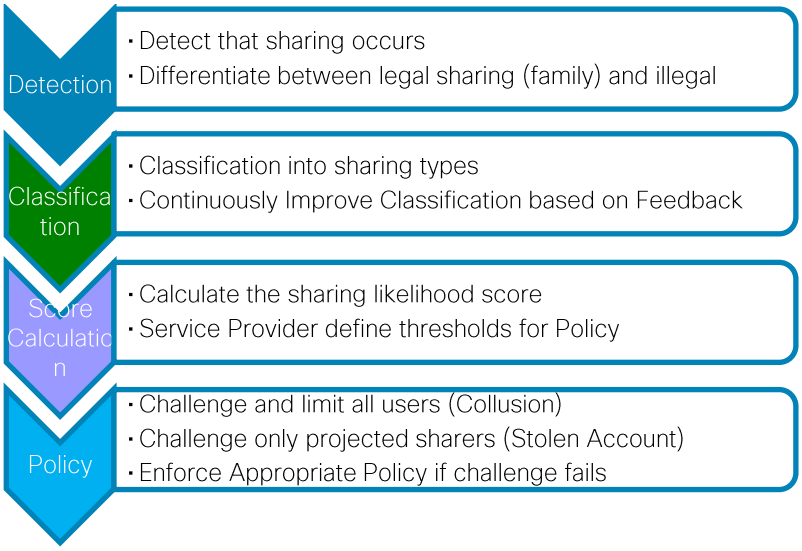

The data science routines obtain their input from tracking of the various behavioral profiles per account. Separate users within an account are identified and logged, where a given user is classified by some combination of behavioral features. Data science routines are also used to distinguish between the account owner, who is defined as the person who registered for the account, and all the other projected sharers. Examples of behavioral features that identify a unique user within each subscription account include:

Account Information
• Account ID • Credentials

User Behavioural Fingerprint

Who	When	Where	What	How
• IP Address • Device ID • Device Type	• Session Time • Program Time	• Session Location	• Programs • Genres	• Trick Modes • Duration • Subtitles

Account Events

- Concurrency Failures
- Recommendation Integrity
- Account Status Change
- Credentials Changes

The weighting of the various features (assigning coefficients) are determined by standard data science routines, such as gradient descent or back propagation.

Data science algorithms, such as multi-class classification or decomposition algorithms (discussed further below), are used to distinguish different users based on these weighted, aforementioned features within an account.

The basic Data Science Algorithms that we use are depicted in the following Activity Diagram:

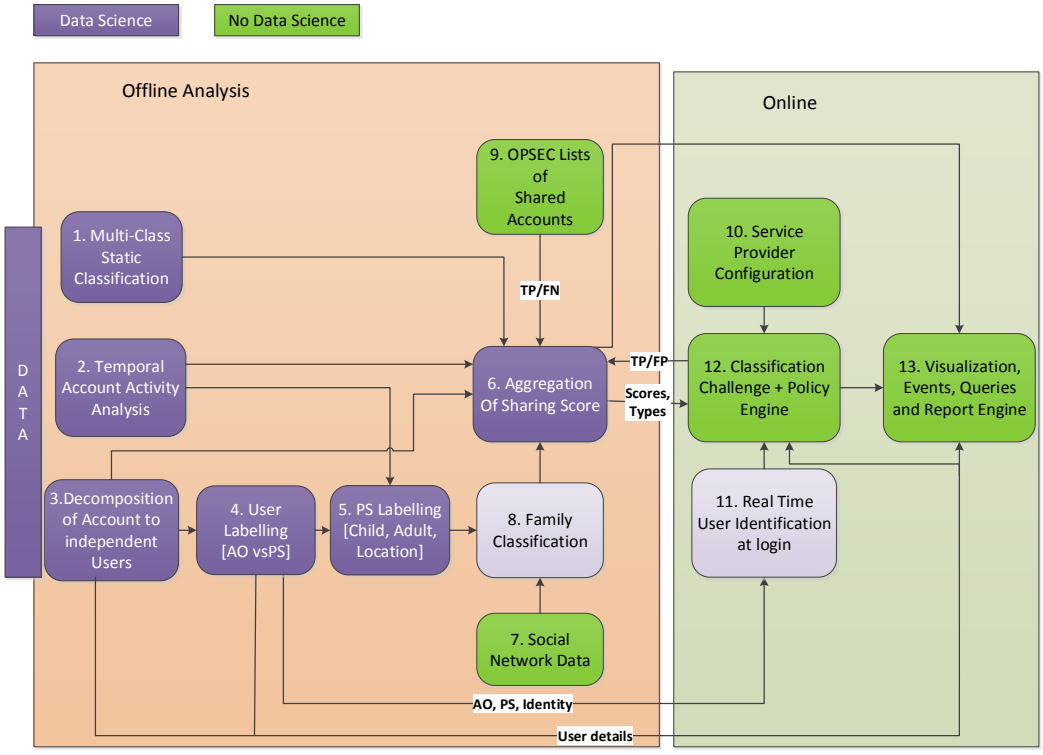

Three different algorithms are used to identify anomalous activity:

1. **Multi-Class Static Classification** is used to identify accounts with anomalous levels of activity, and classify those accounts as either casual sharing accounts, business sharing account or stolen accounts, based various collected features that differentiate the various types. This done by assigning likelihood scores to each sharing type. The likelihood scores are determined based on prediction algorithms using tools such as support vector machines, Bayes networks or decision trees, taking into account the extracted features Such features include, for example:

   - Number of sharers (many for business sharing)
   - Number of concurrency violations (many for business sharing)
   - Number of times same user profile appears across several accounts (typical of stolen accounts)
   - Fixed vs variable number of users (fixed for casual sharing, variable for business sharing or stolen accounts)
   - Constant growth of users (business sharing)
   - Fixed vs variable viewing patterns (fixed for casual and business sharing, variable for stolen accounts)

   The graph below describes using the feature of growth in number of users over time per account to identify typical variations per sharing type.

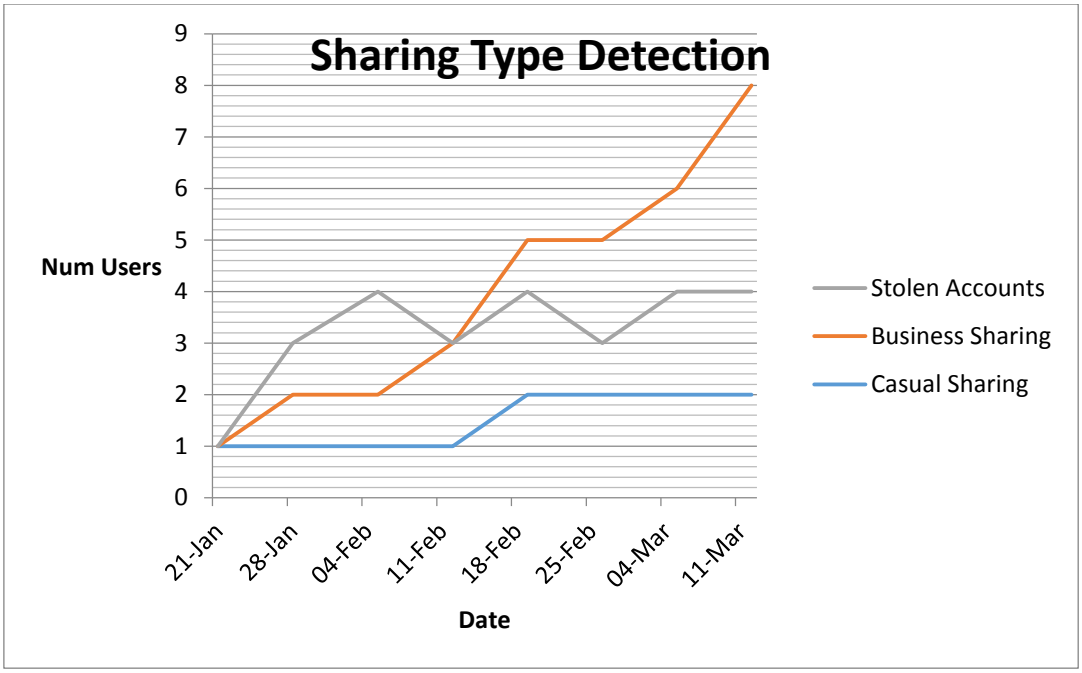

2. **Temporal Account Activity** uses Time Series Analysis to detect permanent rises in activity per account over time which can be used as an indication of increased sharing

3. **Decomposition of an Account into Independent Users** uses advanced machine learning algorithms to actually identify the distinct users in each account and classify them as either the account owner or a projected sharer. An account owner, for instance, is recognized by various means, such as comparing the device and location of viewing

with the account registration IP address, device type, location and/or currency. An example of decomposition is shown below where the users both differentiated and labelled.

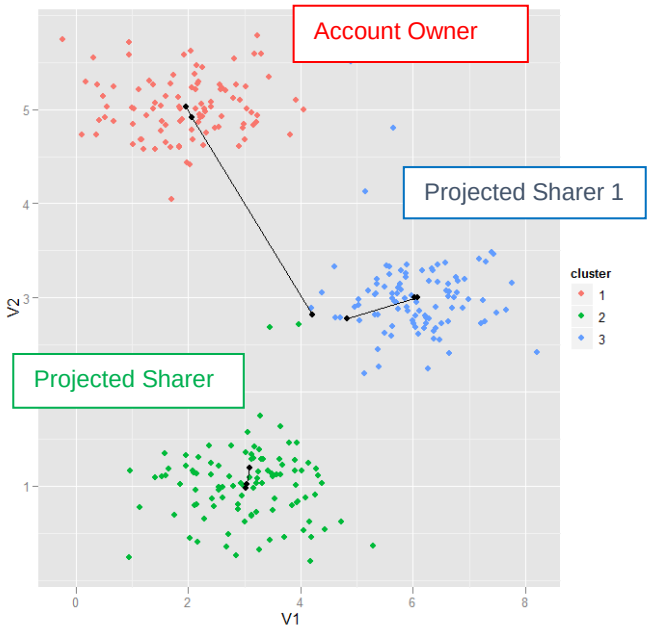

In addition, once each projected user is mapped into a user type such as child or adult using the decomposition technique, and based on the composition of users and their location, a data science algorithm can compute the probability that all users within an account comprise a single family or not in order to differentiate between what may be legal sharing (within a family) or illegal sharing (members external to the family). An example of this form of detection is depicted in the table below, where the account below, projected to comprise two adults and three children, is likely to be legal sharing.

IP address	Device Type	Main Viewing Times	Preferred Genre	Projected User Type
101.234.112.001	iOS Smart Phone	Late Evening	News	Adult
101.234.112.002	Android Smart Phone	Afternoons	Disney	Child
101.234.112.008	Android Tablet	Afternoons	Kids Shows	Child
151.104.042.038	MAC computer	Afternoon	Kids Shows	Child
091.043.242.158	Connected TV	Late Evening	Adult Movies	Adult

In this system, we create an ensemble of the aforementioned data science algorithms in order to compute an aggregated sharing likelihood score that quantifies both the amount of sharing and projected sharing type per

account. This score also takes into account possible feedback from operational security that uses the dark net and other sites to actually inspect and determine sharing transaction taking place in the general Internet

In summary, the sharing classification is based on the large quantity of data that is collected and the static and dynamic activity of the suspected account is compared against other accounts to detect anomalous activity. In addition, the MSP system maps the video sessions and consumption to actual projected users in order to both prove and interpret our sharing hypothesis, surmise that the sharing is illegal, and identify the account owner. The sharing predictive analytics can be summarized by the following class diagram:

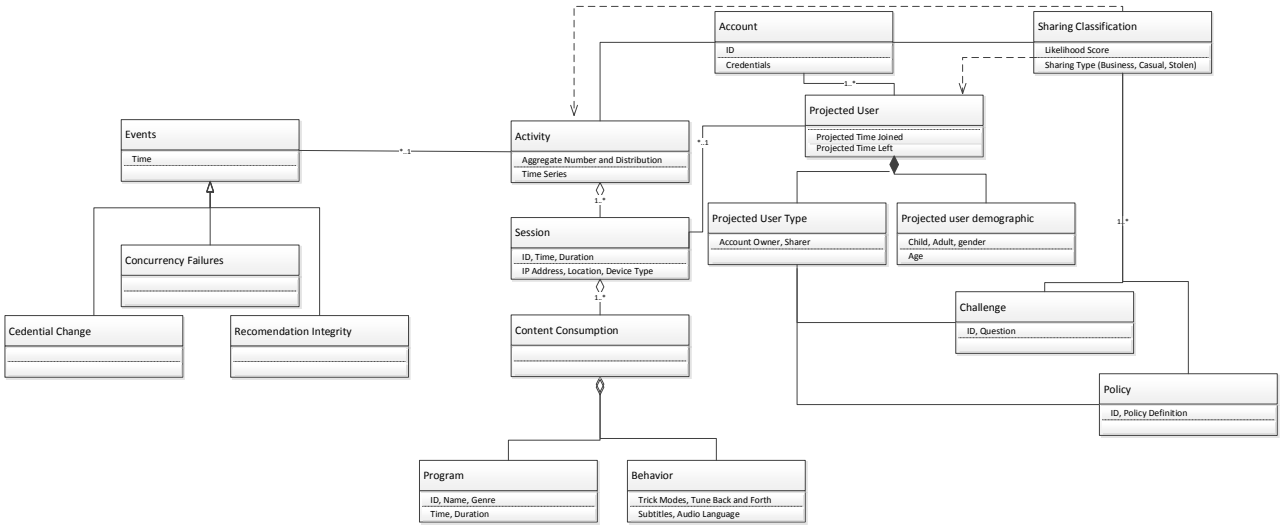

Challenges and Policies

Based on the aforementioned data science analysis, likelihood scores are assigned for each of the sharing types and compared against a service provider-configurable threshold per sharing type. The picture below describes results of the analysis that the system performs, where the red lines are the different thresholds configured by the service provider for each sharing type. Once the threshold is passed, a challenge relevant to that sharing type is enacted.

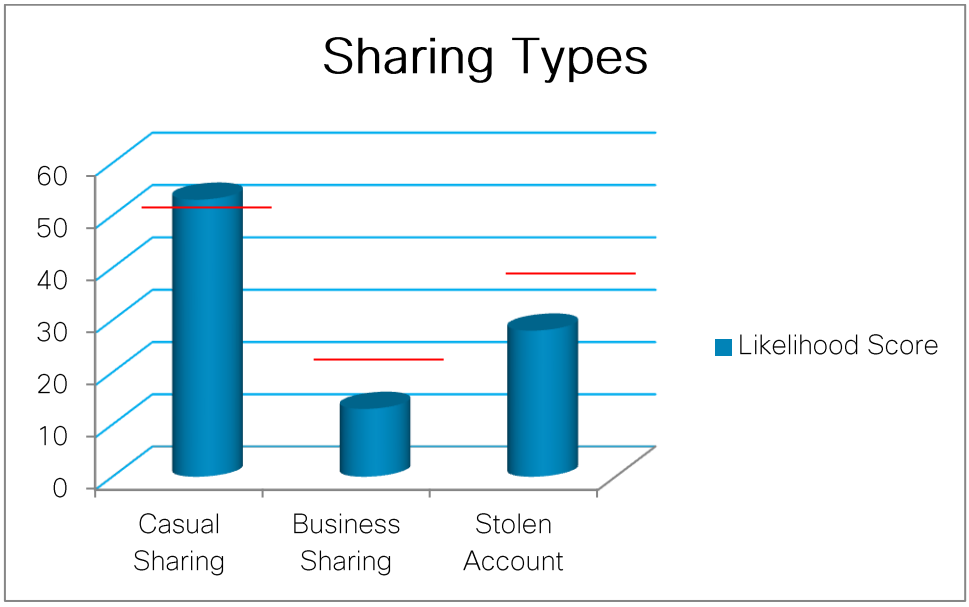

The purpose of this challenge is to improve confidence and verify the predicted results. The system challenges either the account owner or sharer with a question that is relevant to that sharing type. The question is based on knowledge of the nature of the relationship between the account owner and the sharer, which is different per sharing type (described in the initial table). A set of challenges for each sharing type is maintained by the system. A particular challenge is selected based on some combination of service provider preferences, analytic results and random selection. Examples of such challenges are:

- Asking the account owner to identify the shows the projected sharer has watched recently (useful for all casual sharing and business sharing)
- Asking the account owner or the projected sharer what was the last time the other watched online video (useful for business sharing)
- Asking the projected sharer esoteric information about the account owner collected during registration.
  Note: Not all information collected during registration is presented in the account details of the UI (useful for business sharing and stolen account).
- Changing the credentials of the account owner and detecting how many projected sharers still enter the old credentials (stolen account).
- Asking the projected sharer to prove that he is the owner of the mobile phone registered by the account owner via usage of SMS (all sharing).
- In case of a broadcast system, ask projected sharer to enter code or audio code written or played by STB (all sharing).

In addition, each response is timed in order to detect whether the account owner and sharer are colluding on answering the challenge. Note that in the case of stolen accounts, where there is no collusion, a challenge is never posed to the account owner since we do not want to disturb or frighten an innocent person.

If a challenge is answered correctly, the system returns to the initial state, where all projected sharers and account owners are active. If the challenge is not answered correctly, and the system concludes with high certainty that sharing is occurring on a specific account of a specific sharing type, the system recommends to the service provider various anti-sharing policies that manifest typical business goals of the service provider per sharing type. Such policies, as described in the state diagram above, include:

- Blacklist user based on behavioral profiles across all accounts

  Blacklisting is most important for stolen accounts, where the projected sharer cycles multiple credentials of different accounts without the knowledge of the account owners. In such a case, we do not want to risk losing the business of the innocent account owners by revealing to them that their credentials were stolen.

- Warn user of detected sharing

  This is the likely action for casual sharing, where the number of sharers is typically limited and fixed. In such a case, the service provider wants to maintain the business of the account owner even after collusion with sharers.

- Force the account owner to change credentials to highly secure credentials (business sharing or stolen accounts)

  Before terminating the account, mechanisms can be used to remove the sharers while maintaining the account owner, assuming that given the limited relationship between the projected sharer and the account owner, the projected sharer is not be made aware of the new account owner credentials.

- Terminate account (business sharing and casual sharing after multiple warnings)

  Upon gross sharing violations such as business sharing, the service provider typically desires to terminate the account immediately.

  Examples of mappings between likelihood scores per sharing type and policies taken may look as follows:

First Challenge Score (User)	First Policy Policy(User)	Second Challenge Score (User)	Second Policy Policy(User)
Casual Sharing			
60 (AO and PS)	Warn (AO and PS)	90 (AO and PS)	Suspend Account (AO and PS)
Business Sharing			
50 (AO and PS)	Change Credentials (AO)	70 (AO and PS)	Suspend Account |(AO and PS)

Stolen Account			
70 (PS)	Change Credentials (AO)	80 (PS)	Blacklist (PS)

Continuous Learning

Another part of this product uses feedback from the subscriber (such as a correctly-answered challenge) and Cisco's Operational Security team to continuously refine the classification algorithms. For each false positive (i.e. we suspected you but you proved that we were wrong by answering a challenge correctly) we should be able to refine and improve the model, thus decreasing false positives in the future. More importantly, by using some feedback on false negatives, e.g. intelligence gathered by operational security or customer service data detecting stolen accounts that were missed, we can refine the model to catch more bad accounts in the future.

High Level Architecture

The figure below shows the high-level architecture of our Media Sharing Protection system.

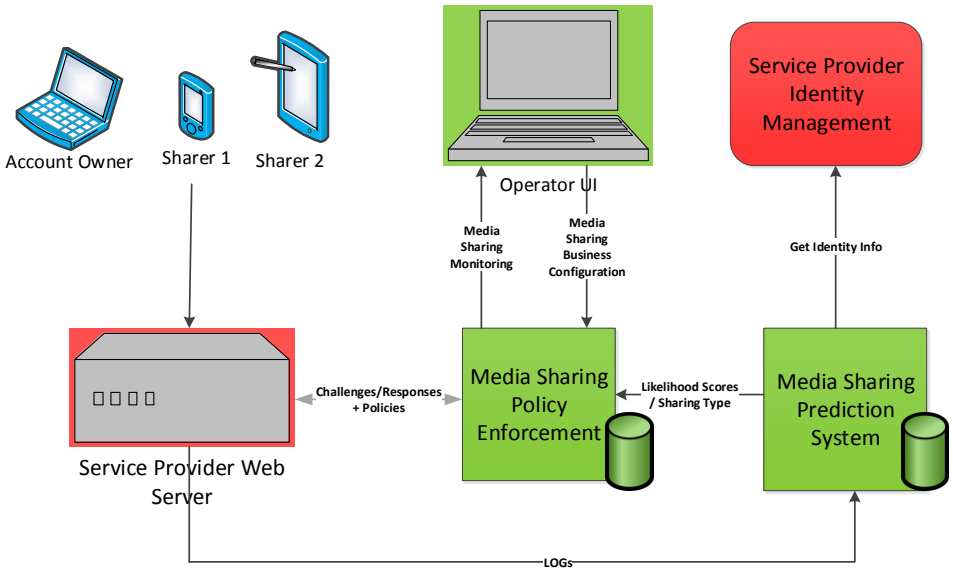

The core architecture of the system is described above. The Media Sharing Protection System is built on the following three components:

Media Sharing Prediction System (MSPS): This component is composed of typical analytics components such as Hadoop Filesystem, Spark, and Kafka, and receives logs of account creation, logins, and content selection. The component has access to the service provider identity system that describes the various users. The MSPS uses data science to determine the

likelihood score for each sharing type by clustering and classifying users based on their behavioral profiles. A diagram depicting the MSPS is shown below:

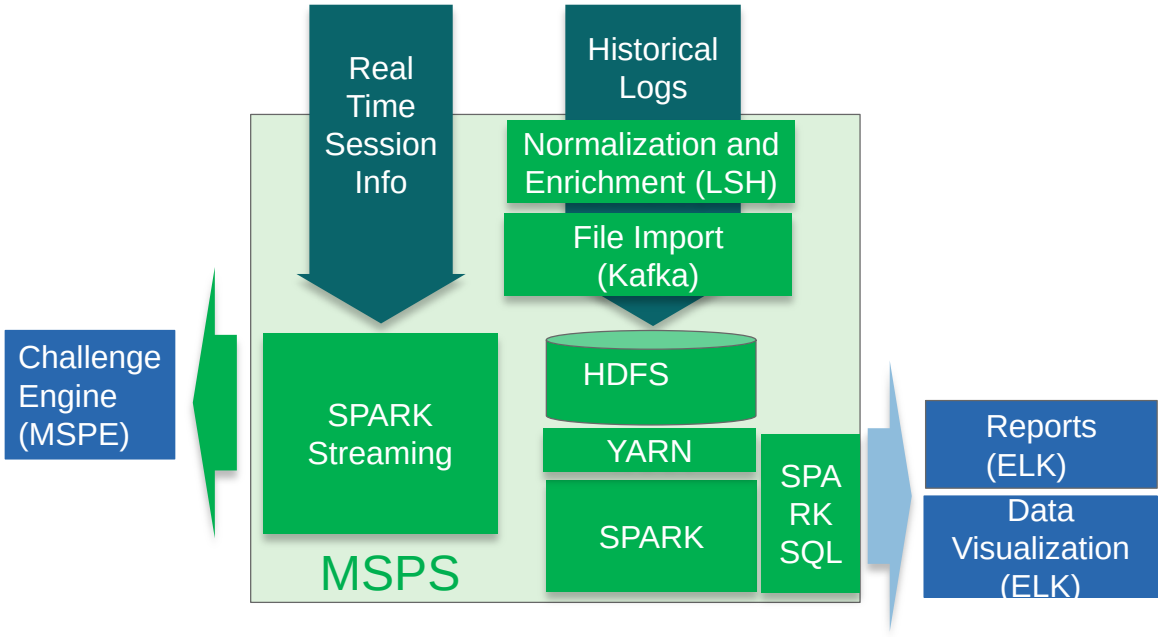

Media Sharing Policy Engine (MSPE): This component receives service provider configuration as input, such as:

- Which sharing type to enforce
- Permitted challenges
- Certainty required before challenge is enacted
- Certainty required before action/response is enacted

Based on this service provider input and the likelihood scores, the MSPE recommends challenges to the Service Provider Web Server, evaluates the response and recommends anti-sharing actions based on an evaluation of the response. The MSPE is also updated by online real time data as shown in the diagram above.

Operator UI: The Operator UI passes the service provider configuration to the MSRS. The Operator UI provides views, statistics and descriptive analytics of all sharing occurring in the system, and elaborates on identifying the projected sharers and the predicted projected sharing types.

