



Botnet Detection

Botnet Detection for Communications Service Providers

Table of Contents

Introduction:.....	3
Beyond the Enterprise Network.....	4
Network-based Botnet Detection.....	5
CnC Communication Identification.....	6
Focusing in on DNS.....	7
Focusing in on DNS (continued...).....	8
Focusing in on DNS (continued...).....	9
Focusing in on DNS (continued...).....	10
Qualifying a Botnet.....	11
What's in CnC?.....	11
Tying Botnet Detection to Alerting.....	14
The Future.....	17

Introduction:

Internet service providers are under increasing pressure to provide „clean pipes – to detect botnets and advanced threats present in their networks, alert or protect their subscribers, and ultimately provide assistance or advice to their subscribers on how to remediate the infection.

With the rapid proliferation of smart phones and other mobile internet computing devices, Telco/mobile/wireless providers now face the same issue, with the additional and potential painful ramifications that malware infections on subscriber devices could result in fraudulent charges appearing on subscriber bills, which will result in a nightmare scenario for their customer service operations.

While there are a variety of technologies designed to help corporations detect, mitigate and remediate botnets and other unauthorized remote control activities within an enterprise network, many of these same technologies are inappropriate for use within Internet service providers, telecom provider networks (wireless and landline), cable and satellite providers, and Internet backbone provider networks (all collectively referred to as Communications Service Providers “CSPs” within this paper).

This paper discusses the unique challenges Communication Service Providers (CSPs) face in protecting their networks and subscribers from advanced malware and botnet threats. It also describes some best practices being adopted by leading CSP s to detect infected “computing devices” (a term used in this paper to collectively refer to PCs, Macs, tablets, smart phones, or other mobile or Internet connected devices), and alternative mechanisms being used to alert the infected customer.

Beyond the Enterprise Network

While many core networking technologies are the same between enterprise networks and CSPs, there are some fundamental differences that impose constraints as to how security technologies can be practically implemented. Some of the most important differences include:

- **Enterprise networks have much greater control over the individual victim computers that lie within their own networks.** Desktop protection and cleanup technologies can be dictated and precisely monitored; meanwhile network-centric mitigation tactics can be applied in an ad hoc manner subject to corporate operational and employee guidelines. **CSPs, on the other hand, have no direct control over the victims within their networks.** Their relationship with the victim is that of service provider to customer and is consequently bound to specific service agreements and regulatory oversight.
- **The size of the network and associated traffic volume is significantly greater within CSP networks.** Many security technologies do not scale well and pose a significant hurdle to overall network performance when attempting to keep pace with the high volumes of streaming network traffic. CSPs need to prioritize network availability and quality of service.
- **The content of the data traversing an enterprise network is deeply important to the corporate security and forensic analysis teams.** The ability to inspect the traffic and retain samples for offline analysis is a typical prerequisite for business. **The precise content of the network traffic being transmitted over the CSP's network is largely unimportant.** Prioritization is given to the most efficient and reliable means of transporting the data between destinations with the highest degree of integrity.

Network-based Botnet Detection

While there is a growing number of network security technologies designed to detect and thwart attempts to compromise a networked computing device and subvert it for botnet use, the vast majority of these technologies are not suited to CSP consumption.

When it comes to the practical detection of botnets and the enumeration of botnet victims within a network, there are two primary categories of positive identification – detection of the presence of communications between the victim’s computing device and the botnet operators command and control (CnC) infrastructure, and the detection of malicious attack traffic being originated from a botnet victim’s computing device. Both techniques can positively identify the presence of botnet agents operating on the victim’s computing device – however, neither of them can reliably identify the precise malware version that has been installed.

The ease in which bot infected computing devices can be identified within a CSPs network is dependent on the precise nature of the botnet and the type of network traffic being consumed by the botnet detection technology.

Several of the classic (albeit “legacy”) criminal botnet threats are prone to producing high volumes of very noisy attack traffic. For example, criminal operators that specialize in spam delivery will issue commands to botnet victims that cause them to pump out high volumes of unsolicited email traffic to a number of targets over the SMTP protocol. Criminal operators that specialize in distributed denial of service (DDoS) campaigns in turn cause the victims under their control to spew forth extremely high volumes of attack traffic at a small number of targets (usually one at a time).

These kinds of attacks are very easy to identify using anomaly detection systems (ADS) – i.e. systems that identify significant changes to network traffic profiles and, based on the port or protocol transporting the traffic, can help identify the nature of the botnet.

Unfortunately, the criminals behind the current generation of botnets tend to be more sophisticated and circumspect in the way they conduct their illegal business. Identity theft, financial fraud, state- endorsed espionage, money laundering, click-fraud, voyeurism, piracy, intellectual property theft, vote rigging, and extortion are but a handful of the criminal campaigns now orchestrated via botnets. These criminal campaigns are unlikely to be even a smidgen of classic “attack traffic” and are likely indistinguishable from typical HTTP-based Web traffic. As such, identification of actual CnC communication plays a critical role in positively identifying the presence of botnet agents on the victim’s computing device, as well as distinguishing between the various botnet operators and their criminal objectives.

CnC Communication Identification

The presence of data communications between a victim's computing device and a criminal's CnC server can be used to good effect in identifying a botnet infection.

There are two primary methods for identifying botnet traffic:

- Interception, parsing and inspection of all data packets traversing the network and the subsequent identification of key data markers associated with known botnet communication profiles. This process is often referred to as Deep Packet Inspection (DPI). For example, some botnet malware families utilize a unique User-Agent variable in their HTTP requests, while others employ a particular structure to the CnC language.
- Enumeration and classification of the remote host to which the victim computing device is communicating. For example, many hosts accessible over the Internet are known to be criminally operated and serve as botnet CnC servers.

Combinations of the two methods are used extensively within enterprise network environments as they provide complementary visibility (and confirmation) of the threat. However, in the context of CSPs (particularly in the USA), DPI is often seen as synonymous with the invasion of privacy, which means that only highly watered-down implementations are used, if at all. As a consequence, CSPs increasingly favor the use of technologies that operate passively and make use of data streams and protocols that contain no personally identifiable information.

From a botnet detection perspective, technologies that can identify that communications to a known bad or suspicious computing device are taking place (or are about to take place) represent a minimal viable perspective on the threat. At its most basic, the detection technology needs to be supplied a list of known bad or malicious servers in the form of IP addresses and/or domain names. There are however a number of limitations to these approaches:

- **IP Addresses:** While all hosts and Internet services must utilize IP addresses, an IP address does not necessarily translate to an individual criminal system or threat. For example, a single Web server accessible via a single IP address may host hundreds of individual Web services and virtual sites – of which only a handful may be affiliated with the criminals. Other factors such as Cloud-based services, virtual hosting and DHCP lease times result in a high churn of IP address translations to a single physical server. Meanwhile, the criminals themselves can make use of fast-fluxing services and hacked “whitelistable” commercial sites to evade IP-centric blacklists.
- **Domain Names:** Host names and domain information make it much easier for Internet systems to locate the specific server they need to communicate with – regardless of the complexities that may be occurring at an IP address level. However, unlike IP addresses, there are an infinite number of address permutations, which means that static lists of known bad domain names are difficult to manage and unwieldy to maintain. In addition, malware CnC discovery and binding techniques that utilize domain fluxing and time-variable algorithm techniques evade domain blacklists.
- **Port/Protocol:** For some classes of legacy botnet threats, the TCP or UDP port and protocol being used for communication can be used to identify the presence of victim computing devices within the network. However, as ADS has become prevalent within both CSP and enterprise networks, such botnet malware has largely been consigned to the annals of history.

Focusing in on DNS

The richest and easiest data stream obtainable for the purpose of botnet detection with a CSP's network lies with DNS. Restricting botnet detection to the use of DNS is obviously a compromise situation. However, it is an efficient protocol that can often be enhanced with additional inspection layers (e.g. proxy log traffic) – depending on the nature (and location) of the CSP's business.

Focusing in on DNS (continued...)

DNS traffic offers a number of critical elements for detecting and enumerating botnets within a CSP environment:

- The vast majority of criminal services rely on DNS to manage and control their botnet victims.
- DNS is a well understood network protocol and is easily accessible within CSP networks.
- DPI technologies are not required to extract actionable intelligence for streaming DNS data.
- DNS data is typically deemed to be public and does not contain any personally identifiable information.
- DNS traffic is generally “low volume” when compared to data-carrying Internet protocols.
- Most CSPs already have extensive operational experience configuring and managing DNS systems.

DNS-based detection also offers a unique opportunity for botnet detection. Before a botnet infected victim can communicate with its CnC server, pass through stolen data, and receive commands, it must first locate the server’s IP address. By monitoring DNS, botnet detections can be made while the victim computing device is waiting for an authoritative DNS server to respond with the IP address – i.e. botnet victims can be detected before they even connect to the CnC server.

Monitoring DNS Traffic

Depending on the CSP and their network configuration, there are three logical places in which to monitor DNS traffic for the presence of botnet communications:

- Below the Recursive. In networks where the CSP hosts their own recursive DNS servers, observing the DNS traffic below the recursive provides a number of distinct advantages:
 - Enumeration of the specific IP address of the device making the DNS lookup.
 - Raw and non-cached visibility of all DNS requests.
 - Ability to detect botnet malware that does not honor DNS TTL information.
 - Ability to analyze the frequency of botnet DNS lookups.
- Above the Recursive. In networks where the CSP hosts their own recursive DNS servers, observing the DNS traffic above the recursive provides a number of distinct advantages:
 - Lower volume of DNS requests to be analyzed due to caching of the server.
 - Analysis of the responses from the Authoritative DNS server to the lookup request.
 - Automatic identification of technologies that modify DNS traffic.

Focusing in on DNS (continued...)

- Spanning port. In networks where the CSP either does not host their own recursive DNS services or allows subscribers to use third-party DNS services, observing the DNS traffic at network egress points or via protocol taps provides a number of opportunities:
 - Visibility of DNS requests and responses that bypass the CSP's recursive DNS infrastructure.
 - Visibility of DNS traffic above the recursive DNS supplied by the CSP.
 - Ability to identify IP addresses of devices that suspiciously mix SP DNS service requests with external third-party DNS requests.
 - The ability to identify IP address and domain name sets that are anomalous to the answers from the CSP managed DNS services.

In almost all past deployments, Damballa has found that access to DNS traffic below the recursive yields the highest botnet detection fidelity. The ability to clearly indicate which IP-enabled device made the DNS lookup request is the most useful element in being able to enumerate the victim and subsequently alert them to the threat. Visibility of spanning port DNS traffic broadens an CSP's ability to encompass victim computing devices that have opted out of utilizing the CSP's specified DNS resolvers – whether they chose to intentionally or because the botnet malware agent was programmed to do so.

Limitations of DNS Monitoring

While DNS traffic is the most economical way of detecting the presence of botnet victims within a CSP's network, it is important to understand the limitations of relying on it as a sole detection technology. These limitations are due in large part to the dynamics of how the protocol is used by the computing device, topology of the network and configuration of the botnet malware itself.

Focusing in on DNS (continued...)

The following limitations apply to DNS-only botnet detection systems:

- **Detection Thresholds.** It is important to understand that a victim computing device will make multiple DNS lookups of its CnC servers. These lookups may or may not be tied to the TTL supplied within the authoritative DNS response and may be cached by recursive DNS servers. For example, many botnet malware families ignore DNS TTLs entirely, while some botnets set TTLs measured in a few seconds as they rapidly flux the IP addresses for a cluster of CnC servers. As a consequence, it is not possible to determine the number of victim computing devices within a CSP's network by merely counting the number of DNS resolutions to known CnC servers.
- **Independent Look-ups.** The DNS resolution of a known CnC domain does not necessarily mean that the host making the lookup request is in fact a victim. There are a number of reasons why a computing device will try to resolve a known CnC server (e.g. the user has chosen to manually investigate a domain after reading about it, or has browsed a Web page that contained lists of malicious domains or links and the Web browser attempted to pre-fetch the content). However, the detection of repeated and cyclical attempts to resolve CnC domains over a finite period of time is a valid way of overcoming these false positives.
- **Network Address Translation (NAT).** Network configurations that make use of NAT tend to obscure the source of the DNS request. For example, while a CSP may service individual subscribers with a single IP-enabled DSL, cable or wireless modem, it is increasingly likely that there will be multiple IP-enabled devices sitting behind the modem. While DNS-based botnet detection will identify a botnet infection (or multiple infections), it is more difficult to clearly identify the specific victim device.

Qualifying a Botnet

When utilizing DNS traffic inspection as the basis for a primary botnet detection technology, it is critical that the knowledge pertaining to the criminal systems being used as the trigger for the alerting is both accurate and timely.

Botnets come in all shapes and sizes. Their infrastructure and business models are as varied as the individuals that control them and their methods of monetization. As such, considerable effort must be expended in qualifying and verifying the threat intelligence surrounding each criminal group and the botnets they operate. Rather than simply applying a long list of known CnC servers, additional metadata covering the dynamics of the botnet infrastructure and its precise threat classification are needed in order to provide accurate botnet victim attribution. Knowing whether a particular domain name is currently a CnC server and not something else is critical to a CSP's botnet detection technology.

What's in CnC?

The Internet infrastructure utilized by cyber-criminals for the distribution, infection, management and coordination of botnet monetization is sophisticated and ever changing. When choosing the best way to detect and enumerate botnet infestations, it is critical that botnet CnC infrastructure is correctly tracked and classified – and kept up to date.

Security teams should be aware of the following characteristics of botnets and their implication to detection strategies:

- **Multi-use Servers:** Some botnet operators host their CnC on the same server as their malware distribution repository and infectious Website. Just because a computing device is observed connecting to the server, it does not necessarily mean that it was compromised and that the malicious payload was installed and the botnet agent was successful in associating with the botnet CnC. Careful classification of domains to CnC-only operations is necessary to limit false positives – in combination with appropriate detection thresholds for DNS lookups by the “infected” computing device.
- **Multiple CnC per Botnet:** Most criminals employ multiple CnC servers for each botnet. For example, it is not uncommon for a single botnet to have hundreds of servers scattered around the globe and to have several thousand domain names associated with them at the same time. The botnet malware will often attempt to reach out and connect to multiple CnC servers – depending on the malware type and objectives of the criminal operator. Care should be taken to correctly associate clusters of domain names to a single botnet in order to provide the right level of alerting and attribution to the same criminal entity.

- **CnC Churn:** As expected with operating multiple CnC servers and a global array of infrastructure, there is often a lot of churn in the CnC servers themselves. For example, some botnets have been running for multiple years and have been associated with tens-of-thousands of domain names during that time – of which only a few thousand are “live” at any point in time. As CnC infrastructure is discovered, blocked, taken down or otherwise removed from the control of their criminal operators, most botnets are capable of handling this kind of infrastructure churn.
- **Legitimate Sites:** Botnet operators are not opposed to exploiting flaws and abusing terms and conditions of legitimate commercial sites. For example, some botnet operators make extensive use of hacked Web servers to host their CnC sites or choose to leverage popular “whitelisted” Websites for CnC channels. Both types of CnC typically get detected in short order and are normally remediated promptly. As a consequence, computing devices may be visiting these sites for legitimate reasons. DPI-based inspection is required to discern these kinds of CnC communications.
- **Domain Fluxing:** Some malware families make use of domain generating algorithms in order to evade static-lists of CnC domains. Depending on the date and time, the malware may seek to connect to hundreds of possible domain names – hoping to locate one that is currently under the criminal’s control – with the objective of receiving new command instructions. DNS-based technologies that identify the repeated DNS lookup and characteristic server responses have proven successful in locating victims whose botnet malware uses this technique.
- **IP-only CnC:** Some botnets employ IP-only CnC servers. DNS-based detection systems will not alert to these servers.

Increasingly Flexible CnC Enumeration Methods

The monitoring of DNS lookups and their authoritative server responses can be used with great effect in the enumeration of botnet infections within a CSP’s network. Botnet detection technologies that appreciate the dynamic nature of DNS and the ways in which cyber-criminal construct their infrastructure are a critical component to accurately detect botnet infestations. The ability to dynamically correlate between multiple threat databases, as well as understand and encompass the behaviors of the actual malware, is of significant value in both increasing detection fidelity and minimizing the possibility of false positives.

Evolving beyond classic blacklist maintenance, some of the characteristics lying behind an efficient DNS-based botnet detection system include:

- **Feedback Loops.** Systems that can observe high volumes of streaming DNS data, identify suspicious domain name resolutions, and feed back the new intelligence to CnC classification engines. For example, being able to detect and respond “congratulations, did you know you’re the first entity in the history of this CSP to ever lookup c2.botcontrol.obscured.org” and then being able to automatically discover that the domain was only registered 4 hours ago, has DNS name servers used by the XYZ Zeus botnet gang, and is pointing to a server that was hacked 18 hours ago and remediated 5 hours ago. This is probably a brand new CnC channel server and should be monitored.
- **DNS Server Response Thresholds.** Botnet malware that utilizes algorithm-based domain name CnC discovery routines produces “random” domain name noise, which is amplified as the recursive DNS servers, in turn respond with their equivalent of a “no such domain”. Frequency analysis and assisted machine learning systems can identify botnet malware that uses these techniques and, most importantly, can enumerate the particular botnet that has control over the victim machine.
- **Domain Reputation.** The historical relationship between domain name, authoritative name server and resolved IP addresses can be used to uncover the links between known botnet CnCs, malware infection sites and other criminal hosting infrastructure. Done in real time, dynamic reputation systems can rapidly identify new botnet infestations and provide associations with previously unclassified CnC infrastructure components – which in turn uncover new aspects of the botnet.
- **Fuzzy Factors.** Using a combination of unique DNS data points related to the recursive DNS response, it is possible to identify suspicious network activities that are indicative of a targeted threat that is yet to develop into a globe-spanning botnet. GeoIP, DNS hosting provider reputation, IP and netblock reputations can be combined to uncover a probable CnC. For example: identifying that the domain name being looked up is hosted as part of a free Dynamic DNS service in Iran, which has a very similar structure to 150 other names used for phishing attacks last month and is currently pointing to a cable modem IP address located in Brazil – while the computing device doing the DNS lookup is located in Canada.

An important aspect of these DNS-related discovery technologies is the ability to group and cluster new or suspicious domain names to particular criminal operators. Failure to provide this degree of association will result in a never-ending list of inactionable domain information. Machine learning algorithms can assist this process and, when backed up with a cloud- intelligence system, can automatically classify the new domains for use in DNS botnet detection systems.

For new domains that have been identified as suspicious but do not have any additional historical or reputational data in order to confirm their use as CnC infrastructure, additional time may be needed to acquire the new data for final “conviction.”

For example, it may be observed that computing devices that look up this particular suspicious domain also look up an additional three domain names within five seconds that are all known to be used for criminal botnet purposes.

Tying Botnet Detection to Alerting

Given the ability to identify victims and ascertain which particular botnet malware they are affiliated with at a particular point in time (using DNS observations alone), the CSP must then link the victim computing devices to a particular customer account and decide how to use the botnet discovery information.

Armed with an IP address, network segment ID and a precise date and time, CSPs can readily identify the customer account affiliated with the botnet victim. However, as mentioned earlier, there may not be a one-to-one match between the customer and a specific computing device.

When deciding on a response strategy to the botnet threat, the CSP must also factor in how long they should measure or track a specific infection. For example, the victim’s computing device may have been infected for two days before their host-based anti-virus software finally recognized the threat and cleaned it up. Meanwhile, the same computing device could have become infected again (probably through the same infection vector – e.g. an infected USB device) two weeks later and consequently reappear as a victim in the CSP’s botnet monitoring system. How long should a CSP track victims in order to ascertain the prevalence of botnets within their entire customer population?

The time needed to positively classify a customer's computing device as „botnet infected' is always going to be dependent on the objectives of the CSP. Some general recommendations include:

- **Overall Infection Metrics.** If the goal of the botnet enumeration is purely to provide information about the breadth and type of botnet infections within a CSP's network (without necessarily needing to correct or remediate the threat), then a time window of one or two weeks is probably sufficient. Tracking the "unique" subscribers (say by IP address but also accounting for DHCP lease renewals) infected over this period will reveal a lower-bound of botnet infections.
- **Remediation Dynamics.** If the goal is to track the success (or otherwise) of various protection and remediation technologies that are being deployed and trialed within the SP environment, then a time window of one to three months is probably precise enough and will remove many of the temporal variations of botnet building campaigns.
- **Customer Alerting.** If the goal is to construct a system designed to alert subscribers to the fact that they have been infected, then a one day-threshold system may prove to be useful. For example, to be designated as "infected" the victim must have been observed trying to connect to the same botnet CnC infrastructure at least once per day for an aggregate of five days within a single month (to account for days in which the computing device is turned off, roaming, etc.). Once the threshold has been reached, the customer is then alerted.
- **Reactionary Blocking.** If the goal is to use botnet detection alerts as a driver for blocking attacks and shaping network traffic, then lower time thresholds are advised. However, a careful balance needs to be struck in relation to the number of new botnet infections being detected per hour or day. CSPs should monitor specific botnets for growth rates on an hourly basis. If a new botnet CnC domain experiences a very high detection rate – but stays steady for several hours – it may be a false positive or a domain related to a newly hacked Website that also hosts legitimate content. If the alerts related to a domain grow in ways not clearly tied to standard usage patterns (e.g. diurnal Internet use patterns), then it is probably related to a new botnet outbreak or building campaign.

Customer Alerting

There are many different mechanisms available to CSPs for the purpose of alerting their subscribers to botnet infection; however CSPs must consider the costs of implementation and the opportunities for malicious abuse.

Some of the alerting (and response) strategies proposed or used by CSPs around the world include the following:

- **Walled Garden.** Subscribers with botnet infected computing devices find that they have limited access to the Internet. In some cases, this limitation may be restricted to the blocking or filtering of particular ports and protocols commonly associated with attack traffic (e.g. SMTP for spam botnets, or port 31337 to servers outside the CSP's network). While in other cases, Internet access may be restricted to Web content and portals hosted by the CSP – or something in-between both walled gardens. The general purpose of a walled garden is to prevent further malicious activity from taking place and provide limited access to tools or information on how to remediate the infected system.
- **In-Session Alerting.** Subscribers with botnet infected systems are presented with messages within their Web browser (or through email) alerting them to the fact that they have been identified as being a botnet victim. These messages typically identify the nature of the threat and provide advice on how best to remediate the infected system. Much like fake-AV alerts, however, this approach can be subject to social engineering attempts by criminals.
- **Blocked.** All Internet traffic to and from the botnet infected system is blocked. The subscriber is expected to contact the customer service department through another system or communication technology in order to understand (and eventually remediate) the infection. This kind of alerting strategy is more popular with mobile and cellular network operators.
- **Quality of Service (QoS) Modification.** The infected subscriber is alerted to the fact that (all or some of) their Internet traffic will be routed using depreciated rules until their systems or network are cleaned. This kind of strategy is more applicable to backbone carriers as they attempt to throttle high volumes of malicious and unwanted traffic.

The Future

Organized crime will continue to tune and optimize the techniques they use to conduct fraud and monetize botnets over the Internet. If the past decade is anything to go by, then the attacks will become more sophisticated and the breadth of fraud techniques will continue to grow. However, given the power and versatility of botnets to serve as the core infrastructure for online crime, we can expect botnets to become even more ubiquitous and further resilient to takedown.

As the devices and platforms used by CSPs and their subscribers continue to change, we can expect the criminal operators to adopt new attacks and fraud vectors that target these very same systems if there is sufficient financial incentive to do so – which is highly likely given current trends. The increasingly diverse range of platforms and vectors for abuse will pose few problems for the botnet malware and CnC infrastructure currently in use by organized crime units. Today's menagerie of cyber-crime tools is already largely sufficient for the future.

It is anticipated that subscribers that fall victim to cybercrime will increasingly look to their CSPs for protection against the threat. The ability to rapidly detect a targeted attack, a breach of the network or a compromised system from within the CSP's cloud will become better understood by their subscribers – providing greater opportunities for those CSPs to satisfy those security concerns with value added services (assuming that the CSP doesn't intend to use this kind of "security" as a differentiator among competing CSP businesses). As a consequence, the opportunity for CSPs to incorporate additional detection technologies that employ DPI or other technologies that may have had some historical concern over the inspection of personal information contained within a subscriber's Internet traffic, could be granted on a per-subscriber basis – with subscribers choosing to opt-in to the enhanced security service offerings.

In addition, knowledge of which subscribers are botnet victims (or are suspected of being victims) will be an increasingly valuable commodity. Many of the organizations and services a CSP's subscriber does business with over the Internet on a daily basis would find this information to be extremely valuable to them. For example, an online bank which is informed that the subscriber's system is likely a botnet victim can increase the level of fraud detection they apply to the online session – optimizing their service delivery and issue fraud watch alerts – quickly detecting fraud attempts and better protecting their subscriber against loss. It is possible that CSP subscribers may choose to opt-in to this kind of data sharing if they felt more secure against fraud or, perhaps, reduced the premiums their online service providers charged them in turn.