Position paper submission for W3C/IAB workshop on Strengthening the Internet Against Pervasive Monitoring (STRINT)[1]

Andrei Robachevsky, Christine Runnegar, Karen O'Donoghue, Mat Ford

Introduction

The ongoing disclosures of pervasive surveillance of Internet users' communications and data by national security agencies have prompted protocol designers, software and hardware vendors, as well as Internet service and content providers, to re-evaluate prevailing security and privacy threat models and to refocus on providing more effective security and confidentiality.

At IETF88, there was consensus to address pervasive monitoring as an attack and to consider the pervasive attack threat model when designing a protocol. One area of work currently being pursued by the IETF is the viability of more widespread encryption. While there are some who believe that widely deployed encryption with strong authentication should be used extensively, many others believe that there are practical obstacles to this approach including a general lack of reasonable tools and user understanding as to how to use the technology, plus significant obstacles to scaling infrastructure and services using existing technologies.

As a result, the discussion within the IETF has principally focused on opportunistic encryption and weak authentication. "Weak authentication" means cryptographically strong authentication between previously unknown parties without relying on trusted third parties. In certain contexts, and by using certain techniques, one can achieve the desired level of security (see, for instance, Arkko, Nikander. Weak Authentication: How to Authenticate Unknown Principals without Trusted Parties, Security Protocols Workshop, volume 2845 of Lecture Notes in Computer Science, page 5-19. Springer, (2002)). "Opportunistic encryption" refers to encryption without authentication. It is a mode of protocol operation where the content of the communication is secure against passive surveillance, but there is no guarantee that the endpoints are reliably identified.

In this paper, we offer a complimentary analysis. We identify some of the components of the Internet architecture that provide attractive opportunities for wholesale monitoring and/or interception, and, therefore, represent architectural vulnerabilities, or choke points. We also suggest possible mitigation strategies and pose some of the questions that need to be considered if the Internet is to evolve to reduce such vulnerabilities. Finally, we identify some significant areas of tension or trade-offs, and we consider possible areas for additional efforts.

The threat model - pervasive attacks

As defined in the Internet Draft "Pervasive Monitoring is an Attack" (http://datatracker.ietf.org/doc/draft-farrell-perpass-attack), "…'pervasive monitoring' means often covert and very widespread intrusive gathering of protocol artefacts including application content, protocol meta-data such as headers, or cryptographic keys used to secure protocols. Active or passive wiretaps, traffic analysis, correlation, timing or measuring packet sizes can also be used as part of pervasive monitoring." The draft also notes that adversaries in this context vary from nation state actors performing surveillance, to commercial enterprises performing legal but privacy-unfriendly traffic analysis, to criminals.

The key aspect that qualitatively differentiates "pervasive monitoring" from other similar attacks (e.g. MITM[2], traditional eavesdropping) is that it is a long-term, untargeted (or indiscriminate) data collection (up to all communication flows available at a particular point) allowing correlation of data flows over a long period of time. Time is an important aspect here as the collection and storage of large data sets means that analysis can be performed at a later time (and maybe with stronger capabilities) should the motivation arise. Notably, the potential scale and reach of the threat is magnitudes greater than more traditional threats.

---

[1] N.B. This submission does not constitute a formal position of the Internet Society and represents joint work of the identified authors only.

[2] Man in the Middle attacks (see RFC 3552 BCP *Guidelines for Writing RFC Text on Security Considerations*)

Moreover, "pervasive monitoring" is a specific threat to individuals' rights and expectations of privacy.

An architectural view: choke points

The Internet was designed to avoid choke points (single points of failure) or to mitigate their impact. Originally, that meant resilience against failures at the IP layer.

However, the concentration and centralization of certain functions at various layers of the Internet architecture has created new choke points and, consequently, facilitated new threats such as pervasive monitoring. Accordingly, it would be useful to consider what are the relevant choke points from an architectural perspective that present the greatest vulnerability to pervasive monitoring and/or offer the greatest access to communications data or metadata.

For the purposes of this paper, we refer to Internet layers that do not necessarily correspond 1:1 to the TCP/IP model.

1) **Datalink layer** – eavesdropping on physical communication links, e.g. optical cables. Intercontinental and inter-regional cables create points of high concentration of data flows. Large Internet Exchanges (IX) providing L2 interconnections between different networks (regional and international) also present choke points at this layer. An adversary can install wiretapping facilities, thus getting access to a high volume of data and a large number of users.

    - Mitigation strategies: a) Encryption of individual communication flows, for instance using IPsec or TLS. This still leaves metadata – protocol headers – accessible to the adversary. b) Encryption of all traffic between the IP forwarding end-points of the link (i.e. L2 encryption between an ISP's routers or switches). This also hides metadata, but requires additional efforts (and costs) from the ISP. c) Diversification of the infrastructure: richer internetwork connectivity, more diverse communication flows (also dependent on diversity of destinations), local traffic exchange (e.g. stimulating traffic exchange at regional IX'es).

    - Challenges (forces against the mitigation strategies): a) Obstacles to achieving economies of scale – inter-regional and intercontinental links are few and high capacity. b) Performance implications – encryption at the wire speed for such links has high performance and cost penalties.

2) **IP layer** – rerouting (redirecting) specific data flows. An adversary can redirect traffic to cause it to be routed via their network (by manipulating BGP announcements – hijacks or route leaks), possibly reinserting the redirected traffic back onto a "clean path" to reach the destination. While, in general, there is no guarantee that a traffic flow will traverse (or not) specific networks, intentional concentration of traffic flows in a single network is a threat. One may observe that this threat exists in most Tier-1 networks.

    - Mitigation strategies: a) Encryption of individual communication flows, for instance using IPsec or TLS. This still leaves metadata – protocol headers – accessible to the adversary. b) Routing (policy) security measures that are usually a combination of best practices of routing hygiene (filtering customer and even peer announcements) and technologies (RPKI, BGPSEC, VPNs).

    - Challenges (forces against the mitigation strategies): a) Encryption may conflict with some ISPs' network management practices (e.g. traffic engineering, troubleshooting). A possible side effect could be that encrypted traffic suffers lower performance, thus discouraging the use of encryption. b) Security and an especially guaranteed path is not part of the normal ISP service offering (connectivity is), so as long as connectivity is not affected there is little interest in mitigating these attacks (a "route leak" by a customer, for instance, may be seen as financially beneficial to the upstream provider (because the customer pays for the traffic) as long as this doesn't cause performance degradation).

3) **Application layer** – several applications/services now attract substantial concentrations of Internet users. Examples: Google services, Facebook, YouTube, Yahoo!, Baidu. "Infrastructural" services are relevant here as well: e.g. Google's public DNS, OpenDNS. This results in very high concentrations of user transactions, related to exchanged data, choices and preferences as well as metadata, that allows service providers, affiliates and adversaries to infer information about the user, their connections, habits, choices, preferences, etc. Additionally, increased features and functionality at the application layer (e.g. in the browser) mean that there are more observable characteristics (such as user agent, headers, fonts, colour depth, screen size, time zone, plug-ins/extensions, sensors) available for fingerprinting, and, therefore, tracking. Mobile device sensors (e.g. GPS receiver, accelerometer, gyroscope, microphone, light, proximity sensors) add further data points for fingerprinting. Also, the reach of some content providers has expanded due to acquisitions of online advertising companies, which makes proliferation of tracking capabilities even higher (see, for instance, Balachander Krishnamurthy "Privacy leakage on the Internet", http://www.ietf.org/proceedings/77/slides/plenaryt-5.pdf). Various types of attacks, from static and dynamic key exfiltration to content exfiltration (see http://datatracker.ietf.org/doc/draft-barnes-pervasive-problem), can be performed by an attacker against these large concentrations of user data.

    - Mitigation strategies: a) End-to-end security. For services where the end-to-end connection is established between the users (e.g. VoIP, P2P) encryption of the communication channel can be an effective mitigation, but only in the case where users are in possession of keys to encrypt the data. b) For non peer-to-peer systems (e.g. mail systems, websites and search engines) object security can mitigate the threat in some cases, allowing the user to encrypt the content of the message. For instance, users can use PGP or S/MIME to encrypt the content of their messages, making it inaccessible to network attackers and intermediate servers. In most cases, though, user metadata (search keywords, visited websites, mail senders, recipients, subject, timestamps, IP addresses, OS and browser fingerprints, etc.) is still accessible. c) Social/content internetworks. The e-mail system was designed to allow users of various administrative domains to communicate with each other. The ability to interconnect diverse content and social networks without significant loss of functionality may improve the diversity and limit the size of administrative domains in these areas.

    - Challenges (forces against the mitigation strategies): a) The network effect is the primary driving force for the enormous growth of some social networks, while others vanish. The bigger the user base, the more attractive the network is for users to join and the more likely it is to attract advertisers that, in effect, cover the financial cost of the service. b) The business model. Many of the services are offered for no financial cost (the so-called “free” service) assuming that people are willing to trade their personal data and privacy for this “free” service. No one forces people to use Gmail[3], or Public DNS, but convenience, accessibility, reliability, usability and cost savings win.

It is important to stress that intermediaries that operate the choke points are not the adversaries (they are observers in the RFC 6973 definition, or witting or unwitting collaborators in “Pervasive Attack: A Threat Model and Problem Statement” (http://datatracker.ietf.org/doc/draft-barnes-pervasive-problem)), but because of their access to high concentrations of user communications, data and metadata, they may be vulnerable to pervasive monitoring attacks.

Trade-offs

Protection against pervasive attacks comes with a cost and often involves making trade-offs, which in turn can influence the deployability of the solution and its effectiveness. Since the pervasive

[3] Gmail offers its users strong security when accessing their accounts that includes 2-factor authentication and the encryption of communication between the user and Google [http://www.google.com/landing/2step/]. But in general the content of messages and the metadata are available to Google. It is possible to use PGP or S/MIME, but not through the web interface.

surveillance threat is global in nature, these trade-offs may potentially affect the whole Internet ecosystem. Difficult decisions will need to be made at each architectural layer and in respect of any proposed mitigation strategy as to whether the risk of pervasive monitoring is worth the trade-off(s). This exercise may prove to be more challenging than developing new technical solutions, as it is likely that: interests will be divergent and sometimes in direct conflict; the whole community (not just the protocol designers) will have to step up and accept shared and collective responsibility for these issues; this is a long-term problem space (as mitigations are deployed, new attack vectors will be discovered). Of course, trade-offs are not necessarily absolutes (e.g. it is possible that usability failures could be avoided if considered in the design phase).

Examples of trade-offs that will need to be considered:

- Privacy/Anonymity vs. Strong authentication/identification
  - Both may be desirable for security purposes, but they can conflict.
- Encryption of content and/or metadata vs. Network operations and management
  - How might mitigations against pervasive monitoring affect intermediaries' ability to legitimately manage their network?
- Encryption of content and/or metadata vs. legitimate law enforcement activities and national security
  - Pervasive use of strong encryption, especially end-to-end encryption may significantly hinder some legitimate law enforcement activities and may have an impact on national security.
- Security vs. Usability
  - This is the well-known trade-off leading to people writing passwords on post-it notes or re-using them. Sending and receiving encrypted email is another example. Using S/MIME or PGP is not "plug-and-play".
- Security vs. Deployability
  - A protocol with built-in strong security may be harder to deploy, since it may require additional operational expertise and costs, or introduce the dimension of time to the operations (e.g. key management).
- Privacy/Security vs. Functionality/Features/Service
  - Could a "free" service be privacy-respecting? How much are users willing to pay, and for what services, to get enhanced security and privacy? What features and functionality are users willing to pay for?

Driving forces for change

Users: They have to be willing and able to use enhanced security options, maybe at the cost of less convenience/performance. Their demand for security services from their ISP, in cases where infrastructure support is required, could also be a very effective driver.

ISPs/network operators: They have to be willing to introduce changes to their infrastructure and incur additional (capital and operational) costs. There may be legal requirements for traffic retention and/or interception that ISPs have to abide. There might be operational practices (e.g. traffic engineering/DPI, accounting) that may conflict with higher security (e.g. encrypted traffic).

Content and service providers (web, CDN, providers of cloud services, mail operators, SIP proxies): They have to be willing to beef up their infrastructure to support higher security (e.g. TLS support, two-factor authentication). They may have to deploy higher privacy constraints, which may conflict with their business model.

Governments: They have to be leaders – supporting the open, inclusive, transparent development of permission-free, interoperable, globally-applicable, privacy and security solutions, encouraging competition in the market, allowing users to choose the services they want (i.e. not prescribing or prohibiting the use of technology), providing financial support to open independent research groups, offering financial incentives for deployment and R&D, etc.

Conclusions

In this paper, we provide an architectural context to frame the workshop discussion, document some vulnerabilities and mitigations based on that framework, discuss some basic areas of tension or trade-offs, and identify some key players capable of driving change.

A possible outcome of the workshop is a framework of potential actions and action owners along with an acknowledgement that the problem is much bigger than just the underlying technical specifications and the resulting products. It is a multi-stakeholder problem with contributing roles for a large number of parties. In conclusion, we would like to offer some examples of more concrete efforts that could be undertaken by various parties to implement some of the mitigation strategies mentioned above. These are some preliminary considerations for further discussions at the workshop, rather than a proposed plan of action.

Effort	**Parties (actors, stakeholders)**
Fix weaknesses in current specifications	IETF and other standards bodies (work ongoing, further examples to come out of workshop)
Investigate users' needs and requirements for usability	Academic community (e.g. PETS / SOUPS), maybe IRTF, education community (like ISOC, EFF, consumer groups), software developers
Reduce deployment and operational costs	NOGs (information sharing, operational practices) IETF (considering "good enough" security – e.g. weak authentication & opportunistic encryption) W3C (e.g. standardizing a JavaScript API for performing basic cryptographic operations in Web applications)
Provide incentives for development and deployment	Government (e.g. through government procurement, financial incentives)
Provide architectural guidance on peer-to-peer communication and infrastructure diversity	IAB
Facilitate the diversification of infrastructure and services	Various capacity building activities, standardization to facilitate interoperability of smaller service domains
Education of users	Academia, media, software developers/vendors
Capacity building – policy makers	ISOC and other Internet community organizations
Increase trust in the cryptography related product lines	Open source community