Joint White Paper

Robust Localization Towards Trust-Enhanced Networking

Data & Security Research Laboratory, Fujitsu Limited and
Cyber Security Research Center @ Ben-Gurion University of the Negev

(Overview) Internet geolocation technology aims to determine the physical (geographic) location of Internet hosts, users, and cloud data. It has been proposed and is currently used for a wide variety of purposes, including targeted marketing, the restriction of digital content sales to authorized jurisdictions, and security applications such as reducing credit card fraud. This raises questions about the authenticity of claims of accurate and reliable geolocation. In this white paper, we present Robust Localization as a key component of the concept of Trust-Enhanced Networking (TEN) in bringing trust to cyberspace through the verification of physical attributes of assets in the real world. We explore the existing technology and identify issues and technical challenges, and from these we outline our direction and strategy for addressing them.

1. Introduction

The Internet connects hosts from all around the world. Sometimes it is valuable to know where, geographically, a particular host or data is located. Informally, Internet geolocation, a.k.a. IP geolocation, is used to solve the problem of determining the physical location of an Internet user or device. The development of Internet geolocation technology is being driven by a number of practical uses such as advertising, zero-trust security, privacy regulations, and location-aware services; targeted advertising is among the most lucrative of these applications. For example, if a web server is able to determine that a visiting user is in Seattle, the server can embed advertisements targeted to Seattle customers in the page served. Also, Websites often tailor content (other than ads) based on geographic location. Other suggested applications of Internet geolocation include automated redirection to nearby servers and web analytics (i.e., analyzing web page access logs to extract marketing data). Besides, we believe it is particularly important for businesses that originate in cyberspace to verify the location of their partners in order to detect impersonation attempts and validate their authenticity from a geographical viewpoint.

In its previous white paper [1], Fujitsu proposed the concept of Trust-Enhanced Networking (TEN), which is a novel networking paradigm. TEN uses geographical information about network entities to provide reliable information about physical attributes of real-world entities and improve users' trust in the overall network. For this purpose, TEN performs step-by-step mapping from physical space to cyberspace. In other words, the verified geographical location of network entities is translated into a trust-aware network map that provides network control functions. In the context of measuring trust, Fujitsu also introduced Quality of Trust (QoT) which is a logical continuation of Quality of Service (QoS) and Quality of Experience (QoE). QoT is computed using information from the verified geographical location of network entities.

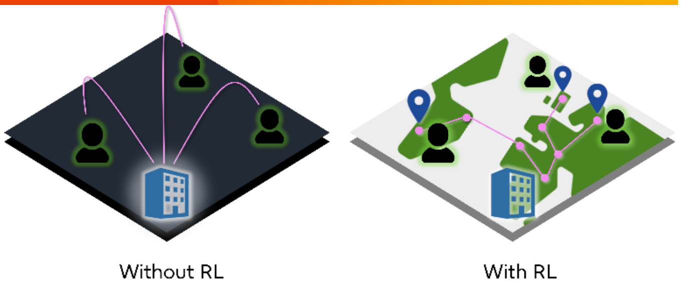

Figure 1: In RL trusted telecommunication is made possible by verified geographical location

In this white paper, we outline the technical challenges associated with network location verification which is one of the significant issues in TEN. Location in the context of the Internet usually means an IP address. However, the granularity of geographic location is often important, with implications on the ability to verify the identity of the target and secure data from theft. Furthermore, while Global Positioning System (GPS) is well-known as a method for tracking geographic location, GPS information is vulnerable to spoofing. There are also a number of approaches for inferring target location based on real-time active measurements via networks (such approaches are referred to as active geolocation methods). One of the promising solutions is Robust Localization (RL), which is also a key technology in TEN. As shown on the left side of Fig. 1, physical attributes are usually obscured when performing activities in cyberspace, however RL sheds light on the physical attributes of entities involved in communication in cyberspace in this setting. RL is one of the active geolocation technologies utilizing network measurements with the potential to provide more reliable results when dealing with malicious attackers and threats. Fujitsu and Ben-Gurion University of the Negev (BGU) have been

Table 1: Literature review

Domain	*Ref.*	*Year*	*Method*
Delay-based geolocation & accuracy enhancements (landmarks, PoP)	[2]	2021	multilayered common routers
	[3]	2019	nearest common router
	[4]	2021	sorting nodes by their subnets
	[5]	2021	using IP webcams as landmark database
	[6]	2019	statistical learning to study localized delay and hop distance correlation
	[7]	2019	boundary nodes
	[8]	2018	closest common router
	[9]	2015	mining Internet forums
	[10]	2019	PoP partition algorithm
	[11]	2018	low-high-low delay distribution
	[12]	2016	IP city-level geolocation
Analysis methods (data mining & machine learning)	[13]	2016	neural networks with only a few landmarks
	[14]	2022	graph neural network combining IP host knowledge and neighborhood relationships
	[15]	2021	using neural networks to cluster landmarks
	[16]	2018	reverse DNS
	[17]	2017	reverse DNS
	[18]	2016	utilizing users' search queries
	[19]	2010	utilizing Facebook friendships
Cloud (data geolocation)	[20]	2020	"weak" and "strong" proofs
	[21]	2020	atomic proof method, which divides the proof into several minimum computation units
	[22]	2015	constraint-based cloud data geolocation using unclonable and tamper-proof device
	[23]	2013	querying an agent that is installed with the data
Geolocation and Proxy/VPN	[24]	2022	measuring and comparing the TCP handshake and TLS handshake
	[25]	2018	verification of the location information of 7 proxy services
Attack	[26]	2022	a system that allows verification of the claimed geolocations of network nodes in a fully decentralized manner
	[27]	2017	methods of hiding a host's real location
	[28]	2014	route hijacking
	[29]	2010	attacks on delay-based IP geolocation techniques
	[30]	2004	hiding IP geolocation using onion routers
RFC	[31]	2013	authorization policy language for controlling access to location information
	[32]	2020	format used by network operators to publish geolocation feeds

performing joint research aimed at exploring reliable geolocation technologies since September 2022, with the goal of developing technologies for trusted communication on the Internet.

2. Technical Survey and Insights

In this section, we present the findings of a survey conducted about geolocation. There are various approaches for improving the accuracy of target node inference. Table 1 lists studies addressing the target IP localization task proposed in the literature.

As seen in Tab. 1 the problem of target IP geolocation has been the subject of extensive research. However, the aspects of "trust" and "robustness" have not yet been considered. Moreover, most of the proposed geolocation solutions are prone to security attacks. As seen in Table 1 the problem of target IP geolocation has been the subject of extensive research. However, the aspects of "trust" and "robustness" have not yet been considered. Moreover, most of the proposed geolocation solutions are prone to security attacks. Given that, delay measurements can be manipulated by a sophisticated adversary, and machine learning methods can be fooled by an adversary demonstrating normal behavior. Moreover, a powerful adversary can hijack routes and damage the network load balance.

The literature review also shows that the delay measurement method, which is the basis for most of the techniques proposed, is unable to perform highly accurate geolocation. This is mainly due to the disadvantages associated with the delay error. As a result, delay-based techniques alone can only achieve district-level accuracy. While machine learning based methods achieve street-level accuracy, they require additional information for training. Similarly, data mining-based techniques achieve the same accuracy as machine learning based techniques but require additional information. In general, passive geolocation methods suffer from coverage and reliability issues. This is because there is no guarantee that a target IP of interest has been previously indexed and that a record of its exact location exists in IP mapping databases. Even if such information is available, it may be inaccurate due to outdated records, intentional misinformation, or a lack of specificity in the location.

On the other hand, active geolocation methods achieve street-level accuracy with less coverage and reliability issues. For instance, the use of landmarks and Point of Presence (PoP) improves the accuracy of basic delay and Round-Trip-Time (RTT) measurements; however, there is the need to identify reliable landmarks or PoP addresses.

Cloud storage services are extremely popular, however data geolocation is still in its early stages. We observe that this topic has not yet been thoroughly investigated, and therefore the very few existing solutions mainly rely on solutions for the well-known and extensively researched IP geolocation problem. Users of cloud services do not trust that their data is safe, and there is a need for research aimed at developing solutions that that geolocation and privacy regulations are being enforced.

Another two commonly used services are the proxy and Virtual Private Network (VPN) which let users "hide" their identity or geolocation. This allows users to fake their identity and consume services that are not typically available to them due to their geographical location. We believe that geolocation techniques can address many of these trust issues.

Lastly, adversarial geolocation is also researched. We believe that handling adversarial scenarios may increase the level of trust. Most of the existing solutions for IP and cloud geolocation are prone to adversarial attacks. However, there are very few solutions for adversarial geolocation scenarios. Thus, there is a need for research aimed at understanding the effect of attacks such as route hijacking, IP spoofing, and route obfuscation on geolocation and trust in general.

It is also worth mentioning two important Request for Comments (RFC) papers. RFC8805 [32] defines a format used by network operators to publish geolocation feeds that include coarse-grained geolocation of IP address prefixes. With a sufficiently high level of granularity, feeds published by authoritative entities could contain trustworthy landmarks. RFC6772 [31] defines an authorization policy language for controlling access to location information. More specifically, it defines condition elements specific to location information in order to restrict access to data based on the current location of the target.

3. Requirements

TEN, which defines trust on the Internet, consists of two factors that serve as the basis of trust in this setting: the trustworthiness of communicating entities and the trustworthiness of data. Trust is integral to communication, and one of the first steps in engaging in communication is to identify the entity that you are communicating with. This rule does not apply to malicious entities impersonating a benign entity in order to steal information or use services without permission.

In addition, users (e.g., service consumers) would like to know that the data they are using is reliable and trustworthy (i.e., that the data has not been tampered with and is safely stored). This requirement does not apply to Cloud Service Provider (CSP) that is relocating data in the cloud or data that has been tampered with by a malicious adversary.

Ensuring trust on the Internet is an issue requiring deep thought and understanding of cyberspace. Users and service providers can easily falsify their identity by performing IP spoofing in IP packets that have a modified source address created. Users can also hide behind a proxy server or VPN, making it very hard to identify them. Geolocation manipulation is also easy to perform and difficult to identify; this can be done by changing physical attributes such as delay measurements. Furthermore, authenticating the location of data that is stored on the cloud is a complicated task, since the clouds' internal architecture is not published.

Trust in the context of geolocation: Geolocation authentication is important for trust. Namely, if organization A verifies the location of Internet resource B, then A trusts that the location of B is correctly reported. Moreover, customer C that trusts A will also trust that B's location is correctly reported.

3.1. Problem Statement

Communication and services in the metaverse are sometimes based on geographical location. For example, the restriction of Japanese TV outside of Japan, a hacker performing reconnaissance from a spoofed location, or a bank account being accessed from a spoofed location. Therefore, users and service providers need to corroborate the location claim of a user against the location claim of that user. In most cases, corroboration is based on delay measurements of the communication between the users. This method has city-level accuracy, and therefore, it can be enhanced by the use of other methods such as machine learning and data mining. Such methods open up new questions about trust.

In this white paper, we provide an overview of existing methods for verifying the geolocation of a target host and propose some new ideas for enhancing the trust of network communication, which is part of the concept of TEN.

We also discuss the main problems in verifying data geolocation; the authentication of data geolocation is becoming a critical issue with the increasing use of cloud services. Data location verification remains an unexplored research domain, and the industry's awareness of the trust implications of these services is quite limited.

3.2. Threat Model

We model trusted geolocation as a three-party problem involving the user, the adversary, and the Internet. The geolocation user, who can also be the victim of the adversary, aims to accurately determine the location of the target using a geolocation algorithm that relies on the target's ground-truth information and measurements of network properties. We assume that: (1) the user has access to a number of landmark machines distributed around the globe to obtain measurements of RTT and network routes, and (2) the user trusts the results of measurements reported by landmarks.

The adversary would like to mislead the user into believing that the target is at a forged location of the adversary's choosing when in reality the target is actually located at the true location. In a similar manner, a CSP can unintentionally or maliciously report a forged data location. More dangerously, a user can hide behind a proxy server or a VPN in impersonation or geolocation-based attacks. Adversaries can manipulate geolocation algorithms and hijack routes. This, of course, requires the adversary to have a different scale of resources.

Such systems are not robust to malicious adversaries that lie about their location or obtained measurements, or to malicious targets that strategically manipulate timing measurements by, e.g., delaying responses to certain timing probes.

The third party in the threat model is the Internet itself. While the Internet is impartial to both the adversary and the user, it introduces additive noise as a result of queuing delays and circuitous routes. These properties introduce some inherent inaccuracy and unpredictability into the results of measurements on which geolocation algorithms rely. In general, an adversary's malicious tampering with network properties (such as adding delay), if done in small amounts, is difficult to distinguish from additive noise introduced by the Internet.

4. Technical Challenges

TEN's requirements for robust localization, as described above, necessitate that multiple technical challenges be addressed. It is obvious that TEN requires a broad and deep understanding of network trust, but there is also a need to understand the vulnerabilities of location-based services. Namely, we need to identify the network's security and trust weaknesses and propose techniques for enhancing the robustness of geolocation-based services. Moreover, a deeper understanding of the technical requirements of new technologies and environments such as the cloud is needed to enhance trust.

The notion of trust. Integrating the notion of trust into existing geolocation technologies requires the integration of authentication and trust propagation mechanics. For instance, ICMP echo replies and "time exceeded messages" commonly used for active geolocation lack any authentication capabilities. Adding such capabilities in order to increase the level of trust in active geolocation is challenging since it requires the revision of protocols.

This situation is alleviated slightly by trusted geolocation feeds from authoritative sources. However, since the granularity of geolocation feeds is insufficient for most use cases, geolocation services must use additional, less trusted sources of information. Thus, the challenge here is twofold: (1) identifying the geolocation data sources of geolocation data as the authoritative geolocation feeds; and (2) quantifying the degradation of trust when using opportunistic information sources.

Privacy and security. When it comes to privacy issues, a self-sovereign mechanism should be considered. Ideally, an innocent person should have ownership and control of location information, such that the user could present their geographical location to other parties or organizations at will. While some mechanisms are designed to control geolocation privacy preferences [31], such mechanisms are designed for application-level geolocation information exchange. Providing similar control for active probing or other location inference techniques is a major challenge.

The most common techniques employed by Internet users to conceal their geolocation include VPN and web proxies. More sophisticated and privacy-concerned users may employ onion-routing infrastructure such as the Tor network. Internet standards should respect the privacy of users reluctant to expose their real geolocation, but at the same time, any location reported by such users should not be trusted. Thus, the effective detection of privacy-enhancing technologies that conceal geolocation is an acute technological challenge faced by both commercial, e.g., video on demand, and cyber security, e.g., zero trust, applications. Furthermore, robust localization technologies should be resilient to all types of attacks against geolocation (see examples in Table 1).

Performance and infrastructure resources. There are many factors affecting the performance and robustness of geolocation. For example, active geolocation requires a wide network of probes deployed around the globe. Before the geolocation of a target host is obtained, the probes must ping it, report the results to a central server, and that server should infer the geolocation. The challenge of timely active geolocation limits its usefulness for use cases that require real-time response such as zero-trust authentication.

Further challenges investigated in the literature include the granularity and accuracy of geolocation. Some of these challenges are imposed by jitter and the natural route dispersion in the Internet. Increasing the location inference granularity to the level of offices or even buildings without trusting the reports of the client device in a robust and efficient way will be challenge despite existing Internet technologies.

Data geolocation. The geolocation of data maintained by cloud-based services is an emerging need and a significant challenge. The difficulty with data location inference arises from the architecture of the cloud service providers themselves (i.e., the separation between the front-end and data storage), multi-level cache servers distributed around the globe, and so forth. While schemes that allow a cloud service provider to prove the location of the customer's data have been proposed, they are rarely used in practice. The development of techniques that allow users to receive proof that their data is located or not located at a specific data center represents a major challenge.

5. Conclusion and Future Directions

In continuation from our previous white paper [1], which introduced the concept of Trust-Enhanced Networking (TEN), this white paper further elaborates on TEN and explores Robust Localization as an important TEN component underpinning the viability of many TEN use cases. From the improvement of services, such as targeted marketing, to the bolstering of security, such as the mitigation of phishing attacks, Internet geolocation has many important real-world applications, and advancements in this domain will thus have far-reaching implications.

From our review of the literature and existing technologies, we observe that despite the plethora of research in this space, few IP geolocation solutions are able to mitigate against subversion by an attacker; thus we conclude that adversarial mitigation of geolocation exercises remains an open research question. Using the insights of the literature review, we formulate a problem statement and threat model to chart our research direction in the area of Robust Localization, with a particular focus on adversarial mitigation and accuracy improvement. We also identified and discussed the technical challenges concerning geolocation in the case of web proxies and VPNs and the geolocation of data itself.

Our approach to exploring the TEN concept consists of two key components: firstly, the goals we would like to achieve, and secondly, the strategy we will use to achieve these goals.

Our overarching TEN goals are threefold: (1) to promote awareness of the need for TEN, (2) to drive the development of new technologies that address this need, and, (3) to facilitate access and adoption of the new technologies for the realization of a trust-enhanced networking paradigm.

Our efforts to encourage discussions about and promote awareness of the need for TEN have already begun through engagement with both academic and industry communities at events such as IETF gatherings, where we have engaged with multiple stakeholders and working groups about the issue. Going forwards, we will continue our engagement with the IETF and explore avenues of engagement (e.g., other relevant communities, bodies, and events)

To advance the development of new technologies, we plan to collaborate with partners in academia and industry, working closely with them to address the issues identified in the literature. We intend to publish the results and findings of our collaborations in peer-reviewed venues to further promote the benefits of TEN and encourage its adoption. Fujitsu and BGU have already begun collaborating in the area of Robust Localization, which is a crucial component of TEN.

To facilitate the adoption of TEN, we plan to produce open-source tools to help accelerate the creation, testing, and subsequent adoption of TEN technologies. We will also work with professional standards bodies to standardize the solutions developed and work with stakeholders to perform proofs of concepts and demonstrate the value and viability of TEN technologies.

With the emergence of metaverses and Web3, and as the activities of companies and society continue to shift to cyberspace, we believe the Trust Enhanced Networking concept will become increasingly important in facilitating trust between stakeholders in the evolving technological era. To realize the TEN paradigm, we will continue strategically, as outlined in this white paper, beginning with the area of Robust Localization.

References

[1] Fujitsu, whitepaper on "Trust-Enhanced Networking," https://www.fujitsu.com/global/about/research/article/202212-trust-enhanced-networking.html, 2022.

[2] Shichang Ding, Fan Zhao and Xiangyang Luo, "A street-level IP geolocation method based on delay-distance correlation and multilayered common routers," *Security and Communication Networks 2021*, 2021.

[3] Fan Zhao, Xiangyang Luo, Yong Gan, Shuodi Zu, Qingfeng Cheng and Fenlin Liu, "IP geolocation based on identification routers and local delay distribution similarity," *Concurrency and Computation: Practice and Experience*, 2019.

[4] Shuodi Zu, Xiangyang Luo and Fan Zhang, "IP-geolocater: a more reliable IP geolocation algorithm based on router error training," *Frontiers of Computer Science*, 2021.

[5] Qiang Li, Zhihao Wang, Dawei Tan, Jinke Song, Haining Wang, Limin Sun and Jiqiang Liu, "Geocam: An IP-based geolocation service through fine-grained and stable webcam landmarks," *IEEE/ACM Transactions on Networking*, 2021.

[6] Zhihao Wang, Hong Li, Qiang Li, Wei Li, Hongsong Zhu and Limin Sun, "Towards IP geolocation with intermediate routers based on topology discovery," *Cybersecurity*, 2019.

[7] Fan Zhao, Rui Xu, Ruixiang Li, Ma Zhu and Xiangyang Luo, "Street-level geolocation based on router multilevel partitioning," *IEEE Access*, 2019.

[8] Jing-ning Chen, Fen-lin Liu, Ya-feng Shi and Xiangyang Luo, "Towards IP location estimation using the nearest common router," *Journal of Internet Technology*, 2018.

[9] Guang Zhu, Xiangyang Luo, Fenlin Liu and Jingning Chen, "An algorithm of city-level landmark mining based on Internet forum," in *Proceedings of 2015 18th International Conference on Network-Based Information Systems, IEEE*, 2015.

[10] Fuxiang Yuan, Fenlin Liu, Donghua Huang, Yan Liu and Xiangyang Luo, "A high completeness PoP partition algorithm for IP geolocation," *IEEE Access*, 2019.

[11] Shuodi Zu, Xiangyang Luo, Siqi Liu, Yan Liu and Fenlin Liu, "City-level IP geolocation algorithm based on PoP network topology," *IEEE Access*, 2018.

[12] Siqi Liu, Fenlin Liu, Fan Zhao, Lixiang Chai and Xiangyang Luo, "IP city-level geolocation based on the pop-level network topology analysis," in *Proceedings of 2016 6th International Conference on Information Communication and Management (ICICM), IEEE*, 2016.

[13] Hao Jiang, Yaoqing Liu and Jeanna N Matthews, "IP geolocation estimation using neural networks with stable landmarks," in *Proceedings of 2016 IEEE Conference on Computer CommunicationsWorkshops (INFOCOM WKSHPS), IEEE*, 2016.

[14] Zhiyuan Wang, Fan Zhou, Wenxuan Zeng, Goce Trajcevski, Chunjing Xiao, Yong Wang and Kai Chen, "Connecting the hosts: Street-level IP geolocation with graph neural networks," in *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022.

[15] Fan Zhang, Fenlin Liu, Rui Xu, Xiangyang Luo, Shichang Ding and Hechan Tian, "Street-level IP geolocation algorithm based on landmarks clustering," *CMCCOMPUTERS MATERIALS & CONTINUA*, 2021.

[16] Ovidiu Dan, Vaibhav Parikh and Brian D Davison, "Distributed reverse DNS geolocation," in *Proceedings of 2018 IEEE International Conference on Big Data (Big Data), IEEE*, 2018.

[17] Quirin Scheitle, Oliver Gasser, Patrick Sattler and Georg Carle, "Hloc: Hintsbased geolocation leveraging multiple measurement frameworks," in *Proceedings of 2017 Network Traffic Measurement and Analysis Conference (TMA), IEEE*, 2017.

[18] Ovidiu Dan, Vaibhav Parikh and Brian D Davison, "Improving IP geolocation using query logs," in *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*, 2016.

[19] Lars Backstrom, Eric Sun and Cameron Marlow, "Find me if you can: improving geographical prediction with social and spatial proximity," in *Proceedings of the 19th international conference on World wide web*, 2010.

[20] Yang Zhang, Dongzheng Jia, Shijie Jia, Limin Liu and Jingqiang Lin, "Splitter: an efficient scheme to determine the geolocation of cloud data publicly," in *Proceedings of 2020 29th International Conference on Computer Communications and Networks (ICCCN), IEEE*, 2020.

[21] Dongzheng Jia, Yang Zhang, Shijie Jia, Limin Liu and Jingqiang Lin, "Dpvgeo: Delay-based public verification of cloud data geolocation," in *Proceedings of 2020 IEEE Symposium on Computers and Communications (ISCC), IEEE*, 2020.

[22] Dong Lai Fu, Xin Guang Peng and Yu Li Yang, "Trusted validation for geolocation of cloud data," *The Computer Journal*, 2015.

[23] Mark Gondree and Zachary NJ Peterson, "Geolocation of data in the cloud," in *Proceedings of the third ACM conference on Data and application security and privacy*, 2013.

[24] Elisa Chiapponi, Marc Dacier, Olivier Thonnard, Mohamed Fangar and Vincent Rigal, "Badpass: Bots taking advantage of proxy as a service," in *Proceedings of Information Security Practice and Experience: 17th International Conference, ISPEC 2022, Springer*, 2022.

[25] Zachary Weinberg, Shinyoung Cho, Nicolas Christin, Vyas Sekar and Phillipa Gill, "How to catch when proxies lie: Verifying the physical locations of network proxies with active geolocation," in *Proceedings of the Internet Measurement Conference 2018*, 2018.

[26] Katharina Kohls and Claudia Diaz, "VerLoc: Verifiable localization in decentralized systems," in *Proceedings of 31st USENIX Security Symposium (USENIX Security 22)*, 2022.

[27] Abdelrahman Abdou, Ashraf Matrawy and Paul Oorschot, "Accurate manipulation of delay-based Internet geolocation," in *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, 2017.

[28] Sharon Goldberg, "Why is it taking so long to secure internet routing? Routing security incidents can still slip past deployed security defenses," *Queue*, 2014.

[29] Phillipa Gill, Yashar Ganjali and Bernard Wong, "Dude, where's that IP? circumventing measurement-based IP geolocation," in *Proceedings of 19th USENIX Security Symposium (USENIX Security 10)*, 2010.

[30] Roger Dingledine, Nick Mathewson and Paul Syverson," Tor: The second-generation onion router," in *Proceedings of 13th USENIX Security Symposium (USENIX Security 04)*, 2004.

[31] H Schulzrinne, H Tschofenig, J Cuellar, J Polk, J Morris and M Thomson, "Geolocation policy: A document format for expressing privacy preferences for location information," *Technical report, RFC 6772 (IETF Stadrads Track)*, 2013.

[32] E Kline, K Duleba, Z Szamonek, S Moser and W Kumari, “A format for selfpublished IP geolocation feeds,” *Technical report, RFC 8805 (Informational)*, 2020.

For further information, please contact:

contact-nwt@cs.jp.fujitsu.com
Data & Security Research Laboratory
Fujitsu Limited