The Emergence of Public Key Infrastructures: Changing Authority Structures and the Automation of Trust in Cyberspace

Dr Milton Mueller, Georgia Institute of Technology, Jimmy and Rosalyn Carter School of Public Policy, Atlanta, GA

Milton@gatech.edu

Abstract

A public key infrastructure (PKI) employs asymmetric cryptography to protect the security of online communications. To implement cybersecurity at Internet scale, the operators of a PKI rely on Certification Authorities (CAs) to sign digital certificates that can be used to authenticate online identities. A PKI is, therefore, a collective authority structure intended to provide secure communications and data integrity. This paper investigates the divergence of PKI authority structures from the hierarchical structure of nation-state sovereignty. It shows empirically that despite early efforts by states to bring encryption technology into conformity with existing governmental power hierarchies, private actors supplanted state actors in the formation of cryptographic trust arrangements, and privatized and globalized infrastructures supplanted territorial ones. Consequently, the authority structure that evolved to meet the cybersecurity needs of the online world diverged from the state-sovereignty structure of the offline world. Understanding how and why this divergence happened contributes to our understanding of how digitization has altered the global information security order and illuminates current debates over "digital sovereignty" (or the lack thereof).

This research was funded by the Internet Society Foundation.

Introduction: PKI as authority structure

A multidisciplinary understanding of digital technology, political economy and international relations renders visible a phenomenon that social scientists would otherwise likely miss: the absence of an isomorphism between the authority structure of sovereign states, and the authority structures set up in cyberspace around the automation of trust and security. That divergence forms the subject of this paper. It tracks the evolution of Public Key Infrastructures (PKIs) to document how transnational institutions formed by business, engineers and civil society supplanted state actors in the formation of global trust hierarchies in cyberspace.

The definition of a PKI and its technical components are explained in greater detail in section 1. A critical aspect of that explanation, however, must be set out here. The ability of a PKI to automate trust among a large number of distributed actors in cyberspace requires authoritative trust anchors, also known as a *root Certification Authority (CA)*, or a *root store* (when the anchor includes multiple CAs). A root is the apex in a hierarchy of digital certificates, where one trusted entity, or a small number of trusted entities, produce security by cryptographically attesting to the identity of the networked entities below them in the hierarchy.

If PKIs followed the sovereign order, territorial states or a delegated agent would sit at the top of a national trust hierarchy. State-operated or state-approved Certification Authority (CA) that would serve as the supreme and exclusive trust root for a given country's digital networks, just as the sovereign occupies the position of the supreme and exclusive political authority in a country. In a state-centric model, a sovereign nation would achieve international interoperability and trust through bilateral negotiations with other sovereigns or use multilateral institutions to negotiate global standards and agreements among states.

This institutional isomorphism, as the paper will show, is what many law and cybersecurity experts expected to happen at the beginning of PKI development in the early 1990s. And the expectation was not unreasonable, because this is what *did* happen to telephone and telegraph networks in the 19th and early 20th centuries. The first electronic communication networks (telegraph and telephone) were monopolized by nation-states in most countries, following the precedent set by national postal monopolies. Postal monopolies themselves were attempts to maintain state sovereignty and security by controlling and surveilling internal and external communications. (Firth, 1898; Mueller, 1986; Davids, 1995) Diffie and Landau (2007), for example, describe how the British and

French governments maintained centralized offices for the monitoring of postal mail in the 19th century. From 1865 to 1990, the International Telecommunication Union (ITU), the world's first formal intergovernmental treaty organization, became the institutional nexus for interconnection, revenue-sharing and security arrangements among state-owned Post, Telegraph and Telephone authorities (PTTs). (Kouli & Laborie, 2023) The global authority structure in telecommunications, in short, was institutionalized in a way that mirrored the sovereign state order.

That is not what happened in cyberspace. The growth of the Internet and the World Wide Web after 1990, coupled with the liberalization of trade in telecommunication services, created a transnational space for commerce and communication. There was a pressing need to secure global computer networks, and this required the use of automated tools like digital certificates organized into PKIs. The emergent security order on the Internet, however, did not mirror the sovereign structure. Early efforts by governments to regulate certification authorities went nowhere; further, governmental efforts to assert control over encryption keys were resisted and met limited success. Instead, strong encryption technology diffused into the private sector, and the Internet industry set up its own PKIs with its own authority structures. States participated in the system - there are, as we will see, CAs run by government agencies – but they are not its rulers, they are merely participants in a largely denationalized global order established through contracting, the market, and non-governmental technical standards and governance entities.

The main goal of this paper is simply to demonstrate empirically the lack of isomorphism between the sovereign order and global PKI trust hierarchies. It explains this deviation by tracking the historical evolution of the Internet, the Web, and cryptographic security methods.

The paper is organized as follows. Section 1 provides a more detailed explanation of PKIs and how they work, so that readers can better understand the nature of trust and authority relations in cyberspace and how they are translated into automated technical infrastructures.

Section 2 provides a three-pronged description of initial governmental efforts to standardize and regulate PKIs and the issuance of digital certificates. It examines the role of the ITU in creating the X.509 certificate standard, and efforts by governments and UN agencies to set legal frameworks for Certification Authorities. It also documents the linkage between governmental efforts to regulate digital certificates and the attempts by states to gain control of users' private keys for law enforcement and surveillance purposes.

Section 3 tracks the rise of commercial businesses around the World Wide Web after 1993. It describes their incentives to apply PKI to the security needs of a commercializing Internet. In this period, we see the rise of a commercial Web browser industry, the rise of a private CA industry, and the movement of standardization activity from the ITU, a multilateral treaty organization, to the Internet Engineering Task Force (IETF), a non-governmental, unincorporated standards development organization. This phase culminates with the implementation of three large-scale open PKIs around private actors: a PKI for the World Wide Web, a PKI for the domain name system (DNSSEC), and a PKI for Internet routing (RPKI).

Section 4 provides a structural analysis of all three digital certificate trust hierarchies (Web PKI, DNS PKI, and RPKI). Drawing on Internet measurement studies, it shows that there is no structural isomorphism between cyberspace's PKI trust hierarchies and the sovereign authority structure, even in countries as determined to assert digital sovereignty as China.

Section 5 elaborates on the idea of authority structures and how they are reflected in technical infrastructures. It attributes the deviation from sovereign state structures to the way key structural features of the Internet protocols interacted with the liberalization of telecommunications and the globalization of e-commerce.

1. PKIs as applied cryptography

PKIs are enabled by asymmetric key cryptography, also known as split-key cryptography. (Diffie and Hellman, 1976) Asymmetric cryptography is a technical-mathematical innovation from the 1970s that came slightly before, and independently of, the Internet protocols. It splits the key used to encrypt communications into two parts: a public key and a private key. The public key can be accessed and used by anyone to *encrypt* messages sent to the public key holder. The message can only be *decrypted* by the recipient's private key. While the two keys are mathematically related, it is computationally impossible to derive the private key from knowing the public key.[1]

Split-key cryptography made it unnecessary for users to transmit encryption keys over the Internet, where they could be intercepted or copied. No matter how remote or impersonal the connection between two parties, online communicators could share a public key over any channel and still achieve secure, encrypted communications by keeping their private key secret. The possession of a unique private key also made it possible for their holders to sign digital objects, such as certificates, documents, or software. Digital signatures can be used by the receiver of a digital object to verify that it came from the right person and that

[1] "Impossible" given existing technology. Quantum computing could reverse this constraint in the future.

the data has not been modified or corrupted in transit. “Public key cryptography combined with one-way hash functions gave rise to documents with digital signatures that can withstand repudiation. By definition, verifying a digital signature automatically proves the authenticity of the signer." (Benantar, 2001, 652).

Splitting the key creates its own security vulnerabilities, however. A public key, being public, can be copied and used by a bad actor to impersonate the actual possessor of the private key in a man-in-the-middle attack. (Diffie and Hellman, 1976) Without some form of trusted verification of the binding between an entity’s name and its public key, the parties communicating cannot be sure that the person announcing a public key is who they say they are.

Digital certificates evolved as the solution to this identity authentication problem. First proposed by Kohnfelder (1978), it took more than a decade for standards and practical methods of implementation to evolve (Gutmann, 2002). The basic result was as follows. Before a user’s public key is disseminated to the public, a high-assurance Certification Authority (CA) uses its own private key to digitally sign a certificate that attests to the link between the user's identity and its public key. The CA is supposed to verify that the subject of the certificate is the real holder of the public key before issuing the certificate. A relying party securely installs the public key of the trusted CA and uses it to verify its signature on each user’s public key certificate. Only after a successful verification of the signature does a relying party initiate a communications channel. (Benantar, 2001; Adams and Lloyd, 2002).

Thus, a PKI is a shared infrastructure for distributing public keys to a large group of users in a trusted manner. A PKI can be implemented within a single, closed organization, or in large-scale, open environments like the World Wide Web. Our focus in this paper is on the open, public implementations, because they involve a form of public governance of an infrastructure that one might associate with states.

A PKI requires participants to place their trust in an authoritative entity, a root-level CA. There are important parallels between the role of a root-level CA and that of the sovereign in political science. A root CA sits at the top of an authority hierarchy. It needs no third party to vouch for its authenticity: it signs its own certificate. By virtue of holding that power, the root entity is also able to compromise the security of everyone else lower in the chain. The root CA can delegate the power to issue certificates to other entities, but those intermediaries are trusted only because they are certified by the root CA. The relying party’s validation software must trace the chain of trust from the intermediate CA back to the root CA. Just as the sovereign has absolute control over the juridical order itself and has the power to decide who is included and who is excluded from this order (Philpott,

2003; Schmitt, 1922), the PKI's root authority has control over which certificates are trusted, and which are not.

A PKI thus requires from its participants collective recognition of certain procedures and decision makers to govern what can be trusted and what cannot. It is, to repeat, an authority structure, a technical manifestation of collective security governance.

In discussing "trust" and "authority" in this context, it is important to remember that a PKI is an attempt to *automate* both. It embeds trust and authority relations into the programmed operation of an information system. While the selection of a root CA and its policies and practices will be manual and reliant on many types of organizational and physical security, once certificates are signed and issued, the authentication process is supposed to be automatic; based on the execution of a program, not the feelings of trust or the personally known reputations of the communicating parties. Automation is necessary because of the large scale of the open global Internet and the heterogeneity of its endpoints and programs (servers, data centers, smartphones, sensors, agents). Trust and security arrangements based on personal knowledge of one's communication partners do not scale to the open online environment and may not even scale in a single large organization. An automatic means of authenticating identity and verifying trust must be built into the information infrastructure.

2. The path not taken: governments and PKI

States and sovereign authority structures were not absent from the initial development of PKIs. Far from it. The ITU, the quintessential multilateral institution in communications, played an important role in standardizing digital certificates. Moreover, the historical record contains several statements by early PKI developers expressing the expectation or possibility that root keys would belong to government agencies. State and national governments proposed and sometimes passed laws intended to govern digital certificates, while multilateral institutions disseminated model laws worldwide. States, especially the powerful United States of America, made concerted efforts to control cryptographic keys and algorithms. States were unable, however, to control the emergence of a digital certificate-issuing industry, nor were they able to position themselves at the top of the resultant trust hierarchies.

This section tracks this early phase of PKI development. Part 2.1 discusses the ITU's digital certificate standard; 2.2 tracks early efforts by the Internet technical community to develop a PKI for email; 2.3 reviews State- and national-level legislative efforts and the diffusion of model legislation by international intergovernmental organizations; 2.4 covers

key escrow programs as an attempt to put states at the center of applications of asymmetric cryptography.

2.1 The ITU and X.509

The most widely used standard defining the format of public key certificates is X.509.[2] This standard originated in the X.500 series developed in the 1980s in an attempt by multilateral institutions to create global computer networking standards.[3] The work was carried out by the Telecommunication Standardization Sector of the International Telecommunication Union (ITU-T), the International Standardization Organization (ISO), and the International Electrotechnical Commission (IEC). Version 1 of X.509 was first approved in 1988; version 2 came in 1993.

With its origins in monopoly telcos and ITU, the X.509 standard's conceptual model for certificate trust and authority closely mirrored the nation-state hierarchy. The standard attempted to incorporate the named subject of a digital certificate (i.e., the holder of the public key) into a global directory administered by national telcos.

Figure 1 (from Gutmann, 2002) shows how the naming structure for certificates was based on a hierarchical database model, with the path through the directory being defined by a series of relative distinguished name (RDN) components that together form a distinguished name (DN). National Certification Authorities would sit at the top of the hierarchy and control access to the directories. Ultimately, X.509 succeeded as a general certificate format, but its assumptions about state-centric certificate authority hierarchies did not survive the mid-1990s.

Figure 1: X.500 Directory and Certificate Naming Model (Gutmann, 2002)

[2] Kohnfelder's 1978 paper, the first documented proposal of what we now call digital certificates, used a structure conceptually similar to what became the X.509 standard a decade later. He proposed the use of certification authorities (CAs) to issue signed statements (i.e., certificates) that associate a public key with a user identity and suggested using a trusted root authority to manage and distribute public keys securely.
[3] That effort was ultimately upstaged by the Internet protocols. See Drake (1993) for a discussion of the standards competition between OSI and the Internet protocols.

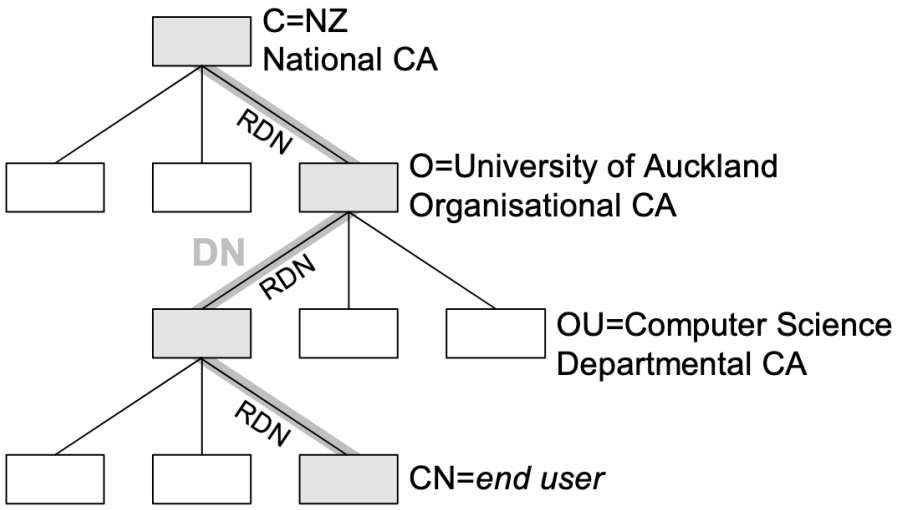

2.2 The Internet and X.509

As the ITU standardized digital certificate format, the internet Engineering Task Force (IETF) was making its first foray into PKI standards: an attempt to protect the privacy of email. In the late 1980s, email was emerging as the “killer app” driving Internet adoption by a wider set of users. Early attempts to create Privacy Enhanced Mail (PEM) (RFCs 1114, 1422) never succeeded in securing that perpetually insecure application, but they did instigate the first negotiations over trust hierarchies and certificate naming standards.

RFC 1114 (Kent and Linn, 1989) proposed a centralized, vendor-controlled trust hierarchy in which RSA Data Security, the owner of the patent protecting the specific mathematical method of creating a public/private key pair,[4] would act as the "co-issuer" for organizational certificates.[5] The trust root for the U.S. was thus conceived to be in the hands of the private sector intellectual property owner - with one notable exception: the U.S. government.

Because the RSA algorithm was based on research sponsored by the U.S. National Science Foundation, the U.S. government retained royalty-free license rights and could

[4] U.S. Patent 4,405,829, commonly known as the RSA patent. The patent was awarded to three inventors: Ronald Rivest, Adi Shamir, and Leonard Adleman. The acronym "RSA" comes from the first letter of each of their last names. The patent was assigned to the Massachusetts Institute of Technology (MIT), where they were working as computer scientists when they developed the algorithm. Their application for the patent was made in 1977, only a year after the publication of Diffie-Hellman's breakthrough work on asymmetric cryptography.

[5] This arrangement was motivated by the need to manage patent licensing for the RSA algorithm and simplify accounting (e.g., a $25 biennial licensing fee for each organization). Organizations essentially shared the CA role with RSADSI, which held the private keys used to sign certificates on their behalf.

establish its own independent certificate generation facilities for national government agencies. A federal government contractor, Trusted Information Systems, developed a closed operational version of privacy-enhanced mail for government agencies, and an open version for the Internet community that would support a “national certificate authority” allowing other users to correspond by email securely. The authority structure projected by RFC 1114 thus seemed predominantly national. It acknowledged that users outside the USA, for example, would not require a license from RSA, and assumed that other countries or government entities would likely establish their own independent CAs, with “procedures for interacting” with them defined as outside the scope of RFC 1114.

2.3 PKI Laws and regulations

The state-centric hierarchy was present not only in X.509 but set the expectations for some of the early work on PKI public policy. Efforts to develop legal rules to support public key infrastructure began in 1992 and by August 1996 had culminated in the Digital Signature Guidelines issued by the Information Security Committee of the American Bar Association's Science and Technology section. (Biddle, 1997; Merrill, 1998) In 1994 prominent U.S. cybersecurity expert Santosh Chokhani, writing in *IEEE Communications*, proposed a “national public key infrastructure” for the United States. It consisted of a “four-layer certification authority for the nation” with a "top-level national node.” (Chokhani, 1994, 72) Froomkin (1996, p 56) noted that a common solution to the problem of the need for a trusted third party to regulate the verification of certificates was:

> …a governmental role in certifying the keys of CAs. The root key would belong to a state or federal agency, and the few CAs that met state licensing requirements would be rewarded with government certification of their root key.

A law passed by the state of Utah in March 1995 gave a government agency the responsibility of being a Root CA. This governmental entity would be charged with policymaking, facilitating implementation of digital signature technology, and providing regulatory oversight of private sector CAs through a comprehensive licensing scheme. (Biddle, 1997, 1233) Digital signature legislation based on the Utah law was proposed in nearly a dozen U.S. states, and passed in Washington, Florida, and California.

The notion spread internationally. The government of Malaysia enacted legislation based upon the Utah Act; similar legislation was considered in Australia, Canada, Germany, Singapore, and the European Union. The cross-border nature of Internet commerce and the appropriate legal regime required was the subject of continuing meetings of the United Nations Commission on International Trade Law's (UNCITRAL) Working Group on Electronic Commerce. (Corwin, 1998)

2.4 Key escrow and PKI regulation

Writing in 1996, it was clear to Biddle (p. 1227) that one reason for the momentum of digital certificate legislation at national and international levels was "its synergy with cryptographic 'key escrow' proposals." For the previous two decades, the U.S. government's national security and law enforcement agencies had fought against the ongoing civilianization, privatization and commercialization of cryptography due to its negative impact on their surveillance capabilities. Their methods, as Banisar (1999) and Diffie and Landau (2001) documented, included attempts to control and suppress the publication of cryptographic science, the subversion of encryption standards, and export controls. As the rise of the Internet intensified demand for encryption and made dissemination of software much harder to control, Federal policies proposed to give government agencies the power to hold private keys "in escrow" and access them when needed. (Froomkin, 1996a) The Clinton administration first proposed its hardware-based key escrow plan – known as the Clipper Chip – in April 1993, part of more than a decade of political contestation between the government and privacy/cybersecurity advocates in business and civil society. (Jarvis and Martin, 2024, Landau, Kent et al, 1994)

In May 1996 the Executive Office of the U.S. President's Office of Management and Budget produced a draft white paper linking CA regulation to its key escrow plan. By setting standards and regulations for what it called "the Key Management Infrastructure (KMI)," the state could "assure timely, lawful, government decryption access." (Dessy, 1997) Other governments had similar intentions. Under a proposed law in the U.K., government-regulated CAs would not only authenticate public key holders but also verify the escrowing of keys. (Biddle, 1996, 1234) Digital signature legislation, in other words, had the potential to create an infrastructure for implementing a key escrow scheme. Rules granting states access to private keys folded into nation-centered PKIs would be the ultimate step in subordinating digital security technology to the sovereign state order.

3. The Web and the rise of global PKIs

This vision of a state-centered authority structure deteriorated very rapidly after 1995. Instead, the commercial promise of a global Internet/Web led to the emergence of Internet browser software, private CAs, revised certificate standards, and global PKIs anchored by non-state actors.

3.1 The Internet Society as "sole, universal trust anchor"

From 1989 to 1993, thinking about a PKI for email moved away from the national model to a more globalized, nonstate actor-based trust hierarchy. This change reflected the growing

maturity and independence of the Internet technical community.[6] RFC 1422 (Kent, 1993) proposed to move the root of the Internet email certification hierarchy to an Internet Policy Registration Authority (IPRA). IPRA would be "the sole, universal trust anchor" for email and set "global policies which apply to all certifications effected under this hierarchy." The IPRA would be operated by the newly incorporated Internet Society,[7] a move that anticipated the Society's later attempt to set itself up as the Internet Assigned Numbers Authority (IANA), the root of the domain name system and policy authority for Internet identifiers.

A single, global root for an email PKI was never implemented, and the approach in RFC 1422 was superseded by the more general (not email-specific) PKI model articulated in RFC 2459 (1999), which is discussed in section 3.5 below. Nevertheless, in RFC 1422 and the discussions around it (Kent, 1993a) one sees a clear departure from a nation-state oriented authority structure.

3.2 Netscape and SSL

Widespread implementation of the World Wide Web protocol after 1990 led to the development of a new software interface for computer networking: the browser. Web browsers made navigating the Internet easy for non-expert members of the public by hyperlinking documents, files and services, making internetworking a true mass medium. (Berners-Lee, 1999; Gillies & Cailliau, 2000)

The first Web browser, Mosaic, was launched by university researchers in 1993. As the economic potential of the new medium became apparent, Netscape Communications commercialized Mosaic in 1994 and moved quickly to realize its potential for commerce over the Web. No one would buy things online if their credit card numbers or bank account data were exchanged in the clear; the channel had to be encrypted. In 1995, Netscape pioneered an encrypted transport protocol, Secure Sockets Layer (SSL), and built it into the browser software.[8] The significance of SSL went beyond the browser itself, however. The lure of Web-enabled commerce converged several categories of actors around the construction of a transnational PKI. Before transport layer encryption would work, e-commerce website operators had to acquire digital certificates for their servers to disseminate their public keys in a trusted manner. This in turn created a market for

[6] See Mueller 2025 for a more complete discussion of the IETF/Internet Society's assertions of autonomy.
[7] ISOC was incorporated on December 10, 1992, as a 501(c)(3) non-profit under the laws of the District of Columbia (Washington, D.C.).
[8] To comply with U.S. export controls (ITAR, 1993) there was a domestic and international version of the browser. The domestic version had 128-bit keys and the international version 40-bit keys. 40-bit keys could at that time be broken with brute force in about 2 days. (Shepherd 1996)

commercial Certification Authorities to issue and sign digital certificates for a fee. The browser software in turn would have to decide which CAs could be trusted and embed them in their software as trust anchors.

The eruption of the World Wide Web created a demand for CAs that state-based institutions, with their slow, fragmented legislative and policy making processes, and their disconnection from economic incentives and operational implementation, were simply not prepared to meet. As Merrill (1998) put it, "The ink was barely dry on the [ABA's] Digital Signature Guidelines [of August 1996] when it became apparent that the deployment of Certification Authorities (CAs) and PKI was proceeding more rapidly than the likely pace of any concerted legislative or regulatory efforts designed to rationalize this fledgling industry." The role of government was also limited by intense pushback against government control of encryption keys generated by the social movement against the Clipper Chip. (Froomkin, 1996a)

3.3 The Trouble with X.509

In their rush to innovate and implement the emerging security infrastructure, the Internet technical community and the Web industry seized upon the ITU's X.509 standard as a readily available format for digital certificates. Their embrace of the ITU's X.509 standard, however, did not replicate its state-centric authority structure.

The ITU had standardized the *format* of certificates – the required data elements and their sequence – but, as one astute analyst of different certification systems wrote at the time (Gerck, 1997), the ITU standard did not establish a naming authority for the assignment of globally unique subject names on a digital certificate.[9] Indeed, by 1996, the PTT monopoly – originally conceived as the gatekeeper for naming users and the issuance of national digital certificates - was becoming a relic of the past thanks to privatization, competition and liberalization in the sector. (Cowhey, 1990; Drake, 2005) Without a nation-state monopoly on certificate issuance, the whole X.509 "Distinguished Name" hierarchy became denationalized, un-standardized and market driven. Neither national governments nor multilateral institutions controlled who would issue certificates, the names assigned to CAs, and the names that CAs assigned to certificate subscribers. Without a national gatekeeper, the X.509 standard could not set up any standardized way to record user

[9] "The same user can have different DNs in different CAs... so different DNs for different CAs do not necessarily mean different users and vice-versa. Further, a DN does not have to contain the user's real-world name or location." (Gerck, 1997)

identity.[10] Detached from any central authority, each CA became a naming authority only for its own certificates and established its own validation procedures.

RFC 1422, the IETF's proposed standard for an email PKI, had already recognized the tensions between the namespaces articulated by the state-centric X.509 standard and the newly emerging Internet namespaces. The author of RFC 1422, Stephen Kent, compared the X.509 standard's RDN approach and the Internet Domain Naming System (DNS), and opted for using the ITU's Distinguished Names (DNs)! He noted that RDNs and DNS were structurally parallel at the top—both used hierarchical, delegated name components, although countries were at the top of the ITU hierarchy while the Internet name space had both country codes and organizational categories (com, net, org, edu) at the top and gave end users the right to freely select where they registered; consequently, most domain name registrations were in generic categories. The key differences were in expressiveness and deployment. X.509 RDNs were attribute-value pairs carrying explicit semantic types (country, state, organization, common name), while DNS labels were caseless ASCII strings typically only 3–6 characters long at the time. X.509 encouraged full descriptive names; DNS encouraged extreme conciseness. Kent argued that shorter domain names increased the risk of name collision and user confusion. However, in terms of deployment at the time, the difference was stark. The DNS had approximately 1.5 million registered hosts in 1993, and thus was serving millions of email users, while X.500 pilot projects had barely over one million registered directory entries and very little client software deployment. But if it was to serve as a digital certificate infrastructure, DNS had shortcomings too. It would require new record types and per-user records.[11] This pragmatic acknowledgment of the infrastructure gap foreshadowed the tensions that the next step in the IETF's PKI work would take in RFC 2459.

3.4 Rise of a Commercial CA industry

In April 1995, as Netscape was distributing browsers with encryption capabilities (SSL 2.0) to millions of Web users, a new company called VeriSign was incorporated in Delaware. Verisign was a spin-off of RSA Data Security, the holder of one of the first patents on a public key encryption algorithm.[12] RSA had been profitably licensing its algorithm to businesses but realized that issuing and managing large numbers of digital certificates

[10] "...regarding validation procedures for the user's identity, Section 11.2.a [of the X.509 standard] states that: "a certification authority shall be satisfied of the identity of a user before creating a certificate for it", which means that identity validation procedures are to be satisfied in the CA's frame of reference by following the CA's own self-defined rules (called CPS Certification Practice Statement)..” (Gerck, 1997).

[11] In fact, Kent proposed that existing DNS and WHOIS databases could store certificates containing true DNs as an interim measure while X.500 matured.

[12]

required a different business model. Spinning off VeriSign allowed RSA to pursue this new market. VeriSign would act as a CA and provide "trust for the Internet and Electronic Commerce through Digital Authentication services and products," according to its mission statement.

VeriSign and other independent commercial CAs had discovered they simply did not need any special legislation to enter the market for issuing certificates. They could allocate risk contractually by using click-through agreements with both subscribers and relying parties. (Biddle, 1997, p. 1239) In early 1996, several other commercial CA services were announced, including GTE, Nortel, and MCI Mail. IBM began to issue certificates for its proprietary software and enterprise solutions to enable businesses to create their own internal PKIs.

By the end of 2000, VeriSign had sold over 485,000 website digital certificates.[13] It managed 1,600 security service solutions for enterprises and had 35 affiliates in its international network. By 2006, VeriSign managed over 3 million SSL certificates. The company sold its certification and authentication business to Symantec in August 2010 for $1.28 billion.

3.5 The IETF PKI X.509 Working Group

In the autumn of 1995, the IETF established the Public-Key Infrastructure X.509 working group (PKIX) to facilitate the use of X.509 certificates within Internet applications. The new standardization effort was begun in response to concerns that, in the newly emerging Web, companies would implement X.509 certificates in ways that would undermine global compatibility. RFC 2459 (Housley, Ford, et al, 1999) developed the X.509 v3 profile, evolving the PKI architecture into a flexible model supporting all entities and applications. It removed structural distinctions between the types of organizations running CAs, disposed of the ITU's RDN concept of naming unique objects, and let trust be configured locally. This approach leveraged several actors' incentives: RSA's desire to distribute its encryption software widely (see 3.5 below); commercial CAs' desire to enter a market by building infrastructure and producing certificates; and organizations' and users' demand for improved, automated security in Web-based applications like browsers, as well as OS and resolver software, which would be instrumental in the success and adoption of Web PKI.

In this architecture, government agencies, private enterprises, and commercial CA vendors operated on equal footing. Policy Object Identifiers and cross-certification agreements allowed both government and private organizations to build independent but interoperable trust networks tailored to their specific needs. It also decentralized trust,

[13] SEC filings for Verisign, 2000 – 2002.

allowing users and organizations (e.g., software providers like Netscape) to determine and aggregate trusted root stores used in applications to validate certificate paths. In 2000, RSA's patent expired, releasing into the public domain the algorithm which served as the primary building block for Secure Sockets Layer (SSL), most email encryption, digital certificates, and virtual private network (VPN) software.

3.6 The Web PKI trust model

Netscape Navigator 2, released in September 1995, included in its software a set of pre-loaded root CAs.[14] By pre-loading root certificates into the browser, Netscape established the trust model that is still the basis for modern Web PKI. The true root authorities are the *browser vendors*, and their decision to trust or not trust specific CAs defines the apex of the Web PKI trust hierarchy.

There were significant weaknesses in the early versions of this rapidly evolved, loosely organized authority structure. As Gutmann (2002) pointed out, the early Web PKI "employ[ed] a form of implicit cross-certification in which all root CAs are equally trusted…" Gutmann wrote that "Many of these CAs are completely unknown, follow dubious practices such as using 512-bit root keys or keys with 40-year lifetimes, appear moribund, or have had their CA keys on-sold to various third parties when the original owners went out of business."[15]

Those chickens came home to roost in July 2011, when Dutch CA Diginotar was compromised by an Iranian hacker, leading to what some called a "security collapse" of Web PKI. (Arnbak, Asghari et al, 2014) Google, Microsoft and Apple were forced to revoke their trust of Diginotar's certificates and remove them from their browsers' root stores.

This crisis led to many calls for governmental intervention, but the industry acted faster. Alerted to the security threat to their own services, the major software vendors and CAs tightened security controls on Web PKI. An industry self-regulatory institution, the CA/Browser Forum, issued a revised set of Baseline Requirements for digital certificates and began to aggressively enforce new standards using the leverage of inclusion in their root stores. The Certificate Transparency (CT) program was introduced by Google in 2013 as a public framework to log all issued Web PKI certificates, allowing the industry to detect the issuance of any unauthorized certificates. (Grindal, Mueller, and Srivastava, 2025) The browsers demonstrated their willingness to kick CAs out of their root stores if they did not

[14] The root certificates included CommerceNet, AT&T, RSA Commercial, RSA Secure Server and MCI Mail.
[15] As the Internet grew, criticism of the inadequacies of the digital certificate system mounted. (Roosa & Schultz, 2010; Vratonjic, Freudigr, et al, 2011)

comply, at one point discarding the leading certificate issuing company in the world, Symantec. (Ma, 2021)

4. The Domain Name and Routing PKIs

Two other global PKIs for the Internet emerged after the turn of the millennium. One new PKI formed around the domain name system (DNS). Later, another was deployed to improve the security of Internet routing. In both cases, implementation and governance relied on nonstate actors, and neither system anchored trust in state-sovereign authority.

The U.S. government, however, did play an important role in supporting research, development and standardization of the domain name and routing PKIs. The U.S. did not, however, impose these standards on the industry, nor did it propose to govern their operation. It conceived of its role as a funder of applied research and as a promotor of more secure standards and architectures for the Internet. To influence cybersecurity under decentralized and privatized conditions, Kuerbis (2011) describes the actions of US government agencies as a form of delegation; it influenced the social networks developing Internet security standards by funding researchers and contracting with developers.

4.1 Domain Name Security Extensions (DNSSEC)

Domain names are alphanumeric identifiers for websites, email addresses, and host computers on the Internet. The DNS standards (RFC 1034 and 1035) define a hierarchical name space with a single root followed by top-level, second-level, third-level, and 50+ name spaces below it. (Figure 7) This naming architecture provides a highly scalable way of assigning globally unique names to computers. It breaks down the coordination of unique name assignments into many smaller, more localized levels, allowing operators of networked computers assign names to their nodes that will not duplicate names assigned to any other computer on the Internet.

However, to function as “addresses;” domain names must be mapped to Internet Protocol (IP) addresses, a process known as “resolution.” Resolving domain names depends on querying a network’s name server, which responds with records that returns the IP address the domain lives at.

This domain name resolution process was one source of security vulnerability. A DNS resolver had no way to verify the authenticity and integrity of the data sent to it by name servers. A cyber-attacker could exploit this vulnerability to misdirect traffic to incorrect or harmful domains.

In his analysis of the early history of the DNSSEC standard, Kuerbis (2011) shows that the Internet technical community was aware of these problems as early as 1990. Discussions

in the IETF about a more secure standard began in 1993, and in 1995, DARPA issued a Broad Area Announcement (#95-15) to "promote redesign of network protocols to remove known security weaknesses." From 1994 to 2005, the IETF revised its DNS protocol to address some of these problems. The new standard was called DNS Security Extensions (DNSSEC).[16]

DNSSEC provides the following contributions to cybersecurity:

- Source authentication: a resolver can cryptographically verify that a response originated from zone's authoritative name server;
- Integrity verification: a resolver can cryptographically verify that the data transferred has not been tampered with in transit
- Authenticated denial of existence: a resolver can verify that a queried domain name does not exist on the authoritative name server.

The public key distributions, digital signatures and hashes involved in a DNSSEC transaction follow the domain name delegation hierarchy. The chain of trust is a strict hierarchy that converges on a single authoritative root zone. The DNS root is administered by the Internet Corporation for Assigned Names and Numbers (ICANN). ICANN first signed the DNS root zone in 2010. In 2014, it began to require all new generic top-level domains it licensed to implement DNSSEC. Currently, 1,523 of the 1,591 generic top-level domains in ICANN's root zone (96%) have implemented DNSSEC. Implementation at the lower levels of the hierarchy is much slower, due to the cost and complexity of signing and verifying DNS zone files. Of the 209.4 million second-level domains analyzed by cybersecurity firm Whisper in 2025, only 9.8 million (4.7%) were cryptographically secured via DNSSEC, though more commonly used domains have a higher rate (12%).[17]

4.2 Routing and addressing: The RPKI

Internet routing is the method by which data packets move from their origin network to their destination network. The process is controlled by a protocol known as Border Gateway Protocol (BGP). In BGP, network operator routers transmit "announcements" of IP address prefixes that tell other networks how to reach it and what other networks it is connected to. The data in these announcements contain two critical forms of numerical

[16] RFC 4033: DNS Security Introduction and Requirements; RFC 4034: Resource Records for the DNS Security Extensions; RFC 4035: Protocol Modifications for the DNS Security Extensions. March 2005. There was an earlier, unsuccessful attempt to standardize DNSSEC from 1997 to 1999.

[17] However, the percentage was 12.6% for the top 1 million website domains, and certain smaller, security-conscious top-level domains like .BANK have 100% implementation of DNSSEC in their second-level domains. Year to year comparisons of 2024 and 2025 indicate that DNSSEC implementation is growing by 8% a year.

information: the unique number (known as Autonomous System Number or ASN) assigned to that network, and numbers indicating which IP address ranges that network holds (known as a prefix). These announcements are propagated across the Internet by other networks. The ongoing BGP conversation among network routers allows any packet originated by any network in the world to find its destination network.

It became evident in the late 1990s that BGP, like DNS, contained serious vulnerabilities. It lacked any built-in way to authenticate whether the networks making routing announcements were the legitimate holders of the IP address ranges they announced. The absence of such authentication made it possible for traffic to be hijacked. Aside from deliberate attacks, inadvertent errors in the configuration of announcements could also divert traffic in damaging ways. (Smith et al 1998. Kent et al, 2000, Mahajan et al 2002) As these problems became known, the US government in 2004 began to provide research and development support for the application of PKI technology to Internet routing. In 2006, the IETF chartered the Secure Inter-Domain Routing (SIDR) Working Group. The complete architecture was not standardized until February 2012, an indication of the complexity of the problem. The new standard was called Resource PKI (RPKI), with the term "resource" referring to Internet protocol number resources. According to RFC 6480, (Lepinski and Kent, 2012) the RPKI standard "enables an entity to verifiably assert that it is the legitimate holder of a set of IP addresses or a set of Autonomous System (AS) numbers." Network operators are increasingly relying on RPKI to validate routing announcements and reduce the spread of BGP hijacks and misconfigurations. (Gouda, Fontugne & Testart, 2025)

RPKI uses X.509 digital certificates to authenticate the binding between a network's route announcements and its ownership of IP address blocks in that route. It does this by adding a cryptographically verifiable "route origin authorization" (ROA) to the announcements.[18] When ISPs receiving the ROA perform a route origin verification (ROV), they can authenticate routing announcements.

RPKI certificate issuance mirrored the IP address allocation hierarchy, just as DNSSEC based its trust hierarchy on the pre-existing domain name delegation hierarchy. When assigning blocks of IP addresses to network operators, each regional internet address registry issues digital certificates attesting to the origin and identity of number resource blocks and their association with a specific ASN. In both cases, the governance regime is run by nonprofit, non-state actors. Regional Internet Registries are essentially trade associations of network operators and other organizations holding address blocks.

[18] "X.509 Extensions for IP Addresses and AS Identifiers" (RFC 3779)

Instead of a global hierarchy converging on a single root, as the DNS does, the RPKI has five roots, one for each regional Internet address registry: North America (ARIN), Europe and the Middle East (RIPE-NCC), the Asia Pacific (APNIC), Latin America and the Caribbean (LACNIC), and Africa (AFRINIC). Because network operators are often multi-national and IP number blocks are traded or shifted over time, the authority structure is not as simple as the one for DNS, but in all cases the trust hierarchy anchors in regional address registries (RIRs). Some nations have formed National Internet Registries (NIRs) which obtain number blocks for ISPs under their jurisdiction. But the NIRs are subordinate to the private RIRs in the trust hierarchy; their trust anchor also converges on the RIR root CAs.

5. Visualizing Authority Structures

This section develops simplified visualizations of authority structures to illustrate graphically how cyberspace PKIs differ from sovereign state authority structures. Authority structures are visualized as a set of hierarchical levels in which an entity's position indicates whether it is above or below other entities in trust authority.

5.1 The Sovereign State Hierarchy

The "National Public Key Infrastructure" proposal published by S. Chokhani in 1994 (Figure 2) provides a perfect illustration of a sovereignty-based authority structure. Each country would have a "top-level national node," just as it has a single sovereign government. The top of the authority structure, which Chokhani called a "Policy Approving Authority (PAA)" would serve as both the trust anchor for a nation, and as a regulatory authority with the power to approve or reject the certificate issuance policies of the Certification Authorities below it. The second level nodes, called Policy Certification Authorities (PCAs), would establish the certificate issuance policy for the community they serve (e.g., hospitals, state governments, industry sectors). PCAs would issue certificates to level 3 nodes, called Certification Authorities (CAs), that would issue leaf certificates to end users. As for an international trust hierarchy, Chokhani recognized two choices. "The national roots (PAAs) could cross-certify each other, or they could be certified by a global root" which "could be managed by a United Nations agency." Both options follow closely the authority structure of state sovereignty. In one case the sovereigns engage in bilateral negotiations, in the other they pursue a universal, multilateral arrangement via an intergovernmental institution.

Figure 3 shows graphically how this authority structure is isomorphic to territorial sovereignty. At both the domestic and international levels, the Chokhani model situates trust authority in nation-states, thus mirroring in cyberspace the idealized structure of political authority in a geographic territory. Each country has its own CA hierarchy, though

space limitations make it possible to illustrate only a few. Interoperability across national CAs is maintained by bilaterally negotiated cross-certification agreements.

Figure 2: National PKI Proposal, 1994

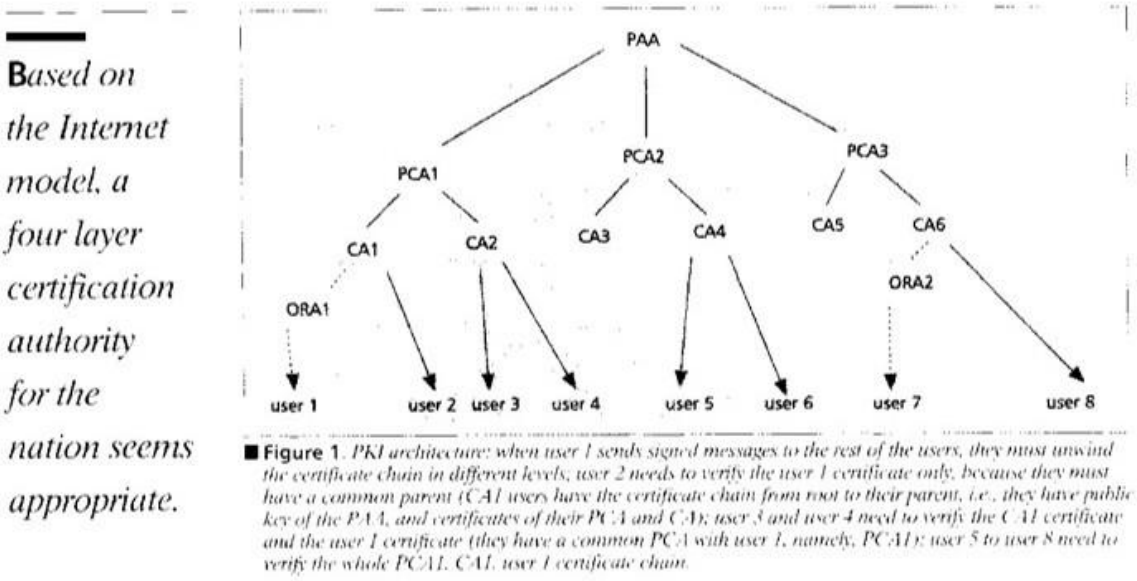

Figure 3: Global System of National-Territorial PKIs

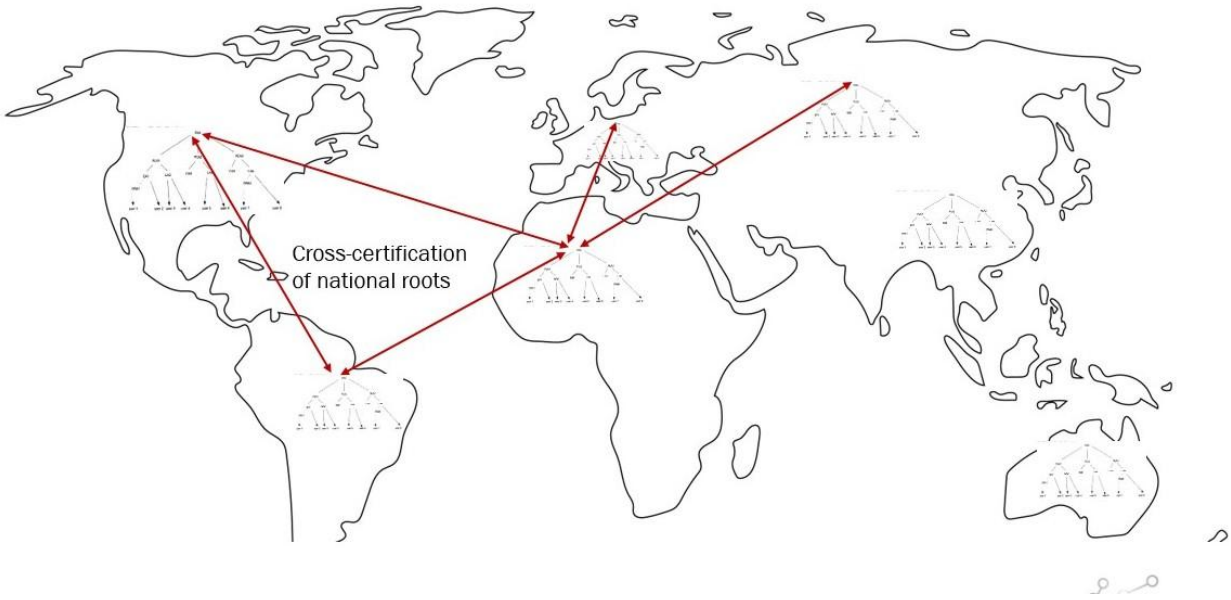

5.2 The WebPKI Hierarchy

The Web PKI trust hierarchy offers a stark contrast with the tidy, state-sovereign order imagined by Chokhani and others before the Internet/Web disruption. Attempts by computer scientists to measure and map Web PKI trust hierarchies uncover a "landscape of certificate validation chains [that] is dynamic, complex, and shaped by evolving

technologies and practices." (Dolbert, et al, 2024 p. 131) The network graph (Figure 4) derived from their longitudinal study of certificate trust chains from 2013 to 2023 shows an intricate web of relations, with many certificates having two or three different paths to a root certificate.

Figure 4, network analysis of WebPKI trust chains from Dolbert et al, 2024

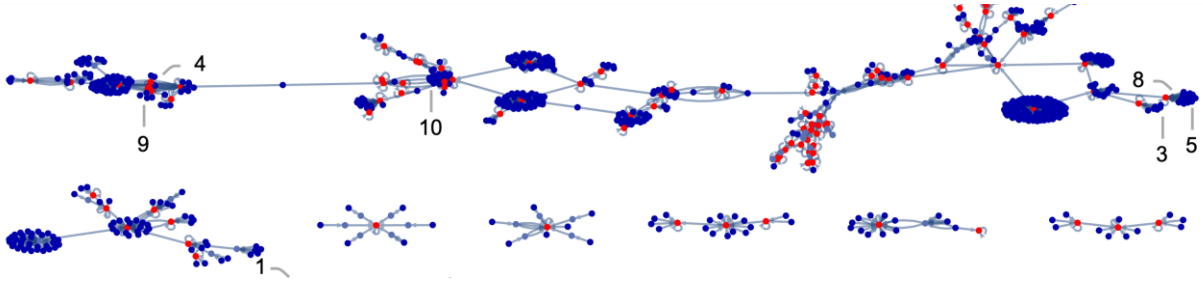

Ma (2021) studied the Web PKI trust hierarchy from 2019 to 2021. He collected the root stores for over 75% of the top HTTP user agents seen at a global content distribution network. This process identified the three independent root store programs that anchor global Web PKI trust, run by Microsoft, Apple, and Mozilla.[19]

Oakes et al (2019) measured Web PKI trust chains from the end user side from July 2017 to January, 2018, using a panel of over 2 million residential participants.[20] This method identified 9.6 million valid certificate chains, but only 6,500 of them account for 90% of the traffic volume. These involved 167,000 Root certificates, 3.5 million intermediate certificates, and 31.57 million leaf certificates. Three percent (3.3%) of the valid certificate chains have two steps from the end user to the root certificate 73% of the valid certificate chains are 3-step; 22% are 4-step; 1% are 5 steps or more. They observed 203 organizations that issued root certificates and 223 that issued intermediate certificates. The top ten organizations issued 89% of the trusted signing certificates.

To simplify a very complex picture, the top of the Web PKI's authority structure consists of four root store programs maintained by transnational, commercial software and service vendors: Microsoft, Apple, Mozilla and Google. (Figure 5) There are about 200 major certificate issuers, predominantly private actors operating globally. The most commonly

[19] In September 2022, after the Ma study was published, Google created its own root store program. *Many other root store providers (e.g., Node, Linux variants) copy their trusted CA list exclusively from Mozilla. NSS dependence, however, "is manually and haphazardly implemented."*

[20] Each time one of the studied end users validated a certificate chain as part of an SSL handshake, the full certificate chain is recorded by the client-side panel software and sent to Comscore storage servers.

Figure 5: Simplified diagram of Web PKI

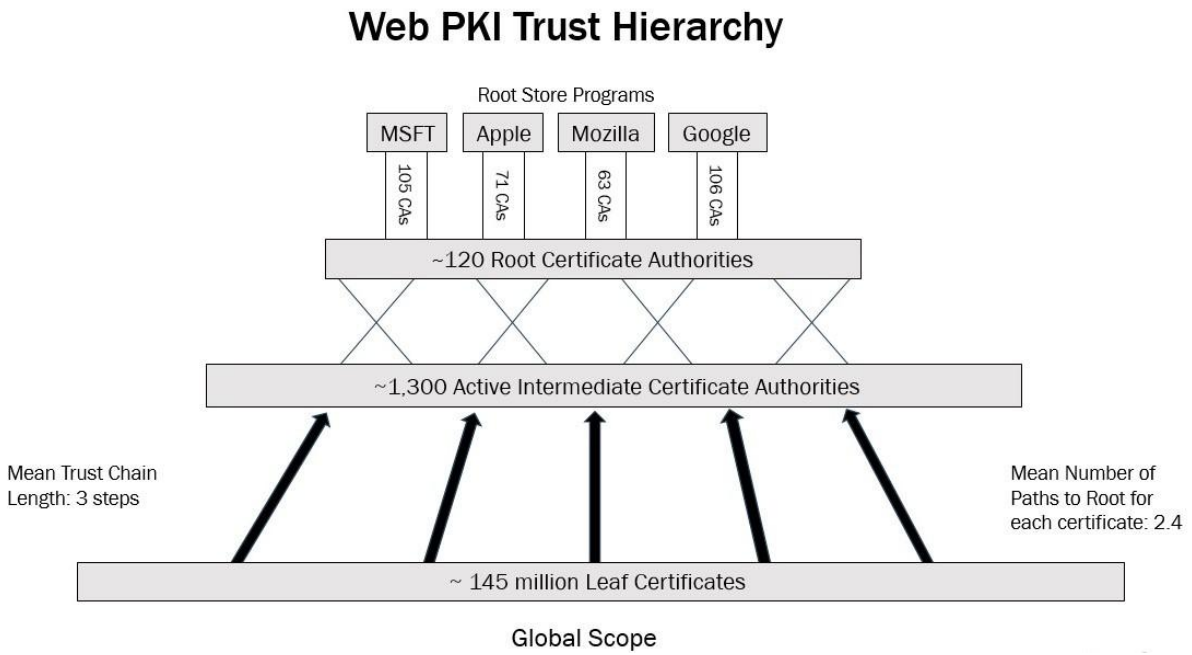

used CA worldwide, Let's Encrypt, is a California nonprofit that uses automated methods to offer domain-validated certificates for free. The companies operating trust stores at the top of the hierarchy have no formal authority over the practices of certificate issuers, but their decision to admit specific CAs into their trusted root stores, or kick them out, is used to enforce security standards and practices. The participants in the Web PKI ecosystem engage in collective action to set security standards and policies in industry institutions such as the CA/Browser Forum and the operation of Certificate Transparency logs.

The degree to which Web PKI certificate patterns don't follow sovereign authority structures is illustrated by an analysis of certificate issuance in China, a country that adheres strongly to norms of state sovereignty on the Internet. We gathered a list of website domains under China's .CN top level domain from the Common Crawl domain vertices dataset (March–May 2025). We then accessed Chinese domains via a proxy server located in Beijing, in case the CAs for internal and external access might differ in China. For each sampled domain, we conducted an active TLS handshake to retrieve the server's presented certificate and extracted the issuing certificate authority from the certificate

metadata.[21] The results show that foreign CAs, especially American ones, issue a dominant share of Web certificates in China.

Figure 6: Certificate Issuer Share of Chinese Websites, 2025.

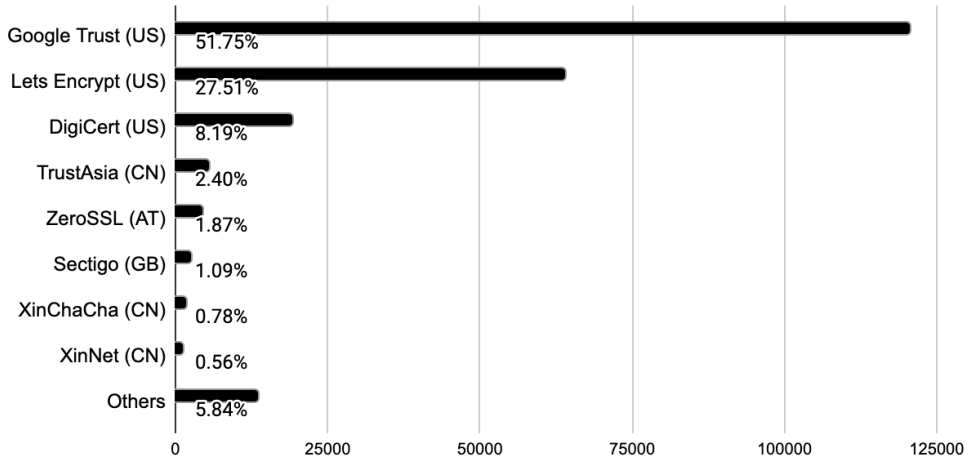

5.3 The DNS PKI Hierarchy

DNSSEC uses a hierarchical chain of trust based on the existing DNS naming structure. Each level of the DNS naming hierarchy, starting with the root, cryptographically signs the records that map a domain name to an IP address. This means that the operators of domain name servers at each level must hash their zone files using their private key and include the hash and the public key-signing key in their answer to DNS queries. Lower-level name servers can then cryptographically verify that the zone file is the correct one.

When DNSSEC is implemented, trust starts with the Root's key-signing key (ICANN), which certifies the resource records of authorized Top-Level Domains. These top-level domains then certify the second-level domains under them that implement DNSSEC, and so on. As a strict hierarchy with a single root operated by a single organization, the DNS PKI hierarchy is easy to visualize (Figure 7).

[21] A single TLS handshake was performed with a timeout of 10 seconds. All measurements were conducted over a single collection window between September 10th and December 5.

Figure 7: Authority Structure of the DNS PKI

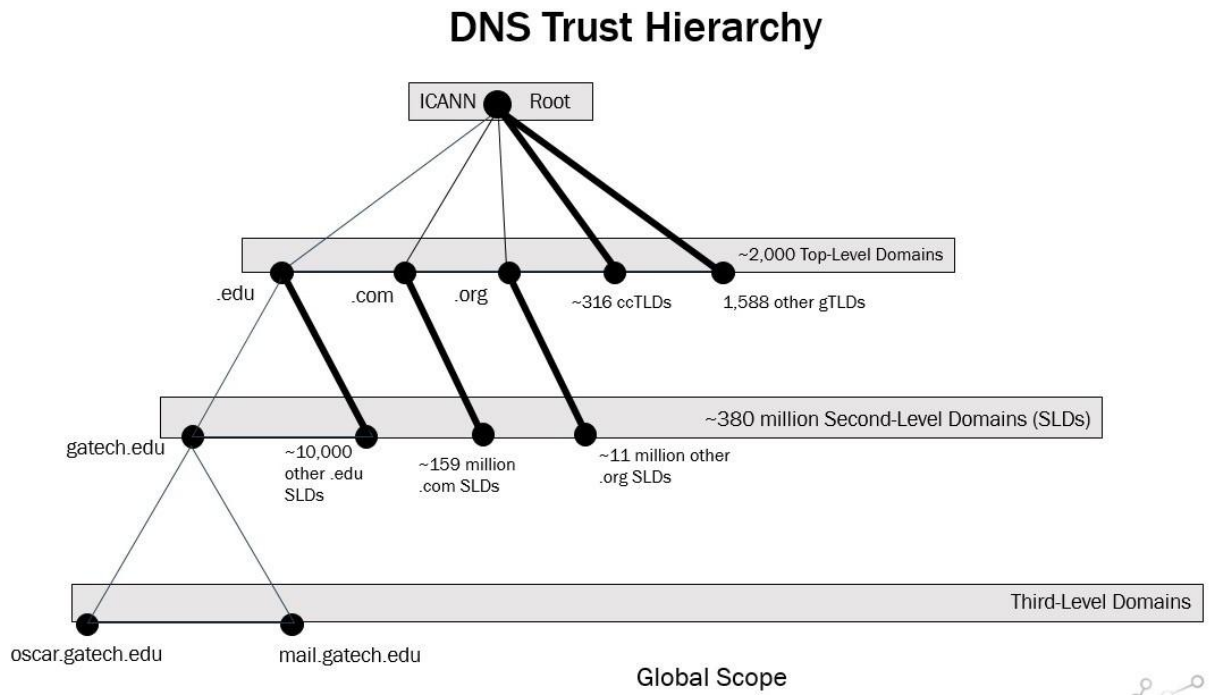

Like the DNS itself, the authority structure of the DNS PKI is global, not territorial, and the apex of the trust hierarchy is a non-state actor, ICANN. ICANN was incorporated in 1998 as a California nonprofit to institutionalize the governance of the DNS root. It was recognized by the U.S. Department of Commerce and other Internet stakeholders as the administrator of the authoritative DNS root and deemed the arena for policy decisions regarding the governance of the DNS. The U.S. relinquished its contractual authority and oversight role over ICANN in October 2016, making it a fully non-governmental governance institution. (Mueller, 2025)

5.4 The RPKI Hierarchy

In the RPKI, certificates are organized in a hierarchy that follows the assignment of IP number resources to network operators. Public Technical Identifiers (PTI),[22] a subsidiary of ICANN, delegates large blocks of IP address numbers to the five regional registries (RIRs). Like ICANN, the RIRs are nonprofit, nongovernmental organizations. As IP numbers are delegated by the RIRs to network operators, certificates are issued that authenticate a network's legitimate possession of the IP number ranges it has been assigned. The

[22] (legacy name IANA

certificate chain roots at each RIR - not at ICANN - and ends with the issuance of ROAs by the network operators' routers. (Figure 7)

The topmost certificates are the RIRs' Trust Anchors, which are self-signed root certificates. (Rodday, Cunha et al, 2023) To protect the security of the trust anchors, the RIR creates one or two layers of intermediate CAs. (Sediqi, Fortugna, et al, 2025; Gouda & Testart, 2024) When the RIR delegates address resources to network operators, its intermediate CAs issue RPKI certificates to network operators or other end-entities, which are used to sign Route Origin Authorizations (ROAs). When included in their routing announcements, ROAs can be used to cryptographically validate the network's ownership of the number ranges it announces.[23]

Figure 6 Trust hierarchy for IP address routing announcements (RPKI)

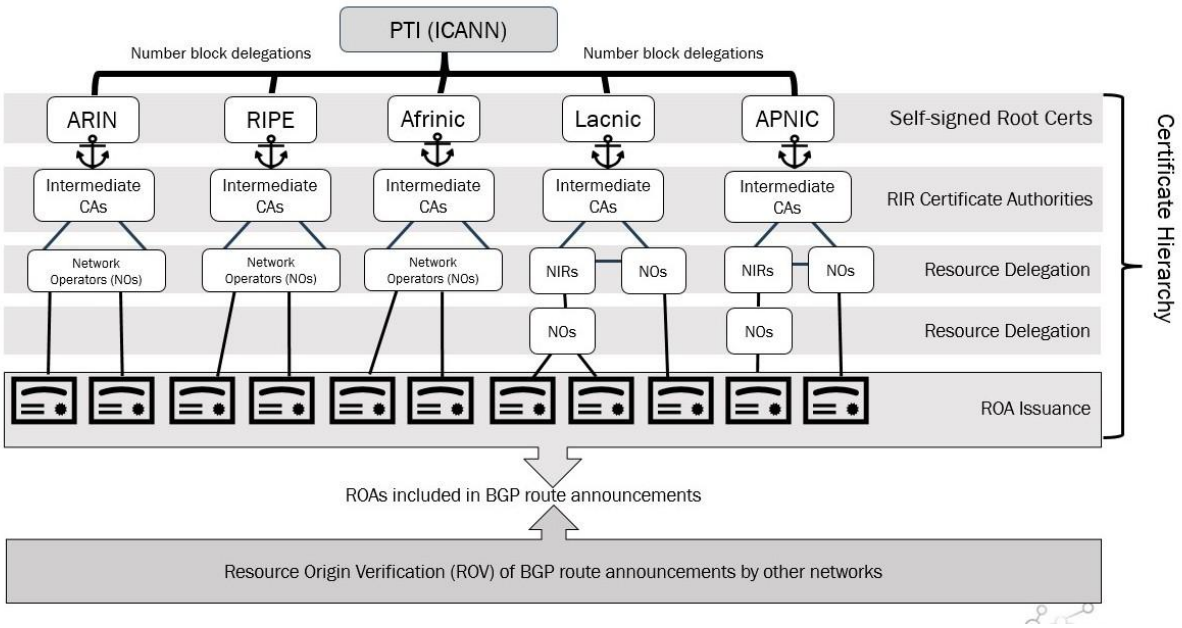

In the Asia-Pacific and Latin American regions, there are National Internet Registries (NIRs) that manage the address space on a national level. NIRs exist in Brazil, Mexico, Vietnam, China, Indonesia, Japan, India and Taiwan. Some of these NIRs (e.g., JPNIC in Japan) are nongovernmental nonprofit membership organizations of network operators, some (e.g. China's CNNIC) are state-run, others (e.g., Brazil's NIC.br) are state-private partnerships.

[23] Each RIR has different policies for sub-delegating certificate issuance to third parties. ARIN, APNIC, and RIPE subdelegate the authority to issue RPKI certificates to 88 different organizations that host and maintain their own RPKI repository. Some of these organizations are large, private sector infrastructure businesses like Amazon. The APNIC repository, however, also contains 4 delegated CAs for National Internet Registries (NIRs) in Japan (JPNIC), Indonesia (IDNIC), Taiwan (TWNIC) and China (CNNIC). The LACNIC registry has a special subdelegation to Brazil, which also operates a NIR. AFRINIC does not have any subdelegations.

NIRs sub-allocate resources to end entities or network operators in their countries, adding another level to the trust chain. Note that while NIRs can be seen as assertions of some degree of national sovereignty over Internet address allocation, NIRs are subordinated in the trust hierarchy to the transnational RIR trust anchors. Their number resources and certificates are delegated to them by the RIR, which sits above them in the hierarchy.

5. Discussion

The construction of a globalized digital infrastructure held together by the Internet protocols produced a revolution in one vital arena of international security relations. It took one of the critical components of interstate relations – the confidentiality and integrity of the public communications infrastructure – and put much of its technical security architecture in the hands of a highly distributed set of nonstate actors. These nonstate actors operate globally and are governed more by market forces, collective action, and engineering principles than geopolitics. This paper documents one important way in which that change manifested itself. It documents, historically and empirically, a divergence between the authority structures of territorial states and the transnational, nonstate actor-based authority structures of the PKIs that support security in cyberspace. This process was not mediated by intergovernmental institutions such as the ITU or other United Nations agencies, but by new institutions of non-state actors.

These PKIs are used to authenticate the identity of websites, encrypt traffic, verify the legitimacy of routing announcements, and protect the integrity of data in domain name records. PKIs are not just technical systems based on applied cryptography; they are shared infrastructures organized around a commonly recognized authority. The users of these infrastructures participate in a hierarchical structure which can automatically, through the exchange of digital signatures and/or certificates, adjudicate trust.

The historical analysis shows how states failed to reproduce their own, sovereign authority structures in applied cryptography. The Internet technical community and business community developed a new, transnational set of authority structures and governance institutions, such as Web PKI and the CA/B Forum, DNSSEC and ICANN, and the RPKI and Regional Internet Registries to oversee them.

The answer to the question of why security in cyberspace did not conform to the established hierarchy of sovereign states is found in the historical evolution of the Internet. In sharp contrast to the prior world of PTT monopolies, the Internet protocols offered any organization in the world the opportunity to declare themselves “a network” and receive globally unique identifiers that would be compatible with any other network in the world.

The networking protocols this regime (TCP/IP, BGP and other IETF standards) were nonproprietary – free software available to anyone in the world. By running these protocols, any actor could communicate globally with few gatekeepers or licensing permissions. Private actors quickly seized upon the opportunities created by this structure to develop globalized communication and commerce relations based on private contracts. This capability rendered State-led efforts to regulate digital certificates irrelevant. The Internet and Web protocols simultaneously privatized and globalized digital networking, putting initiative in the hands of non-state actors interested in globalized commerce and communication. The liberalization of telecommunications and cryptography opened the door to the development of a commercial CA industry and collective action by private actors to produce the public good of online security.

This is not to say that traditional forms of international security have gone away or are no longer important. The development of weapons, armies, navies, and military-strategic intelligence capabilities, including Internet shutdowns, violent invasions and offensive cyber operations, are still options available to nation-states. But states are not in direct control of the basic security architecture of the networks connecting their countries.

References

Ariyapperuma, S., & Mitchell, C. J. (2007, April). Security vulnerabilities in DNS and DNSSEC. In *The Second International Conference on Availability, Reliability and Security (ARES'07)* (pp. 335-342). IEEE.

Arnbak, A, H. Asghari, M. van Eeten, N.A.N.M. van Eijk, Security Collapse in the HTTPS Market, Communications of the ACM, 2014-10, vol. 57, p. 47-55.

Balbi, G., & Fickers, A. (Eds.). (2020). *History of the International Telecommunication Union (ITU): Transnational techno-diplomacy from the telegraph to the Internet* (Vol. 1). Walter de Gruyter GmbH & Co KG.

Berkowits, S. et al., Public Key Infrastructure Study, National Institute of Standards and Technology, MITRE Technical Report, April, 1994

Benantar, M. The Internet public key infrastructure. (2001) IBM Systems Journal, VOL 40, NO 3. 648-665. https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5386926

Berkovits, S., Chokhani, S., Furlong, J. A., Geiter, J. A., & Guild, J. C. (1994). Public key infrastructure study: Final Report. National Institute of Standards and Technology.

Berners-Lee, T. (1999). *Weaving the Web: The original design and ultimate destiny of the World Wide Web by its inventor*. Harper San Francisco.

Biddle, C. B. (1997). Legislating market winners: Digital signature laws and the electronic commerce marketplace. *San Diego L. Rev.*, *34*, 1225.

Chokhani, S. (1994). Toward a national public key infrastructure. *IEEE Communications Magazine*, *32*(9), 70-74.

Cowhey, P. F. (1990). The international telecommunications regime: The political roots of regimes for high technology. *International Organization*, *44*(2), 169-199.

Corwin, P. S. (1998). Electronic authentication: the emerging federal role. Jurimetrics, 38(3), 261-276.

Davids, M. (1995). The Relationship between the State Enterprise for Postal, Telegraph and Telephone Services and the State in the Netherlands in Historical Perspective. *Business and Economic History*, 194-205.

Dessy, (1997). The Code Makers. Trends in Analytical Chemistry, 16(1).

Delignat-Lavaud, A, Abadi, M. 2017. Web PKI: Closing the Gap between Guidelines and Practices. Network and Distributed System Security (NDSS) Symposium. https://www.ndss-symposium.org/wp-content/uploads/2017/09/12_1_1.pdf

Diffie, W. and Hellman, M. 1976. "New Directions in Cryptography." IEEE Transactions on Information Theory 22: 644-654.

Diffie, W., & Landau, S. (2007). Privacy on the Line: The Politics of Wiretapping and Encryption (2nd ed., p. 400). Cambridge, Mass: The MIT Press.

Döberl, M., Von Wangenheim, Y. F., Bruhner, C. M., Hasselquist, D., Arlitt, M., & Carlsson, N. (2024, June). Chain-Sawing: A Longitudinal Analysis of Certificate Chains. In 2024 IFIP Networking Conference (IFIP Networking) (pp. 131-139). IEEE.

Drake, W. J. (2005). The rise and decline of the international telecommunications regime. In Regulating the global information society (pp. 145-196). Routledge.

Drake, W. J. (1993). The Internet religious war. Telecommunications policy, 17(9), 643-649.

Firth, C. H. (1898). Thurloe and the post office. The English Historical Review, 13(51), 527-533.

Froomkin, A. M. (1996). The essential role of trusted third parties in electronic commerce. Or. L. Rev., 75, 49.

Froomkin, A. M. (1996a). It Came From Planet Clipper: The Battle Over Cryptographic Key Escrow. U. Chi. Legal F., 15.

Gerck, E. (1997). Overview of Certification Systems: x. 509, CA, PGP and SKIP. The Black Hat Briefings, 99. https://www.researchgate.net/profile/Ed-Gerck/publication/318700731_original_web_site_Overview_of_Certification_Systems/data/597829520f7e9b277721d8ce/certhtm.pdf

Gillies, J., & Cailliau, R. (2000). How the Web was born: The story of the World Wide Web. Oxford University Press, USA.

Gouda, D., Fontugne, R., & Testart, C. (2025, October). ru-RPKI-ready: the Road Left to Full ROA Adoption. In *Proceedings of the 2025 ACM Internet Measurement Conference* (pp. 415-429).

Greenwood, D. J. (1998). Risk and trust management techniques for an open but bounded public key infrastructure. Jurimetrics, 38(3), 277-294.

Gutmann, P. (2002). PKI: it's not dead, just resting. Computer, 35(8), 41-49. https://people.cs.vt.edu/~kafura/cs6204/Readings/Authentication/PKINotDead.pdf

International Traffic in Arms Regulations, 22 CFR 120-130, Federal Register Vol. 58, No. 139, 22 July 1993.

Jarvis, C., & Martin, K. M. (2024). A milestone in encryption control – what sank the US key-escrow policy? Intelligence and National Security, 39(6), 986-1008.

Kouli, Y., & Laborie, L. (2023). The European Making of National Public Services—Posts and Telegraphs. In The Politics and Policies of European Economic Integration, 1850–1914 (pp. 31-71). Cham: Springer International Publishing.

Kohnfelder, L. 1978 "Towards a Practical Public-Key Cryptosystem." MIT Bachelor's Thesis.

Kuerbis, B. (2011). “Securing critical Internet resources: Influencing Internet governance through social networks and delegation.” Doctoral dissertation, School of Information Studies, Graduate School of Syracuse University.

Landau, S., Kent, S., Brooks, C., Charney, S., Denning, D., Diffie, W., ... & Sobel, D. (1994). Codes, Keys, and Conflicts: Issues in US Crypto Policy. Association for Computing Machinery.

Lepinski, M. and Kent, S. (2012). “An Infrastructure to Support Secure Internet Routing,” RFC 6480. Internet Society.

Linn, J. (2000). Trust models and management in public-key infrastructures. RSA laboratories, 12. https://lasr.cs.ucla.edu/classes/239_1.fall10/papers/trust_pki_rsa.pdf

Lynn, C, Kent, S., Seo, K. (BBN Technologies). "X.509 Extensions for IP Addresses and AS Identifiers," RFC 3779.

Ma, Z. Z. (2021). *Understanding the trust relationships of the web PKI* (Doctoral dissertation, University of Illinois at Urbana-Champaign).

Maeda, A. (2004). 13. PKI solutions for trusted e-commerce: A survey of the de facto standard competition in PKI industries. Information Technology Policy and the Digital Divide, 260.

Merrill, C. R. (1998). The accreditation guidelines: A progress report on work in process of the ABA Information Security Committee. Jurimetrics, 38(3), 345-358.

Moore, Ryan J. (2002) The History, Usage and Implementation of Digital Certificates. Progress Report, Department of Computer Science, University of Warwick, UK.

Philpott, D. (2003). Sovereignty. Stanford Encyclopedia of Philosophy. https://plato.stanford.edu/ENTRIES/sovereignty/

Rodday, N., Cunha, Í., Bush, R., Katz-Bassett, E., Rodosek, G. D., Schmidt, T. C., & Wählisch, M. (2023). The resource public key infrastructure (RPKI): A survey on measurements and future prospects. *IEEE transactions on network and service management*, *21*(2), 2353-2373.

Roosa, S.B., Schultze, S. The “Certificate Authority” trust model for SSL: a defective foundation for encrypted Web traffic and a legal quagmire. Intellectual Property & Technology Law Journal 22. 11 (2010), 3.

Sediqi, K. Z., Fontugne, R., Phokeer, A., Stucchi, M., Candela, M., & Feldmann, A. (2025, June). RPKI Syncing: Delay in Relying Party Synchronization. In *2025 9th Network Traffic Measurement and Analysis Conference (TMA)* (pp. 1-11). IEEE.

Shepherd, S. J. (1996, April). Lessons learned from security weaknesses in the Netscape World Wide Web browser. In IEE Colloquium on Public Uses of Cryptography (pp. 7-1). IET.

Vossen, G., & Hagemann, S. (2007). From Version 1.0 to Version 2.0: A brief history of the web (No. 4). ERCIS Working Paper.

Vratonjic, N., Freudiger, J., Bindschaedler, V. and Hubaux, J.-P. The inconvenient truth about Web certificates. In Proceedings of the Workshop on Economics of Information Security, 2011.

Ye, E. Z., Yuan, Y., & Smith, S. (2002). Web spoofing revisited: SSL and beyond. https://digitalcommons.dartmouth.edu/cs_tr/193/

Yurchenko, S. B. (2017). Logic of order: state hierarchy, law, sovereignty, and war. International Review of Sociology, 27(2), 291-318.