



Reflections on trusting distributed trust

Emma Dauterman*
UC Berkeley

Vivian Fang*
UC Berkeley

Natacha Crooks
UC Berkeley

Raluca Ada Popa
UC Berkeley

Abstract

Many systems today distribute trust across multiple parties such that the system provides certain security properties if a subset of the parties are honest. In the past few years, we have seen an explosion of academic and industrial cryptographic systems built on distributed trust, including secure multi-party computation applications (e.g., private analytics, secure learning, and private key recovery) and blockchains. These systems have great potential for improving security and privacy, but face a significant hurdle on the path to deployment. We initiate study of the following problem: a single organization is, by definition, a single party, and so how can a single organization build a distributed-trust system where corruptions are independent? We instead consider an alternative formulation of the problem: rather than ensuring that a distributed-trust system is set up correctly by design, what if instead, users can audit a distributed-trust deployment? We propose a framework that enables a developer to efficiently and cheaply set up any distributed-trust system in a publicly auditable way. To do this, we identify two application-independent building blocks that we can use to bootstrap arbitrary distributed-trust applications: secure hardware and an append-only log. We show how to leverage existing implementations of these building blocks to deploy distributed-trust systems, and we give recommendations for infrastructure changes that would make it easier to deploy distributed-trust systems in the future.

CCS Concepts

• Security and privacy → Distributed systems security;

Keywords

distributed trust, multi-party computation

*Equal contribution.



This work is licensed under a Creative Commons Attribution International 4.0 License.

HotNets '22, November 14–15, 2022, Austin, TX, USA

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9899-2/22/11.

<https://doi.org/10.1145/3563766.3564089>

ACM Reference Format:

Emma Dauterman, Vivian Fang, Natacha Crooks, and Raluca Ada Popa. 2022. Reflections on trusting distributed trust. In *The 21st ACM Workshop on Hot Topics in Networks (HotNets '22)*, November 14–15, 2022, Austin, TX, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3563766.3564089>

1 Introduction

Distributing trust is a powerful tool for building efficient systems with strong privacy and integrity properties. A distributed-trust system is deployed across n parties (we will subsequently refer to these as trust domains) where, if there are no more than f independent corruptions, the system provides certain security, privacy, and/or integrity properties. In the past few years, we've seen an explosion of academic and industrial cryptographic systems built on distributed trust; some applications are based on secure multi-party computation [9, 34, 86] (e.g. private search [23, 24, 64, 82], private analytics [7, 11, 19, 27, 42] private media delivery [38], private blacklist lookups [46], private DNS [72], anonymous messaging [17, 20, 28, 48, 49, 84], and cryptocurrency wallets [30, 31, 45, 63, 67, 70, 75]), while others are based on Byzantine fault-tolerant consensus and blockchains [6, 8, 15, 25, 36, 37, 47, 50, 56, 88].

In this paper, we initiate the academic study of an often overlooked challenge that distributed-trust systems face on the path to deployment: *bootstrapping a distributed-trust system is surprisingly difficult*. Distributed-trust systems only provide strong security guarantees insofar as there is no central point of attack that allows an attacker to compromise more than an application-specific threshold of parties. Many academic works simply assume the existence of multiple non-colluding servers [20, 23, 24, 28, 38, 38, 46, 82], but an application developer that wants to deploy a distributed-trust system faces difficult questions: who should have administrative control over the different trust domains, and how do you convince another party that you do not control to run your system? We study the difficulties developers have faced when deploying distributed-trust systems in §2.

To make this challenge more concrete, consider a simple application that provides backups for secret keys (e.g., for end-to-end encrypted messaging or cryptocurrency wallets). The user splits its secret key across different trust domains via secret sharing [71]. Therefore, even if the attacker steals secret shares from all but one of the trust domains, the attacker cannot learn users' secret keys (Figure 1). As a strawman,




Figure 1: The application developer sets up a distributed-trust system for secret-key backups. The developer should not be a central point of attack in the system.

the application developer could deploy virtual machines on different cloud providers so that different trust domains correspond to different cloud providers. The problem with this strawman solution is that the application developer still has administrative control over the virtual machines and so is a central point of attack: if the attacker compromises the developer's credentials, the attacker can easily recover every user's secret key.

To achieve strong security guarantees in practice, we would like trust domains to be truly independent so that there is no single point of failure. For security, compromising any system component should always ideally compromise at most one trust domain. For convenience, a developer should ideally be able to set up a distributed-trust deployment cheaply and easily. In reality, absolute separation is very hard to achieve given the architecture of today's systems, and so we must settle for some approximation of independence. One good approximation of independence would be to have the application developer coordinate with and cede administrative control to different organizations that manage servers running on different cloud providers with different hardware, potentially even in different geographic locations [1, 27, 33, 81]. In this way, even if the developer's credentials are compromised, the attacker cannot subvert every trust domain and compromise the entire system. While this first attempt provides security, it does not provide convenience, as it requires time-consuming and expensive human coordination across organizations.

Thus we reach an impasse: *how can an application developer bootstrap a distributed-trust system without herself becoming a central point of attack?*

It is impossible to bootstrap trust out of nothing. We consider an alternative formulation of the bootstrapping problem: instead of ensuring that a system distributes trust correctly by design, what if users can *audit* a distributed-trust deployment? We are inspired by the widely deployed certificate transparency infrastructure [51] where, instead of preventing certificate authorities (CAs) from issuing bad certificates, users can simply detect CA misbehavior. In our setting, we also provide transparency: rather than guaranteeing that the system always distributes trust correctly, we

instead only guarantee that the user will be able to *detect whenever the system does not execute the expected code in different trust domains*. Moreover, the user will obtain a publicly verifiable proof of misbehavior.

We propose a framework that enables a non-expert application developer to deploy a distributed-trust system in a publicly auditable way without human-level cross-organization coordination. To do this, we identify two core, application-independent building blocks from which an application developer can bootstrap any distributed-trust application: secure hardware and an append-only log. While both of these building blocks are themselves challenging to build and introduce additional assumptions, as a community, we have already built reasonable approximations of both (§4) that we can leverage to build arbitrary distributed-trust applications.

To provide public auditability, we use secure hardware to attest to the code running in each trust domain, and we use the append-only log to allow the developer to commit to the application code. Secure hardware such as trusted execution environments (TEEs) [4, 13, 21, 53, 57] allow the client to verify code integrity (i.e. the secure hardware is running the code that the client thinks it should be running). To learn what code each trust domain should be running, the client can check an application-independent, append-only log (we discuss different possible implementations of such a log in §4). The client can query the log for a hash of the code that the application developer has committed to and then check that the code run by that trust domain matches the code in the log. The developer must publish her code to allow clients to inspect it and check that it hashes to the value in the append-only log.

Because we use secure hardware, it might seem like distributed trust is unnecessary now: we can simply run the entire application inside of secure hardware. However, this would make a single type of secure hardware a central point of attack. For example, if we run the application inside Intel SGX and an attacker finds an exploit in SGX, then they can compromise the security of the entire system (this seems plausible given the history of attacks on SGX [18, 35, 65, 68, 76, 77, 79, 80]). Many systems use distributed trust precisely because they do not want security of the entire system to reduce to the security of a TEE. To address this issue, we set up our system to split trust across multiple trust domains and use heterogeneous secure hardware.

We demonstrate how organizations can easily bootstrap a distributed-trust application by building on cloud offerings for secure hardware and deployed certificate transparency infrastructure [51]. While it is possible to use only existing infrastructure, we also describe infrastructure-level changes that would make the distributed-trust ecosystem more efficient and sustainable. To demonstrate the feasibility of our framework, we implement and evaluate a prototype (§5).

Limitations. Our proposal still has several limitations. The first main drawback is the reliance on secure hardware. We can prevent secure hardware from becoming a central point of attack by using heterogeneous secure hardware. The second drawback has to do with performance: running an application inside a TEE is more expensive than running it natively (§5). In §4.2, we describe how secure hardware manufacturers could design TEEs and how cloud providers could offer services specifically tailored to distributed-trust systems to minimize this overhead.

2 Distributed-trust deployments today

We start by examining the challenges organizations face in deploying distributed-trust applications today. We base our discussion of organizations' solutions on published documentation, including whitepapers and blog posts.

Privacy-preserving analytics. Prio [19] splits trust across two servers to compute aggregate statistics without revealing individual users' data and has been deployed for Firefox telemetry [27] and COVID-19 exposure notification analytics [7]. For Firefox telemetry, Firefox runs one server and the ISRG (the public-benefit corporation behind Let's Encrypt, which offers free TLS certificates) runs the other. The COVID-19 exposure notification system computes statistics across iOS and Android users where the ISRG and the National Institute of Health each run a server. To enable other organizations to easily run a Prio system, the ISRG has announced Divvi Up, a service where the ISRG acts as the second trust domain for a Prio deployment (the organization building the application acts as the "first trust domain") [1, 42].

While Divvi Up will make it easier for organizations to deploy private analytics systems [11, 19], it still does not enable general-purpose distributed-trust systems. Moreover, the challenges ISRG faced in setting up Divvi Up illustrate just how hard it is to set up a distributed-trust system correctly [32, 33]. For example, debugging and running integration tests now must take place across organizations that don't have a common release process or deployment system.

Digital advertising. Meta recently announced their Private Lift Measurement solution [60] in which Meta and an advertiser run a multi-party computation protocol [9, 34, 86]. Multi-party computation is a cryptographic tool that allows multiple parties to jointly compute some function over their secret inputs such that the parties only learn the output of the computation (and not the other parties' secret inputs). In Meta's use-case, the advertiser can learn how their campaign is doing without revealing unnecessary information to Meta or the advertiser. Even with Meta's resources, Meta reported that it initially took months to onboard a new advertiser [66].

Private DNS. Cloudflare, Apple, and Fastly authored an IETF draft for oblivious DNS over HTTPS that splits trust between

a proxy and a resolver such that neither learns both the user's IP address and query [44, 72, 81]. Internet service providers PCCW, SURF, and Equinix have committed to launching proxies, enabling the set of organizations running proxies to be disjoint from the set of those operating resolvers.

Permissioned ledgers with infrastructure providers. To run a new ledger, organizations need to start many nodes quickly. Infrastructure providers like Blockdaemon [10], Alchemy [3], or Figment [29] offer nodes as a service. For example, Blockdaemon provides end-to-end nodes for permissioned ledgers like Diem [8] and Hyperledger Fabric [6] (among many others). However, these infrastructure providers are themselves centralized; compromising a provider like Blockdaemon could enable an attacker to compromise a significant fraction of system nodes. The existence of these infrastructure providers illustrates that organizations need a way to easily add nodes to ledger systems.

Financial custody. Users transfer cryptocurrency by signing transactions, and so transaction signing keys can secure millions of dollars. Many financial custody companies deploy solutions where the signing key is split across hardware security modules (HSMs), and the HSMs run a multi-party computation to generate a signature on a transaction [30, 31, 45, 63, 67, 70, 75]. In this way, no HSM ever holds the entire signing key.

A limitation of the financial custody companies we surveyed is that they only provide security if the company is honest at setup time. One company deploys and maintains all of the secure hardware, and the end-user cannot check that the system is set up and distributes trust in the way that the company claims. Furthermore, if the company locks itself out of its machines to defend against post-setup compromise, there is no way to patch bugs or push updates.

2.1 Our findings

The applications we survey fall into one of two categories:

- The application setup was challenging, application-specific, and required cross-organization coordination.
- The application architecture compromises on distributed trust or functionality in some way (beyond what we could reasonably hope for).

In the first category, we have the ISRG's Divvi Up, Meta's Private Lift Measurement, and oblivious DNS. In all of these cases, the system splits trust across organizations correctly, but the developers had to overcome significant hurdles to coordinate across organizations. In the second category, we have some permissioned ledger deployments and financial custody solutions. The challenges we see in these deployments illustrate how difficult it is to truly eliminate a central point of attack, especially as the number of parties grows (e.g. for permissioned ledgers).

Going forward. While applications in the first category provide strong security guarantees, setting up this type of deployment is simply too difficult for many small organizations. In the remainder of this paper, we describe how a developer can set up a distributed-trust application *without expensive, cross-organization coordination*.

3 System overview

We start by describing the building blocks necessary to bootstrap arbitrary distributed-trust applications (§3.1), then describe our system architecture (§3.2), and finally state the properties our system does and does not provide (§3.3).

3.1 Building blocks

Our system requires two core, application-independent building blocks: secure hardware and an append-only log.

Secure hardware. Secure hardware should be able to attest to the code that is running. In particular, the client should be able to verify that it is communicating with a correctly provisioned piece of secure hardware running software that hashes to a particular value. In addition, if a developer is running an integrity-preserving application, the secure hardware should isolate memory or detect tampering, and for privacy-preserving applications, the secure hardware should also encrypt the memory contents. Existing industrial [4, 57] and academic [13, 21, 53] TEEs provide these properties (note that because TEEs generally do not protect memory access patterns [14, 39, 52, 54, 59, 69, 78, 85], developers writing privacy-preserving applications should ensure that memory access patterns do not reveal secret data). In §4.2, we discuss future alternatives to TEEs.

Append-only log. The append-only log should provide integrity: once an entry is added, it cannot be altered or deleted. This log could be implemented as a permissioned ledger [6, 8] or a centrally maintained log that permits public auditing [5, 16, 22, 41, 51, 58, 73, 74]; the only requirement is that the log itself does not become a central point of attack.

Using an append-only log to bootstrap distributed trust initially seems contradictory: if we already have an append-only log built on distributed trust, why do we need to bootstrap distributed trust? The key distinction here is that the append-only log is not application-specific. We want to use some *application-independent* log to bootstrap an arbitrary application easily. We will show how to do this with the currently deployed certificate transparency log [51] (§4.1).

3.2 System architecture

We describe the architecture of a system that distributes trust across n trust domains (Figure 2). Each trust domain runs a server equipped with secure hardware. Ideally trust domains should leverage different types of secure hardware to minimize the chance that an exploit in one type of secure




Figure 2: System architecture. We show two trust domains for simplicity, but in practice the number of trust domains is application-dependent. Trust domain 0 is run by the application owner without any secure hardware.

hardware compromises the entire system. The application developer can run one trust domain on her own without any secure hardware (denoted as “trust domain 0” in Figure 2).

Note that for applications with a small number of trust domains, it is possible to run each trust domain on a different type of secure hardware. For applications with a large number of trust domains, the fact that there are a comparatively small number of secure hardware vendors means that some trust domains might have to use the same type of secure hardware, potentially introducing correlated corruptions (the application only provides security or privacy if the attacker can compromise more than an application-specific threshold of trust domains). We discuss how to improve this shortcoming moving forward in §4.2.

The application developer and client should also be able to interact with an append-only log. The developer can post code updates here, and the client can then check that the trust domains are running the correct code.

3.3 System properties

Given a distributed-trust application with n trust domains that provides security if at least t are honest, we provide the following guarantees.

Auditable by client. Given access to a shared append-only log, the client can check that each of the n trust domains contains a secure hardware element running the code the application developer committed to. The log ensures that the application developer commits to the same code for all clients. Moreover, if the secure hardware is running code that does not correspond to the hash published in the log, the client obtains a publicly verifiable proof that the secure hardware was running the wrong code. The developer must also open-source her code so that clients can inspect the published code to make sure that it does what the developer claims and check that the hash of the published code matches the hash stored in the log. We provide security if the published code is running correctly in at least t trust domains.

Simple for application developer. Our system enables application developers to bootstrap distributed-trust systems using existing cloud resources and infrastructure. Human-level cross-organization coordination is not necessary.

Supports code updates. Application developers can securely update their code. Code updates are necessary to fix security-critical bugs and support new features. Clients learn when the code running in different trust domains is updated, and they can check that the new code matches the hash in the log. Moreover, clients can monitor the log to make sure that the trust domains correctly update to the new code when the developer publishes an update. Because the code is open-source, clients can check the contents of updates.

Non-goals. We do not defend against implementation bugs or backdoors. By examining the application code, the client can gain confidence that the code is doing what the application developer claims, but different implementations or formal verification would be necessary to protect against correlated compromise due to the implementation. Similarly, we can only provide limited protection in the case where the developer is herself the attacker and is allowed to push code updates; the developer can insert backdoors that could be very challenging to detect. Therefore for highly sensitive applications, a developer might consider disabling her ability to push code updates to defend against future compromise.

4 System design

We now describe our system design. We first show how to bootstrap distributed trust today (§4.1). We then outline infrastructure-level changes that would better enable bootstrapping in the future (§4.2).

4.1 Deployment today

To bootstrap distributed trust today, we can build on top of a TEE like Intel SGX or AWS Nitro (for the secure hardware) and certificate transparency (for the append-only log). Cloud providers already provide access to TEEs, with Microsoft Azure supporting Intel SGX and AWS supporting the Nitro enclave. Certificate transparency is a widely deployed system for web certificates [51], and so using the techniques we describe below, an application developer can deploy a distributed-trust application immediately. As a starting point, we first explain how deployment without updates works, and then we show how to layer on support for updates.

Starting point: deployment without updates. Without updates, the system design is straightforward: the developer seals the application code directly inside a TEE in each trust domain. The client can then use the TEE’s attestation mechanism to receive a hash of the sealed code, which it can compare to the hash in the append-only log. If the two match, the client knows that it is communicating with a secure enclave that is running the code that the application developer

committed to. Because the code is sealed onto the enclave, the application developer cannot change the code, which provides security, but also makes code updates impossible.

Supporting updates. Application developers need the ability to fix security-critical bugs and support new features. Supporting updates with current TEEs is challenging because existing TEEs do not support updating the existing code while maintaining the current state of the running application. Moreover, the client needs some way to learn when an update has happened and check that the update was performed correctly. Because we need to defend against malicious updates, we cannot make any assumptions about the behavior of the new code (e.g., we cannot assume that the new code will correctly alert the client that an update occurred).

To address these problems, we add a layer of indirection. Instead of sealing the developer’s code directly on to the enclave, we instead seal an *application-independent framework* on to the TEE. This application-independent framework accepts application code as input and executes it. When the application developer wants to update the code running, she sends the new code to the TEE. Before the TEE starts running the new code, it alerts the client that an update is about to take place and sends the user a hash of the updated code. The client can check this hash against the hash in the log. We can open-source this application-independent framework in order to increase confidence in this framework, and, as before, the client can use attestation to verify that the framework is running correctly on the TEE.

We need to ensure that a malicious update does not prevent the application-independent framework from notifying the client that an update took place. To ensure that the update cannot interfere with the framework, we run the updated application code inside of a sandbox [83, 87]. Sandboxing the application code ensures that the executed code cannot “escape” the sandbox and have an effect on the system outside the sandbox (i.e. the framework). We also need to ensure that the TEE only runs updates from the application developer. We can do this easily by sealing on to the TEE not just the framework, but also a public key. Then each subsequent update needs to be accompanied by a signature that verifies under the original public key.

Leveraging the certificate transparency log. Using the certificate transparency log [51] means that we do not need to deploy a new log, but the certificate transparency log is designed for X.509 web certificates, not application updates. We can use the subject key identifier field in the X.509 certificate format to store the application digest. Prior work has explored transparency logs for application binaries, but in the context of verifying local client code rather than remote server code [2, 40, 61].

4.2 Deployment tomorrow

Infrastructure changes would make deployment even easier and provide greater flexibility in the future.

Expanding cloud provider offerings. Cloud providers should offer services specifically tailored for distributed-trust systems. In particular, cloud providers should allow developers to submit code and code updates and then run the code without allowing the developer to inspect or modify application memory. Clients interacting with the system could get some assurance from the cloud provider that the claimed code is running.

Secure hardware design. TEEs like SGX support sealing code on to the TEE. This design decision requires us to seal our general-purpose framework on to the TEE and then have our framework run the dynamic application code in a software-based sandbox. Changing the hardware design could allow us to support updates much more efficiently. Instead of running the new binaries inside a software sandbox inside the TEE, the hardware could instead isolate the framework from the application binary directly. We simply need the secure hardware to attest to the framework that is running and provide a mechanism for the framework to effectively sandbox the new binary. We hope that our work spurs the development of secure hardware explicitly tailored to bootstrapping distributed-trust systems efficiently.

Building a distributed-trust log. Leveraging the certificate transparency log adds unnecessary overhead for certificate transparency stakeholders. We could instead deploy a log specifically for bootstrapping distributed-trust systems that makes it simple to look up binaries [2, 40, 61].

5 Evaluation

We implemented a prototype of our framework and support execution on AWS Nitro. We use WebAssembly (Wasm) [83] as our sandboxed execution environment. We compile C++ applications into Wasm using Emscripten [26] and run Wasm applications inside Node.js [62]. We implement a BLS threshold signature [12] application on top of our framework using libBLS [55]: each trust domain stores a secret key share, and the trust domains can jointly sign a message.

We evaluate our framework on a single AWS c5.4xlarge instance with a 16-core Intel Xeon 8124M CPU and 32GB RAM. We allocate 4GB RAM and 2 cores for the Nitro TEE.

Certificate transparency log. Although we did not submit certificates to avoid log pollution, we observed that it took on average 35ms to query `ct.googleapis.com` for a certificate (1.7ms RTT) and 1ms to hash our application binary.

Framework overheads. Table 3 shows the threshold signing time for different execution environments. Our baseline measures the processing time of the C++ implementation

Execution Environment	Processing Time	Increase
Baseline	10.2ms	—
Sandbox	14.9ms	46.1%
TEE + Sandbox	15.8ms	54.9%

Table 3: Processing time for producing a BLS threshold signature share under different execution environments. The baseline corresponds to native execution (no TEE and no sandbox).

without a TEE and without sandboxing. Compiling to Wasm and running inside of Node.js imposes a 46.1% overhead (comparable to a previous study on Wasm performance [43]). Running the sandboxed application inside the AWS Nitro TEE increases processing time by 54.9%. This overhead is due to the fact that we need two additional sockets: one to forward request traffic from the client to our framework, and one inside the TEE to communicate between our framework and the sandboxed application.

6 Conclusion

Previously, distributed-trust systems were only an option for organizations that could successfully coordinate with other organizations. However, bootstrapping without cross-organization coordination can enable small organizations to securely deploy distributed-trust systems. For example, end-to-end encrypted messaging applications could use distributed trust to establish a public-key infrastructure or back up secret keys (each trust domain stores a secret key share). We hope that our work motivates the study of building blocks (i.e. secure hardware and append-only logs) tailored specifically to distributed-trust systems. We also leave to future work the question of if these building blocks are necessary for bootstrapping distributed trust, or if there are a completely different set of building blocks we could leverage instead. What other trade-offs can developers make to securely bootstrap distributed-trust systems without requiring cross-organization coordination?

Acknowledgments. We thank the anonymous reviewers for their helpful feedback. We also thank Miles Wada for participating in early stages of this work, as well as Narek Galystan, Jack Humphries, Aurojit Panda, Samyu Yagati, Wen Zhang and students in the Sky security group for feedback that improved the presentation of the paper. This work is supported by NSF CISE Expeditions Award CCF-1730628, NSF CAREER 1943347, and gifts from the Alibaba, Amazon Web Services, Ant Group, Astronomer, Ericsson, Facebook, Futurewei, Google, IBM, Intel, Lacework, Microsoft, Nexla, Nvidia, Samsung, Scotiabank, Splunk, and VMware. This work is also supported by NSF Graduate Research Fellowships and a Microsoft Ada Lovelace Research Fellowship.

References

- [1] John Aas. Project update and new name for ISRG Prio services: Introducing Divvi Up, 2021. <https://divviup.org/blog/prio-services-update/>.
- [2] Mustafa Al-Bassam and Sarah Meiklejohn. Contour: A practical system for binary transparency. In *DPM/CBT@ESORICS*, 2018.
- [3] Alchemy. <https://www.alchemy.com/>.
- [4] Amazon Web Services. Aws nitro enclaves. <https://aws.amazon.com/ec2/nitro/nitro-enclaves/>.
- [5] Michael P Andersen, Sam Kumar, Moustafa AbdelBaky, Gabe Fierro, John Kolb, Hyung-Sin Kim, David E Culler, and Raluca Ada Popa. WAVE: A decentralized authorization framework with transitive delegation. In *USENIX Security*, 2019.
- [6] Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, et al. Hyperledger fabric: a distributed operating system for permissioned blockchains. In *EuroSys*, 2018.
- [7] Apple and Google. Exposure notification privacy-preserving analytics (ENPA) white paper, 2021. https://covid19-static.cdn-apple.com/applications/covid19/current/static/contact-tracing/pdf/ENPA_White_Paper.pdf.
- [8] Mathieu Baudet, Avery Ching, Andrey Chursin, George Danezis, François Garillot, Zekun Li, Dahlia Malkhi, Oded Naor, Dmitri Perelman, and Alberto Sonnino. State machine replication in the libra blockchain. *The Libra Assn., Tech. Rep*, 2019.
- [9] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *STOC*, 1988.
- [10] Blockdaemon. <https://blockdaemon.com/>.
- [11] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. Lightweight techniques for private heavy hitters. In *IEEE S&P*, 2021.
- [12] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the weil pairing. In *ASIACRYPT*, 2001.
- [13] Thomas Bourgeat, Ilia Lebedev, Andrew Wright, Sizhuo Zhang, and Srinivas Devadas. MI6: Secure enclaves in a speculative out-of-order processor. In *IEEE/ACM MICRO*, 2019.
- [14] Ferdinand Brasser, Urs Müller, Alexandra Dmitrienko, Kari Kostianen, Srđjan Capkun, and Ahmad-Reza Sadeghi. Software grand exposure: SGX cache attacks are practical. In *USENIX WOOT*, 2017.
- [15] Miguel Castro, Barbara Liskov, et al. Practical byzantine fault tolerance. In *OSDI*, 1999.
- [16] Melissa Chase, Apoorva Deshpande, Esha Ghosh, and Harjasleen Malvai. Seamless: Secure end-to-end encrypted messaging with less trust. In *CCS*, 2019.
- [17] David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of cryptography*, 1(1):65–75, 1988.
- [18] Guoxing Chen, Sanchuan Chen, Yuan Xiao, Yinqian Zhang, Zhiqiang Lin, and Ten H Lai. SGXPECTRE: Stealing Intel secrets from SGX enclaves via speculative execution. In *IEEE EuroS&P*, 2019.
- [19] Henry Corrigan-Gibbs and Dan Boneh. Prio: Private, robust, and scalable computation of aggregate statistics. In *NSDI*, 2017.
- [20] Henry Corrigan-Gibbs, Dan Boneh, and David Mazières. Riposte: An anonymous messaging system handling millions of users. In *IEEE S&P*, 2015.
- [21] Victor Costan, Ilia Lebedev, and Srinivas Devadas. Sanctum: Minimal hardware extensions for strong software isolation. In *USENIX Security*, 2016.
- [22] Emma Dauterman, Henry Corrigan-Gibbs, and David Mazières. SafetyPin: Encrypted backups with Human-Memorable secrets. In *OSDI*, 2020.
- [23] Emma Dauterman, Eric Feng, Ellen Luo, Raluca Ada Popa, and Ion Stoica. DORY: An encrypted search system with distributed trust. In *OSDI*, 2020.
- [24] Emma Dauterman, Mayank Rathee, Raluca Ada Popa, and Ion Stoica. Waldo: A private time-series database from function secret sharing. In *IEEE S&P*, 2022.
- [25] Sisi Duan, Sean Peisert, and Karl N Levitt. hbft: speculative byzantine fault tolerance with minimum cost. *IEEE Transactions on Dependable and Secure Computing*, 2014.
- [26] Emscripten. <https://emscripten.org>.
- [27] Steven Englehardt. Next steps in privacy-preserving telemetry with Prio, 2019. <https://blog.mozilla.org/security/2019/06/06/next-steps-in-privacy-preserving-telemetry-with-prio/>.
- [28] Saba Eskandarian, Henry Corrigan-Gibbs, Matei Zaharia, and Dan Boneh. Express: Lowering the cost of metadata-hiding communication with cryptographic privacy. In *USENIX Security*, 2021.
- [29] Figment. <https://www.figment.io/>.
- [30] Fireblocks. <https://www.fireblocks.com/platforms/mpc-wallet/>.
- [31] Gemini. Cold storage, keys & crypto: How Gemini keeps assets safe. <https://www.gemini.com/blog/cold-storage-keys-crypto-how-gemini-keeps-assets-safe>.
- [32] Tim Geoghegan. Exposure notifications private analytics: Lessons learned from running secure MPC at scale, 2022. <https://divviup.org/blog/lessons-from-running-mpc-at-scale/>.
- [33] Tim Geoghegan, Mariana Raykova, and Frederic Jacobs. Exposure notifications private analytics. In *Real World Crypto*, 2022.
- [34] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game, or a completeness theorem for protocols with honest majority. In *STOC*, 1987.
- [35] Daniel Gruss, Moritz Lipp, Michael Schwarz, Daniel Genkin, Jonas Juffinger, Sioli O’Connell, Wolfgang Schoecl, and Yuval Yarom. Another flip in the wall of rowhammer defenses. In *IEEE S&P*, 2018.
- [36] Rachid Guerraoui, Nikola Knežević, Vivien Quéma, and Marko Vukolić. The next 700 BFT protocols. In *EuroSys*, 2010.
- [37] Guy Golan Gueta, Ittai Abraham, Shelly Grossman, Dahlia Malkhi, Benny Pinkas, Michael Reiter, Dragos-Adrian Seredinschi, Orr Tamir, and Alin Tomescu. Sbf: a scalable and decentralized trust infrastructure. In *IEEE DSN*, 2019.
- [38] Trinabh Gupta, Natacha Crooks, Whitney Mulhern, Srinath Setty, Lorenzo Alvisi, and Michael Walfish. Scalable and private media consumption with Popcorn. In *NSDI*, 2016.
- [39] Marcus Hähnel, Weidong Cui, and Marcus Peinado. High-resolution side channels for untrusted operating systems. In *USENIX ATC*, 2017.
- [40] Benjamin Hof and Georg Carle. Software distribution transparency and auditability. *arXiv preprint arXiv:1711.07278*, 2017.
- [41] Yuncong Hu, Kian Hooshmand, Harika Kalidhindi, Seung Jin Yang, and Raluca Ada Popa. Merkle²: A low-latency transparency log system. In *IEEE S&P*, 2021.
- [42] ISRG. Introducing ISRG Prio services for privacy respecting metrics. <https://www.abetterinternet.org/post/introducing-prio-services/>.
- [43] Abhinav Jangda, Bobby Powers, Emery D Berger, and Arjun Guha. Not so fast: Analyzing the performance of WebAssembly vs. native code. In *USENIX ATC*, 2019.
- [44] Eric Kinnear, Patrick McManus, Tommy Pauly, and Christopher A Wood. Oblivious DNS over HTTPS. *Internet Engineering Task Force, Internet-Draft*, 2021. <https://www.ietf.org/archive/id/draft-pauly-dprive-oblivious-doh-04.html>.
- [45] Knox. Knox custody. <https://www.knoxcustody.com/security>.
- [46] Dmitry Kogan and Henry Corrigan-Gibbs. Private blocklist lookups with checklist. In *USENIX Security*, 2021.

- [47] Ramakrishna Kotla, Lorenzo Alvisi, Mike Dahlin, Allen Clement, and Edmund Wong. Zyzzyva: Speculative byzantine fault tolerance. *ACM TOCS*, 2010.
- [48] Albert Kwon, Henry Corrigan-Gibbs, Srinivas Devadas, and Bryan Ford. Atom: Horizontally scaling strong anonymity. In *SOSP*, 2017.
- [49] Albert Hyukjae Kwon, David Lazar, Srinivas Devadas, and Bryan Ford. Riffle: An efficient communication system with strong anonymity. In *PoPETs*, 2016.
- [50] Leslie Lamport. Byzantizing Paxos by refinement. In *International symposium on distributed computing*. Springer, 2011.
- [51] Adam Langley, Emilia Kasper, and Ben Laurie. Certificate transparency. *Internet Engineering Task Force*, 2013. <https://tools.ietf.org/html/rfc6962>.
- [52] Dayeol Lee, Dongha Jung, Ian T Fang, Chia-Che Tsai, and Raluca Ada Popa. An off-chip attack on hardware enclaves via the memory bus. In *USENIX Security*, 2020.
- [53] Dayeol Lee, David Kohlbrenner, Shweta Shinde, Krste Asanović, and Dawn Song. Keystone: An open framework for architecting trusted execution environments. In *EuroSys*. ACM, 2020.
- [54] Sangho Lee, Ming-Wei Shih, Prasun Gera, Taesoo Kim, Hyesoon Kim, and Marcus Peinado. Inferring fine-grained control flow inside SGX enclaves with branch shadowing. In *USENIX Security*, 2017.
- [55] libBLS: a C++ library for BLS threshold signatures. <https://github.com/skalenetwork/libBLS>.
- [56] Dahlia Malkhi, Kartik Nayak, and Ling Ren. Flexible byzantine fault tolerance. In *SIGSAC*, 2019.
- [57] Frank McKeen, Ilya Alexandrovich, Alex Berenzon, Carlos V Rozas, Hisham Shafi, Vedvyas Shanbhogue, and Uday R Savagaonkar. Innovative instructions and software model for isolated execution. 2013.
- [58] Marcela S Melara, Aaron Blankstein, Joseph Bonneau, Edward W Felten, and Michael J Freedman. CONIKS: Bringing key transparency to end users. In *USENIX Security*, 2015.
- [59] Ahmad Moghimi, Gorka Irazoqui, and Thomas Eisenbarth. Cachezoom: How SGX amplifies the power of cache attacks. In *CHES*, 2017.
- [60] Graham Mudd. Privacy-enhancing technologies and building for the future, 2022. <https://www.facebook.com/business/news/building-for-the-future>.
- [61] Kirill Nikitin, Eleftherios Kokoris-Kogias, Philipp Jovanovic, Nicolas Gailly, Linus Gasser, Ismail Khoffi, Justin Cappos, and Bryan Ford. CHAINIAC: Proactive Software-Update transparency via collectively signed skipchains and verified builds. In *USENIX Security*, 2017.
- [62] Node.js. <https://nodejs.org>.
- [63] Paxos. <https://paxos.com/crypto-brokerage/>.
- [64] Rishabh Poddar, Sukrit Kalra, Avishay Yanai, Ryan Deng, Raluca Ada Popa, and Joseph M Hellerstein. Senate: A maliciously-secure MPC platform for collaborative analytics. In *USENIX Security*, 2021.
- [65] Hany Ragab, Alyssa Milburn, Kaveh Razavi, Herbert Bos, and Cristiano Giuffrida. Crosstalk: Speculative data leaks across cores are real. In *IEEE S&P*, 2021.
- [66] James Reyes. Building the next generation of digital advertising with mpc. In *Real World Crypto*, 2022. <https://iacr.org/submit/files/slides/2022/rwc/rwc2022/104/slides.pdf>.
- [67] Riddle and code. <https://www.riddleandcode.com/blog-posts/hardware-security-modules-vs-secure-multi-party-computation-in-digital-asset-custody>.
- [68] Michael Schwarz, Moritz Lipp, Daniel Moghimi, Jo Van Bulck, Julian Stecklina, Thomas Prescher, and Daniel Gruss. ZombieLoad: Cross-privilege-boundary data sampling. In *CCS*, 2019.
- [69] Michael Schwarz, Samuel Weiser, Daniel Gruss, Clémentine Maurice, and Stefan Mangard. Malware guard extension: Using SGX to conceal cache attacks. In *DIMVA*. Springer, 2017.
- [70] Sepior. <https://sepior.com/products/advanced-mpc-wallet>.
- [71] Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.
- [72] Sudheesh Singanamalla, Suphanat Chunhapananya, Marek Vavruša, Tanya Verma, Peter Wu, Marwan Fayed, Kurtis Heimerl, Nick Sullivan, and Christopher Wood. Oblivious DNS over HTTPS (ODOH): A practical privacy enhancement to DNS. *PoPETs*, 2021.
- [73] Alin Tomescu, Vivek Bhupatiraju, Dimitrios Papadopoulos, Charalampos Papamanthou, Nikos Triandopoulos, and Srinivas Devadas. Transparency logs via append-only authenticated dictionaries. In *CCS*, 2019.
- [74] Trillian. <https://github.com/google/trillian>.
- [75] Unbound Security. The Unbound CORE MPC key vault.
- [76] Jo Van Bulck, Marina Minkin, Ofir Weiss, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Thomas F Wenisch, Yuval Yarom, and Raoul Strackx. Foreshadow: Extracting the keys to the Intel SGX kingdom with transient out-of-order execution. In *USENIX Security*, 2018.
- [77] Jo Van Bulck, Daniel Moghimi, Michael Schwarz, Moritz Lipp, Marina Minkin, Daniel Genkin, Yarom Yuval, Berk Sunar, Daniel Gruss, and Frank Piessens. LVI: Hijacking Transient Execution through Microarchitectural Load Value Injection. In *IEEE S&P*, 2020.
- [78] Jo Van Bulck, Nico Weichbrodt, Rüdiger Kapitza, Frank Piessens, and Raoul Strackx. Telling your secrets without page faults: Stealthy page table-based attacks on enclaved execution. In *USENIX Security*, 2017.
- [79] Stephan Van Schaik, Alyssa Milburn, Sebastian Österlund, Pietro Frigo, Giorgi Maisuradze, Kaveh Razavi, Herbert Bos, and Cristiano Giuffrida. RIDL: Rogue in-flight data load. In *IEEE S&P*, 2019.
- [80] Stephan van Schaik, Marina Minkin, Andrew Kwong, Daniel Genkin, and Yuval Yarom. CacheOut: Leaking data on Intel CPUs via cache evictions. In *IEEE S&P*, 2021.
- [81] Tanya Verma and Sudheesh Singanamalla. Improving DNS privacy with oblivious DoH in 1.1.1.1, 2020. <https://blog.cloudflare.com/oblivious-dns/>.
- [82] Frank Wang, Catherine Yun, Shafi Goldwasser, Vinod Vaikuntanathan, and Matei Zaharia. Splinter: Practical private queries on public data. In *NSDI*, 2017.
- [83] WebAssembly. <https://webassembly.org>.
- [84] David Isaac Wolinsky, Henry Corrigan-Gibbs, Bryan Ford, and Aaron Johnson. Dissent in numbers: Making strong anonymity scale. In *OSDI*, 2012.
- [85] Yuanzhong Xu, Weidong Cui, and Marcus Peinado. Controlled-channel attacks: Deterministic side channels for untrusted operating systems. In *IEEE S&P*, 2015.
- [86] Andrew C Yao. Protocols for secure computations. In *SFCS*, 1982.
- [87] Bennet Yee, David Sehr, Gregory Dardyk, J Bradley Chen, Robert Muth, Tavis Ormandy, Shiki Okasaka, Neha Narula, and Nicholas Fullagar. Native client: A sandbox for portable, untrusted x86 native code. In *IEEE S&P*, 2009.
- [88] Maofan Yin, Dahlia Malkhi, Michael K Reiter, Guy Golan Gueta, and Ittai Abraham. Hotstuff: BFT consensus with linearity and responsiveness. In *ACM PODC*, 2019.