ISRG Internet Security Research Group

Let's change the *world*
or at least the Internet.

2021 ANNUAL REPORT

A home for critical infrastructure

INTERNET SECURITY RESEARCH GROUP

Internet Security Research Group (ISRG) is the nonprofit behind Let's Encrypt. Since 2013, we've built and fostered Let's Encrypt to be the world's largest Certificate Authority.

This year, we took two major steps forward as an organization: launching two new projects in addition to Let's Encrypt—Prossimo, a project focused on improving memory safety for the Internet's critical software, and Divvi Up, a service to provide privacy-preserving metrics at scale.

ISRG Internet Security Research Group

Contents

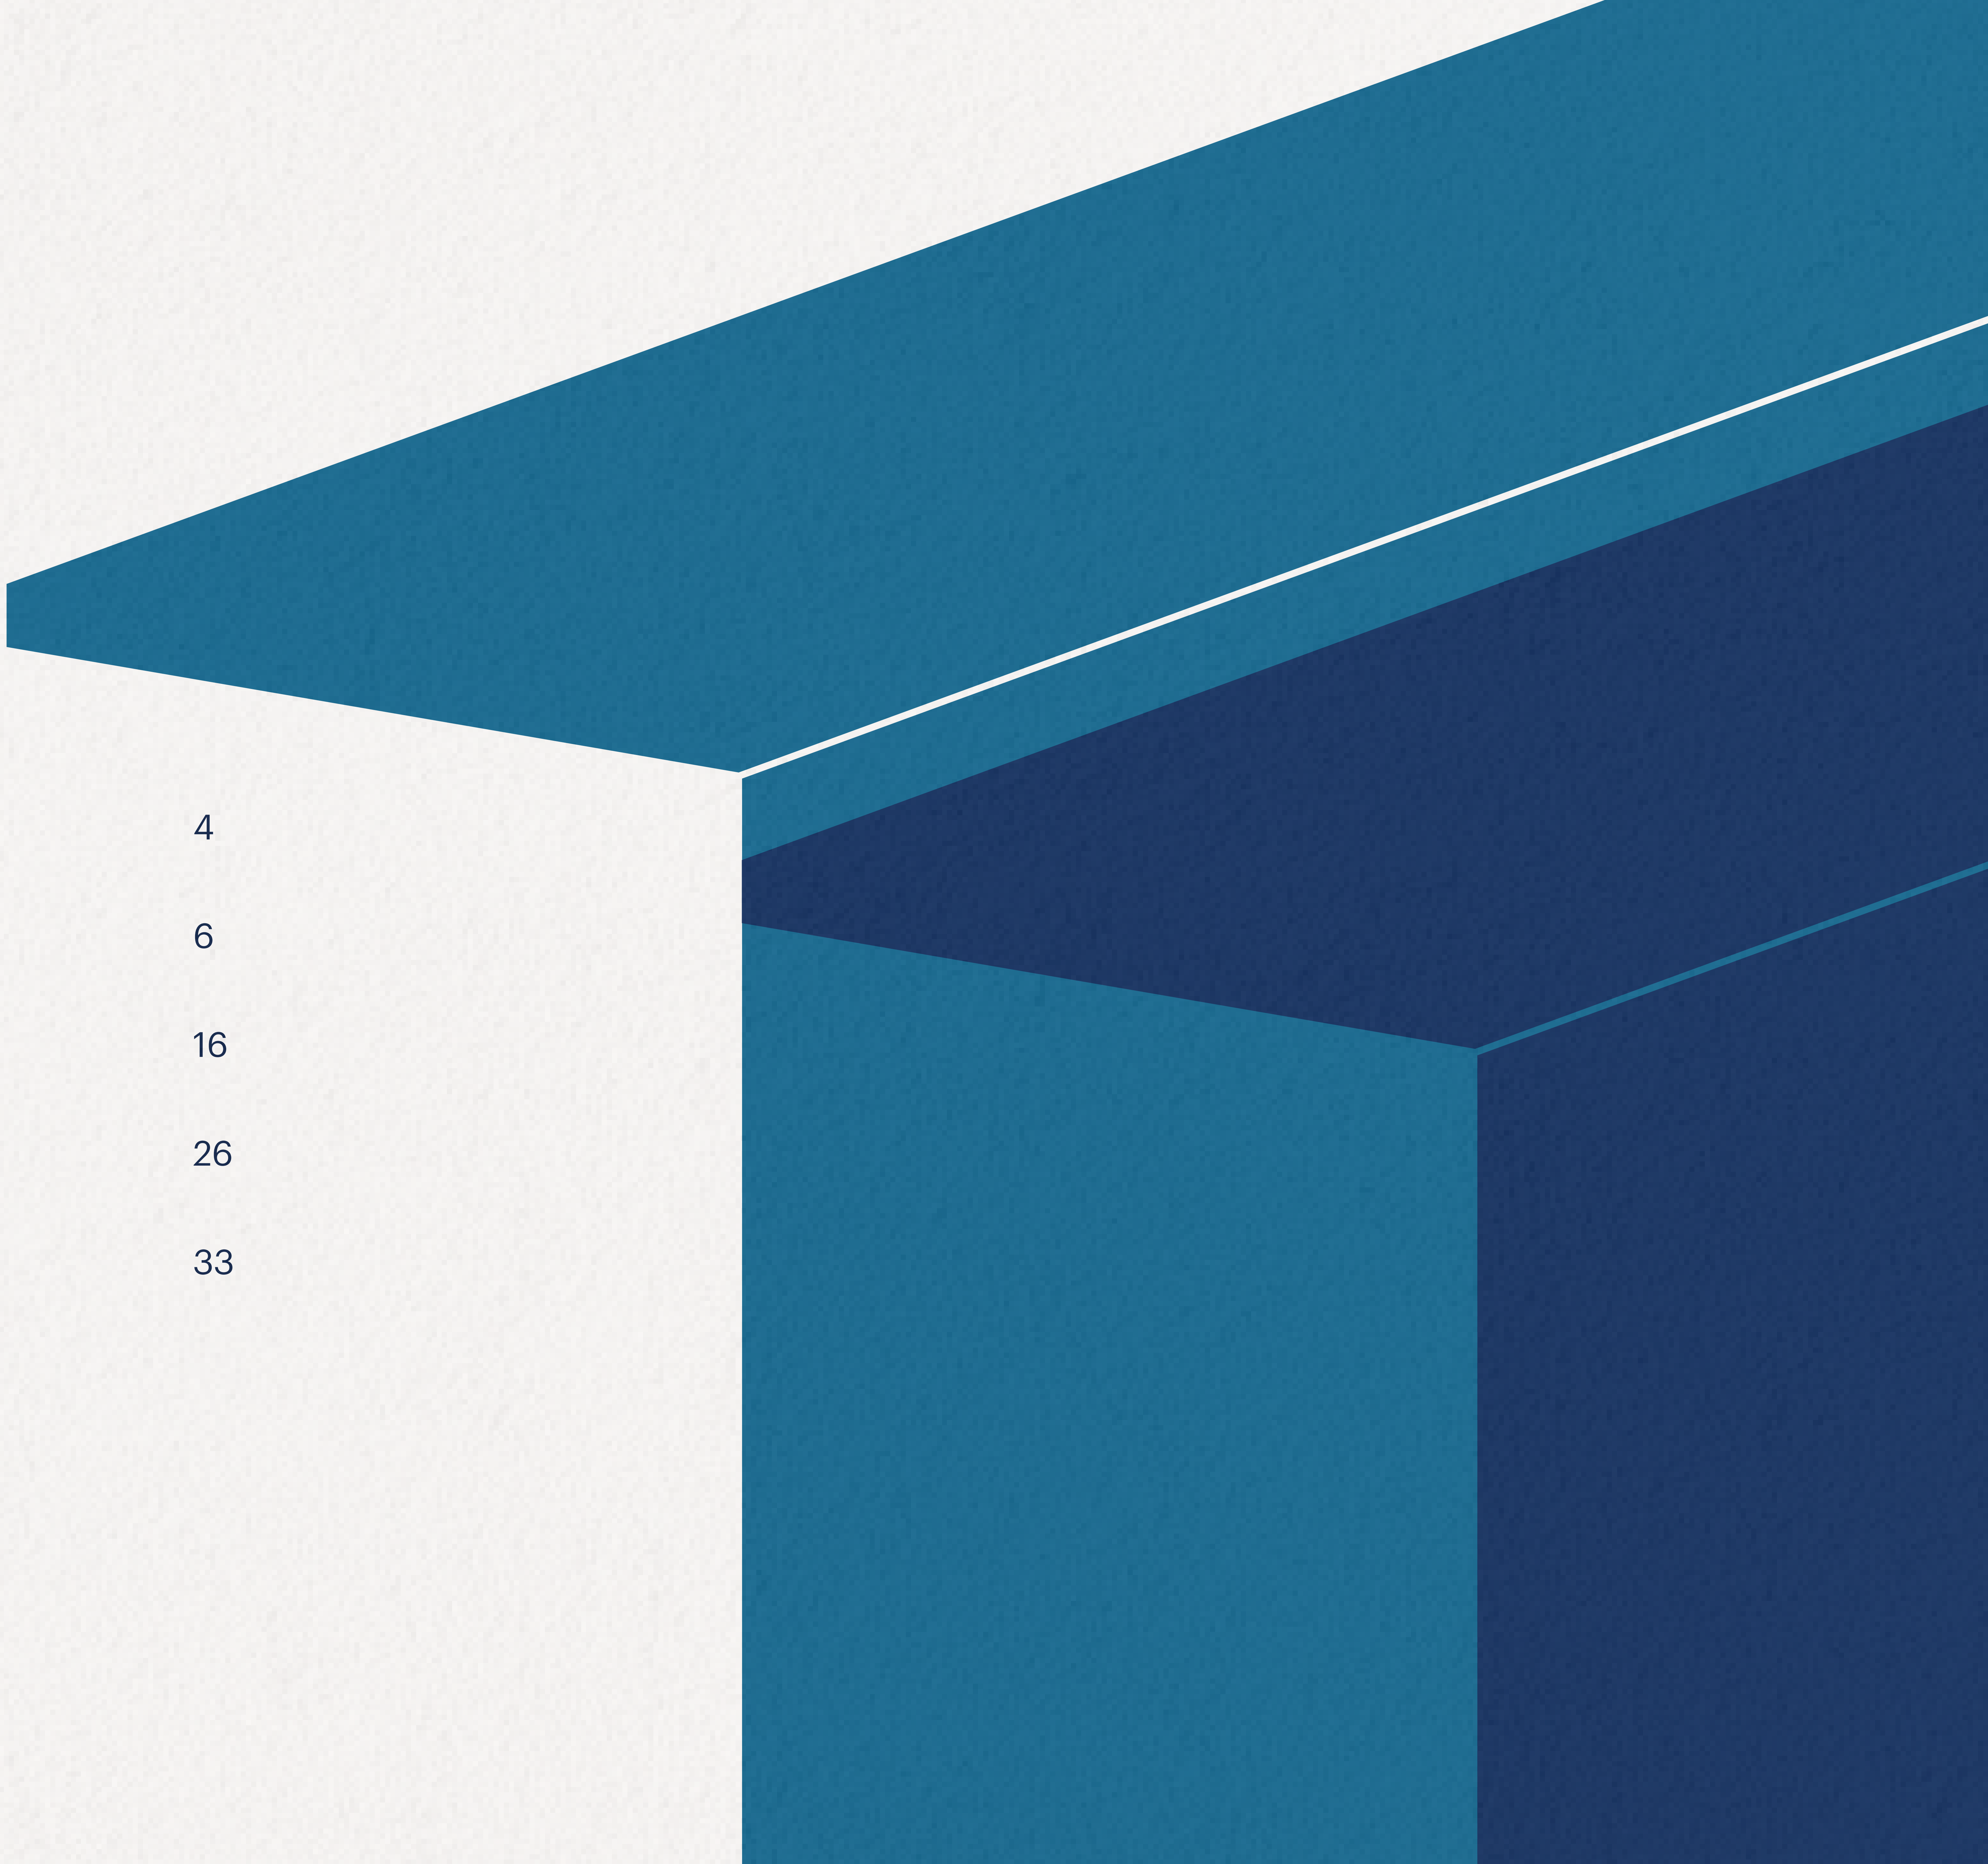

JOSH AAS
EXECUTIVE DIRECTOR

From ideas to implementations

A NOTE FROM OUR EXECUTIVE DIRECTOR

We can do a lot to improve security and privacy on the Internet by taking existing ideas and applying them in ways that benefit the general public at scale. Our work certainly does involve some research, as our name implies, but the success we've had in pursuing our mission largely comes from our ability to go from ideas to implementations that improve the lives of billions of people around the world...

Our first major project, Let's Encrypt, now helps to protect more than 260 million websites by offering free and fully automated TLS certificate issuance and management. Since it launched in 2015, encrypted page loads have gone from under 40% to 92% in the U.S. and 83% globally.

We didn't invent Certificate Authorities. We didn't invent automated issuance and management. We refined those ideas and applied them in ways that benefit the general public at scale.

We launched our Prossimo project in late 2020. We hope that this project will greatly improve security and privacy on the Internet by making memory safety vulnerabilities in the Internet's most critical software a thing of the past. We're bringing a healthy dose of ambition to the table and we're backing it up with effective strategies and strong partnerships.

Again, we didn't invent any memory-safe languages or techniques, and we certainly didn't invent memory safety itself. We're simply taking existing ideas and applying them in ways that benefit the general public at scale. We're getting the work done.

With our latest project, Divvi Up, the core ideas are a bit newer than the ideas behind our other projects, but we didn't invent them either. Over the past decade or so some bright people have come up with a way to resolve the tension between wanting to collect metrics about populations and needing to collect data about individuals.

We believe those ideas have matured enough that it's time to deploy them to the public's benefit. We started by building and deploying a privacy-preserving metrics service for COVID-19 Exposure Notification applications in late 2020, in partnership with Apple, Google, the Bill & Melinda Gates Foundation and the Linux Foundation. We're expanding that service so any application can collect metrics in a privacy-preserving way.

Being ready to bring ideas to life means a few different things.

We need to have an excellent engineering team that knows how to build services at scale. It's not enough to just build something that works - the quality and reliability of our work needs to inspire confidence. People need to be able to rely on us.

We also need to have the experience, perspective, and capacity to effectively consider ideas. We are not an organization that "throws things at the wall to see what sticks." Between our staff, our board of directors, our partners, and our community, we're able to do a great job evaluating opportunities to understand technical feasibility, potential impact, and alignment with our public benefit mission—to reduce financial, technological, and educational barriers to secure communication over the Internet.

Administrative and communications capabilities are essential. From fundraising and accounting to legal and social media, our administrative teams exist in order to support and amplify the critical work that we do. We're proud to run a financially efficient organization that provides services for billions of people on only a few million dollars each year.

Finally, it means having the financial resources we need to function. As a nonprofit, 100% of our funding comes from charitable contributions from people like you and organizations around the world. But global impact doesn't necessarily require million dollar checks: since 2015 tens of thousands of people have given to our work. They've made a case for corporate sponsorship, given through their donor-advised funds, or set up recurring donations, sometimes to give $3 a month. That's all added up to $17M that we've used to change the Internet for nearly everyone using it. I hope you'll join these people and support us financially if you can.

JOSH AAS
EXECUTIVE DIRECTOR

Ambition as big as the Internet.

MOVING THE WEB TO 100% ENCRYPTION

This year saw Let's Encrypt issue 2.5 million certificates every day, on average. Our ability to issue at this scale, routinely and reliably, helped push the percentage of encrypted page loads to its highest level ever.

Let's Encrypt

Let's Encrypt

IMPACT ON THE WEB

As of November 1, 2021, Let's Encrypt serves 266 million websites with 211 million active certificates. Since 2013, HTTPS page loads have grown from 25% to 84% globally; it's close to 92% in the United States.

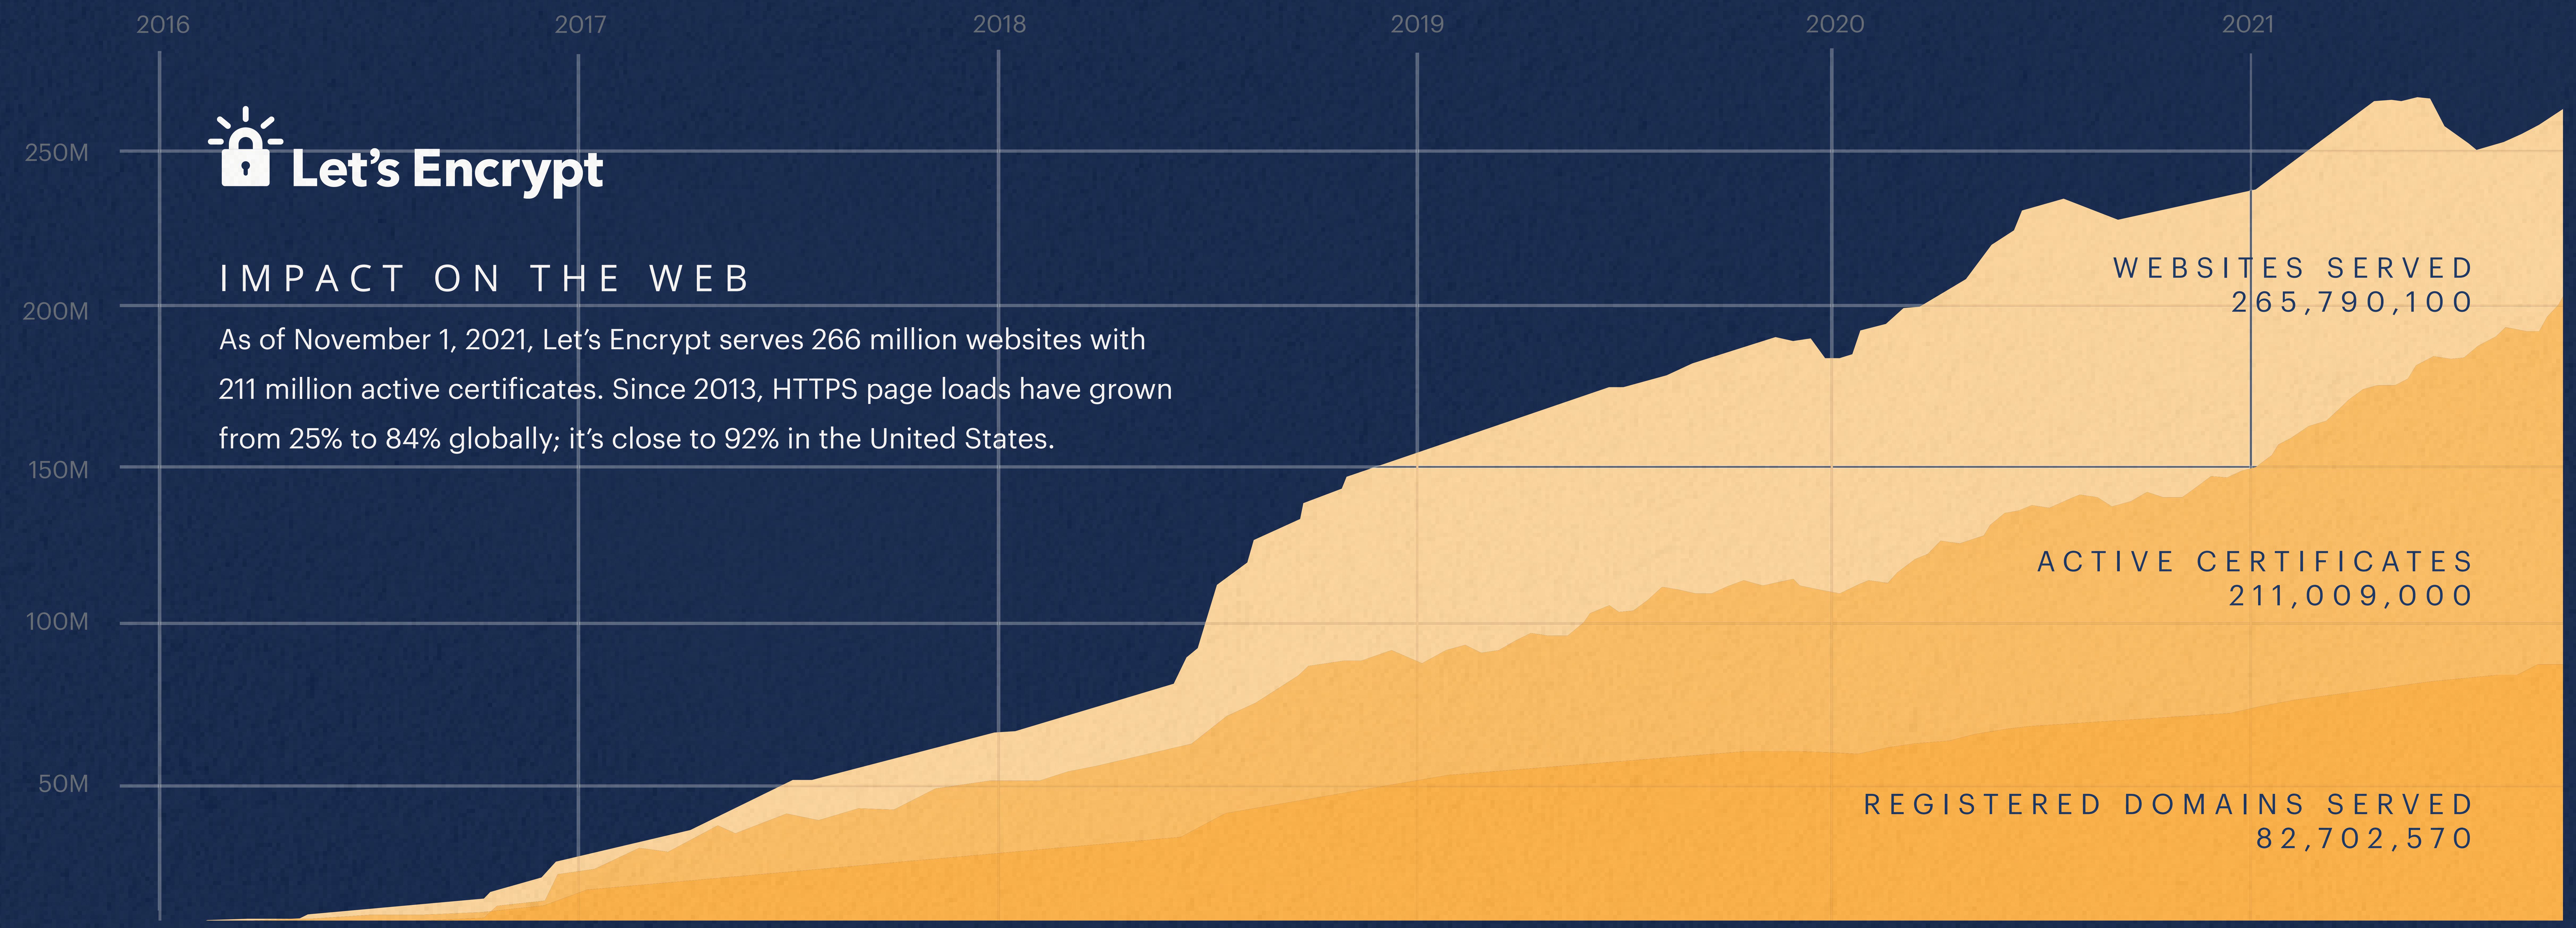

WEBSITES SERVED
265,790,100

ACTIVE CERTIFICATES
211,009,000

REGISTERED DOMAINS SERVED
82,702,570

2021 DAILY ISSUANCE

MIN	AVG	MAX
599K	2.1M	3.2M

TOTAL CERTS ISSUED

2,119,830,000

% OF WEB PAGES LOADED BY FIREFOX USING HTTPS, GLOBALLY

2013	2017	2021
25%	67%	83%

> "LET'S ENCRYPT IS A TRULY SPECIAL NONPROFIT. ITS TLS CERTIFICATES HAVE MADE HTTPS SIMPLE FOR WEBSITE OWNERS EVERYWHERE."

DAVE TEARE
FOUNDER | 1PASSWORD

The journey to two billion certs

THE ORIGIN STORY OF LET'S ENCRYPT

When HTTPS was first introduced by Netscape in the mid 1990s, less than 3% of the world's population had access to the Web. By 2015, 43% of the world's population were on the Internet.

Yet as the world came online, security and privacy for billions of people using the Internet lagged behind. As of 2015, 55-70% of browser page loads used plaintext HTTP. A lack of HTTPS adoption meant that for people using the Web for more and more aspects of their lives, the threat of hostile networks—including mass surveillance and censorship by governments, consumer profiling and ad injections by ISPs, and insertion of malicious code by network devices—was overwhelmingly prevalent.

WHY THE LACK OF HTTPS?

Adopting HTTPS was difficult, even for large, technically astute organizations: obtaining a certificate—we'll just say "cert" from here on out—was complicated and error prone. It took days to provision just one cert. And once you had a cert up and running, it'd be valid for one or more years. So by the time you had to renew it, it's likely most technologists would have to try to remember the process all over again.

Beyond this complexity, certs were expensive. In 2015, the average price for a one-year single-domain cert from the five largest Certificate Authorities (CAs) was $178—a wildcard cert cost $766.

A complex, error-prone process—one that required most engineers to have to talk to procurement—meant that the professional benefits of not adopting TLS greatly outweighed the security benefits of adopting TLS.

SHIFTING THE PARADIGM THROUGH PARTNERSHIP.

What eventually became Let's Encrypt was the merging of two simultaneous efforts to build a fully automated Certificate Authority.

In 2012, a team headed up by Alex Halderman at the University of Michigan and Peter Eckersley at Electronic Frontier Foundation (EFF) was working on a protocol for automatically issuing and renewing certificates. Concurrently, a team at Mozilla led by Josh Aas and Eric Rescorla was working on creating a free and automated certificate authority. The groups learned of each other's efforts and joined forces in May 2013 to create Internet Security Research Group, the nonprofit behind Let's Encrypt.

At first glance, that may not be surprising. But consider the stakes and ambitions of what these groups were attempting:

First, they wanted to overhaul a paradigm that had been the norm for nearly two decades: TLS was expensive and to get it requires very specific knowledge.

More importantly, they wanted to fundamentally disrupt how the Internet works for most people using it, without negatively affecting the way people experienced it.

> "CREATING A NEW KIND OF CERTIFICATE AUTHORITY THAT GIVES OUT FREE CERTIFICATES WAS A CRAZY IDEA...WE HAD TO PROVE THAT THE ECONOMICS WOULD WORK, AND THERE WAS NO WAY TO DO THAT EXCEPT TO JUST BUILD IT."
>
> **ALEX HALDERMAN**
> ISRG BOARD OF DIRECTORS

Leading the way to a more secure Web

LEADING MEANS EVERY DETAIL COUNTS

Let's Encrypt is the largest Certificate Authority (CA) by issuance. It has more currently active certificates than all other browser-trusted CAs combined. Being a critical service for so many organizations and people around the world means every decision and statistic matters. In 2021 that has certainly been true. Here are two stories of how we lead the way towards a more secure Web.

ONE SECOND TOO MANY

Earlier this year, it came to our attention that we were out of line, by one second, with a specific requirement related to certificate lifetime. One second may not seem like much, but leading means every detail counts.

Let’s Encrypt is well known for issuing certificates that are valid for only 90 days. From the beginning, our certificates have been given a 90 day validity period by our CA software by taking the issuance time and adding exactly 2,160 hours to yield the certificate’s “not after” date. However, RFC 5280 defines the validity period of a certificate as being the duration between the “not before” and the “not after” timestamps, *inclusive*. This inclusivity means that Let’s Encrypt’s certificates had all actually been valid for 90 days plus 1 second.

In other words, a hypothetical certificate with a “not before” date of 9 June 2021 at 03:42:01 and a “not after” date of 7 Sept 2021 at 03:42:01 becomes valid at the beginning of the :01 second, and only becomes invalid at the :02 second, a period that is 90 days plus 1 second. The 90-day “not after” time must actually be 03:42:00.

For the lifespan of Let’s Encrypt, we’ve always reasoned about certificate lifespans in terms of hours. Unfortunately, the RFC 5280 definition either requires the CA software to explicitly subtract 1 second from calculated validity periods to account for the inclusivity, or requires the configuration to be defined in seconds rather than the easier-to-analyze hours that we had always relied upon.

This error was not caught by any form of automated certificate linting. Our issuance pipeline uses the ZLint certificate linter project as a mandatory, must-pass step at two different stages of our issuance pipeline: after construction of the “precertificate” for Certificate Transparency logs, and after issuance of the final certificate but before delivery to a subscriber.

Because this error impacted our issuance from the very beginning, it meant it technically impacted every active certificate issued by Let’s Encrypt. At the time, that was 185 million certificates. And although the impact was ubiquitous, we did not stop issuance nor did we revoke any certificates. We deemed either to be an overly aggressive correction given the severity of the issue and the speed at which our team was able to deploy a fix.

TO EVERY ROOT THERE IS A SEASON

In September of this year, the Root certificate originally used by Let's Encrypt expired. Although root expirations are a reality for every CA, our track record of collaboration and innovation made this expiration imperceptible for most users of the Internet around the world.

Since the beginning of Let’s Encrypt, partnering with the CA IdenTrust has been critical. As early as 2013, IdenTrust engineers assisted with the development of Let’s Encrypt, even before ISRG had any full-time staff. In 2014, IdenTrust entered into a long-term agreement with ISRG, creating a cross-signature between ISRG and IdenTrust’s “DST Root CA X3.” This allowed certificates issued by Let’s Encrypt to use the IdenTrust root, making Let’s Encrypt certificates usable for a lot of people, right away, while we submitted our own root “ISRG Root X1” to be trusted by the major software platforms.

The IdenTrust root used by Let’s Encrypt had a 20 year lifespan—from September 30, 2000 to September 30, 2021. All roots trusted by root programs have a lifespan. Eventually, they all expire. That meant Let’s Encrypt needed to work on getting its own root widely trusted before the IdenTrust Root expiration. In 2020, we transitioned issuance from this root to our own.

However, transitioning to our own root did introduce some compatibility problems. Our root isn’t trusted by software that hasn’t been updated since 2016, roughly around the time our root was accepted to many root programs. For many people around the world using old Android devices, this transition could cut off their access to the Internet. In 2020, that could have been as many as 33.8% of Android users.

In late 2020, thanks to innovative thinking from our community and our longtime partners, IdenTrust, we developed a way for older Android devices to retain their ability to visit sites that use Let’s Encrypt certificates after our cross-signed intermediates expired.

We’re proud of the work so many people in our community, staff, and board did to make this transition as smooth as it was. Looking ahead, one of Let's Encrypt's root certificates was issued in 2015, and a second one was issued in 2020. The next expiration for these root certificates will be in 15 years and 20 years, respectively. Who knows, by then our certs might be regularly encrypting communication between the Earth and Mars—we’ll be ready.

> “
> **IN THE END, SOMETHING A LITTLE UNEXPECTED HAS HAPPENED WHICH MIGHT JUST REDUCE THE SERIOUS IMPACT OF THIS EVENT AND MAKE IT A LITTLE MORE PALATABLE. BECAUSE OLD ANDROID DEVICES DON'T CHECK THE EXPIRATION DATE OF A ROOT CERTIFICATE WHEN THEY USE IT, LET'S ENCRYPT MAY BE ABLE TO CONTINUE TO CHAIN DOWN TO THE EXPIRED ROOT CERTIFICATE WITHOUT ANY PROBLEM ON THOSE OLDER DEVICES.”**
>
> **SCOTT HELME**
> SECURITY RESEARCHER

What's next for Let's Encrypt

THE TEAMS PUSHING LET'S ENCRYPT FORWARD

Two teams work together to keep Let’s Encrypt running as the best CA it can be: our Software Engineering team, primarily responsible for maintaining and improving the software behind Let’s Encrypt—Boulder—and the Site Reliability Engineering (SRE) team.

STAYING SHARP: ISRG'S BOOK CLUB

We're proud to have a team of folks dedicated to serving the world with access to free TLS. Perhaps our favorite example of their commitment is Book Club: each week these teams come together to discuss learnings and findings from whatever they happen to be reading—from books about learning emerging languages, to best practices for SRE teams. We love that these teams take the time to advance their knowledge while still leading the world to a more secure and privacy-respecting Web thanks to Let's Encrypt.

The SRE team is tasked with ensuring Let's Encrypt continues to serve at scale, reliably, and stably. In 2021, they tackled work to upgrade our datacenter networking, improve our database performance, officially deprecate Automated Certificate Management Environment v1 (ACME v1), and begin issuance using Elliptic Curve Digital Signature Algorithm (ECDSA). And, they were the team on deck for helping monitor anything related to the root cert transition.

In 2021, the Boulder team moved ahead on multiple fronts. Perhaps most exciting is the work related to ACME Renewal Information (ARI). The plan with ARI is to design and implement a renewal information API as an extension to ACME. ARI seeks to address two issues that affect both ACME and the wider Web PKI.

The first issue is how a Certificate Authority (CA) should inform its subscribers, or a third party, of a CA-initiated certificate revocation event. Today, this is accomplished via email, or other out-of-band notification channels. For subscribers who manually manage their certificates, this method may work fine. However, Subscribers who rely upon automated ACME clients are unlikely to receive this notification in a timely manner. If a subscriber does not act, they may have a revoked certificate until their next renewal.

The second issue is how ACME clients should determine when to renew a regular, non-revoked certificate. Most clients take one of two routes. They either are manually configured to renew at a specific interval (i.e., via `cron` or similar), or parse the issued certificate to determine the expiration date and choose some date preceding it to attempt renewal. While the latter is better, each can cause issues for both the client and the issuing CA. The first option causes significant barriers for the issuing CA changing certificate lifetimes, as the static renewal window makes assumptions about that lifetime that must be manually updated. Both options can cause load clustering for the issuing CA.

As of October 2021, ARI is functional in Let's Encrypt's staging environment. For more on our work to develop ARI, check out the GitHub repo. Because of ARI's potential benefits to the wider Web PKI, we're also in progress on submitting this work to the IETF for standardization.

"Let's Encrypt makes it easy for *everyone* to do the right thing to secure the Internet. We couldn't be happier to give our support to such a great effort."

TOBI LÜTKE

FOUNDER & CEO | SHOPIFY

Introducing Divvi Up

PRIVACY-PRESERVING METRICS AT SCALE—FOR THE PUBLIC BENEFIT

Our goal with Divvi Up is to provide a service that application owners—from public benefit entities, to governments, to private companies—can use to easily collect user population metrics while respecting user privacy.

Divvi Up

Data divided. Data secured.

The privacy problem

THE CHALLENGES OF ENSURING USER PRIVACY

Many types of applications, from mobile and desktop apps to websites, generate metrics about their users. These metrics are valuable for application owners because they are the basis for insights into users’ behaviour. Normally an application would send all of its metrics back to the app owners, either directly or through a middle-entity. Users simply have to trust that owners will respect their stated privacy and security policies.

Stated policy, however, is insufficient as a privacy safeguard. Once an app owner has user data, privacy policies can be violated intentionally or accidentally. The mere possibility of privacy violations can erode trust in applications, discouraging users' engagement. In addition, the presence of personally identifiable information is a liability that more organizations are worrying about, since it can be unintentionally leaked, or stolen in a breach.

Moving commercial metrics collection to more privacy-respecting systems would be hugely beneficial for users because of the sheer amount of data that applications collect. The benefits are even greater for government and non-governmental organizations: by ensuring privacy we can build trust, which is essential when asking populations to participate in public health and safety efforts.

The solution

PRESERVING PRIVACY IS POSSIBLE

We intend to resolve much of the tension between wanting to know information about a population of users and needing to collect information about individuals that might compromise their privacy. Unlike any other system available today, the system we are designing and building has all of the following properties:

- Ability to collect valuable and insight-enabling population metrics
- Ease-of-use for both application owners and users
- Sufficiently low cost, fundamental for adoption of a system that will better protect users
- Excellent privacy protection for users: no application owner or middle-entity—including ISRG—ever possesses complete and identifiable individual user metrics ensuring that user data cannot be mishandled intentionally or unintentionally

This system is all the more valuable in repressive and otherwise difficult operating environments. By not collecting complete and identifiable user metrics, individuals need not worry about compromise by technical or legal means. That safety builds trust, which is invaluable when asking a population to participate in services that can benefit public health and safety. Our service will be built to deliver privacy-preserving metrics through either the “Prio” or “Heavy Hitters” protocols developed by Dan Boneh and Henry Corrigan-Gibbs at MIT. If you’re curious to see the details behind Privacy-Preserving Metrics (PPM), check out the full System Specification. For more on Divvi Up and how you can be one of its first users, email us at divviup@abetterinternet.org.

"

HAVING THE RIGHT DATA IS CRITICAL BUT TOO OFTEN THAT COMES AT A LOSS TO PEOPLE'S PRIVACY. PRIVACY PRESERVING METRICS TECHNOLOGY IS THE NEXT FRONTIER IN MEASUREMENT, ALLOWING US TO SAFELY AND PRIVATELY GATHER IMPORTANT METRICS. WE'RE EXCITED TO BE COLLABORATING WITH ISRG TO DEVELOP THIS IMPORTANT TECHNOLOGY AND LOOK FORWARD TO IT BEING AVAILABLE TO EVERYONE."

ERIC RESCORLA
CTO | FIREFOX

How it works

A SIMPLE SCHEME. COMPLEX MATH.

Divvi Up takes a user-generated metric, from a mobile device, web browser, or other application, and divides the metric into two encrypted shares as it leaves the origin. One half of that metric is sent to a Divvi Up server, the other to a third-party server. When an application owner queries an aggregate statistic of its users, Divvi Up combines the divided metrics from all users and recombines them into a privacy-preserving aggregate.

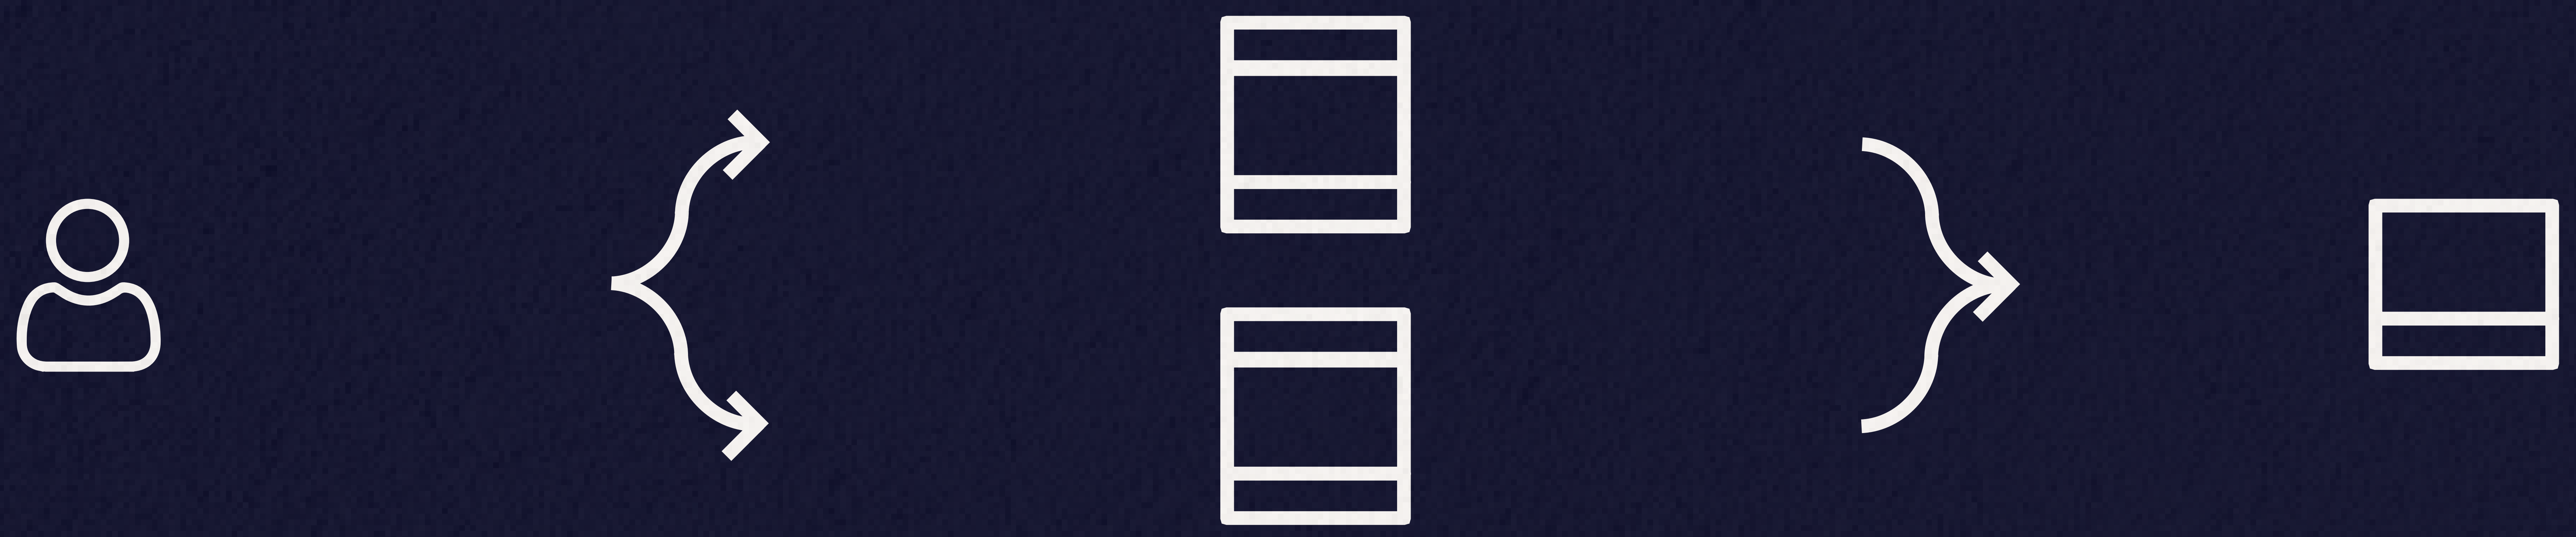

A USER-GENERATED VALUE

DIVIDE THE VALUE

TWO, NON-COLLUDING SERVERS

COMBINE AGGREGATES

ANONYMIZED INSIGHT

Privacy-respecting contact tracing

OUR ROLE IN MITIGATING THE SPREAD OF COVID-19

In early 2020, the advent of the COVID-19 pandemic brought a catalytic moment for the Internet, and those who impact it at scale, to demonstrate its benefit to serve humanity. One urgent example is the proliferation of Apple and Google's COVID-19 Exposure Notification System for contact tracing, to help control and mitigate the spread of COVID-19.

In late 2020, ISRG was approached by Apple and Google with the idea of using Prio as part of this system. ISRG received funding from The Bill & Melinda Gates Foundation and The Linux Foundation to build the first production instance of Prio, to be used as part of the Exposure Notification Private Analytics (ENPA) system. In less than 5 months, ISRG and partners took an idea that existed only on paper and built out this service to begin serving people across the US.

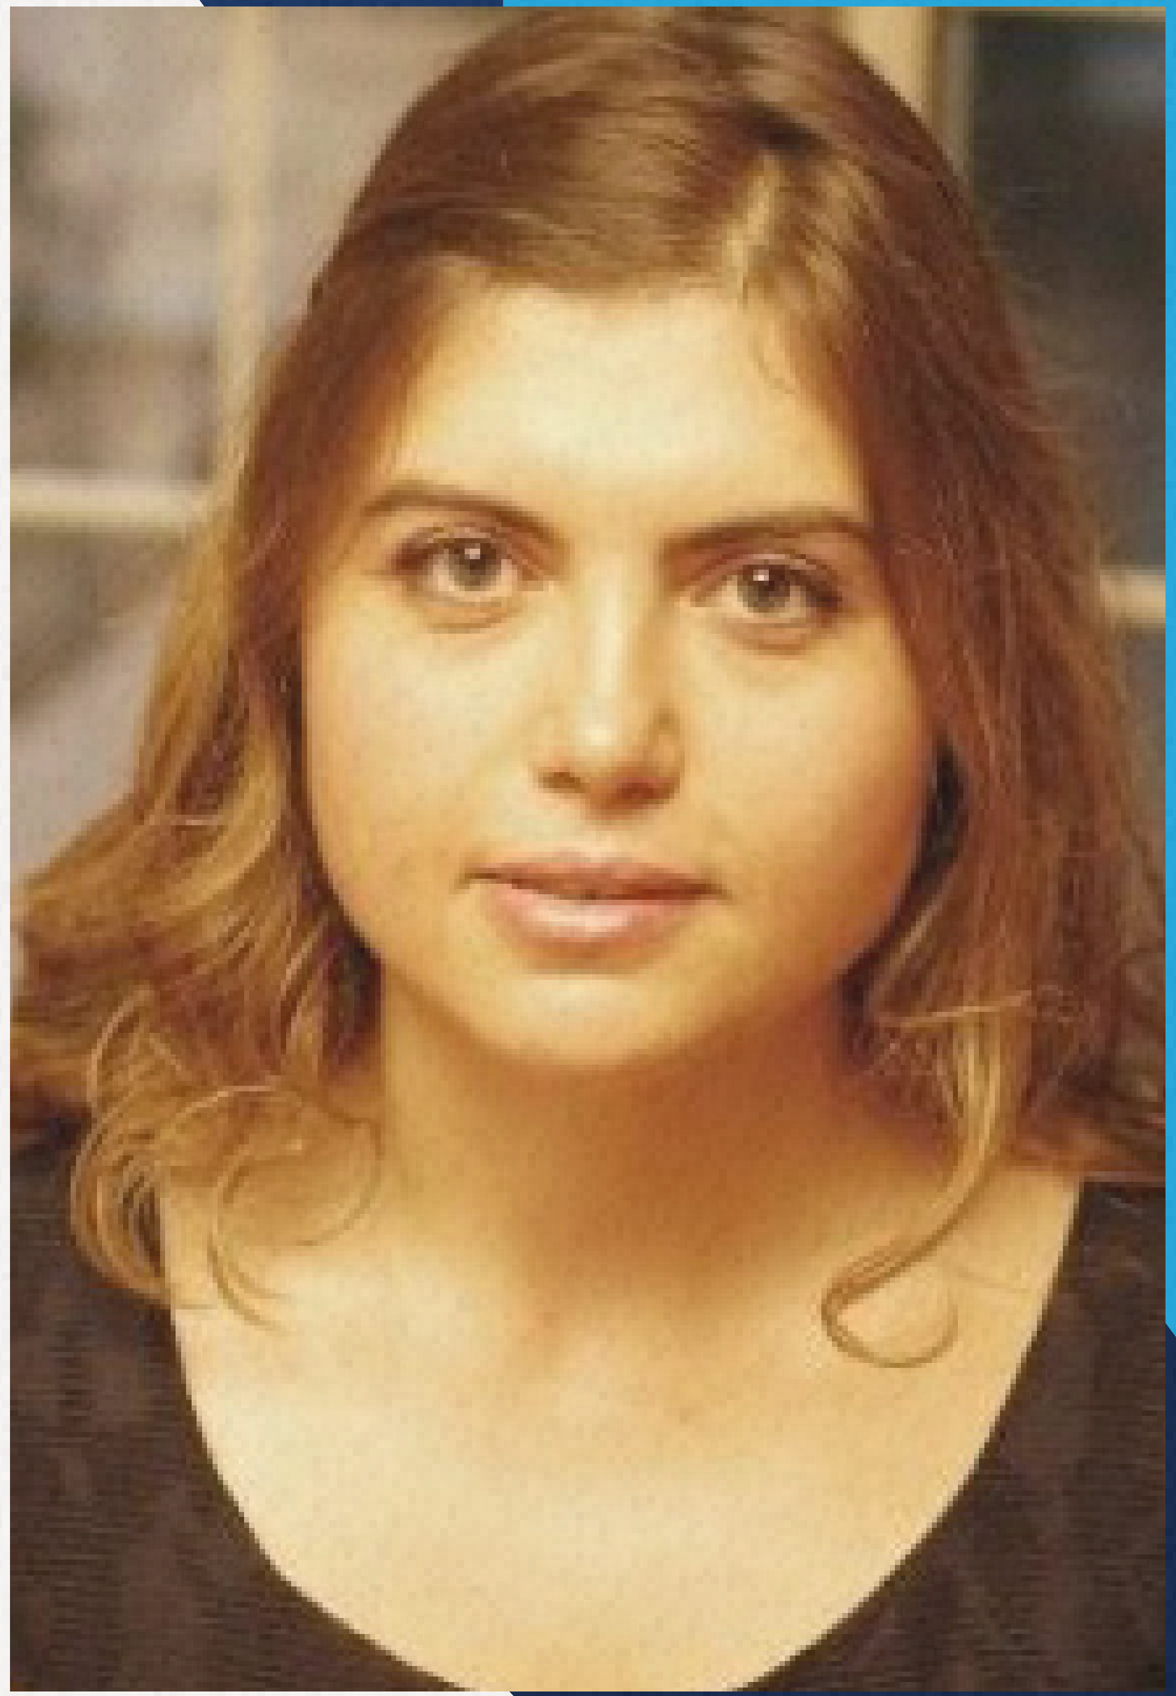

"

PRIVACY PRESERVING MEASUREMENT is an area of growing importance that develops techniques which enable evaluating the efficacy of different services and applications while protecting the privacy of their users. ISRG is developing tools that enable such functionality.

ISRG has been a valuable partner running one of the computation servers in the Exposure Notifications Private Analytics (ENPA) system which provides privacy preserving aggregate metrics to health authorities that help them assess the efficacy of EN. It is also playing a central role in an effort to define and standardize a PPM framework which will be a useful tool for many measurement application scenarios with large user bases."

MARIANA RAYKOVA
RESEARCH SCIENTIST | GOOGLE

PRESERVING PRIVACY FOR COVID-19 RESPONSE

ISRG serves as one of the data processors in the Prio-based privacy-preserving metrics system used by Apple and Google's Exposure Notifications Express (ENX). Along with our partners, we have designed a system that protects privacy while providing useful data, we have created a production-quality implementation, and we operate our service efficiently and reliably. We have been participating in this collaboration since 2020 and see this as a great example of Prio as a useful privacy tool.

There is strong evidence that effective app-based contact tracing can help slow the spread of COVID-19. The role of our system is to enable the Public Health Authorities (PHAs) administering an ENX app to collect aggregate metrics while protecting the privacy of each individual using the app. These metrics can help with the epidemiological response based on new trends in the spread of COVID-19. For example, the aggregate, anonymized data produced by our system can provide the total number of exposure notification alerts displayed to users without exposing how many alerts each individual's device has received, or whether a particular device has displayed alerts at all. This information can help PHAs better understand the effectiveness of the system and adjust the parameters they provide for ENX.

On a person's device, ENX data is divided into two shares in such a way that a person's data is not intelligible without both shares. While the data is still on the device, one share is encrypted using an encryption key from ISRG, and the other is encrypted with a key from the National Institutes of Health (NIH). The data shares are then sent to ingestion servers operated by Google and Apple. The ingestion servers cannot decrypt the data shares, but they can do device authenticity verification and load balancing. Once these functions are performed, the data shares are then passed on to ISRG and NIH. Once we get our share, we sum it into a partial aggregate sum. NIH does the same with their share. The ISRG and NIH partial aggregate sums are then sent to a server operated by MITRE, where they are combined into a complete set of metrics PHAs can view.

This Prio-based process ensures that individual user data is never intelligible once it leaves the user's device, yet useful aggregate metrics are provided to PHAs. Through this collaboration we are pleased to see that Prio can work seamlessly at scale; ISRG has aggregated over two billion device metrics.

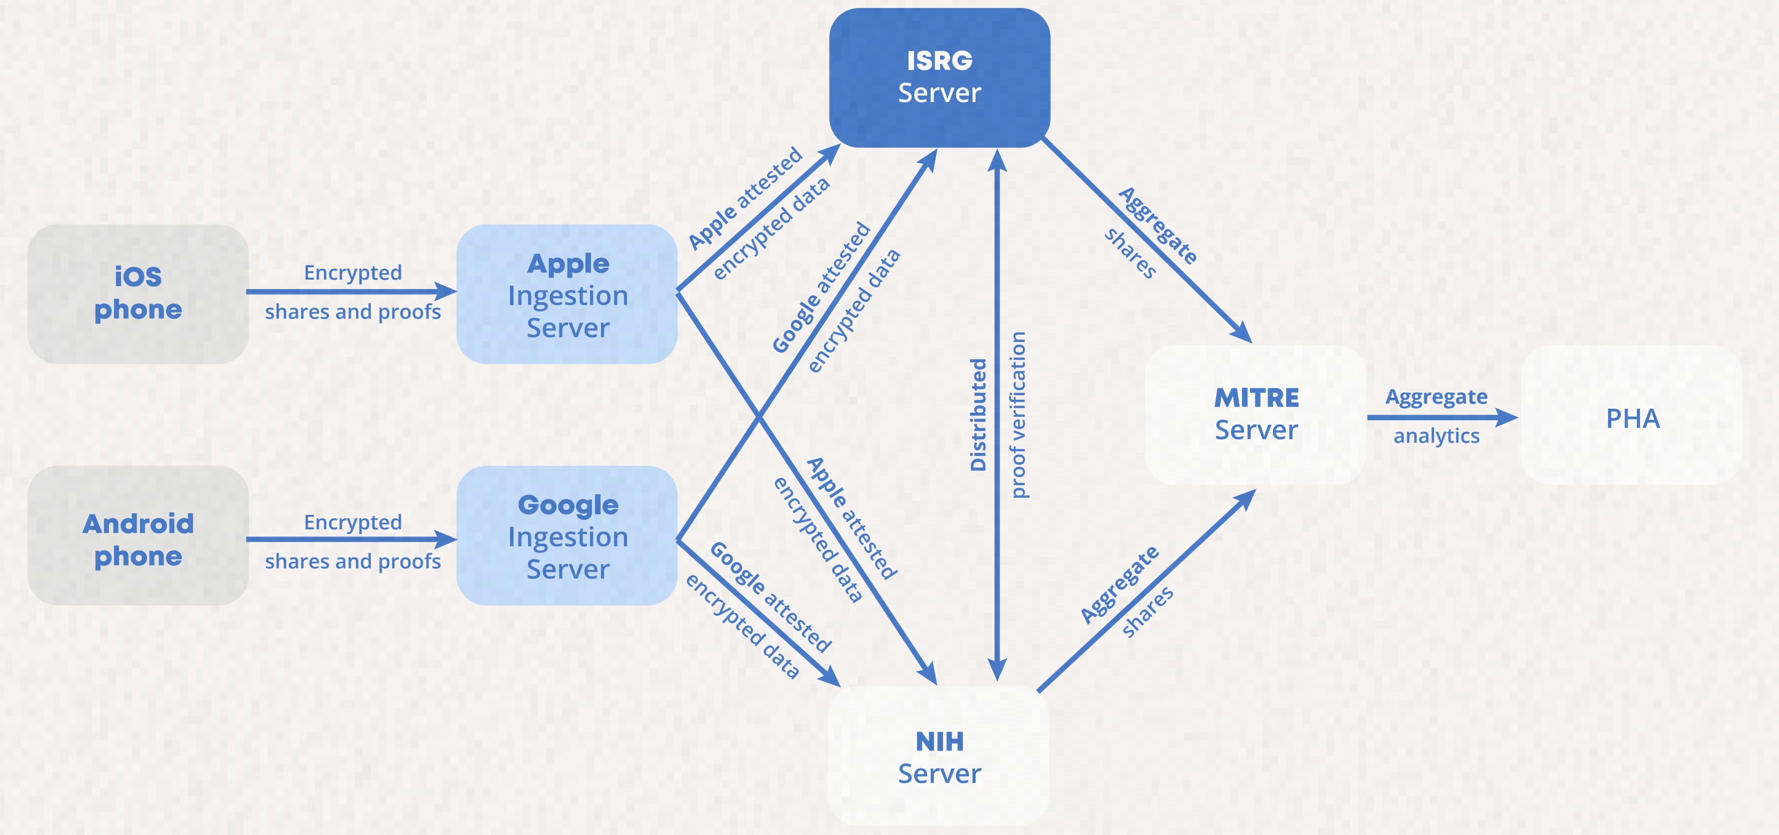
iOS phone
Encrypted shares and proofs
Apple Ingestion Server
Android phone
Encrypted shares and proofs
Google Ingestion Server
Apple attested encrypted data
Google attested encrypted data
Apple attested encrypted data
Google attested encrypted data
ISRG Server
Distributed proof verification
NIH Server
Aggregate shares
Aggregate shares
MITRE Server
Aggregate analytics
PHA

ISRG'S ROLE IN COVID-19 EXPOSURE NOTIFICATION SYSTEM

A quantum leap for privacy

UNMATCHED CRYPTOGRAPHIC GUARANTEES

The algorithms used for Divvi Up are a step-function improvement over what is possible with traditional telemetry reporting schemes. Divvi Up provides cryptographic guarantees that no participant in the system except for the device uploading the report can ever see an individual input. The aggregating servers can only ever see shares of inputs, and the collector only ever sees aggregations over a sufficiently large set of inputs. No server can reveal anything about the original input...

This makes it impossible for server operators to build up privacy-violating profiles of individuals. It also means that even if an aggregator were compelled to reveal their shares, they would be unintelligible unless the other aggregator's shares were also revealed.

This quantum leap in privacy comes with some tradeoffs. First, privacy-preserving metrics (PPM) systems are less flexible. A conventional telemetry system will gather all the data it possibly can and store it indefinitely on the collecting server, allowing arbitrary queries, conversions, aggregations, or visualizations over arbitrary subsets of inputs. PPM integrators, however, must know the aggregations they will want before they design the encoding of inputs in the client. Further, because of the input validity proof that must be computed by the client, transmitted to aggregators and then jointly evaluated, PPM uses more computation time and bandwidth than a conventional system which can decide on the validity of an input by simply examining it.

However, the algorithms currently included in the standard we are developing for PPM (Prio and Heavy Hitters) are significantly more efficient and scalable than any other system providing similar privacy guarantees. Figures 6 and 7 in the Prio paper compare client execution time (which is more precious than server time, since clients are often battery-powered mobile devices) and bandwidth usage against existing systems like Succinct Non-interactive Arguments of Knowledge or Non Interactive Zero Knowledge proofs. Crucially, Prio's bandwidth usage does not grow with submission size. Heavy Hitters, which is optimized for cases where the Prio encoding of an input would be too large, uses a novel construction called incremental distributed point functions whose exciting property is that computation time and bandwidth requirements scale linearly with the size of the input, as illustrated in figures 7 and 8 of the paper, whereas the cost in preceding systems was quadratic.

Differential privacy is another exciting technique used to improve user privacy in data collection systems. The idea is that random noise is added to each individual input such that each reveals less information, but still yields cogent aggregates. The drawbacks to differential privacy are that it can be quite difficult to correctly tune the probability parameters against the size of inputs, the number of clients, or other factors. If a large probability of input permutation is needed to guarantee privacy, then the probability of infrequently occurring events being lost in the noise is high as well. PPM does not have these problems, and PPM deployments can still apply differential privacy in either the client or the aggregation servers. Similarly, submitting telemetry through anonymizing proxies mitigates some risks with correlating inputs together into a profile of a user, and the PPM protocol supports deployments incorporating such proxies.

For more background on different methods of gathering telemetry and why PPM and the algorithms it encompasses are so exciting, we recommend this series of articles by Eric Rescorla, the Chief Technology Officer for the Firefox browser.

Introducing Prossimo

MEMORY SAFETY FOR CRITICAL DIGITAL INFRASTRUCTURE

Launched this year after several years of work, Prossimo seeks to lead the world towards a future where the Internet's critical software infrastructure uses memory safe code, while raising awareness of memory safety and its importance along the way.

PROSSIMO

FOR MEMORY SAFETY

The problem

A UBIQUITOUS, PERSISTENT PROBLEM

The goal of ISRG's Prossimo project is to move the Internet's security sensitive software infrastructure to memory safe code. Many of the most critical software vulnerabilities are memory safety issues in C and C++ code.

While there are ways to mitigate the risk, including fuzzing and static analysis, they do not eliminate the risk. Using memory safe languages eliminates the entire class of issues.

We recognize the amount of work it will take to move significant portions of the Internet's C and C++ software infrastructure to memory safe code, but the Internet will be around for a long time. There is time for ambitious efforts to pay off. By being smart about our initial investments, focusing on the most critical components of the most critical infrastructure, we can start seeing significant security improvements within 1-2 years.

Lack of memory safety is a serious, persistent threat to Internet infrastructure. The problem is mainly C and C++. New problems surface every day, and they appear everywhere you look: Web servers / proxies, utilities, TLS, NTP, DNS, kernels—almost every deployment stack contains many millions of lines of code that is known to be unsafe.

Google estimated that 90% of Android vulnerabilities are due to a lack of memory safety. Microsoft estimated that 70% of all vulnerabilities in their products over the last decade have been caused by a lack of memory safety. An analysis of 0-day vulnerabilities that were discovered being exploited in the wild found that more than 80% of them were due to a lack of memory safety.

But the concern goes beyond the stack; lack of memory safety impacts the public at large and society pays the price—from privacy violations as a result of data breaches, to financial losses from an exploited bug, to denial of public services as a result of a system falling over.

Although this problem is ubiquitous and impacts all of us, we believe Internet-scale change is possible. We know it's possible because we've done it with Let's Encrypt. We aim to do the same with Prossimo.

Our role & focus

INTERNET-SCALE CHANGE IS POSSIBLE

We view ISRG's role as providing strategic planning, facilitation, and communication. We will identify high impact investments, build relationships with maintainers, help develop work plans, and coordinate the work. This includes raising the necessary funds and getting them to the right people. In addition, we will communicate with the public regarding progress and momentum in order to build support for the project and the ideas behind it.

We are focused on bringing memory safety to the Internet’s most critical software. We look for software fitting the following criteria and make plans to move usage to memory safe code:

- Very widely used (e.g. nearly every server and/or client)
- On a network boundary
- Performing a critical function
- Written in languages that are not memory safe, such as C / C++ / assembly

Our approach

PRIORITIZE FOR GREATEST IMPACT

We believe our competencies are well suited to achieving rapid progress in making the Internet's software infrastructure safer for everyone. Initiatives undertaken in the first year of Prossimo's operation have paved the way for improved memory safety in curl, Apache httpd and the Linux kernel, proving that our approach can produce results. The categories of software we are currently working on are:

- Web servers / proxies (2 initiatives)
- Utilities (1 initiative)
- TLS (1 initiative)
- NTP (1 initiative)
- DNS (1 initiative)
- Kernel (1 initiative)

> "AT GOOGLE, WE HAVE FOUND THOUSANDS OF MEMORY CORRUPTION VULNERABILITIES IN CRITICAL OPEN SOURCE PROJECTS BY USING AUTOMATED TESTING METHODS LIKE FUZZING. ISRG IS PLAYING A CRUCIAL ROLE IN MAKING THE INTERNET SAFER BY REWRITING ESSENTIAL OPEN SOURCE SOFTWARE IN MEMORY SAFE LANGUAGES LIKE RUST."
>
> **ABHISHEK ARYA**
> PRINCIPAL ENGINEER & MANAGER | GOOGLE OPEN SOURCE

Our approach operates on the following principles:

WORK WITH MAINTAINERS WHENEVER POSSIBLE.

Maintainers have valuable knowledge and the ability to ship memory safety updates to their existing users. Building competing software and getting users to switch is much more difficult. By working with maintainers, and funding them when it makes sense, we can get safer software into the hands of users more quickly. This approach also strengthens or broadens maintainers' skillsets, which advances the Open Source ecosystem as a whole.

Funding maintainers for the work helps create buy-in and alleviates resource concerns. Some maintainers do not have time or energy to do additional work (e.g., they may have a full-time job), and in those cases we fund contractors with work plans based on maintainer input.

Sometimes it will be necessary to create new software to replace existing software. When we need to do so, we will work as closely as possible with existing communities and attempt to build upon existing high-quality components.

PREFER A MODULAR APPROACH AND INVEST IN LIBRARIES.

We encourage projects to replace libraries or modular functionality with memory safe libraries, rather than embark upon ground-up rewrites. This allows us to break up the work into manageable pieces and deliver value incrementally.

It also allows for build-time configuration to select implementations when existing users need the ability to opt into the older unsafe versions, either because their environment does not support the new language or because there is a functionality difference. This addresses the concern many maintainers have about abandoning certain specialized users.

Since many projects will end up using the same memory safe libraries, this approach also allows us to invest and build confidence in a particular set of libraries. Investments in a library for one project will add value across multiple projects. For example, the curl project will use the Hyper and Rustls libraries. The work we do to build excellent C API wrappers and improve the integration experience will help many projects that will use the libraries in the future.

BUILD TRUST BY PROVIDING ADDITIONAL SUCCESS STORIES.

Some maintainers are understandably hesitant to make fundamental changes to how their projects work, such as adding a new language or replacing important libraries with new ones. It's on us to make the case; we will do that by building up a corpus of success stories and continuing to engage with maintainers about how their concerns can be addressed.

We've started working with more progressive maintainers that need less convincing. As those projects succeed and get positive feedback, other maintainers will come to trust the model that we advocate. Our hope is that over time we can convince more conservative maintainers that moving to memory safe code is a worthwhile endeavor.

What's next

OUR CURRENT & UPCOMING WORK

Prossimo is focused on six different categories where we believe we can have the greatest impact: TLS, NTP (Network Time Protocol), DNS, kernels, web servers / proxies, and utilities.

C

curl: Today curl users can choose to build curl with Hyper and Rustls. Work is being done to make sure full testing is in place to for these new options' long-term stability and success.

A

Apache httpd: Developing mod_tls to replace mod_ssl. The new module uses Rustls for TLS.

D

DNS: Planning a memory-safe high-performance fully recursive DNS resolver.

R

Rustls: Contributing improvements to this memory safe TLS library, positioning it as a suitable replacement for OpenSSL in many applications.

L

Linux kernel: Working with Alex Gaynor and Miguel Ojeda to add support for Rust modules. Getting support from key Linux developers.

N

NTP: Planning a memory-safe Network Time Protocol client and server implementation.

"ISRG has taken on the challenge of changing the *status quo*. I appreciate they are being brave about their selection of high-risk, high-reward projects."

MIGUEL OJEDA

SOFTWARE ENGINEER | RUST FOR LINUX PROJECT

A closer look at the group

PEOPLE, FUNDERS, & FINANCIALS

Founded in 2013, Internet Security Research Group is the nonprofit organization behind Let's Encrypt. Since then, ISRG has grown to a staff and board of thirty, nearly one hundred funders, and tens of thousands of individual donors. Here's a closer look at the people behind ISRG.

ISRG

Internet
Security
Research
Group

Financials

ISRG is proud to run a financially efficient organization. While our growth has been tremendous in terms of impact—it took us five years to issue one billion certificates, but just one year to issue our second billion—our expenses have only grown modestly.

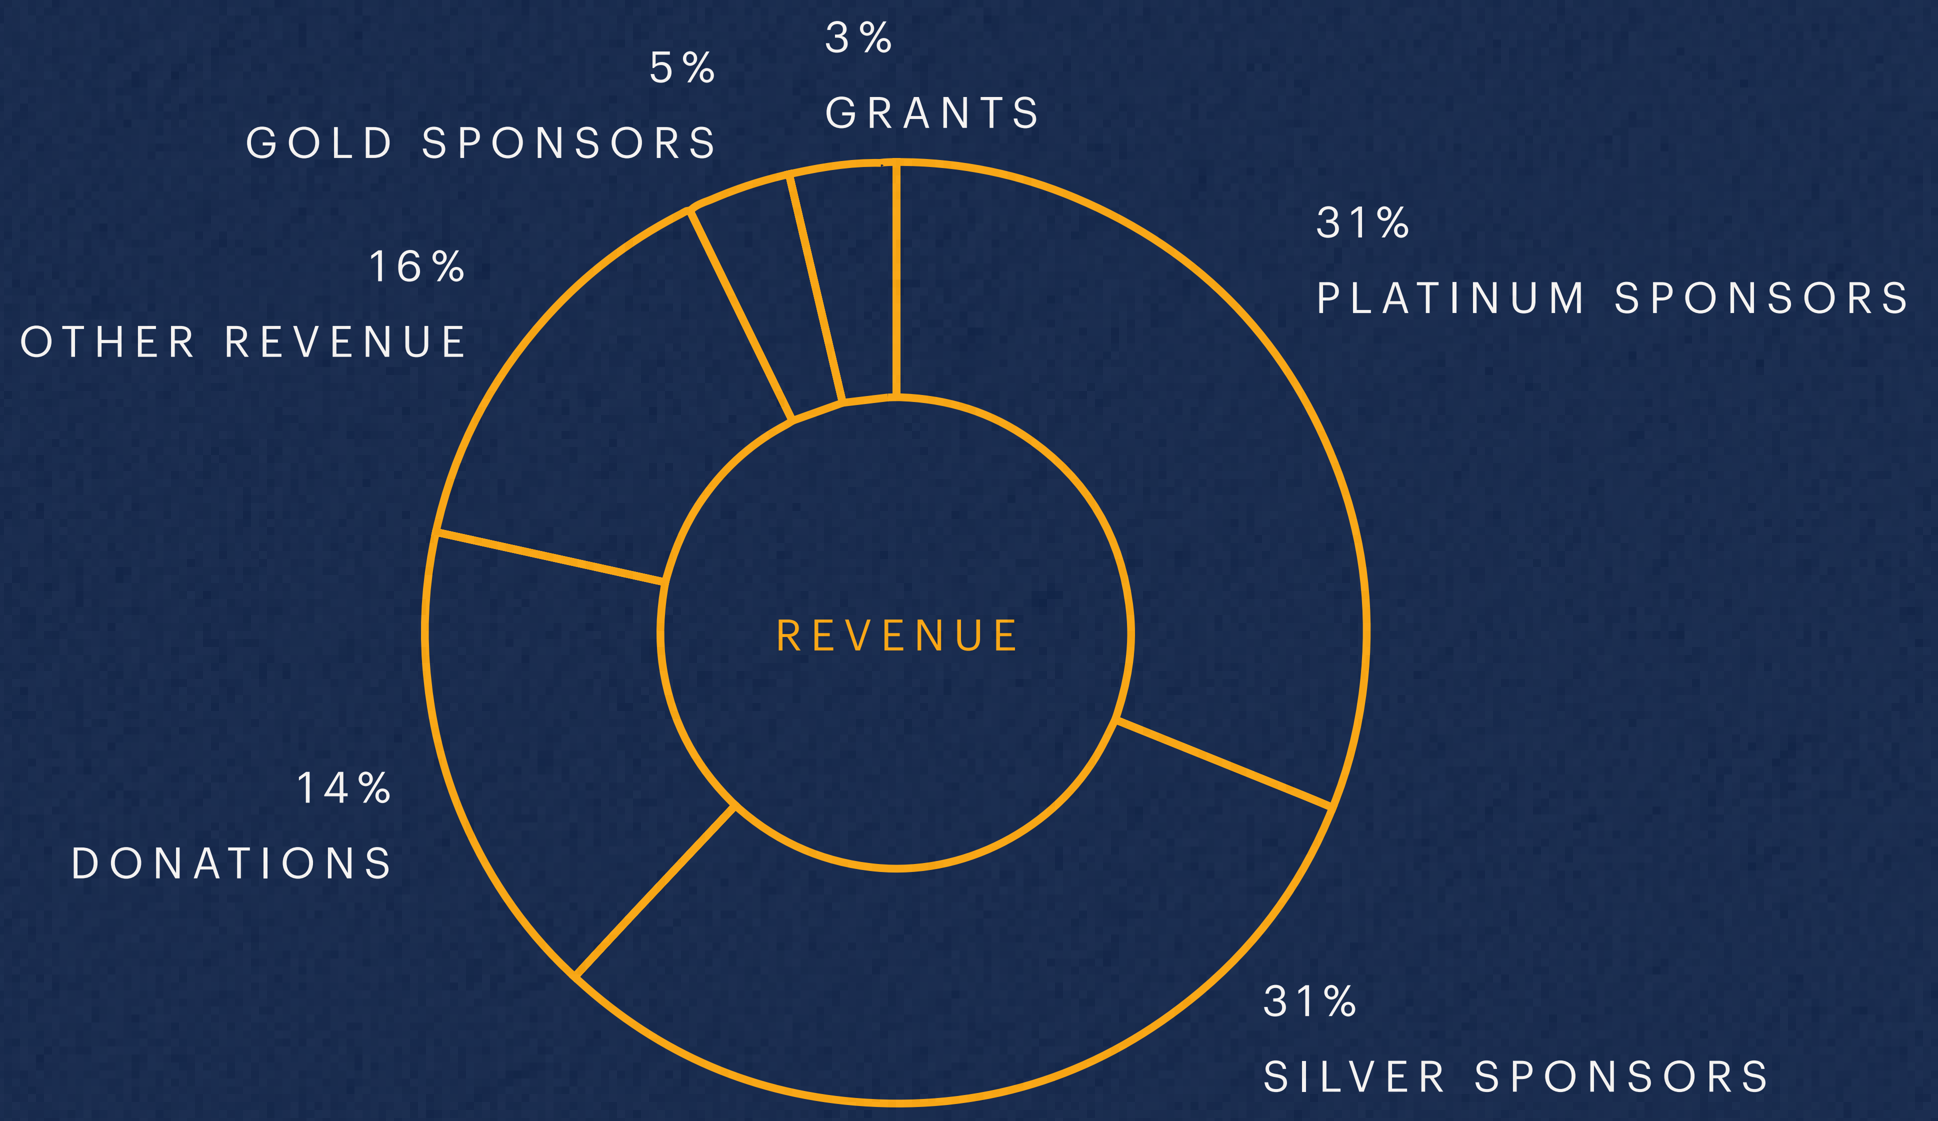

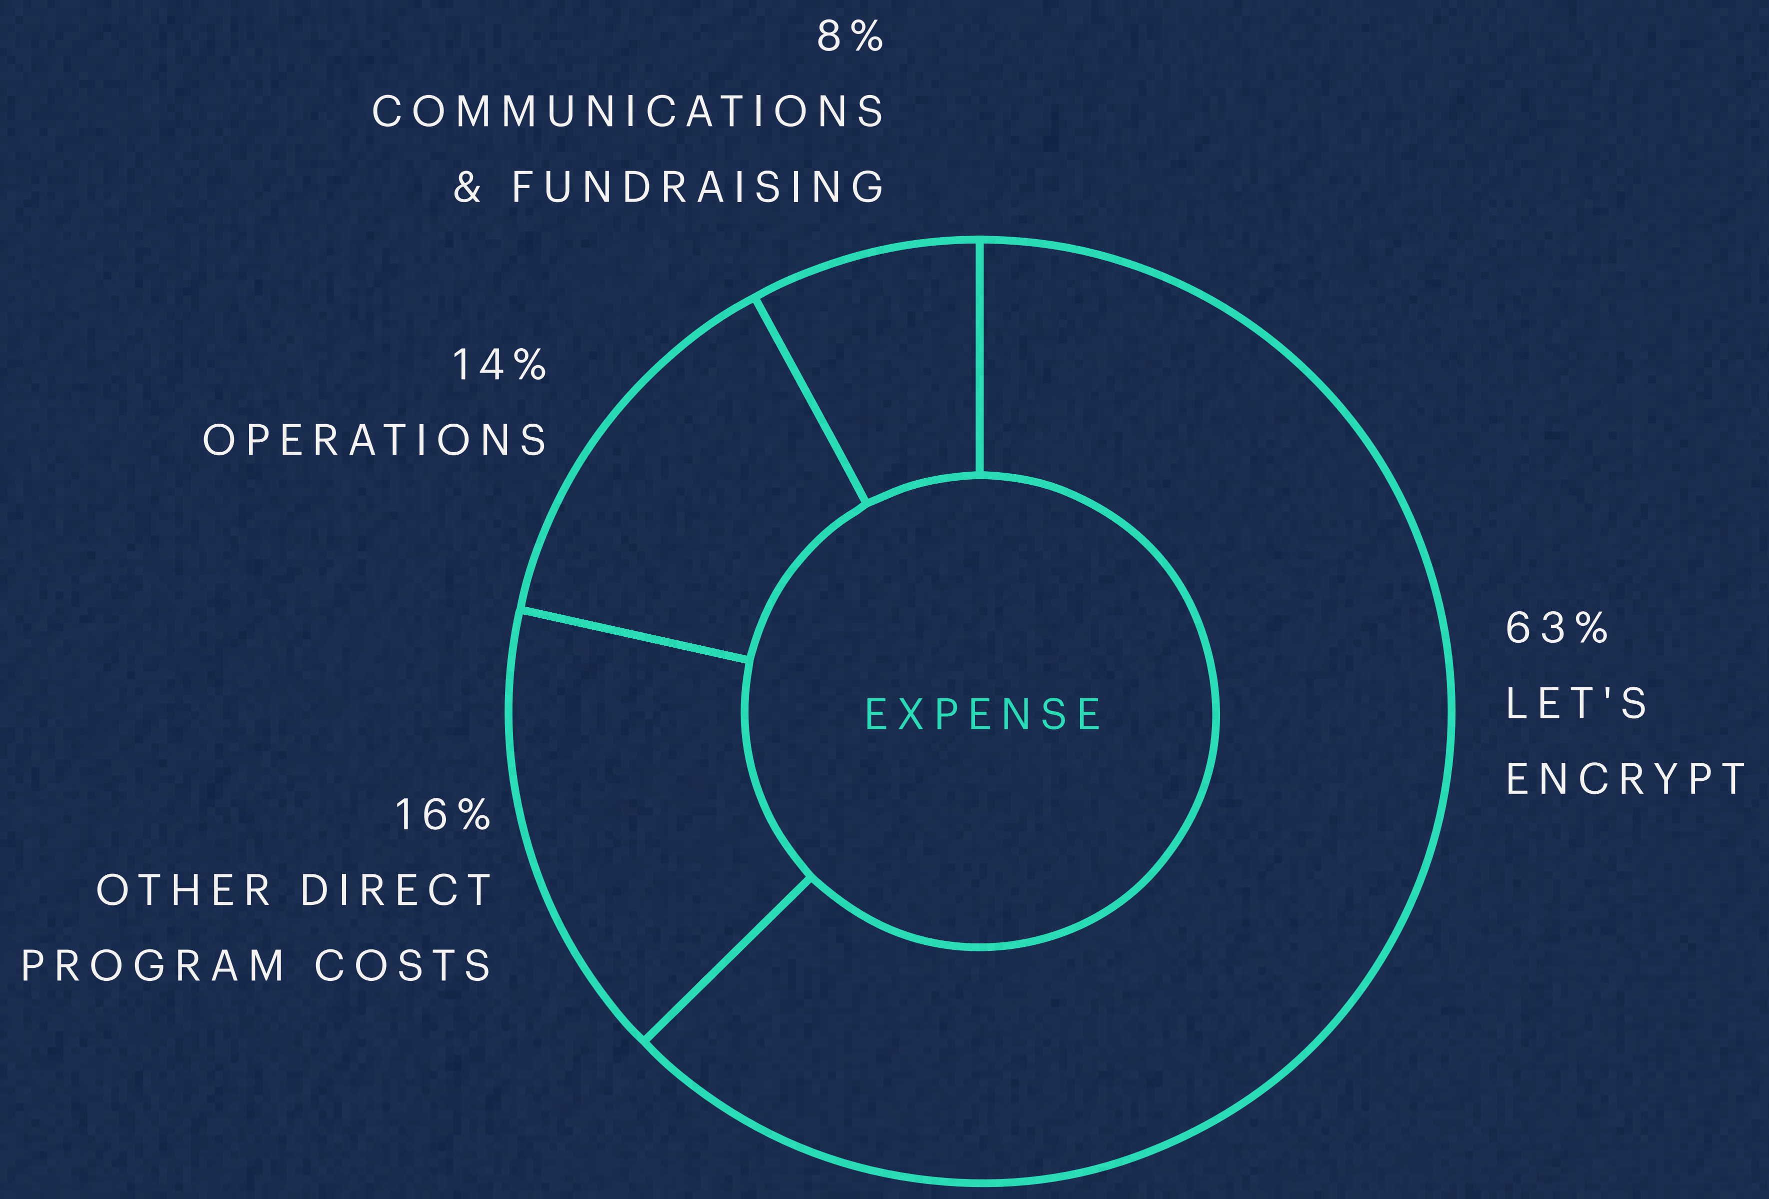

Funders

One hundred three sponsors and funders from more than twenty countries around the world supported us in 2021.

From one-employee shops to thousand-employee companies, we are proud to have support from these organizations who prioritize the importance of investing in a more secure and privacy-respecting Web.

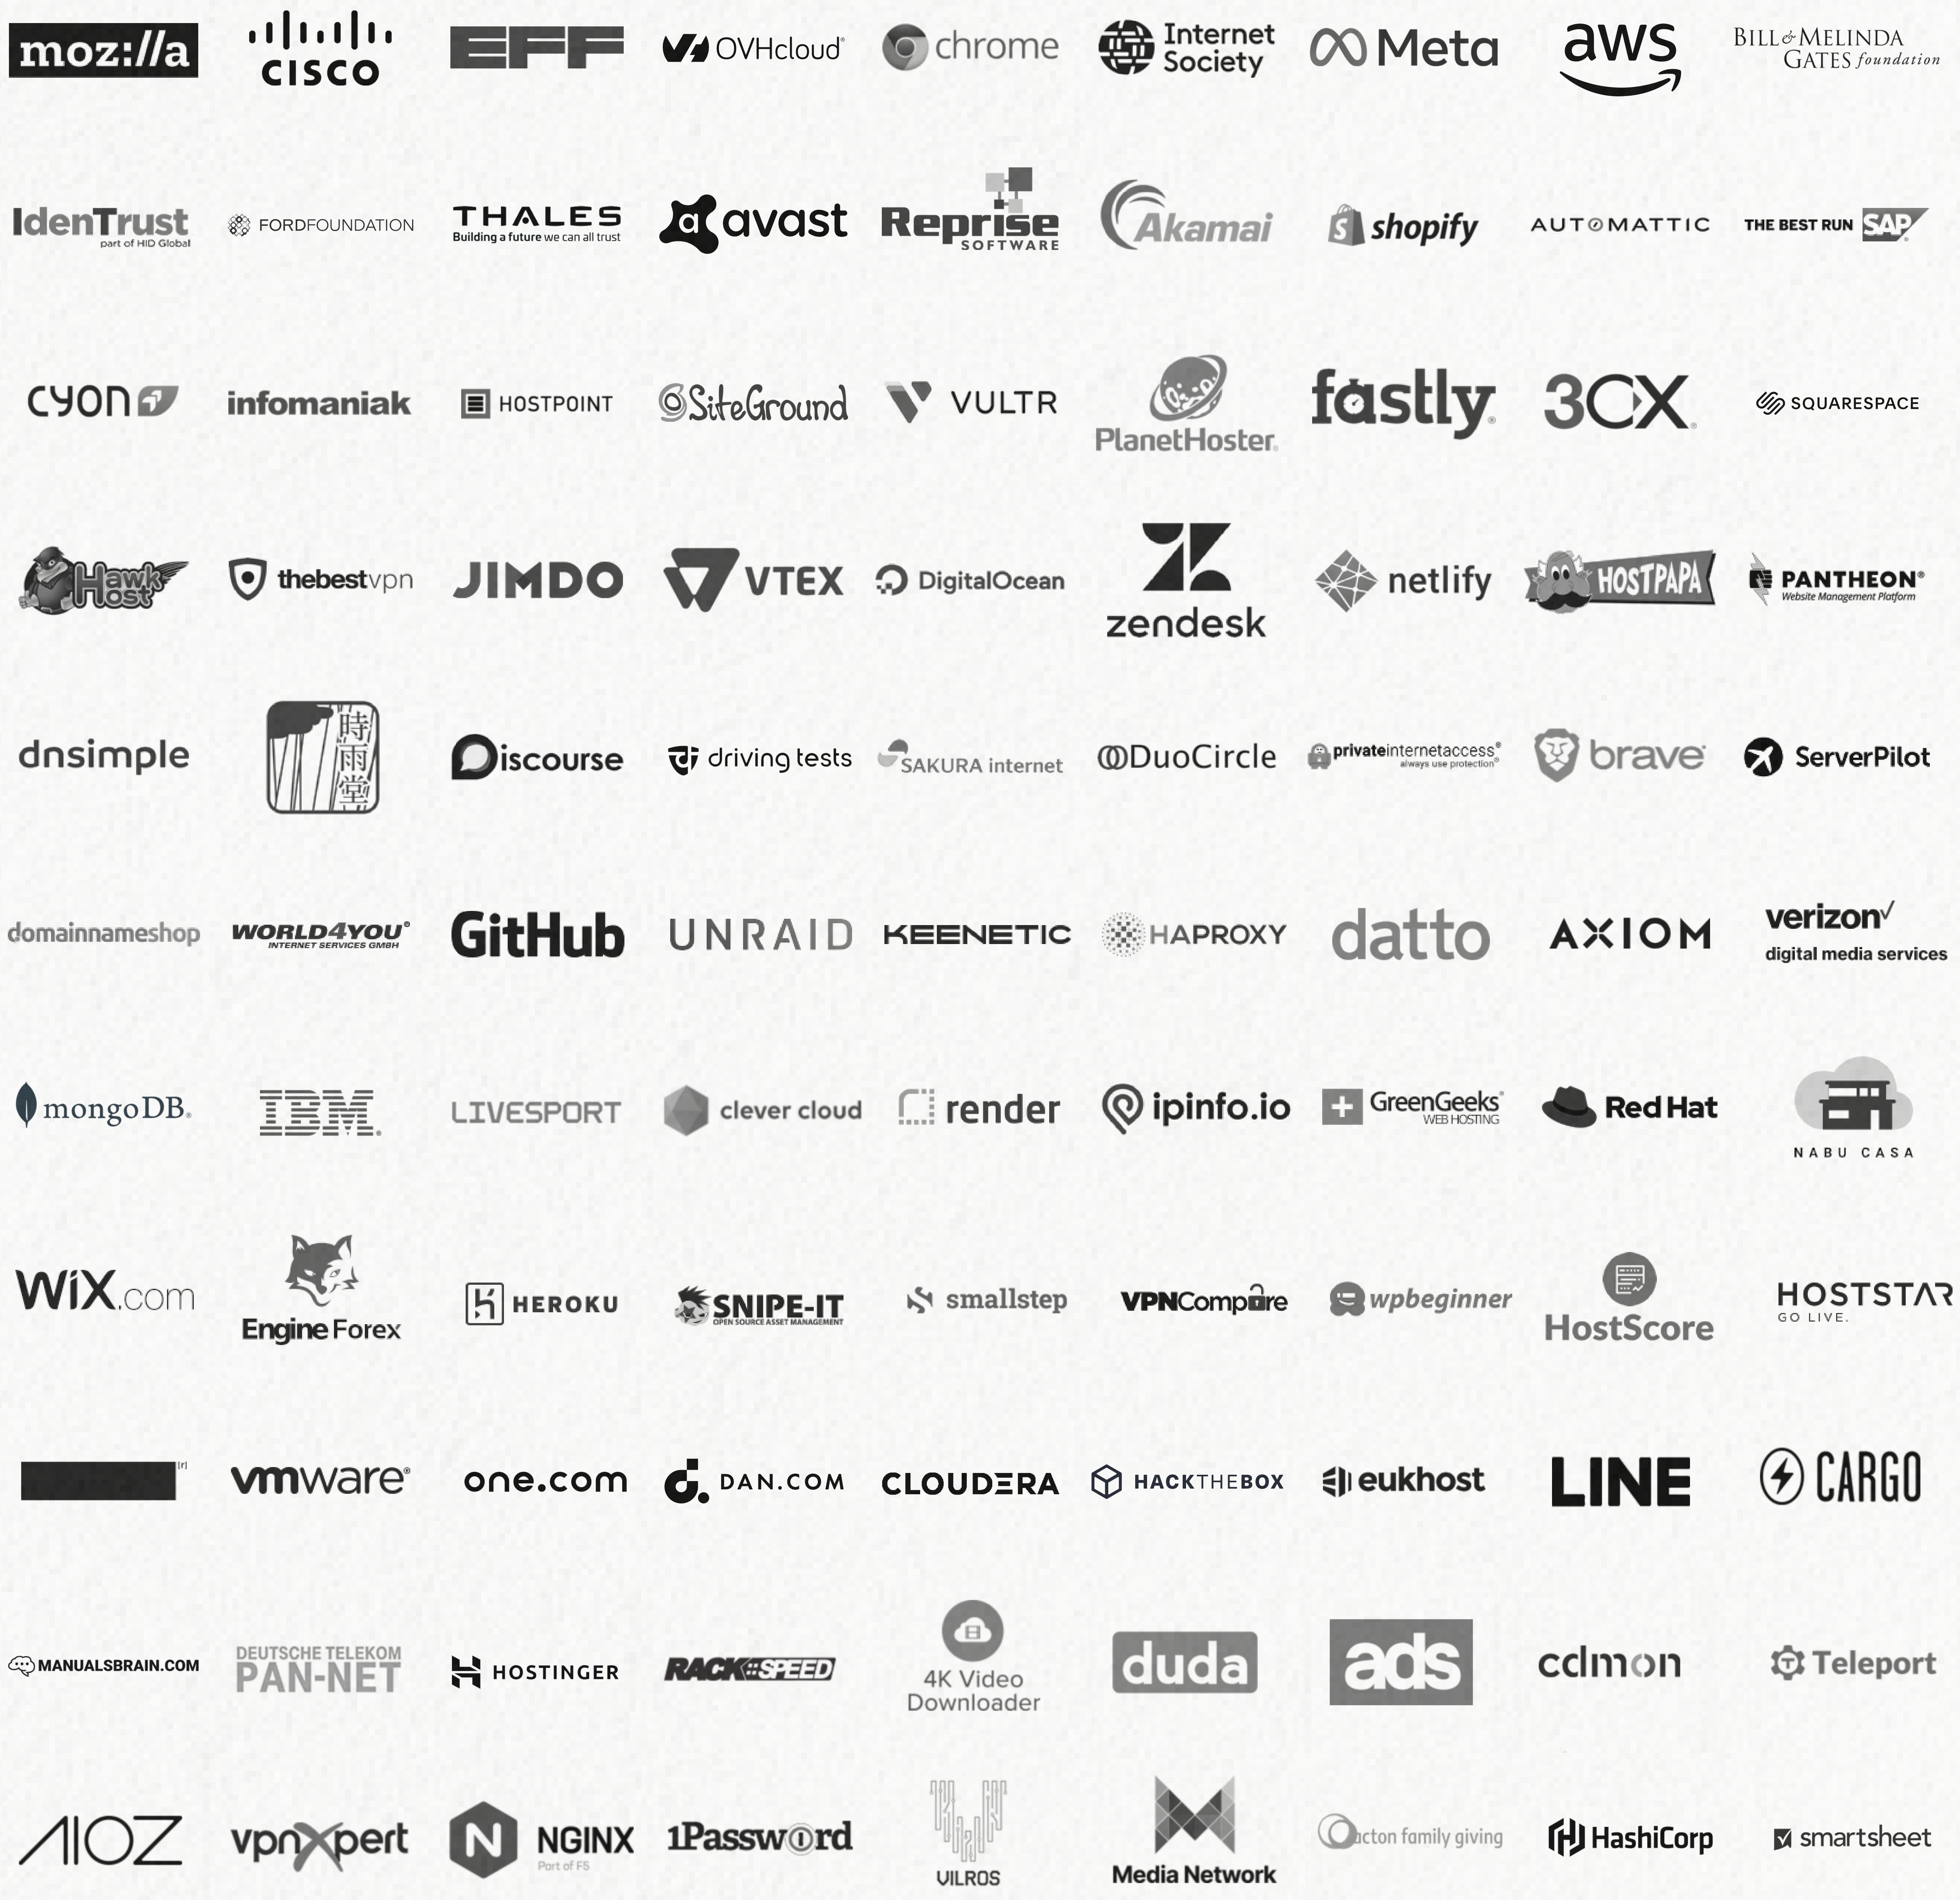

"As a sponsor, Red Hat is proud to support Let's Encrypt and the projects that serve their mission to create a more secure and privacy-respecting internet."

CHRIS WRIGHT
SVP & CTO | RED HAT

Board & Staff

2021 BOARD OF DIRECTORS

AANCHAL GUPTA
INDEPENDENT

CHRISTINE RUNNEGAR
INTERNET SOCIETY

DAVID NALLEY
AMAZON WEB SERVICES

JENNIFER GRANICK
AMERICAN CIVIL LIBERTIES UNION

J. ALEX HALDERMAN
UNIVERSITY OF MICHIGAN

JOSH AAS
INTERNET SECURITY RESEARCH GROUP

ERICA PORTNOY
ELECTRONIC FRONTIER FOUNDATION

PASCAL JAILLON
OVH CLOUD

RICHARD BARNES
CISCO

VICKY CHIN
MOZILLA

ISRG STAFF

AARON | SOFTWARE ENGINEER
AMIR | SITE RELIABILITY ENGINEER
ANDREW | SOFTWARE ENGINEER
BRANDON | SITE RELIABILITY ENGINEER
DAN | SENIOR DEVELOPMENT OFFICER
FINN | COMMS & DEVELOPMENT COORDINATOR
JACOB | SOFTWARE ENGINEER
JAMES | SITE RELIABILITY ENGINEER
J.C. | SITE RELIABILITY ENGINEER
JENESSA | FUNDRAISING SPECIALIST
JILLIAN | SITE RELIABILITY ENGINEER
JOSH | EXECUTIVE DIRECTOR
KIEL | SITE RELIABILITY ENGINEER
OLENA | FINANCE MANAGER
PHIL | SITE RELIABILITY ENGINEER
SAMANTHA | SOFTWARE ENGINEER
SARAH | VP COMMUNICATIONS
SARAH | CHIEF FINANCIAL OFFICER
TIM | SITE RELIABILITY ENGINEER

TECHNICAL ADVISORY BOARD

RICH SALZ | AKAMAI
JOE HILDEBRAND | MOZILLA
JACOB HOFFMAN-ANDREWS | EFF
YUETING LEE | FACEBOOK
RUSS HOUSLEY | INDEPENDENT
RYAN HURST | GOOGLE
STEPHEN KENT | INDEPENDENT
KAREN O'DONOGHUE | INTERNET SOCIETY
IVAN RISTIC | INDEPENDENT

Let's change the world together.

SUPPORT OUR WORK

Thanks to our staff, community, users, sponsors, grantmakers, and individual donors, ISRG and its projects are building a better Internet for everyone, everywhere.

ISRG Internet Security Research Group

Let's Encrypt

PROSSIMO

Divvi Up

The mission of Internet Security Research Group (ISRG) is to reduce financial, technological, and educational barriers to secure communication over the Internet. ISRG is a California public benefit corporation, recognized by the IRS as a tax-exempt organization under Section 501(c)(3).

For more on our work, visit: https://abetterinternet.org

OUR THANKS TO THESE INDIVIDUAL DONORS

This year we received thousands of donations from 50+ countries around the world. Our thanks to these donors for their support.

A2 Engineering Services
ABG
ABHAY SIKARWAR
Abilor
Absolute Genius
Aerial Tour
Ahlin Sodji - immofly.online
Aileen M. Baluyut
AirFlare
Aivaras Stukas
AJ Jordan
Alex
alpet
anagora.org
Andre Medeiros
Andres de Caso (Ensena.com.ar)
Andrew Hedges
Andrew Janke
Andrew King
Ann and Duncan Sterling
annopnod
Antony Lopez
Applex Group
Ardent Management Consulting
Arnaud Bonnefoy
Artan Sinani
artkiddo.pl
aspectra AG
Badassops LLC
Baonetz GbR
Bare Telecom, LLC
bas
Batch.com
Ben Sykes
BenjaminZ
Bertram Paaskesen
Bipin Upadhyay
Brandon Ham
CAHILL WEB WRK
CCM Software Technologies
Certbot is the best
Chad Marshall
Chris Wilson
Christopher Cate
Clearbold, LLC
Cloud9Dynamics
COMPANIA FITNESS
Curiefense
Dami√† Poquet Femenia
Daniel D'Angelo
Daniel Spencer
Danny Apostolov
David Aldana
David M N Bryan
De Bortoli Wines, Australia
Dean Settinelli
Deepak Sharma
Dennis C Fetterly
Diarmuid √ì Briain
Diego Quir√≥s
Digitala
diodonfrost
dohq
Donato Acosta Eusebio
Dongli zhu
Doppelganger
Drake Software Services
elastic items GmbH
Elias Medawar
Epsitec SA
Eric Heydrick
Eric Pannetier
eschool.edu.my
EXG
Felix Tang
fishbeetle
Florian Herzog
Frederic KIEBER
Fraz Ahmed
Gary Gapinski
GILLIBERT
Gladblad
Gopesh Sahu
Greg Hawthorne
Greg Powell
GromHSCR
guenter@loerincz.org
Hans Kappert
Harald Schmidt
Heinz-Josef Colley
Hofstadter, Inc
HopeMedia Italia
Hotsoft Inform√°tica
Huang Meng
Ian Peters
Ikon ehf
impute.me
In memory of Bertrand Might
inGenerator Ltd
IONICA
Istvan Cocron
Ivan Monnier
J Dev - https://jdev.fr
Jacob
jedsada
Jeri Luis Balconi
Jes Drost Nissen
JF Rullier
Johann
Johann Klasek
Jose Monterroso
Joshua Estep
Jozef
JP Pozzi
Juan Martinez Gomez
Jxck
Kevin Kaland, WizOne Solutions
Khurram Khan
Kim Visscher
Kitasato University
Kuhn Computer Solutions
Lasse Dahl
Luca Castiglioni
Lucien Van Elsen
Mahamat Hamid Haggar
Malik Dixon
MANTICORE GIVING
Manuj Dua
Martha Rojas
Martin Lonkwitz
Mathias Zajaczkowski
Matt Jeanes
matt sossi
Maxwell Snyder
Meaulnes Legler
Microspino
midrange.com

Donors are listed alphabetically and recognized exactly as submitted.

> “
>
> WE USED LET'S ENCRYPT IN A STARTUP I FOUNDED, THEN A COUPLE OF DAYS AGO I CAME ACROSS ISRG'S MEMORY SAFETY EFFORT. I THOUGHT IT'S AWESOME, SO I FIGURED IT'S TIME FOR A DONATION.
>
> THANK YOU FOR THE GREAT WORK YOU'RE DOING. I HOPE AS THE STARTUP GROWS WE'LL BE ABLE TO MAKE MORE MEANINGFUL CONTRIBUTIONS."
>
> SINAN TAIFOUR
> CEO | MAQSAM

milordk.ru
Mitch Scobell
Morten Simonsen Motions
Nathan David
netzbetrieb EDV-Herne
Nguyen Ngoc Linh
nhsft.co.uk
Nicolas Sanchez
NICRONICS
Nikolay Konovalov
Niteo
NitsaSoft in Sweden AB
Noor Alasadi & Majd Salloum
Nordine VALLAS
octopy.de
Omni Solution Technologies LLC
one.com
Patrice Delavictoire
Patricia Comesanas
Pavol Poremba
Pedro Luis Gonzalez
Pedro Moya
PeridotCays.com
PnA Digital
Pradyumna Shankar Mahajan
Pyae Phyo Khine
Pyppin
Question Mark Media
quickytools
RainViewer
Raul Fernandez Garcia
Ravishankar Haranath
Ray Maritz
redconfetti
Reliable Web, Inc.
RENCI
resume-template.online
Reva Synergy & Neutral Media
Ricardo Katz
Rickard von Essen
Rober Pankrath
Roger Goudarzi & Nicola Downes
Ross Payne
ruby tseng
Russ Hore
Safety1st
Salah Oka
SALUDCENTER SPA
satyayoga.se
scimsacademy.com
Scott Helme
Sebahattin Celebi
sergey musin
Sergii Rudchenko
Sirichakorn
Suwapinyopas
Snap Surveys Ltd
Sorin Dodenci
Srinivasan Rangaraj
Stefan Foulis
Sudhir Jonathan
Sylvain Parent
takatoshi fujiya
Taylor Hornby
taylordcakes.co.nz
Telesploit
The Old Island Stamp Company
Thorsten Biel
Tobias Ribizel
Torutek
TOTALogistix
Tracy Di Marco White & Jason White
Trident Honda
Tripwire
Tripwire Interactive
Turritopsis Dohrnii Teo En Ming
V
VacationRenter
VAL-U-PRO CONSULTING GROUP & AMAZEBABA
Vasile Iacome
vikto9494
VINN
W.Sanders
wael alttarh
Wally Grotophorst
Warren Volz
Webstel Computer Systems
wemu
Wernilein Eifel
Widda Alata
William Sagoe
www.lastdragon.net
Wysper LLC
Yoshihisa Okamoto
zan.su