

MilagroProposal

Project Proposal: Milagro

Abstract

Milagro is a distributed cryptosystem for cloud computing. Its purpose is to provide an open source alternative to proprietary key management and certificate backed cryptosystems used for secure communication and authentication. The adoption of Milagro will create a secure, free, open source alternative to monolithic certificate authorities and eliminate single points of failure.

Background

The Cloud Computing industry is using 40-year-old cryptographic algorithms and infrastructure, invented for a different era when client-server computing was the dominant paradigm. At the heart of it, is the continued reliance on outdated, and problematic, monolithic cryptographic trust hierarchies such as commercial certificate authorities.

A number of factors are aligning to make this the right time to bring forth an alternative to the Internet's continued reliance on PKI.

The Cloud Infrastructure as a Service (IaaS) industry as a whole encounters friction bringing the largest customers in regulated industries onto their platforms because issues of cryptographic trust, data residency, and data governance prevent total adoption among regulated industries.

Devops teams tasked with running an IaaS provider's datacenter automation encounter challenges scaling and automating data center operations when confronted with the complexities of running encryption, certificate and key management infrastructures built for a client-server era.

Enterprises in regulated industries find challenges to transform entirely into digital businesses because the economics of cloud computing are unavailable to them.

Despite the astounding growth of cloud infrastructure as a service platforms over the last few years, full adoption by organizations with stringent data security requirements won't be achieved until these fundamental capability issues get resolved.

Lastly, the Internet as a whole is suffering from an erosion of trust following incidents with commercial certificate authorities industry, i.e., compromised root keys, and failures in due diligence issuing real domain certificates.

Indeed, mass surveillance, a lack of easy end-user encryption, ~~a growing demand for key escrow under legal oversight~~ [edited 2016-02-03 to strike language open to misinterpretation] http://mail-archives.apache.org/mod_mbox/incubator-milagro-dev/201602.mbox/%3CCAP5eBitsuTQucc8ebWsvd-iGbF0ozso8Ce%3DkbyJrW%2BkdaqE-9w%40mail.gmail.com%3E], and general certificate authority security concerns create the question: How appropriate is the continued dependency on PKI when the goal is to advance the benefits of cloud computing across the technology landscape?

Netcraft is the industry standard for monitoring Active TLS certificates. In May 2015, they stated that "Although the global [TLS] ecosystem is competitive, it is dominated by a handful of major CAs — three certificate authorities (Symantec, Comodo, Godaddy) account for three-quarters of all issued [TLS] certificates on public-facing web servers."

The Internet Security Research Group's (ISRG) "Let's Encrypt" initiative aims to make Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates available for free in an automated fashion. This is a step in the right direction, in that it removes the risk of profit before ethics. The real issue, which is one entity acts as a monolithic trust hierarchy, is not addressed. The monolithic trust hierarchy is a fundamental design flaw within PKI itself.

The rate of attacks against certificate authorities seems to be [increasing] (<http://wiki.cacert.org/Risk/History>) as the obvious single point of compromise design inherent to PKI is becoming a more popular route to carry out attacks.

Proposal

Milagro is an open source, pairing-based cryptographic platform to solve key management, secure communications, data governance and compliance issues that are challenging Cloud Providers and their customers.

It does this without the need for certificate authorities, putting into place a new category of service providers called Distributed Trust Authorities (D-TA's).

The M-Pin protocol, and its existing open-source MIRACL implementation on which Milagro will build, are already in use by Experian, NTT, Odin, Gov.UK and are being rolled out at scale for zero password multi-factor authentication and certificate-less HTTPS / secure channel.

It is proposed that Milagro enter incubation at Apache. At the same time, a draft standard for M-Pin has been prepared and recently submitted to IETF. The standards process at IETF and the platform implementation at Apache will run in parallel.

Why Pairing-Based Cryptography, why now?

Over the last decade, pairings on elliptic curves have been a very active area of research in cryptography. Pairings map pairs of points on an elliptic curve into the multiplicative group of a finite field. Their unique properties have enabled many new cryptographic protocols that had not previously been feasible.

Standards bodies have already begun standardizing various pairing-based schemes. These include the IEEE, ISO, and IETF. Besides identity-based encryption (IBE), the standardized schemes include identity-based signatures, identity-based signcryption, identity-based key establishment mechanisms, and identity-based key distribution for use in multimedia.

NIST has also recommended the standardization and adoption of pairing-based cryptographic systems for government agencies. In the NIST "Report on Pairing-based Cryptography" issued in February 2015, they state:

>"It has been a decade since the first IBE schemes were proposed. These schemes have received sufficient attention from the cryptographic community and no weakness has been identified. IBE is being used commercially, primarily by Voltage Security and Trend Micro. Intel's EPID scheme is another example of pairings being used commercially. > As a result of our study, we believe there is a good case for allowing government agencies to use pairings. Pairings have been shown to have numerous applications, helping to solve problems that are impossible, difficult, or inefficient with traditional public-key cryptography or symmetric encryption."

The biggest beneficiary of these new pairing-based cryptographic protocols will be the Cloud Infrastructure as a Service industry. Pairing-based cryptography can provide real world solutions, right now, to the outstanding issues of cryptographic trust, data security, governance and compliance that create roadblocks to adoption of the Cloud by the industries that can most benefit from it.

Pairing cryptography also makes possible the world in which a fleet of geographically distributed and organizationally independent Distributed Trust Authorities act as multiple private-key generators (PKGs) where trust need not reside in a single entity.

The difference between this new world of Distributed Trust Authorities and the current PKI system will be a landscape that provides secure ease-of-use encryption and authentication, does not rely upon a single trusted third party, and yet allows for limited key escrow subject to an end customer's requirement.

Milagro

The Milagro libraries and tools consist of:

- Distributed Key Management Service API
- Distributed Key Management CLI
- Software Defined Distributed Security Module (SD-DSM) build platform
- Distributed Key Management Endpoints (software)
- Crypto Apps, consisting of:
 - M-Pin Authentication Platform (delivering password-less 2FA)
 - M-Pin Secure Channel (delivering certificate-less TLS-PSK)
 - M-Pin-in-Mobile Client Libraries for iOS, Android and Windows Phone
 - M-Pin-in-Javascript Libraries for Browsers
 - Cloud Encryption Gateway (under nascent development)
 - Distributed Trust Authority Crypto App
 - Generic library for IoT cryptographic library

The startingpoint for these is the existing MIRACL library and tools at <http://github.com/Certivox/>

Distributed Trust Authorities

The Milagro project introduces a service concept called a Distributed Trust Authority, to replace either single-authority certificates or public key infrastructure.

The D-TA splits the functions of a pairing-based key generation server into three services issuing thirds of private keys to distinct identities. The shares of the private keys, received by Crypto App clients or Distributed Key Management Endpoints, become the only entities that possess any knowledge of the whole key created from the shares.

To effect anything resembling a root key compromise that can occur in a traditional PKI or commercial certificate authority, ***ALL*** Distributed Trust Authority servers must be compromised. Cryptographically, one compromise of a Distributed Trust Authority does not yield an attacker any advantage, all Distributed Trust Authority master secrets inside each D-TA providing shares must be compromised. Note that all 3 D-TA's operate independently and are under separate organizational control.

For the following examples, envision a Distributed Trust Authority model consisting of Cloud Provider (D-TA 1), Cloud Provider end customer (D-TA 2) and neutral third party (D-TA 3).

Under this three participant model, where each member is responsible for the security of their D-TA, the Cloud Provider can not subvert the security of the end customer, even with the collusion of the neutral third party. The end customer will not suffer an internal insider attack unless the Cloud Provider and neutral third party also collude.

Distributed Key Management API, CLI, Endpoints

The core infrastructure that consumes these thirds of private keys and is responsible for their distribution is a message bus and API (D-KMS API), a command line interface (CLI) and software (D-KMS Endpoints) which builds the Crypto Applications from source.

Any entity can run any mix or combination of components with other entities, but there is no restriction on configuration. One party may operate all three D-TAs, Endpoints and APIs if they wish.

The D-KMS CLI communicates securely with the API. The API is responsible for either creating cryptographic keys and secrets or protecting existing keys and secrets through cryptographic encapsulation, via a choice of pairing-based protocols. In either case, the API encapsulates the keys and secrets for the identity of particular D-KMS Endpoints.

The D-KMS Endpoints are server operating systems with D-KMS Endpoint software installed. The D-KMS Endpoint software, in conjunction with the D-KMS CLI, has the appropriate pairing-based cryptographic keys to be able to de-encapsulate secrets and keys received from the D-KMS API. These de-encapsulated secrets and keys can be stored, distributed or used in Crypto Applications, such as M-Pin Authentication, Secure Channel or Encryption Gateway.

SD-DSM / Crypto Applications

Software Defined Distributed Security Modules, otherwise known as Crypto Applications "Crypto Apps" get compiled from source files on-demand. Crypto App source files will be hosted on major public repositories such as Github and Apache.

Crypto Applications are scaled across the datacenter through the D-KMS API in conjunction with orchestration tools such as Apache Mesos and consume the de-encapsulated secrets and keys.

M-Pin Authentication and Secure Channel

M-Pin is already deployed by such organizations as NTT and Experian in a two node Distributed Trust Authority model, where MIRACL and its customer each host a D-TA node. In Experian's case, M-Pin was selected to provide authentication for Experian's identity assurance platform, contracted to the UK Government, for secure authentication of online citizens into UK government websites, including HMRC (tax office). M-Pin was selected based on its security efficacy and ability to scale to an Internet scale user population (UK online citizenry).

The M-Pin Authentication Platform serves as an example of what is possible exploiting a pairing based protocol. M-Pin is capable of running in a native browser mode, delivering two-factor authentication. M-Pin binds to any identity (as long as it is worldly unique) and improves the user authentication experience as it can be visualized in a familiar ATM-style pin pad.

It's most unique trait is the exploitation of zero knowledge proof authentication. The M-Pin Client proves to the M-Pin Server it possesses its cryptographic authentication key without revealing it to the server. As a result, the M-Pin Server stores no authentication credentials, eliminating the possibility of credential (i.e., password) smash n' grab attacks.

M-Pin Secure Channel extends the protocol to include authenticated key agreement between server and client and mutual client-server authentication. The 'agreed key' is unique for each session, possessing perfect forward secrecy.

M-Pin Secure Channel takes the agreed key and injects the key into a TLS-PSK session between client and server, providing mutual authentication and perfect forward secrecy without the need for PKI. This cryptographic underpinning can be extended to create secure VPN sessions over various protocols.

In an M-Pin client and server context, clients and servers receive their shares of their private keys from all three Distributed Trust Authorities. In the previously mentioned example, this could be Cloud Provider, end customer and neutral third party or any combination thereof.

M-Pin Client and Server code are already open source, having been previously released under BSD-Clause-3.

The next iteration and revision will be licensed under the Apache License.

Cloud Encryption Gateway

Many proprietary solutions have appeared on the information security market to solve data governance issues about securing data in the cloud with encryption keys managed by an end customer. To date, most of these solutions involve purchasing hardware or virtualized appliances to run in an end customer's datacenter, with nothing more delivered than a single encryption key under control of the end customer, performing sub-optimum deterministic encryption on data sent to the cloud.

The Milagro Cloud Encryption Gateway will be a virtualized or container based software, deployed in an end customer's environment. This CEG will exploit pairing-based capabilities such as attribute-based encryption (anyone in possession of the correct set of attributes can decrypt) and, more generally, predicate-based encryption (anyone in possession of the right set of attributes and a decryption key corresponding to a particular predicate can decrypt).

Doing so increases the flexibility of the solution by being enabled to address data residency and governance requirements such as geo-location while allowing key management and rotation protocols to be enforced.

Rationale

The benefits of a strong authentication, secure channel and cloud encryption via an identity framework for people and things are self-evident, and the plethora of homebrew proprietary solutions and password nightmares seen today is clear evidence of a need for better solutions.

Milagro's distributed trust model is particularly attractive, by virtue of dispensing with need for (and potential for abuse of) any central trust authority without requiring sophistication - such as understanding a Web of Trust - from end users.

A move to incubation at Apache will help the community to grow and take on new members in an environment that guarantees open development and protection of participants.

This is particularly relevant right now as a second corporate team, NTT Data, with its own culture joins as core developers. For the outside world, it offers the strong promise of openness.

Initial Goals

Milagro will seek to integrate the existing projects at Certivox (now MIRACL) and NTT, and will invite participation from a nascent broader community evidenced by the core MIRACL library's 65 watchers and 29 forks at Github.

As well as looking to broaden direct participation, it will seek synergies with relevant Apache projects, for example by providing Milagro plugins for HTTPD and Trafficserver.

The initial software products will be the current standing M-Pin Core platform, client libraries and the SD-DSM and Distributed Key Management API and client CLI (as noted above).

Current Status

Certivox (now MIRACL) has developed open source software at Github since 2014, though the core MIRACL library goes back much further. Projects currently at Github include the M-Pin Authentication Platform and the MIRACL cryptographic libraries under BSD-Clause-3 and AGPL licenses.

These have attracted both community and corporate interest taking them beyond the realm of a single-company project, with NTT being the second corporate team to take a substantial part in development. The project now seeks to transition smoothly to a full Open Development model.

The core team at Certivox (now MIRACL) is geographically dispersed and developers are well-accustomed to using online infrastructure and tools for their everyday work. The team at NTTi3 and NTT DATA and other contributing developers are included amongst the initial committers.

In addition to MIRACL operating a community D-TA, NTT, Experian and Dimension Data have all agreed to host no-charge community D-TAs. Other cloud providers are considering and have been engaged. An open source platform from which to offer these services is a necessary component to finalizing and launching community D-TA's.

Meritocracy and Community

The project is moving from a single (startup) company open source project seeking a wider community, to embrace a second corporate development team and third-party developers. The project is committed to broadening the community through meritocracy, and expects to welcome contributions and recognize contributors.

It is hoped that incubation at Apache will help with this broadening, by providing a widely-recognised and well-understood framework for working collaboratively, growing communities, and protecting contributors.

Core Developers

Dr. Michael Scott, Chief Cryptographer at Certivox (now MIRACL), has been a major open source and standards contributor to the field of elliptic curve cryptography for over twenty-five years.

Others include

Existing team at Certivox/MIRACL:

- Patrick Hilt - CTO
- Kealan Mccusker - Cryptographer
- Stanislav Mihaylov - Architect
- Simeon Aladhem - Developer

Existing team at NTT:

- Go Yamamoto - Cryptographer
- Kenji Takahishi - Developer

Existing ASF Member:

- Nick Kew - Developer

Alignment:

Whereas Milagro has no track record of its own, the Certivox (now MIRACL) team have been working on related projects at Github. Being geographically diverse, the team is well-accustomed to day-to-day working in a similar environment to Apache and with similar tools and processes. The anticipated role of Apache is to help the community to grow without fragmentation of communities, code, or intellectual property.

We are not aware of any link with existing Apache projects. However, it is likely that several Apache projects may be interested in working with Milagro to provide distributed identity services. Plugins for HTTPD and Trafficserver are already anticipated.

Known Risks

Orphaned products

Milagro, as successor to the existing MIRACL and M-Pin software at github, is at the core of Certivox (now MIRACL)'s business and important to NTT, Experian, and other platform adopters who are in the process of coming online.

Interest, and with it both developer and user communities, are expected to grow strongly. There is little risk of the project losing momentum in the foreseeable future.

Experience with Open Source

The software has a history as open source, developed until recently by a geographically distributed team within a single company. Github activity shows some evidence of a wider community. The major new development that leads the proposers to seek incubation at Apache is the coming of new corporate interest: while both corporate teams have open-source experience, their cultures and backgrounds differ.

We hope that incubation at Apache may help the teams collaborate in an environment of mutual benefit, as well as attract independent developers to play a full part.

Homogenous Developers.

The established corporate teams are dispersed across several European countries and Japan. Prospective developers (whose companies are interested in Milagro) are located in other countries, and we anticipate a global community.

Reliance on Salaried Developers

Most of the initial committers are salaried developers from the core corporate teams. Github activity, including 29 forks of the Miracl library, indicates wider community interest, and it is hoped that the developer community will grow substantially at Apache.

Apache Brand

The Apache brand is of course seen as an advantage. However, the project is more directly concerned with the Apache platform and environment to unite diverse teams.

Relationships with Other Apache Products

See Alignment above.

Documentation

Milagro derives from Certivox's existing M-Pin, MIRACL and associated tools at github.com/Certivox/ Documentation at <http://docs.certivox.com/> may also inform and feed into the Milagro project.

Initial Source and Intellectual Property

As soon as Milagro is accepted into the Incubator, Certivox (now MIRACL) will transfer the source code and trademark to the ASF with a Software Grant, and licensed under the Apache License 2.0. Certivox/MIRACL retains rights to its existing MIRACL mark.

External Dependencies

There are no external dependencies and all software is under the sole ownership of Certivox/MIRACL.

Cryptography

This is advanced cryptographic software, and as such may be subject to government interest and red tape in some countries. However, the architecture by which SD-DSM / Crypto Apps are distributed, via open source freely available code repositories, is intentional to exploit the near universal interpretation of the Wassenaar agreement to permit export of open source cryptography without restriction (in most cases).

Required Resources

Mailinglists:

- private
- dev
- users

Git repository (to mirror existing github repo)

- <https://git-wip-us.apache.org/repos/asf/incubator-milagro.git>

Issue Tracking

- JIRA repository to be requested

Trust Authority Service

The podling would like to request a VM at "ta.milagro[.incubator].apache.org" with which to run a Community Trust Authority. It is anticipated that this will serve as a test facility for developers and may become a Trust Authority for the community of ASF committers.

Initial Committers

- Akira Nagai (NTT)
- Brian Spector (Certivox/MIRACL)

- Fuji Hitoshi (NTT)
- Genoveffa Pagano (Certivox/MIRACL)
- Go Yamamoto (NTT)
- Jordan Katserov (Certivox/MIRACL)
- Kealan Mccusker (Certivox/MIRACL)
- Kenji Takahishi (NTT)
- Michael Scott (Certivox/MIRACL)
- Milen Rangelove (Certivox/MIRACL)
- Mitko Yugovski (Certivox/MIRACL)
- Michael Scott (Certivox/MIRACL)
- Nick Kew (Apache)
- Nick Pateman (Certivox/MIRACL)
- Patrick Hilt (Certivox/MIRACL)
- Simeon Aladhem (Certivox/MIRACL)
- Stanislav Mihaylov (Certivox/MIRACL)
- Tetsutaro Kobayashi (NTT)

Sponsors

Champion

- Nick Kew

Mentors

- Sterling Hughes
- Jan Willem Janssen
- Nick Kew

Sponsoring Entity

- The Apache Incubator