ECC in Action

real-world applications of elliptic curve cryptography

The Certicom 'Catch the Curve' White Paper Series

September 2004

certicom™

The *Catch the Curve* White Paper Series

When your cryptography is riding on a curve, it better be an elliptic curve.

Are you riding a crypto roller coaster? Does your ride involve adding strong security to constrained devices? Are you faced with the tradeoff between product differentiation and profit margins? If so, find out about ECC, the next generation public-key cryptosystem.

ECC provides you with:

- longer running battery operated devices that produce less heat
- software applications that run faster and take up less memory
- scalable cryptography for the future

Read the Catch the Curve white paper series to find out why the NSA, Research in Motion, Motorola and other leading organizations have adopted ECC.

The series in detail:

The Certicom Catch the Curve white paper series includes three white papers detailing various areas of ECC.

- White paper 1 provides the foundation for understanding ECC, its strengths and advantages. Available now.
- White paper 2 provides real-world examples of ECC applications, discussing how organizations are using, and benefiting from ECC today. Available now.
- The final white paper in the series concludes with an ROI study on implementing ECC. Available mid-Fall 2004.

For more information on the white paper series, Certicom or our products, please contact Wendy Bissonnette at +1.613.254.9258 or wbissonnette@certicom.com.

www.certicom.com/catchthecurve

certicom™

Making the Right Choices

In securing communications technologies, there are two practical realities: space and time—and neither are infinite.

In fact, in many cases both space and time are quite constrained. In small, portable devices there is little space for the digital circuits that perform cryptographic operations. For remotely-operated devices, speed is a factor as the limited power available to such a device may constrain the amount of time it can devote to cryptographic operations. Added to these challenges is the ever-present expectation for faster and more efficient communications. The result is that the time allocated for encryption is sometimes measured in milliseconds.

These resource constraints have important ramifications for the design of cryptographic systems, equipment, and protocols. However, if the cipher is chosen well, not only can encryption and decryption be executed in very little time; it could be computationally infeasible for hackers to attack.

Asymmetric Cryptography – What It Is and Why We Need It

As discussed in *An Elliptic Curve Cryptography (ECC) Primer*, the first white paper of this series (www.certicom.com/catchthecurve), asymmetric cryptography is an essential family of technologies. The basic, critical feature of asymmetric cryptography—the ability to verify another's identity without gaining the ability to impersonate her—is such a useful property it's widespread in many diverse applications.

Elliptic curve cryptography (ECC) is an efficient method of asymmetric cryptography that offers equivalent security to conventional asymmetric cryptosystems but uses considerably smaller keys. Consequently, ECC can significantly mitigate the demands on chip real estate, processor time, and communications bandwidth made by the need for asymmetric cryptography.

In this paper, the second in the Certicom 'Catch the Curve' white paper series, we will describe how the primitive operations provided by ECC solve real world problems in protocol and system design. We'll talk about the use of elliptic curve-based digital signature schemes in check verification systems and in digital postage marks used in postal system. We'll describe the use of elliptic curve-based key exchange systems in secure web servers and in small embedded control and monitoring devices. Finally, this paper will discuss how certain applications, being planned for the near and distant future, would not be possible without the proper cryptographic implementation.

The Primitives

Cryptographic primitives are the building blocks of cryptographic protocols. The five fundamental primitives are:

- **(1) Encryption and (2) decryption**
- **(3) Signing and (4) signature verification**
- **(5) Key negotiation and exchange**

Encryption is the generating of an opaque ciphertext—a scrambled message from the plaintext. It is a message the user wishes to transfer privately. In asymmetric systems, the user uses the public key and the encryption algorithm to generate the ciphertext from the plaintext.

Decryption is the reverse of encryption—the recovery of the readable plaintext from the opaque ciphertext. In asymmetric cryptosystems, the intended reader of the message uses the private key and the decryption algorithm to recover the plaintext from the ciphertext.

Signing or ***signature generation*** is the procedure whereby a user takes a message (to be signed) and her private key, and produces a special bitstring called the digital signature for that message.

Signature verification is the procedure whereby a user takes a message (to be verified), the digital signature, and the signer's public key, and the output is either 'Signature Valid' or 'Signature Invalid'.

Key negotiation and exchange is the generation—usually with the assistance of asymmetric cryptographic algorithms—of keying material for use in other algorithms. Key negotiation schemes permit parties who, through other means, are already able to verify one another's identity to agree upon keys to be used in symmetric algorithms for later communications. Good key negotiation schemes have the property that they generate appropriate, secure keys, in such a fashion that they do not become available to third parties monitoring the exchange.

These five main primitives enable developers to secure communications and ensure that the four main objectives of information security—confidentiality, integrity, authentication and non-repudiation—are met.

In the remainder of this paper, we will discuss how the elliptic curve-based algorithms offer better versions of each of these primitives, and where these faster primitives can be used.

certicom™

Elliptic Curve Digital Signature Algorithm (ECDSA) and Elliptic Curve Pintsov Vanstone Signature (ECPVS)

Faster Certificates and Signatures

A certificate is a digitally-signed "document" that ties an identity to its public key. An authority (such as a government passport agency or corporate identity server) signs the certificate, establishing that the identity is valid. By verifying the authority's signature on the certificate, the verifying party understands that the identity and public key can be trusted.

The fundamental property that knowledge of the public key—and thus the ability to verify a signature—does not confer knowledge of the private key—and thus, as mentioned, the ability to generate a signature—is what makes asymmetric signature algorithms so useful. Once trust is established the two parties can exchange information (for example, to make a purchase online) or grant access to resources (such as login to a network), or allow admission to a location (for example, immigration at an airport).

Documents or digital files signed by asymmetric digital signatures are provably authentic, unaltered and are not forgeable. These attributes are important in a wide variety of applications, such as the high-volume, high-security business of financial processing.

ECDSA and ECPVS Are Elliptic Curve Signature Schemes

ECDSA has been widely standardized by accredited standards organizations including the American National Standards Institute (ANSI X9.62), the US government's National Institute of Standards and Technology (FIPS 186-2), and the International Organization for Standardization (ISO 15946-2). ECPVS is being adopted into a number of standards including ANSI X9.92 and ISO 9796-3.

ECDSA and Check 21

The United States banking system clears checks by exchanging the actual paper checks. The presenting bank sends a check to the bank that has to honor the debt; the bank honoring the debt disburses the funds.

In the search for ways to cut costs and improve the efficiency of this process, the "Check Clearing for the 21st Century Act" was introduced. Known as Check 21, the Act allows the industry to use a new digital method of presenting checks. In fact, it is estimated that by moving to digital images or paper substitutes that are printed when and where needed, the financial industry could save two billion dollars a year by eliminating paper-based financial instruments such as checks and reducing transportation costs.

However, before reaping this financial reward, the industry must first clear several hurdles; the most significant being quality and security of the image. Signing the images with a digital signature can guarantee the integrity of the image against alteration and allows for authentication. Integrity can be accomplished in two steps: first, by creating a hash value for the image or "hashing" it, and second, by digitally signing the resulting sequence of bits. By using digital certificates, which are provisioned by a trusted third party or the banking institution itself, an ECDSA signature enables non-repudiation: indisputable proof that the institution claiming to have created the check substitute from the original check actually did.

For any Check 21 plan to be successful, it will need to rely on standards-compliant solutions that consider image quality, usability, and image security. ECDSA fits that bill. It has been endorsed by The American National Standards Institute (ANSI), the de facto authority in this regard, to provide the security and efficiency required by applications that convert, process, and store check images for the financial industry.

This particular algorithm was chosen because ECC offers high security using much smaller key sizes than any other known public-key cryptographic scheme available today. This mix of strength and efficiency makes ECC ideal for applications targeting the high security, high volume needs of the financial industry.

Through its exceptionally small key sizes, ECC provides extremely fast digital signing operations—current industry requirements dictate that a solution digitally sign up to 10,000 images per minute. As ECC is the only technology capable of meeting this requirement, major check transport vendors are already using ECDSA.

ECDSA can be used in a range of other applications, including where written signatures have normally been used. Examples include online stores and corporate approval systems, and any application that requires that documents be guaranteed for integrity, authenticity and offers non-repudiation.

And the uses of digital signatures extend beyond the normal realm of signatures. ECDSA can be used to sign digital photographs in such a way that the photograph can be proven to be unaltered. For forensic purposes, this could prove to be critical: a photograph or recording could be admitted as evidence in a hearing, with the assurance that the evidence has not been tampered with or modified in any way.

ECPVS and Digital Postage Marks

While ECDSA offers performance improvements and reduced bandwidth use in signature generation and verification, for some applications ECPVS can provide even better results.

ECPVS is a particularly concise digital signature scheme used in digital postage marks, the digitally signed printed "stamps" used in many metered mail systems. DPMs help to prevent fraud by embedding the paid amount, source and destination address, among other information, ensuring that the proper postage was paid for at the meter.

ECPVS signatures have two advantages over ECDSA:

1. **Signatures can be made even smaller than ECDSA signatures.**

2. **The signature can be designed so that the message itself or portions thereof can be recovered from the signature. This is useful for storing postage details inside the printed stamp.**

An ECDSA signature in a DPM system adds approximately 40 bytes to a signed message and offers the equivalent security as a 1024-bit RSA signature. Use of ECPVS can reduce this appendix still further—to as little as 20 bytes—less than one-sixth the size of an RSA signature.

As mentioned above, ECPVS is also very flexible allowing certain parts of a message to be recoverable from the signature: which parts is determined by the system design. Various parts of the postage mark can also be made available to different parties. The actual security level and signature length of ECPVS is also flexible. A system design can, in very controlled fashion, trade security levels against the bandwidth/envelope space available. But even at the most demanding levels of security, ECPVS is dramatically smaller than RSA.

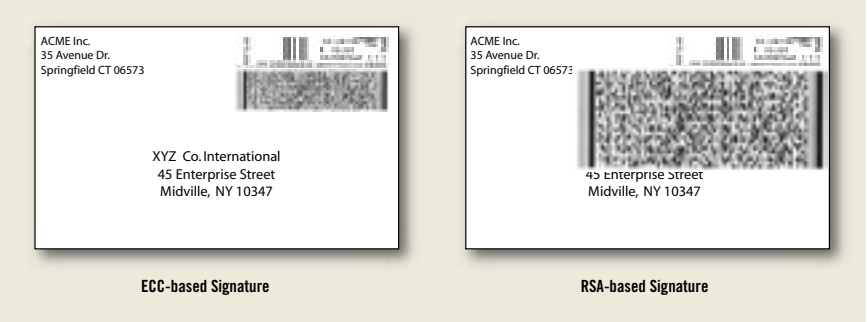

ECC: ROI for Digital Postage Marks
6 times smaller signature with ECPVS

In addition to mail systems, ECPVS can be used to limit forgery in a variety of other applications. Passports and national identity cards can benefit immediately from the use of DPMs. By applying a digital signature to verify an identity's credentials (including a photo), any copies, forgeries or modifications would be immediately detected. Applying an ECPVS mark to a movie or sports ticket or to a subway or bus access card means that copying and alteration of the ticket would render the altered copy invalid, just as the postal system can recognize altered postage marks.

Elliptic Curve Diffie-Hellman (ECDH) and Menezes-Qu-Vanstone (MQV) versus Diffie-Hellman (DH)

Key Exchange for 21st Century Protocols

The DH key exchange protocol is currently one of the most pervasive asymmetric method for generating keying material between parties communicating over an untrusted communications medium. It's conceptually quite simple and well-studied.

In DH, each party generates a key pair—one private, and one public. Each sends the public key to the other, and each party uses the public key received from the other, together with the private key to generate the shared key.[1]

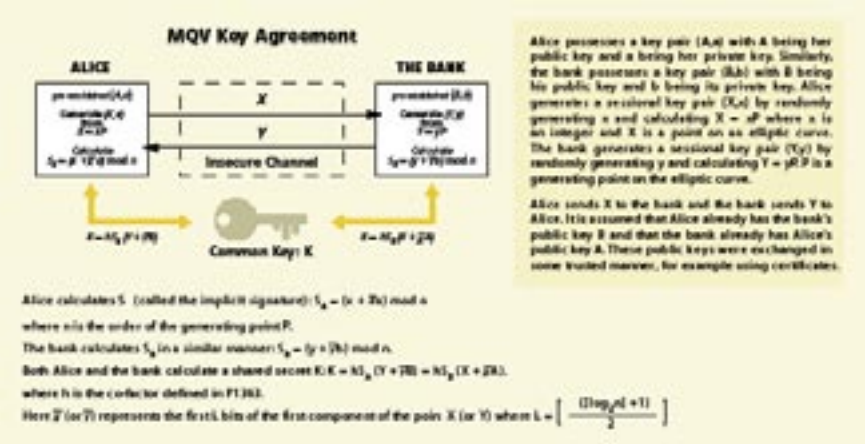

Like all effective key exchange protocols, DH is constructed so that a monitoring party cannot derive the key agreed upon. The public values exchanged by the communicating parties aren't helpful without the private values kept by each communicating party.

But traditional DH over finite fields—the dominant method used until now—has several disadvantages relative to newer elliptic curve-based methods.

[1] The public values are public El Gamal keys—g^x and g^y respectively; the private keys are x and y; the parties calculate $(g^x)^y$ and $(g^y)^x$ to generate the shared key. Since a third party sees only g^x and g^y, they cannot themselves calculate g^{xy}.

One such pressing disadvantage is that it is becoming unwieldy to generate keys long enough for the higher strength symmetric algorithms now in use. To generate a key long enough for 256 bit AES (advanced encryption standard), the communicating parties would each have to generate a 15,360-bit key pair. That can be quite computationally demanding—especially on constrained devices—and significantly increases the bandwidth overhead in the protocol.

Similarly, in relatively low bandwidth communications channels—such as the 9600 bps GSM data service currently in use[2]—sending two 15Kb messages just for key exchange is inconvenient, and close to impractical. For secure web servers that may be less constrained in terms of bandwidth, this type of computation is still impractical because they are expected to process many secure exchanges per second. The additional processor time and storage needed for larger keys can generate significant costs in hardware upgrades.

ECDH in TLS/SSL and IPSec

As has been discussed above but is worth repeating, elliptic curve operations enable you to use smaller values to attain equivalent security. For example, a 512-bit ECDH value is equivalent in security to those massive 15,360-bit DH or RSA values and is secure enough to support generating 256-bit keys for strong AES.

ECDH has been tested in the TLS/SSL protocol family, the encryption layer used to secure web traffic, and has been proposed as an extension to the TLS standard to improve the efficiency of secure web servers. The potential for elliptic curve methods to reduce costs and increase performance in this area of practice is particularly dramatic. For more information on this, see the third white paper in the *Catch the Curve* series to be published in the fall of 2004.

ECDH is being used in IPSec-compliant secure virtual private network (VPN) systems. The benefits for IPSec VPNs are parallel to those of web servers—gateways and clusters of gateways serving large numbers of clients save significantly on processor time, storage, and bandwidth by using the EC groups. For the organization employing the system, the result is the same as for web servers —less hardware servicing more connections.

Elliptic Curve Menezes-Qu-Vanstone (ECMQV)

Additional Security Features with Shorter Keys

As stated, ECDH offers equivalent security with smaller keys. Another refinement—the ECMQV protocol—offers an additional security feature over both DH and ECDH.

As mentioned, DH schemes, for all their utility, do have disadvantages. In traditional DH exchanges, there is very little to guard against a third party doing exchanges with the two original

[2]The CSD, or Cellular Switched Data service

parties, effectively inserting themselves into the secure channel. However, if an independent means of verifying one another's identity is available, such attacks can be prevented.

A unique feature to ECMQV is that a stored copy of the other party's public key is used directly in the generation of the shared secret. The result is a tighter integration between authentication of the other party's identity and the key generation algorithm.

In a sense, ECMQV is key generation and identity verification in a single step. While the identity of parties generating keys with DH, along with the keys generated, can be verified after the fact using various asymmetric methods, ECMQV does this in the same step as key generation. Consequently, ECMQV reduces processor overhead and bandwidth, and makes it possible to prevent man-in-the-middle attacks from even proceeding through the key generation phase.

In addition to this increase in efficiency, ECMQV also prevents a party, which has compromised a communicating party's private key, from thereafter masquerading as a third party to the compromised party. And, ECMQV is still a highly efficient protocol, making very modest computational demands.

ECMQV has been proposed and accepted as a method of key exchange in a number of emerging standards, including IEEE 1363-2000 and ANSI X9.63, and may well be the method of key exchange for the future, particularly in constrained devices such as smart cards, pagers, cell phones, palmtop computing devices, and compact digital music devices.

ECMQV is also in use in encrypted transport mechanisms in audio/video consumer electronics, and may potentially be used in protocols such as ZigBee. ZigBee is an efficient protocol for use in constrained monitoring control devices, such as thermostats and other remote sensors, which may need to run for years on a single battery, or operate on the current from compact photovoltaic solar cells.

In addition, ECMQV is also specified in the NIST Special Publication 800-56 "Recommendation on Key Establishment Schemes (Draft)" that includes key establishment mechanisms recommended for use by the US federal government.

ECMQV for WiMax (802.16)

WiMax is a new wireless standard providing high-throughput broadband connections over long distances. Speeds envisioned are up to 75 Mb/sec, while coverage may be as much as 30 miles from a base station.

It is anticipated that WiMax will be particularly useful in providing 'last-mile' connections to areas

and facilities where appropriate cable infrastructure is not already in place. For cable companies in particular, WiMax is promising because it may allow them to sell internet connectivity more competitively to businesses—currently an issue because at many businesses, television cables are not already in place. WiMax will also enable cable companies to sell to residential users in areas not served by a cable network.

WiMax is promising, but implementing it poses many security issues. WiMax is essentially a wireless transport for DOCSIS, the Data Over Cable Service Interface Specification cable companies use to connect cable modems to the Internet. DOCSIS is ideal for the application in terms of bandwidth and medium—insofar as television cable is also essentially an RF medium—but the major difference between the cable network and WiMax is obvious enough: cable networks are open only to those who have a cable installed in their building, while the wireless application envisioned by WiMax is wide open to anyone with a capable transceiver.

If you were to place a base station on an electrical pole somewhere to service an area of several tens of square kilometers, you would need some way of ensuring that the users connecting through the base station are authorized to do so. Furthermore, if the clients connecting to the network wish to have any presumption of privacy—which becomes an absolute must in such applications as voice over IP (VoIP) and online banking—authentication would become a two way exchange. The base station would need to verify the identity of connecting clients and the clients would need to be certain that the base station they are connecting to is not run by someone interested in stealing their credit card numbers.

Added to this complexity is the fact that by 2006 the WiMax standards hope to support roaming between base stations. The obvious solution to this situation is cryptographic methods, certificates in particular, but both the WiMax base stations and clients are expected to be relatively constrained, compact devices. It is anticipated these stations and clients will have to fit into fairly small, weatherproof packages, running on fairly low power. And with roaming, fast handoff from base station to base station also becomes a must. Every microsecond any authentication and key exchange protocol consumes is going to become critical.

ECMQV is an obvious fit for WiMax. A certificate authority system maintained by or for a consortium of WiMax vendors could provide ECDSA certificates to all devices at key strengths equivalent to 3072-bit RSA but at a fraction of the size. Keys in the ECC certificates need only be 256 bits long. At each initial connection and at each handoff during roaming, ECMQV could be used to do the authentication and key generation in a single implicit step. Employed in combination with 128-bit AES, each connection would be highly secure against sniffing and hijacking, protecting both parties in the exchange—provider and client.

ECMQV and WiFi

Issues raised in WiMax are not unique. WiFi faces similar problems and is now infamous for its security issues. Terms like warchalking and pirate base stations entered the lexicon of network security rather quickly after the widespread proliferation of inexpensive 802.11b cards and hubs. More recent 802.11 standards have made inroads in these areas with such protocols as EAP MD5 and EAP TLS, designed to make things more difficult for hackers.

EAP TLS is reasonably secure, but implementing it is fairly demanding in the constrained cards and embedded hubs typical of 802.11 installations. EAP MD5 is less trouble, but it's basically a password hashing scheme and thus is vulnerable to a dictionary attack.

ECMQV is another good fit here. As in WiMax, a relatively small certificate provides highly secure authentication. A single exchange can authenticate the base station and the card to one another and generate session keys to secure the communications channel. As well, the implementation could have a considerably smaller footprint than any method using RSA or DH primitives—a major advantage when the hardware at one end of the connection has to fit into a single-height PCMCIA (Personal Computer Memory Card International Association) slot.

ECMQV, VoIP and the SIP proxies

ECMQV is also ideally suited for securing VoIP connections, currently in a growth phase.

The relatively low cost of VoIP connections versus traditional circuit-switched telephone lines makes it increasingly attractive to residential users and businesses looking to reduce telecomm costs. But since VoIP traffic moves through packet-switched networks the security concerns involved with VoIP connections are markedly different than those with traditional telephony connections.

The solution many businesses and government agencies using VoIP will turn to is secure VPN technology—technology which secures the VoIP traffic at the IP level. But secure VPN connections are cryptographic connections. And in practical terms, in large meshed networks, asymmetric authentication and key establishment methods are often necessary in the initial phases of each connection.

In VoIP, this can be a challenge due to the number of connections that might need to be established, even to process a single call. VoIP relies on the Session Initial Protocol (SIP) to set up and manage the call. SIP is a workhorse Internet telephony protocol that provides for such features and services as call forwarding, caller/callee authentication, invitations to multicast conferences and the like. SIP may also play a role in billing, in some cases. And in the course of a single call, several SIP proxies may need to be involved and connected to one another.

Using EC methods—such as ECDSA and ECMQV—in the VPN technology used to secure these connections may prove essential to keeping the latency down. Again, as in WiFi and WiMax applications, ECMQV can allow the communicating parties to authenticate to one another and establish session keys in a single step—greatly reducing time, bandwidth demands, and processor time demands.

Conclusion – ECC for efficiency and high security

Throughout this paper, we have discussed how ECC makes so much possible. Without ECC, you can find yourself making difficult choices, balancing security needs with keeping the design feasible in terms of design complexity, power requirements, bandwidth requirements, and cost.

As mentioned, ECC enables devices to run efficiently—producing less heat, using less power, and taking up less real estate on the printed circuit board. Similarly, software applications using ECC run more rapidly and make fewer demands on memory. But what may generate even greater interest, are the future applications that will be made possible because of ECC.

As shown in the Check 21, digital postage mark, WiMax, WiFi and VoIP examples, ECC and its various implementations are being used today and are a good fit for many other applications for the near and distant future. With judicious use of ECC, developers won't have to make that difficult tradeoff. They will be able to have a highly secure design and an efficient one—thanks to the smaller key sizes, and smaller implementations provided by ECC.

About Certicom

Certicom Corp. (TSX: CIC) is the authority for strong, efficient cryptography required by software vendors and device manufacturers to embed security in their products. Adopted by the US Government's National Security Agency (NSA), Certicom technologies for Elliptic Curve Cryptography (ECC) provide the most security per bit of any known public key scheme, making it ideal for constrained environments. Certicom products and services are currently licensed to more than 300 customers including Motorola, Oracle, Research In Motion, Terayon, and Texas Instruments. Founded in 1985, Certicom is headquartered in Mississauga, ON, Canada, with offices in Ottawa, ON; Reston, VA; San Mateo, CA; and London, England.

certicom™

Certicom White Papers

To read other Certicom white papers, visit www.certicom.com/whitepapers.

The Inside Story
Many Happy Returns: The ROI of Embedded Security
Wireless Security Inside Out (authored by Texas Instruments and Certicom)
Welcome to the Real World
Sum Total: Determining the True Cost of Security
Current Public-Key Cryptographic Systems
The Elliptic Curve Cryptosystem for Smart Cards
Elliptic Curve DSA (ECDSA): An Enhanced DSA
Formal Security Proofs for a Signature Scheme with Partial Message Recovery
Postal Revenue Collection in the Digital Age
An Elliptic Curve Cryptography Primer
Good Things Come in Space Packages: An Overview of the Certicom Security Architecture

certicom™

Contact Certicom

Corporate Headquarters

5520 Explorer Drive
Mississauga, Ontario
L4W 5L1
Tel: +1-905-507-4220
Fax: +1-905-507-4230
E-mail: info@certicom.com

Sales Offices

Canada

5520 Explorer Drive
Mississauga, Ontario
L4W 5L1
Tel: 905-507-4220
Fax: 905-507-4230
E-mail: info@certicom.com

Ottawa

84 Hines Road
Ottawa, Ontario
K2K 3G3
Tel: 613-254-9270
Fax: 613-254-9275

U.S. Western Regional Office

1810 Gateway Drive, Suite 220
San Mateo, CA 94404
Tel: 650-655-3950
Fax: 650-655-3951
E-mail: sales@certicom.com

U.S. Eastern Regional Office

1800 Alexander Bell Dr., Suite 400
Herndon, Virginia 20190
Tel: 703-234-2357
Fax: 703-234-2356
E-mail: sales@certicom.com

Europe

Golden Cross House
8 Duncannon Street
London WC2N 4JF UK
Tel: +44 20 7484 5025
Fax: +44 (0)870 7606778

www.certicom.com