DIGITAL CERTIFICATE EXTENSIONS:

SHOULD "BASIC CONSTRAINTS" BE MARKED CRITICAL?

Vivek Kaushik,
Principal Consultant
Netrust Pte Ltd

Date: 2nd September 2007

Scope

The scope of this white paper is to discuss the need for a certificate extension in Digital Certificates called "Basic Constraints" to be marked critical. In this paper, we analyze the rational for this criticality flag with respect to verification, and present a case for a verification process in spite of a non-critical "basic constraints" extension.

Terms Used in this White Paper

CA: Certification Authority, an authority that issues digital certificates to entities. These certificates are used as digital identities in electronic transactions such as authentication and digital signatures.

DN: Distinguished Name, the digital identity of an entity or a CA within the trust infrastructure. Typically, the DN consists of the individual's name and affiliated organization within a CA.

Root CA: A CA that provides trust anchor in a certificate validation by providing a self-signed certificate at the top of the certificate chain.

X.509: An ITU-T standard for the format of a digital certificates, based on RFC 3280.

NIST: **National Institute of Standards and Technology,** a US federal organization tasked with creating standards for interoperability. Standards created by NIST are widely adopted by other countries and organization, and are considered industry standards.

CA and Certificate Extensions

A certification authority (CA) is a body that issues digital certificates to entities that wish to engage in electronic transactions. A CA may issues various types of certificates, each designed to conform to a certain usage scenario. For example, a CA may issue "Type A" certificates for encryption only, and "Type B" certificates for both encryption and digital signature. Also, depending on the type of CA service offered (such as subordinate or root), a CA's own certificate is also customized to its own needs.

For the sake of interoperability as well as to ease that the management of various types of certificates, digital certificates are standardized to X.509 (currently version 3), but customized using optional fields called “certificate extensions”. These extensions are binary fields within a X.509 certificates that are industry standard and can attest to the intended profile and use of the digital certificate.

Criticality and Certificate Verification

An X.509 certificate has two types of extensions: **critical** and **non-critical**. The rational for these extensions is derived from the intention of the CA when a relying party verifies its certificates. By marking an extension as “critical”, the CA that issued that certificate is notifying the relying party that the contents of the extension are critical to the verification of the certificate, that is, the certificate can be accepted only if that critical extension can be processed by the relying party’s certificate validation algorithm.

When a CA marks an extension as non-critical, the relying party has a discretion to accept or reject the certificate, without being able to process the critical extension. If the relying party decides to process the extension, typically this result determines whether the certificate is accepted. For example, if the extension is “key usage” and the value is “encryption” then a digital signature using that certificate can be rejected by the certificate validation algorithm (since it is not the assigned certificate usage).

Critical “Basic Constraints” and Commercial CA

The basic constraints extension identifies whether the subject of the certificate is a CA and the maximum depth of valid certification paths that include this certificate. Certification path processing verifies the binding between the subject DN and subject public key. The binding is limited by constraints which are specified in the certificates which comprise the path. “Basic Constraints” is one of the extensions that allows the certification path processing logic to automate the decision making process.

While the industry standard mandates that the “basic constraints” extension be marked critical for CA certificates (Section 4.2.1.10 of RFC 3280), there is no supporting processing requirement when the certificate is verified by the relying party. The closest standard to certificate path validation is also mentioned in Section 6 of RFC 3280. In this certificate path validation, there is no requirement to have the “basic constraints” to be marked critical for verification. From this, it can be concluded that whether the constraint is critical or not, this does not materially affect the RFC-complaint certificate validation.

A commercial CA by definition must strive to maintain the maximum number of supporting applications for its certificates. It is in the interest of a commercial CA to keep a certificate profile with minimum requirements to support current, former and backward applications.

When a commercial CA considers including an extension in a certificate it does so with the expectation that its intent will be adhered to wherever possible. If a commercial CA would flag an extension critical, this would require that any validation engine that could not process that extension will reject the certificate, thus limiting the set of applications that can verify

the certificate and hence supported by the CA. Therefore, a CA may mark certain extensions non-critical to achieve backward compatibility with validation applications that cannot process the extension. Where the need for backward compatibility and interoperability with validation applications incapable of processing the extensions, is more vital than the ability of the CA to enforce the extensions, then these optionally critical extensions would be marked non-critical.

Also, it is noteworthy to analyze the certificate verification process mentioned in RFC 3280 as well as "NIST Recommendation for X.509 Path Validation". When a verification system encounters a critical "basic constraints" extension on the CA certificate, there are two alternatives; either the validation engine can recognize the extension or it doesn't. If the extension can be evaluated, the verification engine will process the value and make a decision based on the result. If however, the extension cannot be evaluated, the value of the extension is immaterial and the verification engine will proceed to other steps to validate the CA certificate. In these steps, the criticality of the "basic constraints" does not affect the validation logic, that is, the fact that the "basic constraint" was critical or not, does not logically affect the validation process.

Real-world Example

To assess the most common deployments of commercial CAs, the author undertook an analysis of the typical distribution of CA certificates in a end-user setting. Windows XP operating system by Microsoft is the most common client in the world. Windows XP contains a list of "Trusted Root CA Certificates" that have passed Microsoft's requirement for providing CA services to its customers. This is the most common list of commercial CAs in a client setting.

A review of the list of root CA certificates within Windows XP reveals that a substantial number of CAs do not follow the "basic constraints" requirement from RFC 3280. Of a list of two hundred root CA certificates, eighty five were found to be non complaint with respect to the RFC requirement for "basic constraints". Of these eighty-five, forty-two did not have the extension at all, and the remaining forty-three did not mark this extension as critical. Among these CAs are some industry established names such as Verisign (13 CA certificates), Microsoft (3 CA certificates) and Cybertrust (2 CA certificates).

Summary

In summary, the "basic constraints" extension is used to perform certificate validation is a trust chain. Though the industry standard requires the extension to be marked critical for the CA certificate, the criticality flag does not affect the certificate verification process even as specified in the RFC. For a commercial CA, a critical "basic constraints" would considerably limit the number of applications supportable for its products. Therefore, a substantial number of commercial CAs including some established CA service providers, do not flag the "basic constraint" as critical, while still providing trusted digital identities.

References

1. “Internet X.509 Public Key Infrastructure - Certificate and Certificate Revocation List (CRL) Profile”, http://www.ietf.org/rfc/rfc3280.txt

2. “X.509 Certificate Path Discovery Test Suite”, Computer Security Resource Center, NIST, http://csrc.nist.gov/pki/testing/pathdiscovery.html

3. “X.509 Style Guide”, Peter Gutmann, http://www.cs.auckland.ac.nz/~pgut001/pubs/x509guide.txt