Jon Graff, Ph.D.
Director of Architecture and Design
Deloitte & Touche Security Services
60 South Market Street
San Jose, CA 95113

April 7, 1999

Information Technology Laboratory
Attn: AES Candidate Comments, Room 540
National Institute of Standards and Technology
100 Bureau Drive, STOP 8970
Gaithersburg, MD 20899-8970

Dear Information Technology Laboratory:

The following are my comments on the selection of the Advanced Encryption Algorithm Process. They are based on my analysis of submitted data and information gained while I attended the Second AES Candidate Conference in Rome, Italy.

1. I do NOT believe that the selection criteria should be modified. While it is true that more is now known about the prospective algorithms, the rules should not be changed in the middle of the selection process. As part of the submission process, each algorithm creator had the responsibility to balance security against performance.

2. I would desire that only ONE algorithm be selected as the AES. While it is possible to have a multi-algorithm standard, a multi-algorithm will confuse users and serve to aggravate the compatibility problem and possibly slow the implementation of the new standard.

   Only if there is an overwhelming convincing need, should NIST select more than one AES algorithm.

3. I believe that NIST's apprehension against so-called "stealth" patents is a valid concern. These stealth patents are patents that are concealed and only come to light after the selection of the AES. It is my opinion that NIST can mitigate against these stealth patents by prominently posting a notice on the AES website that warns that non-disclosure of possible patent infringement on ANY candidate AES algorithm will be construed as disinterest in upholding the patent. It is widely believed that, according to current patent law as applied to standards, the non-disclosure of patent intent during the standards process is evidence of abrogation of patent interest.

4. The following my analysis of data presented by NIST[i] and Schneier et. al.[ii] on the encryption speeds for the candidate algorithms.

The first figure shows the encryption speeds that NIST reported for the 15 AES candidate algorithms. One can observe that there is a great deal of variability depending on the compiler used. Schneier et.al. made the same observation that an algorithm's encryption speed greatly depended on what compiler was used.

NIST Encryption Speeds with Various Compilers[i]

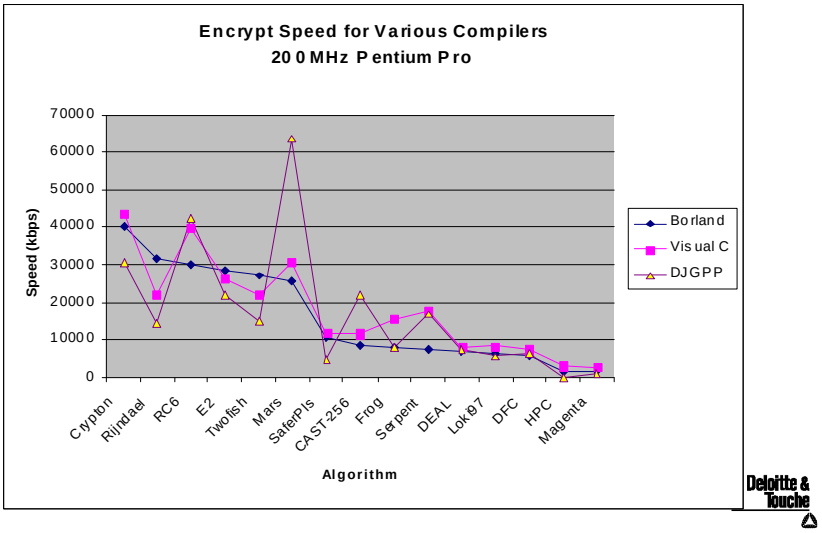

In order to combine the NIST and Schneier et.al. data sets, it was necessary to have a common measurement for speed. NIST presented its data has Mbps and Schneier et. al. presented their data as machine cycles. I chose to "normalize" the data by converting both data sets to a rank ordering of the candidates, with the fastest having a ranking of "1" and the lowest having the ranking of "15".

As shown in the following slide, the combined data allowed for the addition of Schneier et. al.'s dataset to the NIST data. The slide makes evident that the algorithms fall into rather consistant relative rankings. That is, the faster algorithms remain ranked within the faster group regardless of the compiler and the slower algorithms remain ranked with other slower algorithms regardless of the compiler.

Comparative Encryption Speeds

(derived from NIST[i] and Schneier et.al.[ii])

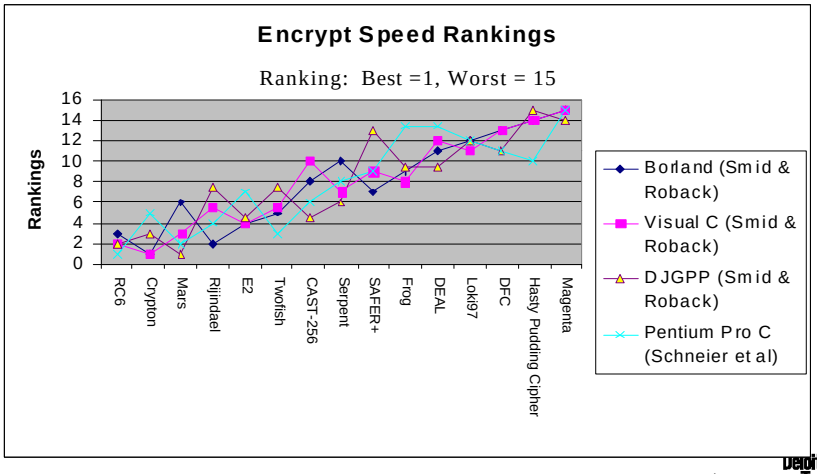

4

At the time I did my analysis[i], the following algorithms had proposed attacks against them: Loki97, Frog, Magenta, DEAL, and Safer+(256). In the following slide these algorithms are enclosed within boxes. All the enclosed algorithms fall in the "slow" group.

Schneier et.al.[ii] did not report more encryption speed data on these slower algorithms.

Comparative Encryption Speeds

(derived from NIST[i] and Schneier et.al.[ii])

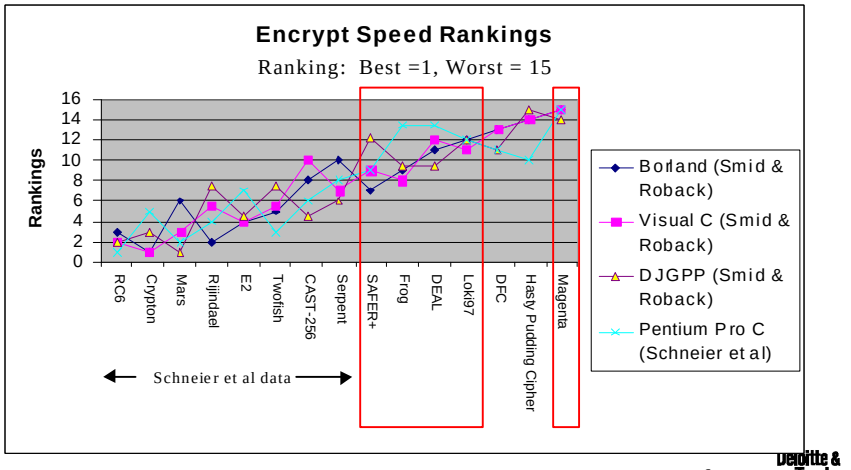

6

The next slide shows all the faster algorithms and compilers that both NIST and Schneier reported. The encryption speeds of 9 algorithms, compiled on 6 compilers

were presented. One can observe that the fastest six algorithms are almost indistinguishable when their speed rankings are shown in the scatter diagram.

Deloitte & Touche

Comparison of Speed Ratings

(derived from NIST[i] and Schneier et.al.[ii])

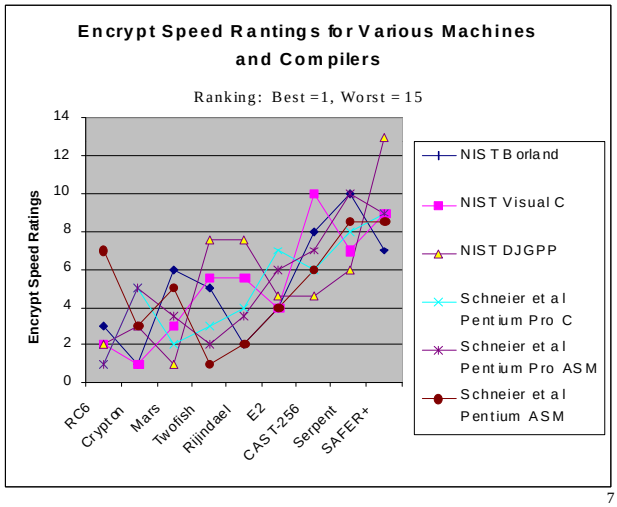

7

Deloitte & Touche

This lack of differences is also demonstrated in the following slide that shows the mean rankings of the 9 candidate algorithms. One should observe that the standard deviation of the rankings is fairly uniform and large, lending credence to the conjecture that the fastest algorithms speeds when compiled with different compilers are almost statistically indistinguishable.

Means and Std. Dev. of Speed Ratings

(derived from NIST[i] and Schneier et. al.[ii])

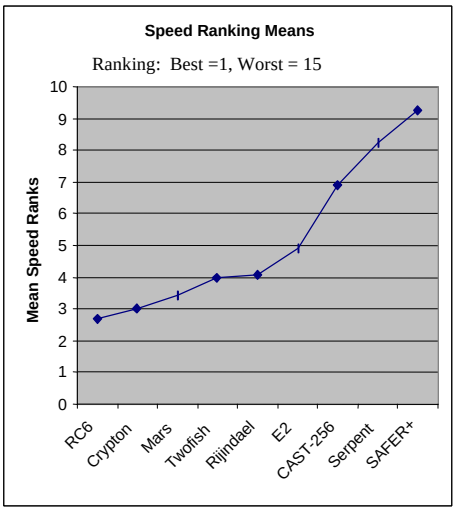

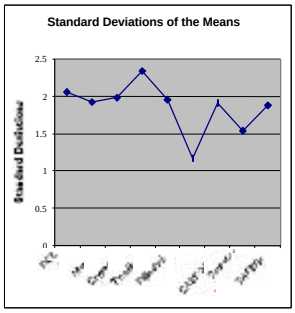

8

One could conclude from the data presented in these two papers that:

1) The algorithms seem to fall into consistent speed groupings. Regardless of the compiler, the faster algorithms consistently fall within the faster group and the slower ones within the slower groups.

2) If the cryptanalysis is to be believed, one could objectively eliminate the "weaker" algorithms, by simply eliminating the slower algorithms.

3) Within the fast set of algorithms, the top 6 do not show significant speed differences. Provided there are no cryptanalytic reasons and these speed results hold up with additional measurement data, measurement and grouping by encryption speed may serve as an objective means to select candidates for the next round.

Sincerely,

Jon Graff, Ph.D.

[i] Miles Smid and Ed Roback, NIST: *Developing the Advanced Encryption Standard*, Presentation at RSA'99, San Jose, CA, January 1999.

[ii] B. Schneier, J. Kelsey, D. Whiting, D. Wagner, C. Hall, & N. Ferguson, *Performance Comparison of the AES Submissions*, posted on the Web