A Cursory Examination of Market Forces Driving the Use of Protection Profiles

Kenneth G. Olthoff
olthoff@earthlink.net

Recently, there has been growing interest in and support for an international framework called the Common Criteria within the Information Assurance community. Much of the U. S. government's Information Assurance strategy for the coming years is based on the assumption of widespread acceptance of the Common Criteria and the National Information Assurance Partnership (NIAP), with positive result to follow. While one can hope that things go as planned, I would like to at least open a discussion on the market forces which may influence the results in unanticipated ways. This is not to say that the future will unfold in any particular way, but merely to caution that we should be prepared for all possible contingencies, even the undesirable ones.

1 Begin with the basics

Let's start with a few terms.

Protection Profile - A reusable and complete combination of Security Objectives, functional and assurance requirements with associated rationale.
Security Function - A part or parts of an Information Technology product or system that is the subject of an evaluation which have to be relied upon for enforcing a closely related subset of the rules from the Security Policy.
Security Policy - A set of rules designed to meet a set of Security Objectives.
Security Target - A complete combination of Security Objectives, functional and assurance requirements, summary specifications and rationale to be used as the basis for evaluation of an identified Information Technology product or system.

Previous efforts in various countries, including the famous "Orange Book" in the U.S., had specific levels or classes of security products, each composed of fixed sets of requirements. These levels were an all or nothing proposition. If a product missed even the tiniest requirement for a given level, they would automatically be dropped to the next level down. There was no opportunity for a customer to indicate that they in fact didn't care about that particular feature, but did care about the rest. The Common Criteria system is an attempt to change this, by allowing the creation of "Protection Profiles" which state the desired requirements of a particular community in almost any combination desired. The Common Criteria also provides for international recognition of the ratings, which becomes interesting in a later discussion.

The Common Criteria/ Protection Profile scheme will be useful only if it is widely used. If the government, or any other interested party, uses this system, they must make the case to both vendors and customers that it will enable the vendors to get more sales; enable the customers to get the products they desire, and a way to make informed purchases; and give the security community a "raising of the bar" in commercial products with investment coming largely from the vendors and customers, rather than being externally funded. In a less obvious detail, the system must ensure the existence of an entire support system, consisting of testing labs, researchers, regulators, Protection Profile writers, and the maintenance of the Common Criteria and its processes.

This assumes that all the parties will have an incentive to participate. and that they will get the "bang for the buck" that they desire. From the government Information Assurance perspective, even full and active participation from the other parties is not success if the net costs of the program do not result in better product, used by a wider range of the target customer base (military, government, and critical infrastructure), than could be achieved by other means.The cost of meeting the residual needs must also be considered. If the Protection Profile scheme is wildly successful by all other measures, but leaves a residual need for higher level products for 50% of the target customer base, it may be more costly overall than a program which is less successful in the general marketplace, but which increases the coverage of the community of interest to 90%,

NOTE: *The views expressed are strictly those of the author, and do not reflect the official policy or views any organization, government, association, religious body or school of economics. Flattery not included. Void where prohibited by law.*

1999 New Security Paradigm Workshop 9/99 Ontario, Canada

file scheme. It remains to be seen if any of these factors will be major components of the cost-benefit calculation, or if they are small enough to be ignored in a reasonable approximation.

With that in mind, let's look at some of the factors which will dictate participation and success for the various parties.

2 The Customer

The Common Criteria/Protection Profile scheme holds out the promise of many benefits for the customer, including the ability to purchase products which have been evaluated to known standards, with the cost spread over the entire market, or perhaps partially underwritten by the vendor, and the ability to tell vendors what products or features are desired. In reality, market effects and other circumstances may place limits on the effective range of Protection Profile benefit for any given consumer.

As an analogy, a product profile for a vehicle could be written at many levels of abstraction. One could specify such factors as minimum speed capacity, a minimum payload capacity, type of engine or motor, passenger safety factors, seating dimensions, etc., and end up with a profile which could be successfully met by the QE II, a 747, the Mir space station, a 67 VW Beetle equipped for off-road racing, or an Abrams tank. Even if we narrow things down to specify a four cylinder automobile to be used as a commuter vehicle for driving to work, going shopping, and running errands, it is still a balancing act to come up with a profile which is generic enough to be met by several competing products, while still providing adequate product differentiating information to the customer.

At the same time, the customer also would like the scheme to have the effect of encouraging competition, innovation, and improvement among the vendors. Unfortunately, we find a counterexample in sports car racing. Different competition classes based on production vehicles exist, and the rules of each are written to allow the use of multiple different makes and models within the class. It rapidly becomes apparent, though, which car provides the best performance under that set of rules, and in many classes, the competitors all end up driving the same make and model of vehicle, even though others are allowed. There is usually one which meets the rules, or the Protection Profile, in a way clearly superior to other available products.

What the customer is unlikely to get is a "Consumer Reports" type rating where one product is compared head-to-head against another product in each specific aspect of the Protection Profile under the same conditions by the same test team. Instead, the customer is likely to get more of a checklist, where each stated requirement of the Protection Profile is either met or not met, with some amount of expository information conveying the details. These details, however, are unlikely to be stated in precisely the same way from test to test, and many of the features are likely to defy easy comparison. The Protection Profile and testing are also unlikely to provide specificity in areas which may be vital to some situations.

In our car example, if our customer community lives in a small mountain town located 400 miles from the next town, they may have some very specific requirements - ability to operate the engine properly in high altitudes, significant resistance to brake fade, gear ratios favorable to driving on steep inclines, higher reliability, and the ability to go over 400 miles on a single tank of gas. Our hypothetical auto profile would not be likely to cover these specific needs, though, because they are not generic enough to represent a customer community large enough to be a market driver.

The Protection Profile process effectively becomes, for most communities, an analogue to a car being approved for sale in the U.S., which requires meeting safety, emissions, and design standards. This in itself provides no product differentiation, however, because the standards are a minimum, and everybody meets them. Even this analogy falls down, though, because as yet there is no mechanism to enforce adherence to a given Protection Profile.

What will drive customer purchase of Protection Profile-meeting products? It may be that particular industry groups may specify compliance for interoperability reasons. It is possible that insurance firms may begin giving risk avoidance discounts and better rates on insurance for loss of data to those who use such products. Nobody knows the minimum size of a community in terms of numbers, market share, or dollar value of potential sales represented which will drive a vendor to meet a given Protection Profile. It is also unclear how many Protection Profiles a vendor will be willing to test against, though the number will probably be very low, and will be limited to those with the broadest market coverage.

A Protection Profile is likely to reflect compromises within the community which produces it. A given community may trade off some aspects of its ideal Protection Profile against the economies of scale represented by a more generic Protection Profile. The more specific, rigorous, unusual, or expensive a product must be to meet a Protection Profile, the less likely it is that vendors will produce a product to do so without significant additional incentives.

On the other hand, a more generic Protection Profile may have many factors which are not of interest to a specific community. This may result in one or more possible non-optimal situations. The first is that the community may end up paying more for products with features beyond their needs. Let's say the widely used Protection Profile calls out Security Functions A,B,C,D, and E, while the particular community only needs A, D, and E. The vendors all produce products to meet the Protection Profile by including all five Security Functions. This particular user community may be paying extra for features B and C, where they might otherwise have been able to purchase a product with only A, D, and E. This is only a problem, though, if the economies of scale do not reduce the price to a point where the community is paying the same to get all five Security Functions that they would otherwise have paid for just the three Security Functions.

That said, there is still some benefit to the customer communities and vendors in simply having a commonly agreed upon language and format in which to express requirements. This is a benefit which should not be overlooked, as specification difficulties have

historically been a source of a great deal of problems in many fields. It is also possible that the Protection Profiles may serve an education purpose. In the example above, the customer community which did not desire B and C might find those sections in the Protection Profile, and be led to re-examine whether they need those Security Functions, or if they would benefit from having them as "nice to have", rather than mandatory, requirements.

3 The Vendor View

The vendor wants product differentiation, which leads to sales and market share. Economies of scale may reduce costs, preferably by more than the reduction of price, thus increasing margin. The vendor also wants clues as to what features the market may want in future products. While the vendor has the opportunity to profit greatly from innovation, there is risk in doing so. Given these constraints, vendors are almost certain to closely examine the market benefits of responding to Protection Profiles, and various strategies will emerge.

Early adopters will hope to gain a head start in the process and market recognition if the idea catches on. Others will wait to see if a significant market develops. A few will engage the process, but only at favorable cost-benefit points. Some will ignore the process altogether, and there may even be those who will deliberately target the markets unlikely to be served by the Protection Profile system, or unlikely to use it. All will be attempting to highlight their strong points.

At least one instance of the early adopter strategy already exists. The company is working closely with key customers to generate a Protection Profile which plays to the strengths of the vendor's product. The vendor is expressly trying to position this Protection Profile as a de facto standard in a particular market. This is a viable strategy, and it benefits both the vendor and the collaborating customers. It remains to be seen, though, if the end result will benefit a wide range of customers, and the effect on competition. In the past there have been many de facto standards which were not technically superior (Beta vs. VHS, Mac vs. PC). Protection Profiles may enable this phenomenon even further.

Another strategy is for the vendors to play along only at minimal cost, designing their products to meet the most generic Protection Profiles, or ones which their products would meet without any further effort on their part. This gives the benefit of the broadest applicability, while not incurring any additional development costs. While this is attractive for the vendors, it tends to drive the vendors to the least common denominator, rather than toward meeting specific customer needs.

In any field, there is a sort of bell curve, with a center peak containing most of the products available. The extremes of the curve on the one side included the older or inferior products, while on the other side are the innovative products using new technology or even whole new theories and models of behavior. Market forces will tend to focus vendor effort on meeting the Protection Profile, but not exceeding it by too much. The customer is far more likely to be able to understand a "meets the Protection Profile" test result, than to do a comparison and trade-off among competing products to see which of the products is superior. While the Protection Profile feature set is supposed to be a minimum, it will tend to become an informal maximum as well.

Let us take a variant on the example we used previously. Let us say that three widely used Protection Profiles in a particular market segment exist. Protection Profile 1 calls out Security Functions A-E, Protection Profile2 calls out Security Functions A-M, and Protection Profile 3 calls out A- H and R-V. Reality is likely to be much more complicated, but to make the model simple, we will assume that each Security Function is independent of all the others, and all Security Functions are equally difficult or costly to implement.

If a vendor wishes to play attention to Protection Profiles in this market, it would be useful to at least meet Protection Profile 1. Future expansion plans might include going after either Protection Profile 2, or Protection Profile 3, but perhaps not immediately both, due to the divergence of Protection Profile 2 and 3. The more Security Functions are added, though, the more the cost increases. It is much more in the vendor's favor to exceed Protection Profile 1 by a small number of Security Functions, than to miss meeting Protection Profile 2 or 3 by a similar amount, because the "standard" testing will focus on existing Protection Profiles. If the vendor tests to Protection Profile 1, any features beyond the requirement of that profile will go untested. If the vendor has the product tested against Protection Profile 2 or 3, the additional features will be tested, but the product will fail to meet the full requirements of the more stringent Protection Profile. The economics and strategy of the situation will tend to cluster products at or just above the levels laid out by the Protection Profiles in existence, while discouraging expansion beyond those clusters in anything other than a step function up to a more rigorous Protection Profile.

In our example, Protection Profile 2 and Protection Profile 3 share Security Functions A-E with Protection Profile 1, and additionally share Security Functions F-H with each other. It might seem logical for a vendor to implement Security Functions A-E as a first step (Profile 1), then
Security Functions F-H as a next step (because those functions are common to Profiles 2 and 3), but unless a new Protection Profile is written which specifies Security Functions A-H, the intrinsic value of those Security Functions F-H will be hard to market. They are not covered in Protection Profile 1, but testing to Protection Profile 2 or 3 will be more costly, while still resulting in a "failing" grade. Making the addition of F-H the next step is thus attractive only if the decision of whether to aim toward Protection Profile 2 or Protection Profile 3 has not been made. Once that decision has been made, the order in which the additional Security Functions are added matters very little. What does matter is that the vendor is driven to avoid incremental improvement.

It should also be noted that there is little incentive for the vendor to work at all on Security Functions W, X, Y or Z in our example. Being the first one to research new Security Functions in the context of a Protection Profile environment incurs both cost and risk. Research is likely to be slowed until a market and/or Protection

Profile appears. The cost of generating a Protection Profile is likely to limit incremental change, and the market forces will tend to limit the speculative research by the vendor within the Protection Profile system. Given the typical life span of de facto standards, it's quite possible that once a core set of "Good Enough" Protection Profiles for products of interest are in existence, creating support for a new, improved, Protection Profile will be sort of like swimming in molasses. This of course ignores the effects of any markets outside the Protection Profile system, which may temper these effects under real market conditions.

To return to the car analogy, the Protection Profile scheme may be the equivalent of Henry Ford's assembly line - the idea which moves the industry from semi-custom production for a specific customer to a commodity product with a few standard models. The good news is that there may be lots more security functionality, even if only at a rudimentary level, on the infobahn. The bad news is that the industry may go to the same extreme which Henry Ford went to at first - your security box can be any color, as long as it's black, and have any features, as long as they are standard. This bucks prevailing trends in much of manufacturing, where flexible automated production and "just in time" inventory allow greatly increased customization of products for specific customers. It also limits the tailoring of security solutions to specific cases.

If my analysis is correct, the vendor has built-in incentives in the Protection Profile marketplace to adhere to Protection Profiles which are easy to implement and have wide acceptance, regardless of appropriateness for a particular customer set, or the actual security provided. The vendor should not be assumed to have any motivation to improve security, to point out weaknesses, or to incorporate incremental advances in their products, except in cases which create product differentiation or enhance the vendor's reputation in ways which outweigh the costs. In this marketplace, the main thing that matters is meeting the Protection Profiles which provide the most market access for the least effort.

Lastly, the vendor has a clear incentive to carefully plan the submission of a product for evaluation, and the advertising of the outcome. Recent history has shown that it is possible to truthfully say that a product has been evaluated and received a given rating or certificate without divulging the fact that the rated version has since been superseded.

The vendor also has incentive to find a testing lab which will be as favorable as possible to the product, which the lab may do for any of a variety of reasons. The vendor may have corporate interactions or alliances with the testing lab's parent company in other areas. The preferences and biases of particular testers may incline them to look upon the vendor's product favorably. Any given tester cannot be expert in all fields, and some testers may miss subtleties which a more experienced person would at least question. An additional factor is the matter of international reciprocal recognition of Common Criteria evaluations. A given nation's system may favor a given type of solution for dogmatic, economic, or political reasons.

The standard testing processes will be followed, and the questions will be asked. There will always, however, be differences in how efficiently this is done, how well the need for further inquiries is recognized, and how vigorous those follow up inquiries are. It it well within the rules for the vendor to do such venue shopping, and the results will be presented to potential customers in the most favorable ways possible. Caveat emptor.

4 Testing, Testing - One, Two, Three...

Moving on to the testing labs, we get a slightly different perspective. The testing labs will have little direct input into Protection Profiles, but they may generate feedback to those writing Protection Profiles about particular difficulties in testing a particular Security Function or feature described in the Protection Profile. As with all the other parties, the testing labs benefit if there are fewer Protection Profiles which are widely accepted, and if they are relatively generic. The testing lab also has an interest in standardizing the tests as much as possible, to increase efficiencies.

It is not in the best interest of the lab to dig extensively into the product beyond the superficial pass, fail, or measure aspect of the testing. They are unlikely to have a close relationship with the customer, because the vendor will be paying for the testing, and will control the distribution of the results. The vendor is probably not going to agree to the release of bad news. For conflict of interest reasons, there will probably be a strict limit on the feedback the testing lab will be allowed to provide to the vendor, as the lab cannot be assumed to be objective when a product comes in to be retested which incorporates design features previously suggested by the lab. It is likely that to preserve the air of impartiality, a separation of the design and evaluation functions will develop, as has been seen in the past history of computer security. The strictness of the limitations will fluctuate periodically, but it is reasonably safe to assume that the lab's function will evolve to a "bring 'em in, check 'em out, and push 'em through" production, with little extraneous postulation or probing. It is not in the lab's economic interest to be sloppy, but neither is it in their interest to do more than is absolutely necessary to verify the vendors claims of meeting the Protection Profile. The sooner that testing can be made a commodity service, the better, from the lab's point of view. It is in the lab's interest to advocate for any change to the system which will result in more products being tested to fewer, simpler, and more generic Protection Profiles.

In short, the testing labs are the vehicle inspectors of the Common Criteria scheme. They don't set the standards, they don't fix your car if there is a violation, and they don't do the work to develop next year's model of any brand of car or truck. They also are unlikely to suggest aftermarket products to improve performance or appearance. They are there to do the tests, check the boxes on the forms, provide cursory notes detailing the areas in which the test subject is deficient, if that is the case, and issuing certificates to the owners of those that pass. Like the vehicle inspectors, they will be licensed by a government authority, and they will have to periodically prove they are competent to be licensed to do inspections, but the reliability and accuracy of any individual inspection they perform will be largely dependent on the personnel involved. Market factors will encourage them to do only what's necessary, with the most cost-effective personnel and equipment possible, and the

regulatory factors will encourage them to adhere to the minimum standards.

5 Profiles as a Cottage Industry

There also needs to be a brief note about the support structure of consultants, training, Protection Profile and Security Target writers, and others which will spring up surrounding the Common Criteria scheme. Many vendors, customers, and regulatory groups will not wish to maintain an internal capability to deal with Common Criteria and Protection Profile issues, and will farm out the actual work to others. As in any contract relationship, the incentive for such service providers will be to make life easier for themselves, all other factors being equal. It is unclear how much overhead will be incurred by this service sector, how efficient it will be, how competent it will be, and what the net effect on security will prove to be. Certainly, the participants will be guided by the wishes and motives of their customers, but they will also have a self-interest which should not be overlooked as a factor in the development of the Protection Profile industry.

6 The Government Perspective

Lastly, let us look at the Protection Profile scheme from the security perspective, as represented by the interests of the government Information Assurance community. As advocates for Information Assurance, the government security community wishes to "raise the bar". It has been postulated that Protection Profiles of some sort will be levied as a minimum requirement on suppliers of security related equipment. The general plan, as I understand it, is to use approved Protection Profiles as an initial filter, with no products considered or recommended which have not been evaluated. The goal is a standardized group of security products which are an improvement, both in inherent quality and in our understanding of them, over that which would be available without the Common Criteria and Protection Profiles.

Those working on the Information Assurance mission and their counterparts in other nations also have an incentive to attempt to drive the standards toward constant incremental improvement. Unfortunately, the incentives already noted will be driving the customers, vendors and test community toward a more generic least common denominator solution, with only infrequent jumps to another, slightly more rigorous plateau of security and quality. It is unclear what mechanism, other than education, the IA community has available to drive the market toward their goal in cases where that goal is in opposition to the market forces. The Government IA community is still recognized as a leader in the security field, but it is now "a" leader - one among several, rather than "the" leader, as in the past. Even if the government experts remain the most knowledgeable in the field (which may no longer be the case in some technologies), there are now others on the scene with significant expertise. Private industry, and even some of the traditional government customers, may embrace the free market sources of knowledge, and any attempts to enforce governmental will by fiat will be tempered by the very cross-recognition that the international Common Criteria partners have set up. If any one nation decides to take their ball and go home, their absence may serve to reduce their influence on the global security marketplace even further.

The Common Criteria/Protection Profile scheme does have advantages for government Information Assurance policy makers. If it becomes widely accepted, it will tend to standardize the language of security, at least in the areas covered by the Common Criteria. It will also tend to give economies of scale, either through competition among interoperable products, or through market domination by one product fitting a given Protection Profile. It will allow government experts to focus on the gaps left uncovered by market forces. It will also provide an educational and public relations vehicle for raising security awareness.

Potential negatives for the Information Assurance mission also come with the package. Increased standardization tends to make for a larger and more attractive target. If an adversary can break a specific product used by a large percentage of the market, or can detect a fundamental flaw in a Protection Profile which most products are built to adhere to, the efficiencies of scale work for the adversary as well. Market forces will tend to push the community toward generic solutions, which means that any given user is less likely to have solutions tailored to their specific needs. The Protection Profile scheme also continues a focus on security as a measurement which is made, a certificate that is issued, or a rating which is given, rather than a total way of thinking and a continuous mode of behavior. If the Common Criteria or a Protection Profile is emphasized out of proportion to the operational and educational aspects of the solution, certified products may become an end in themselves, not a starting point or tool in a continuous process of vigilance.

Of course, each participating nation also has an interest in the Common Criteria from an intelligence collection and/or law enforcement perspective. The potential influences which might be brought to bear, and the strategies behind them, will likely vary more from country to country than is true on the Information Assurance side of the problem. A discussion of the differences between various nations in how the intelligence, law enforcement, Information Assurance, and private industry functionality are divided, and the resulting inter-relationship of the various interested parties, is worthy of an article in itself. The effect of those factors on each nation's Common Criteria strategy is far beyond the scope of this paper. Contemplation of these matters, and their impact, is left as an exercise for the reader.

The government's Information Assurance team is in the odd position of simultaneously being a manufacturer of high performance, limited production products, and also taking on the role of setting standards, administering inspection and licensing programs, and directing research. A potential problem is that these functions are being planned in a way which seems to assume a mature industry, when in fact, the industry and the customer base resemble more closely the automobile culture circa 1909. The available products are still temperamental. The general public does not know how to operate them properly, or even grasp the basic terminology. Good technicians are needed not only for constant maintenance and adjustment, but often for parts fabrication from raw materials.

These technicians are scarce, yet they are generally unrecognized or under compensated for their talents. The laws governing the operations of the products are still largely chaotic, and reflect a significant lack of understanding on the part of legislators. Insurance companies are only beginning to think about the implications of the new technology. And most importantly, the era of long term total maintenance packages to ensure trouble-free operation are a very long way away. This is not to say that this industry cannot or will not evolve differently. It is merely a caution that we may have a lot more work to do than we think, especially in the area of creating an educated consumer, practitioner, designer and operator base.

This paper has attempted to summarize the market forces which have been unleashed within the framework of the Common Criteria initiatives. It is not meant as a prediction of failure. To the contrary, I believe this scheme may possibly be made to work for the benefit of all parties. What I hope I have conveyed, however, is a significant note of caution. I believe that this scheme is far from a "turn it on and forget it" solution. If our aims are to be met, we must monitor the progress of the scheme and actively work with all parties to keep the various driving forces in proper balance. Without significant understanding of the economic, political and technical drivers, and their effect on all the players, the scheme will devolve to a point where it continues to function, but returns little or no security value compared with other possible uses of the same time and resources. Such a system may be kept alive by forces unrelated to security, but its security value and relevance will have been lost.

Unfortunately, I am not confident that there is an adequate understanding of the workings of a market economy at either the macro or micro level among the government advocates of the Protection Profile scheme. I am not convinced that even a theoretically optimal effort would be sufficient to garner the results currently being predicted in some quarters for the Common Criteria/Protection Profile plan. The current strategy may be the best choice, yet still not be a good choice. It only makes matters worse when policies and strategies which rely so heavily on private sector market forces are made by persons who have little or no experience of any kind in private sector, market driven circumstances. Those in the government with oversight and strategy responsibility are metaphorically attempting to coach a game many of them have never actually played, or even watched closely.

6 The "To Do" list

Additionally, further work is needed in understanding the following areas:

The incentives and motives driving the various parties' approach to security research.

The strategy for the appropriate government bodies to attract, maintain, and effectively apply expertise in all areas of security, given reduced budget, the outsourcing of much of the research and design work, and the limited evaluation work which will be retained, once the commercial testing labs are fully operational.

The strategy for meeting the residual needs of customers not met through Protection Profiles.

The precise means by which the government can effectively exert influence on the direction of the commercial use of the Common Criteria and Protection Profiles.

Improvement of the customers' actual administration and operation of systems. The best tools are of little help if they are not operated properly, and in an appropriate configuration.

The legal, economic, and societal forces driving security policies, and the value of security.

The various forces which will drive participation in the Common Criteria/Protection Profile scheme, and the degree of that participation.

Predictions of the percentage and range of customer needs which will be met by the scheme.

The market and regulatory forces motivating the use of Common Criteria certified products as building blocks for solutions to specific problems.

As can be seen from this list, there is an extensive amount of work to be done, even to generate discussion on as basic a level as this paper. In the cases where such work has already been done, it needs to be much more widely disseminated, discussed, dissected, and debated.

Report on the discussion of "A Cursory Examination of Market Forces Driving the Common Criteria"

Kenneth G. Olthoff
olthoff@earthlink.net
kolthoff@radium.ncsc.mil

"In the beginning was the beard, and the beard was Marv [1], and the beard was with Marv, and the beard was good. And lo, the commandments came down from on high, and the commandments were Orange, and the Commandments were good. Can I get an A1?"
<silence from the audience>
"No, and that's the problem. For it's known that we cannot surf both GOTS and mammon, and mammon has the market share, and so it was decided to consort with the gates (and windows) of industry, which is how we got to where we are today[2]."

With this somewhat paraphrased sermonette, delivered in the style one might expect from a fire and brimstone evangelist, Kenneth Olthoff summed up in highly abridged form the history of the computer security marketplace. Having gotten the attendees' attention with his unconventional approach, he then kicked off a discussion of economic issues surrounding the market acceptance of the Common Criteria and its supporting structure of Protection Profiles, Security Targets, and Evaluations.

Mr. Olthoff did not discuss the intrinsic value of the Common Criteria, or any other similar attempt to develop a market for security. Instead, Mr. Olthoff set forth the need to analyze whether the economic model of the Common Criteria will influence the various parties to behave in the desired fashion. Mr. Olthoff's analysis attempted to show that while the outcome was far from certain, a case could be made that the market might in fact lead users, vendors and evaluators to cluster on a few generic, and therefore less effective, Protection Profiles. Mr. Olthoff noted that this varied from the stated intent of creating a market for products that more closely addressed the needs of customer communities. Once the original position was laid out, the discussion opened up, involving most of the attendees.

1999 New Security Paradigm Workshop 9/99 Ontario, Canada

One of the first counter examples raised was the idea of small companies serving niche markets. The various vendors putting out industry-specific applications templates for databases and spreadsheets were offered as examples of instances where the market did not behave as predicted by Mr. Olthoff's analysis. Mr. Olthoff freely acknowledged that his analysis could be incorrect, and that the example was a very viable counter-argument.

Mr. Olthoff indicated that the main goal of his original submission was to get people to consider the economic influences on behavior. He attempted to clarify that the accuracy of his own analysis was of secondary importance, and that given his background and the limited amount of effort put into his analysis, it was assumed that a more skillful investigation of the issue was needed.

The discussion then headed in the direction of open source software, and whether the Common Criteria and similar schemes might provide a vehicle by which open source software might gain a foothold in the security community. While the attendees all seemed kindly disposed toward open source software, a brief discussion led to the conclusion that there were no inherent economic advantages or disadvantages that would lead open source software to fare differently from proprietary software in a marketplace governed by the Common Criteria.

Another topic that arose multiple times during the discussion was a comparison of the Common Criteria to ISO 9000. It was mentioned that in both cases, there is perceived value, but that the generation of paperwork required may add little value to the overall usefulness. It was pointed out that both ISO 9000 and the Common Criteria emphasize specific documentation in a rigorously specified format.

An additional note in the comparison was the difference between the ISO 9000 model and the Common Criteria. It was brought out that one part of becoming ISO 9000 certified is that a firm must have only ISO 9000 certified suppliers. Thus, the bigger firms become accomplices in spreading ISO 9000 to their suppliers,

[1] The reference is to Marvin Schaefer, one of the principal authors of the Trusted Computer System Evaluation Criteria, AKA "The Orange Book", and the possessor of a distinctive beard of the type made famous by the band ZZ Top. Mr. Schaefer was present, which was why he was singled out among the Orange Book authors. No disrespect to the other Orange Book authors or their beards (or lack thereof) is hereby expressed or implied.
[2] If memory serves, the actual wording included references to "making pacts with Beelzebill" and "the gates (and windows) of hell" – ED.

who spread it to their suppliers, etc. The attendees agreed that it would be difficult to spread the demand for the Common Criteria by similar means, given that there is not a hierarchical relationship between CC vendors, users and integrators. It was also noted that while ISO 9000 has been successful in the marketplace, other government instigated mandates such as GOSIP and "C2 by '92" were unsuccessful.

Another question discussed where the true benefits of the Common Criteria might lie. One opinion was that the value and success of the Orange Book was unrelated to, and unaffected by, the underlying economic model, but was instead based on capturing and conveying the state of the art at the time to a wider audience. This brought a response expressing concern about the quality of profiles and evaluations under the Common Criteria, since the Protection Profiles and Security Targets against which evaluations will be done may not be vetted adequately for security value and appropriateness. By contrast, the formulations of the various ratings in the Orange Book went through rigorous peer scrutiny for many years. It was also pointed out that the Common Criteria scheme allows one to separate assurance inherent in the design and development process from the strength of the mechanisms, while those two factors were coupled in the Orange Book.

Getting back to the non-technical drivers, a question was raised as to what factors might drive the acceptance of security products in the marketplace, whether under the Common Criteria scheme, or otherwise. The answers offered included legal liability, insurance requirements for security to gain favorable rates on insurance against loss, guarantees, auditors, and actuaries. One interesting observation was along the lines of "After all the Y2K lawsuits are over, those computer-literate lawyers will be looking for places to put their knowledge to use." There seemed to be consensus among the attendees that some mechanism is needed to create and enforce liability and responsibility for the consequences of security failures. Whatever the mechanism might be, it should apply to those operating the systems, and those designing and selling them.

The general conclusion seemed to be that eventually, security would need to be mandated, either by private means, such as trade associations or the insurance pricing structure, or through government legislation. There seemed little confidence among attendees that security would be a pull function where users demanded it, but that it would instead be a push function, where other agents levied a requirement for security on the users. There were some comments implying that such a push would only work when the awareness among users was sufficient to not actively oppose the imposition of security.

While there seemed to be sufficient interest and opinion to continue discussion of both the specifics of the Common Criteria scheme, and the general concepts of economic forces influencing the security marketplace, the time limitation on the session brought the discussion to a close.