Risk

Summary: an attempt at evaluating risk based on valuation and policy.

The IETF says that risk assessment is:

> A process that systematically identifies valuable system resources and threats to those resources, quantifies loss exposures (i.e., loss potential) based on estimated frequencies and costs of occurrence.

In 1977 at the NCC Robert Courtney published "Security Risk Assessment in Electronic Data Processing Systems", where he calculates risk in the same basic manner as it is done today – estimated value times probability of failure. $100 million in exposure times a 1 percent probability of catastrophe equals $1M of risk. You can either accept the risk or try to implement safeguards to reduce the risk. This is still the basic philosophy or methodology that is most commonly used today.

An ACM report (also in 1977) on Mr. Courtney's new approach said that "both the assignment of dollar value […] and the estimation of the frequency of losses are very imprecise activities.", and went on to say that Courtney himself suggested that "more precision is not only practical, but may be outright undesirable – as this may convey an impression that the methodology is more scientific than it really is."

> I often say that when you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind; it may be the beginning of knowledge, but you have scarcely in your thoughts advanced to the state of Science, whatever the matter may be.
>
> Lord Kelvin, [Popular Lectures and Addresses, 1883-05-03]

While we are going to try to measure and calculate risk, our knowledge is perilously close to the meager kind. We are not going to come up with a reasonable number that equates to the dollars (simoleans, whatever) that is at risk for an organization. However – I think we *could* come up with a number that indicates a tendency of risk, and give some very rough dollar estimates. And especially when tracked over time and used as a *relative* intra-organizational number this could be very useful.

We already have an estimate of value, in "The Dynamic and Depreciating Value of Computers". And the compliance of a computer is at least an indicator on how likely a problem will crop up – at least, if the policy is sound.

This means we can proceed with our risk calculations.

Organizational Value

In its "Security Risk Management Guide"[1] Microsoft puts it well:

> **Quantitative Risk Assessment**
>
> In quantitative risk assessments, the goal is to try to calculate objective numeric values for each of the components gathered during the risk assessment and cost - benefit analysis. For example, you estimate the true value of each business asset in terms of what it would cost to replace it, what it would cost in terms of lost productivity, what it would cost in terms of brand reputation, and other direct and indirect business values. You endeavor to use the same objectivity when computing asset exposure, cost of controls, and all of the other values that you identify during the risk management process [….]
>
> There are some significant weaknesses inherent in this approach that are not easily overcome. First, there is no formal and rigorous way to effectively calculate values for assets and controls. In other words, while it may appear to give you more detail, the financial values actually obscure the fact that the numbers are based on estimates. How can you precisely and accurately calculate the impact that a highly public security incident might have on your brand? If it is available you can examine historical data, but quite often it is not.

Just like with the value of computing, our risk calculation has analogous adjustments for "overall network/computer worth" (except in this case it's "overall network/computer risk"; see "The Dynamic and Depreciating Value of Computers" paper for more on these.) This adjustments must be applied to the risk values found.

Overall network/computer worth

An user has the option (in the configuration section or some other area) to assert that their entire risk is X dollars. The calculations are then performed for all individual systems as per usual, but the individual values are adjusted by dividing the asserted risk by the observed/calculated risk.

For instance, if the risk was estimated to be worth $100M by the user and we found a total risk of $50M, you would multiply the value of individual systems by:

> 100M /50 M = 2

Obviously this could elevate or diminish the risk of both host groups and individual hosts.

---

[1] http://www.microsoft.com/technet/security/topics/policiesandprocedures/secrisk/srsgch02.mspx

Risk Calculation

We'll use the previously described "Loss Potential = Estimated Frequency * Cost Occurrence" calculation.

As an estimate we'll use the calculated value of the individual systems, along with any modifiers (as described above.)

Policy is an indicator of problems, not a true predictor or likelihood. But lacking anything else you might assume it correlates to the likelihood of bad-things-happening. So for a simple transformation we'll divide the lack of compliance as calculated by a native agent by 10 to get an estimated probability. So an agent that has a 75% compliance (e.g. a 25% lack of compliance) will have a estimated frequency of problems of 25/10 or 2.5%.

This has the effect that even if a network were atrociously managed and maintained and had close to a 0% compliance that only 10% of an organization's value is at risk; as a ballpark this might even be high… of course in the real world no one would have such poor compliance, and I'm not aware of any organization of a significant size have it's value cut down substantially more than that due to any incidents. But this should be, of course, configurable and malleable to the user.

Any hosts that have a compliance score of over 90% still have risk (we might simply not recognize it in our initial crude calculation), so give them the same estimated frequency as a system with a 90% compliance value. E.g. an agent with a 97% compliance value will still have a 1% estimated frequency (it maxes out at 90% compliance, which is 10% non-compliance; 10 % / 10 = 1%).

This probability can be modified by two factors. Both these should be options that a user of the system can turn on or off completely, either via a configuration file or in the user interface.

The first considers hosts that a host trusts. Take the host with the highest estimated frequency of problems (e.g. the lowest compliance) that you directly trust (e.g. a host that you trust, not a host that is trusted by a host that you trust via transitivity), divide it by four and add it to the host being examined. This is done because if you trust a host that is risky your own risk increases. So you might have a host with a 1% estimated frequency; if it trusts another host with a 4% estimated frequency the first host's estimated frequency would be:

$$1\% + (4\%/4) = 2\%$$

To calculate this over a network there is a ripple effect as hosts impact hosts not only next to them but less and less as it moves along. The iteration is desirable because you might be influenced by a host not directly trusted by you but still in the list of your

transitively trusted hosts. The hosts you directly trust will affect you more than those farther away.

The second is a global threat level that the user can manually adjust, on a host-group by host-group basis. There are five levels, vaguely analogous to Defcon levels in the military. Threat level one has no effect on risk. This chart shows the effects of the threat levels:

Threat level	Risk multiplier
One	1
Two	1.5
Three	2.0
Four	5
Five	(User Defined)

A host in multiple host groups takes the highest risk multiplier (i.e. they do not multiply with each other!)

Consider a situation with ten hosts; nine are in host group A and have a cost occurrence of $10,000 one is in host group B and has a cost occurrence of $1 million dollars. Assume an overall threat frequency of 2% for all hosts. The total risk is:

9 * .02 * 10,000 + .02 * 1,000,000 = 1,800 + 20,000 = 21,800

If the threat level in host group B was raised to Level Four, the risk would be:

9 * .02 * 10,000 + .02 * **5** * 1,000,000 = 1,800 + 100,000 = 101,800

Changing threat levels can be useful when highlighting dangerous or temporarily highlighted areas of the network (e.g. the DMZ)

Individual hosts as well as host groups should display a risk number.

Non agent hosts

A host without an agent might not have any compliance rating to perform the basic risk calculation. The average compliance of all agents should be used if this is the case.

In addition the estimated frequency should be multiplied by 1.2, for we believe that running an agent decreases your risk profile.

The Final Solution

There will be a number – dollars, whatever. When pressed (clicked) it must show a breakdown of such things as:

- Number of systems involved (including agents & non-agents) in the calculation
- Raw calculated values of systems and the organization
- Breakdown of costs, from high to low (unadjusted and adjusted)
- Organizational values
- Multipliers/adjustments (threat level, host group modifiers, etc.)

When examining individual hosts they must show a complete breakdown of all the factors that entered into its final risk number.

There should be a report showing the historical values for all of these factors as well as the ultimate risk number. Any changes in any of the adjustments/factors should have associated audit records.