

Prioritising Security Expenditure

Security Zone 2006

Anthony Thorn M.A.
AT Systems & Services GmbH

Slide 1

ATSS

FGSec ROSI WG Results

Although this presentation is based on the discussions of the ROSI WG, it conveys my personal opinions, shaped but not always shared by my colleagues 😄 !

The scope is not just ROSI (whatever that means) but:

How to prioritise security expenditure in a medium-large organisation (not a KMU).

Simplified Scenarios

1. Business Project, e.g. **Internet Banking**, which includes security expenditure required by risk management. “Business enabling”.
2. Security Project, e.g. **Single Sign On**, which can be justified by cost savings (help desk call reduction) without quantifying the risk reduction.
3. Security project, e.g. **SPAM Filter**, justified by a ROSI calculation.
4. Security Project, typically addressing **very infrequent threats**, non-financial drivers dominate. ROSI is useful for optimisation.

Slide 3

ATSS

1 Security expenditure is a by-product of a business project.

Risk analysis of a business project, defines a requirement for additional security measures. The security expenditure is approved as a cost of the business project, so there is no special "approval of security expenditure". e.g. e-banking in the 90s.

The expenditure is justified by the security policy implicit in the risk analysis method.

Organisations with a marketing culture will view these expenditures as “business enabling”.

2. Security expenditure is justified without quantifying risk.

A security project, which can be justified by conventional cost savings, without the need to quantify the risk reduction. The justification is the same as for any other project (i.e. ROI, NPV etc.), and the security benefits (=risk reduction) are considered only as a qualitative benefit. e.g. smart card authentication/Single sign on, which was justified by saving users' time and reducing help-desk calls and headcount.

3. Security expenditure is justified by quantifying risk.

ROSI is applied to justify security expenditure. (E.g. IDS, SPAM filter); -more below about how ROSI is calculated. Depending on the organisation's culture ROSI will be used only in situations where information for ROSI (about incident statistics, impact costs and countermeasure) is easily available or more generally.

4 Security expenditure is justified by non-financial factors.

ROSI is not used either because of insufficient information, or because it is too resource intensive, or simply not required (for strategic, policy, personal, culture or political reasons).

The non financial drivers are relevant. e.g. Regulators, Legislation, Due-Care (see below)

Definitions: Risk, Risk Management

Particularly in IT Security, terms are often used differently by different authors.

Risk

By risk we mean **expected** risk, usually calculated as the product of the cost of an incident and the likelihood or frequency of occurrence.

Annualised Loss Expectancy. An incident which would cost us EUR 100'000 is expected to occur twice per year. This gives an ALE of EUR 200'000.

Risk Management

The purpose of risk management is to understand, document and communicate the risks to which an organisation is exposed so that they can be controlled, transferred, accepted or avoided in accordance with the policy of that organisation.

Slide 4

ATSS

“**Expected**” gives a good indication that we are talking about estimating (guessing) future events...

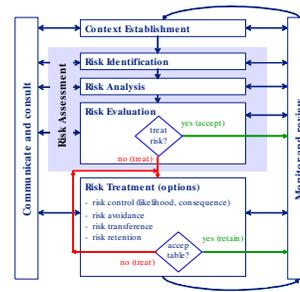
I find the term ALE easier to understand than risk.

In some papers Risk is used to mean threat, and the insurance industry uses risk to refer to risk **exposure** which is the impact or cost of an incident (or of all insured incidents) i.e. ignoring likelihood.

Risk Management is described as “pro-active” as opposed to “reactive”.

Definition: Risk Management Process

is a systematic application of management policies, procedures and practices to the tasks of establishing the context, identifying, analyzing, evaluating, treating, monitoring and communicating risks associated with any activity, function or process in a way that will enable organizations *to control or minimize the risks* defined as “expected losses”. [AS4360]



Slide 5

ATSS

Must be **consistently** applied across the organisation

1. Context Establishment Process
2. Risk Assessment Process
3. Risk Treatment (Options)
4. Communicate and Consult
5. Monitoring and Review

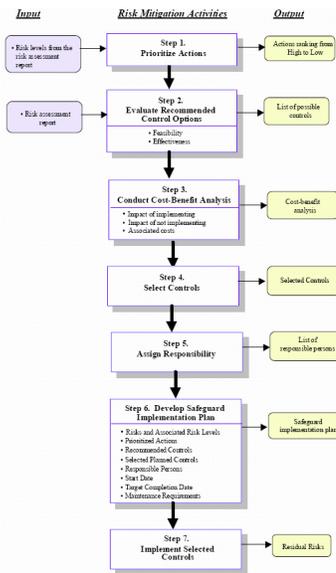
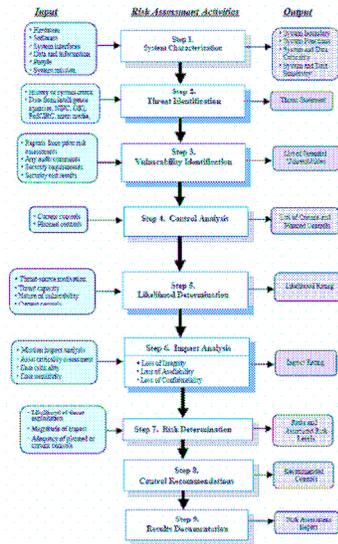
AS4360 is expected to provide the basis for IS27005

High level Risk Analysis – effort required depends on **granularity** – grouping assets is critical

Detailed RA finer granularity,

Tree Models (Risk, Attack...)

NIST SP 800-30



ATSS

See: <http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf>

ISO 27005 "Information Security Risk Management" is expected to be based on BS 7799-3:2005 and the older IS 13335 part 2.

BS 7799-3 Information Security Management Systems - Guidelines for Information Security Risk Management

covers the following:

- Risk assessment
- Risk treatment
- Management decision making
- Risk re-assessment
- Monitoring and reviewing of risk profile
- Information security risk in the context of corporate governance
- Compliance with other risk based standards and regulations

Definition: ROSI

ROSI is normally understood as ROI for security, where ROI is $(\text{benefits} - \text{costs}) / \text{costs}$. ROSI includes Risks in the costs, or Risk-Reduction in the benefits.

We extend the meaning of ROSI to **any financial justification** (e.g. NPV, IRR or Payback time) which includes Risk as a cost or benefit.

By ROSI we mean a cost/benefit analysis where all quantitative and qualitative elements are converted to CHF (EUR or USD).

Slide 7

ATSS

Risk Reduction is the S in ROSI !

Some authors interpret ROSI to include both qualitative and quantitative benefits separately, as opposed to converting everything to CHF (EUR or USD).

You should use whatever formula is standard in your organisation and include risk-reduction.

Return On Investment, Return on Capital Employed, Residual Income, Net Present Value, Internal Rate of Return, Total Cost of Ownership...

Security Benefits / Drivers

1. Financial

This is a financial cost/benefit analysis – ROI or ROSI

2. Address Risks (*like Insurance*)

The risks which we are not prepared to accept.

3. Due-Care

Decisions based on benchmarking & published baselines, also audit reports and external reviews.

4. Legal / Regulatory

Requirements imposed by customers or suppliers also come into this category.

Each of these drivers on its own is potentially sufficient...

Slide 8

ATSS

Financial arguments are not the only ones

1. Financial.

We expect to reduce costs with these measures. I.e. ROI (or NPV or similar). If we include Risk (reduction) in the cost calculation this is ROSI

2. Risk aversion.

We do not want to be surprised by major incidents –disasters or catastrophes (“bumps in the road”). It is important to understand that this is not the same as 1. above. This is a policy issue, and in our model the process is the primary responsibility of Senior Management / Corporate Risk Management. The same motivation as a private person buying insurance.

3. Peer/Due Care

Decisions are based on benchmarking & published baselines. When measures have been adopted by the majority of comparable organisations but not by “our” organisation, the indirect impact of an incident which does occur (even if it is very unlikely) is particularly high – e.g. embarrassing questions by shareholders, if not personal civil or even criminal liability of company officers. The high cost of performing risk analysis for each and every case is a reason for adopting baselines (published or corporate). Audit reports and external reviews come into this category.

4. Legal & Regulatory

Often these factors are simply “no-brainers” but in some cases there is a decision to be made. Also in this category come requirements imposed by customers or suppliers, because they are “costs of doing business”.

Although in many cases more than one of these factors apply, each of these **on its own** is potentially a sufficient justification.

NB. These are only the security-specific factors, the factors common to all projects are omitted

Security

Whatever else you may have been told:

Security **is** managing risks.

Of course it's more than the Risk Management Process, but we "do" security in order to address or reduce risks.

This means that Risk Management (2. Above) is the fundamental motivation for security expenditure.

Impulse

I started the AG-ROSI in order to resolve / discuss controversial (or absurd) published statements like:

"ROSI is easy, you should always do it"

"ROSI is absurd, who ever heard of a ROI for air conditioning"

we soon realised that ROSI and Risk Management were connected, and extended our scope accordingly:

"You cannot measure Risks, Risk Analysis is impossible..."

Slide 10

ATSS

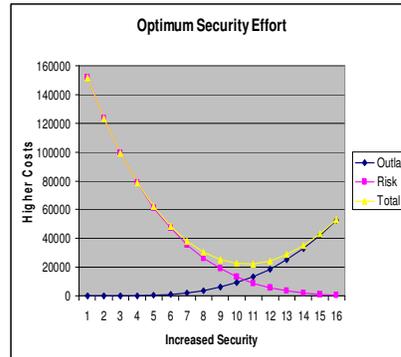
Loosely quoted

Simple Theory I

N.B. Risk is a cost.

If we want to optimise security expenditure we must try to find the point at which increasing expenditure on countermeasures (risk treatment) starts to cost more than the reduction in risk.

This is ROSI.



The lines on the graph do not show the uncertainty in the values

Slide 11

ATSS

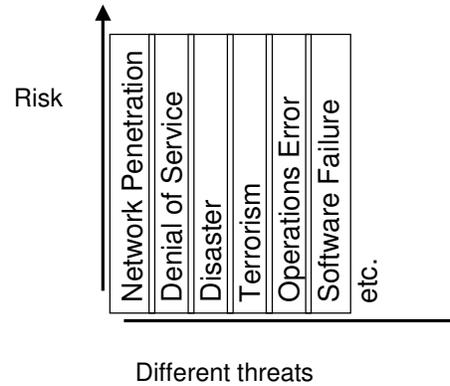
A budgeted cost rather than a bookkeeping entry

Simple Theory II

We need to optimise not only each individual Risk/Risk treatment but across all threats.

This is more difficult and much more onerous.

A “Radar” diagram is often used to show departure from an equable distribution.



Slide 12

ATSS

Diagram is intended to show that we need to obtain **consistent levels** of risk across different threats.

Often shown with RADAR diagram.

Practicalities

Confidence Limits can be very wide (big uncertainty)

Lack of actuarial information we are notorious for not publishing security incidents. Not only the likelihood but also the effectiveness of measures is uncertain.

Indirect costs are uncertain e.g. Loss of profit

Immaterial values are subjective e.g. Reputation

Resource-Intensive

Huge effort is required to perform detailed risk analysis (fine granularity).

Slide 14

ATSS

There are good reasons why organisations do not publish incident information. Often publication would do more damage than the incident itself.

Despite various projects to gather incident information (including anonymous information) results are disappointing.

The effectiveness of measures is also a problem, particularly in a non-standard (proprietary) environment.

Tree models are useful for detailed analysis, especially to show where countermeasures are effective in reducing risk.

Confidence Limits

The confidence limits problem becomes extreme when we address **very unlikely / very high impact incidents** (catastrophes)

Obviously we have a much better estimate of the likelihood of an event which occurs on average every day or every week, than one which has never occurred or only happens every hundred years.

Furthermore Disasters (the only unlikely impacts large enough to interest us) contain a higher proportion of indirect and immaterial costs which are hard to estimate reliably.

We also have much better data on incidents affecting availability than confidentiality (CIA).

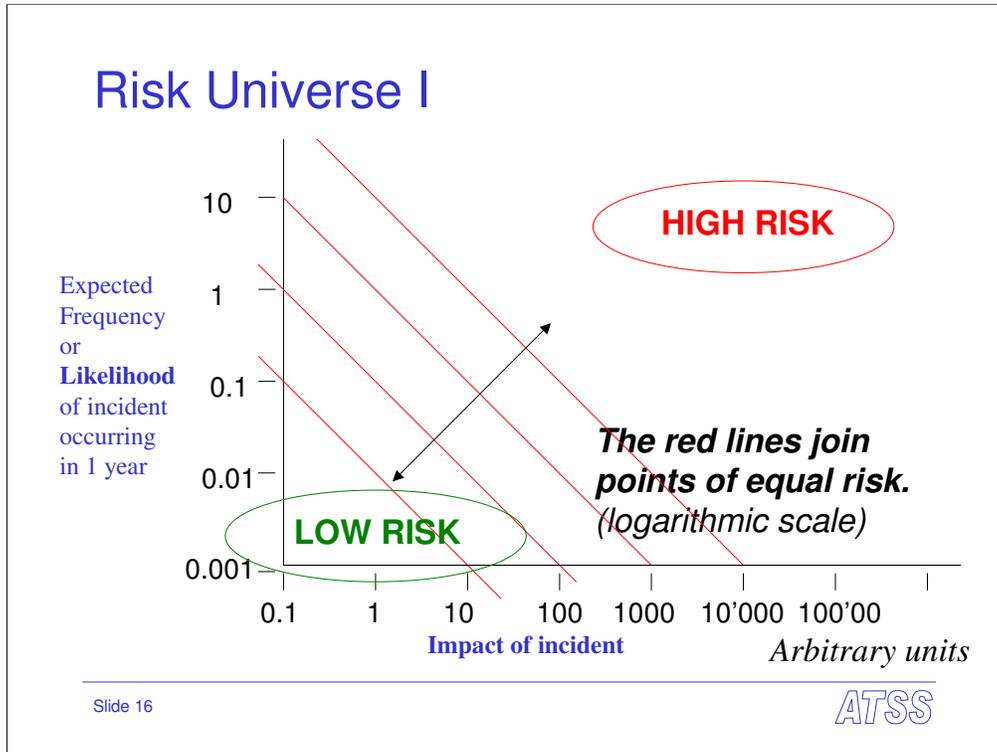
Slide 15

ATSS

Both likelihood and impact are uncertain.

Loss of availability is always detected (implicitly), loss of confidentiality may well remain undetected.

Integrity lies in the middle.



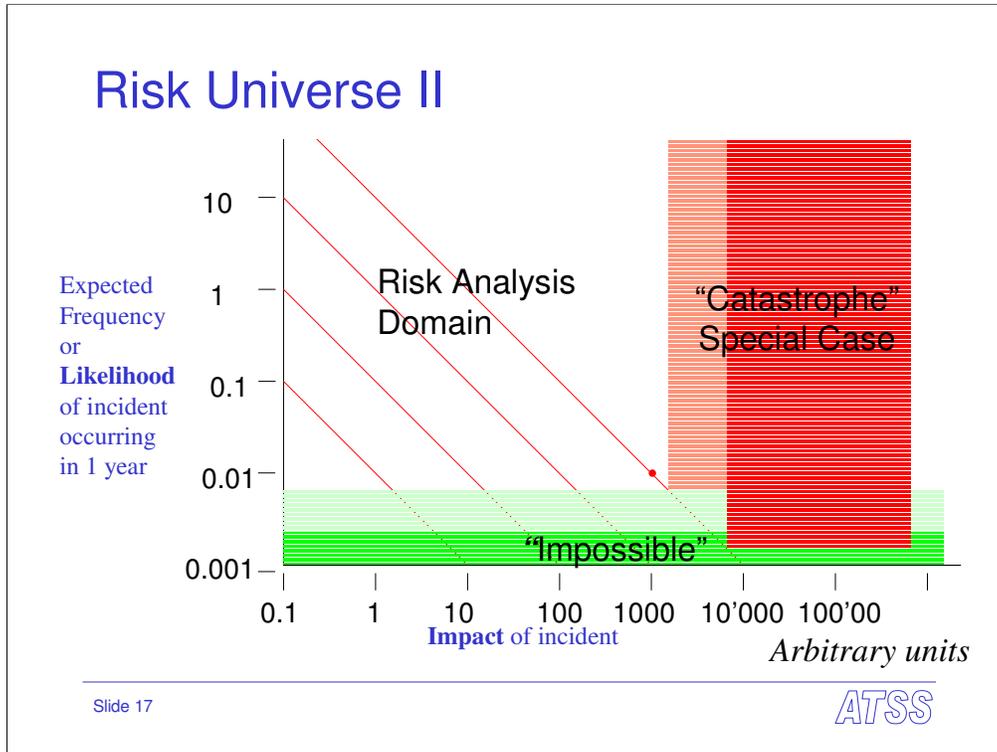
This is how we think about risk.

(The risk contours do not stop at the axes)

Usually there is a level of risk (High Risk) which the policy of the organisation considers “unacceptable”.

Such risks must be treated, avoided or explicitly accepted (signed-off).

Think what it means to equate a low-impact, high frequency risk with a very high impact, highly improbable risk, with the same numeric median value.



Different areas of the Risk Universe have different properties.

The numbers on the diagram are arbitrary, they are intended to remind the reader that the scales are logarithmic.

The boundaries are meant to be “fuzzy” and must be adapted for each organisation according to its policy and scaled to suit its size. The impossible boundary may well not be horizontal.

The Diagram shows that we can ignore incidents so unlikely as to be “impossible”.

It also shows a “catastrophe” area where the risks must be considered specially.

The reasons will be discussed later.

- Lack of information (extremely wide confidence limits)

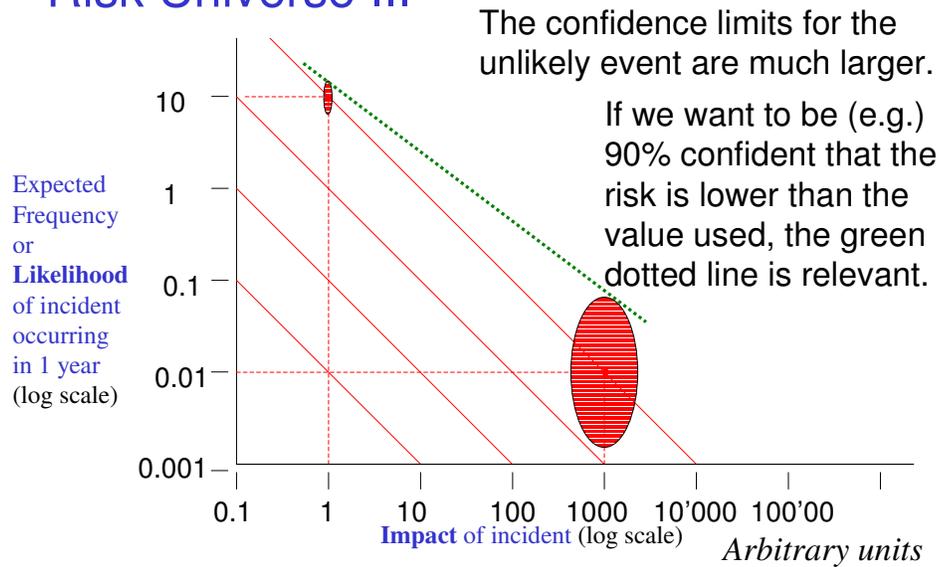
- Risk aversion

- Precautionary principle or MaxiMin

- Risks may affect all customers and can be ignored

In practice due-care, legal/regulatory, apply more than financial and risk drivers here.

Risk Universe III



Slide 18

ATSS

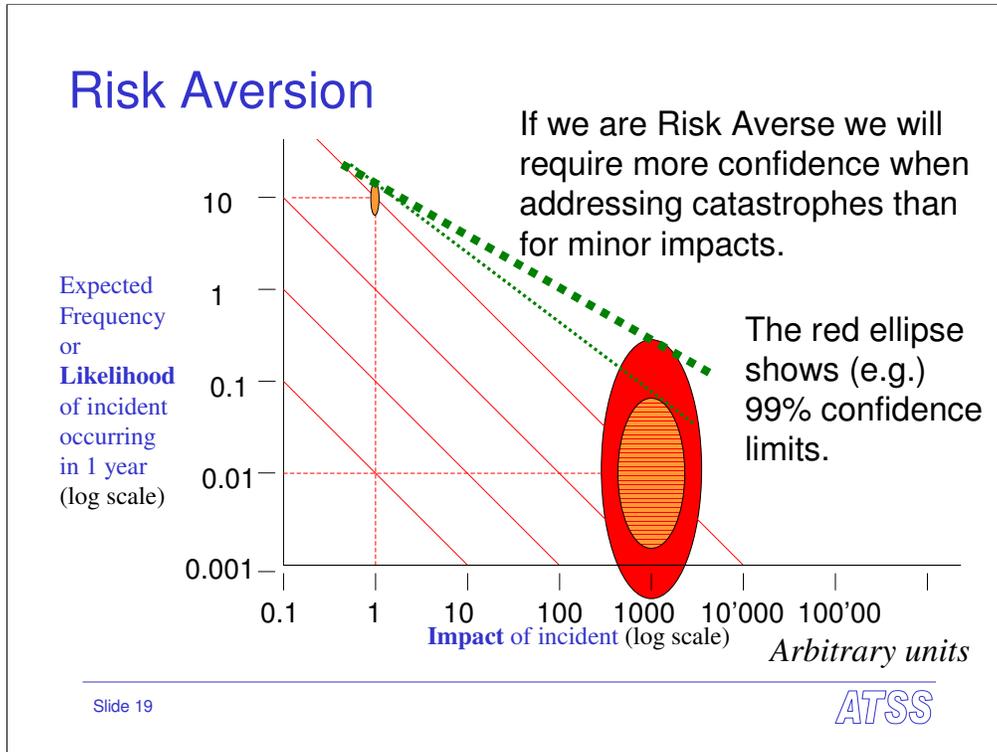
How do you handle risks with very different confidence limits?

This is an example showing possible 90% confidence limits.

If we take the upper risk-boundary, it means “we can be 90% confident that the risk is lower than this value”

Note the uncertainty in impact as well as likelihood.

The figure of 90% is just an illustration.



This example tries to define Risk Aversion.

Risk Aversion (as opposed to Risk Neutrality) means that we **prefer** a residual risk of EUR 1 Mio based on:

“EUR 10'000 100x per year”

to the “same” Risk of EUR 1 Mio based on:

“EUR 100 Mio once in 100 years”

We contend that for frequent small impacts an organisation can work with the median risk value, or very low confidence limits, because there is no stigma and no exceptionally high cost.

For very high impacts this is not true. We do not expect them, and there may be a stigma (due-care) attached to our failure to prepare for this risk. The individual incident is more noticeable - “bump in the road”. Therefore we desire a higher confidence level that the risk is treated.

Obviously the 90 and 99% are arbitrary examples. This is a policy issue.

In practice a weighting factor for big impacts is more efficient, than performing statistical calculations on inadequate data.

When we buy insurance, we are showing risk aversion. In particular we are prepared to pay more than the actuarial risk for “peace of mind”

“EUR 100 Mio once in 100 years” *Who cares about “once in 100 years” incidents?*

But this is misleading. What we really want to be able to say is for example:

“The likelihood that the “100 Mio Incident” will occur in the next year is < 1%.”

The two statements are equivalent. Mathematically the difference is much much smaller than the statistical uncertainty.

Governance Effort

We not only need to "do security" we also need to communicate assurance that "security is being done". i.e. we must document, manage and communicate the results and the process so that it is auditable and reproducible.

What proportion of our resources must we spend on governance as opposed to doing?

Not too much and not too little!

There is no answer to this question.

The point is that if we do a detailed risk analysis of everything, we will have spent too much on analysis.

Slide 20

ATSS

I am not going to talk about SOX!

Baselines

The practical solution to this issue is to adopt a set of Baseline controls which are universally applied across the organisation, and which are adequate for (e.g.) 80% of all the organisation's projects or applications.

This means that we often have "too much security", but we make implementation simpler and governance much easier.

Slide 21

ATSS

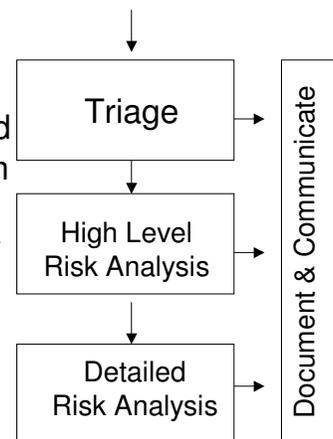
See previous remarks regarding "consistent risk across the organisation"

Triage Example

Each application/project is assessed according to objective criteria, which indicate whether the security requirements are likely to be met by the baseline controls or if it has special security requirements.

Only the special projects are then analysed with a "high level risk analysis" and if necessary a "detailed risk analysis".

The standard projects require no further effort (except documentation and (annual?) review).



Slide 22

ATSS

The triage does not ask about likelihood, and only indirectly about impact.

This information can be provided quite quickly and easily; it could include:

1. number of users,
2. internal or external users,
3. financial transactions involved,
4. confidential/personal/secret data,
5. internet connection, public/private network
6. size of project budget
7. Compliance with standards, (architecture, baselines)
8. etc.

ROSI Main Conclusions

Beware using ROSI where confidence limits are very wide i.e. for **unlikely catastrophes**. Comparison of risks with wide and tight confidence limits is even worse.

You can put a lot of effort into best case/worst case or sensitivity analysis, but when you just do not have enough information:
GIGO Garbage In does not give Gospel Out! 🙄

ROSI is valuable for choosing between alternative countermeasures to address a particular risk, and in this application some of the uncertainties apply to all alternatives and are "cancelled out".

In any case the cost-effectiveness of countermeasures is very important, and it must be questioned even if we are reluctant to write down a number of CHF or USD. i.e. at least a mental or reverse ROSI.

Slide 23

ATSS

Reverse ROSI

A project must make commercial sense.

If a ROSI calculation seems like "fantasy", turn it around!

Your project plans to spend 5 mio.

How will this sum be earned?

Do any of these scenarios seem **plausible**?

- *Might well "save" 100 mio within 20 years?*
- *Might well "save" 15 mio within 3 years?*
- *Might well "save" 5 mio within 1 year?*

for a shorter period you should have some decent information

if all of these are absurd, then question the project!

Baselines

Baselines only address the effort problem.
We have not solved the optimisation issue.
Baseline controls must be practical and cost effective.

Theoretically we could perform ROSI calculations for each alternative version of all potential baseline control candidates.

As discussed above ROSI is useful in order to select the appropriate alternative for a specific control.

However it is hopeless to use ROSI to determine whether to mandate firewalls OR data backup.

Slide 24

ATSS

Most of us do not start from nothing. Some baseline controls are already established, so in practice you are likely to have to decide how “strong” a particular control needs to be.

You can usefully use ROSI to compare different alternatives
e.g. for backup: Tape, SAN remote copy, etc., or frequency ...

Security without Risk !



Alternative approaches which may avoid quantifying risk:

Due Care Baselines

Donn Parker proposed that the baseline controls should be selected in order to meet a "**standard of due care**" as opposed to (attempting) a risk or cost-benefit analysis.

In practice this is not so different from the conventional approach.

Capability Maturity Model

In view of the popularity of the CMM in IT, it is not surprising that it has been applied to IT security.

Process Quality

Balanced Scorecard

Slide 25

ATSS

Due-Care Baselines

In this section we refer to the "Due care baseline" as described by Donn Parker. Parker proposes that the baseline controls should be selected in order to meet a "**standard of due care**" as opposed to (attempting) a risk or cost-benefit analysis. Parker claims not only that it is impossible to perform risk analysis (at least in any useful way) for the reasons indicated above, but also that "avoiding negligence is more important than reducing risk".

Due-Care Baselines are selected by considering published baselines, protection profiles, industry/regulatory standards and considering their applicability and practicality in the actual business and environment. Although we may view this as a "risk analysis" in the widest sense, it makes no attempt to quantify risk (i.e. by estimating impact and likelihood), and this is an important distinction.

Parker also allows that additional "special" controls should be introduced where, there are special risks. For further information see "Fighting Computer Crime" by Donn B. Parker".

One objection to the theory behind the "due-care" approach, is that it is effectively "the blind leading the blind".

In practice the difference between this approach and the conventional "triage" scheme outlined above is smaller than implied by the provocative theory.

When experts try to measure IT security risks in the absence of appropriate actuarial information, this is essentially what they do. They assume that "generally accepted baseline controls" or "current good/best practice" controls are appropriate to normal risks and that additional controls will be required to meet special risks.

A final point is that the "due care baselines" must be kept up to date as formerly exotic controls become economical and practical, and as new threats emerge.

Capability Maturity Model

In view of the popularity of the CMM in IT, it is not surprising that it has been applied to IT security. The starting point is an exhaustive list of controls (IBM calls this the Information Security Framework), from which the required baseline will be defined. Using CMM these controls are classified into 5 maturity levels initial/basic/capable/efficient/optimising. The model can then be used to compare the current capability level against the baseline levels for each control.

The 5 levels certainly provide a more granular description than a simple yes/no, which can foster a "tick-in-the-box" mentality. However (e.g. IBM Security Framework) includes the control "Security Risk Management Framework" and does not attempt (in principle) to eliminate risk measurement.

The CMM provides a useful refinement compared to the conventional model, at the cost of blurring the boundaries between compliance and non-compliance.

Process Quality approaches can be attractive where quality assurance methodology is already pervasive in an organisation.

Balanced scorecard is logically equivalent to converting all costs and benefits to CHF (EUR, USD) i.e. ROSI.

Catastrophes

Ecological issues present the same problems. Not enough information. The **Precautionary Principle** is often invoked but there are a wide variety of interpretations.

MaxiMin

With Maximin we consider the "worst case impacts" of the various alternative courses available, and adopt the one with the **smallest maximum impact**. Impact includes cost of safeguards.

i.e. ignore likelihood above a threshold probability.
applied to impacts above a catastrophe threshold

Do you think this helps?

Slide 26

ATSS

Hard and soft precautionary principles. See EU standard.

Maximin

One expression of the so-called "precautionary principle" used by the environmental and health decision makers is Maximin.

This is applied to risks with high impact and uncertain probability. With Maximin we consider the "worst case impacts" of the various alternative courses available, and adopt the one with the **smallest maximum impact**. In this evaluation the cost of implementing safeguards is part of the impact.

The effectiveness of Maximin (and of other expressions of the precautionary principle depends on two "hidden variables".

Threshold Probability

There must be some probability below which we decide that the incident is **impossible**. Otherwise Maximin / precautionary principle would oblige us to implement safeguards for absurd threats.

Catastrophe Threshold

Similarly there must be criteria for the type and/or magnitude of catastrophe whose risk should be addressed by Maximin.

Both threshold probability and catastrophe threshold are relevant to IT Security.

These thresholds are affected by both objective and subjective factors.

Objectively an organisation might decide that if a nationwide disaster also affects its customers they would not require its services, and that there is no point in safeguarding production under this scenario.

Cognitive Factors

The perception of risks is subjective. This is important because optimisation requires as much objectivity as possible.

Familiarity and Emotion

Familiar and emotional threats are overestimated

Loss aversion / Status quo

People are less ready to relinquish what they have...

System Neglect

Risks which are part of systems are underestimated

Probability Neglect

Probability is ignored altogether particularly for catastrophes

Externalisation

Slide 27

ATSS

We should be aware of some cognitive factors which influence risk perception, when people make decisions on risks. These can affect our analysts in estimating risk, but more importantly they affect the thresholds for "Maximin" and "Impossible" events. The following factors have been identified [Sunstein], and investigated by numerous studies and experiments many involving "Willingness-to-pay (WTP).

Familiarity & Emotion

Risks with which we have some familiarity, either through direct first-hand experience, or through media reports will tend to be over weighted compared to unfamiliar risks. Incidents with emotional impact (horror, fear) and vivid reporting are overweighted.

Loss aversion / Status quo

People are less ready to relinquish something they already have, than the prospect of something of the same value in the future. This tends to reinforce the status quo.

System Neglect

People have difficulty in understanding that risks are part of systems, and that interventions create new risks. This is probably the most important factor generally and certainly for IT.

Probability Neglect

People tend to ignore probability altogether, when the risk is catastrophic. In any case their perception of Risk is not linear with (low) probabilities.

[Schneier] offers a different list with similar content:

1. People exaggerate spectacular but rare risks and downplay common risks.
2. People have trouble estimating risks for anything not exactly like their normal situation. Personified risks are perceived to be greater than anonymous risks.
3. People underestimate risks they willingly take and overestimate risks in situations they can't control.
4. People overestimate risks that are being talked about and remain an object of public scrutiny.

Externalisation

One point, which it is very important to bear in mind is Externalisation. This is where the risk and the costs of the safeguards or the benefits are born by different parties.

There are other factors which are **not really relevant** to our IT security scenarios, such as

The assumption that nature is benign

Aversion to risks which affect underprivileged groups etc.

Conclusions I

Remember **confidence limits** (GIGO). Beware of comparing the risks of frequent and very rare events. Incidents affecting confidentiality are also problematic.

Do use ROSI to compare alternative controls.
Always bear costs in mind. at least "Mental ROSI"

Use whatever cost-benefit methodology is standard in your organisation (NPV, IRR, ROI) and include a risk-reduction term.

Do not base decisions entirely on a single criterion.
Not even ROSI. 😄

Conclusions II

Your Risk management **policy** is critical. It may be "hidden" in your risk management process. It defines the level of risk which is "unacceptable" (risk tolerance), as well as your risk aversion, catastrophe threshold etc.

It assures consistency across your organisation.

Policy can change/be changed over time.

Appropriate Communication, Documentation and Reproducibility (auditability) are vital.

Find the line between too little and too much analysis.

Beware terminology.

What sounds like nonsense may be using a different definition.

Workshop

I have tried to give an impression of the areas we have discussed in our WG in the time available.

We are thinking of organising a small workshop in order to discuss the work of our group in more depth.

Please let me know if you would be interested in taking part. anthony.thorn@atss.ch