

SANDIA REPORT

SAND2004-4233
Unlimited Release
Printed August 2004

A Classification Scheme for Risk Assessment Methods

Philip L. Campbell, Jason E. Stamp

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865)576-8401
Facsimile: (865)576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800)553-6847
Facsimile: (703)605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



A Classification Scheme for Risk Assessment Methods

Philip L. Campbell, Jason E. Stamp

Networked Systems Survivability & Assurance Department
Sandia National Laboratories¹
P. O. Box 5800
Albuquerque, New Mexico 87185-0785

Abstract

This report presents a classification scheme for risk assessment methods. This scheme, like all classification schemes, provides meaning by imposing a structure that identifies relationships. Our scheme is based on two orthogonal aspects—level of detail, and approach. The resulting structure is shown in Table 1 and is explained in the body of the report.

Table 1 Classification Matrix (Shown also as Table 2 on page 13)

Level		Approach		
		Temporal	Functional	Comparative
Abstract	Expert	① Engagement	④ Sequence	⑦ Principles
Mid-Level	Collaborative	② Exercise	⑤ Assistant	⑧ Best Practice
Concrete	Owner	③ Compliance Testing	⑥ Matrix	⑨ Audit

Each cell in the Table represent a different arrangement of strengths and weaknesses. Those arrangements shift gradually as one moves through the table, each cell optimal for a particular situation. The intention of this report is to enable informed use of the methods so that a method chosen is optimal for a situation given.

1. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000.

Table of Contents

1	Introduction.....	7
2	Context	9
3	Background.....	11
4	Our Classification Scheme.....	13
4.1	An Analogy: Risk of Heart Attack.....	13
4.2	Column Headings: Approach.....	14
4.3	Row Headings: Levels.....	16
4.4	Cell Names: Types.....	17
5	Using our Classifications.....	23
6	Conclusions	25
	References	27

List of Tables

Table 1	Classification Matrix (Shown also as Table 2 on page 11).....	3
Table 2	Classification Matrix	13
Table 3	Method Levels.....	17
Table 4	Example Classification.....	23

1 Introduction

This report imposes structure on the set of risk assessment methods in order to reveal their relationships and thus optimize their usage. We present a two-dimensional structure in the form of a matrix, using three abstraction levels for the rows and three approaches for the columns. For each of the nine cells in the matrix we identify the method type by name and example. The matrix helps the user understand

1. what to expect from a given method,
2. how it relates to other methods, and
3. how best to use it.

Each cell in the matrix represent a different arrangement of strengths and weaknesses. Those arrangements shift gradually as one moves through the table, each cell optimal for a particular situation. The intention of this report is to enable informed use of the methods so that a method chosen is optimal for a situation given. The matrix, with type names in the cells, is introduced in Table 2 on page 13 below.

Unless otherwise stated we use the word “method” in this report to refer to a “risk assessment method,” though often times we use the full phrase. The use of the terms “risk assessment” and “risk management” are close enough that we do not attempt to distinguish them in this report.

The remainder of this report is organized as follows.

In Section 2 we provide context for this report—what a “method” is and where it fits.

In Section 3 we present background for our classification scheme—what other schemes we have found, the fundamental nature of methods and their necessary incompleteness.

In Section 4 we present our classification scheme in the form of a matrix, then we present an analogy that should provide an understanding of the scheme, concluding with an explanation of the two dimensions and the nine types in our scheme.

In Section 5 we present examples of each of our classification types.

In Section 6 we present conclusions.

2 Context

Our first task is to address a point of confusion. Both “methodology” and “method” can denote “a particular procedure.”² However, in this report we use the more literal definition, “the study of method,” for methodology. This definition allows us to distinguish between the risk assessment methods we use as examples from the approach we use in this report, which is methodological. We caution the reader that the names given for some risk assessment methods use the word “methodology,” hence the confusion.³

A method imposes order. In fact, the whole point of using a method in this uncertain world is to glean the benefits of that order. Since risk is not objective and no one can grasp the whole picture, the results of a risk assessment are not always repeatable. However, if the assessment is methodical (i.e., it uses a method), then, by definition at least the process can be repeated. This seems to be the most for which we can hope, and thus the focus on method.

A “method” is a way of doing something, “a particular procedure or set of procedures,” as one dictionary puts it. We place “method” the same way that Checkland does, as denoting an approach that is between philosophy and technique (except that Checkland uses the word “methodology” where we use the word “method”):

I take a methodology [we would say “method”] to be intermediate in status between a philosophy...and a technique or method. A philosophy...might be... ‘political action should aim at a redistribution of wealth in society,’...At the other extreme a technique is a precise specific programme of action which will produce a desired result: if you learn the appropriate technique and execute it adequately you can, with certainty, solve a pair of simultaneous equations...A methodology will lack the precision of a technique but will be a firmer guide to action than a philosophy. Where a technique tells you ‘how’ and a philosophy tells you ‘what,’ a methodology will contain elements of both ‘what’ and ‘how.’ ([8], page 162)

Another way to put this is that “philosophy” is just above (i.e., more abstract than) our “abstract” level, and “technique” is just below (i.e., more concrete than) our “concrete” level (see Section 4.3).

In order to avoid missing valuable methods we have adopted a liberal view about what constitutes a method, deciding to include rather than exclude. For example, we include DOE’s “21 steps” [12]: this is not a method per se but an assessment method could easily be constructed from it. In addition some of the items we include as methods themselves call for the use of a risk assessment method, making for a circularity that indicates the leeway afforded in this area. We note that there are more items here than we can absorb.⁴

2. This situation is similar to “bi-annual” which can mean either every other year or twice a year. In this case the latter usage is “disapproved,” as Webster’s dictionary puts it, with the use of the unequivocal “semi-annual” preferred.

3. See also the passage from Checkland below.

4. For example an Internet search on “risk assessment” & “methodology” retrieved 278,000 hits.

3 Background

We present three points in this section. First, current classification schemes are minimal. Second, assessments are comparisons—relative, not absolute. And third, methods cannot completely define either the assessment elements or the processes.

First, the few current classification schemes we have found are minimal. One example is the quantitative vs. qualitative bifurcation. This bifurcation is commonly referenced. AS/NZS 4360 [5] adds a third element to the scheme, making it quantitative vs. semi-quantitative vs. qualitative. Unfortunately it is not clear what this new category provides. Conversely Donn Parker, for example, has argued that the entire quantitative approach has “utterly failed” [26], which would reduce the quantitative/qualitative bifurcation to unity, removing its ability to make any distinctions at all and thus rendering it useless.

Another example classification scheme is von Solms’ traditional assessments vs. baseline controls. BS 7799 [6] is an example of a baseline control. Unfortunately in von Solms’ opinion traditional methods are “tedious, suspect, inconsistent, painful and useless” ([42], page 93), reducing this bifurcation to a matter of separating wheat from chaff, as opposed to being able to make informed use of the methods so that a method chosen is optimal for a situation given.

We want a scheme that tells us what to expect. For example, picking up a book and noting that it is a mystery story tells us what to expect from the book. That one word provides an entire structure. There is room for an unbounded number of plots within that structure, but as we begin the first page we are already wondering if the butler did it, because we have the structure. We would like a similar scheme for risk assessment methods. Otherwise we are somewhat at sea, clinging to a method or two we happened on as we drift, never able to leverage the power available from an informed choice among a range of alternatives.

Second, we claim that assessments are comparisons. The comparative aspect is more explicit in some assessments than others but we claim that all are comparisons. Uncertainty is a necessary feature of risk, so assessment of risk cannot be determined absolutely. The only other option available is to use comparison. The results of assessments are thus relative, not absolute, even if the results are expressed in numbers. Every method has something against which the system under advisement is gauged. That “something” may be completely in the mind of the assessor or part of it may be explicit in a Standard, written down and codified. But all of them have something.

Third, none of the methods we have seen fully describe how to identify an asset or a vulnerability or a threat or a risk or a control, nor do we expect ever to see one that does. Many of the methods give definitions and provide examples but they presume you will know an asset, say, when you see one. This is to be expected. If it were possible to describe assets completely, then it would be possible to identify all of them, something that we do not believe is possible.⁵ The same argument applies to the processes presented in the methods. As a result risk assessment methods can only *help* in the process of risk assessment; they rely upon a person and cannot stand alone.

5. Perhaps such a putative complete list of vulnerabilities, for example, could be used to show that it is not complete since the knowledge of the complete list of vulnerabilities is itself a vulnerability and not on the list.

4 Our Classification Scheme

In this Section we present and explain our classification scheme, shown in Table 2. For ease of reference each cell is numbered (in column-major order). In Section 4.1 we present an analogy to help the reader understand our classification scheme. In Section 4.2, Section 4.3, and Section 4.4 we describe the columns, rows, and cells, respectively, of the matrix.

Table 2 Classification Matrix

Level		Approach		
		Temporal	Functional	Comparative
Abstract	Expert	① Engagement	④ Sequence	⑦ Principles
Mid-Level	Collaborative	② Exercise	⑤ Assistant	⑧ Best Practice
Concrete	Owner	③ Compliance Testing	⑥ Matrix	⑨ Audit

For ease of discussion in the material below we use the term

expert to refer to an outside consultant who is knowledgeable in assessment methods but unfamiliar with the target system.

We use the term

owner to refer to someone who is not knowledgeable in assessment methods but is familiar with the target system.

4.1 An Analogy: Risk of Heart Attack

An analogy can help us explore new territory: we can navigate the new territory by using our map for the old, thereby helping us understand that new territory. The analogy we present in this section is intended to provide navigation capabilities for our classification scheme. The analogy is also intended to support or claim that our classification scheme is neither arbitrary or incomplete.

As a society we have learned that we can reduce the likelihood of heart attacks by taking action. Part of that action is periodically to assess our risk of heart attack. There are several methods we can follow for this assessment: we can use a stress test, a threat analysis, or a lifestyle comparison. We will explain each of these below.

A stress test actually exercises our heart. For example we can use a treadmill and plot our heart rate against the rate at which the treadmill is moving. That curve gives an assessment of our heart health and thus of our risk of heart attack. We presume that there are similar tests extant for particular heart-stressors such as alcohol, emotional stress, and so on. Note that the focus here is on an actual test in real time.

A threat analysis is a review, for example, of the history of heart attacks in our extended family and the kind of life we lead (e.g., what is our purpose in life? how do we handle stress? and what is the nature of our relationships with people?). Our answers are combined in some intuitive way to provide us our risk of heart attack. Note that the focus here is on how we

operate or “function” in our environment: no testing is involved.

A lifestyle comparison compares our lifestyle with a prioritized list of activities, attitudes, and habits that are known (or maybe just believed) to be related to heart attack incidence, such as smoking, diet, and exercise. We determine our risk of heart attack by comparing our life with the list. Such lists can be generated by sifting through what we know about the people who have had heart attacks, looking for commonalities, then checking to see that some statistically significant percentage of those who died of natural causes without suffering heart attacks did not share that commonality. Note that the focus here is on comparison: no testing or examination of function is involved.

Notice that the second approach, threat analysis, can be seen as a blend of the other two. Threat analysis asks how your heart performs—this is where it is similar to the stress test—in its environment—this is where it is similar to the lifestyle comparison. Note also that there is no similar blend of the first and last approaches: the former is in real time; the latter is timeless.

The level of expertise we require from the person administering the assessment can vary as well. For example, we can do a stress test at home, of course, or we can do one with a physician’s assistant, working in partnership with the assistant, pointing out, for example, what we believe are peculiarities we have noted in our heart’s function, or we can hire a specialist to do it all for us while we do little more than obey orders as we occupy our minds with a periodical.

The approaches we have presented above—stress test, threat analysis, and lifestyle comparison—correspond to the three columns of our matrix—the temporal, functional, and comparative approaches. The levels of expertise we have presented correspond to the rows of our matrix.

There are advantages and disadvantages for all these approaches and levels. And of course we can combine them and adjust the mixture to arrive at a new set of advantages and disadvantages.

The results of the application of any of these heart-risk assessment approaches at any level is intended to be the same: a prioritized list of actions that we can take to reduce our risk of heart attacks. It may be that the list contains a single item, such as “Consult a physician: you need to take steps to lower your risk of heart attack.” Given the possible mixes of approaches and levels there are many such possible lists. It may be that the set of items on any pair of such lists is disjoint. However, the intention is that the relative order for any two items that appear on any two lists be the same.

4.2 Column Headings: Approach

In this section we describe the columns of the matrix shown in Table 2.

Risk assessment methods appear to divide into three archetypical approaches which we have named “temporal,” “functional,” and “comparative.” As noted in Section 4.1 above, these three approaches correspond to the stress testing, threat analysis, and lifestyle comparison approaches, respectively, in our analogy for reducing heart attack risk. Each approach is

described below.

4.2.1 Temporal

A temporal method stresses a system: actual tests are applied. These “tests” exercise key components of attacks, subject to some explicit or implicit constraints. The performance of the system as a consequence of the application of those tests is the result of the method.

It may be impractical to apply tests to the system itself. In this case the only choice is to use a model of the system instead. Unfortunately using a model introduces the complicating question of fidelity. An erroneous model confuses matters, possibly providing a false sense of security which is even worse than confusion.

It is not possible to model all possible attacks. It is not possible even to list all possible attacks. The attacks to which we pay primary attention are those about which we are already concerned. But it is important to expand our view by exercising the system (or the model), exploring for additional failure modes, and then we should investigate what attack(s) could arrive at whatever new failure modes we have uncovered. This is an empirical process.

One of the strengths of this approach is that it tests the system itself (or a model), hopefully clearing away misconceptions.

4.2.2 Functional

A functional method balances the temporal approach, described above, and the comparative approach, described below. This is the reason that the functional approach appears in the middle column in Table 2. The functional approach depends less on an understanding of a system (or model) than the temporal approach and it uses more system-specific understanding than the comparative approach. The functional approach focuses on threats and protections. A threat model, a list of vulnerabilities, and the likelihood of success of threats being mounted against those vulnerabilities are weighed against the assets, protections, and the likelihood of success of protections being able to defend those assets against those threats. Aspects of the temporal approach, such as statistical modeling, and aspects of the comparative approach, such as expert systems, can be used in this approach. In von Solms’ scheme, this is the “traditional” approach.

One of the strengths of this approach is that it considers specific threats, vulnerabilities, assets and countermeasures.

4.2.3 Comparative

The comparative approach presents an explicit standard. An owner compares the owner’s system and/or procedures with the standard. Note that there is no explicit system model involved here as there is in the temporal approach. Neither is there an explicit list of threats and assets here as there is in the functional approach. The model and the lists are only implicitly present in generic form. Each standard circumscribes what is considered to be the important aspects for all systems in some broad category such as a particular industry. In each case the

standard is prepared and maintained as the distillation of continually developing expert opinion and experience in the face of a continually changing environment. So instead of every member of a given industry hiring an expert to perform an assessment, the industry members share the generalized results of a small number of such assessments, under the presumption that their systems and activities make the small number of assessments sufficiently applicable to the remainder of them to justify the savings.

One of the strengths of this approach is its simplicity. Comparative methods can be ideal for organizations as they begin to focus attention on security. For example, it is possible for a comparative method at even a concrete level (described below) to fit on a single page.

4.3 Row Headings: Levels

In this section we describe the rows of the matrix shown in Table 2.

In general the successful application of any security or risk method depends on two factors:

- Capability to execute a method and
- Knowledge of the system.

These two factors often pit breadth against depth. Our definitions of “expert” and “owner,” shown at the beginning of Section 4, suggest this conflict. The expert becomes an expert by devoting time to many systems and being able to generalize from them. The owner becomes knowledgeable in his own system by devoting time to that one system. There are exceptions, of course, but generally speaking this conflict appears to be fundamental.

Methods view the landscape from different heights, so to speak, at different altitudes, at different levels. Altitude is a tradeoff between scope and detail: the higher the method, the greater the scope and the coarser the granularity of detail; the lower the method, the smaller the scope and the finer the granularity of detail.

The methods at the higher levels we refer to as “abstract;” the methods at the lower levels we refer to as “concrete.” Abstract methods have a broad application but require a high level of expertise. These methods describe how an expert performs an assessment. Concrete methods have a narrow application but require only a low level of expertise, though they require a user’s knowledge. These methods describe how a system owner can perform an assessment. That is, these methods indicate what to look at and how to evaluate it. We presume that this requires considerable space to describe, which implies that the scope of these methods is narrow. The methods that are a mix of abstract and concrete we refer to as “mid-level.”

The names of the levels can also reflect who “drives” (i.e., directs) the application of the method. The three names we use are “expert,” “collaborative,” and “owner.” Abstract methods are driven by an expert and thus these can be called “expert” methods. The owner is involved to the extent that the expert needs information about the system but the owner does not drive the method. Concrete methods, on the other hand, are driven by the owner and do not require an expert to be explicitly involved during the assessment (though the expert is implicitly involved to the extent that the owner follows the expert’s direction as described by the method). Mid-level methods are driven by both an expert and the owner. Mid-level methods can thus be

referred to as “collaborative.”

To reinforce the concept of altitude implicit in the notion of “levels” we arbitrarily assign numbers to the three levels, as shown in Table 3. In this way additional granularity is available, such as describing a method as at “level 1.5,” say.

Table 3 Method Levels

Number	Description	Driver
3	Abstract	Expert
2	Mid-Level	Collaborative
1	Concrete	Owner

4.4 Cell Names: Types

In this Section we describe each cell name in Table 2 in numerical order (top-to-bottom, left-to-right). At the end of this Section we describe one additional instance of a type 9 method, namely “Construction.”

4.4.1 Engagement

An Engagement consists of experts looking for any way, within given bounds, to compromise assets. An example of this method type in information systems is the Red Team. Some people describe Red Teaming as generating probes that are “representative of actual threats” [30] or as focusing “on finding vulnerabilities and exploiting them as a ‘hacker’ would be expected to do” [19]. The result of this type of method is the cumulative achievement of the experts, given the boundaries and conditions. Engagements can be run on generic system pieces, such as firewalls or routers, independent of any particular system. An example of this type of method in physical systems is the set of tests performed on gates and doors, given certain capabilities (such as tools, explosives, and chemicals) and given certain conditions (such as lighting, weather, and time). This testing attempts to explore all the possible ways to use the tools, explosives, and chemicals within the given constraints. An owner *should* not perform tests of this type and the tests *need* not be done with the owner’s active collaboration, thus this type appears in the Abstract/Expert row.

4.4.2 Exercise

An Exercise links experts and owners together in order to test the protection on assets particular to a given system. Customarily the owner sets boundaries and conditions, possibly providing inside information for and working collaboratively with the experts. An example of this method type in information systems is “Penetration Testing,” for which Fried [15] identifies three sub-types: physical; organizational (or procedural); and electronic. Customarily Penetration Testing requires less expertise than Red Teaming, employing more standard attacks. The list of attacks mounted and the list of systems and functions penetrated, accompanied by the vulnerabilities exploited to effect those penetrations, are the results of using this type. An example of this type of method in physical systems is “force on force,” in which participants enact both attackers and

defenders in a manner that is as realistic as possible without causing injury. This type of test requires the owner's collaboration, thus it is in the Mid-Level/Collaboration row.

4.4.3 Compliance Testing

Compliance Testing is a more formal way of describing "door rattling." The tests included in methods of this type are such that the owner can execute them himself without the aid of an expert. For example, the owner can check that the door locks at a given installation actually do prevent entry; the owner can check that windows are indeed locked; the owner can attempt to enter the facility with a badge bearing the photograph of a jungle primate, say; and the owner can telephone a receptionist at the facility in order to attempt to get a password without providing authentication. Examples of this type of method in information systems are the set of security scripts such as the Security Administrator Tool for Analyzing Networks (SATAN) and Nessus [33]. These provide automated ways to pursue certain attack scenarios. The result of using this type is the output of running the script(s). This type of method is in the Concrete/Owner row since the owner by himself and without the aid of an expert can perform methods of this type.

4.4.4 Sequence

A Sequence method type consists of a series of steps, usually posed as questions, and sometimes in a form as complicated as a flowchart. This type asks the user to follow the steps. The output of the steps is the result of using the type. For example, Kaplan & Garrick present a simple sequence method:

- What can happen? (i.e., What can go wrong?)
- How likely is [it] that that will happen?
- If it does happen, what are the consequences? [27]

Sequence methods are the epitome of abstract methods.

4.4.5 Assistant

An Assistant method type keeps track of things, of details, the way a good human assistant does. In this case the assistant keeps track of combinations of lists such as threats, vulnerabilities, and assets. The best instances of this type "walk" the user through the process, prompting for the input needed to populate and rank each list. The lists are combined and ordered mathematically, or at least in some explicit way, usually defined by the user. The ordered lists, which include primarily a list of vulnerabilities and, hopefully, a list of remedial actions, are the result of using the type.

4.4.6 Matrix

A Matrix method type is a table lookup. This type asks the user to select ranges for n-dimensions. The information in the cells in the corresponding n-dimensional subspace is the result of using the type. An expert system is one implementation of this type and is representative of the functional approach.

4.4.7 Principles

A Principles method type, like all of the Comparative types, is a list. This type asks the user to apply the principles to their system. The application of those principles is the result of using the type. Principles are more abstract than Best Practice (described below) and thus have greater breadth.

4.4.8 Best Practice

A Best Practice method type is a list but it is more specific than a Principles list. A Best Practice list could be based on a standard, such as a Principles list, or it could be its own standard. A Best Practices list consists of directives: Do this, Don't do that. This method type asks the user to compare what they do—their current practice—with the best practice list: the list of differences is the result of using the type. Best Practice may be industry- or application-specific but is usually too abstract to be implementation-specific.

An example of a set of Best Practices is GAO's document "Learning From Leading Organizations" [28]. The GAO identified a number of organizations that it considered to be doing superior or "leading" work in information security management. From those organizations the GAO distilled their practices into 11 Best Practices. Here are the first three Practices:

1. Recognize Information Resources as Essential; Organizational Assets Must Be Protected.
2. Develop Practical Risk Assessment Procedures That Link Security to Business Needs.
3. Hold Program or Business Managers Accountable. [28]

Note the imperative nature of the Practices.

Alexander et al. provide a set of best Practices in the area of architecture design in their book A Pattern Language [3]. The volume consists of 253 "patterns" in 37 groups under three headings: "towns," "buildings," and "construction." The name and description of the first pattern is as follows:

1. Independent Regions: Metropolitan regions will not come to balance until each one is small and autonomous enough to be an independent sphere of culture. ([3], page 11)

Several pages of text describe this pattern and how to use it, indicating the Best Practice that this pattern represents. Without the accompanying text this set of patterns could be considered to be an example of a Principles method.

4.4.9 Audit

An Audit method type is a list but it is more specific than a Best Practices list. An Audit is based on an explicit standard, such as a Best Practice list or a Principles list. This type asks the user to evaluate the effectiveness of the controls in place in fulfilling each item in the standard. The set of evaluations for the set of items in the standard is the result of using the type. Audits can be implementation- or system-specific and can often be applied by an owner. Audits are more

concrete than Best Practice and thus have greater depth.

4.4.10 Construction

Besides Audit, there is a second type that could occupy the Concrete-Comparative cell in the matrix shown in Table 2, namely the “Construction” type. We have not yet found an instance of this type so we do not include it in the classification matrix in Table 2 though we do explain it here.

A Construction method type consists of a set of primitives and a set of rules. This method asks the user to build a model by repeated application of the rules on the set of primitives, unioned with the set of the results of the previous applications of the rules. This is the basis for computer language grammars. For example, the following is a set of three rules for building a “factor,” where “|” separates choices, “:=” indicates that the left side can be replaced by any of the choices on the right side, and DIGIT is a primitive consisting of an element from the set {0, 1, ..., 9} [1]:

```
expr ::= expr + term | expr - term | term
term ::= term * factor | term / factor | factor
factor ::= DIGIT | ( expr )
```

Suppose we want to know if “9+5*2” is a properly constructed expression, based on the primitives and rules shown above. If we can find a parse tree for that expression then it is properly constructed. We start with `expr` and we replace that by either “`expr + term`” or “`expr - term`” or “`term`” and then, we replace each of those recursively using the applicable rule until we have transformed our expression into one that consists of only the following symbols: +, -, *, /, (,), and DIGIT. Here is the parse tree for “9+5*2:”

```
expr ::= expr + term
      ::= term + term * factor
      ::= factor + factor * factor
      ::= DIGIT + DIGIT * DIGIT
```

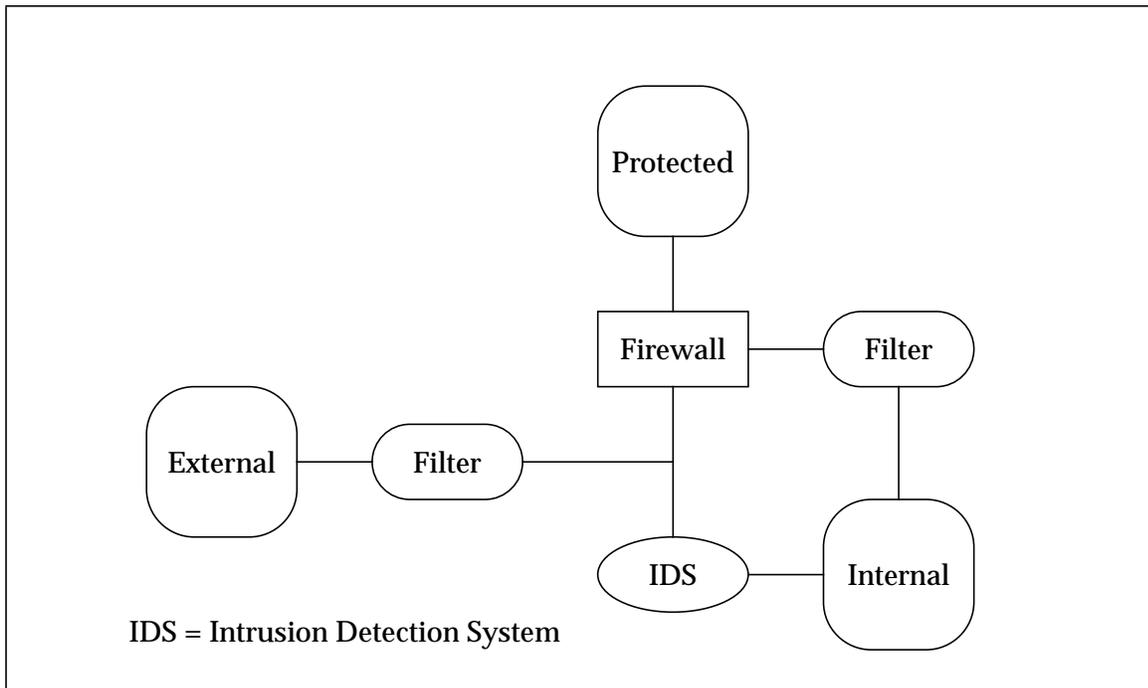
Note that we cannot find a parse tree for “9++5*2” or “+9+5*2,” for example, indicating correctly that these are not properly formed expressions.

Our example above uses the building blocks of mathematical expressions, namely arithmetic operation symbols, parentheses, and the ten digits, but the same approach could use the building blocks of a water production facility, say, using pipes, pumps, valves, tanks, and so on. If the user cannot reproduce the user’s current system via this method, then the user works with the method to find a construction that is “close” and then works to align the current system to the constructed system.

Directives on network security often provide a diagram of a configuration that provides acceptable security. For example, Figure 1 shows one such configuration ([20], Exhibit 131.2, “Perimeter Access Point,” page 1480). We believe that there are variations on this configuration that would be equivalently acceptable but the diagram provides no clues as to their nature. However, a construction method could succinctly express those acceptable variations. Such a

method would be of greater value than the one configuration.

Figure 1 Sample Network Configuration



5 Using our Classifications

In this Section we classify a number of methods in order to clarify our classification scheme and at the same time provide a bridge to the use of actual methods, as shown in Table 4. The methods listed in each category are intended to be illustrative. There is no attempt to be exhaustive.

Table 4 Example Classification^a

Level		Approach		
		Temporal	Functional	Comparative
3	Abstract (Expert)	① Engagement Red Team (e.g., IDART™ [22])	④ Sequence AS/NZS 4360 [5] FIPS PUB 191 [14] IAM [21] IEC/ISO TR 13335 [24] Jelen [25] Kaplan & Garrick [27] NIST 800-30 [31] Schneier [37]	⑦ Principles CoCo [10] Freudenburg [16] GAISP [17] GAPP [18] OECD [33]
2	Mid-Level (Collaborative)	② Exercise Force on Force Penetration Testing [15]	⑤ Assistant Manello [29] OCTAVE [2] RAM-W [35] VSAT™ [43]	⑧ Best Practice DOE's 21 Steps [12] e-Commerce [13] ISF [23] ITIL [7] LfLO [28] NIST 800-53 [32] PoLO [34]
1	Concrete (Owner)	③ Compliance Test- ing security scripts (e.g., SATAN, Nessus) [38] “door rattling”	⑥ Matrix AMSA [4] CRAMM [11] RiskWatch [36] SSAGT [40]	⑨ Audit BS 7799 [6] CobiT® [9] SSAG [39] Trust Services [41]

a. Cell entries are listed alphabetically.

6 Conclusions

This report presented a classification scheme for risk assessment methods, diagrammed in a 3x3 matrix as shown in Table 2 on page 13. Each column, row, and cell of the matrix was described in Section 4. Examples of each cell in the scheme were presented in Table 4 on page 23.

The value of our classification scheme can only be gauged empirically. The scheme is worthwhile if and only if it helps the reader make informed choices so that the method chosen is optimal for the situation given.

References

- [1] A. V. Aho, R. Sethi, J. D. Ullman, "Compilers: Principles, Techniques, and Tools." Addison-Wesley. Reading, Massachusetts. 1986.
- [2] Christopher J. Alberts, Audrey J. Durofee, "An Introduction to the OCTAVESM Method." <http://www.cert.org/octave/methodintro.html>.
- [3] Christopher Alexander, Sara Ishikawa, Murray Silverstein, A Pattern Language: Towns, Buildings, Construction. Oxford University Press. 1977. ISBN 0-19-501919-9.
- [4] AMSA, "Asset Based Vulnerability Checklist for Wastewater Utilities ©." January 2002.
- [5] AS/NZS 4360:1999 Risk Management.
- [6] BS 7799-1:1999 Information security management—Part 1: Code of practice for information security management. BS 7799-2:1999 Information security management—Part 2: Specification for information security management systems.
- [7] Ing. Jacques A. Cazemier, Dr. Ir. Paul L. Overbeek, Drs. Louk M. C. Peters, "Best Practice for Security Management." The ITIL Infrastructure Library. Office of Government Commerce (OCG). 1999. ISBN 0 11 330014 X.
- [8] Peter Checkland, Systems Thinking, Systems Practice. John Wiley & Sons, 1993.
- [9] CobiT 3rd Edition: "Control Objectives for Information and Related Technology (CobiT)." 3rd Edition. July 2000. Published by Information Technology Governance Institute (ITGI). ISBN 1-893209-13-X.
- [10] CoCo: Criteria of Control Board. The Canadian Institute of Chartered Accountants. "Guidance on Control." November 1995. ISBN 0-88800-436-1.
- [11] CRAMM: UK Government Risk Analysis & Management Method. Insight Consulting. <http://www.insight.co.uk/index.htm>.
- [12] Department of Energy (DOE), "21 Steps to Improve Cyber Security of SCADA Networks." (citation unknown)
- [13] e-Commerce Security – Enterprise Best Practices. Information Systems Audit and Control Foundation (ISACF). ISBN 1-893-209-10-5.
- [14] FIPS PUB 191. "Guideline for the Analysis of Local Area Network Security." November 9, 1994.
- [15] Stephen D. Fried, "Penetration Testing." Chapter 15 of the Information Security Management Handbook. Fifth Edition. Tipton & Krause, editors. Auerbach Publishers. Boca Raton, Florida. 2004.
- [16] William R. Freudenburg, "Risky thinking: facts, values and blind spots in societal decisions about risks." *Reliability Engineering and System Safety* 72 (2001), pp. 125-30.
- [17] GAISP: Information Systems Security Association (ISSA), "Generally Accepted Information Security Principles." GAISP V3.0. www.gaisp.org.

- [18] GAPP: Marianne Swanson, Barbara Guttman, "Generally Accepted Principles and Practices for Securing Information Technology Systems." NIST Special Publication 800-14, September 1996.
- [19] J. Todd Hamill, Dr. Richard F. Deckro, Jack M. Kloeber Jr., T.S. Kelso, "Risk Management and The Value of Information in A Defense Computer System." Military Operations Research, V7 N2 2002, pp. 61-81.
- [20] Chris Hare, "Firewalls, Ten Percent of the Solution: A Security Architecture Primer." Chapter 121 of the Information Security Management Handbook. Fifth Edition. Tipton & Krause, editors. Auerbach Publishers. Boca Raton, Florida. 2004.
- [21] IAM: "INFORMATION ASSESSMENT METHODOLOGY. INFOSEC ASSESSMENT METHODOLOGY COURSE. INFOSEC ASSESSMENT TRAINING AND RATING PROGRAM." from www.nsa.gov.
- [22] Information Design Assurance Red Team (IDART™) at Sandia National Laboratories. <http://www.sandia.gov/idart>.
- [23] ISF: Information Security Forum, "The Standard of Good Practice for Information Security." Version 4.0. March 2003.
- [24] ISO/IEC TR 13335 Information Technology--Guidelines for the management of IT Security.
- [25] George F. Jelen, "A New Risk Management Paradigm for INFOSEC Assessments and Evaluations." Computer Security Applications, 11th Annual Conference. 1995. pp. 261-7.
- [26] Ray Kaplan, "A Matter of Trust." Chapter 61 of the Information Security Management Handbook. Fifth Edition. Tipton & Krause, editors. Auerbach Publishers. Boca Raton, Florida. 2004.
- [27] Stanley Kaplan, B. John Garrick, "On the Quantitative Definition of Risk." Risk Analysis, Vol. 1, No. 1, 1981, pp. 11-27.
- [28] LfLO: GAO/AIMD-98-68 Information Security Management. (US General Accounting Office Accounting and Information Management Division, "Executive Guide. Information Security Management. Learning From Leading Organizations.") May 1998.
- [29] Carl Manello, William Rocholl, "Security Evaluation: A Methodology for Risk Assessment." IS Audit and Control Journal. Vol 6, pp. 42-6. 1997.
- [30] George Muncaster, Edward J. Krall, "An enterprise view of defensive information assurance." Mil-com 1999: IEEE Military Communications Conference Proceedings, Atlantic City, NJ, pp. 714-8.
- [31] NIST 800-30: Gary Stoneburner, Alice Goguen, Alexis Feringa, "Risk Management Guide for Information Technology Systems." NIST Special Publication 800-30. October 2001.
- [32] NIST 800-53: Ron Ross, Gary Stoneburner, Stuart Katzke, Arnold Johnson, Marianne Swanson, "Recommended Security Controls for Federal Information Systems." NIST Special Publication 800-53. "Initial Public Draft." October 2003.
- [33] OECD, "OECD Guidelines for the Security of Information Systems and Networks — Towards a Culture of Security." Paris: OECD. July 2002. www.oecd.org.

- [34] PoLO: U.S. General Accounting Office (GAO), "Information Security Risk Assessment: Practices of Leading Organizations." Exposure Draft. GAO/AIMD-99-139. August 1999.
- [35] RAM-W: "Risk Assessment Methodology for Water Utilities," Second Edition. Awwa Research Foundation. Denver, Colorado. 2001.
- [36] RiskWatch, "How to do a Complete Automated Risk Assessment: A Methodology Review." A White Paper available at www.riskwatch.com.
- [37] Bruce Schneier, "Beyond Fear: Thinking Sensibly About Security in an Uncertain World." Copernicus Books, New York. 2003. ISBN 0-387-02620-7.
- [38] Ed Skoudis, "Hacker Tools and Techniques." Chapter 10 of the Information Security Management Handbook. Fifth Edition. Tipton & Krause, editors. Auerbach Publishers. Boca Raton, Florida. 2004.
- [39] SSAG: Marianne Swanson, "Security Self-Assessment Guide for Information Technology Systems." NIST Special Publication 800-26. November 2001.
- [40] SSAGT: Pacific Northwest Laboratories, "Safeguards and Security Survey and Self-Assessment Guide and Toolkit." (This material is available from PNL on a CD by calling (509) 375-4349.)
- [41] Trust Services: American Institute of Public Certified Accountants (AICPA), Canadian Institute of Chartered Accountants (CICA), "Suitable Trust Services Criteria and Illustrations." 2003. www.aicpa.org.
- [42] R. von Solms, "Can Security Baselines replace Risk Analysis?" IFIP TC 11 Conference on Information Security, Research and Business, 14-16 May 1997, Copenhagen, Denmark. pp. 91-98.
- [43] VSAT™ Users Page: <http://www.vsatusers.net/>.