



Research Report no. [], 2009

The Risk-to-Mission Assessment Process (RiskMAP): A Sensitivity Analysis and an Extension to Treat Confidentiality Issues

Jim Watters

The MITRE Corporation

Shaun Morrissey

The MITRE Corporation

Deborah Bodeau

The MITRE Corporation

Sue Cohn Powers

The MITRE Corporation

July 2009

This material is based upon work supported by the U.S. Department of Homeland Security under Grant Award Number 2006-CS-001-000001, under the auspices of the Institute for Information Infrastructure Protection (I3P) research program. The I3P is managed by Dartmouth College. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security, the I3P, or Dartmouth College. Approved for public release; distribution unlimited. MITRE No. 09-xxxx.

PREFACE

The Institute for Information Infrastructure Protection (I3P) was founded in September 2001 by the Department of Homeland Security (DHS) as a consortium of government, academic, and nonprofit institutions to coordinate research and development efforts in information infrastructure protection. The I3P is managed by Dartmouth College with funding from DHS and the National Institute of Standards and Technology. Partners must be institutions that work in the public interest. The 27 such institutions that are current partners include Massachusetts Institute of Technology's Lincoln Laboratory, the MITRE Corporation, Pacific Northwest National Laboratory, Sandia National Laboratories, SRI International, the United States Military Academy at West Point, the University of Illinois Urbana-Champaign, and the University of Tulsa.

The I3P has led a series of research projects investigating ways to advance the security of process control systems (PCS), which are crucial to many critical infrastructures. One of the current projects, Survivability and Recovery of Process Control Systems, extends the work of previous projects by addressing ways to increase the resiliency of PCS systems in the event of a cyber attack. This research report describes the work performed by The MITRE Corporation as part of the project team.

The authors wish to acknowledge the large number of people who contributed both to the advancement of RiskMAP and to the content of this report – in particular, Shimson Berkovits, Sheldon Durrant, Peter Kertzner, Bruce Lamar, Leonard Monk, and Jeff Picciotto. Many thanks go to these people for lending their knowledge, their passion, their insight and their frankness to the sensitivity analysis and the extension of the methodology to treat security issues of confidentiality.

EXECUTIVE SUMMARY

As part of the I3P's Survivability and Recovery of PCS project, The MITRE Corporation conducted a sensitivity analysis of its Risk-to-Mission Assessment Process (RiskMAP) methodology, and developed an extension to RiskMAP, to address Confidentiality as a security issue along with Integrity and Availability.

The initial purpose of the sensitivity analysis was to determine the range of conditions under which RiskMAP's calculation of relative weights for Tasks, Assets and Nodes would behave as order-preserving operations. Over the course of the sensitivity analysis, the RiskMAP team reexamined the methodology's mathematical foundations and the techniques used to generate the primary RiskMAP artifacts: A dependency network and a series of Pareto-style charts that rank-order Mission Objectives, Tasks, Information Assets, and Network Nodes.

While the sensitivity analysis confirmed that the RiskMAP application of Analytic Hierarchy Process (AHP) techniques is sound, the application of Quality Function Deployment (QFD) methods requires care to avoid over-simplification and misinterpretation of the Pareto charts. A number of refinements are developed and described that allow the user to identify and portray the criticality of each Task, Asset, or Node to a single Mission Objective.

The RiskMAP team also developed a methodological extension to enable separate treatment of Confidentiality, Integrity and Availability (C-I-A) within the basic RiskMAP framework. By introducing vectors to represent criticality and risk values with respect to C-I-A, the extension retains the overall character of the current approach. However, the change does increase the complexity and the data input load for the user. The RiskMAP team explored one possible implementation that would limit the added complexity and data input load by a customized MS Excel GUI backed up by a MS Access data base.

The results of the team's work provide improvements that can be applied individually or together in any future RiskMAP application.

ACRONYMS AND ABBREVIATIONS

AHP	Analytic Hierarchy Process
AMP	Assessment Management Platform
C#	C Sharp programming language
CS2SAT	Control System Cyber Security Self-Assessment Tool
C-I-A	Confidentiality, Integrity and/or Availability
COBIT	Control Objectives for Information and related Technology
COM	Component Object Model
CVE	Common Vulnerability Enumeration
DB	Data base
DHS	Department of Homeland Security
DIACAP	Department of Defense Information Assurance Certification and Accreditation Process
DLL	Dynamic Linked Library
DoD	Department of Defense
FIPS	Federal Information Processing Standards Publication
FISMA	Federal Information Security Management Act
GUI	Graphical User Interface
I3P	Institute for Information Infrastructure Protection
IA CAT	Information Assurance Compliance Assessment Tool
ID	Identifier
ISO	International Standards Organization
IT	Information Technology
MAAP	Mission Assurance Analysis Protocol
MORDA	Mission-Oriented Risk and Design Analysis
MS	Microsoft
NIST	National Institute of Standards and Technology
OCTAVE	Operationally Critical Threat, Asset, and Vulnerability Evaluation
PCS	Process Control System
QFD	Quality Function Deployment
RAMCAP	Risk Analysis And Management For Critical Asset Protection
RAPSA	Risk Analysis and Probabilistic Survivability Assessment
RiskMAP	Risk-to-Mission Assessment Process
SEMS	Security and Emergency Management System
UML	Universal Markup Language
VSAT	Vulnerability Self Assessment Tool
XML	eXtensible Markup Language

TABLE OF CONTENTS

Preface	i
Executive Summary	iii
Acronyms and Abbreviations	v
Table Of Contents	vii
List Of Figures	ix
List Of Tables	ix
Section 1: Introduction	1
Section 2: RiskMAP Revisited	3
2.1 Motivation	3
2.2 Matrix 1	3
2.3 Matrices 2 through 4	9
2.4 SUM vs. MAX Methods for Calculating Relative Weights	14
2.5 Handling Overlapping Dependencies through Vectorization	15
Section 3: Adding Confidentiality to the Mix	21
3.1 Motivation	21
3.2 What Are Others Doing?	22
3.3 Extending RiskMAP to Handle Confidentiality	22
3.4 Implementation	28
Section 4: Conclusions	35
Appendix: References and Bibliography	37

LIST OF FIGURES

Figure 1 - Typical case of Matrix 1.....	3
Figure 2 - Two Objectives.....	4
Figure 3 - Three Objectives.....	4
Figure 4 - Four Objectives	4
Figure 5 - Five Objectives	5
Figure 6 - Sample problem – Saaty’s solution	7
Figure 7 - Sample problem - RiskMAP solution.....	8
Figure 8 - Sample problem - ideal solution.....	8
Figure 9 - Pareto view of Mission/Business Objective Weights.....	9
Figure 10 - Matrix 2.....	10
Figure 11 - Pareto 2.....	10
Figure 12 - RiskMAP model	11
Figure 13 - Test Case 1a.....	12
Figure 14 - Test Case 1b.....	12
Figure 15 - Test Case 2a.....	13
Figure 16 - Test Case 2b.....	13
Figure 17 – Node Relative Weights via Summing All Paths.....	14
Figure 18 - Node Relative Weights via Maximum Path	14
Figure 19 - Summing All Paths in a More Complex Case.....	15
Figure 20 – Vector Weights When Summing All Paths	16
Figure 21 - Maximum Path in a More Complex Case	16
Figure 22 - Vector Weights When Taking Maximum Paths.....	16
Figure 23 - Matrix representation of current RiskMAP calculations.....	17
Figure 24 - Introducing Dependency Values	17
Figure 25 – Adding the Second Dimension to Dependency.....	17
Figure 26 - Extended RiskMAP Structure.....	25
Figure 27 - Extended RiskMAP Implementation.....	25
Figure 28 – Single-mode data base structure.....	29
Figure 29 – Single-mode data relationships.....	30
Figure 30 – Triple-mode data base structure.....	31
Figure 31 – RiskMAP Triple Add-In.....	32
Figure 32 - Sorting dialog	33

LIST OF TABLES

Table 1 – Symbols and Definitions.....	18
Table 2 - Survey of Risk Management Methods and Tools.....	23
Table 3 - Criticality of Information Asset to Operational Task.....	26
Table 4 - Criticality of Network Node (or Link) to Information Asset.....	27

SECTION 1: INTRODUCTION

As part of its role in I3P-led research on enhancing PCS security, The MITRE Corporation proposed in 2008 to conduct a sensitivity analysis of its Risk-to-Mission Assessment Process (RiskMAP) methodology and to extend RiskMAP to address Confidentiality as a security issue along with Integrity and Availability. This report details the results of this work.

Previous results reported in (1) have explored scalability and cardinality¹, and much of that work is referenced in this report either explicitly or implicitly. The reader is encouraged to become familiar with the referenced work in order to have the best appreciation of the current report.

This report is organized as follows: Following the introduction in Section 1, the sensitivity analysis is discussed in Section 2. The topic of treating Confidentiality issues with the RiskMAP methodology is addressed in Section 3, and this is followed in Section 4 with a list of conclusions.

The sensitivity analysis and the confidentiality extension were undertaken as parallel efforts. For that reason, Section 3 does not build on Section 2. However, the ideas presented in the two sections can be combined into a single implementation.

¹ The former looked at how well the RiskMAP method handles a large enterprise; the latter examined the scale values being used for evaluating the criticality of Tasks, Information Assets and Network Nodes.

SECTION 2: RISKMAP REVISITED

Within the scope of the I3P PCS Security projects, the RiskMAP methodology has been applied to large and small cases within the Oil and Gas sector of the critical national infrastructure. In addition, RiskMAP has been applied to a number of other cases in the government sector by other project teams. Based on the increasing interest in the methodology, The MITRE Corporation’s RiskMAP development team determined that due diligence called for a new examination of the method. Initially undertaken as a sensitivity analysis, the study grew to include a review of the fundamental mathematics underlying the methodology.

2.1 Motivation

The initial purpose of the sensitivity analysis was to determine the range of conditions under which RiskMAP’s calculation of relative weights for Tasks, Assets and Nodes would behave as order-preserving operations. This is important since an increasing number of users find the Pareto charts to be highly valuable for situational awareness and for decision support. Over the course of the sensitivity analysis, the RiskMAP team reexamined the mathematical foundations for the Pareto charts’ generation and, in fact, for the construction of the dependency paths that are the hallmark of the RiskMAP approach. The following sections trace the team’s examination of Matrix 1’s use of the Analytic Hierarchy Process (AHP) and of the adaptation of Quality Function Deployment (QFD) techniques in Matrices 2 – 4.

2.2 Matrix 1

As explained in (1), the purpose of Matrix 1 is to capture the view of the user regarding the relative importance of Mission/ Business Objectives when taken a pair at a time. Refer to Figure 1 below, showing a typical case of Matrix 1. Having already entered the Mission/Business Objectives to be compared, the user is asked to supply values for their comparison from the scale at the upper left of the figure. As explained by AHP’s creator, Thomas Saaty (2), this “fundamental scale” is used to derive a second, implicit scale which will reveal the comparative weights given to the Mission/Business Objectives. Saaty also observes in (2) that non-integer values may be used to compare items that are very close in value to one another. Such practice has been part of the field uses of RiskMAP as illustrated by the example in Figure 1.

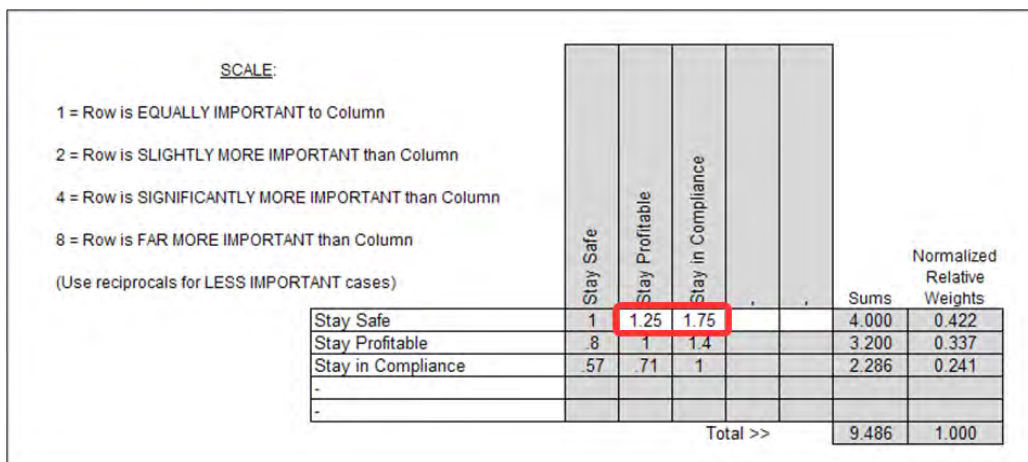


Figure 1 - Typical case of Matrix 1

In an important departure from typical AHP implementations, RiskMAP enforces consistency among the values that fill up Matrix 1. Where normally a user would be asked to provide values for all 16 of the cells used above for comparison scores, RiskMAP only allows user input in the unshaded cells in the first row. The cells along the diagonal contain 1’s because they show comparisons of each Mission/Business Objective

to itself. All entries below the diagonal are reciprocals of their opposites above the diagonal – a standard AHP practice. The departure is that all super diagonal elements below the first row are derived from the entries in the first row. For example, comparing Objective 1 to Objective 2 yields a value of 1.25 while a similar comparison of Objective 1 to Objective 3 yields a value of 1.75. The comparison of Objective 2 to Objective 3 can be found from the ratio of 1.75 to 1.25, or 1.4. This is a condition which must be met in order for the model to be internally consistent. An observable feature of model consistency is that the ratio of any two Objectives' relative weights remains the same as additional Objectives are added to the model. Figure 2 through Figure 5 below illustrate that concept.

SCALE:
 1 = Row is EQUALLY IMPORTANT to Column
 2 = Row is SLIGHTLY MORE IMPORTANT than Column
 4 = Row is SIGNIFICANTLY MORE IMPORTANT than Column
 8 = Row is FAR MORE IMPORTANT than Column
 (Use reciprocals for LESS IMPORTANT cases)

	Objective 1	Objective 2	Objective 3	Objective 4	Objective 5	Relative Weights	Ratios
Objective 1	1	2				3	0.667
Objective 2	0.5	1				1.5	0.333
Objective 3							
Objective 4							
Objective 5							
						4.5	1.000

2.0 (Obj1:Obj2)
 (Obj2:Obj3)
 (Obj3:Obj4)
 (Obj4:Obj5)

Figure 2 - Two Objectives

SCALE:
 1 = Row is EQUALLY IMPORTANT to Column
 2 = Row is SLIGHTLY MORE IMPORTANT than Column
 4 = Row is SIGNIFICANTLY MORE IMPORTANT than Column
 8 = Row is FAR MORE IMPORTANT than Column
 (Use reciprocals for LESS IMPORTANT cases)

	Objective 1	Objective 2	Objective 3	Objective 4	Objective 5	Relative Weights	Ratios
Objective 1	1	2	4			7	0.571
Objective 2	0.5	1	2			3.5	0.286
Objective 3	0.25	0.5	1			1.75	0.143
Objective 4							
Objective 5							
						12.25	1.000

2.0 (Obj1:Obj2)
 2.0 (Obj2:Obj3)
 (Obj3:Obj4)
 (Obj4:Obj5)

Figure 3 - Three Objectives

SCALE:
 1 = Row is EQUALLY IMPORTANT to Column
 2 = Row is SLIGHTLY MORE IMPORTANT than Column
 4 = Row is SIGNIFICANTLY MORE IMPORTANT than Column
 8 = Row is FAR MORE IMPORTANT than Column
 (Use reciprocals for LESS IMPORTANT cases)

	Objective 1	Objective 2	Objective 3	Objective 4	Objective 5	Relative Weights	Ratios
Objective 1	1	2	4	4		11	0.500
Objective 2	0.5	1	2	2		5.5	0.250
Objective 3	0.25	0.5	1	1		2.75	0.125
Objective 4	0.25	0.5	1	1		2.75	0.125
Objective 5							
						22	1.000

2.0 (Obj1:Obj2)
 2.0 (Obj2:Obj3)
 1.0 (Obj3:Obj4)
 (Obj4:Obj5)

Figure 4 - Four Objectives

SCALE:
 1 = Row is EQUALLY IMPORTANT to Column
 2 = Row is SLIGHTLY MORE IMPORTANT than Column
 4 = Row is SIGNIFICANTLY MORE IMPORTANT than Column
 8 = Row is FAR MORE IMPORTANT than Column
 (Use reciprocals for LESS IMPORTANT cases)

	Objective 1	Objective 2	Objective 3	Objective 4	Objective 5	Relative Weights	Ratios
Objective 1	1	2	4	4	8	19	0.471
Objective 2	0.5	1	2	2	4	9.5	0.235
Objective 3	0.25	0.5	1	1	2	4.75	0.118
Objective 4	0.25	0.5	1	1	2	4.75	0.118
Objective 5	0.125	0.25	0.5	0.5	1	2.375	0.059
						40.375	1.000

Figure 5 - Five Objectives

Note that as each new Objective is added, the relative weights do change, but their ratios do not. The condition that must be met to guarantee this outcome can be derived as follows. Let us say that the desired outcome is that the ratio of two Objective weights, $W_{mo_i}; W_{mo_{(i+1)}}$ stays constant. Beginning with a two-Objective set, Matrix 1 would look like this:

	Objective 1	Objective 2	Sum of Row	Relative Weights
Objective 1	1	m_{12}	$1 + m_{12}$	W_{mo_1}
Objective 2	$\frac{1}{m_{12}}$	1	$\frac{1}{m_{12}} + 1$	W_{mo_2}
Total:			$m_{12} + \frac{1}{m_{12}} + 2$	$W_{mo_1} + W_{mo_2} = 1.0$

The Mission Objectives' relative weights W_{mo_1} and W_{mo_2} would be given by:

$$W_{mo_1} = \frac{1 + m_{12}}{m_{12} + \frac{1}{m_{12}} + 2} \quad \text{and} \quad W_{mo_2} = \frac{\frac{1}{m_{12}} + 1}{m_{12} + \frac{1}{m_{12}} + 2}$$

The ratio of the two relative weights would then be:

$$(1) \quad \frac{W_{mo_1}}{W_{mo_2}} = \frac{1 + m_{12}}{\frac{1}{m_{12}} + 1} = m_{12}.$$

If a third Objective is added, Matrix 1 would look like this:

	Objective 1	Objective 2	Objective 3	Sum of Row	Relative Weights
Objective 1	1	m_{12}	m_{13}	$1 + m_{12} + m_{13}$	W_{mo_1}
Objective 2	$\frac{1}{m_{12}}$	1	m_{23}	$\frac{1}{m_{12}} + 1 + m_{23}$	W_{mo_2}
Objective 3	$\frac{1}{m_{13}}$	$\frac{1}{m_{23}}$	1	$\frac{1}{m_{13}} + \frac{1}{m_{23}} + 1$	W_{mo_3}
Total:				$3 + m_{12} + m_{13} + m_{23} + \frac{1}{m_{12}} + \frac{1}{m_{13}} + \frac{1}{m_{23}}$	$W_{mo_1} + W_{mo_2} + W_{mo_3} = 1.0$

Following the same path as for the two-Objective case and canceling out the common denominator, the ratio of W_{mo_1} to W_{mo_2} would be:

$$(2) \quad \frac{W_{mo_1}}{W_{mo_2}} = \frac{1 + m_{12} + m_{13}}{\frac{1}{m_{12}} + 1 + m_{23}} = \frac{1 + m_{12} + m_{13}}{\left(\frac{1 + m_{12} + m_{12}m_{23}}{m_{12}} \right)}$$

For the ratio of W_{mo_1} to W_{mo_2} to be constant between the two-Objective case and the three-Objective case, the expressions in equations (1) and (2) must be equal as stated in equation (3).

$$(3) \quad \frac{1 + m_{12} + m_{13}}{\left(\frac{1 + m_{12} + m_{12}m_{23}}{m_{12}} \right)} = m_{12}$$

Simplifying, this becomes

$$(4) \quad \frac{m_{13}}{m_{12}} = m_{23}.$$

Equation (4) represents the condition which must be met to keep the ratio $W_{m_{01}}:W_{m_{02}}$ constant when the set of Objectives is increased from two to three. The condition expressed in equation (4) can be generalized and applied to all Matrix 1 elements above the diagonal and below the first row:

$$(5) \quad \text{For } k > j \text{ and } i > 1, \quad m_{jk} = \frac{m_{ik}}{m_{ij}} = \frac{m_{1k}}{m_{1j}}.$$

This means that once the first row of values has been entered, the elements in the remaining rows can be calculated -- reducing user input and alleviating the concern about potential input or model inconsistencies. And since the condition in equation (5) will be met, the ratio $W_{m_{0i}}:W_{m_{0j}}$ will remain constant as the number of Mission Objectives varies.

How well does this implementation work? As a test, consider an example from Saaty (2) where AHP is used to derive the relative magnitude of five objects' surface area, given that the objects are all of different shapes. See Figure 6 below. Saaty uses a 1-9 scale for the comparisons and employs an Eigenvector technique to

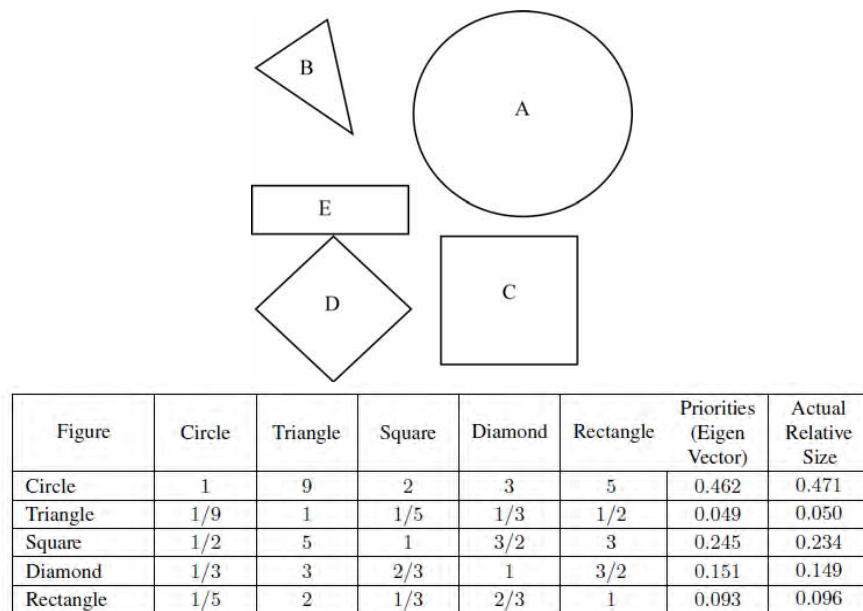


Figure 6 - Sample problem – Saaty's solution

correct for inconsistency among user inputs. (For example, observe that while Circle:Diamond = 3 and Circle:Rectangle = 5, the user inputs Diamond:Rectangle = 3/2, and not 5/3). What answers would the RiskMAP implementation come up with? Adopting the same 1-9 scale values (ignoring the suggested 1, 2, 4, 8 in Figure 7 below), RiskMAP produces values that compare more favorably to the Actual Relative Sizes than do the values produced by the Eigenvector-corrected method. The average error in the RiskMAP values

is .0034, compared to .0052 for the Eigenvector method. Why is there still an error at all? If instead of the integer values 1-9, we were use non-integer values to express more precise estimates that should eliminate some or all of the residual error.

SCALE:

1 = Row is EQUALLY IMPORTANT to Column
 2 = Row is SLIGHTLY MORE IMPORTANT than Column
 4 = Row is SIGNIFICANTLY MORE IMPORTANT than Column
 8 = Row is FAR MORE IMPORTANT than Column
 (Use reciprocals for LESS IMPORTANT cases)

	Circle	Triangle	Square	Diamond	Rectangle	Sums	Normalized Relative Weights
Circle	1	9	2	3	5	20.000	0.466
Triangle	.11	1	.22	.33	0.556	2.222	0.052
Square	.5	4.5	1	1.5	2.5	10.000	0.233
Diamond	.33	3.	.67	1	1.667	6.667	0.155
Rectangle	0.2	1.8	0.4	0.6	1	4.000	0.093
Total >>						42.889	1.000

Figure 7 - Sample problem - RiskMAP solution

To simulate the best possible user inputs, let us use the actual ratios of the shapes' surface areas. Using the actual relative size values from Figure 6, we get the following as precise inputs for Matrix 1.

Circle:Triangle = 0.471/0.050 = 9.420

Circle:Square = 0.471/0.234 = 2.013

Circle:Diamond = 0.471/0.149 = 3.161

Circle:Rectangle = 0.471/0.096 = 4.906

Using these values, Matrix 1 looks like Figure 8 below, yielding the exact values in Saaty's example.

SCALE:

1 = Row is EQUALLY IMPORTANT to Column
 2 = Row is SLIGHTLY MORE IMPORTANT than Column
 4 = Row is SIGNIFICANTLY MORE IMPORTANT than Column
 8 = Row is FAR MORE IMPORTANT than Column
 (Use reciprocals for LESS IMPORTANT cases)

	Circle	Triangle	Square	Diamond	Rectangle	Sums	Normalized Relative Weights
Circle	1.000	9.420	2.013	3.161	4.906	20.500	0.471
Triangle	.106	1.000	.214	.336	.521	2.176	0.050
Square	.497	4.680	1.000	1.570	2.437	10.184	0.234
Diamond	.316	2.980	.637	1.000	1.552	6.485	0.149
Rectangle	.204	1.920	.410	.644	1.000	4.179	0.096
Total >>						43.524	1.000

Figure 8 - Sample problem - ideal solution

This demonstrates that the RiskMAP implementation of AHP does not introduce errors in the generation of relative weights.

What about rank reversals? In the literature are discussions of cases where AHP can actually generate answers like $B > A$ where in reality $B < A$. Triantaphyllou (3) discusses conditions under which ranking irregularities can occur, but these pertain to later steps in the application of AHP; that is, in the evaluation of alternatives once the step of generating relative weights of criteria (or in our parlance, Objectives) has been completed. The RiskMAP implementation does not take AHP this far; only to the point of establishing the relative weights among the Objectives. Thus, the rank reversals are avoided.

If the AHP as implemented in RiskMAP does not add error to user inputs, then the user inputs themselves are the remaining source for potential errors. If the user under- or overestimates an Objective's importance relative to another, the resulting relative weights will faithfully reflect that error. To help minimize the chance of such an error of estimation, RiskMAP includes a validation step using visual feedback. Once Matrix 1 is populated, the relative weights are displayed in Pareto form (see Figure 9 below) so that the user can observe



Figure 9 - Pareto view of Mission/Business Objective Weights

both the trends in importance and the magnitude of the differences. It has been the experience of the RiskMAP development team, both in field uses for the I3P project and with other user groups, that the validation step is both useful and effective. When viewing the Pareto charts, users have regularly been able either to verify the correctness of their inputs or to note errors and make appropriate corrections. What they could not see in the tabular view of Matrix 1 (in this case, the one given in Figure 1), they could see in the graphical view of the Pareto chart. Between the two displays they have been able to arrive at a set of weights that accurately reflects their view of their organization's mission priorities.

Having addressed the potential sources of error – both user-induced and process-induced – and found that the latter is precluded by design and that the former is reasonably controlled through a validation step, the RiskMAP implementation of AHP in Matrix 1 is considered to provide an accurate mapping of the user's view into terms required by the model.

2.3 Matrices 2 through 4

As explained in this project's previous report (1), the Pareto charts provide a view of the Tasks, Assets or Nodes that are most critical to achieving Mission/Business Objectives. As in the case of Matrix 1, the Pareto view provides visual feedback to the user to aid in validating the inputs made to that point. However, the Pareto views for Matrices 2 through 4 are derived quite differently than those for Matrix 1. Matrices 2 through 4 use an adaptation of QFD to generate the Pareto views and while it has been described in (1), some review will be helpful here.

The matrix (or Figure 10) below shows how the Mission/Business Objectives are repeated at the lower left, along with their relative weights as calculated in Matrix 1. As detailed in (1), the Tasks needed to achieve the Objectives are listed across the top and then scores indicating each Task's criticality to each Objective are entered in the intersecting cells. The relative weight of a Task is calculated as the sum of the products of that Task's criticality to each Objective and the Objective's relative weight.

Impact of Task Loss:	Task	Task Weight	Task Criticality													
		0 No impact 30 Work Around 70 Degrade 95 Failure	17.35	10.12	62.53	23.61	23.61	23.61	52.41	23.61	17.35	40.48	10.12	53.13	78.43	78.43
Objective	Weight															
Stay Safe	0.422			70				70						30	70	70
Stay Profitable	0.337	30	30	30	70	70	70		70	30	70	30	70	95	95	
Stay in Compliance	0.241	30		95						95		30	70		70	70

Figure 10 - Matrix 2

The process is the same for Matrix 3, where the relative weight of each Information Asset is calculated as the sum of the products of that Asset's criticality to each Task and the Task's relative weight. In Matrix 4, the relative weight of each Network Node is calculated as the sum of the products of that Node's criticality to each Asset and the Asset's relative weight.

At the completion of each round of scoring, the results are sorted to create a Pareto view for use as visual feedback so the users can judge the accuracy of their inputs. The Pareto view for the Matrix 2 above is shown in Figure 11.




Figure 11 - Pareto 2

The sensitivity analysis began by asking the question: What would an error look like in a Pareto view? The initial answer was that an error would look like an incorrect order from left to right. That is, a Task, Asset or Node appears in a position that is not in keeping with its true importance relative to the others. Since the Pareto views result from a sort by declining relative weight, an error in order directly translates into an error in magnitude². Errors in magnitude are not only a concern because they drive the order in a Pareto view; they are of additional concern because users have consistently used the height of the Pareto bars to infer the absolute importance of the Tasks, Assets and Nodes. This is more information than the RiskMAP method intended to provide, so an examination of the means used to generate the relative weights is certainly in order.

To begin with, what information is being conveyed from the top of the RiskMAP model to the bottom? In the model, shown in Figure 12 below, if one leaves out the Objective relative weights and only considers the lattice depicting the dependency of Objectives on Tasks, of Tasks on Assets and of Assets on Nodes, then it is dependence that is being conveyed from the top layer to the bottom.




Figure 12 - RiskMAP model

The full RiskMAP model, however, also includes the Mission/Business Objectives' relative weight, or priority. Thus the model is conveying two types of information to the bottom: dependence and priority. For a large and complex organization the RiskMAP model can itself be large and complex. The method used to generate the relative weights intentionally jettisons detail in order to provide the user with information that is simple enough to grasp and yet accurate enough to be reliable. But is it reliable? Exactly what are the relationships between the levels of the model? What is the behavior of the mathematical operations under different input conditions?

To examine the potential for inducing an error in the relative weights, we need to look first at some very simple examples. To begin, let us consider a case where an organization has but two Mission/Business Objectives, o1 and o2. Objective o1 is supported by one Task, t1. Task t1 depends on one Information Asset a1 which, in turn, depends on one Node n1. All dependencies are absolute; that is, the criticality scores are all set at 50, which means that degradations will result at any level if disruptions occur in the next level down. Objective o2 depends on ten Tasks, t2 through t11. Each Task depends on one Information Asset so there are Assets a2 through a11. Each Asset depends on one Network Node so there are Nodes n2 through n11. Again, all criticality scores are set at 50. This case is illustrated in below.

² In cases of equal relative weight, the secondary sort criterion is usually by Task, Asset or Node number or name.

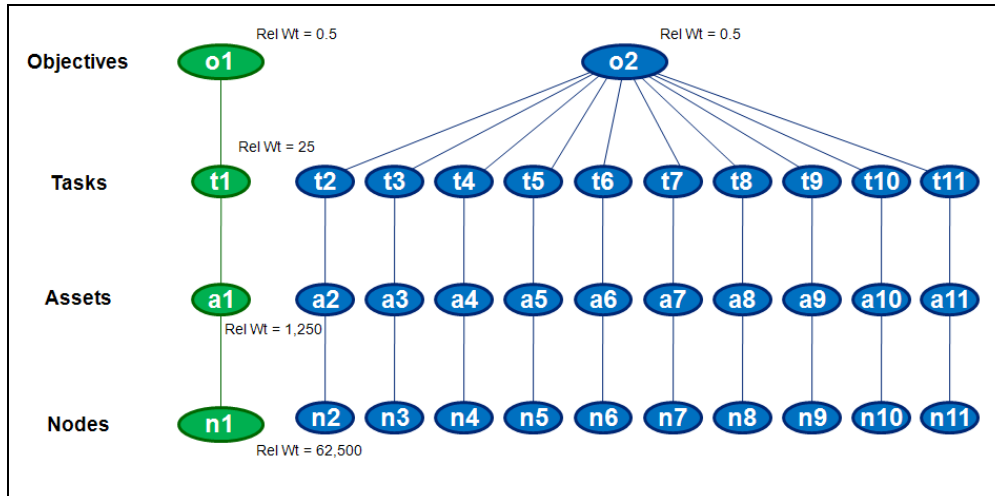


Figure 13 - Test Case 1a

Note that the Objective relative weights are each 0.5; each has equal priority. Each of the Tasks has a weight of $0.5 * 50 = 25$ since there is only one dependency path to each Task so the sum or the products collapses to a single product. Similarly, each Asset is weighted as $25 * 50 = 1,250$ and each Node is weighted as $1,250 * 50 = 62,500$. In the Pareto view, all Nodes would have a bar of equal height. In a minor variation of this case, let us change the Objective relative weights to give higher priority to Objective o1. In Figure 14 below, Objective o1 is now weighted four times as heavily as Objective o2.




Figure 14 - Test Case 1b

Since none of the organizational dependencies has changed, the change in priority among the Mission/Business Objectives is immediately reflected in the weightings of the Nodes that support o1 and o2, respectively. Node n1 is now weighted four times as heavily as each Node n2 through n11. The message implicit in this weighting is that Node n1 is currently the most critical to the organization's overall mission and since Objective o1 is currently of primary importance, the Node weighting seems to make sense. However, let us look at a somewhat different example.

In the next example, the number and relationship of Objectives, Tasks and Assets will be the same as before. However, in this case the Assets supporting Objective o2, meaning Assets a2 through a11, all depend on one Node, n2. As shown in Figure 15 below, all criticality scores remain at 50 and the Objective relative weights are each 0.5.

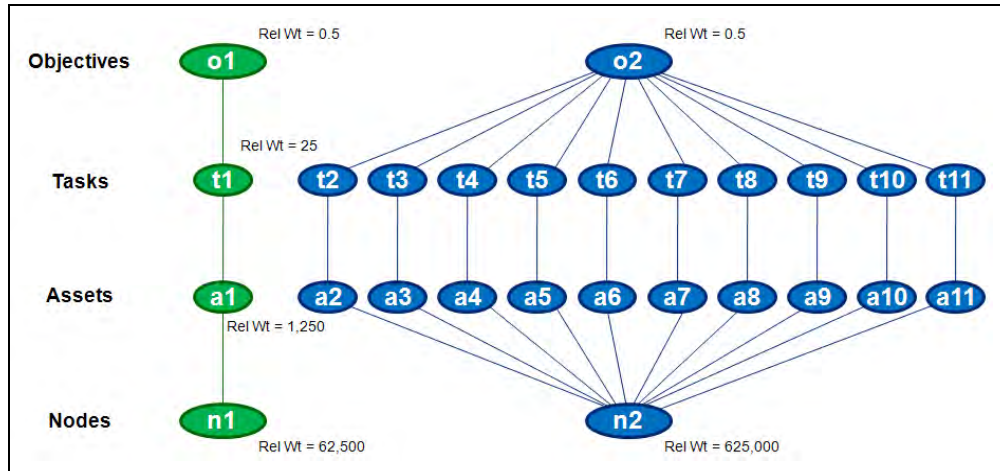


Figure 15 - Test Case 2a

As before, each Task has a relative weight of $0.5 * 50 = 25$ and each Asset has a relative weight of $25 * 50 = 1,250$. However, while Node n1 has a relative weight of $1,250 * 50 = 62,500$, the sum of the products for Node n2 becomes $10 * (1,250 * 50) = 625,000$. On the Pareto chart, Node n2's bar height would be 10 times the size of Node n1's. A common inference is that Node n2 is ten times as important as Node n1, and yet each Node was stated to be indispensable to a Mission Objective. Before commenting further, let us look at one more example.

In the final case, let the relationships in the model be as in Figure 15 above but now, let Objective o1 be weighted four times as heavily as Objective o2. This case is shown in Figure 16 below.

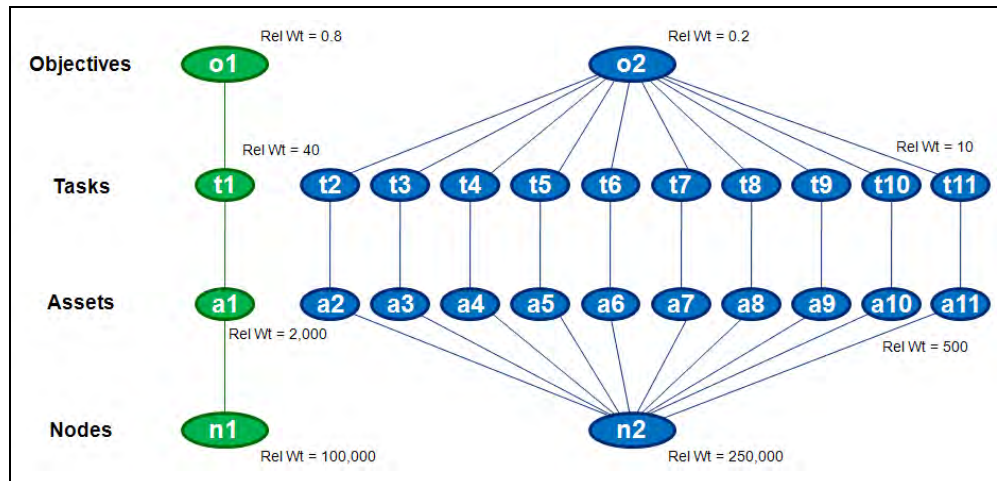


Figure 16 - Test Case 2b

Once again the relative weights for the Tasks and Assets follow the ratio of the Objective relative weights. Task t1, supporting Objective o1, has a relative weight four times that for any Task supporting Objective o2. The same is true for the Asset relative weights. What about the Node relative weights? Node n1 has a relative weight of $2,000 * 50 = 100,000$. But for Node 2, the sum of the products becomes $10 * (500 * 50) = 250,000$. On the Pareto chart, Node n2's bar height would be 2.5 times the size of Node n1's. That is, the 4:1 ratio of Objective relative weights has been overcome by summing ten dependency paths into Node n2. So even though Objective o1 has by far the greatest priority, why does the sole Node supporting it carry less weight? Does that constitute an error? The answer is that it depends on what the user intends to portray in the Pareto charts. The following discussion will examine alternative methods for generating Node relative weights: The "SUM" method illustrated above, and the "MAX" method.

2.4 SUM vs. MAX Methods for Calculating Relative Weights

If one wishes to show both the degree of dependency (the number of dependency paths from Objective to Node and their criticality) and the priority (based on Objective relative weight), then the “SUM” method described above will meet the need. The Node relative weight can be seen as the sum of all dependency paths (from each Objective to the Node), each path multiplied by the relative weight of the Objective that it supports. However, if one wishes to identify the cases where a Node is most critical to a Mission Objective, there is another way to calculate and portray Node relative weight – that being to show only the maximum of the products of dependency path and Objective relative weight. This method will be called the “MAX” method.

To illustrate, let us compare two cases where we can compare solely the effects of summing dependency paths versus taking only the maximum path. In Figure 17 below is a case where two Objectives are dependent on different sets of Tasks, Assets and Nodes. It will be helpful here to use the terms “parent” and “child” where, for instance, a Task will be a child of one or more Objectives and will also be a parent of one or more Assets. Using these terms, it can be seen that Objective 1 has two children, each of which has the same child. The one Asset has but one child. The relative weight of the Node (425) is simply the product of the path value (i.e. criticality score, or 5 in this case) and the relative weight of the Node’s parent (85). However, the relative weight of the Asset is the sum of the products of the path values and their parent weights ($2 * 2$) + ($9 * 9$) = 85. For the Task weights, the sum of the products collapses to single products $2 * 1 = 2$ and $9 * 1 = 9$. The Objectives are given unity weight in this case.




Figure 17 – Node Relative Weights via Summing All Paths

A Pareto view showing Node relative weights would order the Nodes as Node 1, Node 3, and Node 2. To see the difference in using the maximum versus the sum of all paths, let us modify the example as shown in Figure 18.




Figure 18 - Node Relative Weights via Maximum Path

Now each Node's relative weight is taken as the maximum product of a path leading to it and the weight of the parent associated with that path. The heavy lines indicate the maximum paths to each Node from each Objective. In this case, a Pareto view showing Node relative weights would order the Nodes as Node 1, Node 2, and Node 3. Also, note that the relative weights for the Assets change with the mode of calculation.

Which approach is correct? Again, it depends on what is to be portrayed. If the user is interested in each Node's maximum contribution to the overall mission, then the MAX calculation is the better of the two. If one is interested in seeing the breadth, or degree, of dependency on each Node, the SUM calculation will provide that view. During field uses of RiskMAP, some users have felt that the SUM method accurately reflects their view of their organization's dependencies. Other users have been more interested in finding Network Nodes whose failure would cause mission failure – even though the Nodes may not be broadly used. For those cases, the MAX method might be more appropriate. Regardless of whether the SUM or the MAX method is used, the alternatives should be discussed early and the choice clearly stated as an assumption underpinning the RiskMAP analysis.

2.5 Handling Overlapping Dependencies through Vectorization

In the case examined in Figure 17 and Figure 18 it was easy to show each Node's importance to one Objective at a time because the case did not involve any overlapping dependencies. But in the overwhelming cases where Nodes support more than one Objective, the criticalities get mixed together, or confounded. This is an expected result of the simplifying action of QFD. However, it clouds the relationship between each level of the model and prevents the use of the MAX mode of calculation. To remedy this problem it is necessary to keep separate the two kinds of information being conveyed through the model: dependence and priority. In short, the approach is to represent each relative weight not as a scalar but as a vector where each element of the vector is ultimately associated with support to one Mission Objective. To illustrate, let us add an overlapping dependency to the previous example and show how using a vector approach clarifies the picture for both the SUM and MAX methods.




Figure 19 - Summing All Paths in a More Complex Case

Comparing Figure 19 to Figure 17, note that the second Node's weight now reflects added importance due to supporting both Objectives. But how much of that importance goes to each Objective? To see that requires the vector approach illustrated below. In Figure 20, each Node's contributions to each Objective can clearly be seen in the components of its vector of relative weights. The same is true at the Asset and Task levels of the model. By taking these individual components of the relative weights, one can trace the dependency of a single Mission Objective on a Task, an Information Asset, or a Node. If one wishes to use the SUM results, they are still available by summing up the weights for each Task, or Asset, or Node.

Similar benefits can be seen in adding the vector approach to the MAX path calculation. Comparing Figure 21 to Figure 18, the introduction of the overlapping dependency can be seen. In Figure 22 below, the results are expressed as vectors. Again, each Node's maximum contribution to each Objective is readily seen, as is also true for the Asset and Task levels of the model. Again, individual mission dependencies can be seen at any level.




Figure 20 – Vector Weights When Summing All Paths




Figure 21 - Maximum Path in a More Complex Case




Figure 22 - Vector Weights When Taking Maximum Paths

Expressing the relative weights as vectors clearly removes the ambiguity about which Objective(s) a given element of the model supports and to what degree. Implementing this approach in RiskMAP amounts to adding new dependency matrices at each level.

Vector Implementation

In the new formulation, we separate dependency information and priority information, and use a different combining rule than we have in the past. Previously, we had defined and used the following:

$$\begin{aligned}
 trw &= \mathbf{TC} \, orw \\
 arw &= \mathbf{AC} \, trw = \mathbf{AC} \, \mathbf{TC} \, orw \\
 nrw &= \mathbf{NC} \, arw = \mathbf{NC} \, \mathbf{AC} \, \mathbf{TC} \, orw
 \end{aligned}$$

Figure 23 - Matrix representation of current RiskMAP calculations

In Figure 23 above, **NC**, **AC**, and **TC** are the Node Criticality, Asset Criticality, and Task Criticality matrices, respectively. The vectors *orw*, *trw*, *arw*, and *nrw* are Objective, Task, Asset, and Node relative weights.

Treating the RiskMAP model as a directed graph and each of the Objectives, Tasks, Assets and Nodes as vertices on the graph, the new approach treats the dependency information at each vertex in the RiskMAP graph as a vector, with the cardinality of the set of Objectives as the dimension of the vector.

Changes and Constants

First, the *orw* vector containing the AHP output of relative weights is retained. However, it is no longer used as the first vector in the multiplication as shown above. The second issue is separating priority from the value of dependency at each vertex (Objective, Task, Asset, Node). To manage this aspect, we define new entities carrying dependency information at each level. For the four levels, we define *ord*, *trd*, *ard*, and *nrd* as the Objective, Task, Asset, and Node dependency values, respectively. They are shown in Figure 24 below.

$$\begin{aligned}
 trd &= \mathbf{TC} \, ord \\
 ard &= \mathbf{AC} \, trd = \mathbf{AC} \, \mathbf{TC} \, ord \\
 nrd &= \mathbf{NC} \, ard = \mathbf{NC} \, \mathbf{AC} \, \mathbf{TC} \, ord
 \end{aligned}$$

Figure 24 - Introducing Dependency Values

In the existing setup, these characterizations are vectors at each level, with their single index running over Objectives, Tasks, Assets, or Nodes, as appropriate. In the new setup, each vertex carries a vector, indexed by the Objectives list. This can be encoded by replacing each of the vectors shown in Figure 24 with a matrix whose added index would run from 1 to the number of Objectives. See Figure 25 below.

$$\begin{aligned}
 \mathbf{TRD} &= \mathbf{TC} * \mathbf{ORD} \\
 \mathbf{ARD} &= \mathbf{AC} * \mathbf{ARD} = \mathbf{AC} * \mathbf{TC} * \mathbf{ARD} \\
 \mathbf{ARD} &= \mathbf{NC} * \mathbf{ARD} = \mathbf{NC} * \mathbf{AC} * \mathbf{TC} * \mathbf{ARD}
 \end{aligned}$$

Figure 25 – Adding the Second Dimension to Dependency

Here, the **ORD** matrix is a diagonal identity matrix (the dependence of Objective i on Objective j)³. The elements in the **TRD**, **ARD** and **NRD** matrices are the components of each Task's, each Asset's, and each Node's criticality that apply to each Objective. This does beg the issue of what operation is represented by the * shown in the above figure.

The normal operation for matrix multiplication looks like the following:

$$(6) \quad ard_{mj} = \sum_{k=1}^{\#Tasks} [ac_{mk} * trd_{kj}]$$

where the symbols are defined below. If we want to use the MAX approach to formulating relative weights rather than the SUM, this becomes:

$$(7) \quad ard_{mj} = \max_{k=1}^{\#Tasks} [ac_{mk} * trd_{kj}]$$

with similar definitions to be used at each level.

Explicit Formulation

The first step is to develop a consistent set of equations and indices to support implementation.

Symbol	Definition
orw_j	Objective relative weights, vector, index j , size is number of Objectives
trw_k	Task relative weights, vector, size is number of Tasks, content is Pareto ordering weight
arw_m	Asset relative weights, vector, size is number of Assets, content is Pareto ordering weight
nrd_n	Node relative weights, vector, size is number of Nodes, content is Pareto ordering weight
tc_{ki}	Element of TC ; the dependence of Objective i on Task k , from user community
ac_{mk}	Element of AC ; the dependence of Task k on Asset m , from user community
nc_{nm}	Element of NC ; the dependence of Asset m on Node n , from user community
ord_{ij}	Element of ORD ; the dependence of Objective j on Objective i (identity matrix)
trd_{kj}	Element of TRD ; the dependence of Objective j on Task k (non-square matrix)
ard_{mj}	Element of ARD ; the dependence of Objective j on Asset m (non-square matrix)
nrd_{nj}	Element of NRD ; the dependence of Objective j on Node n (non-square matrix)

Table 1 – Symbols and Definitions

The orw vector is the result of the AHP process to determine weights, as described in Section 2. The number of Objectives determines the dimensionality of the vector of dependency values associated with each vertex. With O Objectives, the value O is always one dimension of all of the dependence matrices (which are assembled by stacking the O -dimensional vectors for Objectives, Tasks, Assets, or Nodes into a matrix).

At any level, a set of values for a Pareto-ordering, or a weighting order, can be calculated for the vertices at that level by the inner product of the relative dependence matrix and the orw vector. This takes the form of a

³ This assumes that there are no cross-dependencies among the Objectives. Such cross-dependencies could be indicated by non-zero off-diagonal elements.

standard matrix multiplication of a vector by a matrix. The resulting vector holds the values that should drive the sorting into order. The formulas for calculation at each level are:

$$(8) \quad \begin{aligned} trw_k &= \sum_{j=1}^{\#Objectives} trd_{k,j} * orw_j \\ arw_m &= \sum_{j=1}^{\#Objectives} ard_{m,j} * orw_j \\ nrw_n &= \sum_{j=1}^{\#Objectives} nrd_{n,j} * orw_j \end{aligned}$$

These formulations directly relate the relative weight for a given Task, Asset or Node to a given Objective weight via the dependency associated **only** with the dependencies along the path between the Objective and the Task, Asset or Node.

The matrix values for **TC**, **AC**, and **NC** are determined from user input.

Calculation of Relative Dependence Matrices

As discussed above, if the SUM method is to be use to generate the relative weights, then standard matrix multiplication is used. These expressions are given below, similar to equation (6).

$$(9) \quad \begin{aligned} trd_{kj} &= \sum_{i=1}^{\#Objectives} [tc_{ki} * ord_{ij}] \\ ard_{mj} &= \sum_{k=1}^{\#Tasks} [ac_{mk} * trd_{kj}] \\ nrd_{nj} &= \sum_{m=1}^{\#Assets} [nc_{nm} * ard_{mj}] \end{aligned}$$

If the MAX method is to be used, it will require a modification of the standard matrix multiplication operation, to capture the most significant dependence on each Objective at each vertex. These modified expressions are shown below, similar to equation (7).

$$(10) \quad \begin{aligned} trd_{kj} &= \max_{i=1}^{\#Objectives} [tc_{ki} * ord_{ij}] \\ ard_{mj} &= \max_{k=1}^{\#Tasks} [ac_{mk} * trd_{kj}] \\ nrd_{nj} &= \max_{m=1}^{\#Assets} [nc_{nm} * ard_{mj}] \end{aligned}$$

As discussed above, if the Mission/Business Objectives are orthogonal; i.e., no Objective depends on another, then the **ORD** matrix is an identity-matrix.

$$(11) \quad ord_{ij} = \delta_{ij} = \begin{bmatrix} 1 & | & i = j \\ 0 & | & i \neq j \end{bmatrix}$$

If the Objectives are not orthogonal, the off-diagonal elements would become non-zero. Exploring the treatment of non-orthogonal sets of Objectives is a subject that should be considered in future work. It is likely that such sets will be encountered in the future and that sets already analyzed would, on further inspection, turn out to have dependencies among the Objectives. For the current study, an orthogonal set of Objectives will be assumed.

SECTION 3: ADDING CONFIDENTIALITY TO THE MIX

This section describes the process of extending RiskMAP to consider issues of confidentiality in the course of performing a risk assessment. The reasons for doing this are presented, followed by a brief survey of what other risk practitioners are doing with respect to treating confidentiality. Finally, the extension of the RiskMAP methodology is described.

NOTE: This section does not build on the results of the preceding section. The work described in the preceding section took place in parallel with the work of developing an extension for handling confidentiality. For that reason, the enhancements described in the preceding section are not incorporated into the confidentiality discussion. Once the discussion of adding confidentiality to the RiskMAP methodology is complete, it will be clear that the preceding section's enhancements could be added.

3.1 Motivation

Up until now, RiskMAP has addressed issues of integrity and availability, as these have been the issues identified by PCS owner-operators as paramount during our field uses of RiskMAP on this project. However, future use of RiskMAP could be in sectors where confidentiality is of equal or greater importance, such as in the medical, law enforcement, defense or financial sectors. In fact, such has already become the case since RiskMAP is already being used on projects funded by DoD sponsors. But even within the commercial sector, the RiskMAP development team believes that PCS owner-operators will eventually realize the need to protect the intellectual property represented in their PCS equipment settings, lab results, and other operational data. Witness the fact that while confidentiality issues have not been explicitly addressed in past RiskMAP studies with PCS owner-operators, confidentiality was such a large issue for several owner-operators that they did not even want the existence of the collaboration to be disclosed due to the associated security risks. That certainly demonstrates a sensitivity to some degree at the level of individual companies, but there is also a broader recognition of confidentiality's importance that has been around for some time.

Certainly, when the Federal Information Security Management Act (FISMA) was passed by Congress in 2002, it addressed the issues of confidentiality, integrity and availability (4). Directed under FISMA to produce standards for information security, the National Institute of Standards and Technology (NIST) published a series of Federal Information Processing Standards Publications (FIPS). FIPS 199 (5) defines the security Objectives of confidentiality, integrity and availability and establishes a method for categorizing information (and information systems) in terms of potential mission impact due to a breach of any of the three Objectives. While FIPS 199 is directive upon federal agencies, it recommends that "private sector organizations comprising the critical infrastructure" also follow its guidelines.

Some years ago, ISO 17799 (6) defined information security to include confidentiality along with integrity and availability, and stressed that the confidentiality of operational information could be threatened by developers, testers, and other third-party service providers. Today, a number of PCS owner-operators, e.g. those in sectors such as water or energy, deal with not just the security of operational information but also issues of protecting customer data.

A global security survey completed in 2008 by Ernst & Young (7) found that following, encouraging trends to be true:

- "International information security standards are gaining greater acceptance and adoption"
- "Despite economic pressures, organizations continue to invest in security"
- "Protecting reputation and brand has become a significant driver for information security"

However, the same study also found the following, unsettling trends to be true:

- “Privacy is now a priority, but actions are falling short”
- “Growing third-party risks are not being addressed”
- “Many organizations still struggle to achieve a strategic view of information security”

Clearly, companies recognize the need to address security in large, but there are still areas of shortfall. Potential attackers have a variety of paths available for them to gain access to sensitive data, be it operational, financial, personal etc. To name a few, attackers can employ cyber techniques from afar; they can employ human engineering practices remotely or in person; and they can gain direct access as a third-party service provider.

Examples exist to demonstrate the possibilities of an attacker succeeding in any one of these ploys. In 2000, the remote-controlled sewage equipment run by Hunter Watertech for the Maroochy Shire Council in Queensland, Australia, experienced a series of faults over several months (8). The immediate impact was 800,000 gallons of raw sewage being released, killing marine life and seriously degrading the lives and livelihood of the people in the area. Investigators found that a former third-party contractor had used a copy of software used by the plant, together with compatible wireless equipment, to control system Assets over 40 separate times. He has since been convicted.

NIST has studied the Maroochy incident and determined that had a number of the security controls recommended in NIST Special Publication 800-53 (9) been implemented, the event could have been prevented or at least ameliorated by reducing accessibility of the attacker to operational systems, software, and information.

3.2 What Are Others Doing?

A number of risk management methods and tools have been surveyed as part of this project for comparing their capabilities and approach to those of RiskMAP. Using that same list of tools and methods, a brief survey suggests that of the roughly two dozen surveyed, about half address confidentiality either explicitly or as an implicit part of a broader analysis. These findings are the result of a limited review of publicly-available materials on each method and/or tool. In some cases, a definite indication was not found but an inference could be made that the treatment of confidentiality by the tool or method was either (a) not precluded, or (b) not likely. Obviously, these survey results are not exhaustive but simply intended to provide a “temperature check” of the risk management community’s activity with respect to treating confidentiality. The survey is summarized in Table 2 below.

3.3 Extending RiskMAP to Handle Confidentiality

Development of a RiskMAP model involves the user evaluating the criticality of each Task, Information Asset, and Network Node, in turn. The evaluation is based on a set of scales as described in (1) and when evaluating Information Assets and Network Nodes, involves considerations of integrity and availability. These have been treated together, and represented by a scalar value for each Asset or Node. Representing an Asset or Node’s criticality with respect to integrity and availability by a single number assumes that the integrity and availability issues will have roughly the same import. Making this assumption helps keep the methodology simple for the user. However, the assumption breaks down when trying to represent the impact of confidentiality, integrity and availability (C-I-A) issues by a single score. For example, consider the following Information Asset:

- Details of a new product launch, prior to the event’s occurrence.
 - Need for Confidentiality = high
 - Need for Integrity = Medium
 - Need for Availability = Medium

Table 2 - Survey of Risk Management Methods and Tools

Method/Tool	Source	General Characterization	Comments	Treat Confidentiality?
API-NPRA	Security Vulnerability Assessment (SVA) methodology developed by API-NPRA. www.api.org or www.npra.org	Qualitative five-step process addressing physical and cyber security.	Assigns Asset Criticality based on gross impact of loss.	Yes - explicitly
COBIT	COBIT developed by the IT Governance Institute (ITGI) and available from ISACA www.itgi.org or www.isaca.org	Qualitative tool -- Relates business goals to IT goals and helps manage performance of IT processes	Uses a 0-to-5 maturity scale to assess process attributes	Yes - explicitly
CS2SAT	Control System Cyber Security Self-Assessment Tool (CS2SAT) developed by Idaho National Laboratory (INL) for DHS csrc.inl.gov/Self-Assessment_Tool.html	Quantitative tool -- uses a questionnaire approach and evaluates compliance to applicable standards weighted by the significance of the overall control system compromise	Presents vulnerabilities in terms of non-compliance to standards. Worst-case consequence of a control system compromise is used to determine the Security Assurance Level.	Yes - explicitly
IA CAT	Info Assurance Compliance Assessment Tool (IA CAT) developed by MITRE team led by Daryl Hild. Publicly released presentation by Cathy McCollum (case #07-0560) available at www.mitre.org/news/events/tech07/8.html	Qualitative DIACAP support tool.	Addresses development risks.	Yes - explicitly
MORDA	Mission-Oriented Risk and Design Analysis (MORDA) developed by Innovative Decisions, Inc., Vienna, VA. www.innovatedecisions.com	Quantitative method for designing and implementing secure networks.	Addresses secure network design. Uses "Attack Attractiveness" in place of P(attack).	Yes - explicitly
OCTAVE	Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) developed by CMU. www.cert.org/octave	Qualitative method -- a self-directed, risk-based strategic assessment and planning technique for security.	Focuses on what people <i>know</i> or <i>suspect</i> to be the key assets or their top concerns.	Yes - explicitly
RAPSA	Risk Analysis and Probabilistic Survivability Assessment (RAPSA) developed by the University of Idaho for NIST. www.csd.uidaho.edu/papers/Taylor02a.pdf	Quantitative tool -- combines Survivability System Analysis with Probability Risk Assessment methods.	Method adds quantitative information to the process-oriented Survivability System Analysis method.	Yes - explicitly
AMP	Assessment Management Platform (AMP) developed by SPI Dynamics, Atlanta, GA and subsequently acquired by HP. www.spidynamics.com	Quantitative tool -- a management platform for measuring Web application security risk.	Software defect analysis & security testing.	Implicit part of broader analysis
COBRA	COBRA developed by C&A Systems Security, Ltd. www.riskworld.net	Qualitative tool -- questionnaire-based. ISO 17799 compliance checker.	Identifies system threats, vulnerabilities & exposures.	Implicit part of broader analysis
CORAS	CORAS developed by the Institute of Computer Science. coras.sourceforge.net	Qualitative tool -- combines UML techniques with formal risk assessment methods.	Combines UML techniques with formal risk assessment methods.	Implicit part of broader analysis
CRAMM	UK Government's Risk Analysis and Management Method (CRAMM) developed by Siemens, UK. www.cramm.com	Qualitative comprehensive risk management toolset.	Investigates against ISO/BS standards.	Implicit part of broader analysis
Enterprise Risk Register	Enterprise Risk Register developed by Incom, Roseville, Australia. www.incom.com.au	Qualitative comprehensive risk management toolset.	Uses 5x5 risk matrix for each department, etc.	Implicit part of broader analysis
RiskWatch	RiskWatch developed by RiskWatch Headquarters, Annapolis MD. www.riskwatch.com	Quantitative tool -- deals with broad classes of objects (applications, financial data, system, etc).	Helps meet FFIEC, NERC, GLBA, BSA, NCUA and ISO 17799 & 27001 risk assessment requirements.	Implicit part of broader analysis

Method/Tool	Source	General Characterization	Comments	Treat Confidentiality?
Risk Exposure Analyzer	Risk Exposure Analyzer developed by Skybox Security Inc. www.skyboxsecurity.com	Quantitative tool -- prioritizes vulnerabilities based on vulnerable hosts' exposure. Business impact aggregated from selected host risks.	Focus areas: Risk Lifecycle Management and Network Security Compliance.	Implicit part of broader analysis
@Risk	@Risk developed by Palisade Corporation, Ithaca, NY. www.palisade.com	Quantitative Monte Carlo tool for Business Risks.	Calculates likelihood of outcomes.	Does not seem to be precluded
MAAP	Mission Assurance Analysis Protocol (MAAP) developed by CMU. www.sei.cmu.edu/pub/documents/05_reports/pdf/05tn032.pdf	A protocol (not a method or tool) for assessing operational risks in distributed processes.	Defines degree of Mission Assurance. Recognizes that missions are federated.	Does not seem to be precluded
RAMCAP	Risk Analysis And Management For Critical Asset Protection (RAMCAP) developed by ASME Innovative Technologies Institute for DHS. www.asme-iti.org/RAMCAP	Qualitative method -- a common framework for evaluating & comparing risks. Starts with asset characterization.	Framework for analyzing and managing risks associated with attacks against critical infrastructure.	Does not seem to be precluded
RAM-D (et al)	Risk Assessment Methodologies (RAMs) developed by Sandia National Labs. www.sandia.gov/ram	Family of quantitative tools based on Risk = (Likelihood of Occurrence * Consequence * System Ineffectiveness)	Threat assessment, Consequence Assessment, Vulnerability Assessment.	Does not seem to be precluded
SCAP	Security Content Automation Protocol (SCAP) developed by NIST. nvd.nist.gov/scap.cfm	Quantitative tool -- FISMA and DoD automated policy compliance checker.	Maps high level policy to low level technical security controls and checks compliance.	Does not seem to be precluded
CARVER	Criticality, Accessibility, Recuperability, Vulnerability, Effect, Recognizability (CARVER) developed by the Food Safety and Inspection Service of the US Department of Agriculture.. www.fsis.usda.gov/PDF/CARVER.pdf	Qualitative offensive target prioritization tool. Developed for food industry.	Requires user to think like an attacker.	Does not seem likely
RiskNav	RiskNav developed by MITRE. www.mitre.org/work/sepo/toolkits/risk/ToolsTechniques/RiskNav.html	Qualitative tool -- used to identify, prioritize, and manage project risks.	Provides the means to view the consequence, probability, and status of managing each programmatic risk.	Does not seem likely
RiskOptimizer	RiskOptimizer developed by Palisade Corporation, Ithaca, NY. www.palisade.com/riskoptimizer	Quantitative tool -- combines Monte Carlo simulation and genetic algorithms. Addresses economic or business operations risks.	Replaces uncertain values with risk functions that represent a range of possible values.	Does not seem likely
SEMS	Security and Emergency Management System (SEMS) developed by SEMS Technologies, LLC. www.semstechnologies.com	Qualitative tool -- "total compliance resource for drinking water and waste water utilities".	User identifies assets and assigns priorities.	Does not seem likely
VSAT	Vulnerability Self Assessment Tool (VSAT™) released by the Association of Metropolitan Sewerage Agencies (AMSA). Available at www.VSATusers.net	Qualitative tool -- but uses "Risk Reduction Units" as part of cost-benefit analysis.	Assets paired with threats to determine criticality.	Does not seem likely

To handle the differing scores with respect to C-I-A, it is necessary to track each separately. This replaces the scalar score for any given Asset or Node with a three-element vector. This approach is illustrated below.

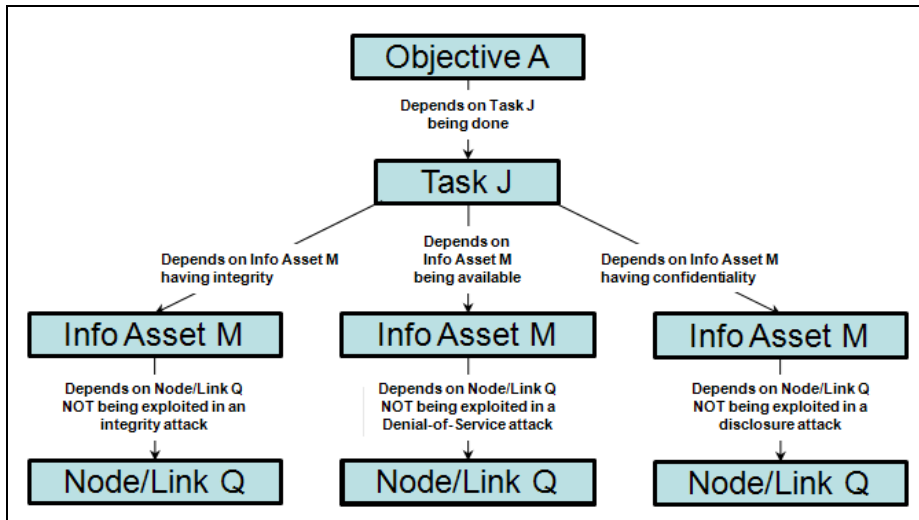


Figure 26 - Extended RiskMAP Structure

Note that in treating confidentiality we must not only consider Network Nodes but also Communications Links that are associated with the Information Asset, since data exfiltration can occur in a Communications Link as well as at a Network Node. The extended structure above implies an extended implementation, shown notionally in Figure 27 where Matrices 3, 4, 5, 4R and 3R have three pages for tracking C-I-A scores.

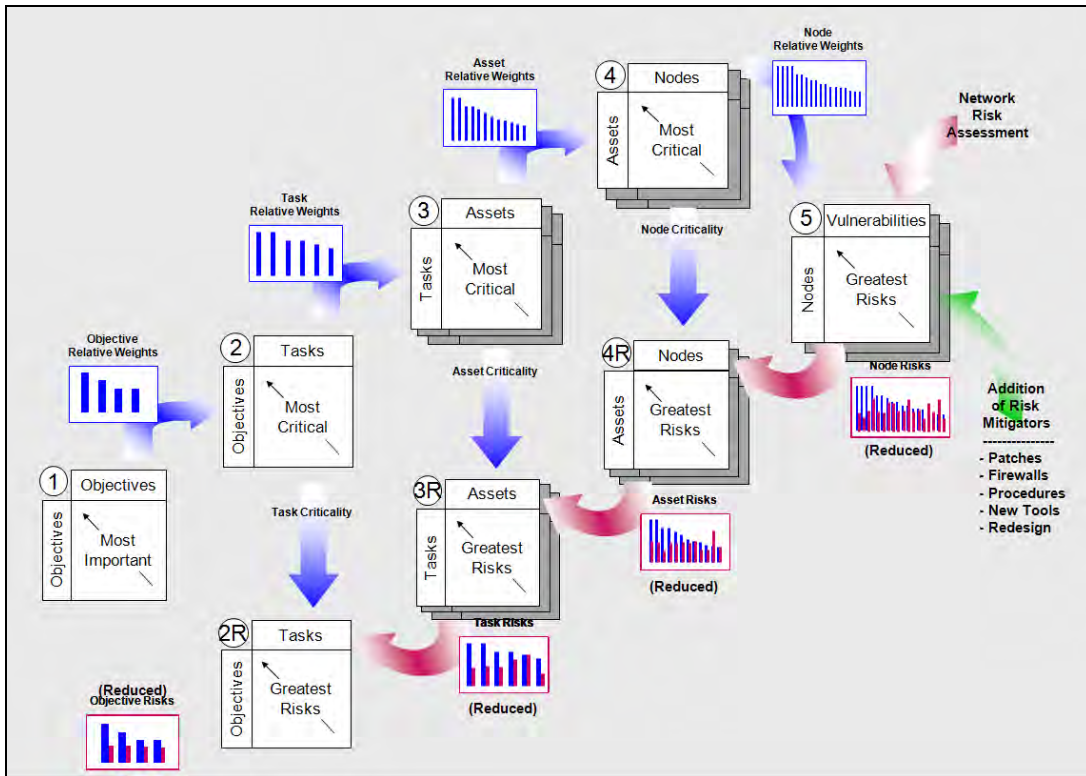


Figure 27 - Extended RiskMAP Implementation

Just as it was necessary to provide scoring tables to guide the user in selecting consistent scores for availability and integrity issues, it is necessary to provide such guidance for confidentiality issues. Table 3 below shows guidance for a user scoring an Information Asset’s criticality to a given Task with respect to confidentiality.

Table 3 - Criticality of Information Asset to Operational Task

Description Based on Impact of Asset Disclosure	Typical or Default Point Score	Score Range
<p>No impact on Task. Intermediate situations include:</p> <ul style="list-style-type: none"> ▪ The disclosure is inconsequential enough that it does not affect the performance of the given Task. Suggested point score: 7 	0	0 – 9
<p>The Task can be performed using an established work-around; i.e., there is no impact on the outcome of the Task if an acceptable work-around can be used to counter the Asset’s disclosure. Intermediate situations include:</p> <ul style="list-style-type: none"> ▪ The work-around uses redundant or backup resources already in place. Suggested point score: 10 ▪ The work-around diverts resources that are or will soon be needed for other purposes. Suggested point score: 40 ▪ In the extreme, the work-around is so costly to the operation as to be nearly untenable. 	30	10 – 49
<p>The performance of the Task is degraded. The degradation may be in terms of timeliness, quality or both. This covers a wide range of situations, for which different point scores may be appropriate. The assignment of the point score should be accompanied by a comment. Intermediate situations include:</p> <ul style="list-style-type: none"> ▪ The outcome of the Task is minimally but noticeably⁴ degraded. Suggested point score: 60 ▪ The outcome of the Task is seriously⁵ degraded, nearly equivalent to the Task’s not having been performed. Suggested point score: 85 ▪ In the extreme, the impact nears that of Task failure. 	70	50 - 89
<p>The Task becomes impossible due to the disclosure of the Asset, because no available work-around can counter the disclosure of the information.</p> <ul style="list-style-type: none"> • In the extreme, the Task cannot be performed later for any possible benefit due to lost opportunity. Suggested point score: 100 	95	90 - 100

Applying this guidance to the “new product launch” example, if the Task to be supported is “Conduct surprise launch of new product” and the Information Asset is “Details of new product launch,” the criticality scores with respect to C-I-A might look like this:

- Criticality (C) – 100 because premature disclosure would cause lost opportunity
- Criticality (I) – 70 because inaccurate information would lead to rework
- Criticality (A) – 70 because non-availability would slow time-critical preparations

⁴ For example, Task performance at the required level of quality is slightly delayed, and/or the quality of Task performance is minimally but noticeably diminished.

⁵ For example, Task performance at the required level of quality is significantly delayed, or the quality of Task performance is seriously diminished.

It should be noted that knowledge of the Asset's harmful disclosure is not required for the disclosure to have an impact on a supported Task. The scores in Table 3 are applicable whether the disclosure is discovered or not. Mitigations applied as a result of the discovery would be scored subsequently.

In a similar fashion, Table 4 below contains guidance for a user scoring Network Node's (or Link's) criticality to a given Information Asset with respect to confidentiality issues.

Table 4 - Criticality of Network Node (or Link) to Information Asset

Description Based on Impact of Node (or Link) Anomaly	Typical or Default Point Score	Score Range
No impact on the Information Asset. Intermediate situations include: <ul style="list-style-type: none"> ▪ The Node or Link displays abnormal behavior that does not affect the confidentiality of the given Asset. Suggested point score: 7 	0	0 – 9
The information's confidentiality can be maintained at the required degree using an established work-around. ⁶ This covers a range of situations, with varying costs or operational impacts of the work-around. Intermediate situations include: <ul style="list-style-type: none"> ▪ The work-around uses redundant or backup resources already in place. Suggested point score: 10 ▪ The work-around diverts resources that are or will soon be needed for other purposes. Suggested point score: 40 ▪ In the extreme, the work-around is so costly to the operation as to be nearly untenable. 	30	10 – 49
The Asset's confidentiality is questionable due to Node/Link anomaly. This covers a wide range of situations, for which different point scores may be appropriate. The assignment of the point score should be accompanied by a comment. Intermediate situations include: <ul style="list-style-type: none"> ▪ The Asset is somewhat likely to be disclosed. Suggested point score: 60 ▪ The Asset is highly likely to be disclosed. Suggested point score: 85 ▪ In the extreme, the impact nears that of certain disclosure. 	70	50 – 89
The Asset's confidentiality is lost due to the Node/Link anomaly, because no available work-around can prevent the Asset's disclosure. <ul style="list-style-type: none"> • In the extreme, the Asset's disclosure causes irreparable harm. Suggested point score: 100 	95	90 - 99

The scoring of each Information Asset and each Node with respect to C-I-A issues can be a daunting effort. In the current Excel implementation, it would be necessary to either (a) complete the scoring for C, I or A and then repeat the whole process for each of the other two; or (b) address C, I and A while scoring each Asset and Node before going on to the next one. This leads to juggling back and forth among spread sheets and can also lead to confusion and input error. For this reason, along with the need to be able to find the worst-case conditions when assessing risk for C-I-A issues, an implementation was chosen that combined MS Excel and MS Access in lieu of the current, Excel-only implementation.

⁶ For example, the information Asset is encrypted for storage or transmission.

3.4 Implementation

This implementation effort built on the results of other work comprising both direct and internally-funded research. That previous work, described in (10), resulted in a prototype that combined a RiskMAP Excel file, a MS Access data base, and the means to import and export data pertaining to the RiskMAP assessment. The capability only pertained to the scalar version of the RiskMAP tool, but it provides the basis for the current implementation. Portions of that prior work are described here, when needed to aid in the reader's understanding. To distinguish between old and new work, the old work product will be discussed as a "single mode" capability while the new work product will be discussed as a "triple mode" capability (for C-I-A assessments).

A new architecture was developed to add a relational database for persistence of RiskMAP's various data elements and the relationships between them. For each RiskMAP Excel workbook, the architecture provides for an associated MS Access database. The RiskMAP GUI remains in Excel and the relational database provides a well-defined data source for data exchanges between RiskMAP and external data producers/consumers.

To describe the architecture, several terms will be used and are defined as follows.

- RiskMAP model – The assembly of elements representing an organization's mission, functions, and dependencies.
- Microsoft (MS) Excel workbook – A MS Excel file, comprising one or more worksheets. An Excel workbook containing a RiskMAP model typically comprises about two dozen worksheets inter-related by formulas and reference calls. For brevity, such a file will be called a RiskMAP workbook.
- MS Access database (DB) – A MS Access file comprising tables and relationships. An Access DB containing a RiskMAP model will comprise tables for each of the elements in the model, along with relationships that mirror those in the companion RiskMAP workbook. For brevity, such a file will be called a companion DB.
- RiskMAP project – The assembly of an Excel workbook and Access DB, each containing the same RiskMAP model, and each linked together.

The software implementation created an Excel COM Add-In to extend the original Excel-based RiskMAP tool. Software development was done using .Net Framework Version 3.5 and C# to create an object-oriented software design that supports diverse data sources/targets including XML, MS Excel and MS Access. The COM Excel Add-in handles the interaction with Excel and maintains synchronicity between a RiskMAP workbook and its companion DB. It also adds features to the Excel GUI that appear under a new "RiskMAP Triple" tab in the Excel Ribbon. A separate .Net class library DLL handles interaction with the MS Access database. With the RiskMAP Add-In active, the user can:

- Enter and persist risk assessments for three security issues (C-I-A)
- View "worst case" assessments from among the three issues
- Use added RiskMAP Excel ribbon buttons to switch between assessment modes

The development of the data base design was founded on the single-mode design. The data structure for that single-mode prototype and the relationships between the data base tables and the RiskMAP Excel structure are shown in Figure 28 and Figure 29 below.

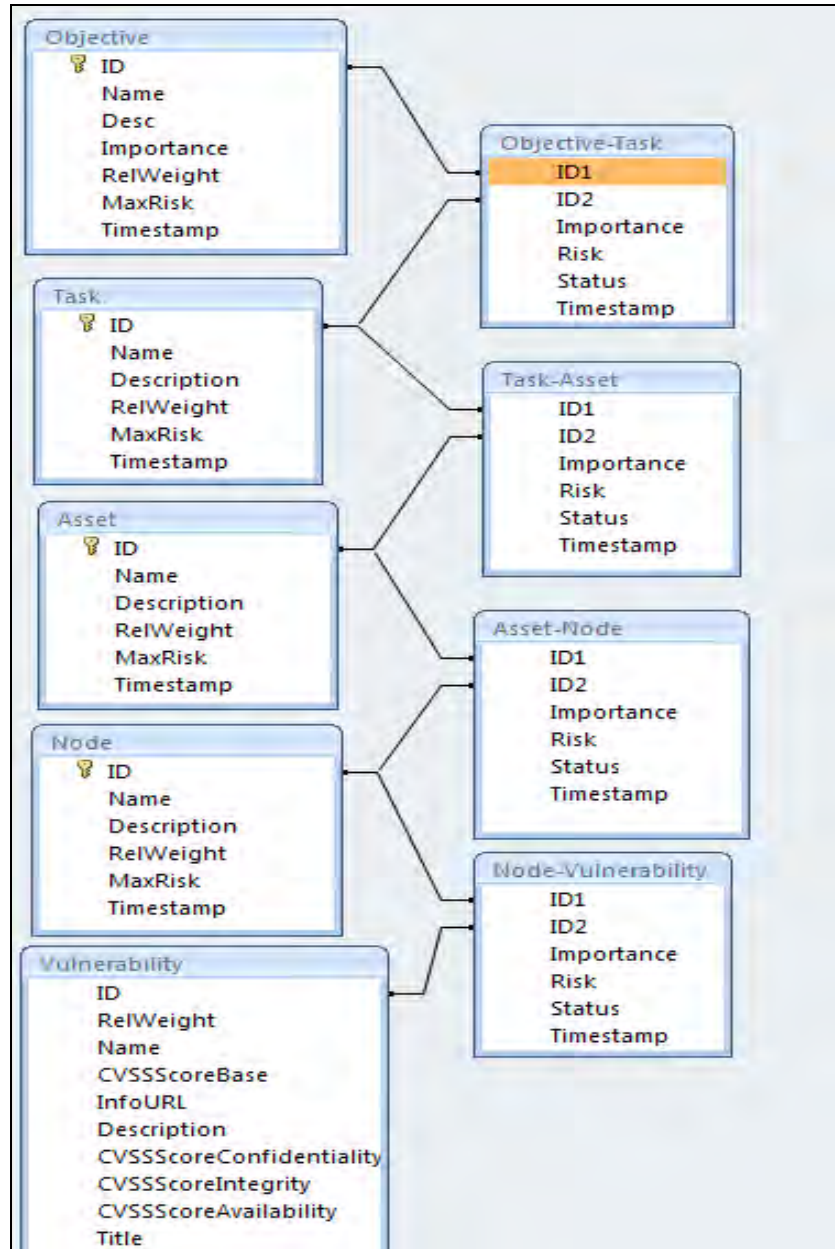


Figure 28 – Single-mode data base structure

One table is provided for each type of RiskMAP element – Objective, Task, Asset, Node and Vulnerability; additional tables capture the user-entered criticality scores as well as risk values as entered by the user and as rolled up to Asset, Task and Mission Objective. The figure below shows an example of the mapping between some of the RiskMAP worksheets and the corresponding tables in the database.




Figure 29 – Single-mode data relationships

The relationships between the data base tables and the RiskMAP Excel structure for the triple-mode design are similar to those shown in Figure 29 but they would be replicated for C-I-A scoring. As seen in Figure 30 below, the database formats are slightly different for the single-mode and triple-mode versions, but the single-mode database was easily upgraded to the triple-mode format.

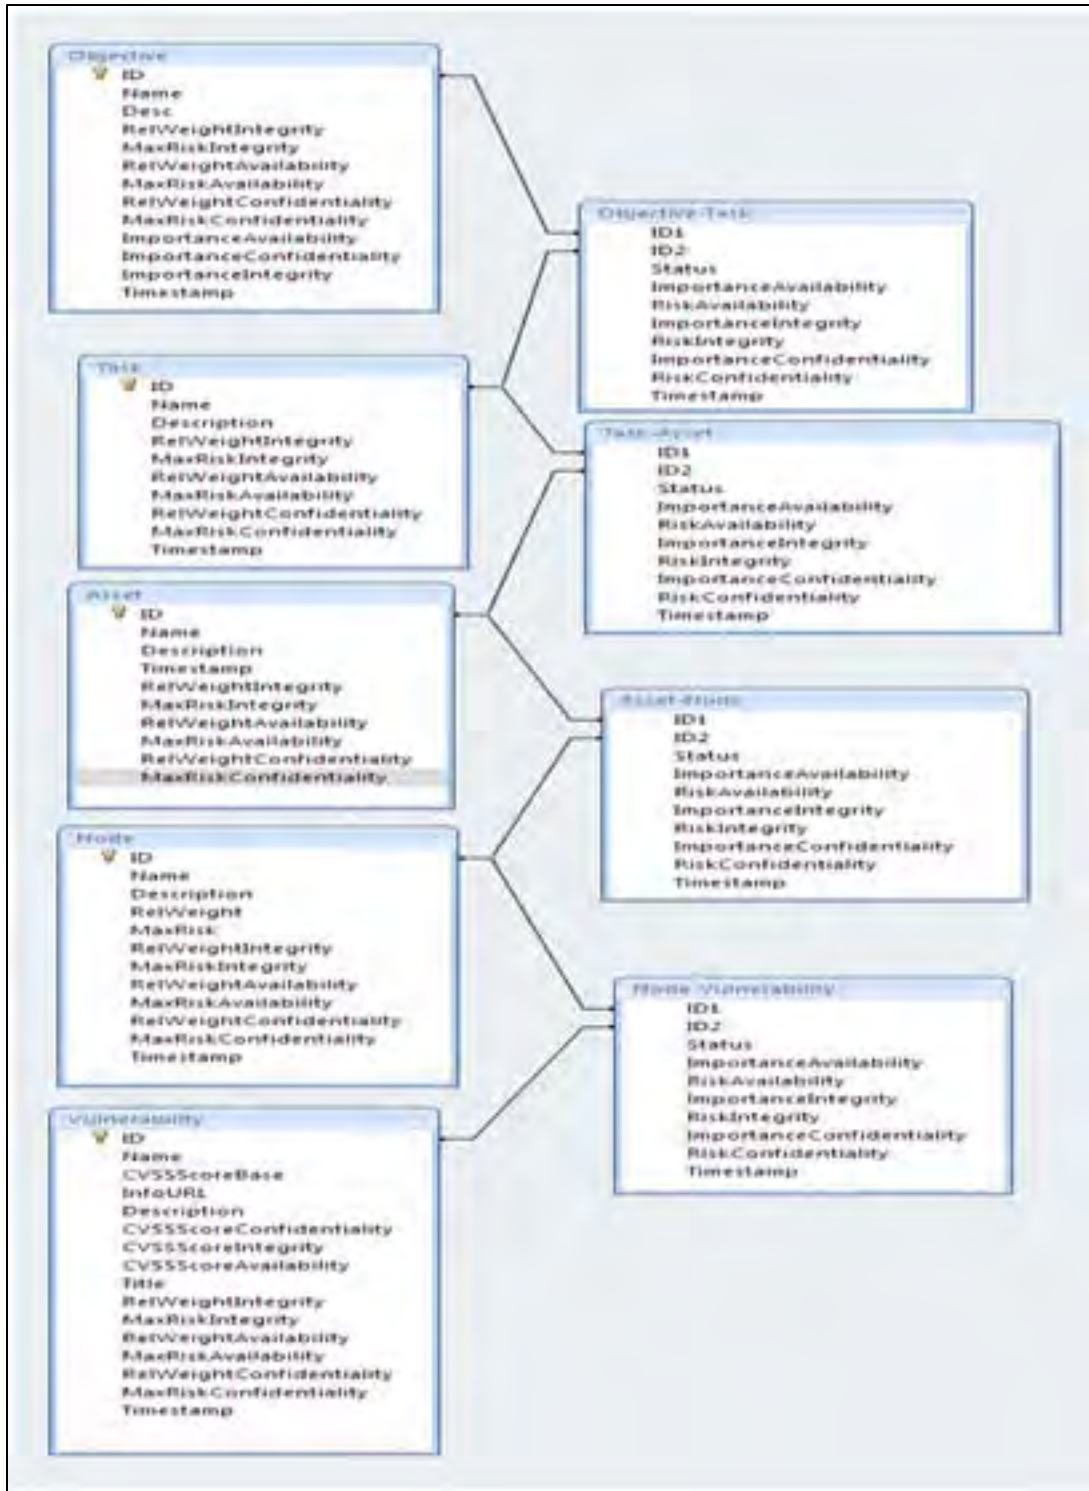


Figure 30 – Triple-mode data base structure

The RiskMAP Add-In adds a new tab to the Excel menu bar. It appears on the far right of the menu bar as seen in Figure 31 below. Selecting the “RiskMAP Triple” tab displays a RiskMAP-specific button bar that provides a user interface for manual database linking as well as other operations described below.

- **Database Link**

- **Load:** updates all worksheets in the open RiskMAP workbook with data read from the companion DB
- **Save:** updates the companion DB with data read from all worksheets in the open RiskMAP workbook

- **Assessment Mode**

Buttons are provided for each of the three assessment modes (C-I-A). Selecting one of these buttons updates all worksheets in the open workbook to display data for the selected assessment mode. All three sets of assessment data are stored in the RiskMAP model database, but only a single set of data is held and displayed in the Excel workbook at one time. This means that both entry and viewing of a RiskMAP workbook with triple assessment data can only be performed with the companion database and the triple assessment add-in active. A fourth ribbon button, 'Max', displays a read-only assessment view showing the maximum of the three importance values for each matrix cell.

In Figure 31 below, the Add-In is in Availability mode. That means the data being presented in the RiskMAP workbook pertains to Availability issues. As illustrated by the red box in the figure, the user is prompted with the active mode. Also, the Availability button in the ribbon contains an asterisk above the word Availability.

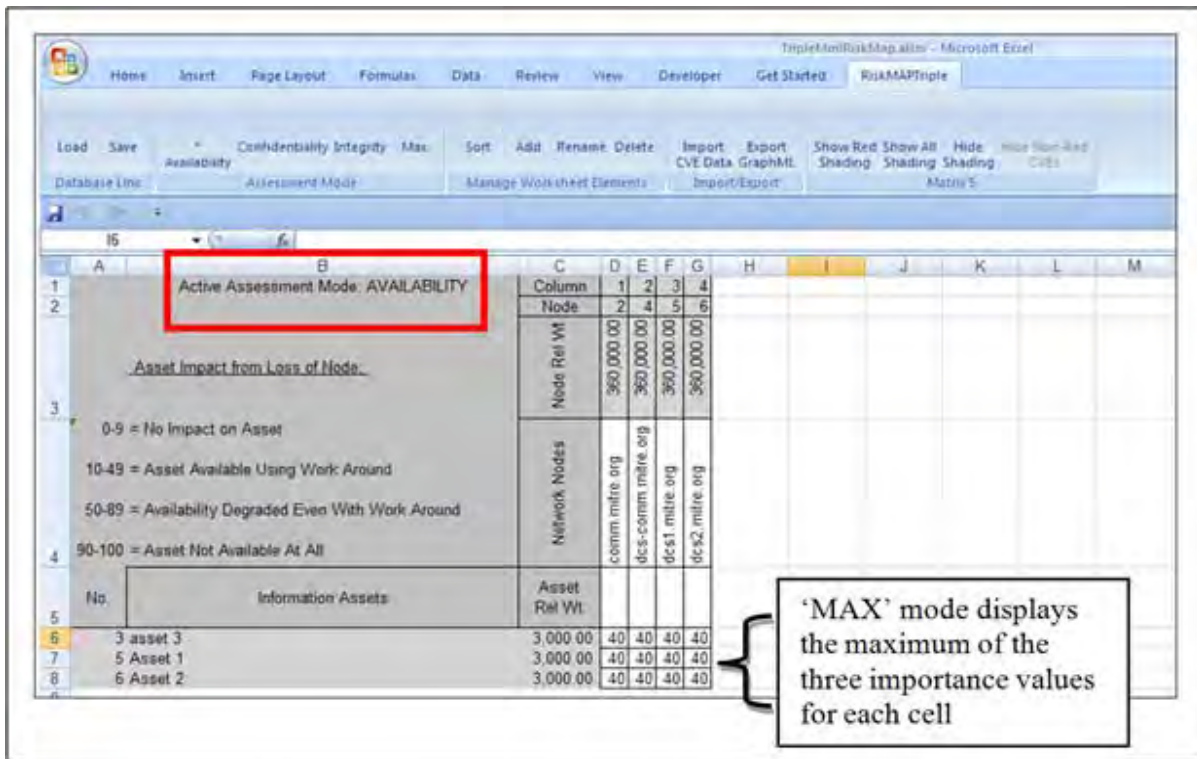


Figure 31 – RiskMAP Triple Add-In

- **Manage Workbook Elements**

These buttons help the user manage the RiskMAP model elements that appear in the columns of the currently displayed worksheet.

- **Sort:** reorders the columns in the currently displayed RiskMAP worksheet. A dialog box like that shown below is displayed to allow the user to select to sort on ID, Name or Relative Weight. The sort operation updates the element order for all worksheets where the particular column element type is displayed. For example, Figure 32 below shows the Matrix 4 worksheet with Nodes as the column elements. A sort operation on Matrix 4 will also reorder the Node elements in Matrix 5 and Matrix 4R. The dialog is context sensitive; i.e., it matches the context of the worksheet in view (Tasks, Information Assets, or Nodes).




Figure 32 - Sorting dialog

- **Add:** creates a new RiskMAP model element – of the type displayed in the columns of the currently displayed worksheet – and adds it as the right-most column. The user is prompted to enter a name for the new element. The element is added to all other worksheets where that particular type of element appears.
- **Rename:** renames a model element on the currently displayed worksheet. Before pressing the ‘Rename’ button, the user must select a single cell within a column element – either Name, Relative Weight or ID – to specify which element is to be renamed. The element is renamed on all worksheets where the element appears.
- **Delete:** deletes a model element on the currently display worksheet. Before pressing the ‘Delete’ button, the user must select a single cell within a column element – either Name, Relative Weight or ID – to specify which element is to be deleted. The Add-In reports the name of the selected element and prompts the user to confirm the desired deletion. The element is then deleted on all worksheets where the element appears.
- **Import/Export**
 - **Import CVE:** This operation carries out the import of Common Vulnerability Enumeration (CVE) data as an aid to assessing Node risk. This data is imported to the RiskMAP companion DB and is automatically presented to the user in Matrix 5 of the current RiskMAP workbook.
 - **Export GraphML:** This operation exports in the currently open RiskMAP workbook in an XML file. The file is stored in the project folder containing the RiskMAP workbook and companion DB.

- **Matrix 5**
 - **Show Red Shading:** When Matrix 5 is displayed, this button suppresses all but those columns containing red-shaded cells (i.e. cases where a vulnerability is present).
 - **Show All Shading:** When Matrix 5 is displayed, this button shows all columns containing green-shaded cells (i.e. cases where a potential vulnerability exists) or red-shaded cells (i.e. cases where a vulnerability is present).
 - **Hide Shading:** When Matrix 5 is displayed, this button suppresses all cell shading.

Development Status

As earlier stated, the confidentiality extension and the sensitivity analysis occurred in parallel. With the sensitivity analysis uncovering several fundamental questions that required addressing, much of the project effort allocated to the confidentiality extension had to be shifted to addressing these questions. As a result, the development of the triple-mode RiskMAP Add-In was curtailed after initial concept demonstration.

If it becomes desirable to include the triple-mode capabilities within the fielded RiskMAP tool (currently version 0.2.00.1231), some functionality should be reviewed and possibly revised and additional testing should be performed. In addition to the status of the separate capabilities provided below, a few general issues should be considered.

- The first is the choice of MS Access. This was initially chosen for its convenience on the desktop. With the addition of the web service export from the database another database engine, maybe MySQL, might be considered.
- A second issue is the existence of two separate Add-Ins: The single-mode version and the triple-mode version. It may be desirable create a single Add-In that meets both sets of requirements.

The single-mode capability has been tested with both small and large RiskMAP models. In its present state, this capability is usable for demonstration purposes but a few issues exist, to include the following:

- Slow performance on very large workbooks
- The add/rename/delete operations provided by the RiskMAP ribbon should be reviewed
- Some users reported problems with the ‘sort’ operation
- The Pareto graphs embedded in the RiskMAP workbook are not handled by the Add-In
- The behavior of the Add-In when a second workbook is opened should be reviewed
- The trigger events for the automatic save operations should be reviewed
- Rounding behavior should be reviewed for consistency with the web service

Since the single-mode version is the basis for the triple-mode version, these issues apply to the triple-mode version as well. For the triple-mode version, additional work is needed in the following areas:

- Sorting behavior for the “Max” mode needs to be defined
- Expansion of the CVE data imported to include details related to the three security issues.

SECTION 4: CONCLUSIONS

The sensitivity analysis described in Section 2 brought forth a number of fundamental questions about the RiskMAP methodology, presenting both an opportunity and a challenge to the RiskMAP team. In the process of answering the questions, the team developed a deeper and more refined understanding of the RiskMAP methodology. Similarly, from the exploration of ways to extend the methodology to address confidentiality as a security issue, the team gained an appreciation of the practical issues in implementing such an extension.

From the sensitivity analysis, several findings emerged.

- Regarding Matrix 1, the adaptation and implementation of AHP in Matrix 1 does not introduce errors in the generation of relative weights for Mission Objectives. Nor does it allow chances for rank-reversals. The potential for errors in user input is mitigated by the use of a visual feedback step which provides the user with a means to spot inconsistencies or biases in prioritization.
- Matrices 2 through 4, which employ QFD to capture user inputs for the network of dependencies from Mission Objectives to Network Nodes and to generate relative weights at the Task, Asset and Node levels, were found to benefit from the use of QFD as a means to manage complexity but also to suffer from the attendant loss of detail. On the positive side, the dependency network and Pareto charts, as generated using the currently-documented RiskMAP methodology, have provided information about the importance of Tasks, Assets and Nodes in the aggregate, which has proven adequate in several field trials. On the negative side, the aggregation technique employed by QFD provides only one view of relative importance (weight) – one that can be misinterpreted unless the viewer has a clear understanding of the techniques used in preparing the view.
- Rather than summing all dependency paths to arrive at the weight of a Network Node (or Information Asset or Task), one can take only the path carrying the maximum value and use this as the relative weight of the Node, Asset or Task. This MAX method will promote those Nodes, etc., that might have limited usage but are nonetheless critical to a Mission Objective. Such a view can be more applicable to certain uses (e.g., quickly finding “crown jewel” or mission-critical Nodes) than would the view created via the SUM method.
- For cases where mission dependencies overlap, as is often the case, the mission dependencies from each Objective on a given Task, Asset or Node can be kept separate for better visibility. By using vector methods to maintain separate track of mission dependencies, one can clearly trace the dependencies from top to bottom and also rank-order the Tasks, Assets and Nodes with respect to a selected Mission Objective.

From the confidentiality extension, the following emerged.

- Separately treating C-I-A issues adds a layer of complexity for the user but it is one that can be managed through a proper GUI.
- The development effort was sufficient to demonstrate a workable implementation of the concepts involved in treating C-I-A issues as part of a RiskMAP assessment. A number of thought-provoking issues were surfaced, which will help shape future efforts.

Overall, the results of the RiskMAP team’s work provide improvements that can be applied individually or together in any future RiskMAP application.

APPENDIX: REFERENCES AND BIBLIOGRAPHY

1. **Kertzner, Peter, et al.** *Process Control System Security Technical Risk Assessment Methodology & Technical Implementation*. Lebanon, NH : The Institute for Information Infrastructure Protection, 2008.
2. **Saaty, Thomas L.** Relative Measurement and Its Generalization in Decision Making. *Revista de la Real Academia de Ciencias. Statistics and Operations Research*, 2008, Vol. 102, 2.
3. **Triantaphyllou, Evangelos.** Two New Cases of Rank Reversals when the AHP and Some of its Additive Variants are Used that do not occur with the Multiplicative AHP. *Journal of Multi-Criteria Decision Analysis*. 2001, Vol. 10.
4. *Federal Information Security Management Act*. Washington, DC : U.S. Congress, 2002.
5. *Standards for Security Categorization of Federal Information and Information Systems*. Gaithersburg MD : National Institute of Standards and Technology, 2004. FIPS PUB 199.
6. *Information technology — Security techniques — Code of practice for information security management*. 2005. ISO/IEC 17799:2005.
7. **Ernst & Young.** *Global Information Security Survey*. London : EYGM, Limited, 2008. EYG No. AU0162.
8. **Abrams, Marshall, et al.** *Malicious Control System Cyber Security Attack Case Study—Maroochy Water Services, Australia*. Gaithersburg, MD : National Institute of Standards and Technology, 2008.
9. **Ross, Ron et al.** *Recommended Security Controls for Federal Information Systems*. Gaithersburg, MD : National Institute of Standards and Technology, 2007. SP 800-53.
10. **Watters, Jim et al.** *RiskMAP Data Import and Export*. Bedford, MA : The MITRE Corporation, 2009. (Limited Distribution).
11. **Ernst & Young.** *Global Information Security Survey*. London : EYGM, 2008.