Security Architecture Smart Guide:

Just Good Enough Risk Rating: A proven method for rapid risk assessment

Brook Schoenfield, Enterprise Security Architect, Autodesk, Inc
Vinay Bansal, Senior Security Architect, Cisco Systems, Inc

1. Problem

Information security and risk departments are typically charged with assessing the information security risk of systems. The result of these assessments is generally some sort of risk rating that can be used to set priorities and make decisions about what exposures to accept and which to mitigate.

However, in the absence of a body of long-term actuarial data covering the many variety of information security incidents and losses, assigning risk values has been fraught with individual interpretations and stylistic variations. Indeed, the risk avoidance appetite of the assessor can greatly influence the rating of risk. This leads to experienced practitioners delivering different, perhaps wildly different risk ratings about the same system. Further, risk ratings for information security risk typically lack repeatability and do not generate comparative values.

There has been a general lack of agreement about which variables must be must be included within an information security risk calculation. This problem is further complicated by products that offer "risk calculations" based on only 1 or 2 of the required terms[1]. For instance, "risk" might be based upon threat and exposure, or vulnerability and hypothetical impact (usually of the worst case possible). The misuse of the term "risk" does not help the practitioner sort through myriad data points to calculate a repeatable and quantified risk measurement.

A number of risk models have been put forward in order to address these problems. Some of these are very easy to use, but do not address the basic problems stated above. Other methods assume that data has been gathered about assets, values, vulnerabilities, threat, and exposure that would be encyclopedic for any but the smallest organization. Such an effort is typically beyond the capability of an enterprise information security team who are typically faced with constantly changing inventories, newly arriving vulnerabilities, incomplete knowledge of threats. Other vendors have offered risk presentation systems. Several that the authors have investigated do not, in fact, rate risk. Rather, these methodologies address some portion of what must constitute a risk rating. For example, a method may only address attack types. Products for storing "risk" often require a significant investment not only to purchase, but to gather and assess risk, as well as requiring significant data entry demands. Once an investment has been made in one of these tools, the effort to keep the assessments updated is a nontrivial and constant task.

What is needed is a reasonably robust, fairly lightweight, inexpensive way to gain repeatability and consistency in risk ratings across practitioners and assessed systems. This guide explains Cisco's Just Good Enough Risk Rating (JGERR)

---

[1] The arithmetic for these calculations is also typically nonexistent or inappropriate. The authors have seen canned ratings for each vulnerability found, simple addition, and even averages over the terms.

tool. JGERR has been in use for 8 years. It is used enterprise-wide at Cisco Systems Incorporated for web infrastructures and applications and for assessing risk for new extranet connections. It is in use by more than 40 practitioners, day in and day out. While far from perfect, JGERR combines knowledge from several leading risk rating methodologies into a methodology that can be used in the absence of a sufficient base of statistical data about information security incidences and losses. JGERR can be implemented for any area within information security. It is lightweight and easy to use.

The methodology presented in this guide does not claim to calculate true risk. It is the authors' belief that the commonly promulgated equation[2] for insurance risk is currently impossible to calculate with any surety in all but the most narrow information security circumstances. JGERR provides instead, a risk rating result that can be used in the absence of an ability to calculate risk accurately.
The rating results are intended for prioritization in order to bring risk within a desired posture. The results should also be useful by management to understand gross risk trends, as well as provide some basis for tracking the work of the assessors[3].

2. Scope

The risk rating methodology described in this guide may be applied to any discreet area for which information security risk assessment and rating is required.

To apply this methodology, it is recommended that the area to which the method is applied be well enough bounded such that the question set that is developed in the implementation is specific enough to deliver meaningful and comparable results within that arena. The questions upon which ratings are based can become so generic as to not capture meaningful results. Further, the more generic the questions get phrased, the more interpretation is required in order to answer them. Implementations must balance generality versus specificity such that results provide an opportunity for sufficient repeatability.

For instance, the example given here does not purport to apply to every area within the broad range of information security risk assessments. Rather, the

---

[2] Risk = Probability * Annualized Loss. There simply is not enough data about the dollar value of losses and how often these occur over a wide enough population in order to calculate the standard risk equation.

[3] On this last point, care must be taken to remember that not every project or system will have a lower risk rating at its end from when it started. There are many reasons for risk rating to increase during the development of a system; the increase can even be dramatic. Despite the foregoing caution, most projects moving through a mature risk practice will have either a flat risk rating from beginning to end or will have ratings lowered through the project development lifecycle as information security controls are applied.

questions were developed specifically to address risk and loss areas of web infrastructures and applications. These questions make assumptions about the security controls that have been built into the local environments on which applications are deployed. Indeed, applications that hew to these standards will generally have a lower residual (unmitigated) risk than applications that have been deployed on one-off environments where security control compromises have been made.

These factors must be taken into account in this methodology. Organizational assumptions about what is more risky and what is less become embedded into the rating form. This is a natural part of capturing practitioner and tribal knowledge into the form's questions and into the ratings.

Implementation of this method is best scoped to a particular area of assessment. Multiple discreet versions of the questionnaire may be applied to the various areas that come under assessment. Using unique, focused instances of the method for each discreet area of assessment still allows the risk results to be comparable across the areas.

2.1. Sector

[X] General (Potentially applies to all sectors)
[ ] Government
[ ] University

2.2. Security Practice Domain

[X] Architecture
[X] Engineering
[ ] Operations
[X] Risk Management

2.3. What functions does it pertain

- Risk assessment and risk rating

3. Known Risks

Any attempt to capture tribal and/or practitioner knowledge is always limited by the capture process, which is, by its very nature, insufficient. The very act of distilling wide experience into a set of questions will always miss something. Further, corner cases are hard to capture. And, decisions must be made about priority, importance, and inclusion/exclusion. Information security is such a broad field, that something will always be left out, either intentionally or unintentionally.

There is a risk when applying this method that risks will be missed. That is why there is a "thumb on the scale" area built in for an experienced risk assessor to add in additional factors not canonized within the questions that have been included in the rating. Still, giving assessors this capability also then brings with it a risk of mis-use in order to force particular result scenarios. That is, if risk is rated high, an engineer or architect might force a set of security controls that overly reflect her/his personal risk appetite. To combat this abuse of power, the authors have insisted in practice that any "thumb on the scale" additional risk receive positive peer review before being used in the resultant risk rating. The details of the peer review process are beyond the scope of this guide.

There is a risk that unconscious assumptions about what is risky and what is preferable can creep in and rig this rating system towards higher or lower ratings results. In fact, the first two versions of this methodology had exactly these problems. Getting these assumptions to be more conscious and apparent in the questions was a part of the refinement of the method.

It is crucial that the ratings from the two values being calculated be kept separate in order for ratings to approximate a risk calculation. Threat, vulnerability, and exposure can be mixed together in questions, but must not be mixed into questions for impact and loss value. These two calculate different values whose arithmetic properties must be clearly separated. That is, the questionnaire dealing with the impact analysis must not include any assumptions about what is less or more vulnerable, what the threat landscape is, nor the controls in place. Impact questions must be solely about loss, though they do not have to be about monetary values only, since the results will ultimately be rated. However, this method will work fine if assumptions about threats, exposure, and vulnerability are conflated into the "technical" analysis questionnaire and rating. The conflation, as long as it's understandable, deliberate, and captures organizational understandings, will not harm the method's results.

4. Prerequisites

It is assumed that any organization contemplating use of this solution already has in place an information security risk assessment practice.

Typically, risk assessments are applied to:

- Applications and similar new projects
- Changes to running systems
- Infrastructures to support operations
- Etc.

It is assumed that the organization has at least a few experienced risk assessors who have practiced within the area to which the method will be applied.

While this methodology does not require any investment in tools or products, it does require an investment in time and attention of those who will distill the local risk knowledge, the organizational information security policies and standards, the risk practice, and the organization's risk posture preferences into the questions to be used during assessment and rating. In our experience this is a part time task which took the originators and practitioners not more than 20-40 hours, stretched over a period of 2 months. In addition, in the two successful implementations with which the authors are familiar, there are periods of peer review, comment, and refinement stretching over a few weeks and additional hours before the methodology is ready for production use. While not trivial, this effort should be within the reach of most information security departments.

Once canonized and put into use, the questions and ratings should be reviewed at least once each year against organizational, threat, and vulnerability changes.

It is assumed that an organization already has the following information security processes:

- An information security team
- A risk assessment practice

5. Required Artifacts

a) A risk rating spreadsheet whose questions capture assessors' understandings about what is risky and what is safer for a particular area. This artifact is developed by the assessment team(s) before being used[4].
b) An architecture diagram of all the components to be used for any system under assessment
c) Data flow diagram across all components of the a system under assessment
d) Data sensitivity rating of the most sensitive data within each flow and to be handled by each component of the a system under assessment
e) Information Security Policies
f) Organizational technical standards to which systems must comply

6. Step-by-Step Problem Solution

The risk rating calculated by this methodology uses 2 terms: an estimated attack vector and a scaled size of impact. These 2 terms are multiplied in order to get the risk rating.

estimated attack vector * *scaled* size of impact = risk rating

[4] Capturing and rating local and industry practice into a series of rated questions is a non-trivial task. Doing so is likely to take several iterations. Senior risk assessment staff, along with project management will probably need to be assigned to this task.

The estimated attack vector is actually a combination of several fundamental components that are typically combined within the term representing probability. Since, as has already been stated, Information risk assessors do not have the statistical data to estimate true probability, we must calculate a number between 0 and 1 to replace true probability term. And that number must represent as many of the factors from which probability might be calculated as we can estimate in a repeatable fashion. The good news is that experienced risk assessors do this every day as mental arithmetic. The difficult part is teasing this information out from mental arithmetic. Gathered attack vector knowledge is transformed into a series of discreet questions, each of which has a multipart, scaled, series of responses: multiple-choice questions, essentially.

For Cisco's web application assessments, a five bucket scale was chosen.

Both terms used to calculate the risk rating are 1 based. There is no 0 loss and there is no 0 probability – there is always some loss value and some risk.

The scale for the attack vector term is:

Attack Vector ("technical exposure") = 0 > N < 1

And the scale for impact is:

Individual impact == 1 >= N <= 10

Impacts are summed for an impact rating of

Total impact = 10 >= N <= 100

Each of the 5 choices is evenly distributed[5] across the range, 1-10, with appropriate arithmetic to calculate the attack vector term, 0-1 (see below for details).

1. 1
2. 3.25
3. 5.50
4. 7.75
5. 10

In the following table, our definitions for each term that must be expressed in the attack vector term for the risk rating. It is the authors' assumptions that each of these situations without the presence of the other conditions is not risk. There seems to be a good deal of confusion amongst security practitioners about what

[5] There is nothing special about this distribution. The scale might be biased low or high, as required. The actual number distribution is less important than the consistency of use. In fact, Cisco's distribution was biased high for years.

risk actually means. Some methodologies equate threat to risk, in others, vulnerability is equated to risk, the assumption being that the very worst exposure and threat preexist simply because the exposure and risk are known to exist somewhere. This of course, is the safest course. But it does not get us closer to calculating risk for a particular situation. That is why we insist upon the existence of an exploit of the vulnerability ("exploit"), a threat agent that is motivated and has sufficient access to exploit the vulnerability, and that the vulnerability is exposed enough for the threat agent to exploit it. In other words, each of these factors is dependent upon the presence the other factors before an information security attack can be promulgated.

Vulnerability	Any weakness, administrative process, or act or physical exposure that makes an asset susceptible to exploit by a threat
Threat	A potential cause of an unwanted impact to a system or organization. (ISO 13335-1)
Exposure	The potential damage to or loss of an asset from a given threat and/or vulnerability after consideration of existing controls
Exploit	A means of using a vulnerability in order to cause a compromise of business activities or information security

Exclude[6] vulnerabilities that are not exposed or do not in the present or the foreseeable future have a motivated threat capable of exploiting the vulnerability. Doing so will help to keep the questionnaire short and focused on the area that is

[6] The authors do not mean to suggest that vulnerabilities that cannot be exploited should not be attended. That is a matter of organization policy and standard. We are suggesting that in order to keep the risk questionnaire lightweight, it is good practice to focus on prioritized attack vectors as defined here.

under assessment. We believe that this is critical to the success of the methodology. It is important to bear in mind that it is quite possible and even preferable to build multiple questionnaires, one for each area of technical expertise and assessment.

In this method, the business impact questionnaire is designed to be completed by a project manager working for the business that is driving a project. The questionnaire should be easily understandable by a project manager or other non-technical person. We removed as much information security jargon as we could. We ask questions that the project manager will likely know or which he/she can easily get answered by the sponsors of the project.

The attack vector, or "technical" questionnaire is designed for a trained information security assessor. This person may not have deep risk assessment skill. But the person will have a working knowledge of the sort of attacks that are currently being promulgated against systems in the domain under consideration. The assessor will have some understanding of the typical controls that are built into underlying infrastructure and the kind of controls that are typically required for applications. Further, the assessor should understand how these typical controls are applied to particular architectural patterns. Due to this pre-requisite training, we were able to generalize the technical questions considerably, as well as to shorten the questionnaire.

In order to build our questionnaires, we use the following process:

- ***Distill* assessor *expertise***
- ***Document typical* exceptions to standards**
- ***Transform* knowledge into assessment questions**
- ***Scale (bucket-ize)* responses**
- ***Gather* impact scenarios**
    - ***Repeat transformation and scale***
    - ***As for attack vector term***

6.1. Distill Assessor Expertise

Interview your current experts who assessed typical projects coming through the area of focus of the questionnaire.

- What typical vulnerabilities do they look for?
- Which vulnerabilities 2 the exclude?
- What threats are considered the most important?
- What are the typical exploits of the most important threats?
- One of the greatest weaknesses contained within local infrastructures in systems?
- What is protected particularly well?

- Which types of vulnerabilities are not exposed to active threats that are currently or projected to attack your systems?

Don't limit yourself to the questions posed in this guide. The object of this process is to discover and document the accumulated knowledge base from your experts. What are they most concerned about? What are they least concerned about because there's enough extant mitigation to discount successful compromise? And, it is important to gather understanding of what lies between these 2 extremes in relation to your systems. Those attack vectors that lie between the two poles (presented above) will become the buckets, the questionnaire choices that you will rate in order to generate the threat, vulnerability, exploit, and exposure combined term. You are trying to distill your attack vectors from worst to best-case scenarios.

6.2. Document Typical Exceptions To Standards

An additional gold mine of technical understanding are the security exceptions written for projects. Repeating exceptions indicate areas of weakness that aren't well defended. Those types of exceptions written only once generally indicate the inability or unwillingness of a particular project to meet organizational standards. These tend to be local. However, exceptions that occur repeatedly can point to areas where compliance is too difficult or expensive. As such, these point to weaknesses in the defense-in-depth capabilities[7]. These can be captured into the questionnaire.
If your infrastructures are fairly mature in that they deliver basic security controls to the hosted systems, you may choose to simply note whether exceptions exist as one of your questions. Your buckets (multiple choice responses) could then list the types of exceptions from none to most risky.

6.3. Transform Knowledge Into Assessment Questions

Each question does not have to encapsulate an entire attack vector's composite terms. One question may model active threats that have got some level of access. While another question might focus on exposed vulnerabilities at which you know there are active attack attempts.

For instance, if you have a population of unremediated cross-site script errors and you have a business driver to protect your web visitors from the effects of these attacks, then one question might focus on the presence of these vulnerabilities. Or, you might combine all input validation attacks into a question. Or, you may simply ask (as we did), "Are there known

[7] Not only should these exception types be included in the questionnaire, these also make a case for a change in security direction to close the holes, to make it easier for systems to comply. After all, an organization creates its standards for good reason, presumably.

vulnerabilities exposed in the application?" In this way, you can generalize all vulnerabilities.

At Cisco, the person answering the attack vector questions is always a security assessor who has had significant training. Hence, we were able to generalize the questions against the assumption that the responder is familiar with typical web based attack patterns. Your questionnaire may have to be more specific depending on what assumptions you can make about the responder/assessor.

Cisco's Web application Attack Vector (technical exposure) questionnaire uses the following questions:

1. How much exposure to attack is there?
2. Are there known vulnerabilities in the application/project or associated infrastructure?
3. Are mitigation or workaround ("hardening") techniques implemented to minimize the risks or vulnerabilities inherent in the infrastructure?
4. To what degree do you suspect deployment of this application/project in its current form would increase the security risk to other systems, applications, resources or projects in the event of a successful compromise?
5. How is entitlement accomplished?
6. How much security review has this application/project been through?

The authors stress that you must not mix impact with attack vector. The two terms' questions must remain distinct in order to generate a value that reasonably equates to risk. As written above, the two terms cannot be mingled or the calculation loses its ability to mimic a statistically based risk calculation. Cisco chose to use separate worksheets, filled in by different persons in order to achieve this separation of terms. None of the questions representing attack vector mention the size of the impact. And the questions describing possible impacts do not describe any of the terms that are combined for an attack vector.

6.4. Scale (Bucket-ize) Responses

In order to derive buckets for your questions, start with binary choice:

- What is the best case situation in relation to the question being asked?
- What is the worst case that can be encountered?

For the sample question, "How much exposure to attack is there?", placement within an environment completely out of control or knowledge of Cisco and its security team was deemed the worst case scenario. This condition was then combined with a situation where serious inability to apply required security controls might exist: where a formal executive signed risk assumption has occurred into the follow worst case response:

- Offsite with an unreviewed ASP or other 3rd party, or contains significant exceptions or risk assumption.

Likewise, the best-case scenario is:

- Purely internal on a InfoSec-approved architecture

Having arrived at best and worst cases, degrees of exposure are filled in to derive the intervening series of buckets that fit the scale.

1. How much exposure to attack is there?
    a. Offsite with an unreviewed ASP or other 3rd party, or contains significant exceptions or risk assumption
    b. Internet Facing on a non-hardened or unknown architecture and/or without layering or with significant exception or risk assumption
    c. Internet Facing on an InfoSec-approved standard architecture
    d. Purely internal on a non-standard architecture or has exceptions
    e. Purely internal on a InfoSec-approved architecture

The responses are then put through the following calculation to create a composite attack vector rating:

$$\frac{{}_1\sum_n (\text{attack vector choices})}{\text{Highest possible score} * \text{Number of questions}}$$

6.5. Gather Impact Scenarios

To build a list of impact assessments, repeat the attack vector process, this time focusing on the types of business impacts that successful security compromises can have. Stress business terms. These must include financial impact. But we believe that the list of impacts must also include dimensions that are difficult to quantify.
Cisco's Web application focused business impact questionnaire contains the following questions:

1. What is the value of this application to company, i.e how much money will it save or bring in to the company in a fiscal year?
2. What critical company systems would the failure of this application impact?
3. What audience would be affected because of an interruption or compromise?
4. Does this project deal with any of the following personally identifying information (i.e. are there any privacy issues)?
5. What is the expected future lifetime of the system

6. Rate the criticality on this application for the continuing operation of your business.
7. Disruption of this application would have what sort of effect on the company's customers?

Again, as in the attack vector questionnaire, decide how many buckets will be useful (Cisco uses 5 for both sets of assessment questions). Choose the best-case scenario as the minimum value and the worst impact as the maximum. Then, fill in the buckets in between as ordinals of impact.

Following is Cisco's first impact question with it associated responses. We ask the project manager or business sponsor to estimate with a best guess the appropriate value[8].

1. What do you estimate the dollar damage to company would be if the data was modified, stolen, destroyed, or subject to unauthorized disclosure?
    a. A) Catastrophic (>=$400 million)
    b. b) Serious loss ($50-400 million)
    c. c) Significant Loss ($5-50million)
    d. d) Loss ($500K to 5 million)
    e. e) Minor loss <= 500,000 USD

As Cisco has grown, so the impact bucket ordinal responses have had to be revised. It is our working theory that as long as the scale remains the same, the bucket sizes can be changed to meet changing circumstances of expansion or decrease.

Arithmetic for deriving a 5 bucket distribution between 1 and 10 is as follows:

> Array of Impact buckets = from 1 to 10, last bucket + ((10-1)/4)

1 is a minimal loss, while 10 defines a maximum loss. Responses are scaled to 5 buckets, giving 4 divisions between 1 and 10, or

> (10-1)/4 = (9/4) = 2.25

The algorithm derives the following scale:

[8] An experienced project assessor can often spot inflated or deflated impact estimates based upon experience with systems of a similar type and size. Choosing the appropriate level of impact is sometimes a dialog between assessor and business sponsor. The responses on the impact sheet should be cross-checked for appropriateness to the overall system under assessment.

1. 1
2. 3.25
3. 5.50
4. 7.75
5. 10

Every loss to a system will have impact, no matter how trivial. The least total impact is 10 (1 * 10 responses). The greatest impact (catastrophic) is 100 (10 * 10 responses).

$$\text{Total Impact} = {}_1\sum_n (\text{responses})$$

6.6. Combine Terms

Estimated Attack Vector * Scaled Size Of Impact = Risk Rating

Attack Vector is “AV”
Impact is “I”

The risk rating is then

$$(0 < AV < 1) * (10 >= I <= 100) = \text{Scaled Risk}$$

Or, precisely in the Cisco spreadsheet:

$$\frac{{}_1\sum_n (AV)}{10 * \text{Number of questions}} * {}_1\sum_n (I) = \text{Scaled Risk}$$

Where:

Attack Vector (“technical exposure”) = 0 > N < 1
Impact == 10 <= N >= 100

In the spreadsheet, the total is color coded to make it easy for the assessor to spot those assessments that will require action and those that signal projects of less information security interest.

Total Scores:	
Technical Exposure	0.55
Impact	62.20
Total Risk	34.21
Additional Security Risk Factor (0-40)	0.00
Composite Risk	34.21

Composite Risk Legend	
Low	0 - 20
Medium-low	15 - 34
medium-high	35 - 64
High	65 - 84
Severe	85+

6.7. Additional Factors Outside the Questionnaire

Obviously, no short questionnaire can cover all of a complex domain for which risk assessment is performed. JGERR has built in a place for risk factors understood by the risk assessor but that fall outside of the questions that are being asked. Additional factors may be discovered during a security review. Or, the system under review may contain components that fall outside the typical scope of analysis. We call these "risk multipliers". In order to account for corner-case situations, there are four generic, open-ended fields into which the risk assessor may place a "finger on the scale" to increase the risk rating.

Are there any factors which are specific to this project that raise the risk significantly?
Factor 1 -please attach documentation about this factor
10) Factors cause this project to have additional very high risk.
8.4) Factors cause this project to have additional high risk.
6.4) Factors cause this project to have additional medium risk.
3.4) Factors cause this project to have additional low risk.
0) There are no factors

Including a space for the unforeseen and the corner case came from literally thousands of applications of the methodology. Over time, it became apparent that an experienced assessor has a more holistic understanding of risk that any

automation. Somehow, that experience had to be incorporated into the risk rating

However, including an open-ended space with which to influence risk ratings is a tremendous responsibility that cannot be taken lightly. Therefore, only trained assessors are allowed to add additional factors. Further, the following conditions must be met before an additional risk factor can be added to the score:

- Additional risk factors must be formally documented
- Additional risk factors must be justified in writing
- Additional risk factors require peer review of at least 2 other assessors
    - One of the peer reviewers must be at the most senior level
    - One of the peer reviewers must be outside the sub-team directly involved in the business area of the project

If there is consensus that there is an additional factor that cannot be scored through the questionnaire and there is consensus about the size of the addition, then the additional factor is accepted.

Total Scores:	
Technical Exposure	0.85
Impact	66.00
Total Risk	56.32
Additional Security Risk Factor (0-40)	14.80
Composite Risk	71.12

7. Recommended Project Gate Initiation Point

At Cisco, risk assessments are performed upon Security Architecture engagement and just prior to go-live. The engagement may begin either before[9] the Execute Commit project gate or just after, when formal architecture and design work begins.

8. Other Smart Guides Referenced

None at this time.

[9] On particularly complex, critical and high-risk projects, earlier security engagement is preferred in order to identify security issues as early as possible. But such early engagement is typically not required for more standard projects.

9. Industry Standards Referenced

- ISO 13335-1

10. Public Domain Tools Referenced

- None at this time.

11. Taxonomy/Definitions

Vulnerability	Any weakness, administrative process, or access or physical exposure that makes an asset susceptible to exploit by a threat
Threat	A potential cause of an unwanted impact to a system or organization. (ISO 13335-1)
Exposure	The potential damage to or loss of an asset from a given threat and/or vulnerability after consideration of existing controls
Exploit	A means of using a vulnerability in order to cause a compromise of business activities or information security
Attack vector	A credible threat exercising an exploit on an exposed vulnerability

Impact	The overall (worst case scenario) loss expected when a threat exploits a vulnerability against an asset

12. Other Materials Referenced

- CVSS: Vulnerability, exposure, some impact: technically focused
- Microsoft Threat Modeling: Exploit, exposure, some threat modeling
- Bruce Schnier's attack tree methodology
- FAIR: Contributes significantly to Just Good Enough Risk Rating.
- The Security-specific Eight Stage Risk Assessment Methodology
- LAVA

13. Publication Date, Version, Authors, and URL where to find document

Date	Version	Author
2011/09/11	0.1d	Brook S.E. Schoenfield Vinay Bansal
2012/02/07	1.1b	Brook S. E. Schoenfield Vinay Bansal Michele Guel

1. **Other Contributors to the Methodology**
   - **Rakesh Bharania & Catherine Blackadder Nelson, concept originators**
   - **Vinay Bansal & the Cisco "Web Arch" team**
   - **Jack Jones & FAIR**
   - **John and Ann-Marie Borrelli & the KnowledgeConnect Security Sharing Forum participants**
   - **An unnamed trust researcher at RSA 2002**
   - **The Information Security Risk Methodology working group at Cisco Systems, Inc:**
     - **Doug Dexter**
     - **Marc Passey**
     - **Richard Puckett**
     - **Jim Borne**
     - **Brook Schoenfield**