



TRUST
Over **IP**
FOUNDATION

Risk Assessment Companion Guide

Version 1.0
24 August 2021

This publicly available guide was approved by the ToIP Foundation Steering Committee on August 24, 2021.

The mission of the [Trust over IP \(ToIP\) Foundation](#) is to define a complete architecture for Internet-scale digital trust that combines cryptographic assurance at the machine layer with human accountability at the business, legal, and social layers. Founded in May 2020 as a non-profit hosted by the Linux Foundation, the ToIP Foundation has over 300 organizational and 100 individual members from around the world.

Please see the end page for licensing information and how to get involved with the Trust Over IP Foundation.

Table of Contents

Table of Contents	2
Document Information	3
Author	3
Contributors.....	3
Revision History.....	3
Terms of Use	3
RFC 2119	4
Executive Summary.....	6
Introduction	7
Purpose	8
1. Risk Assessment Overview	9
1.1 ISO 27005.....	9
1.2 NIST 800-30	10
1.3 Risk Concepts and Terminology.....	10
1.4 Risk Management Hierarchy.....	13
2. ToIP Risk Assessment Process.....	15
2.1 Phase 1 - Establish Context.....	15
2.2 Phase 2 - Identify Risks	15
2.3 Phase 3 - Analyze Risk.....	16
2.3.1 Step 3.1 - Analyze and Determine Risk Likelihood.....	16
2.3.2 Step 3.2 - Analyze and Determine Risk Severity	17
2.3.3 Step 3.3 - Calculate Risk Impact.....	18
2.3.4 Step 3.4 - Perform Risk Triage	19
2.3.5 Step 3.5 - Prepare for Risk Treatment	20
2.4 Phase 4 - Treat Risks.....	20
2.5 Phase 5 - Determine Residual Risk	22
2.6 Phase 6 - Update Risk Assessment	22
Concluding Summary.....	24

Document Information

Author

Scott Perry — Scott S. Perry CPA, PLLC

Contributors

Rieks Joosten

Eric Drury

Sankarshan Mukhopahyay — Dhiway

Karla McKenna — GLEIF

Jan Lindquist

Karen Hand — Precision Strategic Solutions

Sumiran Garg

Judith Fleenor, Trust Over IP Foundation

Revision History

Version	Date Approved	Revisions
1.0	24 August 2021	Initial Publication

Terms of Use

These materials are made available under and are subject to the Creative Commons Attribution 4.0 International license (<http://creativecommons.org/licenses/by/4.0/legalcode>).

THESE MATERIALS ARE PROVIDED "AS IS." The Trust Over IP Foundation, established as the Joint Development Foundation Projects, LLC, Trust Over IP Foundation Series ("ToIP"), and its members and contributors (each of ToIP, its members and contributors, a "ToIP Party") expressly disclaim any warranties (express, implied, or otherwise), including implied warranties of merchantability, non-infringement, fitness for a particular purpose, or title, related to the materials. The entire risk as to implementing or otherwise using the materials is assumed by the implementer and user.

IN NO EVENT WILL ANY ToIP PARTY BE LIABLE TO ANY OTHER PARTY FOR LOST PROFITS OR ANY FORM OF INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THESE MATERIALS, ANY DELIVERABLE OR THE ToIP GOVERNING AGREEMENT, WHETHER BASED ON BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, AND WHETHER OR NOT THE OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

RFC 2119






The Internet Engineering Task Force (IETF) is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and to ensure maximal efficiency in operation. IETF has been operating since the advent of the Internet using a Request for Comments (RFC) to convey “current best practice” to those organizations seeking its guidance for conformance purposes.

The IETF uses RFC 2119 to define keywords for use in RFC documents; these keywords are used to signify applicability of requirements. ToIP has adapted the IETF RFC 2119 for use in the Risk Assessment Companion Guide, and therefore its applicable use in ToIP-compliant governance frameworks.




The RFC 2119¹¹ keyword definitions and interpretation have been adopted by users of the Risk Assessment Companion Guide who follow these guidelines SHOULD incorporate the following phrase near the beginning of their document:

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#).

RFC 2119 defines these keywords as follows:

-  **MUST:** This word, or the terms "REQUIRED" or "SHALL", mean that the definition is an absolute requirement of the specification.
-  **MUST NOT:** This phrase, or the phrase "SHALL NOT", means that the definition is an absolute prohibition of the specification.
-  **SHOULD:** This word, or the adjective "RECOMMENDED", means that there MAY exist valid reasons in particular circumstances to ignore a particular item, but the full implications MUST be understood and carefully weighed before choosing a different course.
-  **SHOULD NOT:** This phrase, or the phrase "NOT RECOMMENDED" means that there MAY exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications SHOULD be understood, and the case carefully weighed before implementing any behavior described with this label.
-  **MAY:** This word, or the adjective "OPTIONAL", means that an item is truly optional. One vendor MAY choose to include the item because a particular marketplace requires it or because the vendor feels that it enhances the product while another vendor MAY omit the same item.

Requirements include any combination of Machine-Testable Requirements and Human-Auditable Requirements. Unless otherwise stated, all Requirements MUST be expressed as defined in [RFC 2119](#).






-  **Mandates** are Requirements that use a MUST, MUST NOT, SHALL, SHALL NOT or REQUIRED keyword.
-  **Recommendations** are Requirements that use a SHOULD, SHOULD NOT, or RECOMMENDED keyword.
-  **Options** are Requirements that use a MAY or OPTIONAL keyword.

¹ <https://datatracker.ietf.org/doc/html/rfc2119>. Accessed June, 2021.

An implementation which does not include a particular option **MUST** be prepared to interoperate with other implementations which include the option, recognizing the potential for reduced functionality. As well, implementations which include a particular option **MUST** be prepared to interoperate with implementations which do not include the option and the subsequent lack of function the feature provides.

Executive Summary

Governance frameworks operate in environments where risks occur. A key objective of governance frameworks is to transparently convey the requirements necessary to mitigate perceived risks to the governing domain (level of the ToIP stack). Ensuring the successful development of a robust governance framework requires thoughtful consideration, broad input from all domain members, and the appropriate knowledge and tools. The Trust Over IP Risk Assessment Companion Guide (RACG) combined with the associated Risk Assessment Worksheet template provides ecosystems and governance framework architects with the knowledge and tools to perform a risk assessment grounded in generally accepted global standards and techniques, that allow for:

-  Proper consideration and identification of potential risks,
-  Critical analysis of potential risks in terms of likelihood and severity needed to calculate a systematic risk impact score,
-  Triage of risks for further treatment,
-  Treatment of the risks using a variety of options that include creation of risk mitigation requirements as part of the governance framework, and
-  Performance of an annual review of risks to ensure criticality of current risks and the consideration of new or emerging risks.

Governance frameworks provide governance authorities with the operational architecture to ensure the continued well-being of trusted digital ecosystems.

Introduction

The Trust Over IP Foundation (ToIP), by defining a complete architecture for Internet-scale digital trust, seeks to enable trusted ecosystems comprised of individuals and organizations - to leverage collective intelligence and expertise, enable innovative business opportunities, and innovative solutions to societal challenges related to our environment, health, productivity, and resource allocation. To succeed, emerging ecosystems require governance authorities and a robust governance framework to identify and mitigate risks with the potential to harm individuals, the organizations, and the overall well-being of the network.

Governing authorities depend on information technology, infrastructure, and systems to achieve their operating objectives. Governance frameworks are a construct of business, technical and information trust requirements; clearly defined, stated, and monitored in order to control risk. Through the exploitation of both known and unknown vulnerabilities, governance frameworks are subject to serious threats that can compromise the ability of a governing authority and participating roles to maintain system level confidentiality, integrity, and availability of the information being processed, stored, and transmitted.

Threats to information systems can include purposeful attacks, environmental disruptions, human/machine errors, and structural failures, and can result in harm to the governing authority's ability to achieve its objectives, undermine trust in the system and organizations, and significantly impact reputation. Therefore, it is imperative that all participants in a governance framework understand their responsibilities and are held accountable for managing information security risk—that is, the risk associated with the operation and use of information systems that support the missions and business functions of their organizations.

Risk assessment is one of the fundamental components of the ToIP Governance Metamodel (i.e., governance blueprint). Risk assessments are used to identify, estimate, and prioritize risk to governance operations and provide subsequent justification for the governance requirements. A comprehensive/systematic risk assessment needs to consider:

1. Relevant threats to the scope of a governance framework including the artifacts and the roles entrusted to process them,
2. Vulnerabilities within the scope of the governance framework and external threats related to transitive trust,
3. Impact (i.e., harm) to governed parties that may occur given the potential for threats exploiting vulnerabilities, and
4. Likelihood that harm will occur.

The result is a determination of risk (typically a function of the degree and likelihood of harm) to inform intelligent mitigation requirements/processes for effective risk management (e.g., disclosure) and continued reliance on ecosystem operations.

Purpose

This document provides emerging ecosystems, organizations, and individuals with the guidance to conduct a risk assessment as an integral component of a governance framework and is intended to serve a diverse group of professionals including:

- Architects of governance frameworks,
- Governing authority leaders,
- Functional governed party leaders with responsibilities for conducting organizational missions/business functions (e.g., mission/business owners, information owners/stewards, authorizing officials),
- Individuals with responsibilities for acquiring information technology products, services, or information systems (e.g., acquisition officials, procurement officers, contracting officers),
- Individuals with information system/security design, development, and implementation responsibilities (e.g., program managers, enterprise architects, information security architects, information system/security engineers, information systems integrators),
- Individuals with information security oversight, management, and operational responsibilities (e.g., chief information officers, senior information security officers, information security managers, information system owners, common control providers), and
- Individuals with information security/risk assessment and monitoring responsibilities (e.g., system evaluators, penetration testers, security control assessors, risk assessors, independent verifiers/validators, inspectors general, auditors).

Risk assessment provides governance architects and participants justification to determine appropriate course of action in response to identified risks. Importantly, this document aligns with the ideals of ToIP governance models and meets the generally accepted criteria for risk assessment processes. This document is intended to be used in conjunction with the Risk Assessment Worksheet (RAW) template in order to capture the output and directives identified in the assessment process.

1. Risk Assessment Overview

1.1 ISO 27005

The ISO 27005 standard on information security risk management within the 27001 Information Security Management System family of standards is globally recognized as the de facto information security management system standard and provides the basis for this ToIP recommended approach to risk assessment. ISO 27005, derived from a broad ISO Risk Management Standard (ISO 31000), applied to information trust provides the foundation for ToIP governance frameworks.

There are many acceptable methods of risk assessment, and the ISO 27005 standard is method agnostic (i.e., it does not specify or recommend any one specific risk assessment method). Therefore, governance framework architects need to ensure they select an appropriate methodology (e.g., as provided by ISO 31010), that can be operationalized and administered within the context of their governance scope. Regardless of methodology, the ISO 27005 standard endorses the following process of continual (and sometimes iterative) structured sequence of activities:

- ✚ Establish the risk management context (e.g., the scope, compliance obligations, approaches/methods to be used and relevant policies and criteria such as the governance authority's risk tolerance or appetite),
- ✚ Quantitatively or qualitatively assess (i.e., identify, analyze and evaluate) relevant information risks, considering the information assets, threats, existing controls, and vulnerabilities to determine the likelihood of incidents or incident scenarios, and the predicted business consequences if they were to occur, to determine a 'level of risk',
- ✚ Treat (i.e., modify [use information security controls], retain [accept], avoid and/or share [with third parties]) the risks appropriately, using those 'levels of risk' for prioritization,
- ✚ Keep stakeholders informed throughout the process; and
- ✚ Monitor and review risks, risk treatments, obligations, and criteria on an ongoing basis, identifying and responding appropriately to significant changes.

These processes describe the general approach to ISO 27005 risk assessment outlined in Figure 1.

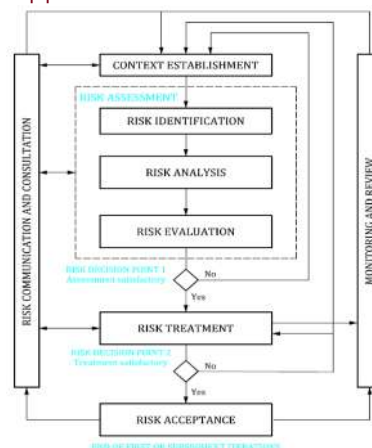


Figure 1. ISO 27005 Risk Assessment Process

1.2 NIST 800-30

NIST 800-30, is a guide for conducting a risk assessment as an application of the ISO 27005 risk management framework. Key components of the ToIP risk assessment methodology were adapted from the (U.S. Department of Commerce) National Institute of Standards (NIST) Special Publication 800-30 [Guide for Conducting Risk Assessments](#). In developing standards and guidelines, NIST consults with other U.S. federal agencies and offices as well as the private sector to improve information security, avoid unnecessary and costly duplication of effort, and ensure NIST publications are complementary with the standards and guidelines employed for the protection of national security systems.

The RECOMMENDED ToIP risk assessment methodology includes:

1. A risk assessment process (as described in this document),
2. A recommended risk model, defining key terms, assessable risk factors, and relationships among the factors,
3. A highly qualitative approach, classifying risk factors based on their assigned values in order to readily assess patterns and interactions; thereby providing a formulaic approach to risk evaluation, and
4. An analytical approach for the identification and risk factor interactions to ensure adequate coverage of the problem space at a consistent level of detail.

1.3 Risk Concepts and Terminology

The following section refers to a number of publicly available NIST definitions and descriptions.² **Risk** can be defined both conceptually and operationally. ISO 31000 defines risk as the “effect of uncertainty on objectives”³, and for organizations, it is the deviation from expected outcomes (whether positive or negative). (NIST 800-30) adopts a more traditional/operational definition of risk “a measure of the extent to which an entity is threatened by a potential circumstance or event”⁴. The NIST definition is more traditional and typically a function of: (i) the adverse impacts that would arise if the circumstance or event occurs; and (ii) the likelihood of occurrence. Information security risks are those risks that arise from the loss of confidentiality, integrity, or availability of information or information systems and reflect the potential adverse impacts to governing authority operations (i.e., mission, functions, image, or reputation), governance assets, governed parties, other organizations, and relying parties.

Risk assessment (NIST 800-30) is the process of identifying, estimating, and prioritizing risks. Assessing risk requires the careful analysis of threat and vulnerability information to determine the extent to which circumstances or events could adversely impact an organization and the likelihood that such circumstances or events will occur.

A **threat** (NIST 800-30) is any circumstance or event with the potential to adversely impact governing authority or governed party operations and assets, other organizations, or relying

² <https://csrc.nist.gov/glossary>. Accessed June 2021

³ <https://www.iso.org/news/ref2263.html#:~:text=Risk%20is%20now%20defined%20as,on%20an%20organization's%20decision%20making>. Accessed June 2021

⁴ <https://csrc.nist.gov/glossary/term/risk>. Accessed June 2021

parties through an information system via unauthorized access, destruction, disclosure, or modification of information, and/or denial of service. Threat events are created from threat sources. A threat source is characterized as:

- 👤 the intent and method targeted at the exploitation of a vulnerability; or
- 👤 a situation and method that may accidentally exploit a vulnerability.

In general, types of threat sources include:

- 👤 hostile cyber or physical attacks,
- 👤 human errors of omission or commission,
- 👤 structural failures of organization-controlled resources (e.g., hardware, software, environmental controls), and
- 👤 natural and man-made disasters, accidents, and failures beyond the control of the organization.

A **vulnerability** (NIST 800-30) is a weakness in an information system, system security procedures, internal controls, or implementation that could be exploited by a threat source. Most information system vulnerabilities can be associated with security controls that either have not been applied (either intentionally or unintentionally), or have been applied, but retain some weakness. However, it is also important to allow for the possibility of emergent vulnerabilities that can arise naturally over time as organizational missions/business functions evolve, environments of operation change, new technologies proliferate, and new threats emerge. In the context of such change, existing security controls may become inadequate and may need to be reassessed for effectiveness.

The **likelihood of occurrence** (NIST 800-30) is a weighted risk factor based on an analysis of the probability that a given threat is capable of exploiting a given vulnerability (or set of vulnerabilities). The likelihood risk factor combines an estimate of the likelihood that the threat event will be initiated with an estimate of the likelihood of impact (i.e., the likelihood that the threat event results in adverse impacts). Note, the likelihood that a threat event will be initiated or will occur is assessed with respect to a specific time frame (e.g., the next six months, the next year, or the period until a specified milestone is reached). The likelihood of threat occurrence can also be based on the state of the organization (including for example, its core mission/business processes, enterprise architecture, information security architecture, information systems, and environments in which those systems operate)—taking into consideration predisposing conditions and the presence and effectiveness of deployed security controls to protect against unauthorized/undesirable behavior, detect and limit damage, and/or maintain or restore mission/business capabilities. The likelihood of impact addresses the probability (or possibility) that the threat event will result in an adverse impact, regardless of the magnitude of harm that can be expected.

The **level of impact** (NIST 800-30) from a threat event is the magnitude of harm that can be expected to result from the consequences of a risk. Such harm can be experienced by a variety of organizational and non-organizational stakeholders including, for example, heads of governing authorities, governed party owners, information owners/stewards, business process owners, information system owners, or individuals/groups in the public or private sectors relying on the governance framework—in essence, anyone with a vested interest in the ecosystem’s operations, assets, or individuals, including other organizations in partnership with the organization.

The figure below illustrates an example of a **risk model** including the key risk factors discussed above and the relationship among the factors. Each of the risk factors is used in the risk assessment process described later in this document.

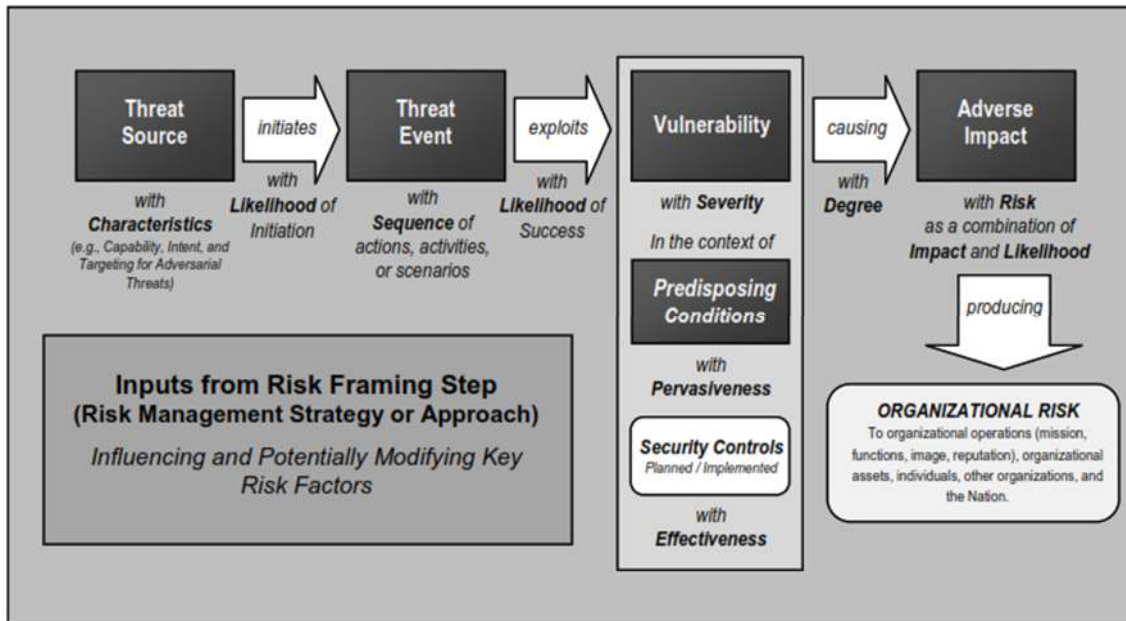






Figure 2: Generic Risk Model with Key Risk Factors (NIST 800-30)

Governing authorities may use **risk aggregation** (NIST 800-30) to roll up several discrete risks (e.g., risk associated with a single information system supporting a well-defined mission/business process) or lower-level risks into a generalized or higher-level risk. Risk aggregation assesses the overall risk to ecosystem operations, assets, and individuals given the entire discrete risk set, often using the worst-case impact scenario to establish the risk upper bound and subsequent risk capacity. One challenge of using worst-case discrete risk impacts in risk aggregation is the possibility of underestimating the organizational risk potential. Consider the following two scenarios; multiple risks that occur concurrently or the same risk occurs repeatedly over a period of time. In either scenario, it is possible that the overall risk incurred is greater than the determined upper bound and organizational risk capacity; unfortunately, impacting the governing authority's organizational operations and assets beyond expected or acceptable levels.

Uncertainty (NIST 800-30) is inherent in the evaluation of risk, due to such considerations as:

-  limitations on the extent to which the future will resemble the past,
-  imperfect or incomplete knowledge of the threat (e.g., characteristics of adversaries including tactics, techniques, and procedures),
-  undiscovered vulnerabilities in technologies or products, and
-  unrecognized dependencies, which can lead to unforeseen impacts.

Uncertainty can be due to incomplete knowledge of the risks associated with other information systems, mission/ business processes, services, common infrastructures, and/or organizations. There is inherent uncertainty associated with estimated risk values. Therefore, it is logical to express risk using categories that encompass a range of values.

Risk Tolerance is determined as part of the ecosystem risk management strategy to ensure ecosystem-wide consistency. Each ecosystem determines its own levels and types of risk that are acceptable and provides guidance on how to identify uncertainty (which can accrue and impact the overall risk potential). Risk tolerance is defined as the level of risk and degree of uncertainty acceptable to the ecosystem.

1.4 Risk Management Hierarchy

An often-anticipated question when considering a risk assessment is, “*What is the scope of the risk assessment?*”. Typically, the scope of the assessment aligns with the governance framework used to mitigate the risks. In a Governance Stack based on the ToIP Governance Metamodel, this can occur ecosystem-wide or within a specific stack level. The following figure illustrates the risk management hierarchy adapted from NIST Special Publication 800-39, which provides multiple risk perspectives from the strategic to tactical level. Traditional risk assessments generally focus on the Tier 3 level (i.e., information system level) and as a result, tend to overlook significant risk factors associated with higher levels.

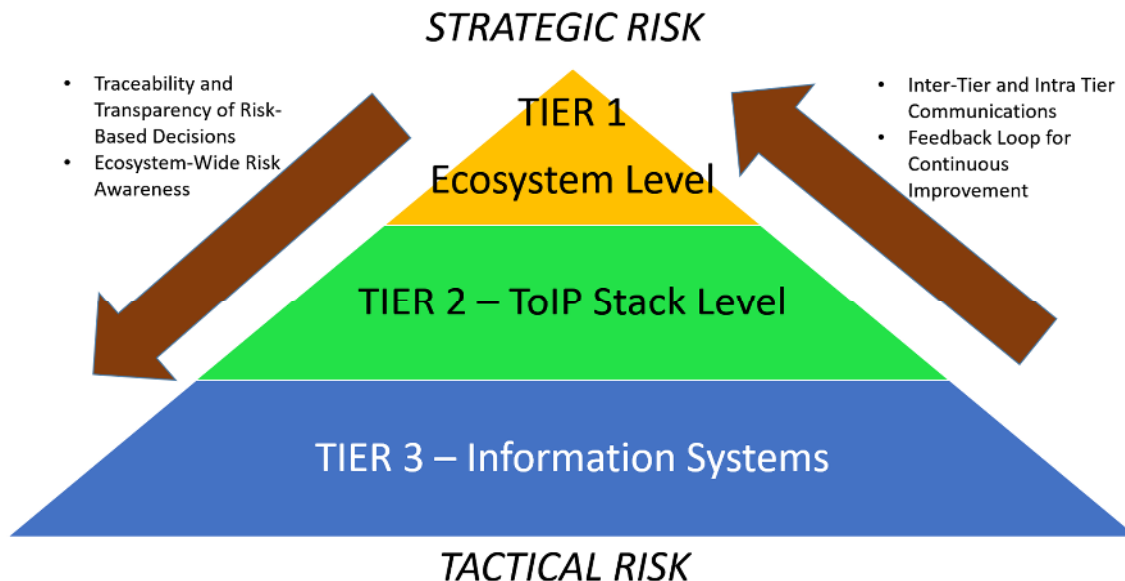


Figure 3. Risk Management Hierarchy

Risk assessments support risk response decisions at the different tiers of the risk management hierarchy.

At the Ecosystem level, risk assessments can potentially affect:

1. Ecosystem-wide information security programs, policies, procedures, and guidance,
2. Appropriate risk responses (i.e., risk acceptance, avoidance, mitigation, sharing, or transfer),
3. Investment decisions for information technologies/systems,
4. Ecosystem-scale procurements,
5. Broad minimum ecosystem-wide security controls,
6. Ecosystem-wide conformance to enterprise/security architectures, and

7. Ecosystem-wide monitoring strategies and ongoing authorizations of information systems and common controls.

At the Stack level, risk assessments can potentially affect:

1. Stack level architecture/security architecture design decisions,
2. Selection of common controls at the stack level,
3. Selection of suppliers, services, and contractors to support stack missions/business functions,
4. Development of risk-aware mission/business processes operating at the stack level, and
5. Interpretation of ecosystem driven requirements with respect to organizational information systems and environments in which those systems operate.

Finally, at an Information system level, risk assessments can potentially affect:

1. Design decisions (including the selection, tailoring, and supplementation of security controls and the selection of information technology products for organizational information systems),
2. Implementation decisions (including whether specific information technology products or product configurations meet security control requirements); and
3. Operational decisions (including the requisite level of monitoring activity, the frequency of ongoing information system authorizations, and system maintenance decisions).

2. ToIP Risk Assessment Process





ToIP RECOMMENDS a generally accepted, systematic approach to risk assessments that occurs in six, waterfall⁵ model driven phases:

1. Establish Context,
2. Identify Risks,
3. Analyze Risks,
4. Treat Risks,
5. Determine Residual Risks, and
6. Update Risk Assessment




2.1 Phase 1 - Establish Context

The first step in the risk assessment process is to prepare for the assessment and establish a context for the risk assessment. The context is established in conjunction with the development and guidance of the Primary Document section of the Governance Framework.

Specifically, the following sections drive the risk assessment context:










-  Introduction and overview - (governing sphere of influence involving risks),
-  Purpose - (of Governing Authority to manage risks),
-  Scope - (roles, data, artifacts (tangible constructs of ecosystem involving risk), and
-  Objectives - (purposeful outcome of governing authority's treatment of risk).

In addition, the following should be considered in preparing for the assessment:

-  Identify the assumptions and constraints associated with the assessment,
-  Identify the sources of information to be used as inputs to the assessment, and
-  Identify the risk model and analytic approaches (i.e., assessment and analysis approaches) to be employed during the assessment.

2.2 Phase 2 - Identify Risks

After establishing context, the next step in a risk assessment is identifying candidate risks to consider. Candidate risks can be derived from many sources, both technical and non-technical. The following types of risk should be considered (not a complete list):

-  Governance,
-  Business,
-  Technical,
-  Information trust (e.g., availability, security, confidentiality, privacy, processing integrity),
-  Reputational,
-  Financial,
-  Social and humanitarian (including guardianship and inclusivity),
-  Adversarial (with capability, intent and targeting),
-  Accidental, and

⁵ https://en.wikipedia.org/wiki/Waterfall_model

Environmental.

The Risk Assessment Worksheet (RAW) Template, to be used in conjunction with this companion guide, contains a starter set of candidate risks to consider based on ToIP member expertise working with governing authorities and their ecosystems. However, if a governance framework has been developed in the absence of a risk assessment, a reverse engineering process can generate candidate risks by inventorying the MUST statements in the governance framework and determining the risks these requirements are seeking to mitigate. For example, if a governance framework has a requirement that all authoritative issuers MUST verify a set of data before issuing credential, it is apparent that it is attempting to mitigate the risk that credentials will be issued without data verification.

Risks may be different from the perspective of varying ecosystem roles. Therefore, it is appropriate for risks during risk assessments to be grouped by role. The RAW template has adopted this approach.

As ToIP ecosystems include data and verifiable credentials within its scope, risk assessors MUST consider the sensitivity of data flow and storage throughout the network. Does data carry risks of disclosure, transmission, storage, or custodianship? How is privacy considered? Is data subject adequately of their privacy rights in a consent notice? Is there an understanding of the privacy risks and compliance risks for emerging privacy regulations? ToIP ecosystems are intended to be built with *Privacy by Design* and data protection by default to mitigate privacy risks.

To finalize a set of candidate risks for further analysis, it is critical to consider potential risks from a variety of perspectives. Therefore, it is RECOMMENDED that a broader community of stakeholders, at minimum, comprising legal, financial, quality management, privacy, security, system development and potential users, SHOULD contribute their perspective on potential risks.




At the end of this phase, a completed set of candidate risks, categorized by role potential risks completing Column B in the associated RAW template is completed. It is important to note, additional potential risks can always be amended to the worksheet after an initial set of risks are analyzed and finalized in later steps.

2.3 Phase 3 - Analyze Risk

Once all candidate risks are plotted on the associated RAW template, then they SHOULD be analyzed by likelihood and severity, to determine the risk impact and triaged for further treatment.

2.3.1 Step 3.1 - Analyze and Determine Risk Likelihood

In analyzing the likelihood of the risk occurring, one must determine the likelihood that threat events of concern result in adverse impacts, by considering:

-  Characteristics of the threat sources that could initiate the events,
-  Vulnerabilities/predisposing conditions identified, and
-  The tier's (ecosystem, stack level or information system) susceptibility, reflecting the safeguards/countermeasures planned or implemented to impede such events.

To maximize the efficacy of countermeasures planned in Phase 4, the risk score should be qualitatively assessed under the assumption of **no countermeasures in place**. This is the **inherent risk**. This often results in conservative estimates of the likelihood; however, this approach will help justify the countermeasures (to the relying parties of the governance framework). Furthermore, appropriate countermeasures will offset the severity of the risk score in the residual risk (determined in Phase 5).




The associated worksheet has a column for likelihood with the following embedded options:

1. Highly unlikely,
2. Unlikely,
3. Possible,
4. Likely, and
5. Highly likely.

It is expected that risk assessors will have different perspectives on risk likelihood and differing levels of what is acceptable. Different perspectives generate discussion and debate; in the risk assessment process, this should be encouraged. Some risk assessors may choose to reduce the number of likelihood categories to three (likely, possible, and unlikely) to simplify discussion and not wrangle over categorization. If risk assessment participants are unfamiliar with the process, it is RECOMMENDED to start simple, effectively execute and then they move to a more mature/complex framework once they have gained confidence in risk management. However, it is important to have agreement on a normative scale. The generally accepted risk posture is to err on the side of being conservative.

2.3.2 Step 3.2 - Analyze and Determine Risk Severity

In analyzing the severity of the risk occurring, one must determine the most likely level of adverse impacts from threat events of concern, considering:

-  Characteristics of the threat sources that could initiate the events,
-  Vulnerabilities/predisposing conditions identified, and
-  The tier's (ecosystem, stack level or information system) susceptibility reflects the safeguards/countermeasures planned or implemented to impede such events.

To maximize efficacy of countermeasures planned in Phase 4, the risk score should be qualitatively assessed under the assumption of **no countermeasures in place**. This is considered **inherent risk**. This often results in conservative estimates of the risk score; however, this approach will help justify the countermeasures (to the relying parties of the eventual governance framework). Furthermore, appropriate countermeasures will offset the severity of the risk score in the residual risk (determined in Phase 5).

The associated worksheet has a column for severity with the following embedded options:






1. Negligible,
2. Minor,
3. Moderate,
4. Major, and
5. Critical.

It is expected that risk assessors will have different perspectives on risk severity and differing levels of what is acceptable. Different perspectives generate discussion and debate; in the risk assessment process, this should be encouraged. Some risk assessors may choose to reduce the number of severity categories to three (minor, moderate and major) to simplify discussion and not wrangle over categorization. If risk assessment participants are unfamiliar with the process, it is RECOMMENDED to start simple, effectively execute and then they move to a more mature/complex framework once they have gained confidence in risk management. However, it is important to have agreement on a normative scale. The generally accepted risk posture is to err on the side of being conservative.

2.3.3 Step 3.3 - Calculate Risk Impact

In this ToIP RECOMMENDED approach, the risk impact score is a calculated result considering the factors of risk likelihood and risk severity. While other models extend impact in quantifiable terms (e.g., making estimates in monetary terms) this is considered OPTIONAL and therefore, not included in this guide.

The associated RAW template will automatically calculate the risk impact score based on the multiplication of the likelihood and severity rankings resulting in the following ranges of impact categories:

-  Low: 1-3,
-  Low-Medium: 4-7,
-  Medium: 8-12,
-  Medium-High: 13-18, and
-  High: 19-25.

The range of values associated with the severity rankings are derived from the expertise of the ToIP membership and in consideration of the NIST 800-30 guide. It is likely that risk assessors adopting this approach would use a different set of range values to align with their business strategy and risk management guidelines. If a different number of categories are chosen in steps 3.1 and 3.2, a commensurately conservative risk impact scoring should be established.

Using the list of categories in steps 3.1 and 3.2, the interrelationship between risk likelihood, severity and impact is best depicted in the following chart:

		SCALE OF SEVERITY					
		1	2	3	4	5	
		NEGLIGIBLE	MINOR	MODERATE	MAJOR	CRITICAL	
SCALE OF LIKELIHOOD	1	HIGHLY UNLIKELY	LOW	LOW	LOW	LOW - MEDIUM	LOW - MEDIUM
	UNLIKELY	LOW	LOW - MEDIUM	LOW - MEDIUM	MEDIUM	MEDIUM	
	POSSIBLE	LOW	LOW - MEDIUM	MEDIUM	MEDIUM	MEDIUM-HIGH	
	LIKELY	LOW - MEDIUM	MEDIUM	MEDIUM	MEDIUM-HIGH	HIGH	
	HIGHLY LIKELY	LOW - MEDIUM	MEDIUM	MEDIUM-HIGH	HIGH	HIGH	

Figure 4. Risk Management Chart

The calculation of risk impact is not uniformly distributed, on purpose. The approach is to be conservatively biased slanting toward higher risk impact scores. These scores are normalized for further treatment in the next step (Step 3.4).

2.3.4 Step 3.4 - Perform Risk Triage

The resulting range of risk impacts creates a decision point for risk assessors. Every candidate risk on the associated RAW template requires further effort to determine the risk treatment and the calculation of residual risk after treatment. It requires (at minimal) annual monitoring to determine the effectiveness of risk treatment actions. Therefore, it is RECOMMENDED that a triage of the candidate risks be performed in order to eliminate a set from further processing.




Having reached this stage in the RECOMMENDED waterfall approach, one begins to have a better understanding of the preponderance of risk from the variety of sources that need to be considered by ToIP ecosystems. The question is where to draw the line between those risks deemed important for further work and those where no further action is required. Ultimately, the decision “line” lies with the risk assessor; however, it is RECOMMENDED to draw a line between low-medium impact risks and medium impact risks. This will allow all medium and higher impact risks to be included in further risk management processing.

To perform the risk triage, eliminate all low and low-medium impact risk candidates from the main tab of the RAW template. It is RECOMMENDED to move discarded risks to a newly created tab

(e.g., DISCARDED RISKS) for full transparency to all participating roles. Discarded risks can be re-analyzed during annual risk assessments (Phase 6) in the event that changes in the market and ecosystem conditions affect the risk likelihood and severity, resulting in a risk impact score that justifies re-adding the risk to the risk management process.

2.3.5 Step 3.5 - Prepare for Risk Treatment

At this point, a triage of candidate risks has been completed, eliminating those risks that fall below the demarcation line. To prepare for further processing, it is RECOMMENDED that one completes the final columns in the RAW template:

-  Numbering Scheme (column A) - In developing risk treatment options, especially risk mitigation actions evidenced as requirements on a governance framework, it is useful to have a unique numbering scheme in place so risk treatments can be easily linked to risks for reference and completeness.
-  ToIP Layer (column C) - The ToIP Governance Stack is segmented into four distinct layers. One can add to this list or change it to other descriptions as needed. The associated RAW template has the following embedded values to choose from:
 - ❖ Ecosystem,
 - ❖ Credential,
 - ❖ Provider,
 - ❖ Utility, and
 - ❖ Blank.
-  Trust Area Affected (column D) - This column reflects the Information Trust Area or Governance category. One can add to this list or change it to other descriptions as needed. The associated RAW template has the following embedded values to choose from (Please refer to the ToIP Glossary for definitions):
 - ❖ Availability,
 - ❖ Confidentiality,
 - ❖ Processing integrity,
 - ❖ Governance,
 - ❖ Privacy,
 - ❖ Security, and
 - ❖ Other.

2.4 Phase 4 - Treat Risks

In the associated RAW template, refer to the column for risk treatment. The following are provided as options:

1. Mitigate,
2. Avoidance,
3. Accept, and
4. Transfer.

This section will explain each option and the criteria for selection.

One would select the **mitigate** option if a requirement of the impending governance framework will or currently operates to reduce the risk. Typically, this is described as a MUST statement in

the governance framework which drives the mandate to be complied. The degree of compliance is dependent on the trust assurance framework in place to hold governed parties accountable to that requirement and the degree of effectiveness that the trust assurance framework has in reducing the risk to an acceptable residual level. This option will likely cover the majority of risks that have been analyzed thus far.

One would select the **avoid** option if the action taken to treat the risk is to abandon efforts to allow the risk to be realized by the governance framework. Usually, the action taken to avoid risk is to reduce the scope of the governance framework to not include the elements that would perpetuate the risk. For example, a governance framework architect may be considering including roles of both Issuer and Verifier on his credential governance framework. After considering the risks of verification, if the governance architect chooses to avoid verification as a service entirely and focus solely on issuance, they would eliminate verifiers and verification from the scope of the governance framework.

One would select the **accept** option if the action taken to treat the risk is to not take any action against the risk by the governance framework. This typically manifests itself when the risk impact is lower or if there are known inherent risks that cannot be treated by a governance framework. Usually, this path is taken when the perceived value of accepting the risk overwhelmingly outweighs the cost of the risk impact when the risk is realized. Organizations accept risks every day when they can justify the value of going into nascent markets, driving innovative technology, or introducing revolutionary products and services to a market. If a risk is accepted, the residual risk would be the same as the risk impact score since no action would be taken to reduce it.

Finally, one would select the **transfer** option if the risk assessor plans to direct the risk treatment to a party outside the scope of the governance framework. This can happen at all risk management tiers to different degrees. At tier 1 (ecosystem), the risk transference could be to another ecosystem to treat; at Tier 2 (stack level) it could be transferred to another stack level to treat, and so on. However, it is more likely to see risk transference occurring within a stack level risk assessment where there are interoperable dependencies between stack layers. For example, on a credential governance framework (ToIP layer three) risk assessment, a risk of unavailable decentralized identifiers (DIDs) on a utility's decentralized layer technology (DLT) network will be transferred to a utility governance framework (ToIP layer one) to treat (most likely with mitigation). Again, the scope statement on the associated governance framework would be reduced to exclude addressing the risk, and a dependency is built from the layer three governance framework to the layer one governance framework to treat the risk.

Often, the action taken to avoid risk is to reduce the scope of the offered services and the governance framework (e.g., remove elements that perpetuate risk). For example, a governance framework architect may be considering including roles of both Issuer and Verifier on his credential governance framework. After considering the risks of verification, the architect chooses to avoid verification as a service entirely and focus solely on issuance. They would then eliminate verifiers and verification from the scope of the governance framework.

Another generally accepted risk treatment is to **share** the risk. Risk sharing is the distribution of risk to multiple organizations. This is done for a variety of reasons including insurance products

and self-insurance. This option is not included in the associate RAW template since it does not seem to apply in a conceivable use case.

Column L (Notes) can be used to provide more context into the risk treatment decision, especially the justification of any risk acceptance.

2.5 Phase 5 - Determine Residual Risk

Risk treatments do not reduce risks to zero. Even the best risk mitigation programs do not prevent all risks from occurring. The assessment of risk remaining after the risk treatment option taken in Phase 4 is known as **residual risk**. Common residual risks will be realized through ineffectively operating governance requirements.

For example, a university ecosystem governance framework may mandate that all degree requirements must be thoroughly checked through a system of manual and automated controls to mitigate the risk of degree credentials being issued wrongly to individuals who have not satisfied degree requirements. This is dependent on the effectiveness of the manual and automated controls in place. Later, as a result of an annual audit, it was determined that those controls were not in place and operating effectively, resulting in a residual risk that a portion of degree credentials did not support valid satisfaction of degree requirements.

How much residual risk is acceptable? This is a valid question, and the answer is “it depends”. Residual risk can result in significant cost, loss of reputation, lack of confidence in the ecosystem and the governance framework to maintain an acceptable level of trust. It is up to the risk assessors and relying parties to determine the level of residual risk that is acceptable. In Phase 4, experience has shown, some level of risk will be accepted as a treatment option. Bottom line, the effect of residual risk **MUST** be analyzed against the cost and effort of treatments needed to reduce it further in order to determine if those actions are justified. Residual risk should be deemed acceptable if the organization understands and accepts the risk given its risk tolerance. Risk mitigation steps should ensure that residual risk minimally meets that level.

One component concerning the communication of risk assessment results is disclosure of likely residual risk. The process of disclosure effectively shares the residual risk with relying parties so they can appropriately gauge the level of trust they can expect from a governance framework to achieve the objectives defined within its disclosed scope.

Column L (Notes) can be used to provide more context into the Residual Risk assessment including actions to be taken for more than acceptable risk.




2.6 Phase 6 - Update Risk Assessment

The sixth phase in the risk assessment process is to maintain the assessment, in order to keep specific knowledge concerning all identified risks current. The results of risk assessments drive risk management decisions and guide risk response. To support the ongoing review of risk management decisions affecting governance, ecosystems need to maintain risk assessments at least annually (maybe more frequently in a significantly changing risk environment) to incorporate any changes detected through their trust assurance framework. Part of maintaining an effective

trust assurance framework is risk monitoring. Risk monitoring provides governing authorities with the means to, on an ongoing basis:


1. Determine the effectiveness of risk responses,
2. Identify risk-impacting changes to ecosystem information systems and the environments in which those systems operate, and
3. Verify compliance.

Maintaining risk assessments includes the following specific tasks on at least an annual basis:

-  Monitoring risk factors identified in risk assessments on an ongoing basis and understanding subsequent changes to those factors,
-  Updating the components of risk assessments reflecting the monitoring activities carried out by ecosystems current, and
-  Updating the governance framework and associated worksheet to reflect current risk conditions and treatments.

Concluding Summary

The Trust over IP Foundation (ToIP), by defining a complete architecture for Internet-scale digital trust, seeks to enable trusted ecosystems. Ecosystems require governance authorities and a robust governance framework. Risk assessment, a fundamental component of the ToIP Governance Metamodel, provides governance authorities with the operational architecture to mitigate risk - identify, estimate, prioritize, and respond. A well-executed risk assessment provides governance architects with varying perspectives (through broad input), provides justification for specific governance requirements, demonstrates actions required to mitigate risks, and informs the ecosystem of residual risk. The Trust over IP Risk Assessment Companion Guide (RACG) combined with the associated Risk Assessment Worksheet (RAW) template provides ecosystems and governance framework architects with the knowledge and tools required to perform a systematic risk assessment informed by globally accepted standards. Thus, ensuring the continual success of ToIP enabled trusted digital ecosystems, innovative business opportunities, and innovative solutions to societal challenges related to our environment, health, productivity, and resource allocation.



TRUST Over IP FOUNDATION

The Trust Over IP Foundation (ToIP) is hosted by the Linux Foundation under its Joint Development Foundation legal structure. We produce a wide range of tools and deliverables organized into five categories:

- ❖ Specifications to be implemented in code
- ❖ Recommendations to be followed in practice
- ❖ Guides to be executed in operation
- ❖ White Papers to assist in decision making
- ❖ Glossaries to be incorporated in other documents

ToIP is a membership organization with three classes—Contributor, General, and Steering.

The work of the Foundation all takes place in Working Groups, within which there are Task Forces self-organized around specific interests. All ToIP members regardless of membership class may participate in all ToIP Working Groups and Task Forces.

When you join ToIP, you are joining a community of individuals and organizations committed to solving the toughest technical and human centric problems of digital trust. Your involvement will shape the future of how trust is managed across the Internet, in commerce, and throughout our digital lives. The benefits of joining our collaborative community are that together we can tackle issues that no single organization, governmental jurisdiction, or project ecosystem can solve by themselves. The results are lower costs for security, privacy, and compliance; dramatically improved customer experience, accelerated digital transformation, and simplified cross-system integration.

To learn more about the Trust Over IP Foundation please visit our website, <https://trustoverip.org>.

Licensing Information:

All Trust Over IP Foundation deliverables are published under the following licenses:

Copyright mode: Creative Commons Attribution 4.0 International licenses
<http://creativecommons.org/licenses/by/4.0/legalcode>

Patent mode: W3C Mode (based on the W3C Patent Policy)
<http://www.w3.org/Consortium/Patent-Policy-20040205>

Source code: Apache 2.0.
<http://www.apache.org/licenses/LICENSE-2.0.htm>