CITRAP
Counter-Insider Threat
Research and Practice

Stewart, A., & Handy, M. (2024). The Design and Implementation of an Insider Threat Maturity Model. *Counter-Insider Threat Research and Practice*.

Scholarship In Practice

The Design and Implementation of an Insider Threat Maturity Model

Andrew Stewart[1], Mark Handy[1]

[1] Morgan Stanley (United States)

Keywords: insider threats, maturity model, organizational planning, design, implementation, visualization

Counter-Insider Threat Research and Practice

Large organizations face the practical challenge of assessing then communicating the status of their insider threat program to internal stakeholders and external regulators. This task is complicated by the absence of a definitive standard against which to perform such an assessment. Several best practice documents exist, but because those documents are authored by different institutions the content is uneven and often duplicative. The objective of this article is to describe the design and implementation of an insider threat maturity model that attempts to address these issues. Design science was selected as the theoretical framework because the goal was to create an artifact that is useful in a practical setting. In accordance with design science, both the process of designing the model and the experience of using the model in a real-world environment are detailed. The model is created by synthesizing 568 recommendations from 8 separate best practice documents published between 2013 and 2023. Based on feedback from stakeholders, both the model and a corresponding data visualization have been found to be effective tools when deployed within a large multinational investment bank.

1. Introduction

Morgan Stanley is a leading global financial services firm with offices in 42 countries and more than 80,000 employees. The nature of the business requires the insider threat team to regularly assess the insider threat program and report the status to multiple stakeholders including internally to upper management, business units, and auditors and externally to groups such as banking industry regulators. These various parties can possess different levels of knowledge regarding insider threats and may require information to be presented at different levels of technical detail. These requirements are not unique to the firm or to the financial services industry, as any large organization operating in a regulated industry will likely face similar challenges.

A maturity model is a framework for measuring an organization's capabilities with regards to a specific discipline (Wendler, 2012). The basic concept is that organizational attributes develop through a number of logical stages (or levels) from an initial level to a more mature level (Gottschalk, 2008). Organizations can therefore use a maturity model to evaluate their capabilities against an external standard, and to obtain guidance regarding how they might improve. As such, a maturity model is an appropriate tool for the task of assessing and communicating the status of an insider threat program.

A number of maturity models have been published for use in areas such as process improvement and software engineering. The basic components of a maturity model are a set of focus areas, a set of maturity levels, and descriptions of the capabilities of the organization at each level for each focus area. Commonly, maturity models define between three and six levels. For example, the Capability Maturity Model (CMM) focuses on software development practices, and defines five levels: 'initial,' 'repeatable,' 'defined,' 'managed,' and 'optimizing' (Paulk et al., 1993).

In 2018, the National Insider Threat Task Force (NITTF) published their Insider Threat Program Maturity Framework (NITTF, 2018b). That document describes nineteen "maturity elements" (meaning capabilities or attributes) of insider threat programs. Those elements are organized into seven topic areas such as "program personnel" and "access to information." However, the NITTF document does not define levels as in a traditional maturity model such as the CMM. Rather, the FAQ for the NITTF document describes the users of the model as being able to "choose among the maturity elements for those that best fit with their workplace environment, technology infrastructure, and mission" (NITTF, 2018a). In 2021, the Cybersecurity and Infrastructure Security Agency (CISA) in collaboration with Carnegie Mellon University's Software Engineering Institute released an Insider Risk Mitigation Program Evaluation (IRMPE) tool to assist organizations with evaluating the maturity of their insider threat programs. That tool uses a fillable PDF that generates a report, and is intended primarily for "small and mid-sized [organizations] that may not have in-house security departments" (CISA, 2022).

While these existing models offer valuable insights and tools for organizations, a fundamental issue with using any *single* maturity model is that the content of that model inevitably reflects the views and preferences of the authors. The possibility also exists that the content might contain biases, omissions, or even errors. Therefore, the decision

Table 1. Seven Guidelines for Design Science Research

Nr.	Guideline	Description
DS1	Design as an Artifact	Design science research must produce a viable artifact in the form of a construct, a model, a method, or an instantiation.
DS2	Problem Relevance	The objective of design science research is to develop technology-based solutions to important and relevant business problems.
DS3	Design Evaluation	The utility, quality, and efficacy of a design artifact must be rigorously demonstrated via well-executed evaluation methods.
DS4	Research Contributions	Effective design science research must be rigorously demonstrated via well-executed methods.
DS5	Research Rigor	Design science research relies upon the application of rigorous methods in both the construction and evaluation of the design artifact.
DS6	Design as a Search Process	The search for an effective artifact requires utilizing available means to reach desired ends while satisfying laws in the problem environment.
DS7	Communication of Research	Design science research must be presented effectively both to technology-oriented as well as management-oriented audiences.

was made to design a maturity model that consolidates the recommendations from *multiple* best practice documents.

2. Design

Design science is an appropriate research paradigm when the goal is to create a specific artifact such as a maturity model in a real-world setting (Gregor & Hevner, 2013). Table 1 lists seven guidelines for design science research as specified by Hevner et al. (2004). These guidelines have been used by other researchers when developing maturity models (Wendler, 2012).

Of the seven guidelines, DS1 and DS2 focus on preparatory activities before the design phase; DS3 and DS4 are centered on evaluating the outcomes post-design; DS5 and DS6 pertain to the design phase, guiding the development and refinement of the model, and DS7 emphasizes the importance of communication after the design is completed. DS5 and DS6 are therefore the two guidelines specifically focused on the topic of design. DS6 refers to design as being a "search process," with the goal of discovering an effective solution to the underlying problem. This requires knowledge of the requirements of the organization – which are described in section 1. DS5 refers to the need for "research rigor," meaning the research should be conducted using appropriate theoretical foundations and research methodologies.

The research methodology selected for this work was the Design Science Research Methodology by Peffers et al. (2007). As shown in Figure 1, Peffers et al. provide a process model with six activities and four possible entry points. The six activities are intended to be carried out in nominally sequential order. Because an organizational need and specific requirements had already been established, and because design had not yet begun, the entry point was the 'Design & Development' entry point, leading to activity 3.

The first step in delivering activity 3 was to select the best practice documents that would be used to populate the model. A number of such documents exist. For this exercise, the eight documents listed in Table 2 were selected. These documents were chosen because the reputable nature of each source was likely to produce content of an acceptable quality, because each source was unique, and because time and resource constraints limited the number of source documents to approximately this number. The documents from FINRA and SIFMA are specific to the financial services sector, which was appropriate for Morgan Stanley. Other organizations might choose to select different source documents that are specific to their sector.

The eight source documents were reviewed to identify all recommendations – meaning all declarative statements regarding courses of action that organizations should take with regards to insider threat programs. 568 recommendations were identified and placed into a spreadsheet. Open coding was then performed, in the style of Grounded Theory (Glaser & Strauss, 1967). Coding of this type is commonly used in the analysis of qualitative data for security research (e.g., Alsowail & Al-Shehari, 2022; Krombholz et al., 2017). This exercise resulted in the creation of 20 categories relating to insider threat programs.

The categories were then consolidated where appropriate in order to reduce the overall number. This was accomplished by bundling conceptually related categories together, such as by bundling the 'management,' 'governance,' and 'metrics' categories together, and by bundling the 'off-boarding' and 'terminations' categories together. The resulting list of 13 categories was as follows:

1. Asset Management
2. Backups
3. Cloud Security, Network Security, and DLP (Data Loss Prevention)
4. End-User Reporting
5. Management, Governance, and Metrics
6. Identity & Access Management
7. Incident Response
8. Monitoring
9. Off-boarding and Terminations
10. Risk Management
11. Threat Intelligence and Information Sharing
12. Training & Awareness, and EAPs (Employee Assistance Programs)
13. Vetting & Onboarding

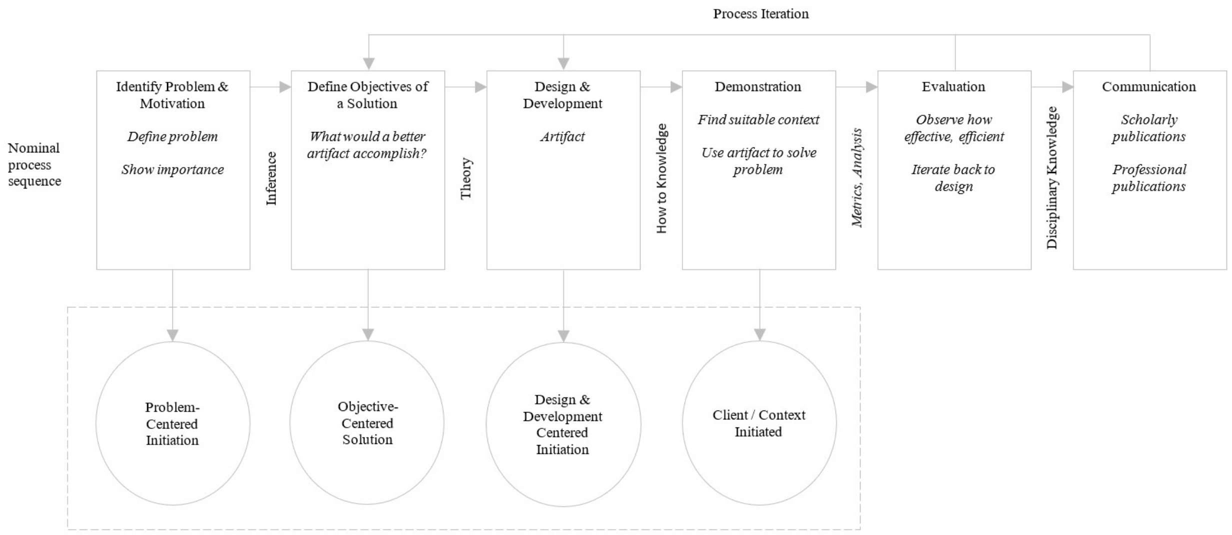
Process Iteration
Nominal process sequence
Identify Problem & Motivation
Define problem
Show importance
Inference
Define Objectives of a Solution
What would a better artifact accomplish?
Theory
Design & Development
Artifact
How to Knowledge
Demonstration
Find suitable context
Use artifact to solve problem
Metrics, Analysis
Evaluation
Observe how effective, efficient
Iterate back to design
Disciplinary Knowledge
Communication
Scholarly publications
Professional publications
Problem-Centered Initiation
Objective-Centered Solution
Design & Development Centered Initiation
Client / Context Initiated

Figure 1. Process Model

Table 2. Source Documents

Source	Title	Publication Year
Financial Industry Regulatory Authority (FINRA)	*Insider Threats – Effective Controls and Practices*	2023
Carnegie Mellon University – Software Engineering Institute (SEI-CERT)	*Common Sense Guide to Mitigating Insider Threats – 7th Edition*	2022
Counter-Insider Threat Research and Practice (CITRAP)	*Seven (Science-Based) Commandments for Understanding and Countering Insider Threats*	2022
Cybersecurity and Infrastructure Security Agency (CISA)	*Insider Threat Mitigation Guide*	2020
Confederation of European Security Services (CoESS)	*Insider Threat Program Development Manual*	2019
Securities Industry and Financial Markets Association (SIFMA)	*Insider Threat Best Practice Guide – 2nd Edition*	2018
National Insider Threat Task Force (NITTF)	*Insider Threat Program Maturity Framework*	2018
Intelligence and National Security Alliance (INSA)	*A Preliminary Examination of Insider Threat Programs in the US Private Sector*	2013

Given the above 13 categories, the next step was to select the maturity levels into which the 568 recommendations from the 8 source documents would be placed. Because of the relevance to the financial services industry in which Morgan Stanley operates, the maturity levels defined in the Cybersecurity Assessment Tool created by the Federal Financial Institutions Examination Council (FFIEC) were selected (FFIEC, 2017). The FFIEC Cybersecurity Assessment Tool defines five levels: 'baseline,' 'evolving,' 'intermediate,' 'advanced,' and 'innovative.' The description of each level is summarized in Table 3.

As the final step on the creation of the maturity model, the 568 recommendations in the 13 categories were evaluated against the 5 FFIEC maturity levels and placed into tables in a document with one table per category. Table 4 shows an example table from the model for the vetting & onboarding category, and Table 5 shows an example table from the model for the monitoring category.

This evaluation process was carried out by multiple individuals to reduce the possibility of individual bias. Where duplicate recommendations existed because of the use of multiple source documents, a single recommendation that best captured the spirit of the collective guidance was selected. As a result, the number of recommendations was reduced from 568 to 127. It was noted that as the recommendations from each source document were processed, the number of novel requirements presented by each additional source document declined substantially. As such, it appeared that the point of saturation was reached – meaning that little to no new recommendations would be forthcoming if additional source documents were added (Saunders et al., 2018).

Table 3. FFIEC Maturity Levels

FFIEC Maturity Level	Characterized by
Innovative	*Driving innovation in people, process, and technology for the institution and the industry.*
Advanced	*Cybersecurity practices and analytics that are integrated across lines of business.*
Intermediate	*Detailed, formal processes.*
Evolving	*Additional formality of documented procedures and policies that are not already required.*
Baseline	*Minimum expectations required by law and regulations or recommended in supervisory guidance.*

Table 4. Vetting & Onboarding Category

Vetting & Onboarding	
FFIEC Maturity Level	Recommendations
Innovative	(a) Initial personnel screening is critical but not sufficient. Thou shalt focus on improving comprehensive, fair, and effective continuous vetting.[a] (b) For continuous employee evaluation (infinity vetting), having a statement in individual working contracts or employee handbooks, explicitly stating that this is a continuous process is essential.
Advanced	(c) During a merger or acquisition, perform background checks on all workforce members to be acquired, at a level commensurate with your organization's policies. (d) Verify the cloud service provider's hiring practices to ensure it conducts thorough background security investigations on its workforce members (e.g., operations staff, technical staff, janitorial staff) before they are hired. (e) Ensure that business partners have conducted background investigations on employees with access to the Firm's information systems or data.
Intermediate	(f) Employers should develop a clear policy towards the use of social media for recruitment purposes, in consultation with employees or their representatives. (g) Perform social media checks behind an "ethics wall," a neutral company employee (or outsourced) who does not make final hiring decisions, who filters out material using described standardized procedures, and who only provides job-related data to the final decision maker.
Evolving	(h) It is strongly recommended that potential employers contact referees to obtain references before offering a position to a candidate employee. (i) Establish personnel security vetting procedures commensurate with an individual's level of information system access.
Baseline	(j) Ensure that potential workforce members undergo a thorough background check, which, at a minimum, should include a criminal background check and credit check. (k) Personal data collected during the recruitment process should not be kept for more than the period allowed under EU GDPR laws.

[a] This playful language is present in the source document. The source document language was generally not altered when populating the model, in order to avoid potentially altering the meaning.

The final document containing the completed maturity model is twelve pages long, which is a tractable length and considerably shorter and thus more manageable than many of the source documents. A visual representation of the end-to-end process used to design the maturity model is shown in Figure 2.

Proceeding with the subsequent activities in the Design Science Research Methodology, the description of activity 4 ("Demonstration") is provided in section 3. The description of activity 5 ("Evaluation") is provided in sections 4 and 5. Activity 6 ("Communication") is delivered by this article as a whole.

3. Implementation

Distinct from the creation of the model itself is the assessment of the insider threat program within Morgan Stanley against the model. A streamlined approach with two phases was adopted for the assessment, with the goal of reducing the resource demand on both the insider threat team and the broader organization. This lightweight approach could potentially be expanded in the future, such as by acquiring material evidence of the status of each recommendation.

In the first phase, the insider threat team reviewed each recommendation in the model, and assigned a score (or label) of 'high,' 'medium,' or 'low' to depict the organization's current state. A high score represented close parity between the description in the recommendation and the current state within the organization. For example, if a recommendation said to "regularly test backup and recovery processes," and the backup and recovery processes inside the organization were indeed tested regularly, then the item would receive a 'high' score. Similarly, if the backup and recovery processes were tested but only infrequently, the item might receive a 'medium' score. A 'low' score

Table 5. Monitoring Category

Monitoring	
FFIEC Maturity Level	Recommendations
Innovative	(a) Install accurate sentiment, attitude and affect analysis for text-based data sources. (b) Use surveys of employees and data collection in order to catalog life events and changes as they occur. (c) Continually monitor employee's suitability to hold positions involving access to sensitive information by monitoring their digital footprint and activities on the internet. (d) Ensure that someone working in an insider risk capacity regularly monitors the SIEM solution to hunt for critical incidents that might not make it into the highest priority alerts.
Advanced	(e) Generate data visualizations to aid in analysis. (f) Prioritize analysis on data being gathered in the form of rule-based alerting, statistical anomaly detection or both, and prioritize these alerts. (g) Enhance the monitoring of workforce members who have an impending or ongoing personnel issue in accordance with the organization's policy and laws.
Intermediate	(h) Use appropriate tools to monitor network and workforce member activity for a period of time to establish a baseline of normal behaviors and trends. (i) Monitor the use of printers, copiers, scanners, and fax machines.
Evolving	(j) Establish a list of indicators that could tip investigators to suspicious behaviors.
Baseline	(k) Use a log correlation engine or SIEM (Security Information and Event Management) system to log, monitor, and audit employee actions. (l) Legal counsel should ensure that all monitoring activities are within the bounds of the law.

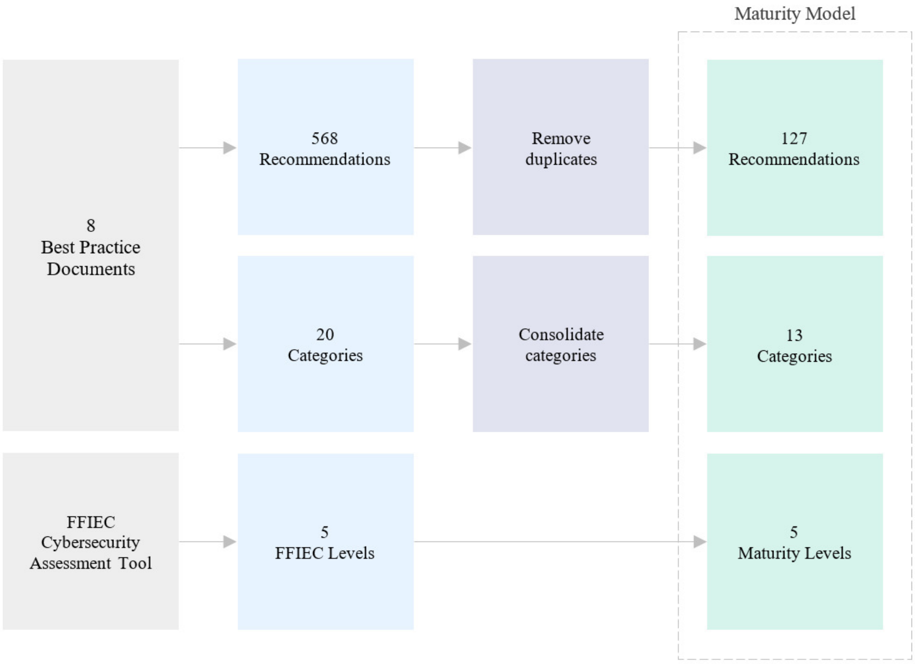

Figure 2. End-to-End Maturity Model Creation Process

would be assigned if no regular testing took place. This type of scoring is subjective, but was rooted in the knowledge and experience that the insider threat team held collectively regarding the organization. When performing this scoring, the spirit of each recommendation was engaged rather than the letter. For example, the recommendation in the 'baseline' level of the vetting & onboarding category in Table 4 refers specifically to EU GDPR (General Data Protection Regulation) laws, but the spirit of that recommendation was taken to mean compliance with all applicable laws.

In the second phase, subject matter experts (SMEs) within Morgan Stanley with specific knowledge of each category in the model then reviewed the scores assigned in the first phase. Depending on the availability of each SME, in some cases their scores were gathered in an interview setting with a member of the insider threat team present, otherwise the scores were provided by the SME via email. 10

scores (7.9%) were changed as a result of the review by the SMEs in the second phase.

A spreadsheet marking each of 127 recommendations as either high, medium, or low cannot easily convey the big picture of the maturity level of an insider threat program. In order to deliver that view, the visualization in Figure 3 was devised. In this visualization each colored box is a recommendation within the maturity model, with the color representing the score. The colored boxes are organized by row – which is the category, and by column – which is the maturity level. The categories and levels often contain a different quantity of recommendations because the underlying best practice documents do not offer the same amount of guidance on every topic, nor at every maturity level.

Figure 3 contains simulated data, but looking at the visualization immediately conveys the sense that the organization's insider threat program operates at the upper-intermediate level overall. In terms of the costs versus the benefits received from spending on security controls, an organization would likely prefer to be assessed at the upper-intermediate or lower-advanced level. If the assessed maturity level was lower, this would mean that the organization had not yet implemented security controls and practices that have proven to be effective (hence their placement in the lower levels of the maturity model). But if the assessed maturity level was higher, this would mean that that the organization was an early adopter of emerging technologies and practices where the cost-benefit of those has not yet proven. An organization can therefore use the results of the assessment to calibrate their program.

With regards to the level achieved within each individual category, there are a number of possible interpretations. A strict interpretation might select the highest level at which all scores at that level and below are scored 'high.' A more permissive interpretation might select the level at which the inflection point is most pronounced.

4. Validation

The visualization in Figure 3 was presented to several stakeholders within Morgan Stanley. Positive feedback was received regarding the ability of the visualization to communicate both the maturity model and the assessment in a manner that was accessible to both technical and non-technical audiences. The visualization also facilitated detailed discussions, as a separate key was used to explore specific recommendations in depth. While it is true that the subjective satisfaction of stakeholders is a qualitative measure of validation, the overall approach was judged to have successfully met the requirements described in the introduction. Namely, to create a framework that could enable the insider threat team to regularly assess the insider threat program, and to enable the reporting of that status to various stakeholders who possess different levels of technical knowledge regarding insider threats.

A second way in which the assessment against the model has proven to be useful is for planning. Recommendations scored low or medium within the model are obvious candidates for future investments. However, an insider threat team within a large organization is unlikely to have direct functional responsibility for all the categories in the model. For example, backups, asset management, and identity & access management are all likely to be distinct and separate functions within a large organization.

To address this aspect, the 127 recommendations in the model were systematically classified using a Responsibility Assignment Matrix (or RACI) (Project Management Institute, 2013, p. 262). 66 recommendations were identified across 8 categories for which the insider threat team was either responsible (R) or accountable (A). 39 recommendations scored high were then removed, since those recommendations represent less opportunity for improvement. The remaining 27 items were then sorted into three buckets.

The first bucket contains recommendations where the recommendation is currently not implemented, and there are no plans to do so. For example, for legal and compliance reasons an organization might choose not to follow the recommendation to monitor the internet footprint of workers. The contents of this bucket are periodically revisited in light of changing circumstances. The second bucket contains recommendations that are actively being worked on. For example, an organization might be in the process of working to improve security monitoring capabilities. The contents of this bucket are reviewed on an ongoing basis to ensure that forward progress is being made. The third bucket contains recommendations that are novel, meaning work items that could potentially be undertaken. For example, an organization might have implemented personnel screening but not yet implemented continuous vetting. The contents of this bucket are reviewed as possible work items when carrying out future planning. Table 6 provides an example of the output of this planning process with simulated data.

5. Limitations

A fundamental limitation is that useful recommendations might exist outside of the best practice documents that were used to populate the maturity model. Best practices are conventional wisdom, by definition. There might therefore be useful guidance – perhaps relating to a fast-moving topic such as technology – that has not yet been recognized and codified as a best practice. In that sense, a maturity model based on current best practices might equip an organization to "fight the last war." The faster the rate of innovation in the field of insider threats generally, the more pronounced this limitation becomes.

With regards to the assessment against the maturity model, in retrospect it would have been preferable to have the SMEs create their own scores, prior to being provided with the initial scores generated by the insider threat team. This would reduce the possibility of anchoring bias. A related matter is the fact that the scoring is entirely subjective from top to bottom, although the use of multiple knowledgeable parties to perform scoring should reduce response errors. Triangulation (the use of multiple research methods to study the same phenomenon) could be improved by using cognitive interviews to gather all SME

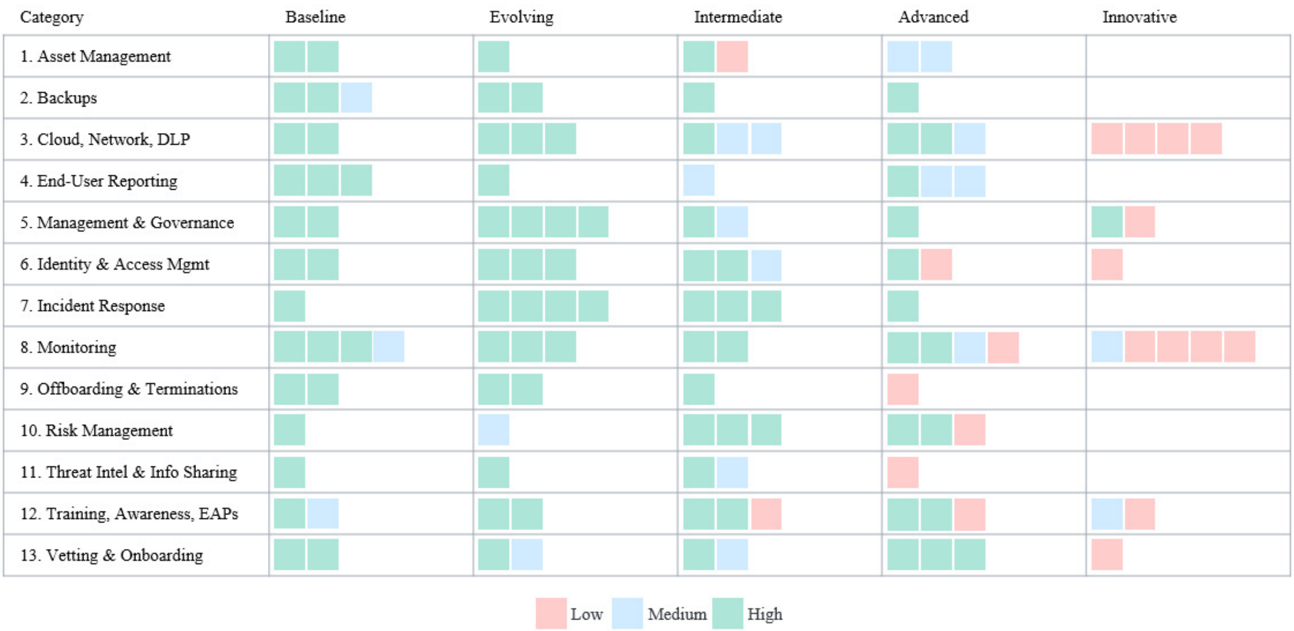

Category	Baseline	Evolving	Intermediate	Advanced	Innovative
1. Asset Management	🟩🟩	🟩	🟩🟥	🟦🟦	
2. Backups	🟩🟩🟦	🟩🟩	🟩	🟩	
3. Cloud, Network, DLP	🟩🟩	🟩🟩🟩	🟩🟦🟦	🟩🟩🟦	🟥🟥🟥🟥
4. End-User Reporting	🟩🟩🟩	🟩	🟦	🟩🟦🟦	
5. Management & Governance	🟩🟩	🟩🟩🟩🟩	🟩🟦	🟩	🟩🟥
6. Identity & Access Mgmt	🟩🟩	🟩🟩🟩	🟩🟩🟦	🟩🟥	🟥
7. Incident Response	🟩	🟩🟩🟩🟩	🟩🟩🟩	🟩	
8. Monitoring	🟩🟩🟩🟦	🟩🟩🟩	🟩🟩	🟩🟩🟦🟥	🟦🟥🟥🟥🟥
9. Offboarding & Terminations	🟩🟩	🟩🟩	🟩	🟥	
10. Risk Management	🟩	🟦	🟩🟩🟩	🟩🟩🟥	
11. Threat Intel & Info Sharing	🟩	🟩	🟩🟦	🟥	
12. Training, Awareness, EAPs	🟩🟦	🟩🟩	🟩🟩🟥	🟩🟩🟥	🟦🟥
13. Vetting & Onboarding	🟩🟩	🟩🟦	🟩🟦	🟩🟩🟩	🟥

🟥 Low 🟦 Medium 🟩 High

Figure 3. Maturity Model Assessment Findings (Simulated Data)

Table 6. Planning Exercise Output (Simulated Data)

Category	Ruled out (No plans to do)	Underway (Are doing)	Novel (Could do)
3. Cloud Security, Network Security, and DLP		Implement Cloud DLP	
4. End-User Reporting	Deliver indicators training to workforce	Deliver indicators training to managers	
8. Monitoring	Monitor the internet footprint of workers	Establish worker behavior baselines	Deploy sentiment analysis
9. Offboarding and Terminations		Establish a physical inventory system	
10. Risk Management	Prohibit the use of personal devices		Consider risk transfer options
13. Vetting & Onboarding		Implement continuous evaluation	

scores, where the interviewee describes their engagement with each question.

Lastly, creating a maturity model using best practice documents and then performing an assessment against that maturity model means that the assessment is in effect being performed against the best practice documents. And it is possible that the recommendations within the best practice documents simply do not fit the requirements of a particular organization. Reviewing the original 568 recommendations reveals that topics such as personnel vetting, security monitoring, and training & awareness received more attention from the authors of the source documents (meaning more content and more detailed content) when compared to topics such as backups and network security. This might partially reflect something of an unconscious bias towards topics that are traditionally the responsibility of insider threat teams. It is not a brute fact that the topics favored by the authors of the best practice documents will necessarily provide the most benefit to an organization when compared to alternatives. Best practices should not always be assumed to be "one size fits all," and each organization should customize its approach. To address this aspect, an assessment against a maturity model would ideally be paired with long-term data regarding operational outcomes such as the frequency and severity of insider incidents, and this would be a useful avenue for future research.

6. Conclusion

The insider problem has been described as "not one problem, but many" (Hunker & Probst, 2011). This aspect emerges in the recommendations provided by best practice documents. Those recommendations span a wide variety of topic areas including management, governance, metrics, network security, and backups, amongst others. Because of the quantity and breadth of these recommendations, and because there is no single definitive best practices document, organizations can face challenges with the task of assessing their insider threat programs. A maturity model is a useful approach for addressing these issues, because the format enables the recommendations from multiple best practice documents to be combined and reduced to a manageable size.

This article has demonstrated that the design and implementation of such a maturity model can be accomplished in a relatively straightforward manner by employing a structured methodology, and without requiring substantial resources. The specific model created by Morgan Stanley is one example of a model created by following that methodology. Other organizations can repeat the process to create models that best fit their needs – such as by selecting source documents that are appropriate to their specific sector. Future research could examine whether a general model can fit organizations in different sectors, or fit different kinds of organization such as government agencies and non-governmental organizations.

A visualization that summarizes the results of an assessment against the maturity model was also presented. In an organizational setting, the manner in which information is communicated can be as important as the information itself. At Morgan Stanley, the visualization has been found to be effective at conveying the results of an assessment to a wide range of audiences. Lastly, a maturity model can be used as an input to planning activities. An organization might set specific goals for the overall maturity level, or for the assessed levels within specific categories.

Maturity models are successfully employed in a variety of fields, including software development, process improvement, and project management. This article is a contribution towards the more prevalent use and study of maturity models for the insider problem.

Submitted: October 24, 2023 PDT, Accepted: September 30, 2024 PDT

This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CCBY-NC-ND-4.0). View this license's legal deed at https://creativecommons.org/licenses/by-nc-nd/4.0 and legal code at https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode for more information.

References

Alsowail, R. A., & Al-Shehari, T. (2022). Techniques and countermeasures for preventing insider threats. *PeerJ Computer Science*, *8*(e938). https://doi.org/10.7717/peerj-cs.938

CISA. (2022). *IR Mitigation Program Evaluation Tool One Pager*. Cybersecurity and Infrastructure Security Agency (CISA). https://www.cisa.gov/sites/default/files/2022-11/IR%20Mitigation%20Program%20Evaluation%20Tool_OnePager.pdf

FFIEC. (2017). *Cybersecurity assessment tool*. Federal Financial Institutions Examination Council (FFIEC). https://www.ffiec.gov/pdf/cybersecurity/ffiec_cat_may_2017.pdf

Glaser, B. G., & Strauss, A. (1967). *The discovery of grounded theory: Strategies for qualitative research*. Routledge.

Gottschalk, P. (2008). Maturity levels for interoperability in digital govrnment. *Government Information Quarterly*, *26*(1), 75–81. https://doi.org/10.1016/j.giq.2008.03.003

Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, *37*(2). https://doi.org/10.25300/MISQ/2013/37.2.01

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, *28*(1), 75–105. https://doi.org/10.2307/25148625

Hunker, J., & Probst, C. W. (2011). Insiders and insider threats: An overview of definitions and mitigation techniques. *Journal of Wireless Mobile Network, Ubiquitous Computing, and Dependable Applications*, *2*(1), 4–27. https://doi.org/10.22667/JOWUA.2011.03.31.004

Krombholz, K., Mayer, W., Schmiedecker, M., & Weippl, E. (2017, August 16–18). I have no idea what I'm doing: On the usability of deploying HTTPS. *26th Usenix Security Symposium*. https://doi.org/10.5555/3241189.3241293

NITTF. (2018a). *Insider threat program (InTP) maturity framework frequently asked questions (FAQs)*. National Insider Threat Task Force (NITTF). https://www.dni.gov/files/NCSC/documents/nittf/20181024_Framework_FAQs_FINAL.pdf

NITTF. (2018b). *Insider Threat Program Maturity Framework*. National Insider Threat Task Force (NITTF). https://www.dni.gov/files/NCSC/documents/nittf/20181024_NITTF_MaturityFramework_web.pdf

Paulk, M. C., Weber, C. V., Curtis, B., & Chrissis, M. B. (1993). *Capability maturity model for software (version 1.1)*. Software Engineering Institute, Carnegie Mellon University. https://doi.org/10.21236/ADA263403

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, *24*(3), 45–77. https://doi.org/10.2753/MIS0742-1222240302

Project Management Institute. (2013). *A guide to the project management body of knowledge (PMBOK guide)* (5th ed.).

Saunders, B., Sim, J., Kingstone, T., Baker, S., Waterfield, J., Bartlam, B., Burroughs, H., & Jinks, C. (2018). Saturation in qualitative research: Exploring its conceptualization and operationalization. *Quality & Quantity*, *52*(4), 1893–1907. https://doi.org/10.1007/s11135-017-0574-8

Wendler, R. (2012). The maturity of maturity model research: A systematic mapping study. *Information and Software Technology*, *54*(12), 1317–1339. https://doi.org/10.1016/j.infsof.2012.07.007