

Current Trends and Advances in Information Assurance Metrics

Nabil Seddigh¹, Peter Piedad¹, Ashraf Matrawy², Biswajit Nandy¹, John Lambadaris², Adam Hatfield³

Abstract—Information Technology (IT) networks are increasingly being used for transfer of mission critical data and information. One of the pressing issues is the development of uniform, deployable measures and metrics that can reliably depict the assurance levels associated with transfer of information over such networks. However, there exist significant challenges to be addressed before such an objective can be realized. This paper captures the current status of research efforts in the area of Information Assurance (IA) Metrics and analyzes their contribution to the field. The challenges of developing such metrics are explored. A novel definition of the term Information Assurance is proposed which takes into account Security, Quality of Service and Availability. Existing taxonomies for IA evaluation are examined and a new taxonomy is proposed which will facilitate future development of a single metric to depict the IA health of an IT network. The objective of this work is to provide a basis and motivation for further research in overcoming the challenges in this area.

Index Terms—Information Assurance Metrics, Security, Quality of Service (QoS), Availability, Rating, Network, Indicator

I. INTRODUCTION

RECENTLY, there has been some effort in utilizing quantitative indicators in a more systematic, coordinated fashion to capture the state of a particular Information Technology infrastructure. Such indicators are intended to reflect the "assurance" of the IT infrastructure to reliably transfer information. These indicators can be used to identify areas of the information infrastructure that require attention. They can also be used by an IT organization as a means of gauging the return on investment for IT Infrastructure equipment purchase.

Despite the existing work that is underway, there is currently no standard or widely accepted method of capturing and presenting the assurance levels associated with a particular IT infrastructure – this includes end-hosts, servers, applications, routers, firewalls and the network that allows

these systems to communicate. A comprehensive, standardized set of quantitative metrics and indicators would be useful in order to identify areas of the IT infrastructure that are candidates for enhancement and pro-actively implement improvements to the IT infrastructure. Analogous indicators exist in the area of financial/commodities markets. These indicators were developed through substantial research and data collection in the area of economics, financing etc.. Nevertheless, there is no indication for such an effort in the area of Information Systems. There have been a number of recent industry and research initiatives to develop standardized ratings that would reflect the Information Assurance associated with a specific product or product development process. For example, the Common Criteria (CC) [6] defines a set of rating levels (EAL 1 to EAL 7) which certify a particular vendor product's security rating. The ratings serve as a comparative platform that the consumer can utilize when comparing security products of various vendors. There has been good progress in developing product ratings, but little effort in developing a rating or indicator for a Network Infrastructure as a whole. Development of such an indicator remains very much a challenge for a number of reasons.

A first challenge is to arrive at a suitable definition of Information Assurance. The definition would then be used to identify key properties of the IT infrastructure that need to be assessed in order to develop a unified indicator or IA metric. Another challenge is that Information Assurance (IA) remains very much a "black art" as evidenced by various research papers, symposiums and panel discussions on this topic [19]. The assurance level of a particular piece of IT infrastructure is arrived at via subjective or intuitive considerations of the experts who are familiar with the protocols, architecture and systems utilized in the network.

This work presents a novel definition of Information Assurance and then uses that definition to develop a taxonomy of IA metric groups that would serve to gauge the IA rating of IT infrastructure. The intention at the present time is to develop quantifiable measures of Information Assurance that allow objective analysis and comparison of a particular IT infrastructure in relation to itself over time. Eventually, it would be desirable to develop standardized metric(s) that can be used to compare the IA rating of an IT infrastructure with itself over time or against other IT infrastructure.

This paper is further divided into five sections. Section 2 proposes a definition of Information Assurance that is based

¹Solana Networks, Suite 210, 120 Robertson Rd, Nepean, Ontario, Canada. K2M 1H7 {nseddigh, ppieda, bnandy}@solananetworks

²Carleton University, 1125 Colonel By Dr., Ottawa ON K1S 5B6 {ioannis, amatrawy}@sce.carleton.ca

³Department of Public Safety and Emergency Preparedness Canada, 2nd Floor Jackson Bldg., 122 Bank St., Ottawa, ON, Canada. K1A 0W6 adam.hatfield@psepc-sppcc.gc.ca.

on the Security, Quality of Service and Availability of the IT infrastructure. Section 3 discusses and evaluates existing work in IA metrics. This includes various measures, methods and approaches from the research community, government and industry. Section 4 presents and analyzes existing taxonomies for IA metrics. Section 5 outlines our proposed taxonomy for IA metrics. Lastly, conclusions are offered in Section 6.

II. DEFINITIONS OF INFORMATION ASSURANCE

The term "Information Assurance" (IA) is widely used in industry and academia, often with widely varying and divergent understanding of its meaning. We suggest three key elements that can constitute the bedrock of a comprehensive yet constrained definition of IA when applied against IT infrastructure. These three elements are Security, Quality of Service and Availability. Figure 1 provides an overview of the three key elements as well as provides examples of metrics or indicators for each element.

Below we describe why these three elements are important when considering a definition of IA:

- Security - Security can be considered as the ability of a system to protect information and system resources with respect to confidentiality, integrity and authentication. Security also includes vulnerability to active attacks by malicious users, viruses, denial-of-service, access lists, non-repudiation, privacy etc. In some parts of the Industry,

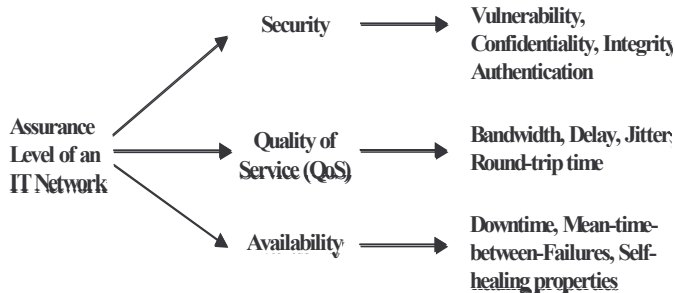


Fig. 1. Elements of a definition of "Information Infrastructure Assurance"

availability is sometimes considered part of security. However, we consider availability as a separate element as its definition could also include other factors apart from denial-of-service impacts on the network. E.g. Network Availability could be impacted by Traffic Engineering.

- Quality of Service (QoS) - QoS is typically used to reflect the perceived performance of a service in relation to the expected performance. For the purpose of this work, the performance properties used to describe an IT infrastructure's Quality of Service include consideration of such indicators as throughput (bandwidth), transit delay (latency), priority over other traffic or users, data loss rate etc. The QoS of a network is a strong contributing factor to its ability to assure the delivery of data from one end of the network to the other.

- Availability - Availability of the IT infrastructure is used to reflect its dependability or reliability. Thus, while QoS

reflects how a network or system performs, availability reflects whether the network or system can even offer the lowest of service qualities. We consider the following attributes and factors as contributing to define the availability of a system: infrastructure downtime, mean-time-between-failures, self-healing properties of a network/system and the ability of the network or system to operate under catastrophic disasters.

We recognize that a number of publications and related work consider only security-related issues as being relevant to IA. In such publications, QoS and Availability are not considered part of the Information Assurance definition. However, we argue that the ability of an IT network to process/deliver information from one user to another, is dependent on all three factors for the reasons mentioned above – hence the need to include them as part of the definition of IA.

Various definitions of IA can be found in the literature [12]. For the purposes of this work, we present the following definition of Information Assurance:

“ Information Assurance is the ability of a network or system to facilitate the timely transfer of information between 2 or more parties in an accurate and secure fashion.”

In a subsequent section of this paper, we will propose a taxonomy of IA metrics groups based on the above definition and three key elements. It will then be possible to use the taxonomy to develop IA metrics that allow organizations to determine their IT infrastructure's IA level based on the network and system's ability to meet their requirements for Security, QoS and Availability.

III. EXISTING STUDIES, TRENDS AND TOOLS

In reviewing the existing related work in the area of Information Assurance (IA) Metrics, we found that there were three distinct types of efforts underway in different environments: government, university and industry.

A. Government Research Agencies

In the late nineties, the United States DARPA (Defense Advanced Research Projects Agency) organization funded a number of exploratory research projects in the area of IA metrics. In [15], Skroch gave a presentation in which he analyzed some of the challenges facing researchers seeking to develop a set of IA metrics. He argued there is a clear need to develop an integrated environment for IA metrics by defining their purpose, meaning, units, range of values, and inherent taxonomies. He also claimed that the primary initial goal is to identify IA metrics that are measurable, testable and useful and then focus on moving as many of these metrics as possible towards the quantitative side of the scale. In [8], Evans and Bush present the result of their DARPA-funded work on applying a fundamental science-based approach to the problem of IA metrics. They propose the use of Kolmogorov

complexity as a basis for depicting the health of a security system. They argue that using a complexity-based method to depict the health of a system has the advantage that it is a fundamental property of information and can thus be applied without detailed knowledge of the system being analyzed. Further work is required to fully gauge the applicability of this approach to IT Networks.

In one of the more thoroughly documented initiatives, the US government's NIST (National Institute of Standards and Technology) has recently developed a guide to security metrics for information technology systems [17]. In this work, the authors provide guidance that an organization can utilize to measure the effectiveness of security-based controls and techniques. The authors propose a seven-step process that an organization can use to develop its own IA security metrics. The most interesting contribution of this work is to outline a sample set of metrics in each of 17 metric groups. The relationships of the metrics are based on a taxonomy developed in a previous work [16]. This work is further analyzed in section 4.

The MITRE Corporation is another government funded organization with previous research in the area of IA metrics. MITRE was the co-sponsor of a workshop on security metrics - WISSRR 2001 (Workshop on Information Security System Scoring and Ranking) [19]. Bodeau [4], Alger [1], Bicknell [3] and Connolly [7] of MITRE all submitted position papers at the workshop. In her paper [4], Bodeau highlighted some of the challenges for the information system security engineering community as they related to the area of IA metrics. The author argued that it is easier to specify an IA assessment function than to describe the corresponding assessment process. Further she suggested that it is important to clarify the IA assessment goals and scope. Examples of goals could be as an input to prediction of future behaviour or alternatively as an input to management decisions. Examples of scope are that it is possible to assess the IA strength of algorithms, specifications, mechanisms, functions, products and systems etc. The approach for a specific set of IA metrics may differ depending on the goals and scope as outlined above. Bodeau also questioned the viability of compressing many factors with complex relationships into a single figure-of-merit.

In his paper [1], Alger focused on the distinction between "measures" and "metrics". In [3], Bicknell described the approach taken by the United States IRS (Internal Revenue Service) to develop security metrics as a basis of conducting an evaluation of its cyber security program. The IRS established a taxonomy of metrics that consisted of 15 categories of security metrics. In [7], Connolly discusses the need to develop IA metrics that focus on operational readiness. She argues that there has been good progress in the area of developing metrics to assess the IA strength of products or systems but limited work in ways of assessing the organizational and operational capacity to react to IA emergencies.

There appear to be a number of initiatives in the area of

"metrics and models" sponsored under the I3P (Institute for Information Infrastructure Protection) [10] umbrella. To date, we have not found any publicly available material summarizing the findings of the researchers in the above area.

The AFRL (United States Air Force Research Lab) [2] appears to have recently funded work by Decision Science Associates, Trusted Computer Associates and Lockheed Martin in the area of IA metrics. The project proposes to use Bayesian statistics and Multi-Attribute Analysis as the basis for developing a prototype IA metrics decision support tool. We have not yet found any publicly available material detailing the results of this project.

B. University Research

The Information Assurance group at Mississippi State University has been conducting research in the area of IA metrics for a number of years. Their work has been captured in a number of publications including [20] and [21]. In [20], Vaughn et al summarize the findings and observations of the WISSRR 2001 workshop [19] on security metrics. Among the findings that they highlighted was the different uses that government and commercial sectors have for IA metrics. The workshop participants came to the conclusion that there does not exist a single set of IA metrics that are applicable across various systems. They also observed that the number of quantitative IA metrics were at present in short supply.

In [21], Vaughn et al propose a taxonomy of IA metrics that can serve to provide a framework for development of metrics and their inclusion in an evaluation or assessment framework. The authors suggest that there are five ways of viewing specific IA metrics. This includes: (i) objective/subjective (ii) quantitative/qualitative (iii) static/dynamic (iv) absolute/relative (v) direct/indirect. The most interesting part of this paper would appear to be the sections that focus on the proposed IA metrics taxonomy as this provides a basis for commentary and future work by other researchers. We review and analyze the authors proposed taxonomy in section 4.

The security group (NISLab) at Gjøvik University College in Norway is also working in the area of Information Assurance Metrics. Some of the projects underway include [11] [13] [14] [18].

C. Industry - Common Criteria

There have been a number of collaborative industrial efforts to define standards for the rating and assessment of security systems. To date, most such efforts have resulted in standards that assist in rating specific vendor products. They do not yet have standards which rate complicated systems such as an IT network. This is to be expected, as it is difficult enough to gain widespread acceptance of a standard means of evaluating a set of products. We expect that in the future such efforts may focus on standard means of evaluating and rating the IA of IT networks as a whole.

One such industrial effort which is slowly gaining momentum in terms of practical deployment is the Common Criteria (CC) [5] initiative. The CC currently provides a

means by which a vendor product can be evaluated for certain security capabilities. The CC developed through a combination of several other national standards for security, including TCSEC (Trusted Computer System Evaluation Criteria), ITSEC (Information Technology Security Evaluation Criteria), CTCPEC (Canadian Trusted Computer Product Evaluation Criteria), and FC (Federal Criteria for Information Technology Security). In particular, the CC has defined EAL (Evaluation Assurance Levels) certification that provides the consumer with a basis on which to compare security ratings of various vendor products. Nortel Networks Alteon switch, Entrust's PKI product, Microsoft Windows 2000 and Sun's Solaris 8.0 are among the products that have received EAL certification. Though use of CC is still not widespread, it appears to be gaining momentum. At present it is unclear whether CC will eventually achieve widespread industry rollout and adoption.

D. Industry – Tools

Multiple inputs are required in order to derive a set of indicators that capture the IA posture of an IT network. Due to the complex nature of an IT network and the many inputs, it would be useful to have a single software-based data fusion tool that would serve to receive the inputs, undertake the transformations and arrive at the quantitative indicator that would capture the IA health of the network. We have not found many examples of such a tool and certainly, there are no tools available today that can undertake the scope of IA evaluation that we are proposing. However, there is some encouraging work underway in this area. In [9], Fox et al describe a prototype next generation IA tool that correlates and fuses results from multiple IA tools into a single assessment of a network's security posture.

Though there is limited work underway in developing data fusion tools, there has been significant progress in developing automated tools to routinely and systematically undertake measurements and tests on an IT network. Such tools would provide the input to the data fusion tools.

With regard to such automated tools, there are a wide variety of security related assessment tools. These tools probe a network for security weakness, reporting what security hazards are present. It should be mentioned that there are still no standards in the area of what type of tests to be conducted though there are many types of tests that are commonly used across a variety of tools. There are also a number of tools available that measure another aspect of IA as we define it - the QoS performance of a network. The tools test such things as QoS performance metrics including Internet delay, round trip time, packet drops, and application quality from the users perspective. There are fewer tools available to measure the third aspect of IA as we defined it - availability and reliability. Availability of an IT network is in general, harder to measure than other aspects of the network. If a network seems to be always available to its users and has never failed, how would one classify its level of availability and reliability? The tools focusing on these areas cover issues such as network segment

outages, associated services affected that outage, associated services that were not impacted by that outage, data storage problems, bit error rates (in case links of the network are satellite or wireless), and so forth.

In general, due to the many inputs that need to be taken into consideration to ascertain the IA posture of an IT network, it is likely that a software tool of some sort is valuable to receive the inputs and arrange them into a coherent, limited set of quantitative health indicators.

IV. EXISTING TAXONOMIES FOR INFORMATION ASSURANCE METRICS

We suggest that development of IA metrics that depict the IA status of an IT network can best be done in the context of a taxonomy. The taxonomy provides logical groupings for IA metrics and illustrates the relationships between these groupings. In this section, we review and analyze three existing taxonomies for IA metrics that we have found to be the most advanced and comprehensive of the ones that are publicly available. For ease of reference in subsequent sections, we use the terms WISSRR, Vaughn and NIST to describe each taxonomy.

A. The WISSRR Taxonomy

The WISSRR 2001 workshop [19] provided a key venue for researchers in the area of IA metrics to present position papers and hold detailed discussions on a variety of issues related to IA. It is apparent from the conference proceedings that there is widespread and divergent understanding with regards to many issues related to IA. The common term IS* (Information Security) was adopted in this workshop to avoid disagreements on terminology. The asterisk (*) was intended to cover such terms as metric, score, rating, rank, assessment result etc. In the workshop, an IS* was defined as a “*value, selected from a partially ordered set by some assessment procedure, that represents an IS-related quality of some object of concern. It provides, or is used to create, a description, prediction, or comparison, with some degree of confidence.*”

The workshop did not propose any specific taxonomy to order the relationships or grouping of IA metrics. However, presentations and talks were organized into three tracks based on the interests of participants: the technical track, the organizational track, and the operational track. In general, there would seem to be an intuitive understanding among workshop participants that the three themes would provide a useful basis around which to organize a taxonomy of IA metrics.

B. Vaughn et al Taxonomy

In [21], Vaughn et al. propose a taxonomy for IA metrics as depicted in Figure 2. They divide their taxonomy into two distinct categories of metrics. The first category (organizational security) of metrics aims at assessing the “IA posture” of an organization while the second category of metrics aims at assessing the IA capabilities of a product or

system (Technical Target of Assessment - TTOA).

The first category of metrics are those for organizational security that are used to provide feedback to improve the IA status of the organization. The authors further classify metrics for organizational security into four sub-categories based on

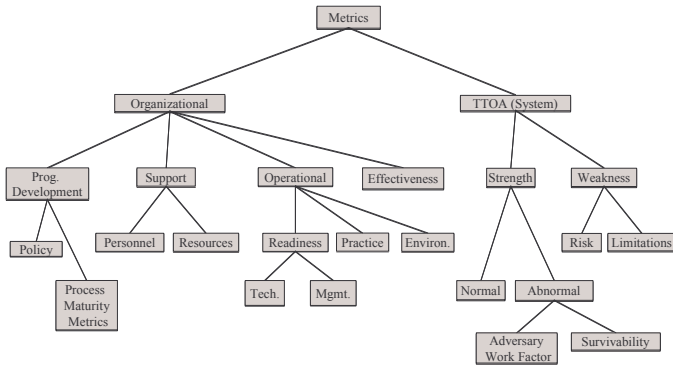


Fig. 2. Taxonomy by Vaughn et al. [21]

what they measure: (i) IA Program Developmental Metrics measure the organization's strength in IA policies and process for developing security-capable products (ii) Support Metrics measure the organization's resources committed towards security and the quality of personnel working on security (iii) Operational Metrics measure operational readiness for security incidents and operational security-related practices (iv) Effectiveness Metrics reflect how effective the organization's practices actually are by counting various security incidents.

The authors have made a valuable contribution to the field of IA metrics by putting forward a taxonomy. Further work may be required to refine the taxonomy in order to make it applicable to an IT organization. Some of the metrics categories seemed more applicable to an organization responsible for developing a hardware or software product while other categories seemed applicable to an IT organization.

The second category of metrics are those referred to as a Technical Target Of Assessment (TTOA). These metrics are intended to measure how much a technical object, system or product is capable of providing security in terms of protection, detection and response. The metrics are often used in comparing or differentiating between alternative and competing TTOA, e.g. the EAL ratings of the Common Criteria. The authors further divide metrics for TTOAs in two sub-categories - metrics for measuring TTOA's strengths and metrics for measuring a TTOA's weaknesses.

There would appear to be limited debate with regard to the second category of the taxonomy. It is clearly intended to provide a means of evaluating a particular vendor product. At this point it is not entirely clear how it can be applied against typical IT networks with multiple different products from a diversity of vendors. We further revisit this issue in section 5.

C. The NIST (National Institute for Standards and Technology) Taxonomy

In [16] and [17], NIST presents a comprehensive taxonomy or set of categories for IA metrics. A diagram outlining NIST's taxonomy is presented in Figure 3. The work is very complete and each category is accompanied by detailed examples of metrics for that category along with a template of attributes for each metric. These attributes include some questions that are asked to help assign a value to the metric. Among the attributes is a formula on how to calculate

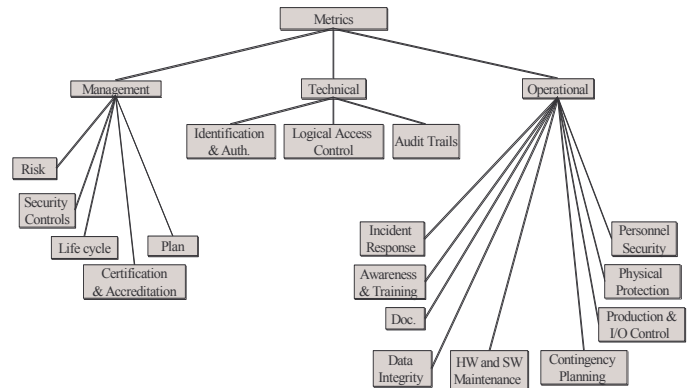


Fig. 3. NIST Taxonomy

a value for the metric from the answers to these questions. We found that going through the exercise of developing sample metrics and filling out the associated per-metric template helped in the task of refining the taxonomy to be of greater practical value.

In all, NIST proposes 17 categories of security metrics groups. The taxonomy divides the metric space into three different categories (management, technical, and operational) - the division that appeared to intuitively emerge from the WISSRR workshop. We note however that all the categories appear intended to be applied against an organization. In other words, the "technical" category of metrics has the objective of assessing the level of technical security controls put in place by an IT organization. Unlike the Vaughn taxonomy's TTOA category, NIST's technical category of metrics is not intended to apply against a specific product. We also note that the NIST taxonomy did not specifically indicate the criteria by which a sub-category was placed under organizational, technical and operational. The reason behind this argument is that in a few cases, we wondered if a particular metric sub-category would better fit under a different category than the one illustrated in




Fig. 4. Proposed IA Metrics Taxonomy for IT Networks

the Figure 3.

V. PROPOSED TAXONOMY TO EVALUATE IA OF IT NETWORKS

In this section, we propose a new taxonomy for IA metrics that can be applied to assess an IT network. In future work, we intend to investigate the viability of utilizing the taxonomy to combine the various metrics categories into two or three indicators that depict the IA health of an entire IT network. We believe that there is much that can be learnt from the actual exercise of developing the taxonomy and associated framework, complete with sample metrics and methods for generating the metrics, and a tool that can fuse all the inputs into a limited set of IA health indicators.

The proposed taxonomy is based on the definition of IA presented in section 2. We believe that the evaluation of an IT network's IA health should include an assessment of its network Availability and QoS aspects in addition to aspects of information security. Figure 4 captures the highest levels of the IA taxonomy tree.

We divide the metrics space into three categories: Security, QoS, and Availability. Under each of these three we consider the different technical, organizational, and operational aspects. Each of these sub-categories are further divided into another level of this tree – more details can be found in sections 5A, 5B, 5C and 5D.

The top level of the taxonomy tree has been deliberately organized as in Figure 4 in order to allow the assessment of a subset of the IA aspects. For example, an organization that is only interested in security can utilize only the security sub-tree. Furthermore, the organization can exercise finer granularity control over the measurement efforts. If the organization only wishes to focus on the operational aspects of security, it can do so without having to undertake measurements related to QoS or even technical security.

A. Sub-Categories of the Taxonomy Tree

Our proposed definitions and use of the three sub-categories in the second level of the IA taxonomy tree differs slightly from those used in the NIST, WISSRR and Vaughn taxonomies. Below we present our intended meanings for each sub-category:

- **Organizational Management:** this group of metrics evaluates an IT organization's emphasis on IA (in terms of goals and policies) and its commitment to IA (in terms of allocated resources).

- **Technical Elements:** this group of metrics evaluates how the technical components of an IA network are capable of providing IA. A subset of this group is static which provides a rating of the technical components capabilities in relation to IA. The remaining metrics under this group are dynamic as they are measured at different points in time after the IT infrastructure is deployed.

- **Operational Practices:** this group of metrics evaluates the operations of an IT organization in terms of complying with the IA goals and policies set by that organization.

As we reviewed the NIST and Vaughn taxonomies, we recognized that in some cases, it is possible to argue that certain metrics could go under more than one category or sub-category of the taxonomy metric tree. Though it is possible to set out objective criteria that determines which part of the taxonomy tree a metric falls under, we do not feel that this is one of the most pressing challenges to be addressed at the present time. However, it is an aspect of future work that could be investigated. As mentioned previously, each of the three categories below Security, QoS and Availability is further broken up into metric sub-categories. The full taxonomy for security is present but due to space limitations, only the Technical Elements sub-tree of the other categories will be discussed. The full IA taxonomy along with sample metrics is available in [22]. We focus on the technical elements sub-tree because the organizational and operational taxonomies are similar across all three groups (Security, QoS, Availability), with Security being the more involved as indicative of industry's focus on that aspect of our definition of Information Assurance Metrics.



Fig. 5. (a) Security Metric Taxonomy: Organizational Management




Fig. 5. (b) Security Metric Taxonomy: Technical Elements

B. Security – Taxonomy

These security metrics in Figure 5 attempt to quantify the process of evaluating each sub-category of characteristics for the IT network – namely Organizational Management, Technical Elements, and Operational Practices. The figures depict the parts of the security taxonomy and captures how each of the metrics sub-categories are organized.

The Organizational Management Program Development metrics evaluate the development of the organization's security program. The Security Plan checks whether the




Fig. 5. (c) Security Metric Taxonomy: Operational Practice

organization has a security plan in place and how often it is updated. Risk Management checks whether risk is periodically assessed and how many departments have risk assessment procedures in place. Policy checks the organization's security policy by considering the following subgroups. Review of Security Controls checks the levels of personnel access controls to different resources and how often security controls are reviewed. Personnel Security checks the required level of personnel background checking before hiring them. This issue will have different levels of relevance according to the type of infrastructure and the organization's activities. Certification checks the certification level requirements set by the organization. This includes the certification required for the personnel's technical capabilities as well as the certification required.

The Resources sub-category deals with the quality and the quantity of the resources allocated to the organization's security, both human and technical. Human metrics evaluate the resources allocated for the developing the organization's security human resources. The Technical metrics evaluate the resources allocated towards the organization's security technical resources (software, hardware, and networking).

The Technical Element Product Rating sub-category is intended to assess the security rating of the products used in the infrastructure. Software metrics assesses the security ratings of the software products used on the network. Network metrics assesses the security ratings of networking equipment such as routers or switches either based on recognized standards or a pre-established set of required features. Hardware metrics check the security ratings of the hardware systems such as servers and desktops.

The Incident Statistics sub-category is intended to assess security incident statistics collected from the deployed systems. These statistics should be indicative of the effectiveness of the technical elements.

The Security Tests sub-category is intended to assess how the IT network copes with a variety of security-related tests. The Penetration Test sub-category is intended to actively test the capability of the network to keep out malicious users who do not have the requisite privileges to access or control parts of the IT network. The Service Denial Test sub-category is intended to assess the level of service provided by a network during active tests to bring down parts of the it due to security problems.

The Operational Practice Susceptibility group of metrics

assesses the infrastructure security vulnerability due to its existence in a certain environment. Effectiveness checks the effectiveness of the security operational practices.

Incident Response checks the capabilities of responding to security incidents. Past Records checks the archiving process for past records of operational practices and measurements. It also evaluates the examination of system logs to check the system effectiveness. Maintenance checks the effectiveness of the security maintenance of software, hardware, and networking equipment.

Technical Readiness checks the technical readiness of the security operations. Documentation checks and evaluates documentation of the security operations. Data Integrity checks and evaluates regulations for insuring the integrity of the data. Contingency Planning checks and evaluates the contingency plans of the security operations.

C. QoS Technical




Fig. 6. QoS Technical Elements Metric Taxonomy

Quality of Service Technical Elements metrics are intended to assess QoS performance characteristics of the network and its associated elements. These metrics evaluate the taxonomy leaf properties in a quantitative fashion to facilitate unambiguous and clear indications of QoS network performance. Figure 6 depicts this part of the proposed IA taxonomy sub-tree.

The Product Capabilities sub-category is intended to assess whether products deployed in the network have the requisite features to allow the identified QoS characteristics to be achieved. In this regard, QoS/SLA Network Management metrics grade the QoS features of the network management software. These include such properties as configuration and monitoring, policy enforcement, and specification of service level agreement (SLA) support. QoS Building Blocks metrics ensure that the network devices include various required QoS building blocks such as multiple queues, priority scheduling, traffic conditioning and policing, differentiated services support, and buffer management.

The Network Capabilities sub-category is intended to assess whether the network as an end-to-end entity contains the requisite capabilities to provide QoS. A network that contains both QoS capable and non-QoS capable devices may only be able to offer a certain level of QoS.

The QoS Tests sub-category is intended to assess the extent to which the network manifests the desired level of QoS.

Routine Operation metrics assess the characteristics of the network under normal operating traffic loads. They are intended to reflect the extent to which the actual QoS behavior compares against the required QoS characteristics. These metrics include such things as amount of data lost through congestion, end-to-end delay etc. Stress Tests metrics are




Fig. 7. Availability Technical Elements Metric Taxonomy

intended to assess the network's capability to handle situations when it is stressed outside of normal operating conditions. The metrics are based on executed tests that help identify weak spots in the network.

D. Availability – Technical

The Network Availability technical metrics are intended to assess the availability characteristics of the network. Network Availability technical elements part of the IA taxonomy subtree is shown in Figure 7.

The Redundancy sub-category is intended to assess the degree of Network Availability that is achievable with the current network infrastructure. The Product metrics assess the product properties of each network device such that the required network availability is achievable. These properties include such features as high reliability, load balancing, fault detection, and protection switching. Network metrics determine whether the specified degree of network availability for the network as a whole is achievable. A network may have devices that can do protection switching, but the network may not be designed such that they can be used.

The Availability Tests sub-category assesses the level of Network Availability attained under active tests. Routine Operation metrics are intended to reflect the characteristics of the network under normal operating traffic loads. These metrics compare the actual Network Availability versus the desired Network Availability. These metrics include downtime (frequency, length, and severity), uptime, and end-user feedback. Stress Test metrics are intended to reflect availability characteristics of the network under active stress testing. i.e. when the network is stressed outside of normal operating conditions. Tests are executed to help identify weak spots in the network and ensure that the availability is guaranteed.

VI. CONCLUSION

Our survey in the area of IA metrics indicated that it is an emerging area which should be given considerable attention in the coming years as our reliance on IT infrastructure continues

to grow and permeate every aspect of our lives. Most of the research undertaken so far has focused on the security aspect of IA metrics.

This work evaluated and analyzed three previously proposed specific taxonomies and sets of IA metrics. This was done with a view towards determining if the prior taxonomies could form a basis for our work, which was to apply them to the assessment of IT networks. We found that the taxonomies had their strengths and weaknesses, but overall none were sufficient to be used to assess the IA posture of an IT network.

This work has proposed a novel definition of “Information Assurance” that was based on three key attributes related to IT infrastructure: (i) security (ii) quality of service and (iii) availability. We argued that the ability of an IT network to deliver information from one user to another is dependent on all three factors.

Therefore, we proposed a new taxonomy and category of metrics that could be utilized for the specific purpose of capturing the Information Assurance rating of an IT network. This taxonomy divided the Information Assurance Metrics aspects of a network into three categories: Security, QoS, and Availability. Each of these categories in turn was subdivided into three sub-categories: Organizational measures, Technical Elements, and Operational Practices.

The area of Information Assurance evaluation and assessment is quite complex. In future work, we propose to develop an IA framework along with sample metrics and measurement methodology, preferably integrated in an expert system tool. Through this framework we propose to fuse multiple inputs to arrive at a limited set of two or three indicators that can mirror the IA health of the IT network. We are currently unaware of any such efforts in the area of Information System Assurance.

REFERENCES

- [1] Alger J, “On Assurance, Measures and Metrics: Definitions and Approaches”, Proceedings of WISSR 2001, May 21-23, 2001, Virginia
- [2] Air Force Research Labs, www.rl.af.mil/tech/programs/ia/ia12.html
- [3] Bicknell P, “Security Assertions, Criteria and Metrics Developed for the IRS”, MITRE Technical Report, May 2001
- [4] Bodeau Deborah, “Information Assurance Assessment: Lessons Learned and Challenges”, Proceedings of WISSR 2001, Williamsburg, Va, May 2001
- [5] Common Criteria. <http://csrc.nist.gov/cc/index.html>
- [6] Common Classification Criteria <http://www.commoncriteria.org>
- [7] Connolly J, “Information Assurance Assessment: Lessons-Learned and Challenges”, Proceedings of WISSR 2001, May 21-23, 2001, Virginia
- [8] Evans S, Bush S and Hershey J, “Information Assurance through Kolmogorov Complexity”, DARPA Information Survivability Conference and Exposition II (DISCEX-II-2001) 12-14 June 2001, Anaheim, California
- [9] Fox K, Henning R, Farrell J, and Vaughn R, “A Prototype Next Generation Information Assurance Tool - A DataFusion Model for Information Systems Defense”, White Paper, Harris Corporation. SCI (Systemics, Cybernetics and Informatics) 2003.
- [10] Institute for Information Infrastructure Protection, www.thei3p.org/

- [11] Mathisen J, "Measuring the effect of an information security awareness drive", MSc Project Plan, Gjøvik University College, Norway, December 2003
- [12] McKnight W, "What is Information Assurance", CROSSTALK, The Journal of Defense Software Engineering, July 2002.
- [13] Nygard Arne R., "Security Metrics in SCADA Networks", Thesis Proposal. 2003.
- [14] Nygard Arne R., "Metrics for Software Resistance Against Trojan Horse Attacks", Project Proposal, Norwegian Information Security laboratory. 2003
- [15] Skroch M, McHugh J and Williams JM, "Information Assurance Metrics: Prophecy, Process or Pipedream", Panel Workshop, National Information Systems Security Conference (NISSC 2000), Baltimore, October 2000
- [16] Swanson M, "Security Metrics Guide for Information Technology Systems", National Institute of Standards and Technology Special Publication #800-26. November 2001.
- [17] Swanson M, Nadya B, Sabato J, Hash J and Graffo L, "Security Metrics Guide for Information Technology Systems", National Institute of Standards and Technology Special Publication #800-55. July 2003.
- [18] Wold G, "Is use of security metrics expedient to measure performance of an implemented organizational security policy", MSc Project Plan, Gjøvik University College, Norway, December 2003
- [19] Workshop on Information, Security System Scoring and Ranking (WISSSR, 2001) Information System Security Attribute Quantification or Ordering (Commonly but improperly known as Security Metrics) – Workshop Proceedings - May 21-23, 2001, Williamsburg, VA.
- [20] Vaughn, Rayford, Ambareen Siraj, and David Dampier, "Information Security System Rating and Ranking," CROSSTALK, The Journal of Defense Software Engineering, May 2002, pp. 30-32.
- [21] Vaughn R, Henning R, and Siraj A, "Information Assurance Measures and Metrics: State of Practice and Proposed Taxonomy", Proceedings of 36th Hawaii International Conference on System Sciences (HICSS 03). 2003.
- [22] Nandy B, Pineda P, Seddigh N, Lambadaris J, Matrawy A, Hatfield A. "Information Assurance Metrics". Technical Report prepared for the Canadian Federal Government Department Office of Critical Infrastructure Protection and Emergency Preparedness (OCIPEP). March 2004.