SSC20-P1-03

Scoring Trust Across Hybrid-Space: A Quantitative Framework Designed to Calculate Cybersecurity Ratings, Measures, and Metrics to Inform a Trust Score

Sean Kinser, Pete de Graaf, Matthew Stein, Frank Hughey, Rob Roller
The MITRE Corporation
1155 Academy Park Loop, Colorado Springs, CO 80910, 719-572-8400
skinser@mitre.org

David Voss, Amanda Salmoiraghi
United States Space Force
7250 Getting Heights, Colorado Springs, CO 80916

DISTRIBUTION A: Approved for public release; distribution unlimited.
Public Release Case Number 20-1454.

ABSTRACT

Multi-national commercial space offerings continue to flourish in a domain once singularly dominated by governments. The same governments are now seeking methods to leverage commercial space in a bid to realize cost savings, leverage additional capacity, and augment missions. The outcome is a hybrid environment of commercial, civil, defense industrial base, and other government space vehicles, networks, ground segments, and data. Within this digitally reliant, hybrid-space environment, cybersecurity remains a key focal point of systems-of-systems design and implementation. Owners and stakeholders seek assurance in the command and control of their assets while customers desire a level of trust in the data or service being provided. We present a method to quantitatively evaluate the overall trust of space-based services as related to the core cybersecurity principles of confidentiality, availability, and integrity. This method considers both qualitative and quantitative assessments generated through phases categorized as Compliance Assessment, Performance of day to day cybersecurity operations (cyber-hygiene), and Incident Response. The inputs of these phases are used to generate a quantitative metric that indicates an organization's ability to securely deliver data and services. This metric is referred to as the Architecture Score Index.

INTRODUCTION

Space is a domain no longer reserved for the world's leading nations as commercial activities have experienced a renaissance in launch, imagery, and other traditionally agency-level, final frontier pursuits. In 2018, the Space Foundation reported that the commercial space industry accounted for 79% of the $414.8 billion in revenue generated globally by space-related activities.[1] Recognizing the clear advantages of additional capacity and capability offered by commercial space, governments are now posturing to accept and integrate these capabilities into their mission architectures for national defense, space exploration, and research.

For military and government leaders, outsourcing the production of space-based information sources or communications infrastructures that are key to national defense is a tough pill to swallow. To convince leaders that a shift to a hybrid-space architecture is in the nation's best interest, a distinguishable level of trust must be achieved, demonstrated, and accepted. The added benefits of resilience, capacity, and speed must exceed the uncertainties associated with trust of a service or data produced outside of their control. The assessment of trust in opportunistic, "good enough" commercial systems should be done in a quantitative, non-subjective manner. The same assessment could be applied to government, civil, or other space mission architectures, providing a clear metric of the overall trust in the data and services received. We present a method of producing such a metric, hereafter referred to as the Architecture Score Index (ASI).

Defining Trust

Trust is a complex topic involving judgments and opinions in the reliability, truth, competence, etc. of a person or a thing. A popular example of trust in action was conceived by Morton Deutsch in 1962 which states: *(a) an individual is confronted with an ambiguous path, a path that can lead to an event perceived to be beneficial or to an event perceived to be harmful; (b) he perceives that the occurrence of these events is contingent on the behavior of another person; and (c) he perceives the strength of a harmful event to be greater than the strength of a beneficial event. If he chooses to take an ambiguous path with such properties, he makes a trusting choice; else he makes a distrustful choice.*[2]

This definition provides a foundation, but as with many other examples of trust, relies too heavily on individual belief, leaving room for subjectivity. Further work by Grandison and Sloman[3] provides a definition of trust more suitable to the ASI. This definition categorizes trust as *confidence in the competence of an entity to act dependably, securely, and reliably within a specified context.*

In the context of the ASI framework, trust is determined as an outcome of cybersecurity compliance, performance considerations (patch cadence and incident response), and calculations, as illustrated in Figure 1. Trust does not exist in this setting without participation from the seller (trustee) and buyer (trustor), enabling an evidence-based version of trust, applied within a specified context.

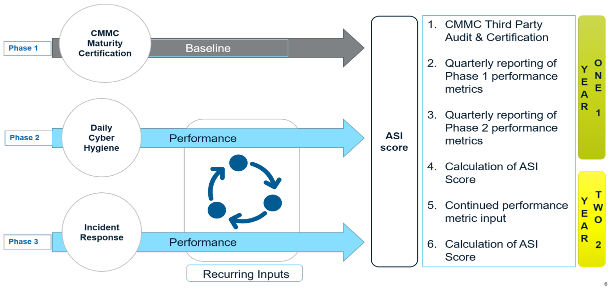
Phase 1
CMMC Maturity Certification
Baseline
Phase 2
Daily Cyber Hygiene
Performance
Phase 3
Incident Response
Performance
Recurring Inputs
ASI score
1. CMMC Third Party Audit & Certification
2. Quarterly reporting of Phase 1 performance metrics
3. Quarterly reporting of Phase 2 performance metrics
4. Calculation of ASI Score
5. Continued performance metric input
6. Calculation of ASI Score
YEAR ONE 1
YEAR TWO 2
6

Figure 1: ASI Framework: a Combination of CMMC, Cyber Hygiene, and Incident Response.

Assumptions

We base the ASI assessment on three principal assumptions:

1. Trust is an enabler and a required element of success between trustor and trustee, whether government, civil, or commercial. Trust, through verification, facilitates the trustor's relinquishment of control, establishing an understood acceptance of quantifiable risk.

2. For the ASI to be calculated, stakeholders must participate by submitting cybersecurity standards compliance information, quarterly performance data, and additional reputation information. Participation in this framework will serve both the reporting organization as well as the end user of the ASI score by creating a mutually beneficial marketplace for hybrid-space products and their consumption.

3. The time standards used in calculating patch cadence and other metrics are based on reasonable industry best-practices and are consistent for every calculation. Each system will be assessed against the same criteria for patching and incident response timelines.

FRAMEWORK APPROACH

Our approach considers three cybersecurity-focused phases with quantitative outcomes from each used to calculate the ASI. Phase 1 leverages the Undersecretary of Defense Acquisition and Sustainment (OUSD (A&S)) Office's Cyber Maturity Model Certification (CMMC)[7] to provide a cybersecurity maturity baseline for an organization. Phase 2 considers reported cybersecurity performance metrics that identify an organization's skill and consistency at implementing daily cyber-hygiene[8] tasks. Phase 3 focuses on incident response metrics to understand how well and how rapidly a company addresses cybersecurity incidents.

All phase criteria are quantified, weighted, and combined to produce the ASI value used as a variable measure of trust for ingestion into architecture modeling tools such as Acropolis[4]. The Acropolis software application provides a quantitative evaluation of the overall performance of any space architecture from traditional monolithic systems to fully hybrid systems-of-systems that blend contributions from multiple providers with distinct capabilities. The results can inform operational, planning, and acquisition decisions for both current and future space missions. Including the ASI allows such tools to trade risk and performance.

Phase 1: Cybersecurity Maturity Model Certification

A certified level of compliance, as the result of third-party auditing, provides an initial baseline and is a key indicator of system security investment and preparedness. Within the CMMC, a higher compliance level indicates a greater level of preparedness and an ability to reliably transport sensitive data while providing a viable defense against an advanced persistent threat (APT). The potential compliance levels range from one to five, with five garnering the highest possible rating of trust in this phase.

Phase 2: Cyber Hygiene

Phase 2 accepts multiple key indicators that reflect performance of daily cyber hygiene tasks. Our current approach uses patch cadence, the periodic cycle over which patch installation protocols are executed. Most companies follow protocols to install patches on their systems. These protocols account for system dependencies (as in a multi-tiered system), urgency associated with each patch, collateral effects (if known), downtime and system usage and availability, impacts to operations, verification of patch installation and resulting performance, configuration management, roll-back or recovery plan (if needed), system backup and image capture, automated maintenance routines, and compatibility with other applications.

Timely patching of critical systems has long been recognized as a measure of performance. Failure to patch such systems in a timely manner has been noted as a primary cause of network breaches. A 2015 Microsoft Security Intelligence Report stated that most of its customers' systems were breached via vulnerabilities for which patches were released but not installed.[5]

Table 1 shows example patch cadence data for a system. Each row is associated with the urgency of the applied patches, and each column is associated the timeliness with which patches are applied, referenced to a 90-day period. Timeliness is categorized as "Best" (patches applied in the first 10% of the period), "Better" (patches applied in the second 10% of the period), "Good" (patches applied in the third 10% of the period), "Average" (patches applied in the fourth 10% of the period), and "Bad" (patches applied in the last 60% of the period). In this example, 15 critical patches, 2 important patches, 6 moderate patches, and 15 low patches were available for the system.

Future work will consider additional metrics as contributors to the Phase 2 scoring approach to include vulnerability scanning, email phishing, credential security, and self-assessment (penetration testing, auditing) outcomes.

Table 1: Example of Patch Cadence Data.

	Best	Better	Good	Avg	Bad
Critical	5	4	3	2	1
Important	1	0	0	0	1
Moderate	2	1	1	1	1
Low	1	2	3	4	5

Phase 3: Incident Response

Phase 3 focuses on incident response performance using a combination of mean time to detect (MTTD) and mean time to resolve (MTTR) to provide a Breach to Resolution Gap (B2RG) calculation. Increased dwell time in a network provides an adversary the luxury of executing the cyber kill chain[9] multiple times as they maneuver to carry out their intent. The amount of time an adversary remained undetected in a network averaged 56 days in 2019, according to the 2020 Mandiant M-Trends Report.[6] Resolution of the incident is also an important aspect of incident response performance, with MTTR including the additional tasks of containment, eradication, and restoration of services.

Mean Time to Detect

MTTD is the average time it takes to detect a cybersecurity incident within an system. The incident occurrence to detection time is the elapsed time from when an incident occurs (normally indicated in an event log or other monitoring mechanism) to when the incident is detected. Figure 2 represents the concept of MTTD used in this framework. In an organization with capable people, processes, and technologies, detection of anomalous events should occur within minutes.

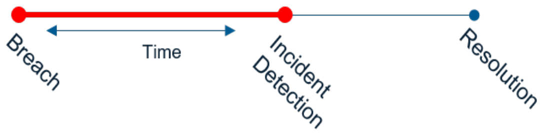

Figure 2: Mean Time to Detect.

An example scenario would be a user opening a malicious attachment received via email. Once opened, the malicious attachment exploits unpatched software on the user's computer, providing attackers an internal beachhead into the network. Subsequent actions by the adversary, in this case connecting from the infected host back to a command and control server, are detected by network intrusion detection systems three days later. The log event indicating when the user executed the malicious attachment is the start time. The detection of adversary communications three days later is the detection time. Time to detect in this case is three days. MTTD is the average of this time over the reporting period for a given system.

Table 2 shows example incident detection data for a system. Each row is associated with the severity of the incident, and each column shows the timeliness of when the incident was detected. Timeliness is categorized as "Best" (incidents detected in the first 10% of the time period), "Better" (incidents detected in the second 10% of the time period), "Good" (incidents detected in the third 10% of the time period), "Average" (incidents detected in the fourth 10% of the time period), and "Bad" (incidents detected in the last 60% of the time period). In this example, 6 severe level incidents, 4 high level incidents, 15 medium level incidents, and 5 low level incidents were detected.

Table 2: Example of Incident Detection Data.

	Best	Better	Good	Avg	Bad
Severe	1	2	3	0	0
High	2	1	0	0	1
Medium	5	4	3	2	1
Low	1	0	1	1	2

Mean Time to Resolution

MTTR is the average time from detection to resolution of an incident as represented in Figure 3. The MTTR metric covers the time from detecting the incident through containment and eradication of the threat, including restoration of services. In an organization with capable people, processes, and technologies, resolution of an incident is possible within hours to days, depending on the severity of the incident. Using the example provided for MTTD, the resolution process begins after the intrusion detection systems alert to the intrusion caused by the malicious email attachment. This process is completed two days later when containment, eradication, and restoration of services have occurred. Time to resolution in this case is two days. MTTR is the average of this time over the reporting period for a given system.

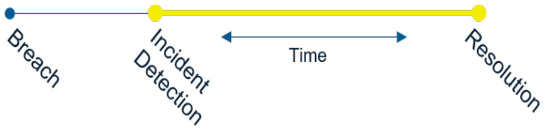

Figure 3: Mean Time to Resolution.

The format for reporting incident resolution data is the same as shown for the incident detection data in Table 2.

Breach to Resolution Gap

B2RG is the sum of the two measurements, MTTD and MTTR, as illustrated in Figure 4. Shortening the time from the incident to detection impacts an adversary's ability to fully execute their kill-chain and/or impact operations. Shortening the time from detection to resolution hastens the containment/eradication of the threat and restoration of all impacted services. The combination of both measured periods of time indicate an organization's comprehensive skill to monitor, detect, and contain an adversary while working to resolve and restore services. From our example case, B2RG is 5 days.

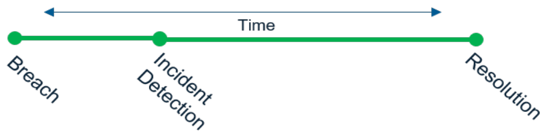

Figure 4: Breach to Resolution Gap.

Based on our research, 7 days is a common time period within which a timely response should detect and resolve various types and levels of incidents.

MATHEMATICAL APPROACH

Each of the cybersecurity criteria/metrics for a system are quantified, weighted, and normalized. They are then used in an algorithm that generates a single numeric value that represents the trust score, ASI, for the system. Shown in Figure 5, this Multi-Criteria

Decision Analysis (MCDA) process accounts for preferences (through weighting) which indicate the significance of each metric.

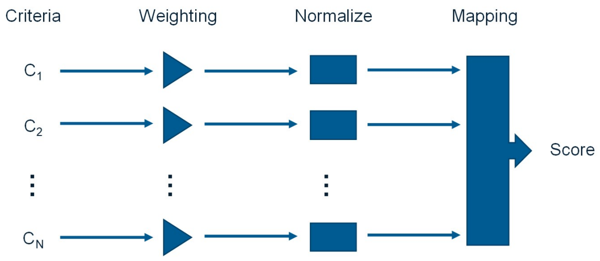

Figure 5: Multi-Criteria Decision Analysis.

Calculation of the ASI is based on CMMC compliance level (Phase 1), patch cadence (Phase 2), and incident detection and resolution (Phase 3). This approach produces metrics from Phases 2 and 3 and maps these to a trust score, with a mapping based on the CMMC compliance level (Phase 1). A system with a higher compliance level receives a more favorable mapping from the performance metrics to the ASI, as illustrated in Figure 6. The ASI ranges from zero (minimum risk) to one (maximum risk). Normalizing each score facilitates the weighting and combination processes. Weighting allows consideration of preferences for each of the criteria.

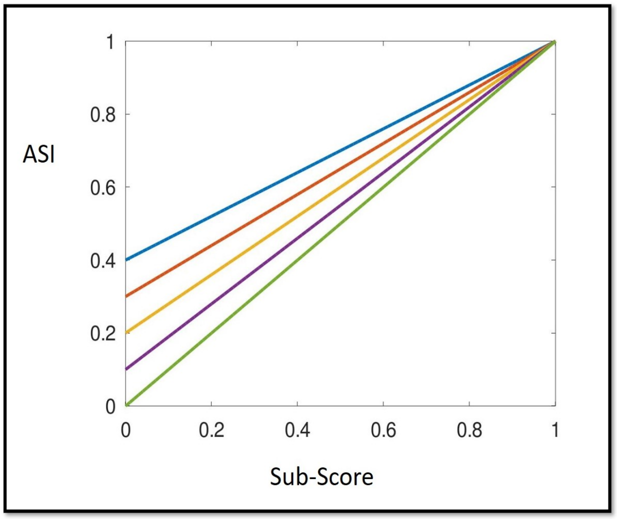

Figure 6: ASI Mapping. The green line represents the mapping for systems with a compliance level equal to five. The blue line represents the mapping for systems with a compliance level equal to one.

Since the significance of any one metric is dependent on the scenario, and metrics and scenarios will vary over time and mission, there is no single best process that produces an optimal ASI. Thus, an iterative application of an MCDA process is used to "tune" the ASI generation. Various weights and mappings are exercised in differing scenarios to produce associated ASI scores. An iterative process of trial runs, with predetermined system architectures and performance factors, must be employed to determine weights and mappings that produce valid ASI scores for each baseline scenario.

The following describes the process used to generate the sub-score for each cybersecurity criterion. The Phase 2 (Patch Cadence) criterion is used in this example to help illustrate the calculation.

The number of patches to be installed for each of the levels of severity (critical, important, moderate, and low) are represented by p_1, p_2, p_3, and p_4, respectively. There are five time-intervals associated with the timeliness of when each of these patches are applied, as defined in Table 1. Table 3 shows how many of the available patches are applied in each of these intervals.

Table 3: Patch Application by Time Interval.

	Best	Better	Good	Avg	Bad
Critical	d_{11}	d_{12}	d_{13}	d_{14}	d_{15}
Important	d_{21}	d_{22}	d_{23}	d_{24}	d_{25}
Moderate	d_{31}	d_{32}	d_{33}	d_{34}	d_{35}
Low	d_{41}	d_{42}	d_{43}	d_{44}	d_{45}

The first step is to normalize the completed patch data to the total number of patches in each category. This produces the data as a fraction of the total.

$$f_{nm} = d_{nm}/p_n$$

The normalized values are then weighted by coefficients w_1 through w_5 according to the timeliness of the patch completion.

$$e_{nm} = w_m f_{nm}$$

The weighting factors (w_m) are chosen in the range of [0,1] to maintain the normalization of the data. These weighting factors account for the preference associated with the timeliness of applying each patch. The weighted values (e_{n1} through e_{n5}) are summed to produce a raw value for each category.

$$r_n = \sum_{i=1}^{5} e_{ni}$$

These normalized values (r_1 through r_4) are the raw sub-scores associated with each category of patch. A second set of weights (c_1 through c_4) is applied to these values according to the preference associated with the criticality of the patches. These weights are chosen so that they sum to a value of 1.0, thus maintaining the normalization of the values.

$$x_n = c_n r_n$$

The representative sub-score (y) is produced from the sum of these normalized values.

$$y = \sum_{i=1}^{4} x_i$$

Finally, this representative value is subtracted from 1.0 to produce the total sub-score associated with this criterion.

$$s = 1 - y$$

This final step produces a total sub-score where a lower score represents a better performance – a score of zero being the best, and a score of one being the worst. The sub-scores are weighted and summed. The weights used for the sub-scores are z_1 for the patch cadence sub-score, z_2 for the detection (MTTD) sub-score, and z_3 for the resolution (MTTR) sub-score. The summed value is then used with the mappings in Figure 6 to produce the associated ASI.

SIMULATED COMPANIES: A PERFORMANCE COMPARISON

Consider three companies, “A,” “B,” and “C” to be evaluated for cybersecurity trust over a 90-day period. Beginning with Phase 1, each company is certified by a third-party auditor and awarded a CMMC compliance level certification. Company A earns a compliance level five rating, and both Company B and Company C earn a compliance level four rating. Each company is now mapped to the ASI scale (Figure 6) based on their associated compliance level. This mapping assigns a baseline score which will be adjusted by application of Phase 2 and Phase 3 performance data, producing a final ASI.

Table 4: Company A - Patch Application.

	Best	Better	Good	Avg	Bad
Critical	6	2	0	0	0
Important	8	3	0	0	0
Moderate	10	3	2	0	0
Low	14	7	1	0	0

Table 5: Company B - Patch Application.

	Best	Better	Good	Avg	Bad
Critical	3	2	2	1	0
Important	3	6	1	1	0
Moderate	2	3	0	10	0
Low	0	10	6	6	0

Table 6: Company C - Patch Application.

	Best	Better	Good	Avg	Bad
Critical	0	2	1	3	2
Important	2	3	1	4	1
Moderate	1	1	5	4	4
Low	3	6	3	6	4

The weighting factors used in this example to compute sub-scores for all three companies are $[w_1, w_2, w_3, w_4, w_5] = [1.0, 0.75, 0.5, 0.25, 0.1]$ and $[c_1, c_2, c_3, c_4] = [0.4, 0.3, 0.2, 0.1]$. The sub-scores are combined with weights $[z_1, z_2, z_3] = [0.4, 0.3, 0.2]$.

For Phase 2 in this example scenario, each company received 56 patches total. Using the priority scale first identified in Table 1, there were 8 Critical, 11 Important, 15 Moderate, and 22 Low patches issued during this 90-day period. The Phase 2 patch cadence performance for these companies is shown in Tables 4 through 6. Company A performed the best out of all 3 organizations with a patch cadence sub-score of 0.08, while companies B and C have patch cadence sub-scores of 0.34 and 0.57, respectively.

For Phase 3 in this example, each organization experienced 26 total incidents. Using the severity scale first identified in Table 2, there were 2 Severe, 5 High, 8 Medium, and 11 Low incidents detected. The Phase 3 MTTD performance for these companies is shown in Tables 7 through 9. Company A performed the best out of all 3 organizations with an MTTD subscore of 0.08, while companies B and C have MTTD sub-scores of 0.34 and 0.57, respectively.

Table 7: Company A - MTTD Data.

	Best	Better	Good	Avg	Bad
Severe	2	0	0	0	0
High	3	2	0	0	0
Medium	6	2	0	0	0
Low	7	3	1	0	0

Table 8: Company B - MTTD Data.

	Best	Better	Good	Avg	Bad
Severe	0	1	0	1	0
High	2	1	1	1	0
Medium	5	1	1	1	0
Low	5	1	1	3	1

Table 9: Company C - MTTD Data.

	Best	Better	Good	Avg	Bad
Severe	0	0	1	0	1
High	1	0	1	1	2
Medium	3	1	1	1	2
Low	2	1	3	2	3

The second portion of the Phase 3 performance includes the MTTR data. Using the same 26 incidents and severity scale as above, the Phase 3 MTTR performance for these companies is shown in Tables 10 through 12. Company A performed the best out of all 3 organizations with an MTTR sub-score of 0.08, while companies B and C have MTTR sub-scores of 0.34 and 0.57, respectively.

The resulting ASI values for these companies are shown in Table 13.

Table 10: Company A - MTTR Data.

	Best	Better	Good	Avg	Bad
Severe	1	1	0	0	0
High	3	1	1	0	0
Medium	4	3	0	0	1
Low	6	3	2	0	0

Table 11: Company B - MTTR Data.

	Best	Better	Good	Avg	Bad
Severe	0	1	0	1	0
High	1	2	1	1	0
Medium	3	1	2	1	1
Low	6	1	1	1	2

Table 12: Company C - MTTR Data.

	Best	Better	Good	Avg	Bad
Severe	0	1	1	0	1
High	1	0	1	1	2
Medium	2	1	0	1	4
Low	4	3	1	1	2

Table 13: Comparison of ASI Values for Companies A, B, and C.

	ASI
Company A	0.09
Company B	0.43
Company C	0.63

Due to a higher CMMC compliance level and better Phase 2 and Phase 3 performance, Company A has the best ASI and is considered a lower cybersecurity risk. Company B and Company C both have a CMMC compliance level of four, but Company B's better performance for Phases 2 and 3 produced a better ASI than Company C. Overall, Company A is considered the lowest risk, followed by Company B and then Company C.

SUMMARY AND FUTURE WORK

Trust is a critical consideration when relying on digital environments outside of the control of the trustor. Quantifying that trust relies heavily on scoring the cybersecurity preparedness, practices, and performance of an organization to ensure a level of defense commensurate with the threat. The ASI is built on a framework designed to accept multiple metrics to accommodate for unique environments, focusing heavily on performance over time to capture variations in the performance data provided.

As organizations participate and performance data is provided for ingestion into the framework, their ASI may adjust to reflect improvements or deterioration in their ability to defend their digital systems. Regardless of the outcome, contributing organizations will have an opportunity to control their ASI score while customers will benefit from a quantifiable method to understand the risk associated with consuming a digitally produced and delivered product.

Future work includes continuing to refine the ASI framework with real-world data from participating organizations. Additionally, adding a reputation phase to capture outliers that may influence the cybersecurity posture is being considered.

Acknowledgments

This technical data was produced for the U. S. Government under Contract No. FA8702-19-C-0001, and is subject to the Rights in Technical Data-Noncommercial Items Clause DFARS 252.227-7013 (FEB 2014). © 2020 The MITRE Corporation. All Rights Reserved.

References

1. Space Foundation Editorial Team. The space report reveals 2018 global space economy exceeded $400 billion for the first time. *Space Foundation*, July 2019.

2. Morton Deutsch. *Cooperation and trust: Some theoretical notes.*, pages 275–320. Nebraska Symposium on Motivation, 1962. Univer. Nebraska Press, Oxford, England, 1962. `https://psycnet.apa.org/record/1964-01869-002`.

3. Tyrone Grandison and Morri Sloman. A survey of trust in internet applications. *IEEE Communications Surveys & Tutorials*, 3(4):2–16, 2000.

4. Matthew Stein, Pete de Graaf, Sean Kinser, John Muhonen, Robert Roller, Amanda Salmoiraghi, Keith Scott, and David Voss. Hybrid Architecture Performance and Evaluation for Quantitative and Comparative Analysis. *(Submitted to 34th Annual Small Satellite Conference Proceedings)*, August 2020.

5. Microsoft Security Team. Microsoft security intelligence report. *Microsoft Security*, Feb 2019.

6. FireEye Journal, 2020 *FireEye annual threat report.* Mandiant Pub., 2020. `https://www.fireeye.com/current-threats/annual-threat-report/mtrends.html`

7. Cybersecurity Maturity Model Certification (CMMC) *Office of the Under Secretary of Defense for Acquisition and Sustainment (OUSD(A&S))* Jan 2020. `https://www.acq.osd.mil/cmmc/`

8. Cyber Hygiene: 11 Essential Practices *Carnegie Mellon University, Software Engineering Institute* Nov 2017. `https://insights.sei.cmu.edu/insider-threat/2017/11/cyber-hygiene-11-essential-practices.html`

9. The Cyber Kill Chain, A Lockheed Martin Overview *The Lockheed Martin Journal* 2020. `https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html`