A Risk Assessment and Reduction Approach for National Critical Infrastructure

Jason Reinhardt[a], Merideth Secor[b], Lindsey Miles[b], Ron Lafond[b], Derek Koolman, II[b], Lauren Wind[c], Ray Ludwig[c], and Jeff Munns[c]

[a] Sandia National Laboratories, Albuquerque, New Mexico, jcreinh@sandia.gov

[b] Cybersecurity and Infrastructure Security Agency (CISA)CISA, Arlington, Virginia

[c] Systems Planning and Analysis Inc, Alexandria, Virginia

Abstract: The Cybersecurity and Infrastructure Security Agency (CISA) leads the national effort to understand, manage, and reduce risk to our cyber and physical infrastructure. CISA must assess risks that cover a broad range of scenarios over a complex set of interdependent critical infrastructure (CI) systems. While many threat and hazard impact models and data sets exist, there is no overarching analytic structure that organizes and integrates these disparate sources into a unified risk assessment. CISA is building capabilities that will address these challenges to support stakeholders across all levels of government and the private sector. First, CISA has developed a National Critical Functions (NCFs) data structure to organize and describe critical infrastructure. This data set provides a set of decompositions structured as directed graphs that break down each identified function into enabling sub-functions that detail the operation and interdependencies across disparate CI systems. The functional description of NCFs serves as a complementary lens to the sector-based organization of CI and better facilitates systemic and cross-sector risk analysis. Additionally, CISA has begun developing the Risk Architecture, a technology-enabled analytic tool that contains a set of standards, scenarios, visualizations, and workflows that leverage the NCF and other integrated CI data sets. This paper describes the need for an integrated approach to CI risk assessment, the NCF decomposition structure, the principles and concepts behind the Risk Architecture, and the approaches to functional interdependency analysis while also providing initial use examples.

1. INTRODUCTION

The Cybersecurity and Infrastructure Security Agency (CISA) leads the national effort to understand, manage, and reduce risk to our cyber and physical infrastructure [1]. To do this, CISA connects public and private sector critical infrastructure (CI) stakeholders to resources, analyses, and tools to help reduce risk and build resilience. Within CISA, the National Risk Management Center (NRMC) works with government and industry to identify, analyze, prioritize, and manage the most significant and systemic strategic risks to the nation's CI [2]. In the current rapidly evolving threat environment, sources of strategic risk are widespread and include cyber and physical attacks; supply chain vulnerabilities; malicious exploits of emerging technology; nation-state aggression; insider threats; pandemics and natural disasters; and the convergence of previously discrete threats and vulnerabilities [3].

The cyber and physical infrastructure across the United States is a complex network of interdependent systems and assets that are categorized under the sixteen CI sectors [4]. These systems and assets are also part of the functions that underpin modern life and span the lens of the sixteen CI sectors construct. Systemic critical infrastructure risks generally implicate multiple sectors simultaneously and pose a threat to national security, national economic security, and national public health or safety [5]. These systemic risks are what CISA has been charged to manage and ultimately reduce.

2. IMPROVING ANALYSIS OF CRITICAL INFRASTRUCTURE RISK

As critical infrastructure becomes more complex and interdependent, CISA has adapted to help stakeholders and operators of CI assess, understand, and manage risk nationwide (for an explanation of

DHS's definition of risk, see [5]). The scale and complexity of this challenge is daunting. An uncountable number of assets (hardware, software, services, etc.) perform the tasks that enable critical functions. While CI system models exist (see for example [6]), they are generally focused on specific sectors and often on specific threats. There is a need for models to be better integrated and aggregated to understand the cascading impacts to CI at the national level. Further, in many cases data that describes specific systems is simply unavailable to analysts working to rapidly respond to unfolding crises. The nation's CI is constantly being upgraded by government and private entities as new technologies come online and new capacities are developed. Data that is available may be rapidly out of date as the systems, policies, and technologies that make up the nation's CI evolve. The span of domains and expertise required to understand all aspects of CI creates additional difficulties as disparate communities, fields, and lexicons must be brought together. Finally, the scenario space that describes the range of possible threats, vulnerabilities, and consequences of concern, from extreme weather, cyber-attacks, and pandemics, to technical failures, supply chain disruptions, and terrorism, is vast, creating scaling challenges for any structured risk analysis.

A standard, high-level model is needed that describes the nation's CI along with a common set of tools that leverage that model to develop risk assessments that inform CISA, the Sector Risk Management Agencies (or SRMAs), and the broader CI community as they address challenging questions. CISA determined that by building capabilities that are functionally focused, the CI community will have a richer understanding of how organizations, technologies, policies, and other factors come together to connect, distribute, manage, and supply critical systems and services. A functional understanding of CI can also enable analysts to understand how failures in the key systems, assets, components, and technologies may cascade across sectors and industries. A functional description of CI can improve analysis by promoting more complete, systematic, and repeatable assessments of risk and by providing a common standard for describing CI threats, vulnerabilities, and consequences. A standard functional description also provides an integrating framework for the disparate models and data sets by setting common definitions for the functions and their disruption.

Over the past two years, the CISA has embarked on a set of initiatives to develop a function-based approach to assessing and managing CI risk. The CISA began by establishing the National Critical Functions (NCFs) framework. The NCF framework serves as a complementary view to the sector organization of CI and provides a description of the CI system that can be considered in addition to asset-based, geographic, and organizational descriptions. That is, a functional representation provides a "what does it do" view of CI, in addition to the "what is it," "where is it," and "who's responsible for it" views. Each of these views plays an important role in assessing, understanding, and managing risk to CI. Specifically, the value of the NCF framework is in its ability to convey the complexities and interdependencies of CI and how they operate at a system level. CISA will continue developing this information in coordination with critical infrastructure stakeholders, endeavoring to deepen the understanding of who and what is required for the successful, sustained, and resilient operation of individual NCFs.

As a second initiative, the CISA has begun the process of developing new tool sets and analytic capabilities that can evaluate CI risk using the NCF framework as a core model, while integrating a broad set of existing and future data and tools. This constellation of capabilities is known as the Risk Architecture. The Risk Architecture operates on top of the NCF framework and facilitates data integration, model coupling, and decision-support visualization for high-priority CI risk analysis. CISA has assembled a proof-of-concept of the Risk Architecture and is currently building upon that foundation to fully implement a NCF and Risk Architecture approach. The Risk Architecture, in combination with the NCF framework, offers a new way of conducting dependency analysis between CI components and systemic critical infrastructure risk, but will require continued effort to develop and apply novel risk analysis solutions. By conducting analysis through the prism of interlocking functions and systems, the Risk Architecture provides a transformative new approach to critical infrastructure risk management.

3. NATIONAL CRITICAL FUNCTIONS

The NCFs are defined as the functions of government and the private sector so vital to the U.S. that their disruption, corruption, or dysfunction would have a debilitating effect on security, national economic security, national public health or safety [7]. Currently, there are fifty-five NCFs that encompass the nation's CI network to distribute, maintain, supply, and connect the services and operations the nation relies on. Each NCF describes a high-level core function such as "Distribute Electricity," "Manage Wastewater," or "Provide Positioning, Navigation, and Timing Services." While CISA's understanding of NCFs continues to evolve and revisions may be made over time, this initial set provides a foundation for assessing and managing systemic risk within CI [8].

3.1. Functional Decomposition

The NCFs themselves provide only the first step in achieving a functional understanding of disruptions, impacts, and interdependencies of critical infrastructure. Each NCF is broad in scope and more detailed functional descriptions will likely be required for more focused and meaningful analysis. Therefore, each NCF can be further broken down into child subfunctions, potentially across several layers. Each set of subfunctional nodes that are decomposed from a function node detail all the necessary functions for the functions operation. As a notional (and incomplete) example, if an NCF was to Supply Potable Water, its first level subfunctions might be Source Fresh Water, Store Water, Treat Water, and Distribute Water. Each of those subfunctions can then be decomposed into narrower subfunctions perhaps adding details about pumping, contaminant testing, maintenance, or other supporting functions. The resulting structure is a tree of functional nodes where each level describes the NCF in increasing detail and specificity. This is represented schematically in Figure 1. As the scope of functional nodes narrows at lower and lower levels of decomposition, classes of physical or virtual assets that enable those subfunctions can be tagged. In our example, we might start identifying where pumping facilities, reservoirs, and treatment centers service specific functions. This structure is extensible, allowing for as much detail as is needed for analysis provided the data exists.

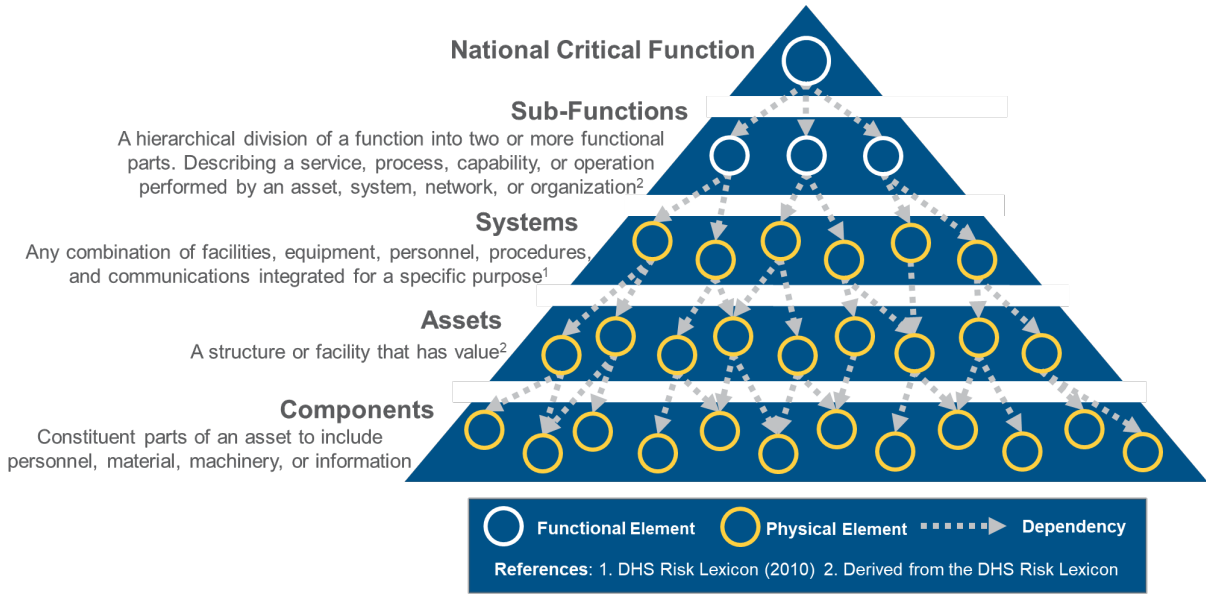

In Fiscal Year (FY) 2021, CISA and its supporting researchers developed the first set of NCF decomposition data which includes over 3,500 functional nodes across all 55 NCFs.

3.2. Functional Interdependencies

Functions do not operate independently; they rely on the provision of other functions to operate successfully. Servers do not run without electricity or network access, water treatment facilities cannot

operate without certain chemicals or personnel, and first responders will have difficulty ensuring the health and safety of the community without communications capabilities. Capturing these dependencies within the set of NCFs and their subfunctions is critical to evaluating risk to the nation's critical infrastructure.

Dependencies between nodes within a NCF are referred to as intrafunctional dependencies. An intrafunctional dependency may indicate that a function is dependent on its subfunctions and is created during the NCF decomposition process. Intrafunctional dependencies may also exist between subfunctions across different parts of the decomposition tree and indicate some kind of non-decompositional dependency. In our example, it could be that a single pumping station not only provides system flow to distribution, but also provides the pressurization required for filtration. Functionally, there might be two subfunction nodes representing the need for pumping stations in the tree with one existing in both the Treat Water lineage and the Distribute Water lineage. But there may be a dependency between the two that indicates that they are both a single instance of the same pumping station.

Dependencies that exist between nodes in two different NCFs are referred to as interfunctional dependencies. These dependencies represent how functions and subfunctions of one NCF may depend on the functions or subfunctions of another NCF. For example, a disruption to a petroleum pipeline's operations may impact the ability to get raw products to a refinery. This may cause a regional shortage over time as fuel reserves are used up, eventually causing shortages that can impact vehicle fleets critical to first responders, city operations, and the delivery of goods and services.

3.3. Functions to Assets

While a functional view of CI represents a significantly different approach than a sector or asset-based view, the different approaches must be complementary and integrated. Large data sets that list types, or even instances of assets that support CI function exist, as do categorization and coding schemas. In some instances, CISA has been able to connect the functional level data to the asset level data through orchestrated use cases, such as the Risk Architecture version 1 tool that connected functional-to-asset data for "Transport Material by Pipeline". Additional efforts are underway to make linkages from the NCF structure to the modernization of the Infrastructure Data Taxonomy [9] connect to systemically important entities [previously referred to as systemically important critical infrastructure, as discussed in 10] and map to existing industrial data sources.

3.4. Graph Representation

Researchers have proposed utilizing graph structures to represent and analyze dependencies in CI [11, 12]. The functions identified in the NCF decomposition efforts, and their interdependencies can be structured as a graph, $G = (V, E)$, where the vertices of the graph, V, represent an NCF and their subfunctions, and the edges, E, represent the dependencies between them. Nodes and edges may be of different types, but each one represents a discrete function that must be accomplished to fulfill the objective of an NCF. Edges may be of different types with each type describing different kinds of dependencies such as a flow of different commodities or information, common constituent components, geographical colocations, or logical connections. Edges are directed and point towards the dependent node – if node B depends on node A in some way, then the edge describing that dependency points from A to B. Because functions depend on their subfunctions, intrafunctional dependencies that indicate a functional decomposition point from child subfunctions to parent functions. An illustrative representation of this graph formulation is shown in Figure 2.

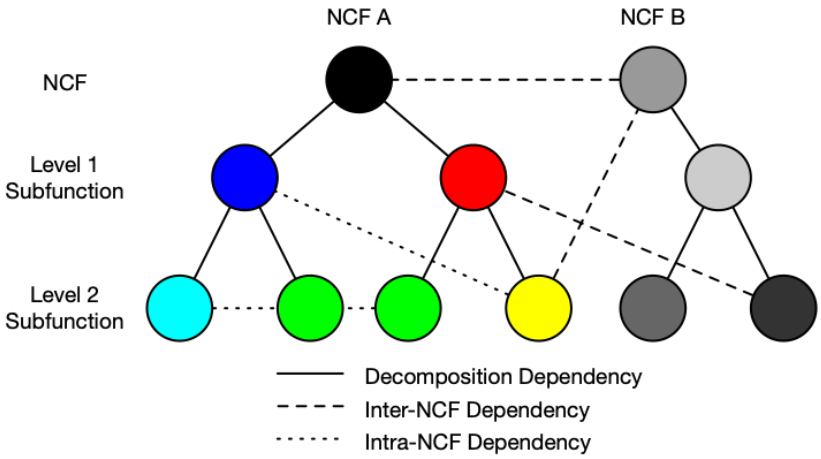

Figure 2: Example Graph Representation of NCF Decompositions and Interdependencies

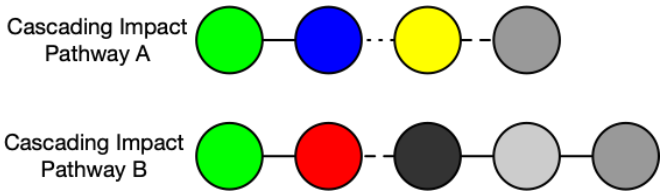

Figure 3: Example Functional Impact Cascade Pathways

By capturing the dependent relationships between functions that enable critical infrastructure, and organizing those relationships in a systematic manner, the graph representation of NCFs provides several benefits. First, the graph itself provides a foundational common understanding of the NCF structure that can serve as a basis for joint analysis activities. Second, the functional representation of NCFs enables analysts to understand potential impacts resulting from dependencies even when as-built data on physical or virtual infrastructure and dependencies is unavailable or impossible to access. Third, the graph structure provides a natural structure for enumerating and evaluating scenarios of concern and organizing analysis around specific functions and disruptions. For example, Figure 3 illustrates how different pathways of cascading impacts may be determined from such a directed multi-graph structure.

Rendering the NCF structure in a graph with standardized definitions also enables systematic evaluations across all three elements of risk: threats, vulnerabilities, and consequences. When considering threats, an analyst may pose questions that consider types of systems or functions that are being targeted by an adversary, or those that share properties that make them likely targets, such as geographies or common components, producing a set of possible scenarios to explore. Using a common set of function definitions and capturing those commonalities in a standard form ensures that these threat analyses are complete and systematic in identifying the attack surface of the NCFs. Further, vulnerability assessments may focus on assets that serve the functions that enable the NCFs, allowing known vulnerabilities to be mapped from assets to a complete set of functions over all NCFs and systemic impacts to be identified. Finally, from a consequence perspective, the graph representation of NCFs allows an analyst to understand how impacts from scenarios that capitalize on those vulnerabilities ripple through the NCFs and create cascading impacts to multiple systems. Further, these assessments can then be integrated end-to-end in scenarios to create aggregate risk assessments. Doing so requires a set of tools, standards, and methods that can operate on the NCF functional decomposition and asset data.

4. RISK ARCHITECTURE

The Risk Architecture is a technology-enabled analytic tool for performing cross-cutting risk analysis for interdependent CI that operates on the NCF structure. While the NCFs represent a systematic and

scalable model of those functions that are vital to enable the operation of the nation's CI, analysts need a toolkit to provide meaningful analysis of the resulting graph. This toolkit must also enable the integration of advanced models and additional data sets into end-to-end risk assessment and analysis results. Additionally, it must provide a common interface and environment for a disparate set of analysts to perform risk analytic operations in a systematic, principled, and repeatable way. The Risk Architecture must be a modular enabling structure, providing toolsets, standards, interfaces, templates, data repositories, and links to external resources for analysis, but also establishing clear handoffs between analysts, consumers, and others involved in the risk assessment process.

The Risk Architecture must address several challenges to properly inform and enable the reduction of risks to the nation's critical infrastructure. It must be built on sound risk analytic principles and enable systematic analyses that are repeatable. Further, the risk architecture must be pragmatic in providing the best assessments given the available data, time, and resources. It must also be scalable so that as new data, models, methods, and resources are made available the quality and capacity for risk analysis increases. Finally, the Risk Architecture must be accessible in that analysts can easily integrate these tools and concepts into their daily workflows and activities. Additionally, it is clear from building out the NCF data set and the associated Risk Architecture methodology that there needs to be clear data connections and linkages. This will enable CISA to fully harmonize various data constructs and leverage existing tools and methodologies.

4.1. Risk Architecture Capability Layers

In trying to achieve these design goals, the Risk Architecture must develop capabilities across a set of four interacting layers, as illustrated in Figure 4. The top layer represents the Risk Architecture Environment, which provides an interface for users, a set of standards that allows for interoperability between the components, and the infrastructure upon which the Risk Architecture operates. The second layer is the Assessment Frameworks and Tools layer, where common assessment tasks and workflows are captured along with lightweight processing tools that analysts can use to examine existing risk data, scenarios, and results. The third layer is the Integrated Dataset layer, which holds common data sets, such as the NCF decomposition graph, scenario libraries, and other pre-build results. These results must already be processed and ready for common risk assessment requests in order to provide timely answers, even if approximate for some situations. The final layer is Analytical Models & Data, which contains analysis tools where modeling and simulation of specific critical infrastructure systems is performed, and system data is collected in detail and processed into data sets that can exist in the Integrated Dataset layer. Note that there is a distinction between assessment and analysis drawn in these layers for the purpose of explanation of the Risk Architecture capabilities. Existing integrated data sets are used to perform rapid responses to pressing risk questions and lightweight tools are referred to as part of the risk assessment process. Deeper development of new data from exploratory or deliberative analytic activities that create integrated datasets are referred to as analysis.

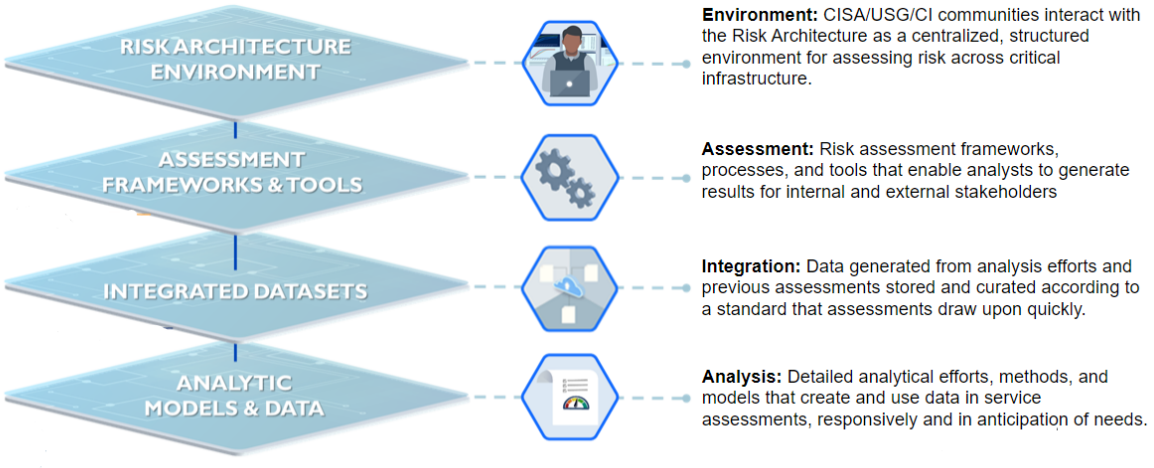

Figure 4: Operational Capability Layers within the Risk Architecture

4.2. Scenarios, Matrices, and Pinch Points

One of the most challenging methodological issues in creating the Risk Architecture is the addressing the sprawling scenario spaces that are within CISA's scope. Considering all hazards and threats, all known vulnerabilities and consequences would rapidly create a probabilistic event tree that is so large that it is not computable. This is not a new problem and the trade-off between scenario space complexity and computational pragmatism is a constant issue in risk analysis. Navigating this trade-off at scale requires careful choices of how to treat classes of scenarios as approximations of many specific scenarios and to develop a rigorous structure and standards that can integrate many disparate analyses.

A convenient set of tools for organizing these types of structures are matrix formalisms, which have been successfully applied, like many of the tools in probabilistic risk assessment, in the cases of nuclear power plant risk analysis [13, 14]. The approach provides a way of organizing highly complex system models in risk analysis using successive matrix multiplication, describing an "initiating event vector" (threat/hazard), a "plant model matrix" and "containment model matrix" (vulnerability), and a "site model matrix" (consequence). Multiplying the initiating event vector by the plant model matrix gets you a vector that produces the various achievable plant states. Multiplying those plant states by successive matrices produces vectors that capture the state of the system along the causal chain. This multiplication of each element is referred to as the global assembly equation and mimics the function of an event tree. While used successfully in focused applications such as nuclear power, such approaches have also been proposed for use in global risk analysis efforts [13].

The Risk Architecture strives to organize risk analysis of CI along a similar means. CI analysts will maintain a threat vector describing a distribution over potential classes of initiating events, denoted ϕ^I, and capturing m possible scenario classes. The ability for those events to present some subset of n potential compromises to CI systems is captured by the $m \times n$ matrix, T. The potential for those n compromises to result in a specific set of q prompt disruptions to specific CI systems or functions is captured by the $n \times q$ matrix, V. Finally, the potential for those prompt disruptions to manifest r types of consequences through cascading impacts in the interdependent CI system is given by the $q \times r$ matrix, C. Each of the matrices, T, V, and C, is a stochastic matrix (rows sum to one) and each element represents estimated conditional probabilities that an initiating event manifests a potential system compromise, that a potential system compromise creates a specific disruption condition, or that a disruption creates a specific consequence condition, respectively. Finally, each matrix can be further decomposed into additional matrices that are multiplied together to get the corresponding integrated matrix for T, V, or C. This system is illustrated in Figure 5.

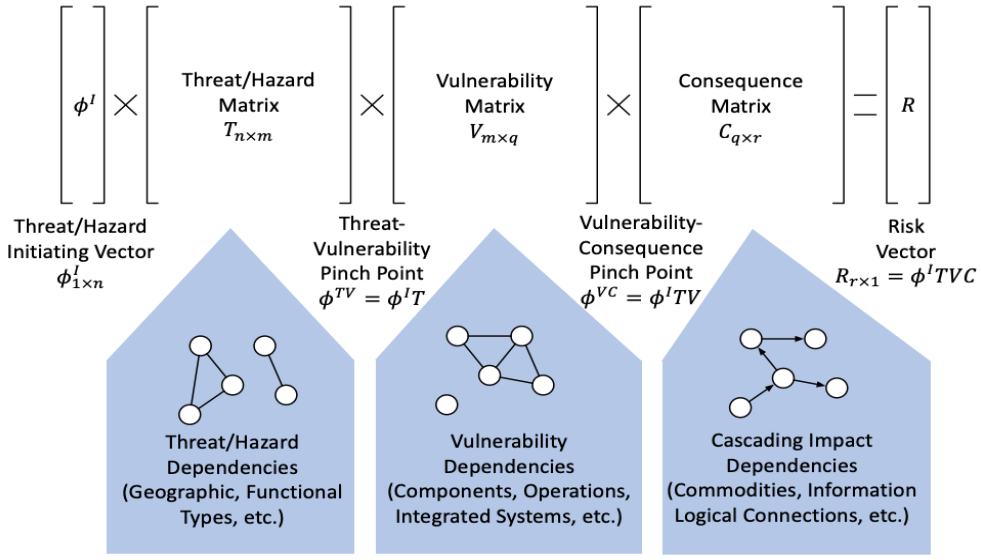

Figure 5: Schematic Representation of Matrix Formalism, Pinch Points, and the Role of Dependency Graphs in the Risk Architecture

The specific situations that are enumerated along the dimensions of the T, V, or C matrices must be determined and agreed upon as a standard for CISA to use to evaluate risk. Instead of allowing the values of m, n, q, and r to grow exponentially as the event tree grows, constraining these dimensions to a reasonably small set is necessary because of the limits to which data and resources are available to characterize different matrix elements. This provides a compact and tractable scenario set over which to evaluate risk but comes with a cost. Specifically, limiting the dimensions of the matrices also limits the degree to which conditionality in the system can be accounted for from end-to-end because it compresses a set of scenarios to pass through specific states rather than accounting for all. This is an approximation that distorts the risk assessment but is likely a necessary and pragmatic one. Methods to choose those compression states has been the subject of some research already [15].

The intermediate vectors determined from incrementally multiplying the T, V, or C matrices represent the distributions over the compressed scenario states. Multiplying the threat/hazard initiating vector, ϕ^I, by the threat/hazard matrix, T, to get the $1 \times m$ vector $\phi^{TV} = \phi^I T$. The vector ϕ^{TV} describes the distribution over the potential system compromises given the initiating vector, ϕ^I. The vector ϕ^{VC} describes a distribution over specific disruption conditions. These intermediate vectors are referred to as pinch points. Kaplan described pinch points as follows "…that every initiating event could also be considered as a pinch point in a larger tree. Thus, this initiating event is caused by an earlier event, which results from a still earlier one, and so on…" [13]. Each of the intermediate pinch point vectors in a matrix formalism has a common property: given the intermediate vector itself one does not need the upstream models to do the remaining analysis. This means that analysts can work in separable spaces to answer some interesting questions. For example, an analyst may search a subspace of the $\phi^{VC}C$ vector space to find consequence maximizing pinch points producing an ordered set of priority system disruptions. As another example an analyst may produce an alternative version of the vulnerability matrix, V', that accounts for proposed changes to the system, such as the deployment of new technologies, changes in policy, or other propositions, and compare the results to a baseline in order to evaluate risk reduction performance.

5. EXAMPLE APPLICATION

A use-case scenario was developed to test the foundational functional analysis capabilities, analytic frameworks, and user interface of the proof-of-concept Risk Architecture Tool developed in FY 21 by the CISA. The original use-case scenario focused on a notional example in which the Risk Architecture develops, analyzes, and communicates the functional dependencies and cascading consequences to CI related to a cyber-attack on a specific NCF system in the U.S. The NCF targeted in early FY 2021 was "Transport Materials by Pipeline" and the system chosen was the Colonial Pipeline. It is important to state that the Colonial Pipeline use-case scenario was developed and approved prior to the ransomware attack on the Colonial Pipeline in May 2021, which resulted in a depletion of 4.6 million barrels along the East Coast over six days and a $4.4 million USD payout [18, 19]. The Colonial Pipeline is the largest U.S. refined petroleum products pipeline system connecting the U.S. Gulf Coast and the New York Harbor Area. It transports over 100 million barrels of fuel daily through 5,500 miles of pipe and provides fuel to over 50 million people in 14 states and seven major airports. Pipeline systems are well studied by CISA analysts through Requests for Information (RFIs) and asked analyses. As a result, data and consequence estimates for disruptions are available for use. Additionally, pipeline CI is represented by a NCF which was among the first to be functionally decomposed by CISA. It exemplifies a collection of CI assets whose disruption could have significant national impacts and is directly relevant to the CISA's mission, which is why it was chosen for the Risk Architecture Tool's proof-of-concept use-case scenario.

5.1. Analytic Questions

The proof-of-concept Risk Architecture Tool was built using foundational capabilities which leverage the CISA's growing understanding of the NCF functional decompositions into associated sub-functions, systems, assets, and components. Starting with NCF functional decomposition data from the "Transport Material by Pipeline" NCF, the analyst can ask and analyze strategic risk questions related to within-NCF dependencies—i.e., how systems, assets, and components interact to successfully provision NCF subfunctions. A summary analysis of intra- and inter-dependencies is produced which an analyst can use to identify the 1st, 2nd and 3rd order consequences derived from a notional pipeline disruption by using simple graph traversal and path count algorithms derived from the Risk Architecture Tool. Additional metrics can help to identify critical nodes (i.e., network centrality) [20]. Importantly, available threat, vulnerability, and consequence information can be stored internally within the Risk Architecture Tool, allowing the analyst to add context to their analysis.

There were seven foundational analytical questions developed by analysts, data scientists, and CI subject matter experts to analyze functional to asset-level dependencies. These questions were used to scope and guide the proof-of-concept Risk Architecture Tool developed. However, the capabilities developed to address these questions can also be used interchangeably for any NCF future scenario and are intended to be expanded upon in future Risk Architecture Tool iterations. The proof-of-concept Risk Architecture Tool analytic questions are:

1. Functionally, how are materials transported by pipeline?
2. Which other NCFs have "Transport Materials by Pipeline" as a critical input?
3. What second and third order effects result from a disruption to the "Transport Materials by Pipeline" NCF?
4. Which other NCFs are critical to the "Transport Materials by Pipeline" NCF?
5. What Colonial Pipeline related assets could be a vector of compromise for "Transport Materials by Pipeline"?
6. Which infrastructure assets collectively enable the Colonial Pipeline's ability to transport materials?
7. Where is the network of infrastructure enabling the Colonial Pipeline located?

5.2. Results

Using the proof-of-concept Risk Architecture Tool, analysts determined the primary and secondary cascading consequences to the "Transport Material by Pipeline" NCF because of a cyber disruption to the Colonial Pipeline. Working through the seven analytic questions, an analyst can leverage the Risk Architecture Tool to explore the NCF network and linked data. This enables analysts to address each question as the RFIs increase in level of detail. For the first time in CI analysis, the functional lens is used to explore how an NCF (i.e., "Transport Materials by Pipeline') operates. There are 91 subfunctions that compose "Transport Materials by Pipeline", which all branch from either transport liquids or transport gases at the first subfunction level.

The analyst can now visualize a decomposed NCF to its own subfunctions and linked to sector-based data via the Infrastructure Data Taxonomy, which is widely used across CISA and other government agencies. The analyst can also identify the upstream and downstream NCF dependencies given the disruption to the NCF. For the use-case scenario, an analyst will discover that there are 18 first order (primary) dependent NCFs to "Transport Materials by Pipeline." Through the dependency cascade analysis developed within the Risk Architecture Tool, the second order (secondary) dependent NCFs, an analyst will find that all 55 NCFs are dependent on "Transport materials by Pipeline" through the highly connected network within three dependency links.

An analyst can evaluate the network metrics associated with the cascade pathways and understand the number of discrete pathways between nodes. A high number of path counts may indicate that a dependent node can be disrupted multiple ways because of a "Transport Material by Pipeline" impact.

Dependency mapping allows exploration of how those cascades happen and will enable an analyst to dig into more and more detail as the NCF decomposition and dependency data continues to mature. The Risk Architecture Tool ultimately allows policy makers to see how interconnected the NCFs are at the functional level. For example, "Manage Wastewater" is a secondary dependency on "Transport Materials by Pipeline" by means of the "Store Fuel and Maintain Reserves" primary dependency cascade. The Risk Architecture Tool also enables the analyst to follow dependencies in the reverse direction and identify NCFs upon which "Transport Material by Pipeline" is dependent. For this use-case scenario, "Transport Materials by Pipeline" is dependent upon 33 NCFs.

Moving on down the NCF decomposition lens from functional to physical (i.e., asset) levels, the Risk Architecture Tool provides an in-depth capability to evaluate the physical representation of the NCF. Using relational asset-level datasets, "Transport Materials by Pipeline" was found to be composed of over 70,000 assets via the Infrastructure Data Taxonomy, and filter specifically to highlight the hundreds of Colonial Pipeline-owned assets. The Risk Architecture Tool produces a visualization that shows the network from NCF node to subfunctions, with the Infrastructure Data Taxonomy linked nodes to the connected assets. Figure 6 provides a screen shot of the Risk Architecture Tool user interface, depicting the physical assets linked to the "Transport Material by Pipeline" NCF functional analysis and the matching geographic location of the assets owned by the Colonial Pipeline to support the use-case scenario.

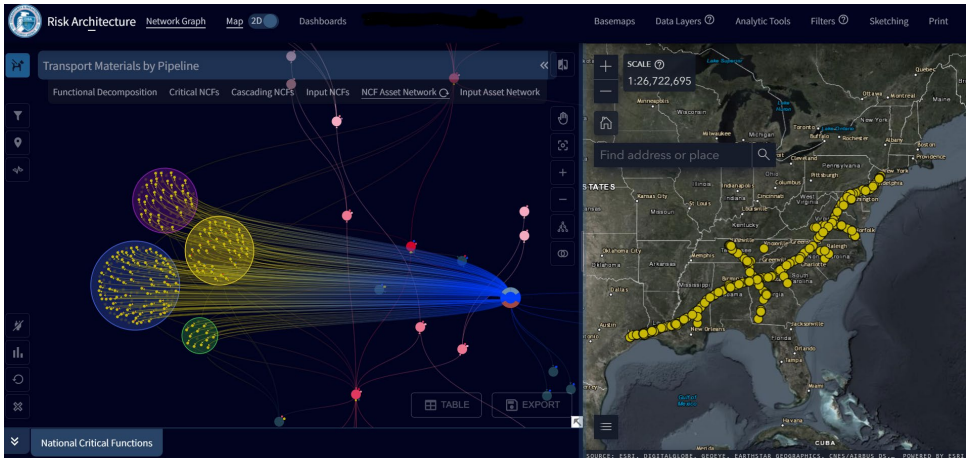

Figure 6: Screen Shot of Proof-of-Concept Risk Architecture Tool

5.3. Lessons Learned

The Risk Architecture Tool capabilities, use-case scenario analysis, and challenges (both analytical and technical) have been demonstrated across the CISA with various stakeholders. Without having a complete NCF decomposition dataset, the Risk Architecture Tool has proven that key functional CI dependencies can be identified and analyzed across the NCF network. The benefit of the proof-of-concept Risk Architecture Tool is that it is: 1) reproducible, 2) additive with future datasets (e.g., event likelihood, vulnerability, consequence) and integrated models, and 3) allows an analyst to answer foundational analytical questions or dive deep into their own analysis with the Risk Architecture Tool dependency capabilities. As a result, the current capability provides a systematic and standardized method for analysts to examine cascading NCF dependencies and the systems and assets associated with those impacts.

6. NEXT STEPS

The main challenge to building a system like the Risk Architecture is that it needs to meet the needs of critical infrastructure risk analysts while also ensuring the tool can evolve with the NCFs and other emerging data sets. The data will need to be regularly validated and maintained, and the Risk Architecture tool itself must integrate new models, workflows, user stories, and methodologies as it

grows over time. Perhaps most importantly, the Risk Architecture needs to have a prospective, forward-looking lens to ensure CISA's investments will add value and support risk-informed decision making.

The CISA has already successfully piloted a use-case scenario around a disruption to the Colonial Pipeline and is in the process of developing a second major use-case scenario around a disruption to water treatment systems. Ultimately, use case scenarios will be used to develop additional analytic capability, to validate systems-level data, and to improve the logic structures needed to translate impacts to assets and systems into functional impacts to the NCFs. Through these efforts, CISA will be able to better contextualize the risk landscape (including a national risk baseline) and identify potential mitigation options that can be used to buy down risk over time.

As the CISA continues to develop and mature the Risk Architecture, it will broaden its focus from the current concentration on event consequences to include vulnerability and event likelihood data. CISA's broad mission focus makes it critical that the CISA partner with CISA subcomponents focused on physical and cyber risk to infrastructure to gain access to other available data feeds (especially regarding cyber threat data feeds) as it builds out the Risk Architecture, thus ensuring that the Risk Architecture is responsive to the entirety of the CISA mission space.

7. CONCLUSION

The National Risk Management Center is developing the Risk Architecture to enable better identification and analysis of risk through the analytic lens of the NCFs. The CISA Risk Architecture will be an innovative, adaptable system that takes advantage of the latest advances in data architecture, simulation and modeling, neural networking, and other data and risk analysis fields to develop a deeper understanding of risk to the nation's CI. This risk analysis information system environment is being built to provide risk analysts and government and private sector decision makers with timely, relevant, analytically rigorous, and defensible insights into current and emerging risks to the CI that they are responsible for overseeing. Ultimately, the goal is to develop a layered information environment that leverages and integrates various types of CI data, tools and models, and analytic approaches to develop holistic understanding of event and strategic risks and provide analytic value to the CI community and the American public.

Funding Disclosure

This research paper was supported by Sandia National Laboratory (IAA no. 70RCSA20K00000044) and Systems Planning and Analysis, Inc. (contract no. GS00Q14OADU104).[*]

[*] The underlying research discussed in this paper was further supported by the following organizations: Argonne National Laboratory, Idaho National Laboratory, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, Pacific Northwest National Laboratory, and RAND.

References

[1] U.S. Department of Homeland Security Cybersecurity and Infrastructure Security Agency. (n.d.). *About CISA*. CISA.Gov. Accessed March 20, 2022, from https://www.cisa.gov/about-cisa

[2] U.S. Department of Homeland Security Cybersecurity and Infrastructure Security Agency. (n.d.). *National Risk Management Center*. CISA.Gov. Accessed March 20, 2022, from https://www.cisa.gov/nrmc

[3] U.S. Department of Homeland Security. (2011, April). *Risk Management Fundamentals*. DHS.Gov. Accessed March 20, 2022, from https://www.dhs.gov/publication/risk-management-fundamentals

[4] U.S. Department of Homeland Security Cybersecurity and Infrastructure Security Agency. (n.d.). *Critical Infrastructure Sectors*. CISA.Gov. Accessed March 20, 2022, from https://www.cisa.gov/critical-infrastructure-sectors

[5] U.S. Department of Homeland Security. (2010, September). *Risk Lexicon*. DHS.Gov. Accessed March 20, 2022, from https://www.cisa.gov/dhs-risk-lexicon

[6] U.S. Environmental Protection Agency. (n.d.) *EPAnet: Application for modeling drinking water distribution systems.* Accessed March 21, 2022, from https://www.epa.gov/water-research/epanet

[7] U.S. Department of Homeland Security Cybersecurity and Infrastructure Security Agency. (n.d.). *National Critical Functions Set*. CISA.Gov. Accessed March 20, 2022, from https://www.cisa.gov/national-critical-functions-set

[8] U.S. Department of Homeland Security Cybersecurity and Infrastructure Security Agency. (2021, December). National Critical Functions: Status Update to the Critical Infrastructure Community. https://www.cisa.gov/sites/default/files/publications/2021_ncf-status_update_508.pdf

[9] U.S. Department of Homeland Security Cybersecurity and Infrastructure Security Agency. (n.d.). Infrastructure Data Taxonomy. CISA.Gov. Accessed March 20, 2022, from https://www.cisa.gov/cisa/infrastructure-data-taxonomy

[10] Cyberspace Solarium Commission. Solarium.Gov. (2020, March). *Cyberspace Solarium Commission Official Report.* Accessed March 20, 2022 from https://www.solarium.gov/report, 5.

[11] Svendsen, N. K., & Wolthusen, S. D. (2007, June). Graph models of critical infrastructure interdependencies. In IFIP International Conference on Autonomous Infrastructure, Management and Security (pp. 208-211). Springer, Berlin, Heidelberg.

[12] Svendsen, N., & Wolthusen, S. (2007, March). Multigraph dependency models for heterogeneous infrastructures. In International Conference on Critical Infrastructure Protection (pp. 337-350). Springer, Boston, MA.

[13] Kaplan, S. (1982). Matrix theory formalism for event tree analysis: application to nuclear-risk analysis. Risk Analysis, 2(1), 9-18.

[14] Sancaktar, S. (1982). An Illustration of Matrix Formulation for a Probabilistic Risk-Assessment Study. Risk Analysis, 2(3), 137-147.

[15] Tiller, M. H. (1989). Global Risk Assessment. In Risk Assessment in Setting National Priorities (pp. 297-306). Springer, Boston, MA.

[16] Iman, R. L., Helton, J. C., & Johnson, J. D. (1990). A methodology for grouping source terms for consequence calculations in probabilistic risk assessments. Risk Analysis, 10(4), 507-520.

[17] Kaplan, S. (1991). Risk Assessment and Risk Management: Basic Concepts and Terminology. Risk management: Expanding horizons in nuclear power and other industries, 11.

[18] Eaton, C. and Volz, D. (2021). “Colonial Pipeline CEO Tells Why he Paid Hackers a $4.4 Million Ransom”. The Wall Street Journal. Accessed: March 17, 2022.

[19] Suderman, A. and Tucker, E. (2021). "Major US pipeline halts operations after ransomware attack". Tech Xplore. Accessed: March 17, 2022.

[20] Barthelemy, M. (2004). Betweenness centrality in large complex networks. The European physical journal B, 38(2), 163-168.