

Vincent Tidwell, Thomas Lowry, William Peplinski[†], and Roger Mitchell
Sandia National Laboratories, Albuquerque, NM 87185

David Binning and Jenny Meszaros
AEM Corporation, Herndon, VA 20171

Abstract

Drinking water supply involves a complex network of natural and man-made infrastructure necessary to capture, store, convey, treat, and discharge this necessary resource. Each component in this “cloud to tap” supply chain faces a host of threats such as systemic decay, population change, natural disaster, cyber and physical attacks, and/or contamination incidents. Evaluating and quantifying the near- and long-term implications of these stressors on the risk and resilience of the current water infrastructure system has historically been implemented at the local utility level. The Drinking Water Resilience Project (DWRP), a collaboration between the Department of Homeland Security, Oak Ridge National Laboratory (ORNL), Sandia National Laboratories, the University of Colorado, Colorado Springs (UCCS), and the University of Tennessee (UT) is aimed at providing a more comprehensive view of water utility risk and resilience. Specifically, effort will develop an infrastructure risk assessment tool to support self-assessment by the asset owner/operator using an interactive, data-rich, web-based application that guides the user through the analysis. The associated analysis is intended to be simple, consistent and comparable. This facilitates the sharing of results, in a secure environment, across multiple levels of government as the need requires. This sharing helps place individual utility results in the broader context of risk borne by similar utilities across the U.S. Shared analysis also helps identify and address issues with assets and resources shared across multiple utilities. Most importantly, risks and mitigating measures can be prioritized across different geographic scales to aid funding decisions made at levels beyond the capacity of a single utility. A demonstration of the framework will be given along with demonstration results from several public utilities.

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Introduction

“There are 16 critical infrastructure sectors whose assets, systems, and networks, whether physical or virtual, are considered so vital to the United States that their incapacitation or destruction would have a debilitating effect on security, national economic security, national public health or safety, or any combination thereof (DHS 2016).”

The value of our nation’s critical infrastructure is particularly evident when it is compromised by disaster. The U.S. has sustained 196 weather and climate disasters since 1980 where overall damages/costs reached or exceeded \$1 billion (including CPI adjustment to 2016). The total cost of these 196 events exceeds \$1.1 trillion (NOAA 2016). Beyond such disasters, failure to maintain and grow our infrastructure is estimated to cost our economy in excess of \$195 billion per year in lost efficiencies (Sherridan and Henry, 2011). Safeguarding infrastructure from internal and external attack also looms large as indicated by cyber-attacks costing the average American firm \$15.4 million per year (Hewlett Packard 2015). Toward these challenges, President Clinton established the President’s Commission on Critical Infrastructure Protection with the task to review the vulnerabilities and threats facing U.S. infrastructures, assess the risks and propose a long-term strategy to assure the nation’s critical infrastructure in the coming decades (1997). The events of 9/11 intensified urgency in this matter with President Bush establishing, in 2003, a national strategy for the *Physical Protection of Critical Infrastructures and Key Assets* (Bush 2003).

Central to protecting our nation’s critical infrastructure is the development of methodologies for prioritizing action and supporting resource allocation decisions associated with risk-reduction initiatives. The foundation for such analysis was established in the early 1970’s in the context of the nuclear power industry (Apostolakis 2004). Following the attacks of September 11, 2001 the American Society of Mechanical Engineers was requested to develop a consistent risk assessment methodology to permit direct comparison within and across industry sectors. The Risk Analysis and Management for Critical Asset Protection (RAMCAP) process was developed around a seven-step methodology that enables asset managers to analyze their risk and risk-reduction options. Consistent with the RAMCAP framework, sector-specific applications soon followed for nuclear power plants, radioactive waste transportation and storage, petroleum refineries, chemical manufacturing plants, LNG off-loading terminal, dams and locks, and water and wastewater systems. The *2002 Public Health Security and Bioterrorism Preparedness and Response Act*, which required all water utilities serving more than 3,300 people to perform security vulnerability assessments, accelerated application within the drinking water sector. A sector specific framework was developed by the American Water Works Association, J100 standard for Risk and Resilience Management of Water and Wastewater Systems. Specific applications were also adapted including the Vulnerability Self-Assessment Tool (VSATTM), the Security and Environmental Management System (SEMSTTM), and the Risk Assessment Methodology-Water (RAM-W).

The difficulties in developing a risk-based framework for prioritizing risk-reduction actions are daunting, largely due to the great uncertainties in understanding the suite of threats. In 2008 the U.S. Congress asked the National Research Council (NRC) of the National Academies to review and assess the activities of the Department of Homeland Security related to risk analysis (P.L. 110-161, *Consolidated Appropriations Act of 2008*). While they found the conceptual framework for

risk analysis (risk is a function of threat (T), vulnerability (V), and consequence (C), or $R = f(T, V, C)$) to be appropriate for decomposing risk and organizing information they questioned its ability for supporting decision making because its validity and reliability were untested (National Research Council, 2010). One area in which such testing is lacking is related to evaluation of the inter-comparability of results across different utilities. Specifically, to what extent does bias in the utility-centric risk assessment impact the inter-comparability of results with other utilities. Bias is introduced through over or underestimating the probability of a threat, the vulnerability to a given threat, failing to identify a critical threat, or through differences in calculating associated consequences. Such bias could significantly skew results leading to mis-prioritization of action.

Here we present a web-based risk assessment framework that promotes the anonymous sharing of results among utilities of similar character. The constructed framework was demonstrated for three water utilities. Results were compared across utilities and were also combined with risk assessment results from four other utilities collected using a different risk assessment application by a different set of analysts. Comparison of the results identifies five values realized by a shared risk assessment framework:

1. Helps recognize and correct bias in analyses,
2. Helps recognize “unknown, unknowns”, that is, helps analysts identify threat, vulnerabilities, or consequences they would otherwise have overlooked,
3. Provides a means of self-assessment and benchmarking for the local utility
4. Provides opportunity to expand analysis to include shared assets and/or threats across multiple utilities, and
5. Helps prioritize actions beyond the scale of a single utility.

Below we discuss the basic approach taken to developing the framework and conducting the utility demonstrations. Results are then reviewed, focusing largely in understanding the differences and biases across the seven utility risk assessments. Finally, the results are discussed with respect to the value of the shared framework as well as potential issues with sharing of sensitive data.

Methods

Framework

The DWRP framework is based on systems theory (e.g., Maani and Cavana, 2002) and uses the concepts of impacts, systems, and threats as its basis for implementation. This is in contrast to the J-100 standard that uses consequences, assets, and threats. The systems’ approach helps promote the idea of a water utility as a system of systems, where the underlying systems are a collection of parts that interact with one another to function as a ‘whole’. It is a top-down approach that changes the fundamental question being asked when performing a RA from “what happens if ‘Asset A’ fails” to “what impacts are we most worried about and what systems need to fail to produce that impact?” (Figure 1). This approach also supports rapid execution of the RA by reducing the tendency to focus on details that are not important for a RA of this type.

While the concepts of impacts, systems, and threats are analogous to consequences, assets, and threats, there are important distinctions. Impacts within the DWRP framework are calculated using a minimal amount of input from the user. This is to standardize and normalize the impacts across different utilities and regions to allow for meaningful comparisons and identification of bias. Systems can be thought of a collection of assets that collectively perform a vital role in the collection, treatment, and delivery of water to the consumer. Likewise, threats within the DWRP

framework are consolidated at a coarser resolution and are meant to describe the broad categories of threats to a utility as opposed to specific descriptors that are part of the J-100 standard.

The intent is to remain consistent with the J-100 standard while offering a means of normalizing data so that results from different utilities can be consistently compared. Ideally, one would have already completed a J-100 RA and then port results of the highest risk features into the DWRP framework. It is important to note that the DWRP framework is not meant to be a substitute for a J-100 RA but rather a complement to a J-100 RA that will provide greater insight into the underlying risks and aid the user in identifying previously unknown risks or risks that are the result of shared systems and/or shared threats.

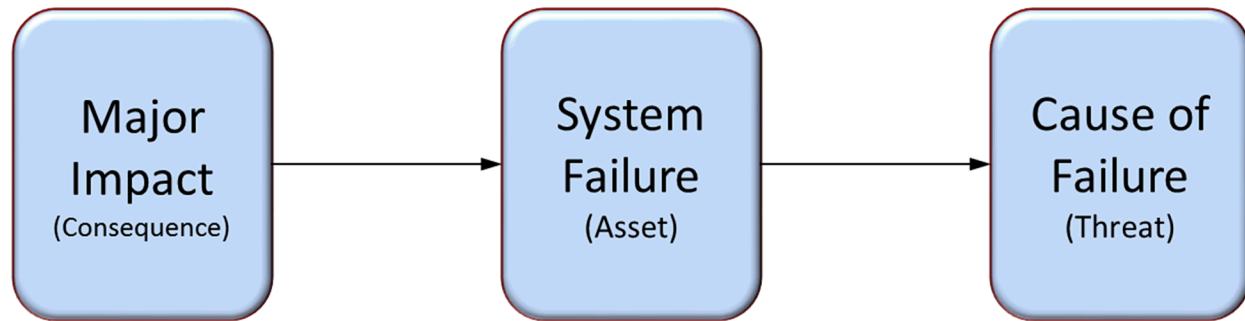


Figure 1. The systems approach is used to put an emphasis on high level assessments that examine the consequences and impacted systems first.

Systems

The systems defined in the DWRP framework are listed in Table 1. This framework was implemented within a user-friendly, interactive software application. Within the list is the entire working structure of a water utility, including the non-physical aspects such as employees and/or knowledge base. When a user enters a system into the framework, they are presented with a dropdown list of Table 1 to categorize the system as well as a blank line to enter a more specific description. For instance, a user may select 'Treatment' as the system, but enter the name of a key pump within the treatment plant as the descriptor. This aids in comparing across utilities in that it is the system outage and the impacts of that outage that are important as opposed to the details of the failure within a system (i.e., at the comparison level, it is more relevant to know that a treatment plant is out for 7 days as opposed to a specific pump within the treatment plant is out). It also provides another level of protection for the individual utilities in that vulnerable assets cannot be identified by users outside their own utility. From the utility point of view, placing the focus at the systems level allows for a more rapid execution of the RA by avoiding the details of a J-100 assessment while also encouraging the user to look at their utility from a different perspective to better identify and prioritize the highest risks.

Table 1. List of systems defined in the framework.

System Name	Description
Employees	Ability to get to work, labor market with adequate skill set
Finished Water Distribution	Distribution between final treatment and the consumer
Information Technologies	SCADA systems, analysis, monitoring, data

Knowledge Base	Experienced employees, institutional knowledge
Maintenance and Administration	Administration services, building maintenance, computer systems
Operations and Maintenance	Water system operations and maintenance
Raw Water Conveyance	Conveyance between source water intake and treatment plant
Source Water	Water quality and availability
Source Water Infrastructure	Systems to collect water from source
Storage	Reservoirs and tanks for storage of raw, treated, or used water
Treatment	Treatment facilities for treating raw water
Other	Everything else

Impacts

Impacts within the DWRP framework are split into three categories; community disruption, health and safety, and financial (Table 2). Community disruption costs are the costs borne by the local community due to a failure of the utility to deliver water. The user is responsible for four inputs when describing the community disruption costs; their total number of connections, the outage time, the percent of their total demand served by the failed system, and the percent of the unmet demand caused by the failure. The calculation is based on the metropolitan gross domestic product (expressed as GDP/person/day), which is predefined in the framework by metropolitan area. The use of the metropolitan GDP normalizes the community impacts across regional and national scales.

Health and safety impacts are impacts that cause deaths or illnesses to consumers, employees, and/or other users of the utilities systems (e.g., boaters on a utility-owned reservoir). The H&S impact calculation is the number of deaths times the value of a statistical life (VSL) plus the number of illnesses times the value of a statistical injury/illness (VSI). The VSL and VSI (\$9.4 million and \$0.94 million, respectively) are the most recent values recommended by the U.S. Department of Transportation (Thomson and Monje, 2015).

Financial impacts are those impacts borne directly by the utility and include lost revenue, repair costs, other costs. Lost revenue is revenue for the utility that is lost from not being able to deliver water or from an inability to bill the customer (such as might occur if a customer database became corrupted or a billing system failed). Repair costs are the total costs for bringing the failed system back to working order, including equipment costs and labor costs. Other costs are meant to capture things like legal costs, loss of customer confidence, and the like. To calculate the financial impacts, customers are required to enter the total number of connections, the outage time, the percent of the total demand served by the failed system, the percent of the unmet demand caused by the failure, the average daily water service (in millions of gallons per day), and the average water rate (usually dollars per 1000 gallons).

Table 2. Impact categories. Variables refer to: t_{out} = outage time [days], n_{cust} = number of hookups, D_S = % of total demand served, D_U = % unmet demand, GDP = metropolitan GDP [GDP/person/day], n_D = number of deaths, VSL = value of a statistical life, n_I = number of illnesses or injuries, VSI = value of a statistical illness/injury, S = average daily service [MGD], r = average

water rate [\$/1000 gal], R_c = repair costs [\$], O_c = other costs [\$.]. Variables in red are supplied by the user. Other variables are provided by the framework. Note that the inputs needed to calculate the community disruption impact are also used to calculate the financial impact.

Impact Category	Description	Calculation
Community Disruption	Costs borne by the local community	$I_{CD} = t_{out} \times n_{cust} \times D_s \times D_u \times GDP$
Health and Safety	Deaths and illness or injury	$I_{HS} = n_D \times VSL + n_I \times VSI$
Financial	Costs borne by the utility	$I_f = t_{out} \times n_{cust} \times D_s \times D_u \times \bar{S} \times r + R_c + O_c$

A key difference between the DWRP framework and the J-100 standard is the manner in which vulnerability is handled. Vulnerability within J-100 can be described as the percentage of time a threat creates an impact once it occurs. A classic example is a direct attack on a facility where the probability of the attack is different than the probability of an attack being successful. A heavily fortified facility may have a high probability of attack (i.e., it often gets attacked) but a low vulnerability (the probability of an attack being successful is low). The DWRP framework assumes that for natural threats, the vulnerability will be one. In other words, we are only interested in natural threats that produce some kind of impact. For other threats, such as direct attack, the vulnerability is included in the D_u variable, the percent unmet demand.

Threats

Like the list of systems, threats have been consolidated to a coarser resolution than used in the J-100 standard to aid in the rapid execution of the RA and to keep the focus at a higher level. Within the DWRP framework, the user chooses a threat based on a dropdown list of the threats listed in Table 3 and then provides a more detailed description if desired. This again helps protect individual utilities since it is only the threat category and not the detailed description that is compared across utilities (e.g., there is a big difference between knowing the risk is from a direct attack on the water distribution system, and that it is from an armed attack at a specific pump station).

Table 3. List of threats used in the DWRP framework.

Aging Infrastructure	Natural – Drought
Contamination	Natural – Earthquake
Direct Attack	Natural – Flood
Human Error	Natural – Hurricane
Loss of Customers	Natural – Ice Storm
Loss of Employees	Natural – Tornado
Loss of Suppliers	Natural – Tsunami

Loss of Utilities	Natural – Wildfire
Other	Sabotage – Cyber
	Sabotage – Physical

Probabilities for the natural threats are automatically provided to the user based on the location of the utility while the other threats use default values based on historical values. The user can override the probability of a threat occurring if the default value seems unreasonable or to perform sensitivity analysis or scenario testing.

Implementation

To execute the DWRP framework, a user begins by inputting descriptive information about their utility, including the location, number of customers (n_{cust}), average daily service (S), and the average charge rate (r). Then, the user inputs the system-threat pairs. When inputting the system-threat pairs, it is important to note that a single system may have many threats and that a single threat may impact many systems. The framework is setup to handle this. For each system-threat pair, the variables needed to calculate the impacts are also entered.

Once entered, the system calculates the system-threat risk within each impact category and then uses a matrix approach to calculate other risk such as the total risk to a system across all threats, total risk from a threat across all systems, or total risk to the utility across all systems and threats. Mathematically, the system threat risks for the community disruption, health and safety, and financial impact categories are calculated using the appropriate version of equation [1]:

$$\begin{aligned} (R_{CD})_{i,j} &= (I_{CD})_{i,j} \times p_j \\ (R_{HS})_{i,j} &= (I_{HS})_{i,j} \times p_j \\ (R_F)_{i,j} &= (I_F)_{i,j} \times p_j \end{aligned} \quad [1]$$

where R is the risk, I are the impacts calculated from Table 2, and p is the probability of the threat occurring. The subscripts, i,j , refer to the system and threat, thus R represents the risk to system i from threat j .

The total risk to a system is then calculated by summing across all risks for that system:

$$R_i = \sum_{j=1}^{n_T} (R_{CD})_{i,j} + (R_{HS})_{i,j} + (R_F)_{i,j} \quad [2]$$

where R_i is the total risk to a system across all threats, and n_T is the number of threats. Likewise, the risk to a utility from a single threat is calculated as:

$$R_j = \sum_{i=1}^{n_S} (R_{CD})_{i,j} + (R_{HS})_{i,j} + (R_F)_{i,j} \quad [3]$$

where R_j is the total risk to the utility from a single threat and n_S is the number of systems. The total risk to the utility across all system threat pairs is calculated using:

$$R_{tot} = \sum_{i=1}^{n_S} R_i \quad [4]$$

As will be discussed in the results section, this ability to look at risk from the system, threat, or utility level provides a means of identifying bias when compared to other utilities and helps the user better understand their key vulnerabilities and risks.

Utility Demonstrations

In person utility demonstrations were conducted using the DWRP Framework to get feedback on the value of the approach, the usability of the interface, and to obtain real world data to run the comparisons. The demonstrations were conducted for three different water utilities, one situated in the Central U.S., one in the South, and one in the West. The size of the utilities varied from 65,000 to 900,000 customers served and average daily deliveries ranged from 32 to 100 MGD. Two of the participating utilities had previously completed a full-scale J-100 risk assessment of their system using VSAT. The third utility has a designated team that continuously performs risk assessments as part of their asset management process. The goal was not to replicate their past risk assessments but rather to map the highest level concerns into a uniform framework that allowed comparison with other utilities.

During the in-person demonstrations, data from their past risk assessment exercises such as individual threat-asset pairs, were entered into the framework. Not all asset-threat pairs were entered, rather only those that are of greatest concern to the utility—only those that “keep the managers up at night.” As described above, consequence and risk are automatically calculated by the framework. The face-to-face demonstrations took less than a day to complete in each case.

To supplement the face-to-face demonstrations, data were obtained from four other J-100 risk assessments. Two utilities were from the Central U.S., one from the South and one from the West. Populations served and average daily deliveries ranged 135,000 to 460,000 and 33 to 120 MGD, respectively. All four assessments were conducted by the same consulting firm working directly with personnel from each of the four utilities and were provided to the research team in a sanitized manner to maintain anonymity. The data (e.g., threat-asset pairs, disruption times, damage costs) were input into the framework where the risks were calculated.

Results

The three in-person demonstrations and the four anonymous data sets provide seven real-world estimations of risk using the DWRP Framework. In the discussions and figures that follow, the four anonymous data sets are referred to as East 1, East 2, Central 1, and West 1, while the three in-person demonstration utilities are referred to as Central 2, South 1, and West 2.

Consistency

The face-to-face demonstrations provided insight into how familiar each utility was with the risk assessment process as well as the results of their own risk assessment. For the two utilities that had completed a J-100 RA, it was apparent that the personnel in the room were generally familiar with the RA process but not so familiar with the details of their results. In both cases, the people who actually performed the RA were present but were not able to quickly and easily find the data necessary for input into the DWRP Framework. The third face-to-face demonstration had the

benefit of several employees who worked full-time doing RA and RA related tasks so their knowledge of the required data was much greater than the other utilities.

The lack of familiarity of the results of their own RA's offers an interesting conclusion: for utilities that lack the resources to address and consider risk on a regular basis, the value of performing a RA may be lost or greatly diminished. This is not to say that their VSAT RA's were not completed well, they were (and one could argue that they were completed very well), but rather that if the utility lacks the means to analyze the results and translate those results into useful action, the value of that RA is lost.

Another interesting aspect of the risk assessments was the variability in how the probabilities and consequences were calculated. Two utilities categorized the probability of a threat occurring (e.g., low, medium, high, very high) and one of those also used set criteria to categorize the consequence on a scale of 1 to 5, with 5 being 'extreme' consequences. The latter approach allows for the creation of a risk table that ranks risk based on where it falls in the table (Table 4). While this approach is useful for identifying assets with high risk, it does not allow for prioritizing within a risk category or for comparing risk from different utilities. For the utilities that used this approach, a 'translation' between each category and its numerical value had to be agreed upon before their data could be put in the DWRP Framework.

A final observation from the face-to-face demonstrations is that the perception of risk, and more specifically, the perceptions of the threats and their likelihood, varies considerably from utility to utility. As a rough generalization, one utility emphasized aging infrastructure and community disruption risk, another emphasized operations and financial risk, and the third emphasized security and financial risk. The reason for this stems from the set of issues that a utility is facing at the moment or has faced in the recent past. For instance, the utility that emphasized security had documented an attempt to hack their computer systems. Their work to address cyber security creates a natural tendency to look at other security issues within their system. This type of dynamic played out in varying degrees across all three demonstrations.

Table 4. Example of risk assessment ranking table using categories of probability (likelihood) and consequence.

		Consequence				
		1	2	3	4	5
Likelihood	1	Very Low	Very Low	Low	Medium	Medium
	2	Very Low	Very Low	Medium	Medium	High
	3	Low	Low	High	High	Extreme
	4	Low	Medium	High	Extreme	Extreme
	5	Low	Medium	High	Extreme	Extreme

Data Comparison

Error! Reference source not found. shows the cumulative probability distribution of risk for each utility across each impact category as well as for the total risk. This shows that probability that the risk associated with a system-threat pair will be lower than a given risk. The plot for the

health and safety risk has only three utilities because the other utilities reported no risk in that category. Note that the x-axis (risk axis) is logarithmic, indicating a wide range of reported risks.

From the plots, it is evident that the distribution of risk is similar across the utilities although outliers can be seen. West 1 has the highest maximum risk for a single system-threat pair at \$209 billion, which was for a direct attack on an administration building that caused a significant number of deaths and injuries. This is reflected by the high maximum value for West 1 in the Health and Safety plot. The lowest maximum total risk is \$216,000 for Central 2, for physical-sabotage on their source water intake that causes a community disruption.

South 1 and West 2 show most of their risk is to the right of the blue dotted line, which represents the average across all utilities, indicating that they may be over estimating their impacts and/or probabilities as compared to the other utilities. This is clearly driven by the financial risk as South 1 and West 2 are also to the right of the 'All Utilities' line in that plot. On the other end of the spectrum, East 2 is to the left of the 'All Utilities' line in the Financial Risk plot indicating that they may be underestimating their repair costs or other costs. Again, the goal with the comparisons is to allow a utility to look at other risk assessments in a manner that may cause them to rethink, or at least, re-check their assessments.

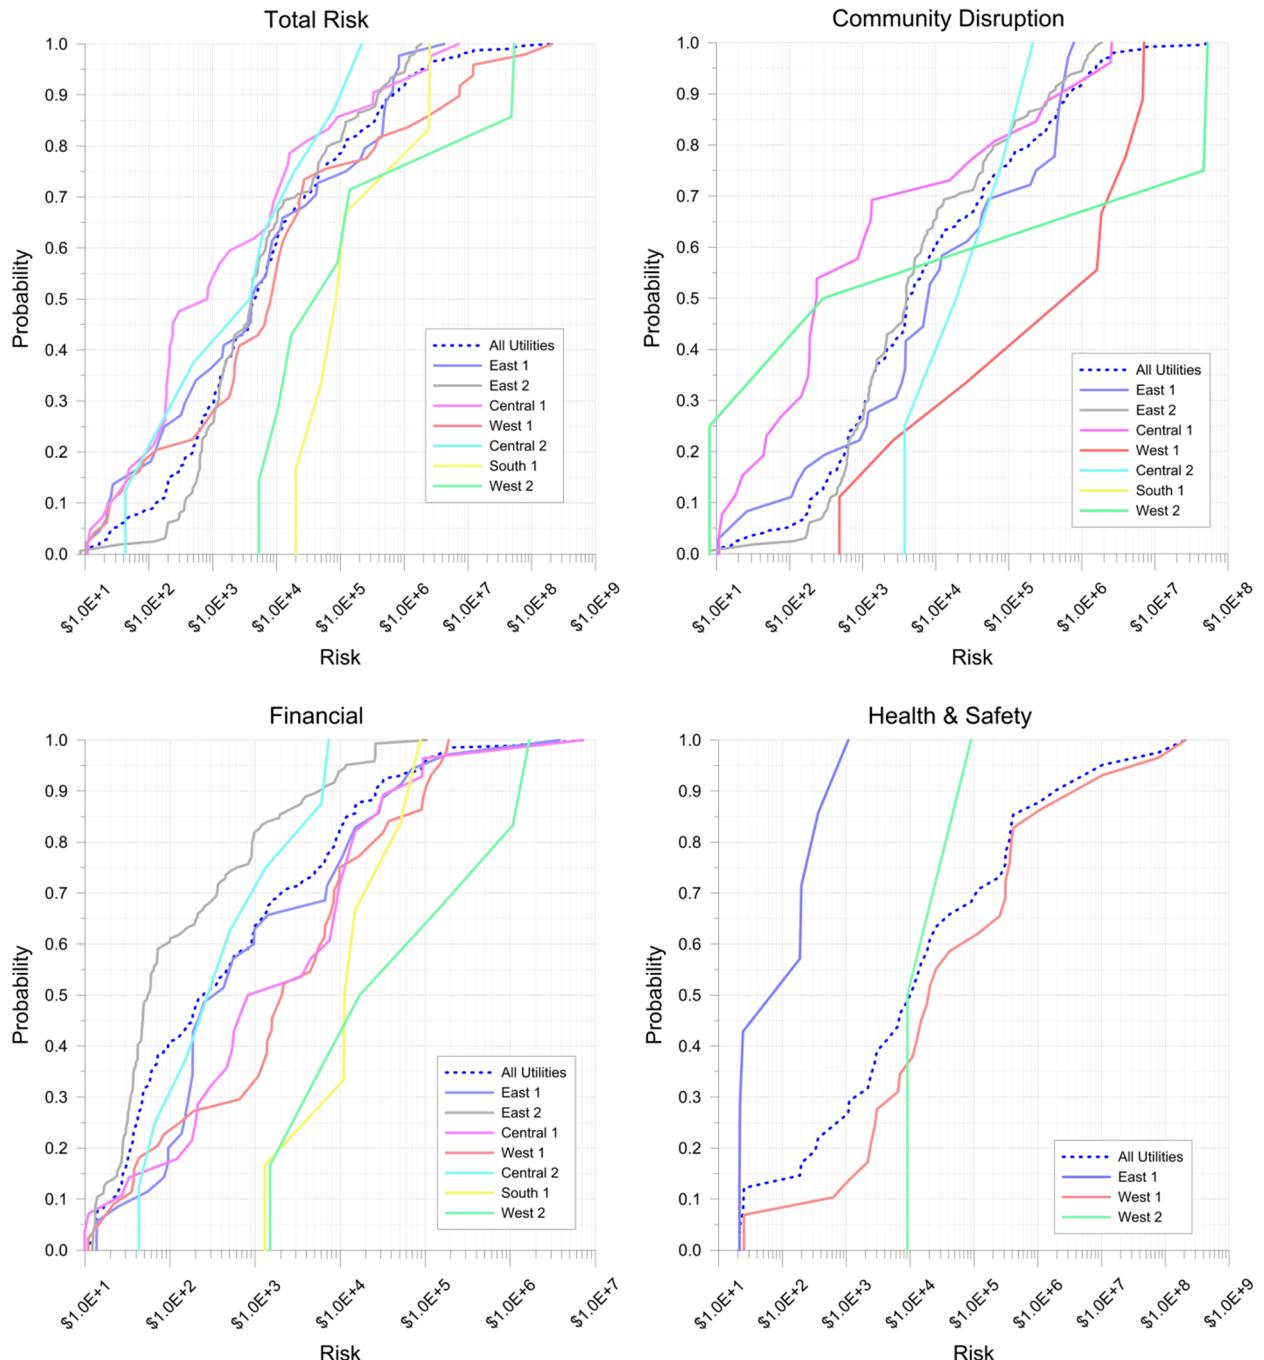


Figure 2 - Cumulative probability distributions for the risk associated with each impact category and the total risk. The 'All Utilities' line represents the average risk across all utilities.

Another way to compare the utilities is to look at the percentage of the total number of occurrences across all utilities that a utility addresses a particular system or threat (Figure 3). The percentages are the percentages of the total. For example the first bar in the left hand plot of Figure 3 shows that the risk assessment for West 1 contains approximately 67% of the system-threat pairs that include employees as the system, with West 2 containing the other 33%. The other utilities do not address risk for the employee system category. This concept is the same for the right hand plot too.

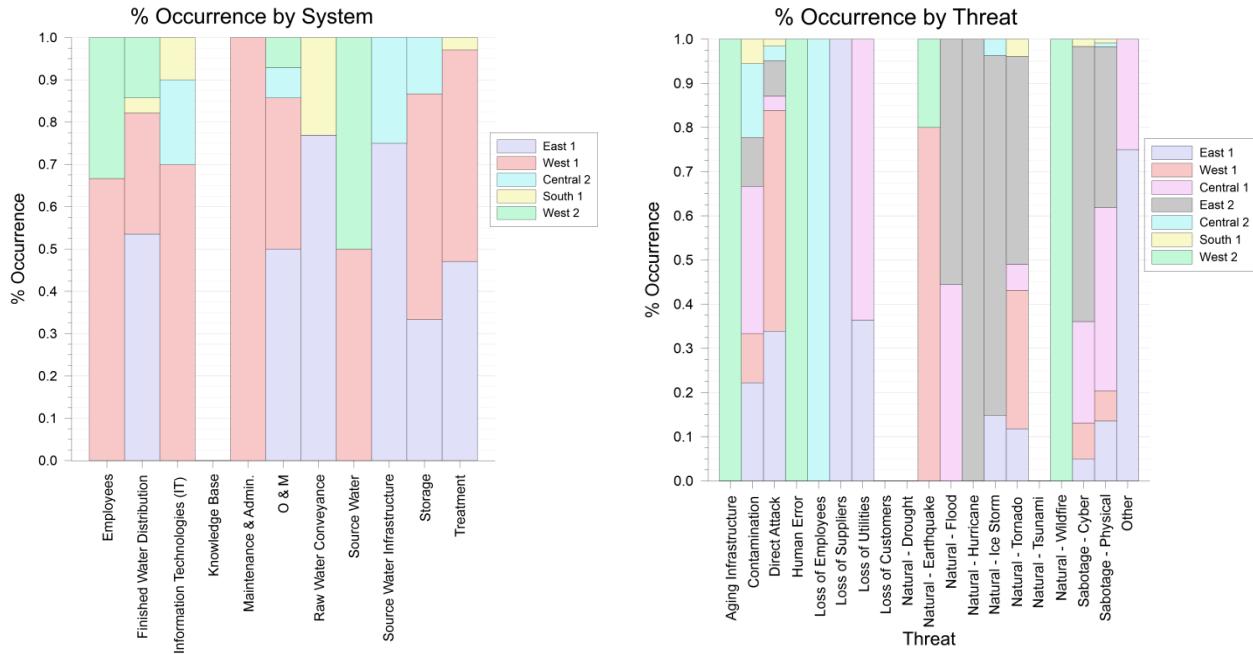


Figure 3. Percentage of occurrence for systems and threats. The percentage is based on the total number of occurrences across all utilities.

By looking at these plots, a utility may realize that they have missed a key system or threat. For instance, on the threat side, East 1 is the only utility to address the risk associated with a 'Loss of Suppliers'. Noting this may cause another utility to address this when they otherwise might have overlooked it.

Figure 4 is the same as Figure 3 but the percentages are based on the total risk across all utilities and not the number of occurrences. Referring back to the West 1 and West 2 case for the Employee system category, one can note that while West 1 contains 67% of the occurrences of the Employee category, West 2 (from Figure 4) contains 98% of the risk. A similar dynamic can be seen with respect to the Flood threat where East 2 contains approximately 55% of the instances where Flood is listed as a threat, but only 2% of the risk associated with flooding.

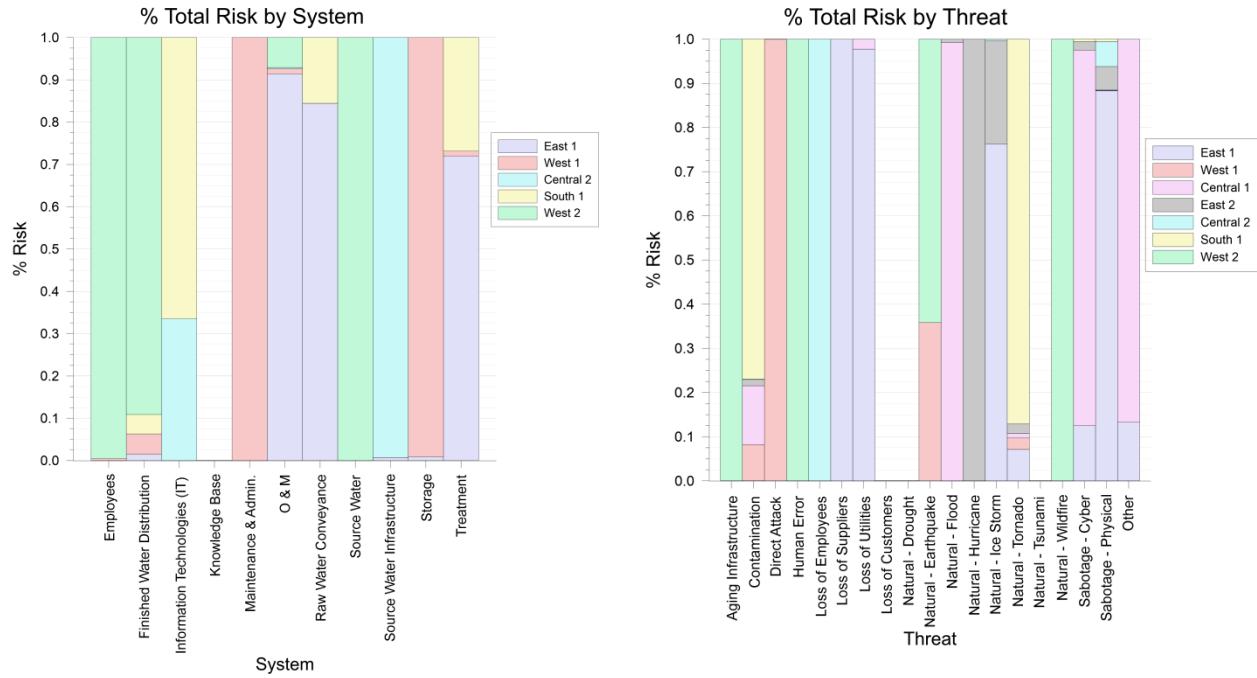


Figure 4. Percentage of total risk for systems and threats. The percentage is based on the sum of total risk across all utilities for each system and threat.

Discussion

Bias

The risk assessments conducted for the seven water utilities yielded comparatively different results. Noted were differences in the overall distribution of risk (comparison of the CDFs), the mix of assets and threats considered and the priority risks identified by each utility (highest ranking threat-asset pairs). There are two basic reasons for the noted differences. First, some of the differences are real given the unique age, design and complexity of their infrastructure, their geographic location and thus exposure to different natural threats, prior efforts to harden their utility against identified threats, and the risk culture of the utility. Second, some difference was introduced into the analysis by bias. Bias resulted from the use of different analytical platforms, different analyst with differing perceptions of risk, and the limited knowledge and experience of the analysis team resulting in the failure to capture the full spectrum of risk.

Ideally, statistical analysis could be used to distinguish bias from the actual risk profile; unfortunately, a population of seven utilities was insufficient to accomplish this. Although there was insufficient data to statistically distinguish bias from the actual risk profile, results bear out the fact that bias exists. For example, the number of threat-asset pairs significantly differs across the seven utility assessments, thus biasing the lower tails of the distributions. Comparison of the threat-asset pairs for each of the utility RAs also reveals bias. Most systems (e.g., asset) and threats were included in three or fewer of the utility RAs although these systems and threats (except in the case of specific natural disasters) are largely common to all seven utilities. Also, several high priority risks identified by a single utility are not recognized by other utilities even though they are likely to have application. Examples include the threat to employees from disgruntled customers and the passage of trunk water lines below important transportation corridors (i.e., highways or rail lines) that would have significant economic implications if the pipe failed.

Value of Shared Analysis

A web-based risk assessment framework that promotes the anonymous sharing of results among utilities of similar character offers a number of potential advantages, among these are assistance in recognizing and correcting bias; identification of “unknown, unknowns”; self-assessment and benchmarking for the local utility; treatment of shared assets and/or threats across multiple utilities; and prioritization of actions beyond the scale of a single utility.

Comparison of results across the seven water utilities clearly indicated the influence of bias on the analysis. This bias was the product of using different risk analysis applications as well as differences in risk perception among the various participants in the analysis. While it is recognized that bias is inherent to the process, there are steps that can be taken to minimize its effects. A single standardized web portal for guiding the RA, while being compatible with existing tools and processes (e.g., consistent with the J100 process and associated applications like VSAT or SEMS), would provide a consistent and unbiased procedure for managing the analysis process. It would also guarantee consistency in assumptions concerning threat probabilities, vulnerability estimates and quantification of consequences (much like the VSAT or SEMS applications provide for the water sector). Anonymous sharing of RA results would also provide a means of comparison, allowing analysts to identify where their results deviate from that of other similar utilities and thus evaluate whether the discrepancies are real or bias. As the database of utilities grows standardized statistical tests could be developed for automated identification of potential instances of bias.

Another value of shared RA is the identification of unknown, unknowns; that is, threats, vulnerabilities, or consequences that would otherwise have been overlooked due to lack of experience or perception of the local team of analysts. An example involves the threat of a disgruntled customer attacking the billing office. While this is a threat that most all utilities would face only one utility identified this threat because they had experienced an upset customer bringing a gun into the administrative offices. Anonymous sharing of risk profiles would help utilities identify these unknown, unknowns by learning from the experiences of other utilities. As the database grows search algorithms could be developed to target and share such events with particular utilities, much like Amazon pairs potential customers with new products based on their buying habits.

Shared RA also provides an effective means of self-assessment and benchmarking for the local utility. While the individual utility can get an indication of high versus low risk, utility leadership needs to place their results in context before deciding what, if anything, they need to do to change their situation. A national-normed view of risk would permit them to gauge how well (or poorly) they are doing relative to others in their cohort. This desire to know “where they stand” is not only an interesting fact, but can become a powerful driver when used by the utility to brief decision-makers when asking for the resources to reduce their high-priority risks. Decision makers understand high risk and typically become clearly motivated when this risk places them in an unfavorable light with respect to their peers. As actions are taken to reduce risk such nationally-normed comparison gives the utility a benchmark against which they can measure improvement.

Shared RA provides benefits beyond the scale of the single utility. Risk profiles aggregated at the metropolitan, county, state and national level would assist planning, prioritization and emergency preparedness beyond the community level. Such cooperation would help identify shared assets such as reservoirs or raw water conveyance systems. Associated risk profiles would change as the values of multiple utilities are aggregated, more accurately reflecting the value of the asset. Aggregation of

risk across shared threats, particularly natural disasters, would likewise add value to regional emergency preparedness efforts. Equally important, sharing of risk data across different geographic regions would assist in effective prioritization of risk reduction actions. This is particularly important in the case of actions requiring resources beyond the capacity of the local utility, providing a means of targeting state and national assistance to the greatest need.

Challenges to Broad Implementation

While the value of a shared RA framework is apparent, there are several important concerns related to implementation. In particular is the perceived need for this tool relative to that of other initiatives. Threats to infrastructure systems are known and well understood by the utilities. Infrastructure managers deal with risk on a daily basis and are accustomed to the mental gymnastics to effectively deal with these risks and continue to protect the public. It is always the emerging risks that keep the sector awake at night. Thus any new tool needs to advance the research and application development of risk if it is to add value to the risk and resilience understanding and sustainability improvement to the utility. Utility leaders universally are supportive of investments that enhance the quality and safety of their systems. But history has taught that protection demands constant search for the next threat that must be handled. Without a clear sense that a new process or procedure is forward looking and adds to the basic mission of the utility, it will receive little enthusiasm or support.

There must be a clear benefit to the utility to participate. Most utilities today are under-staffed and over-worked. The flood of daily operating problems is exacerbated by ever growing governmental regulation, rising customer expectations, and an unrealistically low rate structure that cannot keep pace with the deterioration of the infrastructure. The post-911 vulnerability assessments that were mandated by EPA of all water systems with over 3,300 connections carried with it a \$200,000 stipend. Today we are asking water utilities to complete a probability-based risk assessment using their own funding. Thus, for the average utility to invest the time and resources necessary to produce a meaningful, thoughtful risk assessment, there must be something tangible in it for them.

Utilities have a long and vocal history of refusing to share information, especially with governmental entities that may not have the most robust of track records for protecting this sensitive information. For a utility to even consider participating in a risk assessment program in which their data would be exposed in any form, evokes questions such as:

- a) What is the clear and convincing process that their data is to be protected?
- b) Who will have access to the data?
- c) Will the utilities have the option to control who has access to their data?
- d) Will the utility be notified whenever anyone accesses their data?

Without convincing reassurances, there is only limited hope that utilities will feel comfortable to sharing their data. There is little trust that any governmental entity is in the position to protect this information from political misuse. If there is to be a repository of risk assessment information that is specifically tied to any utility, it must be with a trusted agent. Proper protocols must be developed and then translated through strong software design to assure the guarded access to this information. When accessed by state or federal governments, the agency accessing the data must be authorized and then be limited to only that data specifically needed to meet an agreed program. Finally, the owner of the information must have a way to be informed of who accessed their data, when it was accessed, how much was accessed, and to what purpose the data will be applied.

Summary

Here we demonstrate a web-based risk assessment framework that promotes the anonymous sharing of results among utilities of similar character. The constructed framework was demonstrated for three water utilities. Results were compared across utilities and were also combined with risk assessment results from four other utilities collected using a different risk assessment application by a different set of analysts. Although there was insufficient data to statistically distinguish bias from the actual risk profile, results bear out the fact that bias exists among the seven risk assessments. Bias resulted from the use of different analytical platforms, different analyst with differing perceptions of risk, and the limited knowledge and experience of the analysis team resulting in the failure to capture the full spectrum of risk.

A web-based risk assessment framework that promotes the anonymous sharing of results among utilities of similar character offers a number of potential advantages, among these are assistance in recognizing and correcting bias; identification of “unknown, unknowns”; self-assessment and benchmarking for the local utility; treatment of shared assets and/or threats across multiple utilities; and prioritization of actions beyond the scale of a single utility. While the value of a shared RA framework is apparent, there are several important concerns related to implementation. Without a clear sense that shared RA is forward looking and adds to the basic mission of the utility, it will receive little enthusiasm or support as utilities are constantly bombarded by new regulation and risk reduction initiatives. Beyond this there must be a clear benefit to the utility to participate, as utilities today are under-staffed and over-worked. And, finally there must be demonstrated ability and assurance that utility data will be protected from unintended access or misuse.

Acknowledgements

This work was funded by the Science and Technology directorate of the Department of Homeland Security. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

References

Apostolakis, G.E., 2004. How useful is quantitative risk assessment, *Risk Analysis*, 24(3), 515-520.

Hewlett Packard, 2015. 2015 Cost of Cyber Crime Study: Global, 29 pages.

Maani, K.E., and R.Y. Cavana. 2002. *Systems Thinking and Modelling: Understanding Change and Complexity*. Auckland, New Zealand: Pearson Education New Zealand Ltd.

National Research Council, 2010. Review of the Department of Homeland Security's Approach to Risk Analysis, The National Academies Press DOI 10.17226/12972.

NOAA National Centers for Environmental Information, 2016. Billion-Dollar Weather and Climate Disasters: Overview, at <https://www.ncdc.noaa.gov/billions/>

President's Council on Infrastructure Protection, 1997. The Report of the President's Commission on Critical Infrastructure Protection.

Sherraden, S. and Henry, S., 2011. New America, Costs of the Infrastructure Deficit, March 2, 2011, at <https://www.newamerica.org/economic-growth/policy-papers/costs-of-the-infrastructure-deficit/>

Thomson, K., and C. Monje. 2015. Guidance on Treatment of the Economic Value of a Statistical Life (Vsl) in U.S. Department of Transportation Analyses - 2015 Adjustment. edited by U.S. Department of Transportation. Washington, DC: Office of the Secretary of Transportation.