

Digital Sandbox Whitepaper

Risk Analysis and Capabilities Assessment

Risk analysis is an integral part of an overall risk management program, but understanding risk does not necessarily suggest a recommended course of action to mitigate it. Digital Sandbox employs a “capability build-up” methodology to using risk to identify gaps and inform investment priorities.

RISK BUY-DOWN

The primary policy-based approach used today is to align investments against risks. It is the approach taken by FEMA in allocating funds for the Homeland Security Grant Program and other risk-based grant programs. While useful for high-level direction setting, the risk buy-down approach suffers from two practical issues. First, it requires a risk analysis methodology (e.g., a risk equation) that is sensitive to the effects brought about by risk mitigation investments, which is not the case in current FEMA formulas. Its second practical issue is that creating a mitigation-sensitive risk formula is exceedingly difficult. While this approach may work for deciding between two or three similar proposed investments, it is probably not a cost-effective method for setting overall risk management priorities.

CAPABILITY BUILD-UP

A better approach than risk buy-down is to use risk analysis to drive risk management through an analysis of capabilities. The risks facing a jurisdiction or sector can (and do) suggest a prioritized set of capabilities that should prove most effective at mitigating those risks. Investments that build up those key capabilities should therefore have the highest “risk return.” Coupled with an ability to measure the overall amount of capability one has (capability assessment), this approach can be an effective tool for managing risk.

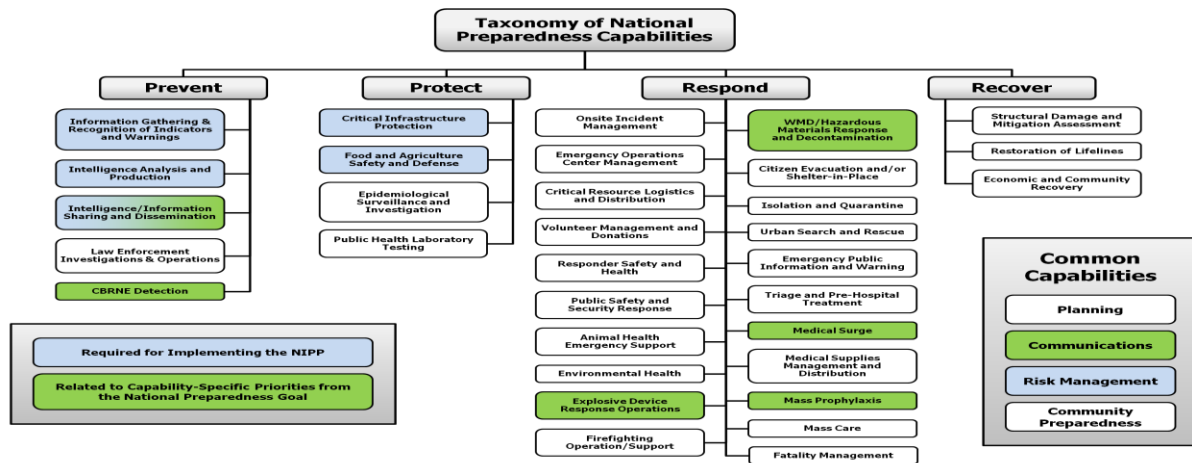
CAPABILITY TERMINOLOGY

This section lists key capability assessment terms along with their working definitions,

CAPABILITIES

The *National Preparedness Guidelines (NPG)* document, published by DHS, calls for the establishment of a *Target Capabilities List (TCL)* that lists those capabilities a jurisdiction might possess to mitigate terrorist risks. According to the NPG, “Capabilities provide the means to accomplish a mission and

achieve desired outcomes by performing critical tasks, under specified conditions, to target levels of performance.” The latest version of the TCL, published on September 13, 2007, includes 37 capabilities grouped into four *Mission Areas* plus a fifth set of capabilities common to all mission areas. **Error! Reference source not found.** shows all 37 capabilities on the TCL grouped into the four mission areas and the common capabilities. The colors show those capabilities that are specifically called out in the National Infrastructure Protection Plan (NIPP) and the NPG.



Target Capabilities List

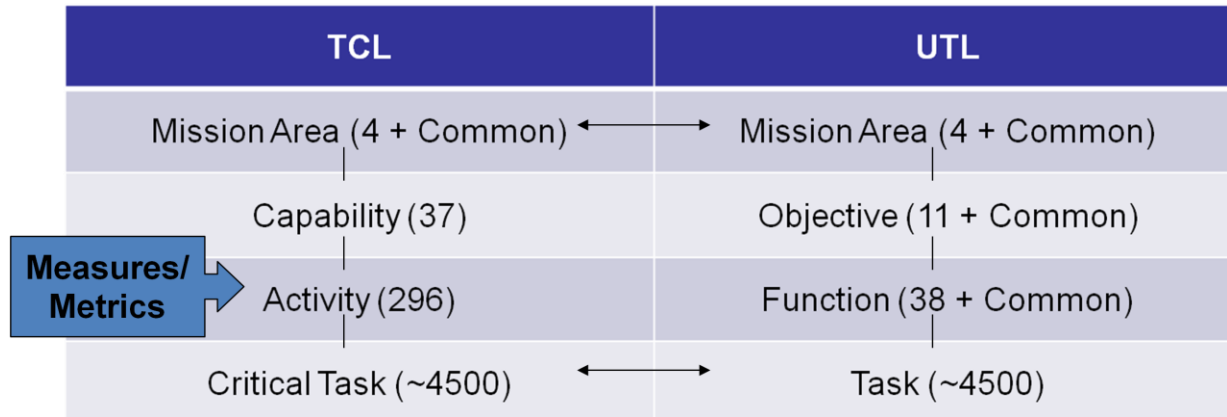
CAPABILITY LEVEL

Capability level refers to the amount of proficiency in a given capability, which is given in units that are appropriate for that capability (e.g., quality, capacity, number, response speed, etc.). The capability level may be specified for an individual asset, a sector, a jurisdiction, or any larger entity. For convenience, “higher” levels of capability refer to greater proficiency (i.e., more units, better-equipped teams, lower response times, etc.), and “lower” levels of capability refer to lesser proficiency.

CAPABILITY COMPLIANCE

Capability compliance is a proxy for measuring capability levels, using the metrics from the TCL. The TCL defines a taxonomy of concepts, with a hierarchy that starts at *mission areas* and goes all the way down

to individual *tasks* to achieve the mission. There is a related document, also called for by the NPG, called the Universal Task List (UTL). The UTL also defines a taxonomy of activities beginning at mission areas and ending at individual tasks, but the intermediate steps have no connection to the intermediate levels in the TCL. **Error! Reference source not found.** shows the taxonomies of the TCL and UTL, showing how they align at the highest and lowest levels, but not in the intermediate levels. The important thing to note is that only the TCL defines measures and metrics, and those are only defined for *activities*. These metrics may be aggregated up to capabilities and mission areas, but they cannot be taken down to the task level. Therefore, the TCL metrics cannot be applied to the functions or objectives of the UTL, only to the activities and capabilities of the TCL.



Sources: TCL (9/2007), UTL v2.1 (2/2007)

TCL and UTL Taxonomies

Each TCL activity has a series of measures and metrics associated with it. Many of these are in the form of questions with yes/no answers. In principle, for each asset it should be possible to evaluate every measure and metric for every activity in the TCL. A simple way to express the result of such an exercise would be to quote the percentage of questions that were answered “satisfactorily” (e.g., answered yes to a yes/no question, achieved the specified metric, more than 50% complete, etc.) for each capability. This metric is defined as the capability compliance score. Every asset should have a set of 37 capability scores, each ranging from 0-100%.

The TCL defines approximately 1800 measures and metrics, or around 50 per capability. An individual asset owner has very little influence over most of these, since many of the capabilities are supplied by

organizations with responsibility for providing services to the asset, such as State and local law enforcement agencies, fire departments, etc. It therefore seems reasonable to define the concept of the *organizational capability compliance*, which is the level of compliance an organization supplies to all assets within its sphere of influence. An asset inherits a certain capability compliance from all the organizations that have responsibility for it. For a given jurisdiction, the *jurisdictional capability compliance* is the minimum compliance score that any asset in the jurisdiction could receive, based on the organizations operating in that jurisdiction. This jurisdictional compliance is inherited by every asset in the jurisdiction, and the asset compliance goes up from there depending on asset-specific capabilities.

One may also compute a *risk-weighted average compliance* for all assets in a jurisdiction. Imagine a jurisdiction with modest jurisdictional compliance, but one in which most of the assets have significant asset-specific capabilities that raise their compliance scores above the minimum inherited from the organizations operating in the jurisdiction. The southern tip of Manhattan might be an example, where the banks may have significantly enhanced the native capabilities provided by civic organizations (hypothetically). The jurisdictional compliance is a measure of the capabilities provided by the government (usually), and is a useful management metric, but it does not tell the whole story. A risk-weighted average asset compliance gives a measure of the actual state of the capabilities in the region. In southern Manhattan, the jurisdictional compliance may be only 10%, but if the individual owners of the highest-risk assets take additional measures to bring their assets up to 90% compliance, the risk-weighted average compliance for the region may be 80%. In this case, it may not be necessary (nor particularly cost-effective) for the local government to invest in additional capabilities itself.

CAPABILITY REQUIREMENT

Capability requirements are the levels of capability that a jurisdiction should have to mitigate its risks. Given a jurisdiction with a set of assets and a risk profile, there ought to be a minimum capability level for each of the 37 TCL capabilities that the jurisdiction should strive to maintain. Unfortunately, just as capability levels remain undefined, so are absolute standards for capability requirements. Without the right tools, capability requirements cannot be implemented.

Instead of absolute capability requirements, a proxy method is to construct a relative measure of the importance of various capabilities to a jurisdiction, based upon its risk. That quantity is called capability importance, and it will be defined below.

CAPABILITY RELEVANCE

Capability relevance refers to the strength of the relationship between a capability (one of the 37 on the TCL) and a hypothetical scenario (hazard type combined with a generic asset from a particular sector and subsector). For each type of scenario, the relevance score for a particular capability reflects the

degree to which an increase in the level of that capability could influence (mitigate) the risk for that scenario. A high capability relevance does not necessarily imply that the capability should be acquired by a jurisdiction, only that for a particular scenario, that capability could help mitigate risk.

Capability relevance is only used as an intermediate product in a risk management system such as Digital Sandbox's Risk Analysis Center (RAC), and is not generally meaningful in itself for use in decision making. There should be relevance scores stored in a lookup table that has dimensions equal to the number of hazard types (16 terrorist attack types in the RAC) times the number of different asset types (125 subsectors in the RAC) times the number of capabilities in the TCL (37). In the RAC implementation, there are 74,000 cells in this lookup table, and each one represents a different capability relevance. If natural hazards were included, the matrix would approximately double in size. This 3-dimensional matrix should be reviewed periodically to make sure the values reflect the best judgment of homeland security practitioners.

In the RAC, the capability relevance tables use three-valued assignments of coupling strength (hi/med/lo). The cardinal values for these are 1.0, 0.3, and 0.01, respectively. The values (relevance scores) were determined by expert elicitation of Digital Sandbox, client, and partner analysts as part of our work with the State of South Carolina.

CAPABILITY IMPORTANCE

The capability importance is simply the risk-weighted capability relevance for a particular capability and scenario. Without risk, capability relevance cannot help a decision-maker set policy. The introduction of risk customizes capability relevance to a particular jurisdiction and makes capability importance a useful risk management measure. A relatively high importance score implies that having a high level of that capability (relatively) effectively mitigates risk. It is capability importance that enables the link between risk analysis and risk management, as will be outlined later.

CAPABILITY GAP

A capability gap is the size of the deficit between one's capability requirement and one's capability level. The concept of a capability gap should be meaningful at any level of aggregation (e.g., asset, sector, jurisdiction). Unfortunately, neither capability requirements nor capability levels are well-defined quantities, so a quantitative proxy methodology is required to make them useful as management metrics. Since this is the language that is often used in risk management, it is important to construct viable proxies for requirements and levels so that a capability gap can be approximated.

CAPABILITY ATTENTION DEFICIT

Capability attention highlights the capabilities with the highest capability compliance scores in the system, and capability importance highlights those capabilities that ought to have the highest levels, given a particular risk profile. Since these two variables can both be given as a percentage of the total in the system, comparing the two variables can yield useful risk management insights. One major emphasis of any risk management program ought to be to align one's capability levels where the risk suggests it would be most effective. In other words, one should focus one's resources and attention on capabilities that are the most important to mitigating one's individual risks.

The simple difference between capability importance and capability attention is called the capability attention deficit. Larger attention deficits should receive the most funding priority in a jurisdiction.

Attention deficit is a relative variable. Large values of attention deficit suggest that a jurisdiction has not placed enough effort (e.g., resourcing, manpower, or other "attention") on that capability at the expense of other capabilities. So, while it is good at pointing out inequities in the distribution of attention, it cannot indicate whether overall too much or too little resourcing is being applied to reducing risk. That decision is reserved for policy-makers to decide, and the principal management metric for that is capability compliance, as described earlier.

CAPABILITY REPORTING AND RISK MANAGEMENT

The two most important capability metrics to use in managing risks are capability compliance and capability attention deficit. Compliance is a measureable proxy for capability level, and attention deficit is a measureable proxy for capability gap. A risk manager should manage investments to minimize the largest attention deficits (ideally bring all 37 to zero) and should manage compliance to be always increasing toward a pre-defined target compliance level (or at least increasing year-over-year compliance).

A report that highlights the largest capability gaps can readily be understood by a risk manager, and a strategy to close those gaps should be relatively easy to develop. For a chosen aggregation level (e.g., sector, municipality, county, urban area, state, region, etc.) the gaps should be summarized across all assets at that aggregation level. An executive summary report will thus show 37 gaps (74 if one includes natural hazards), one for each capability, where the largest gaps should receive the most immediate management (or policy) focus. The organization of the report itself remains an exercise in artistic display of information, but one could imagine a bar chart, dashboard dials, color-coded (hi-med-lo) labels, or other devices to guide the eyes of the risk manager to his or her largest gaps.

There are currently no guidelines to suggest a priori what level (compliance) is an appropriate target for a large state, smaller states, Tier I urban area, Tier II urban area, or any other entity. One benefit of widely adopting this framework will be to benchmark the capability levels (compliance scores) of peer jurisdictions to provide some of this guidance. This would be a valuable activity for DHS to coordinate.

Questions? Comments? Let us know.

sales@dsbox.com

1-877-442-4553

Digital Sandbox has been developing and deploying leading software and service-based risk management solutions since 1998.

We specialize in providing government agencies at every level – Federal, state, and local – with the strategic insight, technical support, and quantitative analysis they need to optimize investments for the risks they face.

Digital Sandbox offerings provide rigor and repeatability for risk and capability assessment and situational awareness, helping customers effectively plan and budget, advocate for, and deploy risk mitigation resources. Our solutions allow government agencies to maximize mission success by minimizing risks to our nation’s citizens, operations, and infrastructure.