



NATIONAL CYBER INCIDENT SCORING SYSTEM



DEFEND TODAY,
SECURE TOMORROW

OVERVIEW

Many incident taxonomies and classification schemes provide excellent guidance within the scope of a single enterprise's security operations center (SOC). However, such systems do not address incident prioritization or risk assessment from a nationwide perspective, which may involve large numbers of diverse enterprises. Large-scale, national cybersecurity operations centers like the Cybersecurity and Infrastructure Security Agency (CISA) need to assess risk while accommodating a diverse set of private critical infrastructure asset owners and operators and U.S. Government departments and agencies. The National Cyber Incident Scoring System (NCISS) is designed to provide a repeatable and consistent mechanism for estimating the risk of an incident in this context.

NCISS is based on the National Institute of Standards and Technology (NIST) Special Publication 800-61 Rev. 2, Computer Security Incident Handling Guide, and tailored to include entity-specific potential impact categories that allow CISA personnel to evaluate risk severity and incident priority from a nationwide perspective. NCISS permits a similar incident experienced by two different stakeholders to have significantly different scores based on the national-level potential impact of each affected entity. The system is not intended to be an absolute scoring of the risk associated with an incident.

NCISS uses a weighted arithmetic mean to produce a score from zero to 100. This score drives CISA incident triage and escalation processes and assists in determining the prioritization of limited incident response resources and the necessary level of support for each incident. The system is not currently designed to support cases where multiple correlated incidents may increase overall risk, such as multiple simultaneous compromises of organizations in a specific sector or region. However, such events can still be readily escalated with expert human intervention.

The inputs to the scoring system are a mixture of discrete and analytical assessments. While every attempt is made to minimize individual biases via training and exercise, different individual scorers will inevitably have slightly different perspectives on their responses to some of the scoring questions. The use of several discrete, verifiable inputs lessens the impact from any individual analytical factor, increasing the overall reliability of the system.

The NCISS aligns with the [Cyber Incident Severity Schema](#) (CISS) so that severity levels in the NCISS map directly to CISS levels.

FORMULA

The NCISS uses the following weighted arithmetic mean to arrive at a score between zero and 100:


- Each category has a weight, and the response to each category has an associated score. The categories are: Functional Impact,
 - Observed Activity,
 - Location of Observed Activity,
 - Actor Characterization,
 - Information Impact,
 - Recoverability,
 - Cross-Sector Dependency, and
 - Potential Impact.

CISA | DEFEND TODAY, SECURE TOMORROW

- Each response score is multiplied by the category weight, and the weighted scores are summed.
- Calculate the minimum possible weighted score sum and subtract this number from the previously calculated sum of the weighted scores. Divide the result by the range: the difference between the maximum possible weighted score sum and the minimum possible weighted score sum. Finally, multiply the resulting fraction by 100 to produce the final result.
- Weights and values are specific to an individual organization’s risk assessment process. Accompanying this document is a representative tool that demonstrates a reference implementation of the concepts outlined in this system.

PRIORITY LEVELS

After an incident is scored, it is assigned a priority level. The six levels listed below are aligned with CISA, the Department of Homeland Security (DHS), and the CISS to help provide a common lexicon when discussing incidents. This priority assignment drives CISA urgency, pre-approved incident response offerings, reporting requirements, and recommendations for leadership escalation. Generally, incident priority distribution should follow a similar pattern to the graph below.



EMERGENCY (BLACK)

An Emergency priority incident poses an imminent threat to the provision of wide-scale critical infrastructure services, national government stability, or the lives of U.S. persons.

SEVERE (RED)

A Severe priority incident is likely to result in a significant impact to public health or safety, national security, economic security, foreign relations, or civil liberties.

HIGH (ORANGE)

A High priority incident is likely to result in a demonstrable impact to public health or safety, national security, economic security, foreign relations, civil liberties, or public confidence.

MEDIUM (YELLOW)

A Medium priority incident may affect public health or safety, national security, economic security, foreign relations, civil liberties, or public confidence.

LOW (GREEN)

A Low priority incident is unlikely to affect public health or safety, national security, economic security, foreign relations, civil liberties, or public confidence.

BASELINE

A baseline priority incident is highly unlikely to affect public health or safety, national security, economic security, foreign relations, civil liberties, or public confidence. The bulk of incidents will likely fall into the baseline priority level with many of them being routine data losses or incidents that may be immediately resolved. However, some incidents may require closer scrutiny as they may have the potential to escalate after additional research is completed. In order to differentiate between these two types of baseline incidents, and seamlessly integrate with the CISS, the NCISS separates baseline incidents into Baseline–Minor (Blue) and Baseline–Negligible (White).

BASELINE – MINOR (BLUE)

A Baseline–Minor priority incident is an incident that is highly unlikely to affect public health or safety, national security, economic security, foreign relations, civil liberties, or public confidence. The potential for impact, however, exists and warrants additional scrutiny.

BASELINE – NEGLIGIBLE (WHITE)

A Baseline–Negligible priority incident is an incident that is highly unlikely to affect public health or safety, national security, economic security, foreign relations, civil liberties, or public confidence. The potential for impact, however, exists and warrants additional scrutiny.

MULTIPLE CONNECTED INCIDENTS

Currently, when a series of connected incidents, or campaign, is evaluated, the overall campaign is given the same priority level as the high water mark of any associated component incident. This does not account for a campaign that may have a more significant total impact than any individual component incident. To take into account incident aggregation when evaluating a campaign, the following rule is applied: If three or more component incidents have the same high water mark, the overall campaign's priority level is raised to the next level.

For example if a campaign has three “Low (Green)” and two “Baseline – Minor (Blue)” component incidents the overall campaign would be set to a “Medium (Yellow)” priority level.

CATEGORY DESCRIPTIONS

FUNCTIONAL IMPACT

Functional impact is a measure of the actual, ongoing impact to the organization. In many cases (e.g., scans and probes or a successfully defended attack), little or no impact may be experienced due to the incident.

OBSERVED ACTIVITY

Observed activity describes what is known about threat actor activity on the network. These options are normalized upon guidance issued by the Office of the Director of National Intelligence (ODNI) and used by the intelligence community. Although the [ODNI guidance document](#) goes into more detail, observed activity is sorted into the following general categories: Prepare, Engage, Presence, and Effect.

Prepare actions are actions taken to establish objectives, intent, and strategy; identify potential targets and attack vectors; identify resource requirements; and develop capabilities.

Engage activities are actions taken against a specific target or target set prior to gaining, but with the intent to gain access to the victim's physical or virtual computer or information systems, networks, and data stores.

Presence is the set of actions taken by the threat actor once access to the target physical or virtual computer

or information system has been achieved. These actions establish and maintain conditions for the threat actor to perform intended actions or operate at will against the host physical or virtual computer or information system, network, or data stores.

Effects are outcomes of a threat actor's actions on a victim's physical or virtual computer or information systems, networks, and data stores.

LOCATION OF OBSERVED ACTIVITY

The location of observed activity describes where the observed activity was detected in the network. The options for observed activity are based on a modified version of the Purdue Enterprise Reference Architecture. ¹ A flexible set of definitions was chosen for this category because each affected entity will likely have a different perspective on what systems are critical to its enterprise. The location of observed activity is likely to change during the course of an incident and should be updated as new information becomes available.

LEVEL 0 – UNSUCCESSFUL

Existing network defenses repelled all observed activity.

LEVEL 1 – BUSINESS DEMILITARIZED ZONE

Activity was observed in the business network's demilitarized zone (DMZ). These systems are generally untrusted and are designed to be exposed to the Internet. Examples are a company's Web server or email server.

LEVEL 2 – BUSINESS NETWORK

Activity was observed in the business or corporate network of the victim. These systems would be corporate user workstations, application servers, and other non-core management systems.

LEVEL 3 – BUSINESS NETWORK MANAGEMENT

Activity was observed in business network management systems such as administrative user workstations, active directory servers, or other trust stores.

LEVEL 4 – CRITICAL SYSTEM DMZ

Activity was observed in the DMZ that exists between the business network and a critical system network. These systems may be internally facing services such as SharePoint sites, financial systems, or relay "jump" boxes into more critical systems.

LEVEL 5 – CRITICAL SYSTEM MANAGEMENT

Activity was observed in high-level critical systems management such as human-machine interfaces (HMIs) in industrial control systems.

LEVEL 6 – CRITICAL SYSTEMS

Activity was observed in the critical systems that operate critical processes, such as programmable logic controllers in industrial control system environments.

LEVEL 7 – SAFETY SYSTEMS

Activity was observed in critical safety systems that ensure the safe operation of an environment. One example of a critical safety system is a fire suppression system.

¹ <http://www.pera.net/>

UNKNOWN

Activity was observed, but the network segment could not be identified.

ACTOR CHARACTERIZATION

One of the greatest challenges in incident response is attributing an incident to a particular actor set and understanding the skill levels and intentions of that actor. CISA may leverage its own analytic body of knowledge as well as that of other mission partners to determine an actor's capabilities with regard to specific target systems such as industrial control environments.

INFORMATION IMPACT

In addition to functional impact, incidents may also affect the confidentiality and integrity of the information stored or processed by various systems. The information impact category is used to describe the type of information lost, compromised, or corrupted.

RECOVERABILITY

Recoverability represents the scope of resources needed to recover from the incident. In many cases, an entity's internal computer network defense staff will be able to handle an incident without external support, resulting in a recoverability classification of Regular. An example of a Regular recovery would be a phishing email that was automatically blocked by a mail server. In Extended recoverability cases, significant efforts such as a multi-agency, multi-organizational response task force may be needed for recovery. For example, if an entity requests support from CISA, the incident is by its nature an Extended recovery. Lastly, it may not be feasible to recover from some types of incidents, such as significant confidentiality or privacy compromises.

REGULAR

Time to recovery is predictable with existing resources.

SUPPLEMENTED

Time to recover is predictable with additional resources.

EXTENDED

Time to recovery is unpredictable; additional resources and outside assistance may be required.

NOT RECOVERABLE

Recovery from the incident is not possible (e.g., sensitive data was exfiltrated and posted publicly, investigation launched).

CROSS-SECTOR DEPENDENCY

Cross-sector dependency is a weighting factor that is determined based on cross-sector analyses conducted by CISA.

POTENTIAL IMPACT

The potential impact category estimates the overall national impact resulting from a total loss of service from the affected entity. Other existing standards for rating cybersecurity incident risk lack consideration for the unique and diverse critical infrastructure assets of the owners and operators and U.S. Government departments and agencies that CISA is tasked with helping to protect. A similar incident at two separate stakeholder facilities might have a significantly different impact to operations at a national level. Therefore, each incident will be scored differently relative to the risk it presents in a nationwide context.

The potential impact value is calculated in advance wherever possible, based on known statistics about the entity in question. Some example statistics that may be used include:

- number of authorized users in the organization,
- reported annual revenue or total annual budget, and
- size of customer base or population served.

Several factors are considered in calculating the potential impact value for individual entities. Certain factors applicable for utility companies, healthcare firms, or financial services institutions are not applicable for Federal Government agencies, so the weighted factors for each type of entity will differ. In developing NCISS, many possible factors were considered for inclusion in potential impact calculations. This particular facet of the scoring system is the subject of continued research and evaluation.

Lastly, due to the inherent difficulties in accounting for all the various circumstances involved in determining the true potential impact, this value in particular should be treated as a best guess estimate for incident response prioritization purposes, and not as a comprehensive illustration of an entity's importance to the national welfare.

CONCLUSION

NCISS is designed to provide a repeatable and consistent mechanism for objectively evaluating the risk of a cybersecurity incident in the national context. Having this system in place has already allowed CISA to provide objective assessments of national-level risk for routine and high-risk cybersecurity events via a repeatable process, facilitating better prioritization and more timely responses to the needs of CISA's constituents and mission partners.