Automated Indicator Sharing (AIS) Scoring Framework Used for Indicator Enrichment

V1.0

Publication: November 2021
Cybersecurity and Infrastructure Security Agency

Contents

WHAT IS THE AIS SCORING FRAMEWORK?

The AIS Scoring Framework defines an algorithm by which organizations can enrich Structured Threat Information Expression (STIX) Indicator objects,[1] shared via AIS, with (1) an opinion value that provides an assessment of whether or not the information can be corroborated with other sources available to the entity submitting the opinion and (2) a confidence score that states the publisher's confidence in the correctness of information they submit into AIS. Together, these enrichments can help those receiving information from AIS prioritize actioning and investigating Indicator objects. CISA uses the AIS Scoring Framework to provide an opinion value for each Indicator object submitted to AIS, and will also use it when CISA provides a confidence score on CISA-published Indicator objects. CISA is sharing the methodology to explain how CISA develops opinion values and confidence scores for individual Indicator objects when provided, and so other organizations can decide whether to use the same methodology in developing their own opinion values and confidence scores.

To identify the appropriate inputs from the STIX Opinion object's predefined vocabulary enumeration -- or confidence property numerical score (described further below) -- CISA developed the AIS Scoring Framework based on the "Evaluation of Information Content" framework in the Admiralty system for evaluating human intelligence. The framework analyzes three characteristics about the data: Confirmed, Logical, and Consistent.[2] Table 1 lists the Admiralty system's possible scoring framework values and their description based on these characteristics, and the associated description used for those framework values in the AIS Scoring Framework. Please note that Indicator objects can be marked as benign via the **indicator_types** property and that the AIS Scoring Framework considers them in the context in which they are marked; Indicator objects marked with any value other than benign (including those marked as malicious-activity, those marked with another option from the `indicator-type-ov`, and those with another or no value) will be evaluated through the scoring framework as if they are marked malicious-activity.[3] That means Indicator objects marked only as benign will be evaluated based on the likelihood that they are benign and all Indicator objects not marked as benign, including those marked as malicious-activity, will be evaluated based on the likelihood that they are malicious. If an Indicator object is marked as both benign and another indicator type, it will not be evaluated with the AIS Scoring Framework, but will still be processed and distributed appropriately through AIS.

Table 1: Scoring Framework Values and Descriptions

Scoring Framework Values	Admiralty Description (Confirmed; Logical; Consistent)	AIS Scoring Framework Description
Improbable	Not confirmed; not logical in itself; contradicted by other information on the subject	Not confirmed; likely not malicious (if marked or otherwise evaluated as malicious-activity) or benign (if marked as benign); contradicted by other information on the subject known to the opinion author.

[1] The AIS Scoring Framework described in this paper was developed specifically for use with Indicator objects only. However, the general approach described in this paper might be adaptable for other STIX Domain Objects (SDOs) and STIX Relationship Objects (SROs). At this time, CISA will only use the AIS Scoring Framework to enrich Indicator Objects, and its use is described here and in other AIS documents consistent with this use.

[2] https://armypubs.army.mil/epubs/DR_pubs/DR_a/pdf/web/fm2_22x3.pdf

[3] https://docs.oasis-open.org/cti/stix/v2.1/os/stix-v2.1-os.html#_muftrcpnf89v

Scoring Framework Values	Admiralty Description (Confirmed; Logical; Consistent)	AIS Scoring Framework Description
Doubtfully True	Not confirmed; possible but not logical; no other information on the subject	Not confirmed; possibly malicious (if marked or otherwise evaluated as malicious-activity) or benign (if marked as benign); no other information on the subject known to the opinion author.
Possibly True	Not confirmed; reasonably logical in itself; agrees with some other information on the subject	Not confirmed; possibly malicious (if marked or otherwise evaluated as malicious-activity) or benign (if marked as benign); agrees with some other information on the subject known to the opinion author.
Probably True	Not confirmed; logical in itself; consistent with other information on the subject	Not confirmed; likely malicious (if marked or otherwise evaluated as malicious-activity) or benign (if marked as benign); consistent with other information on the subject known to the opinion author.
Confirmed	Confirmed by other independent sources; logical in itself; consistent with other information on the subject	Confirmed; likely malicious (if marked or otherwise evaluated as malicious-activity) or benign (if marked as benign); consistent with other information on the subject known to the opinion author.

Of note, the AIS Scoring Framework should be applied by recipients without regard for the specific reputation or sophistication of the submitter and should only represent an assessment of information submitted through AIS against other information available to recipient organizations (e.g., CISA or a later Indicator recipient). At CISA, Indicators are assessed through the AIS Scoring Framework through automated means and only against information available for automated processing.

HOW DOES THE AIS SCORING ALGORITHM WORK?

Prior to submission, submitting organizations can use the AIS Scoring Algorithm to assign a confidence score to an Indicator object. Once the Indicator object is published to AIS participants, the AIS Scoring Algorithm may be used by other organizations to evaluate the pattern property of the Indicator object against various sources to determine the opinion value Appendix A depicts the AIS Scoring Algorithm.

Step 1*:* Is Indicator object present in known-good list? The pattern property of the Indicator object is compared to a list of known non-malicious indicators. If the pattern is present in the list and the Indicator object is marked as benign via the **indicator_types** property, the Indicator object is assigned an AIS Scoring Framework value of *Confirmed*. If the pattern is present in the list and the Indicator object is marked or otherwise evaluated as malicious-activity, the Indicator object is assigned an AIS Scoring Framework value of *Improbable*. If the pattern is not present in the list, the Indicator object is evaluated against Step 2.

Step 2: Has Indicator object been observed by the organization? In this step, the pattern property of the Indicator object is cross-referenced with organization-centric intelligence (e.g., the indicator was observed by the organization or detected by a sensor used to protect the organization's networks). If there is a match and the Indicator object is marked as benign, it is assigned an AIS Scoring Framework value of *Improbable* because signatures are not typically written for benign indicators. If there is a match and the Indicator object is marked or otherwise evaluated as malicious-activity, it is assigned an AIS Scoring Framework value of *Confirmed*. Otherwise, the Indicator object is evaluated against Step 3.

Step 3: Has Indicator object been verified by an analyst? The pattern property of the Indicator object is compared against indicators in the organization's threat intelligence platform (TIP). If the TIP shows that the indicator has previously been verified by an analyst, the Indicator object is assigned an AIS Scoring Framework value of *Probably True*. If the Indicator object has not been verified, it is evaluated against Step 4.

Step 4: Has Indicator object been confirmed by other sources? The pattern property of the Indicator object is compared against other cyber threat intelligence from sources available to the organization. If the indicator is confirmed by another source, the Indicator object is assigned an AIS Scoring Framework value of *Possibly True*. Otherwise, the Indicator object is assigned an AIS Scoring Framework value of *Doubtfully True*.

HOW DO I CAPTURE AN OPINION VALUE IN STIX?

The STIX Opinion object can be used to represent an organization's opinion for any Indicator objects.[4] The following properties (see Table 2) in the Opinion object may be used to represent the opinion value.

Table 2: STIX Opinion Object Properties

Property Name	Type	Description
type **(required)**	string	The value of this property **MUST** be opinion.
explanation **(optional)**	string	An explanation of why the producer has this Opinion. For example, if an Opinion of strongly-disagree is given, the explanation can contain an explanation of why the Opinion producer disagrees and what evidence they have for their disagreement.
opinion **(required)**	enum	The opinion that the producer has about all of the STIX Object(s) listed in the **object_refs** property. The values of this property **MUST** come from the opinion-enum enumeration.
object_refs **(required)**	list of type identifier	The STIX Objects that the Opinion is being applied to.

[4] https://docs.oasis-open.org/cti/stix/v2.1/os/stix-v2.1-os.html#_ht1vtzfbtzda

Property Name	Type	Description
external_references (optional)	list of type external-reference	The **external_reference** property specifies a list of external references which refers to non-STIX information. This property is used to provide one or more URLs, descriptions, or IDs to records in other systems.

The STIX Opinion object includes a predefined vocabulary enumeration (*opinion-enum*) in the underlying **opinion** property, consisting of five values: Strongly-agree, Agree, Neutral, Disagree, and Strongly-disagree. Each value in the predefined object vocabulary is mapped to an AIS Scoring Framework value as shown in Table 3.

Table 3: Scoring Framework Values and Mapping to Opinion Values

AIS Scoring Framework Values	STIX Opinion Property Value
Improbable	Strongly-disagree
Doubtfully True	Disagree
Possibly True	Neutral
Probably True	Agree
Confirmed	Strongly-agree

Figure 1 is an example of an Indicator object with an opinion value assigned by CISA (opinion values by other organizations may look different). The confidence score, as determined by the creator, is captured in the **confidence** property of the Indicator object. The opinion value, as determined by CISA in this example, is captured in the **opinion** property and the description of the opinion value is captured in the explanation property. The **external_references** property may also be used to provide users with more information about how the opinion value is generated.


```
{¶
··"type":·"indicator",¶
··"spec_version":·"2.1",¶
··"id":·"indicator--e9455434-be0c-4c14-a5f4-a7cd51d547a3",¶
··"created_by_ref":·"identity--a7163454-a268-494b-ba07-557023d88014",¶
··"created":·"2020-10-30T08:17:27.000Z",¶
··"modified":·"2020-10-30T08:17:27.000Z",¶
··"object_marking_refs":·["marking-definition--613f2e26-407d-48c7-9eca-b8e91df99dc9"],¶
··"name":·"Malicious·IP·Address",¶
··"description":·"This·IP·address·is·associated·with·known·malicious·activity.",¶
··"indicator_types":·["malicious-activity"],¶
··"pattern":·"[ipv4-addr:value·=·'198.51.100.3']",¶
··"pattern_type":·"stix",¶
··"valid_from":·"2020-10-30T00:00:00Z",·¶
··"confidence":·75¶
}¶
```

```
{
  "type": "opinion",
  "spec_version": "2.1",
  "id": "opinion--c18c72f9-abdd-490a-9781-aaf3b9c50eaa",
  "created_by_ref": "identity--b3bca3c2-1f3d-4b54-b44f-dac42c3a8f01",
  "created": "2020-10-30T12:17:27.000Z",
  "modified": "2020-10-30T12:17:27.000Z",
  "object_marking_refs": ["marking-definition--613f2e26-407d-48c7-9eca-b8e91df99dc9"],
  "object_refs": ["indicator--e9455434-be0c-4c14-a5f4-a7cd51d547a3"],
  "opinion": "strongly-agree",
  "explanation": "Confirmed; likely malicious (if marked or otherwise evaluated as malicious) or
                  benign (if marked as benign); consistent with other information on the subject
                  known to the opinion author. Please see AIS Scoring Framework Used for
                  Indicator Enrichment at https://www.cisa.gov/publication/automated
                  indicator-sharing-ais-documentation for more information on how the
                  opinion is derived.",
  "external_references": [
    {
      "source_name": "AIS Scoring Framework Used for Indicator Enrichment",
      "description": "This reference provides more information about the AIS Scoring Framework
                      and how the opinion is derived.",
      "url": "https://www.cisa.gov/publication/automated-indicator-sharing-ais-documentation"
    }
  ]
}
```


Figure 1: Example STIX Content with Opinion

HOW DO I USE A POPULATED OPINION VALUE?

AIS participants can apply automated or manual triage against the populated opinion value to identify Indicator objects meeting or exceeding designated criteria and filter out the remaining data. This is possible because an opinion value is applied by CISA to all Indicator objects shared within AIS (except those marked as both benign and another indicator type). More information about how to filter objects based on the opinion value can found in *Filtering AIS Content Based on Specified Criteria.*[5]

HOW DO I CAPTURE A CONFIDENCE SCORE IN STIX?

The optional **confidence** property allows the AIS participant to denote the confidence that they have in the correctness of data they produce. STIX permits the use of an integer value from 0–100 to represent the confidence score, which can be mapped to a variety of different measurement scales.[6] When CISA provides confidence scores, it will include a value for the optional **confidence** property on Indicator objects that it publishes based on the AIS Scoring Algorithm described above, and other organizations may do the same with Indicator objects they publish.[7] Table 4 outlines the possible confidence scores that may be present on an Indicator object if the AIS Scoring Framework is used. The **external_references** property may also be used to provide users with more information about how the confidence score is generated.

[5] https://www.cisa.gov/publication/automated-indicator-sharing-ais-documentation
[6] https://docs.oasis-open.org/cti/stix/v2.1/os/stix-v2.1-os.html#_xzbicbtscatx
[7] https://docs.oasis-open.org/cti/stix/v2.1/os/stix-v2.1-os.html#_xzbicbtscatx

Table 4: AIS Scoring Framework Values and Mapping to Confidence Scores

AIS Scoring Framework Values	Confidence Score
Improbable	10 (Represents range 0 – 19)
Doubtfully True	30 (Represents range 20 – 39)
Possibly True	50 (Represents range 40 – 59)
Probably True	70 (Represents range 60 – 79)
Confirmed	90 (Represents range 80 – 100)

Figure 2 provides an example of a confidence score generated by the AIS Scoring Algorithm for an Indicator object.

```
{
  "type": "indicator",
  "spec_version": "2.1",
  "id": "indicator--8fbae920-329a-4e56-8f46-7e4def27d678",
  "created_by_ref": "identity--c2e745b4-b5ae-40a7-925c-bbc3cb1c148b",
  "created": "2020-11-07T09:23:43.000Z",
  "modified": "2020-11-07T09:23:43.000Z",
  "object_marking_refs": ["marking-definition--613f2e26-407d-48c7-9eca-b8e91df99dc9"],
  "name": "Malicious Website",
  "description": "This URL represents a malicious website that is used to host malware.",
  "indicator_types": ["malicious-activity"],
  "pattern": "[url:value = 'http://badurl.biz/372/']",
  "pattern_type": "stix",
  "valid_from": "2020-11-07T00:00:00Z",
  "confidence": 90,
  "external_references": [
    {
      "source_name": "AIS Scoring Framework Used for Indicator Enrichment",
      "description": "This reference provides more information about how the confidence score
                      is derived.",
      "url": "https://www.cisa.gov/publication/automated-indicator-sharing-ais-documentation"
    }
  ]
}
```

Figure 2: Example STIX Content with Confidence Score

HOW DO I USE THE CONFIDENCE SCORE?

AIS participants can apply automated or manual triage against Indicator objects meeting or exceeding a specific confidence score and filter out the remaining data. However, the confidence score is an optional property and will not necessarily be populated by all AIS participants. Of course, AIS participants may still find value in utilizing the confidence score (if present) and the opinion value to understand whether any difference between the publisher and CISA (or other organizations submitting opinions using the AIS Scoring Framework) exists. More information about how to filter objects based on confidence score can found in *Filtering AIS Content Based on Specified Criteria.*[8]

[8] https://www.cisa.gov/publication/automated-indicator-sharing-ais-documentation

APPENDIX A - SCORING FRAMEWORK ALGORITHM

The following diagram shows the AIS Scoring Framework algorithm.

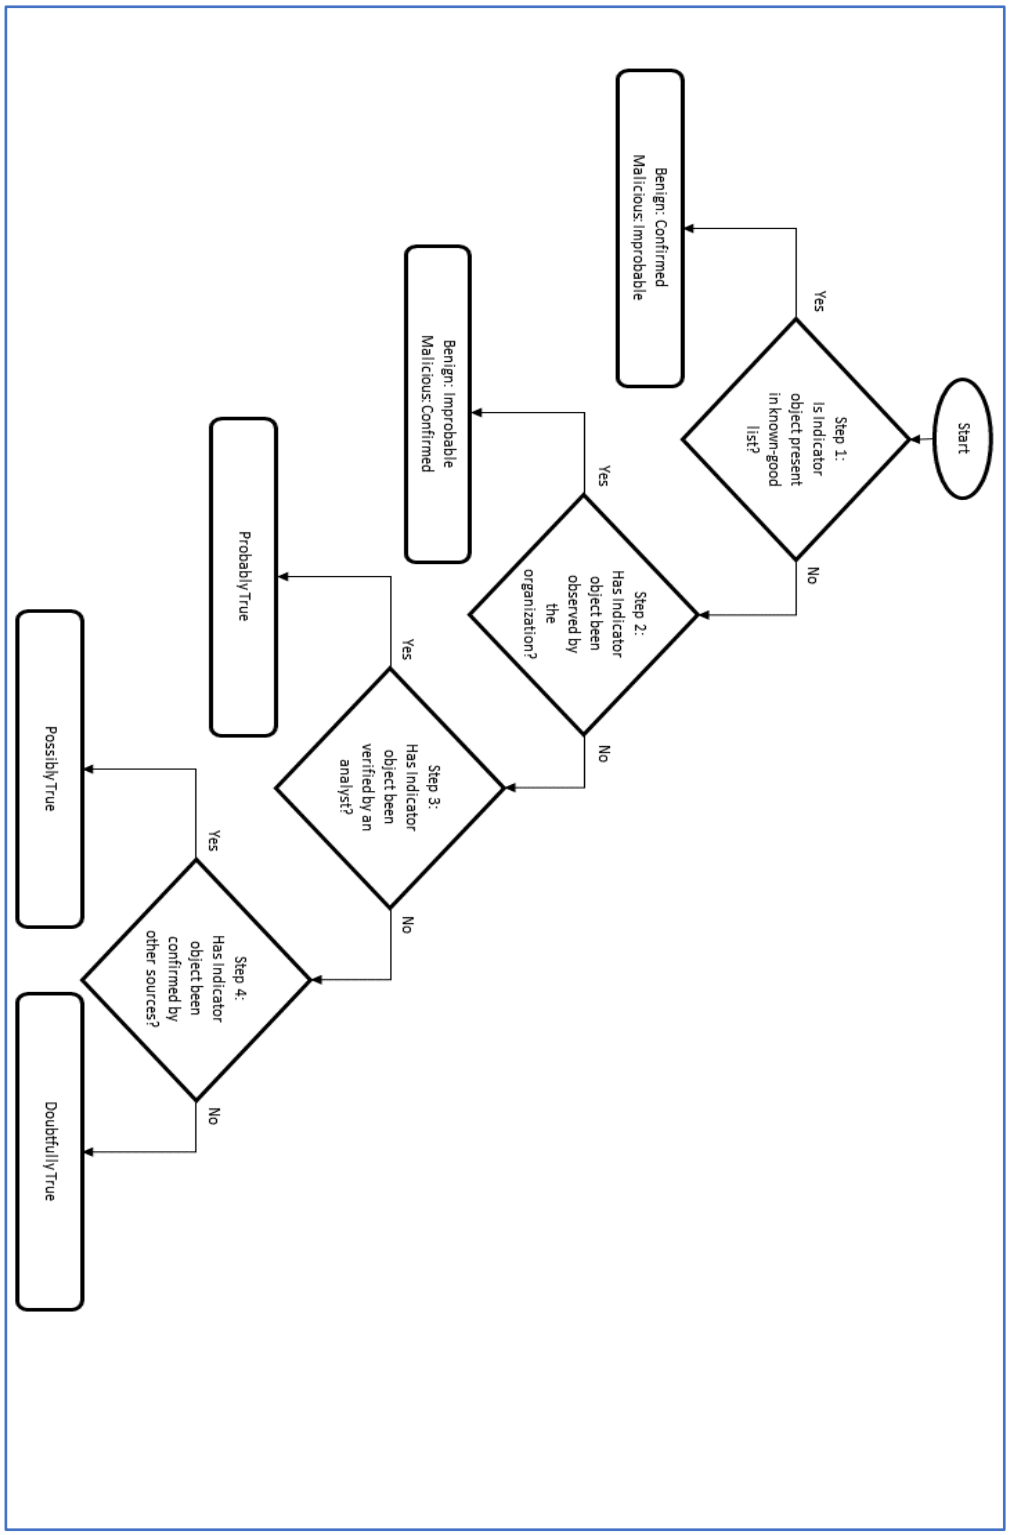

Figure 3: AIS Scoring Algorithm

APPENDIX B - SCORING FRAMEWORK VALUES AND MAPPING TO CONFIDENCE SCORE AND OPINION VALUES

Table 5 shows the mapping between the AIS Scoring Framework values and the confidence score and opinion value property.

Table 5: Scoring Framework Values and Mapping to Confidence Score and Opinion Values

Scoring Framework Values	Confidence Score	Opinion Value Property
Improbable	10	Strongly-disagree
Doubtfully True	30	Disagree
Possibly True	50	Neutral
Probably True	70	Agree
Confirmed	90	Strongly-agree