Balisage: **The Markup Conference**

Balisage Paper: The Model Made Me Do It! A Cautionary Tale from a Security Control Baseline Tool Developer

Joshua Lubell
National Institute of Standards and Technology

Abstract

Even the best written specifications can be complicated documents to read and understand. Normative prose is often supported by tables and diagrams intended to clarify the specification. What happens when a tool developer interprets those clarifying features as a different model than the prose intends? What does this say about relying on derived data models in tools that support the specification? A cautionary tale involving security control baselines from National Institute of Standards and Technology Special Publication 800-53 provides some answers — and insights.

▼ Table of Contents

Note

These opinions, recommendations, findings, and conclusions do not necessarily reflect the views or policies of NIST or the United States Government.

Introduction

This paper relates my experience developing Baseline Tailor [BT], a tool for using the United States government's Cybersecurity Framework [NIST18] and National Institute of Standards and Technology (NIST) Special Publication (SP) 800-53 security control baselines [JTF20] [JTF20B]. Baseline Tailor makes it easier for security practitioners to use these specifications together. However, despite its name, Baseline Tailor is less effective for tailoring security control baselines.

To help understand what “baseline” and “tailoring” mean in the context of cybersecurity, this discussion begins with an overview of SP 800-53 and how organizations use it to manage information security and privacy risks. SP 800-53 defines a collection of hundreds of security controls, each of which protects the confidentiality, integrity, or availability (CIA) of a system and the information it processes or transmits. Each control has zero or more control enhancements (“sub-controls”) that add functionality to or increases the strength of its parent control. This catalog of controls and enhancements is highly detailed and comprehensive, yet implementation-agnostic. Controls are grouped into families, where each family relates to a specific topic, for example, access control or configuration management. Most of these topics correspond to a security requirement specified in the Federal Information Processing Standards Publication (FIPS) 200 [NIST06].

SP 800-53 is the foundation of a Risk Management Framework (RMF) [NIST18] [Lubell20], a system life cycle-based process for managing security and privacy risk. SP 800-53 controls play a role in much of the RMF process, but these three RMF steps are the ones most pertinent to our discussion:

***Categorize* step**	Describes the system and classifies the impact of loss of CIA of its information.
***Select* step**	Chooses an initial set of controls for the system, tailoring them as needed to reduce risk to an acceptable level.
***Implement* step**	Implements the controls and describes their deployment.

A system's *impact*, as determined in the RMF *Categorize* step, may be either “low” (limited), “moderate” (serious), or “high” (catastrophic). Determining system impact relies on the system's high water mark with respect to a loss of CIA. For example, suppose a system has three categories of information: proprietary information where a compromise in confidentiality would be serious, vital information where even a temporary loss of access would be catastrophic to business operations, and information where a compromise would be no more than a nuisance. The system's impact in this case would be equal to the highest CIA impact value over all information types. Since the highest of these impact values (impact of a loss of access to the vital information) is “high”, the company's system impact would be classified as “high”. The underlying logic of the high water mark concept is that confidentiality, integrity, and availability are often interdependent, so a compromise to one is likely to affect the others [NIST06].

Given the size and complexity of the SP 800-53 catalog, the RMF *Select* step can be daunting. To make it easier, SP 800-53 specifies three control baselines for low, moderate, and high impact information systems as starting points for security control selection. For example, an organization selecting security controls for a low impact system might begin with the controls in the baseline for the low impact level (or more succinctly, the low baseline) and tailor them as appropriate. *Tailoring* is the process of modifying a baseline. Examples include identification of common controls (controls that can be inherited by multiple systems to reduce IT infrastructure complexity and save money), assigning values to organization-defined control parameters (e.g., minimum password length), adding additional controls or enhancements, and providing additional guidance. A common example of additional guidance is specification of how a control is (or should be) implemented based on knowledge gained during the RMF *Implement* step.

My study of SP 800-53 and experience working on a project developing implementation methods, metrics, and tools to secure advanced digital manufacturing systems spurred my idea to develop a software tool for tailoring baselines. Such a tool would be helpful for creating manufacturing-specific baselines based on SP 800-53 and also could benefit would-be baseline authors in other industry sectors. The rest of this paper describes Baseline Tailor's implementation and my interpretation of SP 800-53 from the perspective of a tool developer and markup language enthusiast.

A central theme is that SP 800-53 implies multiple conceptual models of tailoring. Supplementing these implied models are derivative structured digital data models that are not officially part of the documentary standard. Collectively, these implied and derivative models provide two distinct approaches to tailoring. One approach is control-centric, which specifies tailoring in terms of the allocation of each control and control enhancement to the SP 800-53 low, moderate, and high baselines. The other approach is baseline-centric, which specifies tailoring in terms of a pre-existing baseline. This initial baseline can be one of the three SP 800-53 baselines, or it can be a pre-existing tailoring of one of the SP 800-53 baselines. As a tool developer lacking firsthand experience tailoring baselines, I failed to grasp this distinction when developing Baseline Tailor. As a result, although successful from a Cybersecurity Framework/SP 800-53 integration standpoint, Baseline Tailor achieved only mixed success with baseline authors.

The paper proceeds as follows. The second section provides examples illustrating the multiple baseline models that SP 800-53 prose and tables imply. It also contrasts the recently-released SP 800-53 revision 5 [JTF20] with the previous revision 4 [JTF13]. The third section compares two distinct structured digital derivations of SP 800-53 implied models: the SP 800-53 database Extensible Markup Language (XML) schema and the newer Open Security Controls Assessment Language (OSCAL) [OSCAL] profile model for baselines. The fourth section describes the design of Baseline Tailor and how it was influenced by implied and derivative control-centric SP 800-53 models. The fifth section discusses how a new version of Baseline Tailor will adopt a new approach to tailoring informed by both OSCAL and the tailoring criteria used in NIST SP 800-171 [Ross], a guidance publication for protecting controlled unclassified information. The sixth section concludes the paper.

Implied Models

SP 800-53 is a documentary standard that describes structured information using prose and tables (that is, a Small Arcane Non-trivial Dataset [Lubell14]). This section uses examples to describe the different baseline models that SP 800-53 prose and tables imply. It also describes an implied model for tailoring baselines from SP 800-82 revision 2, NIST's Guide to Industrial Control Systems (ICS) Security [Stouffer]. I use SP 800-53 security control AU-3 (title: “Content of Audit Records”) as the source for many examples in this and subsequent sections. AU-3 ensures that for each system event logged, the audit record contains the event type, time the event occurred, when and where it occurred, the event's source and outcome, and the identity of anyone or anything associated with the event.

Previous and Current SP 800-53 Revisions

AU-3 has the following control enhancements:

AU-3(1) Additional Audit Information	Audit records must contain additional (organization-defined) information.
AU-3(2) Centralized Management of Planned Audit Record Content	Configuration of audit record content must be centrally managed. This control enhancement was withdrawn in revision 5 and incorporated into control PL-9 (Central Management).
AU-3(3) Limit	Limits Personally Identifiable Information (PII) in audit records to an

Personally Identifiable Information Elements organizationally-defined list of allowable information items (to reduce privacy risk). This control enhancement was added in revision 5.

Table I shows the control-centric AU-3 baseline summary entry from Table D-2 in revision 4. Table D-2 summarizes the baselines for all control families. This table entry conveys that AU-3 is allocated to the low baseline, control enhancement AU-3(1) is allocated to the moderate and high baselines, and AU-3(2) is allocated only to the high baseline. Curiously, Table I does not explicitly list AU-3 as allocated to the moderate and high baselines, even though any control allocated to the low baseline is by definition allocated to medium and high as well.

Table I

AU-3 security control baseline summary (SP 800-53 rev. 4. Table D-2).

Cntl No.	Control Name	Initial Control Baselines		
		Low	Mod	High
AU-3	Content of Audit Records	AU-3	AU-3 (1)	AU-3 (1) (2)

Table D-5 in revision 4 shows a more detailed baseline summary for the Audit and Accountability (AU) family. This table's entry for AU-3 (Table II) contains three rows: one for the control and one for each control enhancement. An “x” in a baseline column indicates inclusion in the baseline. Unlike Table I, Table II explicitly conveys that AU-3 is not only in the low baseline, but also in the moderate and high baselines. Also, having separate rows for each control enhancement reinforces the concept of controls and control enhancements belonging to baselines, as opposed to baseline allocations belonging to a control.

Table II

A more detailed AU-3 security control baseline summary (SP 800-53 rev. 4, Table D-5).

Cntl No.	**Control Name** Control Enhancement Name	Initial Control Baselines		
		Low	Mod	High
AU-3	**Content of Audit Records**	x	x	x
AU-3(1)	Content of Audit Records | Additional Audit Information		x	x
AU-3(2)	Content of Audit Records | Centralized Management of Planned Audit Record Content			x

Now let us move on to SP 800-53 revision 5, whose two major changes from revision 4 include:

Numerous additions, withdrawals, and revisions to the control catalog The changes pertaining to AU-3 were mentioned at the beginning of this subsection.

Baselines in a separate document Instead of being published in appendices, the baselines are documented in a separate publication, SP 800-53B. Decoupling the controls and baselines simplifies document management by enabling revisions to one without having to revise the other. Moreover, it further promotes the baselines as first class citizens rather than “allocations” of controls.

SP 800-53B introduces a fourth baseline: a *privacy* baseline. Although privacy and confidentiality overlap, they are not synonymous. For example, protection of a company's trade secrets is a confidentiality concern, but not a privacy concern. However, ensuring medical patients' ability to review their healthcare records is a privacy concern, but pertains more to information access than to confidentiality. Corrupted PII stored on a system compromises both personal privacy and information integrity. Additionally, a control or control enhancement may be a member of the privacy baseline without supporting management of risks arising from a loss of confidentiality, integrity, or availability. For example, as shown in Table III, the new revision 5 control enhancement AU-3(3) is assigned to the privacy baseline but none of the other baselines.

Table III

AU-3 security control baseline summary (SP 800-53B).

Cntl No.	**Control Name** Control Enhancement Name	Privacy Control Baseline	Security Control Baselines		
			Low	Mod	High
AU-3	**Content of Audit Records**		x	x	x
AU-3(1)	Additional Audit Information			x	x
AU-3(3)	Limit Personally Identifiable Information Elements	x			

The preceding examples point to a fundamental difference between the privacy baseline versus the low, moderate, and high baselines. In nearly all cases, a control enhancement in the low baseline, will also be in the moderate baseline, and a control enhancement in the moderate baseline will also be in the high baseline. The rare exception to this rule is when two control enhancements are mutually exclusive, as is the case with two mutually exclusive enhancements of CM-7 (Least Functionality) from the Configuration Management family. The mutually exclusive control enhancements, CM-7(4) (Unauthorized Software: Deny-by-exception) and CM-7(5) (Unauthorized Software: Allow-by-exception) cannot both be present in the same baseline. The reason is that CM-7(5) is a stronger security measure than CM-7(4), and thus it makes no sense to implement CM-7(4) for a system where CM-7(5) has already been implemented. In other words, if you are already disallowing the use of all unauthorized software by default, there is no need to implement a policy disallowing the use of a specific software product.

Unlike the low, moderate, and high baselines, where a subset relationship prevails in nearly all cases, the privacy baseline is cross-cutting, yet also targeted. It is cross-cutting in that privacy is not specific to confidentiality, integrity, or availability. However, privacy is targeted specifically to the protection of PII. Whereas a loss of confidentiality, integrity, or availability may directly impact people, inanimate assets (“things”), or the environment, a loss of privacy primarily impacts people. This fundamental difference between the privacy baseline and the low/moderate/high baselines is yet another reason to adopt a more annotation-based and baseline-centric approach as discussed in the fifth section.

Industrial Control Systems Overlay

The NIST ICS overlay, based on the SP 800-53 revision 4 baselines and documented in Appendix G of SP 800-82 Revision 2 [Stouffer], provides another implied baseline model. An ICS is a collection of control components (e.g., sensors, actuators, programmable logic controllers, mechanical devices, computers) that act together to achieve an industrial purpose (e.g., manufacturing, transportation, energy production, energy transmission). A security control overlay is conceptually like a tailored baseline or set of baselines, except that overlays are typically intended to be shared by a community of interest rather than meet a single organization's unique requirements. One can think of the ICS overlay as a variant of the SP 800-53 revision 4 low, moderate, and high baselines, with partial tailoring to address basic ICS security assumptions.

As a practical matter, the biggest differentiator between an overlay versus a garden-variety tailored baseline is who the developer is likely to be. Developers of overlays are likely to possess deep knowledge of the SP 800-53 controls and the unique security requirements of their community of interest [SCOR]. Tailored baselines, on the other hand, are usually developed within an organization where the developers may be well acquainted with system-specific or organization-specific security requirements but less familiar with the detailed guidance of SP 800-53[1]. Therefore, most people tasked with tailoring a baseline are inclined to start with one of the low/moderate/high SP 800-53 baselines, each of which contains a smaller subset of the catalog[2], or with an applicable overlay if one exists.

The ICS overlay is specified in a tabular format based on the detailed baseline summary format shown in Table II. Table IV shows how the ICS overlay specifies AU-3. The baseline allocation is unchanged from revision 4. Thus, the Low, Moderate, and High columns show “Selected” in the boxes wherever an “x” occurs in Table I. If a control or enhancement were removed from a baseline, the box would specify “Removed” instead of “Selected”. If a control or enhancement were added to a baseline, the box would specify “Added” instead of being blank. The ICS overlay provides a prose rationale statement justifying each baseline addition or removal. The rationale plus any additional prose guidance appears below the control's baseline allocation. As shown in Table IV, AU-3 has ICS supplemental guidance but neither AU-3(1) nor AU-3(2) have additional guidance. Since the ICS overlay does not modify AU-3's baseline allocation, Table I provides no rationale statement.

Table IV

AU-3 ICS overlay specification (SP 800-82 rev. 2).

Cntl No.	**Control Name**	Initial Control Baselines		
	Control Enhancement Name	Low	Mod	High
AU-3	**Content of Audit Records**	Selected	Selected	Selected
AU-3(1)	Content of Audit Records | Additional Audit Information		Selected	Selected
AU-3(2)	Content of Audit Records | Centralized Management of Planned Audit Record Content			Selected

ICS Supplemental Guidance	Example compensating controls include providing an auditing capability on a separate information system.
Control Enhancement: (1,2)	No ICS Supplemental Guidance.

Table V shows the ICS overlay specification for AU-4 (Audit Storage Capacity), which requires that audit record storage capacity meet organization-defined record retention requirements. The ICS overlay adds the control enhancement AU-4(1) to all three baselines. AU-4(1) requires audit records to be periodically off-loaded from the system being audited onto another system or external storage media. AU-4(1) is not included in any of the SP 800-53 baselines because most IT systems are configured to store audit data locally. However, as the AU-4(1) ICS supplemental guidance and rationale statement explain, this assumption does not hold for ICS. The AU-4 ICS overlay specification is a good example of tailoring that requires specialized expertise in multiple areas: awareness of less-commonly used security controls, knowledge of ICS configuration requirements, and regulatory requirements for some ICS.

Table V

AU-4 ICS overlay specification (SP 800-82 rev. 2).

Cntl No.	**Control Name**	Initial Control Baselines		
	Control Enhancement Name	Low	Mod	High
AU-4	**Audit Storage Capacity**	Selected	Selected	Selected
AU-4(1)	Audit Storage Capacity | Transfer to Alternate Storage	Added	Added	Added

ICS Supplemental Guidance	None.
Control Enhancement: (1) ICS Supplemental Guidance	Legacy ICS are typically configured with remote storage on a separate information system (e.g., the historian accumulates historical operational ICS data and is backed up for storage at a different site). ICS are currently using online backup services and increasingly moving to cloud based and virtualized services. Retention of some data (e.g., SCADA[3] telemetry) may be required by regulatory authorities.
Rationale for adding AU-4 (1) to all baselines	Legacy ICS components typically do not have capacity to store or analyze audit data. The retention periods for some data, particularly compliance data, may require large volumes of storage.

Derivative Digital Models

The latest published SP 800-53 and SP 800-53B documents are considered the normative sources for controls and baselines. However, NIST also provides a non-normative searchable online database version of the controls and baselines. The SP 800-53 database [SP800-53-DB] uses an XML schema representing the SP 800-53 revision 4 and 5 catalogs[4]. This schema is a derivative of the SP 800-53 documents and associates controls with baselines using the control-centric approach of Table I. Figure 1 shows how AU-3 from revision 4 is represented in this derivative model. The `baseline-impact` element allocates AU-3 and its two enhancements to the low, moderate, and high baselines in a manner mirroring Table I, except that allocation of AU-3 to the moderate and high baselines is explicit rather than implicit.

Figure 1

```
<control>...
    <number>AU-3</number>...
    <baseline-impact>LOW</baseline-impact>
    <baseline-impact>MODERATE</baseline-impact>
    <baseline-impact>HIGH</baseline-impact>...
    <control-enhancements>
        <control-enhancement>
            <number>AU-3(1)</number>...
            <baseline-impact>MODERATE</baseline-impact>
            <baseline-impact>HIGH</baseline-impact>...
        </control-enhancement>
        <control-enhancement>
            <number>AU-3(2)</number>...
            <baseline-impact>HIGH</baseline-impact>...
        </control-enhancement>
    </control-enhancements>
</control>
```

AU-3 as represented in the derivative model underlying the NIST SP 800-53 database.

NIST's Open Security Controls Assessment Language (OSCAL) project [OSCAL] [Piez] has developed additional non-normative derivative digital models for representing security control catalogs and baselines (called “profiles” in OSCAL terminology). Like the low, moderate, and high baselines defined in the SP 800-53B document, an OSCAL profile references a set of controls, rather than controls referencing profiles to which they are assigned as in the SP 800-53 database schema. Figure 2 shows AU-3 as included in an OSCAL XML representation of the moderate baseline from SP 800-53B. The `with-id` element specifies that AU-3 and its control enhancement AU-3(1) are both in the moderate baseline. Because control enhancement inclusions are not nested within their parents, OSCAL can represent baselines such as the privacy baseline where AU-3(3) is included but its parent control is not.

Figure 2

```
<profile xmlns="http://csrc.nist.gov/ns/oscal/1.0">
   <metadata>...</metadata>
   <import href="NIST_SP-800-53_rev5_catalog.xml">
      <include-controls>
         ...
         <with-id>au-3</with-id>
         <with-id>au-3.1</with-id>
         ...
      </include-controls>
   </import>
   ...
   <modify>...</modify>...
</profile>
```

AU-3 inclusion in the OSCAL SP 800-53 rev. 5 moderate baseline profile.

Unlike the SP 800-53 database schema, which does not provide a way to represent arbitrary tailoring of controls, the OSCAL profile model can represent modifications to a control beyond assignment to a baseline using the `modify` element shown in Figure 2. Such tailoring can include setting parameters specified in control definitions, supplementing baseline with additional controls or control enhancements, and many other possibilities. In fact, OSCAL profiles could be defined to specify the ICS overlay low, moderate, and high baselines. For the AU-4 ICS overlay specification shown in Table V, doing so would require adding a `with-id` element for AU-4(1) and the necessary markup and content within the `modify` element to represent the ICS supplemental guidance and rationale statements.

The next two sections discuss specific types of tailoring in greater detail. The next section, which describes the current version of Baseline Tailor, provide examples of baseline modification with the addition of supplemental guidance. The following section discusses an annotation-based tailoring approach, using a controlled vocabulary, planned for the next version of Baseline Tailor.

The Baseline Tailor Experience

I developed the initial version of Baseline Tailor six years ago, prior to the existence of SP 800-53 revision 5, SP 800-53B, or OSCAL. As originally conceived, Baseline Tailor's scope was limited to tailoring the SP 800-53 revision 4 security controls. The software design's early influences included the control catalog as documented in revision 4, and the SP 800-53 database XML representation of the control catalog. The revision 4 control-centric baseline allocation summaries and XML (as shown for AU-3 in Table I and Figure 1) together reinforced in my mind the idea of a user interface centered around individual controls rather than baselines. Another early influence of Baseline Tailor's user interface was the NIST ICS overlay discussed in the second section.

I envisioned the typical user of Baseline Tailor to be someone tasked by their organization with creating a tailored baseline but without intimate knowledge of the SP 800-53 catalog or tailoring methodology. To better understand my intentions, consider the following “user story” [Cohn]. Suppose Alice provides IT support for a manufacturing firm. Alice's boss Bob tasks Alice with developing a security plan for a 3D printer the company uses to produce replacement parts for internal use. Alice has a rudimentary understanding of the cyber-risk management process described in [JTF18] but has minimal experience selecting and tailoring security controls. Therefore, she decides to tailor an existing baseline, such as one of the SP 800-53 or ICS overlay baselines. Baseline Tailor is intended for users like Alice who would benefit from partial automation of the security control selection process.

I implemented Baseline Tailor in XForms [XForms], an XML application for specifying forms for the web. A user interface coded using XForms contains a `model` element whose content includes a collection of data “instances” that taken together express the “model” portion of the XForms model-view controller software pattern [Krasner]. Using transformation stylesheets [XSLT] to strip out markup and content from the SP 800-53 database XML download not needed for the Baseline Tailor user interface (the “. . .” portions of Figure 1), I could create much of my implementation's `model` element. Therefore, using the SP 800-53 database XML to implement Baseline Tailor had strong appeal. Figure 3 shows the Baseline Tailor security control editor user interface after a user has selected AU-3 from the Audit and Accountability family to edit. Figure 3 bears a strong resemblance to AU-3's presentation in the ICS overlay as shown in Table IV but without the ICS supplemental guidance.

Figure 3

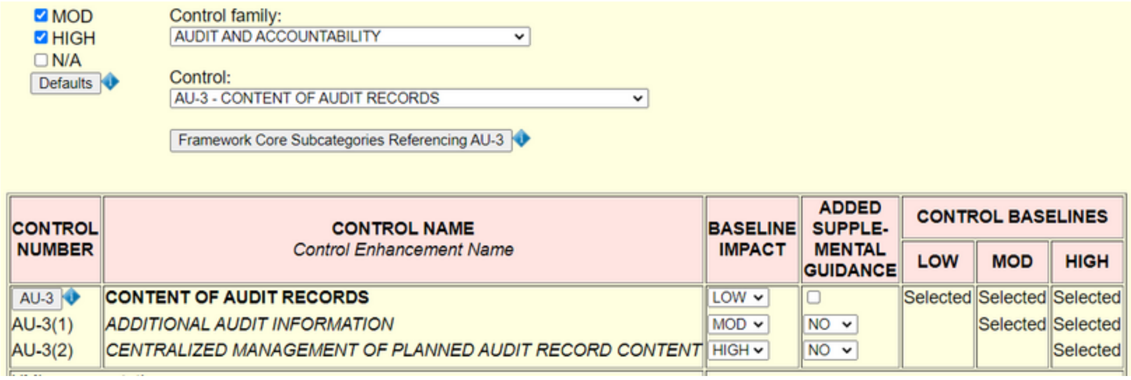

☑ MOD
☑ HIGH
☐ N/A

Control family: AUDIT AND ACCOUNTABILITY

Control: AU-3 - CONTENT OF AUDIT RECORDS

CONTROL NUMBER	CONTROL NAME *Control Enhancement Name*	BASELINE IMPACT	ADDED SUPPLE-MENTAL GUIDANCE	CONTROL BASELINES		
				LOW	MOD	HIGH
AU-3	**CONTENT OF AUDIT RECORDS**	LOW	☐	Selected	Selected	Selected
AU-3(1)	*ADDITIONAL AUDIT INFORMATION*	MOD	NO		Selected	Selected
AU-3(2)	*CENTRALIZED MANAGEMENT OF PLANNED AUDIT RECORD CONTENT*	HIGH	NO			Selected

AU-3 prior to tailoring in Baseline Tailor security control editor UI.

One portion of the Baseline Tailor XForms model I could not adapt from the SP 800-53 database XML was the model instance for representing a tailored control. This model instance, in keeping with the control-centric design of Baseline Tailor, consists of a top-level `tailoredControl` element representing the current state of the security control editor. Figure 4 shows the `tailoredControl` instance when the security control editor appears as in Figure 3. The default element represents baseline impact prior to tailoring, and the impact element represents baseline impact after tailoring. A value of 1 means the control or control enhancement is allocated to the low baseline, which implies moderate and high as well. A value of 2 indicates allocation to the moderate and high baselines, but not to the low baseline. A value of 3 indicates allocation only to the high baseline. Since no tailoring has occurred, the impact values equal the default values. The rationale element and guidance elements contain prose added by the user to explain a change to a control's baseline allocation or add additional guidance to a control or enhancement. Both elements are empty prior to tailoring. The rationale element and guidance elements each have a Boolean valued @flag attribute, which if true indicates that a resizable text box should appear for the user to provide or edit prose text.

Figure 4

```
<tailoredControl>
  <family>AUDIT AND ACCOUNTABILITY</family>
  <rationale flag="false"/>
  <control number="AU-3">
    <title>CONTENT OF AUDIT RECORDS</title>
    <default value="1"/><impact value="1"/><guidance flag="false"/>
  </control>
  <enhancement number="1">
    <title>ADDITIONAL AUDIT INFORMATION</title>
    <default value="2"/><impact value="2"/><guidance flag="false"/>
  </enhancement>
  <enhancement number="2">
    <title>CENTRALIZED MANAGEMENT OF PLANNED AUDIT RECORD CONTENT</title>
    <default value="3"/><impact value="3"/>guidance flag="false"/>
  </enhancement>
</tailoredControl>
```

AU-3 prior to tailoring as represented in Baseline Tailor model instance.

Returning to our user story, suppose Alice uses the high water mark method mentioned in the Introduction section to categorize the 3D printer as a low-impact system and uses the ICS overlay's low baseline as a starting point for selecting security controls. Suppose also that Bob has warned Alice that one of the firm's employees, Charlene, is untrustworthy and might possibly be using the 3D printer for unauthorized purposes. Bob requests that Alice configure the printer's logging system to generate audit records that include more-detailed printer usage information, but only for employees under suspicion. Using Baseline Tailor, Alice modifies AU-3 from the SP 800-53 revision 4 default to (1) add the ICS supplemental guidance for AU-3 (from Table IV) and (2) change AU-3(1)'s baseline impact from moderate to low, thus adding AU-3(1) to the low baseline.

Figure 5 shows the security control editor after Alice's modifications. Her actions cause “Added” to appear in the “LOW” column and text boxes to appear in which she added the ICS overlay supplemental guidance and her own rationale justifying the baseline allocation change. The “XML representation” text box on the left displays current state of the tailoredControl element. This XML captures the change in a machine-readable format.

Figure 5

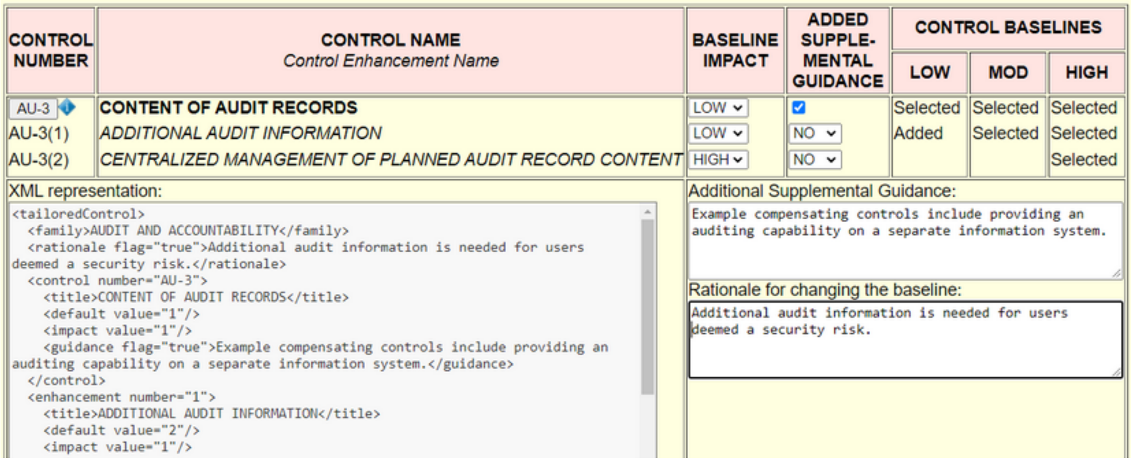

CONTROL NUMBER	CONTROL NAME *Control Enhancement Name*	BASELINE IMPACT	ADDED SUPPLE-MENTAL GUIDANCE	CONTROL BASELINES		
				LOW	MOD	HIGH
AU-3	**CONTENT OF AUDIT RECORDS**	LOW	☑	Selected	Selected	Selected
AU-3(1)	*ADDITIONAL AUDIT INFORMATION*	LOW	NO	Added	Selected	Selected
AU-3(2)	*CENTRALIZED MANAGEMENT OF PLANNED AUDIT RECORD CONTENT*	HIGH	NO			Selected

XML representation:

```
<tailoredControl>
  <family>AUDIT AND ACCOUNTABILITY</family>
  <rationale flag="true">Additional audit information is needed for users
deemed a security risk.</rationale>
  <control number="AU-3">
    <title>CONTENT OF AUDIT RECORDS</title>
    <default value="1"/>
    <impact value="1"/>
    <guidance flag="true">Example compensating controls include providing an
auditing capability on a separate information system.</guidance>
  </control>
  <enhancement number="1">
    <title>ADDITIONAL AUDIT INFORMATION</title>
    <default value="2"/>
    <impact value="1"/>
```

Additional Supplemental Guidance:

Example compensating controls include providing an auditing capability on a separate information system.

Rationale for changing the baseline:

Additional audit information is needed for users deemed a security risk.

AU-3 after adding supplemental guidance and altering AU-3(1)'s baseline impact.

My first major update to Baseline Tailor, implemented within a year of the initial version, added the ability to browse the Cybersecurity Framework [NIST18], a hierarchically-organized taxonomy of security requirements intended to facilitate communication among stakeholders of an organization's current or target security posture. This update integrated the dense, information-rich SP 800-53 catalog model with the Cybersecurity Framework's top-down, easy-to-navigate model [Lubell16]. The result enabled users to easily use the Cybersecurity Framework and SP 800-53 together — without having to leaf back and forth between two page-oriented documentary specifications or deal with spreadsheets. Integrating the Cybersecurity Framework with SP 800-53 turned out to be a bigger success for Baseline Tailor [Kanowitz] than the control-centric tailoring user interface.

In retrospect, multiple factors led to Baseline Tailor's representation of a tailored control being a mismatch for the needs of baseline authors. I made the unfortunate choice of favoring the control-centric implied model shown in Table I and implemented in the SP 800-53 database over the more baseline-centric implied model exemplified in Table II. Neither SP 800-53B nor OSCAL existed at the time, so I was left to my own devices to create a `tailoredControl` model instance. The SP 800-53 database XML was an attractive option for me because XForms implementation is easiest when the underlying data already exists in XML. However, the SP 800-53 database XML is structured for search and traversing a database of controls and not for representing a baseline.

Baseline Tailor 2.0: Tailoring as Annotation

Since releasing Baseline Tailor to the public, users suggested that the security control editor user interface should be more baseline-centric. Specifically, users want to be able to tailor an existing baseline and save the results as a single XML (or JSON) document, rather than having to generate XML snippets from multiple tailoring actions and piece them together in an XML text editor. They also wanted to be able to export a partially or fully tailored baseline to a file and import it later. For a while, I was stuck on how best to revise the “tailoring” part of Baseline Tailor's XForms model. I then had an epiphany that shifted my way of thinking about tailoring away from allocation and instead toward annotation. Three things contributed to my epiphany. The first, of course, was the OSCAL baseline-centric profile model. OSCAL is an appealing format for importing or exporting baselines, given its ability to represent tailoring. Also, since the low, moderate, and high baselines from SP 800-53 revision 4 and SP 800-53B are already available in OSCAL format, I could easily give users a choice of these six baselines as a starting point for tailoring.

The second contributor was SP 800-171 (Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations) [Ross]. SP 800-171, unlike SP 800-53, is targeted to a specific user community: nonfederal organizations that receive United States government funding. And SP 800-171's scope is limited to controlled unclassified information (CUI). CUI includes PII as well as proprietary and other sensitive (but not classified) information. SP 800-171 has gotten more attention recently resulting from increased awareness of the need to protect information pertaining to critical infrastructure coupled with a recent uptick in critical infrastructure cyber-attacks [Slonka].

SP 800-171 consists of a set of 110 security requirements for protecting CUI, each of which maps to one or more of a set of SP 800-53 revision 4 controls. This set of controls is specified using an annotated tailoring of the SP 800-53 revision 4 moderate baseline, defined in the SP 800-171 document as tables for each control family. Every control and enhancement in the moderate baseline has one of four annotation symbols indicating whether it:

- Does not directly protect the confidentiality of CUI.
- Applies only to federal agencies and not to nonfederal entities.
- Is expected to be routinely implemented irrespective of SP 800-171 requirements.
- Is traceable to one of the 110 SP 800-171 requirements.

My implementation of Baseline Tailor 2.0, currently a work in progress, includes a user interface for defining tailoring actions to be applied to a baseline as a set of ordered pairs of tailoring symbols and corresponding criteria. Figure 6 shows this user interface populated with the four SP 800-171 tailoring symbols and corresponding criteria. Figure 7 shows the model instance data representing the user-entered tailoring actions in Figure 6. Note that this user interface is sufficiently general enough to define any set of ordered pairs for any annotation-based tailoring scenario, not just SP 800-171. For example, it could be used to define an arbitrary-length set of tailoring actions for an organization-specific tailoring of the revision 5 low baseline.

Figure 6

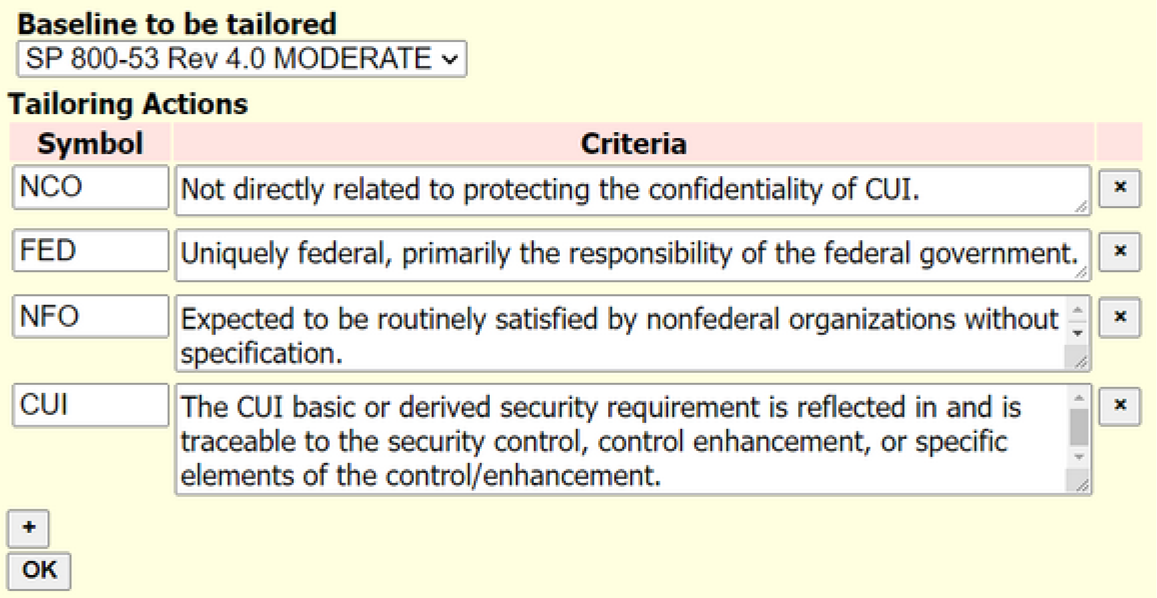

Baseline to be tailored

SP 800-53 Rev 4.0 MODERATE

Tailoring Actions

Symbol	Criteria
NCO	Not directly related to protecting the confidentiality of CUI.
FED	Uniquely federal, primarily the responsibility of the federal government.
NFO	Expected to be routinely satisfied by nonfederal organizations without specification.
CUI	The CUI basic or derived security requirement is reflected in and is traceable to the security control, control enhancement, or specific elements of the control/enhancement.

New tailoring approach applied to revision 4 moderate baseline as in SP 800-171.

Figure 7

```
<action>
  <symbol>NCO</symbol>
  <criteria>Not directly related to protecting the confidentiality of CUI.</criteria>
</action>
<action>
  <symbol>FED</symbol>
  <criteria>Uniquely federal, primarily the responsibility of the federal government.
</criteria>
</action>
<action>
  <symbol>NFO</symbol>
  <criteria>Expected to be routinely satisfied by nonfederal organizations without
specification.</criteria>
</action>
<action>
  <symbol>CUI</symbol>
  <criteria>The CUI basic or derived security requirement is reflected in and is
traceable to the security control, control enhancement, or specific elements of the
control/enhancement.</criteria>
</action>
```

Tailoring actions represented in an XForms model instance.

Figure 8 lists each AU control and enhancement in the revision 4 moderate baseline and specifies via SP 800-171 tailoring action symbols which are traceable to the 110 CUI protection requirements and which are not. For example, AU-3 and AU-3(1) both have the “CUI” tailoring symbol and therefore are traceable to a requirement. AU-1 has the “NFO” tailoring symbol, indicating it is expected to be implemented (satisfied) even though no SP 800-171 requirement maps to it. AU-4, AU-6(1), AU-7(1), and AU-11 are not directly related to protecting CUI confidentiality and no SP 800-171 requirements map to them. The Baseline Tailor 2.0 user interface, in addition to enabling users to add supplemental guidance and rationale statements as does the current user interface shown in Figure 5, will also allow users to associate an annotation symbol from the set of tailoring actions they define. In this way, the user interface will be able to generate a structured digital representation of information that a baseline developer would normally have to record manually in a table within a document, or in a spreadsheet.

Figure 8

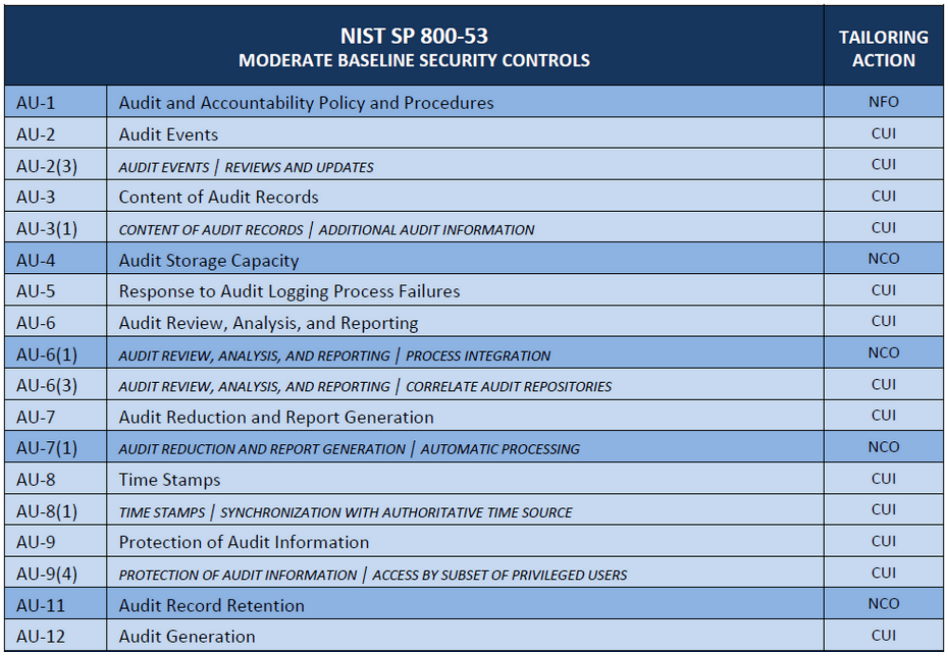

NIST SP 800-53 MODERATE BASELINE SECURITY CONTROLS		TAILORING ACTION
AU-1	Audit and Accountability Policy and Procedures	NFO
AU-2	Audit Events	CUI
AU-2(3)	*AUDIT EVENTS | REVIEWS AND UPDATES*	CUI
AU-3	Content of Audit Records	CUI
AU-3(1)	*CONTENT OF AUDIT RECORDS | ADDITIONAL AUDIT INFORMATION*	CUI
AU-4	Audit Storage Capacity	NCO
AU-5	Response to Audit Logging Process Failures	CUI
AU-6	Audit Review, Analysis, and Reporting	CUI
AU-6(1)	*AUDIT REVIEW, ANALYSIS, AND REPORTING | PROCESS INTEGRATION*	NCO
AU-6(3)	*AUDIT REVIEW, ANALYSIS, AND REPORTING | CORRELATE AUDIT REPOSITORIES*	CUI
AU-7	Audit Reduction and Report Generation	CUI
AU-7(1)	*AUDIT REDUCTION AND REPORT GENERATION | AUTOMATIC PROCESSING*	NCO
AU-8	Time Stamps	CUI
AU-8(1)	*TIME STAMPS | SYNCHRONIZATION WITH AUTHORITATIVE TIME SOURCE*	CUI
AU-9	Protection of Audit Information	CUI
AU-9(4)	*PROTECTION OF AUDIT INFORMATION | ACCESS BY SUBSET OF PRIVILEGED USERS*	CUI
AU-11	Audit Record Retention	NCO
AU-12	Audit Generation	CUI

AU family moderate baseline tailoring (SP 800-171, Table E-3)

The third contributor to my shift in thinking about tailoring was a realization that an annotation-based approach obviates the need for Baseline Tailor 2.0 to explicitly support removal of a control or control enhancement from a baseline. The current Baseline Tailor user interface provides drop-down widgets in the "BASELINE IMPACT" for additions and removals, as shown in Figure 3 and Figure 5. However, as SP 800-171 shows, the annotation symbols "NCO", "FED", and "NFO" can be interpreted as removal plus rationale. Thus, a Baseline 2.0 user will be able to "remove" a control from a baseline simply by assigning it an appropriate tailoring annotation from a user-defined set of tailoring actions.

Adding a control to a baseline will still require more than assigning an annotation symbol. The user will need to select a control or enhancement using "Control family" and "Control" drop-downs (as in the top part of Figure 3), with choices filtered to include only those controls/enhancements in the control catalog that are not already included in the baseline being tailored. The user would then use a self-authored tailoring action to provide a rationale for adding the control/enhancement. For example, returning once again to our user story, with Baseline Tailor 2.0 Alice could define a tailoring action with symbol "THREAT" and criteria "Added to the baseline to protect against a threat specific to the system or organization." After adding AU-3(1) to the low baseline, Baseline Tailor would prompt her to assign an annotation symbol. Alice would choose "THREAT" from the drop-down list and could add her rationale statement ("Additional audit information is needed for users deemed a security risk.") as added supplemental guidance to further justify the tailoring action.

Conclusion

This paper discusses my multi-year education in the nuances of SP 800-53 security control baselines. While developing Baseline Tailor I discovered that SP 800-53 has more than one implied model for tailoring, and that multiple derivative digital models exist outside the normative standard. Subsequently, my mental model of a baseline evolved from a collection of individual controls with assigned impacts to an annotation of a pre-existing baseline. The upcoming version of Baseline Tailor will benefit from the lessons I learned during this journey.

No single baseline model is *the* correct one. The right model for a tool developer depends on the type of tool being implemented. A control-centric model such as the SP 800-53 database schema may be optimal for someone developing a user interface for browsing the catalog and its published baselines. A baseline-centric model makes the most sense for a tool that consumes and/or produces a baseline. Computer-readable digital models are more developer-friendly when the model's data format is a good fit for the implementation language they are using. Models defined as prose in standards establish a ground truth from which implementation-friendly models can be derived.

Like other successful standards, SP 800-53 exists as part of a larger ecosystem of stakeholders. The SP 800-53 ecosystem includes developers of SP 800-53, other standards supporting the RMF process, communities of interest who develop overlays and other discipline-specific guidance, policymakers who advocate for and may mandate the use of these standards, tool developers, and last but not least, those responsible for performing the RMF steps necessary to secure information systems in a risk-informed manner. As a tool developer, I found it essential to be aware of the concerns of all these stakeholder groups. I also learned not to let my predisposition for a single model or implementation technology distract me from considering alternatives that might be a better fit for my tool's purpose.

Note

I am grateful to the OSCAL project for creating digital representations of the SP 800-53 revision 4 and SP 800-53B baselines. These profiles not only contributed to my evolving understanding of baselines, but also proved handy for computing control and control enhancement counts[1] [2].

References

[BT] Baseline Tailor. Available at https://www.nist.gov/services-resources/software/baseline-tailor

[Cohn] Cohn M (2004) *User Stories Applied: For Agile Software Development* (Addison-Wesley).

[JTF13] Joint Task Force Transformation Initiative (2013) Security and Privacy Controls for Federal Information Systems and Organizations. (National Institute of Standards and Technology), Revision 4. `https://doi.org/10.6028/NIST.SP.800-53r4`

[JTF18] Joint Task Force Transformation Initiative (2018) Risk management framework for information systems and organizations:: a system life cycle approach for security and privacy. (National Institute of Standards and Technology, Gaithersburg, MD), NIST SP 800-37r2. `https://doi.org/10.6028/NIST.SP.800-37r2`

[JTF20] Joint Task Force Interagency Working Group (2020) Security and Privacy Controls for Information Systems and Organizations. (National Institute of Standards and Technology), Revision 5. `https://doi.org/10.6028/NIST.SP.800-53r5`

[JTF20B] Joint Task Force (2020) Control Baselines for Information Systems and Organizations. (National Institute of Standards and Technology). `https://doi.org/10.6028/NIST.SP.800-53B`

[Kanowitz] Kanowitz S (2017) A better way to build on the NIST framework. *GCN* 36(6). Available at https://gcn.com/articles/2017/10/05/dig-it-cyber-baseline-tailor.aspx

[Krasner] Krasner GE, Pope ST (1988) A cookbook for using the model-view controller user interface paradigm in Smalltalk-80. *Journal of Object-Oriented Programming* 1(3).

[Lubell14] Lubell J (2014) XForms User Interfaces for Small Arcane Nontrivial Datasets. Proceedings of Balisage: The Markup Conference, Balisage Series on Markup Technologies, vol. 13. (Washington, DC). `https://doi.org/10.4242/BalisageVol13.Lubell01`

[Lubell16] Lubell J (2016) Integrating Top-down and Bottom-up Cybersecurity Guidance using XML. Proceedings of Balisage: The Markup Conference, Balisage Series on Markup Technologies, vol. 17. (Washington, DC). `https://doi.org/10.4242/BalisageVol17.Lubell01`

[Lubell20] Lubell J (2020) A Document-based View of the Risk Management Framework. Proceedings of Balisage: The Markup Conference, Balisage Series on Markup Technologies, vol. 25. (Washington, DC). `https://doi.org/10.4242/BalisageVol25.Lubell01`

[NIST06] National Institute of Standards and Technology (2006) Minimum security requirements for federal information and information systems. (National Institute of Standards and Technology, Gaithersburg, MD), NIST FIPS 200. `https://doi.org/10.6028/NIST.FIPS.200`

[NIST18] National Institute of Standards and Technology (2018) Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Cybersecurity White Paper. `https://doi.org/10.6028/NIST.CSWP.02122014`

[OSCAL] OSCAL: the Open Security Controls Assessment Language. Available at https://pages.nist.gov/OSCAL/

[Ross] Ross R, Pillitteri V, Dempsey K, Riddle M, Guissanie G (2020) Protecting controlled unclassified information in nonfederal systems and organizations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST SP 800-171r2. `https://doi.org/10.6028/NIST.SP.800-171r2`

[Piez] Piez W (2019) The Open Security Controls Assessment Language (OSCAL): schema and metaschema. (Balisage Series on Markup Technologies, Washington, DC), Vol. 23. `https://doi.org/10.4242/BalisageVol23.Piez01`

[SCOR] Security and Privacy Control Overlay Repository (SCOR) *NIST Risk Management Framework*. Available at https://csrc.nist.gov/projects/risk-management/sp800-53-controls/overlay-repository

[Slonka] Slonka K (2020) Managing Cyber Security Compliance Across Business Sectors. *Issues In Information Systems* 21(1). `https://doi.org/10.48009/1_iis_2020_22-29`

[SP800-53-DB] SP 800-53 Release Search *NIST Risk Management Framework*. Available at https://csrc.nist.gov/Projects/risk-management/sp800-53-controls/release-search

[Stouffer] Stouffer K, Pillitteri V, Lightman S, Abrams M, Hahn A (2015) Guide to Industrial Control Systems (ICS) Security. (National Institute of Standards and Technology), NIST SP 800-82r2. `https://doi.org/10.6028/NIST.SP.800-82r2`

[XForms] XForms 1.1 (2009). W3C Recommendation. Available at http://www.w3.org/TR/xforms11/

[XSLT] XSL Transformations (XSLT) Version 2.0 (2007), W3C Recommendation. Available at https://www.w3.org/TR/xslt20/

[1] The SP 800-53 revision 4 catalog has 256 security controls and 666 control enhancements. SP 800-53 revision 5 has an even bigger catalog consisting of 322 controls and 867 enhancements.

[2] For example, the SP 800-53 revision 5 low baseline contains just 149 controls and control enhancements combined. Even the revision 5 high baseline, with 370 controls plus control enhancements, contains less than half of the catalog.

[3] The acronym SCADA stands for Supervisory Control and Data Acquisition. SCADA systems gather and process data and apply operational controls over long distances [Stouffer].

[4] At the time of this writing (July 2021), the schema is available at http://scap.nist.gov/schema/sp800-53/feed/2.0/sp800-53-feed_2.0.xsd.