ISAC ANNUAL SUMMIT 2026 • LOCAL GOVERNMENT TRACK

Building Defensible *Cyber Maturity* in Resource-Constrained Governments

Build the methodology. Own what compounds.

Garrett Ragland-Helf
CISO Advisory Analyst & Group Facilitator, MS-ISAC Leadership Mentoring Program

THE THESIS

Compliance is an output.
Foundation is where the system starts.

COMPLIANCE CYCLE

Full price every cycle

- Produce a number, file the report.
- Nothing operationally changes.
- Starts from scratch every audit.
- Satisfies a requirement. Doesn't drive improvement.

FOUNDATION CYCLE

Compounds year over year

- Year one builds the artifact.
- Year two inherits it. Year three updates it.
- When the auditor arrives, you report.
- The framework changes. The system stays.

The real difference shows up in year three, not year one.

WHAT MAKES IT DEFENSIBLE

Four properties. *Without all four, you have a scoring exercise.*

1

Consistent

Two qualified assessors land on substantially the same scores. The methodology does the work, not the individual.

2

Repeatable

You can run it again in 18 months and see if you moved. Without this, you have a snapshot, not a program.

3

Evidence-Backed

Every score traces to something concrete. A blank evidence column means the score won't survive an auditor.

4

Decision-Driving

Same data, three views. Operators, executives, and your board each get the version they can act on.

If you don't have all four, you have a scoring exercise. Not a methodology.

THE METHODOLOGY

Five stages. One artifact.

Each stage is where defensibility is won or lost — and where the four properties are built or broken.

1	2	3	4	5
Scope	**Evidence**	**Score**	**Weight**	**Report**
Right-size the framework to the entity	Interview, observe, collect artifacts	Apply rubric uniformly across objectives	Renormalize for entity size and risk	Communicate to three audiences

STAGE 1: SCOPE

Scope is where defensibility starts.

Lock all four with leadership at scope, in writing. Everything downstream depends on this slide.

1

Entity profile

Size, mission, sector, resource posture. Drives which control objectives apply at what level of rigor.

2

Regulatory drivers

CJIS, HIPAA, FERPA, IRS 1075, NERC CIP, state law. Determines which compliance audiences your evidence has to satisfy.

3

Strategic drivers

Board priorities, planned grants, insurance renewals, recent incidents. What does leadership most need this to inform?

4

Time and resource budget

Realistic hours from your team. A six-week assessment that finishes beats a six-month one that drifts.

STAGE 2: EVIDENCE

Be your own toughest interviewer.

ASK YOUR TEAM, GET A YES

"Do we have a backup policy?"

"Yes" tells you nothing about whether backups work.

"Do we train users on phishing?"

"Yes" can mean an annual 10-minute video.

"Is our IR plan tested?"

"Yes" can mean it was reviewed in 2019.

ASK YOUR TEAM, GET EVIDENCE

"Walk me through our last successful restore."

Forces detail; surfaces whether anyone has actually tested.

"Show me the most recent training campaign and click rates."

Reveals frequency, content, and outcomes.

"When did we last exercise the IR plan, and what changed after?"

Tests both currency and learning loop.

In self-assessment, the temptation to skip evidence is highest. "I know we do that" is the enemy.

STAGE 3: SCORE

A rubric that survives second-guessing.

MATURITY LEVELS (NIST-DERIVED)

0	**Non-existent**	Not in place; not even informally addressed.
1	**Initial**	Ad hoc, reactive, individual-dependent.
2	**Repeatable**	Some documentation, applied inconsistently.
3	**Defined**	Documented, understood, applied consistently.
4	**Managed**	Measured, reviewed, improved.
5	**Optimized**	Continuous improvement, embedded in culture.

RUBRIC RULES THAT HOLD UP

- Each level has written, observable criteria. If you can't tell me what evidence proves Level 3, you don't have a rubric.
- Score the lowest defensible level. Aspirational scoring inflates trends.
- Half-points are noise. Stay on integers; force decisions.
- Document your evidence next to the score. The score and the why live together.

RUBRIC IN PRACTICE

What defensible scoring looks like.

Example control: Backup and Recovery (PR.IR / RC.RP family)

Lvl	Observable criteria	Evidence required to score this level
1	Ad hoc backups. No documented policy.	Conversation only. No artifact required.
2	Backup policy exists. Backups run on schedule. No regular restore testing.	Backup software config. Schedule documentation.
3	Policy approved by leadership. Restores tested at least quarterly. RPO/RTO defined.	Approved policy. Quarterly test schedule. Most recent restore log.
4	Test results tracked over time. RPO/RTO measured against policy. Failures trigger improvements.	Test result history. Variance reports. Improvement tickets closed.
5	Immutable backups. Continuous integrity monitoring. Automated DR exercises.	Immutable storage config. Monitoring dashboard. DR runbook + after-action.

Two principles: every level has observable criteria, and the score doesn't move without the evidence column filled in.

STAGE 4: WEIGHT

Equal weight is not fair weight.

Why weight at all?

If a school district scores Level 4 on "Acceptable Use Policy" and Level 1 on "Backup and Recovery," treating those equally produces a misleading average. Weighting tells the score what actually matters most for this entity, given its size, mission, and risk profile.

Three weighting principles

- Weights are set BEFORE scores. Otherwise, you're tuning to a desired result.
- Weights are documented and approved by the entity at scope, not after.
- When entity size changes, weights renormalize so total still sums to 100%.

EQUAL WEIGHTING

RISK-WEIGHTED (ILLUSTRATIVE FOR A SMALL DISTRICT)

Same five functional areas. Different signal about where to invest.

ID = Identify • PR = Protect • DE = Detect • RS = Respond • RC = Recover

WORKED EXAMPLE

When the methodology has to flex.

Functional area	Larger entity	Smaller entity	Why the shift
Identify	20%	15%	Smaller asset surface; less governance overhead applies
Protect	25%	35%	Highest leverage for under-resourced entities; preventative > detective
Detect	20%	15%	Limited tooling and staff make high detect maturity unrealistic
Respond	20%	20%	IR planning is high-leverage at any size; weight stays flat
Recover	15%	15%	Backup discipline matters equally regardless of entity size
Total	**100%**	**100%**	*Renormalized; nothing dropped, redistribution explicit*

Illustrative weights only. Yours should be defined and approved with the entity at scope.

STAGE 5: REPORT

One assessment. Three audiences.

Operators

IT STAFF, SYSADMINS, SECURITY ANALYSTS

Specific, prioritized findings tied to control objectives. What to fix first, second, third. Concrete enough to scope work.

Executives

CIO, CISO, CITY MANAGER

Maturity scores by function, trend vs prior cycle, top three risks in plain language, budget asks tied to risk reduction.

Governing body

SCHOOL BOARD, COUNCIL, COMMISSIONERS

One-page narrative. Are we safer than last year? What is the entity doing well, what isn't, what does leadership need.

Same data. Three views. The mistake is sending one report to all three audiences and hoping each finds their part.

AVOIDING THE TRAPS

Common failure points in assessments.

Scope drift

Self-assessment starts as a maturity check and morphs into active remediation. Scope tightly or split the work.

Score inflation

Pressure to give your own program 'good news.' In self-assessment this is the single biggest risk.

Assertion in place of evidence

Accepting 'we do that' from a colleague you trust. Self-assessment needs MORE documentation, not less.

No remediation owner

Findings without an owner, due date, and budget line are wishes. Build owner/date into the workbook.

Weighting after the fact

Adjusting weights once you see scores. Lock weights with leadership at scope, in writing.

Disappearing baseline

Next person inherits no documentation, rebuilds from scratch. Outlast the assessor.

WHY THIS WORKS

The artifact compounds.

Year one is the investment. Every cycle after pays back. This is what makes proactive practical for resource-constrained orgs.

YEAR 1

Build the foundation

- Workbook structure
- NIST CSF → compliance crosswalk
- Rubric with observable criteria
- First scoring cycle
- Initial evidence repository

YEAR 2

The artifact inherits

- Year 1 workbook reused
- Crosswalk extends to new mandates
- Score deltas vs baseline visible
- Policy library v1 published
- Trend column added

YEAR 3

Retrieval, not creation

- Three cycles of trend data
- Mature crosswalk and policy library
- Audit answer: "pull the workbook"
- Grant application: "pull the workbook"
- Board update: "pull the workbook"

The reactive cycle costs full price every year. The proactive system pays back from year two onward.

When the auditor shows up, *you're not starting.* *You're reporting.*

Mandated audit?	*Cyber insurance renewal?*	*Grant application?*
Pull the workbook.	**Pull the workbook.**	**Pull the workbook.**

WHAT TO TAKE HOME

Five things to do Monday morning.

1. Start the crosswalk. NIST CSF in column A, every compliance audience you owe in columns B onward. One source of truth.
2. Audit your scoring rubric. Each maturity level needs written, observable criteria. If it doesn't, that's the first fix.
3. Replace yes/no questions even with your own team. "Show me" or "walk me through" beats "do we have" every time.
4. Lock weights at scope, in writing, with leadership. Never adjust after seeing scores.
5. Split the report. Operators, executives, and your governing body each get their own artifact. Same data, three views.

THE GOAL

The goal isn't a defensible assessment.

It's a proactive practice.

The artifact that compounds, the rubric that survives turnover, the cadence that doesn't depend on anyone else's funding cycle.

Garrett Ragland-Helf
garrett.helf@apollo-is.com
CISO Advisory Analyst, Apollo Information Systems • Group Facilitator, MS-ISAC Leadership Mentoring Program