

The Current State of Cybersecurity in GRC

Have our investments made us safe?



EAST BAY

November 8th,
2024
ISC2 East Bay Fall
Conference
@Las Positas
Livermore, CA

**GRC – why nobody
likes it anymore.**



Alternate Discussion Title: Enterprise Governance, Risk and Compliance; A New Approach.

Robin Basham, M.Ed., M.IT, ITSM, CISA, CGEIT, CRISC, Archer Certified Consultant, CCM Working Group Regular, Managing Partner, CEO, CISO EnterpriseGRC Solutions, Inc., Board Member ISC2 East Bay 2017 to current - Contact info: robin@enterprisegrc.com

Most of this presentation appeared as more optimistic content in the OCEG RedBook v1 2003, as contributions to OASIS in 2005, at ISACA Silicon Valley Conferences from 2009 to 2014, SF ISACA, ISACA LA, CISO Forums, and to ISC2 East Bay 2017 to Present

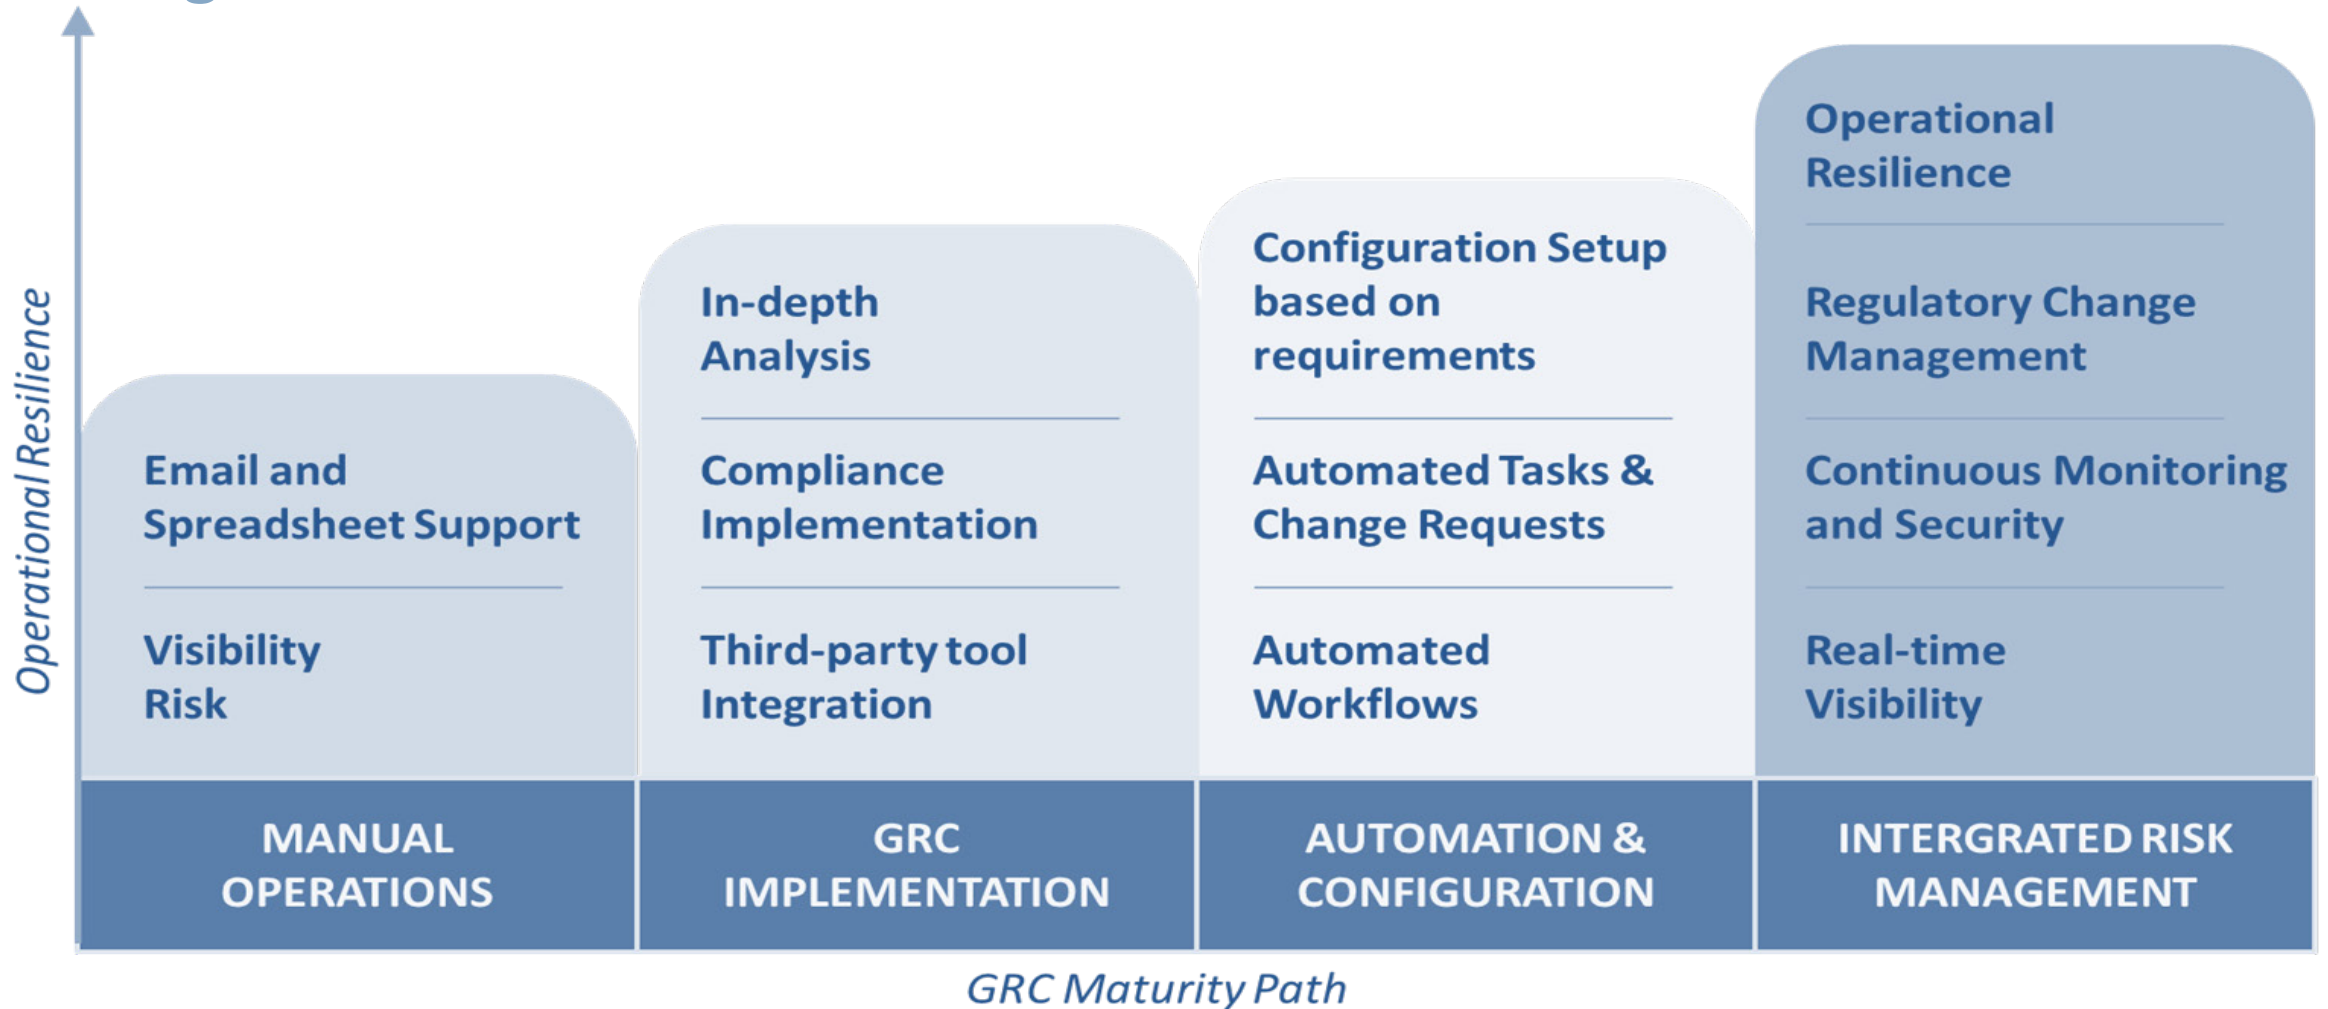
I hope to eventually get something right.

Packaging a Typical Global Enterprise Assessment

- Company represents “Secure” Products and Services
- Desires externally validated Certifications that expand entry to regulated markets
- To meet objectives for half dozen assessments (for example, PCI, SOC, ISO-27n, NIST-800-53, HITRUST) InfoSec & Assurance will involve the BU in 6000+ touch points annually (evidenced by sampled process and technology)
- Interfacing externally to make 1800+ control assertions per BU



What are GRC Systems SUPPOSED TO DO? Does this align with the investment made?



Governance Risk and Compliance – 25 years of failing

Corporations are accountable for compliance and risk management across at least four domains: Legal, Technical, Financial, and Operational
In the last 20 years, we've increased Privacy, Social and Environmental Requirements.



In the previous five years, we've added zero trust requirements.

We are currently adding AI compliance to the mix.

Here's what I want to know.

How come every time we fail, we add more compliance?

Compliance is not Security.





True or False – No one wants to be safe.

Stop bad ideas while encouraging new and better ideas

Define and capture failure so we can adapt policy

Reward initiative while calling out what's wrong and why it matters

Foster a culture of commitment, collaboration, and knowledge transfer

We are agile through rapid status meetings and common metrics (the infamous GRC Dashboard)

Tie every service to revenue, configuration to the customer, the process to product, and people to controls -> Build a well-architected GRC.

Could common assurance methods be better served by AI?




Process & Policy Assignment

- Primary Process Owner and Policy Owner
- Sub process inputs and outputs, assets, owners



Map many to one common Control Universe

- Match intent of overall objectives – map their evidence outputs
- Align test steps across this and other implemented test methodologies



Associated Question Bank

- Consistent method to confirm coverage
- Document increased scrutiny

What if I told you
compliance is the easy
part?



Is Risk the hard part?



Write down everything you would include in your responsibilities for risk. Are you adequately empowered to manage it? Can you do it?

The R in GRC – Risk Management

(As Reported by Ernst & Young and PwC, IIA, ISACA 25 years ago)



nistspecialpublication800-39.pdf

The R in GRC is a VERB Risk Management Process (As Prescribed by NIST SP 800-39 and ISACA's The Risk IT Framework)

Establish the context

Identify the risks

Analysis of the risks

Evaluate the risks

Treat the risks

Monitor and review

Communicate and consult

NIST Special Publication 800-39

NIST
National Institute of
Standards and Technology
U.S. Department of Commerce

Managing Information
Security Risk
*Organization, Mission, and Information
System View*

JOINT TASK FORCE
TRANSFORMATION INITIATIVE

INFORMATION SECURITY

Computer Security Division
Information Technology Laboratory
National Institute of Standards and Technology
Gaithersburg, MD 20899-8933

March 2011



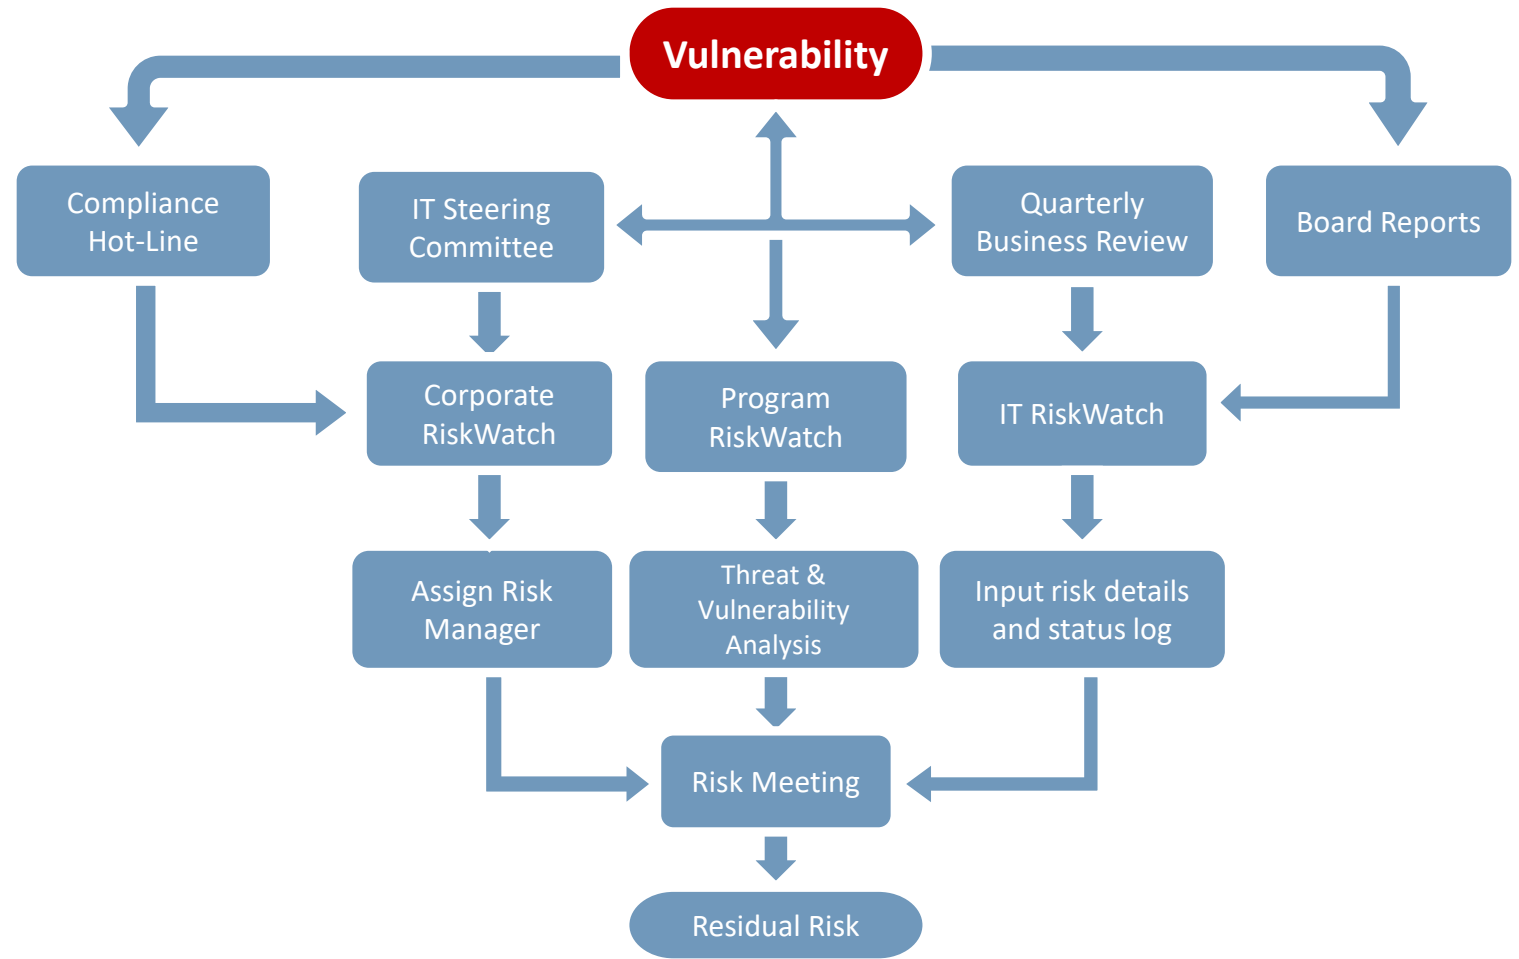
U.S. Department of Commerce
Gary Locke, Secretary

National Institute of Standards and Technology
Patrick D. Gallagher, Director

Table of Contents

CHAPTER ONE INTRODUCTION.....	1
1.1 PURPOSE AND APPLICABILITY	3
1.2 TARGET AUDIENCE.....	3
1.3 RELATED PUBLICATIONS.....	4
1.4 ORGANIZATION OF THIS SPECIAL PUBLICATION.....	5
CHAPTER TWO THE FUNDAMENTALS.....	6
2.1 COMPONENTS OF RISK MANAGEMENT	6
2.2 MULTITIERED RISK MANAGEMENT.....	9
2.3 TIER ONE—ORGANIZATION VIEW.....	11
2.4 TIER TWO—MISSION/BUSINESS PROCESS VIEW	17
2.5 TIER THREE—INFORMATION SYSTEMS VIEW	21
2.6 TRUST AND TRUSTWORTHINESS	23
2.7 ORGANIZATIONAL CULTURE	28
2.8 RELATIONSHIP AMONG KEY RISK CONCEPTS	29
CHAPTER THREE THE PROCESS.....	32
3.1 FRAMING RISK	33
3.2 ASSESSING RISK	37
3.3 RESPONDING TO RISK	41
3.4 MONITORING RISK.....	45
APPENDIX A REFERENCES.....	A-1
APPENDIX B GLOSSARY	B-1
APPENDIX C ACRONYMS.....	C-1
APPENDIX D ROLES AND RESPONSIBILITIES.....	D-1
APPENDIX E RISK MANAGEMENT PROCESS TASKS	E-1
APPENDIX F GOVERNANCE MODELS.....	F-1
APPENDIX G TRUST MODELS	G-1
APPENDIX H RISK RESPONSE STRATEGIES	H-1

Yes or No; Are vulnerabilities the main input for Enterprise Risk?



Sample Cyber Security Model

Why would the goals for Cybersecurity be different from the goals for GRC?

LEVEL 1 Initial	LEVEL 2 Advanced	LEVEL 3 Self-Assessed	LEVEL 4 Integrated	LEVEL 5 Vanguard
<ul style="list-style-type: none"> Minimal cyber awareness Minimal cyber info sharing Minimal cyber assessments and policy & procedure evaluations Little inclusion of cyber into Continuity of Operations Plan (COOP) 	<ul style="list-style-type: none"> Leadership aware of cyber threats, issues and imperatives for cyber security and community cooperative cyber training Informal info sharing/communication in community; working groups established; ad-hoc analysis, little fusion or metrics; professional orgs established or engaged No assessments, but aware of requirement; initial evaluation of policies & procedures Aware of need to integrate cyber security into COOP 	<ul style="list-style-type: none"> Leaders promote org security awareness; formal community cooperative training Formal local info sharing/cyber analysis. initial cyber-physical fusion; informal external info sharing/ cyber analysis and metrics gathering Autonomous tabletop cyber exercises with assessments of info sharing, policies & procedures, and fusion; routine audit program; mentor externals on policies & procedures, auditing and training Include cyber in COOP; formal cyber incident response/recovery 	<ul style="list-style-type: none"> Leaders and orgs promote awareness; citizens aware of cyber security issues Formal info sharing/analysis, internal and external to community; formal local fusion and metrics, initial external efforts Autonomous cyber exercises with assessments of formal info sharing/local fusion; exercises involve live play/metrics assessments Integrate cyber in COOP; mentor externals on COOP integration; formal blended incident response and recovery 	<ul style="list-style-type: none"> Awareness a business imperative Fully integrated fusion /analysis center, combining all-source physical and cyber info; create and disseminate near real world picture Accomplish full-scale blended exercises and assess complete fusion capability; involve/ mentor other communities/entities Continue to integrate cyber in COOP; mentor externals on COOP integration; formal blended incident response and recovery

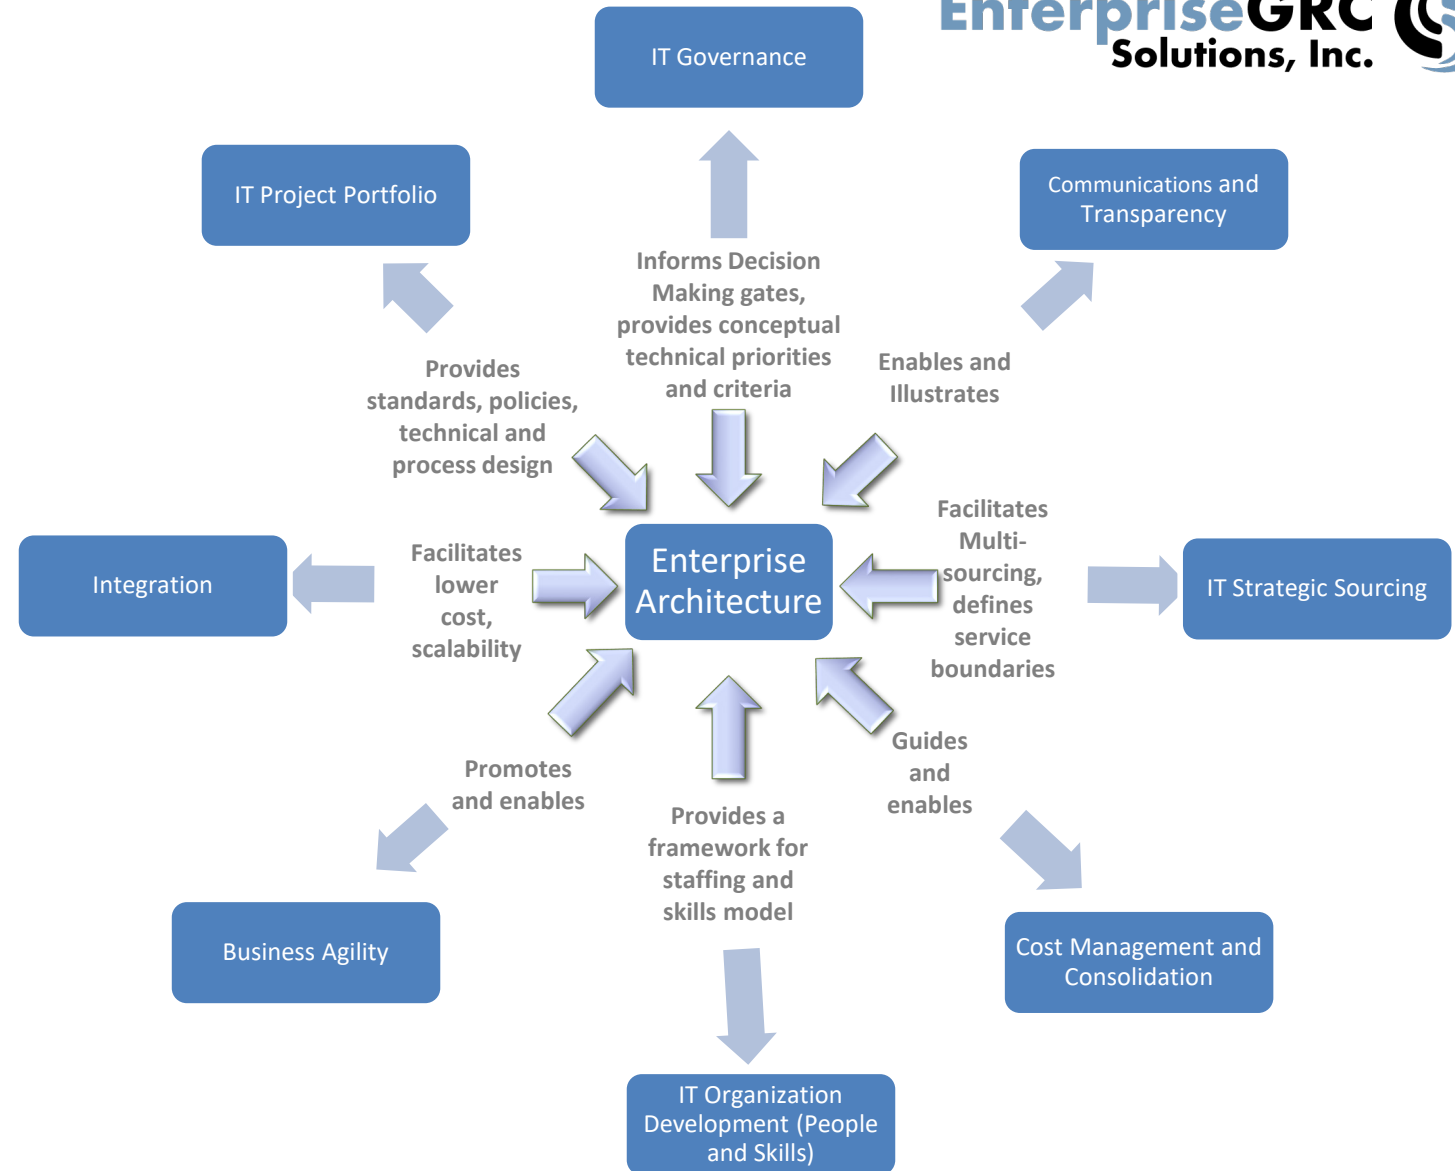
Is GRC Everything or Nothing?



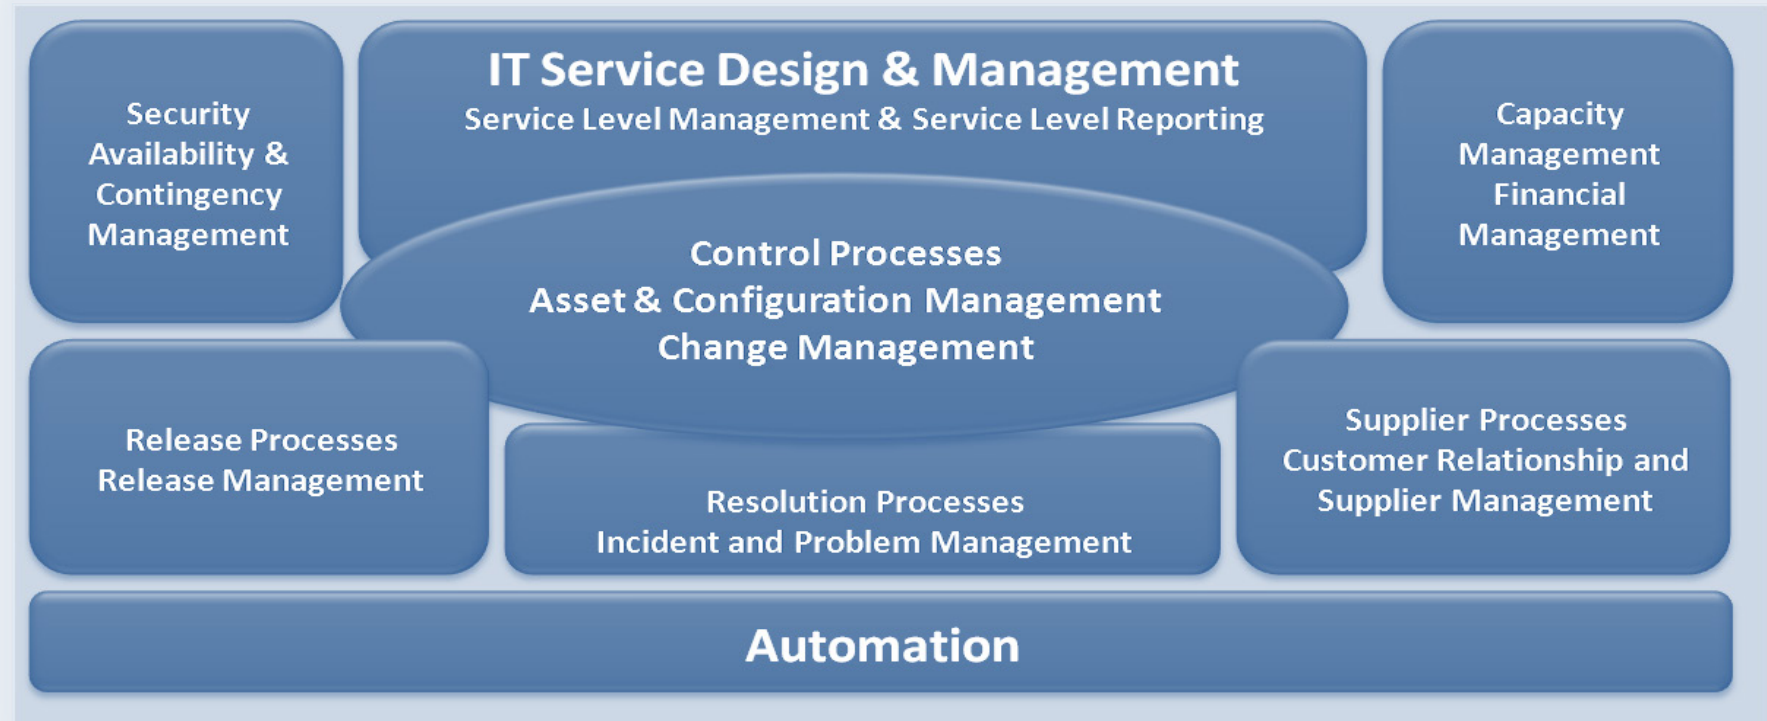
“GRC is the integrated collection of capabilities that enable an organization to achieve Principled Performance - the ability to reliably achieve objectives, address uncertainty, and act with integrity.”

<https://www.oceg.org/ideas/what-is-grc/>

Enterprise GRC Architecture is lost in the extensive scope of the entire EA. People will fail in GRC because of the inevitable land grab and the pointless attempt to do too many things.



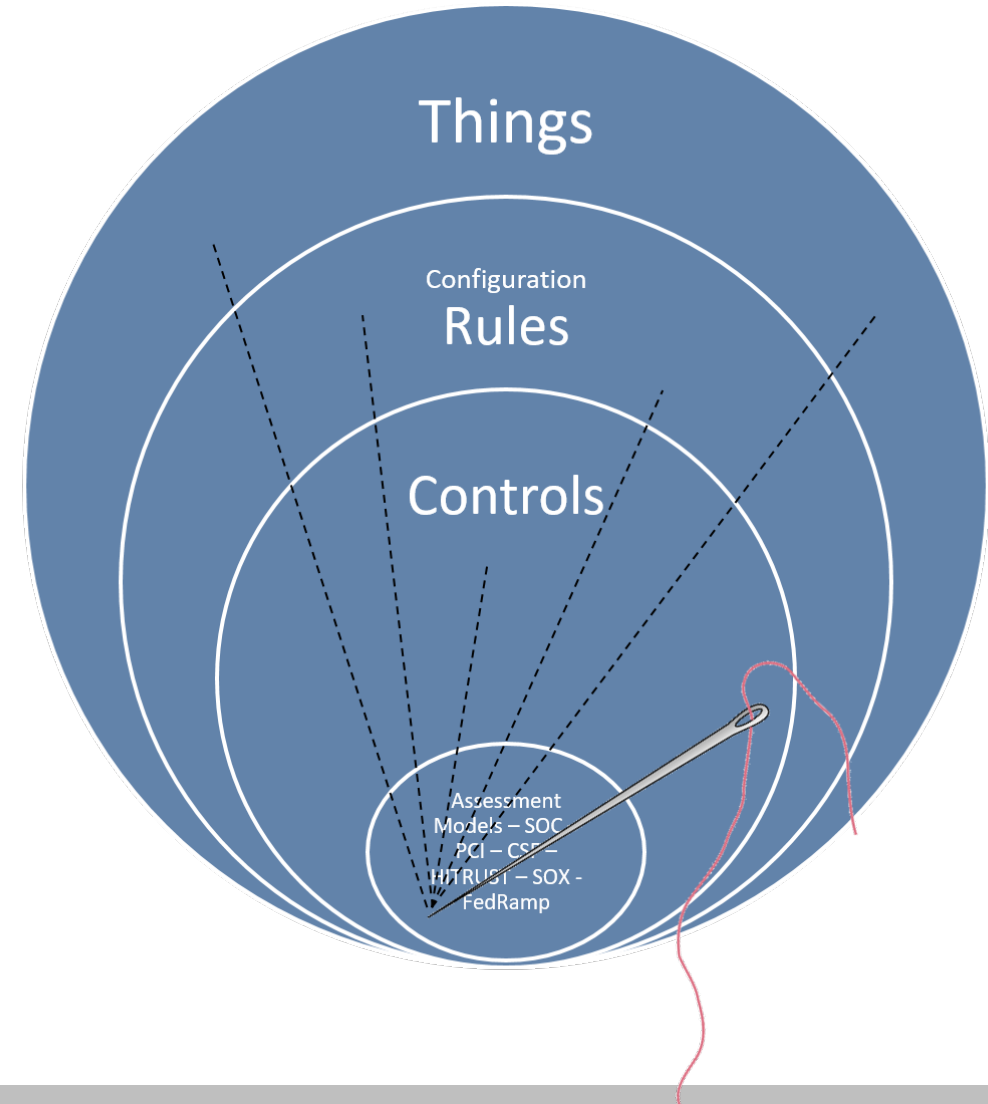
What if we refocus GRC on Technology Operations Risk? What will happen to cybersecurity if we do this?



Is this Integrated GRC?

What are the main problems GRC intends to solve?

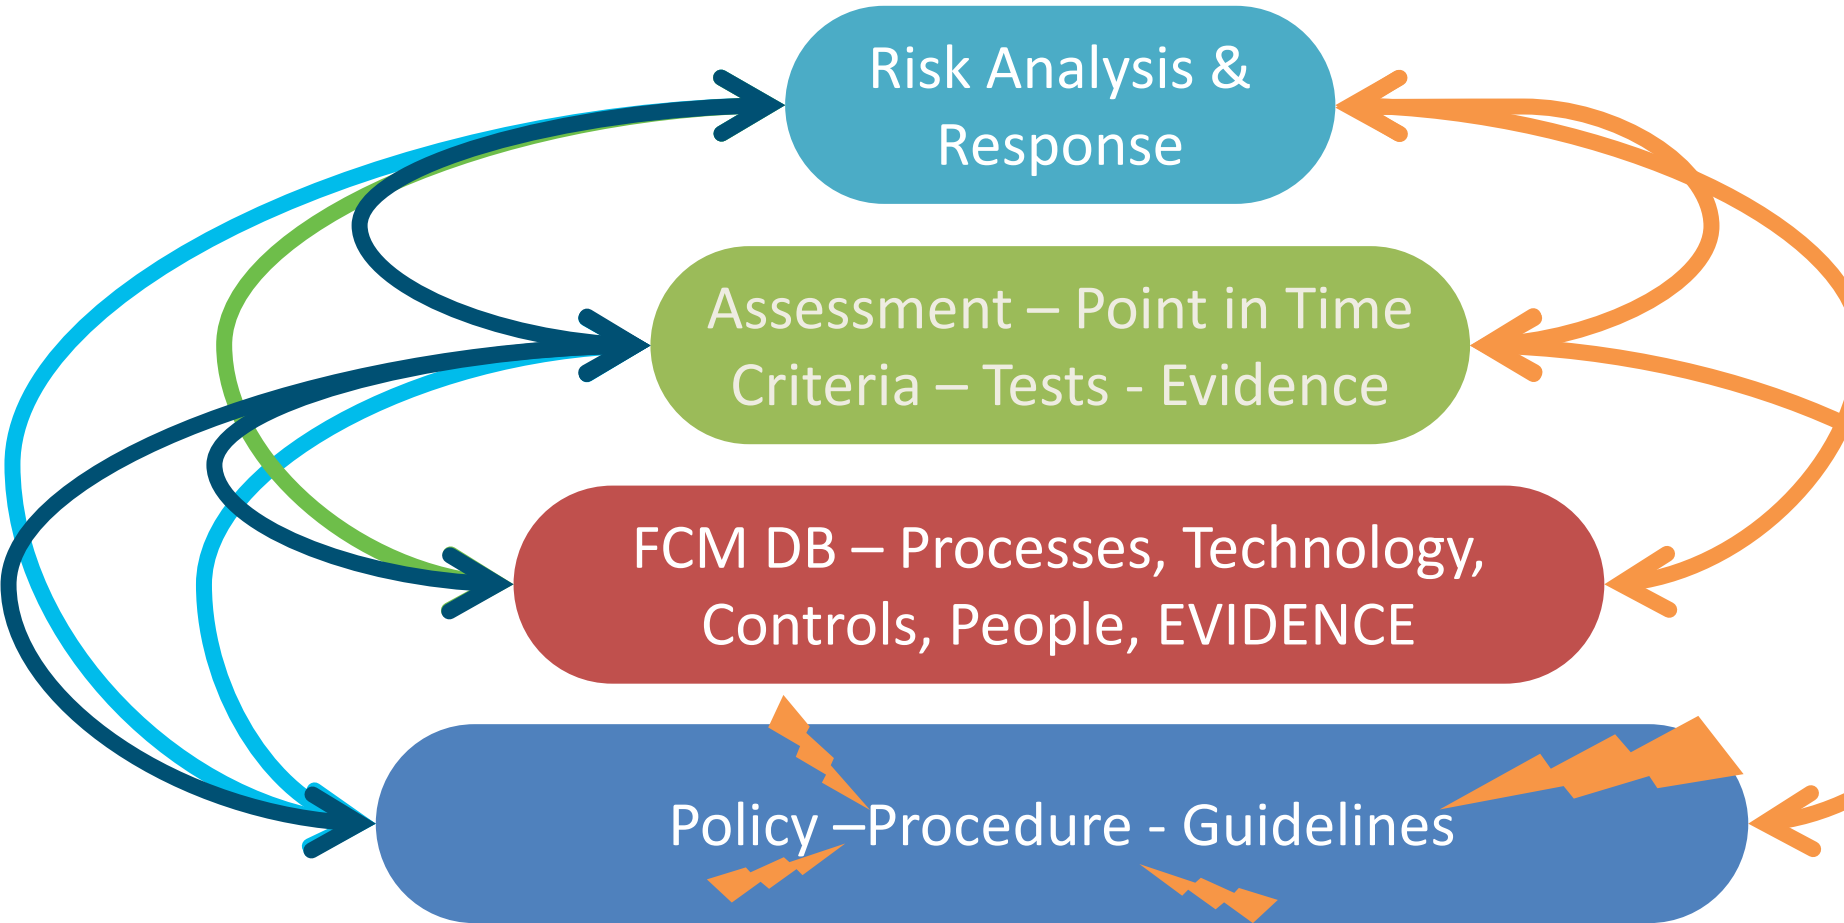
Is GRC part of Cybersecurity, or is Cybersecurity a part of GRC?




We will never have enough time and resources

Here's what's putting cracks in our foundation:

- Too many policies
- Out-of-date or incorrect procedures
- Redundant efforts and no feedback loop
- Wrong level of information per policy, standard, guideline v. SOP affecting classification and use on the document
- Lack of evidence mapping



Could AI propose additions and modifications to existing policies, procedures, and guidelines while ensuring stakeholder adoption and creating context-based risk and exception tracking?



Zero Trust Implementation Goals



Decrease

Decrease infrastructure and perimeter complexity.



Ensure

Ensure policy enforcement works across hybrid physical and cloud environments.



Ensure

Ensure that policy enforcement is applied universally to all assets (i.e., devices) and locations independently.



Ensure

Ensure the solution is compliant with organizational policy and government/industry standards.

Zero Trust Security Objectives – Moving from Attack Surface to Protect Surface

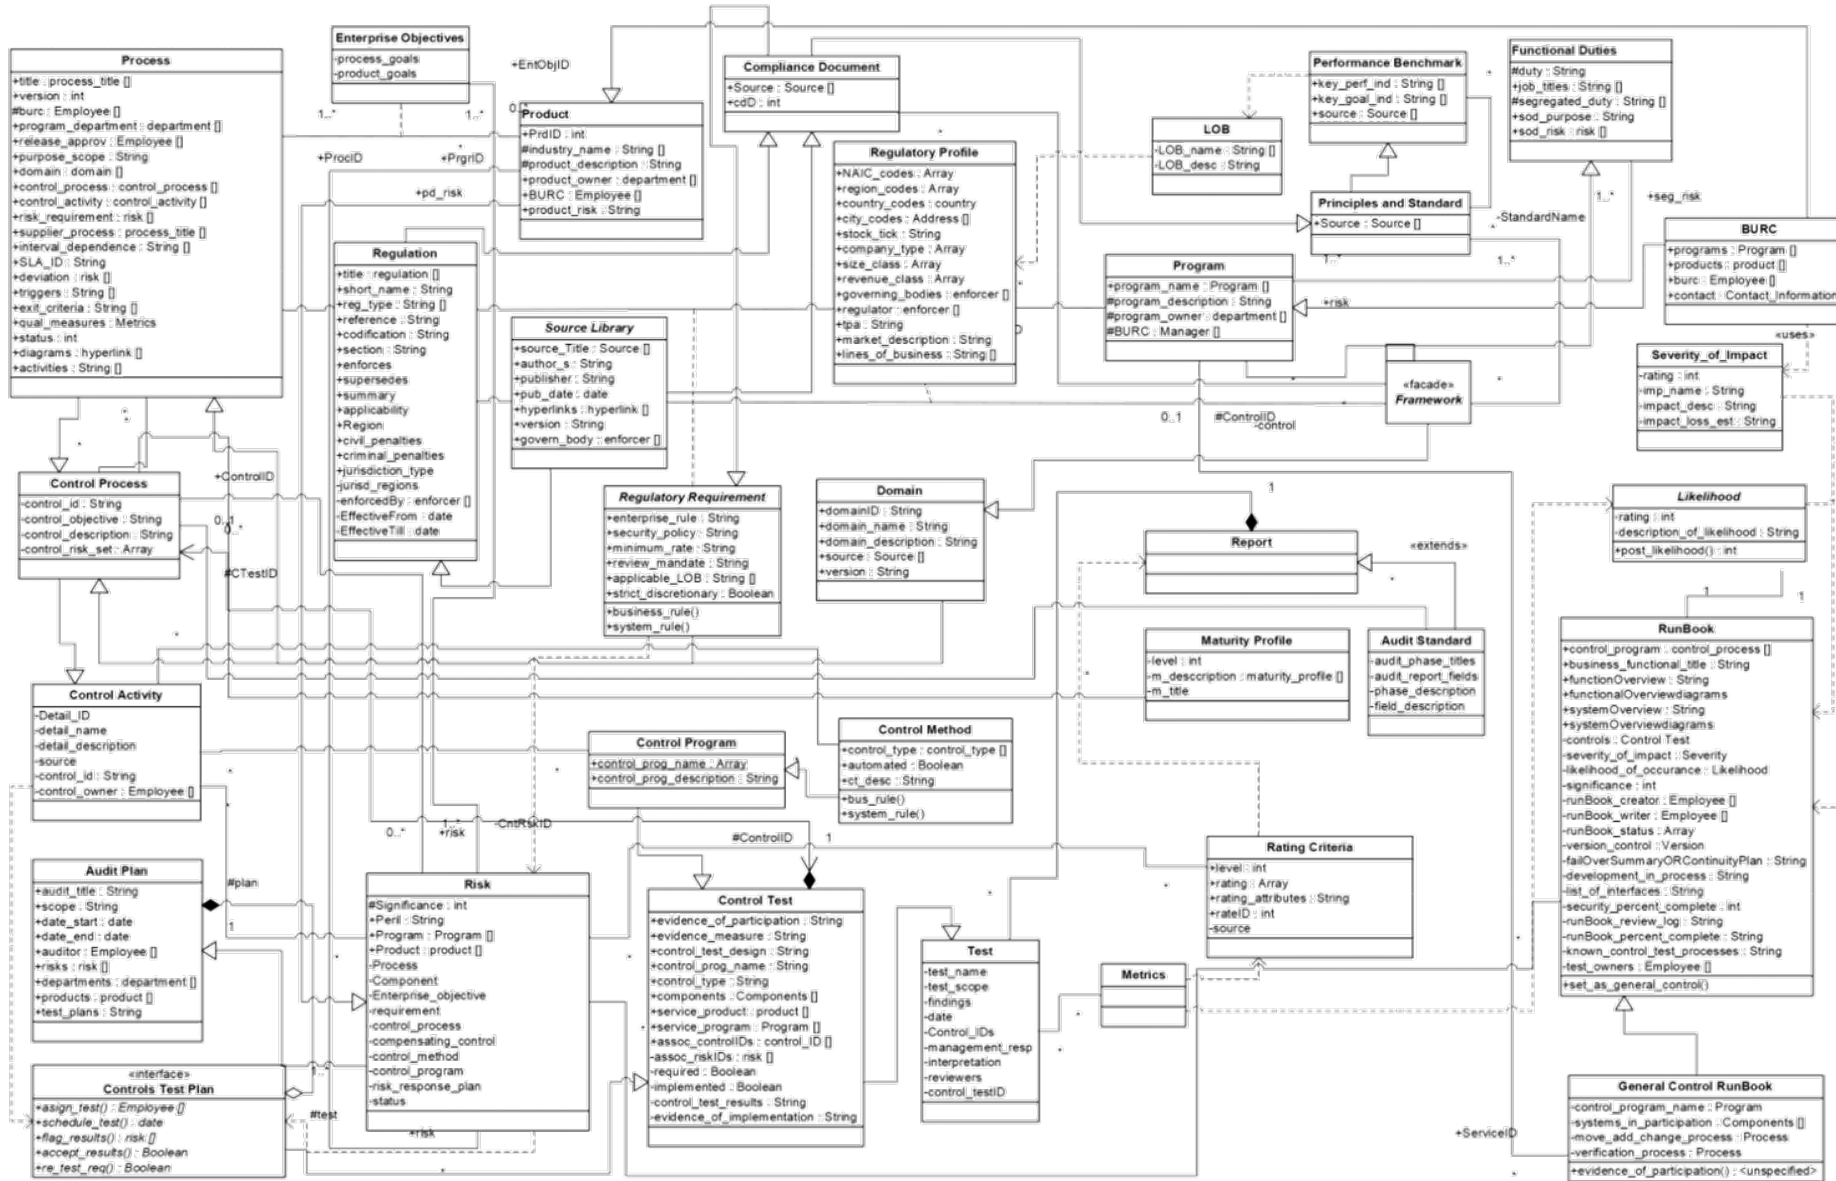
Component	Zero Trust Capability
Authentication System	The explicit ability to verify the identity of a process or device.
Authorization System	The ability to grant or deny device access to data, assets, applications, or services by a policy enforcement point.
Privileged Access Management	The ability to secure, control, and manage privileged access to critical assets and applications.
Software-Defined Perimeter or Networking	The ability to provision and control network components using code.
Device Compliance	The ability to validate that policy engine decisions are enforced on device endpoints.
Network Segmentation	Network traffic can be segmented at either the macro or micro level depending upon the organization's application and data resources.
Data Loss Prevention Systems	The ability to inspect network traffic and application-based traffic and apply rules to allow or deny it. This capability works in tandem with the policy engine.
Security Information and Event Management Systems	A security information and event management system provides network and application traffic visibility and supports the notion of continuous monitoring and reporting on the success and failure of the enforcement of policy engine rules.

GRC maturity is lagging the CISA zero trust maturity model

https://www.cisa.gov/sites/default/files/2023-04/zero_trust_maturity_model_v2_508.pdf

Identity - Risk Assessments - Optimal	Agency determines identity risk in real-time based on continuous analysis and dynamic rules to deliver ongoing protection.
Data - Governance Capability - Advanced	Agency begins integration of data lifecycle policy enforcement across the enterprise, enabling more unified definitions for data governance policies.
Cross Cutting - Automation and Orchestration - Advanced	Agency automates orchestration and response activities enterprise-wide, leveraging contextual information from multiple sources to inform decisions.
Cross Cutting - Automation and Orchestration - Optimal	Agency orchestration and response activities dynamically respond to enterprise-wide changing requirements and environmental changes.
Cross Cutting - Governance - Advanced	Agency implements tiered, tailored policies enterprise-wide and leverages automation to support enforcement. Access policy decisions incorporate contextual information from multiple sources.
Cross Cutting - Governance - Optimal	Agency implements and fully automates enterprise-wide policies that enable tailored local controls with continuous enforcement and dynamic updates .

Consider password policies in the tiered requirements. If we mature, are we allowed to remove out-of-date controls? Why would better controls have to be managed as compensating?




A good GRC Architect knows a lot.

A great GRC Architect knows what problems the GRC should be tasked to solve and then solves them.

Hey Alexa, am I ready to let in the auditor?

Tell me if we have everything, we need to meet with an external auditor on this *<date>*, for this *<audit>*, for this *<product>*, and for a set of *<processes>* and their assessment specific *<controls>*.

Answer by showing the set of *<tests>* and their *<artifacts>* as evidence to support our statement that a control, for a particular product, and for a period of time, is READY to pass audit.



The first step of the Revolution is admitting that we already lost the war.



Towards a Standard for Identifying and Managing Bias in Artificial Intelligence - NIST Special Publication 1270

- NIST SP 1270 intends to surface the salient issues in the challenging area of AI bias and to provide a first step on the roadmap for developing detailed socio-technical guidance for identifying and managing AI bias. Specifically, this special publication:
- describes the stakes and challenge of bias in artificial intelligence and provides examples of how and why it can chip away at public trust;
- identifies three categories of bias in AI — systemic, statistical, and human — and describes how and where they contribute to harm;
- describes three broad challenges for mitigating bias — datasets, testing and evaluation, and human factors — and introduces preliminary guidance for addressing them.

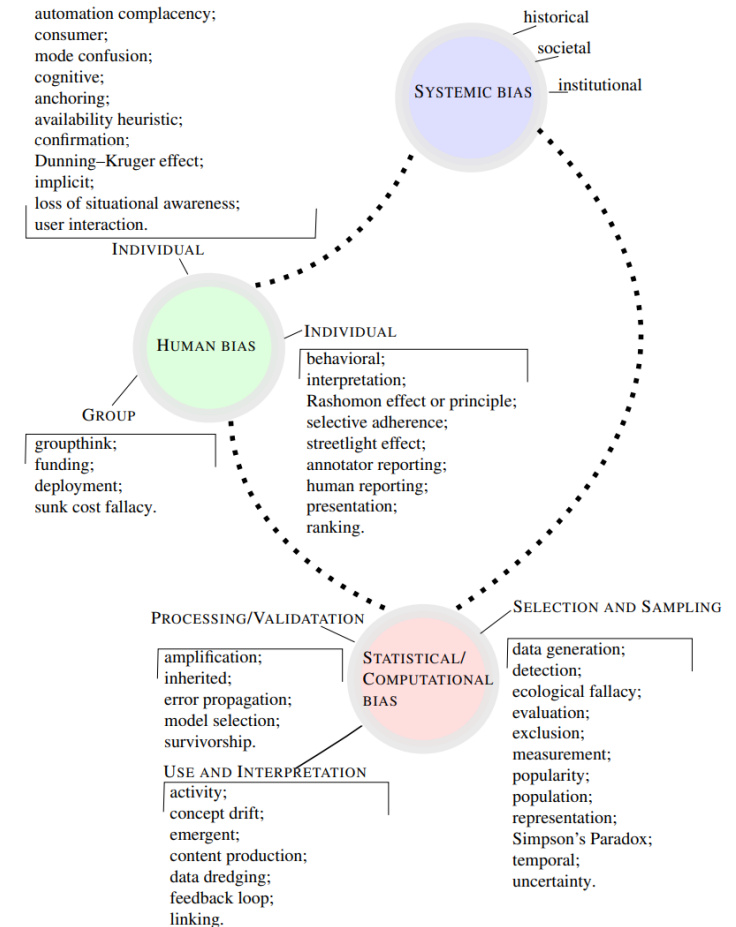


Fig. 2. Categories of AI Bias. The leaf node terms in each subcategory in the picture are hyperlinked to the GLOSSARY. Clicking them will bring up the definition in the Glossary. To return, click on the current page number (8) printed right after the glossary definition.




Unintended Consequences

- Why didn't product development and marketing consider that a hostile Ex would use AI to locate a past partner living in a safe house via Facebook, Google, or X, with or without the programmatic use of the facial recognition API?
- Why didn't medical AI include that bodies vary tremendously by gender, age, ethnicity, race, and skin tone, building that into the Bayesian logic behind prescriptive techniques used to administer medication or to diagnose symptoms? (More from Jodi)
- Why didn't IBM, Microsoft, or Amazon predict that policing would circumvent due process if they could get what they wanted by using a product's API? (Sales suspended in 2020)
- Why has the United States taken so much longer to propose rules for the use of private information in AI systems than the EU has done via GDPR and the recent Artificial Intelligence Act? Will the release of the AI Bill of Rights be enough?
- Consider the following recommendations made by The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation. Brundage, M. et al. (2018, February). Future of Humanity Institute. Retrieved from <https://arxiv.org/ftp/arxiv/papers/1802/1802.07228.pdf>
- Policymakers should collaborate closely with technical researchers to investigate, prevent, and mitigate potential malicious uses of AI.
- Researchers and engineers in artificial intelligence should take the dual-use nature of their work seriously, allowing misuse-related considerations to influence research priorities and norms and proactively reaching out to relevant actors when harmful applications are foreseeable.
- Best practices should be identified in research areas with more mature methods for addressing dual-use concerns, such as computer security, and imported where applicable to the case of AI.
- Actively seek to expand the range of stakeholders and domain experts involved in discussions of these challenges

How biases contribute to harms

From a Bayesian inference perspective, this can be seen as updating the prior of the model to help avoid issues that may arise from using stale prior probability distribution. Organizations are recommended to employ appropriate governance procedures to adequately capture this cross-organizational need and ensure no negative impacts from using AI technology.

Transparency, datasets, and test, evaluation, validation, and verification (TEVV)

	Systemic Biases	Statistical and Computational Biases	Human Biases
 <p>Datasets <i>Who is counted, and who is not counted?</i></p>	<ul style="list-style-type: none"> Issues with latent variables Underrepresentation of marginalized groups 	<ul style="list-style-type: none"> Sampling and selection bias Using proxy variables because they are easier to measure Automation bias 	<ul style="list-style-type: none"> Observational bias (streetlight effect) Availability bias (anchoring) McNamara fallacy
 <p>Processes and Human Factors <i>What is important?</i></p>	<ul style="list-style-type: none"> Automation of inequalities Underrepresentation in determining utility function Processes that favor the majority/minority Cultural bias in the objective function (best for individuals vs best for the group) 	<ul style="list-style-type: none"> Likert scale (categorical to ordinal to cardinal) Nonlinear vs linear Ecological fallacy Minimizing the L1 vs. L2 norm General difficulty in quantifying contextual phenomena 	<ul style="list-style-type: none"> Groupthink leads to narrow choices Rashomon effect leads to subjective advocacy Difficulty in quantifying objectives may lead to McNamara fallacy
 <p>TEVV <i>How do we know what is right?</i></p>	<ul style="list-style-type: none"> Reinforcement of inequalities (groups are impacted more with higher use of AI) Predictive policing more negatively impacted Widespread adoption of ridesharing/self-driving cars/etc. may change policies that impact population based on use 	<ul style="list-style-type: none"> Lack of adequate cross-validation Survivorship bias Difficulty with fairness 	<ul style="list-style-type: none"> Confirmation bias Automation bias

Source: [Towards a Standard for Identifying and Managing Bias in Artificial Intelligence \(nist.gov\)](https://www.nist.gov)