ACG Discussion - 6/15/2022

Distribution/Dissemination Control:
POC: djshepard@sei.cmu.edu

Dave Shepard

Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

DM22-0555

Story arc for this talk:

1. Risk
2. RMF
3. ATO
4. cATO
5. Misunderstandings and Limitations

First, let’s establish a common understanding.

What is risk?

What is risk?

- The possibility of losing something that matters to you.

What is risk?

- The possibility of losing something that matters to you.

What is risk management?

What is risk?

- The possibility of losing something that matters to you.

What is risk management?

- The act of protecting something that matters to you.

What is risk?

- The possibility of losing something that matters to you.

What is risk management?

- The act of protecting something that matters to you.

How does one “manage” the possibility of loss?

What is risk?

- The possibility of losing something that matters to you.

What is risk management?

- The act of protecting something that matters to you.

How does one “manage” the possibility of loss?

- By Purchasing Insurance
  - Protects against a financial loss
  - Not every loss is financial in nature

What is risk?

- The possibility of losing something that matters to you.

What is risk management?

- The act of protecting something that matters to you.

How does one “manage” the possibility of loss?

- By Purchasing Insurance
  - Protects against a financial loss
  - Not every loss is financial in nature
- By Keeping Secrets
  - Protects against loss of important information
  - Secrets are fickle things

What is risk?

- The possibility of losing something that matters to you.

What is risk management?

- The act of protecting something that matters to you.

How does one “manage” the possibility of loss?

- By Purchasing Insurance
  - Protects against a financial loss
  - Not every loss is financial in nature
- By Keeping Secrets
  - Protects against loss of important information
  - Secrets are fickle things
- By Defending Life
  - Protects your friends, family, and country
  - A loss of life is the most precious kind of loss

Risk management frameworks formalize the process of protecting that which matters to you.

Many sectors have compliance and legal requirements and processes and frameworks for managing risk. We will focus on the government sector use-cases, which diverge from the concerns of many other sectors.

GDPR: General Data Protection Regulation
FISMA :Federal Information Security Management
SOX : Sarbanes–Oxley
HIPAA : Health Insurance Portability and Accountability
PCI DSS: Payment Card Industry Data Security Standard
NIST :National Institute of Standards and Technology,
And many more..

Understanding Risk is Hard

In cybersecurity alone, one taxonomy identifies ~70 categories of risk.

AccessControlRisk
AccountabilityRisk
AdvancedPersistentThreatRisk
AssetRisk
AttackVectorRisk
AuthenticationRisk
AuthenticityRisk
AvailabilityRisk
CampaignRisk
CompromisedAssetRisk
CompromiseRisk
ConfidentialityRisk
CourseofActionRisk
CyberAdvisoryRisk
CyberAlertRisk
CyberAttackRisk
CyberEventRisk
CyberIncidentResponsePlanRisk
CyberIncidentRisk
CyberResilienceRisk
CyberRisk
CyberRunRisk
CyberSecurityRisk
CyberThreatRisk
DataBreachesListRisk
DataBreachRisk
DefenceinDepthRisk
DenialofServiceRisk
DetectFunctionRisk
DistributedDenialofServiceRisk
EnvironmentalHazardsRisk
ExploitRisk
HackingRisk
IdentifyFunctionRisk
IdentityAndAccessManagementRisk
IncidentResponseTeamRisk
IndicatorsofCompromiseRisk
InformationSharingRisk
InformationSystemRisk
InformationTheftRisk
IntegrityRisk
ITRisk
MalwareRisk
MissionRisk
MisuseRisk
MultiFactorAuthenticationRisk
NonRepudiationRisk
OperationalRisk
PatchManagementRisk
PenetrationTestingRisk
PhysicalActionRisk
ProtectFunctionRisk
RecoverFunctionRisk
ReliabilityRisk
RespondFunctionRisk
Risk
SituationalAwarenessRisk
SocialEngineeringRisk
SystemicCyberRiskRisk
TacticsTechniquesAndProceduresRisk
TechnicalErrorRisk
ThreatActionRisk
ThreatActorRisk
ThreatAssessmentRisk
ThreatIntelligenceRisk
ThreatLedPenetrationTestingRisk
ThreatVectorRisk
TrafficLightProtocolRisk
VerificationRisk
VulnerabilityAssessmentRisk
VulnerabilityRisk

I must give credit to the original creators of this particular risk taxonomy, which I had readily available from use on a different project:

https://www.openriskmanual.org/wiki/Risk_Taxonomy

Understanding Risk is Hard

In cybersecurity alone, one taxonomy identifies ~70 categories of risk.

AccessControlRisk
AccountabilityRisk
AdvancedPersistentThreatRisk
AssetRisk
AttackVectorRisk
AuthenticationRisk
AuthenticityRisk
AvailabilityRisk
CampaignRisk
CompromisedAssetRisk
CompromiseRisk
ConfidentialityRisk
CourseofActionRisk
CyberAdvisoryRisk
CyberAlertRisk
CyberAttackRisk
CyberEventRisk
CyberIncidentResponsePlanRisk
CyberIncidentRisk
CyberResilienceRisk
CyberRisk
CyberRunRisk
CyberSecurityRisk
CyberThreatRisk
DataBreachesListRisk
DataBreachRisk
DefenceinDepthRisk
DenialofServiceRisk
DetectFunctionRisk
DistributedDenialofServiceRisk
EnvironmentalHazardsRisk
ExploitRisk
HackingRisk
IdentifyFunctionRisk
IdentityAndAccessManagementRisk
IncidentResponseTeamRisk
IndicatorsofCompromiseRisk
InformationSharingRisk
InformationSystemRisk
InformationTheftRisk
IntegrityRisk
ITRisk
MalwareRisk
MissionRisk
MisuseRisk
MultiFactorAuthenticationRisk
NonRepudiationRisk
OperationalRisk
PatchManagementRisk
PenetrationTestingRisk
PhysicalActionRisk
ProtectFunctionRisk
RecoverFunctionRisk
ReliabilityRisk
RespondFunctionRisk
Risk
SituationalAwarenessRisk
SocialEngineeringRisk
SystemicCyberRiskRisk
TacticsTechniquesAndProceduresRisk
TechnicalErrorRisk
ThreatActionRisk
ThreatActorRisk
ThreatAssessmentRisk
ThreatIntelligenceRisk
ThreatLedPenetrationTestingRisk
ThreatVectorRisk
TrafficLightProtocolRisk
VerificationRisk
VulnerabilityAssessmentRisk
VulnerabilityRisk

I must give credit to the original creators of this particular risk taxonomy, which I had readily available from use on a different project:

https://www.openriskmanual.org/wiki/Risk_Taxonomy

Being able to categorize the risks your systems face is really one of the first steps you'll need to do, in order to prioritize your risk mitigation needs.

The NIST Risk Management Framework Helps to Prioritize your Risks.

What is the NIST Risk Management Framework (RMF)?

- Information security framework for the entire federal government that replaces legacy Certification and Accreditation (C&A) Processes applied to information systems
- RMF is a key component of an organization's information security program used in the overall management of organizational risk
- NIST Special Publication 800-37, "Guide for Applying the Risk Management Framework to Federal Information Systems", transforms the traditional Certification and Accreditation(C&A) process into the six-step Risk Management Framework (RMF).
- The Risk Management Framework (RMF) provides a disciplined and structured process that integrates information security and risk management activities into the system development lifecycle

1. ***Categorize*** the information system and the *information processed, stored, and transmitted* by that system based on an impact analysis
2. ***Select*** an initial set of baseline security controls for the information system based on the security categorization; tailoring and supplementing the security control baseline as needed based on an *organizational assessment of risk and local conditions*.
3. ***Implement*** the security controls and describe how *the controls are employed within the information system and its environment of operation*.
4. ***Assess*** the security controls using appropriate assessment procedures to determine the extent to which the *controls are implemented correctly*, operating as intended, and producing the desired outcome with respect to meeting the security requirements for the system
5. ***Authorize*** information system operation based on a *determination* of *the risk to organizational operations and assets, individuals, other organizations*, and the Nation resulting from the operation of the information system and the decision that this risk is acceptable.
6. ***Monitor*** the security controls in the information system on an *ongoing basis including assessing control effectiveness, documenting changes to the system or its environment of operation*, conducting security impact analyses of the associated changes, and reporting the security state of the system to designated organizational officials.

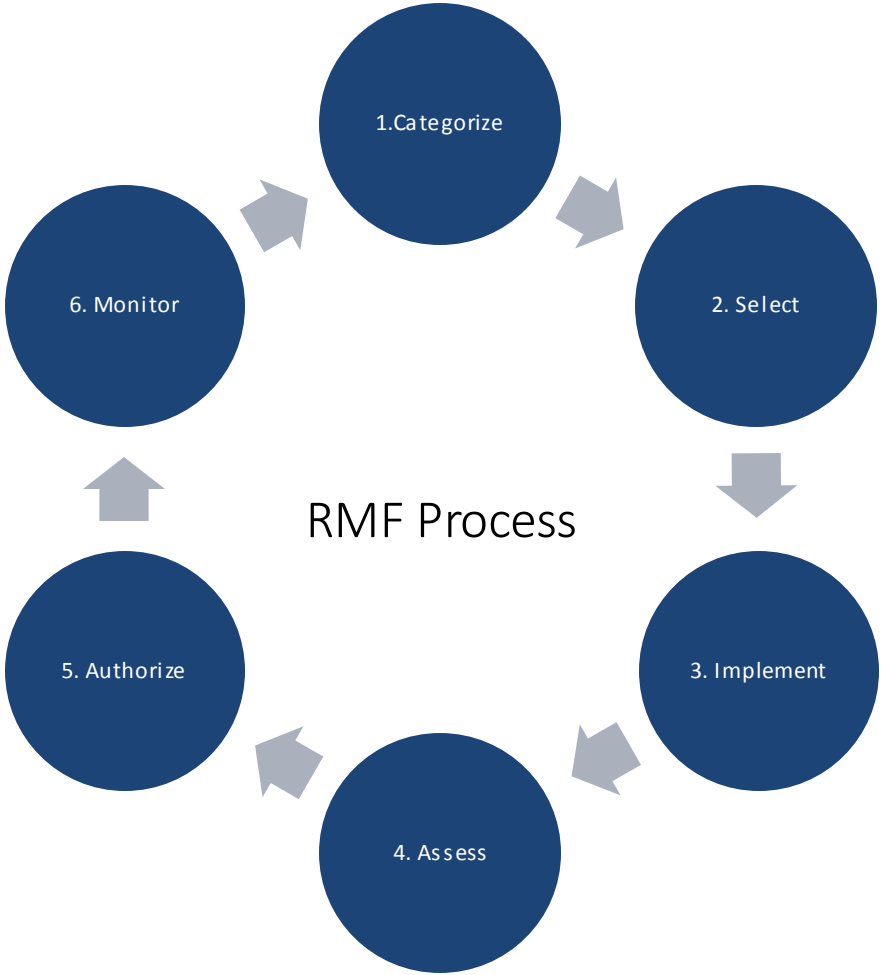

RMF is the starting point for most implementers of ATO

ATO is:

The official management decision given by a senior organizational official to authorize operation of an information system and to explicitly accept the risk to organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, other organizations, and the Nation based on the implementation of an agreed-upon set of security controls.

--NIST

Developing Authority to Operate(ATO)

Typical security controls scenario: **924** controls selected for implementation and all are in Excel format;

Control Number	Family	Control Title	Control Text	Confidentiality	Integrity	Availability	Supplemental Guidance
		TOOLS					
SI-4 (10)	SI	INFORMATION SYSTEM MONITORING | VISIBILITY OF ENCRYPTED COMMUNICATIONS	The organization makes provisions so that [Assignment: organization-defined encrypted communications traffic] is visible to [Assignment: organization-defined information system monitoring tools].	High Moderate	High Moderate	High Moderate	Supplemental Guidance: Organizations balance the potentially conflicting needs for encrypting communications traffic and for having insight into such traffic from a monitoring perspective. For some organizations, the need to ensure the confidentiality of communications traffic is paramount; for others, mission-assurance is of greater concern. Organizations determine whether the visibility requirement applies to internal encrypted traffic, encrypted traffic intended for external destinations, or a subset of the traffic types.
SI-4 (11)	SI	INFORMATION SYSTEM MONITORING | ANALYZE COMMUNICATIONS TRAFFIC ANOMALIES	The organization analyzes outbound communications traffic at the external boundary of the information system and selected [Assignment: organization-defined interior points within the system (e.g., subnetworks, subsystems)] to discover anomalies.	High Moderate Low	High Moderate Low	High Moderate Low	Supplemental Guidance: Anomalies within organizational information systems include, for example, large file transfers, long-time persistent connections, unusual protocols and ports in use, and attempted communications with suspected malicious external addresses.
SI-4 (12)	SI	INFORMATION SYSTEM MONITORING | AUTOMATED ALERTS	The organization employs automated mechanisms to alert security personnel of the following inappropriate or unusual activities with security implications: [Assignment: organization-defined activities that trigger alerts].	High Moderate Low	High Moderate Low	High Moderate Low	Supplemental Guidance: This control enhancement focuses on the security alerts generated by organizations and transmitted using automated means. In contrast to the alerts generated by information systems in SI-4 (5), which tend to focus on information sources internal to the systems (e.g., audit records), the sources of information for this enhancement can include other entities as well (e.g., suspicious activity reports, reports on potential insider threats). Related controls: AC-18, IA-3.
SI-4 (13)	SI	INFORMATION SYSTEM MONITORING | ANALYZE TRAFFIC / EVENT PATTERNS	The organization: (a) Analyzes communications traffic/event patterns for the information system; (b) Develops profiles representing common traffic patterns and/or events; and (c) Uses the traffic/event profiles in tuning system-monitoring devices to reduce the number of false positives and the number of false negatives.				

Obviously, developing and maintaining ATO is time-consuming, expensive. It also typically causes delays to schedules, cost overruns, and resentment for process frameworks.

"The muddled, bureaucratic process to obtain an ATO and launch an IT system inside government is widely maligned — but beyond that, it has become a pervasive threat to system security. The longer government takes to launch a new-and-improved system, the longer an old and potentially insecure system remains in operation."

— Mary Lazzeri, ATO ASAP: Let's finally fix the security compliance problem

Without attribution…

“The time and cost of implementing a continuous ATO was the same as implementing ATO just once. It’s a no brainer.”

RMF is the starting point for most implementers of cATO

NIST has no formal definition for cATO, even though it shows up in their document library.

A definition of cATO, or sometimes C-ATO:

"…a continuous authority to operate (cATO) is the continuous authorization of software components such as containers by building security into the entire development lifecycle using DevOps technologies and processes ensuring that the application and its components meet security levels equal to or greater than what an ATO requires."

cATO authorizes the platform, the development process, and the team that produces the product under a continuous monitoring process that maintains the residual risk within the risk tolerance of the Authorizing Official (AO)

DevSecOps provides a useful framework of automation and process with the people necessary to develop secure software that can facilitate the development of a continuous ATO.

DevSecOps is:

"DevSecOps is a set of principles and practices that provide faster delivery of secure software capabilities by improving the collaboration and communication between software development teams, IT operations, and security staff within an organization, as well as with acquirers, suppliers, and other stakeholders in the life of a software system."

--SEI

NIST 800-137 (September 2011) defines “Information Security Continuous Monitoring”:

Information security continuous monitoring (ISCM) is defined as maintaining ongoing awareness of information security, vulnerabilities, and threats to support organizational risk management decisions.

This is a related publication to RMF (800-37), covering the continuous monitoring requirements that are necessary for cATO.

Steps required by ISCM:
Define an ISCM strategy;
Establish an ISCM program;
Implement an ISCM program;
Analyze data and Report findings;
Respond to findings; and
Review and Update the ISCM strategy and program.

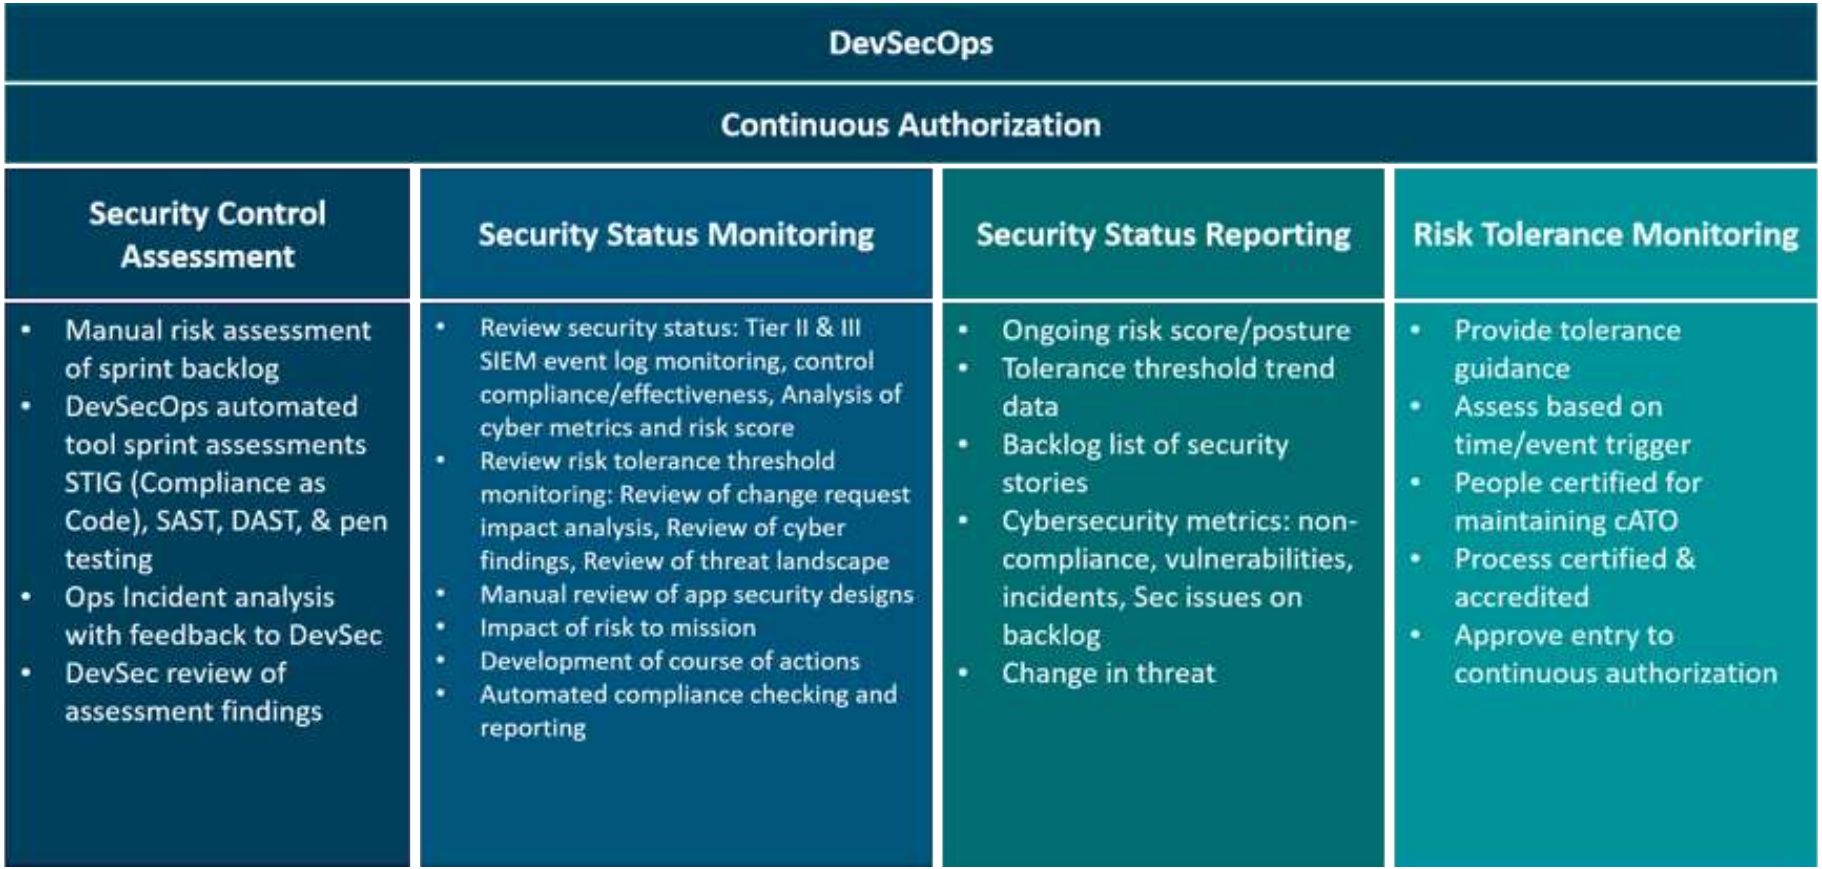

Achieving Continuous ATO

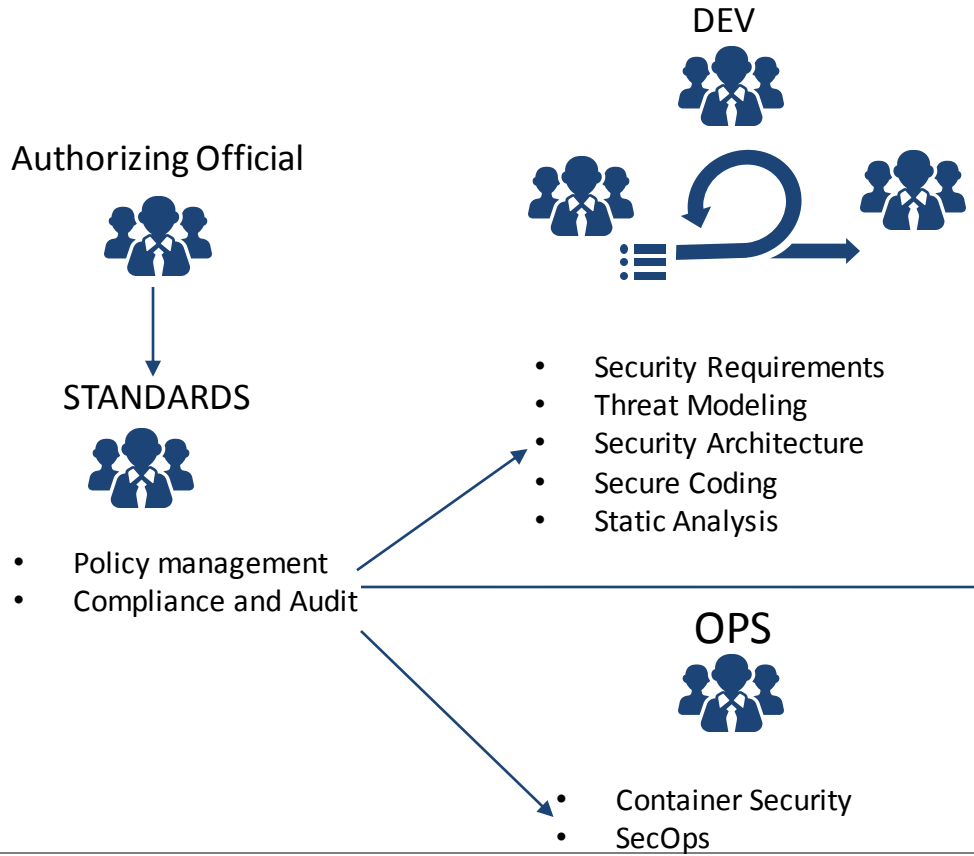

In Agile and DevSecOps fashion, all of this works better when all of these stakeholders work together to achieve the goal.

Governance (Measuring Effectiveness)

	SPEED TO MARKET METRICS
SECURITY AUTOMATION	• Time to delivery • # high severity bugs found in production
THREAT MODELING	• # of unmitigated high severity attacks
COMPLIANCE AND REGULATORY	• # of unmitigated high severity clauses
SECURE DEVELOPMENT	• Code scan results
SECURITY OPERATIONS	• # of security incidents in production

The state of practice for cATO is evolving. There are limitations to what can be achieved through automation and there are misunderstandings about what is currently possible.

Misunderstandings

If I use security-hardened containers, I'm inheriting their RMF controls for my cATO.

- Yes, but your AO gets to decide which controls are required for your cATO.
- Systems are not the sum of isolated software deployments. Systems have complex interactions between components and emergent behaviors that must be understood and accounted for in the cATO decision.
- Security-hardened containers must often be modified, in order to make them useful for your environment. Modification of the deployment package means controls must be re-validated.

Fixing The Misunderstandings About Hardened Containers

Hardened containers, as distributed by a repository such as the Iron Bank, could come with a documentation package containing:

- Which controls are being implemented and how
- Which features have been modified and/or removed from the vendor's baseline
- Additional steps that may/will be necessary when deployed
- Deployment assumptions, including how environment variables, configuration files, ephemeral storage containers, etc., will be passed/captured/required by the container
- A manifest containing the last set of scan results for the container, including details about the scanner(s)

I'm certain there are others I haven't thought about.

Limitations

Effectiveness of your controls will only be as good as your ability to assess them.

- This has always been true of ATO. It's still true of cATO, only now you rely on automated checks to validate your control assumptions are not violated. Control checks are written as automated policy, by developers, who make mistakes in their logic. Can you validate their logic? (Maybe. Hopefully.)
- Mistakes in automated policy checks can allow checks to be violated/circumvented, intentionally and accidentally. Your audit logs may or may not be helpful in detecting policy violations.
- A mistake doesn't necessarily mean the check itself is wrong. It can also mean the check doesn't validate assumptions about constraints under which the check is valid. Again, poor coding can undermine the effectiveness of a constraint validation check.
- Even if you do everything right in developing a control check, there are limitations to what the controls will be able to prevent/detect, due to system complexity.