

Continuous Control Monitoring Implementation Series

#1 - The Case for CCM

White Paper

Presented by David Vohradsky
CEO and Founder

Introduction

Organizations are facing unprecedented challenges in managing cyber risks and compliance obligations. Cyber threats are accelerating in frequency and sophistication, while regulators and industry standards are imposing stricter requirements on how organizations must govern information security. A stark example: studies show it takes organizations more than 277 days to identify and contain a data breach. During such a prolonged gap, attackers roam freely and security controls remain ineffective. This reality has spurred a fundamental shift in approach – from periodic, point-in-time security assessments to Continuous Controls Monitoring (CCM).

This white paper explores why continuous monitoring of security controls has become essential for effective cyber risk management and compliance. We will discuss the drivers behind CCM, define what it entails, examine alignment with key compliance frameworks (ISO 27001, AESCSF, APRA CPS 234, ASD Essential 8, ISM, and RFFR), and outline a practical path to implementing CCM across various security domains. This is Part 1 of a 10-part series providing detailed guidance for practitioners on building a robust CCM capability.

The Evolving Compliance Landscape

In recent years, regulators have raised the bar on what constitutes adequate security governance. Traditional compliance regimes relied on infrequent audits or attestation. Today, however, laws and standards increasingly mandate ongoing assurance that security controls are effective:

- **Security of Critical Infrastructure (SOCI) Act** – Australia's SOCI Act and its associated rules require critical infrastructure entities to implement a Risk Management Program, including processes to "minimise risks... as far as reasonably practicable" and to detect and respond to threats in real-time. This implies continuous hazard monitoring and rapid control of emerging risks, rather than annual reviews. Non-compliance can lead to government intervention, fines, and reputational damage
- **ISO/IEC 27001:2022** – The leading international standard on information security was updated in 2022 to emphasize continuous improvement and monitoring. Annex A control 8.16 ("Monitoring Activities") explicitly requires that networks, systems and applications should be monitored for anomalous behavior and that appropriate actions be taken to investigate potential incidents. Additionally, Clause 9.1 of ISO 27001 obliges organizations to evaluate the performance of security controls regularly.
- **APRA CPS 234** – The Australian Prudential Regulation Authority's CPS 234 (2022) Standard for financial institutions underscores continuous oversight. It requires entities to "test the effectiveness of information security controls through a systematic testing program" at a frequency commensurate with risk. Meeting these expectations practically necessitates automated, continuous monitoring mechanisms to rapidly identify control failures or breaches.
- **Australian Energy Sector Cyber Security Framework (AESCSF)** – Developed by AEMO for the energy industry, AESCSF is a comprehensive maturity model that includes 11 domains. Continuous monitoring is woven throughout. For instance, in the Identity & Access Management domain, one objective is to "continuously monitor access to critical systems and data to detect unauthorized access attempts and anomalies".
- **Australian Government Information Security Manual (ISM) and Essential Eight** – The ISM (by ASD) and the Essential Eight both highlight the need for ongoing control of systems. Essential Eight, in particular, defines maturity levels for eight critical controls (application whitelisting, patching, macros settings, user privileges, etc.). Achieving maturity Level 3 (the highest) typically requires automation and continuous enforcement.
- **Right Fit For Risk (RFFR)** – RFFR is a cybersecurity accreditation framework used by Australian government agencies (like the Department of Employment and Workplace Relations) to assess service providers. It combines ISO 27001 and ISM controls. A core requirement in RFFR is that an organization have documents and processes for continuous monitoring and improvement. In fact, to achieve RFFR certification, an organization must produce "a Continuous Monitoring Plan" as well as an "Improvement Plan for monitoring security objectives and targets"

Across these examples, a common theme emerges: security is not a point-in-time concern. Frameworks are converging on the expectation that organizations know – at any given moment – the status of their critical controls and security posture. Non-compliance is increasingly tied to failure to detect issues in a timely manner (e.g. undetected vulnerabilities, misconfigurations, or access violations that persist). Therefore, aligning with these frameworks compels the adoption of continuous monitoring capabilities.

Defining Continuous Controls Monitoring

Given these drivers, what exactly do we mean by Continuous Controls Monitoring in a practical sense? At its core, CCM is the practice of constantly evaluating whether security controls are present and functioning correctly, using automated data feeds and building on metrics. This goes beyond traditional continuous security monitoring (which often refers to watching for attacks or intrusions) – CCM is about monitoring the controls themselves. For example, rather than just monitoring for malware infections, CCM would monitor that antivirus protections are installed, running, and up-to-date on all required systems at all times.

NIST defines “Information Security Continuous Monitoring” (ISCM) as maintaining ongoing awareness of information security, vulnerabilities, and threats to support organizational risk management decisions.”. Gartner Research, on the other hand, emphasizes the outcome of CCM: reducing business losses and audit costs by automating control auditing. They describe CCM as a set of technologies that enable continuous auditing of risk controls, as opposed to the periodic audits of the past. Both perspectives capture important aspects: NIST highlights continuous awareness (focused on security posture), while Gartner highlights continuous audit and business benefits.

Key characteristics of CCM include:

- **Automated data collection:** CCM relies on integrations with systems (cloud platforms, endpoints, IAM systems, etc.) to gather evidence of control status. For example, pulling a list of all user accounts and access privileges daily to detect any policy violations, or using an API to query cloud resource configurations against security benchmarks continuously.
- **Predefined metrics or indicators:** Organizations implementing CCM define what “control healthy” vs “control failure” looks like in measurable terms. These are often called Key Risk Indicators (KRIs) or control metrics. For instance, “% of devices with the latest security patches applied” could be a metric. If that percentage falls below a threshold, it triggers an alert as a control weakness.
- **Continuous or high-frequency evaluation:** Controls are checked on a regular schedule (hourly, daily, etc.) or in real-time where possible. The frequency is determined by the nature of the control and risk – e.g., monitoring user access changes might be near-real-time, whereas scanning configuration compliance might be daily or weekly. In my ISACA Journal Article¹ I suggest running tests “at a point in time close to when transactions or processes occur,” effectively meaning as fast as the process being monitored.
- **Alerting and workflow for exceptions:** When a control is found to be out of compliance (an “assertion fails”), the CCM system generates an alert or report. There must be a defined process to triage these alerts, investigate, and remediate the issue. This often integrates with ticketing systems or dashboards for the security team. Part of CCM's process value is ensuring that identified control gaps are communicated and addressed promptly – before they escalate into incidents or audit findings.
- **Auditable logging:** Continuous monitoring systems typically log all checks and results. This creates an evidence trail that auditors and executives can review, providing assurance that controls were continuously monitored, and any issues were managed. In essence, CCM can feed into continuous compliance reporting, making it easier to demonstrate adherence to regulatory frameworks on an ongoing basis.

In summary, CCM operationalizes the dictum “trust, but verify” at all times for your security controls. It overlaps with concepts like IT compliance automation, Security Posture Management, and DevSecOps monitoring – all of which revolve around using automation to ensure your organization's security measures are never left unchecked for long.

Benefits of Continuous Controls Monitoring

Adopting CCM yields significant benefits for an organization's security and risk posture:

- **Timely Risk Detection:** The most obvious benefit is catching security control failures or gaps quickly. For example, CCM might detect that a critical server has been misconfigured or left unpatched within a day of the issue arising, allowing the team to fix it before attackers exploit it. This is crucial given that studies have shown a majority of breaches involve known vulnerabilities or misconfigurations that went unaddressed. By continuously monitoring such things as patch status and system configurations, organizations can drastically reduce their window of exposure.

1. Vohradsky, D. (2015), “A Practical Approach to Continuous Control Monitoring”, ISACA Journal, 2015(2).

- **Improved Compliance and Audit Readiness:** Continuous evidence collection means an organization is always “audit-ready.” Instead of panicking before an audit to gather proof that, say, backups were taken or user access was reviewed, the CCM system has been reviewing it all along. This can reduce the effort of compliance reporting. Gartner pointed out that continuous control monitoring greatly reduces reliance on labor-intensive manual audits and testing. Many firms initially pursue CCM to ease the pain of meeting compliance requirements like CPS 234, Essential 8, ISO 27001, and find that it not only makes audits easier but also improves the actual security outcomes.
- **Efficiency and Cost Savings:** Automation is cheaper and more reliable at scale than manual control assurance. Consider an enterprise with thousands of cloud resources – manually checking each for security configuration (or even audit grade sampling) would be untenable on a frequent basis. A CCM capability (such as a GRC integrated with a Cloud Security Posture Management system) can automatically assess all resources against policies continuously. This yields productivity gains: staff can be redeployed to high-value analysis and remediation rather than repetitive data gathering.
- **Increased Visibility and Assurance:** Continuous monitoring provides management and the board with up-to-date insight into security. For example, a board risk dashboard might show key metrics like “99% of laptops have up-to-date anti-malware” or “All high-risk cloud buckets are properly encrypted as of this morning.” This fosters confidence that security is under control. It also supports better decision-making – trends can be spotted (e.g. a slowly declining patch compliance percentage might indicate a resourcing issue in IT that can be corrected before it becomes a crisis). In governance terms, CCM enables continuous assurance. Audit and risk committees are increasingly interested in this level of visibility, moving away from relying solely on point-in-time audit opinions.
- **Supports Rapid Response:** When continuous monitoring detects an issue, it often can be tied into automated response or at least fast-tracked manual response. Some CCM implementations integrate with Security Orchestration, Automation and Response (SOAR) tools to not just alert on a control failure but also initiate a fix. For instance, if CCM finds an unauthorized port open on a server, it could automatically trigger a script to close it or quarantine that server from the network until reviewed. This tight integration from detection to response helps close the loop quickly, embodying the DevOps ideal of automation in IT operations. Even when full automation isn’t possible, having immediate alerts with contextual data means responders can act faster. This can dramatically reduce the dwell time of threats and the window in which an attacker could exploit a lapse in controls.

Challenges and Considerations

Implementing CCM is not without challenges. It’s important to acknowledge these to set realistic expectations and plan appropriately:

- **Data Integration:** Aggregating data from many disparate systems (cloud platforms, on-premises systems, SaaS applications, identity providers, etc.) can be complex. Organizations often need to deploy or build a centralized monitoring platform or use a combination of tools. Integration and normalization of this data (so that it can be analyzed uniformly) is a significant upfront effort.
- **Defining Metrics and Thresholds:** Deciding what to monitor continuously and what the criteria for “failure” are requires expertise. If thresholds are too lax, you might miss important issues; too strict and you could flood the team with false positives. For example, is an account with 45 days of inactivity a concern or 90 days? These definitions should align with risk appetite and policy. It may take tuning over time to get right.
- **Alert Fatigue and Workflows:** Continuous monitoring can quickly produce a large volume of alerts if not tuned well. Teams must be prepared with processes (and sufficient staffing or automation) to handle the alerts. Otherwise, there’s a risk of important warnings being ignored. Prioritization is key – e.g., failing controls that pose a high risk should be highlighted over minor compliance deviations. One must incorporate CCM into the operational workflow of security, IT, and risk teams, including clear ownership of who investigates what.
- **Tooling and Costs:** While CCM can save money in the long run, it often requires investment in tools up front – whether a Governance, Risk, and Compliance (GRC) platform with CCM capabilities, cloud security monitoring tools, SIEM enhancements, or custom scripts and development. There is also a learning curve and ongoing maintenance to ensure these tools remain effective as systems change. Organizations should budget for these and possibly consider managed services if in-house capacity is limited.
- **Cultural Change:** Moving to continuous monitoring may require a cultural shift. Teams that are used to a project-based audit mentality might need to adapt to a steady operational tempo of “monitor, fix, improve, repeat.” There can be initial resistance, especially if CCM surfaces a lot of issues that were previously hidden – it can feel like things are getting worse when in fact you’re just now shining a light on them. Executive support is important to reinforce that discovering and fixing control gaps continuously is positive, not punitive.

Despite these challenges, the trend is clear: the benefits and necessity of CCM far outweigh the difficulties. As Gartner observed, CCM is fast becoming “essential tech for large, heavily regulated organizations” in the near future

Conclusion

Continuous Controls Monitoring represents a paradigm shift in cybersecurity and compliance management. Instead of point-in-time verification, it establishes an ongoing, almost real-time, validation of an organization's security defenses. For industries under strict regulations or facing advanced threats, CCM is not just a "nice to have" – it's rapidly becoming a must. Frameworks like ISO 27001, AECSEF, APRA CPS 234, Essential 8, and RFFR are all signaling the same direction: keep your controls in check continuously, or risk falling out of compliance and exposing yourself to breaches.

This white paper (Part 1) has set the stage by explaining why CCM is critical and what it entails at a high level.

In Part 2 "Designing a Continuous Controls Monitoring Program – Frameworks, Scope, and Strategy" we delve into how to design and kick-start a CCM program – from selecting controls and aligning with frameworks, to choosing tools and metrics.

Subsequent parts provide practical, step-by-step guidance on implementing continuous monitoring in specific domains like asset management, access control, incident response, and more. Our focus throughout is on actionable advice, lessons learned from the field, and how to overcome the common challenges identified in Part 1. By following this series, a cyber risk practitioner will gain the knowledge to move beyond theory and start continuously monitoring and managing controls in their environment, ultimately achieving stronger security outcomes and peace of mind with regulators.



Solution Spotlight: MyRISK + Oracle Integration Cloud for CCM Enablement

To move from theory to practical execution, organizations need platforms that can operationalize continuous monitoring at scale. Our MyRISK HyperGRC platform is designed specifically for this purpose, providing an integrated solution that ingests control data, automates responses, and generates continuous evidence of control effectiveness.

Oracle Integration Cloud (OIC) forms the backbone of our integration layer. With hundreds of prebuilt adapters for common enterprise systems (like cloud services, HR platforms, directories, databases, GRC tools), OIC allows HyperGRC to ingest control data from virtually any source — including identity systems, cloud APIs, configuration databases, and logging services. This wide integration reach ensures that our platform has the inputs needed to support continuous testing across diverse environments.

Once ingested, this data can feed into AI agents deployed in Oracle Cloud Infrastructure (OCI) Data Science. These can be pre-built or organization-specific models and agents to detect anomalies, predict control failures, or classify risks. For example, HyperGRC AI agents could analyze change management logs to identify patterns that often precede failed controls, or continuously score risk based on access, configuration drift, or missing control evidence.

Control exceptions, once identified, can be automatically handled through our built-in workflow automation. For instance, if a critical control fails (such as missing audit logs, or a cloud misconfiguration), MyRISK could trigger an incident workflow, create a task in ServiceNow (via OIC), notify the responsible team, and track resolution status until closure — providing a full, auditable lifecycle for control failures.

Finally, MyRISK uses a publish-subscribe (pub-sub) messaging architecture to distribute events asynchronously across its components and external systems. When a control failure is detected, it could simultaneously notify compliance teams, update dashboards, log evidence, and initiate workflows — all triggered from a single control event message. This architecture is essential for scalability and modularity, allowing large organizations to scale their CCM efforts without bottlenecks.

Together, MyRISK + Oracle provide a future-ready, continuously operating backbone for risk and control assurance — turning CCM from a concept into a living, actionable program.

Smarter Cyber Risk.
Automated. Empowered. Secured.