Google Cloud

Cyber-Physical Resilience and the Cloud

Putting the White House PCAST report into practice.

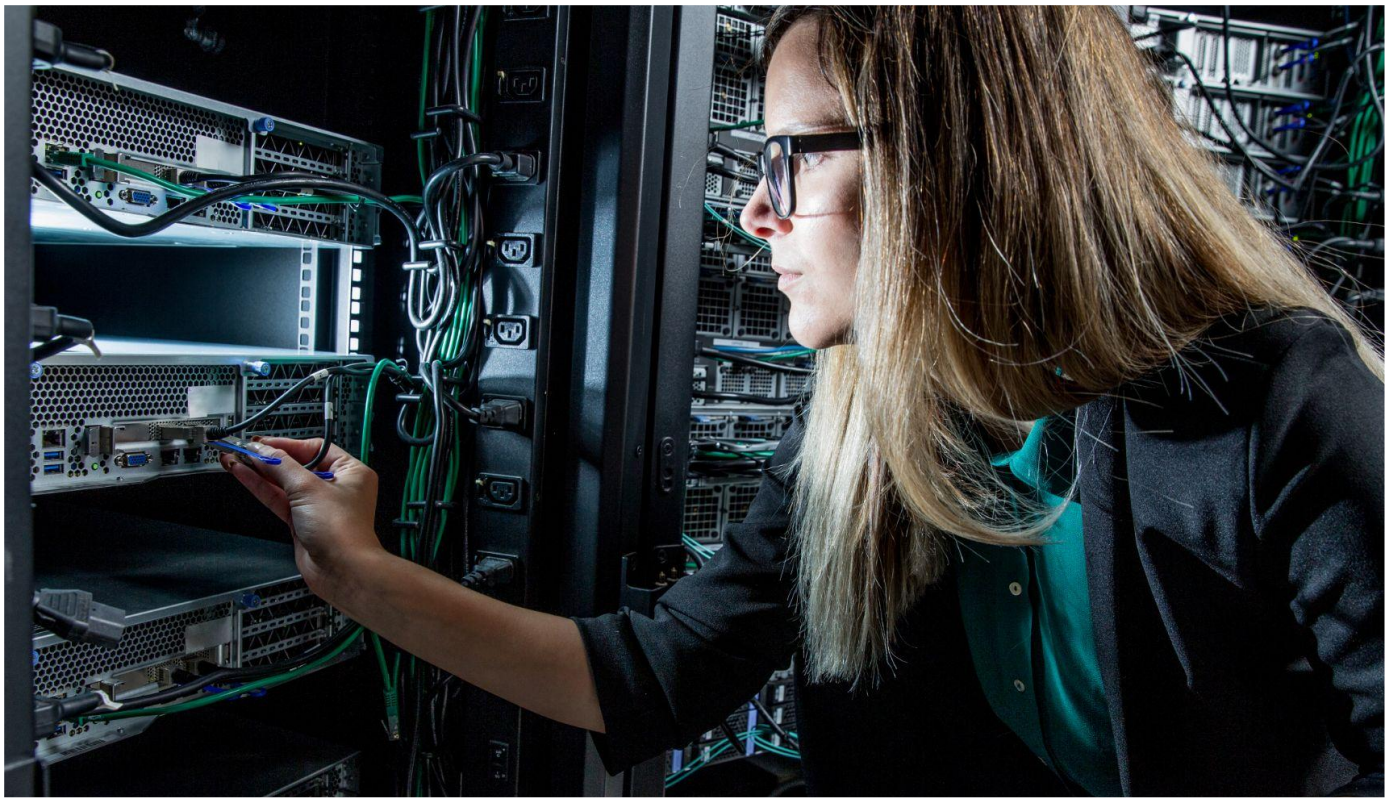

November 2024

Table of Contents

Introduction

Organizations often rely on lagging indicators to measure the success of their security and resilience efforts. These indicators, such as the number of successful breaches or the financial impact of a cyberattack, provide valuable insights into past performance. However, they are limited prognosticators: only rarely do they help prevent future incidents or guide proactive improvements.

Leading indicators, on the other hand, focus on proactive measures and predictive factors that can help organizations identify potential risks and vulnerabilities before they are exploited. By shifting the focus to leading indicators, organizations can adopt a more proactive approach to security and resilience, enabling them to anticipate and mitigate threats more effectively.

The recent President's Council of Advisors on Science and Technology (PCAST) report, "Strategy for Cyber-Physical Resilience: Fortifying Our Critical Infrastructure for a Digital World," holds significant importance for our society beyond its implications for IT and business. However, at 51 pages, it is unlikely that many people will read and study it in its entirety.

While lengthy, the PCAST report crucially emphasizes the importance of considering broader technology metrics rather than solely relying on traditional security-specific measures. This is because building technology better, with a focus on classic software security and software supply chain transparency, inherently leads to improved security and resilience outcomes.

For example, robust software development practices, such as those outlined in the NIST's Secure Software Development Framework (SSDF), can significantly enhance the overall security posture of an organization. By focusing on building secure and resilient technology from the ground up, organizations can achieve a wide range of commercial outcomes, including improved security and resilience.

The central theme of the report emphasizes the need to enhance efforts in safeguarding critical infrastructure. Cyberattacks against the systems that govern critical infrastructure can have far-reaching consequences, including potential disruptions to power supply, water availability, and transportation networks.

We're reviewing the 10 "leading indicators" presented in the PCAST report to better help organizations develop their cyber-physical resilience. In this paper, we evaluate how the 10 indicators from the PCAST report can be used to improve organizational resilience.

01

Hard-Restart Recovery Time: Rebuilding from the Ground Up

The report states that a **hard-restart recovery time** is the duration it takes *"to reconstitute/rebuild a system from scratch (as distinct from backup and recovery time objectives). This metric is intended to assess an organization's ability to remove circular dependencies during a restart, to assure backups can survive fully destructive shutdowns or attacks, and that software and data can be restored to service. The application of this concept will vary across different sectors."*

What does hard-restart recovery time mean?

Hard-restart recovery time is the amount of time needed to completely rebuild a system from the ground up (likely excluding hardware). This includes installing the operating system, required software, restoring data backups, and reconfiguring settings.

Circular dependencies are situations where components or services rely on each other in such a way that it's difficult to determine the order to bring them back online (backup required identity provider, but it needs to be restored first from a backup).

Destructive shutdowns or attacks refer to events that completely wipe out a system, potentially including backups, requiring a full restoration.

Why should hard-restart recovery time matter to your org?

- **Worst-case scenario planning**: Establishing Hard-Restart Recovery Time helps organizations prepare for catastrophic failures where typical backup and restore measures might be insufficient.
- **Cyber resilience**: Planning for hard restarts improves an organization's ability to recover from destructive cyberattacks like ransomware that can render backups unusable.
- **Minimizing downtime**: By addressing circular dependencies and streamlining the rebuilding process, organizations reduce the time critical systems are offline in the event of a full-scale disruption.

How can you make hard-restart recovery time better?

Maintain documentation for system configuration, software installations, and data restoration procedures. Make sure that the detailed documentation does not rely on the system in question for access and that it is kept current.

Adopt collaboration tools and methods for the teams involved in the restart to actually do the work. Ensure that they do not rely on the system being restarted. Defining concrete recovery steps to clearly identify the order of operations is key.

Map dependencies by analyzing and untangling circular dependencies in systems to determine the optimal order for bringing components back online during a rebuild.

Store offline backups or air-gapped backups that are isolated from the network and therefore protected from destructive attacks.

Test hard-restart scenarios regularly to verify that procedures work as expected. Identify potential bottlenecks, and ensure that testing spans the range of disruptive “live” testing and a range of ongoing tests. Rely on regular tools, avoid test-only tooling that tends to atrophy over time.

These are your first steps

01

Begin by inventorying all hardware, software, and configuration settings that would be required for a hard restart of critical systems.

02

Start small with testing, which means running the iterative process of gradually increasing the complexity and scope of resilience tests, and incorporating the lessons learned.

How you can benefit

- **Enhanced disaster recovery preparedness:** Reduce your vulnerability to catastrophic events and ensure you can rebuild systems even in the worst scenarios.
- **Time-bound response:** Streamline the hard restart process for reliable recovery in case of severe emergencies. Test the response time against requirements and SLOs to fit in the desired time boundary.
- **Improved security posture:** Reduce your reliance on potentially compromised backup methods, which can strengthen your cyber resilience.
- **Adjacent benefits**: Improved understanding of the system for day-to-day operations and troubleshooting. Early identification of points of failure single or fragile. Improved employee preparedness and education. Increased employee engagement and morale, understanding the importance of the system to national well-being. Signal to the public that action is being taken to protect their interests and well-being. Early identification of components that are close to obsolescence permitting controlled improvement.

How should you measure success?

Reduction in recovery time

Track your hard-restart recovery time during regular drills over time to measure improvement.

Successful recovery

Evaluate whether a hard restart can fully restore critical systems to operational state.

Dependency management

Assess how successfully your organization has addressed circular dependencies during the hard restart planning and testing.

How Google Cloud can help

01

Cloud backups for resilience: Modern cloud offers immutable object storage, ensuring that backups can not be altered or deleted, even by ransomware. This protects critical data and provides a reliable foundation for recovery.

02

Automated recovery playbooks: Cloud provider automation tools enable the creation and testing of recovery playbooks that orchestrate the restoration of systems and data from immutable backups. These playbooks streamline the recovery process and reduce the risk of errors.

03

Faster recovery times: Cloud providers can typically restore service much faster than on-premises systems because they have a larger pool of resources that they can use to quickly spin up new instances of your applications and data.

04

Reduced downtime: Cloud providers have a number of features that can help to reduce downtime, such as automatic failover and load balancing. If one instance of your application fails, another instance can automatically take over, without any interruption to service.

05

Regular hard-restart drills: Cloud providers emphasize the importance of conducting periodic drills that simulate complete system outages. These drills help identify bottlenecks, test recovery procedures, and ensure that systems can be rebuilt from scratch using cloud-based resources.

02

Cyber-Physical Modularity

The report states that ***cyber-physical modularity*** *is, "a system-wide measure, computed as the mean of the impact of single points of failure. The measure considers the additional failures and impacts that come as cascades via dependencies on each initial single-point failure. This can be captured by a summary measure of the impacts of each primary point of failure. Alternatively, the measure can be computed as the mean operational capability of the service, even if degraded, summed over all single points of failure. For example, given that each single point of failure may cause a temporary outage or reduction of quality of service, what is the consequent median tested recovery time for all single points of failure to be repaired/recovered? For how many single points of failure is the defined minimum viable operating delivery objective sustained?"*

What does cyber-physical modularity mean?

Cyber-physical modularity: The design of systems where interconnected physical and digital components are organized into well-defined modules with limited dependencies on each other. This approach aims to contain the impact of failures.

Single points of failure: Components or connections whose malfunction could cascade into a system-wide issue, impacting a broader set of functions.

Mean operational capability: How well the system can operate, potentially in degraded mode, after single points of failure. This is measured against a minimum viable delivery objective.

Why does it matter to your org?

- **Reduced cascading failures:** High modularity limits the spread of issues arising from single points of failure, making the system more resilient.
- **Faster recovery:** Isolated problems are easier to identify and fix, improving median recovery time (such as how quickly issues can be repaired).
- **Maintain essential services:** Modularity helps ensure the sustainability of critical functions defined in the minimum viable operational delivery objective, even in the event of disruptions.

How can you make cyber-physical modularity happen?

- **Component identification:** Analyze your cyber-physical systems to map out individual components and their interdependencies.
- **Dependency analysis:** Assess how the failure of each component could potentially impact other parts of the system.
- **Modular design:** Where possible, redesign system architecture to group components into self-contained modules with minimal external dependencies.
- **Failure isolation mechanisms:** Implement safeguards like circuit breakers or redundant pathways designed to prevent cascading effects when single points of failure occur.

This is your first step

01

Begin creating a visual representation (such as a diagram) of your cyber-physical systems, especially their dependencies that clearly identify components and connections.

How would you benefit?

- **Improved system reliability:** Modular systems are less likely to experience widespread failure due to problems in a single area.
- **Operational efficiency:** Maintain critical functions and overall system stability even when individual components fail.
- **Easier maintenance:** Troubleshooting, repairs, and upgrades can be done in a more targeted manner, reducing overall downtime.

How can you measure success?

Frequency of cascading failures

Track incidents before and after modularity improvements. A decrease would indicate success.

Recovery time

Measure time from initial failure to the restoration of the minimum viable operational delivery objective.

Number of modules affected

For each single point of failure, count the total number of modules impacted within the system.

How Google Cloud can help

01

Improved resource isolation: Cloud computing can improve resource isolation by providing virtualized environments that are separated from each other. This can help to prevent failures in one part of the system from cascading to other parts of the system.

02

Increased redundancy: Cloud computing can increase redundancy by providing multiple copies of data and applications. This can help to ensure that the system can continue to operate even if one or more components fail.

03

Enhanced security: Cloud computing can enhance security by providing a variety of security features, such as network controls, IAM, and data encryption.

04

Design systems with redundancy and diversity to avoid single points of failure: Cloud providers such as Google understand the importance of redundancy. We also advocate for diversity in system design to avoid correlated failures.

05

Implement graceful degradation to maintain partial functionality during disruptions: Site Reliability Engineers (SRE) design the systems for graceful degradation, stating that services should produce reasonable but suboptimal results if overloaded.

03

Internet Denial/Communications Resilience

The report states that ***Internet Denial/Communications Resilience*** *focuses on items such as "Internet denial testing: consider loss of Internet connectivity as a special notable point of failure," "Explicitly test the impacts, nature of degraded service, and disruptions vs. operational continuity in the face of internet disconnection." "Some services are so critical that they should be operable safely, even in some degraded state, in the absence of network connectivity," and "Consider backup communication channels to diverse modes of communication in the event of Internet failure."*

What does it mean?

Internet denial refers to the loss or significant disruption of internet connectivity. This can occur due to technical failures, natural disasters, and cyberattacks.

Communications resilience ensures business continuity by maintaining the ability to communicate effectively even if the primary internet connection is severed.

Why does it matter to your org?

- Since nearly all organizations now rely on the internet for communication, data access, and critical operations, an unexpected loss of connectivity could halt work, damage relationships with clients, and have a significant negative financial impact.

- Communications resilience planning can help prepare your organization to continue functioning through internet disruptions and reduce potential losses.

How can you make communications resilience happen?

Conduct risk assessments: Thoroughly analyze your operations to identify which processes and services absolutely rely on internet connectivity. Prioritize these for resilience planning.

Develop failover plans: Designate alternative communication channels that can operate independently of your primary internet connection (such as cellular networks, satellite connections, and dedicated communication lines).

Test regularly: Perform drills to simulate loss of connectivity. Ensure that backup communications systems work as intended and that staff are trained to use them effectively.

Decentralize communication infrastructure: If possible, consider distributing critical communication resources across multiple physical locations to increase redundancy.

This is your first step

01

Start by mapping your organization's critical dependencies on internet connectivity. This will lay the foundation for your resilience planning.

How would you benefit?

- **Reduced downtime:** You minimize interruption to critical operations and services during internet outages.
- **Enhanced reputation:** Demonstrate reliability and preparedness to clients and stakeholders.
- **Avoid financial loss:** Prevent lost revenue and productivity due to communications disruptions.
- **Increased adaptability:** Create a more flexible organization, ready to respond to unexpected challenges.

How can you measure success?

Metrics

Track the duration and impact of any internet outages, both before and after resilience measures are implemented.

Reduced reliance

Measure the proportion of critical services or functions that continue operating successfully during a simulated internet denial scenario.

Staff feedback

Survey employees on their confidence level and speed of response during internet disruptions.

How Google Cloud can help

01

Cloud provides redundancy: While cloud will not help if the entire internet goes down, cloud increases resilience, provided you have redundant ways to connect to it, from multiple internet links to private connections.

02

Geographically-dispersed data centers: Cloud computing can help to improve DDoS resilience by providing data centers around the world. This means that if one data center is attacked or goes offline, traffic can be routed to another data center, ensuring that critical services remain available.

03

Redundant network connections: Cloud computing can also help to improve internet Denial/Communications Resilience by providing redundant network connections. If one network connection is attacked or goes offline, traffic can be routed over another network connection, ensuring that critical services remain available. Design connectivity with redundancy and diversity to avoid single points of connection failure.

04

Cloud redundancy: Cloud increases resilience provided you have redundant ways to connect to it, from multiple internet links to private connections. Private connectivity services add to the resiliency here.

05

Superior DDoS defense: Cloud providers has extensive experience in mitigating DDoS attacks. Their nfrastructure has massive bandwidth capacity and advanced mitigation techniques to absorb and deflect large DDoS attacks.

04

Manual Ops

The report states that ***fail-over to manual operations*** *is for "physically actuated systems typically controlled by cyber operational technology, what is the degree of local manual control that can sustain a minimum viable operational delivery objective when automation is lost? How frequently is manual control practiced to sustain organization muscle-memory of its use? Additionally, to what extent is there a broader primary or back-up analog "control plane" to the system or components? While digitization is inevitable and valuable, maintaining some degree of analog control may be necessary and warranted for certain highly critical systems."*

What does it mean?

Fail-over to manual operations is the process of switching from automated systems (often controlled by cyber-operational technology, or OT) to direct human control when the primary systems fail or become compromised.

Physically-actuated systems are systems involving mechanical components or processes that are physically moved or manipulated.

Minimum viable operational delivery objective defines the core functions and level of service that must be maintained even in the event of disruptions.

Analog control plane refers to alternative systems, which are often non-digital and less complex, that provide a basic operational mode.

Why does it matter to your org?

- **Safety:** In critical systems, manual operation may be necessary to prevent accidents or damage when automation fails.
- **Cybersecurity:** Maintaining a degree of manual control reduces an organization's reliance on potentially vulnerable automated systems, lowering the risk of cyberattacks disrupting critical operations.
- **Business continuity:** Ensures that critical operations can continue, even in degraded mode, during failures of the automated systems.

How can you make a fail-over to manual operations happen?

Identify critical systems: Evaluate your physical systems to determine which have a high potential impact on safety and operations if automated control is lost.

Develop manual procedures: Create clear, well-documented procedures for manual operations of critical systems.

Train personnel: Regularly train individuals responsible for backup procedures to maintain proficiency and "muscle memory" for manual operations.

Maintain analog backup: Where possible, consider maintaining analog control planes as a last resort fallback in the event that both primary and digital backup systems fail.

This is your first step

01

Conduct an audit of your critical physically actuated systems, map their reliance on automation, and identify where manual takeover potential exists.

How would you benefit?

- **Enhanced resilience:** Reduce the risk of severe disruptions or shutdowns caused by automation failures.
- **Improved safety:** Mitigate risks related to critical system malfunctions.
- **Regulatory compliance:** Demonstrate ability to operate safely even in adverse conditions, potentially a requirement in certain industries.

How can you measure success?

Time to manual override

Track how quickly personnel can initiate manual control during unannounced drills.

Success of manual operations

During drills, monitor the extent to which minimum viable operational objectives are met when running manually.

Frequency of automation failures

Log the frequency and impact of automation failures, establishing a baseline to see if resilience measures yield improvement.

How Google Cloud can help

01

SRE processes for resilience: Cloud Service Reliability Engineering (SRE) practices emphasize building resilient systems that can withstand failures. SRE principles including error budgeting, blameless postmortems, and continuous improvement can be applied to manual operations to enhance their reliability and effectiveness.

02

Break glass process: Cloud platforms often support "break glass" processes, which provide a secure way to bypass normal access controls in emergency situations when controls planes don't work. This can be crucial when automated systems fail and immediate manual intervention is required. Specific to access, in case of system failure, critical systems may offer a break glass process for direct access.

03

Private connectivity as a control path: Establishing a private connection to a cloud provider, independent of your primary infrastructure, can serve as a fully independent control path. This ensures that even if your primary systems and networks are compromised, you can still maintain control through the cloud.

04

Customer Reliability Engineering (CRE) teams: CRE teams partner with customers to engineer their services for reliability.

05

Control Pressure Index (CPI)

The report states *that **the control pressure index** is "the extent to which defense-in-depth is applied by measuring how much of a critical security or resilience objective is carried by a single control (that if failed would put the whole system at risk)."*

What does it mean?

Control Pressure Index (CPI): A metric that quantifies the degree to which a system relies on individual security or resilience controls. A high CPI indicates over-reliance, which means a single control failure could have serious consequences.

Defense-in-depth: A cybersecurity strategy emphasizing multiple layers of overlapping controls to protect against a range of threats. This redundancy means that even if one control fails, others can still maintain protection.

Critical Security or Resilience Objective: A key goal or requirement related to maintaining a secure and resilient system (such as data confidentiality, system availability, and disaster recovery.)

Why does it matter to your org?

- **Risk assessment:** The CPI helps pinpoint potential vulnerabilities where the failure of a single control could have a disproportionate effect.
- **Resource allocation:** Understanding CPI guides the prioritization of resources to reinforce critical controls with high pressure indices.
- **Building a robust defense:** Low CPI scores signify a greater degree of defense-in-depth, resulting in a more resilient and secure system overall.

How can you make a control pressure index happen?

Map controls to objectives: List critical security/resilience objectives. Identify all controls that contribute to protecting each objective.

Assess control criticality: Analyze each control to determine how much of an objective's defense it bears, and assign it a rating – high, medium, low.

Calculate CPI: For each objective, count the number of high-criticality controls. If only one "high" control exists, it signals a high CPI and higher risk. Ideally, you want multiple high or medium controls.

Mitigate high CPI: Where possible, introduce additional controls, improve the strength of a control, and distribute the responsibility across several controls to lower the CPI.

This is your first step

01

Choose a critical security or resilience objective (such as protecting confidential data) and begin listing all the controls within your organization that contribute to fulfilling that objective.

How would you benefit?

- **Identify single points of failure:** A clear CPI highlights areas of weakness and vulnerability in your defenses.
- **Informed decision-making:** Guides security spending and effort to strengthen the most crucial aspects of your defenses.
- **Demonstrate due diligence:** Calculating CPI is a way to show regulators and stakeholders that you are proactively managing risk.

How should you measure success?

CPI over time

Recalculate your CPI regularly, focusing on critical objectives. Aim to see a downward trend as you implement measures to strengthen controls and reduce single-point dependencies.

Number of high CPI objectives

Track how many critical objectives have a high-pressure index. A decreasing number highlights progress.

Incident analysis

When incidents do occur, note if they were caused by the failure of single, over-relied-upon controls.

How Google Cloud can help

01

Cloud as immune system: Cloud platforms offer built-in redundancy and distributed architectures. This means if one control fails, others can immediately take over, preventing a single point of failure from compromising the entire system.

02

Immutable infrastructure: Cloud computing enables the concept of immutable infrastructure, where servers and components are replaced rather than modified. This can help prevent configuration drift and ensure that security controls remain consistent and effective.

03

Rapid patching and updates: Cloud environments often allow for faster patching and updates compared to traditional on-premises systems. This means vulnerabilities can be remediated more quickly, reducing the window of time during which a single control is exposed to risk.

06

Software Reproducibility

The report states that ***software reproducibility*** *is the "extent of software in a particular system that can be repeatedly and continuously built and distributed while maintaining conformance with the National Institute of Standards and Technology (NIST) June 2023 Secure Software Development Framework (SSDF) requirements, including disclosing software bill of materials (SBOMs) and supply chain levels for software artifacts (SLSA) conformance levels. It is especially critical for vendors of software to critical infrastructure sectors to provide vital patches in a timely manner that will work with an infrastructure organization's updated IT environment and for software providers to assure the continuity of the build environments throughout that software's supported life. Critical infrastructure (e.g., hospitals, water) must be able to update legacy systems without losing additional software tools. A software reproducibility metric could be contextualized as "time to support" surrounding updates. Modern software lifecycle management practices in DevSecOps approaches are highly applicable."*

What does software reproducibility mean?

- **Software reproducibility:** The ability to consistently rebuild and distribute software that remains compliant with security standards and can integrate updates smoothly with minimal disruptions in critical infrastructure.
- **NIST's Secure Software Development Framework (SSDF):** A set of guidelines for secure software development practices issued by the U.S. Office of Management and Budget.
- **Software Bill of Materials (SBOM):** A comprehensive list of a software's components, including third-party dependencies, to identify potential supply chain risks.
- **Supply Chain Levels for Software Artifacts (SLSA):** A framework for evaluating security in the software supply chain.
- **DevSecOps:** An approach that integrates security (Sec) into the software development (Dev) and operations (Ops) lifecycle.

Why does software reproducibility matter to your org?

- **Timely security updates:** Enables organizations in critical infrastructure to rapidly apply security patches to mitigate vulnerabilities.
- **Supply chain trust:** Building reproducible software and providing SBOMs improve transparency and trust in the software, minimizing supply chain risks.
- **Legacy system support:** Ensures long-term compatibility of software updates within complex and often outdated critical infrastructure environments.
- **Compliance:** Fulfilling OMB's SSDF and SLSA compliance guidelines demonstrates commitment to security and builds trust.

How can you make software reproducibility happen?

Source Code Management: Use version control systems (e.g., Git) to track every change to the codebase meticulously.

Build Automation: Implement automated build and deployment pipelines (CI/CD) for reliable and repeatable processes.

Containerization: Package software and its dependencies into containers, promoting reproducible environments regardless of the underlying infrastructure.

SBOM Generation: Produce SBOMs automatically during the build process to track software components and their origins.

DevSecOps Integration: Shift security left within the development process, addressing vulnerabilities early on.

This is your first step

01

Assess your current software development process. Identify areas where you can improve reproducibility, such as source code management and build automation.

How would you benefit?

- **Reduced risk:** Minimize vulnerabilities and expedite security patch deployment.
- **Operational efficiency:** Streamline software updates for critical systems, decreasing operational downtime and the risk of disruptions.
- **Enhanced reputation:** Demonstrate a commitment to secure software practices to clients and regulators.

How can you measure success?

Time to support

Track the time it takes to roll out a security update from release to successful deployment within critical infrastructure systems.

SBOM completeness

Measure the percentage of software for which accurate and up-to-date SBOMs are generated.

Vulnerabilities patched

Track the number of critical vulnerabilities your reproducibility approach helps address in a given period.

How does Google Cloud help?

01

Robust and verifiable software supply chain: Modern cloud computing delivers on the secure supply chain for software, including Managed and Secure Open Source.

02

Cloud way as a roadmap even if some things are local: cloud-style practices increase resilience even if used on-premise or for private cloud. Cloud platforms facilitate DevSecOps practices, where security is integrated throughout the software development lifecycle. This ensures that security controls are built into the software from the start and are consistently applied, supporting secure software development.

03

Infrastructure as Code (IaC): Cloud platforms excel in IaC support. This allows you to define your entire infrastructure in code, version it, and replicate it consistently across environments. This ensures your software runs on identical foundations, minimizing variability.

04

Immutable Infrastructure: Cloud promotes the concept of immutable infrastructure, where servers or containers are replaced rather than modified. This eliminates configuration drift and ensures each deployment starts from a known, reproducible state.

05

Global Distribution: Cloud's global reach allows you to replicate your software across multiple regions. This not only improves disaster recovery but also lets you test reproducibility in different geographic locations, exposing potential environment-specific issues.

07

Preventative Maintenance Levels

The report states that ***preventative maintenance levels*** *are the "percentage of the overall cost of systems operations that is devoted to preventative maintenance (e.g., upgrades, security patching, reducing technical debt)."*

What does it mean?

Preventative maintenance (PM): Refers to proactive maintenance activities performed on equipment, systems, and software to prevent breakdowns and failures before they occur. This includes upgrades, security patches, and general upkeep to reduce technical debt.

Technical debt: The concept that cutting corners in software development for short-term gains can lead to longer-term problems, increased complexity, and higher maintenance costs.

Preventative maintenance levels: Are measured as the percentage of operational costs allocated explicitly to carry out proactive maintenance tasks.

Why does it matter to your org?

- **Reduced downtime**: Regular PM can significantly reduce the risk of unexpected disruptions and outages, resulting in greater operational stability.
- **Cost savings**: Preventative maintenance can be more cost-effective in the long run compared to the costs of emergency repairs or system replacements.
- **Longer asset lifespan**: PM helps extend the useful life of assets and delays the need for expensive replacements.
- **Improved safety**: PM can contribute to a safer working environment by mitigating potential hazards and ensuring equipment and systems operate as intended.

How to make preventative maintenance happen?

Develop PM schedules: Define clear preventative maintenance plans and schedules for each critical system based on manufacturer recommendations, historical data, and risk assessments.

Track maintenance activities: Implement a system (software tools or even simple logs) to record when maintenance occurs and to track any emerging issues.

Budget allocation: Devote a specific portion of your operating budget to preventative maintenance activities, ensuring they remain a priority.

Training: Provide staff with the necessary training and resources for the types of PM to be performed.

This is your first step

01

Start by creating an inventory of all critical equipment and systems within your organization that require regular maintenance.

How would you benefit?

- **Increased system reliability:** Experience fewer unexpected failures and breakdowns.
- **Lower operational costs:** Achieve potential cost savings by addressing issues proactively and avoiding larger repairs down the line.
- **Regulatory compliance:** In some industries, a strong PM program may be a regulatory requirement.

How can you measure success?

Downtime reduction

Track unplanned downtime before and after implementing a PM program.

Return on investment (ROI)

Calculate the cost savings (in avoided repairs, etc.) compared to the investment in PM.

Asset lifespan

Track the lifespan of equipment and compare those under a PM program versus those without.

How Google Cloud can help?

01

Automated patching and updates: Cloud platforms excel at automating the deployment of security patches and software updates. This ensures your systems are up-to-date, reducing vulnerabilities and minimizing the risk of unexpected failures.

02

Remote monitoring and diagnostics: Cloud providers offer comprehensive monitoring tools that enable proactive identification and resolution of potential issues before they escalate into critical failures. This allows for timely maintenance and prevents costly downtime.

03

Elastic scaling for maintenance windows: Cloud resources can be easily scaled up or down to accommodate maintenance activities without impacting system availability. This flexibility allows for seamless maintenance without disrupting user experience.

04

Version control and rollback capabilities: Cloud environments provide robust version control and rollback mechanisms. If issues arise after updates or changes, you can easily revert to a previous stable version, minimizing downtime and data loss.

08

Inventory Completeness

The report states that ***inventory completeness*** *is the "extent of the universe of an organization's operations – including information technology (IT), operations technology (OT), and supply chain (to 4th party as well as 3rd party)—that is encapsulated in a validated and managed inventory or asset register."*

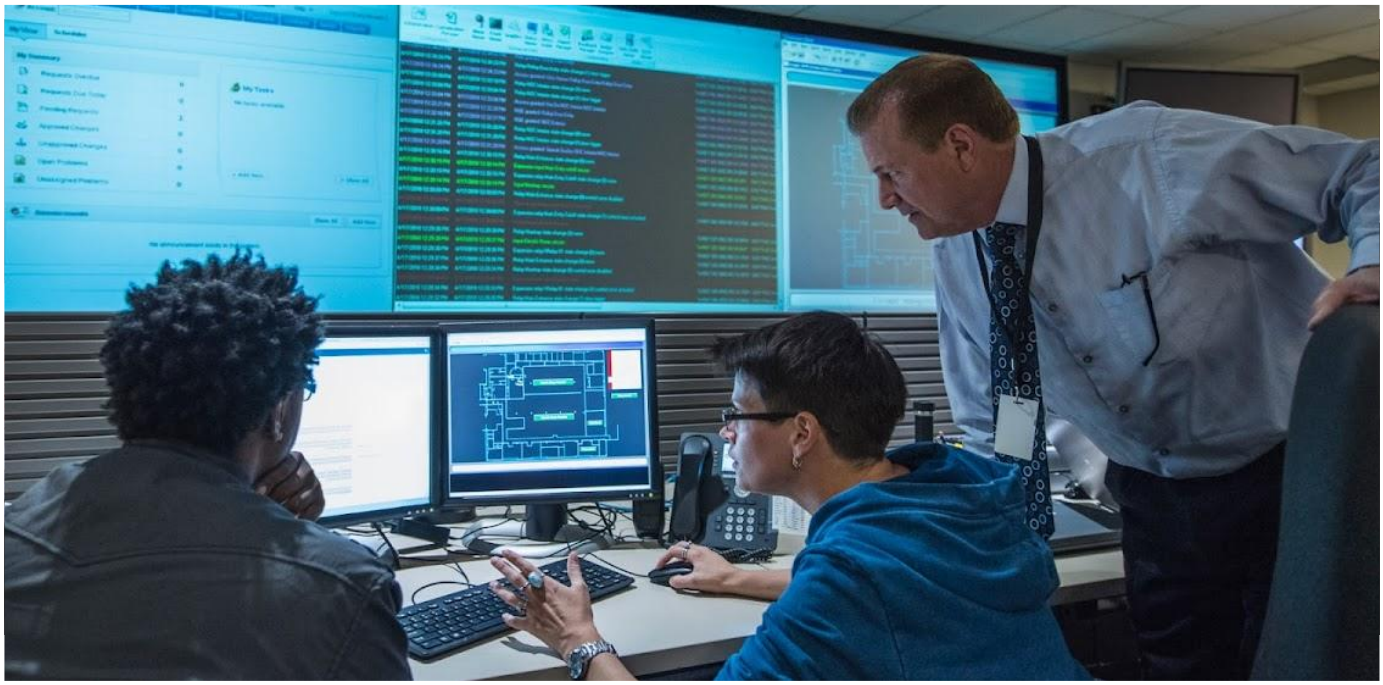

What does it mean?

- **Inventory completeness:** Measures how well your organization has identified and documented all relevant assets, encompassing IT, OT, and supply chain elements.
- **Information technology (IT):** The traditional computers, networks, software, and data used within an organization for handling information.
- **Operational technology (OT):** The hardware and software components directly responsible for monitoring and controlling physical devices and processes (such as industrial machinery).
- **Supply chain:** The network of vendors, suppliers, and distributors involved in creating and delivering goods or services, including 3rd party (direct) and 4th party (indirect) relationships.
- **Asset register:** A centralized database or record-keeping system that houses comprehensive information about all assets owned or controlled by the organization.

Why does it matter to your org?

- **Cybersecurity risk management:** Incomplete inventories leave blind spots, increasing the likelihood that unmanaged and unpatched assets create vulnerabilities.
- **Incident response:** Rapid and effective response to incidents depends on knowing exactly what systems are affected and where they reside in the network.
- **Supply chain transparency:** Complete inventories enhance your understanding of supply chain dependencies and potential risks beyond immediate (third-party) vendors.
- **Regulatory compliance:** Certain industries have mandates for maintaining comprehensive asset inventories as part of compliance standards.

How to make inventory completeness happen?

Automated discovery tools: Use network scanning and asset discovery tools to identify and catalog IT and OT components.

Manual processes: Supplement with manual data collection, especially for OT equipment or legacy systems that might not be easily detected by automated means.

Supply chain mapping: Collaborate with suppliers and vendors to gather information about components and materials they provide, including their own supply chain dependencies (4th party).

Regular updates: Schedule periodic reviews to ensure the inventory remains up-to-date as assets are added, removed, or modified.

This is your first step

01

Determine who within your organization will own and maintain the asset register. This will provide a clear point of accountability.

How can you benefit?

- **Improved threat visibility:** A complete inventory exposes potential attack surfaces and vulnerabilities that might otherwise go unnoticed.
- **Operational efficiency:** Maintain better control of assets for software licensing, updates, and support purposes.
- **Informed decision making:** Make better strategic decisions about technology investments and risk management with accurate asset data.

How can you measure success?

Percentage of assets registered

Calculate the percentage of known IT, OT, and supply chain components tracked in your asset register. Aim to close any gaps.

Vulnerabilities identified

Track how many vulnerabilities are uncovered as a result of improvements to your inventory completeness.

Data accuracy

Regularly audit the inventory to ensure asset details are correct and up to date.

How Google Cloud can help?

01

Improved visibility: Cloud platforms often provide robust asset discovery and tracking capabilities. This can help you identify and catalog all your cloud resources, reducing the chances of overlooking critical components in your inventory.

02

Automation: Cloud services can automate many inventory management tasks, such as data collection, reconciliation, and reporting. This reduces the potential for human error and ensures more accurate and up-to-date inventory records.

03

Integration with third-party tools: Cloud platforms often integrate seamlessly with specialized inventory management and asset tracking solutions. This allows you to use advanced features and maintain a single source of truth for your inventory data.

04

Scalability and flexibility: Cloud infrastructure can easily scale to accommodate growing inventory needs. This ensures you have the necessary resources to track and manage your assets effectively, even as your environment expands.

09

Stress-Testing Vibrancy

The report states that ***stress-testing vibrancy, also known as red teaming,*** *tests the "extent of systems that have been subjected to an extreme offensive, adversarial security test (possibly AI augmented), to test defenses against reliable operation (this could be against an especially constructed "cyber range" and might be achieved with "chaos engineering" principles. This should include explicit testing against multi-point attacks—where an adversary is coming after multiple points in a system with multiple tactics, potentially both physical and cyber."*

What does it mean?

Stress-testing vibrancy (red teaming): Refers to a comprehensive security assessment approach where a team acts like an adversary (red team) using advanced techniques (potentially AI-powered) to identify vulnerabilities in a system's ability to maintain reliable operation under extreme attack.

Cyber range or virtual red teaming: A simulated environment designed to replicate real-world systems for safe and controlled security testing. Attack simulations and impact assessment may use GenAI for hypothesizing the impact on the environment.

Chaos engineering: Practices that involve deliberately introducing controlled disruptions to systems to understand how they respond and ensure resilience.

Multi-point attacks: Coordinated attacks targeting various parts of a system simultaneously using diverse tactics, potentially including both physical and cyber elements.

Why does it matter to your org?

- **Identifying unforeseen weaknesses:** Red teaming uncovers vulnerabilities traditional penetration testing might miss, revealing how systems respond under sustained pressure.
- **Prepare for real-world threats:** Provides a realistic picture of what your organization might face from a sophisticated adversary, aiding in developing appropriate defenses.
- **Test resilience:** Evaluates your ability to withstand multi-pronged attacks and recover efficiently to maintain critical operations.

How to make stress-tested vibrancy happen?

Define scope and objectives: Clearly define the systems under test, the level of attack intensity, and the desired learning outcomes of the red teaming exercise.

Utilize cyber ranges and red teams: Consider using a secure cyber range environment to simulate realistic attacks without risking damage to production systems.

Chaos engineering integration: Incorporate chaos engineering principles to test how systems respond to unexpected disruptions and failures.

This is your first step

01

Conduct a basic security assessment of your systems to identify potential targets for red teaming exercises.

How would you benefit?

- **Enhanced security posture:** Proactively discover vulnerabilities before attackers exploit them, leading to a more robust security posture.
- **Improved incident response:** Red teaming helps identify weaknesses in your incident response plans and allows teams to practice effective procedures.
- **Increased confidence:** Strengthens your confidence in your organization's ability to respond effectively to complex and multi-faceted attacks.

How to measure success?

Vulnerabilities discovered

Track the number and severity of vulnerabilities identified through the red teaming exercise.

Improved incident response

Evaluate team performance during the exercise, assessing their ability to detect, contain, and recover from attacks.

Lessons learned

Measure the effectiveness of the exercise in identifying critical security gaps and implementing corrective actions.

Determining the frequency and scope

Red teaming exercises will need to be weighed against available resources and risk tolerance.

How Google Cloud can help?

01

Scalability: Cloud platforms allow for on-demand provisioning of resources, enabling you to simulate large-scale, real-world attacks that thoroughly test your system's limits. This provides a more realistic testing environment compared to traditional on-premises setups.

02

Cost-effectiveness: Cloud resources can be spun up and down quickly, allowing you to create temporary, isolated environments for red teaming exercises without significant capital investment. This makes it easier to conduct frequent and comprehensive testing.

03

Diverse deployments: Cloud providers offer a variety of configurations and geographic regions. This allows you to test your system's resilience across different environments, exposing potential vulnerabilities specific to certain setups or locations.

04

Rapid replication: Cloud environments can be easily replicated and reset, enabling you to quickly iterate on testing scenarios and reproduce results. This accelerates the feedback loop and allows for more efficient and effective red teaming exercises.

10

Common Mode Failures

The report states that ***common-mode failures and dependencies*** *can help identify organizations (and others in their supply chain) that in the event of failure would represent significant harm to a whole sector – because of the concentration they represent. As part of this, finding and eliminating circular dependencies is vital i.e., organization X depends on Y to cover and vice versa."*

What does it mean?

Common-mode failures (CMFs): Simultaneous or near-simultaneous failures of multiple, seemingly independent systems as a result of a single shared cause.

Dependencies: Reliance one organization or component has upon another within a system or supply chain.

Significant harm to a sector: Potential negative consequences on an entire industry or sector stemming from a concentrated disruption or failure within a small number of key entities.

Circular dependencies: Situations where two or more organizations are mutually reliant upon each other, such that a failure in one could cascade and lead to the failure of others.

Why does it matter to your org?

- **Sector-wide risk management:** Understanding CMFs and critical dependencies helps manage systemic risk across an entire industry sector, even beyond your organization's direct control.
- **Supply chain resilience:** Reveals hidden dependencies and vulnerabilities within your supply chain, where the failure of a single third-party supplier could cause cascading issues.
- **Business continuity:** Proactively identifying CMFs and circular dependencies enables implementing mitigation strategies to maintain operations in the event of disruptions caused by failures outside your organization.

How to make common-mode failure resiliency happen?

Mapping critical dependencies: Map out interdependencies within your sector, highlighting key suppliers, shared resources, and infrastructure providers.

Vulnerability assessment: Evaluate the potential impact of known common mode failure risks (e.g., natural disasters, cyberattacks, technology failures) on critical dependencies.

Analyze circular dependencies: Systematically examine relationships for circular dependencies, identifying those that pose the highest risks.

Mitigation strategies: Develop plans to decouple circular dependencies, establish redundancies, or develop alternative suppliers/resources to lessen the impact of single-point failures.

This is your first step

01

Identify critical external providers (utilities, software vendors, key suppliers, etc.) whose failure would dramatically impact your operations.

How can you benefit?

- **Increase sector resilience:** Contributing to a more robust industry sector less likely to experience systemic failures due to common mode problems.
- **Prevent cascading failures:** Reduce the risk of operational disruption caused by failures in interdependent organizations within your sector.
- **Regulatory compliance:** In some critical sectors, demonstrating proactive identification of CMFs and dependencies may be a regulatory requirement.

How can you measure success?

Number of circular dependencies eliminated

Track how many circular dependencies you've successfully resolved.

Potential disruptions avoided

Monitor if you've reduced the risk of cascading failures during actual disruption events.

Regulatory appraisal

In industries with strict compliance, positive assessments during audits can indicate success.

How can Google Cloud help?

01

Diverse deployments: Cloud providers offer geographically distributed data centers and redundant resources. This allows you to spread your applications across multiple locations, reducing the impact of common mode failures (such as natural disasters or power outages) that could take down an entire on-premises system.

02

Fault isolation: Cloud platforms facilitate containerization and microservice architectures. By breaking down applications into smaller, isolated components, failures are localized, preventing cascading effects that could cripple the entire system.

03

Automated redundancy management: Cloud services can automatically manage redundant instances and failover mechanisms. This ensures service continuity even during failures, as the system can seamlessly switch to healthy resources without manual intervention.

04

Monitoring and rapid response: Cloud providers offer robust monitoring tools that enable early detection of common mode failures and performance bottlenecks. This allows for faster response and recovery, minimizing downtime and data loss.

05

Multicloud focus: Using a multicloud strategy can significantly mitigate the risks associated with common mode failures. By distributing applications and data across multiple cloud providers, organizations reduce the likelihood of a single incident or outage impacting their entire system.

11

Conclusion

The PCAST report serves as a critical roadmap for enhancing the cyber-physical resilience of our nation's infrastructure. The 10 leading indicators provide a framework for organizations to assess their current resilience posture and identify areas for improvement. By embracing these indicators and taking proactive steps, organizations can significantly reduce their vulnerability to cyber-physical attacks and ensure the continuity of essential services.

The future of cyber-physical resilience lies in a collective effort that spans across government agencies, private sector organizations, and individuals. By working together to implement the PCAST recommendations, we can create a more secure and resilient infrastructure that can withstand the evolving threat landscape. This will require ongoing investment in research and development, as well as a commitment to fostering a culture of resilience across all levels of society.

Ultimately, the goal is to build a future where cyber-physical systems are not only secure but also adaptable and capable of recovering quickly from disruptions. This will require a continuous effort to identify and mitigate vulnerabilities, as well as a willingness to embrace new technologies and approaches that can enhance resilience. By taking proactive steps today, we can ensure a safer and more resilient future for all.