



**RedLock**  
Cloud Threat Defense



Cloud Security Buyer's Guide  
Based on the

**NIST Cybersecurity Framework**

<b>Overview</b>	<b>3</b>
<b>01 - Function: Identify</b>	<b>5</b>
● <b>Asset Management</b>	<b>5</b>
● <b>Risk Assessment</b>	<b>6</b>
<b>02 - Function: Protect</b>	<b>7</b>
● <b>Access Control</b>	<b>7</b>
● <b>Data Security</b>	<b>7</b>
● <b>Information Protection Processes and Procedures</b>	<b>8</b>
● <b>Protective Technology</b>	<b>8</b>
<b>03 - Function: Detect</b>	<b>10</b>
● <b>Anomalies and Events</b>	<b>10</b>
● <b>Security Continuous Monitoring</b>	<b>11</b>
<b>04 - Function: Respond</b>	<b>13</b>
● <b>Analysis</b>	<b>13</b>
● <b>Mitigation</b>	<b>13</b>
<b>Summary</b>	<b>15</b>
<b>Checklist</b>	<b>16</b>
<b>About RedLock</b>	<b>18</b>

# Overview

Public cloud computing adoption is outpacing cybersecurity defenses. The absence of a physical network boundary to the internet combined with the risk of accidental exposure by users with limited security expertise, increases the attack surface in the cloud by orders of magnitude. It is imperative for organizations to develop an effective strategy to protect their Amazon Web Services (AWS), Microsoft Azure, and Google Cloud environments.

On February 12 2013, the President of the United States issued Executive Order 13636, "Improving Critical Infrastructure Security", which called for the development of a set of best practices to help organizations manage cybersecurity risks. In response, the National Institute of Standards and Technology (NIST) created the Cybersecurity Framework (CSF) in 2014 through collaboration between government and the private sector. The framework does not place additional regulatory requirements, rather it references globally recognized standards for cybersecurity such as ISO/IEC 27001, ISA/IEC 62443, and COBIT 5.

While the foundations of the NIST CSF framework are based on standards and security best practices for traditional on-premise infrastructure, it can easily be extended to manage risks across cloud computing environments such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud. Applying this framework, enables organizations to meet their obligations in the shared responsibility model and secure systems, networks, and users within the cloud computing environment.

The CSF framework outlines the set of policy, business, and technological requirements for managing risk. It consists of five concurrent and continuous functions: Identify, Protect, Detect, Respond, and Recover. This guide specifically focuses on key technological requirements for four of the functions (Identify, Protect, Detect, Respond) and explains how they translate for managing risks within a cloud computing environment.

The guide also contains a handy checklist of requirements that can be used to evaluate security solutions to protect your cloud computing environment.

## IDENTIFY

### Asset Management

- Resource identification
- Application identification
- Data flow mapping

### Risk Assessment

- Third party feed ingestion
- Resource risk scoring and prioritization

## DETECT

### Anomalies and Events

- User and network behavior baseline
- Event correlation
- Alerts

### Security Continuous Monitoring

- Network monitoring
- User activity monitoring
- Resource monitoring
- Vulnerability monitoring

## PROTECT

### Access Control

- Access permissions management
- Network integrity protection

### Data Security

- Data-at-rest is protected
- Data-in-transit is protected

### Information Protection & Procedures

- Configuration baseline created and maintained

### Protective Technology

- Audit logs maintained and reviewed

## RESPOND

### Analysis

- Impact analysis
- Network forensics

### Mitigation

- Incident contained
- Incident mitigated

# 01

## Function: Identify

The goal of the Identify function of the NIST CSF framework is to develop the organizational understanding necessary to manage cybersecurity risk. This requires identifying data, personnel, devices, systems, and facilities within the environment and understanding the risks associated with them.

### Asset Management

#### **ID.AM-1: Physical devices and systems within the organization are inventoried**

Since organizations do not manage physical devices and systems in public cloud computing environments, this specific requirement can be translated to inventorying cloud resources instead. In order to be able to effectively protect your environment, it is important to have visibility into compute instances, managed databases, storage services, virtual networks, and users within it. For example, if you know that one or more of the Google Compute Engine instances in your environment is a database, you can create a policy that monitors databases for direct connections to the internet since that is a bad security practice.

Identifying cloud resources can be challenging due to their ephemeral nature and rate of change. According to [research](#) by the RedLock Cloud Security Intelligence (CSI) team, the average lifespan of a cloud resource is only 127 minutes. Given the rapid rate of change, manually auditing the environment is infeasible. The only way to maintain an accurate inventory in a constantly changing environment is by automating resource discovery using APIs.

#### **ID.AM-2: Software platforms and applications within the organization are inventoried**

The ephemeral nature of the environment poses similar challenges for inventorying software platforms and applications. Since no application identification is provided by default, the only way to do this is by applying AI on disparate data sets to determine the applications running in cloud compute instances.

#### **ID.AM-3: Organizational communication and data flows are mapped**

Mapping organizational data flows not only involves monitoring north-south network traffic (between the internet and the cloud environment), but also east-west traffic (within applications in the cloud environment). This is especially important with the growing adoption of containers and microservice architectures. In cloud computing environments

# 01

## Function: Identify

where workload IP addresses are constantly changing, simply analyzing netflow logs will not provide the necessary context. In order to build an accurate data flow between applications, netflow logs must be correlated with configuration data.

### Risk Assessment

#### **ID.RA-2: Threat and vulnerability information is received from information sharing forums and sources**

Numerous third party feeds that provide network threat intelligence and host vulnerability information exist for your on-premise environments. However, using those feeds in a silo to assess the risk posture of cloud workloads does not work since they lack context on the cloud environment. For example, receiving a weekly Qualys or Nessus report on missing patches for specific IP addresses is inadequate when you do not know what types of applications are running on those IPs, when the scans were run, if they were accepting connections from the internet, and whether they had access to sensitive data.

To properly solve this, the feeds from traditional threat and vulnerability management solutions should be continuously ingested and augmented with cloud-specific insights to automatically assess the risk posture of your cloud computing environment.

#### **ID.RA-5: Threats, vulnerabilities, likelihoods, and impacts are used to determine risk**

Context is important to be able to accurately assess risk. In cloud computing environments, context is obtained by correlating a resource's configurations, associated user activities, network communications, vulnerability information, and threat intelligence. Automated risk scoring based on the correlated data is required to prioritize resource risk and enable faster remediation. For example, a resource with a known vulnerability that is receiving suspicious traffic from the internet, should receive a higher risk score than one that is vulnerable but not exposed to the internet. For environments with hundreds of resources across numerous cloud accounts, correlating these disparate data sets can be challenging.

## 02 Function: Protect

The goal of the Protect function of the NIST CSF framework is to implement safeguards that ensure delivery of services. This includes controlling access, segregating networks, and establishing configuration baselines.

### Access Control

#### **PR.AC-4: Access permissions are managed, incorporating the principles of least privilege and separation of duties**

IAM roles are used to control access to cloud computing environments. Quite often, users are given overly permissive access than is necessary to perform their duties, simply because administrators lack a reasonable understanding of the job role and access requirements. It is important to implement a solution that monitors IAM role configurations, and compares them against actual usage patterns to recommend limiting of access to services not actively in use.

#### **PR.AC-5: Network integrity is protected, incorporating network segregation where appropriate**

Unlike on-premise environments, the network in cloud computing environments is software defined. In cloud environments, it is not only important to segment network traffic using security groups and IP tables, but to also continuously monitor the traffic to ensure that the applications are only receiving intended traffic and automatically remediate any issues. Research from the RedLock CSI team revealed that 93% of resources do not restrict outbound traffic at all which makes accidental data loss or data exfiltration in the event of a breach very easy.

### Data Security

#### **PR.DS-1: Data-at-rest is protected**

Encrypting data at rest is vital for regulatory compliance to ensure that sensitive data saved on disks is not readable by any user or application without a valid key. Cloud service providers typically offer data-at-rest encryption with customer controlled keys. Organizations should enable the cloud-native encryption, monitor the environments for data stores

## 02 Function: Protect

without encryption enabled, and remediate any issues. Surprisingly, research from the the RedLock CSI team revealed that 82% of databases in the cloud are not encrypted which goes against established best practices and risks non-compliance with regulatory mandates such as GDPR.

### **PR.DS-2: Data-in-transit is protected**

All traffic to and from workloads running critical applications must be encrypted in-transit. This is critical to protect data breaches from man-in-the-middle attacks. This should be achieved by enforcing appropriate TLS configurations on managed cloud services such as load balancers, storage buckets, and databases. Finally, customers should have security tools in place to monitor traffic to and from these workloads and automatically remediate any issues.

## Information Protection Processes and Procedures

### **PR.IP-1: Baseline configuration is created and maintained**

Configuration drift is a common issue in cloud computing environments. Such basic misconfigurations occur because users can change resource configurations on-demand without going through a security change control process. For example, a user can make an Amazon S3 bucket accessible to anyone on the internet through a simple configuration change, without understanding the risks associated with the action.

The problem is amplified when there are a large number of users making changes across hundreds of resources in dozens of cloud environments. As a result, it is necessary to automate configuration baseline creation and monitoring in cloud computing environments. This process can be streamlined using industry standard benchmarks such as [CIS Foundations](#).

## Protective Technology

### **PR.PT-1: Audit/log records are determined, documented, implemented, and reviewed in accordance with policy**

Cloud services can log user activity to maintain an audit trail. However, it is an optional configuration and up to



## 02 Function: Protect

individual organizations to enable. As a security best practice, organizations should enable audit logging, monitor the configuration, and automatically enable it if it is disabled.

Manually reviewing audit logs for suspicious activities can be very tedious, especially when there are a large number of users making changes across hundreds of resources in dozens of cloud environments. The only feasible way to achieve this involves applying AI to baseline activity and identify deviations for each user that could indicate account hijacking or malicious insiders.

The goal of the Detect function of the NIST CSF framework is to identify the occurrence of a cybersecurity event. This involves continuous monitoring for anomalous user behavior, suspicious network traffic, and vulnerable resources.

## Anomalies and Events

### **DE.AE-1: A baseline of network operations and expected data flows for users and systems is established and managed**

Baselining user and network behavior is challenging in cloud computing environments due to the sheer volume and velocity of change. With the shift to DevOps, scripts are being used to automate many activities which makes it impossible to manually track behavior, especially across a large volume of cloud resources.

To exacerbate matters, environments where users are working across multiple cloud accounts requires correlation of logs across accounts which is non-trivial.

Additionally, microservice-architectures in the cloud are often short-lived, and generate vast amount of network traffic that must be analyzed. As a result, an AI-based approach is necessary in order to baseline user and network behavior in such dynamic environments.

### **DE.AE-3: Event data are aggregated and correlated from multiple sources**

Data aggregation and correlation is important for creating context. Specifically, resource configurations, user activity, network traffic, and threat intelligence data must be correlated in order to paint a comprehensive picture of risk.

Accessing this data in cloud computing environments can be challenging for a few reasons. The fragmented architecture of cloud environments across multiple accounts makes data collection from multiple sources tedious. Moreover, organizations do not have physical access to the infrastructure for popular API-driven services such as Amazon S3. The only way to access the necessary data for these resources is using APIs. Lastly, the volume of network traffic logs is often quite large which requires expertise in data science to derive meaningful insights. An AI-based approach that correlates massive volumes of disparate data sets including configuration, user activity, network traffic, and threat intelligence data is required to create context and assess true risk.

**DE.AE-5: Incident alert thresholds are established**

Traditionally, alert thresholds are defined by policy and they are triggered when these thresholds are crossed. Quite often this can lead to a lot of false positives or negatives, causing alert fatigue.

Alerts that provide context are far more actionable. So for example, an alert that indicates that a security group is open to the internet definitely highlights a poor security practice. A contextual alert on the other hand, will actually indicate if there is an active threat. Using a similar example: a contextual alert is raised if a security group is open to the internet and associated with a database that is receiving traffic from a suspicious IP address, which indicates an active threat and enables you to quickly respond. As discussed earlier, context is gleaned from correlating configuration, user activity, network traffic, and threat intelligence data, reinforcing the need for AI-driven data correlation.

## Security Continuous Monitoring

**DE.CM-1: The network is monitored to detect potential cybersecurity events**

Unlike on-premise environments, the network in cloud computing environments is software defined, which creates greater risk of exposure due to logic errors. For example, a common human error of associating an overly permissive security group with an Amazon RDS service can lead to significant data exposure.

It is imperative to implement a network monitoring solution for cloud computing environments. Traditional on-premise solutions such as proxies and agents create blind spots since they cannot be implemented for popular API-driven services such as Amazon ELB and Amazon RDS. An API-based approach that ingests netflow logs from a cloud computing environment is the most effective way to achieve comprehensive network monitoring and detect potential cybersecurity events.

However, analyzing netflow logs alone is not effective since resource IP addresses are constantly changing in cloud computing environments. In order to build an accurate map, netflow logs must be correlated with configuration and threat intelligence data. This enables mapping flows to associated resources, determining traffic directionality, as well as identifying the type of traffic (north-south, east-west, suspicious traffic).

**DE.CM-3: Personnel activity is monitored to detect potential cybersecurity events**

While the shift to DevOps is enabling rapid innovation, it creates an additional dimension of risk of exposure from privileged users with limited security expertise. These users are typically provisioned with broader access than usually necessary for their role, since granting granular access in the cloud is complex and unwieldy.

Monitoring user activities to identify malicious insiders and account compromises is necessary to detect threats within your cloud computing environment. An AI-based approach is required to detect issues such as insider threats and account compromises. Recent high profile breaches involving compromised credentials and API access keys (e.g., OneLogin) demonstrate the importance of this capability.

**DE.CM-7: Monitoring for unauthorized personnel, connections, devices, and software is performed**

DevOps teams have substantial autonomy to configure and operate cloud environments depending on business needs. Often developers can be found running “non-standard” images, network configurations, and software without realizing the risks this can introduce and without the knowledge of security teams. It is important for security teams to be able to rapidly detect any deviations from organizational policies, and identify any unauthorized users or cloud services.

**DE.CM-8: Vulnerability scans are performed**

Identifying and remediating vulnerabilities in a cloud computing environment is far more tedious than in on-premise networks. The challenge lies with the fact that resources are created on-demand without any context on their identity which makes it tough to determine which patches are relevant to a specific resource. Automated resource discovery alleviates this issue (ID.AM-1) which then enables you to ingest your existing vulnerability data feeds to scan your cloud computing environment for any unpatched resources.

## 04 Function: Respond

The goal of the Respond function of the NIST CSF framework is to develop and implement measures to take action when a cybersecurity incident occurs. This involves responding, analyzing, and mitigating incidents.

### Analysis

#### **RS.AN-2: The impact of the incident is understood**

Understanding the impact of an incident in a cloud computing environment can be difficult, especially if it is constantly changing. For example, if you determine that there was traffic to an unpatched MongoDB resource from a suspicious IP address a month ago, you will want to drill down and determine what other resources the database was connected to at that time in order to assess downstream impact. Due to the ephemeral nature of cloud computing environments, that database may have already been terminated, which makes it difficult to perform the analysis. In order to be able to perform impact analysis of an incident at a point in time, historical snapshots of the environment must be captured and available for incident response teams.

#### **RS.AN-3: Forensics are performed**

Performing forensics in a cloud computing environment is challenging. Consider a scenario where unpatched MongoDBs are being targeted in a ransomware attack. You will want to prioritize identifying unpatched MongoDB resources in your environment and determine if there is any network traffic from these resources to the internet. Having a map that highlights these instances and visualizes the network flows aids in quickly determining your potential exposure. Context is required to build this map which is derived from using AI to correlate configuration, user activity, network traffic, host vulnerability, and threat intelligence data.

### Mitigation

#### **RS.MI-1: Incidents are contained**

When an incident occurs, taking measures to quickly contain the issue is critical. For example, if there is an alert warning about an Amazon S3 bucket that has been configured to allow public access, your cloud security solution should contain the issue by disabling public access for the bucket.

## 04 Function: Respond

### **RS.MI-2: Incidents are mitigated**

Timely mitigation in cloud computing poses challenges due to the lack of a central entity that governs the entire environment. Instead, incidents have to be dealt with on an individual basis with the corresponding DevOps team that owns the impacted resources. In addition, full context is required to be able to effectively resolve the issue (DE.AE-3).

Chances are that you already have remediation workflows in place for on-premise incidents. Integrating your cloud security solution with your enterprise workflow management tools allows you to maximize your existing investments. Alternately, your cloud security solution could automatically remediate the issue. So in the example above, as soon as the publicly exposed Azure blob storage is discovered, the cloud security solution should automatically change the configuration to private.











## Summary

Public cloud computing presents new security challenges. Current security architectures consisting of point solutions address discrete problems, but are ineffective for detecting advanced threats in the cloud. Moreover, the operational challenges created by this approach are exacerbated in distributed cloud computing environments.













Threat defense in the cloud requires a new AI-driven approach that correlates disparate security data sets including network traffic, user activities, risky configurations and threat intelligence, to provide a unified view of risks across fragmented cloud environments. The RedLock 360 platform is built on these exact principles to help organizations achieve effective security governance across AWS, Azure, and Google Cloud.

A checklist is included in this guide to serve as a quick reference when you evaluate security solutions to protect your cloud computing environment.

# Checklist

NIST CSF Subcategory	Recommendation	RedLock	Solution #2	Solution #3
<b>IDENTIFY</b>				
<b>ID.AM-1:</b> Resources identification	Correlate data sets to identify resources			
<b>ID.AM-2:</b> Application identification	Correlate data sets to identify applications			
<b>ID.AM-3:</b> Data flow mapping	Correlate data sets to build flow map			
<b>ID.RA-2:</b> Third party feed ingestion	Ingest third party feeds such as vulnerability data to enrich risk models			
<b>ID.RA-5:</b> Resource risk scoring & prioritization	Apply automated risk scoring to prioritize resource risk			
<b>PROTECT</b>				
<b>PR.AC-4:</b> Access permissions management	Monitor IAM role configurations and immediately auto-remediate issues			
<b>PR.AC-5:</b> Network integrity protection	Monitor network configurations and immediately auto-remediate issues			
<b>PR.DS-1:</b> Data-at-rest is protected	Monitor data-at-rest encryption configurations and immediately auto-remediate issues			
<b>PR.DS-2:</b> Data-in-transit is protected	Monitor data-in-transit encryption configurations and immediately auto-remediate issues			
<b>PR.IP-1:</b> Configuration baseline creation/monitoring	Automate configuration baseline creation and monitoring			



NIST CSF Subcategory	Recommendation	RedLock	Solution #2	Solution #3
<b>PR.PT-1:</b> Audit logs maintained and reviewed	Enable audit logs and re-enable if the configuration is disabled			
<b>DETECT</b>				
<b>DE.AE-1:</b> User & network behavior baseline	Leverage AI to create user and network behavior baselines			
<b>DE.AE-3:</b> Event correlation	Correlate configuration, user activity, network traffic, and threat intelligence data			
<b>DE.AE-5:</b> Alerts	Ensure alerts provide context on risk			
<b>DE.CM-1:</b> Network monitoring	Ingest netflow logs and correlate with other data sets to interpret context			
<b>DE.CM-3:</b> User activity monitoring	Ingest audit logs and leverage AI to detect user behavior deviations from baselines			
<b>DE.CM-7:</b> Resource monitoring	Monitor resources and alert if unauthorized entity is discovered			
<b>DE.CM-8:</b> Vulnerability monitoring	Ingest vulnerability feeds and scan environment for issues			
<b>RESPOND</b>				
<b>RS.AN-2:</b> Impact analysis	Drill down using interactive risk map for downstream impact analysis			
<b>RS.AN-3:</b> Network forensics	Perform incident investigation using interactive risk map			
<b>RS.MI-1:</b> Incident contained	Get recommendations to contain the incident			
<b>RS.MI-2:</b> Incident mitigated	Integrate with existing enterprise workflow management tools, or leverage APIs to auto-remediate issues			

## About RedLock

RedLock enables effective security governance across Amazon Web Services, Microsoft Azure, and Google Cloud environments. The RedLock Cloud 360™ platform takes a new AI-driven approach that correlates disparate security data sets including network traffic, user activities, risky configurations, and threat intelligence, to provide a unified view of risks across fragmented cloud environments. With RedLock, organizations can manage risks, validate architecture, and enable security operations across cloud computing environments.

Global brands across a variety of verticals trust RedLock to secure their public cloud computing environments. RedLock has received a number of industry accolades including finalist for Most Innovative Startup at RSA 2017, CRN Emerging Vendors in Security 2017, and TiE50 Winner 2017.

To learn more:  
Call: +1.650.665.9480, Visit: [www.redlock.io](http://www.redlock.io)  
© 2017 RedLock Inc. All rights reserved.

RedLock, RedLock logo, and RedLock Cloud 360 are trademarks of RedLock Inc.  
All other registered trademarks are the properties of their respective owners.