CIMCOR

The Authoritative Guide to System Integrity Assurance

SYSTEM INTEGRITY ASSURANCE is the confidence and certainty that the appropriate security controls and compliance requirements are in place to ensure the accuracy and consistency of data throughout its entire life-cycle of operation.

Contents

Introduction

The Importance of Change Management

Why FIM Hasn't Solved Cybersecurity Problems

Conclusion

Introduction

Executive Summary

Despite constantly rising cybersecurity spending, data breaches and security incidents are rising by the year. Despite vendor-prompted calls to invest in more flashy tools and solutions, this is not a problem that organizations can spend their way out of.

Instead, organizations should focus on getting the basics right and maintaining integrity across their entire IT environment. This is what File Integrity Monitoring (FIM) tools were designed to help with—but their promise was never realized.

FIM has a bad name in the cybersecurity industry, mainly because FIM tools don't deliver on their claims or promises. Instead of maintaining integrity, they have become 'shelfware' for most organizations as they are simply too noisy and cumbersome to be useful.

In short, these tools fail to deliver integrity, instead providing nothing more than change monitoring detection.

This white paper will examine how moving away from change monitoring and towards system integrity assurance can significantly help organizations improve cybersecurity outcomes, proactively respond to security threats, reduce the time and effort needed to maintain and demonstrate compliance and employ a zero trust security practice.

Key Learning Points

- Cybersecurity teams are stuck in reactive mode, drowning under a constantly growing pile of alerts—and the prevailing approach to cybersecurity isn't doing anything to help.
- To reverse the current trends, organizations need to re-evaluate the fundamental principles of cybersecurity.
- While flashy tools get more attention, industry experts understand that fundamentals like system integrity, configuration, and change management are more important.
- Change management, a core IT practice, is critical to cybersecurity. Tracking and (where necessary) remediating change in real-time cuts off many security incidents at the source.
- By focusing on maintaining integrity across an IT environment, organizations can drastically improve cybersecurity outcomes while cutting out 95% of change noise.
- When cybersecurity controls and policies are effective, well-enforced, verifiable, and regularly reported, demonstrating compliance ceases to be a drain on time and resources.

Security and Compliance: We're Going the Wrong Way

In 2011, the cybersecurity market was valued at around $60 billion[1] in annual spending. In 2021, it's expected to reach $150.4 billion.[2] That's a Compound Annual Growth Rate (CAGR) of 9.63% over a decade, and there's no sign of spending slowing down.

From 2020 to 2027/28, analysts expect the CAGR of global cybersecurity spending to continue at a rate of 9.4%[3], 10.9%[4], or 12.5%[5], depending on which source you trust.

With all that spending, you'd expect the rate of security incidents and data breaches to fall—but they haven't. The number of recorded breaches is rising year by year. The number of breached records hit a new high during Q1 2021[6], and nobody expects them to fall in the coming years.

When it comes to our ability to identify and contain breaches, there's more bad news.

Between 2015 – 2020, the Mean Time To Identify (MTTI) security breaches remained static at 206 days, while the Mean Time To Contain (MTTC) a breach rose from 69 days to 73 days. That makes the average time needed to identify and contain a security breach an incredible 279 days.[7]

What Can We Learn From This?

Despite a huge rise in cybersecurity spending, threat actors are getting better, faster than we are.

From this, we can deduce two lessons:

1. Today's approach to cybersecurity isn't working.
2. Organizations can't spend their way out of the problem.

And, perhaps the situation is even worse. Increasing cybersecurity budgets and spending creates a false sense of security that comes crashing down when an organization is inevitably breached.

You've probably heard the oft-repeated phrase, "it's not **if** but **when** your organization is breached." While it may seem self-serving for cybersecurity vendors to repeat this over-and-over, it's a truism—and the data above makes it abundantly clear.

[1] https://www.ifsecglobal.com/uncategorized/pwc-report-global-spending-on-cyber-security-to-hit-60-billion-by-year-end/

[2] https://www.gartner.com/en/newsroom/press-releases/2021-05-17-gartner-forecasts-worldwide-security-and-risk-management

[3] https://www.alliedmarketresearch.com/cyber-security-market

[4] https://www.grandviewresearch.com/industry-analysis/cyber-security-market

[5] https://www.globenewswire.com/news-release/2021/03/17/2194254/0/en/Global-Cybersecurity-Market-Size-to-Grow-at-a-CAGR-of-12-5-from-2021-to-2028.html

[6] https://www.itgovernance.co.uk/blog/data-breaches-and-cyber-attacks-quarterly-review-q1-2021

[7] https://www.ibm.com/uk-en/security/data-breach

The 'Fog of More'

We can all agree that no organization can do everything when it comes to cybersecurity. The available systems, controls, and processes are simply too expansive (and expensive) to even contemplate the idea. This leaves organizations trying to figure out which controls to implement with their limited human and budget resources.

This is where we run into a serious problem that most organizations haven't yet solved. Tony Sager, SVP and Chief Evangelist at The Center for Internet Security (CIS), explains[8]:

> *Defenders lose because they are overwhelmed. There's too much advice and too many consultants, tools, compliance requirements, and marketing messages to process. They don't know where to start, and that makes them susceptible to any message or tool that claims to solve their problems.*

With so much choice, many cybersecurity leaders (and their teams) become paralyzed. They do their best to prioritize budgets and energy, but the outcomes don't match their efforts.

Security is Stuck in 'Reactive Mode'

When you're at war, reacting to your enemy is the worst position to be in. However, that's how most cybersecurity teams are forced to operate.

Perimeter defense tools like firewalls and IDS/IPS tools do an essential but incomplete job. The predominant approach to cybersecurity relies heavily on reactive monitoring and incident response, hoping to head off each threat before it does serious harm.

Worse, many cybersecurity teams are over-reliant on individual 'security heroes' to fight threats. This is a poor use of resources, and it's also a dangerous and potentially costly position. Being reliant on individuals creates a huge weakness—what if that person isn't in the office or leaves the organization for a new opportunity?

The fact is that no cybersecurity team should be reliant on individuals—and everybody knows it. What they really need is the proper machinery in place to prevent threats at their source with only limited human involvement.

All of this brings us to an inevitable conclusion:

To reverse the current trends surrounding cybersecurity spending and outcomes, we need to re-evaluate the fundamental principles of cybersecurity.

> *Never permit your enemy to gain an advantage over you in any way. You can be sure your enemy is thinking likewise; either you lead the enemy, or he will lead you.*
>
> **— Miyamoto Musashi,**
> The Book of Five Rings

[8] https://www.youtube.com/watch?v=OZLO-xekp3o

What Can Cybersecurity Learn from IT?

Historically, there has been plenty of negativity and friction between IT operations and cybersecurity teams. However, as an industry, we need to accept a simple fact. IT departments have been around a lot longer than cybersecurity teams, and their processes are more mature.

Consider one of the most prominent frameworks for IT service management (ITSM): the IT Infrastructure Library (ITIL). Developed in the 1980s by the UK Government, ITIL has evolved into the most comprehensive set of IT practices ever devised. It's more widely[9] used than any other framework, and even Microsoft used it as the basis[10] for its Microsoft Operations Framework (MOF).

To see what cybersecurity teams can learn from it, consider the ITIL Life Cycle's five principles:

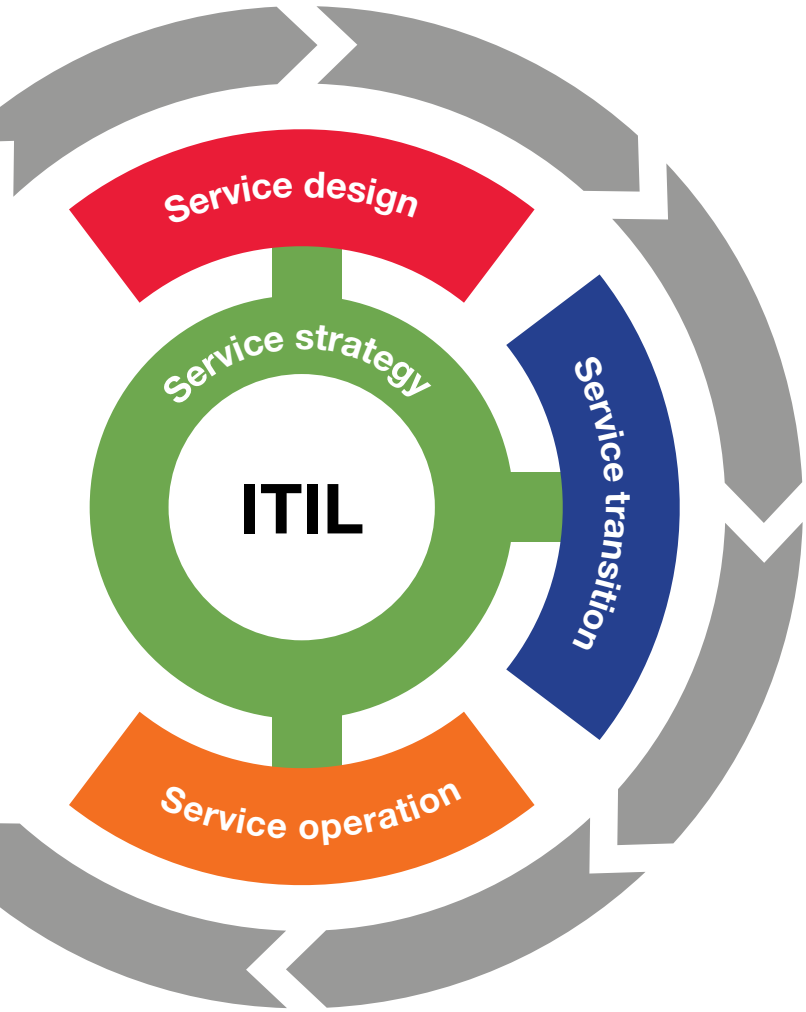

Service strategy defines the perspective, position, plans, and patterns a service provider needs to execute in order to meet business outcomes.

Service design includes the design of services, governing practices, processes, and policies required to realize the strategy.

Service transition ensures new, modified, or retired services meet business expectations.

Service operation coordinates actions and processes to deliver and manage services at agreed levels. It also manages the technology used to deliver and support services.

Continual service improvement ensures services are aligned with changing business needs by identifying and implementing IT service improvements.

Note[11]: These definitions have been slightly condensed for brevity. You can find full definitions in the ITIL® glossary and abbreviations, 2011. Source: AXELOS, ITIL v.3 (2011)

[9] https://www.peoplecert.org/itil-certification-family

[10] https://www.itilnews.com/index.php?pagename=ITIL_V3_and_Microsoft_Operational_Framework_4_MOF_4

[11] https://www.axelos.com/corporate/media/files/glossaries/itil_2011_glossary_gb-v1-0.pdf

Protecting Operations

Notice how ITIL doesn't focus on individual systems or processes but rather on meeting business expectations at a pre-agreed level. IT operations teams have known for years that downtime is inevitable, and all they can do is limit its length and frequency. This is the whole purpose of SLAs—to ensure downtime is kept to an acceptable minimum. This is vital, and it's in stark contrast to how the cybersecurity industry portrays its function.

ITIL emphasizes the importance of getting the basics right and having the systems and processes to achieve the most important objective: minimizing downtime.

Bringing this into the cybersecurity domain, we can assume that some level of 'failure' is inevitable. Almost all organizations will be breached at some point, so the important consideration is how to minimize their frequency, impact, and duration. Scott Alldridge, President at the IT Process Institute (ITPI) and MSSP IP Services, explains:

> *Use a scorched earth approach. Assume you've already been breached and need to recover. What recovery point are you comfortable with, and how long can it take? Once you have your answers, reverse engineer security controls from there just like an IT department would.*

The Importance of Change Management

To ensure business expectations are met, one of the most critical components of ITIL is change management, which is the core function of service transition. For many years, IT departments have understood the importance of change management to maintain SLAs at an acceptable level.

In *The Visible Ops Handbook*, the authors explain (emphasis ours):

> *High-performing IT organizations **eliminate change as a causal factor for an outage as early as possible in the repair cycle.** They identify the assets directly involved in the service outage and examine all changes made on those assets in the previous 72 hours. This information is [compared to] all authorized and scheduled changes. [...] When issues are escalated to problem managers, they have all relevant and causal evidence at hand and [...] can successfully diagnose issues without logging into any infrastructure over 50% of the time!*

This approach is directly applicable to cybersecurity. By setting objectives (service strategy), a baseline for acceptable service levels and activities (service design), and managing changes from that baseline (service transition), cybersecurity teams can achieve the same level of operational success (service operation) as IT departments. Think about it. When was the last time your organization's IT systems went offline for a non-security reason—and how long did it last?

The Importance of Change Management

Cybersecurity Fundamentals: Less is More

How can we apply the ITIL mindset to cybersecurity? The first thing we can do is eliminate complexity and focus on a small number of basic principles. It's telling that just a handful of software vendors dominate the ITSM market.

By contrast, the cybersecurity market has hundreds or thousands of software vendors competing for budget, all with different solutions to different perceived problems. Simply, cybersecurity teams face a huge challenge just to understand their options—let alone to make effective decisions.

In The Paradox of Choice, American psychologist Barry Schwartz argues that eliminating consumer choices can greatly reduce anxiety for shoppers. Bring that into the cybersecurity world, and you can add an extra dimension. Limiting choice for cybersecurity leaders doesn't just minimize anxiety—it also improves results, as measured by the maintenance of acceptable service levels.

The Compliance Problem

Of course, cybersecurity teams face a complicating factor. Unlike traditional IT departments, they are subject to a complicated web of cybersecurity frameworks and regulatory requirements that aim to ensure organizations implement appropriate security controls. These requirements all have slightly different recommendations and priorities, adding to the confusion. Tony Sager, SVP and Chief Evangelist at The Center for Internet Security (CIS), explains:

> *Compliance requirements are what I call cosmic frameworks. They proclaim 'thou shalt achieve this,' but aren't prescriptive about how to do that. It creates an industry of tea leaf readers trying to interpret requirements, which is great for job security but very poor for business outcomes.*

To put this another way, most frameworks take a descriptive approach—they tell organizations what to achieve, but not how to achieve it. Tony explains that this approach creates a 'special snowflake' approach that forces each organization to find its own solution to each requirement. This alone creates a huge amount of work for cybersecurity teams, reducing the resources available to protect against threats.

However, this approach is fundamentally flawed. While there are undeniable differences between organizations, most are more similar than they are different. Worse, the vague nature of requirements creates a 'Wild West' approach to cybersecurity, where thousands of vendors spring up to fill organizations' perceived security and compliance needs—which are often contrary to the simple objective of minimizing the frequency, severity, and duration of security breaches.

Start at the Beginning

While it's rarely discussed, key players in the cybersecurity industry fully understand the importance of solid fundamentaly—it also improves results, as measured by the maintenance of acceptable service levels.

Notice how some of the industry's most widely discussed solutions are considered less critical and even ***optional***. These include:

- » Anti-malware
- » EDR
- » HIPS
- » TDR
- » Vulnerability shielding
- » Behaviour monitoring

These solutions are so widely marketed that you'd think they are ***critical*** to securing sensitive assets—but this Gartner report indicates quite the contrary.

Risk-Based Hierarchy of Workload Protection Controls

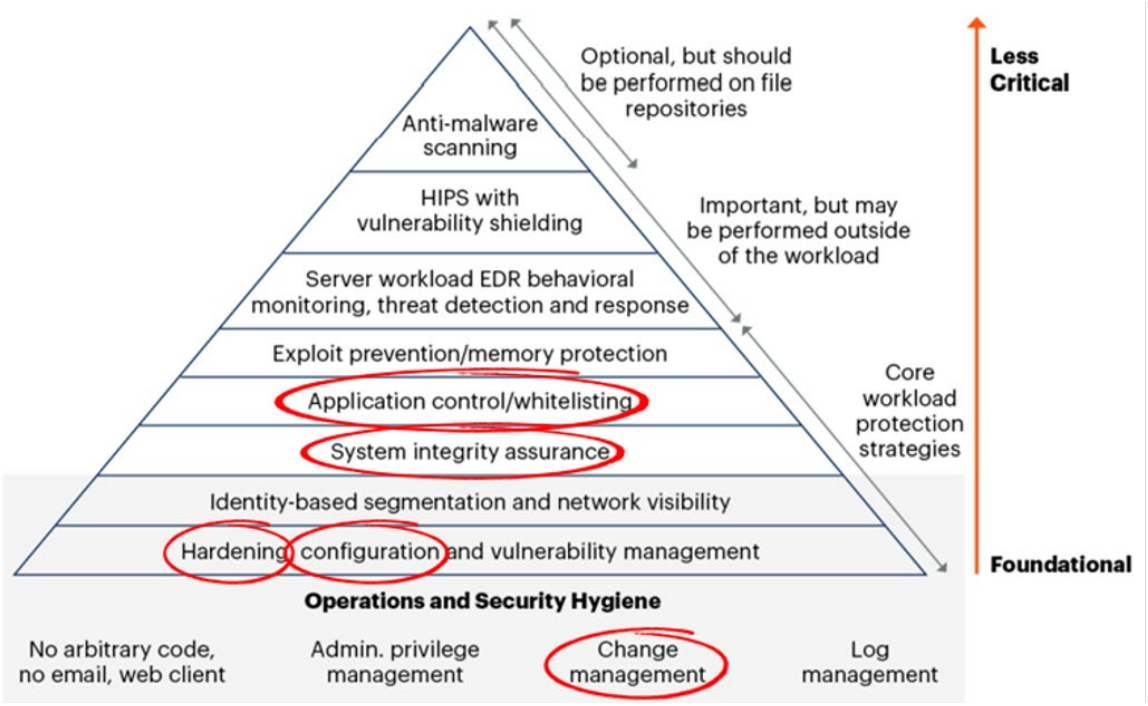

Meanwhile, more 'boring' controls like change management, system integrity assurance, application controls, segmentation, and configuration management are considered ***foundational*** and should be solidified before even considering controls further up the pyramid.

Going a stage deeper, are change management, system integrity assurance, application controls, hardening, and configuration management ***really even security controls?*** Or are they IT considerations? Either way, Gartner acknowledges them as some of the most critical requirements for a secure cloud environment.

Despite this, these controls are rarely discussed or publicized by vendors or analysts. As a result, they are rarely a cornerstone of an organizations' cybersecurity strategy. To find a genuine discussion (and recommendation) for such foundational cybersecurity controls, we have to turn to one of the industry's only ***prescriptive*** (i.e., tells you what to do, not simply what outcomes you need to achieve) frameworks: The CIS Controls.

Prescriptive Controls in Action

In 2001, the NSA released its security guides into the public domain, prompted by several high-profile breaches of its commercial partners. The boundaries between government agencies and their partners were disappearing, and there was a pressing need to help those partners ensure the security and integrity of government data.

There was a problem, though. The NSA guides were extremely thorough and provided far more guidance than a partner organization could implement in a short period. Further guidance was available from organizations like NIST, but this suffered from the same challenge—too many controls, too little time and money.

This resulted in a conversation at the NSA—led by lifelong NSA security expert Tony Sager—about producing a small, prioritized list of essential security controls. After many additions and several changes of ownership, this list became the CIS Controls, a list of just 18 best practice controls (referred to as Control Families).[12]

Unlike typical ***descriptive*** frameworks, the CIS Controls take a prescriptive approach, telling organizations exactly what to do to protect against pervasive cyber threats. To give an idea of their effectiveness, several independent studies of version 6.1 of the CIS Controls found that just the first five controls protected against 85% of all cyber attacks.[13]

Today, the first five controls are:

1. **Inventory and Control of Enterprise Assets**
2. **Inventory and Control of Software Assets**
3. **Data Protection**
4. **Secure Configuration of Enterprise Assets and Software**
5. **Account Management**

Do you see a correlation here to the ITIL controls discussed earlier? At a basic level, they require an organization to establish a continually updated baseline for hardware, software, data, user accounts, and asset configuration—and then track and remediate changes from that baseline.

Contributing to that baseline, CIS also maintains the CIS Benchmarks, a set of 140+ configuration guides to help organizations establish hardened systems to protect against evolving cyber threats. In line with ITIL's service design principle, the Benchmarks provide a baseline for the asset configuration. If the baseline remains current, it's easy to identify activity that isn't acceptable—i.e., unauthorized changes that negatively affect configuration—and block it at the source.

of all security incidents can be auto-detected with three detective controls…**configuration**, **change** and **release management**.

- IT Process Institute

[12] https://www.cisecurity.org/controls/cis-controls-list/

[13] https://static1.squarespace.com/static/555f9696e4b0767a7f0769b3/t/5c885173f4e1fcb114e1e2dd/1552437623967/The_First_Five_Guide_v1.1.pdf

Why Establish a Trusted Baseline?

> *When you're in an earthquake on a unicycle, juggling chainsaws, the only way to survive is to tack down everything you can tack down, so you can deal with what you can't.*
>
> **— Stephen Chakwin**

One of the strangest things about cybersecurity compared to other disciplines is the focus on finding bad things and preventing them. Think about how you would manage physical security for a building, e.g., a government office. How would you stop the wrong people from getting in?

Most likely, you wouldn't try to track ***every single person*** who isn't supposed to be in the building. That would quickly exhaust your resources and achieve essentially nothing. Instead, you'd build and maintain a list (baseline) of everybody who ***should*** be there and use a control system (probably ID cards and security guards) to ensure ***only*** those people are allowed in.

Of course, this system isn't perfect. Sometimes, someone who was supposed to have access isn't allowed in. This is easy to manage. The blocked individual simply tells the guard why they should be allowed in, and it's quickly verified (or not). This process is called 'managing by exception.'

Alternatively, some people will try to force their way in. Again, this is easy to manage by exception. The security guard will see the problem and apprehend them.

This approach runs contrary to most public discussions of cybersecurity principles.

Most cybersecurity controls use blacklists to try to identify all possible 'bad things' and prevent them. Instead of maintaining a small database of things that are allowed, cybersecurity teams maintain a monstrous database of things that ***aren't*** allowed and constantly monitor for them.

This approach is reactive, slow, and misses threats simply because they haven't been seen before.

Imagine how life would be for cybersecurity teams if we followed in the footsteps of traditional IT operations and service management. Consider this ITIL-inspired, basic approach to cybersecurity:

- **Service strategy:** Determine objectives for the security function
- **Service design:** Set a trusted, authoritative baseline of what you have (software, hardware, services, etc.) and what is allowed to ***be*** and what ***happen*** in your environment.
- **Service transition:** Enforce the baseline by monitoring changes in the environment and blocking anything that isn't explicitly allowed.
- **Service operation:** Carry out normal security operations to identify any threats or issues to make it past baseline enforcement.
- **Continual service improvement:** Learn from mistakes and make changes to the baseline.

As we'll see shortly, this approach is very achievable—and with far better results than most cybersecurity teams have come to expect.

Why is Change Important in Cybersecurity?

Once you have a trusted, authoritative baseline, you have a place to start from.

However, there's an obvious argument against the system described above. Even if your baseline is set to a mythical 'perfectly secure state,' one change could create a huge weakness.

Change is the nemesis of IT and cybersecurity professionals who need to maintain a secure and available environment. Unauthorized, unexpected, and unwanted changes to critical files, systems, and devices can quickly open a gaping hole in an organization's cybersecurity posture. At that point, it doesn't matter how good the rest of its controls are—a breach may be imminent.

What is a Change?

Everything that happens in an IT environment (good or bad) starts with a change: a file, configuration setting, or device is altered, deleted, added to, or even just read by a user or service.

Every bad thing in an organization's environment begins with change… but so does every good thing. The challenge lies in determining the difference between good and bad. This is also where our baseline comes into play. Anything not included in the baseline can be assumed bad until proven otherwise. For each change, an organization can follow a simple process:

1. Determine precisely what changed in the environment.

2. Check whether the change is authorized under the baseline.

3. Allow, block, or roll back the change as appropriate.

> *"You need to know what you have and what changes are acceptable. Then you have to stop everything else and manage by exception where necessary. It's not necessarily easy because there's lots of change. If you have the machinery to control this, you have the basis of integrity."*
>
> **— Tony Sager,**
> SVP and Chief Evangelist at The Center for Internet Security (CIS)

As of March 2020, the total number of new malware detections worldwide amounted to **677.66 million programs**, up from 661 million new malware detections at the end of January 2020. These malicious programs intend to add, modify or delete files, which can be mitigated through a closed-loop change control process.

Why is Change So Important?

As we've seen, all bad things in an IT environment begin with change. This fact is clarified by IDC research, which found that a huge proportion of IT outages are caused by human error, including failure to conform to change management processes. In other words, failure to properly manage change in an IT environment is among the largest causes of unplanned downtime.

Operations Errors
40%

People and Process

- Hiring, Training, Procedures
- IT Process Maturity
- Automation & Ops Arch
- Change & Problem Mgmt.
- Integration and ProdIT-DRM Modernization Validation Testing

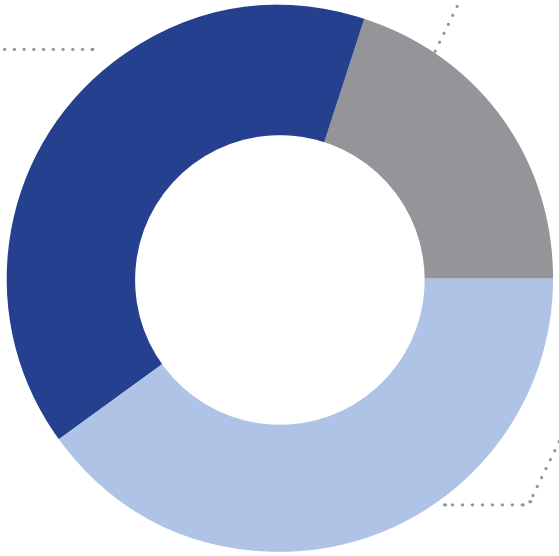

Environmental Factors, HW, OS, Power, Disasters
20%

Externals

- Redundancy
- Service Contracts
- Proactive Monitoring
- Business Support

Application Failure
40%

People and Process

- App. Architecture/Design
- SDLC Enhancements
- Change & Problem Mgmt.
- Configuration Management
- Performance/Capacity Planning

Model is based on Gartner IDC Study on Causes of Network Downtime

Donna Scott, VP and Research Director at Gartner, goes a step further by stating that:

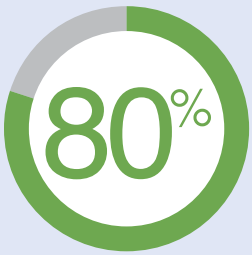

> 80% of unplanned downtime is caused by people and process issues, including poor change management practices, while the remainder is caused by technology failures and disasters.

Based on their experience working with hundreds of IT organizations, the authors of The Visible Ops Handbook further note that even once an incident has occurred, 80% of Mean Time To Recovery (MTTR) is wasted on non-productive activities. Most notably, determining which change is responsible for the outage.[14]

[14] https://itpi.org/the-visible-ops-book-series/visible-ops-handbook-review/

Managing Change as a Cybersecurity Function

A study commissioned by the U.S. Department of Defense (DoD) determined that:

> *Information security hinges on the effectiveness of the change management process. As a result, we need to implement a detective control to verify compliance [with an authoritative baseline] and take decisive action when the process is not followed.*

Source: File integrity monitoring tools: Issues, challenges, and solutions, Applied Research Center, Florida International University, 2020[15]

Notice the wording. Change management isn't just important—it's the lynchpin of the entire information security function. With all this in mind, why hasn't there been an attempt by cybersecurity vendors to address change management?

As it turns out, there has: File Integrity Monitoring (FIM) tools.

> *All security start with a change or a need for change. For this reason, change control becomes the ultimate security backstop regardless if it's on-prem, VM's or in the cloud.*

[15] https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.5825

Why FIM Hasn't Solved Cybersecurity Problems

If change is so important, how do you monitor change in an IT environment?

Simple: use a monitoring tool that tells you every time something changes. This is what FIM tools were initially designed to do—detect changes in all files across an IT environment and alert the cybersecurity team. This approach is approved by the DoD study mentioned above.

Note the use of integrity. If a file has integrity, it is in precisely the right state and only affected by permitted changes. The state of a file is determined using a cryptographic checksum—known as a file hash—and other checks such as file size, version, modified by, creation date, modified date, cache operations, and configuration values.

The study goes on to note that a FIM tool "*[...] compares and verifies the current state and baseline of files [...] to detect unauthorized file operations in a system.*"

If FIM tools did what they were supposed to, they would help cybersecurity teams identify and prevent most attacks—at least those that rely on file changes or access. But, as most security professionals already know, FIM tools don't do what they are meant to do. Here's why:

> *Ensuring integrity of sensitive files in file systems is imperative to computer systems. The vast majority of attacks work through unapproved or unauthorized access to sensitive files to take secret data like secret keys, passwords, credit card numbers, and so on. After that, attackers generally conceal their traces by subverting critical files like system logs.*
>
> Source: File integrity monitoring tools: Issues, challenges, and solutions[15], Applied Research Center, Florida International University, 2020

Problem #1: Noise

A typical FIM tool simply monitors files for change and produces alerts—lots of alerts. They produce so many alerts they have become 'shelfware' for most cybersecurity teams. They are theoretically valuable but useless in the real world because they produce too many alerts to manage with no context or verification.

Noise is a ubiquitous issue across many cybersecurity tools. Security Operations Centers (SOC) and Incident Response (IR) teams are already buried under more alerts than they can manage, so they simply shut off or ignore alerts from their FIM tool. They keep the tool for compliance purposes, of course—they just don't use it.

[15] https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.5825

Problem #2: FIM Isn't FIM

At the heart of this problem lies a simple fact:

FIM tools don't provide FIM (File Integrity Monitoring) at all. They provide basic File Monitoring or in many cases just the simple action of detecting change.

The following graphic shows the difference between File Monitoring and System Integrity Assurance.

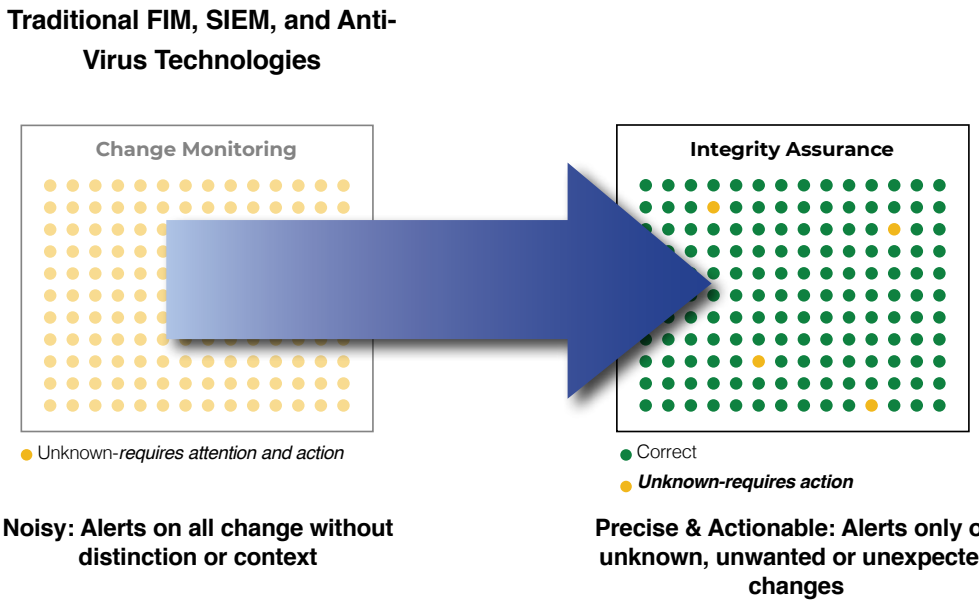

On the left, we see what typical 'FIM' tools provide: a massive list of changes without any context or distinction. This list is too large to triage, so cybersecurity teams ignore these change alerts.

On the right, we see what a FIM tool should provide: a small list of unverified changes to be signed off, prevented, or rolled back. This list is easy for cybersecurity teams to process and helps to maintain the integrity of files or directories. To use our ITIL-inspired objectives, it also helps to minimize the frequency, severity, and duration of incidents and breaches.

Imagine the difference between File Monitoring and FIM after a big update, e.g., Patch Tuesday.

There could be hundreds of thousands of changes in the IT environment, and a File Monitoring tool would create an alert for every single one. There could be a hundred unauthorized, dangerous changes in that list of alerts, but nobody would know because they don't have time to check. At best, the tool might integrate with some blacklist resources to identify changes known to be malicious. Still, though, the list of unverified changes is far too large to manage.

On the other hand, a genuine FIM tool (i.e. a System Integrity Assurance platform) can identify every change that is allowed, including those made by vendor-verified patches, and exclude those from its alerts but still securely stored for audit evidence. By highlighting only changes that aren't explicitly allowed, FIM tools could enable cybersecurity teams to manage by exception—and FIM then becomes a cornerstone to what its initial intentions and objectives were designed for.

Problem #3: Too Resource-Intensive

Most FIM tools identify change by completing daily polling scans of all files in an IT environment. This process is hugely resource-intensive, so it usually happens overnight. While it would be more valuable to scan the environment continuously, this is simply impossible, as it would interfere with other IT operations.

Bringing Integrity to Your Environment (Not Just Files)

Integrity is the accuracy and completeness of data throughout its entire life cycle. That means no matter what service, device, or user accesses, stores, processes, transmits, or receives data, it remains accurate and complete. For this to be possible, four things are needed:

1. **An authoritative baseline of what data should look like.**
2. **A way to identify and protect data from unauthorized change.**
3. **A way to roll back unauthorized changes not blocked at the source.**
4. **A way to verify that controls 1 – 3 are in place and working correctly.**

Notice we're talking about ***data***, not just files. To have integrity, you need to protect all of the data in your environment—including data held in configuration files, network devices, endpoints, directory services, cloud instances, and more. We'll cover this in more detail in the next section.

Other Tools that provide 'FIM'

Many cybersecurity solutions like AV and SIEM tools claim to provide FIM. However, these tools suffer from the same problems—they provide change alerts without context or verification. Once again, this is just FM or Simple Change Monitoring posing as FIM.

> *Integrity means that data is protected from unauthorized changes to ensure that it is reliable and correct.*
>
> **— Mike Chapple,**
> Professor of IT, Analytics & Operations, University of Notre Dame

Working From a Trusted Baseline

System integrity assurance works under the same principle as physical security. It establishes a known, trusted, and authoritative baseline of what is allowed and then prevents, limits, or rolls back everything else. Whenever an unknown change occurs, it's managed by exception so that acceptable changes are added to the baseline while unacceptable changes are prevented.

Closed-Loop Integrity Assurance can be demonstrated to work as follows in the real world.

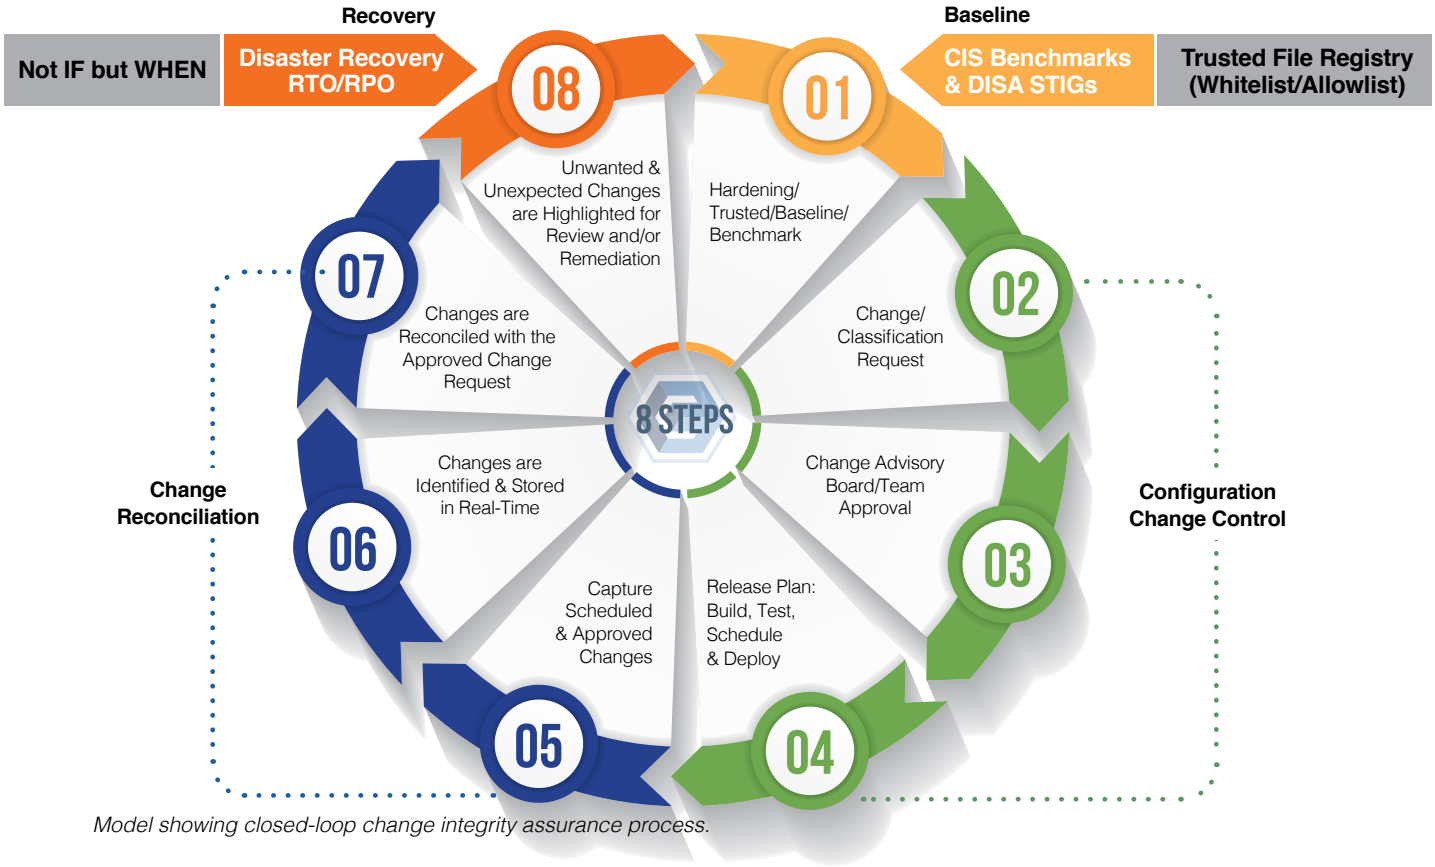

Model showing closed-loop change integrity assurance process.

This is a closed loop process for managing changes from a trusted baseline. Similar to the change management procedures laid out by ITIL, the loop covers all stages needed to ensure only acceptable changes are allowed to proceed, while others are prevented or rolled back.

There's an obvious elephant in the room: **This looks way too time-consuming.** What cybersecurity (or IT) team has time to run through this entire process for every single change?

So long as there is a trusted and authoritative baseline in place that includes everything that should ***be*** and ***happen*** in the environment, the loop only needs to occur for unknown changes that need to be verified. I.e., changes that aren't known for sure to be good or bad.

Further, the majority of this loop can be automated. With the right technology—one that continually updates the baseline to reflect changes known to be good or bad—human intervention is only needed for a small percentage of changes. In the next section, we'll see FIM in action, including how it can suppress traditional 'change noise' by up to 95%. First, there's another elephant in the room.

Why is Nobody Talking About Integrity?

Beyond 'lip service,' almost nobody in cybersecurity talks about integrity, least of all vendors. Some compliance frameworks include integrity as a requirement but include no guidance on how to achieve it—and most of the time, genuine integrity isn't required to pass a compliance audit.

The obvious reason is that **integrity is boring**. It's more fun to focus on the latest shiny tools than to steadfastly stick to the fundamentals. However, this theory puts the blame on the shoulders of cybersecurity teams and their leaders—when in fact, it should be somewhere else.

Tony Sager, SVP and Chief Evangelist at The Center for Internet Security (CIS), explains it like this:

> *Cybersecurity has been treated like wizardry. If you treat it like wizardry, the only defense is more wizardry. You need flashy tools and insight into what some hacker is doing in another country. Honestly, most of this stuff is overblown in terms of its real value. Wizardry is great for job security but bad for corporate success. You can't have a program based on wizardry. You need to have discipline and management and repeatability and data and science behind it.*

Dr. Ian Levy, chief technical director of GCHQ's National Cyber Security Centre, took a more direct approach to the problem during a 2019 talk:

> *We are allowing massively incentivized companies to define the public perception of the problem. If you call it an advanced persistent threat, you end up with a narrative that basically says, 'you lot are too stupid to understand this, and only I can possibly help you. Buy my magic amulet, and you'll be fine.' It's medieval witchcraft. It's genuinely medieval witchcraft.*

The System Integrity Assurance Platform(s)

A system integrity assurance platform enforces the Integrity Assurance Loop explained in the previous section. It enforces a trusted baseline across an entire IT environment to allow expected, legitimate changes to go ahead, block or roll back changes known to be dangerous, and alert on unexpected changes that aren't known to be good or bad. A system integrity assurance platform is very different than a simple change monitoring/detection tool. A system integrity assurance platform focuses on encompassing the entire workflow while interfacing with a variety of external tools, in order to achieve compliance.

To achieve this, system integrity assurance platforms must rely on three critical components:

1. **Maintain and secure a complete inventory and register of all critical files** throughout the network. This includes those held by hardware and software assets, along with their correct states, configurations, and settings.
2. **Access to whitelist/allowlist database of known and trusted file hashes** containing metadata, and configuration settings to validate and verify the integrity and authenticity of data, no matter where it is.

Remember our physical security analogy? The best way to ensure bad things don't happen in an IT environment is to *only* allow good assets and changes.

However, there is still a benefit to the traditional approach of identifying and categorizing malicious files and activity: it reduces noise. If you have a current list of files and activities known to be malicious, there's no need for a human to manually investigate them when they turn up in an IT environment. This is why integrity assurance platforms employ a fifth component:

3. **A blacklist/denied list** of known bad file hashes from external intelligence from threat feeds, file reputation services, and malware data repositories. Ideally, an integrity assurance platform would be able to alert if any blacklisted/denied list files were ever resident on any device or system during it entire operating life-cycle.

What's in a whitelist/allowlist?

A whitelist/allowlist must include all known and trusted file hashes to validate and verify the authenticity and integrity of individual files, configuration settings, etc. This requires a massive database constantly updated with the latest OS updates, software patches, etc., in all languages, customized by country. This amounts to a list of billions of hashes, and the list must grow in real-time to keep change noise to an absolute minimum.

By including these three components, system integrity assurance platforms can further enforce integrity across the IT environment while suppressing traditional change noise by 95%. Instead of triaging a torrent of unknown changes (which no one can do), cybersecurity teams only triage a small number of potentially harmful changes. Once these changes are categorized, the baseline is automatically updated so that the platform can handle future changes of the same type without human intervention.

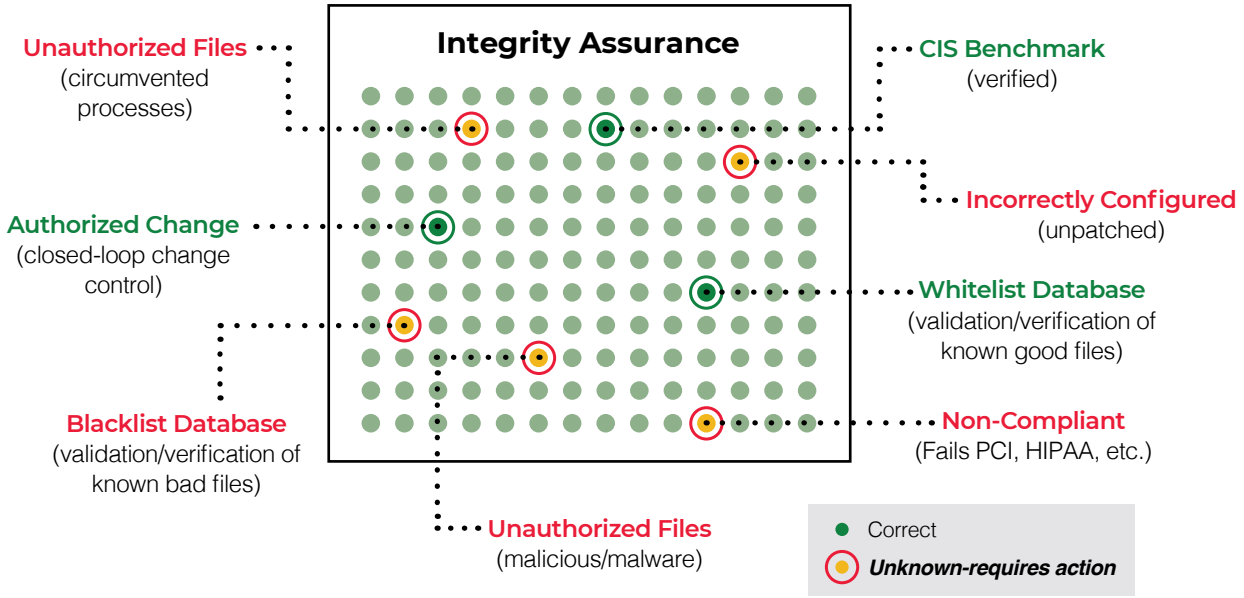

This brings us to the ultimate question:

What happens if you know every time something changes, and you stop anything that isn't authorized… and you expand this capability across all asset classes?

- Ransomware and other malware can't run in the environment.
- Attackers can't traverse the network or exfiltrate data.
- Nobody can add, modify or delete files or configurations to make them non-compliant or introduce new risks or vulnerabilities.
- Users can't accidentally run malicious attachments.
- Nobody (even privileged administrators) can alter critical system files.
- Mitigates software supply chain security issues and risks.

This approach doesn't solve cybersecurity entirely. The field is too large and complex for that to be possible. But it does take away a massive proportion of the risk and threats that can arise in an IT environment with minimal human involvement.

It's Not Just About Files (or Monitoring)

As we have already alluded, FIM is crucial, but it's not enough to ensure integrity across an IT environment. File changes are important, but what's vital is changes to data—wherever it exists. This covers a wide range of hardware and software assets, including:

- Files
- Configuration
- Databases
- Hypervisors
- Network devices (e.g., firewalls, switches)
- Metadata (e.g., of databases)
- Directory services
- Endpoints
- Cloud instances and containers

At the same time, integrity is about far more than monitoring change. A system integrity assurance platform must include ten critical capabilities to enforce integrity across an IT environment:

CIS or DISA STIG benchmark support and integration.

Real-Time change monitoring and detection to identify all changes within the environment.

Collection and storage of forensic evidence and detail for every change, including the source IP, user, time, and process.

Reconciliation and curation between observed changes against authorized/approved changes.

Categorization (i.e. whitelist/allowlist and black list/deny list) of changes as good, bad, or unknown.

Alerting for unknown changes that require human intervention.

Prevention of disallowed changes to sensitive assets.*

Roll back and remediation (A.K.A. 'self-healing' or resiliency) of disallowed changes to other asset groups.*

Baseline updates to include new file hashes and configurations categorized as good.

Embedded ticketing functionality to enable workflow automation and control or integration with traditional ITSM tools.

What's in a trusted, authoritative baseline?

A trusted baseline includes all of the assets, file hashes, configuration settings, etc., allowed to exist in an environment. In addition to information determined by the organization, an integrity assurance platform leverages best practices from authoritative sources like CIS Benchmarks and DISA STIGs to establish a known and trusted baseline that can be restored at any point in time.

*Some assets (e.g., critical system files) should never be changed, so these changes are blocked. Other changes are categorized as bad after the event and immediately rolled back to a trusted state.

A system integrity assurance platform completes these actions automatically and in real-time. A human only gets involved with unknown and unexpected changes to decide if they are acceptable. Even then, the platform should note the decision to deal with similar changes automatically in the future.

This provides several key benefits:

- Blocks most threats at their source.
- Instantly 'heals' files and settings to their trusted state.
- Provides deep insight into the state of any asset or system.
- Decreases incident response time by providing thorough forensic evidence.
- Reduces remediation time and costs.
- Drastically improves compliance audit (preparation and inspection).
- Decreases MTTI and MTTC to just seconds…rather than months.

A system integrity assurance platform drastically improves cybersecurity posture and outcomes without the need for flashy gimmicks, advanced CTI, or Artificial Intelligence. It might be boring, but it works.

> *Based on our extensive testing across the DoD, Intelligence Community, and rest of government, it was clear to me that secure configuration management is a foundational, must-do element of any successful security management program.*
>
> **— Tony Sager,**
> Senior VP and Chief Evangelist at the Center for Internet Security (CIS)

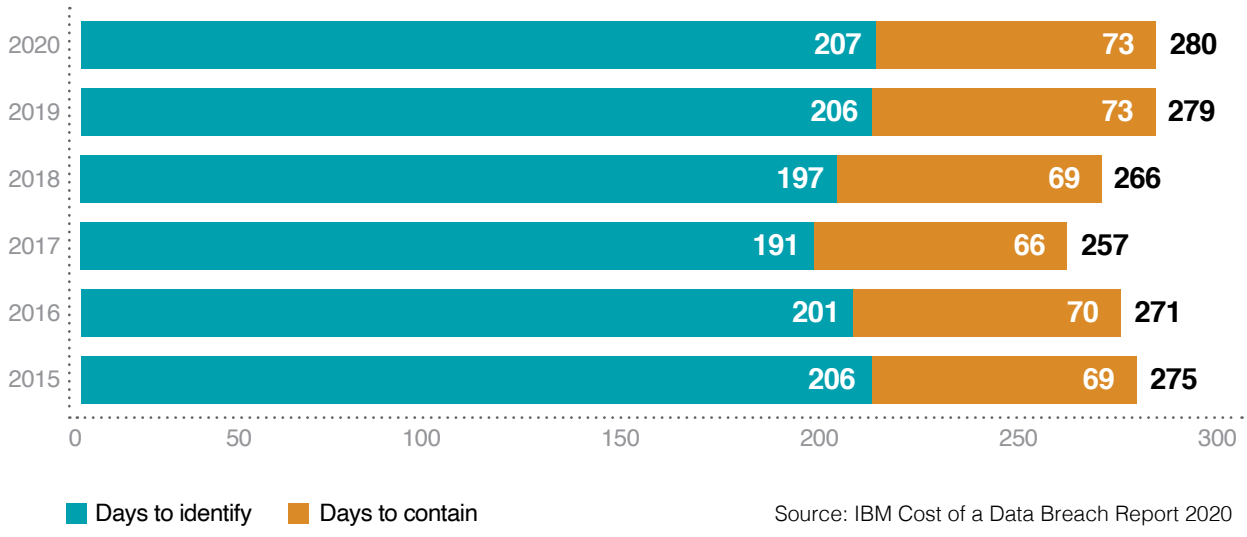

Source: IBM Cost of a Data Breach Report 2020

Why Verify Integrity in Real-Time?

There's little point in reaching a trusted baseline if you can't maintain it.

A major failing of traditional FIM tools is their lack of real-time monitoring. Most FIM tools run daily polling scans that drain system resources and fail to identify any harmful changes in between. This renders the organization unable to respond to attacks (and even mistakes) for as much as 24 hours, giving attackers ample time to cause damage, traverse the network, or steal sensitive data.

Even then, a typical FIM tool only provides monitoring, leaving the organization to identify malicious changes and complete remediation efforts.

Real-Time Verification Prevent Breaches

Only by monitoring change in real-time can an organization respond instantly to unexpected and unwanted changes. This is the only way to proactively prevent cyberattacks at their source without restricting operations to reactive threat feeds.

There's nothing wrong with ***using*** threat feeds, of course—but they should be a supplement to integrity assurance, not a replacement.

A major benefit of integrity assurance platforms is the ability to 'self-heal' files and settings to a trusted state. For example, if a server configuration setting is changed in a way that makes it non-compliant with the appropriate CIS Benchmark or DISA STIG, the integrity assurance platform can instantly reverse the change before it causes harm.

For sensitive files and assets, the platform takes this a stage further by enabling cybersecurity teams to block all changes at the source. Now, even a privileged administrator will be unable to make changes unless the block is lifted.

All of this is only possible with an integrity assurance platform.

Improved Performance

We've established that the polling scans conducted by typical FIM tools are resource-intensive. Doesn't that mean real-time monitoring should be even more resource-intensive… *all the time?*

No. System integrity assurance platforms don't scan the environment constantly. They scan it once to establish a baseline, then receive change data from agents and modules across the environment, often in real-time. If an asset or file doesn't match the baseline, the tool knows a change has occurred. This process is highly efficient and barely registers on the resource monitor.

Note: Agents and modules harvest data at the kernel level with higher privileges than the pure user-mode only solutions employed by other FIM tools. This enables system integrity assurance platforms to gather deeper forensic evidence, adding more value to change management and incident response.

Mastering Compliance (Without Wasting Resources)

Perhaps the biggest challenge organizations face is maintaining compliance with regulatory and partner requirements. They manage it briefly for their annual audits, but maintaining compliance year-round seems out of the question. This is another area where cybersecurity can learn from IT.

In *The Visible Ops Handbook*, the authors note that high-performing IT organizations have a trusting relationship between operations and auditors. Controls and policies are effective, well-enforced, verifiable, and regularly reported on. As a result, they spend very little time on compliance activities and audit preparation—and they also have fewer audit findings and repeat findings.

When you have outstanding systems and processes—including automatic audit trail capture—the proof element of compliance becomes easy to manage.

> *Compiling evidence for compliance isn't a good use of time. You only do it because the regulator says you have to. Most organizations have lots of regulators to satisfy, so it becomes a very repetitive and painful process. If you have good machinery and operations, it provides almost all the proof they need, and those resources are freed up for more useful activities.*
>
> **— Tony Sager,**
> Senior VP and Chief Evangelist at the Center for Internet Security (CIS)

System Integrity Assurance for Compliance

A system integrity assurance platform automates the process of achieving and maintaining compliance with frameworks like PCI-DSS, HIPAA, NIST 800-171, CMMC, and many more. It does this by:

1. **Building the requirements of all applicable frameworks into the trusted baseline.**
2. **Continually monitoring all files and configurations against the baseline.**
3. **Raising an alert when it finds an issue or misconfiguration and providing clear evidence and guidance on how to resolve it.**

Armed with this information, it's easy for cybersecurity teams, asset owners, and IT operations teams to quickly identify and resolve any issues that bring the organization out of compliance.

Critically, this mostly automated process provides the monitoring, enforcement, and audit trail needed to demonstrate compliance at any time—not just during an audit. This drastically reduces the amount of time and resources spent on compliance activities, freeing them up for more valuable, security-oriented functions.

Note: This is another area where roll back or 'self-healing' capabilities come to the fore. Most of the time, changes that lead to non-compliance are unintentional, and the files or configurations involved shouldn't be changed. An integrity assurance platform can automatically roll back these changes, ensuring ongoing compliance and reducing the compliance workload.

How Compliance Frameworks Map to Integrity Assurance

Integrity assurance can help any organization reach and maintain compliance with any framework. All that's required is for the cybersecurity team to update its trusted baseline to include all relevant compliance requirements and then action any alerts the system raises.

To give an idea of how valuable integrity assurance can be to a compliance program, the image below shows how it maps to eleven common frameworks.

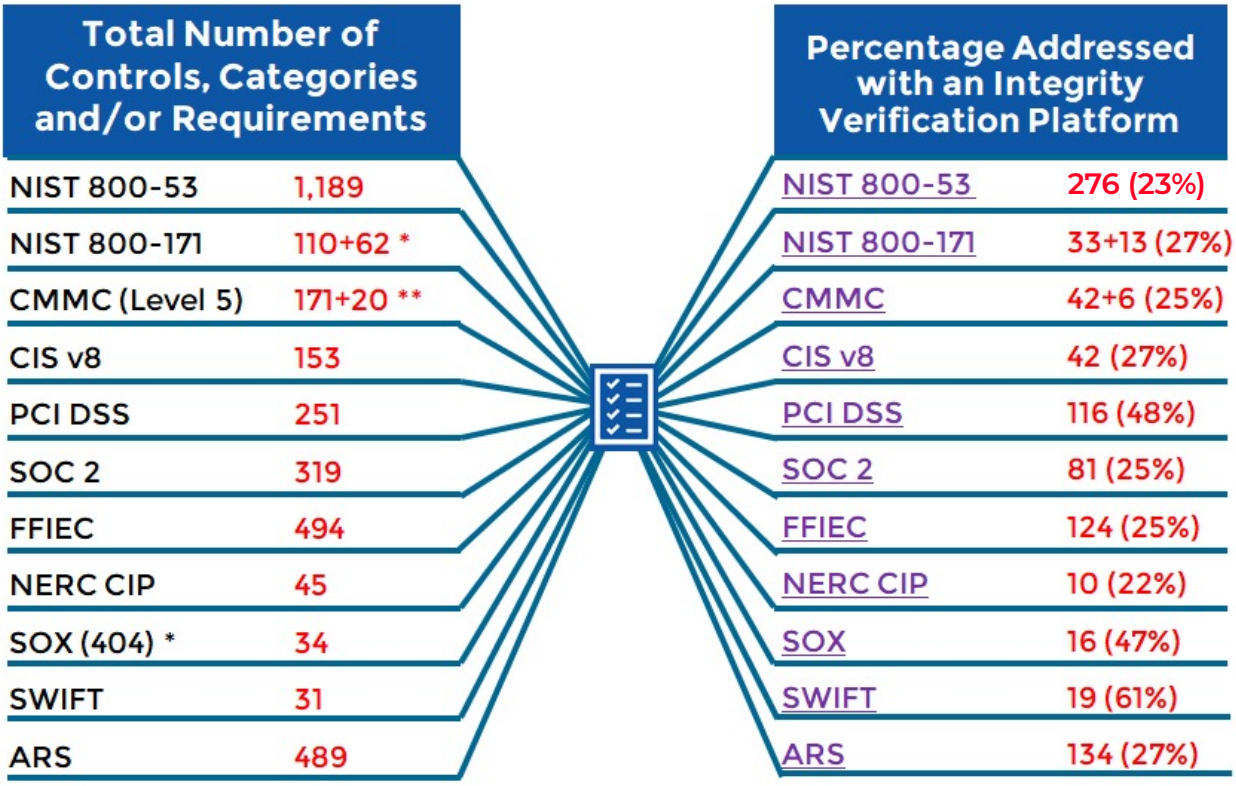

Integrity Assurance and PCI-DSS

PCI-DSS includes two sections that require a change detection capability: 10.5.5 and 11.5. An integrity assurance tool comfortably satisfies these requirements—but it goes much further. An integrity assurance tool covers 116 (48%) of the framework's controls, drastically reducing the amount of resources needed to reach and maintain compliance.

How Does System Integrity Assurance Align With NIST?

The National Institute of Standards and Technology (NIST) develops cybersecurity standards, guidelines, best practices, and other resources to meet the needs of U.S. industry, federal agencies, and the broader public.

One of those best practices frameworks is the Cybersecurity Framework (CSF) which helps organizations understand and address risks with a common approach and language to improving critical infrastructure cybersecurity. The Framework Core is a set of cybersecurity activities and principals aligned with a desired outcome common to all critical infrastructures and verticals industries. The Framework Core consists of five functional areas that considers a lifecycle approach to an organization's management of cybersecurity risk—Identify, Protect, Detect, Respond, Recover.

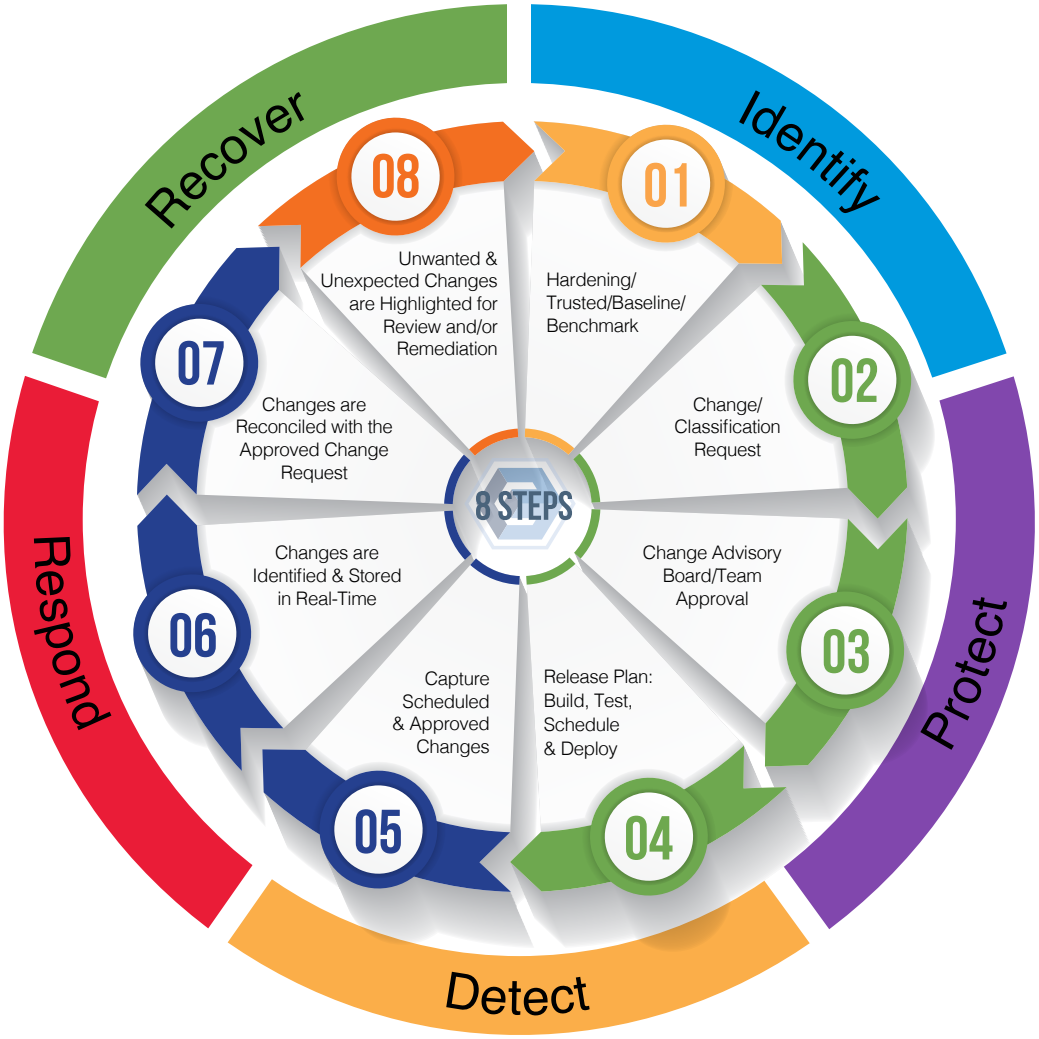

As you can see, NIST's five core functional areas align directly with that of the eight steps of a closed-loop integrity assurance model. This alignment provides for a consistent and uniform strategy when implementing an integrity strategy that incorporates not just one or two of the NIST functional areas but all five.

Zero Trust IS System Integrity Assurance

Zero Trust is a strategic initiative that helps identify and prevent successful data breaches by eliminating the concept of trust within an organization's network architecture and replacing it with integrity assurance. This is further supported by SANS CIA which includes three key security principals that every system should adopt as a standard and best practices—Confidentiality, Integrity and Availability.

There are three stages of a Zero Trust security model—Assessment, Control, and Recovery operations. The premise for a Zero Trust solution requires an approach to never trust but to always verify. This means that every user, device, application, workload and data flow should be treated as untrusted. Zero Trust is fundamentally a shift in how we approach security. To date, we've tried to identify malicious activity through a methodology of searching for the bad as opposed to managing form a known good and verified state of operation. When changes to a system happen, they ALL must be considered untrustworthy until a workflow process validates and verifies its integrity of those changes by determining if they were approved and authorized by an authoritative person or board. Only until this happens will the concept of Zero Trust become a reality.

Conclusion

It All Comes Down to This

In exercise science, there's a common expression:

> *Complicate to profit, simplify for results.*

The implication being that if you want to lose weight or get fitter, you don't need a Navy SEAL inspired workout or an expensive exercise machine. Instead, you should focus on the basics: eating better and getting some exercise. Most people intuitively know how to do these things—but there's no money in teaching people how to do the basics.

Parallels to the cybersecurity industry abound.

With literally thousands of vendors selling a cacophony of products, services, solutions, and advice, it's hard for cybersecurity teams to know where to focus their efforts and resources.

This white paper has made a case for focusing not on flashy toys, but on getting the basics right. It's not the most exciting approach to cybersecurity, but it's far more effective than chasing the latest trends and threats.

System integrity assurance is a path to create and maintain an IT environment that is resilient to dangerous change—both accidental and malicious. As mentioned multiple times throughout this paper, the end goal is the same as it is for IT operations:

To maintain a secure, available IT environment that supports business objectives.

If you have questions about system integrity assurance—or anything else related to this white paper—you can contact us at info@cimcor.com

The System Integrity Assurance Platform Selection Checklist

System Components:

- Asset discovery capabilities to identify, inventory, and collect information about all physical assets connected to the network. Includes routers, switches, servers, hosts, and firewalls.
- A comprehensive register and storage of all critical files in the environment, including those held in hardware and software assets, cloud instances, containers, etc.
- Built-in ticketing system, or at least two-way integration with an existing ticketing system.
- Compliance assessment, reporting, and remediation guidance.
- Integration with tools (e.g., SIEMs) to provide the enhanced analysis, correlation, and forensic data needed to mitigate attacks and detect anomalies.
- Ability to query both whitelist and blacklist databases to validate and verify file trust and authenticity.
- Comprehensive dashboards and reporting for all security and compliance needs.
- A full change management workflow that covers identification, investigation, triage, assigning tasks to engineers, final remediation, and confirmation.

Note: An unalterable audit trail avoids the danger of an administrator disabling the solution or monitoring of specific files/configurations.

Security Requirements:

- Encrypted communications between system components.
- Encrypted and compressed storage of file hashes and settings.
- Encrypted audit logs that are unalterable, even by system administrators.
- Monitoring of actions taken by solution administrators and users.

Questions to ask:

- Is the solution capable of true real-time change detection?
- Does it provide all the functions needed for integrity assurance (or is just file monitoring)?
- Is it easy to install, configure, and use?
- Can it be set up to meet your needs? (e.g., agent or agentless, on-premise or virtualized)
- Does it collect critical change information such as the user, process, and originating IP?
- Can it show precisely how a file changed with a side-by-side comparison to the original file?
- Does it integrate with other security solutions such as SIEMs?
- What inherent security does the solution have?
- Does it require costly training or professional services to implement and maintain?
- Is it scalable to meet your integrity assurance needs?

Bring Integrity to Your Environment with CimTrak

CimTrak is the industry's only genuine system integrity assurance platform. It combines all the capabilities required for "real" FIM, plus everything else discussed in this paper.

That includes:

- **CIS of DISA STIG benchmark** support and integration.
- **Real-time change monitoring and detection** across the entire IT environment.
- **Collection of forensic details** and evidence for every change.
- **Reconciliation and curation** between observed changes against authorized and approved changes.
- **Alerting** for unknown changes that require human intervention.
- **Comprehensive whitelist/allowlist and blacklist/deny list** to categorize changes and reduce noise by 95%.
- **Prevention, roll back, and remediation** of disallowed changes.
- **Automatic baseline updates** to include accepted changes to hashes and configurations.
- **Embedded Ticketing** or support ITSM integrations to assign and track authorized work orders and remediation if necessary.
- **Seamless integration with leading SIEMs**, helpdesk, incident, and ticketing systems.
- **Full encryption** of all communications, data, and audit logs.
- **24/7/365 compliance** enforcement, benchmarks, and reporting.

CimTrak also uses the **Trusted File Registry™** to identify all changes caused by vendor-verified patches and updates. This enables the tool to automatically categorize hundreds of thousands of changes as good, ensuring analysts remain free to focus on changes that pose a real danger.

To see what CimTrak can do for your organization, arrange a free demo today.

CIMTRAK

Supported Platforms

CimTrak for Servers, Critical Workstations & POS Systems

WINDOWS: XP, Vista, 7, 8, 10, Embedded for Point of Service (WEPOS), POSReady, Windows 10 IoT Enterprise
WINDOWS SERVER: 2003, 2008, 2012, 2016, 2019
LINUX: Amazon, CentOS, ClearOS, Debian, Fedora, Oracle
SUN SOLARIS: x86, SPARC Red Hat, SUSE, Ubuntu, others
MAC: Intel, Power PC
HP-UX: Itanium, PA-RISC
AIX

Windows Parameters Monitored

FILE ADDITIONS, DELETIONS, MODIFICATIONS, AND READS
ATTRIBUTES: compressed, hidden, offline, read-only, archive, reparse point
Creation time, DACL information, Drivers, File opened/read, File Size, File type, Group security information, Installed software, Local groups, Local security policy, Modify time, Registry (keys and values), Services, User groups

UNIX Parameters Monitored

FILE ADDITIONS, DELETIONS, AND MODIFICATIONS
Access Control List, Attributes: read-only, archive, Creation time, File Size, File type, Modify time, User and Group ID

Supported Platforms CimTrak For Network Devices

Cisco, Check Point, Extreme, F5, Fortinet, HP, Juniper, Netgear, NetScreen, Palo Alto, Others

Supported Platforms CimTrak For Databases

Oracle, IBM DB2, Microsoft SQL Server
MySQL PARAMETERS MONITORED, Default rules, Full-text indexes, Functions, Groups, Index definitions, Roles, Stored procedures, Table definitions, Triggers, User defined data types, Users, Views

Supported Hypervisors

Microsoft Hyper-V, VMware ESXi 3x, 4x, 5x, 6x, 7x

Supported Cloud Platforms

Google Cloud, Amazon AWS, Microsoft Azure

Supported Container & Orchestration Integrations

Docker, Docker Enterprise, Kubernetes, Google Kubernetes Engine (GKE), Amazon Elastic Kubernetes Service (EKS)

Supported Ticketing Integrations

CA ServiceDesk, Atlassian Jira, ServiceNow, BMC Remedy

Supported SIEM Integrations

IBM QRadar, McAfee Event Security Manager, Splunk, LogRhythm, Microfocus Arcsight, and others