



# Is the S in SRE for "Security"?

John Benninghoff  
Security Differently



Hi, I'm John Benninghoff. I started my consulting company, Security Differently, with a goal of making security less scary, and as much a part of technology engineering as safety is part of mechanical or structural engineering.

This is a story about how security and SRE overlap. I'll have a QR code at the end for you to download the slides with notes and links to all the references.



### "About Me"

My grandfather, circa 1940: the pilot of ~65 years (15-80), always used his pre-flight checklist, started my interest in aviation safety, *The Checklist Manifesto*. I asked, "Can we use this for security?" This led me to the book *Engineering a Safer World* by Nancy Leveson, the "new view" of safety, and the STAMP/STPA Workshop @ MIT. Later, I started a Masters degree at Trinity College Dublin, studying safety science and how we can apply it to security and reliability.  
<https://psas.scripts.mit.edu/home/>  
<https://www.tcd.ie/courses/postgraduate/courses/managing-risk-and-system-change-msconline/>


Images: MIT, John Phelan,

[https://commons.wikimedia.org/wiki/File:MIT\\_Building\\_10\\_and\\_the\\_Great\\_Dome,\\_Cambridge\\_MA.jpg](https://commons.wikimedia.org/wiki/File:MIT_Building_10_and_the_Great_Dome,_Cambridge_MA.jpg)

TCD, Patrick Theiner,

[https://commons.wikimedia.org/wiki/File:Trinity\\_college\\_library.jpg](https://commons.wikimedia.org/wiki/File:Trinity_college_library.jpg)

Image Description: Photo of my grandfather in his American Airlines uniform, the Great Dome at MIT, the Old Library at TCD.



Erik Hollnagel, the scientist who created Safety-II, observed that in safety, much like security, we tend to focus on only the bad outcomes. Success happens when we avoid the bad outcomes, shown on this normal curve in red.


Hollnagel, E. (2014). Is safety a subject for science? *Safety Science*, 67, 21-24.

<https://doi.org/10.1016/j.ssci.2013.07.025>

R source code for visualizations:

<https://jabenninghoff.github.io/security/analysis/constraints.html>


Image Description: normal curve with red area under lower left curve



He argued that we can't have a science based on the non-occurrence of bad events – you can't study something that doesn't happen – instead, we must consider and study the whole range of outcomes, including good (green) and bad outcomes, and the “normal” outcomes between.


I've observed that there are two ways of reducing bad outcomes.

Image Description: normal curve with red area under lower left curve and larger green area under lower right curve



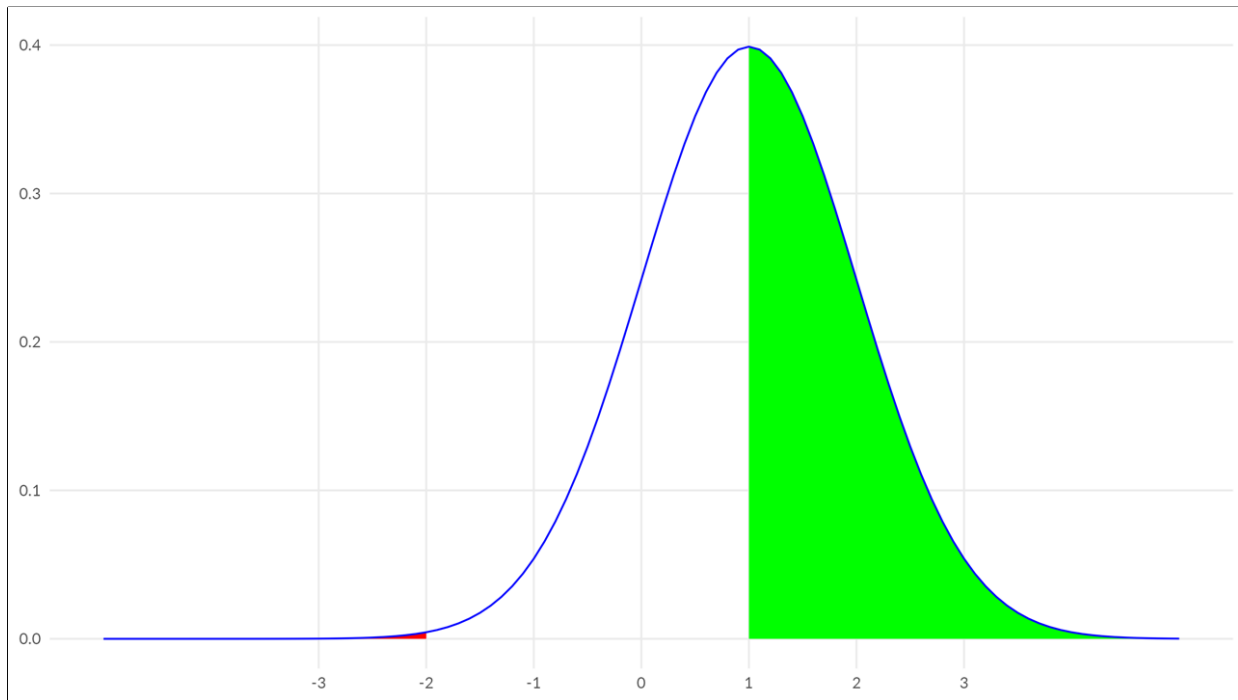
One way is to constrain behavior – introducing policies, controls and procedures that protect against negative outcomes – making the curve narrower.

Image Description: narrower normal curve with red area under lower left curve and larger green area under lower right curve




However, constraints reduce both unexpected negative *and* positive outcomes, shown here. This shows the downside of controls.

Image Description: original and narrower normal curves with red area under lower left curve and larger green area under lower right curve



Another, better, way to reduce bad outcomes is by improving performance – shifting the curve to the right.


Image Description: normal curve shifted right with red area under lower left curve and larger green area under lower right curve



Focusing on improving performance means that security is no longer an expense, since it both reduces bad outcomes and increases good outcomes.

Image Description: original and right shifted normal curve with red area under lower left curve and larger green area under lower right curve





Bringing all three together, this model shows how improving performance is a better strategy – which raises the question, how do we improve security performance?

Part of this is perspective: "I have good security because my home network has never been breached by the Russian or Chinese Governments" vs "We have good security because we successfully defended against an internal red team attack" - preventing occurrence vs performing when exposed to threats.

Image Description: original, narrowed and right shifted normal curve with red area under lower left curve and larger green area under lower right curve

**Cybersecurity performance is highly correlated with software delivery and reliability performance**

Findings from the data, 3 reports, all from 2019 (!)

## 2019 Accelerate State of DevOps Report

Metric	Low	Medium	High	Elite
Deployment frequency	1/month – 2/year	1/week – 1/month	1/day – 1/week	On-Demand
Lead Time for Change	1 month – 6 months	1 week – 1 month	1 day – 1 week	< 1 day
Time to Restore Service	1 week – 1 month	< 1 day	< 1 day	< 1 hour
Change failure rate	46-60%	0-15%	0-15%	0-15%

Google DORA Research started by Nicole Forsgren shows how performance in productivity, reliability, availability and security tend to move together.

Forsgren, N., Smith, D., Humble, J., & Frazelle, J. (2019). *2019 Accelerate State of DevOps Report*. DORA & Google Cloud.

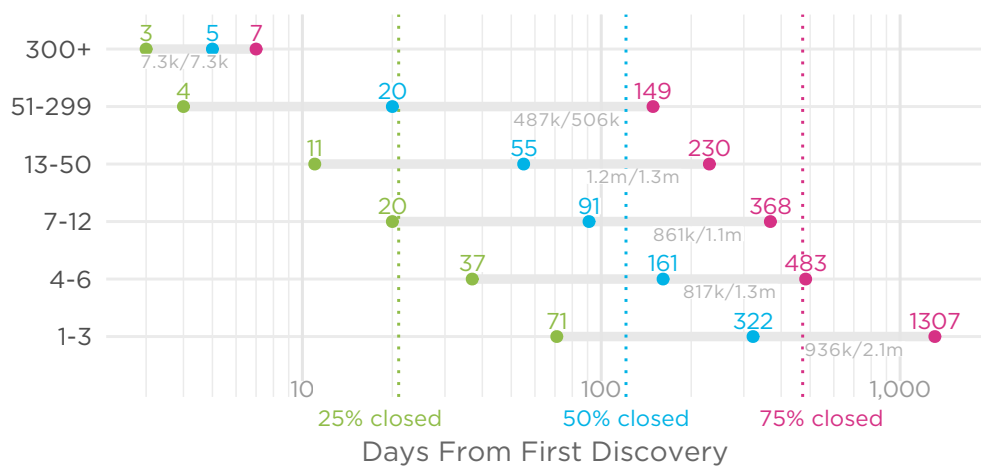
<https://research.google/pubs/pub48455/>

Forsgren, N., Humble, J., & Kim, G. (2018). *Accelerate : the science behind DevOps : building and scaling high performing technology organizations* (First edition. ed.). IT Revolution.

<https://dora.dev>

<https://www.information-safety.org/2022/07/09/definitive-dora-introduction/>

**FIGURE 44: EFFECT OF SCAN FREQUENCY ON FLAW PERSISTENCE INTERVALS**



Source: Veracode SOSS Volume 9

Also in 2019, Veracode and Cyentia published the State of Software Security Volume 9. This is the most significant finding from that report (the “wow” plot) – organizations that scan 300 or more times per year fix security bugs orders of magnitude faster.

Figure used with permission.

Veracode. (2019). State of Software Security Volume 9.

<https://www.veracode.com/sites/default/files/pdf/resources/ipapers/state-of-software-security-volume-9/index.html>

Image Description: Figure showing 25%, 50%, and 75% of security bugs closed (in days, on a log scale) on x-axis and number of scans per year on y-axis.

## 2019 State of the Software Supply Chain

“Most projects stay secure by staying up to date.”


“Projects that update dependencies more frequently are generally more secure.”

Research conducted by Stephen Magill and Gene Kim examined data from the Java Maven Central Repository. They found that most projects stayed secure (remediated security vulnerabilities) by staying up to date (updating their dependencies). In other words, they simply updated dependencies instead of making security a separate task. Want to be secure? Update dependencies frequently.

Sonatype, Galois, & IT Revolution. (2019). 2019 State of the Software Supply Chain. <https://www.sonatype.com/en-us/2019ssc>

Magill, S., & Kim, G. (2019). A data-driven look at practices behind exemplar open source projects. <https://www.youtube.com/watch?v=YoWkuFzEYFs>

## Three Modes of Security Performance



My model of security performance. Today we'll be talking about Mode 1 and Mode 2.

Mode 1: Security is entirely contained within general performance


Mode 2: Security is partly outside of general performance

Mode 3: Security is entirely outside of general performance

<https://www.information-safety.org/2022/05/30/secure360-2022/>

Image Description: A large circle labeled "General" with a smaller circle labeled "Security" contained inside.

## Three Modes of Security Performance



My model of security performance. Today we'll be talking about Mode 1 and Mode 2.

Mode 1: Security is entirely contained within general performance

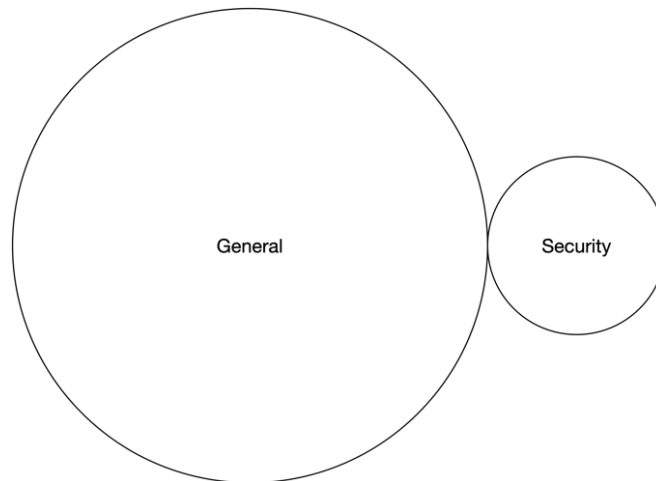
Mode 2: Security is partly outside of general performance

Mode 3: Security is entirely outside of general performance

<https://www.information-safety.org/2022/05/30/secure360-2022/>

Image Description: A large circle labeled "General" overlapping with a smaller circle labeled "Security".

## Three Modes of Security Performance



My model of security performance. Today we'll be talking about Mode 1 and Mode 2.

Mode 1: Security is entirely contained within general performance

Mode 2: Security is partly outside of general performance

Mode 3: Security is entirely outside of general performance

<https://www.information-safety.org/2022/05/30/secure360-2022/>

Image Description: A large circle labeled "General" with a smaller circle labeled "Security" just outside.



# Evidence-based cybersecurity policy

## Top security controls

- **Attack surface management**
- **Patching cadence**
- Cloud-based email
- Avoiding specific VPNs
- Multi-Factor Authentication

JOURNAL OF CYBER POLICY  
<https://doi.org/10.1080/23738871.2024.2335461>

 **Routledge**  
Taylor & Francis Group

 OPEN ACCESS  Check for updates

**Evidence-based cybersecurity policy? A meta-review of security control effectiveness**

Daniel W. Woods <sup>a,b</sup> and Sezaneh Seymour<sup>b</sup>

<sup>a</sup>School of Informatics, University of Edinburgh, Edinburgh, United Kingdom; <sup>b</sup>Coalition Inc., San Francisco, United States of America

Academic paper published in 2024, a meta-review of industry and academic studies primarily looking at the impact of controls on insurance claims and security incidents. The top 2 security measures were Attack surface management (configuration and hardening) and Patch cadence (patch faster). Note on use of MFA: passkeys are probably better, we don't have as much empirical data compared to MFA.

Woods, D. W., & Seymour, S. (2024). Evidence-based cybersecurity policy? A meta-review of security control effectiveness. *Journal of Cyber Policy*, 1-19. <https://doi.org/10.1080/23738871.2024.2335461>

Image Description: A picture of the title section of the “Evidence-based cybersecurity policy?” paper.

## Security performance contained in SRE

### Top security controls

- **Attack surface management**
- **Patching cadence**

### Operational activity

- Inventory and configuration management
- Software and Dependency management

The top two security controls are core operational (SRE) activities – managing inventory and configuration, updating software, both basic technology maintenance activities. In past talks, I've argued that we shouldn't have a vulnerability management security team, except as an independent audit of maintenance performance.

## Security capabilities that overlap with SRE

- Observability
- Incident Response
- Post-Incident Investigations
- Testing

While some security capabilities are fully contained in SRE, most are overlapping, including observability, incident response and investigation, and testing.

## Observability is Observability.

My first security project, in 1999, was to build a security monitoring tool (Network Intrusion Detection System). SHADOW was an open-source tool that used tcpdump and perl to build web pages to report on suspicious activity at the protocol level, which we later replaced with a more capable commercial tool. One of the things that happened was that the network team and I would call each other – I would see network problems and they would see security issues. Both SRE and Security benefit from visibility into what's going on “below the line”, although they do care about different things; SRE is more focused on performance, Security more on anomaly detection.

## Observability – SIEM Tools


<p>General Purpose Security Specific Security Specific</p>	<p>General Purpose General Purpose Security Specific Security Specific Security Specific</p>
<p>Omitted</p>	<p>Security Specific General Purpose Security Specific General Purpose General Purpose</p>

SIEM = Security Information and Event Management. “SIEM systems provide a single interface for gathering security data from information systems and presenting it as actionable intelligence.” I’ve personally seen a common trend with SIEM: organizations replacing a security-specific SIEM solution with an add-on to a general-purpose solution. This makes sense for a couple of reasons: first, saves cost over purchasing two tools, second, in my experience general purpose solutions scale much better.

The chart here shows a certain analyst’s “Quadrant”, listing both security-specific and general purpose SIEM/observability tools. General purpose tools account for about half, and the top 2 tools are general purpose.

[https://en.wikipedia.org/wiki/Security\\_information\\_and\\_event\\_management](https://en.wikipedia.org/wiki/Security_information_and_event_management)  
<https://cloud.google.com/blog/products/identity-security/google-is-named-a-visionary-in-the-2024-gartner-magic-quadrant-for-siem/>

## Incident Response & Investigation



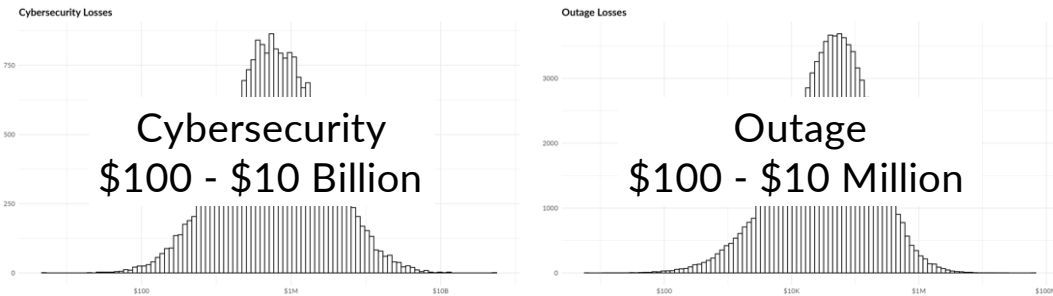
Security incidents are larger but less frequent than outages: security incidents occur about once every 3-10 years (or longer), depending on size. A typical (median) loss event is just over \$250K, 95th percentile is \$52M. This is reflected in these histograms from an artificial but realistic simulation (100,000 runs). The shape is similar as the simulation uses similar distributions for each, but the magnitude is much different. (Outage \$100-\$10M, \$100K typical, Cybersecurity \$100-\$10B, \$1M typical)

Security Incident data from: Cyentia Institute. (2022). Information Risk Insights Study (IRIS) 2022. [https://www.cyentia.com/wp-content/uploads/IRIS-2022\\_Cyentia.pdf](https://www.cyentia.com/wp-content/uploads/IRIS-2022_Cyentia.pdf)

Histograms: <https://jabenninghoff.github.io/security/analysis/rq-demo.html#individual-histograms>

Image Descriptions: (Left) A normal-shaped histogram of cybersecurity losses on a log scale. (Right) A normal-shaped histogram of outage losses on a log scale. (Both are log-normal distributions)

## Incident Response & Investigation



Security incidents are larger but less frequent than outages: security incidents occur about once every 3-10 years (or longer), depending on size. A typical (median) loss event is just over \$250K, 95th percentile is \$52M. This is reflected in these histograms from an artificial but realistic simulation (100,000 runs). The shape is similar as the simulation uses similar distributions for each, but the magnitude is much different. (Outage \$100-\$10M, \$100K typical, Cybersecurity \$100-\$10B, \$1M typical)

Security Incident data from: Cyentia Institute. (2022). Information Risk Insights Study (IRIS) 2022. [https://www.cyentia.com/wp-content/uploads/IRIS-2022\\_Cyentia.pdf](https://www.cyentia.com/wp-content/uploads/IRIS-2022_Cyentia.pdf)

Histograms: <https://jabenninghoff.github.io/security/analysis/rq-demo.html#individual-histograms>

Image Descriptions: (Left) A normal-shaped histogram of cybersecurity losses on a log scale. (Right) A normal-shaped histogram of outage losses on a log scale. (Both are log-normal distributions)

## Incident Response & Investigation

### Security Team

- Incidents last days or weeks
- Threat Hunting
- Forensics & Investigation
- Packet Analysis & Memory Dumps
- **Information restricted**

### SRE Team

- Incidents last hours
- On-Call Incident Response
- Post-Incident Review
- Packet Analysis & Memory Dumps
- **Information shared**

The differences between incidents develops overlapping but different skills on the security and SRE teams. What I've seen is that security incident response plans are longer and larger, and security spends more time looking for problems, and analyzing the damage through forensics and automation. SRE incident response is higher tempo and more frequent, with more practice on communication and coordination. Both teams still have experts in packet analysis and memory dumps! These skill-sets are complementary and are useful for both breaches and outages.

Unfortunately, I've too often seen a difference in information flow: SRE teams tend to share information broadly to improve response, but security teams tend to over-restrict information, which I believe hurts security response.



# Does the system behave the way we expect?

(when things aren't going well)

Testing is a large and difficult subject to cover. In my experience, the core question is this. Both Security and SRE extend this question just a bit. While we will never have certainty, we can make it less likely the system will behave in unexpected ways. This includes different forms of testing but also static code analysis, formal methods, chaos engineering, language constraints (memory safe languages, parameterized SQL queries), and documentation, including writing down assumptions. Continual testing with automation is key; the manager of one of the highest performing development teams once said to me “We don’t have bugs.” (He meant *unresolved* bugs) While Security and SRE tests are focused on different things (malicious behavior vs accidents), bugs are bugs, and finding and fixing bugs improves both security and reliability. Human nature is to test for the happy path, both teams contribute to success by testing for different kinds of unhappiness.

## Security Level Objectives

"Security Level Objectives". Unlike Service Level Objectives, it's unlikely that there is a tolerable level of security breach loss, especially for smaller companies - a small shop that brings in \$100K per year could lose nearly its entire annual earnings in a typical loss event! Instead, we can leverage indicators that are correlated with lower security risk: attack surface, patch cadence, use of MFA. My experience: I started a vulnerability management program in 2002, and in retrospect, I realized we had an SLO. When there were "too many" vulnerabilities on the report, our head of infrastructure would send out an email asking our teams to clean it up (patch things).

Cyentia Institute. (2022). Information Risk Insights Study (IRIS) 2022.  
[https://www.cyentia.com/wp-content/uploads/IRIS-2022\\_Cyentia.pdf](https://www.cyentia.com/wp-content/uploads/IRIS-2022_Cyentia.pdf)

## Potential Security Level Objectives

- Vulnerability rate (# of vulns per endpoint)
- Open ports per public service
- % of endpoints not using MFA
- % of orphaned accounts
- Login failure rate
- Login success rate

What might an SLO look like?

Like traditional SLOs, falling below the agreed metric results in diversion of resources to improve security. While 100% is probably the wrong target, in some cases 100% may be the right target: by eliminating all remote code execution vulns, we were able to successfully defend against an internal pen-test in 2004.

## Practical Take-aways

- SRE work supports core security performance
- Extend SRE capabilities to support security and security capabilities to support SRE
- Security and SRE will go further together

SRE and Security have overlapping and complementary skills

## Slides, Connect & Resources



**Connect:**  
[linkedin.com/in/jbenninghoff/](https://www.linkedin.com/in/jbenninghoff/)

**Website:**  
[jbenninghoff.com](https://jbenninghoff.com)  
[security-differently.com](https://security-differently.com)

**Resources:**  
[erikhollnagel.com](https://erikhollnagel.com)  
[dora.dev](https://dora.dev)  
[cyentia.com](https://cyentia.com)



Scan the QR code for slides and more! Questions?

Image Description: A QR code linking to <https://bento.me/jbenninghoff>