Developing ITIL - Mature Security Incident Response With SIEM

A Plan for CSIRT Maturity Models
via monitoring-driven Kanban

Part 1 of 3

Foreword

Information security stands as one of the newest, formally adopted spheres within Information Technology, only reaching acceptance as a standard component of IT Service delivery within the last decade. However, as a field within IT that is not directly driven by the advancement of market demand, the evolution of technology, or balanced competitive advantage (driven instead by competitive advantage driven by unregulated hostile actors), the path to process maturity for Information Security has been a difficult one.

Today, Information Security still carries many stigmas because of its difficult in attaining a maturity model as quickly as some other aspects of the IT sphere have. We have one significant difference in security however, compared to standard ITIL-capable services such as Technical Support.

Security Response Is Not an Internally Demand-Driven Service

Specifically to illustrate, the tech support model adapts so quickly to ITIL, because of the obvious demand-driver.. people calling the helpdesk to fix problems with their computers. Let's flip this on its head for a second here.. how quickly would tech support have adapted to the ITIL model, if every tech support person had to cold-call around the company, asking them “are you having problems with your computer?”

And here we get to the crux of the matter, that security is a discovery-based service for the most part, classically, driven by audit to discover issues to drive workload, since security issues usually do not directly make themselves known.

This is not to say that security is without demand-driven aspects, password resets, requests to open firewall ports; however many of these things are more in the realm of usability than true security, certainly their demand driver is one of usability, just like tech support is.

And thus we come to the realm of uncertainty that makes true security: continued monitoring and response, so difficult to drive with ITIL, that vague area where response action is driven only by discovery, and discovery is an unpredictable result.

The goal of this document is to demonstrate using SIEM and log correlation to:

1. Automate Discovery to produce a demand-driven workflow
2. Plan out a workflow solution that can enable real metrics and gap analysis for your security program
3. Develop an Incident Response process that includes a significant portion of repeatable, measurable and instructable processes.
4. Develop an Service Catalog for a CSIRT that provides value to the business outside of crisis management.

Table of Contents

Definitions and Myths

Firstly, let's take a look at some of the existing models within Information Security Departments, and lay down some of the delineations between them, and some common mis-perceptions about those delineations that are encountered.

(Now, depending upon the size of any real-world security department and its maturity level, many of these roles may overlap out of necessity of resource limitations.)

Security Programs should ideally follow the same multi staged lifecycle of any other IT service, with clear separation of duties along its course.

Architecture → Deployment → Operation → Usage

The most common combination observed in actual security departments is that of operation and end-usage being combined: it is not uncommon for the same people administering the intrusion detection systems to be the people analyzing the output from it. In an ideal level of organization development, resourcing and process maturity, these two roles would be separated; ideals aside however, it is this common combination of duties that produces and oft-encountered mis-perception, that Security Operations and Security Response are the same functional role..

In practice however, this unison has played out almost universally, with the demands of the operational side of the role consuming all available resources to the exclusion of the monitoring and response aspect. Combined with the loosely-driven aspect of monitoring, discovery and response commonly used for security incident response, leads us to the situation we have today; where organizations that have been breached, are forced to disclose that they only discovered the breach week, even months, after it occurred.

It becomes immediately apparent that operations and monitoring must be separated roles, within any successful security program..Having a SOC cannot replace the need for a separate response functional role, either internally or via integration with an MSSP.

SOC vs CSIRT

A commonly-encountered mis-perception is that Security Operations and Security Monitoring and Response are an identical service domain. Fundamentally however the two exist across a mirror of separation of duties; complementary to one another but ideally separated into specific services.

(Inevitably, the two functions will be found being executed side by side, but this is a mark of limited service level maturity and resourcing)

SECURITY OPERATIONS

Implementation
Deploys and maintains security controls, adapting them to the functional requirements of the business.,

Maintains Systems Compliance
Maintains compliant configurations and controls on deployed systems in a scheduled fashion

Administrative Access
General administrative access for configuration and maintenance.

Planned Workload
Workload driven by staged business planning and change control. Unpredictability is driven by Exception planning.

Tasked Externally
Change requests come from external business units to configure security controls to enable business continuity.

SECURITY RESPONSE

Validation
Observes the functioning of those controls to continually validate they are presenting the intended security posture.

Locates Non-compliant systems
Identifies systems that are not under compliant governance (re: 'Shadow IT') as they are discovered by their activity.

Audit Access
Limited Read-Only access where possible to maintain integrity of information

Discovered Workload
Workload driven by Discovery and Alerting. Unpredictability is an integral aspect of capacity planning.

Tasks Externally
Tasks the business with Remediation Changes to ensure business continuity.

Ultimately, a clear delineation between security operations and security monitoring, prevents the generation of 'selective blindness' to issues and visibility; even more so when ultimately, the response team is the source of the data that validates the efforts of the operations team.

The relationship between Operations and Response is a symbiotic one, each side fueling the other, but without that separation, producing effective service drivers and metrics is all but impossible.

The State of Incident Response Today

Because of Incident Response's assumption of being a functional role driven by manual discovery from monitoring systems, it remains a fundamentally ad-hoc role that seems a poor fit for being process-driven standards like ITIL.

This is not the case however, and individually, organizations are discovering ways to bring process and metrics to incident response, discovering much additional value in the side-effects of doing so.

Commonly-Observed levels of Incident Response Implementations

The Emergency Team

The weakest of all incident response plans, and historically, where the vast majority of organizations have started from. A breach is discovered by an observant sysadmin and a response team is pulled from SME's in the organization to build an ad-hoc investigation team. The classic “Panic Model” reigns as process here, perhaps with some limited guidance from Business Continuity Planning.

The “Review during Downtime” model

Commonly seen once the organization has a limited security operations program in place; where a small portion of time is set aside from operational and deployment workloads to review logs Methodology is usually left to the experience and skill level of the individual performing review, usually with no direct process for obtaining metrics or intelligence from the results beyond where something requiring investigation is discovered.

The Linear Event model

Events to review are generated by log correlation, or just derived directly from dedicated security controls. Each event is handled as an atomic unit of data and response work, extrapolation of individual events and alarms into larger patterns is usually left to the instincts and intuition of the individual analyst. Workflow tends to be split amongst several disparate systems.

This third stage, of linear event and response, has become the de facto modern standard for most incident response functions; commonly tagged as the 'firefighting' model, it is reactive, and generates poor metrics drivers beyond uncontexted numbers. When compared to other process driven services, it's obviously still an immature service model.

Reactive Models are difficult to escape from, the classic trap of “We're too busy fixing problems to build new solutions” is not a creation of the IT world, but certainly seems symbolic of it.

Information security itself, has a reputation for being a cost center that produces very little value-added back to the organization, a necessary evil that is the cost of doing business over an increasingly hostile Internet.

However, Information Security operations often have access to a level of information about infrastructure usage and activity within the enterprise that lies at a greater level of granularity and metadata contextuality than almost any other functional unit in the organization; information that, if derived from a well-matured process-driven service offering, could provide a great deal of value to business intelligence and process-steering efforts through the enterprise.

The key is, retooling incident response workflow and information gathering, to produce something that can be packaged for consumption by the rest of the enterprise.

A Word on Metrics.

Mention “Security Metrics” in a crowded room of security practitioners, and watch a heated discussion of conflicting opinions arise on-call.

Information Security's rapid rise to significance over the past two decade and rapidly-evolving theater and technology have left almost no stable ground on which to build a solid foundation of fundamentals that have been tested over time. The biggest casualty in this hyper-accelerated evolution of the field has been the establishment of accepted Key Performance Indicators.

A core tenet of any proper metrics program, are the Key Performance Indicators, figures which demonstrate the rate of improvement (or the inverse) of a business operation over time, which can be correlated back to specific activities and decisions to directly prove their effectiveness (or lack of). This is the key purpose of any metrics program: to identify and learn from previous mistakes and successes, in a process of informed continuous improvement.

ITIL does define a section on “Incident Management', this section is loosely-defined however, into a catch-all “IT Incident” definition, derived around business continuity and external input (ie. The assumption that incident responders are working on incidents only once they are identified by parties outside the incident response group). Fore-mostly however, it makes the assumption of linear, granular tasking, in the same model as the IT Service Desk.

(Later in this document, I will be making the case that incident response is fundamentally not a linear operation, and that all metrics models that treat it as such, are broken on first principles)

Vectors and Scalars

In the field of physics, there are two essential types of values in all calculations: *Scalars* and *Vectors.*

Scalars are isolated values, with no outside context. Indeed They remain the same regardless of any context. A common example would be *mass*. An object has a mass of 1 Kilogram no matter where it is, or how much physical space it occupies. The context of the object cannot change the scalar value of its mass.

Vectors are contexted values, and can change depending on that context. An object has *weight*, dependent on both the mass value and gravity context of the object. An object with high mass, may still have no weight in the corresponding context of gravity.

But most specifically, vector values allow the calculations of *change over time.*

Simply put, numbers (scalars) without context (vectors) are not true metrics; It is meaningless to say that costs to operate the IT Service Desk have doubled within the last ten years, without also showing how the number of employees has tripled in the same time.

And here, metrics and information security analysis share a similar tenet, that is is not the raw data that is important, but the *relationships amongst that data.*

Information Security and Scalar Metrics

Obtaining viable metrics from any security program has been an essentially ad-hoc process almost since inception, largely due to the mistake of confusing values with metrics (or even worse, of mistaking point-in-time compliance results for metrics).

Some common examples of 'security metrics' delivered

> *"This quarter we blocked 200 Million Spam Emails"*
>
> (..but how many were *not* blocked?)
>
> *"This quarter we saw a rise in the number of viruses detected on our systems by 20%"*
>
> *(..what was the percentage change in those not detected?)*
>
> *"Attacks against our perimeter increased by 50%"*
>
> *(..so were the current threats just increasing their attacks? Or are more people attacking us now?)*

In short, the examples demonstrate the two fundamentals of true metrics, relational change over time and attribution of change factor. Building Gap Analysis is one of the key purposes of a good metrics program, to enable and inform resourcing and demand management.

This informed gap analysis can be derived however, through the implementation of a security monitoring and response program, and the base assumption that:

Everything the Incident Response Program does,
is a gap indicator of where the rest of
the security program has failed

It sounds like strong wording at first, but operates on the (impossible in the real world assumption) that, were the controls deployed and operated by the rest of the security program, perfect, the Incident Response team would never see anything that required investigation.

By realizing incident response as the primary source of gap analysis metrics, the potential for creating true metrics and sources of business intelligence from the entire Information Security program, becomes a possibility.

Incident Response is not Tech Support

Inevitably, every incident response department evolves to the place where it needs to put a workflow tracking system into place; the solution that gets implemented is almost always, some variant of a tech support trouble-ticket system, usually slightly retooled to try and match incident response work, rarely successfully.

These implementations bring to light some basic differences in incident response workflow versus support incident workflow

Incident Response is not a granular process.

In security work, each alert triggered by controls, cannot be taken purely as a single isolated datapoint, Those seven different IDS alarms from different areas of the infrastructure may be directly related. Without some indication that they are related (and no direction of workflow to address related alerts first) it remains purely up to the attention and intuition of the analyst to track any patterns they may see. If significant numbers of other, unrelated alerts are investigated between those alerts, no relation may even be possible to be discerned.

Alerts Should not be Prioritized on First-In First-Out

In a perfect world, alerts would be prioritized on their impending threat to the enterprise, not the order they arrived in. Indeed, this prioritization would be adaptive and recalculate priorities as further events happen and alarms trigger.

Tradition Escalation models map poorly to Incident Response.

The model of 'tier 1', 'tier 2' levels of complexity are difficult to assign to security events, until well into the investigation. Some simple levels of initial assignment are possible based on the technology involved, however, it is more appropriate to map things based on specific domains than on pre-assumed complexity. A good example is that of quality assurance and bugfixes; in that model, there is no equivalent of a 'tier 1 bug', but instead a 'database bug', or a 'GUI bug', which are then routed to directly to analysts with the relevant domain expertise.

Standard Methods of creating efficiency metrics are poor indicators

Just as Average Call Time has been much-maligned as a single indicator for service desk effectiveness (sacrificing quality for quantity), the same caveat applies to incident response: a single alert may uncover a hornet's nest of related data and tasking that may take days to cover to completeness.

Potential New Models for Incident Response Workflow

Earlier we covered the three most commonly-observed models for Incident Response Workflow; let's now cover two possible stages of evolution of those models, more closely aligned with ITIL service maturity.

Intelligence-driven Adaptive Incident Response.

Over the past few years, incident response has found itself becoming aligned to two well-established fields that have a wealth of experience to draw from in building a new maturity model; Business Intelligence and Open-Source Intelligence (nb: “open source” here does not refer to software, but to publicly-available information). The BI and OSINT communities have many highly detailed and adaptive technologies and methodologies that can map directly to Information Security service programs.

This model essentially works by building a security monitoring and incident response program that attempts to directly map its workflow and service offerings to those of BI and OSINT; the goal is to not only effectively investigate and remediate security events, but to use them as a primary driver of attribution, metrics, modeling, attribution and prediction to the security program and business as a whole by focusing on large-scale patterns and interactions between business operations and the workload encou..ntered by incident responders.

The fundamental pivot of this model is that incident response creates a wide swathe of data that can be leveraged into overall business intelligence efforts, once BI-style techniques are applied to it. Fundamentals such as data reduction through metadata aggregation techniques, OLAP cube modeling and analysis provide ways to reduce the immediate workload and in turn bring larger contexts to individual units of work, allowing long-term patterns to emerge.

At the introspective level, metrics for the Incident Response group itself become far more relatable to the enterprise, moving away from raw numbers of 'incidents handled', towards a more contexted via of what those incidents represented to coverage and service levels toward individual business units. The linear first-come-first-served workflow becomes a more demand-driven adaptive flow, re-prioritized dynamically by the actual potential business impact to events.

The 'Just-In-Time' (Kanban) Adaptive Reaction Model

This evolution of the prior model, takes the enhanced awareness and dynamic re-prioritization of that model, adapting the Incident Response process from being a corrective control at the end of the security lifecycle, into being a fundamental process driver for the entire security program.

The model works on the grounding that “A vulnerability that is not exposed to a threat, is not a risk”, and builds dynamic exposure models of the infrastructure from state and activity data, providing a prioritization queue for remediation work based upon imminent threat and moving windows of exposure, rather than monolith remediation plans (which by definition, do not scale to match the pace of the business, being outdated as soon as the plan is finalized).

This 'just in time' model attempts to prioritize work in the most resource-effective way for enterprise security, identifying where the least work will have the most benefit. A key concept here is the edict that “As a defender, you don't get to define what are the highest-value targets on your infrastructure; your attackers dictate that”; to wit, the most important system on your network is, from the viewpoint of an attacker, the ones that let them get a foothold on your infrastructure that can remain undisturbed and undetected for as long as they need; by the time they trigger alerts on the systems you have prioritized as high value, the damage has already been done.

Adaptively re-prioritizing these lower-level systems in response to their exposure to threat as it happen is the key to the kanban Incident Response model; that recent firewall change that exposed a lab VM that was previously unreachable remotely, suddenly makes the patching and monitoring of that VM a one-time priority, closing the window of exposure as it opens.

This “exposure response” evolution of Incident Response is where we can truly say we are approaching proactive workflow; while still reactively event-driven, we focus more on reacting to events that pre-indicate security breaches.

Both of these models are implementable today, though the requisite technologies are rarely combined and the execution processes rarely in use within Incident Response. The rest of this document illustrates potential implementations of these two models and their components.

Conclusion to Part 1

- Our current incident response workflow models are born from reactionary necessity
- We are drowning in data, yet refuse to adapt models for maximizing efficiency of handling this data that have provably worked in other service models.
- Linear workflows do not scale well for incident response.
- Incident Response is an excellent source of Metrics, but requires contexting with business process and change-over-time to make meaningful metrics from the numbers.
- There will always be more work than there are resources to handle; new prioritization systems that adapt to emerging events instead of arbitrary assessments are vital to maximizing the effectiveness of resourcing and workload.

In Part 2

- Selecting a Service Catalog to build workflows and metrics from.
- Suggested Data Fields to acquire and track.
- Configuring SIEM to drive process and procedure.
- Building Key Performance Indicators for Incident Response that demonstrate value to the Business.