

An Operational Framework for Incident Handling

Giovanni Bottazzi, Giuseppe F. Italiano, Giuseppe G. Rutigliano
University of Rome “Tor Vergata”, Rome, Italy
gbottazzi73@gmail.com, giuseppe.italiano@uniroma2.it,
rutigliano@ing.uniroma2.it

Abstract

The information security management is a widely discussed topic in recent years, due to the increasing number of attacks and the growth of the damage they can cause to the daily life of a society. In this context, new emerging paradigms, such as IoT, the CPS and Critical Infrastructure, converge towards common technologies, resulting in a dangerous interconnection and interdependence of worlds formerly separated, or even isolated. For this purpose, numerous cybersecurity frameworks have been defined, identifying organizational methodologies, mainly process-oriented, for managing a security infrastructure. This article is rather oriented to define a framework with a special attention to the management of the IT incidents, describing some minimal arrangements that need to be adopted in order to respond effectively and efficiently to a cyberattack, to mitigate the damages suffered and to limit the analysis and the recovery time.

1 Introduction

The threat posed by the criminal use of the Internet is constantly growing and causes, at the moment, in most cases, a steady rise of phenomena whose purpose is mainly to carry out criminal acts, for economic purposes, to the detriment of millions of users (victims). Without dwelling on the difference existing between common crimes perpetrated by means of the Internet (e.g., online sale of drugs, counterfeited goods, etc.) and crimes born and developed thanks to it (Distributed Denial of Service, cyber extortion, etc.), it can be asserted, without fear of contradiction, that cybercrime is a growth industry.

The revenues are great, considering that the Internet economy is capable of generating 2 to 3 trillion dollars per year. It has been estimated (June 2014) that the likely annual cost to the global economy from cybercrime is more than \$400 billion. A conservative estimate would be \$375 billion in losses, while the maximum could be as much as \$575 billion (McAfee, 2014). Criminals still have difficulty turning stolen data into financial gain, but the constant stream of news contributes to a growing sense that cybercrime is out of control.

Moreover, the current trends suggest considerable increases in the scope, sophistication, number and types of attacks, number of victims and economic damage (Europol IOCTA, 2015). In addition to the partiality of information available on the exact incidence of cybercrime on the Internet economy (or on the global economy), we must consider the vagueness of information related to what could be funded with the cybercrime industry revenues. Are they used just for self-funding or for supporting terroristic actions/organizations?

There are some important factors worth highlighting in this context: the widespread of unsecured targets and the low cost barriers for the newly emerging “Crime as a service” rental model. The term “unsecured targets” must be related not only to consumer devices (e.g., smartphones, tablets, smart TVs, etc.), but also to technological (and critical) infrastructures facing for the first time the Internet, e.g., Smart Buildings, Cyber Physical Systems, Industrial Control Systems and all the devices related to the paradigm of the “Internet of Things”. The attention of industry is yet not fully focused on cyber security or on privacy-by-design. Many of the so-called smart devices are actually quite dumb when it comes to their security posture, being unaware of the fact that they are part of a botnet or being used for criminal attacks.

The security of Critical Infrastructures (CI) needs to be addressed in a holistic and effective manner in order to protect economies and societies, mainly because it relates to complex systems composed of various hardware and software components, often deployed by different vendors and designed with little consideration for network and software security. Cybercrime is already able to move huge amounts of capital from the licit economy to the illegal one. It is very likely that, in the near future, the data will be even more alarming, because it is hard to imagine that the trend will decrease, and because today we are witnessing many examples that can highlight as the sophistication of the attacks can further degenerate into "terrorist use" of the Internet or, worse, into cyber terrorism.

The invasiveness of the Information and Communication Technology in various aspects of the life of a civil society, naturally brings, on the one hand, all the benefits related to automation, monitoring and remote controlling, a.k.a. genuine innovation, but brings also, on the other hand, all the side effects related to its misuse, improper or, worse, criminal (Bottazzi, Me, 2016).

Moreover, the emergent industrial paradigm is convergent toward common software technologies, inheriting all the issues already available for the “pure” IT ecosystem, but involving infrastructures much more critical than the usual Internet blogs, with easily predictable domino effects.

In this context, a number of initiatives have been realized, with specific reference to the well-known frameworks for cyber security aimed at providing to organizations a homogeneous and volunteer approach to face up cyber security, in order to reduce the risk linked to cyber threats.

These frameworks, use the business drivers to guide cybersecurity activities and consider cybersecurity risks as part of the organization’s risk management processes, with important fittings relating, for example, the critical infrastructures, the small and medium enterprises or the incident handling (to be considered as the reacting phase of a cybersecurity framework).

However, they still suffer from a high-level perspective in both the prevention and reaction stages, and need a much deeper analysis, aiming at providing practical guidelines to security designers, operators and officers.

2 Current and future Internet-related threats

Many of the Internet-related tools already available, whose business model can amplify their capabilities, as already said previously, have been used in the past for actions that can hardly be labeled as mere crime (e.g., the well-known “Estonia 2007” and STUXNET). It was not introduced any technological novelty (a massive DDoS in first case and a number of Zero-day vulnerabilities in the second). Late last year, a wave of cyber-attacks hit several critical sectors in Ukraine (Kaspersky, 2016).

Widely discussed in the media, the attacks took advantage of known *BlackEnergy* Trojan as well as several new modules.

BlackEnergy is a Trojan created by a hacker known as “Cr4sh”. In 2007, he reportedly stopped working on it and sold the source code for an estimated \$700. Around 2014, a specific user group of BlackEnergy attackers began deploying SCADA-related plugins to victims in the Industrial Control Systems (ICS) and energy sectors around the world. The attackers have simply decided to shift the focus towards new targets, using spear-phishing emails carrying malicious Excel/Word documents with macros to infect computers in a targeted network (example in Figure 1). It should be clear at this point that, given the maturity of both the technology and the business model, a criminal attack can become terroristic, or something worse, just considering “who is opposed to who” and not by means of the tools used (almost always the same).

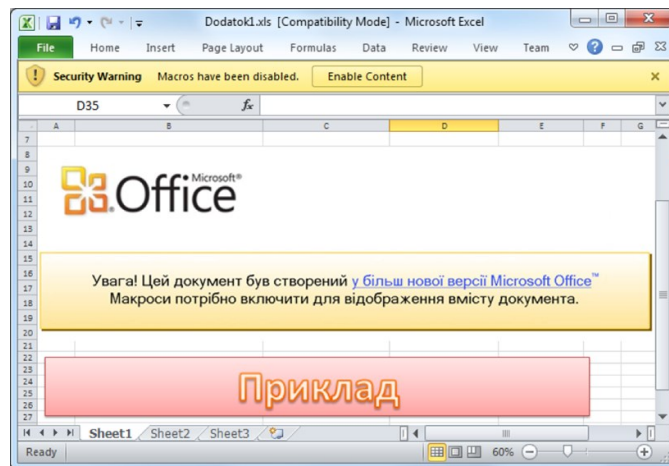


Figure 1: Excel file with macros used in the BlackEnergy campaign.

With reference to CIs, despite the numerous attempts made so far, there is still no universally recognized definition, or at least a definition that provides a classification fitting the characteristics of each nation. A critical infrastructure is often identified as that infrastructure whose incorrect working, even for a limited time period, may negatively affect the economy of individual subjects or groups, involving economic losses and/or even expose people and things to a safety and security risk (TENACE project, 2014).

CIs are at the heart of any advanced civilized country. These infrastructures include among others: finance and insurance, transportation (e.g. mass transit, railways and aircrafts), public services (e.g., law enforcement, fire and emergency services), energy, health care. All surveys from leading security stakeholders indicate that attacks are expected to increase in scale, to become more accurate and precise, and therefore to become real cyber weapons. This raises security concerns (and threats), because of the interconnection of two previously isolated worlds, the Internet and CI systems. Interestingly, the Internet is itself an underlying critical asset of modern CIs, because their controlling systems are often distributed over remote, Internet-connected locations. Moreover, these two worlds are not only interconnected, but also interdependent with an increased risk of a domino effect. Unfortunately, we already witnessed a number of events that must be perceived not just as the “*weaponization of the coffee pot*”.

In addition to the mentioned BlackEnergy event, we can also cite the accident suffered by the Israeli power grid on last February, labeled as the largest cyber assault that the country has experienced (The Jerusalem Post, 2016). Moreover, the year 2016 started with quite a number of security incidents related to hacks on hospitals and medical equipment (Securelist Blog, 2016). They include a ransomware attack

on a Los Angeles hospital, the same in two German hospitals, a case of researchers hacking a patient monitor and drug dispense system, an attack on a Melbourne hospital and so on – in just two months of 2016!

Finally, a recent research (Wendzel et al., 2014) envisioned that a new class of botnets deployed on Building Automation Systems (BAS) and used for “novel scenarios” like remote access to sensor data for mass surveillance, or remotely locking the building and holding the people inside for ransom.

Citing the WatchDogs video game, “we are no more individuals, we are data clusters”.

3 Cybersecurity frameworks (a.k.a. related works)

Cyber attacks on CIs are no longer a theoretical, but a real problem probably since the discovery of the STUXNET worm in July 2010. Criminal gangs, hence, can benefit from a relatively new business capacity (widest sense of the term), in addition to the already vast possibilities offered by the most common cybercrime. Consequently, security researchers envisioned the goal of the next generation of malware to stop quietly production at a utility, to affect production of a rival, to sell the shares of a company or to extort money under the threat of a disruption.

On the other hand, cyber threats certainly cannot be faced only by giving up the potentials offered by the IT systems and their interconnection within the network, thus losing the increase of productivity and efficiency linked with computerization. The answer should be systematic, aimed at raising the citizens’ awareness, the “duty of care” of companies and the International “due diligence” of the country about the cyber threat.

It is crucial that in this process of collective raising awareness, we shift from an idea of “IT system security” or “IT security” to that of “cyber threat management”. This means, among other things, to define a process that respects the Constitution principles regarding, for example, the business activity management in order neither to contrast the social benefit nor to affect safety, freedom and human dignity. This consideration implies that the cyber security perspective is not to be seen just in technologic terms, but rather requires taking into account the overall legal and formal duties and the principles of social interest, into which the public and private framework need to converge. For this reason, the duty of protection should become part of the top management responsibility of an organization, as it requires a specific and accurate evaluation by the ones who have the direction and management power.

In this context, the research community felt the need to define a series of coordinated actions to be taken in order to manage the cyber risk. Such actions involve the organization and the technology departments of the company, in addition to the financial management of the risk, also through the establishment of a residual risk management strategy and a strategy to protect the company balance. Furthermore, the cyber risk is intrinsically highly dynamic. It changes as threats, technology and regulations change. In order to start approaching this issue in a way which is useful for the country system (State, enterprises and citizens), common grounds (a.k.a. Frameworks) have been designed, in which the various production sectors, government agencies and regulated sectors can recognize their business, so to align their cyber security policies in a steadily developing process. To reach this aim, a common Framework should be first neutral both in terms of business-risk management policies and in terms of technology, so that each player could keep on using its own risk management tools, managing its technology assets while monitoring at the same time the compliance with sector standards. A Framework may help an enterprise to plan a cyber-risk management strategy, developed over the time according to its business, size and other distinguishing and specific elements of the enterprise.

Starting from the “Framework for Improving Critical Infrastructure Cybersecurity” issued by the NIST (NIST, framework, 2014), a number of related/linked documents have been issued. For instance, the Italian National Cybersecurity Framework (Baldoni, Montanari, 2015), while based on the one

developed by NIST, has shaped the overall approach to the Italian scenario, in terms of legal and enterprise frameworks.

Moreover, each key concept contained in the abovementioned Frameworks, may have been a topic of study in other guidelines, such as the “Computer Security Incident Handling Guide” (NIST, guide, 2012), to be considered as a deepening of the Core activity labeled as “Respond” in the Cybersecurity Framework.

With specific reference to the actions that must be performed during/after a cyber-attack, further research documents have been produced for better identifying skills, organization and collaboration of the teams involved in managing security incidents (Computer Security Incident Response Teams) or for improving their response efficiency (Ruefle et al., 2014; Steinke et al., 2015).

None of the above documents, regardless of their depth level, helps in outlining what should be the security measures to implement, from a purely practical point of view.

We believe that the guidelines for the prevention, management and reaction to a cybersecurity incident, with a strong practical orientation (to be understood as the technological arrangements to implement), should have two main targets. On the one hand, they should represent a valuable tool for inflecting the concepts contained within the frameworks developed so far, and on the other hand, they should be able to spread awareness and greater technical knowledge regarding the burden to bear for deploying a specific IT service.

For example, the adoption of https certificates whose trustworthiness is always verifiable, could help many organizations to significantly reduce the attack surface and to protect them from attack techniques (e.g., Man In The Middle) which, although easy to find on the black market, are not yet easily countered. In the next section, we will describe a framework for limiting the effects of cyber attacks (incident response) in the time, space and data domains.

4 A practical framework for incident handling

In the abovementioned scenario, the well-known paradigm where “the attacker is getting stronger” (Brady et al., 1999), or, alternatively, the attacker has thermodynamics on his side, results to be reinforced and, without proper and very challenging strategies, the phenomenon is expected to grow and hardly there will be a turnaround. Starting from the characterization of the threat, in terms of organization, objectives, operating procedures and action times, we can assert that it is possible to identify common elements, even if it may originate from heterogeneous subjects, with significantly different resources and capabilities.

The group that materially performs a coordinated attack has usually a lean structure and counts a number of people that can vary from a few units to a few tens. The command chain is short, the skills and the tactical objectives are clear and quite common (unlike those strategic): DoS, defacement, data exfiltration, etc. The actions are carried out in subsequent stages, or parallel when possible, by highly specialized working groups (e.g., system analysts, database experts, network experts, crackers, etc.), with a previously scheduled timing and through very “easy-to-find” tools (usually on the black market).

The time, thus, is a crucial success factor for the attacks, often carried out in periods and times carefully chosen, in order to find the lowest defensive power as possible. The attackers have the advantage of surprise on their side, and of course, they try to maximize the effects.

The defensive organization, instead, involves a very large number of individuals in the ICT management. In addition, the common users may be subject to social engineering or phishing attacks. In this context, the number of people involved increase dramatically.

The defensive organization, then, turns out to be much less streamlined than the attacker's one, in terms of tasks and responsibilities. There is no best approach, because the absolute separation of duties

and the use of dedicated resources, that might seem the best choice from a conceptual and procedural point of view, still poses problems on group cohesion (internals and outsourcers) and quick responses.

Moreover, if not all operators are fully trusted, there is even more the need to define tasks clearly identified, easy to implement (because previously tested) and not subject to misinterpretation. The command chain (defensive) is further complicated by the absence of clear operational procedures and can become an obstacle to the operations of containment and remediation mostly when, for example, there is the need to bring together the appropriate committees for identifying actions to be undertaken. Moreover, the procurement of human resources, technologies and services, is composed by articulated procedures and bureaucratic processes, subject to appropriate authorizations. In other words, the defensive processes result to be much slower compared to those involved in the threat evolution and there is the need of organizational, technological and operational changes.

We can imagine a large IT organization as the one described in Figure 2. Typically, we have Headquarters, hosting datacenter and executive workstations, a number of branch offices and mobile devices. The Internet resource is shared through Headquarters or directly through the branch offices (rare). In order to measure the cost of a compromise, we need to identify the assets that can be involved, their single value, the time needed for analyzing, identifying and recovering a compromise. Moreover, a compromise can start from an asset and evolve to other assets through privilege escalation and/or lateral movement. The final goal of a cyber attack is always the compromise of an infrastructural service (e.g., DNS, internal Domain, etc.), giving the attackers the full capacity to act.

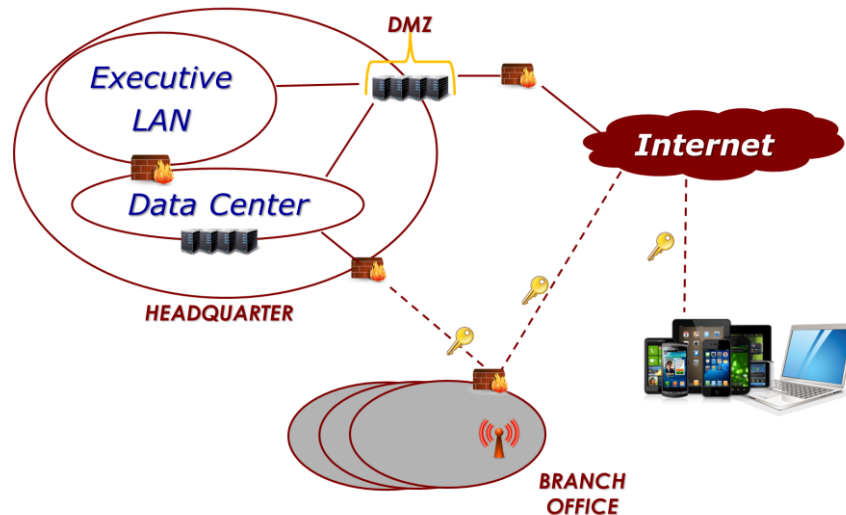


Figure 2: IT architecture of a large corporate organization

In this context, the total cost of compromise (of course just in the success case) can be seen as the sum of the costs of all the “objects” compromised in each asset, plus the time needed for detecting, analyzing and recovering each compromised asset. Thus, we have:

$$\forall \text{ Asset}, \text{Cost}(\text{Asset}) = \sum_{i=1}^{i=n} (V_i + AT_i + RT_i)$$

where n is the number of “objects” composing each asset, V is the Value of each object composing the specific Asset, to be considered as the criticality of the object, AT is the Analysis Time required for analyzing and finding the compromise, RT is the Recovery Time required for recovering the original behavior of the “object”. We can describe the typical assets deployed within a large corporate network,

as the one described in Figure 2, resumed in Table I, together with possible values for V_i , AT_i and RT_i (ranging from 1 to 1,000).

ASSET	V_i	AT_i	RT_i
Common Workstation (mobile device)	1	1	1
Admin Workstation (mobile device)	10	1	50
Executive Workstation (mobile device)	100	1	10
Application Service	300	500	100
Infrastructural Service	1,000	800	1,000

Table I: distribution of V_i , AT_i and RT_i , within a large corporate network.

The values in Table I have been defined considering low AT and RT values for almost all workstations. The only exception regards the RT value of the Admin workstations, which usually require more starting configurations. The V value, instead, has been considered greater for the Executive Workstations, considering all the business strategic data stored within. Moreover, the application and infrastructure services obviously have a high V value, but above all, they have the AT and RT values much higher than the workstations. Finding a compromise inside the e-mail servers or inside the domain infrastructure could be very challenging and recovering from it may result in reinstalling a brand new system.

At this point, it is easy to verify how the total cost of an incident increases, during time, because of the escalation to assets with higher value, and higher costs, resulting in an increased burden for the organization, in terms of analysis and recovery (Figure 3). In Figure 3 we reported the examples related to the compromise of an infrastructural service (e.g., through an Advanced Persistent Threat - APT). We supposed the APT started from a workstation within a branch office (e.g., mail-attached malware) and from an application service (e.g., SQL injection and subsequent reverse shell), described by the blue and the red series respectively (the green series will be illustrated later).

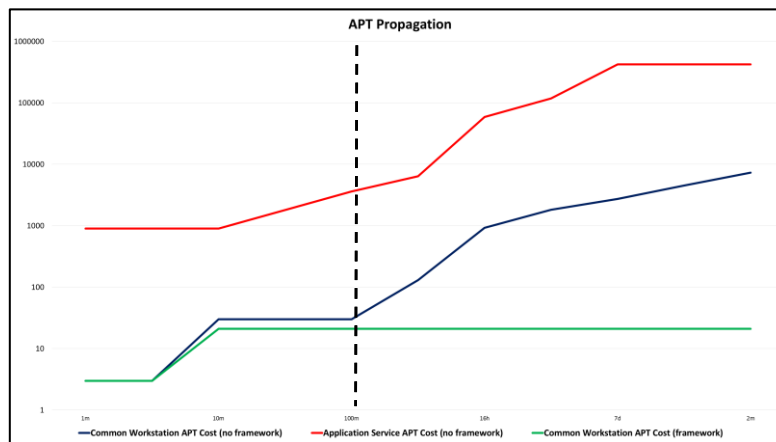


Figure 3: Compromise escalation during time

In the first case, the attack escalation has been imagined as moving first through the common workstations and the application services, before reaching the infrastructural services. In the second case, the escalation path, starting from the application services, is definitely reduced. In both cases, the cost of each escalation step has been computed using the previous formula, thus summing the costs of all the objects involved in the attack. We hypothesized the IT infrastructure composed by 100 application services, 10 infrastructural services, 100,000 common workstations, 50 admin workstations and 30 executive workstations. As we can see, after about 2 hours (the vertical dotted line), assuming

that the two attacks are successfully scaled or addressed toward the “Data Center”, the global cost of the compromise is completely different (displayed in logarithmic scale).

Of course, we did not consider the probability that an attack is successfully completed, which depends on numerous factors such as the security systems in operation, the type of threat, etc. In any case, it would be limited to the identification of a scale factor of the trend described in Figure 3.

In order to reduce the speed gap between attackers and defenders is necessary to adopt procedural and technical measures, designed to streamline the defense system. The arrangements we are going to describe, do not replace the requirements and recommendations contained in other security frameworks, but are meant to improve their effectiveness, with easier and faster defense operations or, in some cases, slowing down the attackers' operation.

Firstly, the initial perimeter of the infection must be reduced. This objective can only be achieved by acting on IT assets with an "incident-handling" perspective and thus implementing, during the design and the management stages, all the prescriptions related to the network segmentation and data segregation. For instance, if an attacker can break into a server, located inside a wide data center protected by perimeter defenses but no internal segregation/segmentation, he can easily extend the attack on the entire server farm. Thus, the reduction of the attack surface must be coupled with the increase of the compartmentalization of internal services.

Domain	Actions	
	Ex-ante	Ex-post
Time	Minimize the Internet border gateways.	Close all or part of the Internet border gateways.
	Arrange only interesting logs in usable format.	Mine the logs starting from those concerning the compromised sector.
Space	Prepare hardware/software probes up to high level of depth.	Limit or interrupt traffic starting from the compromise sector.
	Segment the network infrastructure.	Close the network links starting from the compromise ones.
	Segregate the application data.	Inhibit the communications to and from the compromised application service.
	Implement Content Delivery Networks.	Switch to CDNs.
	Access the application services via terminals and credentials limited and well identified (admins and not users).	Isolate the access to application services.
Data	Access the infrastructural services via terminals and credentials limited and well identified.	Isolate the access to infrastructural services.
	Backup or service-snapshots (off line).	Recovery (on line).
	Data encryption (in order to prevent that the compromise of the server will automatically expand to the data).	Reduce data access to a few entities (prior verification of the credentials).
	Prepare workstations never used (off line) and ready for the emergencies.	Use new workstations previously prepared.

Table II: A framework for incident handling.

In addition, the operators dedicated to monitoring and analyzing possible cyber attacks, usually do not have sufficient Rules of Engagements (RoE) for implementing timely corrective and/or mitigation actions. Hence, the second objective to be pursued concerns the organizational structure, always for

managing security emergencies (incident-handling perspective). It is necessary to make efforts for delegating as much as possible at every command-chain level the responses to attacks, identifying a priori the possible scenarios and preauthorize the responses, giving to operators clear RoE that can be put in place immediately, without fear of incurring in personal liabilities and without waiting for the appropriate authorizations. This approach, borrowed from the military affairs is, in our opinion, the only way to make effective a tool that otherwise would risk being essentially a system for observing the evolution of a phenomenon and not for contrasting it.

The third objective is to prepare in advance the data necessary for the incident analysis. Usually, huge amounts of logs are kept, sometimes even in excess, without a well-focused objective. For instance, if we consider the Internet gateway (the same for all users) of Figure 2, it is very likely that the vast majority of the logs, without any pre-processing, do not provide any useful information to the identification of an attack (e.g., misconfigurations, authentication methods not supported by the application software, etc.). All these logs should be separated from the others in the "peacetime", in order to reduce the amount of data to be inspected in the "wartime".

Moreover, the firewalling policies, usually record all the failed attempts to access network resources (drop logs). Many of these logs refer to mere attempts made by unaware users trying to access to resources for which they do not own the related grant (e.g. network shares). This is quite different from an automatic port scan of an entire class of IP addresses.

Hence, the paradigm must be reversed, focusing on what are the most useful logs for the analysis, select them in accordance with objective criteria, organize and prepare them (pre-process) in order to have a repository easy to inspect. The current tendency is to store as many logs as possible without figuring, again, any incident-handling scenario.

The framework, we propose encompasses a whole series of minimal preventive technical arrangements (ex-ante) as well as a series of practical actions, again minimal, delegated to the operators (ex-post), but well defined and tested (Table II) in three main domains (Time, Space and Data). Everything reported in the ex-post column, must be simulated and approved in the "peacetime", in order to allow operators to act with a fast freedom of action.

For instance, the NIST "Computer Security Incident Handling Guide", with specific reference to the Chapter named "Handling an Incident", provides advices regarding logistics, organization and tools, fundamental to the success of incident response programs. In particular, in line with the examples so far treated, the section related to "Incident Analysis", considers the "Log Retention Policies" as extremely helpful and suggests frequent log reviews for improving knowledge. However, this does not give any practical measure for extracting knowledge about an attack, as better described in the following example.

The green series shown in Figure 3, describes how the infection trend described by the blue series can change by applying some of the elements in Table II. We imagine that the attack starts with a phishing email received by a workstation of a branch office (series blue in Figure 3), in order to install a malware that tries to contact its command-and-control center on the Internet. We imagine also that the LAN of the branch office is equipped with a probe (ex-ante measure in the space domain) and that there is only one Internet gateway whose logs can highlight the Internet domains contacted by the internal network (two ex-ante measures in the time domain).

At some point, the security infrastructure/organization will detect the presence of an infection. Assuming that there is still been no escalation, retrieving the Internet domains contacted by the infected machines and acting on the probes, we will be able to quickly identify infections and successful malicious actions, as well as isolate the workstations/LAN (green series in Figure 3).

In the absence of the three above-mentioned measures (one in the space domain and two in the time domain), the time needed for localize and isolate the infection will be longer and will be higher the probability that the attack can scale to higher-value objects. Of course, the use of one or more of the measures provided by the framework in Table II, must be preceded by a trade-off evaluation between the cost of the measure and the cost resulting from its non-application.

5 Conclusions

Our proposals (minimal), intend to shift attention on the incident response stage, preparing first the rules, regulations, techniques, procedures, skills and tools to avoid having to approach the issue only after an incident. The work done, or at least organized, before, lengthens the time of the attack and reduces the time of response, bringing balance on the field, but also can be used as a "gym" for forming the awareness, developing the know-how and raising up the motivation of people involved in cyberdefence. The framework proposed in this document requires efforts and investments at various levels that may appear superfluous in normal conditions (peacetime), but assumes importance if contextualized with the recovery difficulties of a cyber incident, managed with general-purpose plans, with inevitable impacts on the exponential growth of time and costs.

References

- Baldoni R., Montanari L. (2016, February). *A National Cyber Security Framework*. Available via: <http://www.cybersecurityframework.it/en>
- Brady Robert M., Anderson Ross J., Ball Robin C. (1999, September). *Murphy's law, the fitness of evolving species, and the limits of software reliability*.
- Bottazzi G., Me G., (2016, June). *Cybercrime-funded terrorism and the threats posed by future technologies*. NATO Advanced Research Workshop, Terrorists' Use of the Internet: Assessment and Response. Available via www.cyberterrorism-project.org
- Europol, Report (2015). *The Internet Organised Crime Threat Assessment (IOCTA)*.
- Kaspersky (2016, January). *Newly discovered BlackEnergy spear-phishing campaign targets Ukrainian entities*. Retrieved from <http://usa.kaspersky.com/about-us/press-center/press-releases/2016/newly-discovered-blackenergy-spear-phishing-campaign-targets-uk>
- McAfee Report (2014, June). *Net Losses: Estimating the Global Cost of Cybercrime*.
- NIST (2012). *Computer Security Incident Handling Guide*. Retrieved from <http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf>
- NIST (2014). *Framework for Improving Critical Infrastructure Cybersecurity*. Available via: <https://www.nist.gov/cyberframework>
- Ruefle R., Dorofee A., Mundie D., Householder A. D., Murray M., Perl S. J. (2014, Sept./Oct.). *Computer Security Incident Response Team Development and Evolution*. IEEE Security & Privacy (Volume: 12, Issue: 5, Sept.-Oct. 2014).
- Securlist blog (2016, March). *Hospitals are under attack in 2016*. Retrieved from <https://securelist.com/blog/research/74249/hospitals-are-under-attack-in-2016/>
- Steinke J., Bolunmez B., Fletcher L., Wang V., Tomassetti A. J., Repchick K. M., Zaccaro S. J., Dalal R. S., Tetrick L. E. (2015, Jul./Aug.). *Improving Cybersecurity Incident Response Team Effectiveness Using Teams-Based Research*. IEEE Security & Privacy, vol. 13, no. , pp. 20-29, July-Aug. 2015.
- TENACE Project (2014, March). *Critical Infrastructure Protection: Threats, Attacks and Countermeasures*. Available via: <http://www.dis.uniroma1.it/~tenace/index.php?lang=eng§ion=0>
- The Hacker News (2016, January). *602 Gbps! This May Have Been the Largest DDoS Attack in History*. Retrieved from <http://thehackernews.com/2016/01/biggest-ddos-attack.html>
- The Jerusalem Post (2016, January). *Israel's electrical grid attacked in massive cyber attack*. Retrieved from <http://www.jpost.com/printarticle.aspx?id=442844>
- Wendzel S., Zwanger V., Meier M., Szlosarczyk S. (2014, October). *Envisioning Smart Building Botnets*. "Hack In The Box" Security Conference 2014.