

Adapting Enterprise Security Approaches for Evolving Cloud Processing and Networking Models

Andrew Hutchison 

Department of Computer Science, University of Cape Town,
Cape Town 7700, Republic of South Africa
hutch@cs.uct.ac.za

Abstract. With the advent of public cloud services, enterprises are moving to adopt the lower cost, more flexible, scalable public cloud offerings like OTC, AWS and Azure. Simultaneously they are adapting their network models to move away from centralized enterprise QoS networks (with internet breakout from a single or few large enterprise gateways) in favor of lower cost, direct offloading of corporate traffic from company locations via local and distributed Internet service providers. Using this model enterprises are also accessing cloud services from multiple entry points, and this completely changes the enterprise security deployment landscape. As an additional ongoing trend, the networking of physical devices is bringing a whole new ‘operational technology’ domain to the enterprise space, and a new approach to enterprise security is therefore required. In this paper the drivers of change in approaches to security for public cloud computing are presented, considering also the responsibilities of the customer and of the cloud service provider and the component which enterprises still need to focus on. In addition, the network model for security is explained and considered, with the new distributed deployment zone for security as described. Cyber physical/IoT type systems are also then discussed as an additional security landscape over which enterprises increasingly need to take special care.

Keywords: Enterprise security · Hybrid cloud · Local internet breakout

1 Introduction

1.1 The Changing World

It is clear that there is a large momentum in enterprises to embrace cloud based processing models, in contrast to having their own infrastructure. With different layers being virtualized, organizations are embracing software, platforms and infrastructure as services from various providers. Since many of these cloud services are accessed via Internet paths, often with multiple entry points, there is increasingly less imperative to route all corporate traffic back to main data center locations, or processing hubs of the organization, since increasing parts of the workload are serviced directly to distributed locations. In this sense there is a trend to ‘offloading’ corporate traffic from more expensive, Quality of Service based MPLS networks and instead to route some of the processing requirement directly to cloud providers and services via local Internet links.

Another trend which is growing is the digitization of processes, and the connecting of all sorts of devices and sensors into the enterprise landscape. This introduces a further class of traffic and processing requirement, which is also typically in line with the processing and network model described.

From an enterprise security point of view, this evolving picture changes the enterprise security landscape quite considerably. Current centralization of processing has meant that most enterprise traffic has been routed to few processing locations, and large security hubs have typically also been co-located to ensure that incoming and outgoing traffic is inspected and marshalled in various ways to achieve security objectives and organizational integrity. The emerging situation described means that traffic is likely to depart (and enter) the enterprise from many different points – and this changes the way that security needs to be considered in this vastly expanded and distributed landscape.

In terms of the actual migration of processing to cloud based services, there are also additional and new security requirements for organizations. Cloud providers do not necessarily, for example, provide secure Operating System images or basic security management beyond the raw virtual machines which are provided. Organizations also need to link the cloud processing models into their application architecture, so topics like identity management, access control and confidentiality/integrity still need to be realized in a holistic way across these new landscapes as well.

With the addition of new types of device (increasingly including cyber-physical systems, likely representing the operation technology areas of an organization) there is a whole new class of device and connected entities to consider in the security space too.

This paper is structured such that each of the considerations (processing model/network model and expanded processing components) is discussed further and considered in terms of security implications.

2 The New Cloud World

2.1 Hybrid Cloud

With local virtualization having existed for some time, the next step in our computing evolution has been remote virtualization through private, and increasingly public, cloud services. The implicit security and availability of cloud services is increasingly considered adequate by enterprises for their processing requirements. With regard to private cloud services, there have been service providers who over the last decade or more have already been providing shared services accessed by open networks – although often for closed user communities. These so called private clouds have the advantage that customers can to some extent tailor the requirements, and have more participation in the configuration, establishment and operation of the cloud service. With public clouds becoming more scalable, flexible and cheaper than private clouds, organizations have started to embrace this model of ‘market services’ as opposed to having their own cloud communities or customized environments. On the one hand this is understandable, as it provides endless scalability and dynamic addition or removal of capacity based on the large economies of scale of the cloud providers. But on the other hand it introduces a new and de-coupled architecture for cloud based applications. It is widely acknowledged

that to really achieve the benefit of cloud based applications there should be a (re-)architecture to support this and generally just ‘migrating’ applications to the cloud is not the most effective approach for leveraging the full benefits of cloud processing.

It is not the intention of this paper to focus on the security of clouds *per se*, but rather to consider the technological and organizational implications on enterprises which may be moving towards this mode of processing.

In Fig. 1 the AWS cloud service is used as an example to illustrate processing responsibilities which are provided in the cloud service – in contrast to those which need to be addressed by customers. It is clear that there is an extensive customer responsibility for dealing with different aspects of customer data, platform & application management, OS/network/firewall configuration and both client and server side encryption, integrity and authentication. In addition, network traffic protection needs to be incorporated, as applicable. And in terms of identity and access management, this task still needs to be managed by customers of the cloud service as well.

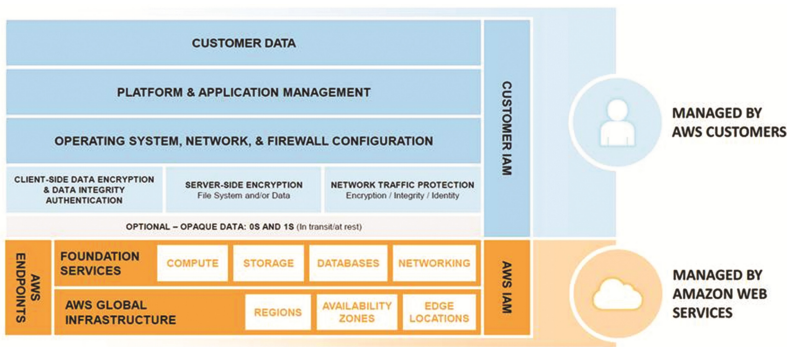


Fig. 1. AWS cloud service areas of customer and provider responsibility (Source: Amazon).

The intention here is not to provide answers or approaches for each of these items in particular, but rather to reinforce the point that enterprise responsibilities for security do not just vanish with the adoption of cloud processing – particularly at the level of Infrastructure as a Service (IaaS). With Platform- and Software as a service solutions there may be more consideration of the security concepts and built in mechanisms, but at an infrastructure and processing level there is still a lot of augmentation and integration which is required. It is considered to be the case that most current cyber-attacks are against the “blue” boxes and not the “orange” boxes of Fig. 1 – which is in fact the area which is *not* under the cloud service provider’s responsibility.

In [1], for example, the situation regarding “AWS Customer Security Responsibilities” is made very clear in the following important text (*italics added for emphasis of key points*): “With the AWS cloud, you can provision virtual servers, storage, databases, and desktops in minutes instead of weeks. You can also use cloud-based analytics and workflow tools to process your data as you need it, and then store it in the cloud or in your own data centers. Which AWS services you use will determine how much configuration work you have to perform as part of your security responsibilities.

AWS products that fall into the well-understood category of Infrastructure as a Service (IaaS), such as Amazon EC2 and Amazon VPC, are completely under your control and *require you to perform all of the necessary security configuration and management tasks*. For example, for EC2 instances, you're responsible for *management of the guest OS* (including *updates* and *security patches*), any application software or utilities you install on the instances, and the *configuration of the AWS-provided firewall (called a security group) on each instance*. These are basically the same security tasks that you're used to performing no matter where your servers are located.

AWS managed services like Amazon RDS or Amazon Redshift provide all of the resources you need in order to perform a specific task, but without the configuration work that can come with them. With managed services, you don't have to worry about launching and maintaining instances, patching the guest OS or database, or replicating databases - AWS handles that for you. However, as with all services, *you should protect your AWS Account credentials and set up individual user accounts with Amazon Identity and Access Management (IAM) so that each of your users has their own credentials and you can implement segregation of duties*. We also recommend using *multi-factor authentication (MFA) with each account, requiring the use of SSL/TLS to communicate with your AWS resources, and setting up API/user activity logging with AWS Cloud-Trail*. For more information about additional measures you can take, refer to the AWS Security Resources webpage".

In the above guidance to AWS cloud customers, it is made very clear that there are still many security tasks and activities to fulfil.

The picture becomes even more complex in the case that multiple, or hybrid, clouds are used. In this case there also needs to be a harmonized view to ensure that organizational security policies, requirements and architecture are preserved by the arms-length processing, storage and access approach.

In Fig. 2 a target hybrid architecture is depicted, while the list of security services on the right hand side shows some of the security aspects which need to be extrapolated and integrated for a multi-cloud, hybrid processing model.

Seeing enterprises confronted with this complexity, service provider organizations are advocating models such as shown in Fig. 3 whereby a common security framework is achieved across a collective of Cloud Provider specific security frameworks. Offerings such as a "Cloud Integration Center" show the evolving role of current outsource and private cloud providers, who are showing flexibility in fulfilling the potentially tricky customer requirements and heterogeneous integration tasks which the new world requires.

Overall, organizations embarking on a cloud based strategy need to do a careful business case to ensure that they are not missing important tasks and responsibilities within a cloud eco-system. Simultaneously, technical solutions for the identified tasks need to be defined and the cloud approach established within the identity, authentication, encryption and integrity regimes which are applicable. The Common Security Framework of the enterprise has to be expanded to include the approaches of the different cloud providers, and an assessment should be done on whether the hybrid approach is still consistent with the organizational security objectives and requirements.

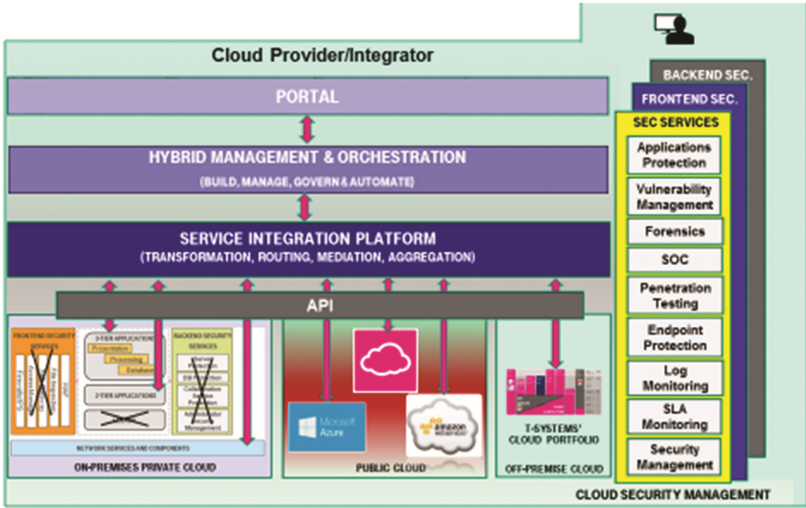


Fig. 2. Cloud security management issues across multiple cloud providers

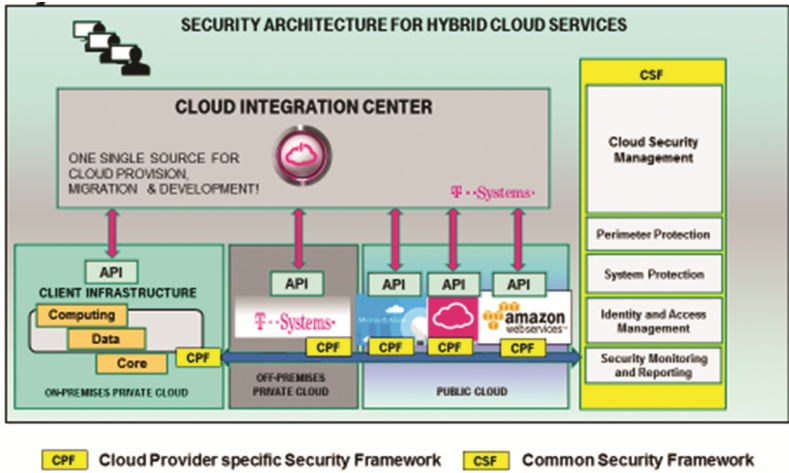


Fig. 3. Security architecture and services to enable a hybrid architecture

3 Pathway to the Clouds

3.1 Offloading and Going Direct

While enterprise Global Area Networks (GANs) and Wide Area Networks (WANs) have existed for many years, the associated requirements of: getting traffic to a central processing point; and ensuring quality of service have resulted in any-to-any type MPLS and other QoS networks being the typical enterprise connectivity model.

With the changing processing landscape, whereby distributed access points to cloud based services become the norm, there is potentially less traffic which needs to go to a central location. It is attractive to ‘offload’ traffic from a corporate QoS/MPLS network, since lower cost local network access can be obtained at each business location with relevant traffic being directed immediately to the (cloud) service access point.

It is clear that the ‘site’ for security shifts from a single, highly centralized concentration of traffic to multiple, decentralized Internet access points. This suggests that an associated security model is required, to ensure that Internet bound (and originating) traffic is inspected and filtered/managed in the same way that it would have been had the centralized single breakout model been implemented.

In Fig. 4 a hybrid IP WAN architecture is reflected, categorizing source locations in terms of criticality. While high and medium criticality locations are connected via both Internet and MPLS, the low criticality locations connect only via best-effort Internet. The paths to cloud providers are also reflected on the right hand side of the figure showing how connections can either be made directly from source locations to cloud providers (and the enterprise data center), or the cloud provider could also be invoked as a processing and storage engine from the enterprise data center.




Fig. 4. Typical hybrid IP WAN architecture [2]

In the Figs. 5 and 6 more details of a typical large, global enterprise are shown with respect to their connectivity models. In Fig. 5 there is a strong emphasis on primary and secondary sites – in line with a centralized processing model, but Fig. 6 reflects how a more direct access to cloud based services can occur – with the associated security elements being identified. These are typically the services which can also be ‘virtualized’ and provided as cloud services to organizations so that the communication paths are filtered and secured.

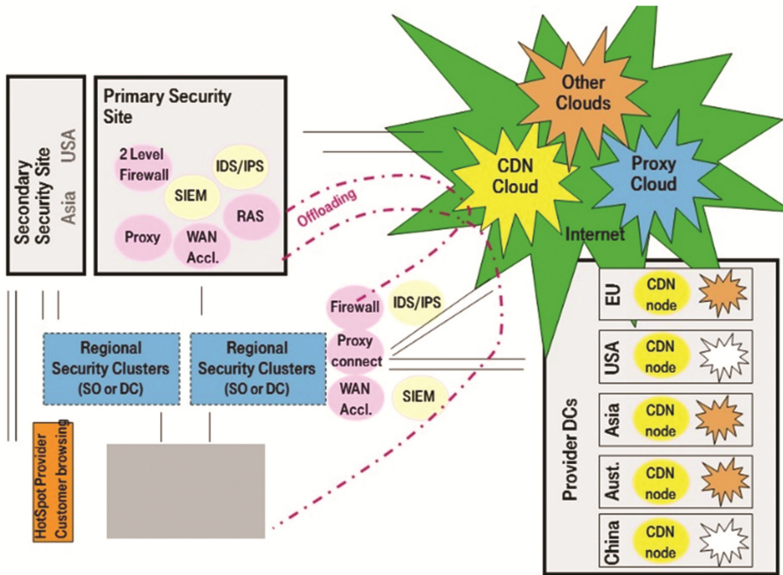


Fig. 5. Typical large enterprise with primary site and regional sites

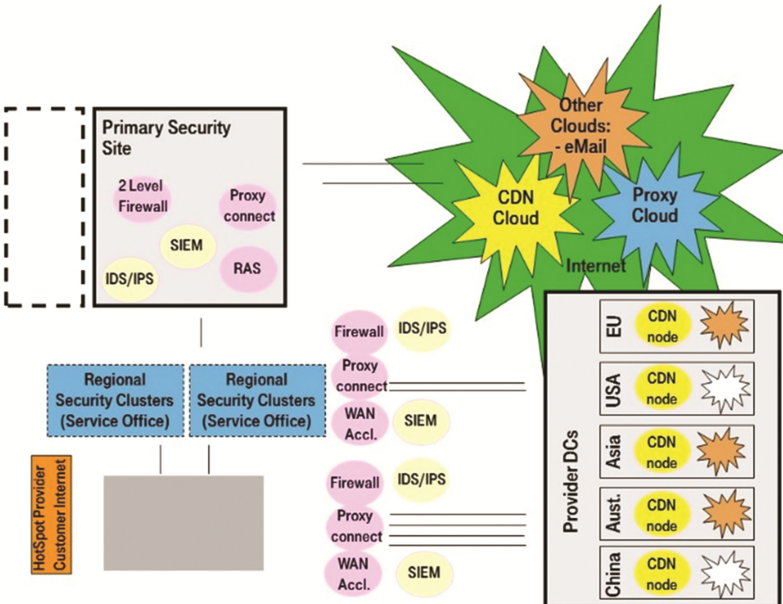


Fig. 6. Example large enterprise with primary site and regional access to cloud services

Content Distribution Networks (CDNs) are also shown in Figs. 5 and 6 and it is also worth mentioning the role which these can play in optimizing and caching content with a great performance enhancement impact. This approach can further reduce the actual amount of unique processing which may be required of processing systems – ensuring that content which is used frequently and widely is accessible directly from local points.

In Fig. 4 it should be evident that the location of the security services has moved from the primary and secondary sites in Fig. 4 to a model whereby security becomes more localized. This situation lends itself well to a cloud based security service, whereby the distributed network paths are secured via a distributed (possibly cloud based) security service which is then ‘inline’ with the revised connectivity and communication model. It is this model which enterprises need to embrace and support as they move toward cloud based processing with associated ‘local’ breakout connectivity.

4 More Things to Deal with

4.1 IoT

The mega-trend to digitization of businesses introduces a whole new class of components into the enterprise landscape. In addition to the typical ‘information technology’ (IT) there is now an additional frontier of ‘operational technology’ (OT) which is either directly or indirectly part of the technology landscape. By direct and indirect is meant the distinction between those organizations which operate specific machinery, vehicles, equipment or processes with physical entities (and therefore process control and other automation activities are implicit) and those organizations which may indirectly monitor physical aspects of their buildings, equipment, supply chains etc. Those organizations with direct machinery and OT environments have the opportunity to manage and monitor these areas in increasingly connected ways. With respect to supporting the indirect business environment, many organizations are deploying sensors and using control software to perform smart management of buildings, vehicles, assets (for example stock, temperature, power, refrigerators, furnaces etc.). With both direct and indirect modes of OT, management and efficiency possibilities are enhanced. Unfortunately a side-effect of this increasingly connected mode is that security consequences are also introduced.

In very many cases physical systems and machines have been designed as standalone, autonomous systems. With hyper connectivity and networking of numerous new types of device the possibility is introduced for external connection to the physical systems. Although this can be very useful for monitoring and managing systems, it introduces the possibility that systems could also be manipulated if an unintended party is able to connect to, and communicate with/control, the physical system. Security principles of authentication, access control, confidentiality and integrity of system interaction are all required – but not necessarily provided – in the inter-connected mode.

The challenge for enterprises is to look very carefully at the landscape of IT and OT environments, and to isolate these via zoning mechanisms which at least ensure a reduced locus of control [3].

Protection of specific cyber-physical communication channels is an emerging area of activity and, for example, SCADA interaction is one particular area of investigation.

But the plethora of sensors and access channels to all manner of devices from cameras to display screens to doors to elevators to medical equipment mean that there is a chance for these devices to be interrogated or controlled by adversaries. Botnet attacks from unexpected sources like security cameras have already been observed.

4.2 End-to-End Cyber-Physical Security

One of the key challenges of security is to provide an end-to-end scope for transactions and interactions which occur. In Fig. 7 the IoT landscape of a large service provider is shown as a sample approach. Different layers from connectivity to service are shown as components of the IoT approach. What is most relevant for this discussion is the positioning of security, which is shown as an end-to-end theme spanning all the other layers of enablement. From the figure it should be evident that the cloud processing and network model discussions of the preceding sections of this paper are also building blocks of this emerging IoT world, and therefore this domain brings together all of the issues and requirements which we have discussed in the previous deliberations on evolving enterprise requirements for cloud and network.



Fig. 7. Example IoT approach and organization (Source: Deutsche Telekom)

With many organizations utilizing Security Information and Event Management (SIEM) platforms to monitor and manage their overall security landscape, it is essential that IoT devices can also be managed within such SIEM approaches. In Fig. 8, the integration of SCADA devices using the SIEM platform AlienVault is illustrated. The general security challenge of such sources and cyber-physical systems in general has been discussed in [4].

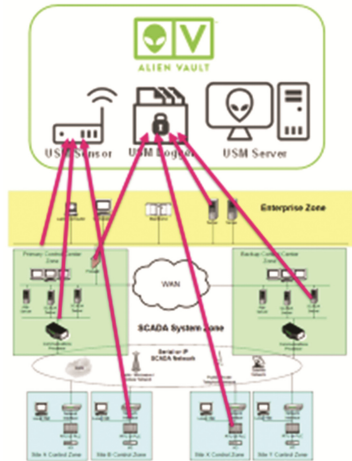


Fig. 8. Example of use of AlienVault SIEM to monitor SCADA system input feeds

Some of the key SCADA security requirements with respect to SIEM integration of events include:

- Asset discovery and management
- Vulnerability Management
- Network and Host intrusion detection
- Log collection, aggregation, correlation and storage

Collection and tracking of these attributes supports the implementation of selected ISA99/IEC62443 requirements as shown in Table 1.

Table 1. Supporting the implementation of ISA99/IEC62443 requirements:

| | |
|---|---|
| FR 2 – Use Control, including | SR 2.8 – Auditable events |
| | SR 2.9 – Audit storage capacity |
| | SR 2.10 – Response to audit processing failures |
| | SR 2.11 – Timestamps |
| FR 3 – System Integrity, including | SR 3.2 – Malicious code protection |
| | SR 3.3 – Security functionality verification |
| FR 5 – Restricted data flow, including | SR 5.2 – Zone boundary protection |
| FR 6 – Timely response to events, including | SR 6.1 – Audit Log accessibility |
| | SR 6.2 – Continuous Monitoring |

5 Conclusion

5.1 Key Findings

It is evident that to achieve a high level of security in a hybrid cloud environment, the aspects of People, Processes and Technology need to be re-visited and re-defined.

While additional options like next generation firewalling/IDS/IPS/DDOS protection/data encryption/log correlation/enhanced incident management with escalation, etc., can be obtained via the ‘marketplaces of services such as AWS and Azure, they typically require acceptance of expensive extended SLAs or purchase of third-party options from their marketplace. And even in this case, many such options are self-managed, i.e. the enterprise still needs to fulfil various aspects or functions.

Based on some infrastructure analyses, it has been approximated that for some large enterprise environments, a complete hybrid cloud landscape - with the same level of security as the current on-premises solution - would cost roughly the same in the public cloud. In addition, many virtual security appliances do not scale as well as hardware appliances, which in some cases may further increase the costs. Many enterprise customers are not really aware of this situation and are attracted by the cheap VM unit costs and basic services without perhaps looking at the full picture.

Another area which is cost driven, is that of the corporate network approach of many enterprises. Instead of long-hauling all traffic to centralized data centers, where processing is done, it is in many instances possible to ‘offload’ traffic from the enterprise Global Area Network to local, direct links to cloud service providers and platforms. Since the opportunity for central policy and rule management is lost, alternative solutions for the local break out traffic are required. Various security service providers have cloud based security solutions which can be used to take over the functions which the central firewall and gateway infrastructure may have been providing in the past. The network strategy of enterprises should be consistent with their evolving processing model, and consider that the “frontier” for security may well be widely distributed across their branches or enterprise locations.

Internet of Things (IoT) devices introduce a whole new security dimension to organizations, since now cyber-physical systems become components of the security landscape. For enterprises involved in any kind of manufacturing, monitoring, production, automation, transport activity etc. there are an increasing number of devices and sensors which are being “connected” and this introduces new attack paths and threat vectors. Enterprises should ensure that systems are adequately protected, and integrated from the outset, for example into the Security Information and Event Management (SIEM) and other Security Operation Centre (SOC) monitoring functions.

5.2 Future Work

More work needs to be done on the identification and development of suitable solutions and approaches for the processing, networking and smart-connection of cyber-physical devices. In this paper we provided a motivation and problem statement, and in this context organizations need to assess specific solutions to see how a distributed, but harmonized, security solution can be implemented to harness the benefits within an orderly and well planned security context.

5.3 Closing Remarks

It is likely that organizations will obtain an increasing proportion of their services from generic infrastructure engines such as AWS, and via myriad networking paths as described. With numerous devices and sensors being incorporated, the landscape has the potential to grow and become even more complex and sophisticated. Considerations such as ‘jurisdiction of processing and storage’ can also play a role in which public cloud provider to select, and this should be yet another consideration when selecting a security partner.

References

1. Amazon Web Services, Inc.: Overview of Security Processes (2016)
2. Munch, B., Rickard, N.: Cloud Adoption is Driving Hybrid WAN Architectures. Gartner (2017)
3. ENISA: Ad-hoc and sensor networking for M2 M Communications – Threat Landscape and Good Practice Guide (2017)
4. Hutchison, A., Rieke, R.: Management of security information and events in future internet. In: Proceedings of Cyber Security and Global Affairs, Conference on Cyber Security, Budapest (2011)

Computer Network Security

7th International Conference on Mathematical
Methods, Models, and Architectures for Computer
Network Security, MMM-ACNS 2017, Warsaw, Poland,
August 28-30, 2017, Proceedings

Rak, J.; Bay, J.; Kotenko, I.; Popyack, L.; Skormin, V.;
Szczypiorski, K. (Eds.)

2017, XIII, 362 p. 141 illus., Softcover

ISBN: 978-3-319-65126-2