

Gaining an Edge in Cyberspace with Advanced Situational Awareness

Vincent Lenders | armasuisse

Axel Tanner | IBM Research

Albert Blarer | Trivadis AG

Organizations that rely on cyberspace as a mission-critical asset require advanced situational awareness to maintain a tactical advantage over emerging threats. A new cyber-situational awareness framework relies on the OODA (observe, orient, decide, act) cycle to provide near real-time cognitive mapping for corporate environments.

Cyberattacks are considered a major corporate and even national threat. Our dependence on cyberinfrastructures is so omnipresent that security incidents can lead to disastrous effects. In 2007, a series of coordinated cyberattacks on Estonian banks, parliament, ministries, newspapers, and TV stations showed that critical parts of an entire country's cyberinfrastructure can be rendered completely unavailable for days.¹ Companies such as Amazon and eBay that offer online services to customers can expect to lose up to millions of dollars per day when their services are down. Perhaps even more dramatic are attacks aiming to steal sensitive data or sabotage critical infrastructure. The 2010 Stuxnet attack resulted in significant physical damage to Iran's nuclear facilities, destroying years of work.

Despite several decades of research on intrusion detection and prevention and billions of dollars of annual worldwide investments in IT security technologies, the threat landscape hasn't changed significantly. Most recent attack reports, such as Red October in 2013, reiterate the high asymmetry between attackers and defenders and the fact that government organizations and companies still can't deal with cyberthreats appropriately.² Even large IT security solution vendors, such as RSA, have experienced cyberattacks that have led to severe damage to finances and reputation.³

Almost all nations have realized that the current cyberthreats must be addressed decisively with a more

holistic approach, and governments have recently come up with national cyberdefense strategies to reduce their vulnerability. Although differences exist among these strategies, especially in the use of offensive countermeasures, all identify the lack of cyber-situational awareness as a key problem in IT infrastructure operation. Unfortunately, many organizations view cybersecurity as a necessary evil and address security by achieving point-in-time compliance to industry and government standards, which still mostly lack the notion of cyber-situational awareness. Therefore, enterprises implement only minimum requirements to pass annual certification. However, to effectively handle cyberthreats, organizations should go further by continuously taking into account current threats, vulnerabilities, risks, and their potential business impact.

Situational awareness has long been a fundamental military capability in warfare; thus, it's not surprising that the first concepts for cyber-situational awareness originated in the military domain.^{4,5} However, basic cyber-situational awareness capability is necessary to master the complex cyberthreats in nonmilitary domains as well. Private companies and critical infrastructure operators are now challenged to implement similar capabilities.

In this article, we propose a cyber-situational awareness framework based on the OODA (observe, orient, decide, act) cycle. Originally developed by Colonel

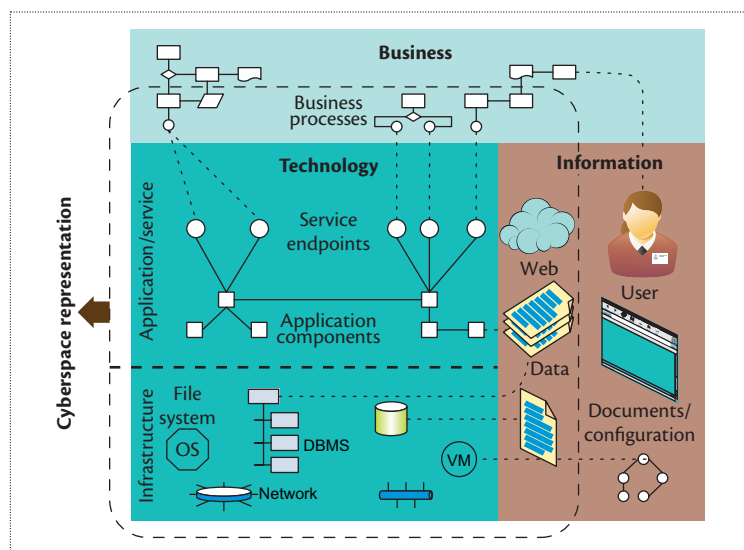


Figure 1. The cyberspace domain consists of technology, business, and information domains. Note that elements of the business and information domains might represent part of cyberspace as we define it, as a great deal of business processes and important information assets are implemented and respectively contained in software elements that might become exposed to unwanted manipulations.

John Boyd of the US Air Force,⁶ the OODA cycle is the classical decision support model used in military operations, providing cognitive mapping from the lowest state level, such as sensor-derived raw information, to the highest level of inference, such as comprehensive analysis of threats and vulnerabilities requiring remediation by human actions. In a typical military system, data sensors observe hostile activity, such as vehicle motion, wireless emissions, or approaching missiles. In cyberspace, the sensors are different, but the concept is similar. Cybersensors observe information flowing in networks from sniffers and extract information from system log files, network management tools, user profile databases, system messages, and operator commands. Human analysts interested in malicious intruders' identity, attack rate, threat, and impact must have a comprehensible view of these activities.

Our framework's novelty lies in its technical realization. Although the OODA cycle has been conceptually suggested to achieve cyber-situational awareness for humans,^{4,5} existing frameworks don't provide a technical implementation at all stages of the OODA loop and thus delegate most of the required tasks, such as sensor data fusion, inference, reasoning, and decision making, to humans. In most large organizations, the technology, business, and information assets of the cyberspace domain are too complex to be captured and understood without technical assistance, and automation tools are necessary to assist human decision makers. Our

framework relies on the use of Semantic Web technologies and formal ontologies to model an organization's cyberassets and their interdependences. We support reasoning with an integrated decision support system that accounts for tradeoffs and uncertainties.

Scope and Objective

Imagine a large company that operates a private cloud infrastructure to provide critical IT services to its business divisions. The cloud infrastructure uses thousands of servers, and millions of clients access those services from the Internet every day. The infrastructure is constantly attacked by malware, malicious outsiders, or insiders aiming to steal sensitive data, corrupt business operations, or achieve financial gains. These attacks might manifest in denial of service, software exploitation, code injections, buffer overflow, phishing, data exfiltration, and so on. This scenario is today's reality for many corporations, which need to detect and react to attacks to mitigate business impact and guarantee uninterrupted, safe, and secure business operations.

The Cyberdomain

To achieve cyber-situational awareness, we must first address the cyberspace domain's scope and boundary: what are the key elements and components to consider? The term "cyberspace" is used broadly to describe the virtual world of computers and the networks connecting them. In this article, cyberspace refers to three highly interconnected domains, as Figure 1 depicts.

At its core, the *technology* domain corresponds to the physical infrastructure, such as networks, switches, and routers, as well as the software infrastructure, including operating systems, virtualization environments, and file or database management systems. Middleware, applications, and services complete this domain. The other two domains are *business* and *information*. Business processes are increasingly implemented as software that can be transformed automatically into executable service code. Therefore, business process models might become exposed to unwanted manipulations like any other software element in cyberspace. Similarly, information assets are increasingly becoming targets of external attacks and manipulations. Status and usage of documents, structured data in databases or data warehouses, knowledge modeled with ontologies, and human expertise contribute to situational awareness and should thus be included in our cyberspace representation. However, not all elements of the business and information domains are part of the cyberspace domain (indicated by the dashed border in Figure 1); there will always be important physical assets and processes outside the cyberspace domain and essential knowledge internalized only in people.

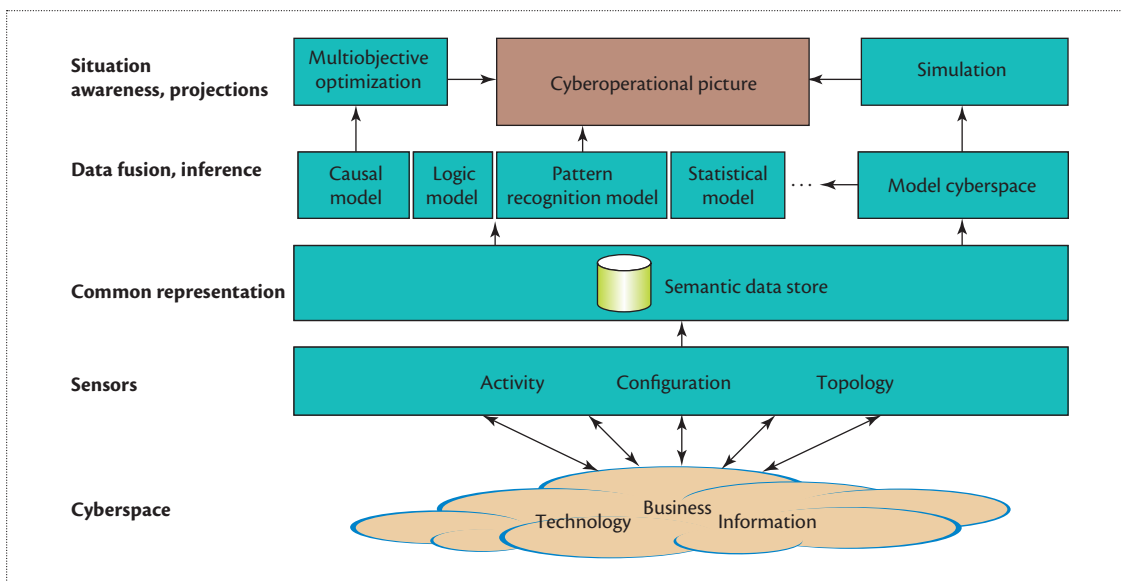


Figure 2. Information process chain to achieve cyber-situational awareness. The state of cyberspace is assessed with sensors that continuously or reactively monitor the environment. Situational awareness is achieved by means of data fusion and inference on the data that is stored and expressed in a common representation format, allowing optimizations, simulations, and visual representation of the overall state of the system to support decision makers.

From Sensors to Situational Awareness

According to Mica R. Endsley, situational awareness comprises three steps: perception (level 1), comprehension (level 2), and projection (level 3).⁷ Level 1 is the perception of elements in time and space that are important to a particular decision maker. Level 2 involves a synthesis of disjointed elements from level 1 that need to be understood in the context of the decision maker's role to make a sound decision. Level 3 is the projection of that understanding into the future to predict the impact of those elements in the context of the decision maker's current situation.

In general, situational awareness implies human awareness. However, full awareness of the state of today's typical complex cyberspace is impossible without supporting technology and automation. It's in this sense that our framework aims to support the system's perception, comprehension, and projection for human decision makers.

Figure 2 shows our framework's information flow model. At the lowest level, sensors pick up activity, configuration, and topology information from the technology, business, and information domains to achieve level 1 situational awareness. The challenge is combining the data from the variety of sensors in different formats, ranging from network flow records to usage statistics and topology graphs. Sensor data must be normalized, cleansed, and transformed into a suitable common representation format on the syntactic level. This data is stored in a semantic data store as basic facts about the cyberinfrastructure's history and the current state.

To support level 2 situational awareness, elements from this common representation are linked in the semantic data store to provide semantic meaning and inference capabilities on the cyberinfrastructure components and assets. Data fusion therefore creates an overall model of the cyberspace domain, including the state of and dependencies among its components, using the data in the semantic data store. Further inferencing capabilities act on incomplete or conflicting information to propagate the components' problem state to higher levels of dependent services and form a cyberspace model that correctly represents the state and dependencies over time.

Finally, level 3 situational awareness enables projections into the future using the representation cyberspace model in combination with specialized abstract models. Visualization of the cyberinfrastructure in a cyberoperational picture allows stakeholders quick understanding, analysis, and decision making. Various underlying models are also used for more complex, multiobjective optimization problems and projected simulations of the cyberinfrastructure's future development to quickly and efficiently support decision makers in achieving situational awareness.

Situational Awareness Framework

As previously stated, the OODA loop is the classical decision support model used in military combat operations. According to Boyd, an entity that can process this cycle more quickly than its opponent can get

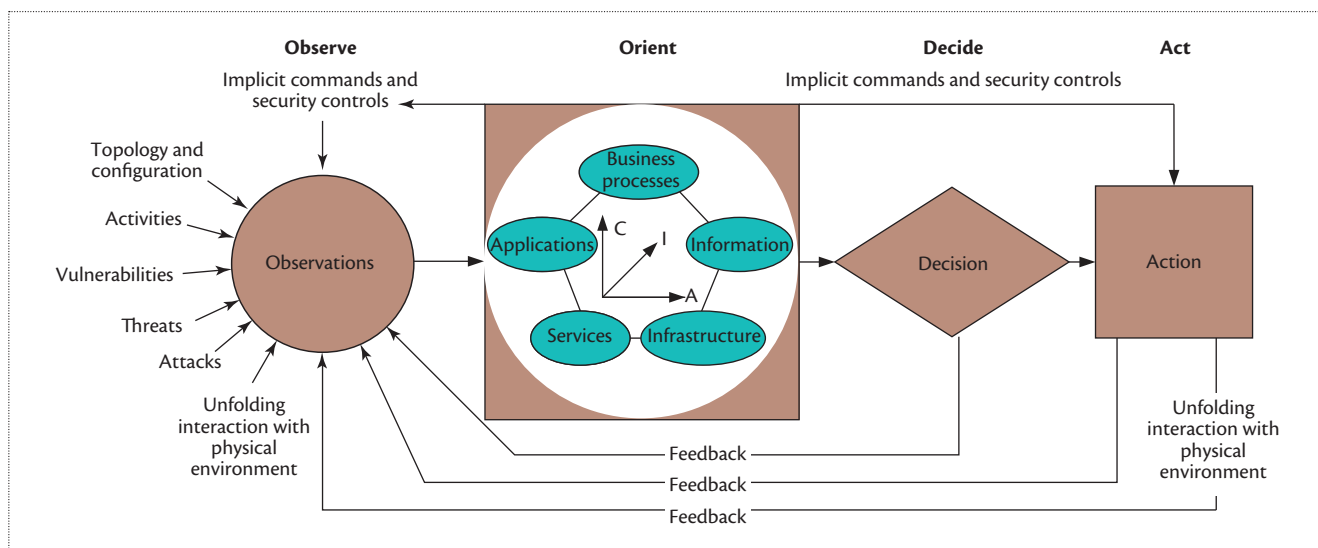


Figure 3. OODA (observe, orient, decide, act) loop for advanced situational awareness in cyberspace (inspired by Christian Sorensen’s “Cyber OODA: Towards a Conceptual Cyberspace Framework”⁴). The goal is to process this cycle quickly by observing and reacting to unfolding events more rapidly than an opponent and therefore gain the advantage.

ahead of the opponent’s decision cycle and thus gain the advantage.

Figure 3 shows the four steps of the OODA loop. Note that the “loop” is actually a set of interacting loops in continuous operation. Therefore, in complex processes with multiple decision makers, the situation is much more complex than the figure indicates, as decisions involve a team of people observing and orienting simultaneously on different system levels and timescales.

For example, typical decisions are needed to perform remediation actions in service outages or system compromises. According to the model, observations are first made by collecting raw information from sensors, such as cyberspace topology, activity, or security sensors. This information is then processed during orientation to provide the status of the cyberassets, including infrastructure, services, applications, information, and business processes. Feedback is possible at all stages of the loop if more observations are necessary for better orientation and more informed decision making.

Cyberspace is closely coupled with the physical environment; actions in the cyberspace domain might directly impact the external environment, and vice versa. In Endsley’s definition of situational awareness, perception can be viewed as part of the observe phase, whereas comprehension and projection are more related to the orient and decide phases of the OODA loop, respectively.

Observe

The observables of interest to achieve cyber-situational awareness far exceed assessing point-in-time compliance and must include continual and detailed insights from

technology, information, and business domains. At a minimum, we need a holistic view of the following factors:

- topology and configuration,
- user and administrator activity,
- implicit commands and security controls,
- software and service vulnerabilities,
- current threats,
- ongoing attacks, and
- unfolding interaction with the physical environment.

These features must be monitored continuously at the device, service, and application levels using a range of cybersensing and automated auditing technologies. For example, deployed topology and configuration are discovered through configuration management databases (CMDBs), Simple Network Management Protocol traps, access policies, process models, and dynamic discovery tools such as nmap or traceroute. User activities are captured by logs including logon times, service access statistics, and network monitoring. Host-based intrusion detection systems (antivirus software); network-based intrusion detection systems, such as Snort (<http://snort.org>); and incident tickets report attacks and threats. Automated scanners, such as Nessus (www.tenable.com/products/nessus), discover vulnerabilities.

Information from sensors can be represented in different semantics and produce huge data volumes. Therefore, sensed information must be fused and aggregated, reducing the initial data while maintaining the ability to drill down and retrieve information on demand, which is the goal of the orient phase.

Orient

In the orient step, we focus on analyzing the meaning of the observed activities with the overall goal to determine the impact on the security state in the near future. The key dimensions of cybersecurity for this purpose are confidentiality, integrity, and availability (CIA). Information collected in the observe phase is combined in the orient phase to provide suitable metrics for the individual components across these three dimensions.

Although the cybercomponents' availability is generally much easier to assess than confidentiality or integrity, security sensors, such as virus scanners and intrusion detection systems, might provide appropriate indicators of individual hosts' and services' security states. For example, when an antivirus software raises security alarms, we might infer that the host's integrity and confidentiality are compromised. Security state is generally transitive—that is, when a host's operating system is compromised, the files and processes running on that host might be compromised. Similarly, all encryption keys this host uses to access remote services might be compromised.

To enable continuous security tracking, the system must constantly fuse sensor information from the observe phase. In this context, there's a tradeoff between pulling sensor data on demand and having the sensors push the most recent updates. With the multitude of sensors and the large amount of data, the latter might result in a huge overhead. However, we can combine both strategies for optimal orientation depending on the deployment context.

Fusing sensor information is challenging as it requires a suitable real-time model of the cyberspace domain. Over the years, standards and software frameworks have been proposed to model IT infrastructure elements and their dependencies. Of these efforts, the Distributed Management Task Force's Common Information Model (CIM; <http://dmtf.org/standards/cim>) is attractive because it has high development maturity, is in active development, and has a broad reach across the industry. However, CIM is class based, subject to a very strict structure, and not easily extensible.

A more fluid way to capture structured data is the Semantic Web approach, which allows very flexible data modeling.⁸ The Semantic Web's format and technologies have largely been standardized by the W3C Consortium (www.w3.org/standards/semanticweb). Most important, it aims to capture the enormous amount of very general data currently available on webpages for human consumption and convert it into machine-readable formats. Many tools for handling extremely large (Internet-scale) amounts of data are freely available.

We combine the best parts of CIM and the Semantic Web, modeling the cyberspace domain according to a Semantic Web representation of CIM. In the Semantic

Web model, resources are expressed with the very general structure of triples $\{subject\}-\{predicate\}-\{object\}$ to describe all facts and relationships. These triples are standardized as Resource Description Framework (RDF; www.w3.org/RDF) in which subjects and predicates are defined as Uniform Resource Identifiers (URIs) and objects are defined as either URIs relating to another named component or literal values to capture data. This results in a data model expressed in a graph structure. The following example RDF graph shows protocol services used in our reference implementation (in N3 notation):

```
csa:atm_al_svc
  cim:ElementName "ATM Service A1";
  rdf:type cim:ProtocolService;
  cim:ServiceUsesSecurityService
    csa:kdc_atm_svc.
csa:kdc_atm_svc
  cim:ElementName "KDC Service";
  rdf:type cim:KerberosKeyDistributionService;
  cim:ServiceAffectsElement csa:kdc_atm_key.
```

These RDF statements describe an air-traffic monitoring (ATM) service and its dependencies on a Kerberos key distribution service.

In addition to the basic RDF triples representing the data, higher-level standards such as RDFS (www.w3.org/TR/rdf-schema) and OWL (www.w3.org/TR/owl-features) allow the definition of schemas and classes with constraints—that is, ontologies. A standardized query language for these graphs, called SPARQL (www.w3.org/TR/rdfsparql-query), is similar to the database query language SQL but adapted for graph-matching queries. For example, a SPARQL query of the form

```
SELECT ?service
WHERE {
  ?service rdf:type cim:ProtocolService.
  ?service cim:ServiceUsesSecurityService
    ?secservice.
  ?secservice cim:ServiceAffectsElement
    csa:kdc_atm_key. }
```

looks for a pattern of a service (variable `?service`), which is itself a `cim:ProtocolService` and depends on a security service (variable `?secservice`), which in turn depends on the specific Kerberos key `csa:kdc_atm_key`. The query looks for subgraphs matching this pattern and will output all services depending on this Kerberos key as specified by the `SELECT`

statement. In this example, the only result is the service `csa:atm_al_svc`.

This type of structured knowledge representation enables advanced inferencing and reasoning on the data as required for orientation. For example, knowing that a Kerberos key has been compromised, the system can infer that a service using this Kerberos service is also compromised. This type of inferencing can be performed by a simple SPARQL query in which the reasoner reports all services depending on a `cim:KerberosKeyDistributionService` using a compromised Kerberos key. In other words, we can detect the compromise of services indirectly by identifying compromised services on which they depend. We can formulate advanced queries like this to answer questions for the orient and decide phases.

Decide

Decision making is the process of identifying and selecting a course of action to solve a specific problem. In the OODA loop, a decision is based largely on observations of the evolving situation, tempered with classification and implicit filtering of the problems being addressed. Observations and their meanings derived from SPARQL queries of the previous phases establish the input and the derived knowledge of situational awareness for decision making.

Decisions can be complex, with two major obstacles potentially constraining an optimal decision: *tradeoffs*—for instance, among confidentiality, integrity, and availability—and *information uncertainty*.

Assume we observe anomalous network traffic patterns indicating a malware infection on a central network node. Should we isolate this node, endorsing the node data's integrity and confidentiality at the cost of reduced or disrupted availability? What if the intrusion detection turns out to be a false positive? A detailed cost representation of alternative actions and a quantitative credibility measure for the observed situation would be highly beneficial for the decision process.

Tradeoffs. The goal of security is to protect critical information infrastructures' CIA. Tradeoffs among CIA have been described extensively in the literature.⁹ The traditional view of tradeoffs is an economical one: a cost function can determine an optimal proportion of resource allocation and minimize the total cost. Handling several allocation components—three in the case of CIA—results in a multiobjective optimization problem, also called Pareto optimization.^{10,11}

Optimal tradeoff solutions also depend on the system's domain: an online retail business might place high value on availability, whereas retail banking might promote data integrity and confidentiality.

Note that optimal resource allocation might change dynamically if the cyberspace situation alters. An external event—such as the malware attack in our earlier example—might drive an immediate decision to reallocate resources.

In addition, awareness of cross-domain activities might be required when moving from an operational to a more strategic level. For example, the chief financial officer of a company selling online goods will be interested in how an event in the cyberspace domain influences the ability to sell products and make profits.

Information uncertainty. A low degree of belief in observed data is cumbersome in any decision process. Important information is often missing when sensors don't capture the entire cyberspace state; a belief function revealing the plausibility of specific incidences would facilitate the decision process. In recent years, many models have been created to tackle the problem of information uncertainty and improve the degree of belief in sensory data, including Bayesian networks, Dempster-Shafer theory, and fuzzy logic. The majority of these models use Bayesian statistics.

In traditional statistics, probabilities are interpreted as the certainty, or uncertainty, that an event of a random experiment occurs. By contrast, in Bayesian statistics, probabilities reveal the plausibility that an event, such as a security incident, occurs. Most Bayesian inference models allow fusion of data from many different sensory channels, which makes them convenient for sensor inputs.¹² The main problem lies in selecting the appropriate model for a given problem. Model selection depends largely on the knowledge of the probability measures used for the inference¹³ and the different types of uncertainty originating from data incompleteness, data inaccuracy, inconsistency, ambiguity, inaccessibility, and malicious activities.

Act

Two types of actions result from a decision. On one hand, a decision maker might choose to address a problem by implementing direct countermeasures in cyberspace, such as blocking ports on a firewall. These countermeasures will impact the system state and be visible through direct observations from cybersensors.

On the other hand, decisions might lead to countermeasures that result in changes in the physical world. For example, a company might choose to react to phishing attacks by asking its personnel not to open any emails with unknown origin. Such actions can indirectly impact cybersystems if the action in the physical world influences the cyberspace domain. However, it's also possible that such an action has no influence on the cyberspace state. In this case, the OODA loop is

effectively closed via the unfolding interactions with the physical environment at the observation level, as Figure 3 shows.

Reference Implementation

To test our framework, we designed and implemented a demonstrator with the goal of creating a reasonably complex but still manageable functional playground for testing the ideas and benefits of a cyber-situational awareness system. In this article, we focus on an air-traffic monitoring service that's part of a larger cyberinfrastructure with many interdependent services.

To make the infrastructure realistic from an IT perspective, its components—servers, network devices, applications, sensors, and monitoring tools—are fully implemented, though limited in scope, as virtual machines. This ATM service consists of multiple (simulated) radar systems—redundant data-gathering servers running the main ATM service application, accessible to users through load-balancing middleware. Kerberization of the corresponding services ensures strong authentication.

Observe

To achieve situational awareness on all system levels, we deployed continuous and periodic sensors. We achieved continuous activity monitoring via host monitoring using the open source tool Munin (<http://munin-monitoring.org>) and network flow monitoring with the IBM Tivoli Netcool Performance Flow Analyzer. We monitored the configuration setup using a CMDB and observed system log files for state monitoring (for example, of load-balancing components). We used Nessus to periodically scan all hosts and services for vulnerabilities. Deployed malware scanners and intrusion detection systems sent alerts to a central collector to report security incidents or indicated threats through persistent attacks. Together, these sensors provided the base for a holistic view of the current cyberspace environment.

Orient

Most operational cybersystems lack a central collection and systematic evaluation of monitored information. Thus, decision makers might be unable to make optimal decisions because important pieces of information are unavailable to them. In our framework, sensor data is translated to a common syntax and semantics during the orient phase.

To achieve this, we developed adapters that translate all observed data dynamically into Semantic Web RDF triples and made them available as “Linked Data” (www.w3.org/standards/semanticweb/data) for easy accessibility. We used the CIM's semantic terms for a common vocabulary and added missing terms where necessary. Dependencies among different parts and services are also captured in CIM, represented as RDF triples, and enacted by adding the corresponding logical model in the form of explicit rules in the triple store that represents the “model cyberspace” in Figure 2.

On top of this model, our framework provides visual high-level views as well as drill-down views exposing more details of the infrastructure for analysis of lower-level components and dependencies, if necessary. For this, we implemented a hierarchical representation of the overall system from top-level services and applica-

tions to components involved with all their interdependencies. Figure 4 shows selected views of the demonstrator's cyber-situational awareness console.

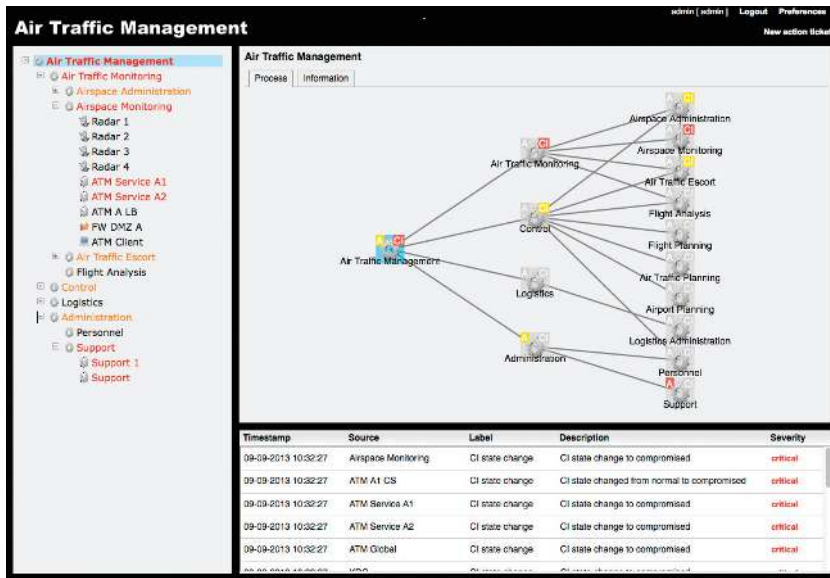
Figure 4a shows the top-level view of

all high-level services and applications; users can drill down by clicking the components. For each component—be it an application, service, or business process in the different views—we assigned the corresponding security state for CIA. Information about the components' dependencies and their security states is inferred through rules and SPARQL queries. The lower right shows a list of recent events related to the selected components' security.

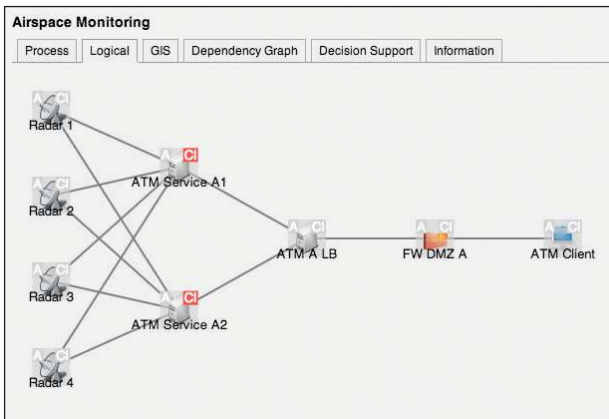
One major difference in traditional monitoring tools is that views are generated by automated discovery, whereas conventional systems require manual work. Hence, views will automatically update when the infrastructure changes. In addition, our RDF-based representation of the cyberspace domain is vendor independent and extensible; therefore, any component can be represented in the model without restrictions.

In addition, we can see the state of services and applications in the CIA security dimensions; for example, Figure 4 shows the status after a key compromise in the ATM service and the unavailability of a support service server. Corresponding services are marked in red in all views to highlight the problems, with “C,” “I,” and “A” indicating problems with the component's confidentiality, integrity, and availability, respectively. Using the dependencies among the services, the underlying logical model spreads these problems to dependent services, which also change to red or yellow according to their importance. These

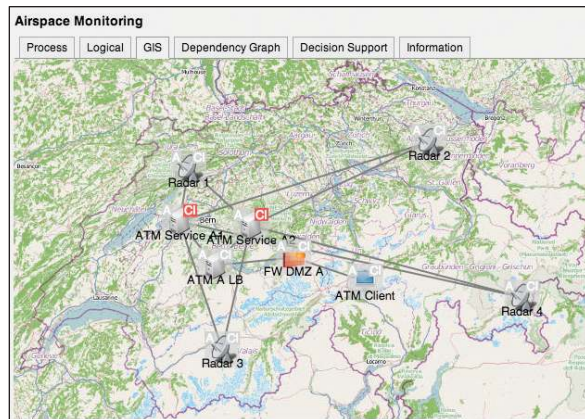
“Awareness of cross-domain activities might be required when moving from an operational to a more strategic level.”



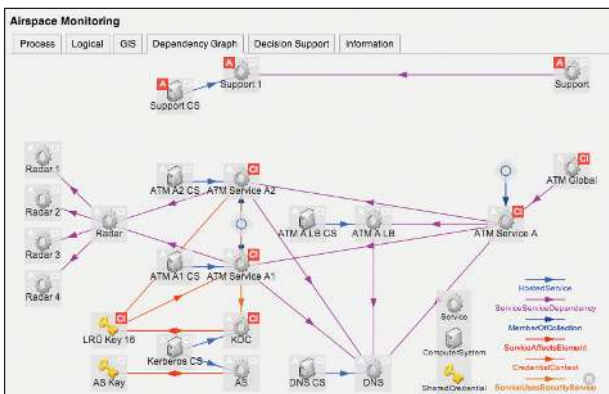
(a)



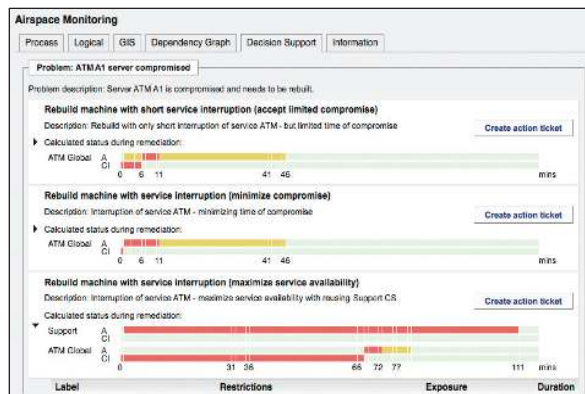
(b)



(c)



(d)



(e)

Figure 4. Cyber-situational awareness demonstrator. (a) Interactive top-level view of services and applications with color codes and confidentiality, integrity, and availability (CIA) icons indicating all services' security states and interactive drill-down views for the Airspace Monitoring service's (b) logical component dependencies, (c) geographic location, (d) dependency graph and CIA compromise chain, and (e) decision support view offering different remediation plans and their security-cost tradeoffs.

dependencies are usually so complex that administrators and decision makers can't appropriately account for the impact of specific problems across the cybersystem. With a formal representation of the cybersystem and their state representation with RDF triples, this inference can be performed using rules and SPARQL queries and thus accomplished by machines instead of humans.

Figures 4b–e show exemplary views (corresponding to different detail tabs in the upper right of Figure 4a) when the Airspace Monitoring element is selected. Figure 4b shows the logical context and connections among the constituting components, and Figure 4c shows the components' location. Again, these views are generated dynamically based on the formal model representation of the cyberspace domain. The benefit of this approach is even more evident in Figure 4d, which shows how the compromise of one ATM server (ATM A1 CS) leads to a compromise of the corresponding Kerberos key and thus the entire ATM service. This type of reasoning enables a rapid assessment of the impact of attacks on other system components to determine business risk.

Decide

As an example of support in the decide step, the demonstrator projects actions of different remedial workflows into future system states. This is particularly important for critical decisions that have alternate remedial plans. Starting with the compromised system state, this decision support function uses the dependency model to calculate the system state during the remedial workflows.

Figure 4e shows the result of such a calculation for an ATM service compromise. The decision support system evaluates workflows created by human experts in the cybermodel and summarizes the projected state of CIA over time during the different workflows, visualizing the impact of each as input for decision makers. In this case, three possible workflows restore the ATM service's full trusted state: the first minimizes the availability interruption (without involving other services), the second minimizes the time of service compromise, and the third aims to maximize availability during remediation by temporarily involving the support service infrastructure (resulting in a longer compromise time for the ATM service and an additional outage of the support service).

During this step, human decision makers must handle tradeoffs and information uncertainty. By visualizing the state over time as projected during remediation, our framework helps decision makers understand the tradeoffs among different action paths, weighing the importance of a highly available service against acceptable business risk resulting from using a compromised service. Although a traditional IT business risk analysis

might give some guidance on which decisions are better for risk minimization, our framework maps higher-layer business decisions to detailed infrastructure components, providing more advanced situational awareness, including cost and security modeling.

Act

By selecting a remedial workflow, decision makers trigger the act step of the OODA loop—either starting automated workflows or delegating tasks to system experts, administrators, or users.

In contrast to existing OODA-based proposals, which leave much of the orientation and decision function to humans, we provide a technical implementation to support all phases of the loop. Although maintaining an up-to-date and valid ontological representation of cyberspace across technology, business, and information over time certainly remains nontrivial, coordinated effort by industry and academia toward this goal is a key prerequisite for achieving situational awareness and gaining an edge in cyberspace over current and future threats. ■

Acknowledgments

We thank Stefan Engel and Marc Stoecklin for their valuable contributions during the project and the reviewers for their very constructive feedback to the article. IBM, Tivoli, and Netcool are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product or service names may be trademarks or service marks of IBM or other companies.

References

1. I. Traynor, "Russia Accused of Unleashing Cyberwar to Disable Estonia," *Guardian*, 17 May 2007; www.theguardian.com/world/2007/may/17/topstories3.russia.
2. GReAT, "Red October' Diplomatic Cyber Attacks Investigation," *SecureList*, 14 Jan. 2013; <http://securelist.com/analysis/publications/36740/red-october-diplomatic-cyber-attacks-investigation>.
3. "Anatomy of an Attack," RSA Fraud Action Research Labs, 1 Apr. 2011; <https://blogs.rsa.com/anatomy-of-an-attack>.
4. C.L.B. Sorensen, "Cyber OODA: Towards a Conceptual Cyberspace Framework," PhD dissertation, School of Advanced Air and Space Studies, Air Univ., 2010.
5. *Cyber Situational Awareness—Issues and Research*, S. Jajodia et al., eds., vol. 46, Springer, 2010.
6. J.R. Boyd, "The Essence of Winning and Losing," Danford, 1995; www.danford.net/boyd/essence.htm.
7. M.R. Endsley, "Toward a Theory of Situation Awareness in Dynamic Systems: Situation Awareness," *Human Factors*, vol. 37, no. 1, 1995, pp. 32–64.



Executive Committee (ExCom) Members: Christian Hansen, President; Dennis Hoffman, Jr. Past President; Jeffrey Voas, Sr. Past President; W. Eric Wong, VP Publications; Alfred Stevens, VP Meetings and Conferences; Marsha Abramo, VP Membership; Shiuhyng Winston Shieh, VP Technical Activities; Scott Abrams, Secretary; Robert Loomis, Treasurer

Administrative Committee (AdCom) Members: Marsha Abramo, Scott Abrams, Loretta Arellano, Lon Chase, Joseph Childs, Pierre Dersin, Lance Fiondella, Carole Graas, Lou Gullo, Christian Hansen, Dennis Hoffman, Samuel Keene, Pradeep Lall, Zhaojun (Steven) Li, Robert Loomis, Pradeep Ramuhalli, Rex Sallade, Shiuhyng Shieh, Alfred Stevens, Jeffrey Voas, and W. Eric Wong

<http://rs.ieee.org>

The IEEE Reliability Society (RS) is a technical society within the IEEE, which is the world's leading professional association for the advancement of technology. The RS is engaged in the engineering disciplines of hardware, software, and human factors. Its focus on the broad aspects of reliability allows the RS to be seen as the IEEE Specialty Engineering organization. The IEEE Reliability Society is concerned with attaining and sustaining these design attributes throughout the total life cycle. The Reliability Society has the management, resources, and administrative and technical structures to develop and to provide technical information via publications, training, conferences, and technical library (IEEE Xplore) data to its members and the Specialty Engineering community. The IEEE Reliability Society has 23 chapters and members in 60 countries worldwide.

The Reliability Society is the IEEE professional society for Reliability Engineering, along with other Specialty Engineering disciplines. These disciplines are design engineering fields that apply scientific knowledge so that their specific attributes are designed into the system / product / device / process to assure that it will perform its intended function for the required duration within a given environment, including the ability to test and support it throughout its total life cycle. This is accomplished concurrently with other design disciplines by contributing to the planning and selection of the system architecture, design implementation, materials, processes, and components; followed by verifying the selections made by thorough analysis and test and then sustainment.

Visit the IEEE Reliability Society website as it is the gateway to the many resources that the RS makes available to its members and others interested in the broad aspects of Reliability and Specialty Engineering.



8. T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific American*, vol. 284, no. 5, 2001, pp. 34–43.
9. C. Ioannidis, D. Pym, and J. Williams, "Investments and Trade-Offs in the Economics of Information Security," *Financial Cryptography and Data Security*, LNCS 5628, Springer, 2009, pp. 148–166.
10. *Multiobjective Decision Making Theory and Methodology*, V. Chankong and Y.Y. Haimes, eds., Elsevier Science, 2008.
11. *Vector Optimization—Theory, Applications and Extensions*, J. Jahn, ed., Springer, 2010.
12. *Data Fusion: Concepts and Ideas*, H.B. Mitchell, ed., Springer, 2012.
13. P. Smets, "What Is Dempster-Shafer's Model?" *Advances in the Dempster-Shafer Theory of Evidence*, R.R. Yager et al., eds., John Wiley & Sons, 1994, pp. 5–34.

Vincent Lenders is a research program manager at armasuisse. He developed the IT security and privacy concepts of the operational C4 and ISTAR systems at the Swiss Air Force. Lenders received a PhD in electrical engineering from ETH Zurich, Switzerland. He's the industrial director of the Zurich Information Security and Privacy Center at ETH Zurich and a member of IEEE and the ACM. Contact him at vincent.lenders@armasuisse.ch.

Axel Tanner is a research staff member in the security group of IBM Research, Zurich. His research interests include the integration of data from multiple heterogeneous sources to enable awareness of the overall state of an ICT environment as well as secure digital identities. Tanner received a PhD in physics from the University of Zurich, Switzerland. He's a member of the ACM. Contact him at axs@zurich.ibm.com.

Albert Blarer is a principal consultant at Trivadis AG, Switzerland. His research interests include software architectures of COMINT and OSINT systems and the analytical problems of such intelligence platforms. Blarer received a PhD in biomathematics from the University of Basel, Switzerland. Contact him at albert.blarer@trivadis.com.

