Cognitive and Autonomic Cyber Defense

Fred Maymir-Ducharme, Ph.D.,
6710 Rockledge Drive
Bethesda, MD 20854
UNITED STATES OF AMERICA

FredMD@us.ibm.com

Lee A Angelelli
6710 Rockledge Drive
Bethesda, MD 20854
UNITED STATES OF AMERICA
LAngelel@us.ibm.com

Douglas W Stapleton
PO Box 506
Bowral NSW 2576
AUSTRALIA
Doug.Stapleton@HInfoSec.com.au

ABSTRACT

This paper describes how Cognitive Computing can be applied to Cybersecurity and extensions of traditional Cybersecurity to address the evolving cyber challenges faced by today's Network-Centric Warfare. Cybersecurity has evolved to address new security challenges introduced by the ubiquitous nature of Cloud Computing, and the multi-variate attributes of Mobile Computing. The authors will further explore related challenges and solutions as the scope of Cybersecurity is extended to defend the military's Situational Awareness communications grid (e.g., Data Links and Satellite Communications) and to protect the mission critical components of Military Platforms used to defend land, sea, air and space (e.g., sensors and weaponry.) This paper will describe novel approaches extending traditional cybersecurity technology used to identify known threats (e.g., signatures and models used to protect against unauthorized access, denial of service, and the corruption of the system's integrity), to also identify "unknown-unknown" threats and vulnerabilities (e.g., False Negatives). The authors will expand on cognitive computing techniques aimed at improving or enabling machines to "learn," "understand," and when appropriate "autonomically act" to defend against nefarious attacks and actions – including Cybernetics (e.g., Decision Support Tools and Automation) required to protect Unmanned Systems. The above discussions with be framed around NATO Standardized Agreements (STANAGS) and Network-Centric Warfare Cyber Defence.

1.0 INTRODUCTION

The military concept of C4ISR (Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance) represent the key elements of National Defense. At its simplest level, the C4ISR model illustrated in Figure 1 below depicts the Command & Control managing the military's Land, Naval, Air & Space components, exploiting Intelligence from its Information, Surveillance and Reconnaissance platforms. Implicit in the model is the reliance on Communications; and of course, as the size and complexity of military operations have increased, so has the dependence on Computers. And another observation is that as

messaging technology matured (e.g., transitioned from analog to digital messaging, networks capable of transmitting audio, video and all types of data, wireless networks, etc.), Communication also came to rely heavily on networks and computing.
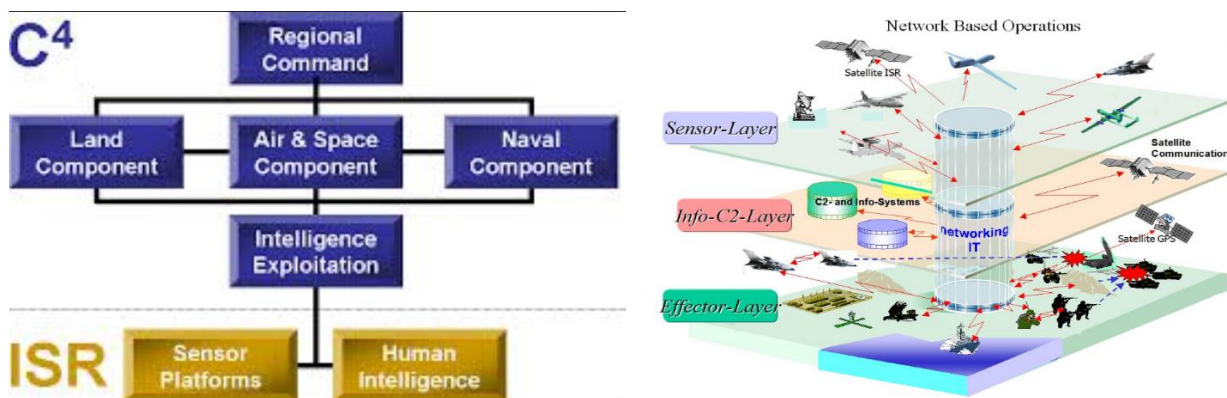
Figure 1: C4ISR Model and Network-Centric Warfare

The term Cyberspace notionally represents the various environments that have evolved to support networked computing across the globe. And the military doctrine of Network-Centric Warfare further highlights the emphasis and reliance on networked computing. But as these new capabilities and innovations have emerged, so have the threats and vulnerabilities.

The traditional approach to security (often referred to as Information Security) relies on mechanisms such as authorization, access controls, system integrity, and assurance of service, which are applied across the three traditional security levels of defense : (1) Prevention to stop threats when and where possible ; (2) Detection & Recovery to dynamically recognize threats that permeate the preventive mechanisms and take immediate action; and (3) Logs & Audits to identify the threats that by-passed the first two lines of defense and reflect existing vulnerabilities requiring stronger security enhancements. Cybersecurity builds on traditional information security to deal with the evolution of Cyberspace as it grows to include very large and complex systems, mobile computing platforms, cloud computing platforms, and an array of sensors and actuators.

> *"The cyber threat that we are seeing is very real and very serious. Sophisticated attacks on networks that underpin our society can be carried out from anywhere in the world. The most effective defences for governments and private sector alike are through enhanced collaboration and trust. A robust NATO Industry Cyber Partnership is key to counter threats we face and increase our collective resilience."* [Mr Koen Gijsbers, General Manager, NCI Agency, September 2014]

Securing a nation's network-centric C4ISR system is very challenging. NATO's cybersecurity is considerably more challenging, especially considering the number of systems from multiple nations that require interoperability and information sharing. This paper explores a variety of evolving cybersecurity capabilities for these very large and complex systems, as well as advanced cybernetic approaches that can then be built on top of these capabilities. And ultimately, the authors of this paper also point to the need to integrate Cybersecurity, Situational Awareness, and the available Intelligence as the figure below illustrates. These systems are traditionally siloed and are managed by different organizations. Integrating these effectively requires advanced analytics capabilities that are able to understand information in different

contexts, continuously and incrementally learn, and cognitively provide decision support or autonomically act in real-time to the tsunami of threats ahead of us. Today's technological advances in informatics and cybernetics make the notion of a Threat Intelligence Platform possible and a very powerful defense solution.
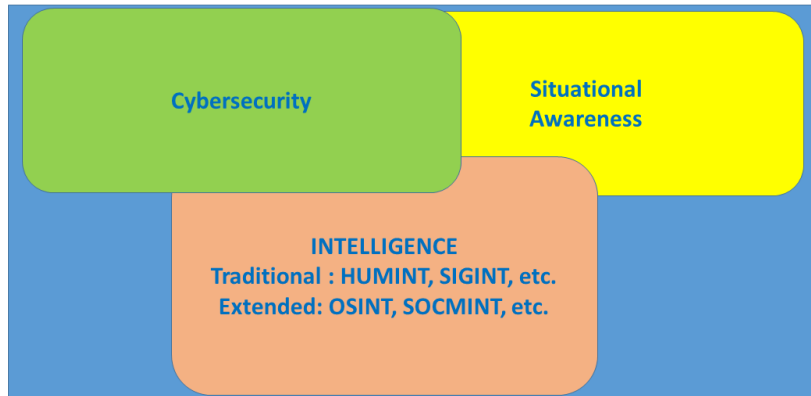
Threat Intelligence Analysis

Figure 2: Threat Intelligence Analysis

Issues in Cyber Defence have gained an increasing share of the research focus of NATO over the last decade. Cyber Defence Situational Awareness (CDSA) is now an urgent need for all NATO Nations. NATO 2020 [14] states that NATO must "*accelerate efforts to respond to the danger of cyber-attacks by protecting its own communications and command systems, helping Allies to improve their ability to prevent and recover from attacks, and developing an array of cyber defence capabilities aimed at effective detection and deterrence.*"

NATO endorsed an enhanced Cyber Defence policy as one of the priorities of a Defence Planning Package at the summit in Cardiff, Wales 4-5 September 2014. The policy reiterated "*that the fundamental cyber defence responsibility of NATO is to defend its own networks, and that assistance to Allies should be addressed in accordance with the spirit of solidarity, emphasizing the responsibility of Allies to develop the relevant capabilities for the protection of national networks.*" Whilst further acknowledging "*that international law, including international humanitarian law and the UN Charter, applies in cyberspace.*" [13]

Recent activity in this area includes:

- An Industry Workshop on Cyber Security Capabilities was held in The Hague, Netherlands, on 24-25 April 2012 under the auspices of IST-096.

- *Future Cyber Defence Concepts and Tools*, a whitepaper by the IST/ET-066 [15].

2.0   CYBERSECURITY IS EVOLVING

2.1   Cybersecurity Today

The United States Computer Emergency Readiness Team (US-CERT) published a report in 2012, stating that there was a 782% increase in cyber-attack incidents over a six-year period. The study cited a steady increase of cyber-attack incidents reported by Federal Agencies, from the 5,500 reported in 2006, to the 48,600 reported in 2012.
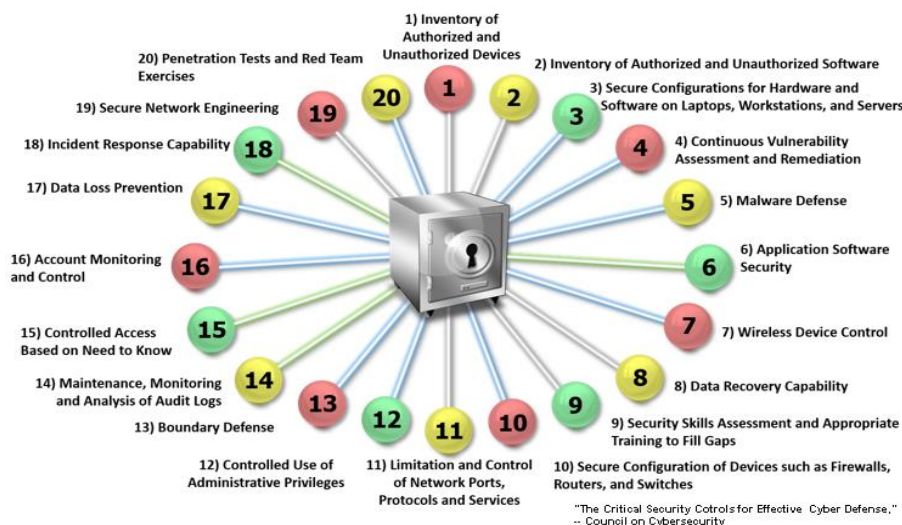
Figure 3: Critical Security Controls for Effective Cyber Defense

The figure above lists many popular security components (mechanisms and solutions) available today. As discussed in the Introduction, there are numerous security mechanisms that can be applied across the traditional three levels of defense to address a variety of vulnerabilities. For example, authentication and access controls can be applied at the protection level to enforce passwords and manage access to data or application assets (e.g., discretionary, mandatory or role-based access) – while an intrusion detection system may monitor at the system level to identify instances in which an entity accesses (or attempts to access) an asset without the proper credentials. Interoperability is key across many of these security components; but interoperability alone is insufficient in addressing the numerous possible vulnerabilities in a system. The authors define a vulnerability as a specific threat to an individual asset; hence, given the broad number of assets (e.g., computers, applications, data, and individuals) and threats (e.g., denial of service attacks, viruses, malicious software, password hacking), one has to consider adding more holistic security solutions that understand the numerous relationships and dependencies between the many security system elements.

The figure below represents a system that is able to interoperate between the federation of security solutions across a system (or across an enterprise), with the ability to exchange data, as well as take action (e.g., set off an alarm or send a system command to the relevant security components). The example system below is also able to recognize known security patterns (e.g., patterns of incidents and events). The system below codifies the various relationships and dependencies between security components and is capable of recognizing and reacting to threats that cannot be secured against by individual components.
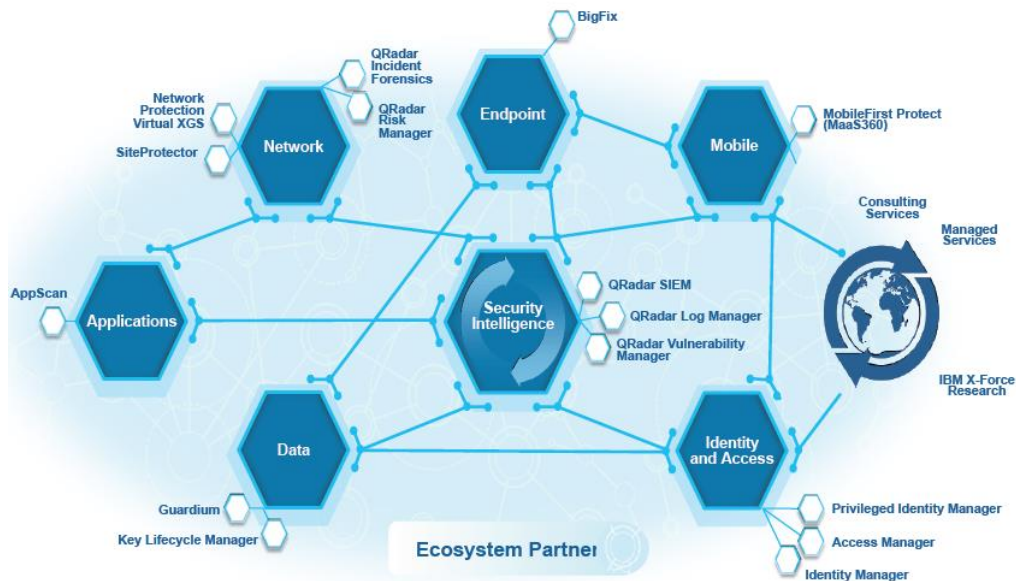
Figure 4: Security Ecosystem

In addition to the security informatics described above, another important focus of cybersecurity is in cybernetics. The authors of this paper define cybersecurity cybernetics as the automation of communications or actions that can be built on informatics. In particular, automating security actions for cases in which a system identifies one of the three levels has been breached, or providing decision support to humans (or interfacing systems) with respect to the cybersecurity informatics.

The figure below represents a military cybersecurity cybernetics system that assesses threats and specifies appropriate actions. In various military organizations, this is an example of the "Observe, Orient, Decide, and Act (OODA) Loop" – which was first described by US Air Force Col. John Boyd.
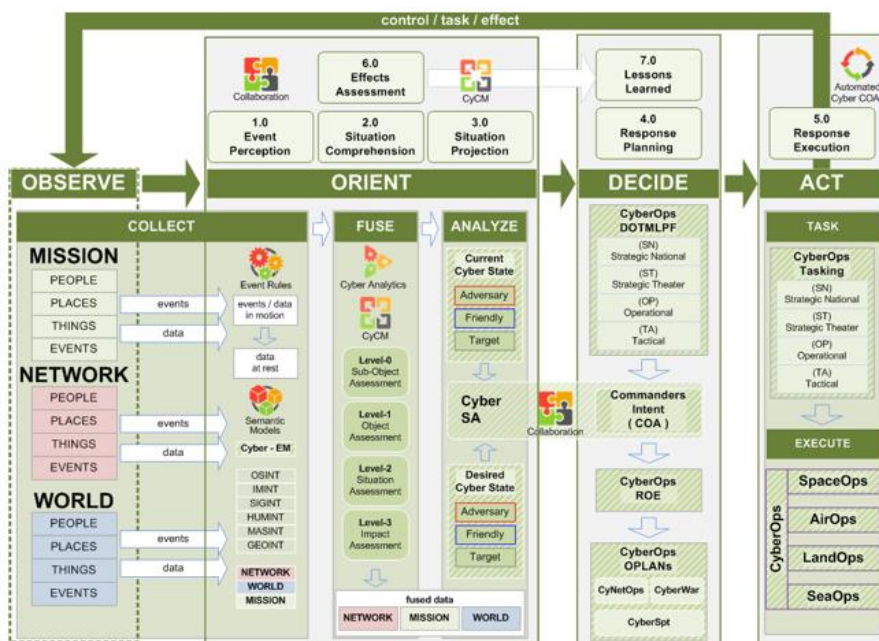
Figure 5: A Military Cybersecurity Cybernetics

2.2 Next Generation Cybersecurity

The next generation cybersecurity will leverage relevant information that's typically external to cybersecurity systems. We discuss exploiting situational awareness information, as well as intelligence from the numerous "X-INT" providers. The Defense industry's notional and evolving Threat Intelligence Platform (TIP) represents such a system. This platform must support a Federation of Systems (FOS), given the dynamic nature of both, the available technology, as well as the new data and data sources. And given NATO's perpetually changing participation and country sponsorship, a federated approach is optimal.
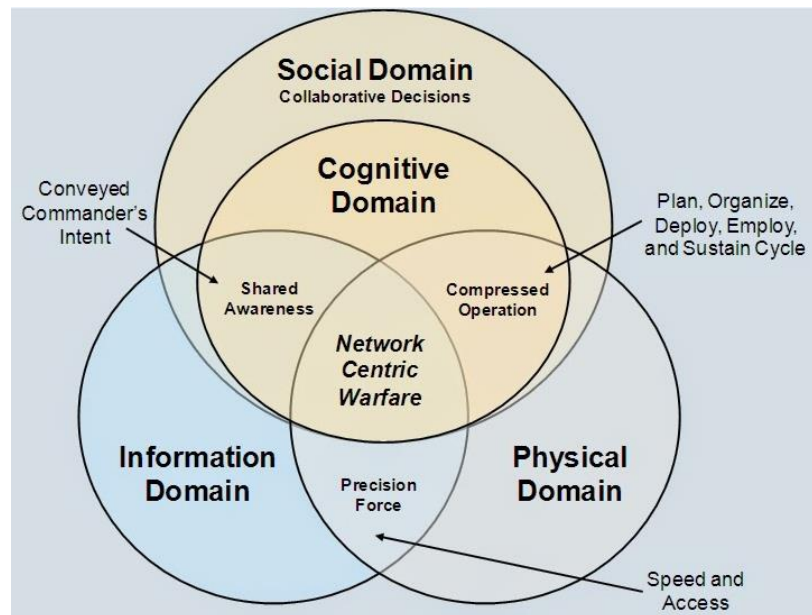
Figure 6: Cyber Threat Intelligence

Cybersecurity scope is evolving and extending to Cyber Threat Intelligence (OSINT, SOCMINT, HUMINT, etc.) which itself is extended to Threat Intelligence Analysis

A Threat Intelligence Platform (TIP) is an emerging technology discipline that helps organizations aggregate, correlate, and analyze threat data from multiple sources in real-time to support defensive actions. TIPs have evolved to address the growing amount of threat data generated by a variety of internal and external resources (such as system logs and threat intelligence feeds) and help security teams identify the threats that are relevant to their organization. By importing threat data from multiple sources and formats, correlating that data, and then exporting it into an organization's existing security systems or ticketing systems, a TIP automates proactive threat management and mitigation. A true TIP differs from typical enterprise security products in that it is a system that can be programmed by outside developers, in particular, users of the platform [Wikipedia]

2.3 NATO Approach to Cyber Resilience

NATO is committed to building up its cyber resilience with an approach being developed in the current SAS-116 *Military Strategic Level Decision Making within a (future) framework of Cyber Resilience*. The preliminary discussions concluded that there is gap in understanding between the technical reality of cyber and the strategic and operational levels of decision making. Cyber activities have increasing impact in all aspects of military operations. These activities can no longer be regarded as solely a technical issue. Also, an uncontested cyber environment (cyberspace) is no longer a credible assumption. While cyber defence will never be complete, NATO is striving to achieve cyber resilience.

There is an urgent need for accurately modelling and simulating cyber events and this was undertaken in 2012 through NMSG-117 (*Modelling & Simulation for Cyber Defence*) which was tasked with investigating the aspects of Cyber Defence that can be supported with Modelling and Simulation (M&S). There was a particular focus on M&S environments for cyber threats, synthetic environments, standards and processes [16]. Current approaches for an exercise coming under 'cyber-attack' are generally to turn off portions of the C2 network in an arbitrary manner which would be wholly unrealistic in a real cyber-attack.

3.0   THREAT INTELLIGENCE ANALYTICS

Cybersecurity is plagued by advanced persistent threats. Today's solutions can barely keep pace with the rate and new variations of emerging threats, which span multiple domains. Intelligence threat analysis needs to deal with very large and complex set of rapidly changing threats and technology our adversaries have at hand. This explosion of threat vectors and exploitable vulnerabilities has forced security experts to extend the traditional scope of cybersecurity across multiple domains and multiple disciplines.

The Defense Industry has experienced a data explosion over the last couple of decades, noting that approximately 20% of the data is structured and 80% is unstructured (unstructured text, semi-structured text, and multimedia). The internet, mobile devices and sensors (e.g., audio & video), are but a few of the "Big Data" catalysts. The US Department of Defense (DoD) collectively has experienced growth in data, with organizations that dealt with petabytes 10 years ago, now having to deal with exabytes. The traditional approach to multi-media analysis typically involves humans manually "tagging" the data after a full or partial analysis. As data grows exponentially, the manual analysis and tagging approach becomes less effective (due to increased human errors) and efficient (due to lack of human resources). In addition to the traditional multimedia (e.g., images, audio and video), there is non-traditional data such as sonar, radar and electro-optics. As the size, volumes, and complexity of these data sets increase, so do the challenges associated with the analysis (e.g., multi-source analysis) and visualization of the data (e.g., aggregations such as mash-ups).

Threat Intelligence Analysis focuses on automated analytics that can be applied to the various types of data (structured & unstructured text, multimedia, and proprietary) – to augment the human analytics and tagging. For example, audio technology can be used to transform speech to text, and one can then apply advanced search and text analytics to the textual information in the audio. Audio technology also provides the ability to distinguish speakers in the voice-to-text transformation – and can authenticate or identify the speaker if the speaker's voice has been captured in a biometric identity management system. Trans lingual technologies (semantic & statistical machine translation, morphological analyses, multi-lingual search, dynamic language identification, etc.) extends these capabilities to support new and evolving languages. In addition to the audio capabilities above, advanced video capabilities such as image recognition, OCR, and a variety of imagery analytics can be applied to exploit information contained in videos. The ability to apply analytics to multimedia dynamically or "on demand" (beyond the static tagging approach) is crucial because of the dynamic nature of "topics of interest" and the challenges posed by massive data volumes and non-linear complexity. Threat Intelligence Analysis also deals with advances in applying analytics to "data in motion," which provides the ability to collect, process, exploit and disseminate information much faster than ever before.

This section describes a variety of advanced technologies (which we refer to as cybersecurity informatics) that provide essential threat intelligence analytic capabilities.

3.1   Stream Processing

The notion of "Big Data" reflects today's situation, in which we now have access to petabytes and exabytes of data – a plethora of information that can be exploited by cybersecurity solutions. These large volumes of data come with a variety of challenges, which we sometimes categorize into "the 4 V's:"

Volume : Today's systems (particularly Cybersecurity systems) can exceed terabytes and petabytes – into exabytes… Deep packet inspection systems can capture and analyze terabytes of data, some easily processing over a petabyte in less than a month. Most of these systems don't store the data for long, recognizing the associated tremendous data storage that is otherwise required.

Velocity : Data is not only being created at a much faster pace than ever before, it is becoming available at a much faster pace. A Distributed Denial of Service (DDoS) attack is virtually limited only by the network

bandwidth. Products advertise the ability to protect against 500GB/s, but an OC-192 pipe can transmit twice that volume in one second.

Variety : The variety of data extends beyond the set of "structured & unstructured text, and multimedia" that most organizations are dealing with today. Marine biologists ingest and analyze petabytes of sonar, creating beam forms and other electronic products that can each be in the terabyte range. And referring back to the smarter utilities data volume example, that industry also has to deal with a variety of different data types as utility metering is transformed from analog meters to digital data for measurement, analysis, and ultimately – optimized management.

Veracity : The accuracy, as well as the provenance of data also plays a major role in analytics. The more important the decision (e.g., mission critical systems, or those with lives at stake) require much higher fidelity and accuracy, as well as assurance (or prioritization) based on the pedigree of the data (e.g., where it came from, who created it, and who has modified it).


Figure 7: Stream Processing

There are other data issues to consider, but one in particular is worth noting. A new paradigm has evolved over the last decade. Stream computing differs from traditional systems in that it provides the ability to process (analyze) **data in motion (on the network)**, rather than **data at rest (in storage)**. Figure 7 above illustrates the parallel nature of stream processing, as well as the "pipe and filter" architecture that allows a variety of analytics to be applied in parallel to data in motion.

Stream computing is meant to augment current data at rest analytic systems. The best stream processing systems have been built with a data centric model that works with traditional structured data and unstructured data, including video, image, and digital signal processing. Stream processing is especially suitable for applications that exhibit three characteristics: compute intensity (high ratio of operations to I/O), data parallelism allowing for parallel processing, and ability to apply data pipelining where data is continuously fed from producers to downstream consumers. As the number of intelligent devices gathering and generating data have grown rapidly alongside numerous social platforms in the last decade, the volume and velocity of data that organizations can exploit has mushroomed. By leveraging Stream Processing practitioners can make decisions in real-time based on

a complete analysis of information as it arrives from monitors and equipment (measurements and events) as well as text, voice transmissions, and video feeds.

Stream processing provides the ability to analyse a variety of data feeds in motion, generating alarms before the data is stored for processing, and transforming the data to the various formats or lexicons that different cybersecurity components' APIs may require. The figure above highlights many other capabilities that we won't belabour in this paper.

3.2 Machine Learning

As described earlier, the authors view vulnerabilities as a function of threats and assets. Today's cybersecurity systems have a reasonable handle on known vulnerabilities – i.e., known threats that have been used against known assets. IPS systems apply various rules and signature analyses to prevent "known-knowns." And IDS systems apply signature-less and other analyses to detect new threats on assets requiring protection, e.g., "unknown-knowns." The notion of "persistent threats analysis (PTA)" leads us to the bigger challenge: "unknown-unknowns."

Machine learning PTA systems such as IBM's Cognitive Cyber Defense (CCD) system have demonstrated the ability to discover unknown-unknowns, and subsequently then updating the IPS and IDS systems accordingly. These tools employ a combination of descriptive (e.g., discovery through affinity analysis) and prescriptive (supervised machine learning) techniques, alongside predictive and behavioural models to identify persistent threats.
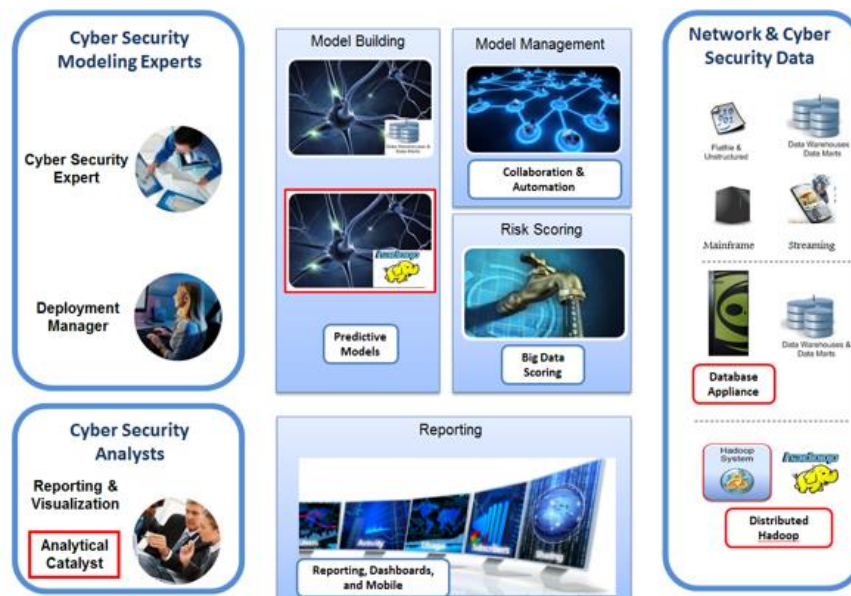
Figure 8: Stream Processing and Machine Learning PTA System

The CCD exemplar system above is a self-contained PTA detector that sits at the gateway, processing Domain Named Server (DNS) response and Netflow data. Feature extraction and model scoring is applied to categorize predicted benign domains (what the system believes is safe), and predicted infected domains (what

the system believes is infected). Forensic analyses are applied to discover false-positives and false-negatives in both domains, which are then used to incrementally fine tune the models.

The CCD system is built on Stream Processing, which allows it to apply a variety of filters such as Black-Lists and White-Lists, Proxy Log mappings and transformations, etc., while going through the Ingestion, Filter & Enrich, Extraction, Classification, and Sink & Visualization phases that apply data-at-rest heuristics.

3.3 Cognitive Computing Theory

A military researcher might pose questions around cyber threats. We need computers to interact and reason over natural-language content in the same way that humans do. That development has occurred over a long period of time and has primarily advanced through techniques of open-domain question-answering (QA) [1, 2, 3].

To have computers reason like a human has required advances in many areas of computer science and artificial intelligence (AI), including information retrieval (IR), natural-language processing (NLP), knowledge representation and reasoning (KR&R), machine learning, and human-computer interfaces (HCIs) [4].

Essentially there is no one computer program that is able to reproduce the subtleties of the human mind and the way it understands language interactions. There are multiple analytical paths, each of which contributes part of the solution.

The challenge in computerising these multiple analytic paths was addressed by an architecture framework for integrating diverse collections of text, speech and image analytics called Unstructured Information Management Architecture (UIMA) [5] developed by IBM and later contributed to the Apache foundation [6] and is in use by industry and academia today.
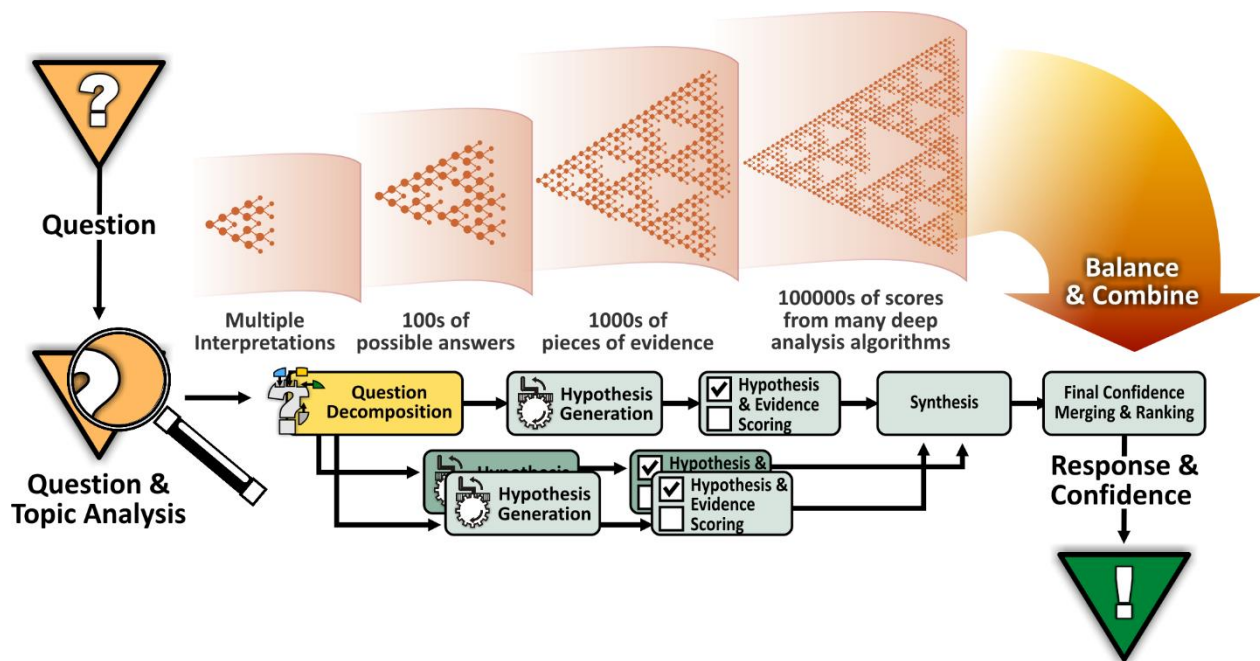

Figure 9: Question and Answer Processing Pipeline

Building on the UIMA architecture, a set of parallel processing pipelines is used to analyse the question and independently pursue possible candidate answers by searching many different resources. Evidence is then gathered for each alternative answer until a final weighting gives a confidence score for presentation of the preferred answer as illustrated in Figure 9.

A feedback loop where the system is trained by experts in their field improves the correctness of the answering process. Context is a key part of that training process; for example, a reference to 'Atlas' would need disambiguation by context to either a member of the Atlas rocket family; a mountain range which stretches across north-western Africa extending about 2,500 km (1,600 mi) through Algeria, Morocco and Tunisia; or a collection of geographic maps. A key advantage provided by that expert training is making the results available to all subsequent queries. Future users interacting with the system will benefit from that expert training, which the system will 'remember' and apply to future questions.

Understanding the question using natural language parsing techniques is the first part of the processing pipeline. The class of thing being asked for is referred to as the *Lexical Answer Type* or LAT [7, 8]. Questions can be further decomposed into sub questions that can be independently answered [9].

The final stage in the pipeline architecture is to bring together the evidence for possible answers and to produce a ranking of the likelihood or confidence in a particular answer being the "correct" one with machine learning assisting in improving the overall ability of the system to rank correct answers [10].

To be practical, we need the performance of cognitive computing to be reasonable. An interesting example is when IBM built a computer system called Watson to compete on the U.S. game show Jeopardy! using the cognitive computing processes outlined above, the first computer attempts took two hours to answer each single question. Although correctly answered, it was hardly a champion performance level. The solution as part of the UIMA architecture was to introduce massive parallelism in the computations [11]. Ultimately 2,880 processors working in parallel on a single question brought the response time down under 3 seconds allowing the Watson computer to challenge and finally win Jeopardy! in 2011 against the two highest ranked human players [12].

3.4   Contextual and Cognitive Analytics

The cyber security technologies of tomorrow that will protect governments from cyberattacks will leverage cognitive computing technologies. These cognitive cyber security solutions combine three main technologies that enable human cognitive thinking: NLP/Information Extraction; hypothesis generation and evaluation; and dynamic learning computing to effectively harness the explosion of unstructured data. These cognitive technologies enable cognitive cyber security systems to ingest and extract entities and relationships from cyber security data sources and are stored in a cyber intelligence corpus that represents that the real-world cyber security domain. Cognitive cyber security systems should apply graph computing to build a knowledge graph from the entities and their relations extracted from unstructured text. The knowledge graph provides the ability to apply graph modeling algorithms to perform multi-inferencing to identify and computationally infer non-obvious relationships spanning over time to understand and correlate events occurring beyond one's observation space. Research work in graph computing is being used to extract structured data into a graph model that represents the values (nodes), entities (child nodes) and relationships (edges) that represent the ontology used to represent the security domain we use for information extraction.

Humans and Cognitive Systems Interactions - "The goal isn't to replicate human brains, though. This isn't about replacing human thinking with machine thinking. Rather, in the era of cognitive systems, humans and machines will collaborate to produce better results, each bringing their own superior skills to the partnership. The machines will be more rational and analytic—and, of course, possess encyclopedic memories and tremendous computational abilities. People will provide expertise, judgment, intuition, empathy, a moral

compass, and human creativity. … cognitive systems will be designed to draw inferences from data and pursue the objectives they were given". To enable this cognitive assistant and human interaction – cognitive cyber security solutions need to go beyond either a simple question-answering, or a keyword search paradigm. Like Watson, they need to use deep natural language processing to shift the human–machine interface to further enhance human abilities.

Use case A typical use case for a cognitive cyber security system is its ability to detect multiple non-obvious cyberattack vectors occurring over time across the cyber security kill chain in different geographic regions using its corpora to hypothesize and drawl conclusions based on evidence that a larger cyber campaign (e.g. Dyre Wolf's spear phishing emails (attack delivery), Upatre establishing command and control (C2), and distributed denial of service (DDOS) attacks (deception) while Dyer Wolf (extracting money from victims) is being launched against an organization(s). Cognitive cyber security solutions are able to understand who are the threat actors, their TTPs, exploited targets, and intended effect(s). Cognitive cyber security solutions are able to understand an organizations IT infrastructure and known related vulnerabilities to identify vulnerable systems and impacts to recommend a prioritized list of course of actions.

A cognitive cyber security system can bring cyber security defense in depth awareness and intelligence to a level previously unattainable by classical security systems. By exploiting deep semantic reasoning and natural language machine learning technologies, enable cognitive cyber security assistants to understand human natural languages. Cognitive cyber security assistants can ingest cybersecurity documents (security lab reports, news feeds, Wikipedia, etc.); then normalize, extract, and represent the ontologies (form) and relationships (function) of entities into a continuous accumulation of cybersecurity intelligence. For cognitive systems to understand real-world cyber security problems and reason to derive the best answer – there are many entities, behaviors and interactions that involve human thinking. A cognitive cyber security assistant uses these integrated technologies listed below to mimic human thinking in its DeepQA pipeline to enable security professionals to ask questions in human natural language and receive direct, confidence-based responses.

- Natural language processing by helping to understand the complexities of unstructured data. Apply advanced NLP parsing and Part of Speech Tagging techniques to determine how entities influence and affect one another in specific situations.

- Hypothesis generation and evaluation (probabilistic computing) by applying advanced analytics to weigh and evaluate a panel of responses based on only relevant evidence. Apply network reasoning algorithms to computationally derive obvious and non-obvious hypothesis, patterns and relationships

- Dynamic learning (machine learning (ML)) by helping to improve learning based on outcomes to get smarter with each iteration and interaction.

These technologies are foundational, without which neither computers nor humans can determine the correct correlation between questions and answers. These technologies provide cognitive cyber security assistants the ability to continuously learn and adapt, and can span tremendous stores of cybersecurity documents (tens of millions). Cognitive cyber security assistants should contain Deep Question and Answering (DeepQA) pipelines like the Watson's Factoid Pipeline. It is referred to as Watson's DeepQA. DeepQA capabilities is the cognitive architecture for answering human language questions. The Factoid pipeline technologies is used by Watson's DeepQA to understand and answer human natural language questions. Once a question is posed, Watson performs the following steps in order to find an answer.
- *Question Analysis* - Analyze the linguistics of the question by decomposing the question using deep parsing and analyzing the question to understand what is being asked and what constraints are being imposed on the answer.

- *Hypothesis Generation* - Generates possible answer candidates by building and running search queries using CAS from QA phase. This phase is also responsible for passage scoring and filtering (Large Aperture Pipeline Framework).

- *Hypothesis Evidence Scoring* - Explores corpus to look for evidence and builds a case for and against each evidence. Performs context dependent and context independent scoring on generated candidate answers.

- *Final Merger & Ranking* - CASs from previous stage are combined. Machine Learning Models are applied to determine best possible answers.

- *Supporting Evidence Merging & Ranking* - This concluding phase retrieves all the evidence that was collected during Primary Search execution. It, then, applies passage model to evidence and derives a ranked list of Answers & Evidence.

3.5   Domain Adaption – Teaching Cognitive Assistants the Cyber Security Domain

As mentioned above, the world of cyber-security can be described as a composition of complex and uncertain inter-connected parts. Architects should apply system thinking (systems dynamics) to understand how to apply cognitive computing to describe and understand, the forces and interrelationships that shape the behavior of cyberattacks and best incident response for detecting, mitigating, and preventing a cyberattacks. **Domain adaption** is a method used to teach cognitive systems like Watson to understand a specific domain. Domain adaptation involves the creation of new unstructured and structured data sources, new ingestion artefacts and training data, and training model tuning, and system processing customizations. Adaptation is an iterative process of experimentation, analysis, and development. The goal of the process is to tailor the system so that it can provide relevant and meaningful information to security professionals in the cyber security domain [28]. In domain adaption, ontologies are used to teach a cognitive system how to understand the concepts and relations of the real-world cyber security domain through ingesting, understanding, and reasoning on unstructured text. As shown in Figure 10 – an ontology represents the cyber security entities, relations, and co-references as they exist in the real-world.
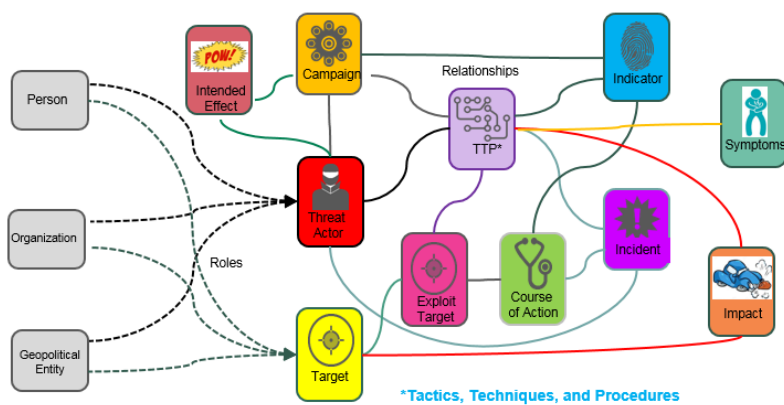

Figure 10 – Cyber Watson Type System

The full cognitive system's DeepQA pipeline discussed above not only includes the ontology (model of the language domain) but the annotators used to recognize mentions, classify them, recognize relationships between individuals, and co-references between like entities of the same instance across the cognitive system's corpus. For cognitive systems – information extraction is a key technology to develop annotators that understand text, as it identifies the important conceptual objects and relations between them in a discovery.

Information extraction provides the key capabilities for cyber security organizations to automate the process of extracting cyber security tacit information from unstructured text. Tacit Knowledge (TK) generally refers to information that is difficult to convey, store, or transfer explicitly [26]. The cognitive cyber security system applies cognitive DeepQA technologies to ingest unstructured data and extract tacit knowledge from cyber security documents to develop a cyber intelligence corpus that represents the real-world cyber security domain ontology (entities and relationships). These technologies allow organization to keep up with the exponentially growth in the amount of unstructured text information and parse the information into structure text that a cognitive system's DeepQA pipeline can reason on the corpus and perform inferencing on computationally derived relationships between entities that's defined in the ontology model.

The information extraction is meant to build classification and/or detection models (i.e. sequential classification model). The distinction here is mostly happening at decoding time: for sequence classification, the classification of an example depends on the classification of the surrounding tokens (i.e. POS tagging, text chunking, Named Entity recognition, mention detection), while for example classification the examples are not dependent on their surrounding examples (e.g. prepositional phrase attachment, medical diagnosis of a patient, etc.). This difference is reflected at decoding time, where search over the possible sequences is needed (e.g., Viterbi or Forward-Backward search) and is not needed for example in classification [27]. Cognitive cyber security assistant should use information extraction technology like IBM's Statistical Information and Relation Extraction (SIRE) which is an UIMA-based system that can perform:

- Mention detection: identify spans that are mentions of targeted ontology's entity types (TTP, ADVERSARY_RESOURCE, VICTIM_TARGETED, INDICATOR, etc.).

- Co-reference resolution: group the mentions of like entity of the same instance across unstructured data sources

- Relation extraction: identify relations between pairs of extracted mentions within the same sentence.

Figure 11 illustrates how the Maximum Entropy classifier (statistic machine translation) is used in IBM's Statistical Information and Relation Extraction (SIRE) toolkit to understand cyber security important entities mentioned, relations between different entities in a sentence, co-references of the same mentions of the same entities within and across documents, and converts textual data into structured data. It shows the WKS machine learning decoded values – where the machine was able to understand a Sasser [TTP] spreads [THREAT_RELATED_TARGET] by exploiting the system [RESOURCE.SYSTEM/VICTIM_TARGETED] through a "stack based buffer overflow" [EXPLOIT_TARGET].
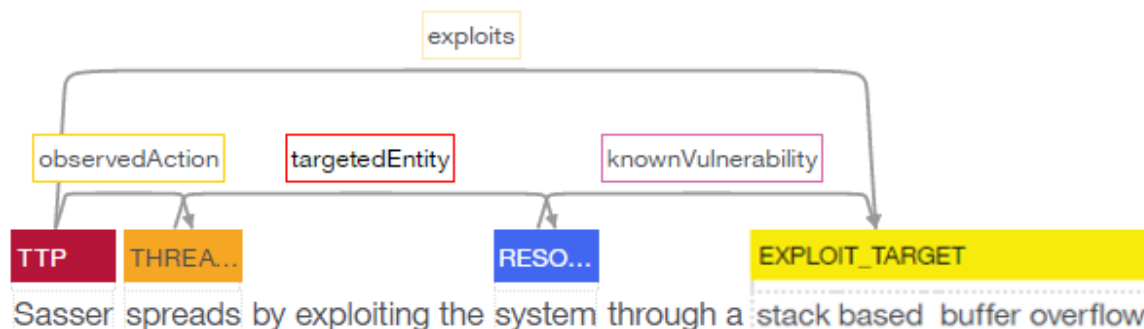

Figure 11 – Information Extraction of Cyber Security Text

Figure 12 shows how these cognitive capabilities enable cognitive cyber security assistants to understand concepts by decomposing expressions of an idea and then combining the results with context and the

probability that certain terms in the passage are being used in a certain way. And, as with humans, cognitive cyber security assistants' confidence is proportional to the evidence that supports those probabilities and the number of reasoning algorithms that it has available to test hypotheses. After Cyber Watson establishes a certain level of understanding, decomposing the problem against its probable intent, Cyber Watson can recompose the elements in various ways, each of which can be tested to provide new concepts. These cognitive capabilities can then be used by security professionals to drive new discovery and insight, helping them to better understand the relationships between malicious intent of an attack, types of attacks involved, actors orchestrating the attack, and the actors' methods-tradecraft or objectives.
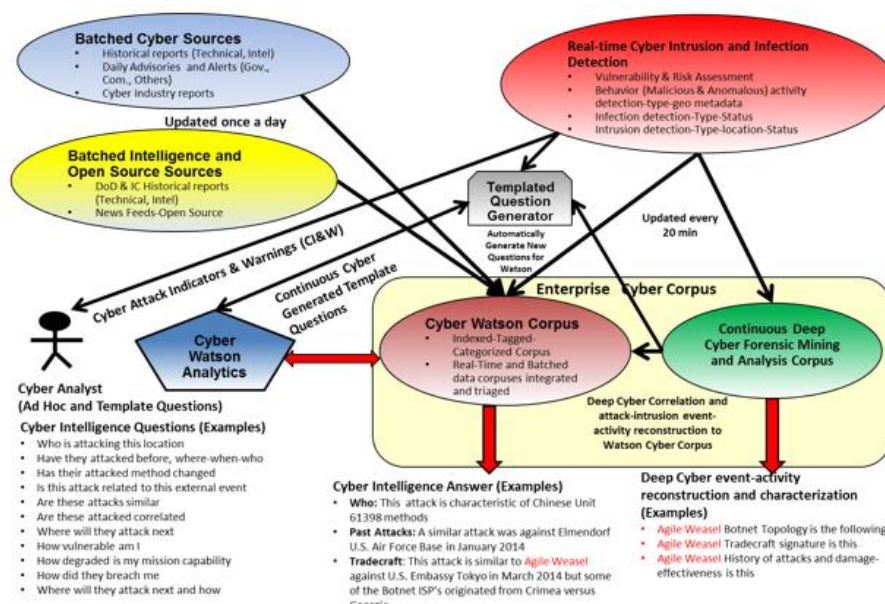

Figure 12: Cognitive Threat Intelligence Analytics

4.0   ADVANCED CYBER DEFENSE CYBERNETICS

Whereas Section 3 described the informatics and advanced technology envisioned in future cybersecurity systems, this section will focus on the associated cybernetics. The authors categorize two different (though not mutually exclusive) cybersecurity cybernetic systems: Automated Cybersecurity Decision Support Systems, and Autonomic Cybersecurity Systems. These are described in more detail in the following sub-sections.

4.1   Automated Cybersecurity Decision Support Systems

Decision Support Systems (DSS) can be described as computerized decision enabling tools that are designed to support the understanding, planning, management, operation, and maintenance of a mission or responsibility. Today's networked computer systems enable individuals from high level executives to lower level operators to use information in radically new ways, to make dramatically more effective and efficient decisions -- and make those decisions more intelligently and rapidly. Given the broad scope of cybersecurity, as well as the size and complexity of cybersecurity systems, DSS play a major role in the design and implementation of these systems.

Traditional DSS systems can be categorized as passive (e.g., make suggestions, but do not make decisions), active (e.g., capable of providing a complete solution or initiating the solution), or cooperative/hybrid (e.g., requiring human intervention at various intervals of the decision cycle). This paper builds on the "cooperative" category, which supports the "human in the loop" requirement in most military systems. Furthermore, traditional DSS can be found that are communication-driven, data-driven, document-driven, knowledge-driven, model-driven, process-driven, or some hybrid combination of these and other drivers. The envisioned cybersecurity DSS take all of the above drivers into consideration.

4.2   Autonomic Cybersecurity Systems

NATO has already begun to exploit autonomous and semi-autonomous technology – e.g., Unmanned Aerial Systems, Unmanned Maritime Systems, and Unmanned Ground Systems. In addition to military capabilities across space, air, land and sea, these systems provide intelligence, surveillance, and reconnaissance (ISR) – which advanced cybersecurity informatics can exploit. As these systems become more and more autonomous, they will require innovative cybersecurity capabilities to deal with a multitude of scenarios (e.g., networked versus stand-alone) and of course, new and advanced persistent threats.

IBM published the Autonomic Computing Manifesto at the turn of this millennium. The original intent of autonomic computing was in the modernization of data centers and enterprise-wide systems; but as the technology evolved alongside computing technology, these concepts were extended to address the IBM Smarter Planet Initiative, recognizing one can monitor and control many sensors and actuators outside of the traditional enterprise systems elements model.

The Autonomic Computing (AC) model has four tenets: Self-Configuring, Self-Healing, Self-Optimizing, and Self-Protecting. These tenets include the following capabilities:

   Self-Configuring – Many of today's corporate data centers have a variety of components from a variety of vendors. Installing, configuring, and integrating components is time consuming and error-prone. The AC vision is to automate the configuration of components and systems according to high-level policies; and the rest of the system adjusts automatically.

   Self-Healing– Determining problems in large, complex systems today can take a team of programmers and system administrators multiple weeks. The AC provides automated detection, diagnosis and resolution to localized software and hardware problems.

   Self-Optimizing – Today's hardware and software contain hundreds of measurements, diagnostics, alarms, and many more tuning options and parameters. The AC strategy is to continuously seek opportunities to improve performance and efficiency.

   Self-Protecting – The majority of today's mechanisms for protecting and detecting system attacks and failures are siloed and are slowly beginning to use standard security naming conventions. The AC approach is to correlate the information from many disparate protection and detection system elements, to automatically identify and defend against malicious attacks and cascading failures, to provide early warning, and to prevent system-wide failures.

The AC tenets can be applied from multiple perspectives on cybersecurity cybernetic systems. These tenets can be applied to support the "survivability" of an autonomous (unmanned or disconnected) system; and these tenets can then also be applied to add autonomy the cybersecurity mission: Protection, Dynamic Detection & Recovery, and Logs & Audits.

A novel aspect of the AC vision is the goal of making decisions that take multiple perspectives into consideration. This can be accomplished with contextual analytics, which provide views of the system elements from various levels of abstraction.
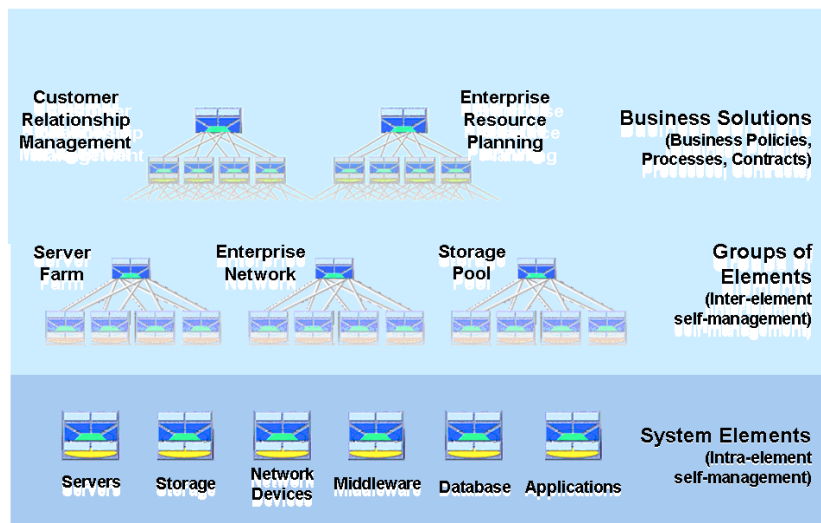
Figure 12: Autonomic Computing Abstraction Levels

Figure 12 illustrates three different views of the same system elements, but with a variety of relationships within different contexts. The lower level abstraction views system elements as a stand-alone entity, and only considers the type of element and vendor specific informatics when diagnosing problems and selecting corrective action. The middle level abstraction looks at the aggregation of similar elements (e.g., a pool of servers) and would extend the diagnostics and problem resolution to consider things such as load balancing, disaster recovery policies, etc.

The top level abstraction takes on a very different perspective. In this example, the top level abstraction places system elements within the context of the corporate missions the system supports. One might place a higher priority on keeping lighter computing loads on mission critical system elements, rather than uniformly spreading computing loads across an enterprise-wide server farm.

From the cybersecurity cybernetics perspective, the elements at the lower level represent the numerous security components of the system (e.g., access controls, IDS, IPS). The mid-level elements may include a number of the advanced cybersecurity informatics described in Section 3. And the top level represents the mission(s) of the system that the cybersecurity system is there to protect.

4.3 Autonomic Computing Cybernetics

The AC design is built on the premise that one must be able to measure, make a decision, and control system elements at all of the abstraction levels (as described in the example of figure 14). This is called the AC control loop. Control loops are designed to provide the self-configuring, self-healing, self-optimizing, and self-protecting capabilities.

Figure 13: Autonomic Computing Control Loop Component

Figure 13 illustrates the composition of the AC control loop components. The monitor sub-component includes mechanisms that collect, aggregate and filter element data (measures) from sensors associated with the specific element being controlled. The analyze sub-component is made up of mechanisms that model or analyze complex situations within the scope of the control loop. The plan sub-component decides the actions needed to achieve the desired goals and objectives. (Note: Analyze and Plan are used to support the "decide" phase of the AC control loop.) The decision logic has embedded "act, monitor and learn" functions, which verify the plan successfully solved the problem or achieved the desired outcome. The execute sub-component executes the actions planned, supporting the control phase of the AC control loop. The knowledge sub-component represents the KM associated with the controlled element, and continuously grows as it stores the information that's been measured, analyzed, planned, and executed – and the ensuing results.


Figure 14: Autonomic Computing Hierarchical Design

Note that the AC control loop structure illustrated in Figure 14 has sensors and effectors at the top and at the bottom. This design provides the ability to manage a single system entity, a pool of entities, or the aggregation of multiple entities -- i.e., in support of the various system levels of abstraction described earlier and illustrated in the hierarchical figure above.

It should be noted that AC is not a product. It is a vision and an approach that can be implemented today at various levels of automation (e.g., self-healing and self-protection), and support many system level abstractions (e.g., resource pools or mission systems). As Information Technology (IT) matures, system elements have improved the information provided through "sensor" interfaces, and have increased the element commands that can be specified through "effector" interfaces to the system elements.

5.0   SUMMARY AND CONCLUSIONS

In the evolution of NATO's collective approach to Cyber Defence, consideration needs to be given to going beyond the traditional approaches to Cybersecurity into the realms of Threat Intelligence Analytics. This will help fulfil NATO's objective of becoming more Cyber Resilient.

This paper provides an overview of advanced cybersecurity informatics (systems), cybersecurity cybernetics (automated decision support and autonomic systems) and the underlying technology (e.g., cognitive computing) that will enable future cybersecurity systems as they evolve to exploit additional data sources and platforms for Threat Intelligence Analyses. The exemplar informatics and cybernetics have been validated with systems varying in size and complexity.

It is recommended that consideration be given to developing a specific NATO Intelligent Threat Analytics platform supporting a federation of cybersecurity systems, allowing configuration & tailoring to meet the perpetually changing missions, world-wide events, and participating country requirements.

6.0 REFERENCES

[1]   R. F. Simmons, "Natural language question-answering systems:1969", *Commun. ACM*, vol. 13, no. 1, pp. 15–30, Jan. 1970.

[2]   M. Maybury, *New Directions in Question-Answering*. Menlo Park, CA: AAAI Press, 2004.

[3]   T. Strzalkowski and S. Harabagiu, *Advances in Open-Domain Question-Answering*. Berlin, Germany: Springer-Verlag, 2006.

[4]   D. A. Ferrucci, Introduction to "This is Watson", *IBM J. Res. & Dev.*, vol. 56, no. 3/4, Paper 1, pp. 1:1, May/Jul. 2012.

[5]   D. Ferrucci and A. Lally, "Building an example application with the unstructured information management architecture," *IBM Syst. J.*, vol. 43, no. 3, pp. 455–475, Jul. 2004.

[6]   Apache UIMA. [Online]. Available: http://uima.apache.org/

[7]   A. Lally, J. M. Prager, M. C. McCord, B. K. Boguraev, S. Patwardhan, J. Fan, P. Fodor, and J. Chu-Carroll, "Question analysis: How Watson reads a clue," *IBM J. Res. & Dev.*, vol. 56, no. 3/4, Paper 2, pp. 2:1–2:14, May/Jul. 2012.

[8] M. C. McCord, J. W. Murdock, and B. K. Boguraev, "Deep parsing in Watson," *IBM J. Res. & Dev.*, vol. 56, no. 3/4, Paper 3, pp. 3:1–3:15, May/Jul. 2012.

[9] A. Kalyanpur, S. Patwardhan, B. K. Boguraev, A. Lally, and J. Chu-Carroll, "Fact-based question decomposition in DeepQA," *IBM J. Res. & Dev.*, vol. 56, no. 3/4, Paper 13, pp. 13:1–13:11, May/Jul. 2012.

[10] D. C. Gondek, A. Lally, A. Kalyanpur, J. W. Murdock, P. Duboue, L. Zhang, Y. Pan, Z. M. Qiu, and C. Welty, "A framework for merging and ranking of answers in DeepQA," *IBM J. Res. & Dev.*, vol. 56, no. 3/4, Paper 14, pp. 14:1–14:12, May/Jul. 2012.

[11] *UIMA Asynchronous Scaleout*, The Apache Software Foundation, Apache UIMA Development Community, ver. 2.3.1, 2011. [Online]. Available: http://uima.apache.org/d/uima-as-2.3.1/uima_async_scaleout.html

[12] E. A. Epstein, M. I. Schor, B. S. Iyer, A. Lally, E. W. Brown, and J. Cwiklik, "Making Watson fast," *IBM J. Res. & Dev.*, vol. 56, no. 3/4, Paper 15, pp. 15:1–15:12, May/Jul. 2012.

[13] NATO Online: http://www.nato.int/cps/ic/natohq/official_texts_112964.htm

[14] NATO 2020: Assured Security; Dynamic Engagement, Online: http://www.nato.int/cps/en/natolive/official_texts_63654.htm?selectedLocale=en , Accessed 2016-09-08

[15] Luiijf, H.A.M. and Smulders, A. (eds), *Future Cyber Defence Concepts and Tools*: A whitepaper by the IST Panel Exploratory Team 066 (IST/ET-066), report STO-TR-IST-ET066, TNO, The Hague, April 2013.

[16] S. Croom-Johnson et al, *Modelling and Simulation for Cyber Defence* MP-MSG-133-09

[17] Maymir-Ducharme, F.A. "CARDS/NSA Unified INFOSEC Architecture (UIA) Expert System," Computer Information Systems Institute (CISI) Workshop, May 1994, National Security Agency, Ft Meade, MD

[18] Kephart, Jeffrey O. and Chess, David M. "The Vision of Autonomic Computing" Journal COMPUTER, Vol. 36, Issue 1, January 2003. IEEE Computer Society Press, Los Alamitos CA

[19] Palmisano, Sam "A Smarter Planet: The Next Leadership Agenda," Speech to the Council on Foreign Relations, New York, NY, 6 November 2008, http://www.ibm.com/ibm/ideasfromibm/ca/en/smartplanet/20090210/sjp_speech.shtml

[20] *Dreier, AS (2012). Strategy, Planning & Litigating to Win. Boston, Massachusetts: Conatus. ISBN 978-0-615-67695-1.* Uses the OODA Loop as a core construct for a litigation strategy system unifying psychology, systems theory, game theory and other concepts from military science.

[21] Maymir-Ducharme, Fred and Angelelli, Lee "Cognitive Analytics: A Step Towards Tacit Knowledge?" Journal on Systemics, Cybernetics and Informatics: JSCI Volume 12 - Number 4 - Year 2014

[22] Maymir-Ducharme, F. and Ernst, R. "Optimizing Distributed and Parallel TCPED Systems," US Geospatial Intelligence Foundation (USGIF) Technical Workshop, Denver CO, July 17-19, 2013

[23] Kephart, Jeffrey O. and Chess, David M. "The Vision of Autonomic Computing" Journal COMPUTER, Vol. 36, Issue 1, January 2003. IEEE Computer Society Press, Los Alamitos CA

[24] Maymir-Ducharme, Fred and Angelelli, Lee "The Smarter Planet: Built on Informatics and Cybernetics," Keynote presentation, proceedings of the 8th International Multi-Conference on Society, Cybernetics and Informatics, July 2014

[25] D. A. Ferrucci, "Introduction to 'This is Watson'," IBM Journal of Research and Development, Vol 56 May/July 2012

[26]   Polanyi, M. (1966) *The Tacit Dimension*, London: Routledge & Kegan Paul

[27]   Radu Florian, "How-to Build a MaxEnt Sequence Classification Model with the SIRE Toolkit," IBM TJ Watson Research Center, Yorktown Heights NY, April 4, 2013

[28]   Imed Zitouni, Xiaoqiang Luo, and Radu Florian. Arabic Computational Linguistics, chapter A Statistical Model for Arabic Mention Detection and Chaining. CSLI Publications, March 2010.