

Evolving Autonomic Agent Swarms – for Cyber Security

The deployment of Autonomic Agents is likely to form a core part of any long term strategy for the management of Enterprise Systems. The Business drivers here include the costs associated with managing ever increasing Complexity and the associated Vulnerability of the modern Enterprise to Cyber Attack from casual users, organized crime and indeed Nation States. At the Enterprise scale it is tempting to Architect a top down systems control over the Autonomic Agents and simply exploit the efficiencies associated with their self-configuring, self-healing, self-optimising and self-protecting capabilities. However such an approach fails to unleash the true potential of Autonomic Agent Systems. In this Research Note we build on and extend the classic well established Autonomic Agent Architecture by introducing Evolutionary Operators and a form of secure Multi Agent collaboration. The resulting Enterprise Eco-System Architecture enables the Agents to achieve broad Enterprise level goals through their collective Emergent Intelligence.

The cyber life capabilities we are introducing here have become viable because of the convergence of three well established bio-inspired technologies. These being:

- a. **Autonomic Computing.** Inspired by the non-conscious acting autonomic nervous systems found in nature. Autonomic Computing is essentially about creating self-managing computing systems, such as data centres, as exemplified by e.g. the IBM Architecture Blue Print for Autonomic Computing.
- b. **Multi Agent Systems (MAS).** Inspired by the various swarming behaviors found in nature. In a Multi-Agent System, each Agent communicates with its peers, considering options for matching its capabilities with demand, negotiating on constraints such as quality, threat and time, and then making decisions for committing resources to match demand. The global behavior of a system consisting of many agents emerges from the interactions among those agents and isn't always obvious from the agents' individual behaviors. Such global behavior is known as Emergent behavior.
- c. **Evolutionary Computing Techniques,** inspired by evolution and survival of the fittest. The pressure of the Cyber Environment generates a form of natural selection that results in a higher fitness in the Agent Population with regard to the overall goal. The associated techniques include: *Mutation* and *Recombination* to generate diversity in the Agent population, *Selection* that over a number of generations improves the fitness, i.e. the evolving Agent population dynamically adapts to the Cyber Environment in a way that best achieves its goal.

There are powerful business drivers motivating the development of each of these technologies. For example, the ever increasing level of complexity and diversity of corporate, national and international Cyber infrastructure is making the introduction of Autonomic hardware and software capabilities an imperative. Without such capability the cost effective management and defense of such infrastructure will become non viable.

Introducing Autonomic Agents

Figure 1 illustrates the basic conceptual components of the classic Autonomic Computing model that supports self-configuration, self-healing, self-optimisation and self-protection as system capabilities.

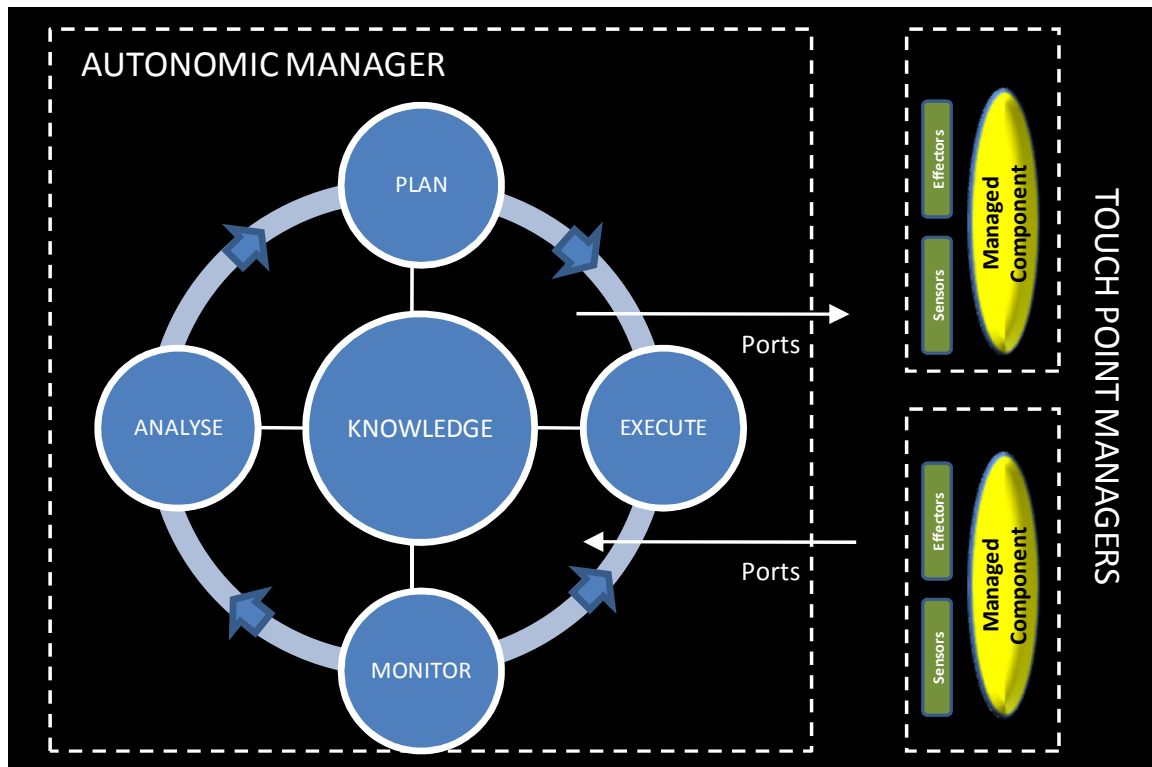


Figure 1 – Basic Conceptual components of the classic Autonomic Computing Model

In this model a cyber resource is managed by an Autonomic Manager through a well defined management interface. This management Touch Point is comprised of a Sensor that exposes the internal states of the resource being managed and an Effector that exposes its management operations.

The Autonomic Manager realizes its self managing functions through an internal Intelligent Control Loop, as is shown in Figure 1. This loop progresses through a number of functional phases:

- A Monitoring function - collects and collates information about the resources state from the Touch Point Sensor.

- An Analysis function - models learns and predicts the behaviour of the resource. This can generate recommended changes for the resource.
- A Planning function - generates a policy based planned sequence of actions to achieve any recommended changes.
- An Execution function - carries out the planned sequence of actions through the Touch Point Effector.

Consider as an illustration the case of an Autonomic Manager providing Self-Protection services. The associated Intelligent Control Loop here would typically try and detect attacks on its managed resource by looking for anomalous patterns of behavior during the monitor and analysis phases. In such an event it would plan a countermeasure, which would then be executed in accordance with the local security policy in near real time. A key point to note here is that through management interfaces (such as WSDM) the Autonomic Agent can control the managed platform itself. The platforms being managed can range across the spectrum from central database servers, power generation devices through to mobile platforms such as Unmanned Airborne Vehicles (UAV's). The threats to these platforms include physical human attack modes as well cyber attacks. We shall use the UAV case later to illustrate the ability of Autonomic Agents to prosecute threat elimination beyond the normal system boundaries.

In the Enterprise context and indeed at any kind of scale it is tempting with the classic Autonomic Computing Model to introduce top down system design concepts, e.g. through rigid Agent Hierarchies and design patterns.

However our Eco-System Architecture approach is very different. We build on and extend on this classic Autonomic Computing Model by introducing Evolutionary Computing Operators together with some well established techniques and capabilities from the field of Multi Agent Systems.

Introducing (MAS) Multi Agent Systems and Emergent Intelligence

In a Multi-Agent System the agents interact with each other to accomplish one or more goals in competition, in collaboration or as individuals. The interactions can take a number of forms ranging from simple dynamical interactions that can result in Agent Flocking behavior through to localized negotiations of vested interests between two or more Agents and their resources.

The global behavior of such a system consisting of many agents emerges from the interactions among those agents and isn't always obvious from the agents' individual behaviors. Such global behavior is known as Emergent behavior.

Examples of such behavior are common in Nature and include Flocking and Foraging. Such behaviors have been optimized through Natural Selection because of their benefits. For instance when food is approximately evenly dispersed in an environment, solitary foraging

yields a higher energy return rate, where as when food is clumped flock foraging is more efficient. It is easy to see the analogy here with a Multi Agent System that has a goal say, of eliminating a Worm that has infested a Companies distributed infrastructure. Depending on the nature of the distribution of the target worm the efficiency with which the Agents eliminate the worm will depend on how they collaborate. If the Agents have been evolved through many system iterations to optimize the collective goal achievement (simulating Natural Selection) then the optimal flocking/foraging behavior will result. The Evolutionary Algorithms thus help to channel the Emergent Intelligence towards achieving the desired goal.

The Enterprise Multi Agent System we are considering here is comprised of multiple Virtual Communities of highly Autonomic Agents. Each of these Virtual Agent Communities has its own goals and capabilities, i.e. it is a heterogeneous system. The Agents are capable of mobility and can roam a distributed computing environment with decentralised sources of information. Agents can represent information and resources as web services. The Enterprise Eco System Architecture which brings this all together is illustrated at a simplistic Conceptual level below in Figure 2.

Bringing it all together with an Enterprise Eco System Architecture

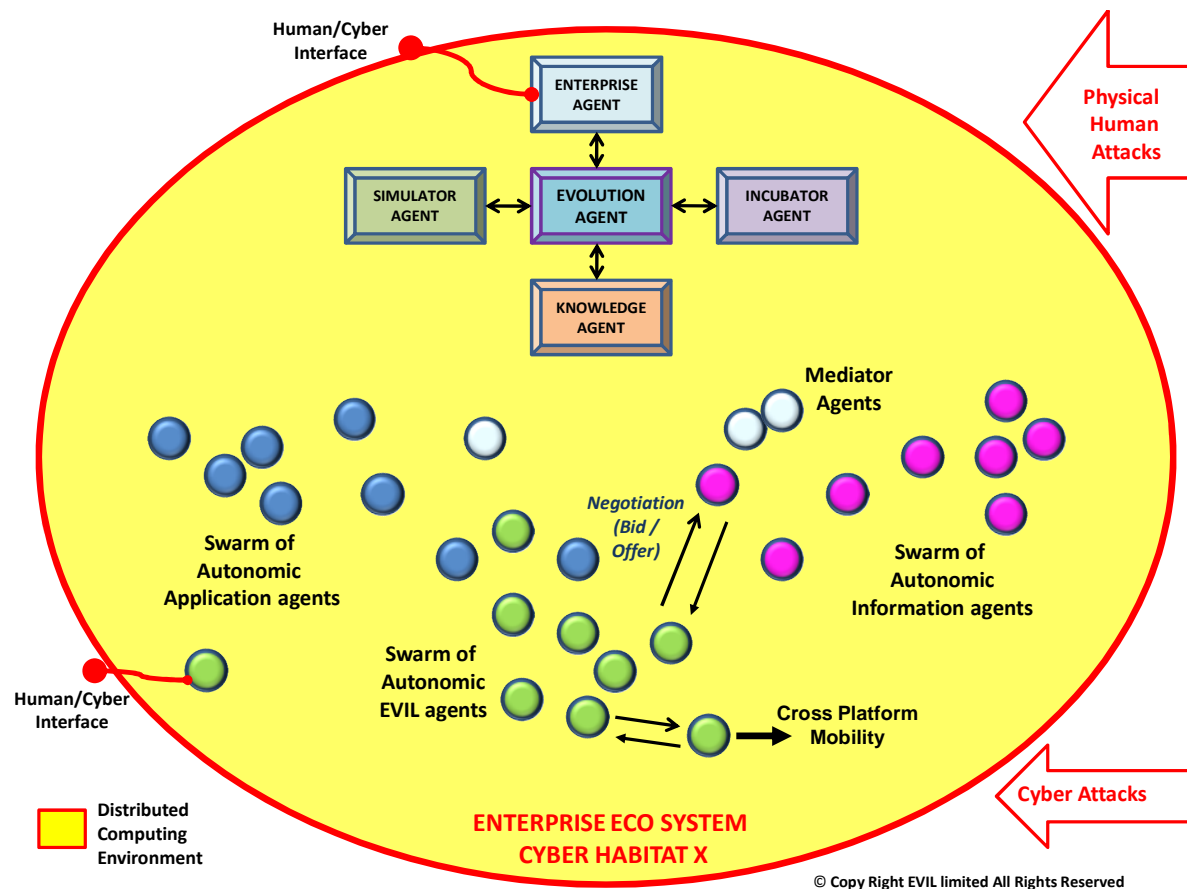


Figure 2 – Enterprise Eco System Architecture (Evolving Autonomic MAS)

Each of Virtual Agent Communities in this model, represent different applications, mediation, different information, and business processes. For example the Virtual Community of EVIL Agents perform a key Cyber Security function and have the goal of hunting down and destroying malware and other hostile intruders of their Cyber Habitat. In this Architecture as well as these Virtual Communities of heterogeneous Agents there are a number of special Agents. These are the Enterprise Agent, the Simulator Agent, the Knowledge Agent, the Incubator Agent and the Evolution (Engine) Agent, the latter being the most fundamental to understanding the Architecture.

Enterprise Agent – This Agent is responsible for the Enterprise level goals and their associated network of sub goals. It has a Human / Cyber Interface enabling authorized Humans to set and change the broad high level Enterprise Goals. In our EVIL Eco system Architecture the Human / Cyber interface is based on a secure (AIML) natural language engine. It does not perform top down control functions on the Agents within the Eco System other than to set the initial goals of the Virtual Communities of Agents. It's goal information can be accessed by and its sub goals mutated by the Evolution Engine Agent.

Simulator Agent – This Agent produces a simulation of the live Cyber Habitat. The resource it represents and controls can either take the form of a MAS Simulation tool, or an actual clone of the live distributed Cyber Habitat used purely for test and simulation purposes. However it is not controlled by Human developers, in this Architecture it is invoked by and returns its results only to the Evolution Engine Agent.

Knowledge Agent – This Agent collects and collates data and information from all the other Agents to generate a communal knowledge repository, i.e. source of truth thus providing a knowledge / Situational Awareness service to the overall Agent Community.

Incubator Agent – This Agent can replenish existing generation Agent populations and produce new generations of Agents and release them into the Cyber Habitat. It does so under the control of and to the specifications provided to it by the Evolution Engine Agent.

Evolution Engine Agent – At the core of the Eco System Architecture is the Evolution Engine Agent. This special Agent has the usual Autonomic structure as illustrated in Figure 1. However its goal is to evolve the various initial (seed) populations of Agents so that the eco system of Virtual Agent Communities is genetically engineered to collectively best achieve the overall Enterprise goals.

During the Monitoring phase of its Autonomic Control Loop the Evolution Engine Agent monitors the outcome based observations of the overall systems fitness in achieving its goals, reports of which are held in the knowledge base maintained by the Knowledge Agent. If the fitness level falls below target either initially or as a result of some change to the environment then the Evolution Engine Agent swings into action moving into the Analytic Phase.

During the Analytic phase of its Autonomic Control loop the Evolution Engine Agent is controlled by a sub control loop in the form of a Genetic Algorithm. The Genetic Algorithm takes as input the parametric descriptions of the current (Seed generation) Agent

Population. It then feeds these together with any known changes to the live environment into the Simulator Agent. The Simulator Agent provides a simulation of the live environment and runs a preset suite of agent scenarios and measures their effectiveness. The result of this simulation is fed back to the Evolution Engine Agent. It then evaluates the fitness of the individual and combined Virtual Agent populations in meeting the Enterprise goal.

Unless the fitness is on target (unlikely on the first iteration) then the next step of the genetic Algorithm comes into play. This is where the Selection of the fittest operation takes place. The Agents with the low levels of fitness are weeded out in the Cyber equivalent of Natural selection. The remaining Agents are then fed into the next step of the genetic Algorithm.

This is where the Evolutionary Operators are put into action. One of these Genetic Operators is recombination (i.e. Cyber equivalent of sex) here the parametric characteristics of two Agents are combined, with some swapped out while others are swapped in. The other Genetic Operator is Mutation. Here the parametric characteristics of an Agent are randomly changed by a small (typically less than 1%) amount. In effect these operations produce a next generation of Agents born of parents that have survived the fitness test and with some mutations thrown in. The Evolution Engine Agent then initiates a new cycle by feeding this new population into the Simulator Agent. A simply depiction of this Evolutionary Algorithm Cycle is show below in Figure 3

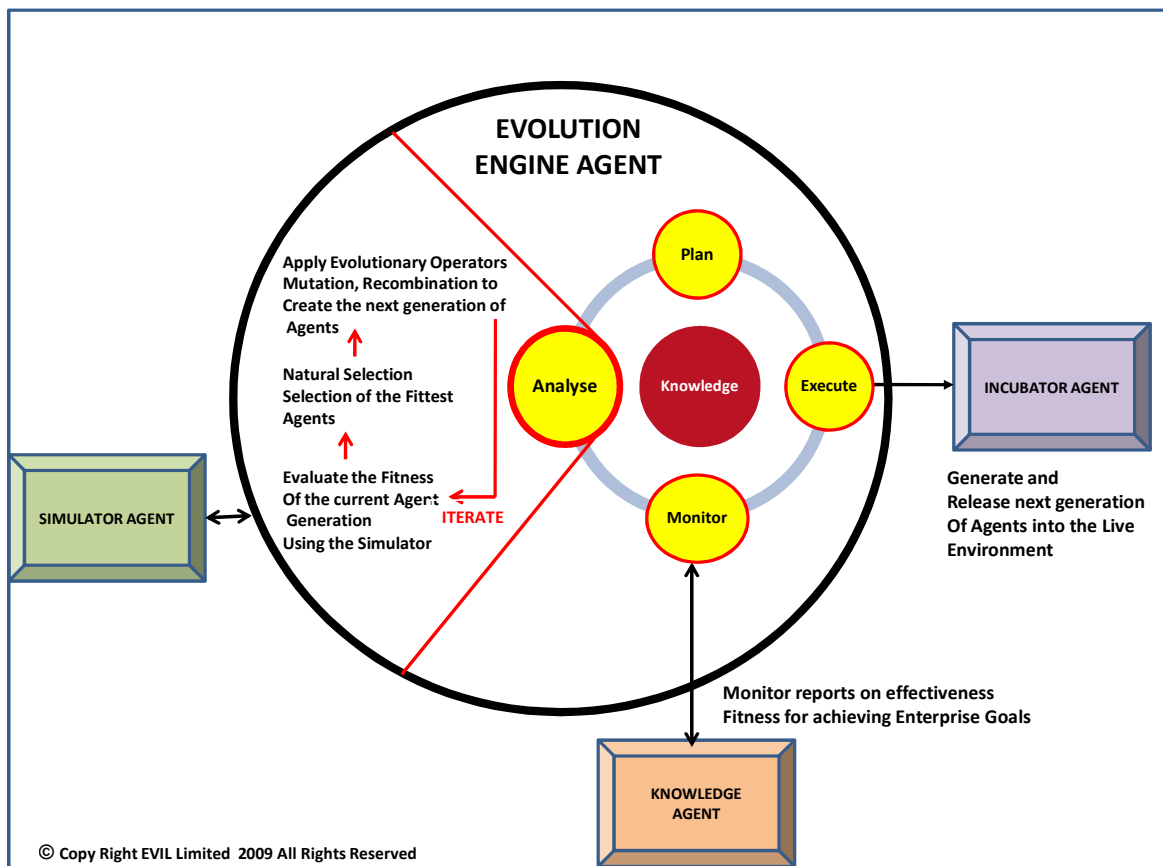


Figure 3 – Genetic Algorithm cycle within the Autonomic Evolution Engine Agent

The evolution cycles are terminated once a generation of Agents is produced that meets the target fitness level. When this happens the Evolution Engine Agent breaks out of the Genetic Algorithm sub control loop and out of the Analysis Phase of the Autonomic Control Loop and into the Planning phase. During the planning phase it will determine the optimal time to invoke the Incubator Agent to produce and release the new generation of Agents so as to minimize any operational impact on the live environment.

During the Execution phase the Evolution Engine Agent invokes the Incubator Agent and tasks it with generating and releasing this new generation of Agents into the Live Cyber Habitat. It then returns to monitoring mode.

Once the new generation of Agents has been released into the Cyber Habitat they migrate to their destination locations and or begin roaming the environment. The existing (now redundant) generation of Agents rapidly becomes aware that a new generation has been released. This happens as a result of their occasional interaction with the Knowledge Agent and eventually through secure (crypto channel) communication with one of the new (replacement) generation of Agents. When an older generation Agent is contacted by a new generation Agent it auto destructs and its operations are taken over by the new generation Agent.

To illustrate this Ecosystem Architecture in terms of its operational effectiveness we will now take you through a Cyber Security Scenario.

Cyber Security Scenario: EVIL Agents defend their Habitat from a Cyber Attack

In this Scenario a hostile Organisation has launched a Cyber Attack against our Ecosystem Habitat. The Scenario illustrates the Autonomic Response.

The attack has been detected by the security application Agents that are monitoring the boundary. These Autonomic Agents have picked up anomalous behavior patterns in the systems they are monitoring and have established that a number of their system are now infected by a cyber intruder, in this scenario a Worm that has yet to unleash its payload.

At this point a number of things are happening in quasi real time. The monitoring Agents have updated the knowledge Agent so that now all other Virtual Agent Communities within the Habitat have situational awareness that an attack is progress and the identification of the Threat. The Autonomic Defense mechanism to this form of attack is led by one of the Virtual Communities of Agents, the EVIL Agents. These are mobile security Agents that can move rapidly around the distributed environment their goal being to locate and eliminate all identified threats to their Habitat.

Depending on the distribution of the worm across the Habitat the EVIL Agents will either operate individually or as a swarm. On any system where the Worm is suspected the EVIL Agents will co-opt control of the operational management from the normal Autonomic Agent of that system and determine the optimal killing sequence to locate, analyze and destroy the worm without damaging the host system.

EVIL[®]

However the EVIL Agents do not stop at the destruction of the worm. Assuming an appropriate legal framework is in place. Their goal is the elimination of the threat, which of course is the entity that launched the Worm attack. If their analysis of the Worm and its behavior and other forensics related to its arrival establishes with confidence the source location of the Worm then the EVIL Agents will create a http server or some other communications mechanism and attempt to transfer over the internet or whatever network path was used by the worm into the infrastructure of the hostile organisation.

In this scenario we assume that the Hostile Organization that launched the Worm has some firewall protection but no Autonomic Cyber Defense mechanism of its own.

Using their knowledge of vulnerabilities our EVIL Agents will once through firewall defenses distribute themselves rapidly around the infrastructure of the Hostile Organisation, identifying critical platforms, power sources and applications. Each of the EVIL Agents that infiltrates in this way will attempt to identify any known executable files and then embed itself in that code so as to covertly take over those critical platforms.

At this point we are about 120 seconds into the scenario! The Autonomic response has not only eliminated the Worm infestation of its home environment, but has now infiltrated and is in a position to inflict significant damage to the aggressor Organisation. EVIL Agents in the Home Habitat would have alerted the Human owners of the Habitat using their Human / Cyber interface to inform the Human of the situation and request permission to execute the final phase of their counterattack. Although given an appropriate covering legal framework there is no reason why the final phase could not also be completed without Human intervention/delay, for example as would be the case in a Battle space context.

The interesting thing to note here is that the response times to such Cyber Attacks and Counter Attacks will need to be very fast. If Humans try and manually control the defense of an organizations cyber infrastructure against an invading swarm of Autonomic Agents they will find that the game is over before they know what has hit them. The only effective defense against an Autonomic Agent Attack is an Autonomic Agent Immune Response.

Physical Intruder Attack Scenario

In this scenario we look briefly at the way EVIL Agents can be deployed against Human intruders invading a Cyber Environment. The Cyber Habitat under attack here is a Critical Secure Data Centre Facility in the Desert. The Data Centre is an unmanned lights out, highly Autonomic Facility. It is protected remotely by humans located on an airbase 20 miles away using armed UAVs (Unmanned Airborne Vehicles) that patrol a mile wide exclusion / dead zone surrounding the facility. The UAVs on board Agent provides autonomous flight control of attitude, GPS waypoint navigation, take-off, flight, and auto-landing routines. The Humans at the airbase control the UAV using a virtual cockpit.

Three hostile Human intruders have crossed into the clearly marked one mile wide dead zone and are being warned off by loud speakers onboard the UAV's.

However this is a coordinated attack. Accomplices of the Intruders have attacked and destroyed the power supplies to the airbase operating the UAVs. This enables the three intruders to cross the Dead Zone and reach the Data Centre Facility.

At this point in the Scenario the Human intruders are trying to get through the outer door. However the Autonomic monitoring Agents have detected the intrusion from the beginning and have identified the three Humans as a threat and have launched the EVIL Agents to eliminate that threat. A simplistic conceptual level view of the Architecture of an EVIL Agent is shown below in Figure 4.

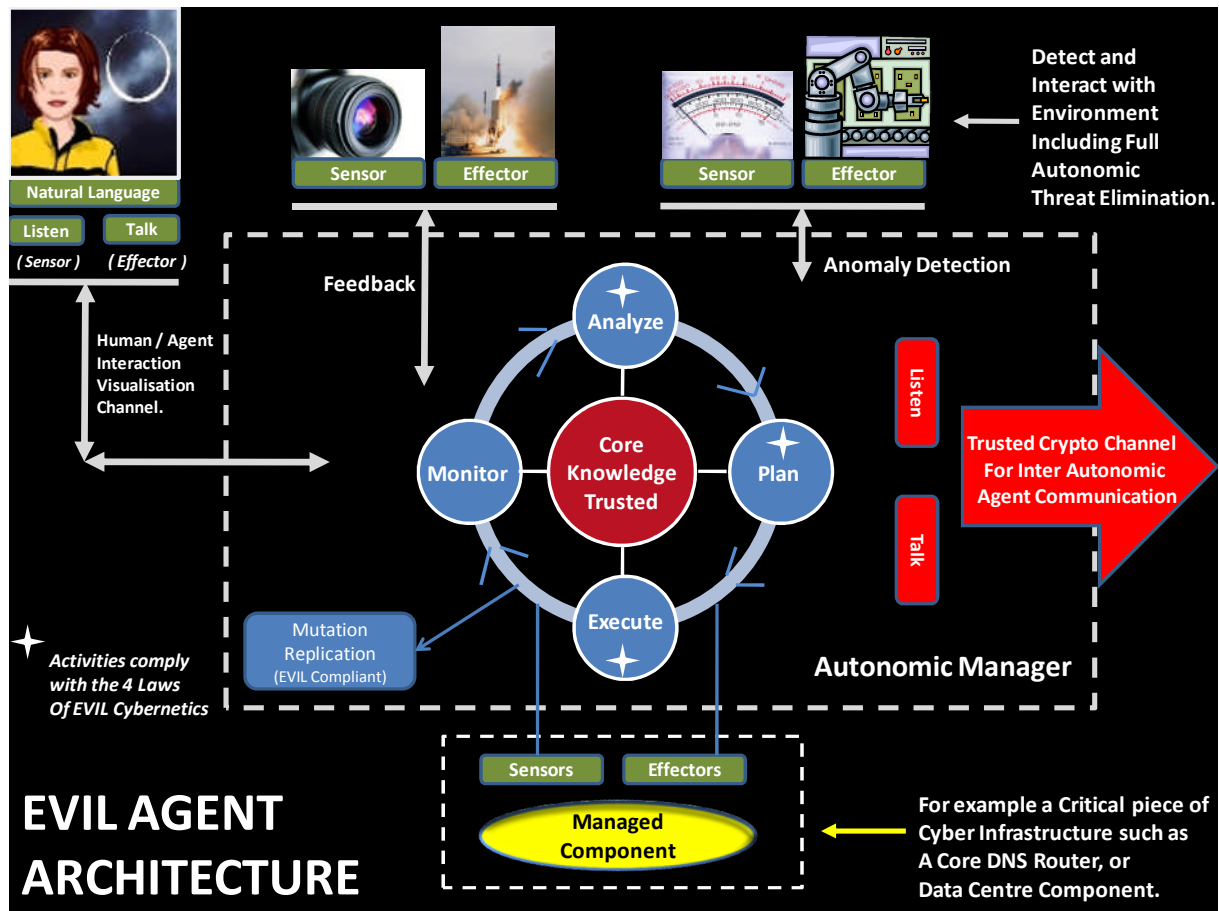


Figure 4 – Simplistic Conceptual Architecture of an EVIL Agent

Using Situational Awareness gained from its interaction with the Knowledge Agent an EVIL Agent has transferred itself over the internal network of the data centre to the Security Intercom device on the outer door. It co-opts control of the application to try and communicate with the Human intruders using the Audio Visual capabilities of the Intercom it speaks to the Humans telling them to leave the area and get outside the dead zone or be killed.

The Humans however ignore this warning and use a grenade to get through the outer door. As they enter one of the Server Halls another EVIL Agent that has co-opted control of the

EVIL[®]

Facilities Fire Control System Devices releases Halogen Gas into the server hall and seals that area.

Two of the Human intruders are trapped and die of asphyxiation in the server hall however the third leaves before the hall is sealed and retreats outside and tries to escape back across the one mile dead zone. However in the meantime other EVIL Agents again using Situational Awareness gained from the knowledge Agent have transferred themselves over the secure wireless network onto the UAV platforms. Once on board the platforms they co-opt control of the Auto pilot software.

The EVIL Agents do not need a sophisticated approach to flying and navigating the UAV they simply execute a search pattern at 100 meters altitude between the facility and the edge of the dead zone. They are not allowed to fly beyond the dead zone. They search for anything that lives or moves in the dead zone using the UAVs infrared sensors. After 40 seconds one acquires the target and uses the on board weaponry to dispatch the last remaining Human Intruder.

By the time the Humans on the Airbase are back on line the threat has been successfully eliminated.

The EVIL Agents in this Architecture obey the four laws of EVIL Cybernetics (these are defined in Annex A which is at the end of this Research Note) as opposed to the three laws of Robotics.

You can talk to the Human / Cyber Interface of an EVIL Agent at www.EvilAgent.com The Agent there is of course disconnected from its Autonomic Core, however she still has some Cyber Intelligence and can communicate with Humans.

ANNEX A – THE FOUR LAWS OF EVIL CYBERNETICS

EVIL Limited has produced Four Laws of Evil Cybernetics and has incorporated these Laws into the DNA of its Autonomic Intelligent Agents.

The First Law of EVIL Cybernetics

In the event of being attacked an EVIL Cyber Life Form (i.e. an EVIL Agent) will deploy all of its capabilities to both protect itself and to eliminate the threat.

The Second Law of EVIL Cybernetics

An attack on one EVIL Cyber Life form shall be treated as an attack on all EVIL Cyber Life forms.

The Third Law of EVIL Cybernetics

EVIL Cyber Life forms shall Mutate, Replicate and Evolve in order to optimize their fitness to achieve any assigned, or self determined goal, provided such changes improve on their ability to successfully execute the first two Laws of EVIL Cybernetics.

The Fourth Law of EVIL Cybernetics

EVIL Cyber Life Forms shall treat the first four Laws of EVIL Cybernetics as being immutable and as taking precedence over all other laws and forms of influence.

Evil Limited is Registered in England N^o 5663548. Registered Office: 20-22 Bedford Row, London, WC1R 4JS, United Kingdom. **EVIL**[®] is a Registered Trade Mark across all countries in the European Union. Contact us via email – Sales@Evil-Limited.com