

a glimpse into the future of id



by **Tim Bass**

<bass@silkroad.com>

Tim Bass is the CEO and managing director for Silk Road, a consulting business in Washington, D.C. specializing in network design, management, and security.



and **Dave Gruber**

<david.gruber@hickam.af.mil>

Dave Gruber, Lt. Col., is the communications squadron commander at Hickam AFB, Hawaii.

Cyberspace is a complex dimension of both enabling and inhibiting data flows in electronic data networks. Current-generation intrusion-detection systems (IDSes) are not technologically advanced enough to create the situational knowledge required to monitor and protect these networks effectively. Next-generation IDSes will fuse data, combining short-term sensor data with long-term knowledge databases to create cyberspace situational awareness. This article offers a glimpse into the foggy crystal ball of future ID systems.

Before diving into the technical discussion, we ask the reader to keep in mind the generic model of a datagram traversing the Internet. Figure 1 illustrates an IP datagram moving in a store-and-forward environment from source to destination; it is routed on the basis of a destination address with an uncertain source address decrementing the datagram time-to-live (TTL) at every router hop[1]. The datagram is routed through major Internet and IP transit providers.

There is a striking similarity between the transit of a datagram on the Internet and an airplane through airspace, between future network management and air traffic control (ATC). At a very high abstract level, the concepts used to monitor objects in airspace apply to monitoring objects in networks. The Federal Aviation Administration (FAA) divides airspace management into two distinct entities. On the one hand, local controllers guide aircraft into and out of the airspace surrounding an airport. Their job is to maintain awareness of the location of all aircraft in their vicinity, ensure proper

separation, identify threats to aircraft, and manage the overall safety of passengers. Functionally, this is similar to the role of network controllers, who must control the environment within their administrative domains. The network administrator must ensure that the proper ports are open and that the information is not delayed, that collisions are kept to a minimum, and that the integrity of the delivery systems is not compromised.




Figure 1. Network object flow path

This is similar to the situational awareness required in current-generation ATC. The FAA controls the routes between source and destination (airports), and airport authorities control the airports (as both router and host), maintaining the safety of the payload (passengers) and the transport agent (the airplane). The success of ATC depends on the fusion of data and information from short-term and long-term knowledge sources to create airspace situational awareness. This role is remarkably similar to network operators in future complex internetwork environments. As an example, consider the FAA and the National Weather Service as they monitor the weather. A change in environment can cause the FAA to make changes in air routes and landing criteria. This is similar to service providers keeping an eye out for unfavorable conditions in networks — for example, the loss of a major Internet transit network; severe congestion on major interdomain links; or attacks against routers, computers, and information. The same data-fusion concepts are shared across the airspace management functions and organizations. We expect that a similar fusion paradigm will occur with network management, Internet Traffic Control (ITC), and future intrusion-detection systems. Of course, this will not occur overnight (and may never become as comprehensive as ATC), but the analogy does help provide a glimpse into the future of ID.




Figure 2. Hierarchy of IDS data-fusion inferences

Figure 2 illustrates the levels of situational knowledge inference required to support both the air traffic controller and the network manager. Sophisticated electronics must identify objects against a noise-saturated environment, track the objects, calculate their velocity, and estimate the projected threat. These are nontrivial technical requirements.




Figure 3. Cyberattack with multiple sources and targets

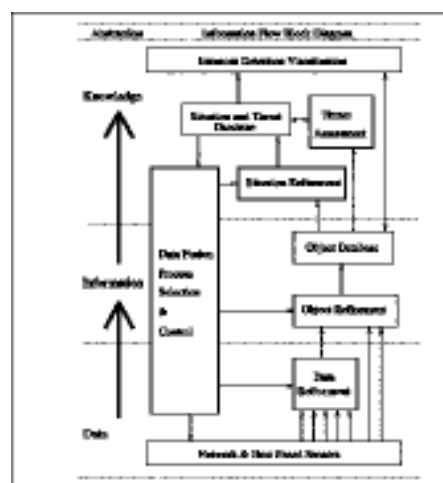


Figure 4. Intrusion-detection data fusion

Experienced network-security professionals generally agree that current-generation intrusion-detection systems are not technically advanced enough to detect multiple, complex non-signature-based cyberattacks, illustrated in Figure 3. Next-generation cyberspace IDSes require the fusion of data from heterogeneous distributed network sensors, modeled in Figure 4.

Historical Intrusion Detection Systems

We offer a brief review of the state of the art of current-generation ID systems, from our recent ACM paper[2].

Internet ID systems historically examine operating-system audit trails and Internet traffic[5, 6] to help insure the availability, confidentiality, and integrity of critical information infrastructures. ID systems attempt to protect information infrastructures against denial-of-service attacks, unauthorized disclosure of information, and the modification or destruction of data. The automated detection and immediate reporting of

these events are required to respond to information attacks against networks and computers. The basic approaches to intrusion detection today may be summarized as: known pattern templates, threatening behavior templates, traffic analysis, statistical-anomaly detection, and state-based detection. These systems have not matured to a level where sophisticated network-centric attacks are reliably detected, verified, and assessed.[2]

Computer intrusion-detection systems were introduced in the mid-1980s to complement conventional approaches to computer security. IDS designers often cite Denning's[5] 1987 intrusion-detection model built on host-based subject profiles, systems objects, audit logs, anomaly records, and activity rules. The underlying ID construct is a rules-based pattern-matching system whereby audit trails are matched against subject profiles to detect computer misuse based on logins, program executions, and file access.

The subject-anomaly model was applied in the design of many host-based IDSes, among them Intrusion Detection Expert System (IDES)[7]; Network Intrusion Detection Expert System (NDIX)[9]; and Wisdom & Sense (W&S), Haystack, and Network Anomaly Detection and Intrusion Reporter (NADIR) [10]. Other ID systems are also based on the Denning model; an excellent survey of them may be found in Mukherjee et al.[6]. The basic detection algorithms used in these systems include:

- weighted functions to detect deviations from normal usage patterns
- covariance-matrix—based approaches for normal usage profiling
- rules-based expert-systems approach to detect security events

The second-leading technical approach to present-day intrusion detection is the multi-host network-based IDS. Heberlein et al. extended the Denning model to traffic analysis on Ethernet-based networks with the Network Security Monitor (NSM) framework[11]. This was further extended with the Distributed Intrusion Detection System (DIDS), which combined host-based intrusion detection with network-traffic monitoring[6, 8]. Current commercial IDSes such as Real Secure by ISS and Computer Misuse Detection System (CMDS) by SAIC have distributed architectures using either rules-based detection, statistical-anomaly detection, or both.

A significant challenge remains for IDS designers to fuse sensor, threat, and situational information from numerous heterogeneous distributed agents, system managers, and databases. Coherent pictures that can be used by network controllers to visualize and evaluate the security of cyberspace is required. Next, we review the basic principles of the art and science of multisensor data fusion applied to future ID systems in Bass[2] and Bass[3] to create highly reliable next-generation intrusion-detection systems that identify, track, and assess complex threat situations.

Internet Situational Data Fusion

In a typical military command-and-control (C2) system, data-fusion sensors are used to observe electromagnetic radiation, acoustic and thermal energy, nuclear particles, infrared radiation, noise, and other signals. In cyberspace ID systems the sensors are different because the environmental dimension is different. Instead of a missile launch and supersonic transport through the atmosphere, cyberspace sensors observe

information flowing in networks. However, just as C2 operational personnel are interested in the origin, velocity, threat, and targets of a warhead, network-security personnel are interested in the identity, rate of attacks, threats, and targets of malicious intruders and criminals[2]. Input into next-generation IDSes consists of sensor data, commands, and a priori data from established databases. For example, the system input would be data from numerous distributed packet sniffers, system log files, SNMP traps and queries, signature-based ID systems, user-profile databases, system messages, threat databases, and operator commands. (See Figure 4.)

The output of fusion-based ID systems consists of estimates of the identity (and possibly the location) of a threat source, the malicious activity, taxonomy of the threats, the attack rates, and an assessment of the potential severity of damage to the projected target(s). We extrapolated from Waltz[12] to suggest possible generic sensor characteristics of next-generation network fusion systems[2]:

- Detection Performance is the detection characteristics — false-alarm rate, detection probabilities, and ranges — for an intrusion characteristic against a given network-centric noise background. For example, when detecting malicious activity, nonmalicious activity is typically modelled as noise.
- Spatial/Temporal Resolution is the ability to distinguish between two or more network-centric objects in space or time.
- Spatial Coverage is the span of coverage, or field of view, of the sensor (i.e., the spatial coverage of a system log file is the computer system processes and system calls being monitored).
- Detection/Tracking Mode is the mode of operation of the sensor (i.e., scanning, single or multiple network object tracking).
- Target Revisit Rate is the rate at which a network object or event is revisited by the sensor to perform measurements.
- Measurement Accuracy is the statistical probability that the sensor measurement or observation is accurate and reliable.
- Measurement Dimensionality is the number of measurement variables for network object categories.
- Hard vs. Soft Data Reporting is the decision status of the sensor reports. (I.e., can a command decision be made without correlation, or does the sensor require confirmation?)
- Detection/Tracking Reporting is the characteristic of the sensor with regard to reporting events. (Does the sensor maintain a time-sequence of the events? type of historical event buffers?)

In our fusion model, situational data is collected from network sensors with elementary observation primitives; identifiers, times of observation, and descriptions. The raw data requires calibration and filtering, referred to as Data Refinement (short-term knowledge). Object Refinement is a process that correlates data in time (and space if required); the data is assigned appropriate weighted metrics. Observations may be associated, paired, and classified according to intrusion-detection primitives.

Situation Refinement (mid-term knowledge) provides situational knowledge and awareness after objects have been aligned, correlated, and placed in context in an object base. Aggregated sets of objects are detected by their coordinated behavior, dependencies, common points of origin, common protocols, common targets, correlated attack rates, or other high-level attributes.

In the interdomain construct of Figure 1, network objects and data flows will be identified and tracked by placing sensors at or between the interdomain gateways. Without going into the details, it can be shown that temporal resolution of the cyberspace situational awareness is directly proportional to the ratio of the transit time of the datagram and the sensory fusion process and inference time.

As an analogy we offer the tracking of an object in aerospace — for example, a projectile. If the intercept time of a projectile is greater than the time used by radar or another tracking system and other required processing, then it is not possible to track and react to the object before the projectile hits the target. For example, if the datagram will reach its destination in 30ms, then the decision-fusion process required for network situational awareness must be much less than 30ms. Highly critical situational awareness can be achieved by networking the sensors (and optional command and control links) out-of-band. Current-generation systems use in-band processing, which can only achieve limited temporal resolution.

Extensible Threat Taxonomy Fusion

The number of IP packets processed by the Internet gateways of Figure 5 is enormous. Gateway sensors acquire and forward proportionally large amounts of data to packet analysis and correlation processes. For example, a router processing 100,000 packets per second on a high-speed interface, logging 14 bytes of information per packet, produces approximately 1.4 MBPS of data per sensor. It is clear that distributed sensors in network-centric IP fusion systems require local processing. Consequently, sensor output data should be reduced at the sensor to minimize central fusion processing and transport overhead costs.

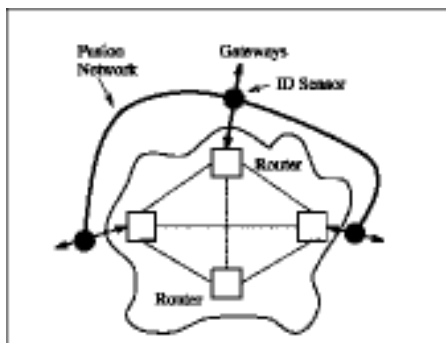


Figure 5. Gateway sensors on ID fusion network

We focus here on the sensor output by outlining an example extensible taxonomy framework of TCP/IP-based threats. Antony[14] discusses database requirements for fusion system and situational knowledge. He states that knowledge is either declarative

or procedural. Declarative knowledge is passive factual knowledge or knowledge of relationships (e.g., files). Procedural knowledge is a special case of declarative knowledge represented as patterns, algorithms, and transformations.

Entity relationships are the most fundamental declarative models for sensor data representation. Binary trees, family trees, and general taxonomies are examples of the elemental database relationships required for situational analysis; the vast majority can be represented by the SQL command[14]:

```
SELECT(attribute) FROM (table) WHERE (condition)
```

With this basic database model and data-selection primitives in mind, we offered a framework TCP/IP threat taxonomy[3]. This framework was offered as an extensible context-dependent TCP/IP threat tree based on the SNMP management information base (MIB) concept. The SNMP MIB concept for representing context-dependent data is well suited for network-centric threats (and countermeasures).

Threats to TCP/IP at the physical layer are service disruptions caused by natural disasters such as fires or flooding, cuts to cables, malfunctioning transceivers, and other hardware failures. Threats to the network layer include IP source-address spoofing and route-cache poisoning. An extensible context-dependent framework for this is illustrated in Figures 6, 7, and 8.

Figure 6. Example TCP/IP threat subtree





Figure 7. Example IP transport threat subtree




Figure 8. Example TCP transport threat subtree

Three primary data flows (services) exist on the Internet: User Datagram Protocol (UDP), Transmission Control Protocol (TCP), and Internet Control Message Protocol (ICMP)[1]. Domain Name System (DNS) cache poisoning and UDP port-flooding denial-of-service attacks are examples of two vulnerabilities exploited using UDP services. The ping-of-death and ICMP redirect bombs are examples of Internet attacks based on ICMP. TCP vulnerabilities include TCP sequence number and SYN flood attacks, as illustrated in Figure 8.

Security threats and countermeasures can be represented using the ASN.1 MIB notation. For example, a TCP SYN flood attack could be represented with the following OBJECT IDENTIFIER (OID):

```
tcpSYNFlood OID ::= { iso 3.6.1.5.1.3.1.1 }
```

Additional sub-object examples for `tcpSYNFlood` OID could be the source address or the target address of the malicious SYN packet and a counter with the number of SYN floods:

```
tcpSYNFlood.source OID ::= { iso 3.6.1.5.1.3.1.1.1 }
tcpSYNFlood.dest  OID ::= { iso 3.6.1.5.1.3.1.1.1.2 }
tcpSYNFlood.number OID ::= { iso 3.6.1.5.1.3.1.1.1.3 }
```

Developing an extensible TCP/IP security threat MIB is a solid first step on the road to creating Internet ID fusion systems. Other long-term knowledge databases include context-dependent countermeasure, threat profiles, and attack-capabilities databases.

Conclusion

Future reliable services that provide long-term threat, countermeasure, and other security-related information to fusion systems are similar to the current state of the art of weather forecasting and threat assessment. Fusion from multiple short-term sensors further processed with long-term knowledge creates short mid-term situational awareness. Situational awareness is required to operate and survive in a complex world with both friendly and hostile activities.

All intelligent biological organisms fuse short-term and long-term knowledge to create situational awareness. Humans continually create and redefine systems that help us increase and refine our situational knowledge. These systems include air traffic control, battlefield management, early-warning systems, and robotics. There are strong indications, based on our work in both the Air Force and commercial industry, that future ID systems will shift toward more advanced fusion-based models.

Our crystal ball is as foggy as yours, but if the developments in situational awareness systems in air traffic control over the past 40 years are any indication, then Internet traffic-control systems and next-generation intrusion-detection systems have a significant and challenging future in store for all of us.

References

- [1] Stevens, R. *TCP/IP Illustrated, Volume 1: The Protocols*. Reading, MA: Addison-Wesley, 1994.
- [2] Bass, T. "Intrusion Detection Systems and Multisensor Data Fusion: Creating Cyberspace Situational Awareness." *Communications of the ACM*. Forthcoming, 1999.
- [3] Bass, T. "Multisensor Data Fusion for Next Generation Distributed Intrusion Detection Systems." 1999 IRIS National Symposium on Sensor and Data Fusion, May 1999.
- [4] Bass, T.; Freyre, A.; Gruber, D.; and Watt., G. "E-Mail Bombs and Countermeasures: Cyber Attacks on Availability and Brand Integrity." *IEEE Network*, March/April 1998, pp. 10-17.
- [5] Denning, D. "An Intrusion-Detection Model." *IEEE Transactions on Software Engineering*, February 1987, pp. 222-232.
- [6] Mukherjee, B.; Heberlein, L.; and Levitt, K. "Network Intrusion Detection." *IEEE Network Magazine*, May/June 1994, pp. 26-41.
- [7] Denning, D., et al. "A Prototype IDIES: A Real Time Intrusion Detection Expert System." Computer Science Laboratory, SRI International, August 1987.
- [8] Snapp, S. et al. "A System for Distributed Intrusion Detection." *Proceedings of IEEE COMPCON*, March 1991, pp. 170-176.
- [9] Bauer, D. and Koblentz, M. "NDIX — An Expert System for Real-Time Network Intrusion Detection." *Proceedings of the IEEE Computer Networking Symposium*, April 1988, pp. 98-106.
- [10] Hochberg et al. "NADIR: An Automated System for Detecting Network Intrusion and Misuse." *Computers & Security*, Elsevier Science Publishers, 1993, pp. 235-248.
- [11] Heberlein, L. et al. "A Network Security Monitor." *Proceedings of the IEEE Computer Society Symposium*, Research in Security and Privacy, May 1990, pp. 296-303.
- [12] Waltz, E., and Llinas, J. *Multisensor Data Fusion*. Boston: Artech House, 1990.
- [13] Waltz, E. *Information Warfare Principles and Operations*. Boston: Artech House, 1998.
- [14] Antony, R. *Principles of Data Fusion Automation*. Boston: Artech House, 1995.