



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Anomaly Detection: How to apply hypothesis testing to IDS events for quality control

Monica D. Sanchez
February 23, 2004

GSEC Practical Version 1.4b Option 1

Abstract

Hypothesis testing is a common statistical technique applied to in the engineering field of quality control to determine if a statistically significant difference exists between two data sets. By applying hypothesis testing as an anomaly detection technique to Intrusion Detection System (IDS) data, a decision can be made between sets of events to determine anomalous behavior. Therefore, the system is enabled to alert with measured confidence.

Terminology

Anomaly detection – detecting activity that is characteristically or statistically different for any given user, system, or network. [8]

False alarms – IDS alert triggered with no security significance or impact. Kevin Timm, a Network Security Engineer at Netsolve Inc., notes that false alarms are generally classified in the following groups:

“Reactionary Traffic alarms: Traffic that is caused by another network event, often non malicious. An example of this would be a NIDS device triggering an ICMP flood alarm when it is really several destination unreachable packets caused by equipment failure somewhere in the Internet cloud.

Equipment-related alarms: Attack alerts that are triggered by odd, unrecognized packets generated by certain network equipment. Load balancers often trigger these types of alarms.

Protocol Violations: Alerts that are caused by unrecognized network traffic often caused by poorly or oddly written client software

True False Positives: Alarms that are generated by an IDS for no apparent reason. These are often caused by IDS software bugs

Non Malicious alarms: Generated through some real occurrence that is non malicious in nature.” [9]

Intrusion Detection Systems (IDS) – hardware and software products that inspect and log network traffic that is perceived as malicious or suspicious at the host (HIDS) or network level (NIDS).

Introduction

With all the exposed security vulnerabilities, threats and available exploits, security has become a primary focus at all levels of industry. Many enterprises have introduced Intrusion Detection Systems (IDS) as a layer in their “Defense in Depth” strategy to provide visibility and control over their networks by detecting attacks, policy violations, resource misuse, and faulty configurations. [5, 7]

Yet, the effectiveness of IDS technologies has been in question recently. This skepticism has been strengthened by recent Gartner reports - "Hype cycle for information security, 2003" and "Intrusion Detection is dead – long live intrusion prevention". [7] The research firm has indicated that “intrusion-detection systems

are a market failure” due to a high rate of false alarms. [4, 7] According to Kevin Timm, a Network Security Engineer at Netsolve Inc., 90% of all IDS alerts are false alarms or “noise”; the remaining 10% are actual pertinent security events. The concern with false alarms is that by generating so many unwarranted alerts the “value and urgency” of real security events is weakened by drowning IT resources in a pool of noise. [4, 9]

So the problem does not seem to lie with whether IDS can detect and categorize traffic appropriately, “but rather its ability to suppress false alarms”. [1, 6] Well, why not approach this problem as a quality control issue? Engineers continue to improve the science of consistent quality delivery- why not leverage those techniques to the data gathered by an IDS. A common technique used for quality control is hypothesis testing.

Hypothesis testing is a statistical analysis technique that assists with finding statistically significant difference between two data sets. By applying this method as an anomaly detection technique to IDS data, the system will be enabled to alert with measurable confidence. Thus, any alerts generated would have a specific statistical meaning.

Statistical Refresher

Before getting immersed in hypothesis testing, a quick statistics refresher is in order.

Random Variables:

Take a six-sided die as an example. If the die is rolled, a value of 1, 2, 3, 4, 5 or 6 is expected. This is a random variable (X). The actual rolled value – the outcome - is considered to be the measured random variable or an event (x_n).

There are two types of random variables: discrete and continuous. A discrete random variable can only take on one specific value (such as the outcome of rolling a die), while a continuous random variable can take on a range of values (for instance, the temperature outside).

The probability that the die will land on 5 is denoted as

$$P(X) = x_n \longrightarrow P(X = 5) = \frac{1}{6} \quad (1)$$

where $P(X)$ is the probability distribution function.

Sample Mean:

The mean is a measure of the “location” of the probability distribution function. For a continuous distribution this is defined as

$$m = \sum_{\text{all } x} xP(x) \quad (2)$$

where \sum denotes summation over all values of x . [2, pg 99] However, in most cases, the probability distribution is not known beforehand; thus, it is common to estimate the mean using the sample mean noted as \bar{x} .

$$\bar{x} = \frac{\sum_{n=0}^N x_n}{N} \quad (3)$$

where N is the total number of sampled outcomes and n is an outcome.

For example, if the die is rolled four times with the outcome of 1, 3, 4, and 6. The sample mean is 3.5.

$$\bar{x} = \frac{1+3+4+6}{4} = 3.5 \quad (4)$$

Sample Variance:

The variance is a measure of the width or spread of the probability distribution function. Its definition is similar to equation 2, and for similar reasons the sample variance is commonly used to estimate the variance. [11]

$$s^2 = \frac{\sum_{n=0}^N (x_n - \bar{x})^2}{N - 1} \quad (5)$$

By using the previous example of rolling a die four times with the outcome of 1, 3, 4, and 6, the variance can be calculated by:

$$\begin{aligned} s^2 &= \frac{(1-3.5)^2 + (3-3.5)^2 + (4-3.5)^2 + (6-3.5)^2}{4-1} \\ &= \frac{6.25 + 0.25 + 0.25 + 6.25}{3} = 4.\bar{3} \end{aligned} \quad (6)$$

Sample Standard Deviation:

The sample standard deviation is the square root of the sample variance. It represents the average error from the mean.

$$s = \sqrt{\frac{\sum_{n=0}^N (x_n - \bar{x})^2}{N-1}} \quad (7)$$

The standard deviation has an important meaning with regards to the normal distribution.

Continuing with the same example of rolling a die four times with the outcome of 1, 3, 4, and 6, the sample standard deviation can be calculated by:

$$\begin{aligned} \sqrt{s^2} &= \sqrt{\frac{(1-3.5)^2 + (3-3.5)^2 + (4-3.5)^2 + (6-3.5)^2}{4-1}} \\ &= \sqrt{\frac{6.25 + 0.25 + 0.25 + 6.25}{3}} = 2.08 \quad (6) \end{aligned}$$

Normal Distribution

The normal distribution, also known as the Gaussian distribution, is the most commonly used continuous probability distribution function. A probability distribution relates an outcome to its probability. It is often associated with a bell-shaped curve. [2, pg 65]

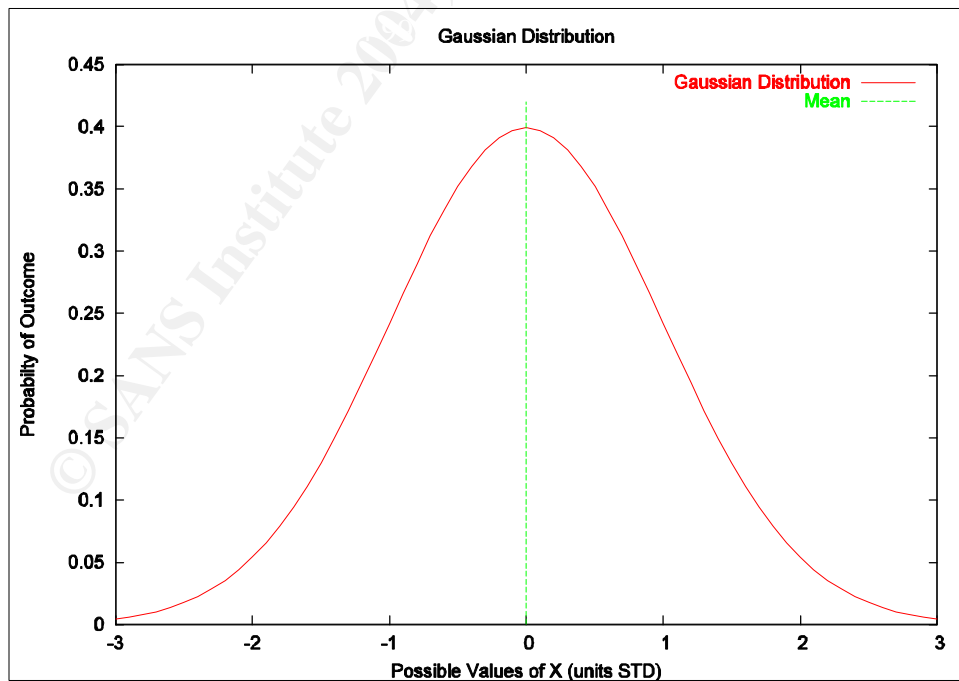


Figure 1 – An example of a normal or Gaussian distribution, also known as, a bell-shaped curve. The x-axis is the value of the random variable in units of standard deviation. The y-axis is the probability of the random event. The further a possible event is from the mean, the less likely it is to occur.

Cumulative Distribution Function

The cumulative distribution function (cdf) is the probability that X takes a value less than or equal to particular value x . The cdf is the area below the probability distribution curve up to the point x . For example, what is the probability that rolling a die will result in rolling a value ≤ 2 . [2, pg 30]

$$P(X \leq x = 2) = \frac{1}{3} \quad (8)$$

An important variable associated with the cdf is z_a . It is the x value with the a cdf probability of $1 - a$. Mathematically, this is

$$z_a \text{ such that } P\{X \leq z_a\} = 1 - a. \quad (9)$$

Student's T-Distribution

The student's t-distribution is often used by professors when grading on a curve. This distribution often replaces the cdf in statistical analysis when sample sizes are small (<30). If T is the student's t-distribution, then

$$T = \frac{\bar{x} - m}{s / \sqrt{n}}, \quad (10)$$

and analogous to z_a , the variable $t_{n-1,a}$ is the x value of t-distribution with probability $1 - a$. [12]

Central Limit Theorem

The central limit theorem states that the mean of a random variable will be normally distributed regardless of the distribution of X . [10]

Hypothesis Testing Example

Hypothesis testing is a technique used to make a decision whether to accept or reject a premise based on observations. [2, pg171] For example, a masonry contractor sells cinder block for structures that requires heavy load sustainability. The contractor states that the cinder blocks can bear 8,000 pounds per square inch (psi). It is crucial for the cinder blocks to meet these minimum requirements, because the structure calculations will be based on the cinder block strength. If the blocks are stronger than 8,000 psi then the stability of the structure will not be compromised. [2, pg172]

So, the contractor's cinder blocks can only be used if it is determined that the blocks meet the minimum requirements. A statistical approach would be to take a random sample of the cinder blocks and make a decision on the outcome obtained by the observed data. Just to note, this does not mean that a wrong

decision can never be made. If a bad sample is taken, the wrong conclusion can be inferred. [2, pg172]

In this example, the null hypothesis, in statistical terms, is that the average strength of the cinder blocks meets the minimum requirements of 8,000 psi ($H_0 : \mu = 8,000 \text{ psi}$). The alternate hypothesis is any situation where the stability of the structure could be compromised. In this case, the alternate hypothesis would be whenever the blocks have an average strength less than 8,000 psi ($H_A : \mu < 8,000 \text{ psi}$). The null hypothesis is “accepted” as truth unless it is statistically “proven” otherwise. [2, pg173]

In summary, hypothesis testing is a systematic technique used to make decisions. It also provided a mechanism to quantify risk – the probability that the wrong decision will be made.

Hypothesis Testing Application

Hypothesis testing was applied to IDS data in a similar way to conduct quality control. In the IDS application of hypothesis testing, the null hypothesis was normal traffic or background noise. The data used for the application was raw IDS summary data. The data was inherently corrupt, because the data consists of both normal noise and abnormal traffic making it difficult to calculate the sample mean and sample standard deviation of normal traffic. For the purpose of this application, the hypothesis testing technique was applied two-fold, as the data is corrupt (with abnormal data) and must be filtered. For the initial application of hypothesis testing, the data was assumed to be Gaussian, and the process was comprised of three steps.

1. Data was divided into moving windows. Then, the sample mean of each window was used to be able to invoke the central limit theorem.
2. Data was removed from the data set to calculate a sample mean and sample standard deviation for normal traffic. If the abnormal data was present, then the sample mean and sample standard deviation would be artificially inflated. As the abnormal data was removed, the data became nearly Gaussian.
3. Sample mean and sample standard deviation was used to calculate the IDS alert criteria, also known as, the decision criteria.

The raw IDS summary data used in this application was obtained from a large e-commerce company that utilizes Internet Security Systems™ (ISS) products. The data was divided into three severities (1-Low, 2-Medium, and 3-High). The severities are assigned by the IDS administrator based on the company's environment. The data set consisted of 6-months worth of IDS data broken down hourly. Here is a snap shot of the data:

Month	Day	Hour (0-23)	Events (Y)	Hour Count (0 to M-1)	Issue Severity
6	1	0	249	0	3
6	1	1	252	1	3
6	1	2	236	2	3
6	1	3	292	3	3
.
.
.
11	30	20	2821	4150	3
11	30	21	3401	4151	3
11	30	22	3152	4142	3
11	30	23	3167	4153	3

Figure 2 – Six months of hourly raw data for all issue severity 3 events. Hours range from 0 to 23, and the hour count begins at 0 to M-1. Events (Y) are the total number of severity 3 events in that hour. The issue severity is assigned by an IDS administrator.

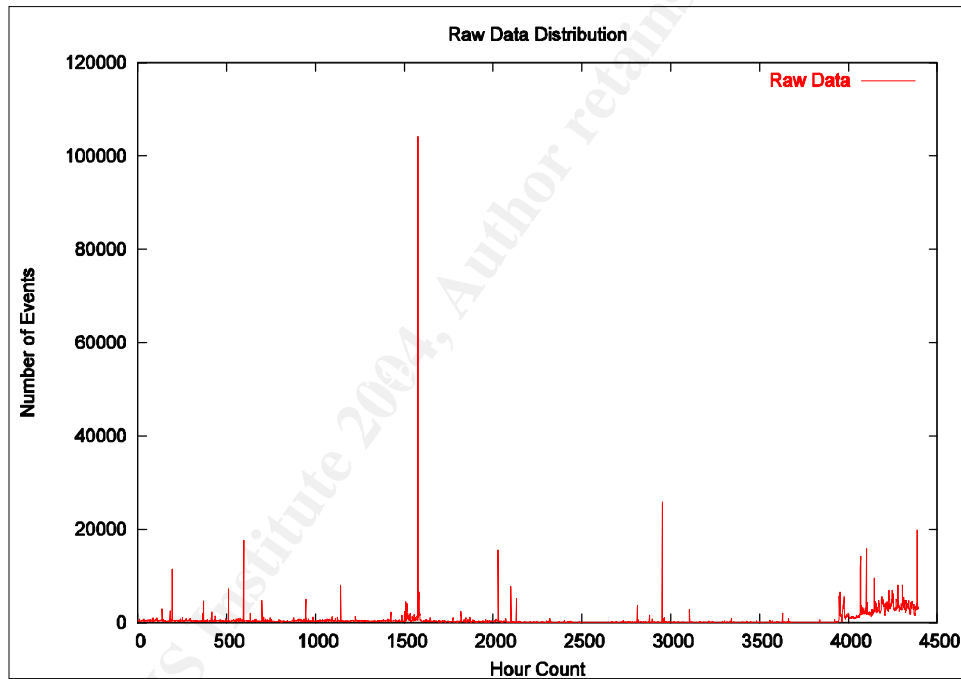


Figure 3 – Data distribution for severity 3 data set. Y-axis is the number of event and X-axis is hour count (0 to M-1).

Initially, the raw data is positioned into moving windows. Consider a moving vehicle, as an analogy for moving window; the vehicle travels from point A to point B. The entire route is considered the data set, so as the vehicle is in motion, the driver's view changes. The driver is observing a "moving window" of the total route as the vehicle is in motion. To accomplish this for this application, the sample size of the moving windows must first be calculated for the desired critical difference. If N is the sample size, then

$$N = \frac{S^2 (z_a + z_b)^2}{(m_1 - m_2)^2} . \quad (11) \text{ [2, pg178]}$$

In equation 11, s is the standard deviation and $m_1 - m_2$ is the critical difference. The critical difference is the minimum detectable change in sample mean. The variables z_a and z_b are the inverses of the standard cdf for the probability of $1-a$ or $1-b$. The variables a and b represent risk. By substituting $a = 0.05$, $b = 0.05$, $s = std(Y)$, and $m_1 - m_2 = std(Y)$, N yields an initial sample size of 11.

Thus, the moving windows will consist of events for N hours. For example,

$$\mathbf{W} = \begin{bmatrix} x_{.0} & x_1 & \dots & x_{M-1} \\ x_{.1} & x_2 & \dots & x_M \\ \vdots & & & \\ \vdots & & & \\ x_{.N-1} & x_N & \dots & x_{M+N-2} \end{bmatrix} \quad (12)$$

where M is the number of windows and N is the sample size; this will be referred to as the window matrix. Thus, each widow can be represented as the following:

$$\mathbf{w}_0 = \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ \vdots \\ x_{N-1} \end{bmatrix}; \mathbf{w}_1 = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ \vdots \\ x_N \end{bmatrix}; \dots; \mathbf{w}_{M-1} = \begin{bmatrix} x_{M-1} \\ x_M \\ \vdots \\ \vdots \\ x_{M+N-2} \end{bmatrix}. \quad (13)$$

These make the columns of the matrix \mathbf{W} .

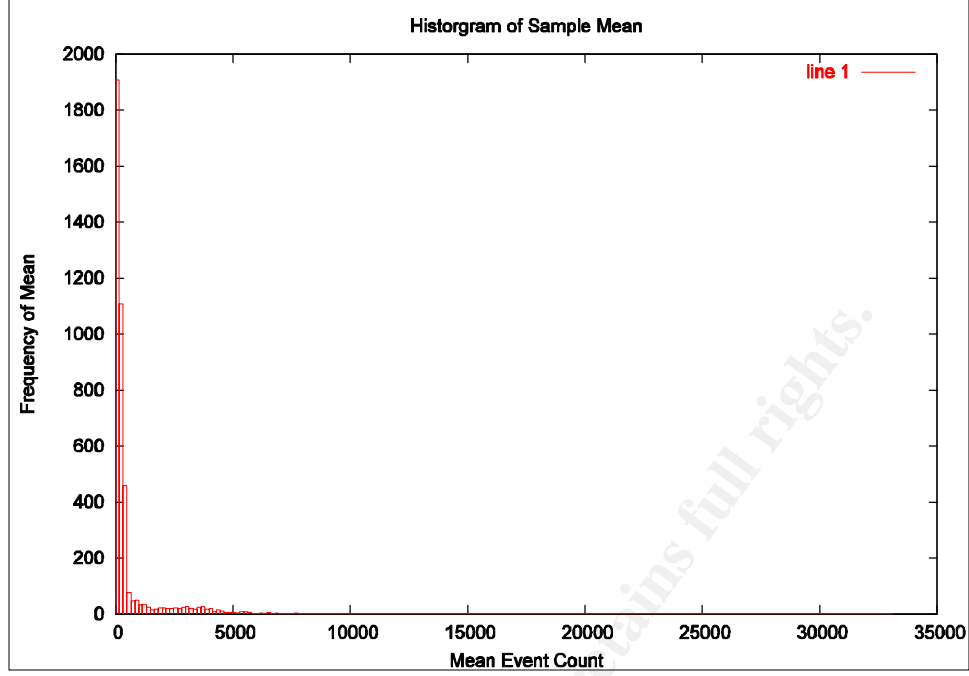


Figure 4 – Histogram of \bar{w} shows the data is not normally distributed.

If \bar{w} is the vector containing the mean of each sliding window, then \bar{w}_m is an element of the vector that is noted as

$$\bar{w}_m = \sum_{n=0}^{N-1} \frac{\mathbf{W}_{1,n}}{N} = \frac{x_0 + \dots + x_{n-1}}{N}. \quad (14)$$

Figure 5 shows the histogram of \bar{w} . It is apparent that the data is not Gaussian, so the first hypothesis test was used as a systematic way to filter abnormal data and obtain a distribution closer to being Gaussian. Now in order to determine abnormal data, the sample mean was taken over all sliding windows. If \bar{W} is the total sample mean of \bar{w} , then

$$\bar{W} = \sum_{m=0}^{M-1} \frac{\bar{w}_m}{M} = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} \frac{\mathbf{W}_{m,n}}{MN}. \quad (15)$$

The sample standard deviation of \bar{w} is also calculated to identify the expected error from the sample mean with

$$S_w = \sum_{m=0}^{M-1} \sqrt{\frac{(\bar{w}_m - \bar{W})^2}{(M-1)}}. \quad (16)$$

These values served to eliminate the abnormal data points. These points must be removed such that the sample mean and sample standard deviation for the normal traffic are not artificially inflated due to the corrupting data. To remove the abnormal data points, a decision criterion must first be calculated to determine what points will be thrown out. If $c1$ and $c2$, are the decision criteria, then

$$c1 = m - z_{\alpha/2} \frac{S}{\sqrt{N}} \quad (17)$$

and

$$c2 = m + z_{\alpha/2} \frac{S}{\sqrt{N}}. \quad (18)$$

Note, the t-distribution replaces the cdf when $N < 30$. By substituting $m = \bar{W}$, $S = S_w$, and $\alpha = 0.05$, then $c1 = 228$ and $c2 = 861$. If \bar{W} was indeed normally distributed, then over time, there is a 5% probability that \bar{W}_m would fall outside of $c1$ and $c2$.

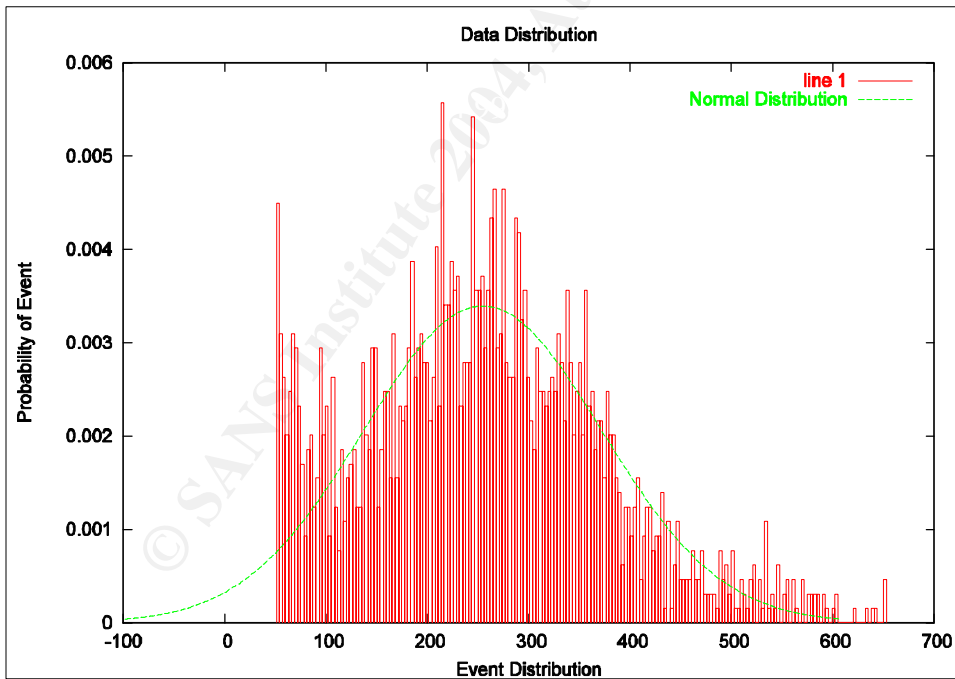


Figure 5 – Histogram of the sliding window, \bar{W}_0 , after abnormal traffic removed with a normal distribution. The data is starting to converge toward a normal distribution.

For the next step, all data points where $c1 \leq \bar{w}_m \leq c2$ were removed from \bar{w} ; this new vector is referred to as \bar{w}_0 . The remaining points are nearly Gaussian, see Figure 6. The sample mean (\bar{x}_0) and sample standard deviation (s_0) of the normal remaining data will be calculated for the basis of the null hypothesis. If \bar{x}_0 is the mean and s_0 is the standard deviation, then

$$\begin{aligned}\bar{x}_0 &= \sum \frac{\bar{w}_m}{M_2} \\ s_0 &= \sqrt{\sum \frac{(\bar{w}_m - \bar{x}_0)^2}{M_2 - 1}}\end{aligned}\tag{19}$$

for $c1 \leq \bar{w}_m \leq c2$ where M_2 is the total number of data remaining after the abnormal data was removed.

After the sample mean and sample standard deviation of normal traffic is calculated, the IDS alert criteria can then be calculated using equations 17 and 18. This is where the second round of the hypothesis testing application appears. Since a new sample standard deviation was determined, a new sample size must be calculated. The sample size of the moving window was re-calculated using equation 11 where $a = 0.001$, $b = 0.001$, $S = s_0$, and $m_1 - m_2 = 3s_0$. The values of a and b , again, represents risk. The equation yields $N = 4$. The window matrix \mathbf{W} was then re-populated with events of sample size N , and \bar{w} , \bar{W} and S_w were re-calculated as well. By making the proper substitutions, equations 17 and 18 yield $c1 = 50$ and $c2 = 657$. Since \bar{w} is nearly Gaussian, and then over time, there is a 0.1% probability that \bar{w}_m would fall outside of $c1$ and $c2$. As there was no concern with events that underperformed, $c1$ was not applied; the only concern was the upper bound, critical value, of $c2$. That is not to say $c1$ is unimportant, not seeing the usual amount of traffic could indicate something is wrong such as the IDS or network is down or IDS has been compromised among other countless possibilities. There may be other mechanism already in place to detect such circumstances, so in this application, $c1$ was ignored. Consequently, any events above the critical value “disprove” the null hypothesis – an IDS alert would ensue. The null and alternate hypothesis can be noted as

$$\begin{aligned}H_0 &= m \leq 657 \text{ events} \\ H_A &= m > 657 \text{ events}\end{aligned}\tag{20}$$

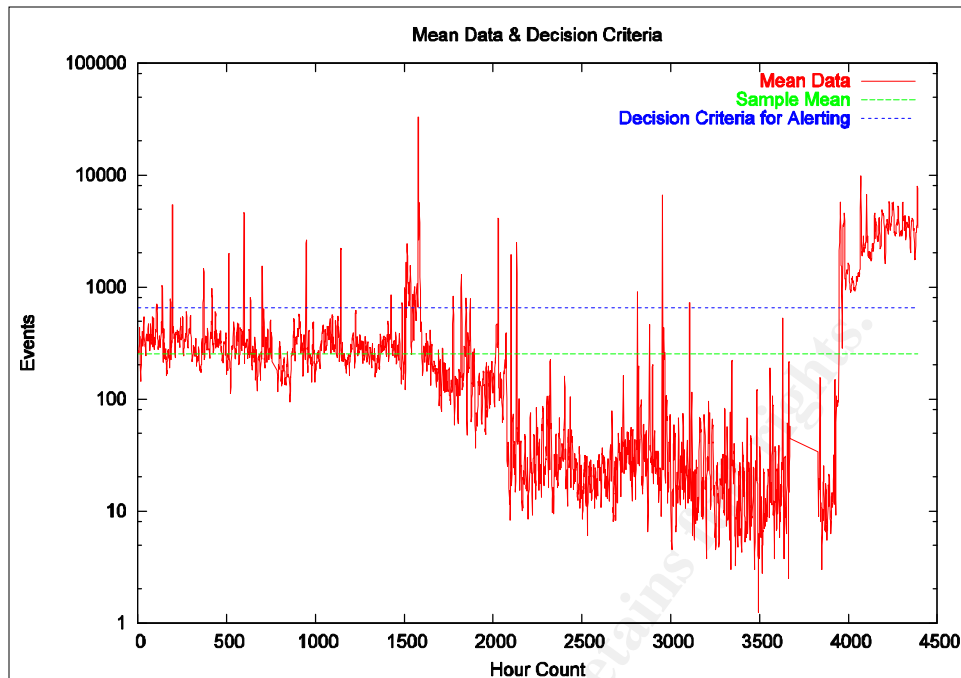


Figure 6 – Mean data, \bar{w} , for severity 3 data set with decision criteria used for null hypothesis and sample mean. The y-axis is the event count in log scale. Any value above *Decision Criteria for Alerting* would “disprove” the null hypothesis. Thus, an IDS alert would be generated.

The calculated alerting criteria can be used over time until the quality of alerts begins to degrade due to change in network activity, since “normal and attack traffic evolve over time”. [3, pg 19] Network traffic is not static; thus, the normal and attack traffic evolves over time. If the quality begins to degrade, the data requires a re-calculation of the alerting criteria, and the entire process discussed in this paper will have to be implemented. Yet, since this process can be easily automated, there is no required manual process.

The hypothesis testing process discussed throughout this paper was applied to the other data sets (severity 2 and 1) as well; the results are shown below.

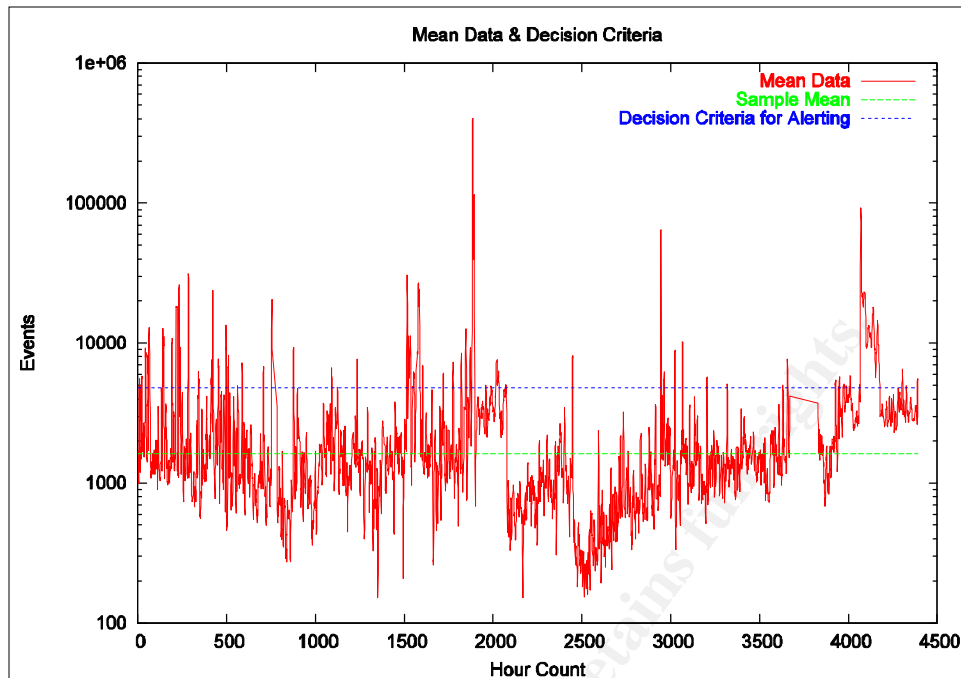


Figure 7 - Mean data, \bar{w} , for severity 2 data set with decision criteria used for null hypothesis and sample mean. The y-axis is the event count in log scale. Any value above *Decision Criteria for Alerting* would “disprove” the null hypothesis. Thus, an IDS alert would be generated.

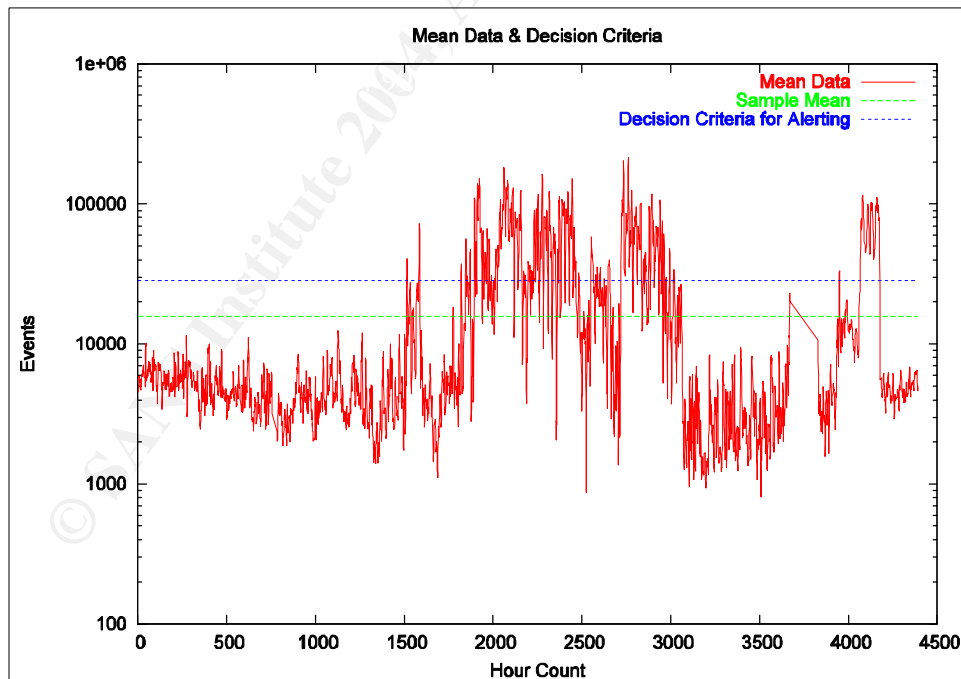


Figure 8 - Mean data, \bar{w} , for severity 1 data set with decision criteria used for null hypothesis and sample mean. The y-axis is the event count in log scale. Any value above *Decision Criteria for Alerting* would “disprove” the null hypothesis. Thus, an IDS alert would be generated.

Conclusion

Hypothesis testing was applied to IDS summary data to determine alert criteria. The hypothesis testing procedure was used twice. Hypothesis testing was initially used to eliminate data from the IDS summary data set to gauge normal traffic. The hypothesis testing was performed using the mean of a moving window, in order, to be able to invoke the central limit theorem. Once the normal traffic statistics were estimated, a hypothesis test was used again to determine an alerting criterion. The alert criterion was then calculated for the null hypothesis. The purpose of this procedure was to detect changes in the sample mean through inherent noise in the measurements. However, there is always a chance that a sample mean will be misleading - this is the risk or the total probability of making a wrong decision. The risk can be calculated using conditional probability by multiplying the risk of the first hypothesis test application (5%) to the second 0.05% (since c^2 was only used) resulting in 0.025% risk or 99.975% confidence.

For the severity 1 data set (see figure 9), the data seems to need further refinement and analysis. It appears that the normal traffic should be more confined and that the decision criterion is high; this is because the abnormal data deviates significantly. The null hypothesis must be refined to eliminate more abnormal data by adjusting risk and/or critical difference terms or introducing more data.

Hypothesis testing provides a scientific way of determining traditional thresholds. Many implementation of security monitoring products require the user to provide a threshold value such that the system can alert if a threshold is exceeded; yet, how are those values determined? Well, administrators mostly just eye-ball the value which is not based on anything of value. Hypothesis testing provides statistical meaning to this value.

In real-world use, the hypothesis testing application discussed in this paper could be applied to a more granular view of the data such as per intruder or IDS signature not necessary merely severity level.

References

- [1] Axelsson, Stefan, "The Base-Rate Fallacy and the Difficulty of Intrusion Detection". ACM Transactions on Information and System Security, Vol. 3, No. 3, August 2000, Pages 186-205.
<<http://delivery.acm.org/10.1145/360000/357849/p186-axelsson.pdf?key1=357849&key2=2948174701&coll=GUIDE&dl=ACM&C FID=11111111&CFTOKEN=2222222>>
- [2] Barnes, J. Wesley. "Statistical Analysis for Engineers and Scientists – A Computer-Based Approach". McGraw-Hill, 1994.
- [3] Mell, Peter, Vincent Hu, Richard Lippman, Josh Haines, and Marc Zissman. National Institute of Standards and Technology – Computer Security Resource Center. <<http://csrc.nist.gov/publications/nistir/nistir-7007.pdf>>
- [4] Messmer, Ellen. "Security debate rages". Network World, October 6, 2003. <<http://www.nwfusion.com/news/2003/1006ids.html>>
- [5] "Maximizing the Value of Network Intrusion Detection". Product Management Group of Intrusion, Inc., <<http://www.intrusion.com/products/download/maximizingvalueIDS.pdf>>
- [6] Snyder, Joel. "False positive remain a major problem. Network World", October 13, 2003. <<http://www.nwfusion.com/reviews/2003/1013idsalert.html>>
- [7] Snyder, Joel, David Newman and Rodney Thayer. "What network IDS can – and can't – do". Network World, October 13, 2003. <<http://www.securityfocus.com/infocus/1463>>
- [8] Tanase, Matthew. "One of These Things is not Like the Others: The State of Anomaly Detection". July 1, 2002. <<http://www.securityfocus.com/printable/infocus/1600>>
- [9] Timm, Kevin. "Strategies to Reduce False Positives and False Negatives in NIDS". September 11, 2001. <<http://www.securityfocus.com/infocus/1463>>
- [10] Weisstein, Eric W. "Central Limit Theorem". From *MathWorld*--A Wolfram Web Resource. <<http://mathworld.wolfram.com/CentralLimitTheorem.html>>
- [11] Weisstein, Eric W. "Sample Variance". From *MathWorld*--A Wolfram Web Resource. <<http://mathworld.wolfram.com/SampleVariance.html>>
- [12] Weisstein, Eric W. "Student t-Distribution". Wolfram Research. <<http://mathworld.wolfram.com/Studentst-Distribution.html>>

Appendix A – Tool used for computations and plotting

GNU Octave was used to calculate and plot all steps in this paper. GNU Octave is [free software](http://www.octave.org/) that can be obtained at <http://www.octave.org/>.

The following script was used:

```
N = samplesize(.05, .05, std(Y), std(Y)); %Est. initial sample size
N = round(N);
U = window(Y,N,N,M-N); %Raw data into moving windows

mu = mean(U); %Take mean of each window
sigma = std(U); %Std. of each window

mu0 = mean(mu); %Mean of mean of all windows
sigma0 = std(mu); %STD of mean of all windows

XX = [mu0-3*sigma0:mu0+3*sigma0]; %Possible Event count
YY = normal_pdf(XX,mu0,sigma0^2); %Probability of event count

[C1,C2] = singlemean_c(mu0,0.05,sigma0,N); %1st critical values used

idx = find((mu>C1)&(mu<C2)); %remove abnormal data

%%ALERTING CRITERIA STARTS HERE!!
mu0_1 = mean(mu(idx)); %Take mean of moving windows
of data without abnormal data

sigma0_1 = std(mu(idx)); %Take std of moving windows
of data without abnormal data

N = samplesize(.001, .001, 3*sigma0_1, sigma0_1); %sample size for
null hypothesis
moving windows
used to determine
null hypothesis
critical values
or alerting
criteria

N = round(N);
U = window(Y,N,N,M-N); %Data in moving windows with
new sample size

mu = mean(U);
sigma = std(U);

[c1,c2] = singlemean_c(mu0_1,0.001,sigma0_1,N); %null hypothesis
critical values
```

```

function sN = samplesize (alpha, beta, cDiff, sigma)
% sN = samplesize (alpha, beta, cDiff, sigma)

z_alpha = normal_inv(1 - alpha);
z_beta = normal_inv(1-beta);

sN = ((sigma^2) * ((z_alpha + z_beta)^2))/(cDiff^2);
endfunction

function [c1, c2] = singlemean_c (u, alpha, sigma, n)
% [c1, c2] = singlemean_c (u, alpha, sigma, n)
% Decision Criteria
%% if n (sample size) is too small (i.e. less than 30)
%% Then use t distribution is used
if(n > 30)
    z_alpha = normal_inv(1 - alpha);
else
    z_alpha = t_inv(1 - alpha, n-1);
end

c1 = u - (z_alpha/2) * (sigma/sqrt(n));
c2 = u + (z_alpha/2) * (sigma/sqrt(n));
endfunction

function U = window(u,I,M,N)
% Creates a set of vectors from input u
% U = [u(I-M),...,u(I+N-M)]
%      [...      ...]
%      [u(I),...,u(I+N)]

U = zeros(M,N);
for i = 1:N
    if (i+I-1) <= M
        U(M-(i+I-2):M,i) = u(max(1,I-M):i+I-1);
    else
        U(:,i) = u(i+I-M:i+I-1);
    end
end
end

```