

Checklist for Deploying an IDS

Andy Cuff 2003-12-30

1.0 Introduction

Installing a Network IDS (NIDS) onto a network requires a significant amount of thought and planning. In addition to the technical issues and product selection there are resource issues, from product cost to manning the sensor feeds and supporting the infrastructure that must also be considered.

The scope of this article considers the worst case scenario, that of deploying a NIDS on a remote network (target). The introduction of an IDS into a organization's network can be sensitive and often has political implications with the network staff, and thus a checklist written from the perspective of an outside consultant (even if the IDS is deployed internally) that appeases all parties can be useful to ensure a successful implementation.

While this topic is broad, there's sufficient information and planning required to form the basis of the checklist. If you are unfamiliar with the terminology in this article, please refer to my previous SecurityFocus articles [[A to H](#)] [[I to Z](#)] on IDS terminology.

2.0 Pre Deployment

2.1 Determining Policy, Junction of Maintenance

When installing an IDS a policy needs to be developed to ensure responsibilities are clearly defined. This is especially important when delivering an IDS capability remotely or to another organisation's network. The Junction Of Maintenance (JOM) defines where your responsibility for the hardware starts and finishes, and this will usually be the network switch port or tap with which the IDS connects to the target network. On the subject of failing hardware, people administering the target network must be made fully aware that if network taps are used, even fail safe taps can take up to a second for the interfaces to re-negotiate and could potentially disrupt services, though recent improvements have reduced this latency considerably. If the network is remote then it is advisable for the policy to reflect that the target network manpower can be called upon for a predefined duration for power resets, etc. Attempting this retrospectively through contractual alteration, if required, can be expensive and time consuming. If you rely on the distant network for support, ensure you have a telephone authentication system in place and don't fall victim to a social engineering attack. It's all too

easy for an attacker or Pen Tester to call the local staff where your IDS is installed and ask them to power it down. Most of these issues can be avoided if you are willing to have your IDS application reside on one of the target network's hosts, though in my experience it can never be completely trusted and raises the question of who maintains the software and OS deployed on the system. If an OS update corrupts the IDS application, then who takes responsibility for fixing it? Finally, discuss and set in policy the rules of engagement for automated response. This is especially important when you are deploying Intrusion Prevention Systems.

2.2 Comparing IPS versus IDS

An Intrusion Prevention System or inline IDS will block packets that meet the criteria of an event signature. These packets could have legitimately been accepted by the firewall and allowed through. As signatures can block packets in a fashion similar to a firewall, there are some that advocate replacing firewalls with IPS. I feel this is a dangerous step. In my opinion an IPS complements the firewall very well and they work well together, but the firewall should be left in place.

This is a good time to mention the dreaded false positives. The myth that an IPS will kill a network through its false positives doesn't have to hold true. It can be set to simply alert rather than block, thus only blocking those packets where the likelihood of a false positive is very low. An IDS can work in a similar fashion to an IPS, though. Rather than blocking packets in line, it can craft various responses: TCP resets to the source or destination (or in some cases both) of the offending packet, crafting unreachable/unauthorized replies and spoofing the border device. A big seller for stateful IPS is preventing the leaking of confidential information from an organisation. For example, you might want to retain corporate knowledge by blocking any document that contains the word "prototype", from leaving the network through the use of an IPS signature.

If the site has a policy for accepted use at their gateways, it is essential to use this to build the policy for your IDS. For instance, there is little point in reporting POP3 usage if it is permitted. There is some value in recommending changes to policies if they are blatantly insecure, but be careful not to oversell the issue and alienate the other network staff. At this stage your priority is to simply get the IDS or IPS in the door. Once the IDS starts chattering you can revisit those "practises dangerous to security". Policy also needs to be defined regarding how you respond to an incident and should include statements that direct forensics and evidence preservation activities. Furthermore, what assistance can you expect to receive from the site itself following

an incident and what actions are they expecting you to complete?

2.3 Gaining IDS mindshare

Gaining the trust of the target network's staff is imperative to a successful installation. I always find that the target network sys admin's concerns regarding intrusion detections are not necessarily focused on external traffic gaining access to their network, but what you will see of their poor practices and, more importantly, how you will react. I find that an amnesty of a predetermined duration, say a month, is a great icebreaker, where all detections are discussed with sys admins before their escalation. Possible exceptions to this must be discussed or you may destroy any trust, such as with high risk intrusions that occur out of office hours and detections of such a nature that they must be reported to the authorities under duty of care (i. e., discovering paedophilia). Winning the hearts and minds of all network staff will pay dividends when it comes to reducing false positives, post install.

2.4 Gather network topology diagrams

The availability of up-to-date network diagrams is essential not only for locating the best site for an IDS, but also post installation. The IDS analysts must be able to understand the events that they are seeing and how they relate to the network. They should include: network devices, bandwidth, transmission media, IP addresses, sub-netting information, default gateways, operating systems, host names, applications, and more. Be prepared to create them yourself as many of the diagrams out there are no more than scribbles on the back of a cigarette packet. Whilst you are making enquiries, ask if there are any planned upgrades over the next 12 months which may affect your choice of sensor, especially bandwidth (replacing a 100mb sensor for a Gigabyte sensor could be very costly if it is only six months down the line).

2.5 Identify physical infrastructure

You have to identify your requirements for the installation such as rack space, switch/hub ports, power outlets, UPS, cooling, taps and any mandatory local requirements like fiber infrastructure and fail over. Where you reserve any of the above for your use it is best to label them, otherwise rest assured, they will be gone when it comes to installation day. Get/make diagrams of everything: rack layouts, room layouts, building layouts, etc. I can't stress this enough, staff move on and it's all too easy to lose your IDS in a large data warehouse. Identify telephones close to where the IDS will be located, which is great for remote reboots. I asked one IDS vendor to include a feature that allows me to blink all the lights on the front of the box when I

wanted distant users to reboot; this was after some bright spark cycled the power on the firewall by mistake. HP has a great solution for remote access with the "Lights-Out" card on their DL servers. This great device then gives you full KVM into the machine so you can do all the normal KVM things from the other side of the world. In addition you can power up from off, even adjust BIOS settings. All that is required is another power socket and an extra IP address.

2.6 Sensor topology discussion

Some considerations need to be made when installing your IDS:

Management. You need to decide whether you manage your IDS Inband, Outband or Pseudo Outband.

NIPS or NIDS. As discussed earlier, this depends on your requirements. My preference is for NIPS though configured to be fully open and not blocking.

Method of Connection. Do you use taps or the span port on a switch? Obviously a NIPS must be inline, though I have seen an IPS tap which supports a backup IPS should the first fail. The tap is configurable to fail open or fail closed. The sensor will need to be installed either in a rack or a desktop PC, depending on the deployed infrastructure. Will you need rack shelving and will your equipment be too deep for the racks?

2.7 Site Access

If this is a remote installation, are there any restrictions in accessing the site? Further, if you are not employed by the target organization, will your staff need to be cleared before arrival?

2.8 Network Name

This is self explanatory but essential when dealing with large data centres.

2.9 Network Function

This is useful for the analysts when responding to an incident.

2.10 Target date

When is it convenient to install the IDS? If downtime is required this will likely need to be

provisioned in advance. Will the install have to be outside of normal hours?

2.11 Points of contact

Identify all key players from network staff to security staff, both network and physical, and have them listed in one place in case of a major incident.

3.0 Pre- Installation Phase

3.1 Procurement

It should go without saying, but shipping delays can occur that setback the installation. Ensure you procure the equipment well in advance, and on arrival lock it away.

3.2 Quarantine

Delays often occur, therefore once the sensor is built put it somewhere safe, well away from those well-intentioned "borrowers". I like to place mine in an IDS crèche, all connected up, where it can report to a manager, receive updates and be soak tested, and best of all nobody can "borrow" any bits unnoticed.

3.3 Building

When you build an IDS/IPS think secure install at all times. Harden it as much as possible as there is nothing worse for your credibility than having an IDS rooted.

3.4 Testing

It is useful to soak test an IDS before deployment, ideally in the same configuration as you will use on site. Make sure you test any taps for failing in the correct manner, either open or closed. If possible, fluctuate temperature and place the IDS under considerable network load (warm it up). Ensure your own router, firewalls and VPN permit access from the remote network; perform simulated remote updates to the IDS and operating system.

4.0 Installing the sensor

If all the above has gone according to plan, the installation should be a dream. Whilst you are still onsite, it is probably worth investigating the initial events with the site admins. If a third

party has installed the IDS try to call the site ASAP to clear off those initial (often many) false positives. Do not try to make assumptions. The site sys admins are the experts on their network and should be able to provide the technical feedback to those first few events.

5.0 Post installation feedback

False positive tuning is essential. The first coarse tuning should have occurred by using the site's policy to define the initial IDS policy. Subsequent fine tuning should be carried out periodically. Rather than attempting this on an event-by-event basis, wait a week and look at the historical information, sorted by count. Rather than adjust the IDS policy to reduce false positives, I find it easier to try and cure the source of the alert, which will require some site interaction. For instance, if you see alerts for SNMP "public" community string ask the site to change away from the default rather than ignore the event. You will need to provide as much information about each false positive as you can in a tabular form and also, if possible, suggest a course of action to the administrators for tuning/patching the relevant system. If this course of action cannot be completed then you will have to take IIDS tuning action, which could include filtering the source or destination address, removing the signature entirely or reducing it's severity. All tuning needs to be fully documented, and do not forget to ensure that there is a column in the documentation for the regular network staff to comment on your recommendations. The second false positive reduction period should occur at around the one month stage. After this point the IDS should be singing sweetly (depending upon your choice of IDS) and false positives could be dealt with on an individual basis.

6.0 Summary

The above information has been compiled after many years of installing IDS on a variety of networks. While this checklist is not comprehensive, it should give the reader an insight into some of the lessons I have learned the hard way.

Author Credit

View [more articles](#) by Andy Cuff on SecurityFocus.

[Privacy Statement](#)

Copyright 2006, SecurityFocus