


Scanning, fixing, and reporting security issues with Security Blanket

TIGHT SHIP



Trusted's Security Blanket lets you analyze security compliance in a few easy steps. **BY KURT SEIFRIED**

I don't know about you, but I spend entirely more time locking down my servers and making sure they stay locked down than I would like. At least once in the past (that I know of), one of my servers has been compromised. I failed to a) upgrade WordPress and b) ensure that my server was locked down so local access didn't allow an attacker to escalate privileges easily. The problem was not that I didn't know how to lock down my servers or have the time to do so; I just had other things to do, and trying to keep up with everything you need to do to secure a server and keep it secured is not my idea of fun.

So, what am I, or any administrator, supposed to do if we have dozens or hundreds of servers to lock down with varying levels of security, to which we are applying updates, installing new software, and generally mucking about with on a daily basis? And what about the administrators that have to deal with compliance issues like PCI-DSS or the various government standards (which are less fun to read than RFCs)?

Enter Security Blanket

Security Blanket [1] is a software package from Trusted Computer Solutions, a company with a long history in the government and compliance space. The basic premise of Security Blanket is that

automated tools make compliance easier, and automated tools that know what you need to be compliant with make it even easier. Security Blanket employs a client-console model (allowing multiple consoles) that allows anything from a single machine to lots of machines (a maximum of 1000 per console is recommended). Conceptually, Security Blanket is very similar to Puppet [2]; it has an encrypted communications channel and a variety of modules on the client that can take actions (e.g., turning things on or off, changing configuration settings, etc.). On the client side of Security Blanket is the dispatcher that listens for commands and sends responses.

Installing Security Blanket

Installation is pretty simple and well documented. Once you unpack the tarball and run the script called *SB_Install*, the script offers some options like installing the client software, the console, or both. If you are installing a stand-alone system, you will need the console and the client. If you plan to have multiple clients and a console, then the con-

sole does not have to have the client software installed.

Next, you will be prompted to run *cert_gen.sh*, usually located in the */usr/share/security-blanket/tools/* directory. Note that a bug in the install script requires you to copy *cacert.pem* and *Disp.pem* manually to the */var/lib/security-blanket/files/certs/* directory (TCS has said they will have this fixed in an upcoming release). Once you have installed the certificates, you'll have to run *SB_Setup* in the */usr/share/security-blanket/tools/* directory. Finally, on the console, you need to install the license key; it's the standard gibberish, so cut and paste from the email they send you, and you'll be good to go.

Prerequisites for Security Blanket aren't too complicated. You'll need Java on the console (it's Tomcat based), and on the clients, you will need the PyXML library (otherwise you get an error saying that it is missing if you try to send commands to the client).

Basic Configuration

Once you have a console and some clients, what happens next? To set up clients, you add them to a group, apply a profile to the group, and that's basically it. You have a choice of eight default profiles (Figure 1), and you can modify or create your own profiles (more on this later). Once

Name	Summary
CIS Benchmarks	Center for Internet Security Benchmarks
DCID	DCID 6/3
DISA UNIX STIG	UNIX Security Technical Implementation Guide
FERC CIP	Critical Infrastructure Protection
JAFAN	JAFAN 6/3
NISPOM	NISPOM Chapter 8
PCI DSS	Payment Card Industry Data Security Standard
Web Services Protection	Web Services Protection (SANS LAMP)

Figure 1: Default profiles support by Security Blanket.




Figure 2: Baseline scan results.

you have applied a profile to a group and added clients to the group, you can scan them (against that profile) and apply that profile to the systems in question.

Scanning

Security Blanket supports three types of scans: a baseline scan, a security scan, and a quick security scan. The baseline scan is not a security scan (Figure 2); it collects information about the host, such as name, distribution, hardware devices, network configuration, and installed packages. The scan and quick scan are the same, except the quick scan doesn't run system-intensive modules or slow modules. However, on anything but a heavily overloaded server, I strongly recommend using the full scan because the quick scan might miss things.

Once you click the *Scan* button, the console sends a command to the dispatcher running on the client(s). Once you have sent a command to a client, it must complete before you can send more commands. Unfortunately, Security Blanket does not currently have a way to display which commands have been sent and are waiting for responses, so if you need to run a command on a host, you might need to wait. Once the scan completes, the client will connect back to the console, give the results of the scan, and create a notification alert in the web interface – a little bit of red text will be added to the top of the interface, letting you know how many notifications are outstanding.

As you can see from the initial scan (Figure 3), a default install of Fedora 11

is not exactly what you would call PCI-DSS compliant (93 failures, 47 passed, and 26 other). I have to admit I was curious as to what exactly failed, and the nice thing about the report is that you get a complete listing of each module that ran along with the output (Figure 4). If you click on the title of the issue, you'll get a description of the problem (e.g., *Disables rsh*) with a description of why it's probably a good idea to fix it and which security standards require this.

Applying a Security Profile

Obviously I have a problem (viz., 93 failures) that should be addressed. The solution is as easy as hitting the “apply” button and waiting a few minutes. When you run the “apply,” it's a lot like run-

ning a scan; the command is sent to the client, and the client executes it and returns the results to the console, which then creates a notification event. If any errors occur (e.g., a module fails to run properly), they will be mentioned on the notification screen and displayed in the report.

As you can see, a scan after the profile apply shows far fewer problems (Figure 6). In my case, it failed (Fedora support isn't finished yet) to deal with *su* properly; also, there was an error with SNMP and something about system log permissions.

Summary

Security Blanket pretty much does what it says on the tin; it's pretty easy to set up and configure, and usage is ... well ... simple (hit *Scan*, hit *Apply*, and customize the profile as needed).

So, why pay the money for this product rather than using a system like Puppet, which is open source and free? A number of features make Security Blanket worth paying for. The following sections highlight some of the reasons why Security Blanket might be well worth the investment.

Profiles

The first and most important of the features that make Security Blanket worth paying for is the prebuilt security profiles [3][4][5][6][7][8][9]. (I couldn't find the SANS LAMP policy online.) Just for laughs, I downloaded the PCI-DSS standard and started reading it. Some

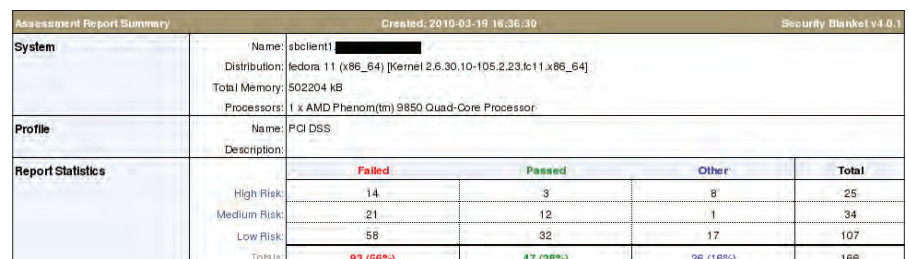


Figure 3: Scan results for a new host (lots of issues).




Figure 4: SSH results, expanded.

parts of it are pretty clear, like Requirement 5: “Use and regularly update anti-virus software or programs.” That part is pretty self-explanatory.

However, in section 8.5, more than a dozen specific issues deal with passwords, ranging from minimum password strength to account lockout time (30 minutes) and idle session logout time (15 minutes). Implementing these password restrictions would mean changing a huge number of settings – from system password policies to screen savers (locking idle sessions) – and specific services that support logins (like FTP). Having prebuilt profiles and modules to implement these changes results in a huge time savings.

Regulatory Compliance

Again, there’s the dreaded “C” word – compliance. The reality is that in most organizations that are driven by compliance or have to play by compliance rules, it doesn’t matter how well you lock down a system unless you can prove that it has been properly locked down with a nice audit report. This requires, first, some sort of mechanism to scan the machines and, second, a list of things to scan for – which is not yet available in products like Puppet, as far as I know.

The Security Blanket reports also categorize vulnerabilities (High, Medium, and Low risk) and provide numerical output, which is something managers love. (The most common business theory I have heard says if you can get numbers out of it, you can measure and control it – shades of 6 Sigma and Total Quality.)

Automation

One of my favorite features in Security Blanket, however, is the ability to schedule actions, especially chains of actions. For example, you can schedule a group of hosts to run a scan, followed by an

Apply command, and then follow this up with a second scan at say 4am every day. This will tell you whether the hosts are changing (i.e., before the apply, there were new security issues) and ensure that system updates and other changes aren’t undoing security (and if they are, they will be fixed and reported).

The reality is that any security or backup task that isn’t automated probably won’t get done (like my failure to upgrade WordPress on a weekend).




Figure 5: Security Blanket client commands.

Undo

I didn’t even know about this feature until I asked one of the support engineers why they disabled programs by simply removing the executable bit (as opposed to removing the file, uninstalling it, etc.). Turns out this is because Security Blanket can undo pretty much anything it does.

So, if you accidentally tighten the security settings too much, or break something on a critical server, you can quickly undo it. This means you’ll have more time to figure out what went wrong, as opposed to having to fix it under pressure and then figure out what went wrong.

What’s Missing

One thing I’ve noticed in a lot of product reviews is that you only hear about the good stuff, and the reviewer neglects to mention what doesn’t work or is missing from a product. So, what is not so great about or missing from Security Blanket? My biggest wish would be diff’ing of scan reports. I really don’t want to see the whole report every single time (my

eyes start to glaze over); instead, I would much rather have a diff of the current scan report to either the previous scan of that host, or against a baseline system. I am told this is on the roadmap, and I hope it gets done, because it would be a great feature (it would really compact the amount of information you need to see).

Conclusion

So, should you spend your money on this product? If you need to deal with compliance and auditing issues, it will definitely help. And, if you’re stuck with government standards and mandatory compliance, then this is probably an especially good idea.

Even if you don’t have to teach to the test, as it were (i.e., you’re not beholden to these standards), I really like the starting base provided by these policies and the ease with which they can be modified to suit a specific installation.

Also, I’m a little surprised that more products don’t have reliable “undo” features (I certainly wouldn’t use a word processor that didn’t have *Undo*, but then I walk on a tight rope with no safety net when I administer my systems all the time). Overall, I like this product – mostly because it actually does what it claims it will do, and it does it pretty painlessly. ■

INFO

- 1] Trusted Computer Solutions – Security Blanket: <http://www.trustedcs.com/Security-Blanket/SecurityBlanket.html>
- 2] Puppet: <http://projects.puppetlabs.com/projects/puppet>
- 3] CIS benchmarks: <http://cisecurity.org/>
- 4] DCID 6/3: http://www.fas.org/irp/offdocs/DCID_6-3_20Manual.htm
- 5] DISA Unix STIG: <http://iase.disa.mil/stigs/stig/unix-stig-v5r1.pdf>
- 6] FERC CIP: <http://www.ferc.gov/industries/electric/indus-act/reliability/cip.asp>
- 7] JAFAN 6/3: http://www.lazarusalliance.com/horsewiki/images/f/fa/JAFAN_6_3.pdf
- 8] NISPOM: http://www.fas.org/sgp/library/nispom/5220_22m2.pdf
- 9] PCI-DSS: <https://www.pcisecuritystandards.org/>

Assessment Report Summary		Created: 2010-03-30 02:04:03		Security Blanket v4.0.1	
System	Name: sbclient1	Distribution: fedora 11 (x86_64) [Kernel 2.6.30.10-105.2.23.fc11.x86_64]	Total Memory: 502204 KB	Processors: 1 x AMD Phenom(tm) 9850 Quad-Core Processor	
Profile	Name: PCI DSS	Description:			
Report Statistics		Failed	Passed	Other	Total
High Risk:	0	17	8	25	
Medium Risk:	1	32	1	34	
Low Risk:	3	87	17	107	
Total:	4 (2%)	136 (82%)	26 (16%)	166	

Figure 6: Scan results for a the new host after the security profile is applied.