Opus One/Network World
SSL VPN Testing Methodology
V3
Joel Snyder/jms@opus1.com

Based on the review criteria, here is our test plan. Filling out this test plan should include screen shots of all relevant areas! Label screen shots with test numbers/evaluation points so they don't get lost.

1. AUTHENTICATION.

Test authentication by connecting to our different authentication servers. For each of the authentication methods, verify proper operation. For each authentication method, verify that authorization is possible (i.e., group retrieval).

1.1: local database authentication (simply verify that this exists)

1.2: RADIUS server. Connect to radius.opus1.com and validate that users can authenticate and that group information from the CLASS attribute can be retrieved. CLASS should be allowed to be a list of groups.

1.3: LDAP server. Connect to ldap.opus1.com and validate that users can authenticate and that group information from memberOf attribute can be retrieved. Should be able to have multiple groups returned.

1.4: Windows server. Connect to the Windows domain citrix.opus1.com (AD server is LENIN). Verify that a user can be a member of multiple groups. NOTE: this is not AD-through-LDAP, but direct to AD connection.

1.5: ACE server via RADIUS. Connect to ACE server on Thrip and verify that SecurID-specific things (new PIN mode) all work properly. Verify that group information is retrieved (same as 1.2, basically)

1.6: List other authentication methods claimed to be supported, such as NTLM. List other authentication systems (such as single sign-on integration) that are supported.

1.7: Digital Certificates. Validate that the PKI implementation allows us to issue a digital certificate from our CA, to retrieve the CRL from the CA, and authenticate users using digital certificates with support for CRLs. How are users mapped from certificate CNs to 'user' objects? How are groups retrieved and mapped? What other options are there for certificate-based authentication (such as certificate+password authentication)? Validate CRL support using revoked certificate.

All further testing is going to be done with the LDAP directory.

LDAP DIRECTORY

Everybody is cn=<name>, o=Opus One
Groups are in memberOf structure
with name cn=<groupname>, ou=Groups, o=Opus One

cn=Joel Snyder	groups=Boys
cn=Jan Trumbo	groups=Girls
cn=Chris Janton	groups=Boys
cn=Romeo Julieta	groups=Girls, Cats
cn=Wanda Rutkiewicz	groups=Girls, Cats
cn=Oliver Mellors	groups=Boys, Cats
cn=Claudia Augusta Tiberius	groups=Girls, Cats

2. FINE GRAINED ACCESS CONTROL

Test fine-grained access control by implementing our security policy in the device. Start by defining resources, and then move to the actual policy. (Note that end-point security is evaluated elsewhere, although how the EPS is used to influence policy should be described here).

RESOURCES:

Application-based:

Citrix server (Stalin,Gorky,Rykov)
Terminal Services server (Rykov)
Telnet server (Bass)
SSH server (Bass)

File-based:

SMB server (Ginger)
FTP server (Bass)

Web-based:

PMDF: Bass:7633; :7433 (http/https)
Webmail: Sloth:80
WUG: Thrip:88
Power: r4b.ap9211:80
Exchange: zinoviev:80
Domino: trotsky:80
KVM: dsr8020:80

IP-based:

VOIP: 192.245.12.176-.181, and .224
"Opus C": 192.245.12.0/24

POLICY:

GROUP CATS has access to all resources, at all times
GROUP BOYS has access to the following resources:

When they pass End Point Security check:
Citrix
Telnet
FTP
all web services
entire Opus C (i.e., including VoIP)

When they fail End Point Security check:
Citrix; DSR8020
FTP GET but not PUSH
only PMDF:7633/monitor/, /doc/, and /dispatcher/
VOIP services, but not the rest of the C

GROUP GIRLS has access to the following resources:

When they pass End Point Security check:
Terminal Services
SSH
SMB GingerData/extranet/ directory
all web services
entire Opus C

When they fail End Point Security check:
Terminal Services; DSR8020
SMB GET but not PUSH
only PMDF:7633/monitor/, /doc/, and /dispatcher/
VOIP services, but not the rest of the C

END POINT SECURITY is defined as:

- if you are on a platform that gets viruses, then you have
  to have a virus scanner running
- if you are on a platform that does NOT get viruses, then
  you are by definition OK

2.1: How well is the product able to match our security policy? Does the policy actually work (i.e., does the product do what we hope it will do when we put in the policy)?

2.2: How many other types of security policy elements are available? What is the granularity of security policy? A partial list from the requirements is:

a) Source IP and Destination IP; Destination DNS
b) Destination port
c) Time-of-day; day of week; limits on date
d) User name
e) User group (e.g., OU information from a DN on a certificate or Windows Groups from the Windows SAM)
f) User identifiers held (e.g., "User can dialup" bit in Windows SAM or Service-Type=Administrative via RADIUS)
g) Application layer information (e.g., http GET versus PUT; URL data; file store versus file retrieval)
h) Status of client integrity check
i) Authentication method
j) Client platform (i.e., browser type; Java-enabled; cookie-enabled)
k) Strength of encryption

3. MANAGEMENT

Evalute the management, high availability, and scalability features of the product.

3.1: Management of configuration. How is the product managed? Is delegated management or partitioned management available? What are the relevant features here? How "easy" is it to manage? Does the management pass the Ned C test of "user friendliness?" What happens when multiple people log on?

3.2: Management of operations and monitoring. What operations and monitoring features are available? Can you see who is on? System load? Load over time? Can you kick someone off?

3.3: High Availability. How does the product achieve high availability? What are the management features related to this? Does this actually work (need to test!!!)? What are the restrictions on this? Single site? Multiple site?

3.4: Scalability. How does the product achieve scalability? How is the management handled? What level of replication and synchronization happens? Does it actually work? What are the restrictions on this? Single site? Multiple site?

3.5: SNMP. What SNMP is available? Traps? MIBs?

4. AUDITING, LOGGING, and ACCOUNTING

Management auditing, logging, and reporting. What logging and auditing features are available? Do the reports that are available make any sense? Do the reports get aggregated from multiple devices?

4.1: Auditing : every successful or unsuccessful user login and access needs to be audited via some logging mechanism; every management action needs to be audited; every system startup and shutdown needs to be audited.

4.2: Accounting: summary statistics on user sessions need to be aggregated and held via either logging or RADIUS accounting. This would include values such as time connected, data transferred in/out, and other relevant counters.

4.3: Logging: the recording of audit and accounting information needs to be managed in a versatile manner. Audit and log information need to be stored locally within some ring buffer, with flexible control if the buffer fills (e.g., overwrite, system halt, email or FTP outbound). Local buffers need to be retrievable from remote systems via common mechanisms, such as TFTP or FTP. SYSLOG and/or SNMP need to be available as options for real-time logging and health-check information. Can we generate local reports?

4.4: Debugging: what features for debugging are available? Can we see a single session? Can we run traceroute and ping? How about tcpdump? Is that a good idea?

4.5: Does the product have breakin detection/evasion? What other IDS-ish things are true? What does NMAP say?

5. CLIENT SUPPORT

Ignoring the question of application-layer support, what kind of client support is available?

Our client platforms are:

Windows 2000+IE5.5 (with administrator password)
Windows XP+IE (no administrator password)
Windows XP+Firefox (no administrator password)
Windows XP+IE (with administrator password)
Windows XP+Firefox (with administrator password)
Mac OS X+Safari
Mac OS X+Firefox
Palm Treo (Blazer)
Symbian Series 80 (Opera)
Random airport systems

5.1: Basic client support. Are there any unrealistic client requirements just to get logged in and working? Verify that our client platforms work.

5.2: Retargeting output. For our Palm and Symbian phones, is there any good small screen support?

5.3: Non-English support: can it handle multiple languages? Simultaneously, or only at once?

5.4: End Point Security requirement related to clients. [perhaps put this in EPS support]

5.5: Port Forwarding requirements related to clients. Can we run port forwarding on all our clients?

5.6: Network Extension requirements. Can we run network extension on all our clients?

6. SSL VPN CORE CAPABILITIES

Evaluate how well the product supports the four access modes/models for SSL VPN. Are there other models supported?

6.1: Reverse Proxy. Does this work with our test applications? What level of access control is available? (Describe!)

Web-based:

PMDF: Bass:7633, :7433 (http/https)
Webmail: Sloth:80
WUG: Thrip:88
Power: r4b.ap9211:80
Exchange: zinoviev:80
Domino: trotsky:80

6.2: Application Translation. Does this work with our test applications? Make sure to test push and pull! Large files! What are the vagaries of the file translation application? What operations are supported? What level of access control is available?

File-based:

SMB server (Ginger)
FTP server (Bass)

6.2.1: Are there other application translations supported? For example, webmail?

6.2.2: Are there other thin-client applications supported? For example, SSH/Telnet? Terminal Services?

6.3: Port Forwarding. Does this work with our test applications?

Application-based:

Citrix server (Stalin,Gorky,Rykov)

Terminal Services server (Rykov)
Telnet server (Bass)
SSH server (Bass)

What level of access control is available? What other features are present? Does this have auto-launch? Does it modify hosts files? (and then fix them?) Can it auto-create mappings for things like Exchange?

6.4: Network Extension: Does this work? How many clients is this compatible with? What is the performance? (how are we going to measure that???) What are the security features that are applied?

IP-based:

VOIP: 192.245.12.176-.181, and .224
"Opus C": 192.245.12.0/24

6.4.1: How is the client deployed, and how is the policy updated? Are normal client-side features such as split tunneling and local LAN access available? What interfaces does the client present to the operating system (i.e., shim, virtual adapter, PPP)? What measures does the client take to avoid TCP-over-TCP issues? What platform support is available in the client? How does the client operate in a “no admin access” environment?

6.5: What other access methods are supported? Clear text (a la Caymas)? SSL-ification of email SMTP/POP/IMAP? Anything else?

7. END POINT SECURITY

This is the hard one. Start by describing how the end point security works. Divide into integrity checking, behavior changes, and protective services.

7.1: How does integrity checking work? Can we check prior to authentication? Can we check during a session (or continuously)? What levels of compatibility do we have with our test platforms? What can be checked for? How does this integrate with off-box policy servers? How difficult is it going to be to keep this updated?

7.2: How does integrity checking change behavior? What are the knobs and options we have to change the behavior of the system if the check fails? Remediation servers? Denial of access? Modification of access? Notification to the user?

7.3: Protective services: What kinds of protective services are available: personal firewall, virtual desktop, cache cleaner?

7.4: Protective services on the device: What kinds of protective services does the device include (the Check Point/Fortinet question)?

8. USER WORKPLACE/PORTAL

We want to evaluate how the portal works.

8.1: Describe the portal model and how customizable this is.

8.1A: Virtualization: by IP? By URL?

8.1B: Login page. Does it work with small screens? Large screens (doh)? How customizable is it?

8.1C: Network manager's ability to customize the graphics, colors, logos on the portal. Ability to skip the portal entirely and redirect to a corporate portal. Is there an API?

8.1D: Network manager's ability to control what is actually shown on the portal

8.1E: End user's ability to control their own portal

8.1F: "portal management" in the context of applications. Does the vendor float a Javascript over the page? (does this work?) Or another page? How does the user get back to the portal?

8.2: End Point Model: what is the end point security model? Are cookies passed through to the end system or cached on the SSL VPN device? If a cookie is used for session persistence, can it be stolen and jumped to another IP address? Can users store passwords on the device or do they have to keep typing them in? Can users control this and clear it out? Does the device support single sign-on by passing through cached credentials? Is the product vulnerable to cross-site scripting attacks?

9. PERFORMANCE

10. EVERYTHING ELSE

10.1: Control of security parameters. Can you turn off SSLv2?
10.2: How does this look to a service provider? CLI? QoS?
10.3: FIPS-ish stuff.
10.4: Built-in IPsec, dynamic routing?
10.5: Disaster recovery (backup/restore)?
10.6: Documentation?
10.7: Appliancification?