Lawrence Berkeley National Laboratory
LBL Publications

Title
Recent Advances in Bro Intrusion Detection System

Permalink
https://escholarship.org/uc/item/35c285d0

Author
Sharma, Aashish

Publication Date
2012-10-01

Copyright Information
This work is made available under the terms of a Creative Commons Attribution-NonCommercial-NoDerivatives License, availalbe at https://creativecommons.org/licenses/by-nc-nd/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Recent Advances in Bro Intrusion Detection System

Aashish Sharma
Jay Krous

Berkeley Lab

Outline

1. Input Framework
2. Finding spear phishing
3. IPv6
4. Decapsulating tunnels
5. Instrumented SSHD
6. Vuln Java/Flash detection
7. Syslog2bro
8. Visual analysis
9. Fast searching
10. 100G
11. Internal Bro

What is input framework

- Framework for importing data
  - Adaptable to different sources
    - digest any kind of data feed
  - Simple, yet flexible user interface
    - allows you to manipulate data line by line eg. http://www.badurl.co.cc/baddir/badexe.jpg

Input framework

- Asynchronous operation
- Real-time operation
  - Handles 50,000 events/sec with increase in system load levels of .2%
  - updates and deletes from data files will get propagated to the right tables automatically
- Data processing outside of main thread
- Injection interleaved with packet processing
- Already deployed and being used

1: Input Framework - Leveraging the intelligence from the Community

Integrating Feeds from

- DoE CIRC, REN-ISAC (CIF), IID
  - IP address
  - Host names
  - URLs
  - MD5
  - Subject
  - Attachment
  - Sender
  - From
  - To (Mail_to, Reply_to)

Input-Framework -> Intel-Framework

Date: Sat, 20 Oct 2012 08:50:01 -0700 (PDT)

From: Big Brother <bro@bro.lbl.gov>

To: alert@lbl.gov

Subject: [IR (stomp) deep-bro] **Intel::SMTP_Malicious_Mailto**

Message: Malicious to:: [indicator=info@info.com, description=/cpp/tippers/JC3-689315-Federal-Tax-Phish/mail], info@info.com

Sub-message: info@info.com

Connection: 198.63.37.121:62484 -> 128.3.41.120:25

Connection uid: **7gw6BLZZNQb**

Email Extensions

---

orig/src hostname: bananabusaruba.com

resp/dst hostname: mailgate.lbl.gov

--

[Automatically generated]

More details

Oct 20 08:49:55 **7gw6BLZZNQb** 198.63.37.121 62484 128.3.41.120 25 1 bananabusaruba.com <test@bananabusaruba.com> <JOHNDOE@lbl.gov> Sat, 20 Oct 2012 15:49:53 +0000 "Mrs. Safia Farkash Gaddafi" <test@bananabusaruba.com> info@info.com mrssafiagaddafi2011libya@yahoo.com.ph

<20121020154954.66309.qmail@bananabusaruba.com> - From Mrs. Safia Farkash Gaddafi.... - (qmail 66310 invoked by uid 9593); 20 Oct 2012 15:49:54 -0000 - 250 ok: Message 88725239 accepted 128.3.41.120,198.63.37.121 MIME::Lite 3.0104 (F2.72; T1.15; A1.47; B3.01; Q3.01) F

Input-Framework -> Intel-Framework

From: Big Brother <bro@stomp.lbl.gov>
Subject: [IR (stomp) deep-bro] **Intel::JC3_SensitiveDNS_Lookup**
To: alerts@lbl.gov

Message: [host_host=XXXXXXX.cloudfront.net, host_added=2012-09-21 21:11:02, host_updated=2012-09-21 21:15:39, indicator_active=Yes, indicator_added=2012-09-21 21:11:02, indicator_changed=0000-00-00 00:00:00, indicator_removed=0000-00-00 00:00:00, indicator_state=block_and_report, threat_level=<uninitialized>, indicator_source=circ__jc3circ_soc__manual_entry] XXXXXXX

Sub-message: d1ros97qkrwjf5.cloudfront.net

Connection: XXXXXXXXXXX:64708 -> 128.3.34.186:53
Connection uid: lsbX50qvxg7

Email Extensions
---
orig/src hostname: XXXXXXXX-xp.dhcp.lbl.gov

resp/dst hostname: ns1.lbl.gov

--
[Automatically generated]

Input-Framework -> Intel-Framework

Oct 22 23:17:31 zePVN7xbtY8 2001:400:613:18::b5d 63687 2620:83:8000:140::3 53 udp Intel::JC3_SensitiveDNS_Lookup

[host_host=beacon-1.newrelic.com, host_added=2012-09-21 21:11:02, host_updated=2012-09-21 21:15:39, indicator_active=Yes, indicator_added=2012-09-21 21:11:02, indicator_changed=0000-00-00 00:00:00, indicator_removed=0000-00-00 00:00:00, indicator_state=block_and_report, threat_level=<uninitialized>, indicator_source=circ__jc3circ_soc__manual_entry]

Software framework

```
redef Software::vulnerable_versions += {
["Flash"] = [$major=11,$minor=4,$minor2=402,$addl="264"],
["Java"] = [$major=1,$minor=6,$minor2=0,$addl="34"],
["adobe_reader"] [$major=9,$minor=4,$minor2=6,$addl=""],
};
```

PHP, Wordpress, Mail clients, Android*, Apache...

(862 Unique applications identified with precise versions in today's logs)

Spear phishing: Looking for the targeted stuff

Spear phishing: Looking for the targeted stuff

- Came across: www.malware-tracker.com
- Ran some of the embedded pdf's and word documents collected over a period of time against it
- Needed some engineering to address volumes
  - python script - multiple buckets for clean, suspicious, infected
- At-least two new signatures issued by sophos after detection

Sample spear phish

Rating: EXPLOIT :1 MALWARE :1 SEVERITY :30 HITS :0 HAS_EXE :1

==========

Time: 1342007556.302290

From: <agulbra@nvg.unit.no>

To: <JOHNDOE@lbl.gov>

when_ts Wed, 11 Jul 2012 12:52:38 +0100

Subject: Re: Is that your document?

1342007557.718325 fPUHmL6ke86 212.58.56.90 64369 128.3.x.x 25 1 **part6.zip** 29832 application/zip **bb19060fde6e92bfaf5c585e56e3cb8e**

/**home/users/bro/extract**/smtp-entity_212.58.56.90:64369-128.3.x.x:25_1.dat (empty)

1342007556.302290 fPUHmL6ke86 212.58.56.90 64369 128.3.x.x 25 1 lbl.gov <agulbra@nvg.unit.no> <johndoe@lbl.gov> Wed, 11 Jul 2012 12:52:38 +0100 agulbra@nvg.unit.no johndoe@lbl.gov -

- - Re: Is that your document? - - - 250 ok: Message 80405646 accepted 128.3.41.146,212.58.56.90

IPv6

IPv6 monitoring

- LBL has Production wireless on IPv6
  - (LBLnet service, ACS, cluster on border
- Ability to monitor IPv6 traffic
  - http, dns smtp irc, ftp, ...
- Monitor IPv6 headers
- ICMPv6
  - icmp_router_solicitation, icmp_router_advertisement, icmp_neighbor_solicitation, icmp_neighbor_advertisement, icmp_redirect
- Constantly developing for for new attacks/trends and capabilities

4: IPV6 - Trying to find unknowns

- IP/mac address binding for dhcp jailing
  - ISC - yes, no, ummm, nah, not really, may be ....
- Tracker Ticket #833 extract the mac-address from the ICMPV6 using events
  - icmp_neighbor_advertisement
  - icmp_neighbor_solicitation

- Alerts on attacks on ICMPv6 protocol eg.
  - Rogue routers for fake router advertisements, build neighbor caches
  - Proactive response where a rogue RA results in another packet injected with lifetime of 0

Tunnel Decapsulation

Decapsulating tunnels via new tunnel framework

- Currently monitor
  - Teredo, AYIYA, IP-in-IP (both IPv4 and IPv6), and SOCKS
- Logs the outer tunnel connections are in both conn.log and tunnel.log,
- Proceed to analyze the inner payload as if it were not tunneled, while maintaining a new tunnel_parents column pointing back to the outer connection(s)

Tunnel.log:

Oct 22 23:59:59 **iqQAPRcFG37** 206.238.131.210 11437 131.243.168.204 65535 Tunnel::TEREDO Tunnel::CLOSE

Connections:

Child:

Oct 22 23:58:59 YEUXT7vEmzl 2001:0:519d:e224:4db2:d352:3111:7c2d 3104 2002:7c0c:5733::7c0c:5733 11447 udp **iqQAPRcFG37** worker-2

Parent: Oct 22 23:58:59 **iqQAPRcFG37** 206.238.131.210 11437 131.243.168.204 65535 udp teredo 0.009993 worker-2

Instrumented SSH

Instrumented SSH

- Keystrokes entered and responses sent with this version of SSHd is sent for analysis to Bro.
  - Sensitive information, such as passwords, is filtered out.

- Using various signatures, some complex and some fairly simple, Bro is able to alert us when an account appears compromised.

- Furthermore, once a compromise is confirmed, the logs from this version of SSH will help us determine the extent of the compromise and what, precisely, the intruder did.

- Code available at: `http://code.google.com/p/auditing-sshd/`

SSL Analysis

SSL Analysis

- Certificate extraction.
- Very flexible certificate validation.
- Full protocol parsing.
- Identification on any port.
- Support for SSL 2.0 - TLS 2.1
- Flag on things like
  - certificate has expired
  - self signed certificate
  - self signed certificate in certificate chain
  - unable to get local issuer certificate

Vulnerability Detection

Bro software framework

- Know what software is running on your network

- Use this to prioritize your patching and vulnerability management
  - Java
  - Flash

Drive-by-download infections

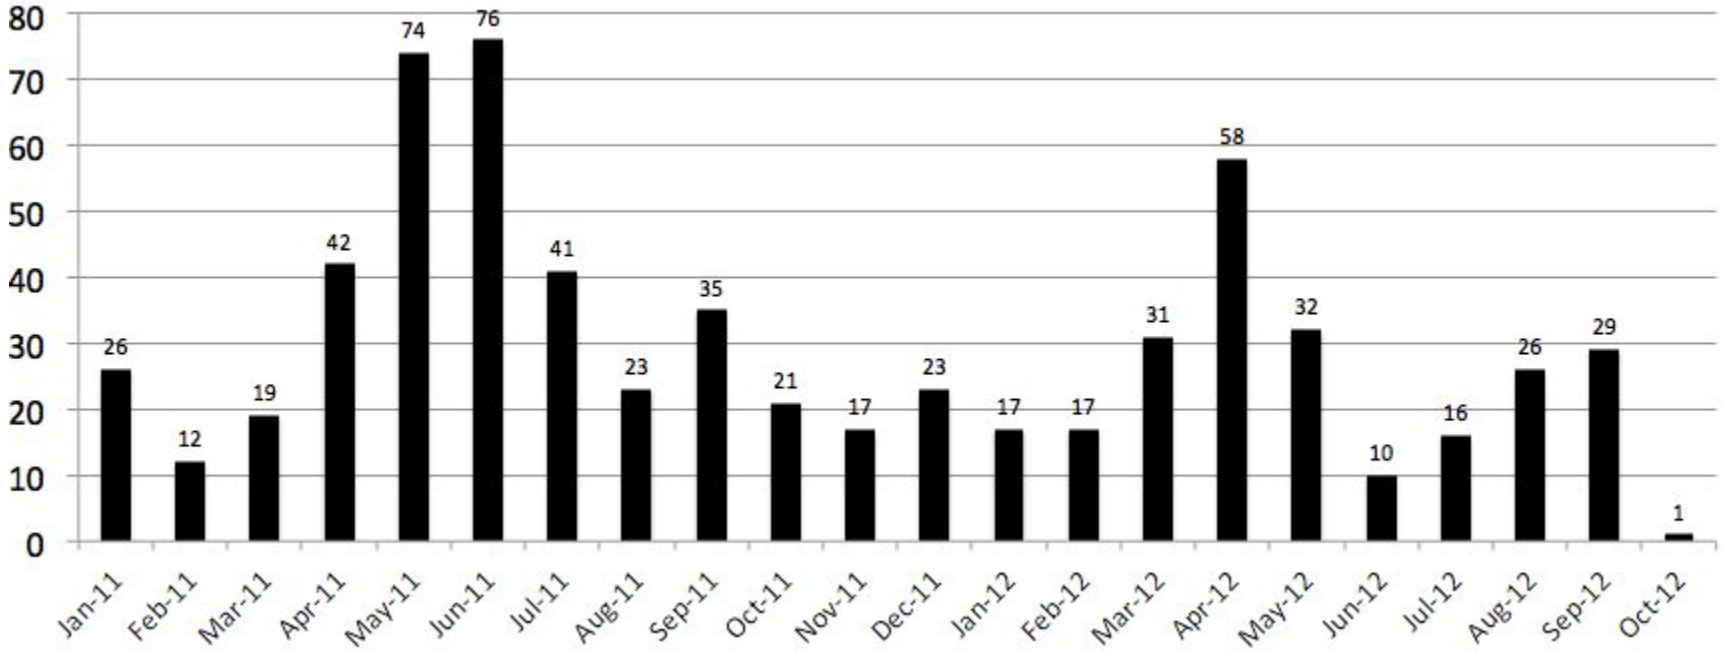

Aggressive patching and additional Influence of external factors:

- Google taking down various domains,
- New version of flash with auto-updates,
- Active blocking, RPZ of malicious domains etc

Example

Oct 23 09:29:48

Software::Vulnerable_Version vulnerable: Flash 10.3.183-18 user_agent: Mozilla/4.0 (compatible; MSIE 7.0; Windows NT 6.1; WOW64; Trident/5.0; SLCC2; .NET CLR 2.0.50727; .NET CLR 3.5.30729; .NET CLR 3.0.30729; .NET4.0C; .NET4.0E; NP06; InfoPath.2) - 128.3.182.219

Oct 23 07:41:18

Software::Vulnerable_Version vulnerable: Java 1.6.0-31 user_agent: Java/1.6.0_31

Java and Flash installations

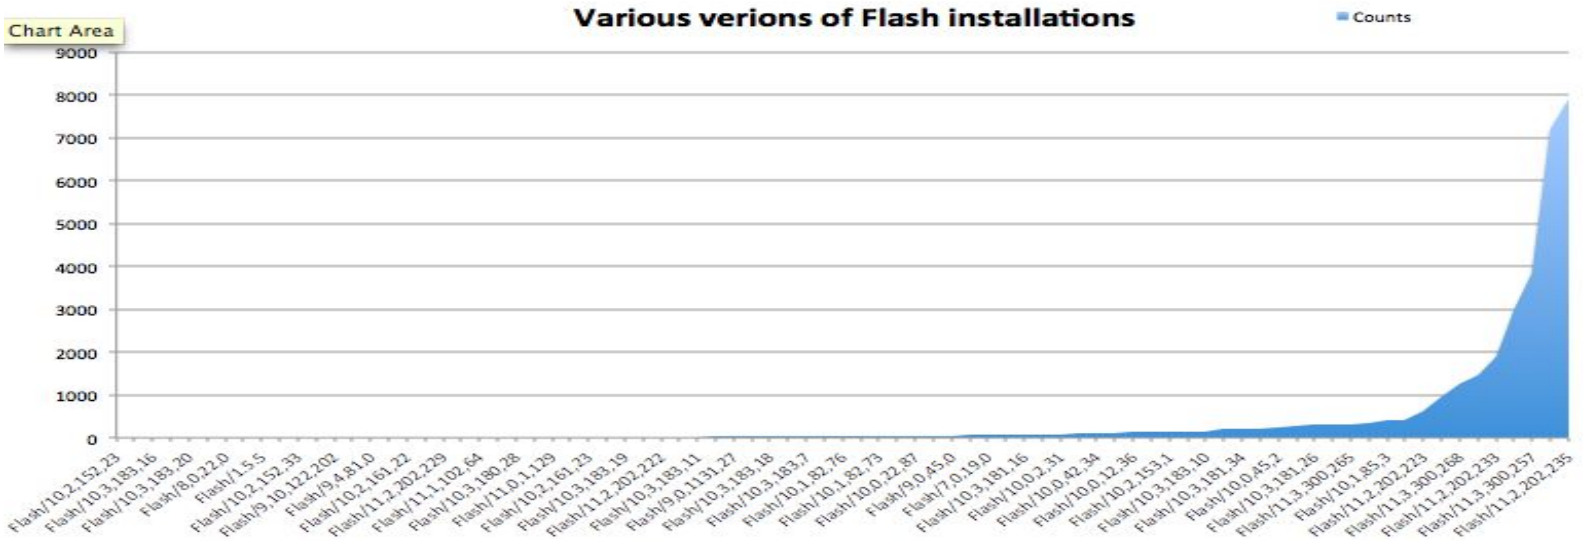

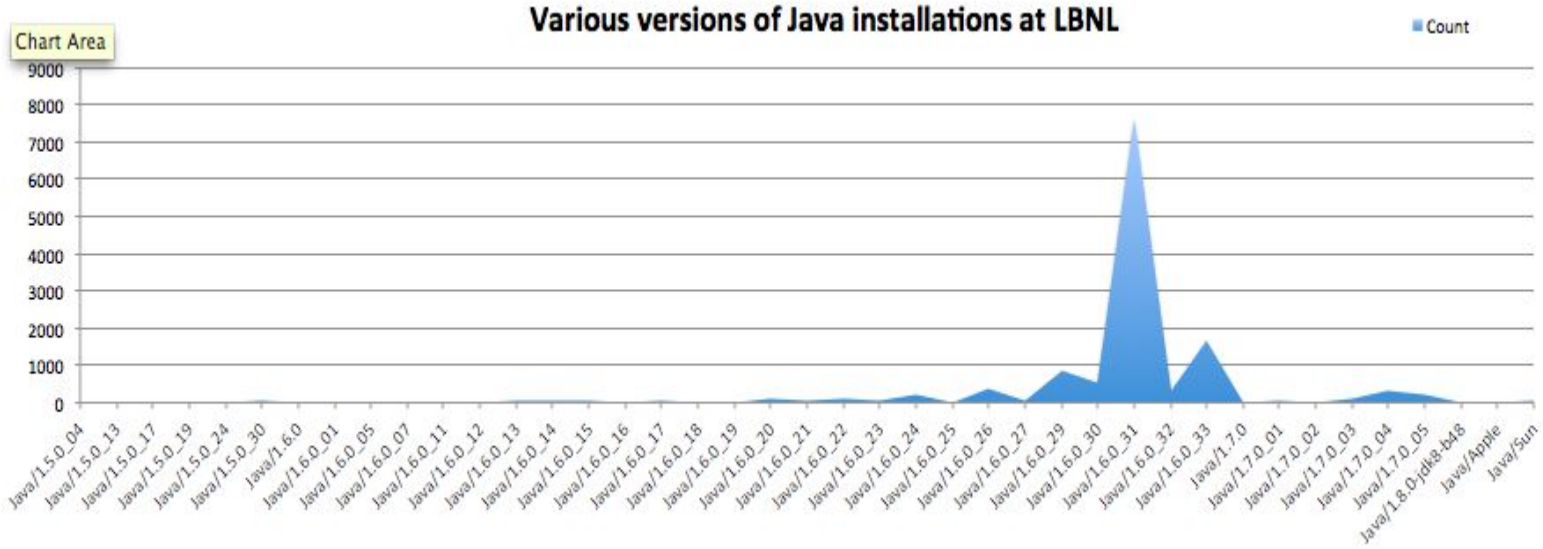

syslog2bro

- Block bruteforce SSH scanner on border
- Slightly more sophisticated attackers use a botnet where they would try one account per host/ip
- feed syslog data to bro - let it make correlations and initial drops
- Bro-2.0 has limited (udp) capability to sniff syslog on the wire - eliminating the need to feed it syslog

Got data, now what ?

7. Lynxeon - visual analysis of Bro

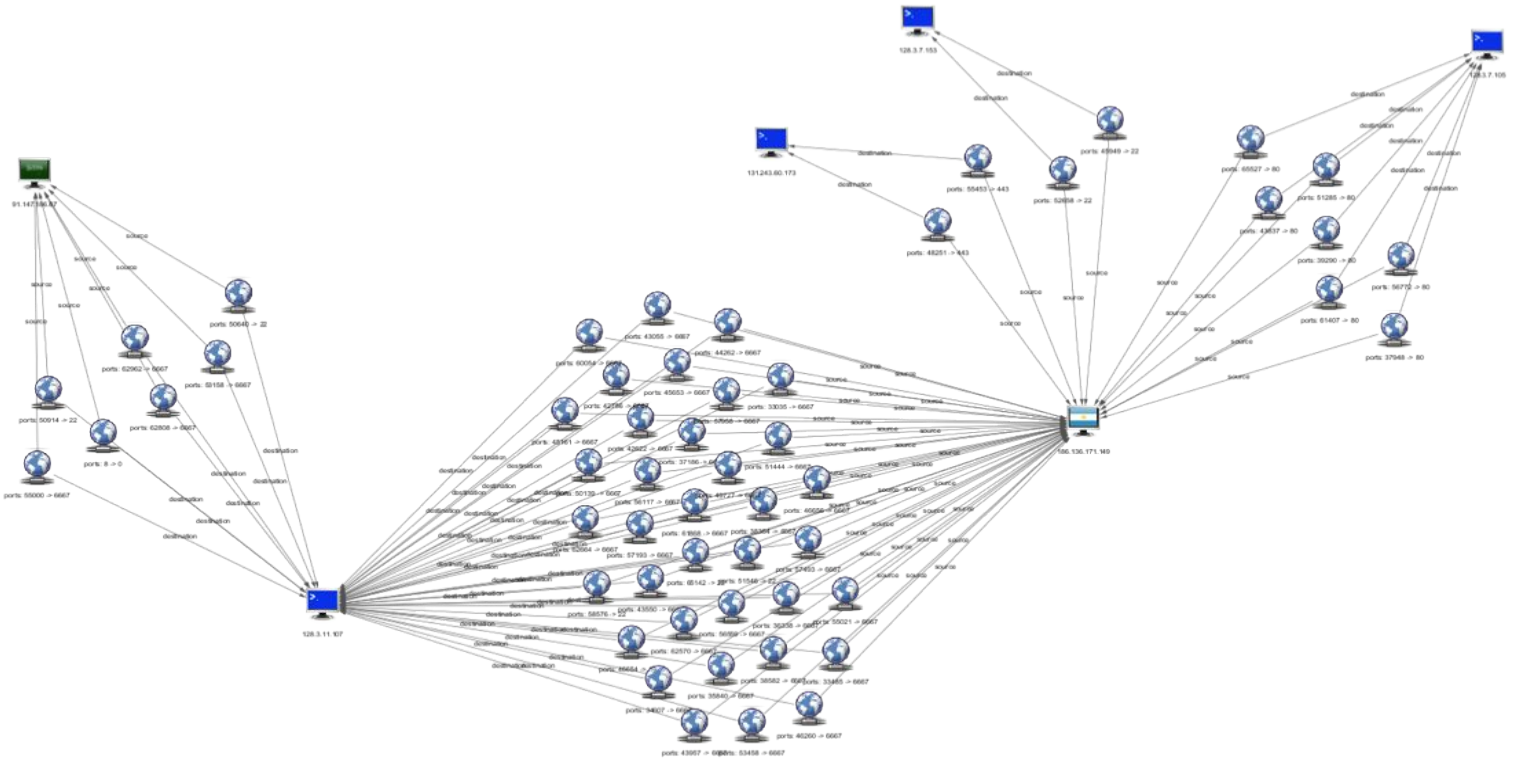

Built-in Analytics +
Timeline +
Connection graphs +
Scriptable - similar to the bro policy + -
Bro policies work on the wire, While lynxeon provides a capability to run the policy on the logs

8. Searching bro logs

Goal, search 6 months in < 10 sec

- fgrep
  - 30+ mins - 2 hours
- GNU parallel - try it +++
  - 2-10 mins
- Hadoop - Fail
- Oracle database - Fail
- Biggest fastest disk array - Fail
- SSD - Fail

Searching bro logs

- Google BigQuery
  - < 10s
  - Presently 10 billion conn logs (6 months)
- Problems
  - new columns in the table ? Reindex ?
  - Multiple-Columns, inner-join/outer-join ?
- Pricing
  - woha! I just ran a $22,143.99 query ?

100 Gb Roadmap

- Using bro cluster approach - solved problem
- Break 100Gb into multiple 10Gb feeds
- Exploring tapping infrastructure capabilities

- Various questions
  - Time machine ?
  - Identify & subsequently ignore very large "big data" transfers or subnets.

Internal Bro cluster

Q: Can we cluster Bro inside of the network

A: Yes

Q: Can we run bro on every subnet?

A: Ummm

Q: How about at-least the important subjects

A: Yes

Pilot: convert the ~4 subnet Bro to become part of a cluster

Initially watching 4 subnets but target is to expand to another 12 subnets this FY

Good defense-in-depth strategy.

Internal Bro

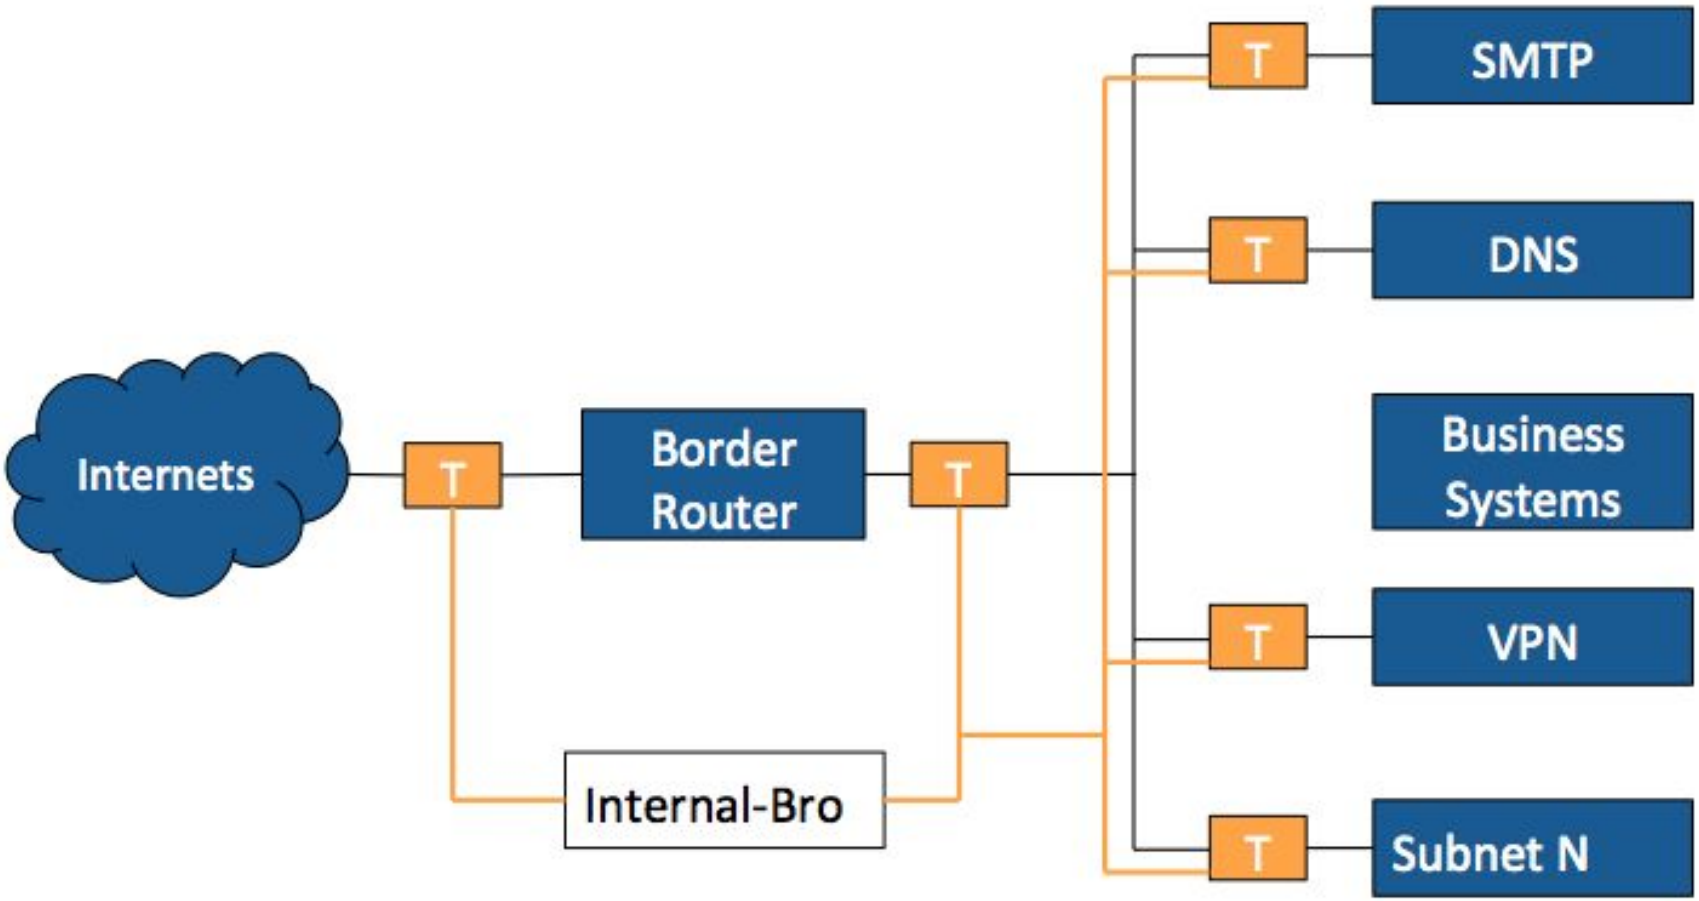

Questions