Threat Modelling and Beyond for Cisco ACI

Frank Block – fblock@ernw.de
Jan Harrie – jharrie@ernw.de

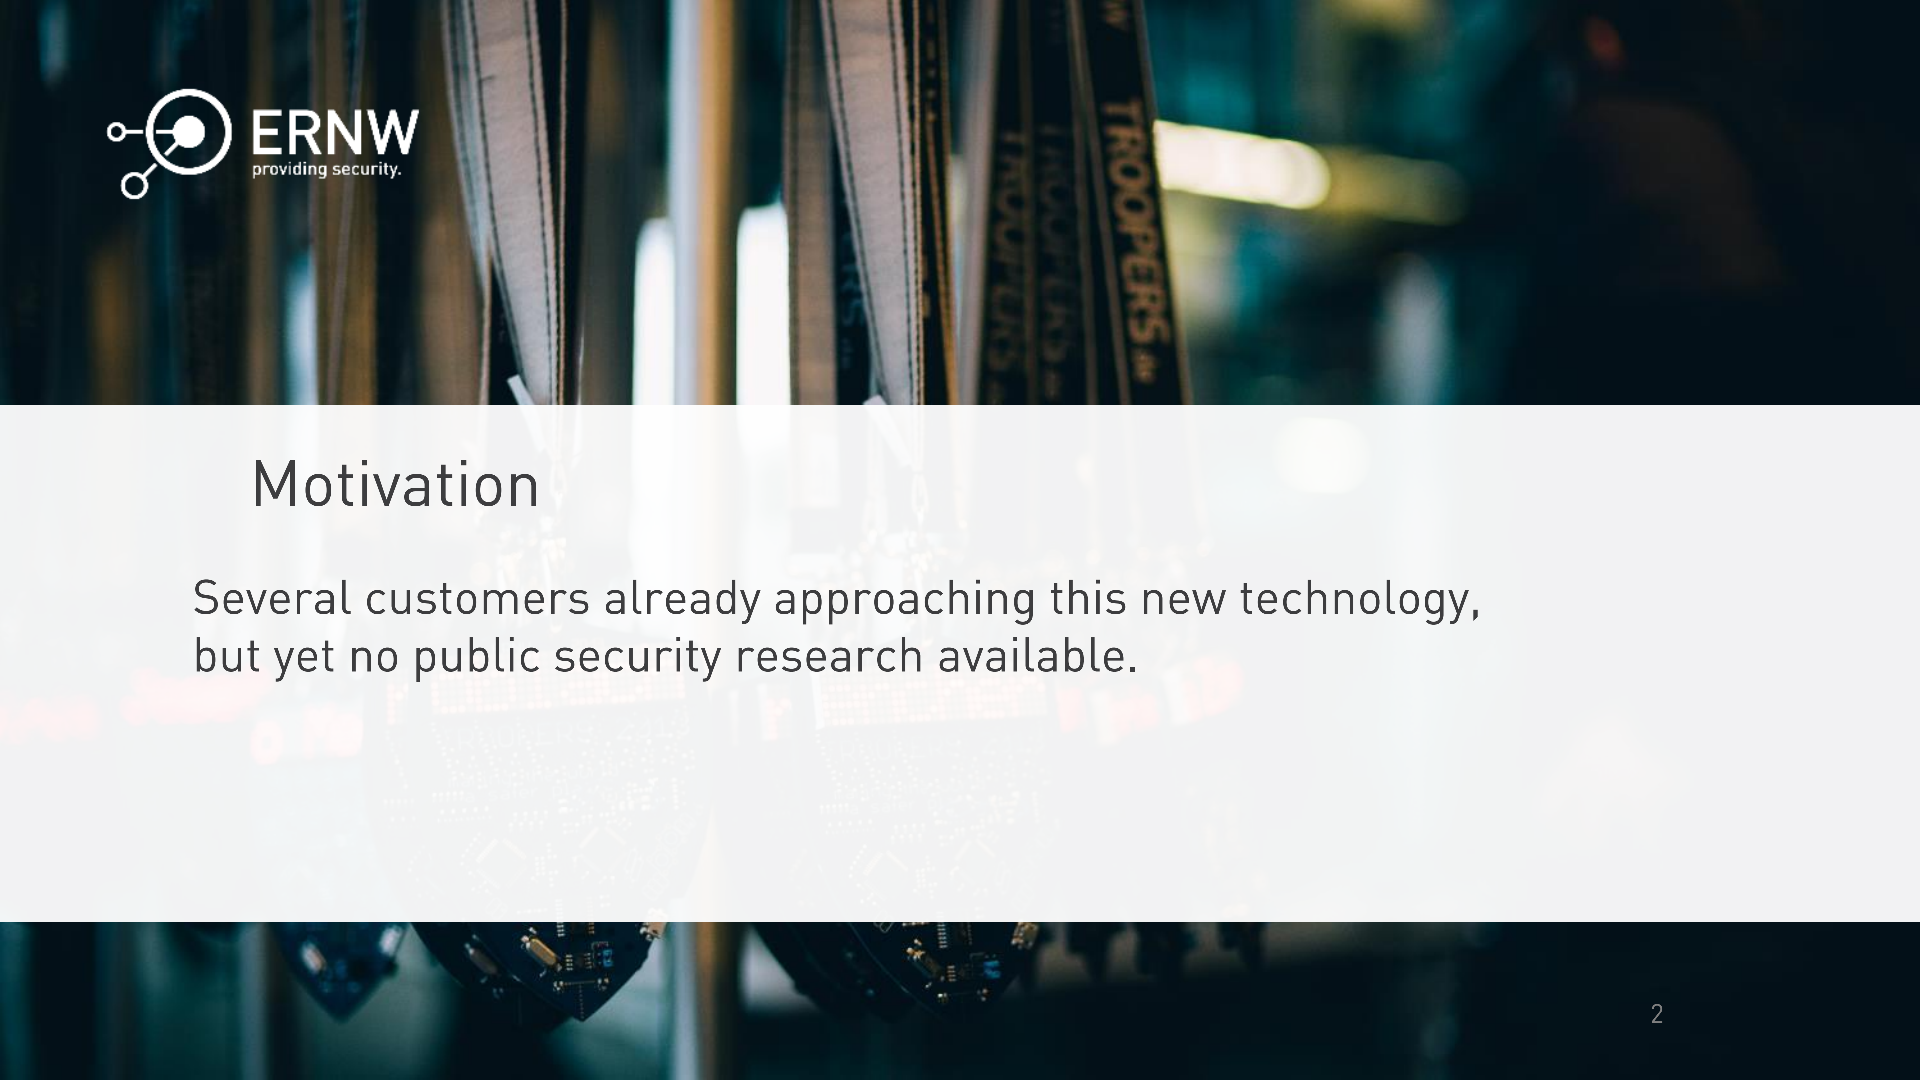

Motivation

Several customers already approaching this new technology, but yet no public security research available.

#whoami - Jan

- Security Consultant @ERNW GmbH
- Former Security Analyst/Pentester/WebApp-Monkey
- M.Sc. IT-Security TU Darmstadt

- Interests:
  - Orchestration Solution,
  - Red Teaming/Social Engineering

#whoami - Frank

- Security Researcher @ERNW Research GmbH
- Pentester/Incident Analyst

- Interests:
  - Hacking stuff and hunting Hackers
  - Memory Forensics
  - Party tonight!

ERNW
providing security.

Agenda

- ACI WTF!?
- Threat Modelling Cisco ACI
- Deep-Dive into Various Threats
- Technical Attack Surface Overview

Cisco ACI

A short introduction

Physical Layout 1/2

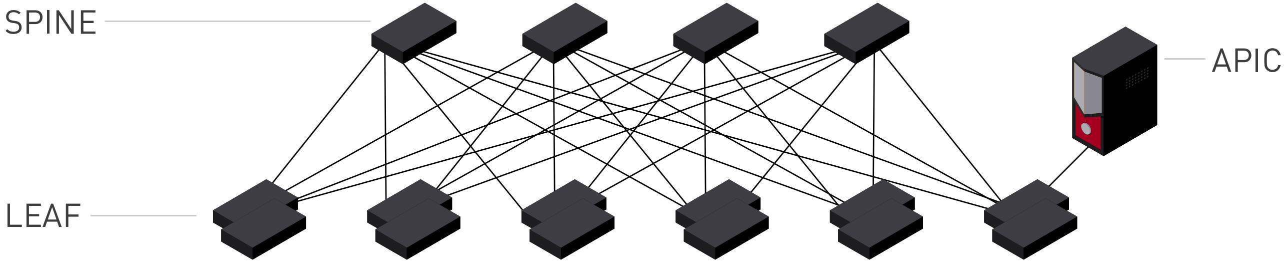

Physical Layout 2/2

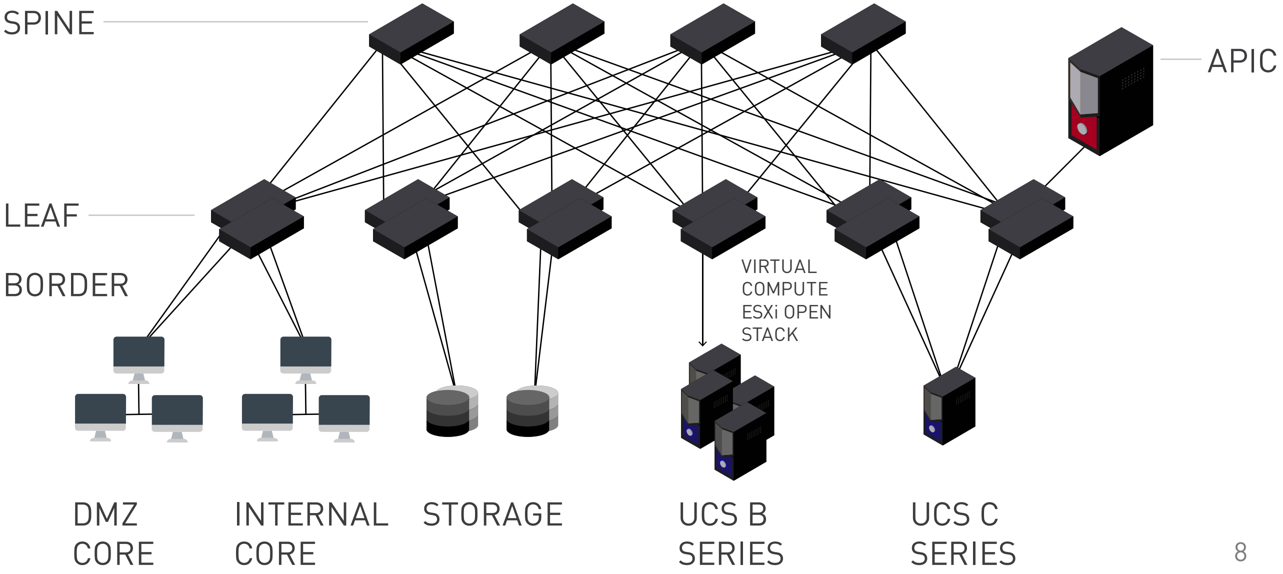

ERNW
providing security.

Background Knowledge

- Systems
  - Cisco APIC Release 4.0(3d) in Feb-19
  - Cisco NX-OS Release 14.0(3) in Feb-19
- Market Challenger to VMware NSX
- SDN Solution based on VXLAN
  - Application Centric Infrastructure
  - Micro Segmentation

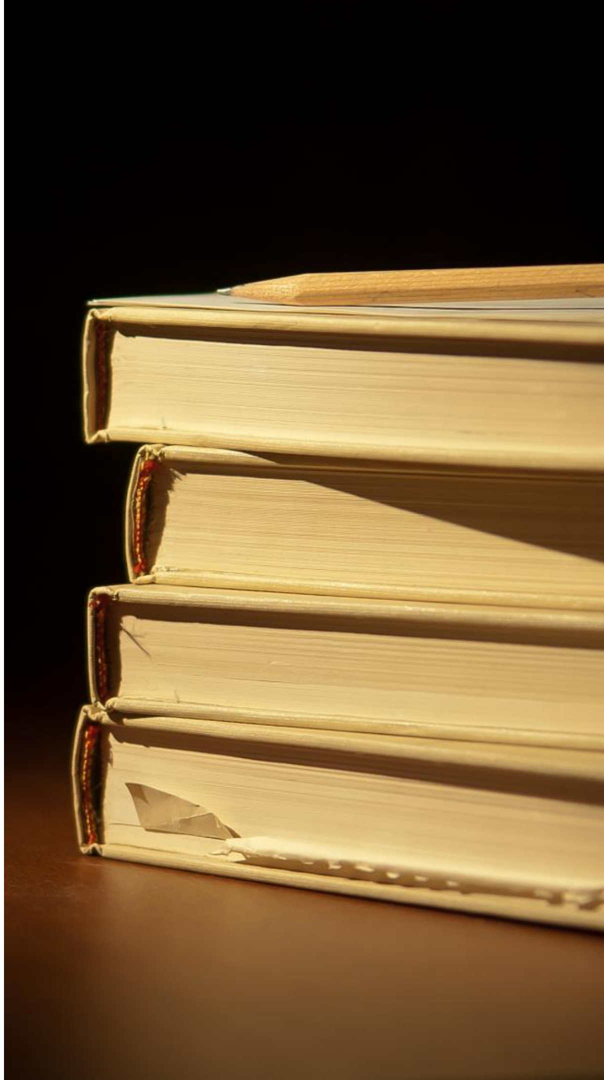

VXLAN – RFC 7348

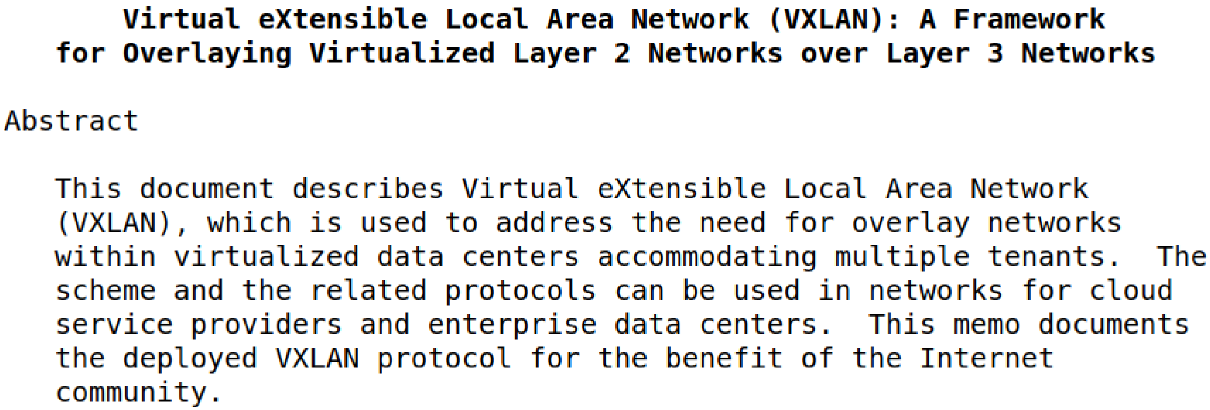

Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks

Abstract

This document describes Virtual eXtensible Local Area Network (VXLAN), which is used to address the need for overlay networks within virtualized data centers accommodating multiple tenants. The scheme and the related protocols can be used in networks for cloud service providers and enterprise data centers. This memo documents the deployed VXLAN protocol for the benefit of the Internet community.

Source: [RFC7348]

Certificates

U.S. Department of Defense – Unified Capabilities Approved Products List (DoD UC APL)

Architecture (1/2)

- Spine-Leaf Hierarchy
- Application Policy Infrastructure Controller (APIC)
- Overlay Transport Virtualization (OTV)
- Virtual Tunneling Endpoint (VTEP)
- Endpoint Groups (EPGs)
- Edge devices use IS-IS
- Fosters zero-trust model implementation
- Inter-EPGs communication whitelisted by contracts (L3/L4)

Architecture (2/2)

- Traffic whitelisting stateless on Layer 3-4, integration external applications possible (L4-7)
- Authentication, Authorization, and Accounting (AAA) over RBAC with LDAP and Microsoft Active Directory, RADIUS, and TACACS+
- APIC maintains current state of fabric and offers REST API
- ACI Virtual Machine Manager (VMM) with ACI Virtual Edge (AVE)
- VXLAN between ESXi/KVM and leaf switch

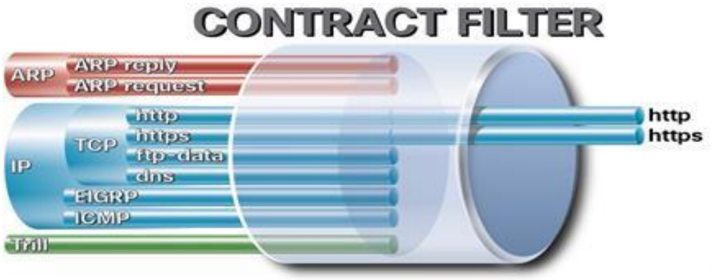

Source: [cfilter]

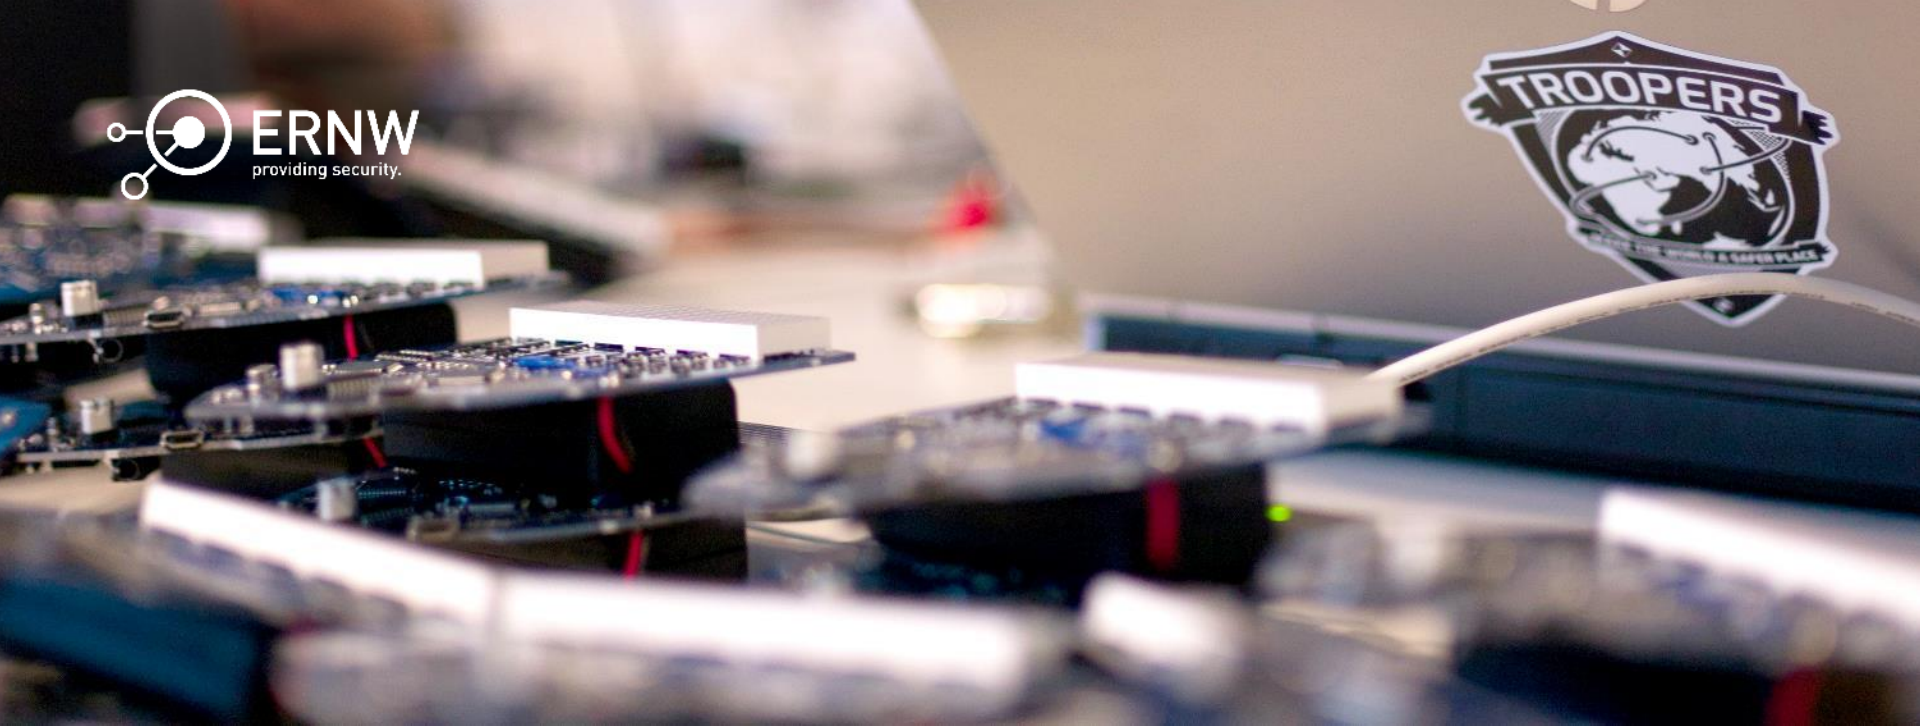

Threat Modelling

Get your hands dirty!

Short recap STRIDE

- Threat model initially developed by Microsoft in the course of their SDLC initiative (~ 2003)

But…

- STRIDE was developed in a specific context (application security) and some of elements might not be easily applicable to infrastructure projects (networks, cloud et al.)

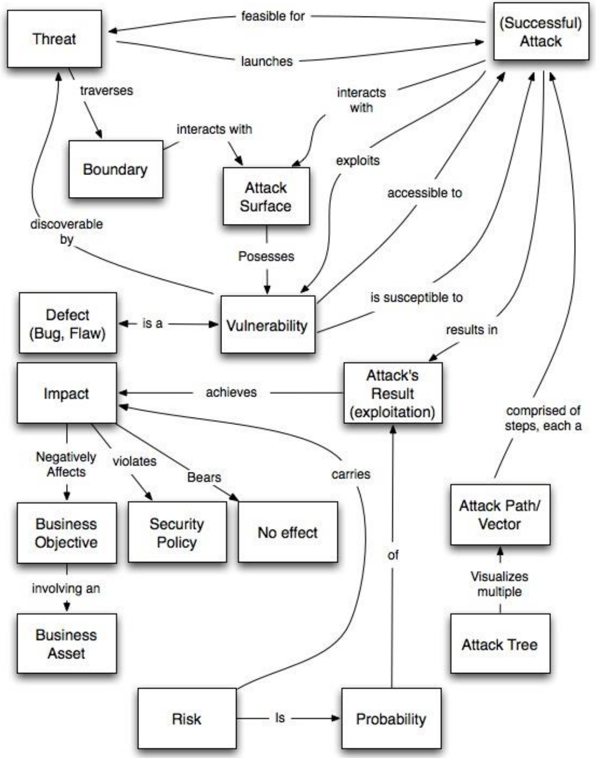

Source: [as09,ms,synops]

Short recap STRIDE

- Threat model initially developed by Microsoft in the course of their SDLC initiative (~ 2003)

But…

- STRIDE was developed in a specific context (application security) and some of elements might not be easily applicable to infrastructure projects (networks, cloud et al.)
  - Some categories might not be overly suited for network context or might have a different meaning (for example “tampering”).
  - Network security has some specific threats (e.g. “sniffing”).
  - While similar when it comes to overall direction, individual categories might differ as for risk profile, impact, attack methods etc. (*Denial-of-Service*)

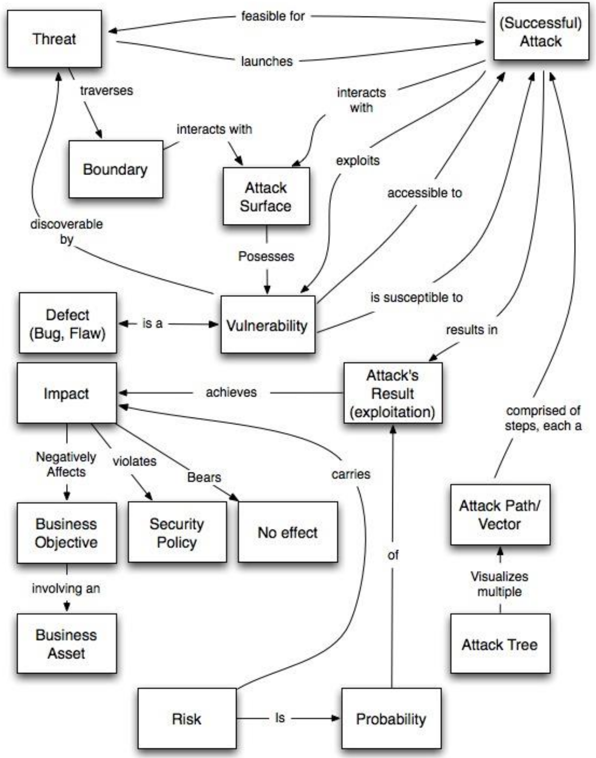

Source: [as09,ms,synops]

ERNW
providing security.

Potential Sources for Network-Related Threats / Standards

- ISO 27000 series, namely ISO 27033-3
- IETF – RFC 4778 *Current Operational Security Practices in Internet Service Provider Environments*
- NIST *Special Publications*
- Others, with more specific context
  - GSMA documents
- Book: Router Security Strategies from Cisco
- And so on …

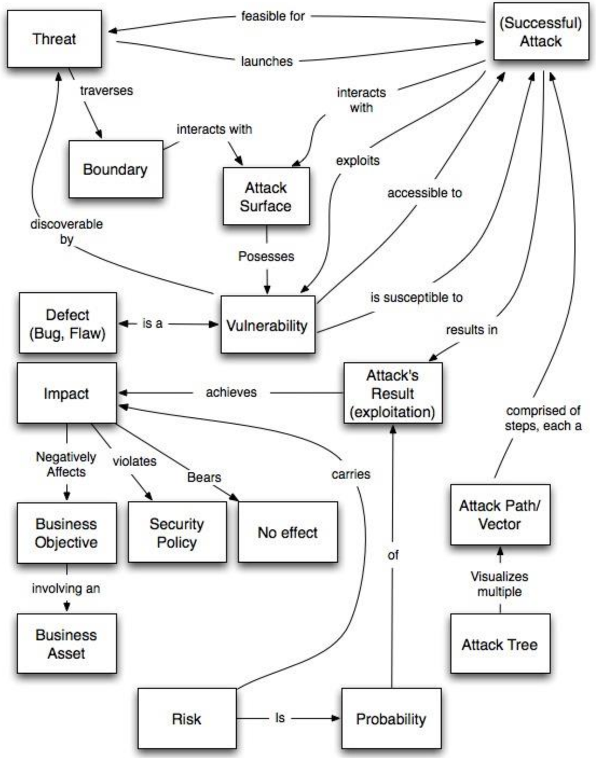

Source: [as09,ms,synops]

Suggested Adaptions of STRIDE Model, Based on the Above Sources

- Rewording of two existing categories in order to better reflect network security landscape
  - Tampering →Interception
  - Elevation of privilege →Unauthorized access
- Removal of one category
  - *Repudiation* (not really suited for network context)
- Addition of one category
  - *Sniffing* (as this one of main threats on nw level)

This Leads To...

STRIDE Category	Description	Applicable on NW infra layer	Overall Rating	Suggested Approach	Result
Spoofing	Provide false identity	x	medium	preserve	Spoofing
Tampering	Malicious modification of data	x	high	Interception	Interception
Repudiation	Make sure action was performed by certain party	-	-	remove	-
Info disclosure	(Sensitive) info is exposed	x	medium	preserve	Info disclosure
Denial of service		x	high	preserve	Denial of service
Elevation of priv.	Get unintended level of access	x	high	Unauthorized access	
				add	Sniffing

Threat Collection

1. VXLAN-Overlay Breakout
2. ACI Filter Mechanism Bypass
3. Switch Spoofing
4. VTEP Spoofing
5. APIC Spoofing
6. EPG Misconfiguration
7. Account Compromise
8. APIC Compromise
9. Switch Compromise

Threat Collection

1. VXLAN-Overlay Breakout
2. ACI Filter Mechanism Bypass
3. Switch Spoofing
4. VTEP Spoofing
5. APIC Spoofing
6. EPG Misconfiguration
7. Account Compromise
8. APIC Compromise
9. Switch Compromise

ACI Filter Mechanism Bypass

- Scenario
  - Various filter mechanisms based on …
    - VXLAN
    - EPG
    - External appliance
  - Attacker can bypass filter
- Threat
  - Unauthorized access
  - Information disclosure
- Counter Measure
  - Switch hardening (as far as possible)
  - Establishing multiple fabrics
  - Network monitoring and anomaly detection

VTEP Spoofing

- Scenario
  - Attacker spoofs VTEP and gets access to VXLAN overlay
- Threat
  - Unauthorized Access
  - Denial of Service
  - Interception
- Counter Measure
  - Switch hardening by
    - First-Hop-Security
    - Dedicated Control Plane Network
    - Data Plane Security
    - Physical isolation of Switches
  - Network monitoring and anomaly detection

Switch Compromise

- Scenario
  - Attacker compromise Spine-/Leaf-Switch and gets full access
  - Manipulation of Control Plane, e.g., Endpoint-Discovery
- Threat
  - Unauthorized Access
  - Denial of Service
  - Information Disclosure
  - Interception
- Counter Measure
  - Restricted management access
  - Classical nw segmentation for sensitive systems
  - Network monitoring and anomaly detection

Technical Attack Surface Overview

Attack Vectors on the APIC

Source: [dp40]

Management Interface

- Separate Out-of-Band Management Interface on the APIC.

- For IPv4, most TCP ports are blocked, except:
  - 22/tcp and 443/tcp

- When looking at IPv6 link-local, the firewall has/had no restrictions (CVE-2019-1690 - Fixed with Version 4.2(0.21c)):
  - 22/tcp, 443/tcp, 12569/tcp and 30865/tcp

isshd – 22/tcp

- “Special” SSH Daemon (2.9MB vs. 813KB)
- Puts connecting user in a chroot environment.

```
-xinetd-+-isshd---isshd---loginshell---scriptcontainer
```

- Special account “admin” is offered for administrative tasks.
- This account is not part of the `/etc/passwd`, but seems only be made available via a special PAM module.

isshd – 22/tcp

- The module references a hardcoded file, which contains the admin user's password hash and which is probably used for the authentication process.

- In general, the SSH service offers a stripped down local access and old school Cisco configuration.
  - `conf t` via uWSGI/HTTP (running as root)

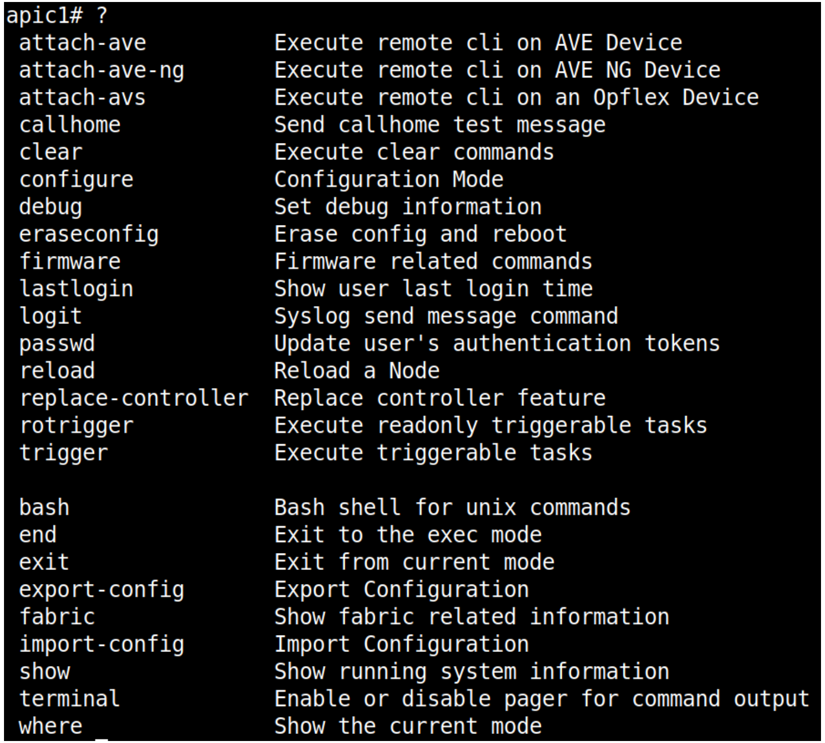

```
apic1# ?
 attach-ave           Execute remote cli on AVE Device
 attach-ave-ng        Execute remote cli on AVE NG Device
 attach-avs           Execute remote cli on an Opflex Device
 callhome             Send callhome test message
 clear                Execute clear commands
 configure            Configuration Mode
 debug                Set debug information
 eraseconfig          Erase config and reboot
 firmware             Firmware related commands
 lastlogin            Show user last login time
 logit                Syslog send message command
 passwd               Update user's authentication tokens
 reload               Reload a Node
 replace-controller   Replace controller feature
 rotrigger            Execute readonly triggerable tasks
 trigger              Execute triggerable tasks

 bash                 Bash shell for unix commands
 end                  Exit to the exec mode
 exit                 Exit from current mode
 export-config        Export Configuration
 fabric               Show fabric related information
 import-config        Import Configuration
 show                 Show running system information
 terminal             Enable or disable pager for command output
 where                Show the current mode
```

```
POST /decoy/exec/help.cli =>
generated 1211 bytes in 18 msecs
(HTTP/1.1 200)
```

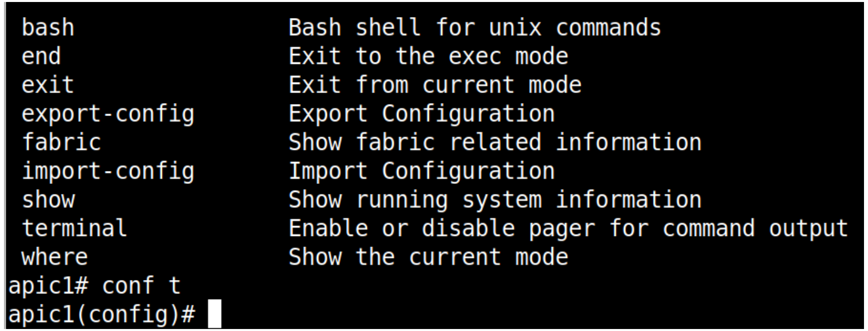

```
 bash                  Bash shell for unix commands
 end                   Exit to the exec mode
 exit                  Exit from current mode
 export-config         Export Configuration
 fabric                Show fabric related information
 import-config         Import Configuration
 show                  Show running system information
 terminal              Enable or disable pager for command output
 where                 Show the current mode
apic1# conf t
apic1(config)# 
```

```
POST /decoy/exec/cmd.cli =>
generated 0 bytes in 42
msecs (HTTP/1.1 200)
```

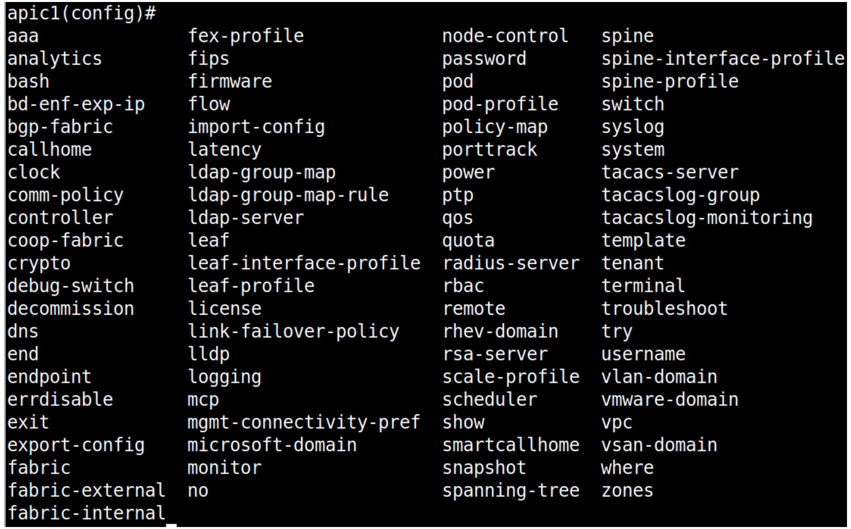

```
apic1(config)#
aaa              fex-profile             node-control   spine
analytics        fips                    password       spine-interface-profile
bash             firmware                pod            spine-profile
bd-enf-exp-ip    flow                    pod-profile    switch
bgp-fabric       import-config           policy-map     syslog
callhome         latency                 porttrack      system
clock            ldap-group-map          power          tacacs-server
comm-policy      ldap-group-map-rule     ptp            tacacslog-group
controller       ldap-server             qos            tacacslog-monitoring
coop-fabric      leaf                    quota          template
crypto           leaf-interface-profile  radius-server  tenant
debug-switch     leaf-profile            rbac           terminal
decommission     license                 remote         troubleshoot
dns              link-failover-policy    rhev-domain    try
end              lldp                    rsa-server     username
endpoint         logging                 scale-profile  vlan-domain
errdisable       mcp                     scheduler      vmware-domain
exit             mgmt-connectivity-pref  show           vpc
export-config    microsoft-domain        smartcallhome  vsan-domain
fabric           monitor                 snapshot       where
fabric-external  no                      spanning-tree  zones
fabric-internal
```

```
POST /decoy/exec/tab.cli => generated 1705 bytes in 1719 msecs (HTTP/1.1 200)
```

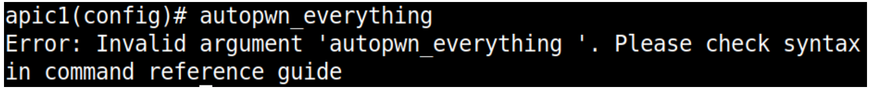

```
apic1(config)# autopwn_everything
Error: Invalid argument 'autopwn_everything '. Please check syntax
in command reference guide
```

```
Process Process-29:
Traceback (most recent call last):
  File
"/usr/lib/python2.7/multiprocessing/process.py",
line 258, in _bootstrap
    self.run()
  File
"/usr/lib/python2.7/multiprocessing/process.py",
line 114, in run
    self._target(*self._args, **self._kwargs)
  File
"/mgmt/opt/controller/decoy/apps/execserver/execapp
.py", line 75, in execCommand
    raise ex
ValueError: Error: Invalid argument
'autopwn_everything '. Please check syntax in
command reference guide

POST /decoy/exec/cmd.cli => generated 0 bytes in 46
msecs (HTTP/1.1 400)
```

Some Challenge Response Functionality

- `/data/challenge.plugin` contains a changing string.
- This path is used by the PAM module, mentioned for isshd.

- The library is also loaded by isshd itself and the nginx and might be used to allow some special local/web access.

Nginx – 443/tcp

- Serves the APIC management GUI.
- Moreover, several paths are configured that are forwarded to locally listening HTTP and uWSGI endpoints.

- Most of them, including the nginx itself, are running as root.

REST API

- The APIC implements a REST API, accessible via `/api`.
- The old school Cisco configuration via SSH and the Management GUI are both based on it.
- The GUI offers functionality to trace requests and responses being made by the GUI.

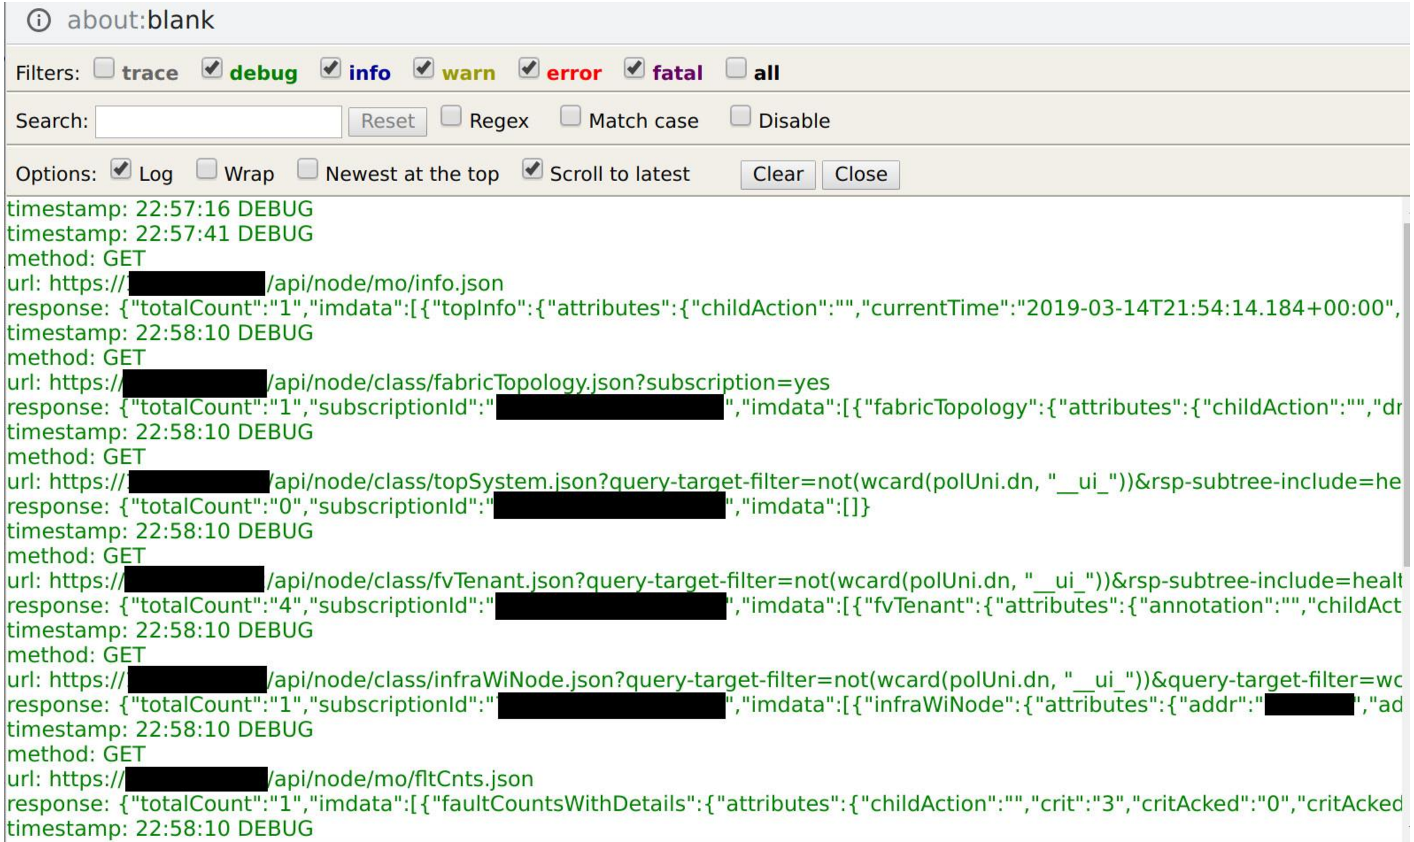

about:blank

Filters: ☐ trace ☑ debug ☑ info ☑ warn ☑ error ☑ fatal ☐ all

Search: Reset ☐ Regex ☐ Match case ☐ Disable

Options: ☑ Log ☐ Wrap ☐ Newest at the top ☑ Scroll to latest Clear Close

```
timestamp: 22:57:16 DEBUG
timestamp: 22:57:41 DEBUG
method: GET
url: https://[redacted]/api/node/mo/info.json
response: {"totalCount":"1","imdata":[{"topInfo":{"attributes":{"childAction":"","currentTime":"2019-03-14T21:54:14.184+00:00",
timestamp: 22:58:10 DEBUG
method: GET
url: https://[redacted]/api/node/class/fabricTopology.json?subscription=yes
response: {"totalCount":"1","subscriptionId":"[redacted]","imdata":[{"fabricTopology":{"attributes":{"childAction":"","dr
timestamp: 22:58:10 DEBUG
method: GET
url: https://[redacted]/api/node/class/topSystem.json?query-target-filter=not(wcard(polUni.dn, "__ui_"))&rsp-subtree-include=he
response: {"totalCount":"0","subscriptionId":"[redacted]","imdata":[]}
timestamp: 22:58:10 DEBUG
method: GET
url: https://[redacted]/api/node/class/fvTenant.json?query-target-filter=not(wcard(polUni.dn, "__ui_"))&rsp-subtree-include=healt
response: {"totalCount":"4","subscriptionId":"[redacted]","imdata":[{"fvTenant":{"attributes":{"annotation":"","childAct
timestamp: 22:58:10 DEBUG
method: GET
url: https://[redacted]/api/node/class/infraWiNode.json?query-target-filter=not(wcard(polUni.dn, "__ui_"))&query-target-filter=wc
response: {"totalCount":"1","subscriptionId":"[redacted]","imdata":[{"infraWiNode":{"attributes":{"addr":"[redacted]","ad
timestamp: 22:58:10 DEBUG
method: GET
url: https://[redacted]/api/node/mo/fltCnts.json
response: {"totalCount":"1","imdata":[{"faultCountsWithDetails":{"attributes":{"childAction":"","crit":"3","critAcked":"0","critAcked
timestamp: 22:58:10 DEBUG
```

Device Packages

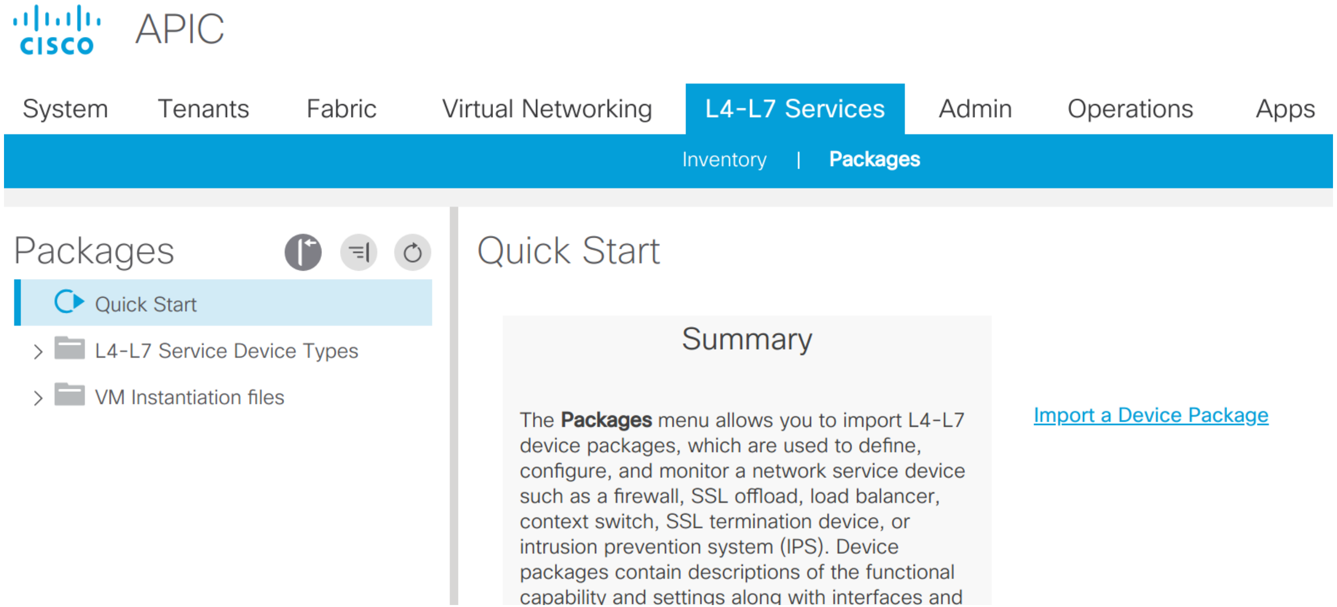

Summary

The **Packages** menu allows you to import L4-L7 device packages, which are used to define, configure, and monitor a network service device such as a firewall, SSL offload, load balancer, context switch, SSL termination device, or intrusion prevention system (IPS). Device packages contain descriptions of the functional capability and settings along with interfaces and

Device Packages

- Enables easy integration of L4-7 devices.
- A .zip file containing an XML file and Python scripts.
- No signatures/signing.

- Once uploaded, the archive is extracted and the Python script executed.
- So far, only an authenticated user can upload a new device package.

Appliance Director – 12569/tcp

- Seems like a custom service which uses TLS with client certificates.
- Yet no communication observed.
- But at least, also runs as root.

csync2 – 30865/tcp

- Open source software https://github.com/LINBIT/csync2
- Essentially rsync for multiple hosts.
- Simple protocol, which transfers the password as-is for authentication.

- The service is configured with a long password.

csync2 Protocol

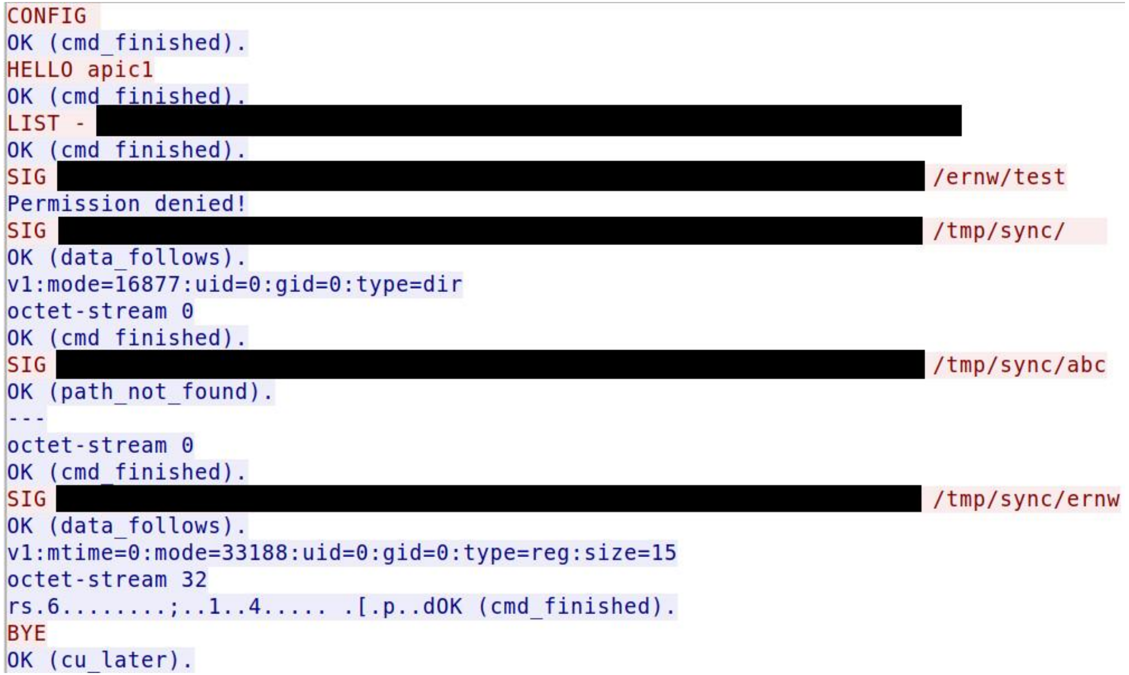

```
CONFIG
OK (cmd_finished).
HELLO apic1
OK (cmd_finished).
LIST - [redacted]
OK (cmd_finished).
SIG [redacted] /ernw/test
Permission denied!
SIG [redacted] /tmp/sync/
OK (data_follows).
v1:mode=16877:uid=0:gid=0:type=dir
octet-stream 0
OK (cmd_finished).
SIG [redacted] /tmp/sync/abc
OK (path_not_found).
---
octet-stream 0
OK (cmd_finished).
SIG [redacted] /tmp/sync/ernw
OK (data_follows).
v1:mtime=0:mode=33188:uid=0:gid=0:type=reg:size=15
octet-stream 32
rs.6........;..1..4..... .[.p..dOK (cmd_finished).
BYE
OK (cu_later).
```

Technical Attack Surface Overview

A quick look to the Leaf Switches

Source: [dp40]

OpFlex Control Protocol

- In order to be able to push policies (basically ACI configuration) to Leaf Switches, the OpFlex Control Protocol is used.

- The protocol is based on JSON and supports several RPC methods (JSON-RPC version 1.0).

- There is an IETF Draft from April 2016 which, based on first comparisons, seems to be conform with the actual implementation.

OpFlex Control Protocol

```
{“id":["echo",40],"method":"echo","params":[12345678]}.
                                   {"id":["echo",40],"result":[12345678]}.

{"id":["send_identity",1],"method":"send_identity", ...}
               {"id":["send_identity",1],"result":{"name":"10.0.0.1:8009",
                   "my_role":["endpoint_registry","policy_repository"], ...
                   "peers":[{"role":[...],"connectivity_info":"10.0.0.1:8009"},
                          {"role":[...],"connectivity_info":"10.0.0.2:8009"}]}}

{... "method": "policy_update", "params": [... }
```

IETF Draft – Security Considerations

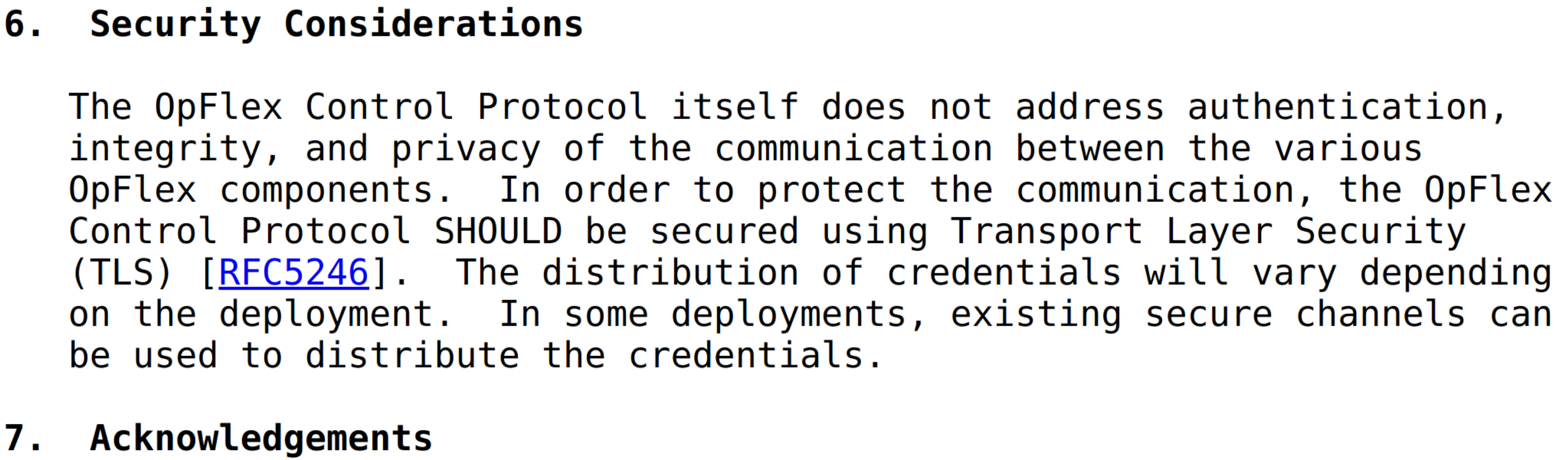

6. Security Considerations

The OpFlex Control Protocol itself does not address authentication, integrity, and privacy of the communication between the various OpFlex components. In order to protect the communication, the OpFlex Control Protocol SHOULD be secured using Transport Layer Security (TLS) [RFC5246]. The distribution of credentials will vary depending on the deployment. In some deployments, existing secure channels can be used to distribute the credentials.

7. Acknowledgements

OpFlex Service

- As far as we have seen, services speaking OpFlex are using TLS.
- The service is accessible in the management network.
- And again, runs as root.

Next Steps

- Attacking/Fuzzing the Protocols.
- Having a closer look at the challenge response functionality!
- Getting our hands on some client certificates (for the Appliance Director).
- Investigating at least the local services for the nginx.
- Assessment of the REST API.
- …
- And, getting remote root might be nice ;-)

Security Considerations

- Restrict Access to the management interface.
- Network monitoring and anomaly detection.
- Watch out for new Updates.
- Do not import Device Packages from Spam/4chan/stackoverflow !

ERNW
providing security.

Thanks for your Attention!

Open Questions?

✉ fblock@ernw.de
jharrie@ernw.de

@WEareTROOPERS
@NodyTweet

www.ernw.de

www.insinuator.net

Sources

[RFC7348] https://tools.ietf.org/html/rfc7348

[as09] https://cloudblogs.microsoft.com/microsoftsecure/2009/08/27/the-threats-to-our-products/

[ms] https://docs.microsoft.com/en-us/previous-versions/commerce-server/ee823878(v=

cs.20)

[synops] https://www.synopsys.com/blogs/software-security/wp-content/uploads/2015/08/threat-modeling-glossary-diagram.jpg

[dp40] https://docplayer.net/docs-images/40/21587129/images/15-0.jpg

[cfilter]
https://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/aci/apic/white_papers/Cisco_IT_ACI_Design.docx/_jcr_content/renditions/Cisco_IT_ACI_Design_13.jpg