



Cisco *live!*

January 29 - February 2, 2018 · Barcelona

BRKSEC-3303

Cisco Web Security Appliance - Best Practices and Performance Troubleshooting

Ana Peric, M.Sc.E.E, Technical Leader Services

Cisco Spark




Questions?

Use Cisco Spark to communicate with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install Spark or go directly to the space
4. Enter messages/questions in the space



cs.co/ciscolivebot#BRKSEC-3303



Abstract

This session will describe the best practices of deploying and configuring Cisco Web Security Appliance (WSA), with the special attention to WSA configuration optimization in order to achieve the optimal system performance.

Session will describe WSA best configuration practices in the first section, but also the most common “pitfalls” when it comes to configuring Web Security Appliance.

Based on experience of Cisco TAC Engineers, we will continue with a deep-dive of troubleshooting WSA performance, that will give Web Security System Administrators more insights into tools and techniques of troubleshooting the most common performance related issues.

Agenda

- Introduction
- Understanding Cisco Web Security Appliance Pipeline
- Configuration Considerations and Best Practices
- Troubleshooting WSA Performance Issues
- Performance Monitoring & Final Thoughts
- Q & A

Introduction – About Me

Ana Perić

- Joined Cisco in 2012
- Based in Munich, Germany
- Technical Leader Services in Cloud Support Organization (aka Cloud TAC)
- M.Sc.E.E (Diploma Engineer of Electrical Engineering and Computer Science), CCIE #39884 R&S
- Passionate about Web/Email Security, Cloud Technologies, Automation, and Innovation
- Proud aunt of three-year-old boy



For your reference symbol

- There is a content in your handouts that is not going to be presented in this session, but is important for further reference
- All the slides that are there for your reference are marked with:



Web Security Appliance Pipeline

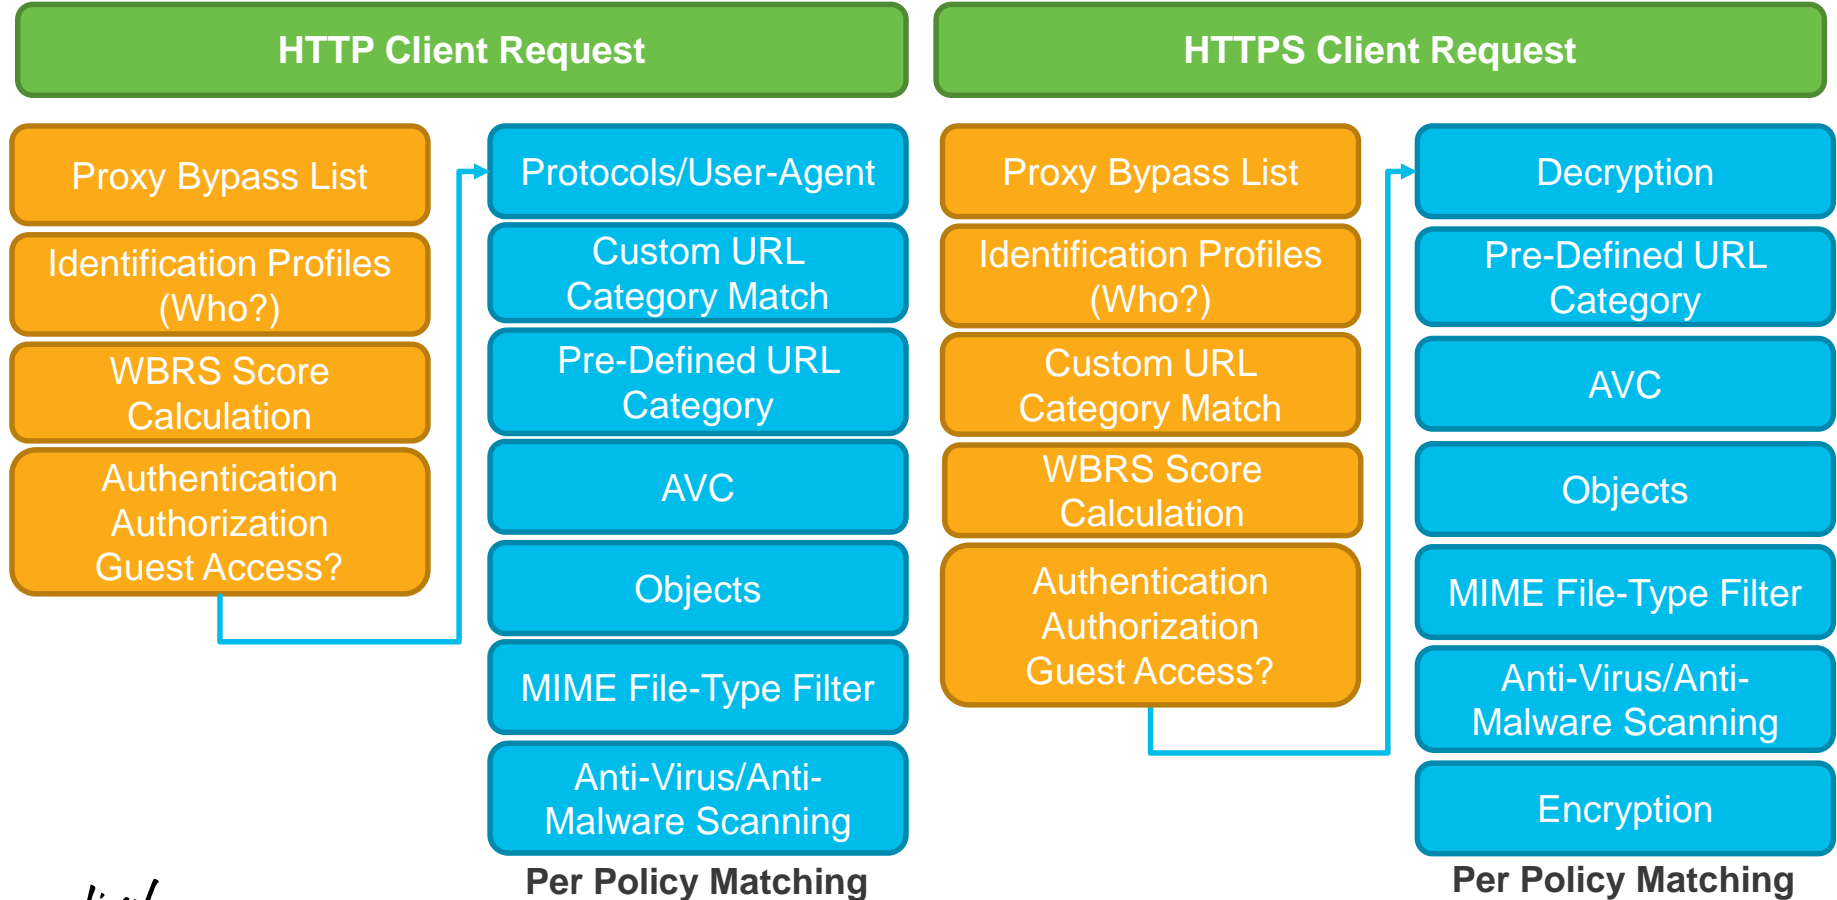
WSA Policy Types

Refresher



- **Identification Policy** (Who? / How? / How do we recognize/categorize the end-user?)
- **Access Policy** (Actions for HTTP / HTTPS decrypted traffic)
- **Decryption Policy** (HTTPS traffic handling / what do we decrypt?)
- **Routing Policy** (Upstream Proxy Handling)
- **Outbound Malware Policy** (Do we permit upload of Malware content)
- **Data Security Policy** (What content type can we upload)
- **Other Policy Types:** SaaS/SOCKS Policies/WTT

Web Security Appliance Pipeline for HTTP/HTTPS



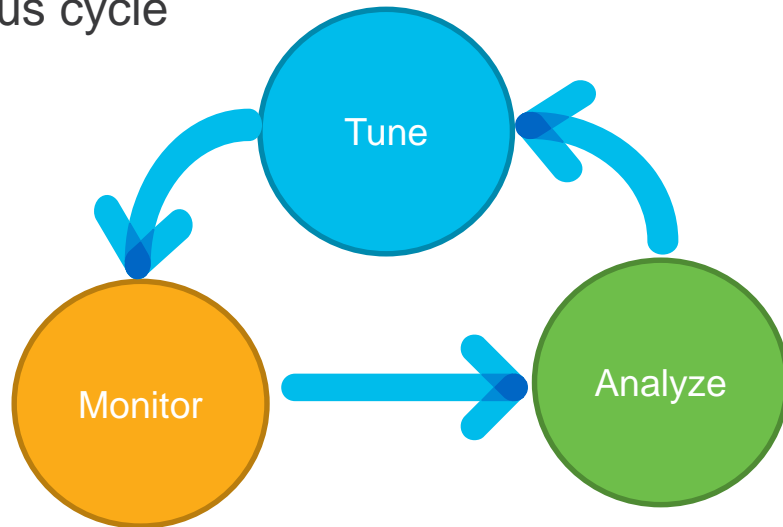
WSA Configuration Considerations & Best Practices

WSA Configuration Considerations & Best Practices

- Access & Decryption Policy Best Practices
- Custom URL Categories
- Network Configuration and Tuning
- Logging

Disclaimer / Before we start

- “One configuration fits all” / ”One Optimization Technique fits all” – is utopia
- “Best Practices” are just general guidelines that should positively impact WSA operation, and ease-up maintenance
- Part of a continuous cycle



WSA Policy Optimization Recommendations

Less is More – Decrease the policy count and unneeded repetition

Order and Collapse - Use one Policy for multiple Identities

Collapse more Policies into one (use logical OR operator)

Position Authentication Exemption Policies to the TOP & Avoid per-user matching

*Authentication exemption by User-Agent is not possible in transparent environment

Position the most frequently used policies as close to the TOP as possible

Decrypt only what is really needed & DROP the same Categories AP would BLOCK

Avoiding unnecessary decryption that is CPU intensive

Try to avoid using “All Identities” – be more specific when possible

WSA Policy Optimization Recommendations

Less is More – Decrease the policy count and unneeded repetition

Order and Collapse - Use one Policy for multiple Identities

Collapse more Policies into one (use logical OR operator)

Position Authentication Exemption Policies to the TOP & Avoid per-user matching

*Authentication exemption by User-Agent is not possible in transparent environment

Position the most frequently used policies as close to the TOP as possible

Decrypt only what is really needed & DROP the same Categories AP would BLOCK

Avoiding unnecessary decryption that is CPU intensive

Try to avoid using “All Identities” – be more specific when possible

WSA Policy Configuration (1)

Identification/Access/Decryption Policies – **Order & Collapse**



- Policy matching is always top-to-bottom
- Identification/Access/Decryption Policies that are **exemption of authentication** should always be **on the top**
- Policies with high probability of being used should be placed **closer to the top**
- Try to keep Access/Decryption policies that match "**All Identities**" to minimum
- Use "logical OR" – match multiple Identities in single access policy – collapsing

“Always design Web Security Policies as if the person who ends up maintaining them is a violent psychopath who knows where you live”

Ana Peric inspired by Unknown artists @9GAG, Feb 2017

WSA Policy Configuration (2)

Identification/Access/Decryption Policies – Policy Count, counts



- Policy processing is very CPU intensive:
=> Less is more!
- Limit the number of Identification Policies
- Limit the number of Access/Decryption Policies (10-15 is most likely quite enough)
- Avoid per-user policy matching – rather match entire LDAP/AD group

Access Policies

| Order | Policy Name | Protocols and User Agents | URL Filtering | Applications | Objects | Anti-Malware and Reputation | Delete |
|-------|---|---------------------------|-----------------|-----------------|------------------|---|--------|
| 1 | ap.Policy20 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 2 | ap.Policy19 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 3 | ap.Policy18 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 4 | ap.Policy17 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 5 | ap.Policy16 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 6 | ap.Policy15 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 7 | ap.Policy14 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 8 | ap.Policy13 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 9 | ap.Policy12 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 10 | ap.Policy11 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 11 | ap.Policy10 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 12 | ap.Policy9 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 13 | ap.Policy8 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 14 | ap.Policy7 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 15 | ap.Policy6 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 16 | ap.Policy5 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 17 | ap.Policy4 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 18 | ap.Policy3 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 19 | ap.Policy2 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| 20 | ap.Policy1 Identification Profile: Global All identified users | (global policy) | (global policy) | (global policy) | (global policy) | (global policy) | * |
| | Global Policy | No blocked items | Not Available | Not Available | No blocked items | Web Reputation: Unavailable Advanced Malware Protection: Unavailable Webroot: Unavailable | |



HTTPS Decryption – Refresher (1)

- WSA can decrypt HTTPS traffic by acting like “Man-in-the-Middle”
- For security reasons, try to use SSL key size that is **> 1024 bits**
- HTTPS decryption is controlled in:
 - Global Setting: HTTPS Proxy Configuration (**Security Services -> HTTPS Proxy**)

HTTPS Proxy

| HTTPS Proxy Settings | |
|--|--|
| HTTPS Proxy: | Enabled |
| HTTPS Ports to Proxy: | 443 |
| Root Certificate and Key for Signing: | Using Generated Certificate: Common name: heka.ironport.local Organization: CC Organizational Unit: IronPort Country: DE Expiration Date: Nov 19 21:20:36 2020 GMT Basic Constraints: Not Critical |
| Decryption Options | |
| Decrypt for Authentication: | Enabled |
| Decrypt for End-User Notification: | Enabled |
| Decrypt for End-User Acknowledgement: | Enabled |
| Decrypt for Application Detection: | Disabled |
| Invalid Certificate Options | |
| Invalid Certificate Handling: | Expired: Monitor Mismatched Hostname: Monitor Unrecognized Root Authority / Issuer: Drop Invalid Signing Certificate: Drop Invalid Leaf Certificate: Drop All other error types: Drop |
| Online Certificate Status Protocol Options | |
| OCSPP Result Handling: | Revoked Certificate: Drop Unknown Certificate: Monitor OCSPP Error: Monitor |



HTTPS Decryption – Refresher (1)

- WSA can decrypt HTTPS traffic by acting like “Man-in-the-Middle”
- For security reasons, try to use SSL key size that is **> 1024 bits**
- HTTPS decryption is controlled in:
 - Global Setting: HTTPS Proxy Configuration (**Security Services -> HTTPS Proxy**)
 - Per-Policy configuration: **Web Security Manager -> Decryption Policies**

Decryption Policies

| Policies | | | | | |
|---|---|--|-----------------|-----------------|--------|
| Add Policy... | | | | | |
| Order | Group | URL Filtering | Web Reputation | Default Action | Delete |
| 1 | dp.subnet10 Identification Profile: id.subnet10 All identified users | Pass Through: 1 Monitor: 61 Decrypt: 4 Drop: 13 | (global policy) | (global policy) | |
| | Global Policy Identification Profile: All | Monitor: 79 | Enabled | Decrypt | |
| | | Country: DE Expiration Date: Nov 19 21:20:36 2020 GMT Basic Constraints: Not Critical | | | |
| Decryption Options | | | | | |
| | Decrypt for Authentication: | Enabled | | | |
| | Decrypt for End-User Notification: | Enabled | | | |
| | Decrypt for End-User Acknowledgement: | Enabled | | | |
| | Decrypt for Application Detection: | Disabled | | | |
| Invalid Certificate Options | | | | | |
| | Invalid Certificate Handling: | Expired: Monitor Mismatched Hostname: Monitor Unrecognized Root Authority / Issuer: Drop Invalid Signing Certificate: Drop Invalid Leaf Certificate: Drop All other error types: Drop | | | |
| Online Certificate Status Protocol Options | | | | | |
| | OCSLP Result Handling: | Revoked Certificate: Drop Unknown Certificate: Monitor OCSLP Error: Monitor | | | |

HTTPS Decryption – Refresher (2)



Decryption Policy Actions

- **Drop** – traffic is dropped / HTTPS connection Terminated
Note: In transparent deployment NO End-User Notification is displayed for dropped connection
- **Decrypt:** Traffic is decrypted, and further matching access policy is used to determine further behavior
- **Pass-Through:** HTTPS connection will not be intercepted – end-user communicates with destination HTTPS server directly **w/o** additional scanning
- **Monitor** – this is NOT an final action, means that we only continue further checks down the pipeline



Decryption Policy Considerations – in more details

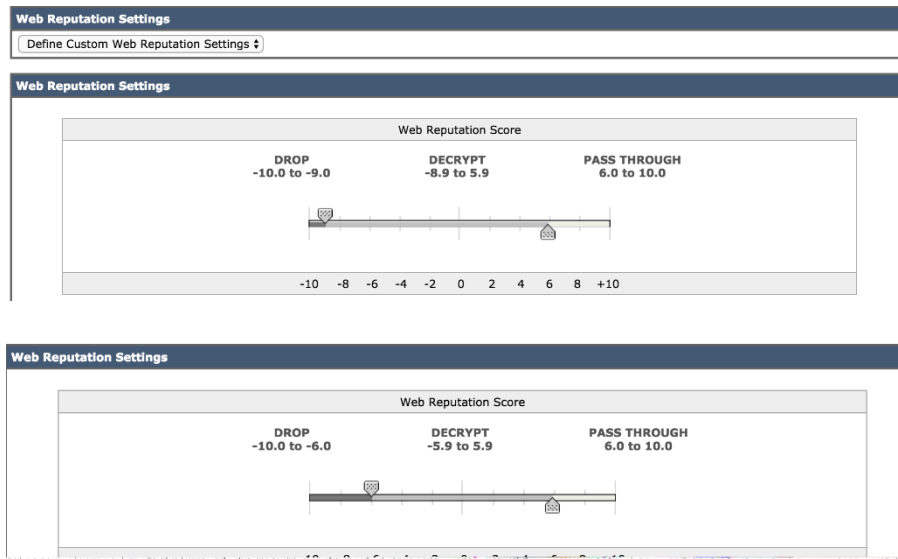
- Decrypt only traffic needed by Company Security Policy & Rely on **WBRs!**
- **What do I need to decrypt, and what not?**
 - **Decrypt** only categories that would need further fine-grained control / access policy processing / Referrer Exemption & AV/AM scanning
 - **Decrypt** for:
 - Authentication
 - End-User-Notification display
 - End-User-Acknowledgements display
 - **Pass-through** traffic that might be confidential (i.e Financial / Banking sites)
 - **Drop** the traffic that would have action **Block** by the corresponding Access Policy
 - **Drop** Categories matching: Illegal, forbidden, and business inappropriate content



HTTPS & WBRS – How do dots connect?

- Unless explicitly specified by Custom or Pre-defined URL category, action will be determined by WBRS score (if WBRS is enabled)
- One can choose default (**less aggressive**)
- Custom WBRS decryption - **more aggressive** values setup

Decryption Policies: Reputation Settings: Global Policy



Custom URL Categories (1)

Regular Expressions – sure, go ahead, but only if you really have to

- Avoid large amount of Custom URL Categories
- Avoid using regular expressions (especially "match any" regex)
 - Try not to use "match any" **.*** - **be more specific**
 - Whenever possible use rather **Sites** field than **"Advanced->Regular Expressions"**

Custom and External URL Categories: Edit Category

| Edit Custom and External URL Category | |
|---------------------------------------|---|
| Category Name: | <input type="text" value="cat.whitelist"/> |
| List Order: | <input type="text" value="1"/> |
| Category Type: | Local Custom Category |
| Sites: ? | <div><input type="text" value=".domain.com, domain.com"/> <small>(e.g. 10.0.0.1, 2001:420:80:1::5, example.com.)</small></div> <div>Sort URLs <small>Click the Sort URLs button to sort all site URLs in Alpha-numerical order.</small></div> |
| ▼ Advanced | <div>Regular Expressions: ? <input type="text" value="[a-z]+\,domain.com.*zip"/> <small>Enter one regular expression per line.</small></div> |

Custom URL Categories (2)

Cisco External & 3rd Party Feeds (Microsoft Office Format)

- Cisco External & 3rd Party Feeds are new way of automatically obtain custom URL category list from external server, using HTTP/HTTPS protocol
- The same recommendation apply for creating Cisco External Feeds:
 - Use feed entry type “site” as much as possible vs Using “regex” entries
 - Try avoiding having more than **1000** lines in per External Feed File





Example of external feed file

```
www.facebook.com,site
```

```
\.xyz,regex
```

```
\.somedomain.com/ana/*.d,regex
```

```
ad2.*\.com,regex
```

```
cisco.com,site
```

```
2000:1:1:11:1:1::200,site
```

```
www.cisco.com/.*/licensing,regex
```

```
([a-z]|[0-9])+\.cisco.com/.*/licensing,regex
```

```
www.duke.edu,site
```

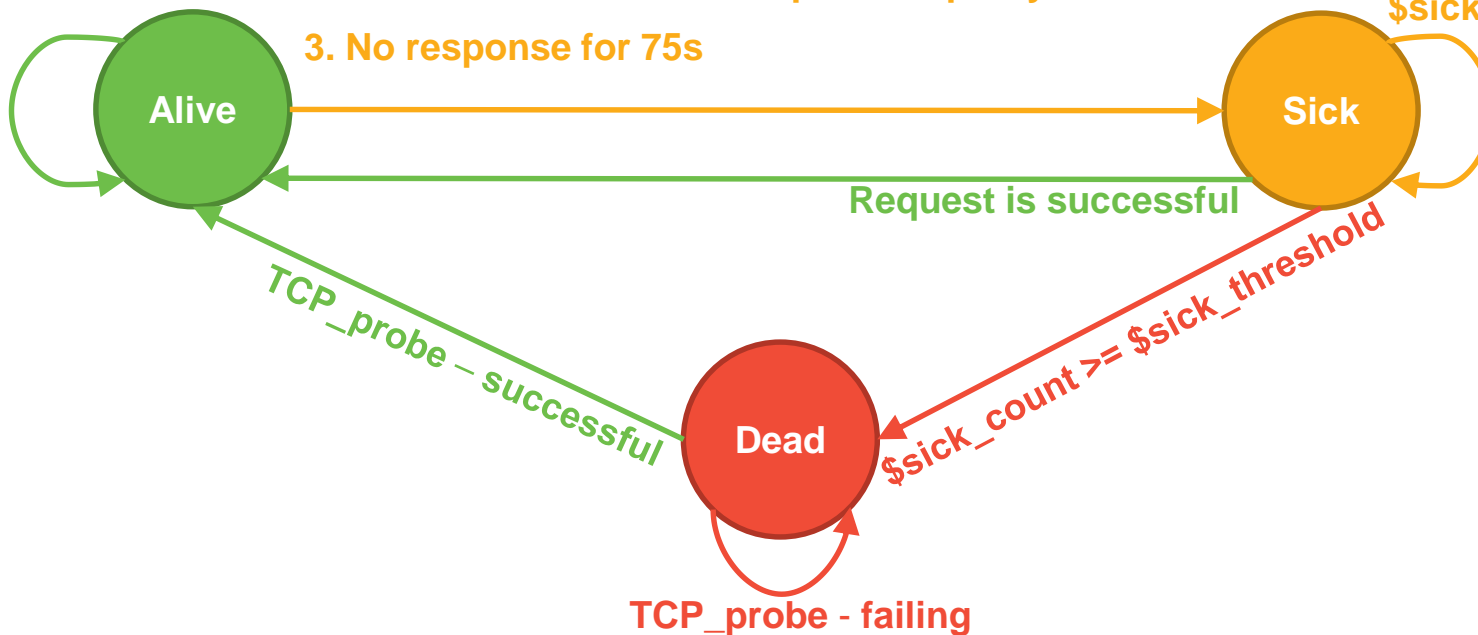


Upstream Proxy – Peer State Machine

Destinations
responsive

1. No HTTP VIA response header & HTTP 50x received
2. TCP RST received from upstream proxy
3. No response for 75s

Request still
failing (1-3)
\$sick_count++





Network Configuration – DNS

- Priority of 0 takes precedence
- If more DNS server entries have the same priority, WSA will chose DNS IP **randomly**
 - If deterministic DNS server choice is needed -> use different DNS priorities
- Use Microsoft AD IP address in DNS list if LDAP/NTLM authentication is used

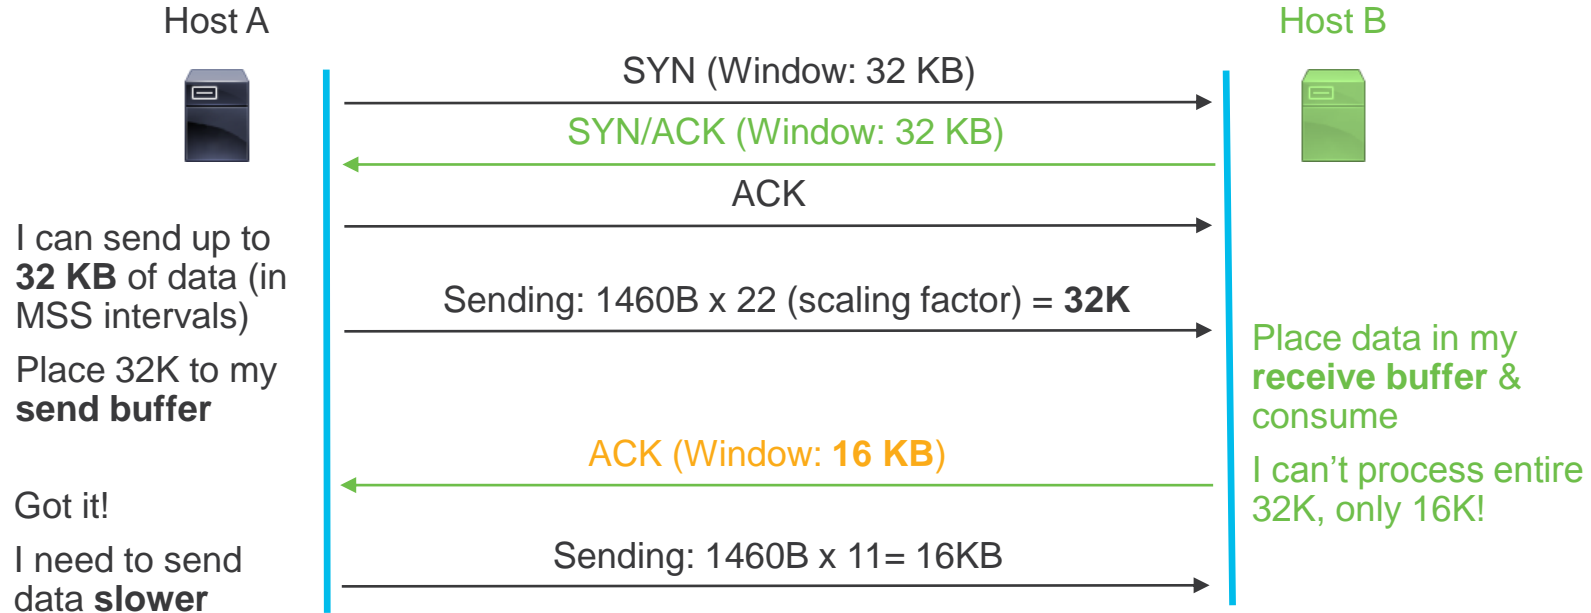
| DNS Server Settings | | |
|---|------------------------|---------------|
| DNS Servers: | Use these DNS Servers: | |
| | Priority | IP Address |
| | 0 | 10.49.222.60 |
| | 1 | 10.49.217.190 |
| Routing Table for DNS traffic: | Management | |
| IP Address Version Preference: | Prefer IPv4 | |
| Wait Before Timing out Reverse DNS Lookups: | 20 seconds | |
| DNS Domain Search List: | None | |



Advanced DNS Lookup Options

- WSA CLI -> **advancedproxyconfig -> DNS**
- **Four modes**
 - Select one of the following options:
 - 0 = Always use DNS answers in order
 - 1 = Use client-supplied address then DNS
 - 2 = Limited DNS usage
 - 3 = Very limited DNS usage
 - Proxy will by default failover to DNS results if upstream proxy is unresponsive
 - By default WSA used client supplied address, and then DNS lookup.
 - If using upstream proxies only, the best practice is to **use option 3**

TCP Windows / Send and Receive Buffers



Q: How does BW correlate to buffer size?

A: $TCP_Buffer_Size = Max_BW \times RTT$

Example

$BW=100\text{ Mbps}, RTT=80\text{ ms}$

$Buffer_Size \geq 1000\text{ KB}$

Network Configuration – Network Tuning Command

- New **networktuning CLI** command introduced in 9.1.1-074 GD
- **Use with care!** And only if there is need to increase individual dl/upload speed

```
wsa.example> networktuning
```

```
Choose the operation you want to perform:
```

- **SENDSPACE** - TCP sendspace (8192-131072) default 16000
- **RECVSPACE** - TCP receivespace (8192-131072) default 32768
- **SEND-AUTO** - TCP send autotuning (ON=1/OFF=0) default OFF
- **RECV-AUTO** - TCP receive autotuning (ON=1/OFF=0) default OFF
- **MBUF CLUSTER COUNT** - number of mbuf clusters(98304, 1572864) default varies as per system memory (4G RAM => 98304) linear scaling is recommended with increasing RAM
- **SENDBUF-MAX** - Maximum send buf, size(131072 - 2097152) default, 1MB=1048576
- **RECVBUF-MAX** - Maximum recv buf, size(131072 - 2097152) default, 1MB=1048576
- **CLEAN-FIB-1** - Remove all M1/M2 entries from Data routing table

Network Tuning Recommendations (more aggressive settings)



| Model Group | Memory | SEND-AUTO & RECV-AUTO | Proxy controlling dynamic windows sizes (send and receive)* | SENDSPACE | RECVSPACE | MBUF CLUSTER COUNT |
|------------------------------|--------|-----------------------|---|---------------|---------------|--------------------|
| S000v,S100v(6),S170, S370(4) | 4 GB | ON = 1 | NO/NO | 32768 - 65536 | 32768 - 65536 | 98304 |
| S370(8), S190(8), S300v 8 | 8 GB | ON = 1 | NO/NO | 65536 | 65536 | 196608 |
| S670(8-16), S380(16) | 16 GB | ON = 1 | NO/NO | 131072 | 131072 | 393216 |
| S680 (32), S390 (32) | 32 GB | ON = 1 | NO/NO | 131072 | 131072 | 786432/1572864 |
| S690, S690X | 64 GB | ON = 1 | NO/NO | 131072 | 131072 | 1572864 |

*Configured in WSA CLI -> **advancedproxyconfig -> MISCELLANEOUS**

Would you like proxy to perform dynamic adjustment of TCP receive window size?[Y]>N

Would you like proxy to perform dynamic adjustment of TCP send window size?[Y]>N

WSA and Authentication

Refresher and configuration considerations

- **Authentication methods supported**
 - LDAP/Basic
 - NTLMSSP/NTLM
 - Kerberos
 - TUI (Transparent User Identification) – CDA/ISE
- Choosing the right authentication surrogate:
 - **Avoid using NO SURROGATES** – too much burden on Auth helper processes
 - Use IP surrogates where applicable (timeout can be reduced – **<15min is not recommended**)
 - If IP surrogates can't fit the deployment -> use session cookies*

Logging and Reporting Tips

- Push one **accesslogs** subscription to SIEM using FTP/SCP/Syslog, rather than keeping them on WSA with large file size
- In order to easily troubleshoot, add the following Optional logging parameters to accesslogs:

Date: %L Dst-IP: %k UserAgent: %u ADGroup: %g AuthMethod: %m TransID: %I

| Log Subscription | |
|---------------------------|---|
| Log Type: | Access Logs |
| Log Name: | <input type="text" value="accesslogs"/> <small>(will be used to name the log directory)</small> |
| Rollover by File Size: | <input type="text" value="10G"/> Maximum <small>(Add a trailing K or M to indicate size units)</small> |
| Rollover by Time: | <input type="text" value="None"/> |
| Log Style: | <input checked="" type="radio"/> Squid <input type="radio"/> Apache <input type="radio"/> Squid Details |
| Custom Fields (optional): | <input type="text" value="Date: %L Dst-IP: %k UserAgent: %u AD"/> Custom Fields Reference |
| File Name: | <input type="text" value="aclog"/> |

Troubleshooting WSA Performance Issues

How many times have you felt like this?



Not sure if Internet is slow

...

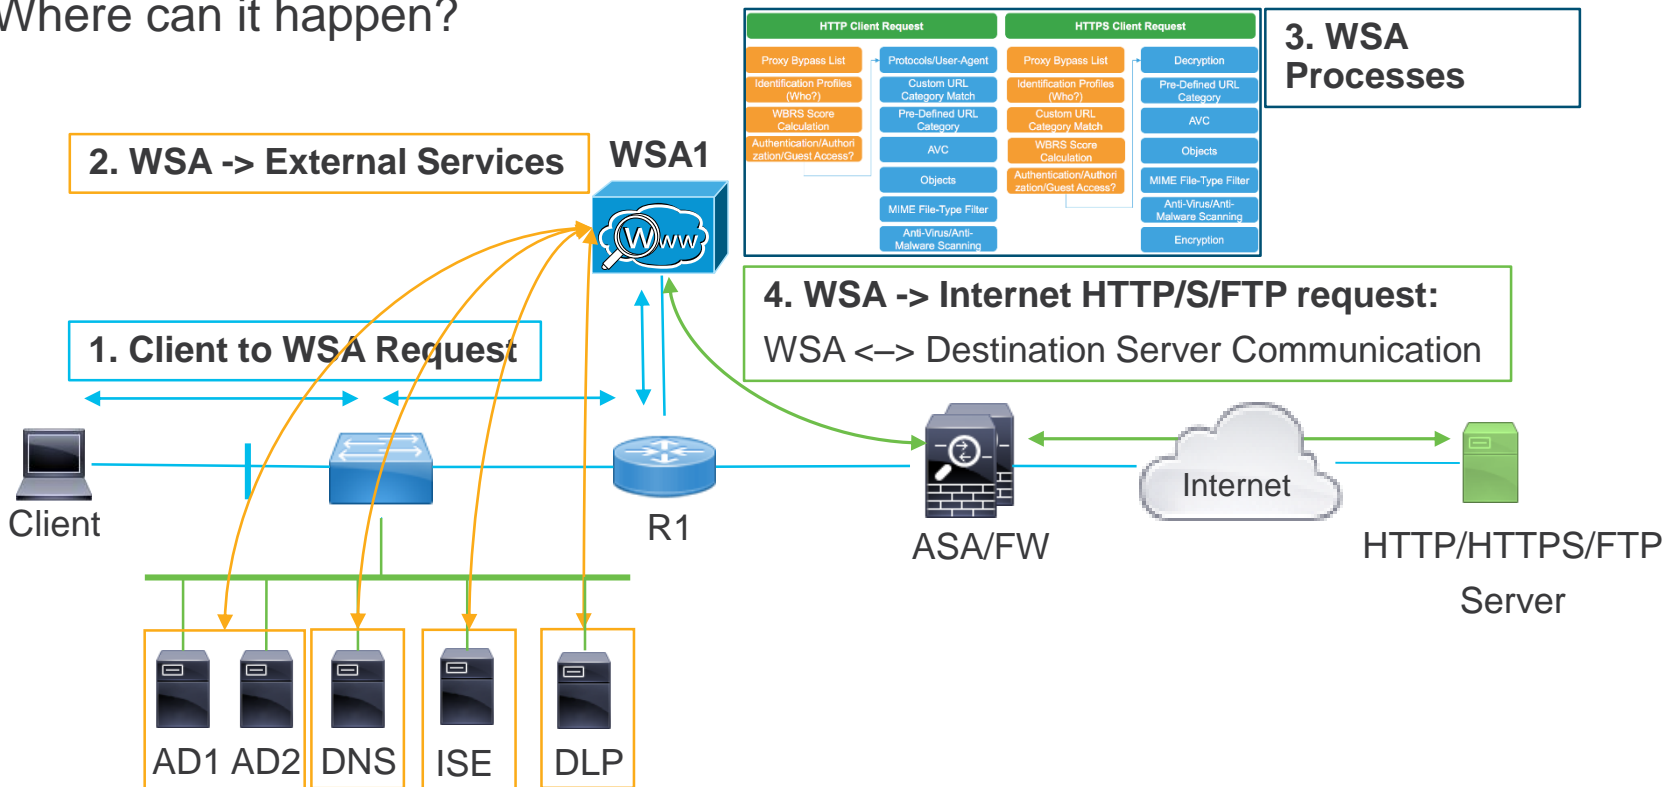
Or my mind is suddenly faster?

Troubleshooting WSA Performance Issues – Usual Causes

- WSA Sizing Issues / WSA is overloaded
- Configuration Complexity
- DNS Issues / Slowness
- Authentication Issues / Slowness
- Network Issues
- Reporting & Logging / WSA Disk Performance Issues

Slow Internet access

Where can it happen?



“Our Internet is slow”

What questions do I need to ask?

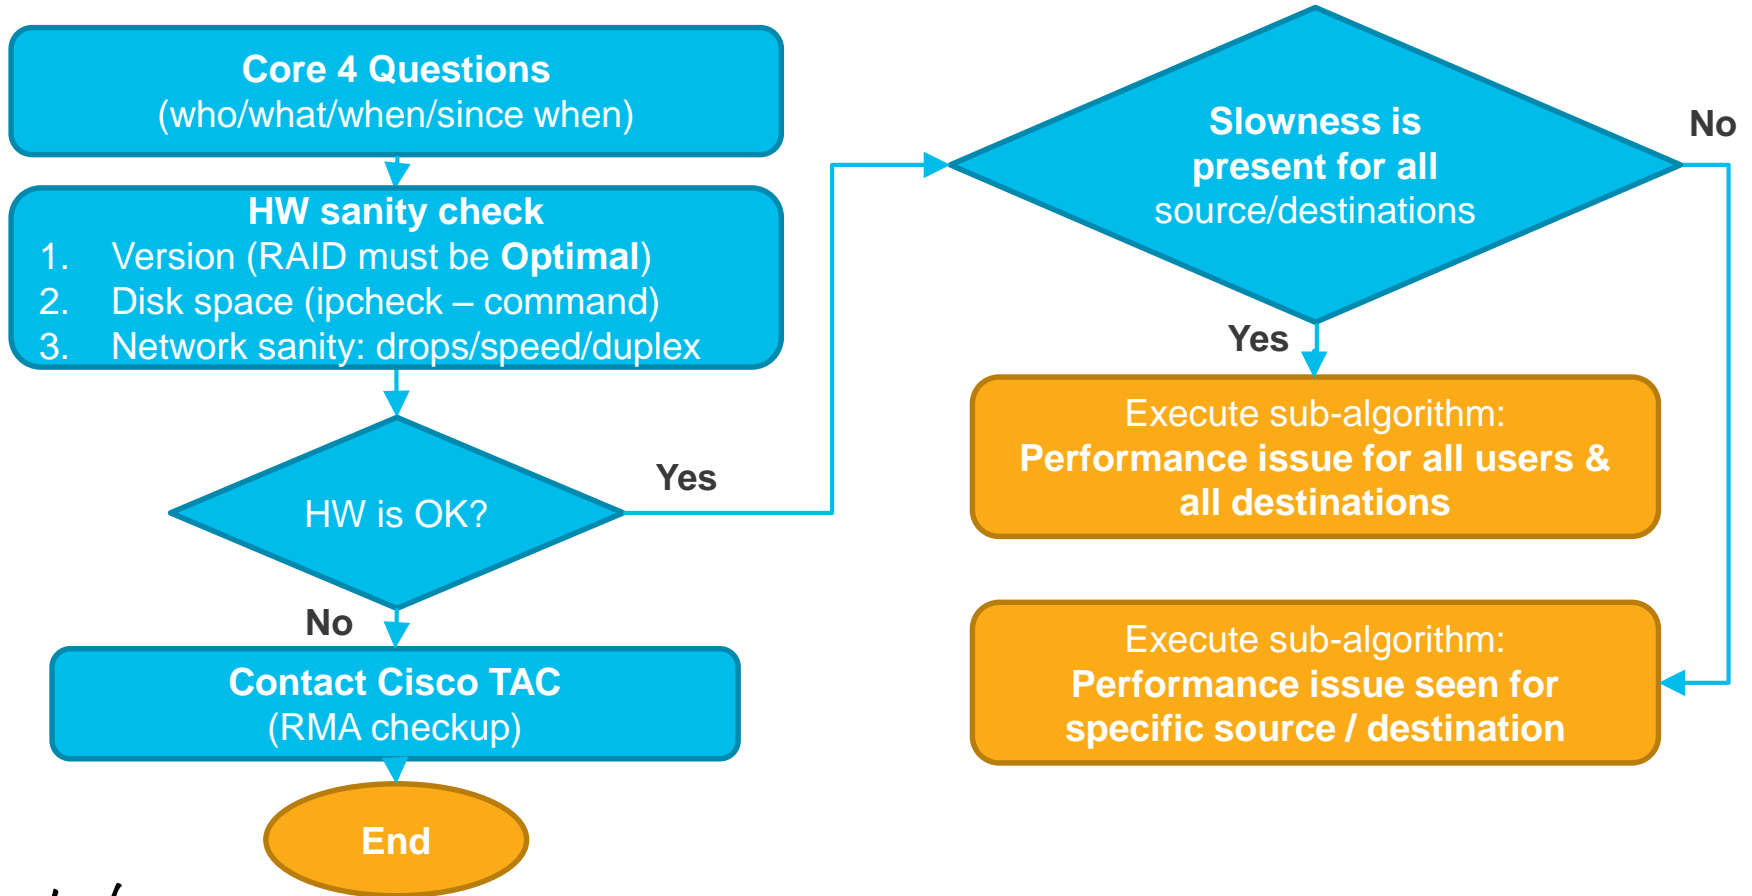
- **WHO** is affected?
 - One user / Group of users vs All users
 - Example: IP/subnet/username(s) of users
- **WHAT** are users searching for?
 - Specific URL / vs All URLs
 - HTTP / HTTPS / FTP? Upload/download?
- **WHEN** is this happening?
 - All the time?
 - Specific time (morning/noon/peak-traffic time)?
- **SINCE WHEN** & did something change?

So, your Internet is SLOW?

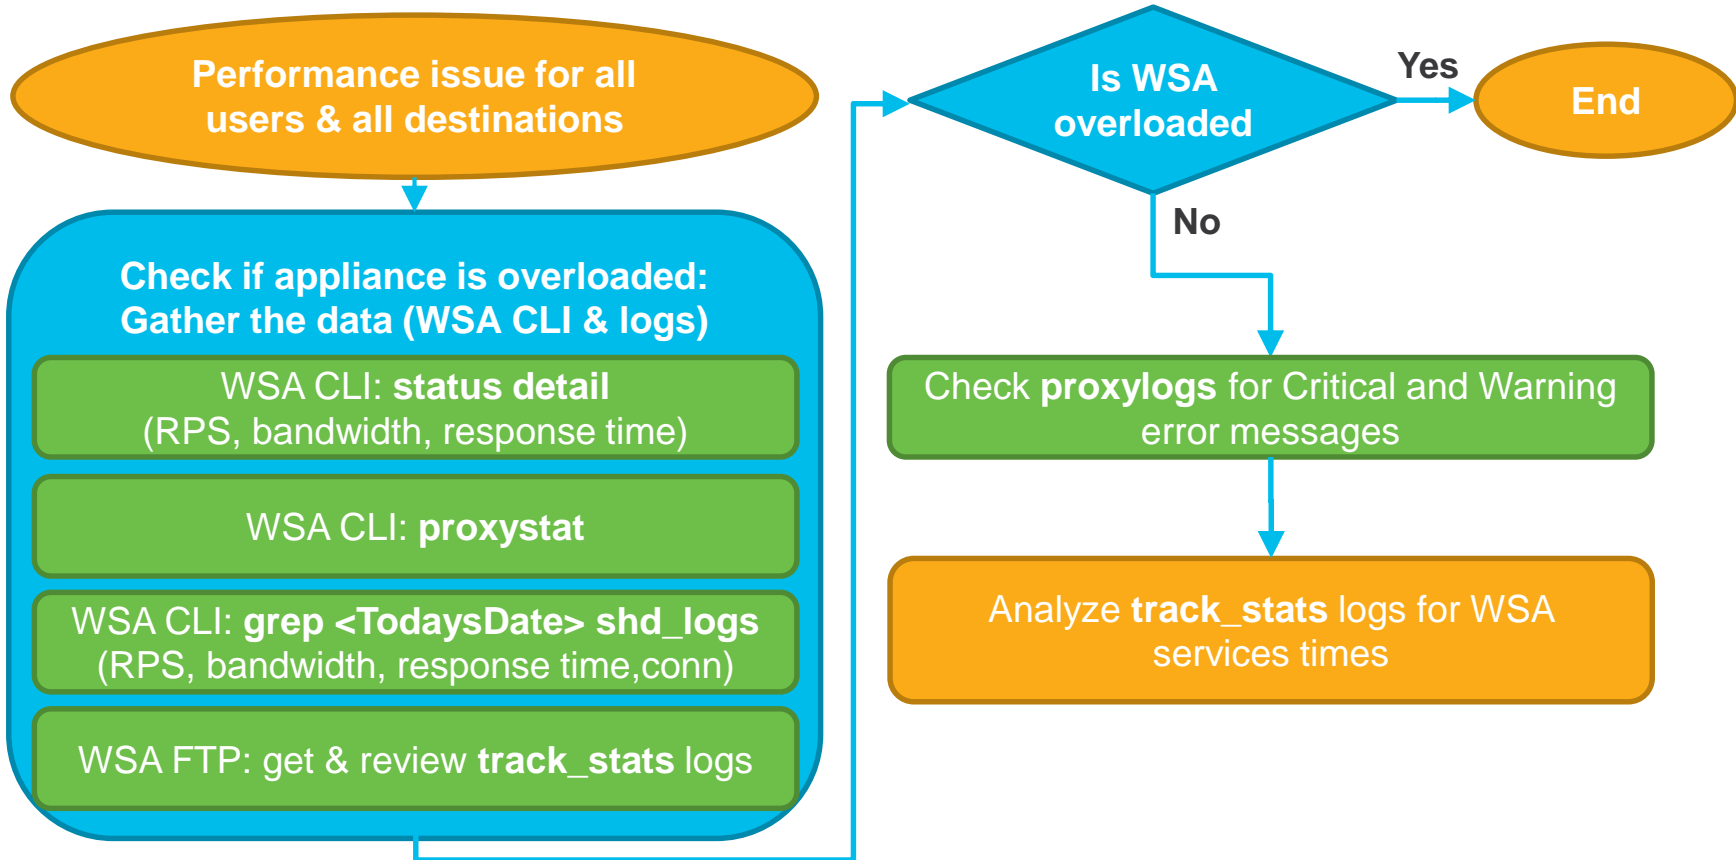


Tell me more about it!

Performance Troubleshooting Workflow (1)



Performance Troubleshooting Workflow (2)



The most common causes of Performance Issues

- WSA Sizing Issues (is your WSA overloaded?)
- Configuration Complexity is too high
- DNS Issues / Slowness
- Authentication Issues / Slowness
- Network Issues (Client to WSA / WSA to remote Server /PMTU Discovery)
- Reporting & Logging / WSA Disk Performance Issues

Internet Traffic Trends



Internet Users

27.1B




Conn & Devices

3.3 ZB



Internet volume
2021



HTTP/S



Video Traffic



Websites



Average Internet traffic
YoY: 2016 to 2021



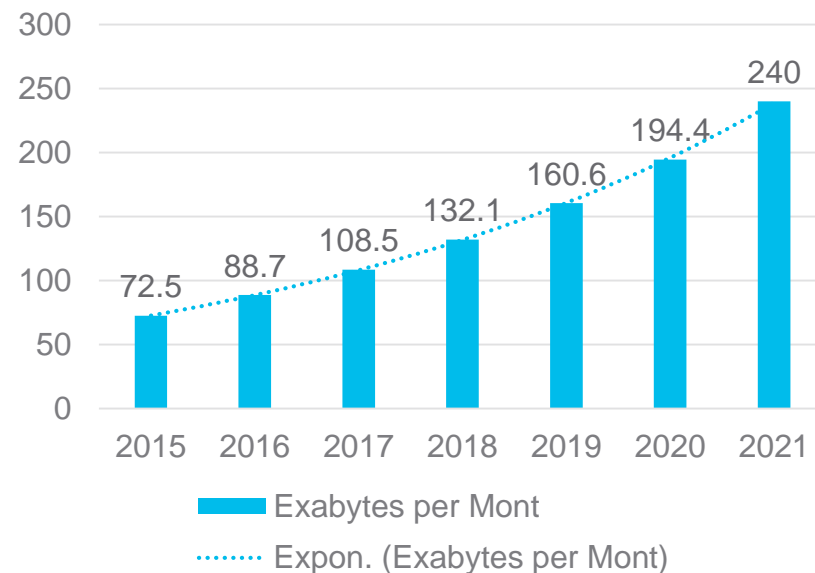
Busy Hour Increase
YoY: 2016 to 2021



Global Internet Traffic Projections

| Year | Global Internet Traffic |
|------|-------------------------|
| 1992 | 100 GB/day |
| 1997 | 100 GB/hour |
| 2002 | 100 GB/s |
| 2007 | 2000 GB/s |
| 2015 | 20,235 GB/s |
| 2020 | 61,386 GB/s |

Exabytes per Mont



Source: Cisco Systems, “The Zettabyte Era - Trends and Analysis“

WSA Sizing related Issues– Why is Sizing Important



Key factors to WSA sizing

- **Number of RPS (requests per second)**
- Percentage of decrypted HTTPS traffic
- Year to Year projected traffic growth



WSA Features Used

- Essential Bundle
- Anti-Malware Bundle
- Premium
- AMP / Threat Grid (on Premises)

WSA Sizing related Issues (2)



Preparations – current system survey

- Find out the RPS / traffic mix of the current system & HTTPS decrypted ratio
- Use Cisco Sizing Calculator to find out the model / and YoY recommendations
- **What if customer doesn't know number of RPS?**
 - Do a manual estimation based on known parameters:
 - Non-heavy users – we estimate 10% concurrently active users
 $\$RPS = \$UserNumber * 0.1$
 - Heavy Internet users – 15%-20 concurrently active users:
 $\$RPS = \$UserCount * 0.15$
 - Example: RPS = 1000 users * 0.15 = 150 Requests/s – in Year 1
 - Take into the account 20-25% traffic growth YoY – i.e:150 RPS in Y1, 188 in Y2...

WSA Sizing Calculator (available for Cisco Partners)

<https://fwm.cisco.com>

WSA Sizing Calculator and
Sizing Guides in PDF

Use ProxyLogAnalyzer Tool

Sizing Calculator

Home - Sizing Calculator

Important: Please note that the sizing for WSA is dependent on the correctness of traffic parameters you enter in the form below. If you are not sure about the exact values, we strongly recommend that you download the ProxyLogAnalyzer tool and analyze your proxy access logs before you proceed with the sizing calculation.
To download the tool use below links :

- [Windows Operating System](#)
- [Macintosh Operating System](#)

Additional Information:

- [Video Tutorial](#)
- [AsyncOS 9.1 Sizing Guide](#)
- [AsyncOS 9.0 Sizing Guide](#)
- [AsyncOS 8.5 Sizing Guide](#)

| | | |
|---|---|--|
| AsyncOS Version | <input type="text" value="AsyncOS 9.1.1v"/> | |
| Total no of Users | <input type="text" value="1000"/> | User Type <input type="text" value="Power Users"/> |
| Unique IP Address Count [?] | <input type="text" value="1500"/> | <i>(Please change if required)</i> |
| Peak RPS | <input type="text" value="150"/> | <i>(Please change if required)</i> |
| Percentage of RPS that is HTTPS Decrypted | <input type="text" value="30"/> | |
| Bundle | <input type="text" value="AMP Threat Grid (Cloud) + Sophos + Webroot"/> | |
| Appliance | <input type="text" value="S390"/> | Suggest |
| YoY Growth Factor (%age) | <input type="text" value="20"/> | |
| Number of years | <input type="text" value="3"/> | |
| Target Proxy CPU Utilization (%age) | <input type="text" value="80"/> | |
| Require Standby (n+1) | <input type="text" value="Yes"/> | |

Is WSA Overloaded?

OK, my sizing was good back then, still - I have slowness issues - now what?

Check if WSA is overloaded:

1. WSA CLI: **status detail**
 - Get CPU value
 - Get Response Time in last min & hour
 - Bandwidth Used
 - Server/Client Connection Count
2. Issue **proxystat** WSA CLI command
3. Check **shd_logs**
4. Check proxy **trackstat** logs

Note

- High Memory Utilization is normal (85%+ is expected)
- CPU usage should be **< 80/85%**
- Response time in last min/hour should be ideally **< 1000ms / to 1500ms**

Is WSA overloaded – Tools: status detail

```
wsa01> status detail
Status as of: Mon May 30 12:12:33 2016 CEST
Up since: Tue May 17 17:46:28 2016 CEST (12d 18h
26m 5s)
System Resource Utilization:
  CPU                2.4%
  RAM                32.9%
  Reporting/Logging Disk 19.8%
Transactions per Second:
  Average in last minute 150
  Maximum in last hour 180
  Average in last hour 145
  Maximum since proxy restart 478
  Average since proxy restart
Bandwidth (Mbps):
  Average in last minute 15.015
  Maximum in last hour 17.694
  Average in last hour 14.254
  Maximum since proxy restart 102.667
  Average since proxy restart 0.058
```

```
Response Time (ms):
  Average in last minute 6081
  Maximum in last hour 14798
  Average in last hour 4618
  Maximum since proxy restart 2146808
  Average since proxy restart 1179
Cache Hit Rate:
  Average in last minute 0
  Maximum in last hour 1
  Average in last hour 0
  Maximum since proxy restart 27
  Average since proxy restart 0
Connections:
  Idle client connections 5
  Idle server connections 1
  Total client connections 7
  Total server connections 11
```



Is WSA Overloaded – Tools: status detail

```
Status as of: Fri Jul 08 12:37:21 2016 CEST
Up since: Thu May 26 15:18:11 2016 CEST (42d 21h
19m 9s)
System Resource Utilization:
  CPU                4.6%
  RAM                37.5%
  Reporting/Logging Disk 10.6%
Transactions per Second:
  Average in last minute    407
  Maximum in last hour     646
  Average in last hour     372
  Maximum since proxy restart 646
  Average since proxy restart 377
Bandwidth (Mbps):
  Average in last minute    56.097
  Maximum in last hour     417.692
  Average in last hour     49.080
  Maximum since proxy restart 417.692
  Average since proxy restart 49.879
```

```
Response Time (ms):
  Average in last minute    5427
  Maximum in last hour     512373
  Average in last hour     3998
  Maximum since proxy restart 512373
  Average since proxy restart 4035
Cache Hit Rate:
  Average in last minute    23
  Maximum in last hour     36
  Average in last hour     18
  Maximum since proxy restart 36
  Average since proxy restart 19
Connections:
  Idle client connections   3284
  Idle server connections   2047
  Total client connections  14515
  Total server connections  13120
```

Is discovered RPS higher than recommended?

- Compare Average RPS from status detail output with Sizing Guide values
- If we see Average RPS in the last Hour more than recommended => we do most likely handle more traffic than WSA sizing recommends.
- **Example**
 - Average RPS in last minute: **150**
 - Average RPS in last hour: **145**
 - Maximum RPS in last hour: **180**
 - Model: **S170** & SW AsyncOS 9.0.x
 - **Feature Set – Premium** (WBRS, NTLM, WUC, Webroot, Sophos, Adaptive Scanning)
 - Sizing Sustained RPS: **100 request/s**
 - **Question: Is this WSA overloaded?**
Answer: Most likely YES, but we for sure need more evidences

Is WSA overloaded – Tools: **proxystat**

- Proxystat WSA CLI tool will show live statistics when it comes to proxy process
- Convenient to be used as quick “sanity check” - output will come each second

| %proxy | reqs | | | | client | server | %bw | disk | disk |
|--------------|-------------|----------|--------------|-------------|---------------|---------------|------------|--------------|-----------|
| CPU | /sec | hits | blocks | misses | kb/sec | kb/sec | saved | wrs | rds |
| 2.00 | 1 | 0 | 0 | 12 | 12 | 12 | 6.2 | 0 | 0 |
| 0.00 | <u>2781</u> | 0 | <u>26416</u> | <u>1381</u> | <u>18083</u> | <u>18083</u> | 0.0 | 715 | 0 |
| 55.00 | <u>2547</u> | 0 | <u>23836</u> | <u>1619</u> | <u>40576</u> | <u>40576</u> | 0.0 | <u>5568</u> | 23 |
| 55.00 | <u>2498</u> | 0 | <u>23092</u> | <u>1875</u> | <u>131414</u> | <u>131414</u> | 0.0 | <u>10446</u> | 350 |
| 61.00 | <u>3545</u> | 0 | <u>33625</u> | <u>1812</u> | <u>124975</u> | <u>124975</u> | 0.0 | <u>9373</u> | 45 |
| 72.00 | <u>3272</u> | 0 | <u>30505</u> | <u>2181</u> | <u>157903</u> | <u>157903</u> | 0.0 | <u>7700</u> | 18 |
| 71.00 | <u>2995</u> | 0 | <u>27575</u> | <u>2361</u> | <u>135113</u> | <u>135113</u> | 0.0 | <u>3600</u> | 57 |
| 62.00 | <u>2500</u> | 0 | <u>22317</u> | <u>2676</u> | <u>141960</u> | <u>141960</u> | 0.0 | <u>1678</u> | 471 |



Is WSA overloaded - Logs: **shd_logs**

- **shd_logs** are generated **each minute** by WSA and contain general performance indicators like:
 - CPU Usage
 - Memory Usage
 - Swap Memory in/out usage
 - RPS
 - Bandwidth
 - Client Connections
 - Server Connections
 - Load of various WSA Services (AM/AV)/Reporting/WBRS
 - Etc...



shd_logs – field description

| SHD log field Name | Explanation |
|--------------------|---|
| CPULd | Percent of CPU used on the system as reported by the OS, 0-100% |
| DskUtil | Space used on the log partition, percentage 0-100% |
| RAMUtil | Percent of memory free, as reported by OS, 0-100% |
| Reqs | Average number of transactions (requests) in past minute |
| Band | Average bandwidth saved in the past minute |
| Latency | Average latency (response time) in the last minute |
| CacheHit | Cache hit average in the past minute |
| CliConn | Total number of current Client Connections |
| SrvConn | Total number of current Server Connections |
| MemBuf | Current total amount of Memory buffers that are free |
| SwpPgOut | Number of pages that were swapped out, as reported by OS |
| xLd entries | CPU utilization by specific WSA service |

Is WSA Overloaded? - Logs: `shd_logs`

Example: `wsa1> grep -e "Nov..9.*CPULd.[2-9][0-9]" shd_logs`

We are searching for dates that contain 9 in the date, and where CPU load was greater or equal than 20%:

```
Wed Nov 9 18:08:38 2016 Info: Status: CPULd 25.2 DskUtil 72.5 RAMUtil 76.3 Reqs
145 Band 5726 Latency 74 CacheHit 14 CliConn 1156 SrvConn 587 MemBuf 0 SwpPgOut
7382983 ProxLd 26 Wbrs_WucLd 1.1 0.0 WebrootLd 0.0 SophosLd 0.0 McafeeLd 0.0
LogLd 1.4 RptLd
Sat Nov 19 11:48:41 2016 Info: Status: CPULd 21.0 DskUtil 70.4 RAMUtil 78.7 Reqs
25 Band 609 Latency 190 CacheHit 0 CliConn 482 SrvConn 385 MemBuf 0 SwpPgOut
593518 ProxLd 57 Wbrs_WucLd 0.0 LogLd 0.0 RptLd 0.0 WebrootLd 0.0 SophosLd 0.0
McafeeLd 0.0
```

Tip: You can be creative when it comes to using `grep` WSA CLI command.

Narrow down what you really need – is it date of Nov.19: `grep "Nov.19" shd_logs`



Is WSA Overloaded? - Logs: `shd_logs`

```
wsa> grep "CPUld.[2-9][0-9].*Reqs.[1-9][0-9][0-9]" shd_logs
```

```
Thu Nov 17 09:15:10 2016 Info: Status: CPUld 66.6 DskUtil 86.4 RAMUtil  
76.3 Reqs 100 Band 6292 Latency 373 CacheHit 6 CliConn 1728 SrvConn 1439  
MemBuf 0 SwpPgOut 320164 ProxLd 11 Wbrs_WucLd 2.7 LogLd 0.0 RptLd 0.0  
WebrootLd 0.0 SophosLd 4.3 McAfeeLd 0.0
```

```
Thu Nov 17 13:26:57 2016 Info: Status: CPUld 81.2 DskUtil 86.4 RAMUtil  
79.8 Reqs 136 Band 43110 Latency 722 CacheHit 23 CliConn 1799 SrvConn 1264  
MemBuf 38398 SwpPgOut 498504 ProxLd 15 Wbrs_WucLd 4.1 LogLd 0.0 RptLd 0.0  
WebrootLd 0.0 SophosLd 9.7 McAfeeLd 0.0
```


Logs: **track_stats** logs

- **Track_stats** logs are one of the most important logs when speaking about “general performance” of WSA
- Generated by default **each 5 minutes**
- Contain valuable information about:
 - CPU usage (system vs User time), Avg/Peak RPS, throttling connection count
 - Memory chunk usage/swap & Memory chunk allocation
 - Cumulative Request Statistical Distribution of number of request delay per type of WSA service (how much we waited for DNS, Remote Server, WSA to Client, AV Scanner etc...) & much more

Logs: track_stats logs – how to get them?

- **Only** available using SCP or FTP access to WSA's Management IP under
- **Where:** Under `/track_stats/` directory

```
scp admin@10.49.222.72:/track_stats/prox_track.log /log_store/prox_track.log
```



1. Getting current track_stats via SCP



Index of /

| Name | Size | Date Modified |
|---------------------|--------|-----------------------|
| configuration/ | | 11/23/16, 4:25:00 PM |
| captures/ | | 8/26/16, 2:29:00 PM |
| diagnostic/ | | 8/26/16, 2:29:00 PM |
| track_stats/ | | 11/22/16, 8:00:00 AM |
| password_words.txt | 1.1 MB | 8/26/16, 2:37:00 PM |
| external_auth_logs/ | | 11/11/16, 11:33:00 AM |
| audit_logs/ | | 11/22/16, 7:54:00 PM |
| upgrade_logs/ | | 11/11/16, 11:33:00 AM |

2. Getting current track_stats via FTP



Logs: track_stats logs

WSA Overloaded issues: What can I look in **track_stats**?

- **Step 1** (Prerequisite): Obtain **track_stats** logs at external machine (xNIX)
- **Step 2**: Use **grep** tool on external machine to extract the following data

```
Linux-machine$ grep -iE "date|traffic over|thrott|(user|system) time" prox_track.log
##snip##
Current Date: Sat, 19 Nov 2016 17:29:06 CST
           user time: 3.422 (1.141%)
           system time: 0.459 (0.153%)
INFO: traffic over past minute - 16.53 reqs/sec
INFO: traffic over past hour - 25.05 peak / 15.29 avg reqs/sec
INFO: traffic over past day - 49.30 peak / 16.33 avg reqs/sec
INFO: traffic over past week - 49.30 peak / 16.33 avg reqs/sec
INFO: traffic over all time - 49.30 peak / 16.48 avg reqs/sec
INFO: percentage of throttled transactions to the total number of transactions over past minute - 0.00 %
INFO: percentage of throttled transactions to the total number of transactions over past hour - 0.00 peak / 0.00 avg %
INFO: percentage of throttled transactions to the total number of transactions over past day - 0.00 peak / 0.00 avg %
INFO: percentage of throttled transactions to the total number of transactions over past week - 0.00 peak / 0.00 avg %
INFO: percentage of throttled transactions to the total number of transactions over all time - 0.00 peak / 0.00 avg %
```



Average & Peak RPS

Logs: track_stats logs – HTTPS Volume

What can I look in track_stats?

- How much HTTPS traffic is WSA processing?

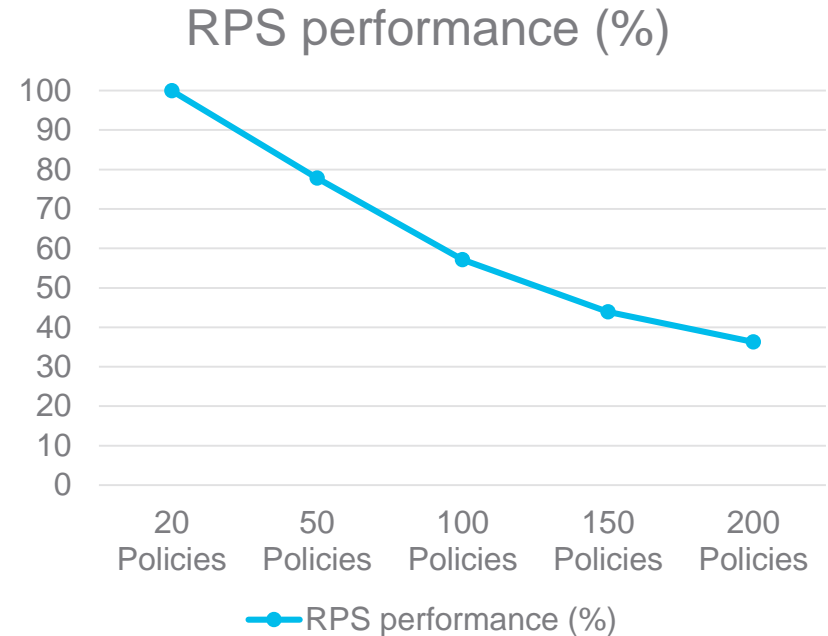
```
Linux-machine$ grep -iE "HTTPS|avg reqs/sec|traffic over|Total SSL" prox_track.log
##snip##
Current Date: Sun, 19 Nov 2017 01:59:49 CET
INFO: HTTPS Passthrough handshake skip count 0
# Traffic Rate                # Total Transactions                # HTTPS                # HTTPS(Passthrough)
                               [peak | avg reqs/sec]            [peak | avg reqs/sec]  [peak | avg reqs/sec]
traffic over past minute      0.22                               0.08                   0.00
traffic over past hour        2.20 | 0.31                        1.00 | 0.12           0.90 | 0.02
traffic over past day         34.70 | 0.09                       14.80 | 0.03          1.00 | 0.00
traffic over past week        34.70 | 0.02                       14.80 | 0.01          1.00 | 0.00
traffic over all time         34.70 | 0.03                       14.80 | 0.01          1.00 | 0.00
INFO: Total SSL Handshakes    : 36
INFO: Total SSL Handshakes Finished : 30
INFO: Total SSL Handshakes Unfinished : 6
```

The most common causes of Performance Issues

- WSA Sizing Issues (is your WSA overloaded?)
- Configuration Complexity is too high
- DNS Issues / Slowness
- Authentication Issues / Slowness
- Network Issues (Client to WSA / WSA to remote Server /PMTU Discovery)
- Reporting & Logging / WSA Disk Performance Issues

WSA Policy Complexity


- WSA performance is dependent on WSA Policy complexity
- More policies -> Less of a Performance
- Pay attention to:
 - Number of Identity Policies
 - **Number of Access, Decryption Policies**
 - Number of Custom Local & External URL Categories (especially with wide regex matching)



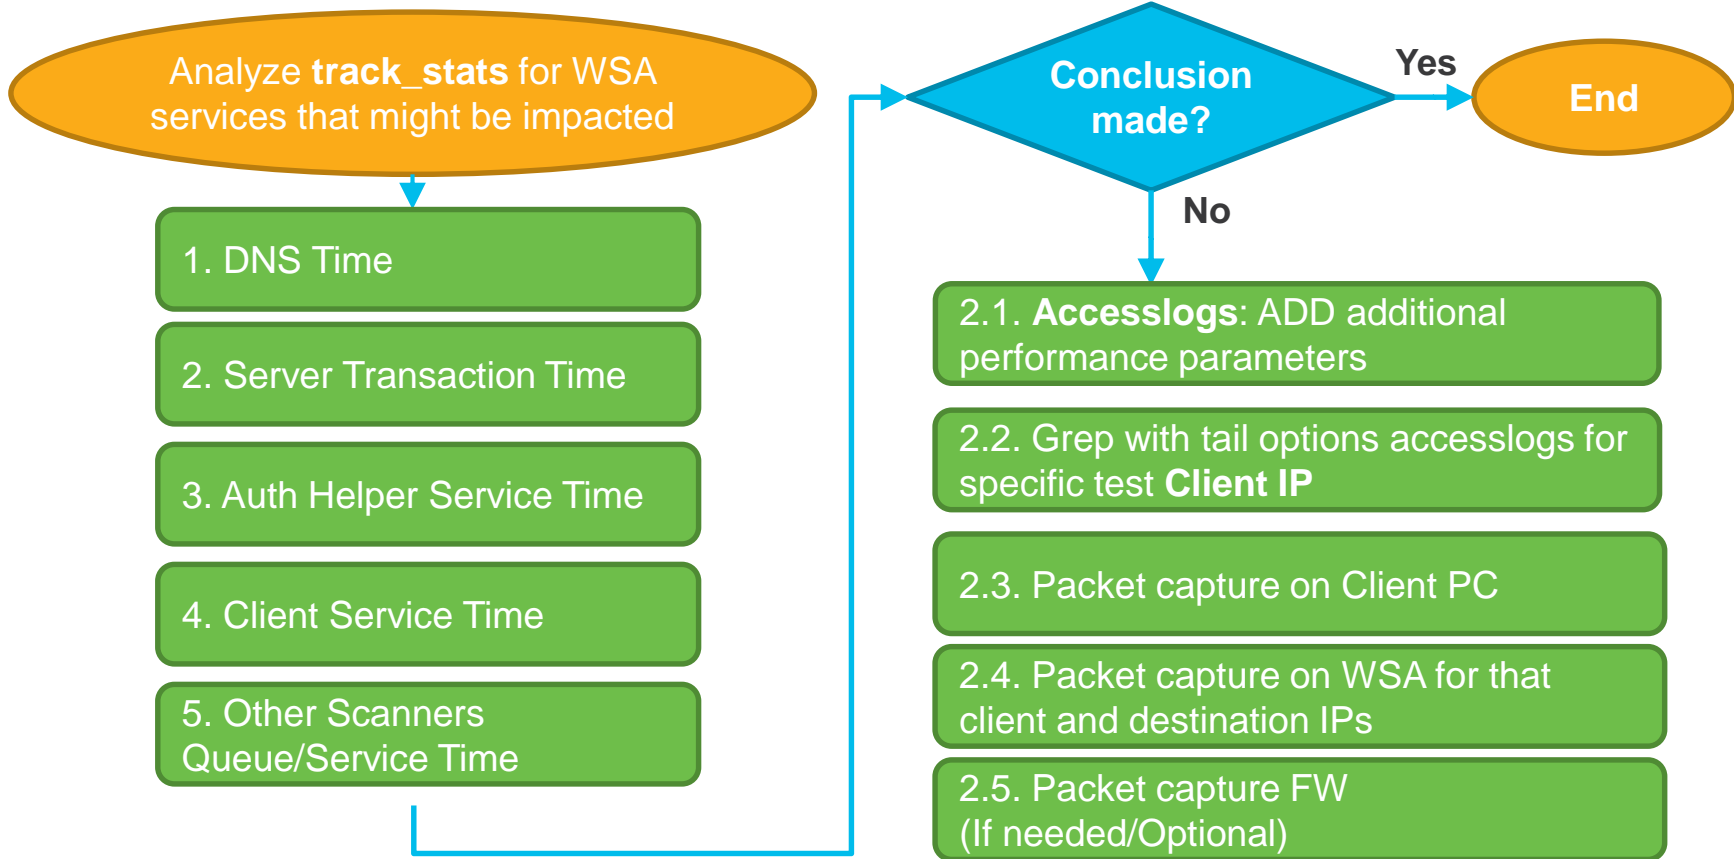
WSA policy complexity – how can I detect it?

- Count the policies / Identities / Custom URL categories / regex expressions
 - Keep them by recommendations from the first chapter
 - Keep the policy complexity **max Medium** (20 Policies/Identities/Custom URL cat)
- Check **track_stats** logs for User/System Time Ratio
 - If **User Time** is constantly much larger than System time - we spend a lot of CPU cycles on “User tasks”: **Configuration is too expensive!**
 - Example of a box with very large User Time (Configuration Complexity):

```
Linux-machine$ grep -iE 'date|(user|system) time' prox_track.log
Current Date: Sat, 1 Nov 2016 12:39:02 CET
                user time: 212.789
                system time: 10.554
```



Performance Troubleshooting Workflow – Trackstats



The most common causes of performance issues

- WSA Sizing Issues (is your WSA overloaded?)
- Configuration Complexity is too high
- DNS Issues / Slowness
- Authentication Issues / Slowness
- Network Issues (Client to WSA / WSA to remote Server /PMTU Discovery)
- Reporting & Logging / WSA Disk Performance Issues

DNS issues as source of performance problems

- The most common causes for slow Internet access with WSA:
 - Slowly responding **DNS** server
 - Network introducing slowness between WSA-DNS server(s)
- Indication of DNS slowness / tools to test:
 - Check DNS Server configuration (**WSA CLI -> dnsconfig**)
 - **nslookup** the FQDN
 - Use **dig** tool to check the response from DNS server

DNS issues – test tools (1)

- **nslookup** - quick DNS sanity test – response should arrive without big delay

```
wsa01> nslookup www.google.com  
A=216.58.204.132 TTL=30m
```

- **dig** – more comprehensive test – allows us to specify what DNS server we want to test the lookup against.
- Example usage of **dig** command:
 - **dig @10.49.222.60 www.cisco.com**
 - **dig @8.8.8.8 www.google.com**



DNS issues – test tools (2)

```
wsa01> dig @10.49.222.60 www.cisco.com

; <<>> DiG 9.8.4-P2 <<>> @10.49.222.60 www.cisco.com A
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->HEADER<<- opcode: QUERY, status: NOERROR, id: 46299
;; flags: qr rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;www.cisco.com.                IN      A

;; ANSWER SECTION:
www.cisco.com.                3596    IN      CNAME   origin-www.cisco.com.
origin-www.cisco.com.        47      IN      A       72.163.4.161


;; Query time: 3698 msec
;; SERVER: 10.49.222.60#53(10.49.222.60)
;; WHEN: Mon Nov 28 17:32:18 2016
;; MSG SIZE rcvd: 72
```

Test tools: track_stats logs – DNS Time

| Current Date: Tue, 10 Sep 2016 09:15:54 BST | | Number of requests |
|---|-----------|--------------------|
| DNS Time | 1.0 ms | 450578 |
| DNS Time | 1.6 ms | 2606 |
| DNS Time | 2.5 ms | 1281442 |
| DNS Time | 4.0 ms | 551855 |
| DNS Time | 6.3 ms | 281121 |
| DNS Time | 10.0 ms | 1602851 |
| DNS Time | 15.8 ms | 1076228 |
| DNS Time | 25.1 ms | 618966 |
| DNS Time | 39.8 ms | 354198 |
| DNS Time | 63.1 ms | 231474 |
| DNS Time | 100.0 ms | 430026 |
| DNS Time | 158.5 ms | 396313 |
| DNS Time | 251.2 ms | 162560 |
| DNS Time | 398.1 ms | 87255 |
| DNS Time | 631.0 ms | 33736 |
| DNS Time | 1000.0 ms | 23888 |
| DNS Time | 1584.9 ms | 17770 |
| DNS Time | 2511.9 ms | 19326 |
| DNS Time | 3981.1 ms | 11155 |
| DNS Time | 6309.6 ms | 358495 |

```
Server$grep -iE 'date|DNS Time' prox_track.log
```

DNS Time



The most common causes of performance issues

- WSA Sizing Issues (is your WSA overloaded?)
- Configuration Complexity is too high
- DNS Issues / Slowness
- Authentication Issues / Slowness
- Network Issues (Client to WSA / WSA to remote Server /PMTU Discovery)
- Reporting & Logging / WSA Disk Performance Issues

Authentication issues as cause of the performance degradation



- Authentication can introduce slowness in the cases when:
 - AD/LDAP server is responding slowly to WSA (or not responding)
 - ISE is responding slowly to WSA (when ISE-WSA integration is used)
 - **How can we test:**
 - WSA CLI: **testauthconfig** for general reachability of AD/LDAP server(s)
 - Check **auth_logs & proxylogs** for critical or warning error messages
- ```
wsa01> grep -e crit -e warn -e netlogon -i proxylogs
```
- ```
wsa01> grep -e crit -e warn -e netlogon -i auth_logs
```
- Check **track_stats** logs for **Auth Service Time** and **Auth Helper Wait Time**

Test tools: track_stats logs – Auth Helper Service/Wait Times

| Current Date: Tue, 10 Sep 2016 09:15:54 CET | Number of requests |
|---|--------------------|
| Auth helper Service Time 1.0 ms | 350578 |
| Auth helper Service Time 1.6 ms | 2606 |
| Auth helper Service Time 2.5 ms | 1281442 |
| Auth helper Service Time 4.0 ms | 551855 |
| Auth helper Service Time 6.3 ms | 281121 |
| Auth helper Service Time 10.0 ms | 1602851 |
| Auth helper Service Time 15.8 ms | 1076228 |
| Auth helper Service Time 25.1 ms | 618966 |
| Auth helper Service Time 39.8 ms | 354198 |
| Auth helper Service Time 63.1 ms | 231474 |
| Auth helper Service Time 100.0 ms | 430026 |
| Auth helper Service Time 158.5 ms | 396313 |
| Auth helper Service Time 251.2 ms | 162560 |
| Auth helper Service Time 398.1 ms | 87255 |
| Auth helper Service Time 631.0 ms | 33736 |
| Auth helper Service Time 1000.0 ms | 23888 |
| Auth helper Service Time 1584.9 ms | 17770 |
| Auth helper Service Time 2511.9 ms | 19326 |
| Auth helper Service Time 3981.1 ms | 11155 |
| Auth helper Service Time 6309.6 ms | 458495 |

```
Server$ grep -iE 'date|Auth' prox_track.log
```

Authentication Helper Service Time

- If authentication helpers are handling too many request, auth request is placed in a **queue**
- Shows distribution of requests by **Queuing Time** in order to get processed by Authentication helper
- If most of the requests are taking $\geq 1.5s$ to reach Authentication helper:

Check authentication surrogate: **Don't use NO SURROGATES**

Authentication Helper Wait Time

- Time since packet came into authentication process, and passed to LDAP/AD for processing – till AD sends a response

Authentication troubleshooting - connection requirements



Protocols & ports used by WSA Authentication process

| Protocol | Port |
|----------|-------------------|
| kinit | 88 |
| net | 88, 389, 445 |
| winbindd | 88, 389, 445, 139 |
| nmbd | 137, 138, 139 |

Protocol / Ports that need to be allowed if FW is between WSA and LDAP/AD

| Port | Protocol | Description |
|------|----------|----------------------------|
| 88 | TCP/UDP | Kerberos 5 tickets traffic |
| 389 | TCP/UDP | LDAP |
| 445 | TCP | SMB |
| 137 | UDP | NetBIOS Name Service |
| 138 | UDP | NetBIOS Datagram Service |
| 139 | TCP | NetBIOS Session Service |

The most common causes of performance issues

- WSA Sizing Issues (is your WSA overloaded?)
- Configuration Complexity is too high
- DNS Issues / Slowness
- Authentication Issues / Slowness
- Network Issues (Client to WSA / WSA to remote Server /PMTU Discovery)
- Reporting & Logging / WSA Disk Performance Issues



Network settings relevant to performance

WSA Model vs **networktuning** CLI settings

- More network memory and buffers for send/receive = more upload/download speed
- **Attention:** more network memory buffers = less concurrency – stay in recommended

Duplex/speed & MTU

- Make sure to either auto negotiate the speed/duplex for proxy interface
- Or – manually set via:
etherconfig CLI command

Path MTU discovery

- Make sure ICMP is enabled from WSA for PMTU discovery to work
- **Enable** it on WSA
- If PMTU discovery can't work, make sure WSA has correct MTU

```
wsa01> etherconfig
```

Choose the operation you want to perform:

- **MEDIA** - View and edit ethernet media settings.
- **VLAN** - View and configure VLANs.
- **MTU** - View and configure MTU

```
wsa01> pathmtudiscovery
```

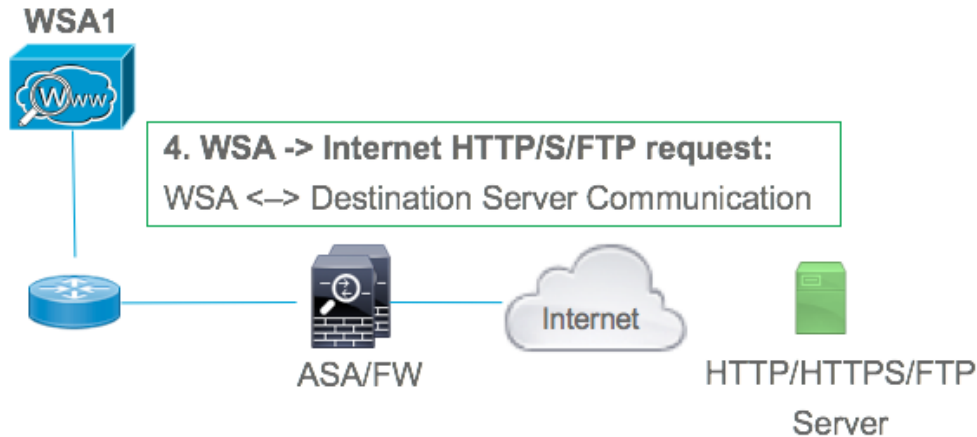
Path MTU discovery is currently enabled

Server Transaction Time



Huston, we have a problem, but it's **not** up to WSA 😊!

- **Server Transaction Time** in `track_stats` logs shows **time WSA waits to receive response from a remote server**
- If large portion of requests fall into buckets of **1,5s** or more -> problem is most likely introduced **upstream from WSA!**



Server Transaction Time – example calculation

| Current Date: Tue, 10 Sep 2016 09:15:54 CET | Number of requests | |
|---|--------------------|--------|
| Server Transaction Time 1.0 ms | 285 | |
| Server Transaction Time 1.6 ms | 57 | |
| Server Transaction Time 2.5 ms | 165 | |
| Server Transaction Time 4.0 ms | 112 | |
| Server Transaction Time 6.3 ms | 4397 | |
| Server Transaction Time 10.0 ms | 6174 | |
| Server Transaction Time 15.8 ms | 10091 | |
| Server Transaction Time 25.1 ms | 8351 | |
| Server Transaction Time 39.8 ms | 8537 | |
| Server Transaction Time 63.1 ms | 5927 | 34,39% |
| Server Transaction Time 100.0 ms | 7577 | |
| Server Transaction Time 158.5 ms | 6015 | |
| Server Transaction Time 251.2 ms | 7186 | |
| Server Transaction Time 398.1 ms | 6948 | |
| Server Transaction Time 631.0 ms | 5386 | 25,82% |
| Server Transaction Time 1000.0 ms | 4081 | |
| Server Transaction Time 1584.9 ms | 3284 | |
| Server Transaction Time 2511.9 ms | 2520 | |
| Server Transaction Time 3981.1 ms | 2254 | |
| Server Transaction Time 6309.6 ms | 38886 | 39,27% |

```
Inx$ grep -iE 'date|Server Transaction' prox_track.log
```

Low / Expected Server Transaction Time

Most of the request volume should go here

Medium / Expected Server Transaction Time

Still OK – not too alarming (medium)

High Server Transaction Time

Indication of network slowness between WSA and remote Site

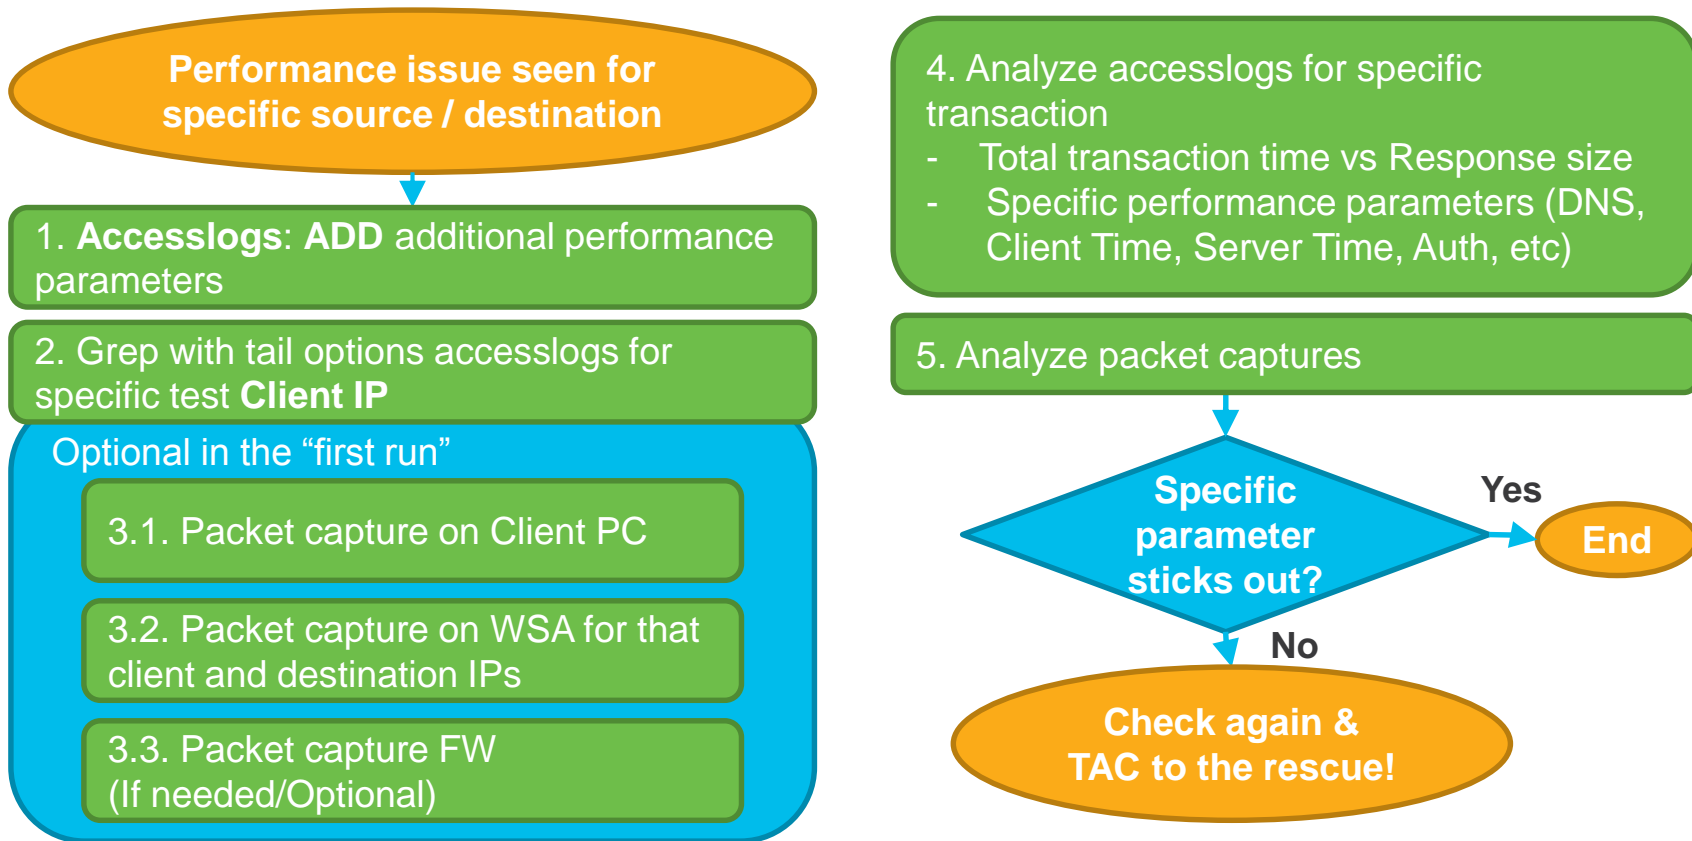
The most common causes of performance issues

- WSA Sizing Issues (is your WSA overloaded?)
- Configuration Complexity is too high
- DNS Issues / Slowness
- Authentication Issues / Slowness
- Network Issues (Client to WSA / WSA to remote Server /PMTU Discovery)
- Reporting & Logging / WSA Disk Performance Issues

Do we have performance issues due to heavy disk RD/WR operations?

- Performance issues might be side effect of:
 1. Lack of disk space
 2. Reporting process contributing to slowness (on WSA side)
 3. WBRS process writing too frequently on platforms with slower disk IO (**S160/S170/Virtual WSA** – depending on storage speed)
- How to verify (and try to implement a workaround):
 1. **ipcheck** CLI command – check for percentage of free disk space
 2. Check reportd_logs, reportqueryd_logs for critical and warning messages, if needed disable reporting and test again:
CLI -> diagnostic -> reporting -> DISABLE
 3. **Increase WBRS update interval** from default 5m to 1h / 8h or in severe slowness cases to 1d

Performance Troubleshooting Flow: Specific Src/Dst



Accesslogs with performance parameters – Style 1

- The best way to detect latency contributors on individual transactions
- How to set it up:
 - **WSA GUI -> System Administration -> Log Subscriptions: accesslogs**
 - Add the following in the Optional log parameters field:

```
Date: %L Dst-IP: %k UsrAgnt: %u ADGroup: %g AuthMethod: %m  
TransID: %I PrfPara: %:<a %:<b %:<d %:<h %:<r %:<s %:>1 %:>a  
%:>b %:>c %:>d %:>h %:>r %:>s %:1< %:1> %:b< %:b> %:h< %:h>  
%:m< %:m> %:w< %:w> %x
```



Accesslogs with performance parameters – Style 2

```
[ Request Details: ID = %I, User Agent = %u, AD Group
Memberships = ( %m ) %g ] [ Tx Wait Times (in ms): 1st byte to
server = %:<1, Request Header = %:<h, Request to Server = %:<b,
1st byte to client = %:1>, Response Header = %:h>, Client Body =
%:b> ] [ Rx Wait Times (in ms): 1st request byte = %:1<, Request
Header = %:h<, Client Body = %:b<, 1st response byte = %:>1,
Response header = %:>h, Server response = %:>b, Disk Cache =
%:>c; Auth response = %:<a, Auth total = %:>a; DNS response =
%:<d, DNS total = %:>d, WBRs response = %:<r, WBRs total = %:>r,
AVC response = %:A>, AVC total = %:A<, DCA response = %:C>, DCA
total = %:C<, McAfee response = %:m>, McAfee total = %:m<,
Sophos response = %:p>, Sophos total = %:p<, Webroot response =
%:w>, Webroot total = %:w<, Anti-Spyware response = %:<s, Anti-
Spyware total = %:>s; Latency = %x; %L ]
```

Accesslogs with performance parameters



Edit Log Subscription

| Log Subscription | |
|---------------------------|---|
| Log Type: | Access Logs |
| Log Name: | <input type="text" value="accesslogs"/> <i>(will be used to name the log directory)</i> |
| Rollover by File Size: | <input type="text" value="100M"/> Maximum <i>(Add a trailing K or M to indicate size units)</i> |
| Rollover by Time: | <input type="text" value="None"/> |
| Log Style: | <input checked="" type="radio"/> Squid <input type="radio"/> Apache <input type="radio"/> Squid Details |
| Custom Fields (optional): | <input type="text" value="Date: %L Dst-IP: %k UsrAgnt: %u AD"/> Custom Fields Reference |
| File Name: | <input type="text" value="aclog"/> |

Accesslogs – Performance Parameters – Style 1

```
1480424375.305 7518 10.49.222.72 TCP_MISS/200 1034 GET
http://www.cisco.com/ - DIRECT/www.cisco.com text/html
DEFAULT_CASE_12-ap.Policy20-id.subnet10-NONE-NONE-NONE-
DefaultGroup <IW_comp,8.7,0,"-",0,0,0,-,"-",,-,-,-,"-",,-,-,"-", "-
",-,-,IW_comp,-,"-",,"-", "Unknown", "Unknown", "-","-",1.10,0,-
,"Unknown", "-","-","-",,-,-,"-", "-",-> - Date:
"20/Nov/2016:13:59:35 +0100" Dst-IP: 72.163.4.161 UserAgent:
"curl/7.31.0" ADGroup: - AuthMethod: NONE TransID: 3106 PrfPara:
0 0 0 0 0 144 0 17 0 7089 0 41 17 0 17 0 0 0 0 0 0 0
```

Total request time [ms]

Request size [B]



Accesslogs – Performance Parameters – Style 2

```
1480974055.775 7408 10.60.71.233 TCP_MISS/200 67118 GET http://www.cisco.com/ -  
DIRECT/www.cisco.com text/html DEFAULT_CASE_12-DefaultGroup-id.ana-NONE-NONE-  
NONE-DefaultGroup <IW_comp,8.7,0,"-",0,0,0,1,"-",-,-,-,"-",1,-,"-","-",-  
,IW_comp,-,"Unknown","-","Unknown","Unknown","-","-",72.48,0,-,"Unknown","-  
",1,"-",-,-,"-","-",-> - Request Details: ID = 52, User Agent = "curl/7.43.0",  
AD Group Memberships = ( NONE ) - ] [ Tx Wait Times (in ms): 1st byte to server  
= 179, Request Header = 0, Request to Server = 0, 1st byte to client = 2056,  
Response Header = 0, Client Body = 544 ] [ Rx Wait Times (in ms): 1st request  
byte = 0, Request Header = 0, Client Body = 0, 1st response byte = 145,  
Response header = 0, Server response = 2056, Disk Cache = 0; Auth response = 0,  
Auth total = 0; DNS response = 0, DNS total = 3363, WBRs response = 0, WBRs  
total = 92, AVC response = 0, AVC total = 0, DCA response = 0, DCA total = 0,  
McAfee response = 0, McAfee total = 0, Sophos response = 0, Sophos total = 0,  
Webroot response = 0, Webroot total = 0, Anti-Spyware response = 0, Anti-  
Spyware total = 29; Latency = 7468; "05/Dec/2016:22:40:55 +0100" ]
```



Accesslogs: performance parameters (1)

| Perf Parameter Number | Log Parameter | Optional Performance Parameter Description |
|-----------------------|---------------|--|
| 1 | %:<a | Wait-time to receive the response from the Web Proxy authentication process, after the Web Proxy sent the request |
| 2 | %:<b | Wait-time to write request body to server after header |
| 3 | %:<d | Wait-time to receive the response from the Web Proxy DNS process, after the Web Proxy sent the request |
| 4 | %:<h | Wait-time to write request header to server after first byte |
| 5 | %:<r | Wait-time to receive the response from the Web Reputation Filters, after the Web Proxy sent the request |
| 6 | %:<s | Wait-time to receive the verdict from the Web Proxy anti-spyware process, after the Web Proxy sent the request |
| 7 | %:> | Wait-time for first response byte from server |
| 8 | %:>a | Wait-time to receive the response from the Web Proxy authentication process, including the time required for the Web Proxy to send the request |



Accesslogs: performance parameters (2)

| Perf Parameter Number | Log Parameter | Optional Performance Parameter Description |
|-----------------------|---------------|---|
| 9 | %:>b | Wait-time for complete response body after header received |
| 10 | %:>c | Time required for the Web Proxy to read a response from the disk cache |
| 11 | %:>d | Wait-time to receive the response from the Web Proxy DNS process, including the time required for the Web Proxy to send the request |
| 12 | %:>h | Wait-time for server header after first response byte |
| 13 | %:>r | Wait-time to receive the verdict from the Web Reputation Filters, including the time required for the Web Proxy to send the request |
| 14 | %:>s | Wait-time to receive the verdict from the Web Proxy anti-spyware process, including the time required for the Web Proxy to send the request |
| 15 | %:1< | Wait-time for first request byte from new client connection |
| 16 | %:1> | Wait-time for first byte written to client |



Accesslogs: performance parameters (3)

| Perf Parameter Number | Log Parameter | Optional Performance Parameter Description |
|-----------------------|---------------|--|
| 17 | :%b< | Wait-time for complete client body |
| 18 | :%b> | Wait-time for complete body written to client |
| 19 | :%h< | Wait-time for complete client header after first byte |
| 20 | :%h> | Wait-time for complete header written to client |
| 21 | :%m< | Wait-time to receive the verdict from the McAfee scanning engine, including the time required for the Web Proxy to send the request |
| 22 | :%m> | Wait-time to receive the response from the McAfee scanning engine, after the Web Proxy sent the request. |
| 23 | :%w< | Wait-time to receive the verdict from the Webroot scanning engine, including the time required for the Web Proxy to send the request |
| 24 | :%w> | Wait-time to receive the response from the Webroot scanning engine, after the Web Proxy sent the request |

Performance Monitoring & Final Thoughts

Monitoring & Log Visibility


Traditional Methods

- SNMP (use v3)
- SNMP Traps
- **Automation** (automation actions using CLI)

SIEM / Logs / External Reporting

- **SMA** – Centralized Reporting & Web Tracking
- **AWSR** (Advanced Web Security Reporting)
- Third Party Solutions

Monitoring Performance using Big Data Solutions?



- ELK
- Splunk, IBM QRadar
- AWS Kinesis
- Why not applying **Machine Learning** on top of your data?

SNMP – What should we monitor

WSA SNMP MIBs

- [AsyncOS Web MIB](#)
 - **WSA Performance MIB**
 - OID: 1.3.6.1.4.1.15497.1.2
- [AsyncOS Mail MIB for WSA](#)
 - Desc: Original MIB
 - OID: 1.3.6.1.4.1.15497.1.1.1
- [AsyncOS SMI MIB for WSA](#)
 - Top-Level SMI MIB

SNMP Traps

- Enable SNMP traps (WSA CLI > **snmpconfig**)
- Enable additionally **CPUUtilizationExceeded** (>85% or 90% of load)
- memoryUtilizationExceeded

WSA Performance related SNMP OIDS - Hardware

| OID | Description |
|--------------------------------|---|
| 1.3.6.1.4.1.15497.1.1.1.18.1.3 | raidID |
| 1.3.6.1.4.1.15497.1.1.1.18.1.2 | raidStatus <ul style="list-style-type: none">- driveHealthy (1)- driveFailure (2)- driveRebuild (3) |
| 1.3.6.1.4.1.15497.1.1.1.18.1.4 | raidLastError |
| 1.3.6.1.4.1.15497.1.1.1.10 | fanTable (shows fan names, rotation speeds) |
| 1.3.6.1.4.1.15497.1.1.1.9.1.2 | degreesCelsius – Temperature in degrees C |

Let us now map **status detail** command to SNMP OIDs (1)

| System Resource Utilization | Example Value |
|------------------------------------|---------------|
| System Resource Utilization | |
| CPU | 14.8% |
| RAM | 81.7% |
| Transactions per Second | |
| Average in last minute | 10 |
| Maximum in last hour | 150 |
| Average in last hour | 80 |
| Maximum since proxy restart | 250 |
| Average since proxy restart | 100 |

| SNMP OID | Name |
|------------------------------------|--------------------------|
| System Resource Utilization | |
| 1.3.6.1.4.1.15497.1.1.1.2.0 | perCentCPUUtilization |
| 1.3.6.1.4.1.15497.1.1.1.1.0 | perCentMemoryUtilization |
| Transactions per Second | |
| 1.3.6.1.4.1.15497.1.2.3.7.1.1.0 | cacheThruputNow |
| 1.3.6.1.4.1.15497.1.2.3.7.1.2.0 | cacheThruput1hrPeak |
| 1.3.6.1.4.1.15497.1.2.3.7.1.3.0 | cacheThruput1hrMean |
| 1.3.6.1.4.1.15497.1.2.3.7.1.8.0 | cacheThruputLifePeak |
| 1.3.6.1.4.1.15497.1.2.3.7.1.9.0 | cacheThruputLifeMean |

status detail to SNMP OIDs – Bandwidth (2)

| System Resource Utilization | Example Value | SNMP OID | Name / Description |
|-----------------------------|---------------|---------------------------------|---------------------------------|
| Bandwidth (Mbps) | | | |
| Average in last minute | 0.005 | 1.3.6.1.4.1.15497.1.2.3.7.4.1.0 | cacheBwidthTotalNow (Kbps) |
| Maximum in last hour | 2.145 | 1.3.6.1.4.1.15497.1.2.3.7.4.2.0 | cacheBwidthTotal1hrPeak (Kbps) |
| Average in last hour | 0.018 | 1.3.6.1.4.1.15497.1.2.3.7.4.3.0 | cacheBwidthTotal1hrMean (Kbps) |
| Maximum since proxy restart | 8.010 | 1.3.6.1.4.1.15497.1.2.3.7.4.8.0 | cacheBwidthTotalLifePeak (Kbps) |
| Average since proxy restart | 0.002 | 1.3.6.1.4.1.15497.1.2.3.7.4.9.0 | cacheBwidthTotalLifeMean |



status detail to SNMP OIDs – Response Time (3)

| System Resource Utilization | Example Value | SNMP OID | Name / Description |
|-----------------------------|---------------|---------------------------------|----------------------------|
| Response Time (ms) | | | |
| Average in last minute | 633 | 1.3.6.1.4.1.15497.1.2.3.7.9.1.0 | cacheTotalRespTimeNow |
| Maximum in last hour | 109214 | 1.3.6.1.4.1.15497.1.2.3.7.9.2.0 | cacheTotalRespTime1hrPeak |
| Average in last hour | 3666 | 1.3.6.1.4.1.15497.1.2.3.7.9.3.0 | cacheTotalRespTime1hrMean |
| Maximum since proxy restart | 109214 | 1.3.6.1.4.1.15497.1.2.3.7.9.8.0 | cacheTotalRespTimeLifePeak |
| Average since proxy restart | 199 | 1.3.6.1.4.1.15497.1.2.3.7.9.9.0 | cacheTotalRespTimeLifeMean |



status detail to SNMP OIDs – Cache Hit Rate (4)

| System Resource Utilization | Example Value | SNMP OID | Name / Description |
|-----------------------------|---------------|---------------------------------|--------------------|
| Cache Hit Rate | | | |
| Average in last minute | | 1.3.6.1.4.1.15497.1.2.3.7.5.1.0 | cacheHitsNow |
| Maximum in last hour | | 1.3.6.1.4.1.15497.1.2.3.7.5.2.0 | cacheHits1hrPeak |
| Average in last hour | | 1.3.6.1.4.1.15497.1.2.3.7.5.3.0 | cacheHits1hrMean |
| Maximum since proxy restart | | 1.3.6.1.4.1.15497.1.2.3.7.5.8.0 | cacheHitsLifePeak |
| Average since proxy restart | | 1.3.6.1.4.1.15497.1.2.3.7.5.9.0 | cacheHitsLifeMean |



status detail to SNMP OIDs – Connection (5)

| System Resource Utilization | Example Value | SNMP OID | Name / Description |
|-----------------------------|---------------|-------------------------------|-----------------------|
| Connections | | | |
| Idle client connections | | 1.3.6.1.4.1.15497.1.2.3.2.7.0 | cacheClientIdleConns |
| Idle server connections | | 1.3.6.1.4.1.15497.1.2.3.3.7.0 | cacheServerIdleConns |
| Total client connections | | 1.3.6.1.4.1.15497.1.2.3.2.8.0 | cacheClientTotalConns |
| Total server connections | | 1.3.6.1.4.1.15497.1.2.3.3.8.0 | cacheServerTotalConns |
| Average since proxy restart | | | |

Other SNMP OIDs to monitor (not a final list)

| OID | Description |
|-------------------------------|--|
| 1.3.6.1.4.1.15497.1.2.3.1.2.0 | cacheCpuUsage (proxLd) – how busy is prox process |
| 1.3.6.1.4.1.15497.1.2.3.1.4.0 | cacheUsedStoragePct |
| 1.3.6.1.4.1.15497.1.2.3.2 | proxyClientSidePerf - Group |
| 1.3.6.1.4.1.15497.1.2.3.3 | proxyServerSidePerf Group |
| | |
| | |
| | |

| OID | Description |
|---------------------------------|----------------------|
| 1.3.6.1.4.1.15497.1.2.3.5.1.1.1 | cacheMedianTime |
| 1.3.6.1.4.1.15497.1.2.3.5.1.1.2 | cacheHTTPClntSvcTime |
| 1.3.6.1.4.1.15497.1.2.3.5.1.1.3 | cacheHTTPMissSvcTime |
| 1.3.6.1.4.1.15497.1.2.3.5.1.1.4 | cacheHTTPHitSvcTime |
| 1.3.6.1.4.1.15497.1.2.3.5.1.1.5 | cacheHTTPSrvSvcTime |
| 1.3.6.1.4.1.15497.1.2.3.5.1.1.6 | cacheDnsSvcTime |
| 1.3.6.1.4.1.15497.1.2.3.5.1.1.7 | cacheHTTPSvcWaitTime |

Demo

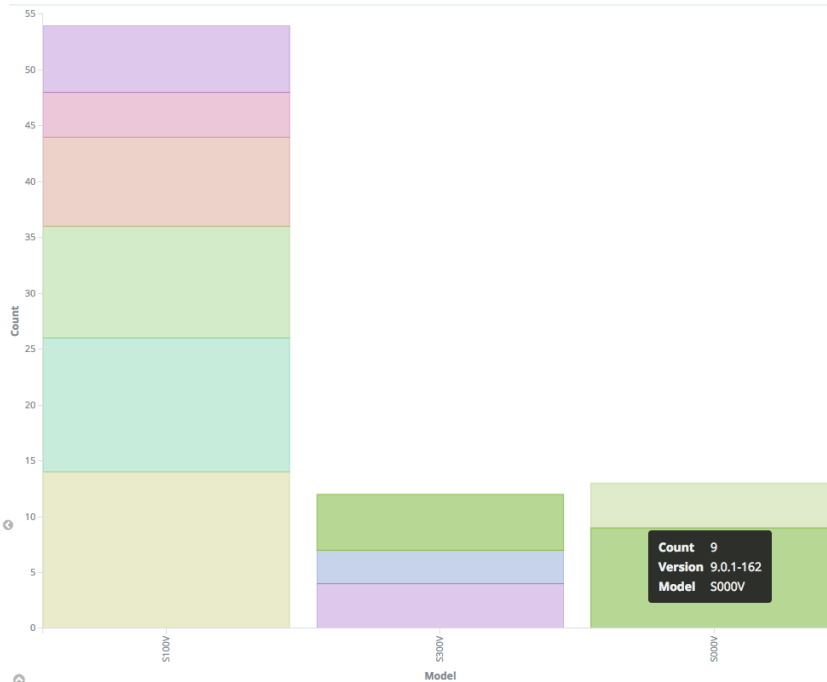

WSA Performance Monitoring using power of Logs and ELK Stack

Using power of Logstash / Elasticsearch & Kibana for WSA Log Management

ELK Stack = Elasticsearch, Logstash & Kibana

Powerful open-source solution that can be used for:

- Data (example logs) intake / processing / parsing & normalization / output – powerful data pipeline (Logstash)
- Scalable, HA, Multitenant, real-time data analytics, full text search engine (Elasticsearch)
- Visualization of data (Kibana)



ELK – WSA Performance Monitoring accesslogs

Agenda

- Introduction
- Understanding Cisco Web Security Appliance Pipeline
- Configuration Considerations and Best Practices
- Troubleshooting WSA Performance Issues
- Performance Monitoring & Final Thoughts
- Q & A

Cisco Spark




Questions?

Use Cisco Spark to communicate with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install Spark or go directly to the space
4. Enter messages/questions in the space



cs.co/ciscolivebot#BRKSEC-3303

- Please complete your Online Session Evaluations after each session
- Complete 4 Session Evaluations & the Overall Conference Evaluation (available from Thursday) to receive your Cisco Live T-shirt
- All surveys can be completed via the Cisco Live Mobile App or the Communication Stations

Don't forget: Cisco Live sessions will be available for viewing on-demand after the event at www.ciscolive.com/global/on-demand-library/.

Complete Your Online Session Evaluation



Continue Your Education


- Demos in the Cisco campus:
 - WSA, Cisco Umbrella
- Walk-in Self-Paced Labs
- Tech Circle
- Meet the Engineer 1:1 meetings
- Related sessions
 - BRKSEC-3015 - TLS Decryption on Cisco Security Devices



Thank you



You're



Cisco *live!*