CND Splunk Customer Success Story

Log Latency, So What?

Client

Public Sector, Government
United Kingdom, October 2023

Client Challenge:

The client operates a private cloud based VDI environment with approximately 900 virtual machines for end users, plus a small number of servers. These endpoints were then monitored with a Splunk UF installed on them and then all output sent to a single Heavy Forwarder (HF), this would then parse the logs and redact any sensitive data before forwarding onwards to Splunk Cloud (SaaS) where the SOC would monitor the logs for threat actor activity etc. The key challenge that the SOC were experiencing was log latency i.e. the delta between the log event being generated on an endpoint and then it reaching the Splunk Cloud instance; in the worst cases this could be up to 4 days.

Key Challenges

- Client's log latency taking up to 4 days to generate.
- Client left at risk of Threat Actor due to the longevity of each log.
- Client lacked a Deployment Server resulting in lack of ability to make required changed to configurations.

Log Latency, so what?

This excessive latency provides a window of potential opportunity for a Threat Actor (TA) to operate within, if logs take between 4hrs and 4 days at worst to reach the SOC then a TA could operate undetected within this time to detonate malware.

Key Results

- Once new HF hardware installed, immediate results seen with every core utilised.
- Large backlog of events streamed forward reducing timings from 4 hours to 7 seconds.
- Client knowledge and understanding increased leaving them prepared for future trouble-spotting.

What Fixes Have Been Tried?

The client architecture team were adamant that their HF was not the cause of the ingestion latency, they cited evidence of average CPU load, RAM utilisation and had informed multiple Splunk consultants who had already assessed this environment. They were unwilling to add resources to the HF. They also lacked a Deployment Server and thus the ability to make changes to configurations, they had to wait until a new formal build image had been generated and pushed to all endpoints which could take several weeks.

Splunk x CND

CND Splunk Customer Success Story

Log Latency, So What?

Client

Public Sector, Government
United Kingdom, October 2023

Approach:

Starting at the log sources on the endpoints, I noted that on the UF's under the following stanza, a low output limit was in place:

```
#limits.conf
[thruput]
maxKBps = 128
* If specified and not zero, this limits the speed through the thruput processor to the specified
rate in kilobytes per second.
```

I persuaded the client to uplift this to 2048 or 2MB/s. This equates to a 16x increase in bandwidth, which could in some environments have unforeseen consequences, however being entirely cloud there was unlikely to be a contention issue on the network and it would help ensure that any of the busier servers were able to immediately despatch logs at scale and remove any backlog.

The HF:

The HF was significantly under-specification at a 4 CPU cores, it should be at 12 physical / 24vCPU and 12GB RAM. Whilst this is a trivial item to identify and any certified administrator could recognise this, the value add here was persuading the client who was unwilling to change their infra and satisfied that it was right-sized. By engaging with the client in a constructive and crucially experienced manner I was able to get them to agree to uplifting the hardware. I also encouraged them to deploy further instances so that their resilience would improve from a single HF to three HF's and employ standard Splunk load-balancing.

The Uplift:

The new HF hardware was ready and the additional vCPU allocated (36vCPU) in their cloud platform, after restarting SplunkD and monitoring the linux CPU usage via 'htop' we could immediately see every core being utilised. The RAM had been set much higher at 60GB and within 30 seconds this had reached the allocated limit, in-fact it started using swap disk and crashed initially. The HF had been critically under-resourced. This had been a client mis-understanding on the number of cores versus the workload requirements. We tuned the hardware settings further and within an hour a huge backlog of events flooded towards their SOC. We then left it over the weekend and their improved hardware lead to ingest latency reducing from 4 days to approximately 7 seconds.

Conclusion

Working with experienced Splunk PS resources who have worked in infrastructure teams for many years before retraining to be a Splunk consultant allows us to talk the same language as the client. This promotes trust and gives opportunities for them to really take on our advice and accept proposed alterations. In this case the client was thrilled at the changes and we went onwards in future engagements to fix other challenging items which have significantly improved their effectiveness at detecting nefarious activity. A simple fix but a huge client impact and onwards Splunk success in their environment.

Splunk x CND