

HTTP Event Collector in Splunk 6.5 More Super Powers!

Itay Neeman

Director of Engineering, Splunk

Shakeel Mohamed

Software Engineer, Splunk

.conf2016

splunk >

Disclaimer

During the course of this presentation, we may make forward looking statements regarding future events or the expected performance of the company. We caution you that such statements reflect our current expectations and estimates based on factors currently known to us and that actual events or results could differ materially. For important factors that may cause actual results to differ from those contained in our forward-looking statements, please review our filings with the SEC. The forward-looking statements made in the this presentation are being made as of the time and date of its live presentation. If reviewed after its live presentation, this presentation may not contain current or accurate information. We do not assume any obligation to update any forward looking statements we may make. In addition, any information about our roadmap outlines our general product direction and is subject to change at any time without notice. It is for informational purposes only and shall not, be incorporated into any contract or other commitment. Splunk undertakes no obligation either to develop the features or functionality described or to include any such feature or functionality in a future release.

What the HEC?


A new token-based JSON API for events

Send events *directly* from anywhere
(servers, mobile devices, IOT)


Easy to configure / works out of the box.

Easy to secure

Highly performant, scalable and
available



In 6.5, HEC returns with even more powers!





In Splunk 6.3:

HEC only accepts data in our JSON event format

HEC RAW

HEC RAW

You can now send data in arbitrary formats to HEC!

Useful for integration with existing systems that can send data over HTTP

HEC RAW – send data in "any" format

```
curl -k https://localhost:8088/services/collector/raw?  
channel=F35D2F91-A751-4823-BDB9-482415B42DD1 -H  
'Authorization: Splunk 46931F1C-352C-4DF6-820C-F2689CF88494'  
-d '09-26-2016 12:20 PST Look MA, no JSON'
```


HEC Raw – override metadata

```
curl -k "https://localhost:8088/services/collector/raw?  
channel=F35D2F91-A751-4823-BDB9-482415B42DD1  
&source=mysource&sourcetype=sourcetype1" -H  
'Authorization: Splunk 46931F1C-352C-4DF6-820C-F2689CF88494'  
-d '09-26-2016 12:20 PST Look MA, no JSON'
```

HEC RAW

- Timestamp extraction and line breaking rules run
- Events must be bounded to a single request / a single event cannot span multiple requests
- Each request must have a channel
- You can override metadata per request

In Splunk 6.3:

Clients do not know if events have been indexed



Indexer ACK

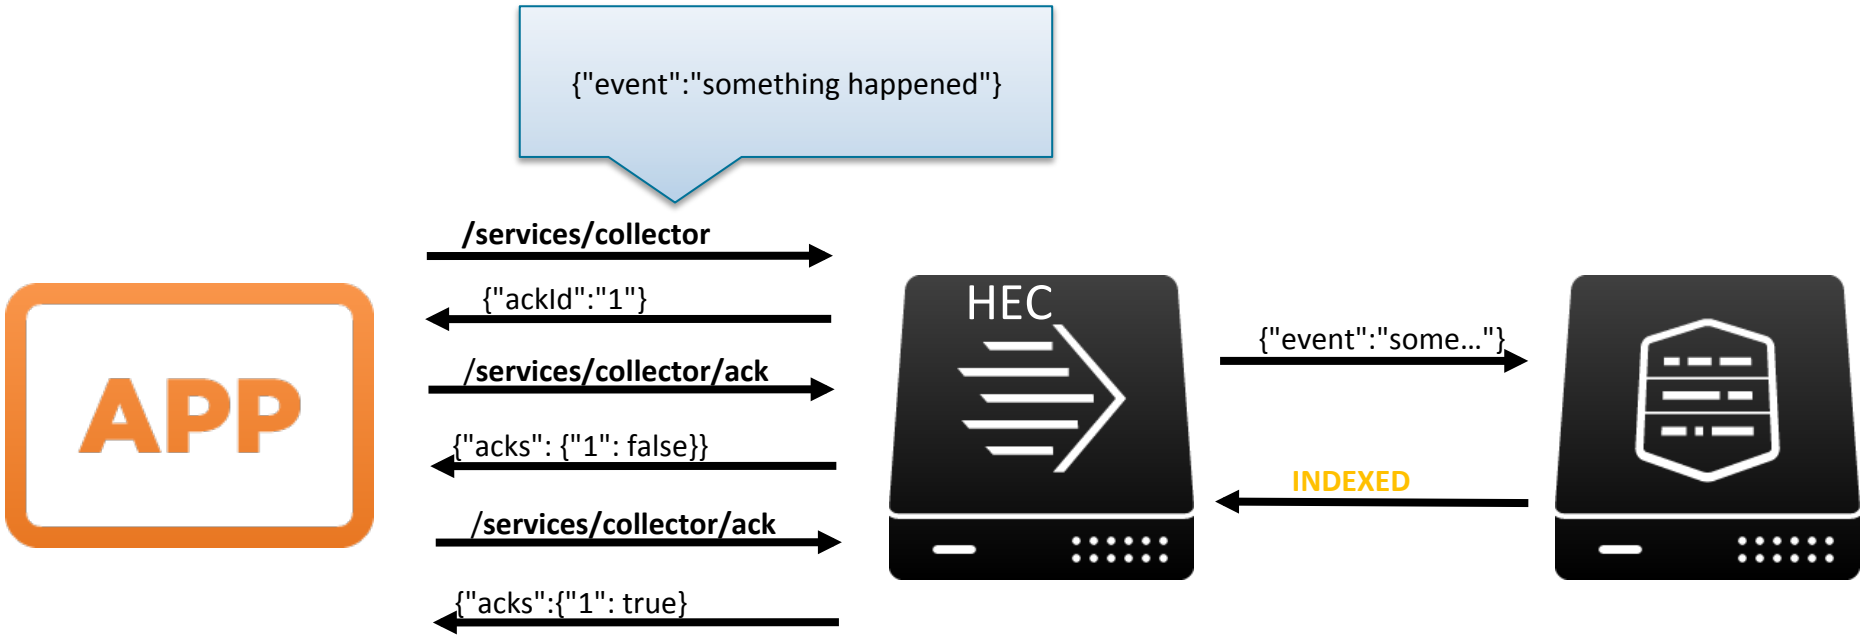
Indexer ACK

You can now receive an acknowledgment when events are indexed*

Useful for reducing data loss during outages

*Flexibility comes at a cost

Indexer ACK



Indexer ACK – sending an event

```
curl -k https://localhost:8088/services/collector?  
channel=F35D2F91-A751-4823-BDB9-482415B42DD1 -H  
'Authorization: Splunk 46931F1C-352C-4DF6-820C-F2689CF88494'  
-d '{"event":"something happened"}'
```

```
{"ackId":1}
```


Indexer ACK – polling ack IDs

```
curl -k https://localhost:8088/services/collector/ack?  
channel=F35D2F91-A751-4823-BDB9-482415B42DD1 -H  
'Authorization: Splunk 46931F1C-352C-4DF6-820C-F2689CF88494'  
-d '{"acks": [1]}'
```

```
{"acks": {"1": true}}
```


Indexer ACK

- ACK is enabled per token
- Supports raw & JSON endpoints
- Each event returns an ACK ID on submission. The ID corresponds to all events within that request
- Client must poll the ACK endpoint to check if there if the event(s) were acknowledged
- Each request for sending events or checking ACK must have a channel
- ACKs are released from memory after they have been received



In Splunk 6.3:

Only regex extractions are supported

Index

field enhancements

Index field enhancements

JSON Field extraction finally works for HEC!

You can now specify additional index fields separate from the "event" payload! (JSON endpoint only)

Useful for search time perf improvements as well as for supplying event custom metadata

Index field enhancements

```
curl -k https://localhost:8088/services/collector? -H  
'Authorization: Splunk 46931F1C-352C-4DF6-820C-F2689CF88494'  
-d '{"event": "something happened", "fields": {"severity": "INFO",  
"category": ["foo", "bar"]}}'
```

Index field enhancements

- You can use sourcetypes with JSON index field extraction enabled
- You can supply arbitrary key-value-pairs to be indexed using "fields"
- Values can only be strings, numerics will be supported in the future
- Array values are supported

In Splunk 6.3



Clients must set the auth header to the Splunk token

Basic Auth

Basic Auth

You can now use Basic Auth to authenticate your HEC requests

Useful for systems that do not support custom auth header values but can support Basic Auth via the URI (Github web hooks)

Basic Auth

```
curl -u x:46931F1C-352C-4DF6-820C-F2689CF88494  
-k https://localhost:8088/services/collector?  
-d '{"event":"something happened"}'
```

Basic Auth

- You can now authenticate to HEC using Basic Auth
- Username can be anything
- Password is the token



Today

Only authenticated clients can send to HEC

Query String Auth

(FUTURE)

Query String Auth

You will be able to set the token via the query string

Useful for web hooks / systems that only allow setting a URI / have no header support. (Atlassian)

Query String Auth

```
curl -k https://localhost:8088/services/collector?  
token=46931F1C-352C-4DF6-820C-F2689CF88494  
-d '{"event": "something happened"}'
```


Query String Auth

- Will allow passing the token as a "token" param in the Query String
- Will be disabled by default
- Will be configurable per token

A lot of new goodness for HEC in 6.5!

- RAW
- ACK
- Index field enhancements
- Basic Auth

And more!



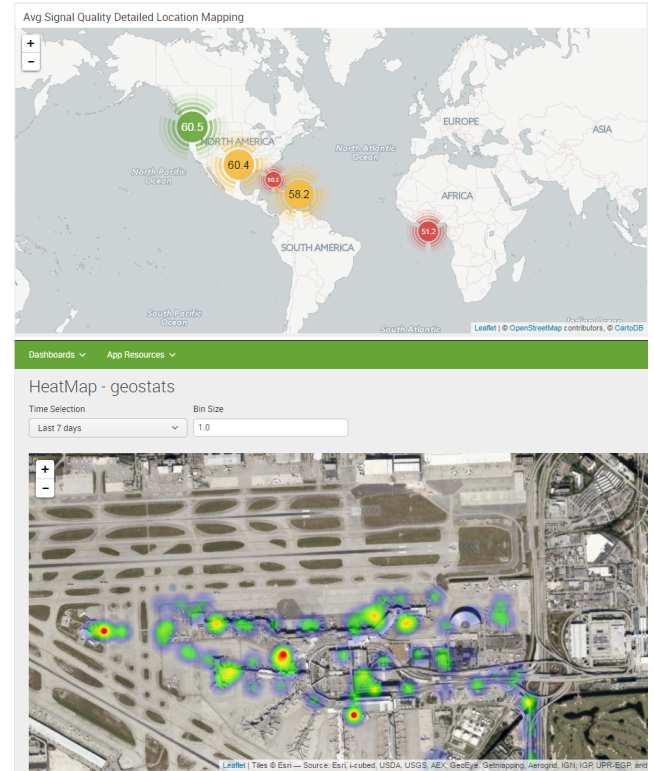
Partner success!



.conf2016

NetMotion Diagnostics®

- Mobile Connectivity Intelligence
 - Location-based connectivity, coverage and end-to-end root-cause analysis of performance, failures and bottlenecks between mobile devices, wireless networks, and cloud or on-premises applications
 - Field service organizations with thousands of devices collecting and transmitting data around the clock



NETMOTION DIAGNOSTICS®

Before

- Using syslog messages to pump data into Splunk

After

- Switched to using Splunk HEC with Splunk Logging for .NET
- Worked with Splunk on enhancements, including sending a PR (Open Source Rocks!)
- 5x+ more events per second, especially with millions of events!

Learn more!

http://splk.it/new_HEC

THANK YOU

.conf2016

