splunk> .conf2017

Running Enterprise Security at Capacity Accurately Thanks to Data Model Acceleration

Gabriel Vasseur, PhD | Senior Cyber Security Analyst

September 2017 | Washington, DC

Forward-Looking Statements

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 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.

Splunk, Splunk>, Listen to Your Data, The Engine for Machine Data, Splunk Cloud, Splunk Light and SPL are trademarks and registered trademarks of Splunk Inc. in the United States and other countries. All other brand names, product names, or trademarks belong to their respective owners.

By attending this talk, you agree to owe a beer to the speaker.

Who is This Guy?

- French
- Lives in England
- Works for THALES UK
- PhD in theoretical physics
- 10 years in the IT security industry
- Currently
  - On paper: Senior Cyber Security Analyst
  - In reality: Resident Data Scientist / Splunk Guru
- Spoke at .conf2016:

https://conf.splunk.com/files/2016/recordings/become-a-regular-expressions-ninja-and-unlock-your-splunk-potential.mp4

What is He Talking About?

1. Our **story** and why we're here giving this talk
2. The best introduction to **Data Models** ever!
3. Dive into the world of **DM Acceleration**
4. The **Big Picture**: accuracy, load and Data Model acceleration
5. How to **optimize** DM Acceleration

What is This Talk Not About?

- tstats tutorial
- Enterprise Security (ES)
- ITSI
- Spelunking
- ...

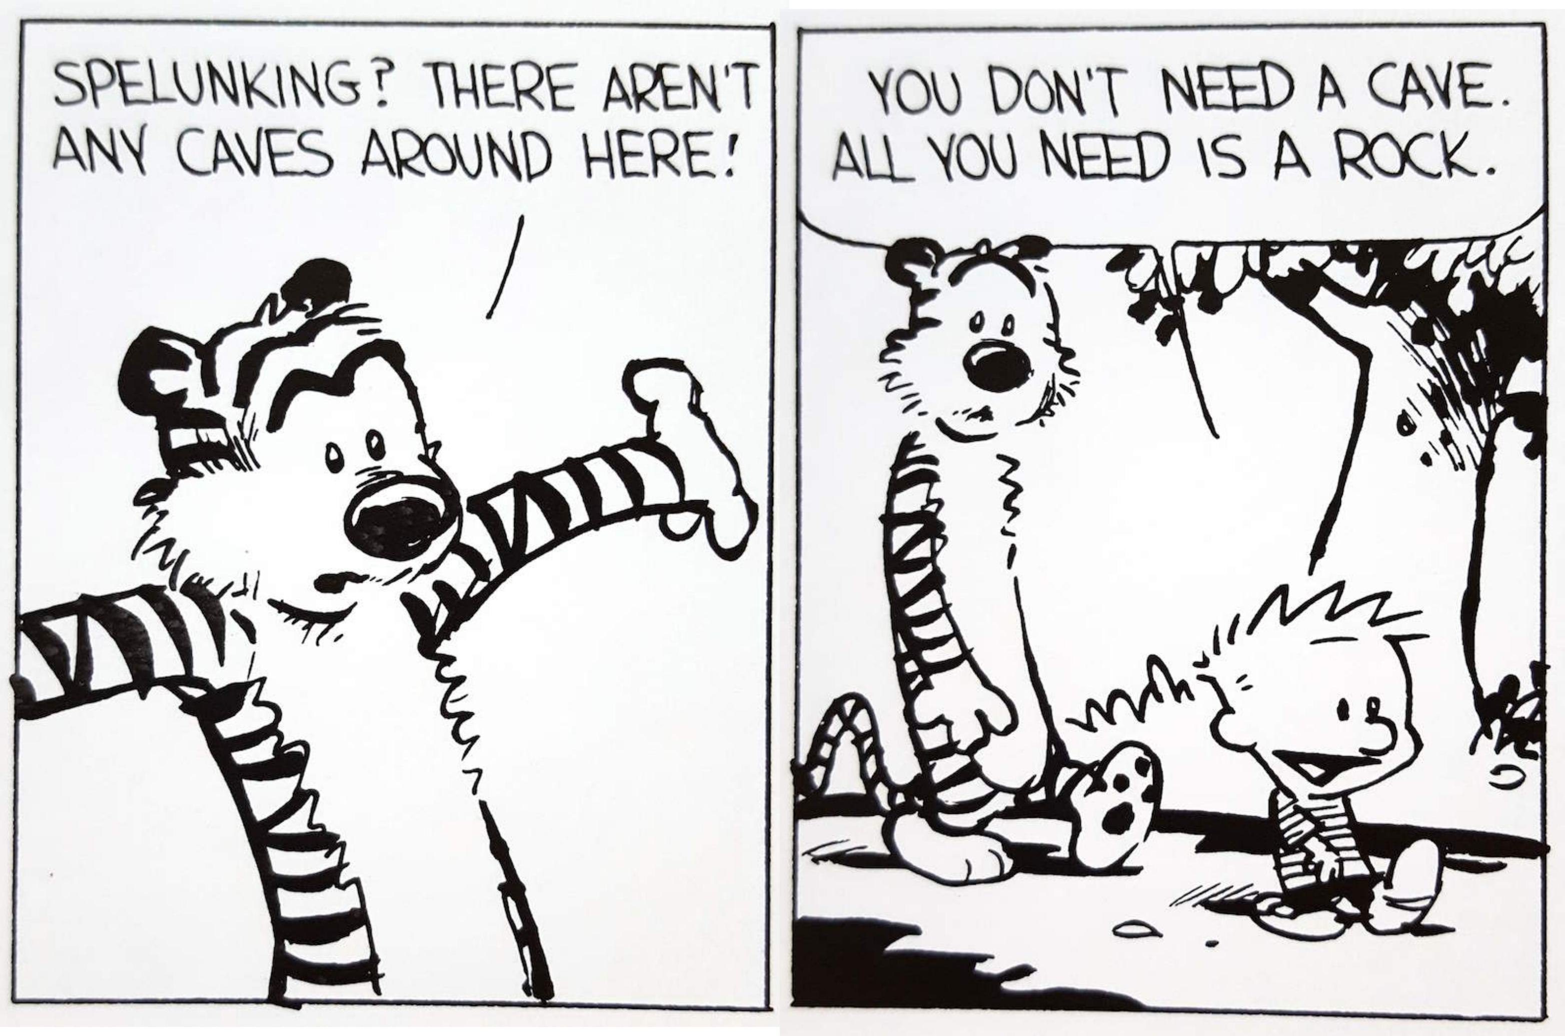
SPELUNKING? THERE AREN'T ANY CAVES AROUND HERE!
YOU DON'T NEED A CAVE. ALL YOU NEED IS A ROCK.

Calvin & Hobbes

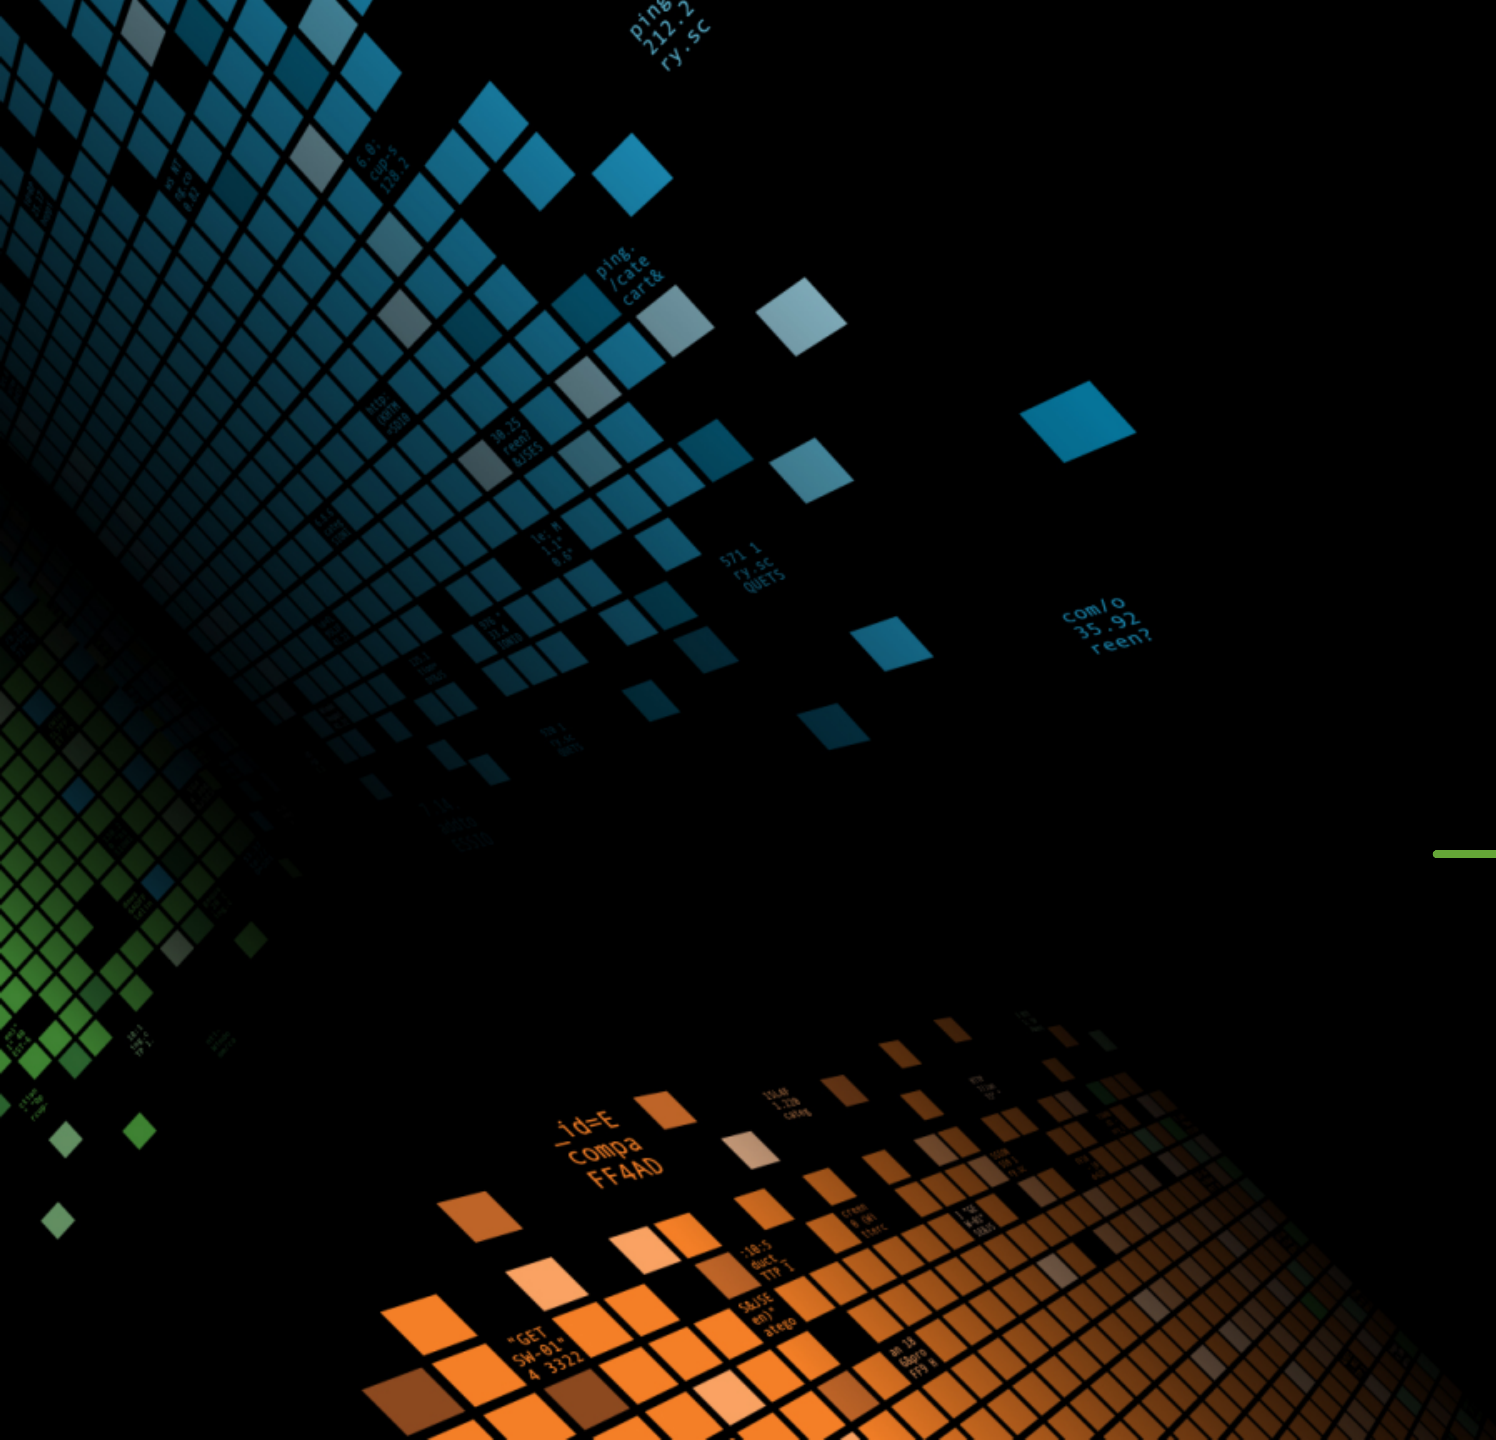
ping
212.2
ry.sc
ping.
/cate
cart&
com/o
35.92
reen?
_id=E
compa
FF4AD
"GET
SW-01"
4 3322

A Love Story

Company Meets Splunk

(2013)

- Assesses data inputs (~50 GB daily) and use cases (Security, ~10 users)
- Looks at minimum requirements for Splunk 6.0.1 ES 3.0 (12 cores & 12GB RAM!)
- Buys **over-spec'd** hardware (16 cores & 32GB RAM)

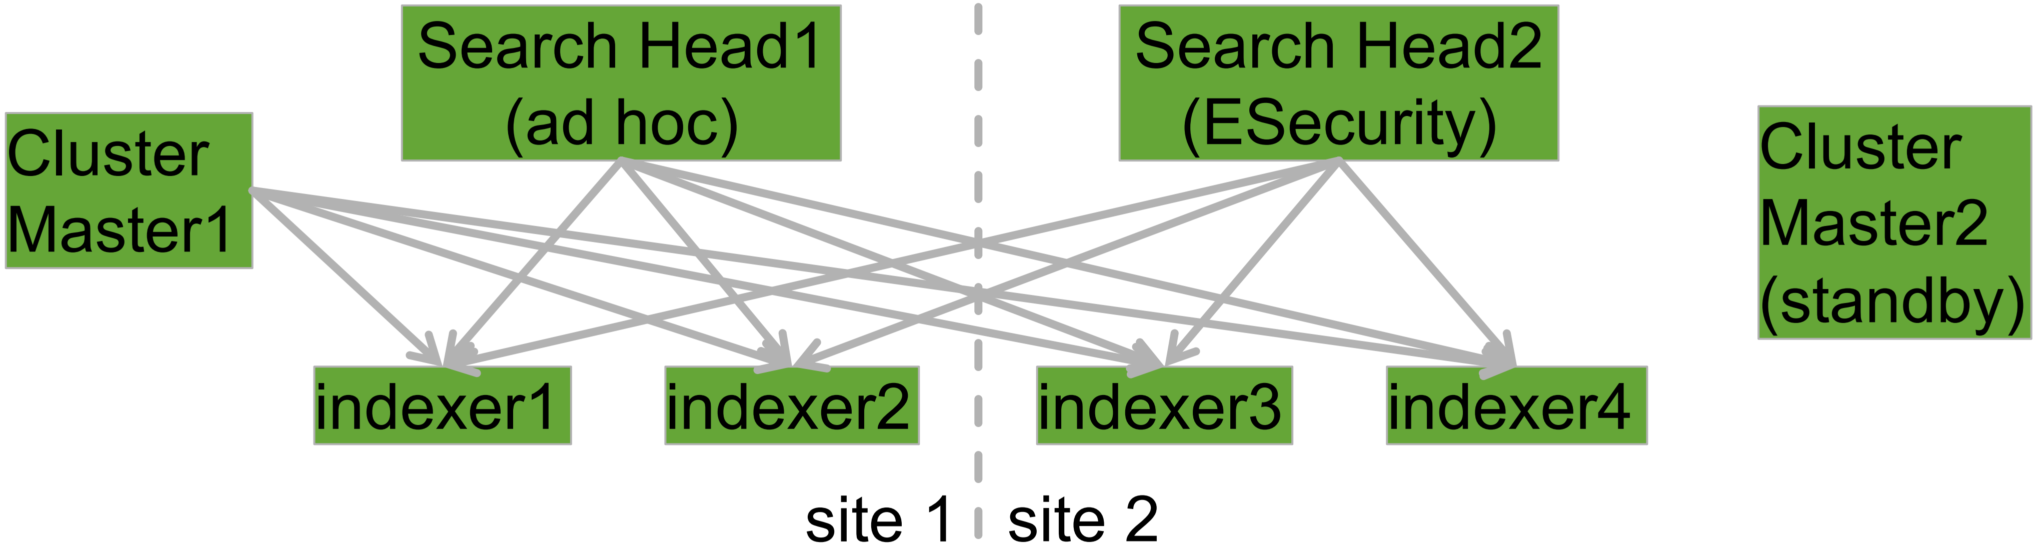
Search Head1 (ad hoc)

Search Head2 (ESecurity)

Cluster Master1

Cluster Master2 (standby)

indexer1

indexer2

indexer3

indexer4

site 1

site 2

- Honeymoon period: the relationship feels fresh and exciting!

Trouble

(2013 -> 2016)

- Many upgrades later (Splunk 6.0 ES 3.0 to Splunk 6.4 ES 4.4), minimum requirements are 16 cores & 32GB RAM, hardware is no longer over spec'd
- More data has been on-boarded (~150GB daily)
- More use cases / searches / alerts have been developed
- Best practices not necessarily followed...

- Splunk gives and gives, but reaches its limits:
  - Skipped scheduled searches
  - Laggy & crashy data model acceleration
  - Risk of false negatives...

Money Doesn't Buy Happiness

(But it helps)

- Throw hardware at the problem and scale up:
  - Indexer cluster (add more indexers)
  - Search head cluster (a bit complex with ES)
  - More RAM
  - SSDs?

- But...
  - Can't afford it and/or will take too long, need a fix NOW
  - Most importantly, will it be enough? Do we really understand what affects accuracy?

Bring in the Counselors

- Health check with Splunk PS
  - Kick-start ideas for tune-up
- Long-term Splunk PS engagement
  - Implemented many improvements
  - Made data model acceleration really key to managing load

This talk is the summary of what we learned

Disclaimer

- This is what works for us, your mileage may vary!

What Is A Data Model?

“A Data Model is a Hierarchically Structured Search-time Mapping of Semantic Knowledge About One or More Datasets.”

Splunk docs

Life Without Data Models

- `index=endpoints sourcetype=calvin:AV`

```
25-sep-2017 4pm OPS threat definition updated to revision 1986
25-sep-2017 5pm DETECTION path=hobbes.exe verdict=BAD threat=water.balloon
```

_time	log_type	path	verdict	threat
Sep-25 4pm	OPS			
Sep-25 5pm	DETECTION	hobbes.exe	BAD	water.balloon

- `index=servers sourcetype=hobbes:malwarebuster`

```
2017/09/25T6pm loglevel=alert state=infected virusname=baseball-cheat path=calvin.exe
2017/09/25T7pm loglevel=alert state=infected boot sector is compromised
```

_time	loglevel	state	virusname	path
Sep-25 6pm	alert	infected	baseball-cheat	calvin.exe
Sep-25 7pm	alert	infected		

The Malware Data Model

As part of the Common Information Model (CIM) all malware events should:

- Be tagged with "malware" and "attack"
- Have a "signature" field
- Have an "action" field
- "action" must be either "allowed" or "blocked"
- ... etc ...

Now, let's write some props and transforms for each vendor to make their data CIM-compliant...

"Please excuse the crudity of this model"

Achieving CIM Compliance

- `index=endpoints sourcetype=calvin:AV`

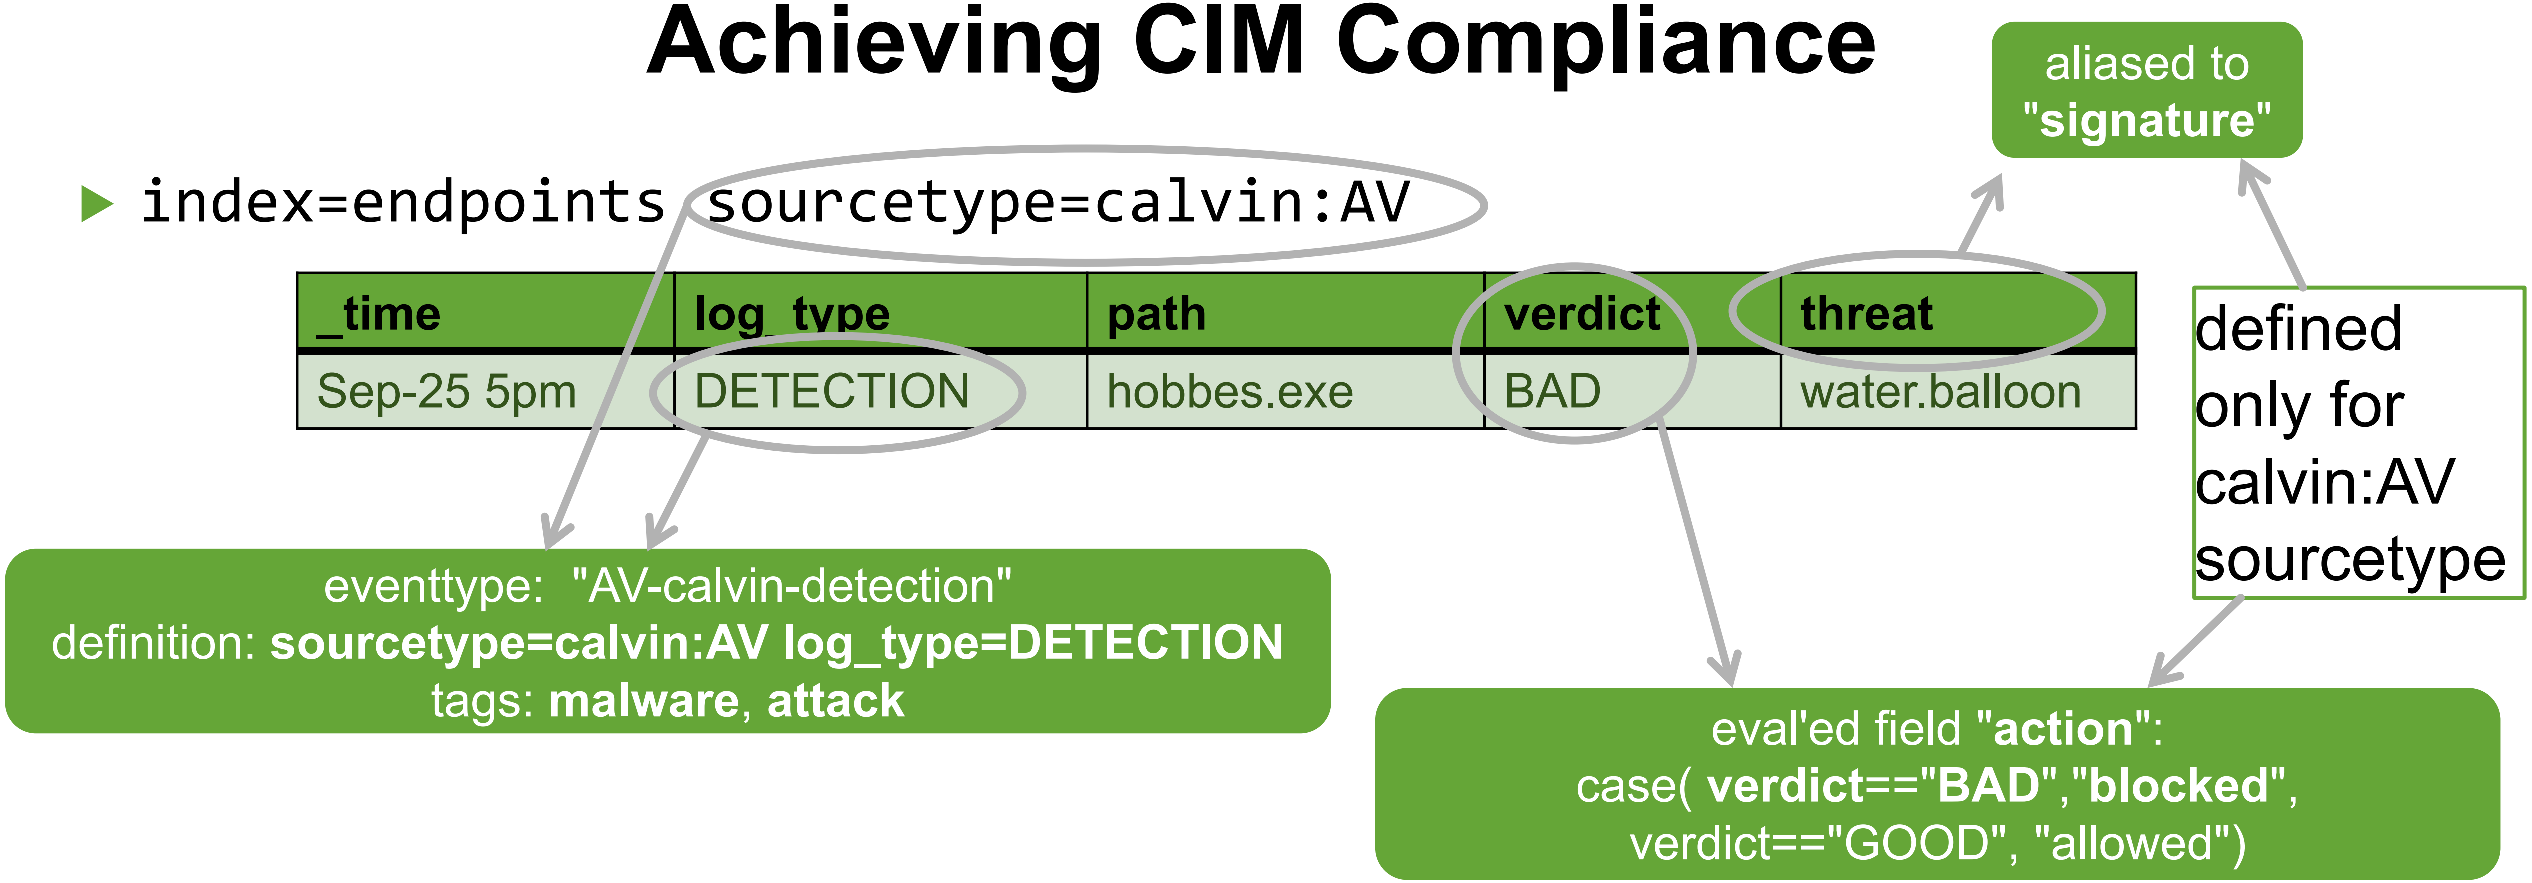

_time	log_type	path	verdict	threat
Sep-25 5pm	DETECTION	hobbes.exe	BAD	water.balloon

aliased to **"signature"**

defined only for calvin:AV sourcetype

eventtype: "AV-calvin-detection"
definition: **sourcetype=calvin:AV log_type=DETECTION**
tags: **malware**, **attack**

eval'ed field "**action**":
case(**verdict=="BAD"**,**"blocked"**,
verdict=="GOOD", "allowed")

- We package this configuration in an "Technical Add-on" (TA) called TA-calvinAV
- Or we get it ready made from Splunkbase
- We do something similar for the Hobbes AV sourcetype

Life With CIM Compliant Data

- `index=endpoints sourcetype=calvin:AV`

```
25-sep-2017 4pm OPS threat definition updated to revision 1986
25-sep-2017 5pm DETECTION path=hobbes.exe verdict=BAD threat=water.balloon
```

_time	log_type	path	verdict	threat	action	signature	tag
Sep-25 4pm	OPS						malware
Sep-25 5pm	DETECTION	hobbes.exe	BAD	water.balloon	blocked	water.balloon	malware, attack

- `index=servers sourcetype=hobbes:malwarebuster`

```
2017/09/25T6pm loglevel=alert state=infected virusname=baseball-cheat path=calvin.exe
2017/09/25T7pm loglevel=alert state=infected boot sector is compromised
```

_time	loglevel	state	virusname	path	action	signature	tag
Sep-25 6pm	alert	infected	baseball-cheat	calvin.exe	blocked	baseball-cheat	malware, attack
Sep-25 7pm	alert	infected			blocked		malware, attack

Life With CIM Compliant Data

- `tag=malware tag=attack`

```
25-sep-2017 5pm DETECTION path=hobbes.exe verdict=BAD threat=water.balloon
2017/09/25T6pm loglevel=alert state=infected virusname=baseball-cheat object=calvin.exe
2017/09/25T7pm loglevel=alert state=infected boot sector is compromised
```

_time	log_type	loglevel	path	state	verdict	virusname	threat	action	signature	tag
Sep-25 5pm	DETECTION		hobbes.exe		BAD		water.balloon	blocked	water.balloon	malware, attack
Sep-25 6pm		alert	calvin.exe	infected		baseball-cheat		blocked	baseball-cheat	malware, attack
Sep-25 7pm		alert		infected				blocked		malware, attack

- CIM = abstraction layer = you can do vendor-agnostic searches

Life With CIM Compliant Data

- `tag=malware tag=attack action=blocked | stats count by signature`

signature	count
baseball-cheat	1
water.balloon	1

- Let's take it to the next level...

The Malware Data Model

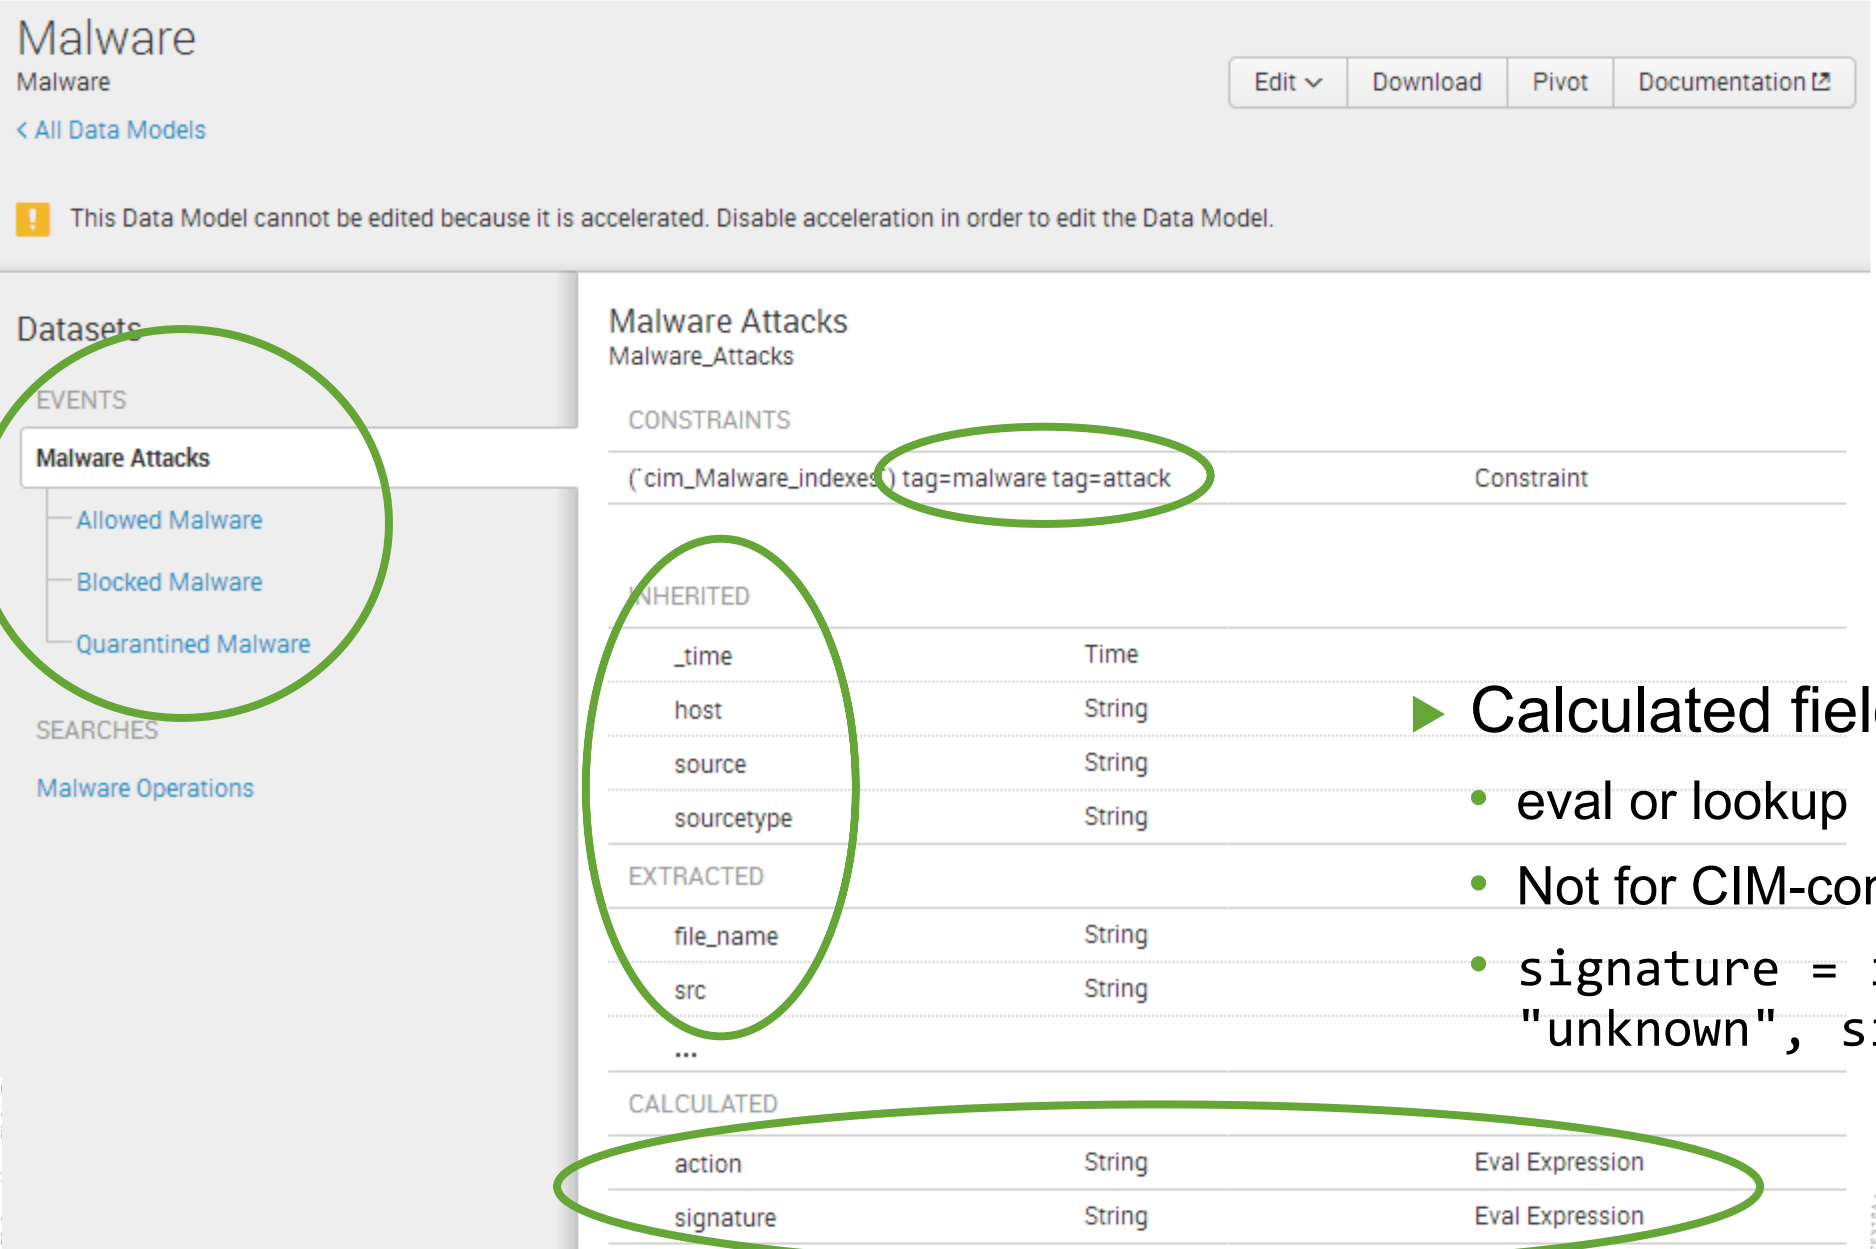

Malware

Malware

< All Data Models

! This Data Model cannot be edited because it is accelerated. Disable acceleration in order to edit the Data Model.

Datasets

EVENTS

- Malware Attacks
  - Allowed Malware
  - Blocked Malware
  - Quarantined Malware

SEARCHES

- Malware Operations

Malware Attacks

Malware_Attacks

CONSTRAINTS		
(`cim_Malware_indexes`) tag=malware tag=attack		Constraint
INHERITED		
_time	Time	
host	String	
source	String	
sourcetype	String	
EXTRACTED		
file_name	String	
src	String	
...		
CALCULATED		
action	String	Eval Expression
signature	String	Eval Expression

- Calculated fields
  - eval or lookup
  - Not for CIM-compliance!
  - `signature = if( isnull(signature), "unknown", signature)`

The datamodel Command

```
| datamodel Malware Malware_Attacks search
```

```
25-sep-2017 5pm DETECTION path=hobbes.exe verdict=BAD threat=water.balloon
2017/09/25T6pm loglevel=alert state=infected virusname=baseball-cheat object=calvin.exe
2017/09/25T7pm loglevel=alert state=infected boot sector is compromised
```

_time	...	...	path	state	verdict	virusname	threat	Malware_Attacks.action	Malware_Attacks.signature	Malware_Attacks.tag
Sep-25 5pm	...	...	hobbes.exe		BAD		water.balloon	blocked	water.balloon	malware, attack
Sep-25 6pm	...	...	calvin.exe	infected		baseball-cheat		blocked	baseball-cheat	malware, attack
Sep-25 7pm	...	...		Infected				blocked	unknown	malware, attack

The datamodel Command

- `| datamodel Malware Malware_Attacks search | search Malware_Attacks.action=blocked | stats count by Malware_Attacks.signature`

Malware_Attacks.signature	count
baseball-cheat	1
water.balloon	1
unknown	1

- Constraint and vendor stuff is abstracted away
- Still access to raw event and all fields

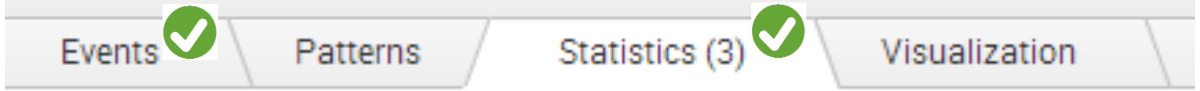

- Freshly brewed: uses latest search-time configuration and DM definition
- Doesn't benefit from DM acceleration :-(
- Note: datamodel command to be replaced with from command

splunk> .conf2017

The tstats Command

```
| tstats count from datamodel=Malware where
Malware_Attacks.action=blocked by Malware_Attacks.signature
```

Malware_Attacks.signature	count
baseball-cheat	1
water.balloon	1
unknown	1

Does the same as the datamodel command BUT:

- No access to _raw or any non-DM fields

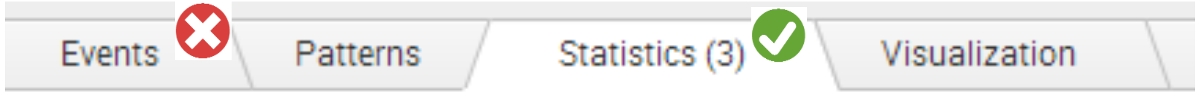

- Stats oriented, no way to just get a table of all the events one by one
- **Will benefit from DM acceleration!!!**
- Can also be used outside DM with indexed fields

The Pivot Command & Datasets UI

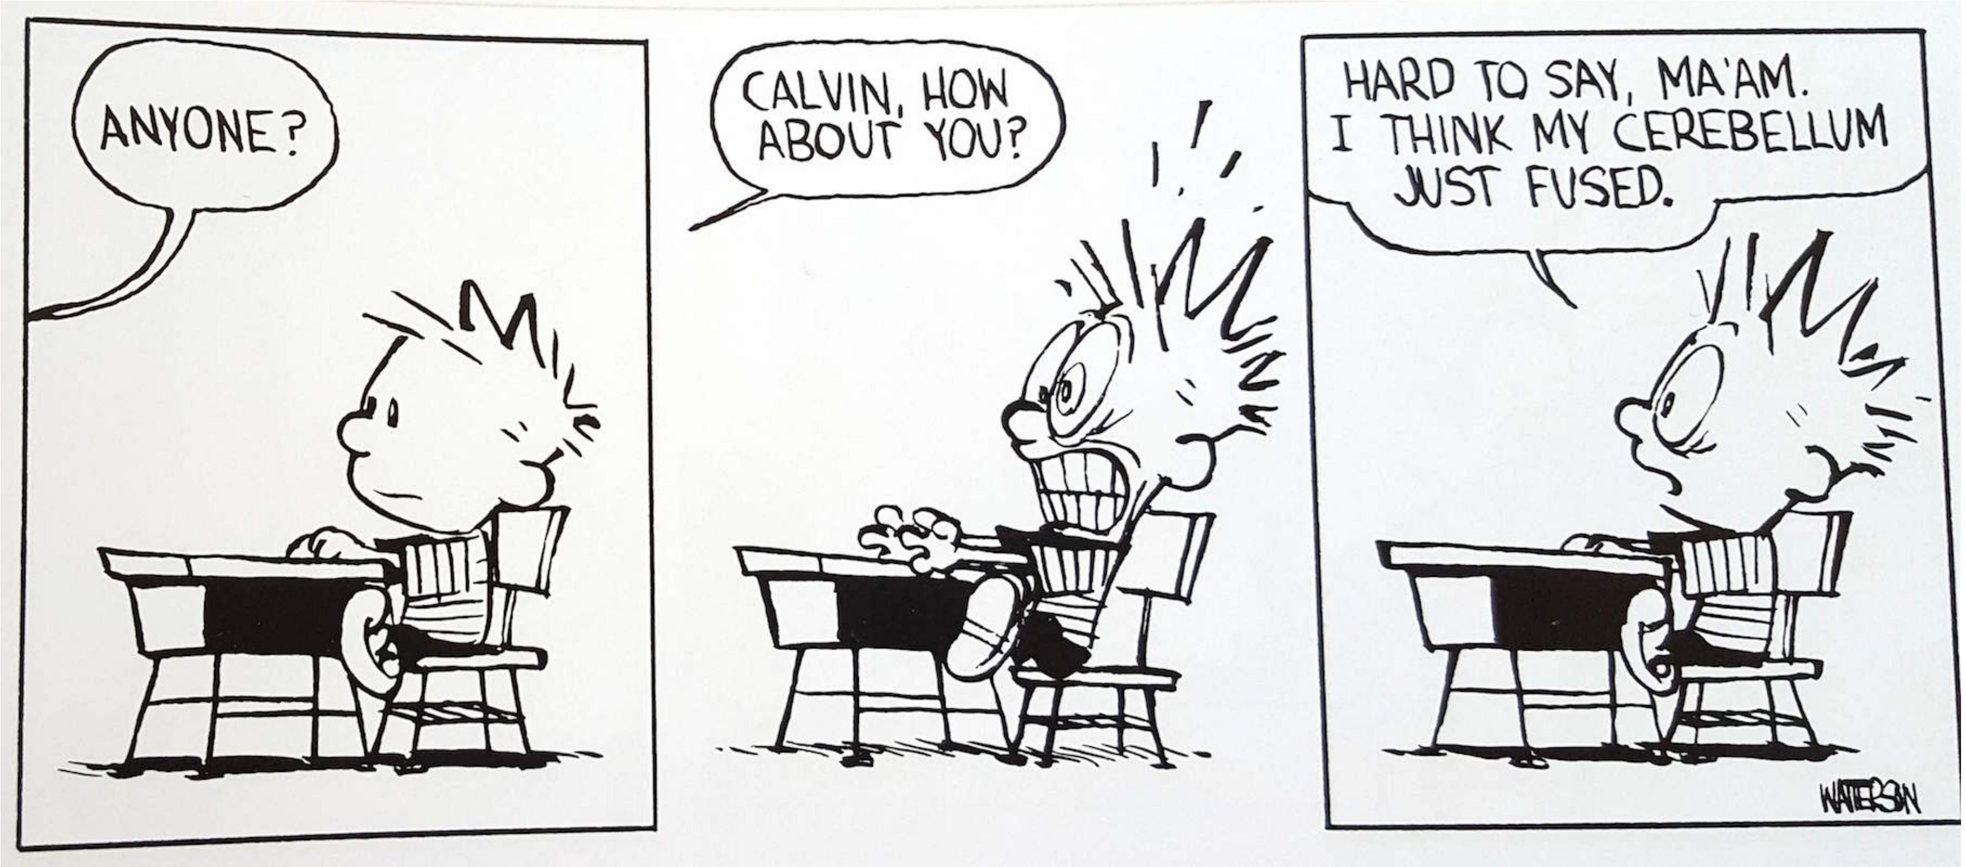

- Don't ask me!
- It's powered by tstats...
- ...so go learn tstats!

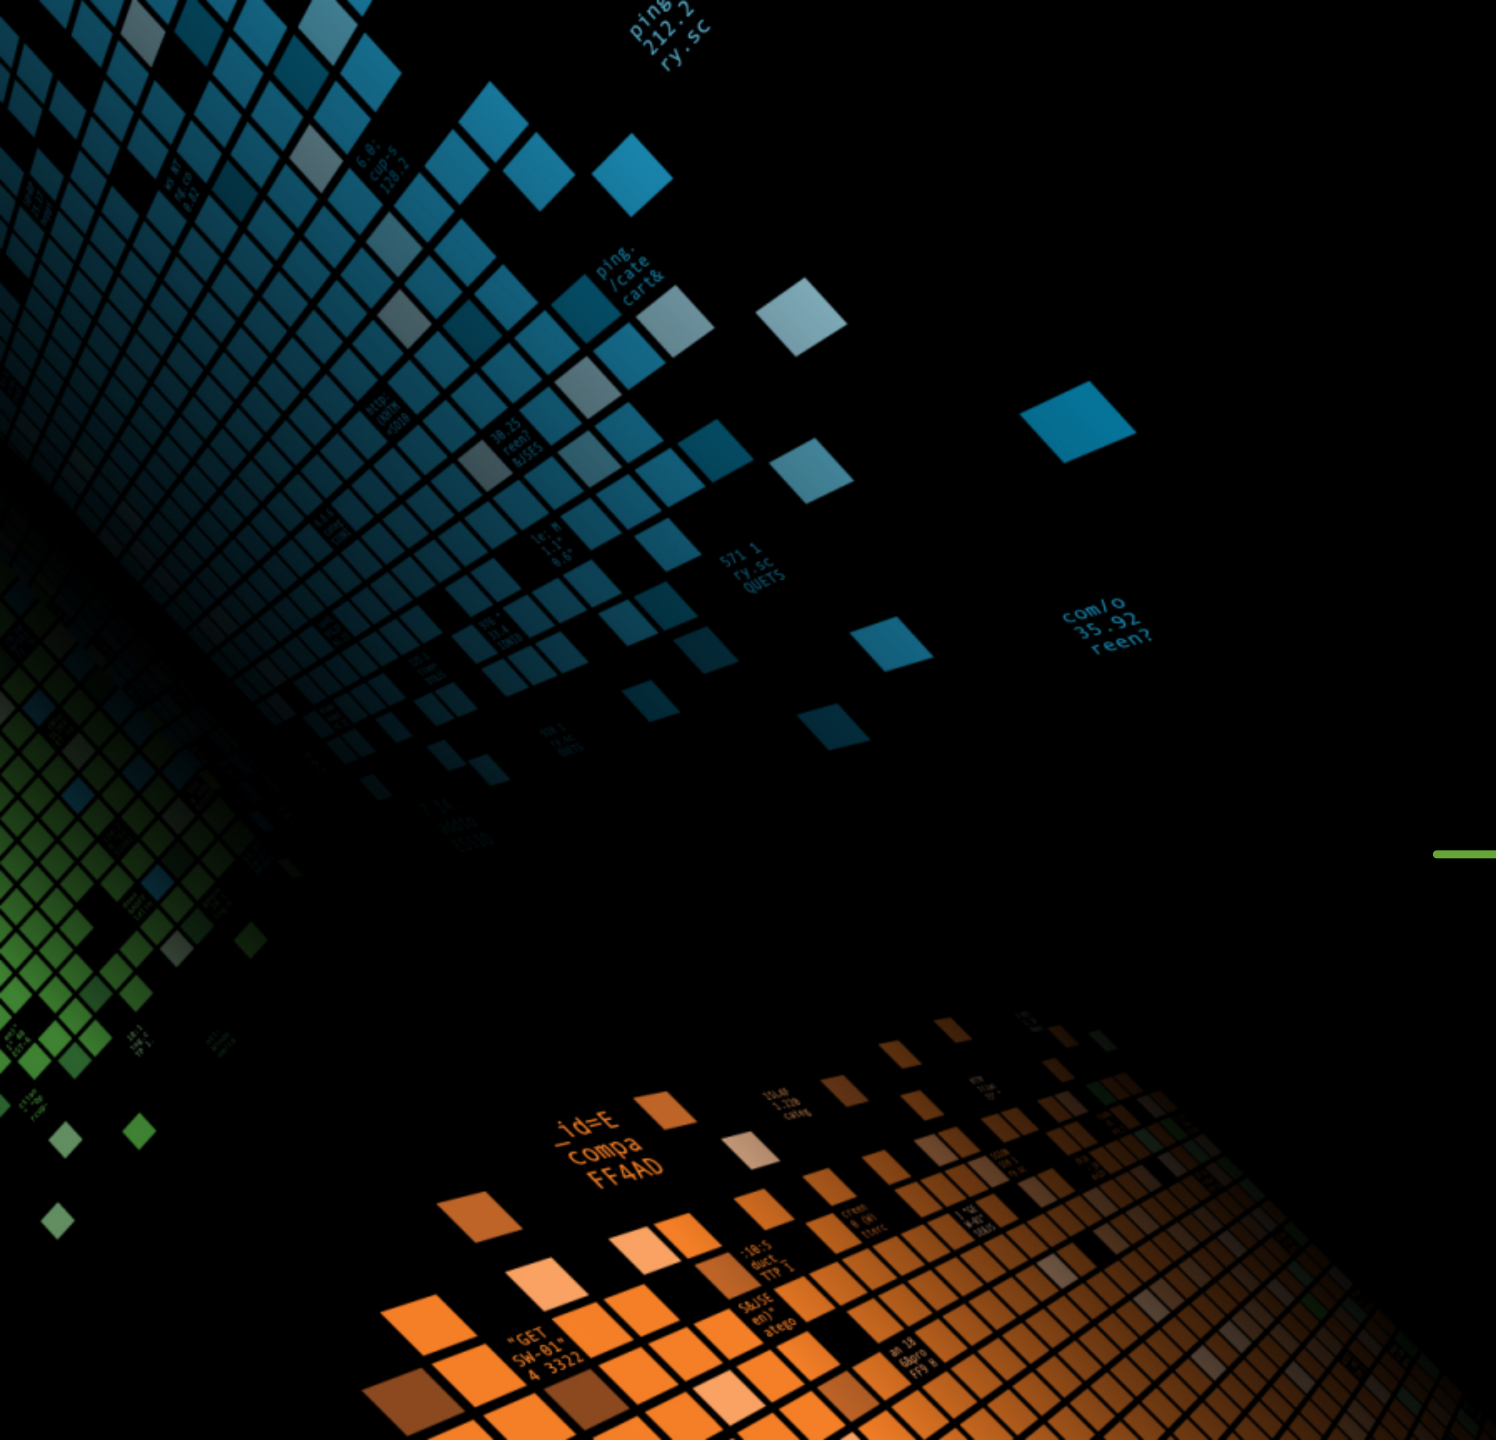

Accelerating Data Models

Data Model Acceleration

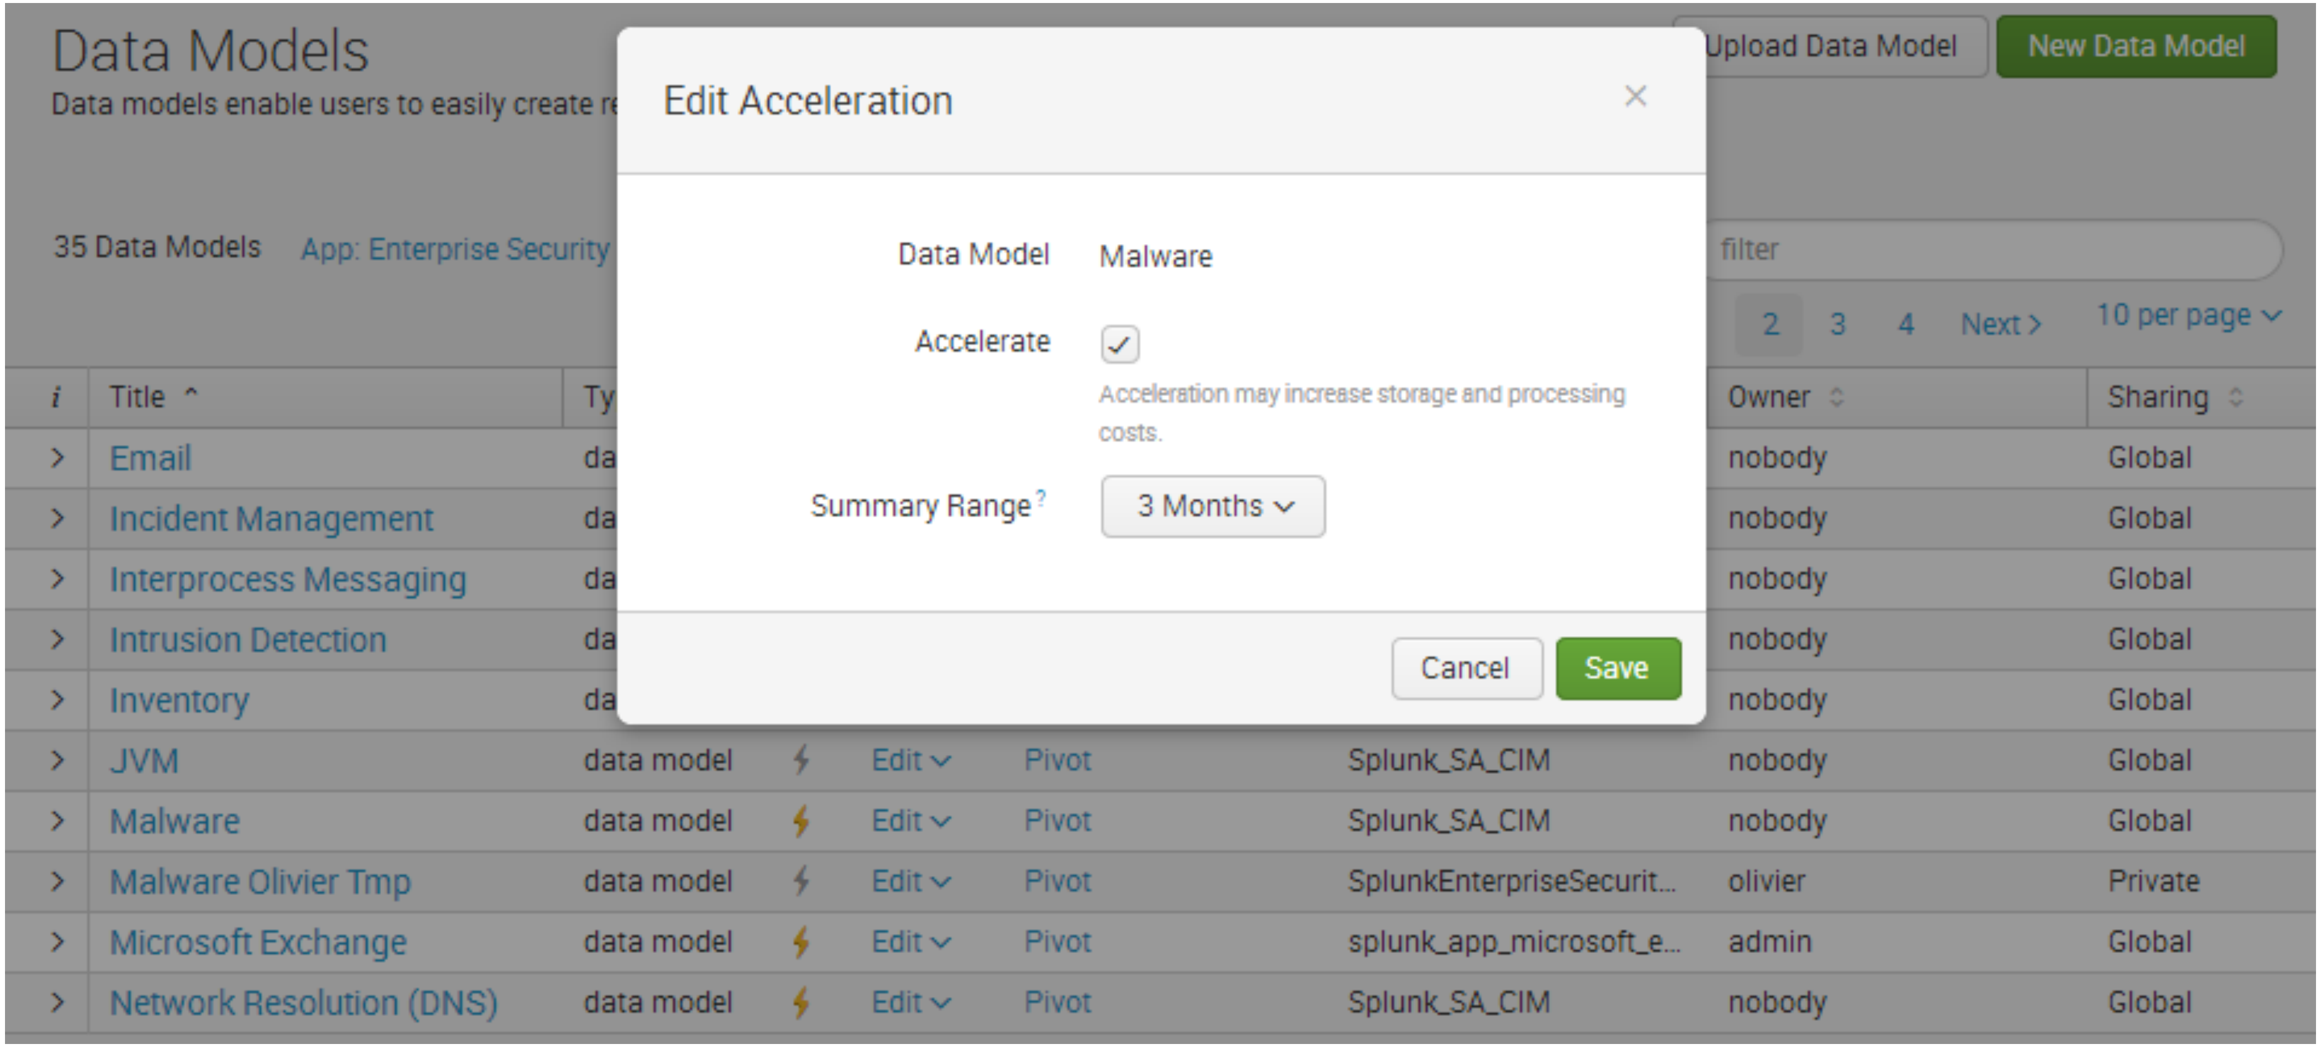

Data Models

Data models enable users to easily create r

Upload Data Model | New Data Model

35 Data Models   App: Enterprise Security

filter

2   3   4   Next >   10 per page

i	Title ^	Ty					Owner	Sharing
>	Email	da					nobody	Global
>	Incident Management	da					nobody	Global
>	Interprocess Messaging	da					nobody	Global
>	Intrusion Detection	da					nobody	Global
>	Inventory	da					nobody	Global
>	JVM	data model	⚡	Edit	Pivot	Splunk_SA_CIM	nobody	Global
>	Malware	data model	⚡	Edit	Pivot	Splunk_SA_CIM	nobody	Global
>	Malware Olivier Tmp	data model	⚡	Edit	Pivot	SplunkEnterpriseSecurit...	olivier	Private
>	Microsoft Exchange	data model	⚡	Edit	Pivot	splunk_app_microsoft_e...	admin	Global
>	Network Resolution (DNS)	data model	⚡	Edit	Pivot	Splunk_SA_CIM	nobody	Global

Edit Acceleration ×

Data Model: Malware

Accelerate: ☑

Acceleration may increase storage and processing costs.

Summary Range ?: 3 Months

Cancel   Save

splunk> .conf2017

How DM Acceleration Works

- Remember this search?
- `| datamodel Malware Malware_Attacks search`

```
25-sep-2017 5pm DETECTION path=hobbes.exe verdict=BAD threat=water.balloon
2017/09/25T6pm loglevel=alert state=infected virusname=baseball-cheat object=calvin.exe
2017/09/25T7pm loglevel=alert state=infected boot sector is compromised
```

_time	...	...	path	state	verdict	virusname	threat	Malware_Attacks.action	Malware_Attacks.signature	Malware_Attacks.tag
Sep-25 5pm	...	...	hobbes.exe		BAD		water.balloon	blocked	water.balloon	malware, attack
Sep-25 6pm	...	...	calvin.exe	infected		baseball-cheat		blocked	baseball-cheat	malware, attack
Sep-25 7pm	...	...		Infected				blocked	unknown	malware, attack

How DM Acceleration Works

- Something similar now runs every 5 minutes
- Any non-DM fields are filtered out

_time	Malware_Attacks.action	Malware_Attacks.signature	Malware_Attacks.tag
Sep-25 5pm	blocked	water.balloon	malware attack
Sep-25 6pm	blocked	baseball-cheat	malware attack
Sep-25 7pm	blocked	unknown	malware attack

- The results are shoved into the **Data Model summary**
- Earliest/latest cover whole retention period
- Only freshly indexed data is considered (late-arriving data included!)
- DM summary = extra .tsidx files in buckets on the indexer
- No _raw included

Data Model Acceleration

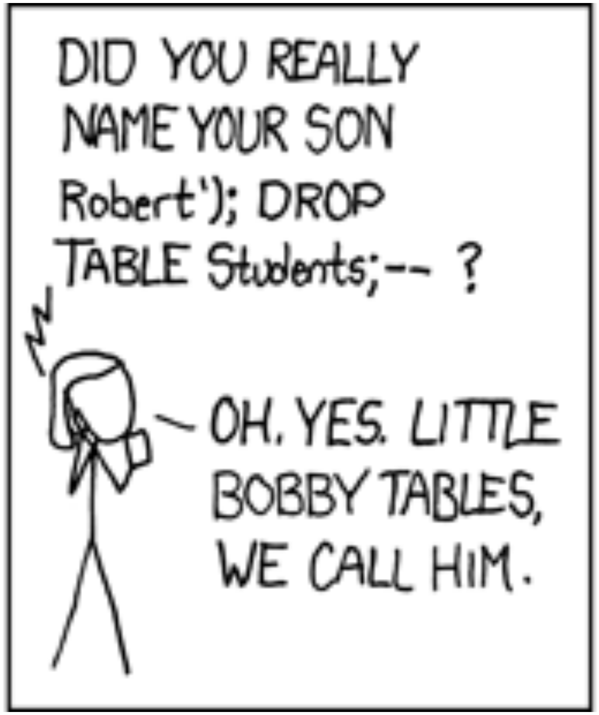
DID YOU REALLY NAME YOUR SON Robert'); DROP TABLE Students;-- ?
OH. YES. LITTLE BOBBY TABLES, WE CALL HIM.

- Think of the DM summary as a traditional DB table:
  - DM definition = schema (list of columns)
  - Constraint is run as a search, each result = one row
- Except:
  - Implementation is really not like a DB table...!
  - No way to just list rows/events one by one

- With DM acceleration, tstats command is automagically much faster!
- **Best of both worlds:**
  - Late binding schema (not much set in stone at index-time, flexible search-time configuration)
  - Speed of rigid schema with accelerated DM

tstats summariesonly=t

- tstats automatically use DM summaries if they are available
- The last few minutes won't be, most likely
- Doesn't sound like much, but impact is big
- Use **summariesonly=t**
- Careful:
  - Make sure the retention period covers your earliest/latest
  - If you have a time picker in a dashboard, make use you tell your users!

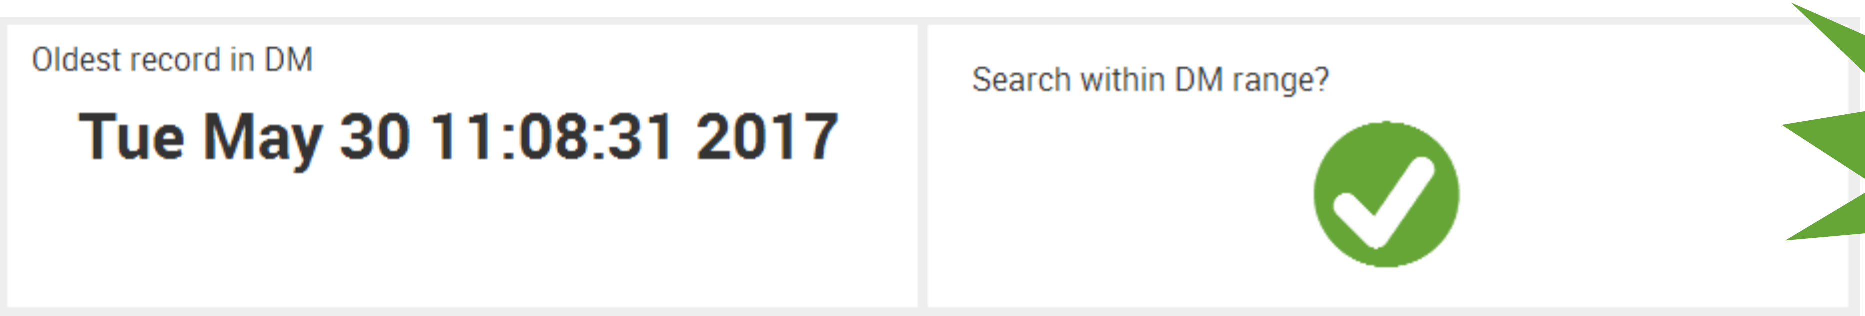

Oldest record in DM

Tue May 30 11:08:31 2017

Search within DM range?

✓

  - Make sure the DM summary is fully populated (backfilled + no lag)

Dashboard Source included!

ES's Data Model Audit Dashboard

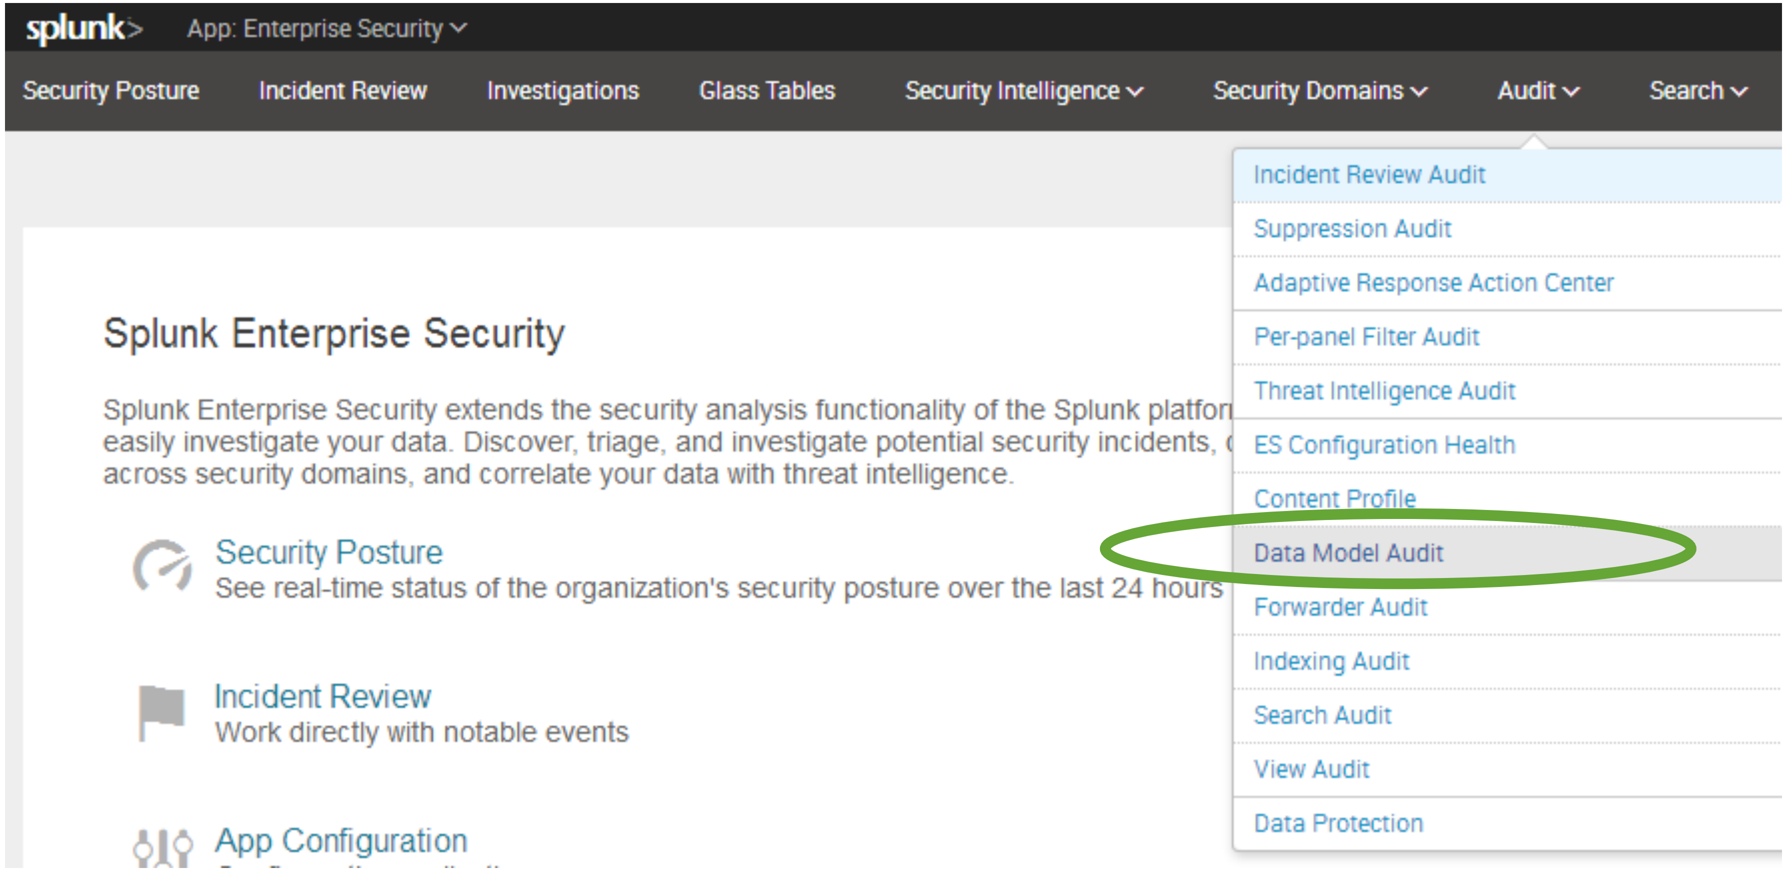

Splunk Enterprise Security

Splunk Enterprise Security extends the security analysis functionality of the Splunk platfor easily investigate your data. Discover, triage, and investigate potential security incidents, across security domains, and correlate your data with threat intelligence.

Security Posture
See real-time status of the organization's security posture over the last 24 hours

Incident Review
Work directly with notable events

App Configuration

ES's Data Model Audit Dashboard

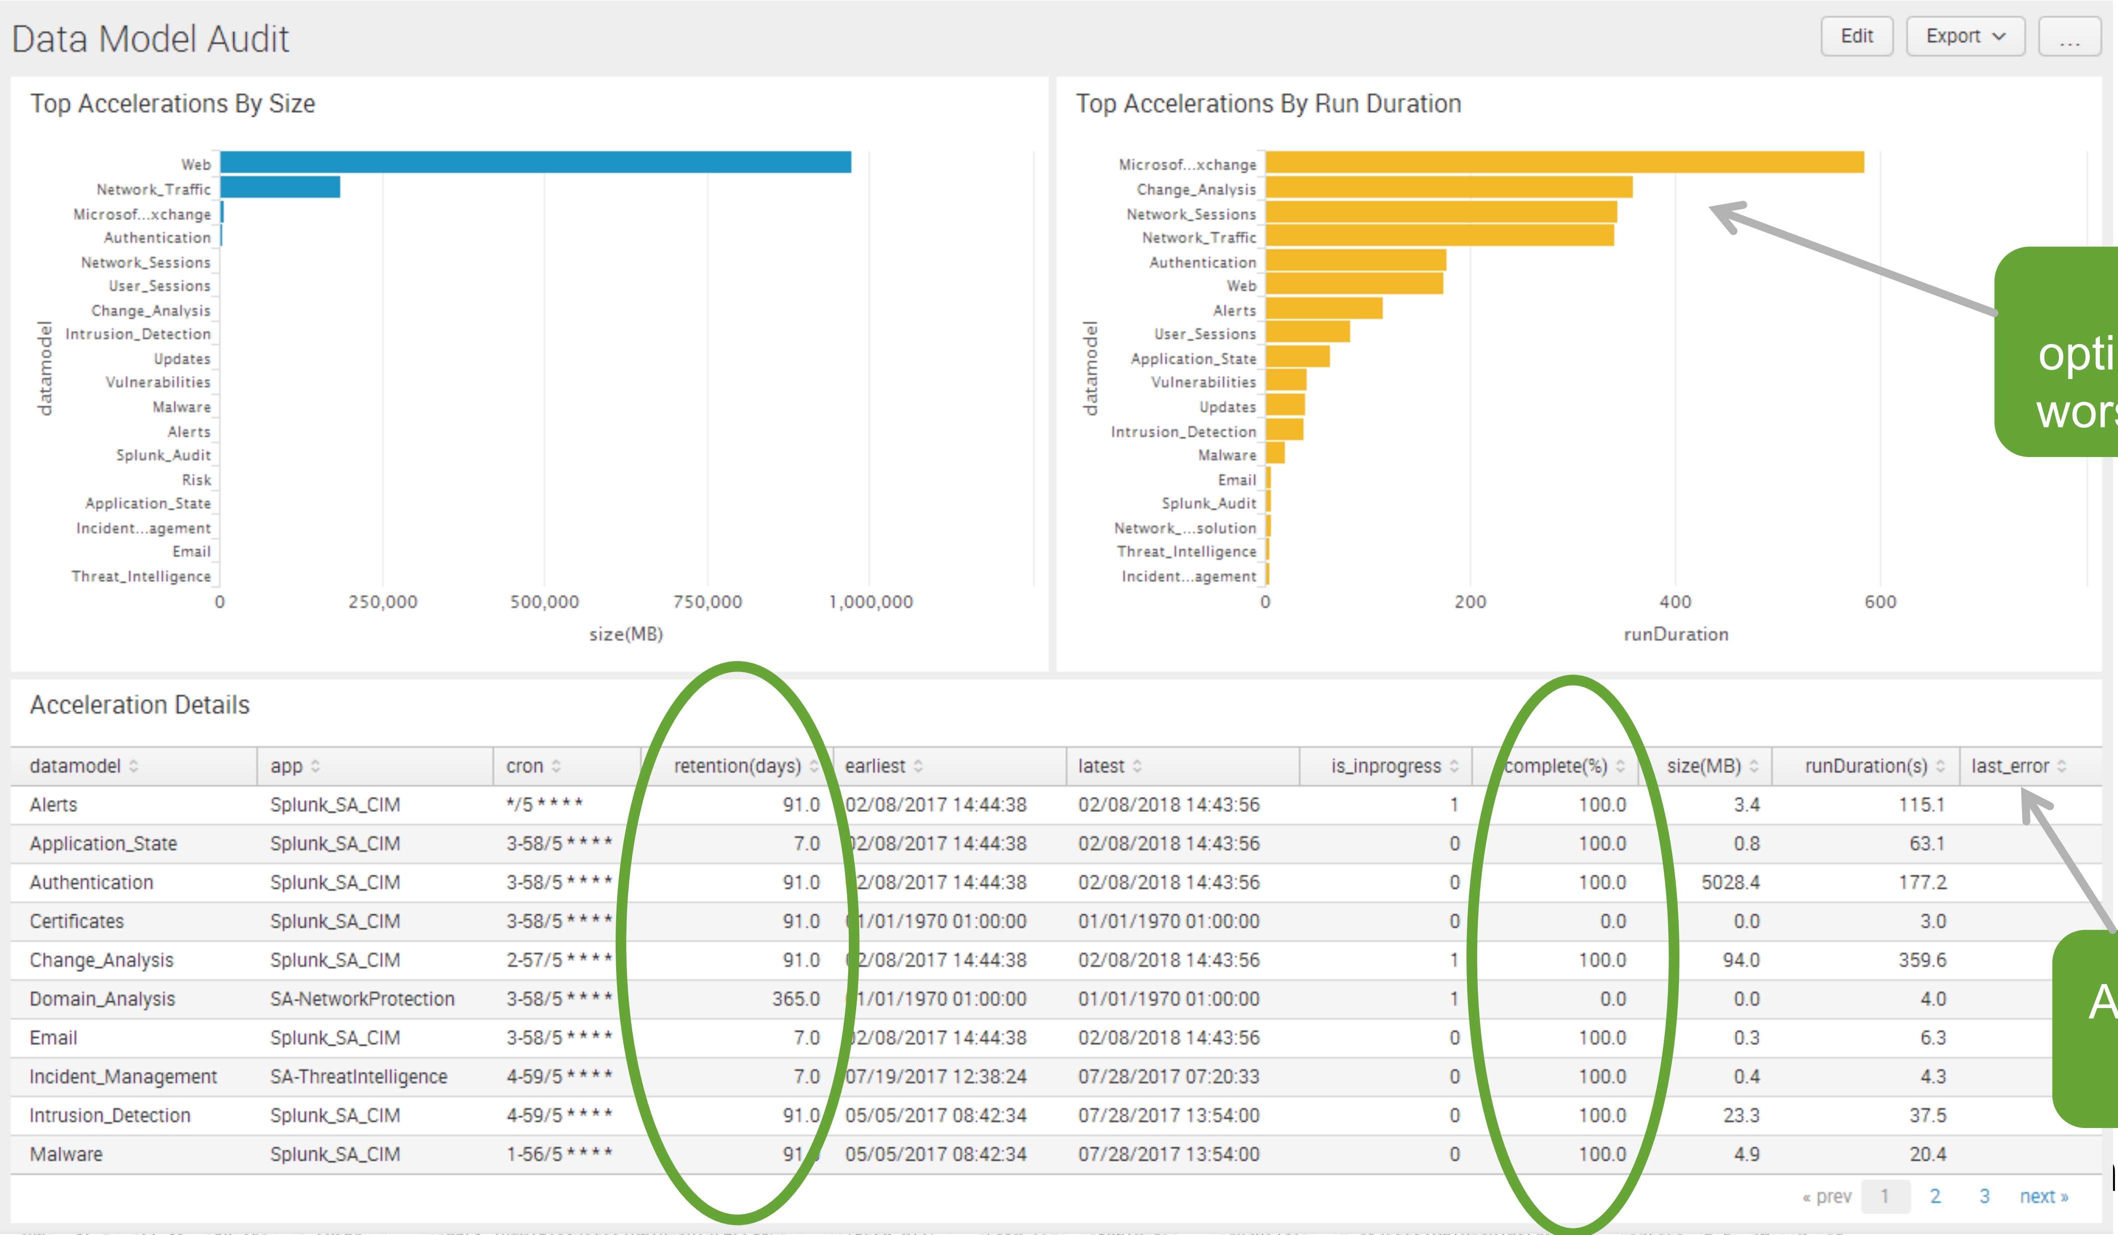

Acceleration Details

datamodel	app	cron	retention(days)	earliest	latest	is_inprogress	complete(%)	size(MB)	runDuration(s)	last_error
Alerts	Splunk_SA_CIM	*/5 * * * *	91.0	02/08/2017 14:44:38	02/08/2018 14:43:56	1	100.0	3.4	115.1	
Application_State	Splunk_SA_CIM	3-58/5 * * * *	7.0	02/08/2017 14:44:38	02/08/2018 14:43:56	0	100.0	0.8	63.1	
Authentication	Splunk_SA_CIM	3-58/5 * * * *	91.0	02/08/2017 14:44:38	02/08/2018 14:43:56	0	100.0	5028.4	177.2	
Certificates	Splunk_SA_CIM	3-58/5 * * * *	91.0	01/01/1970 01:00:00	01/01/1970 01:00:00	0	0.0	0.0	3.0	
Change_Analysis	Splunk_SA_CIM	2-57/5 * * * *	91.0	02/08/2017 14:44:38	02/08/2018 14:43:56	1	100.0	94.0	359.6	
Domain_Analysis	SA-NetworkProtection	3-58/5 * * * *	365.0	01/01/1970 01:00:00	01/01/1970 01:00:00	1	0.0	0.0	4.0	
Email	Splunk_SA_CIM	3-58/5 * * * *	7.0	02/08/2017 14:44:38	02/08/2018 14:43:56	0	100.0	0.3	6.3	
Incident_Management	SA-ThreatIntelligence	4-59/5 * * * *	7.0	07/19/2017 12:38:24	07/28/2017 07:20:33	0	100.0	0.4	4.3	
Intrusion_Detection	Splunk_SA_CIM	4-59/5 * * * *	91.0	05/05/2017 08:42:34	07/28/2017 13:54:00	0	100.0	23.3	37.5	
Malware	Splunk_SA_CIM	1-56/5 * * * *	91	05/05/2017 08:42:34	07/28/2017 13:54:00	0	100.0	4.9	20.4	

« prev 1 2 3 next »

Focus optimization on worst offenders

Anything here is BAAAD!

Monitoring DM Acceleration Lag

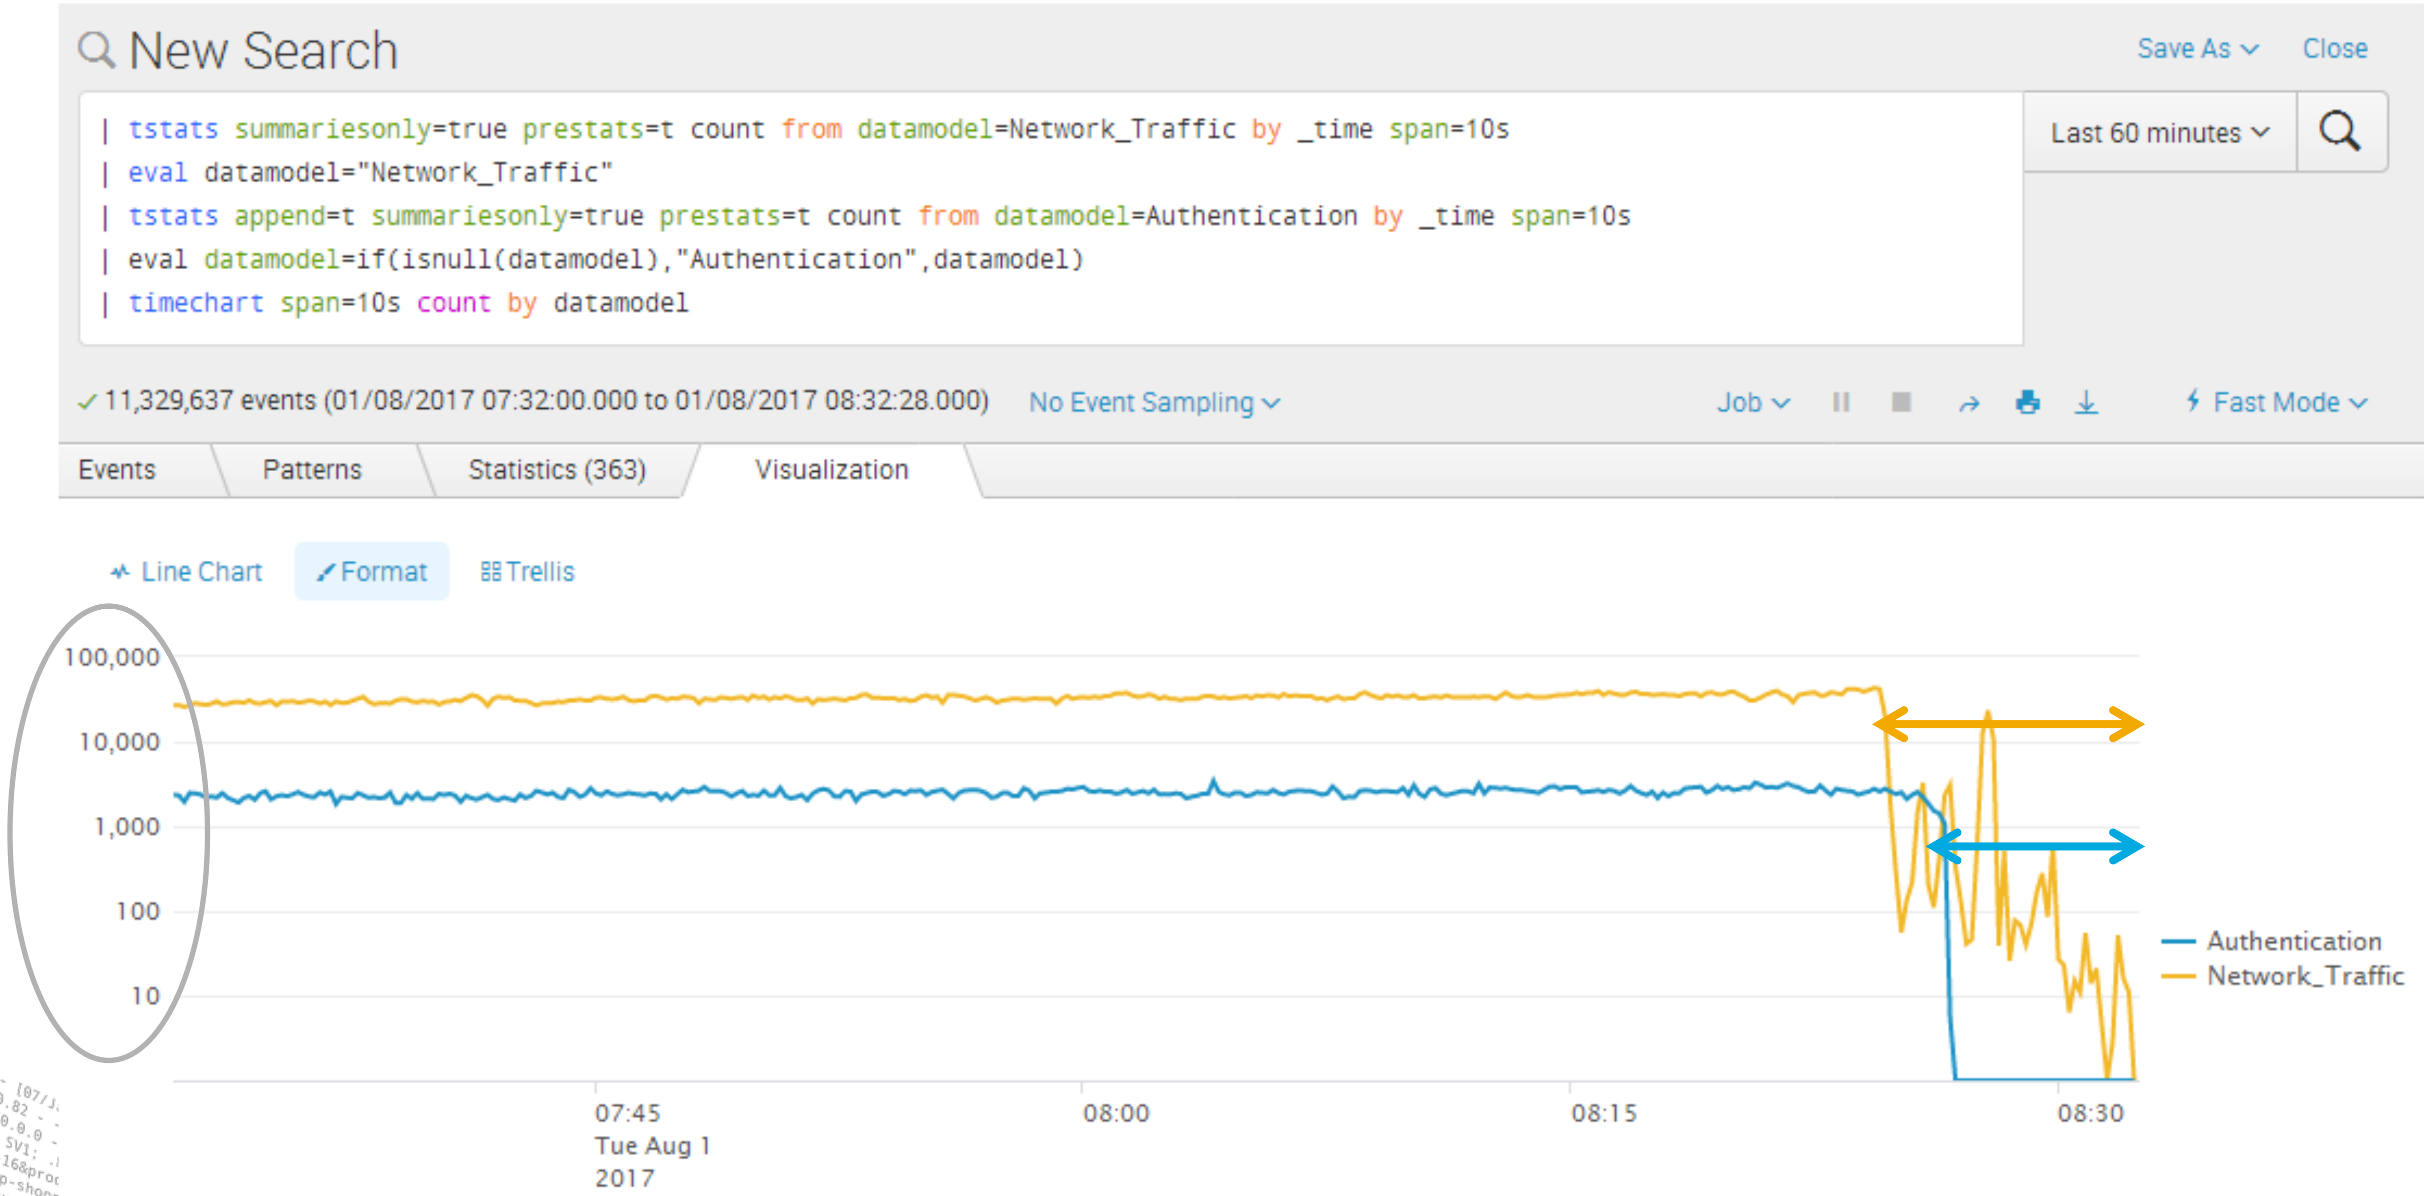

```
| tstats summariesonly=true prestats=t count from datamodel=Network_Traffic by _time span=10s
| eval datamodel="Network_Traffic"
| tstats append=t summariesonly=true prestats=t count from datamodel=Authentication by _time span=10s
| eval datamodel=if(isnull(datamodel),"Authentication",datamodel)
| timechart span=10s count by datamodel
```

✓ 11,329,637 events (01/08/2017 07:32:00.000 to 01/08/2017 08:32:28.000)

Data Model Acceleration

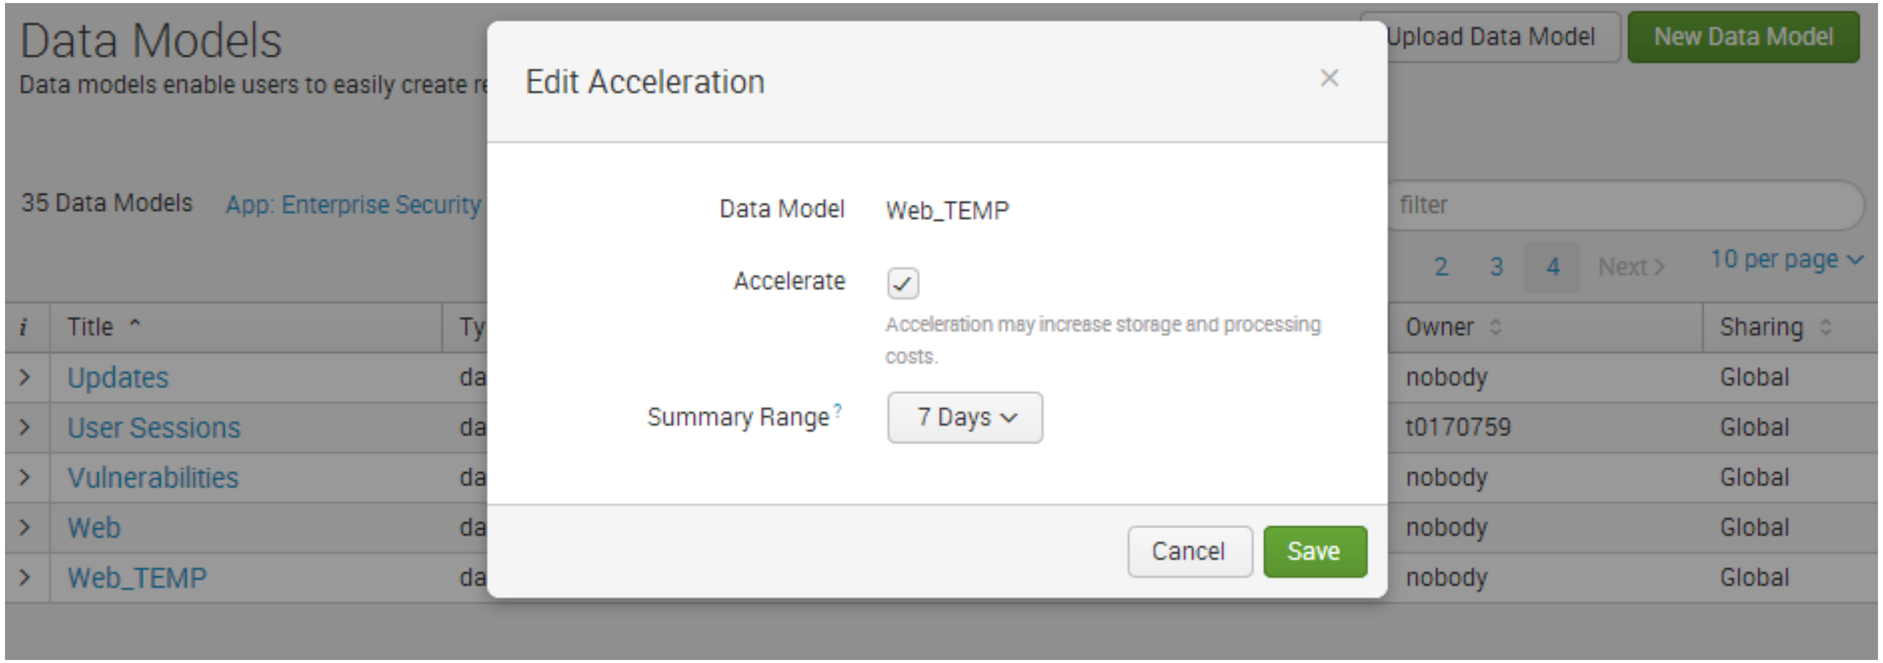

Data Models

Data models enable users to easily create re

35 Data Models App: Enterprise Security

Edit Acceleration

Data Model: Web_TEMP

Accelerate: ☑

Acceleration may increase storage and processing costs.

Summary Range?: 7 Days

i	Title ^	Ty	Owner	Sharing
>	Updates	da	nobody	Global
>	User Sessions	da	t0170759	Global
>	Vulnerabilities	da	nobody	Global
>	Web	da	nobody	Global
>	Web_TEMP	da	nobody	Global

- In datamodels.conf:

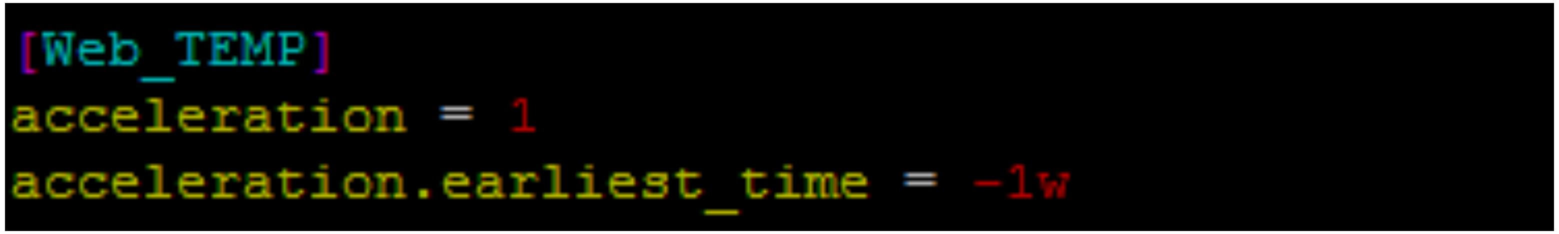

```
[Web_TEMP]
acceleration = 1
acceleration.earliest_time = -1w
```

Understanding DM Backfilling

Dashboard Source included!

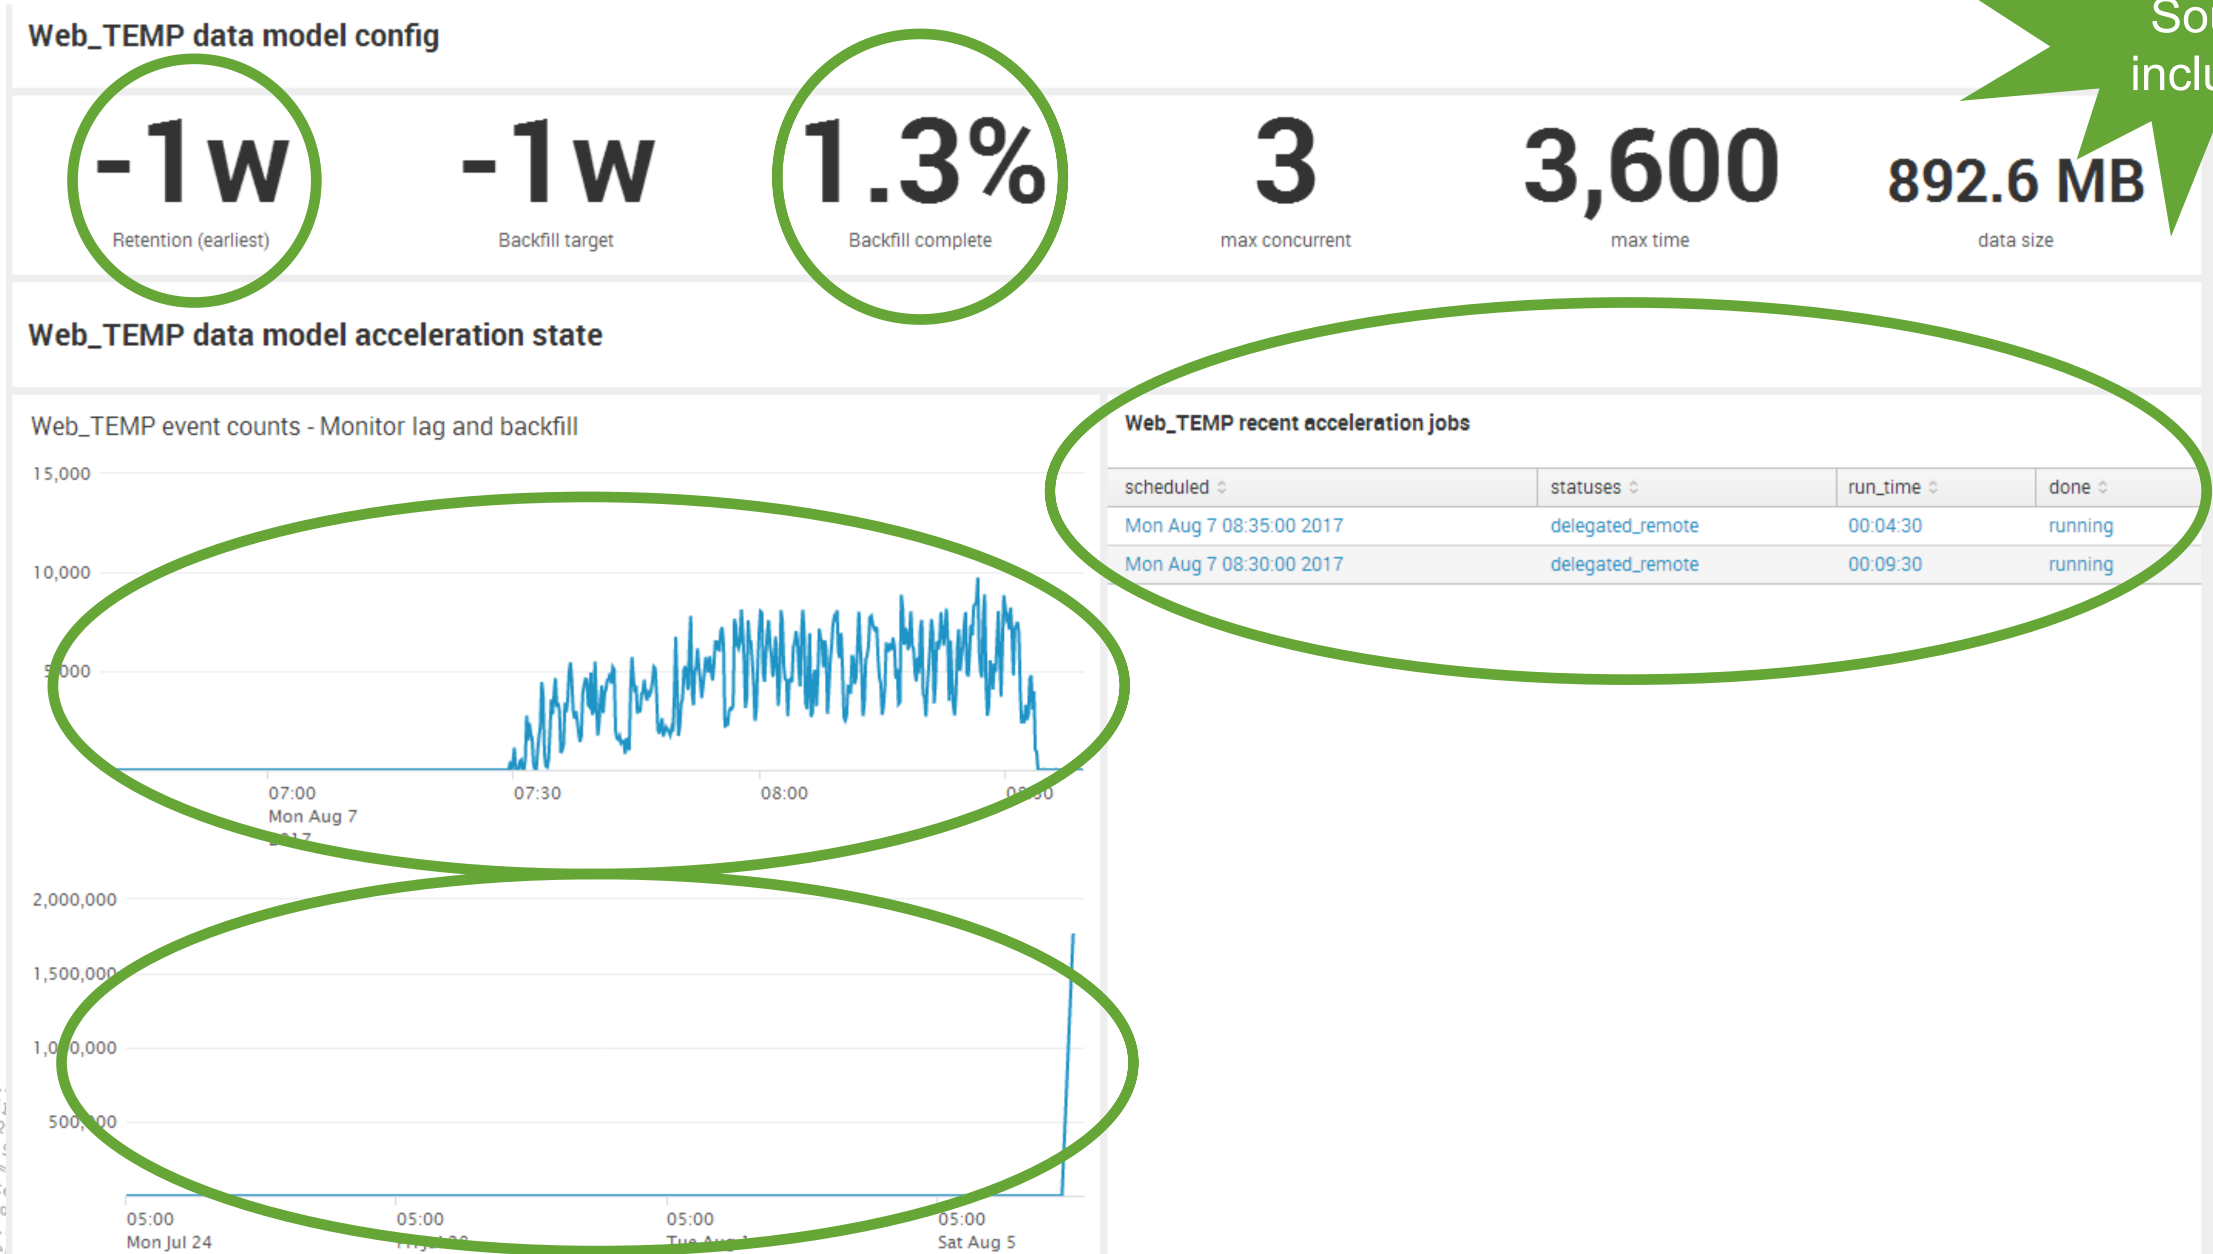

Web_TEMP data model config

-1w	-1w	1.3%	3	3,600	892.6 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Mon Aug 7 08:35:00 2017	delegated_remote	00:04:30	running
Mon Aug 7 08:30:00 2017	delegated_remote	00:09:30	running

Understanding DM Backfilling

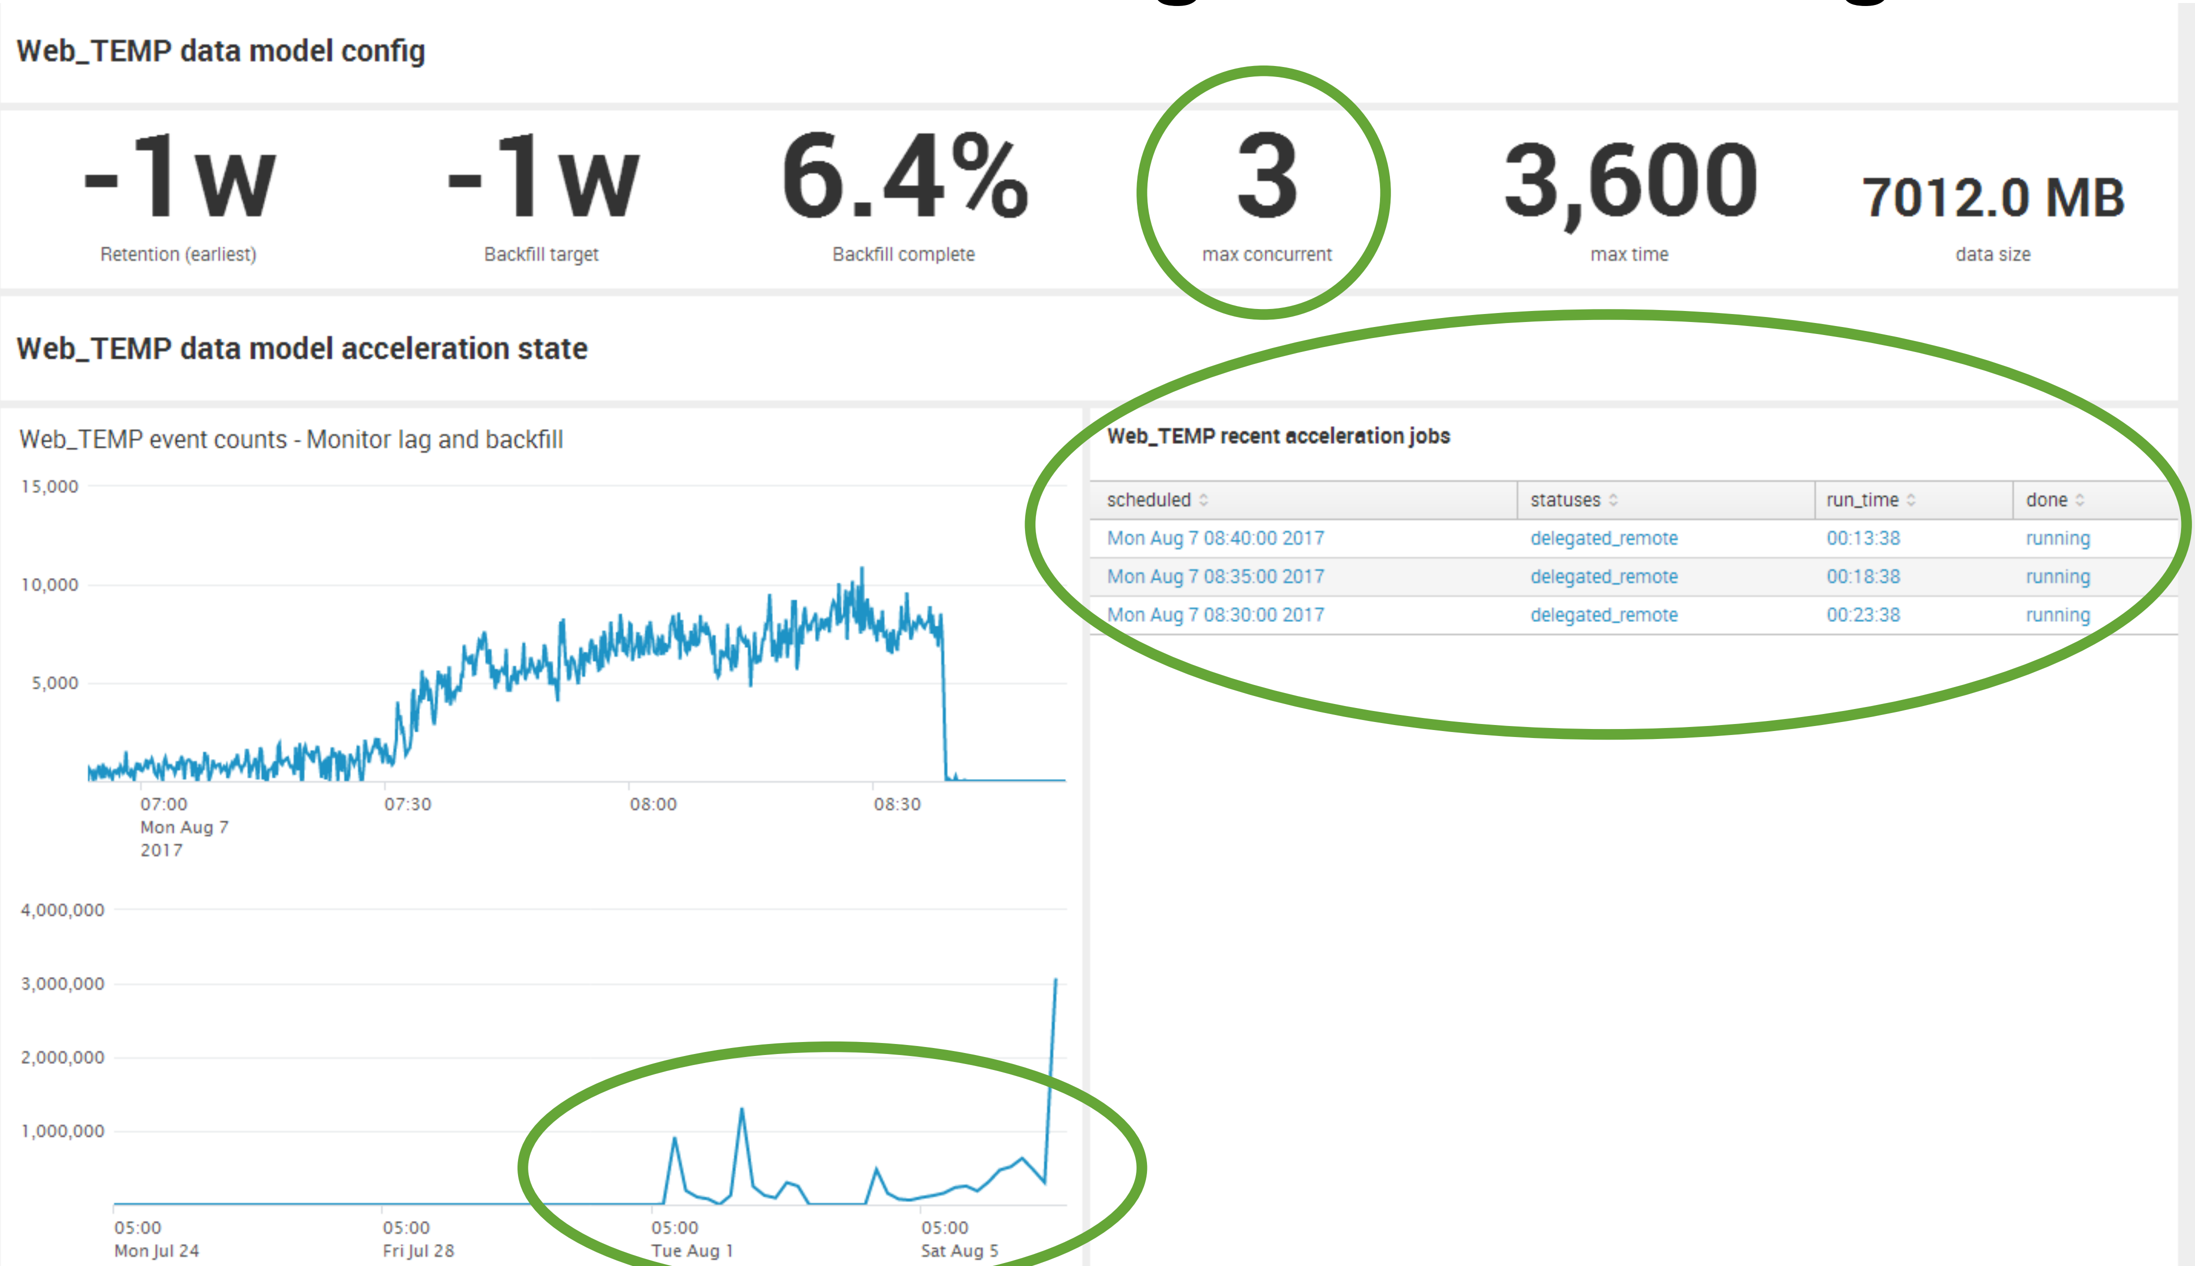

Web_TEMP data model config

-1w	-1w	6.4%	3	3,600	7012.0 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Mon Aug 7 08:40:00 2017	delegated_remote	00:13:38	running
Mon Aug 7 08:35:00 2017	delegated_remote	00:18:38	running
Mon Aug 7 08:30:00 2017	delegated_remote	00:23:38	running

Understanding DM Backfilling

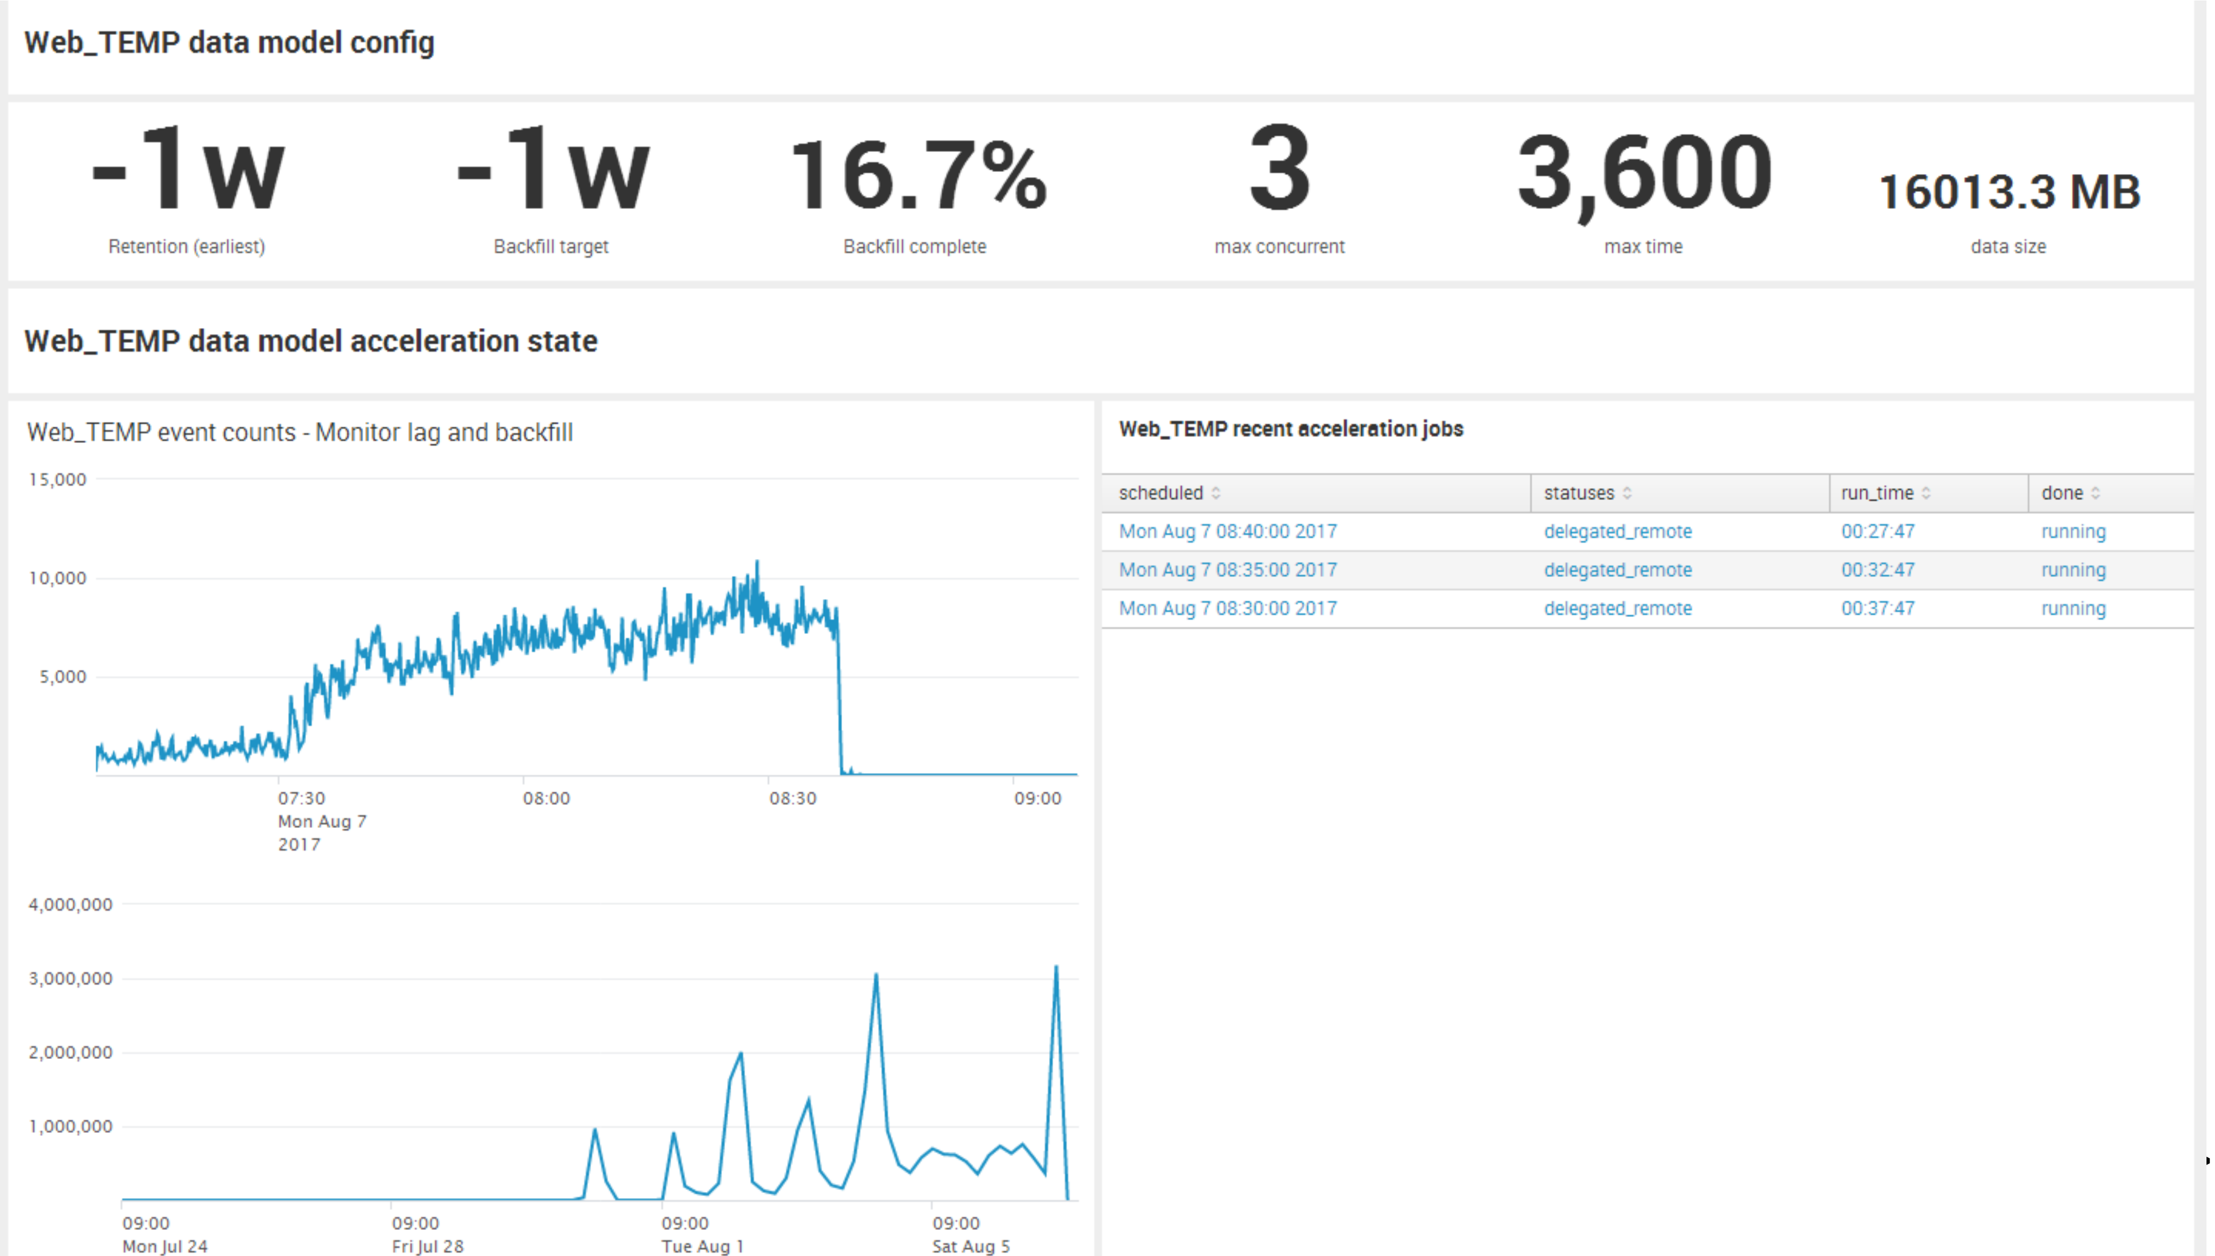

Web_TEMP data model config

-1w	-1w	16.7%	3	3,600	16013.3 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled ⇕	statuses ⇕	run_time ⇕	done ⇕
Mon Aug 7 08:40:00 2017	delegated_remote	00:27:47	running
Mon Aug 7 08:35:00 2017	delegated_remote	00:32:47	running
Mon Aug 7 08:30:00 2017	delegated_remote	00:37:47	running

.conf2017

Understanding DM Backfilling

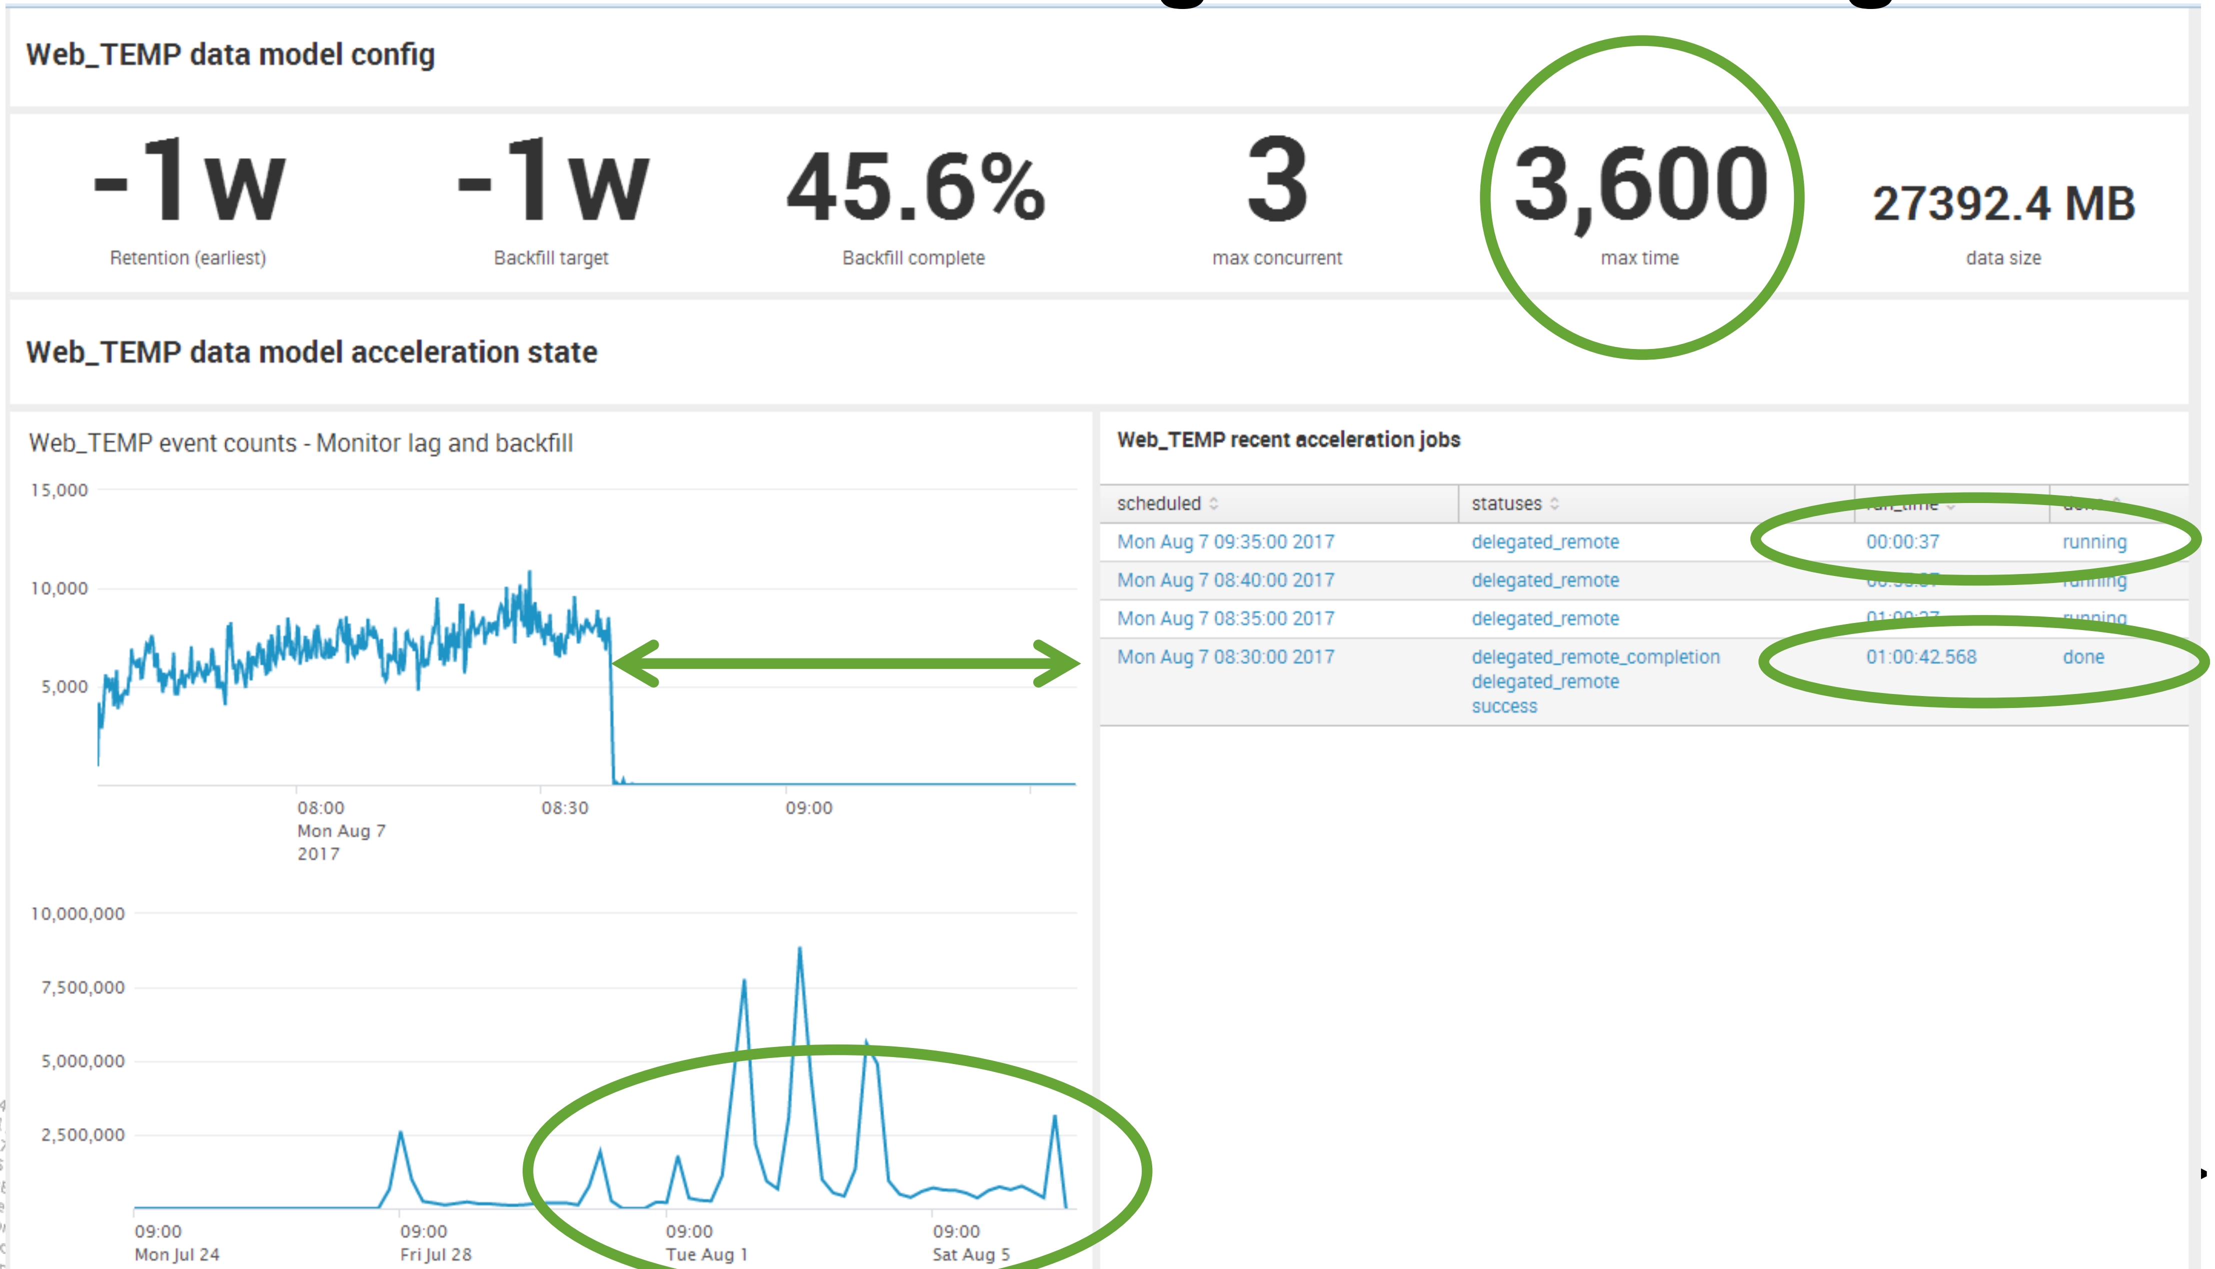

Web_TEMP data model config

-1w	-1w	45.6%	3	3,600	27392.4 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Mon Aug 7 09:35:00 2017	delegated_remote	00:00:37	running
Mon Aug 7 08:40:00 2017	delegated_remote	[illegible]	running
Mon Aug 7 08:35:00 2017	delegated_remote	[illegible]	running
Mon Aug 7 08:30:00 2017	delegated_remote_completion delegated_remote success	01:00:42.568	done

Understanding DM Backfilling

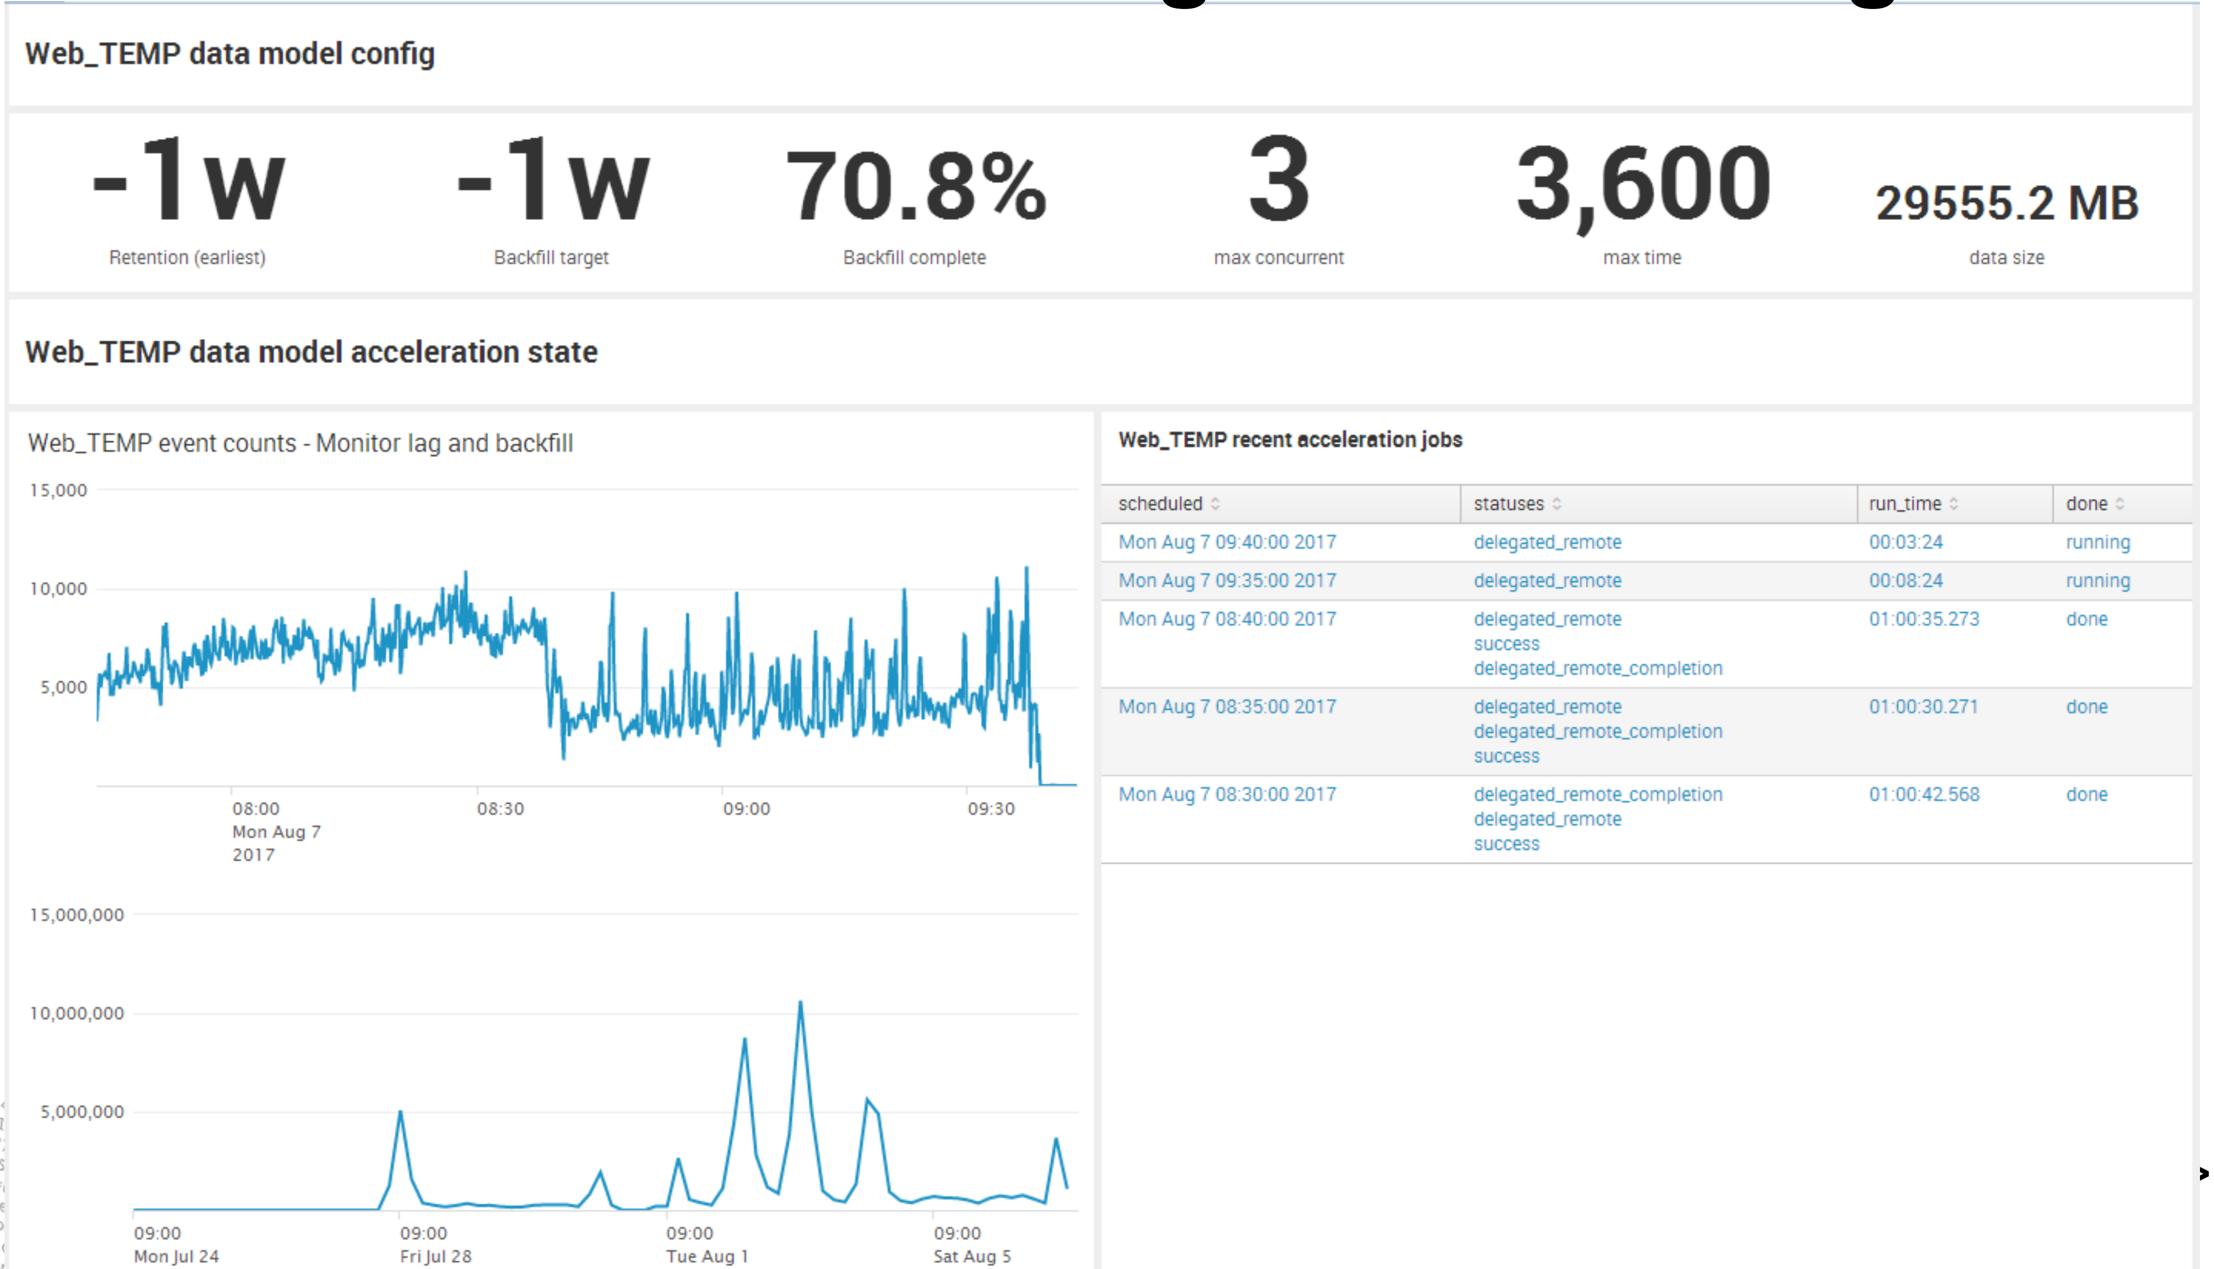

Web_TEMP data model config

-1w	-1w	70.8%	3	3,600	29555.2 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Mon Aug 7 09:40:00 2017	delegated_remote	00:03:24	running
Mon Aug 7 09:35:00 2017	delegated_remote	00:08:24	running
Mon Aug 7 08:40:00 2017	delegated_remote success delegated_remote_completion	01:00:35.273	done
Mon Aug 7 08:35:00 2017	delegated_remote delegated_remote_completion success	01:00:30.271	done
Mon Aug 7 08:30:00 2017	delegated_remote_completion delegated_remote success	01:00:42.568	done

Understanding DM Backfilling

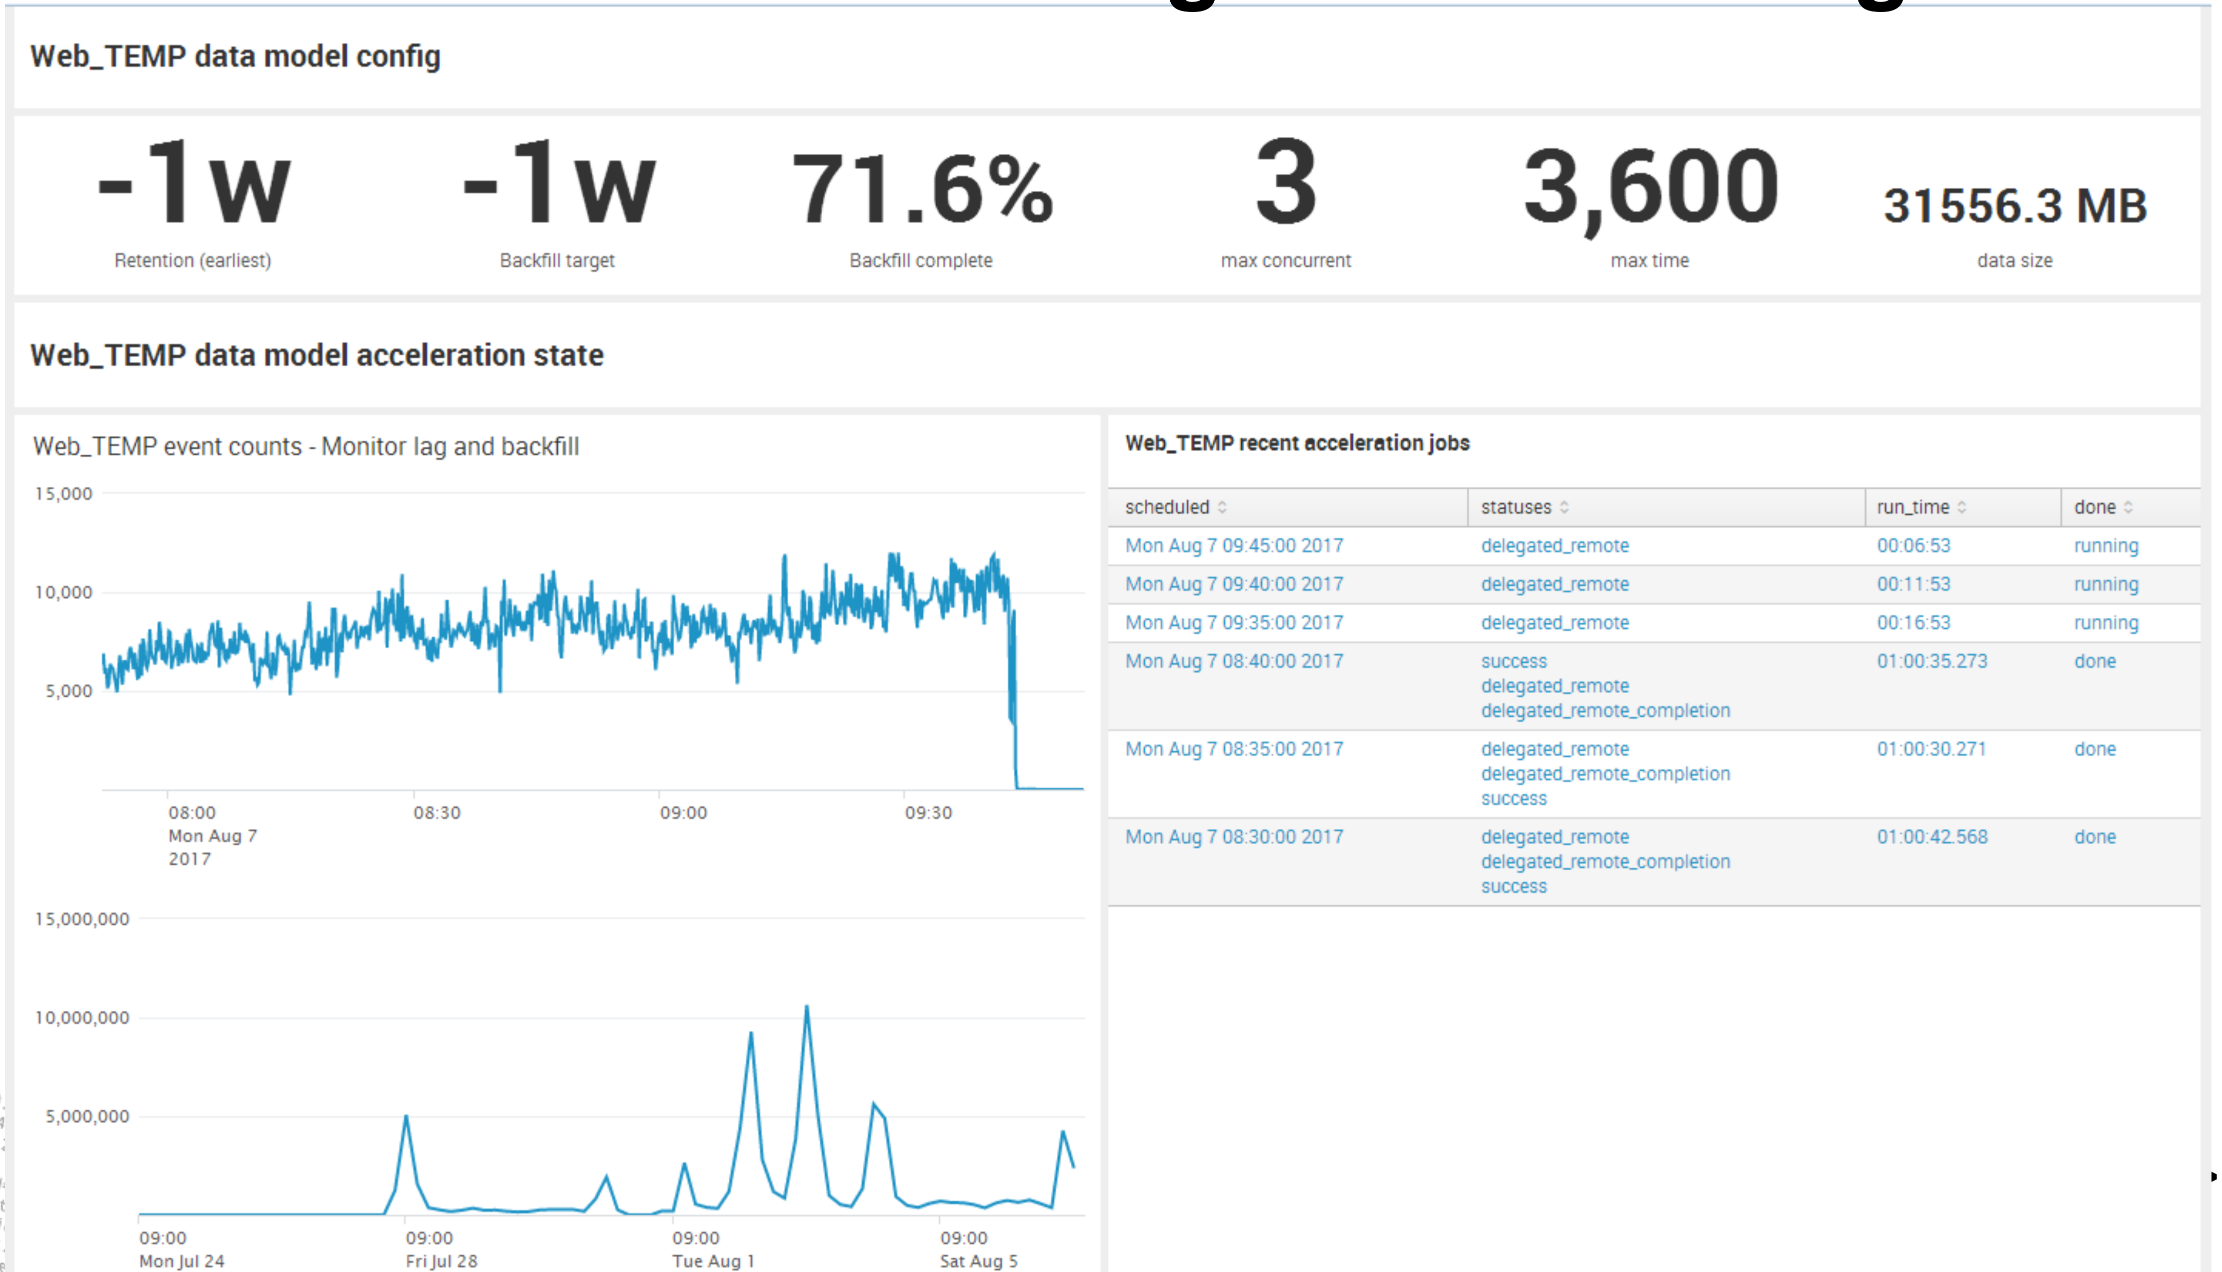

Web_TEMP data model config

-1w	-1w	71.6%	3	3,600	31556.3 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Mon Aug 7 09:45:00 2017	delegated_remote	00:06:53	running
Mon Aug 7 09:40:00 2017	delegated_remote	00:11:53	running
Mon Aug 7 09:35:00 2017	delegated_remote	00:16:53	running
Mon Aug 7 08:40:00 2017	success delegated_remote delegated_remote_completion	01:00:35.273	done
Mon Aug 7 08:35:00 2017	delegated_remote delegated_remote_completion success	01:00:30.271	done
Mon Aug 7 08:30:00 2017	delegated_remote delegated_remote_completion success	01:00:42.568	done

Understanding DM Backfilling

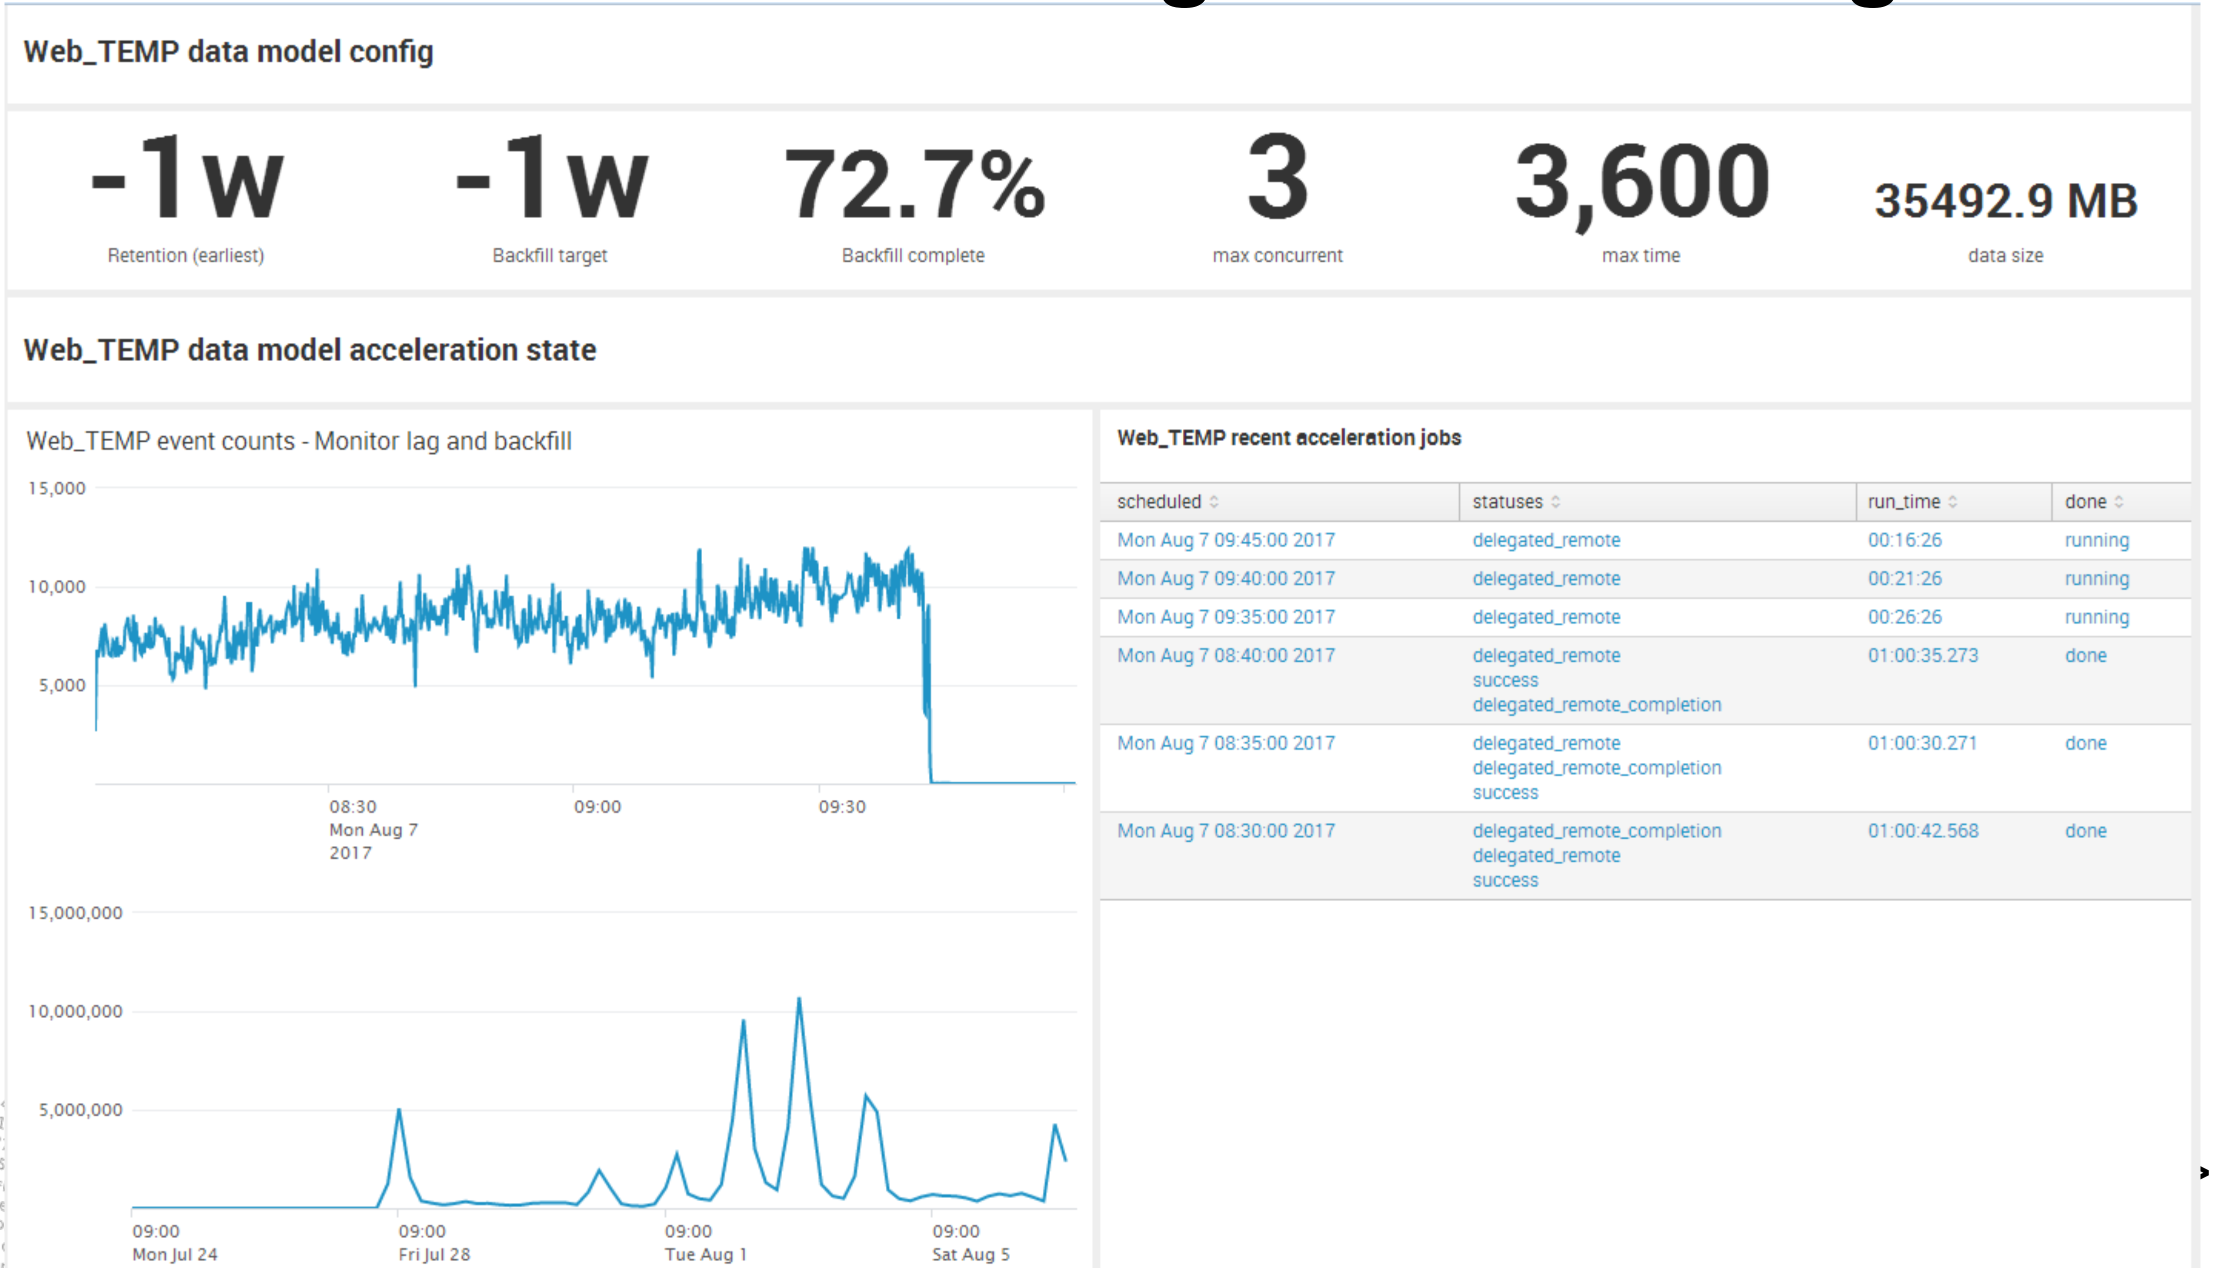

Web_TEMP data model config

-1w	-1w	72.7%	3	3,600	35492.9 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Mon Aug 7 09:45:00 2017	delegated_remote	00:16:26	running
Mon Aug 7 09:40:00 2017	delegated_remote	00:21:26	running
Mon Aug 7 09:35:00 2017	delegated_remote	00:26:26	running
Mon Aug 7 08:40:00 2017	delegated_remote success delegated_remote_completion	01:00:35.273	done
Mon Aug 7 08:35:00 2017	delegated_remote delegated_remote_completion success	01:00:30.271	done
Mon Aug 7 08:30:00 2017	delegated_remote_completion delegated_remote success	01:00:42.568	done

Understanding DM Backfilling

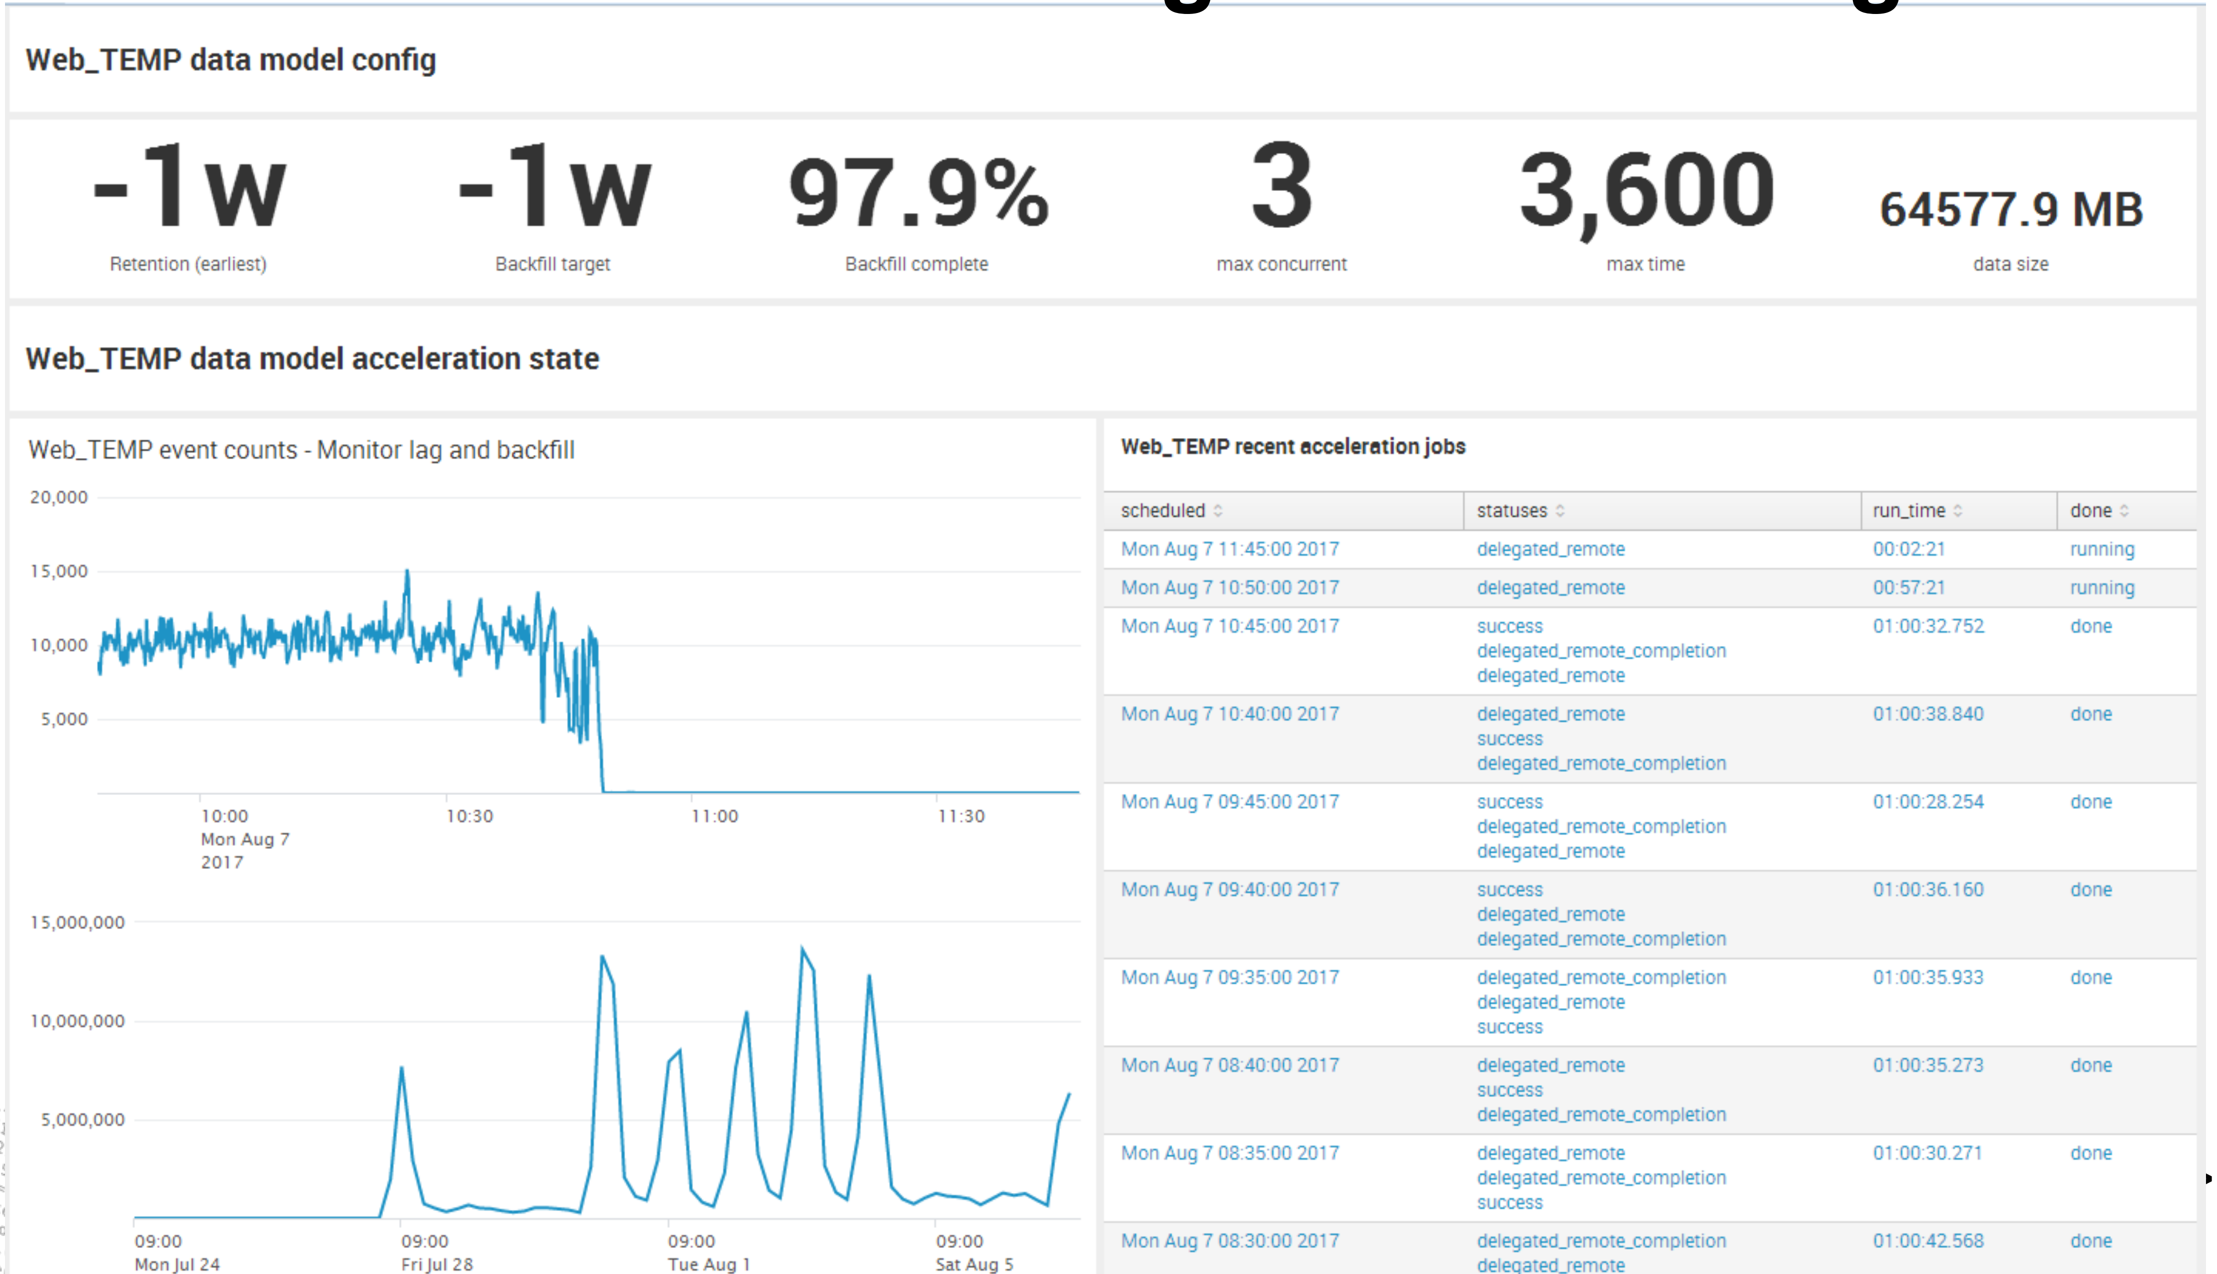

Web_TEMP data model config

-1w	-1w	97.9%	3	3,600	64577.9 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Mon Aug 7 11:45:00 2017	delegated_remote	00:02:21	running
Mon Aug 7 10:50:00 2017	delegated_remote	00:57:21	running
Mon Aug 7 10:45:00 2017	success delegated_remote_completion delegated_remote	01:00:32.752	done
Mon Aug 7 10:40:00 2017	delegated_remote success delegated_remote_completion	01:00:38.840	done
Mon Aug 7 09:45:00 2017	success delegated_remote_completion delegated_remote	01:00:28.254	done
Mon Aug 7 09:40:00 2017	success delegated_remote delegated_remote_completion	01:00:36.160	done
Mon Aug 7 09:35:00 2017	delegated_remote_completion delegated_remote success	01:00:35.933	done
Mon Aug 7 08:40:00 2017	delegated_remote success delegated_remote_completion	01:00:35.273	done
Mon Aug 7 08:35:00 2017	delegated_remote delegated_remote_completion success	01:00:30.271	done
Mon Aug 7 08:30:00 2017	delegated_remote_completion delegated_remote	01:00:42.568	done

Understanding DM Backfilling

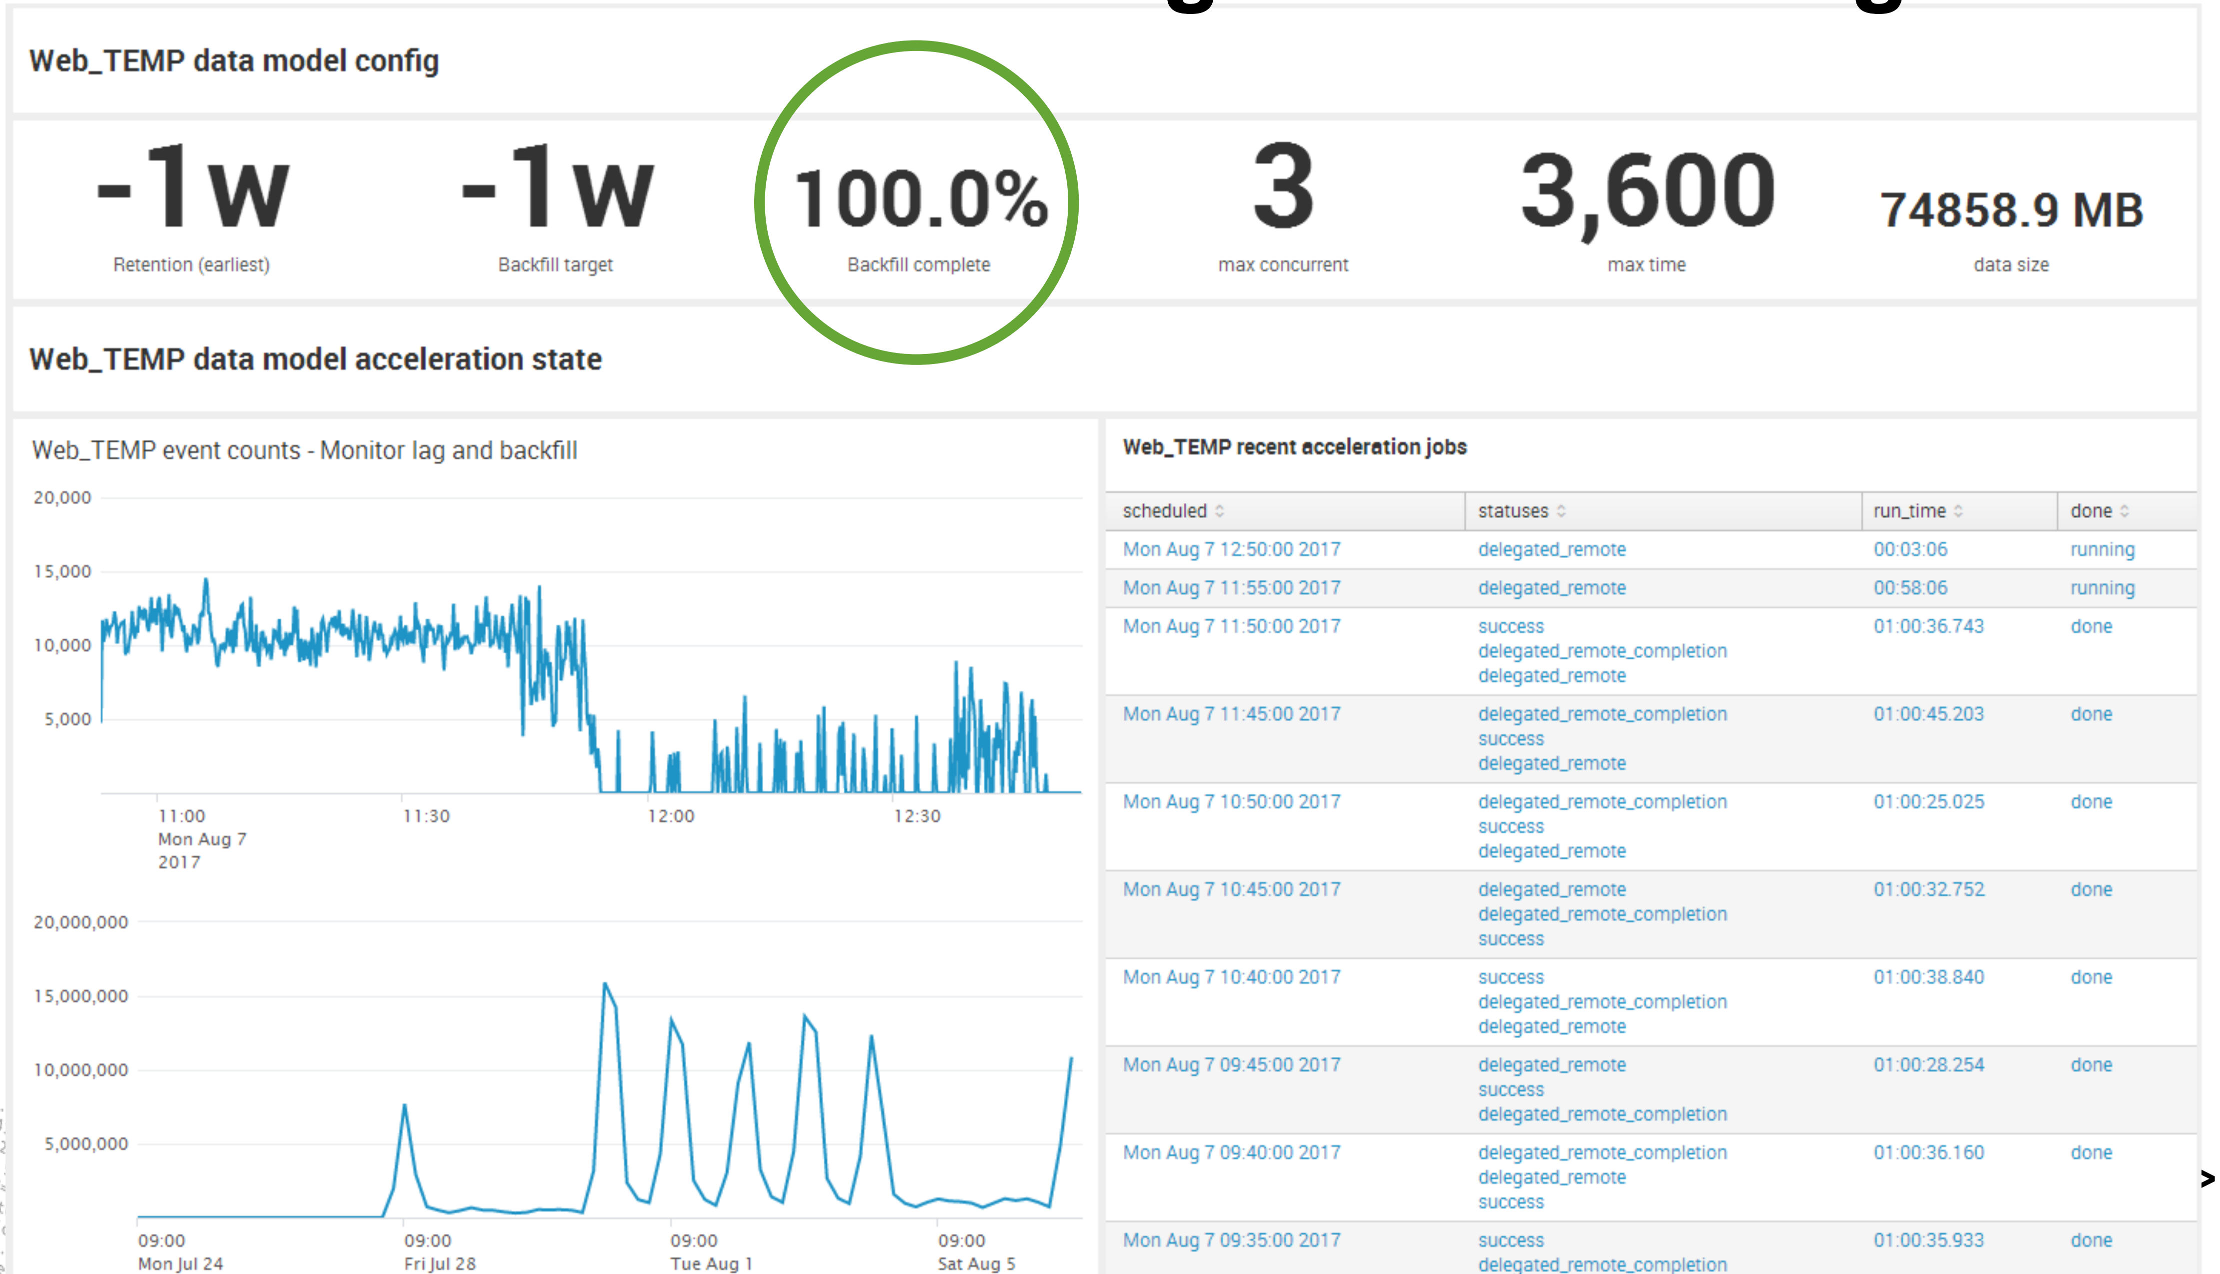

Web_TEMP data model config

-1w	-1w	100.0%	3	3,600	74858.9 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Mon Aug 7 12:50:00 2017	delegated_remote	00:03:06	running
Mon Aug 7 11:55:00 2017	delegated_remote	00:58:06	running
Mon Aug 7 11:50:00 2017	success delegated_remote_completion delegated_remote	01:00:36.743	done
Mon Aug 7 11:45:00 2017	delegated_remote_completion success delegated_remote	01:00:45.203	done
Mon Aug 7 10:50:00 2017	delegated_remote_completion success delegated_remote	01:00:25.025	done
Mon Aug 7 10:45:00 2017	delegated_remote delegated_remote_completion success	01:00:32.752	done
Mon Aug 7 10:40:00 2017	success delegated_remote_completion delegated_remote	01:00:38.840	done
Mon Aug 7 09:45:00 2017	delegated_remote success delegated_remote_completion	01:00:28.254	done
Mon Aug 7 09:40:00 2017	delegated_remote_completion delegated_remote success	01:00:36.160	done
Mon Aug 7 09:35:00 2017	success delegated_remote_completion	01:00:35.933	done

Understanding DM Backfilling

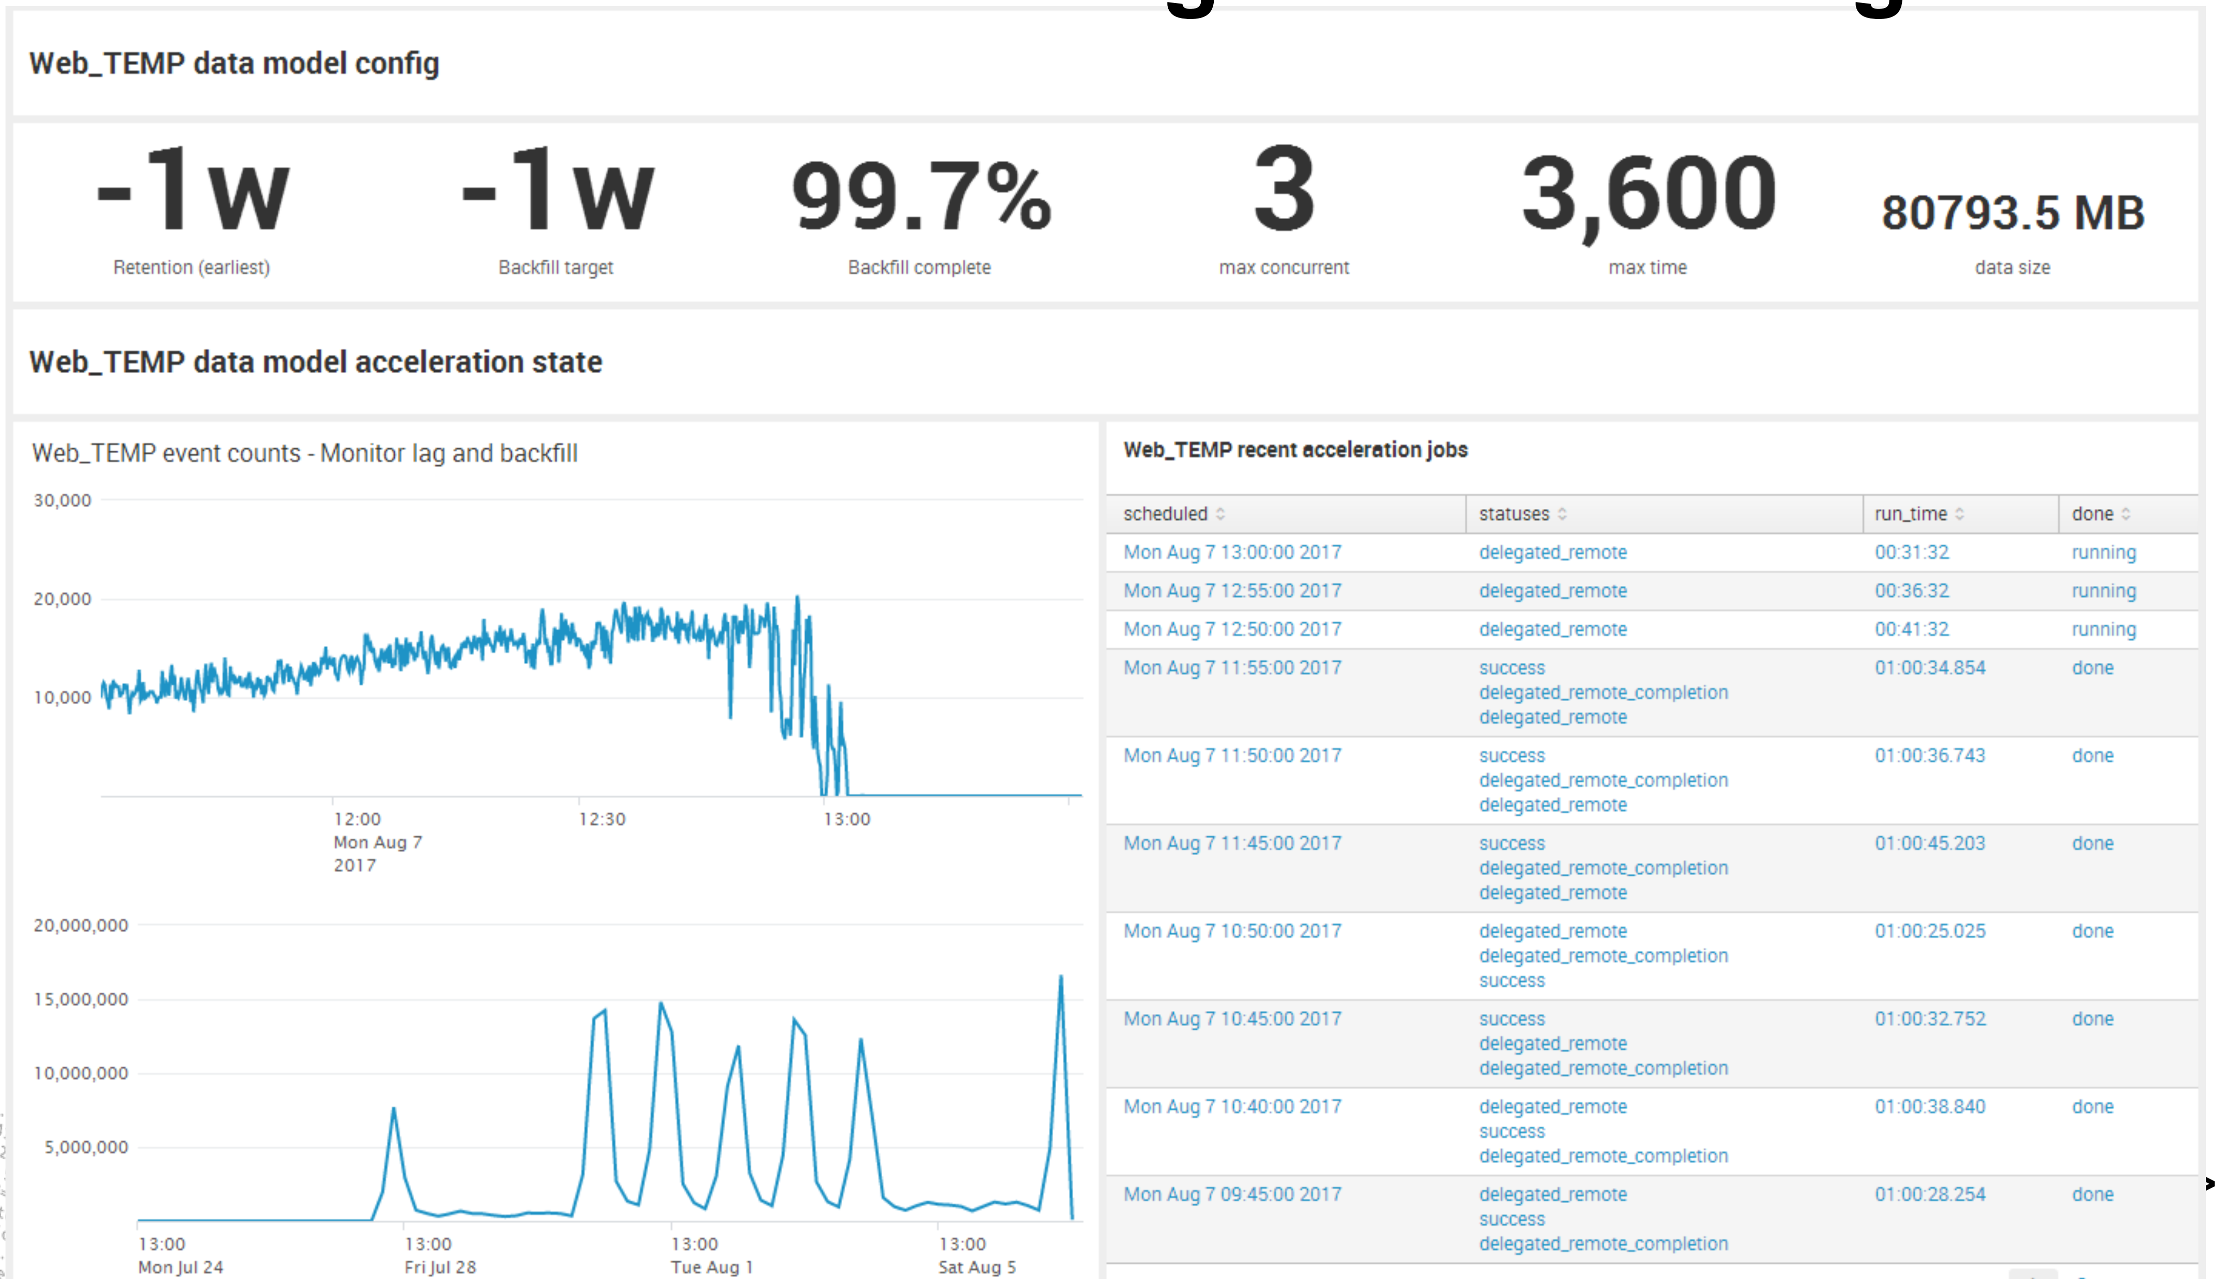

Web_TEMP data model config

-1w	-1w	99.7%	3	3,600	80793.5 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Mon Aug 7 13:00:00 2017	delegated_remote	00:31:32	running
Mon Aug 7 12:55:00 2017	delegated_remote	00:36:32	running
Mon Aug 7 12:50:00 2017	delegated_remote	00:41:32	running
Mon Aug 7 11:55:00 2017	success delegated_remote_completion delegated_remote	01:00:34.854	done
Mon Aug 7 11:50:00 2017	success delegated_remote_completion delegated_remote	01:00:36.743	done
Mon Aug 7 11:45:00 2017	success delegated_remote_completion delegated_remote	01:00:45.203	done
Mon Aug 7 10:50:00 2017	delegated_remote delegated_remote_completion success	01:00:25.025	done
Mon Aug 7 10:45:00 2017	success delegated_remote delegated_remote_completion	01:00:32.752	done
Mon Aug 7 10:40:00 2017	delegated_remote success delegated_remote_completion	01:00:38.840	done
Mon Aug 7 09:45:00 2017	delegated_remote success delegated_remote_completion	01:00:28.254	done

.conf2017

Understanding DM Backfilling

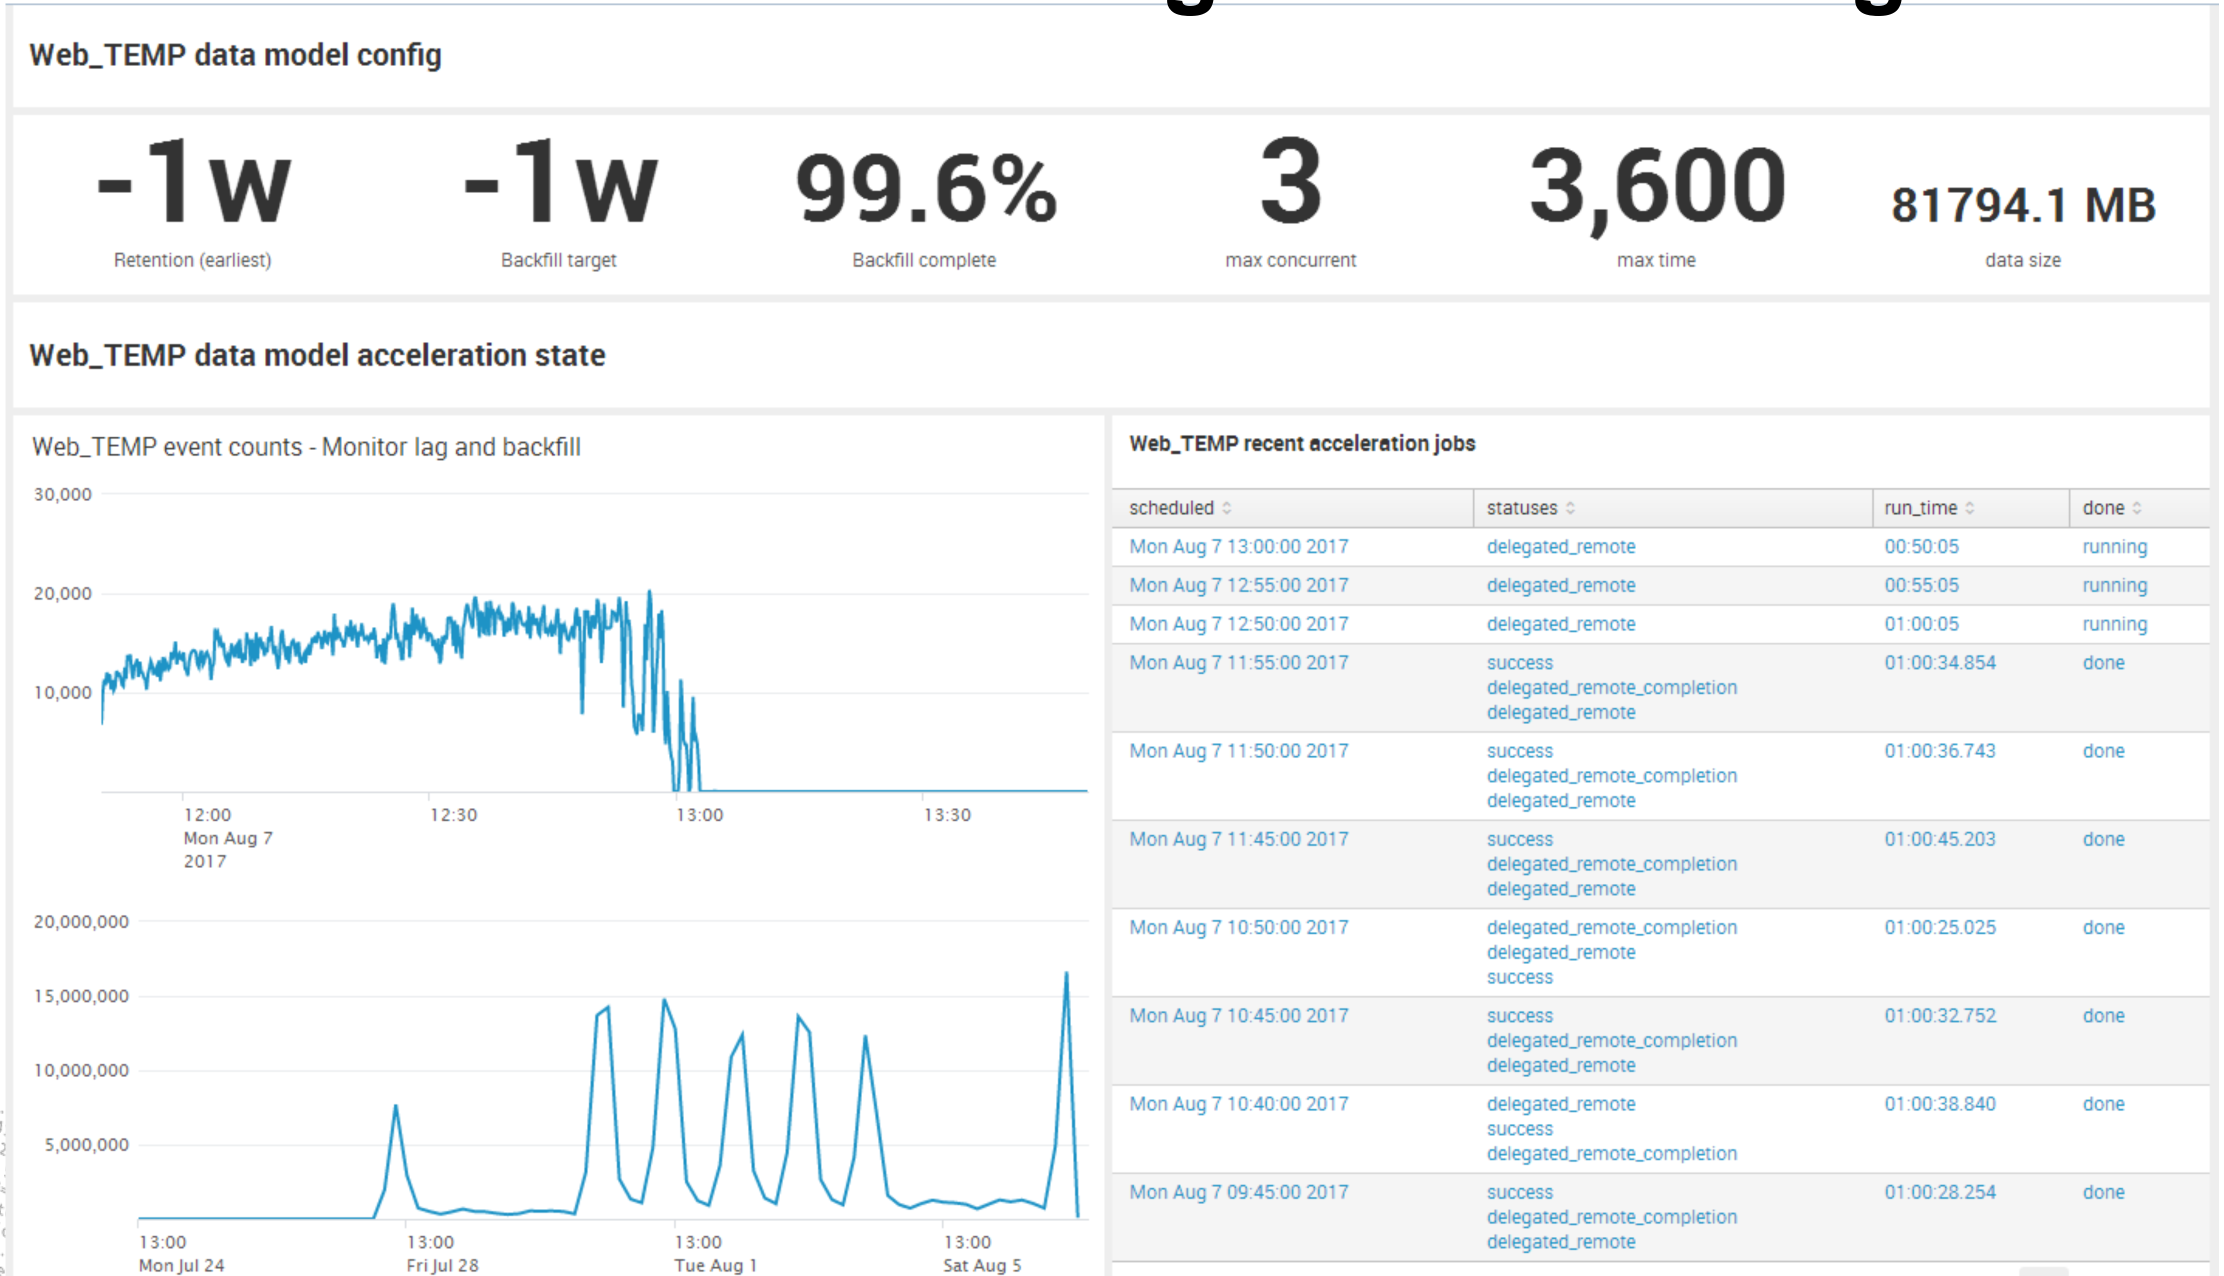

Web_TEMP data model config

-1w	-1w	99.6%	3	3,600	81794.1 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Mon Aug 7 13:00:00 2017	delegated_remote	00:50:05	running
Mon Aug 7 12:55:00 2017	delegated_remote	00:55:05	running
Mon Aug 7 12:50:00 2017	delegated_remote	01:00:05	running
Mon Aug 7 11:55:00 2017	success delegated_remote_completion delegated_remote	01:00:34.854	done
Mon Aug 7 11:50:00 2017	success delegated_remote_completion delegated_remote	01:00:36.743	done
Mon Aug 7 11:45:00 2017	success delegated_remote_completion delegated_remote	01:00:45.203	done
Mon Aug 7 10:50:00 2017	delegated_remote_completion delegated_remote success	01:00:25.025	done
Mon Aug 7 10:45:00 2017	success delegated_remote_completion delegated_remote	01:00:32.752	done
Mon Aug 7 10:40:00 2017	delegated_remote success delegated_remote_completion	01:00:38.840	done
Mon Aug 7 09:45:00 2017	success delegated_remote_completion delegated_remote	01:00:28.254	done

Understanding DM Backfilling

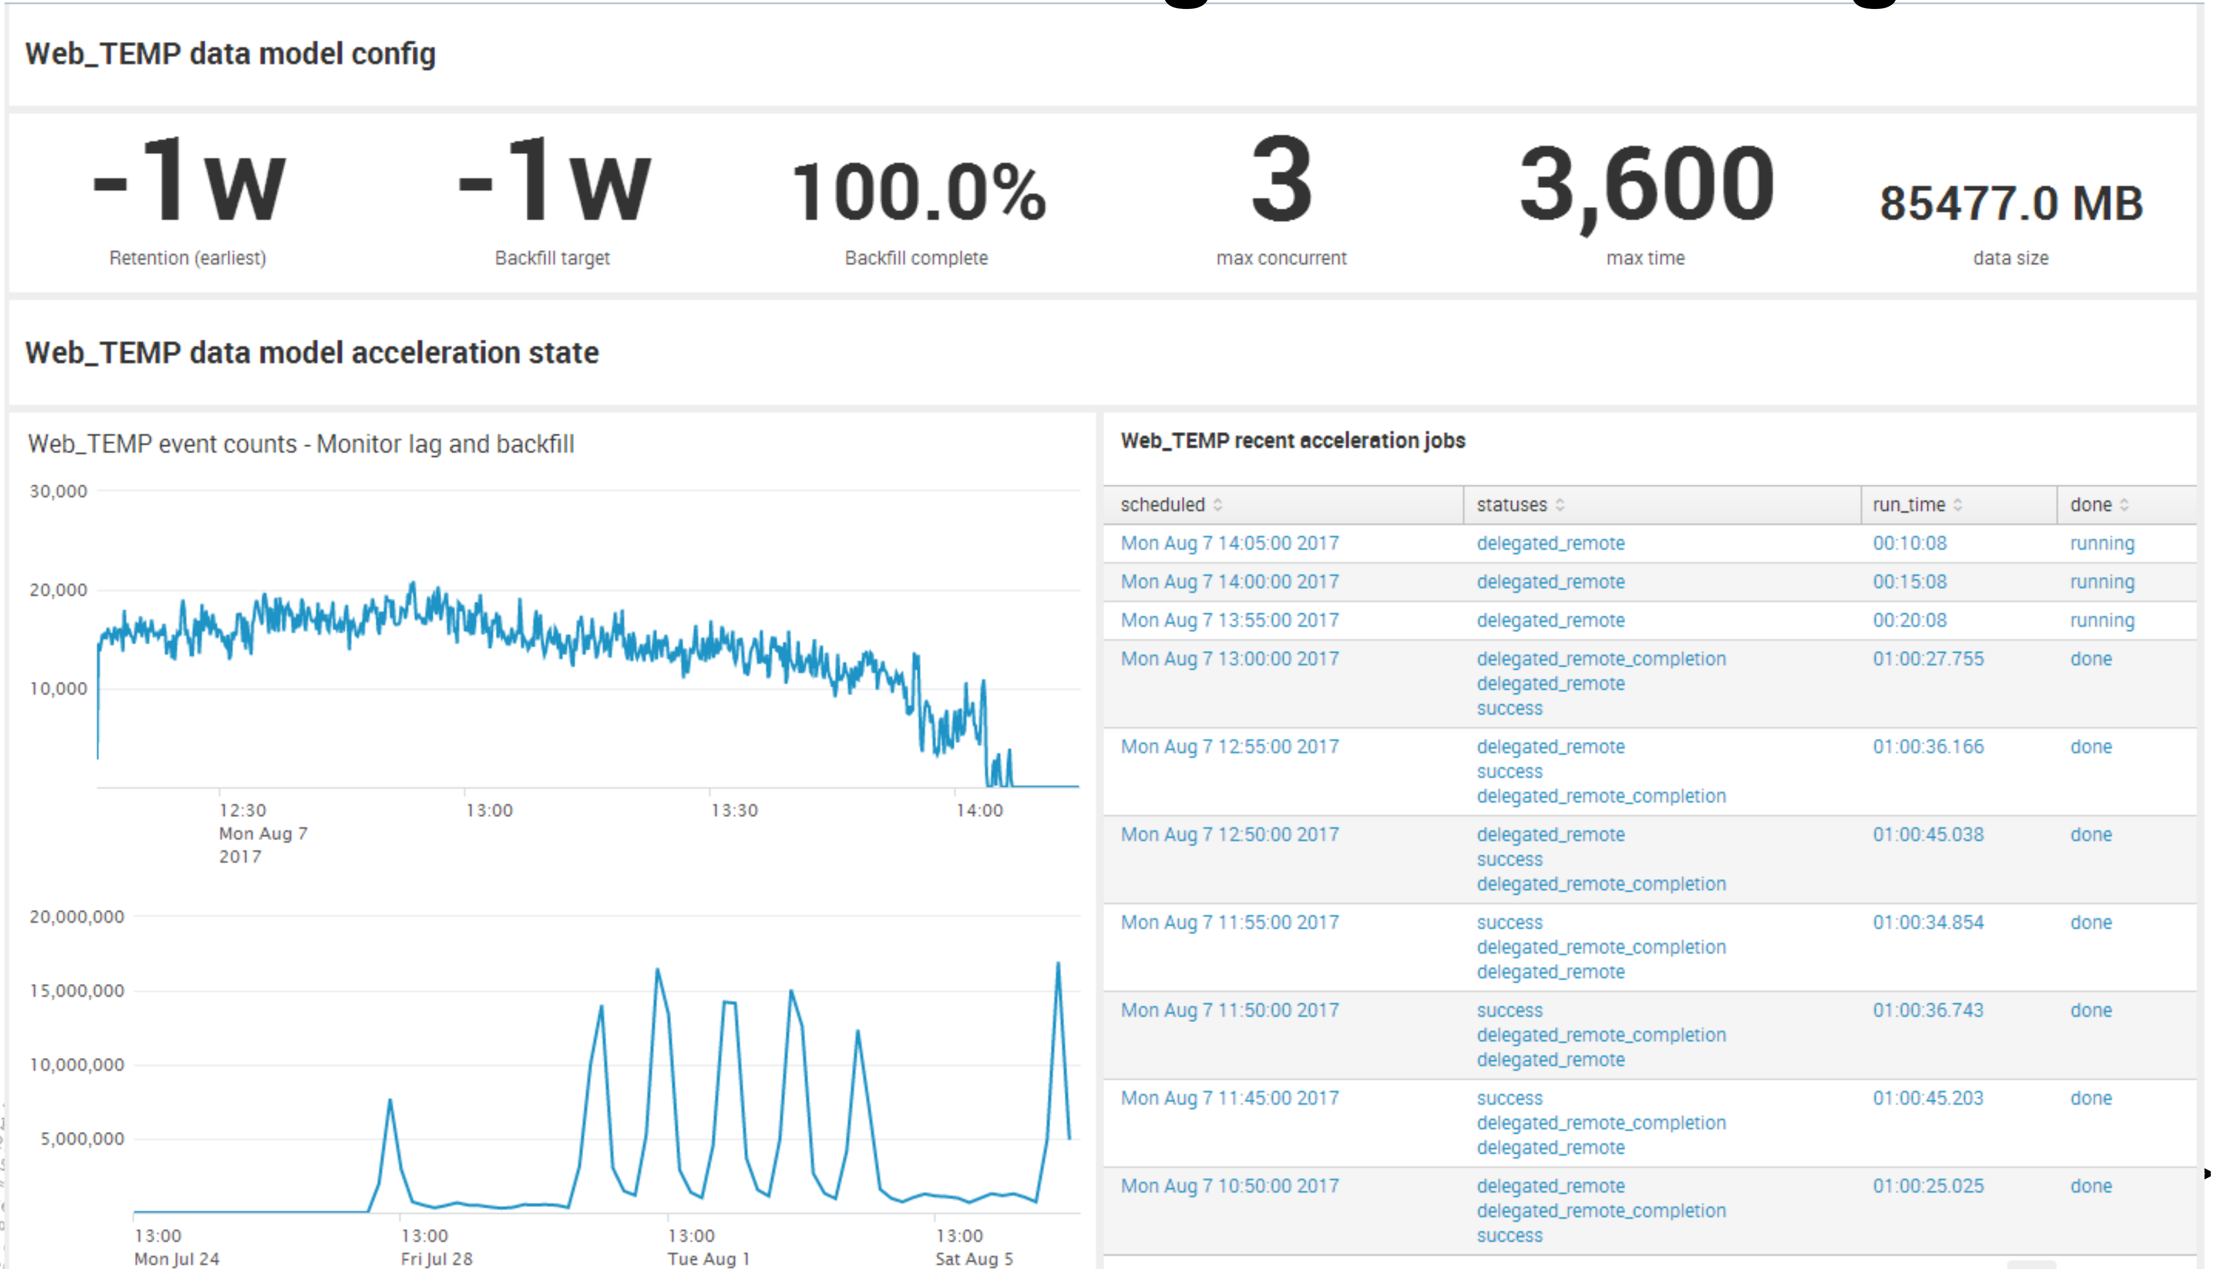

Web_TEMP data model config

-1w	-1w	100.0%	3	3,600	85477.0 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Mon Aug 7 14:05:00 2017	delegated_remote	00:10:08	running
Mon Aug 7 14:00:00 2017	delegated_remote	00:15:08	running
Mon Aug 7 13:55:00 2017	delegated_remote	00:20:08	running
Mon Aug 7 13:00:00 2017	delegated_remote_completion delegated_remote success	01:00:27.755	done
Mon Aug 7 12:55:00 2017	delegated_remote success delegated_remote_completion	01:00:36.166	done
Mon Aug 7 12:50:00 2017	delegated_remote success delegated_remote_completion	01:00:45.038	done
Mon Aug 7 11:55:00 2017	success delegated_remote_completion delegated_remote	01:00:34.854	done
Mon Aug 7 11:50:00 2017	success delegated_remote_completion delegated_remote	01:00:36.743	done
Mon Aug 7 11:45:00 2017	success delegated_remote_completion delegated_remote	01:00:45.203	done
Mon Aug 7 10:50:00 2017	delegated_remote delegated_remote_completion success	01:00:25.025	done

Understanding DM Backfilling

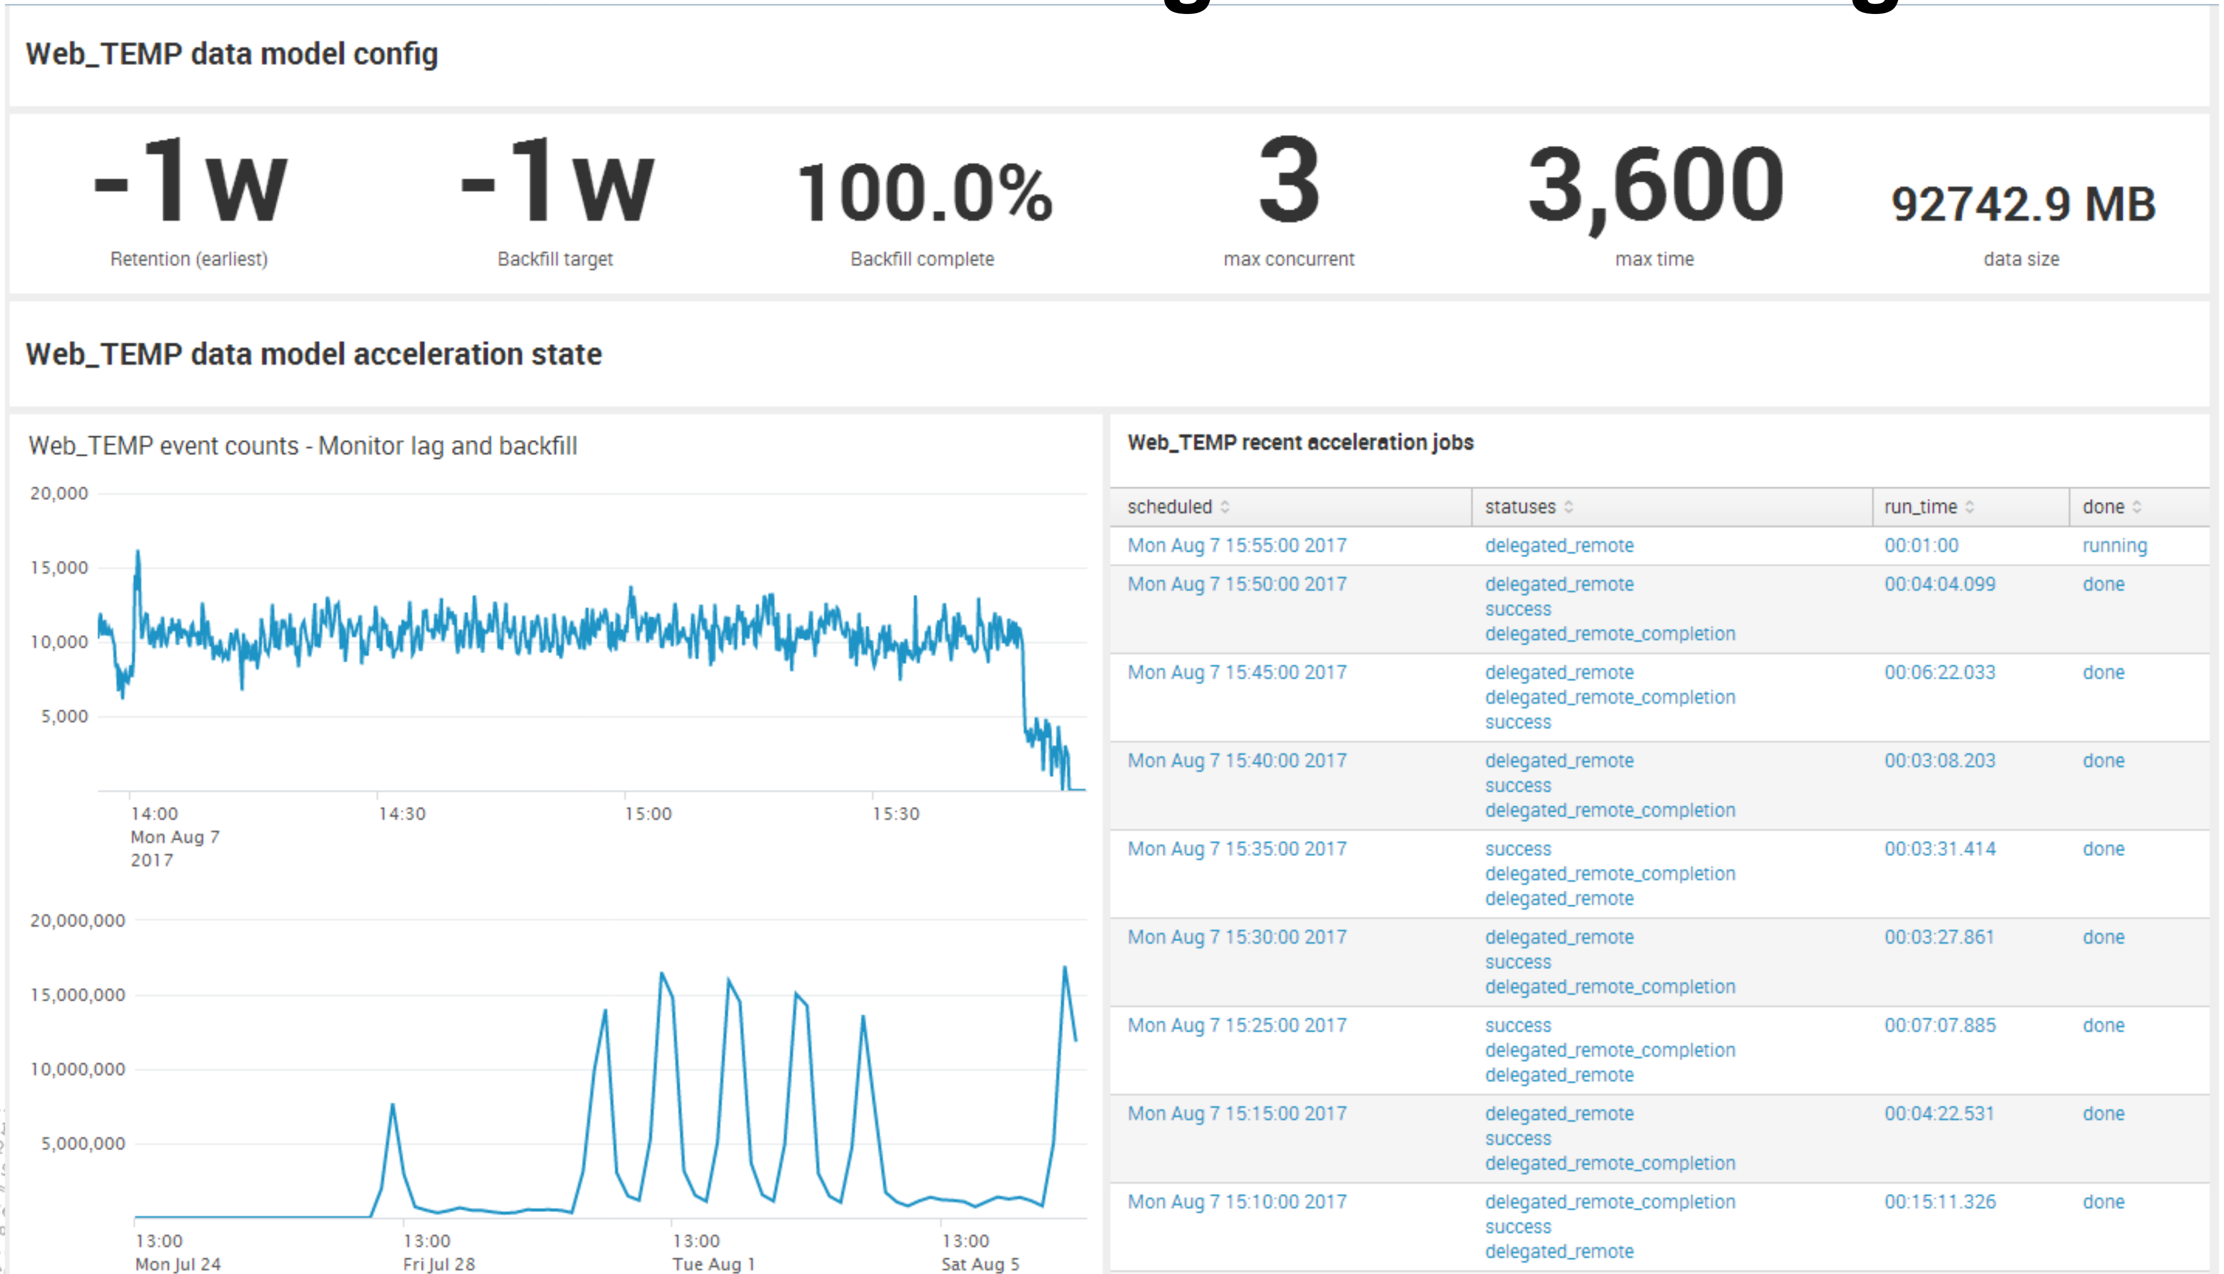

Web_TEMP data model config

-1w	-1w	100.0%	3	3,600	92742.9 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Mon Aug 7 15:55:00 2017	delegated_remote	00:01:00	running
Mon Aug 7 15:50:00 2017	delegated_remote success delegated_remote_completion	00:04:04.099	done
Mon Aug 7 15:45:00 2017	delegated_remote delegated_remote_completion success	00:06:22.033	done
Mon Aug 7 15:40:00 2017	delegated_remote success delegated_remote_completion	00:03:08.203	done
Mon Aug 7 15:35:00 2017	success delegated_remote_completion delegated_remote	00:03:31.414	done
Mon Aug 7 15:30:00 2017	delegated_remote success delegated_remote_completion	00:03:27.861	done
Mon Aug 7 15:25:00 2017	success delegated_remote_completion delegated_remote	00:07:07.885	done
Mon Aug 7 15:15:00 2017	delegated_remote success delegated_remote_completion	00:04:22.531	done
Mon Aug 7 15:10:00 2017	delegated_remote_completion success delegated_remote	00:15:11.326	done

Limiting Data Model Backfilling

- In datamodels.conf:

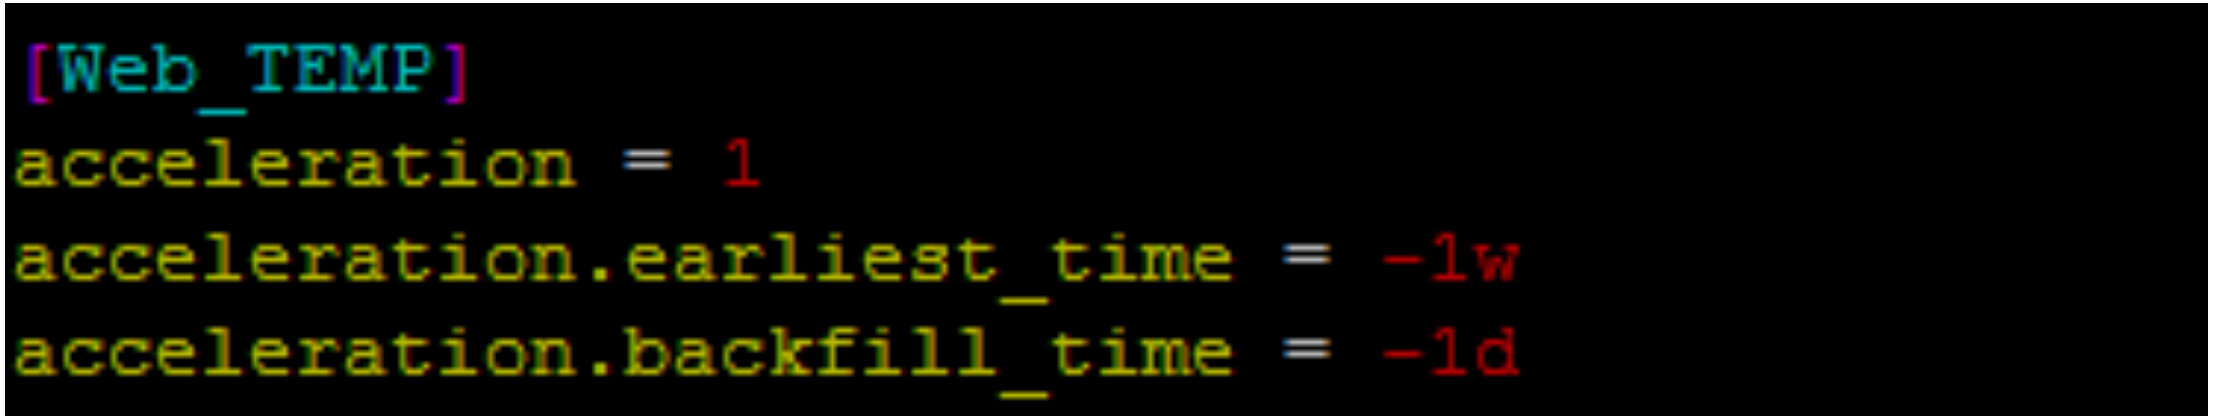

```
[Web_TEMP]
acceleration = 1
acceleration.earliest_time = -1w
acceleration.backfill_time = -1d
```

Limiting Data Model Backfilling

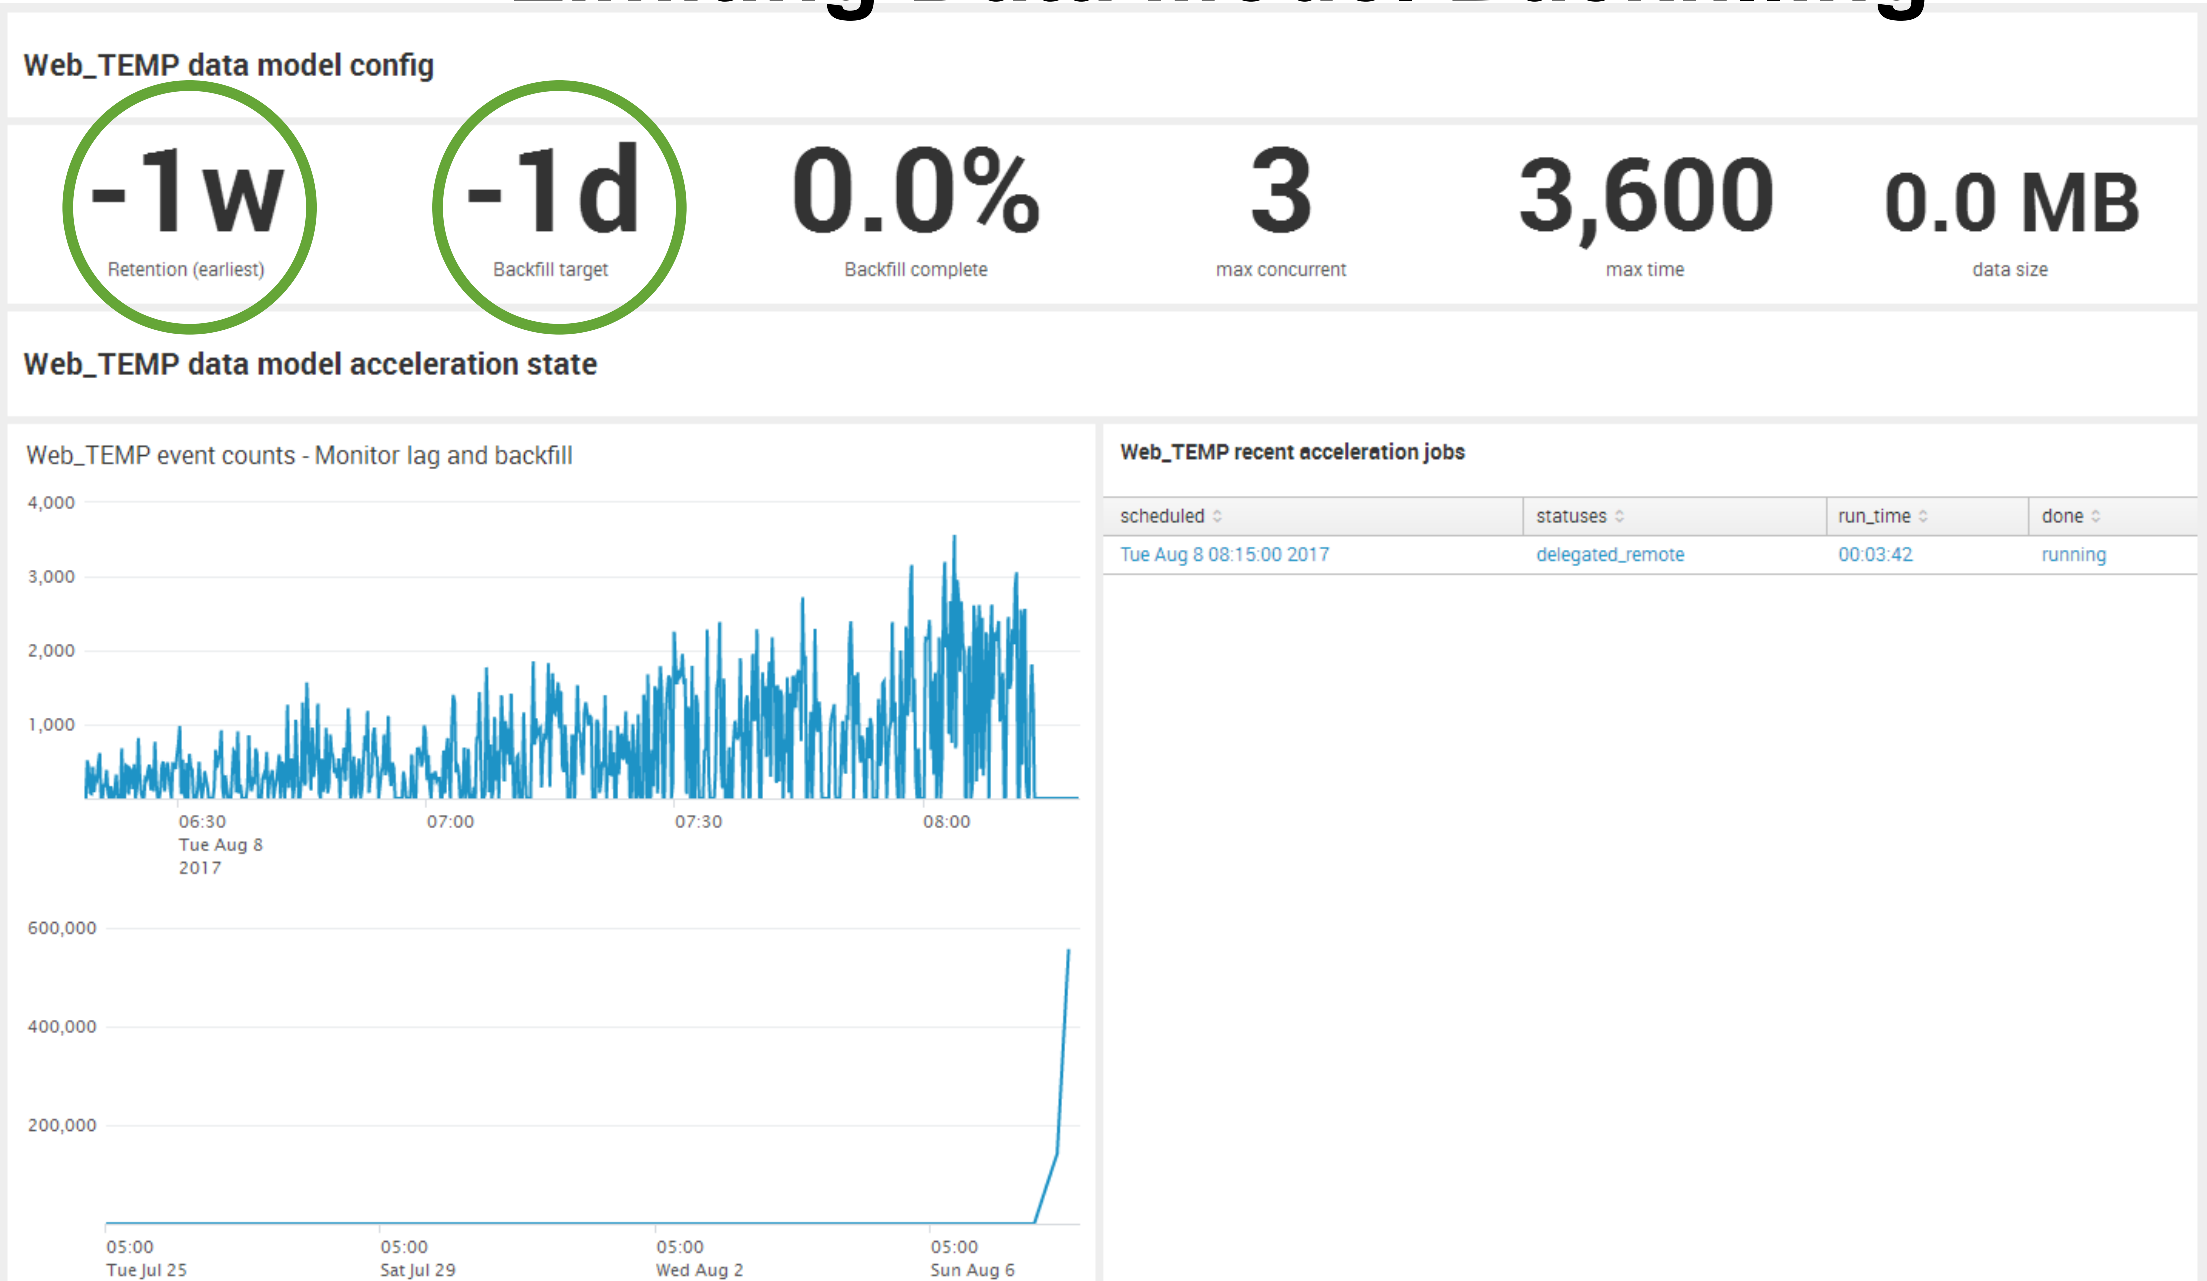

Web_TEMP data model config

-1w	-1d	0.0%	3	3,600	0.0 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 08:15:00 2017	delegated_remote	00:03:42	running

Limiting Data Model Backfilling

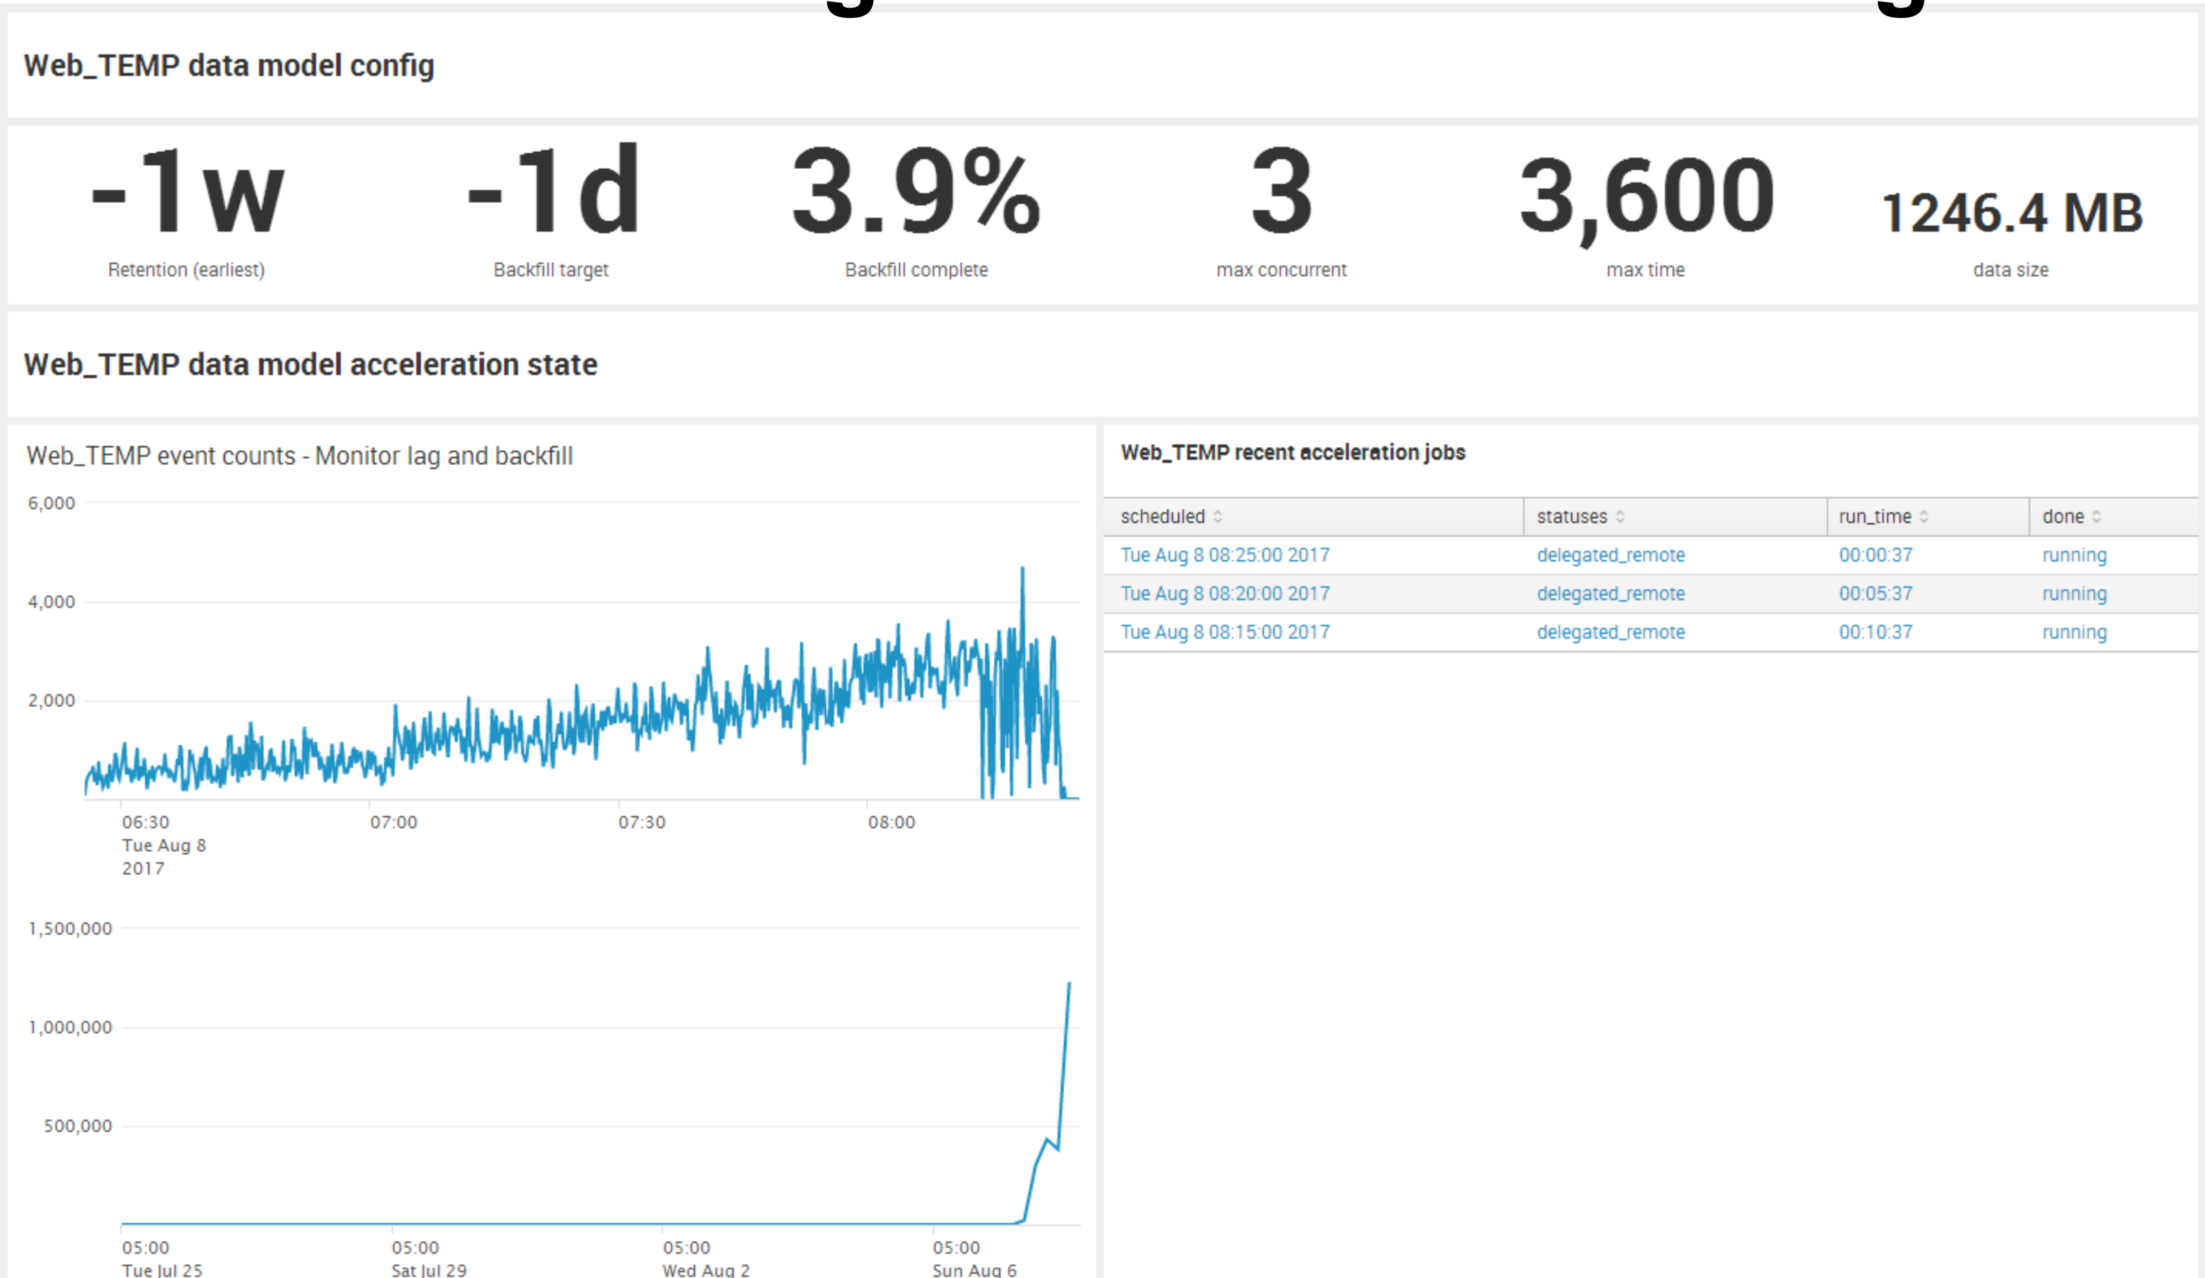

Web_TEMP data model config

-1w	-1d	3.9%	3	3,600	1246.4 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 08:25:00 2017	delegated_remote	00:00:37	running
Tue Aug 8 08:20:00 2017	delegated_remote	00:05:37	running
Tue Aug 8 08:15:00 2017	delegated_remote	00:10:37	running

Limiting Data Model Backfilling

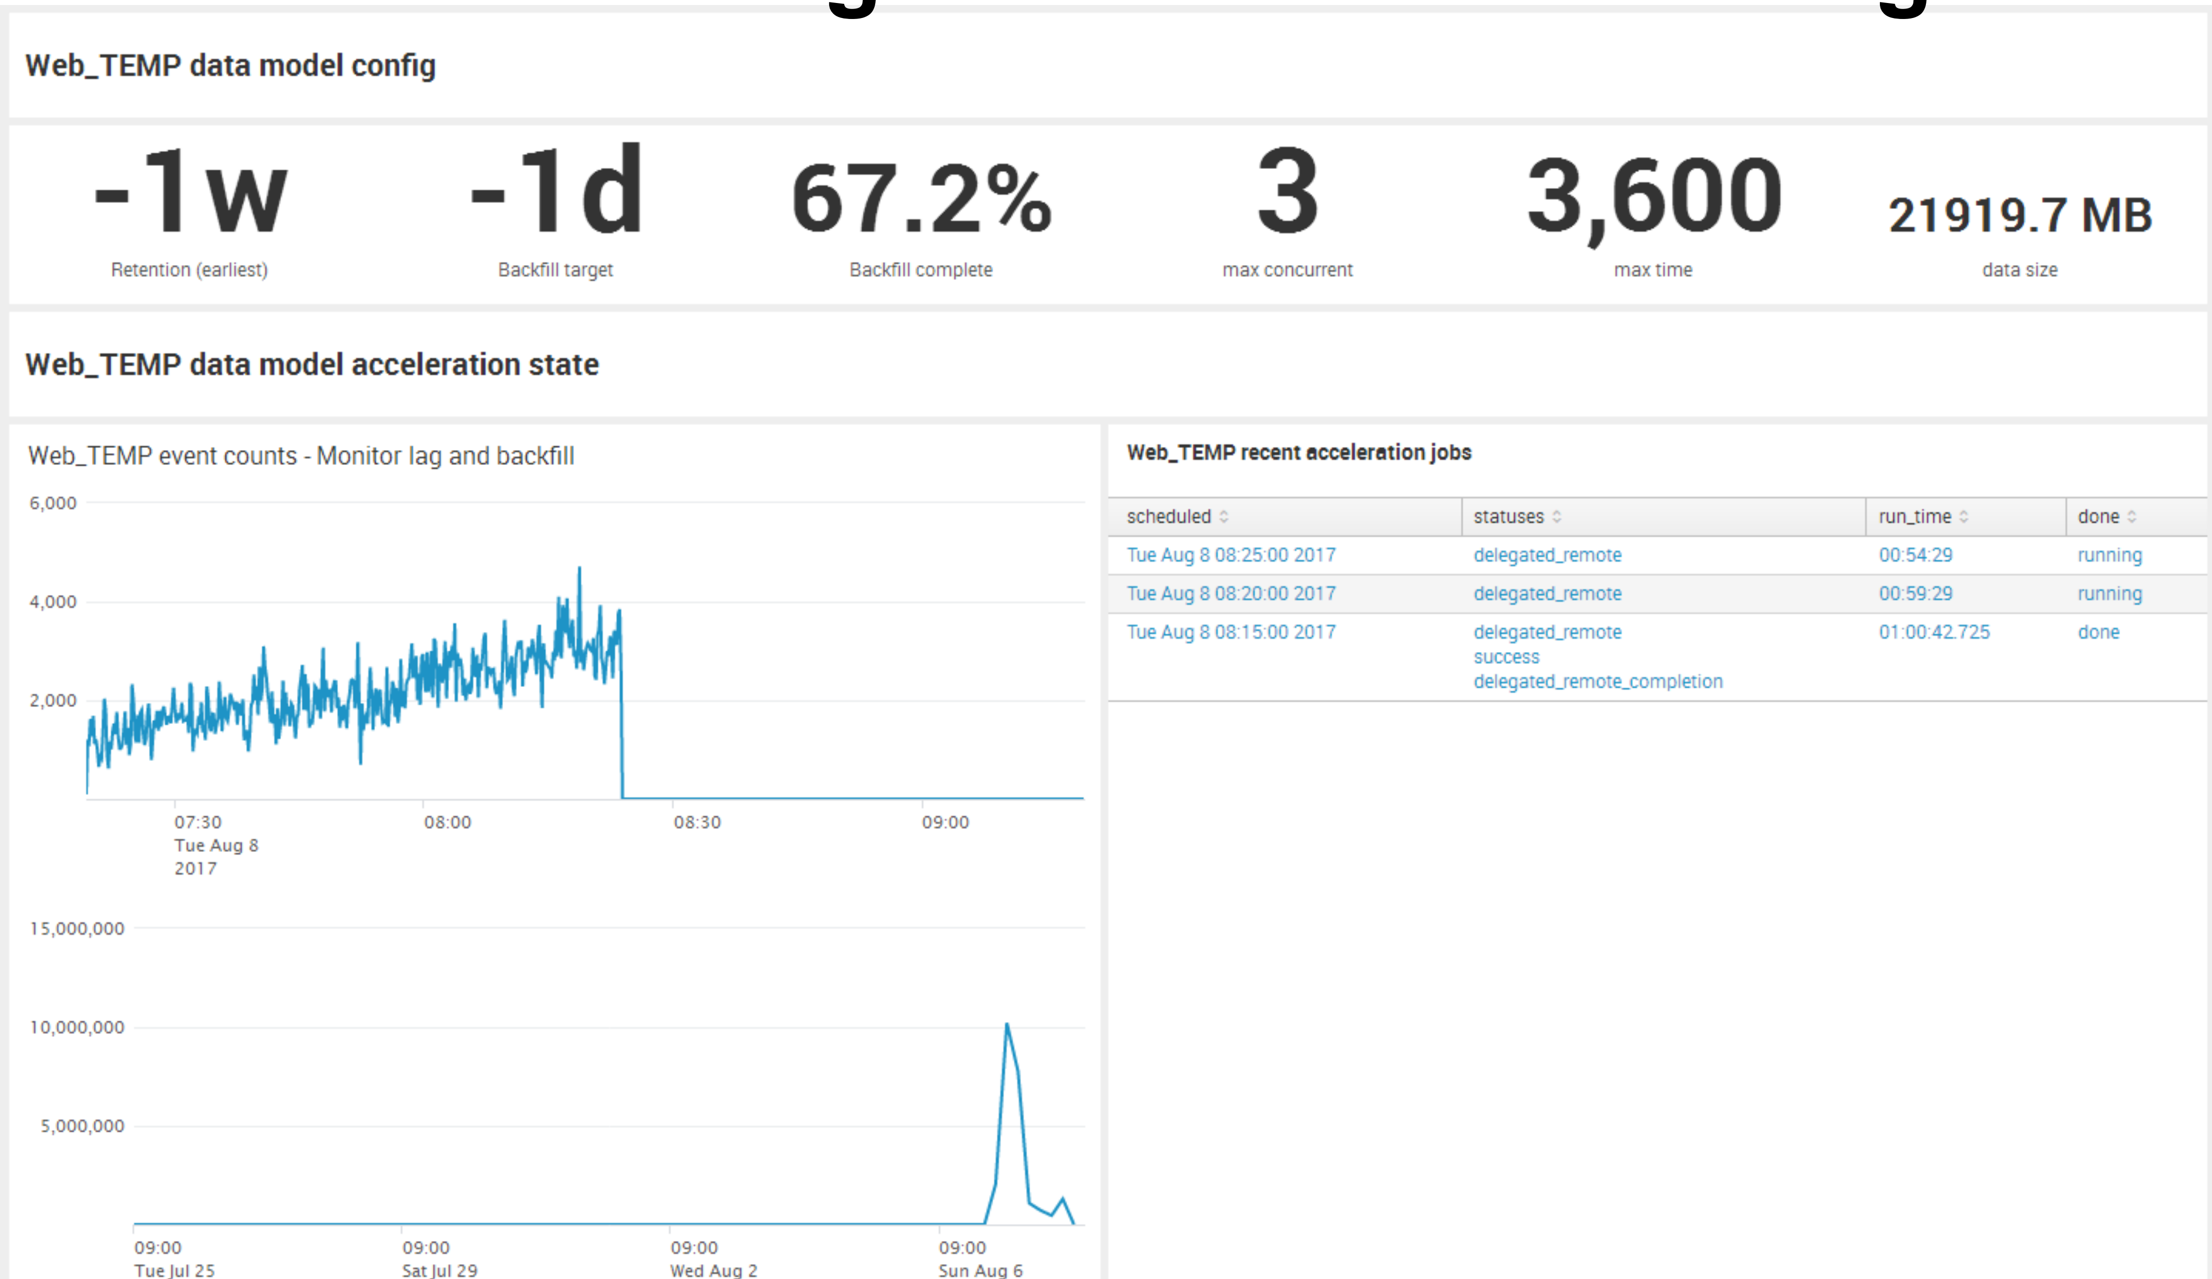

Web_TEMP data model config

-1w	-1d	67.2%	3	3,600	21919.7 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 08:25:00 2017	delegated_remote	00:54:29	running
Tue Aug 8 08:20:00 2017	delegated_remote	00:59:29	running
Tue Aug 8 08:15:00 2017	delegated_remote success delegated_remote_completion	01:00:42.725	done

Limiting Data Model Backfilling

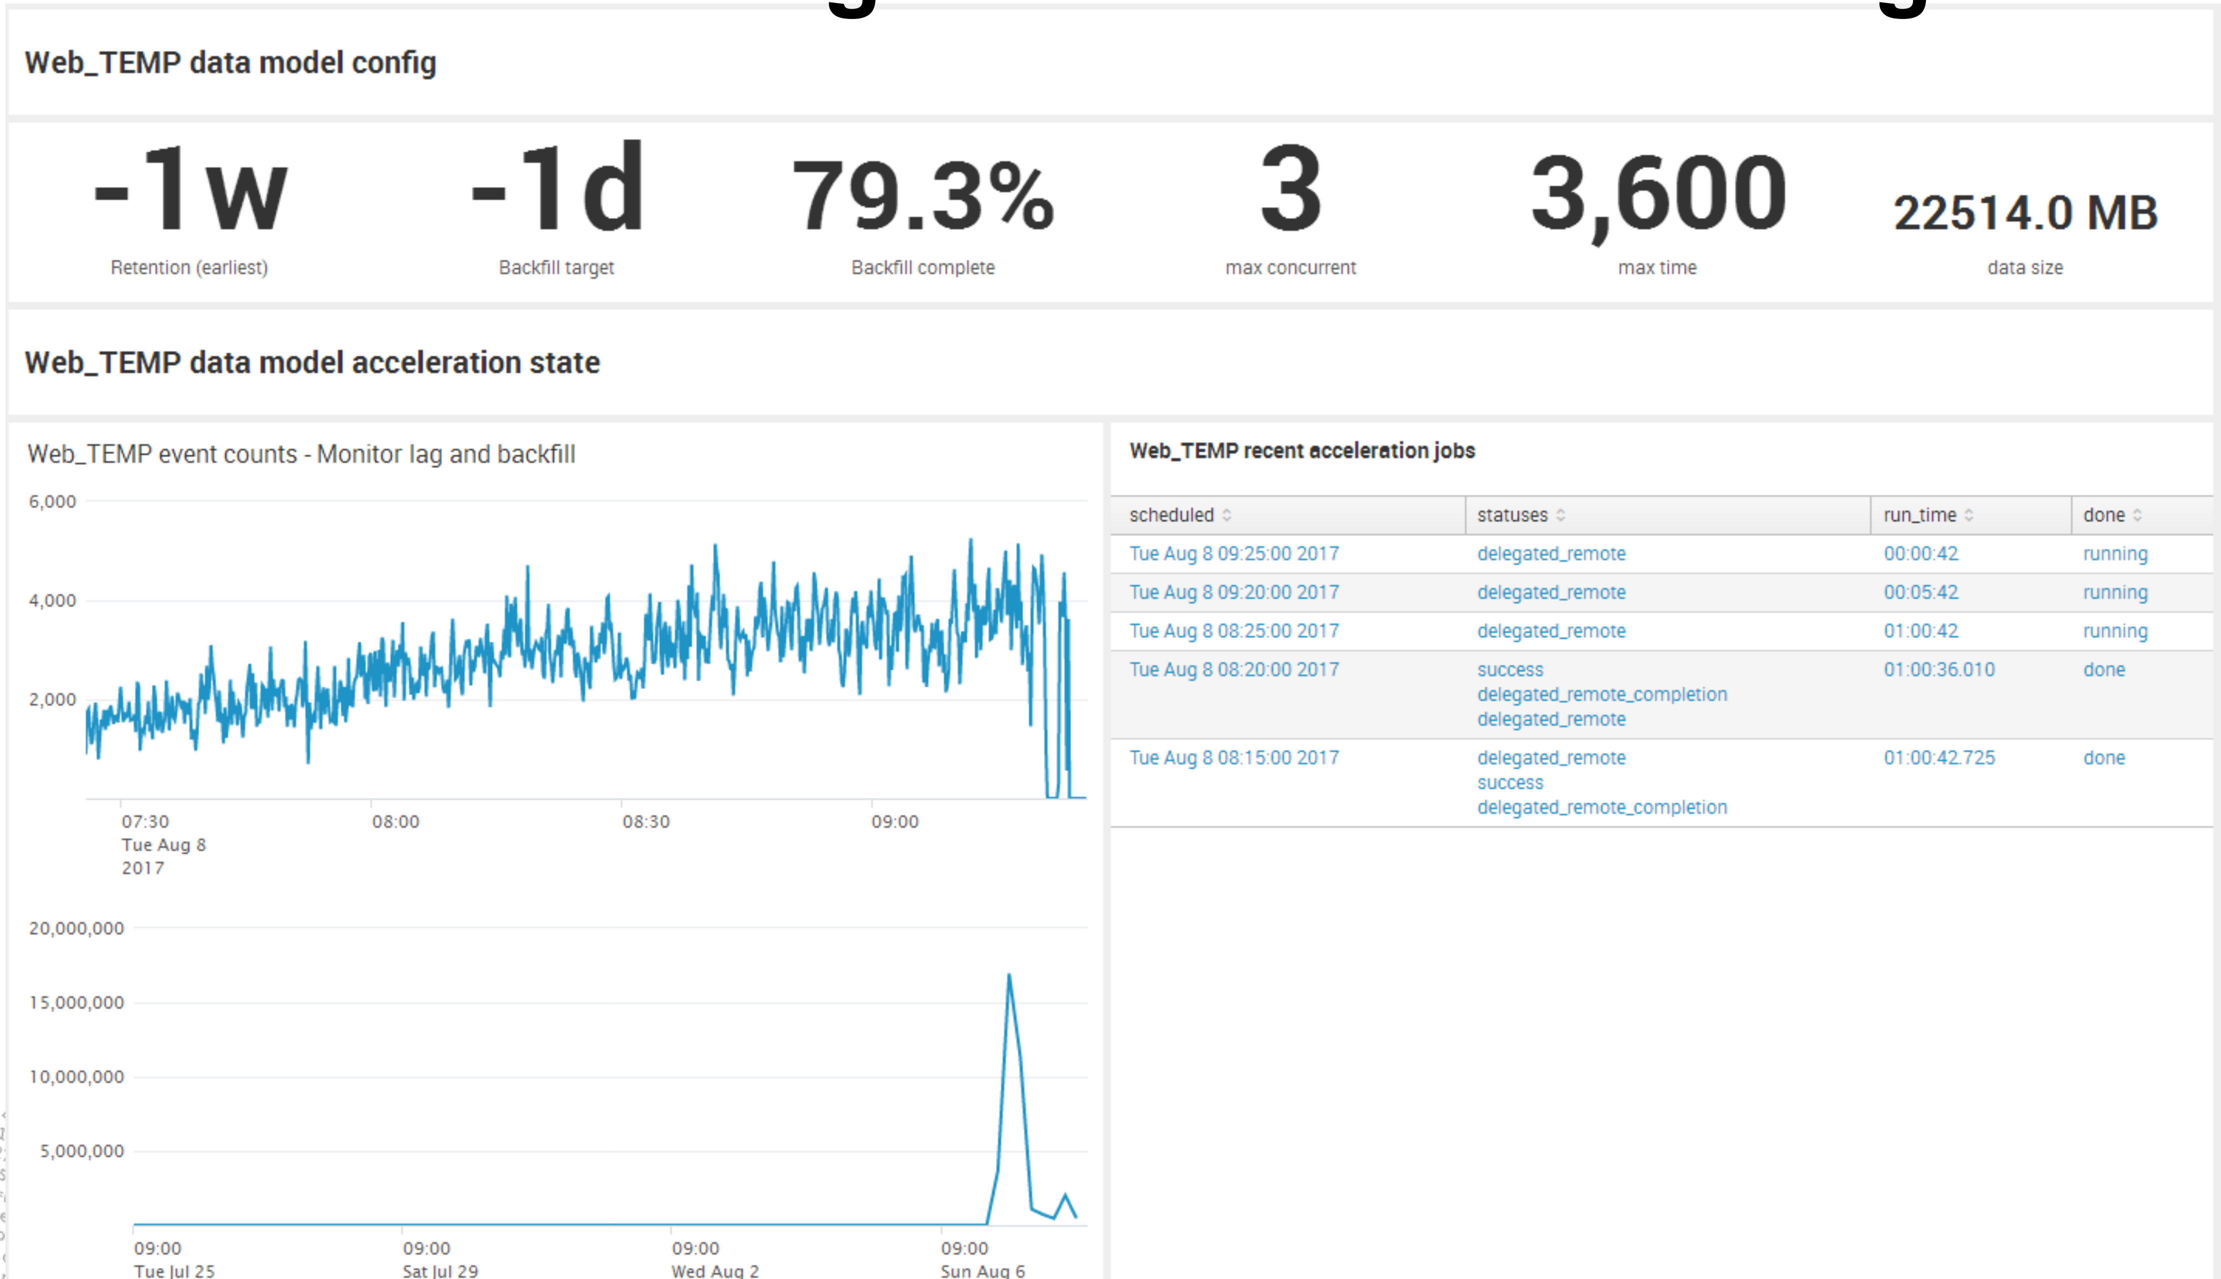

Web_TEMP data model config

-1w	-1d	79.3%	3	3,600	22514.0 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 09:25:00 2017	delegated_remote	00:00:42	running
Tue Aug 8 09:20:00 2017	delegated_remote	00:05:42	running
Tue Aug 8 08:25:00 2017	delegated_remote	01:00:42	running
Tue Aug 8 08:20:00 2017	success delegated_remote_completion delegated_remote	01:00:36.010	done
Tue Aug 8 08:15:00 2017	delegated_remote success delegated_remote_completion	01:00:42.725	done

Limiting Data Model Backfilling

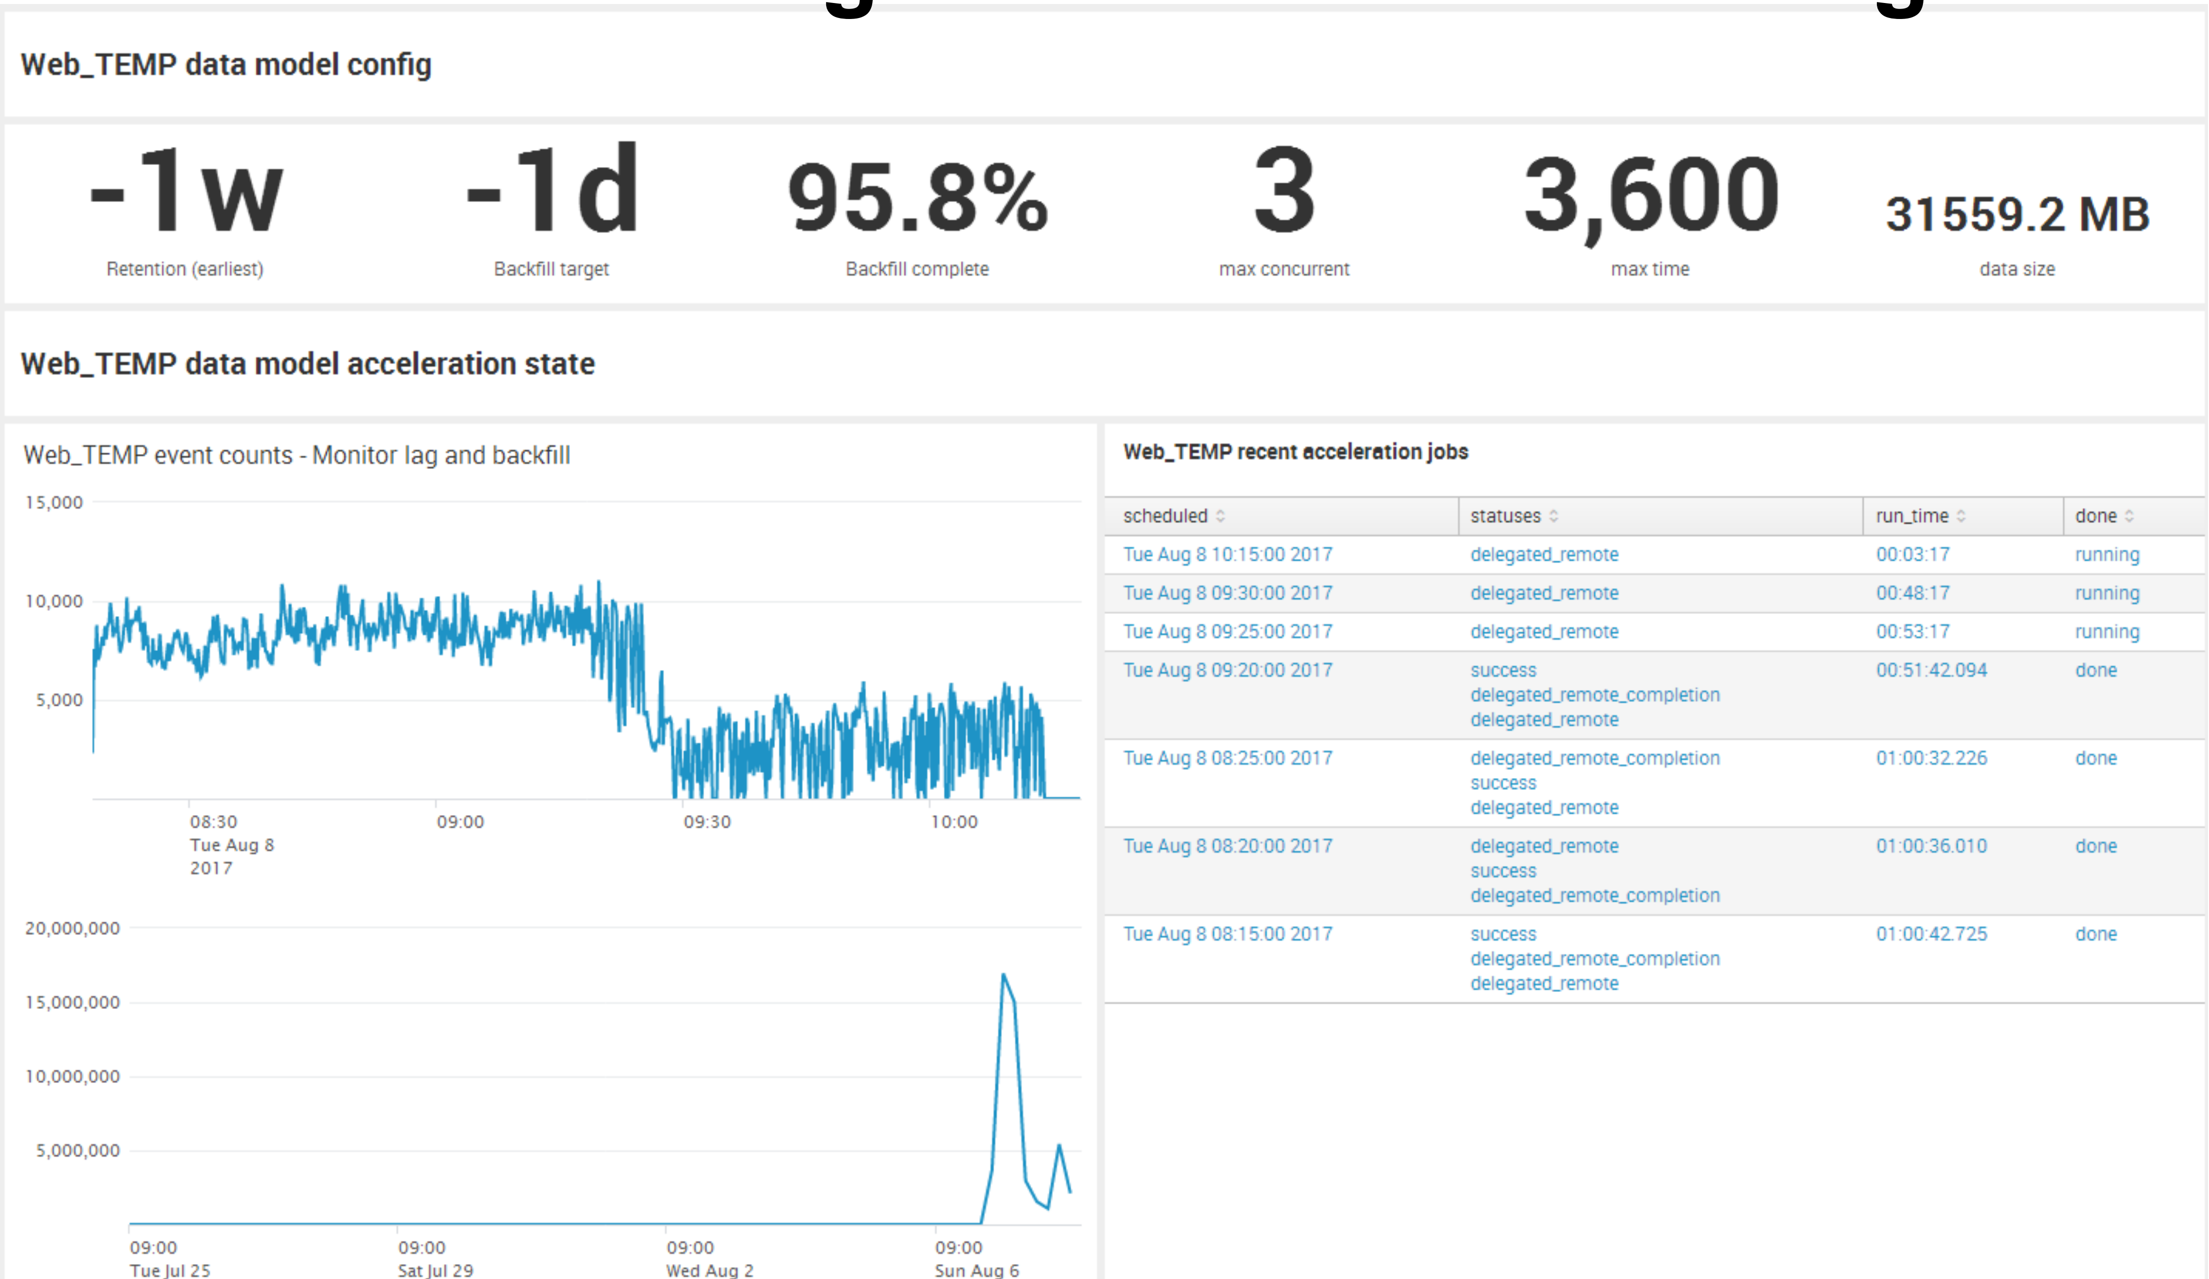

Web_TEMP data model config

-1w	-1d	95.8%	3	3,600	31559.2 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 10:15:00 2017	delegated_remote	00:03:17	running
Tue Aug 8 09:30:00 2017	delegated_remote	00:48:17	running
Tue Aug 8 09:25:00 2017	delegated_remote	00:53:17	running
Tue Aug 8 09:20:00 2017	success delegated_remote_completion delegated_remote	00:51:42.094	done
Tue Aug 8 08:25:00 2017	delegated_remote_completion success delegated_remote	01:00:32.226	done
Tue Aug 8 08:20:00 2017	delegated_remote success delegated_remote_completion	01:00:36.010	done
Tue Aug 8 08:15:00 2017	success delegated_remote_completion delegated_remote	01:00:42.725	done

Limiting Data Model Backfilling

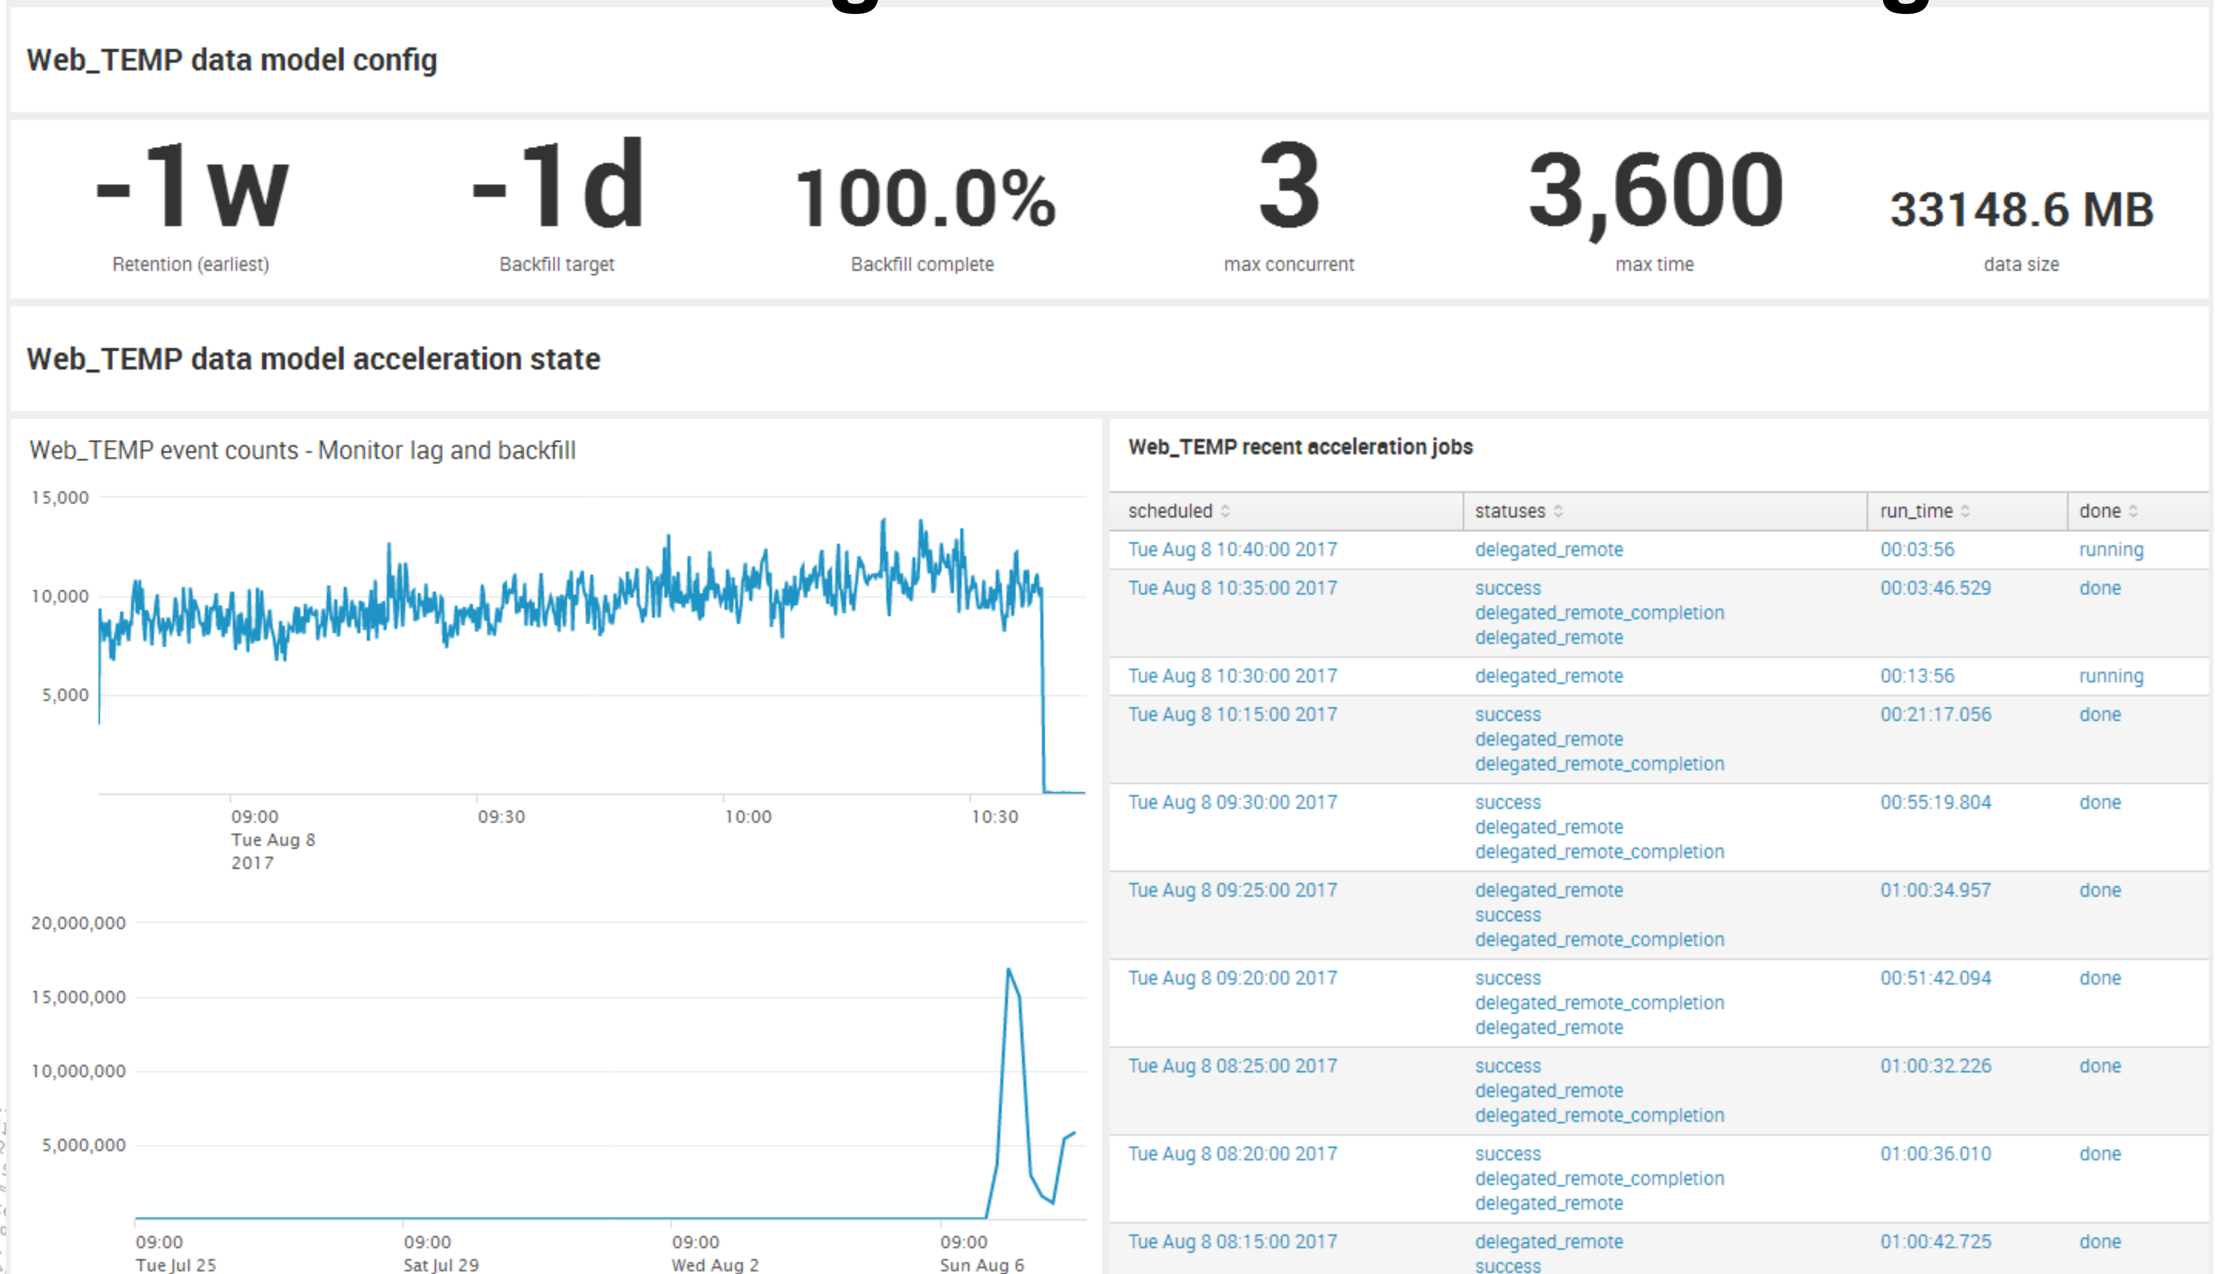

Web_TEMP data model config

-1w	-1d	100.0%	3	3,600	33148.6 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 10:40:00 2017	delegated_remote	00:03:56	running
Tue Aug 8 10:35:00 2017	success delegated_remote_completion delegated_remote	00:03:46.529	done
Tue Aug 8 10:30:00 2017	delegated_remote	00:13:56	running
Tue Aug 8 10:15:00 2017	success delegated_remote delegated_remote_completion	00:21:17.056	done
Tue Aug 8 09:30:00 2017	success delegated_remote delegated_remote_completion	00:55:19.804	done
Tue Aug 8 09:25:00 2017	delegated_remote success delegated_remote_completion	01:00:34.957	done
Tue Aug 8 09:20:00 2017	success delegated_remote_completion delegated_remote	00:51:42.094	done
Tue Aug 8 08:25:00 2017	success delegated_remote delegated_remote_completion	01:00:32.226	done
Tue Aug 8 08:20:00 2017	success delegated_remote_completion delegated_remote	01:00:36.010	done
Tue Aug 8 08:15:00 2017	delegated_remote success	01:00:42.725	done

Backfilling Without Lag

- In datamodels.conf:

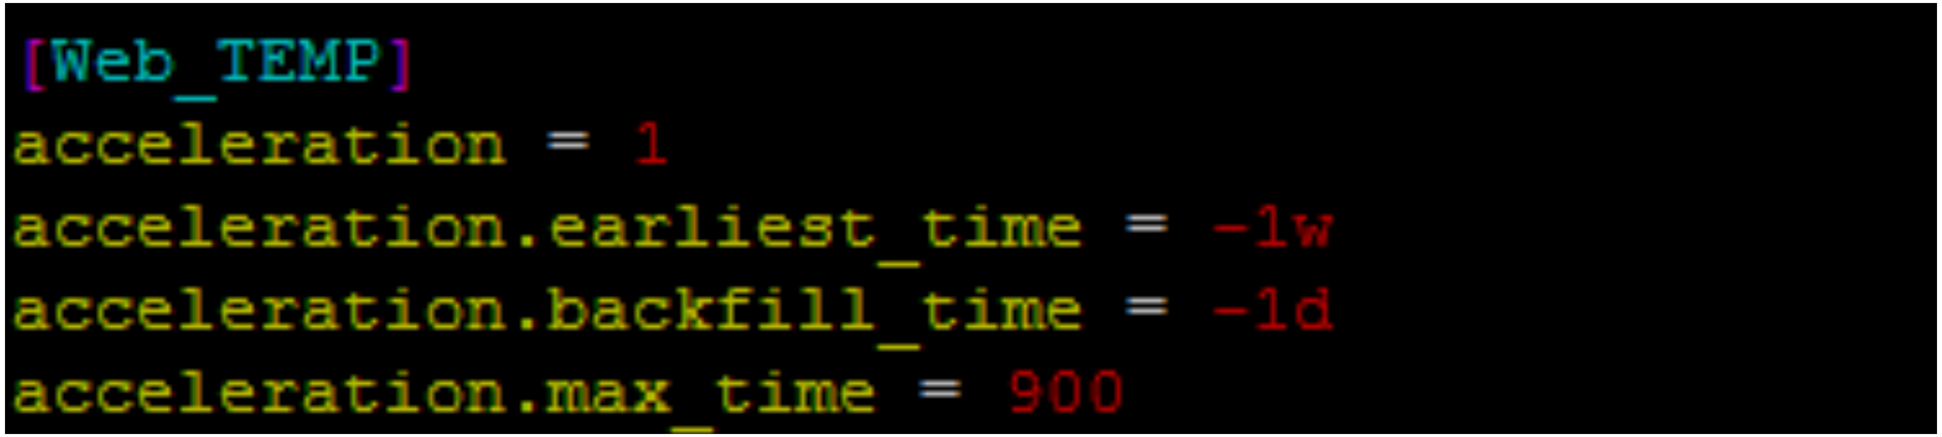

```
[Web_TEMP]
acceleration = 1
acceleration.earliest_time = -1w
acceleration.backfill_time = -1d
acceleration.max_time = 900
```

Backfilling Without Lag

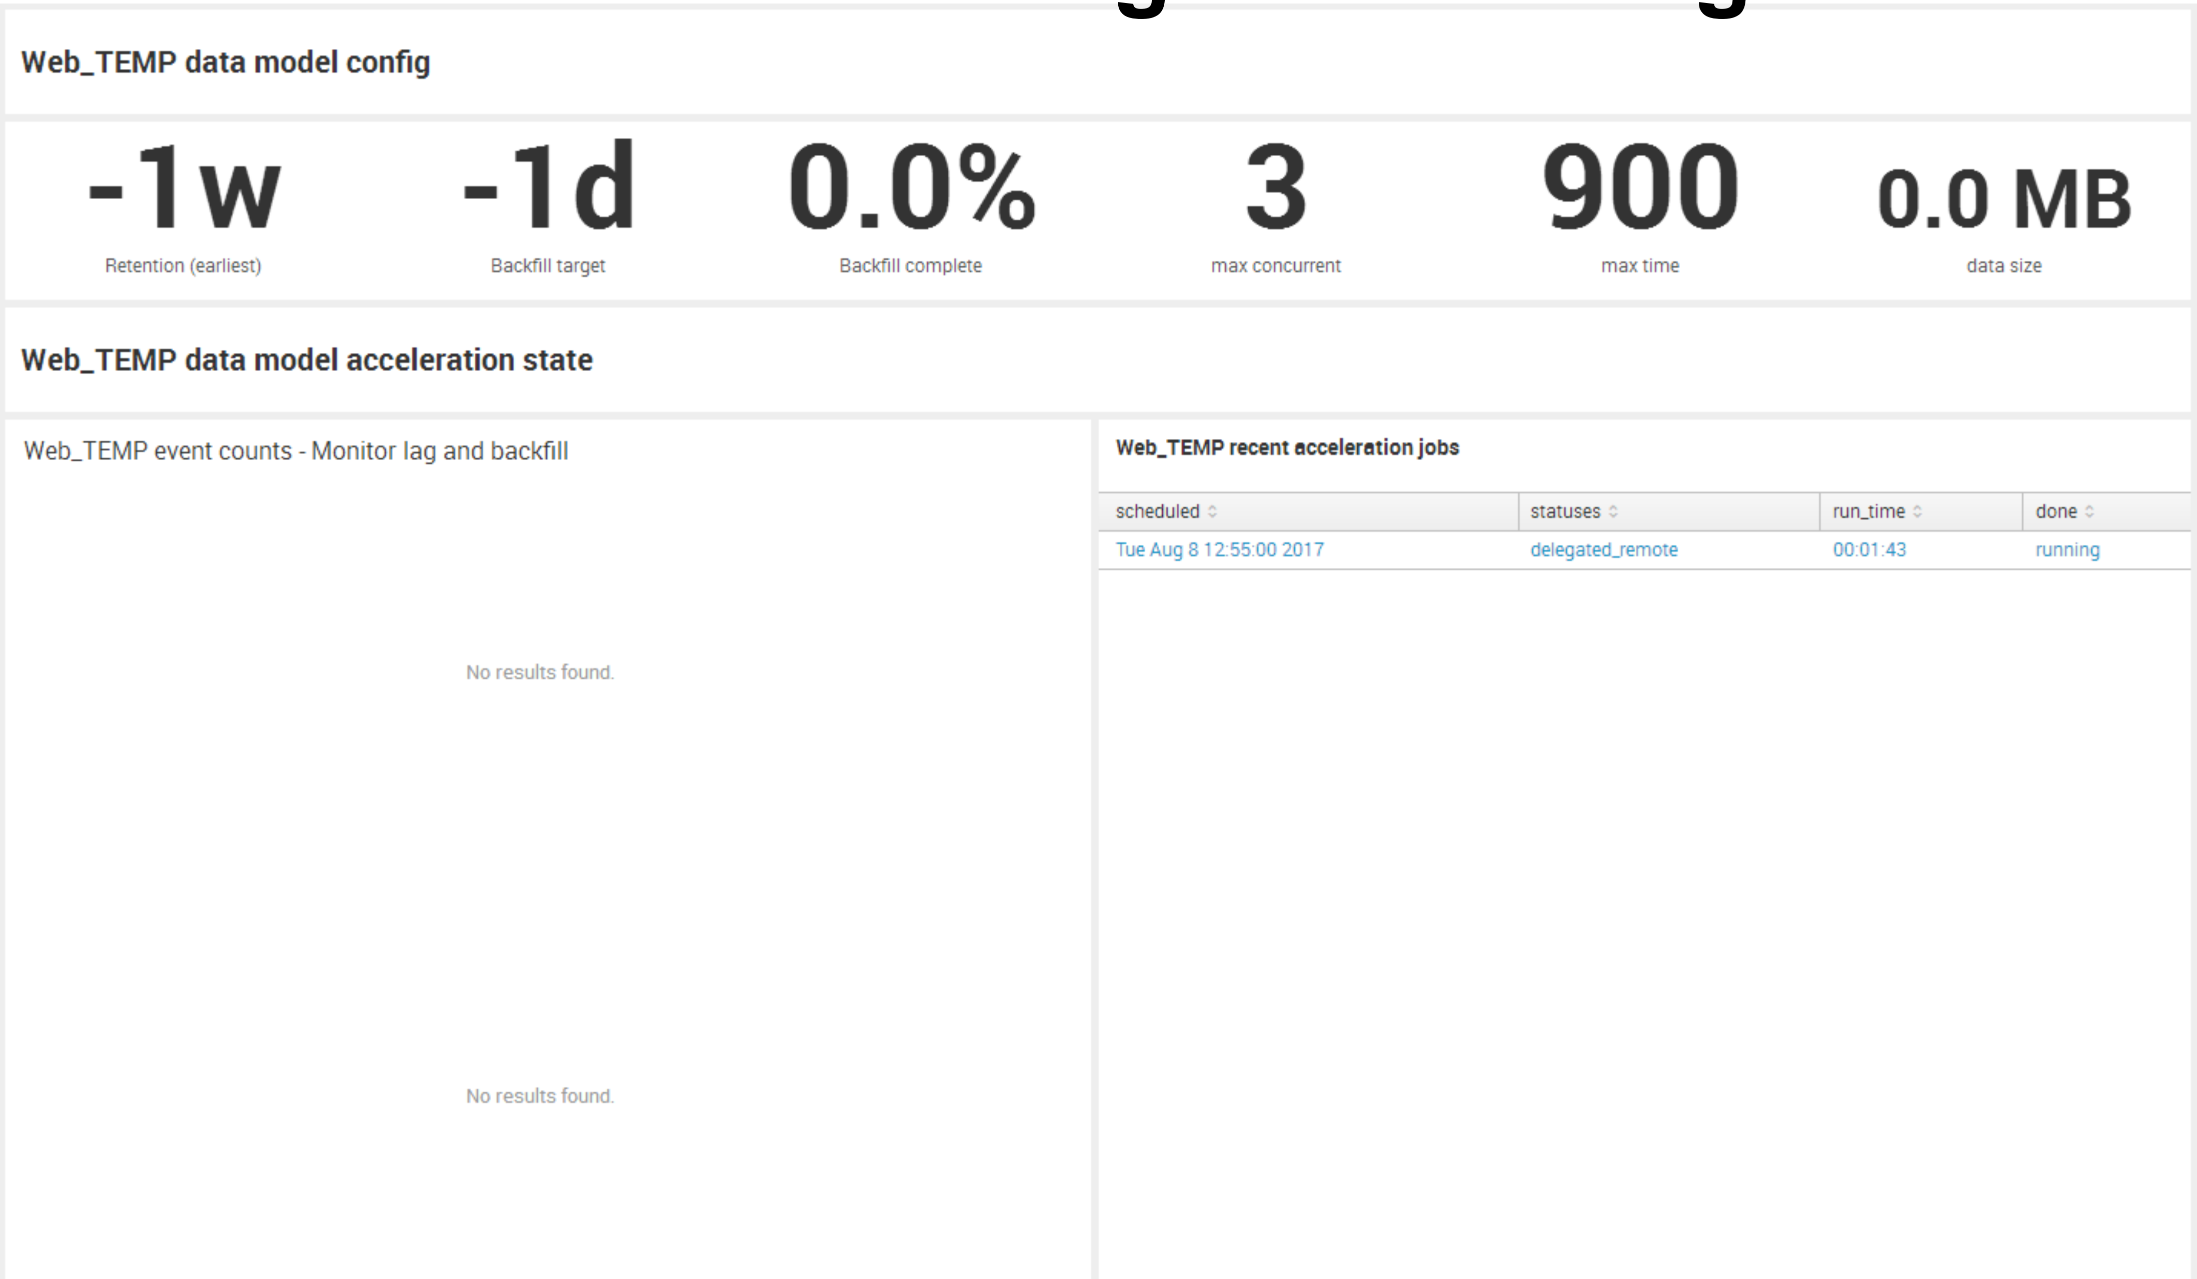

Web_TEMP data model config

-1w
Retention (earliest)

-1d
Backfill target

0.0%
Backfill complete

3
max concurrent

900
max time

0.0 MB
data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

No results found.

No results found.

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 12:55:00 2017	delegated_remote	00:01:43	running

.conf2017

Backfilling Without Lag

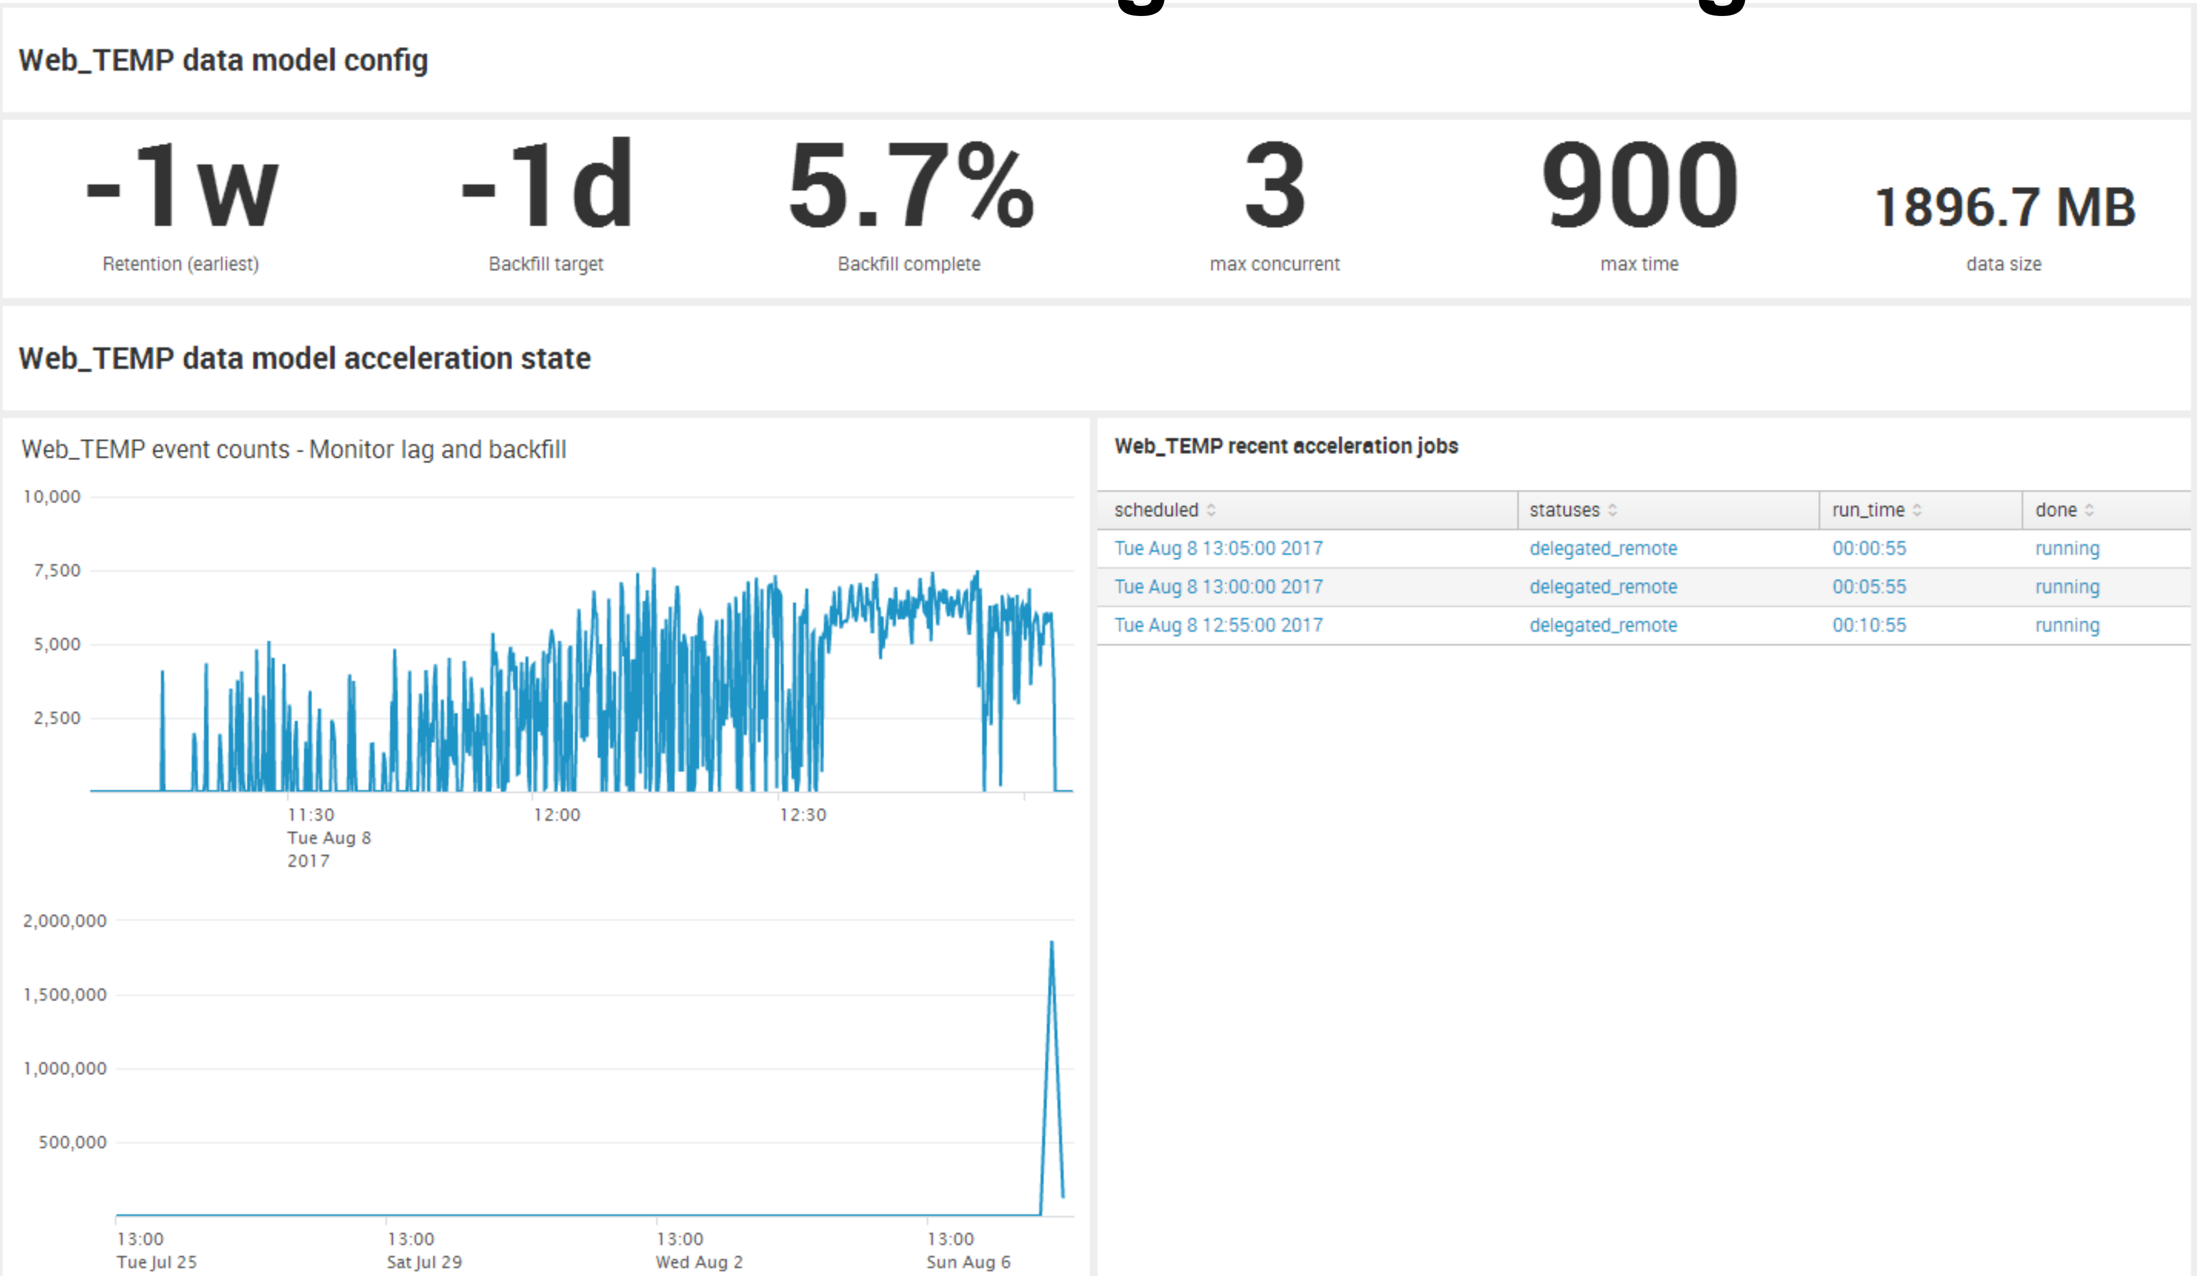

Web_TEMP data model config

-1w	-1d	5.7%	3	900	1896.7 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 13:05:00 2017	delegated_remote	00:00:55	running
Tue Aug 8 13:00:00 2017	delegated_remote	00:05:55	running
Tue Aug 8 12:55:00 2017	delegated_remote	00:10:55	running

Backfilling Without Lag

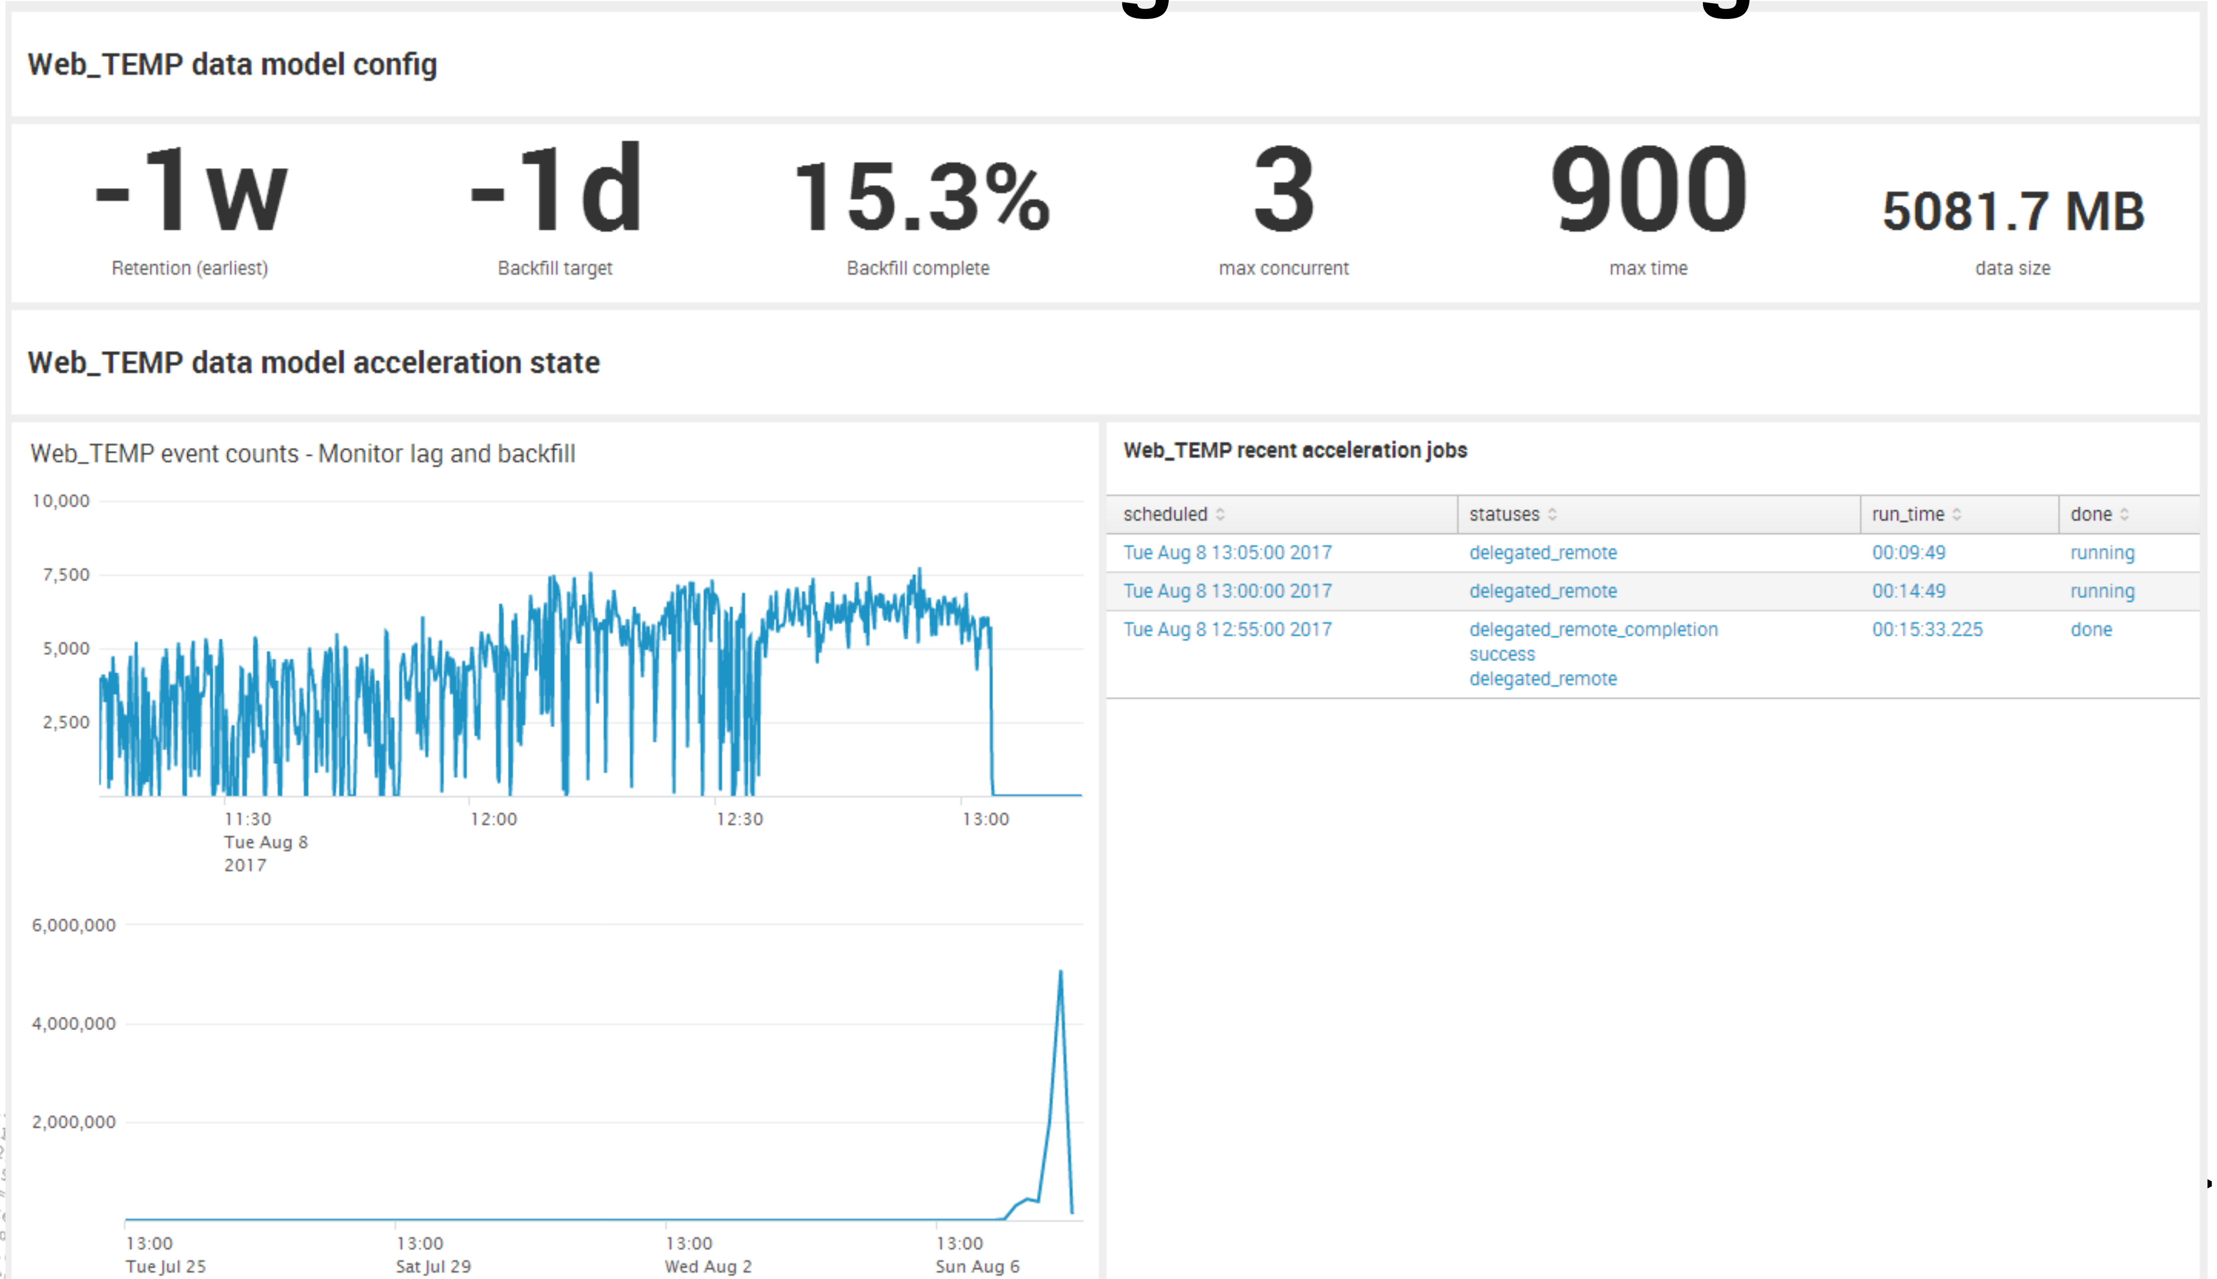

Web_TEMP data model config

-1w	-1d	15.3%	3	900	5081.7 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 13:05:00 2017	delegated_remote	00:09:49	running
Tue Aug 8 13:00:00 2017	delegated_remote	00:14:49	running
Tue Aug 8 12:55:00 2017	delegated_remote_completion success delegated_remote	00:15:33.225	done

Backfilling Without Lag

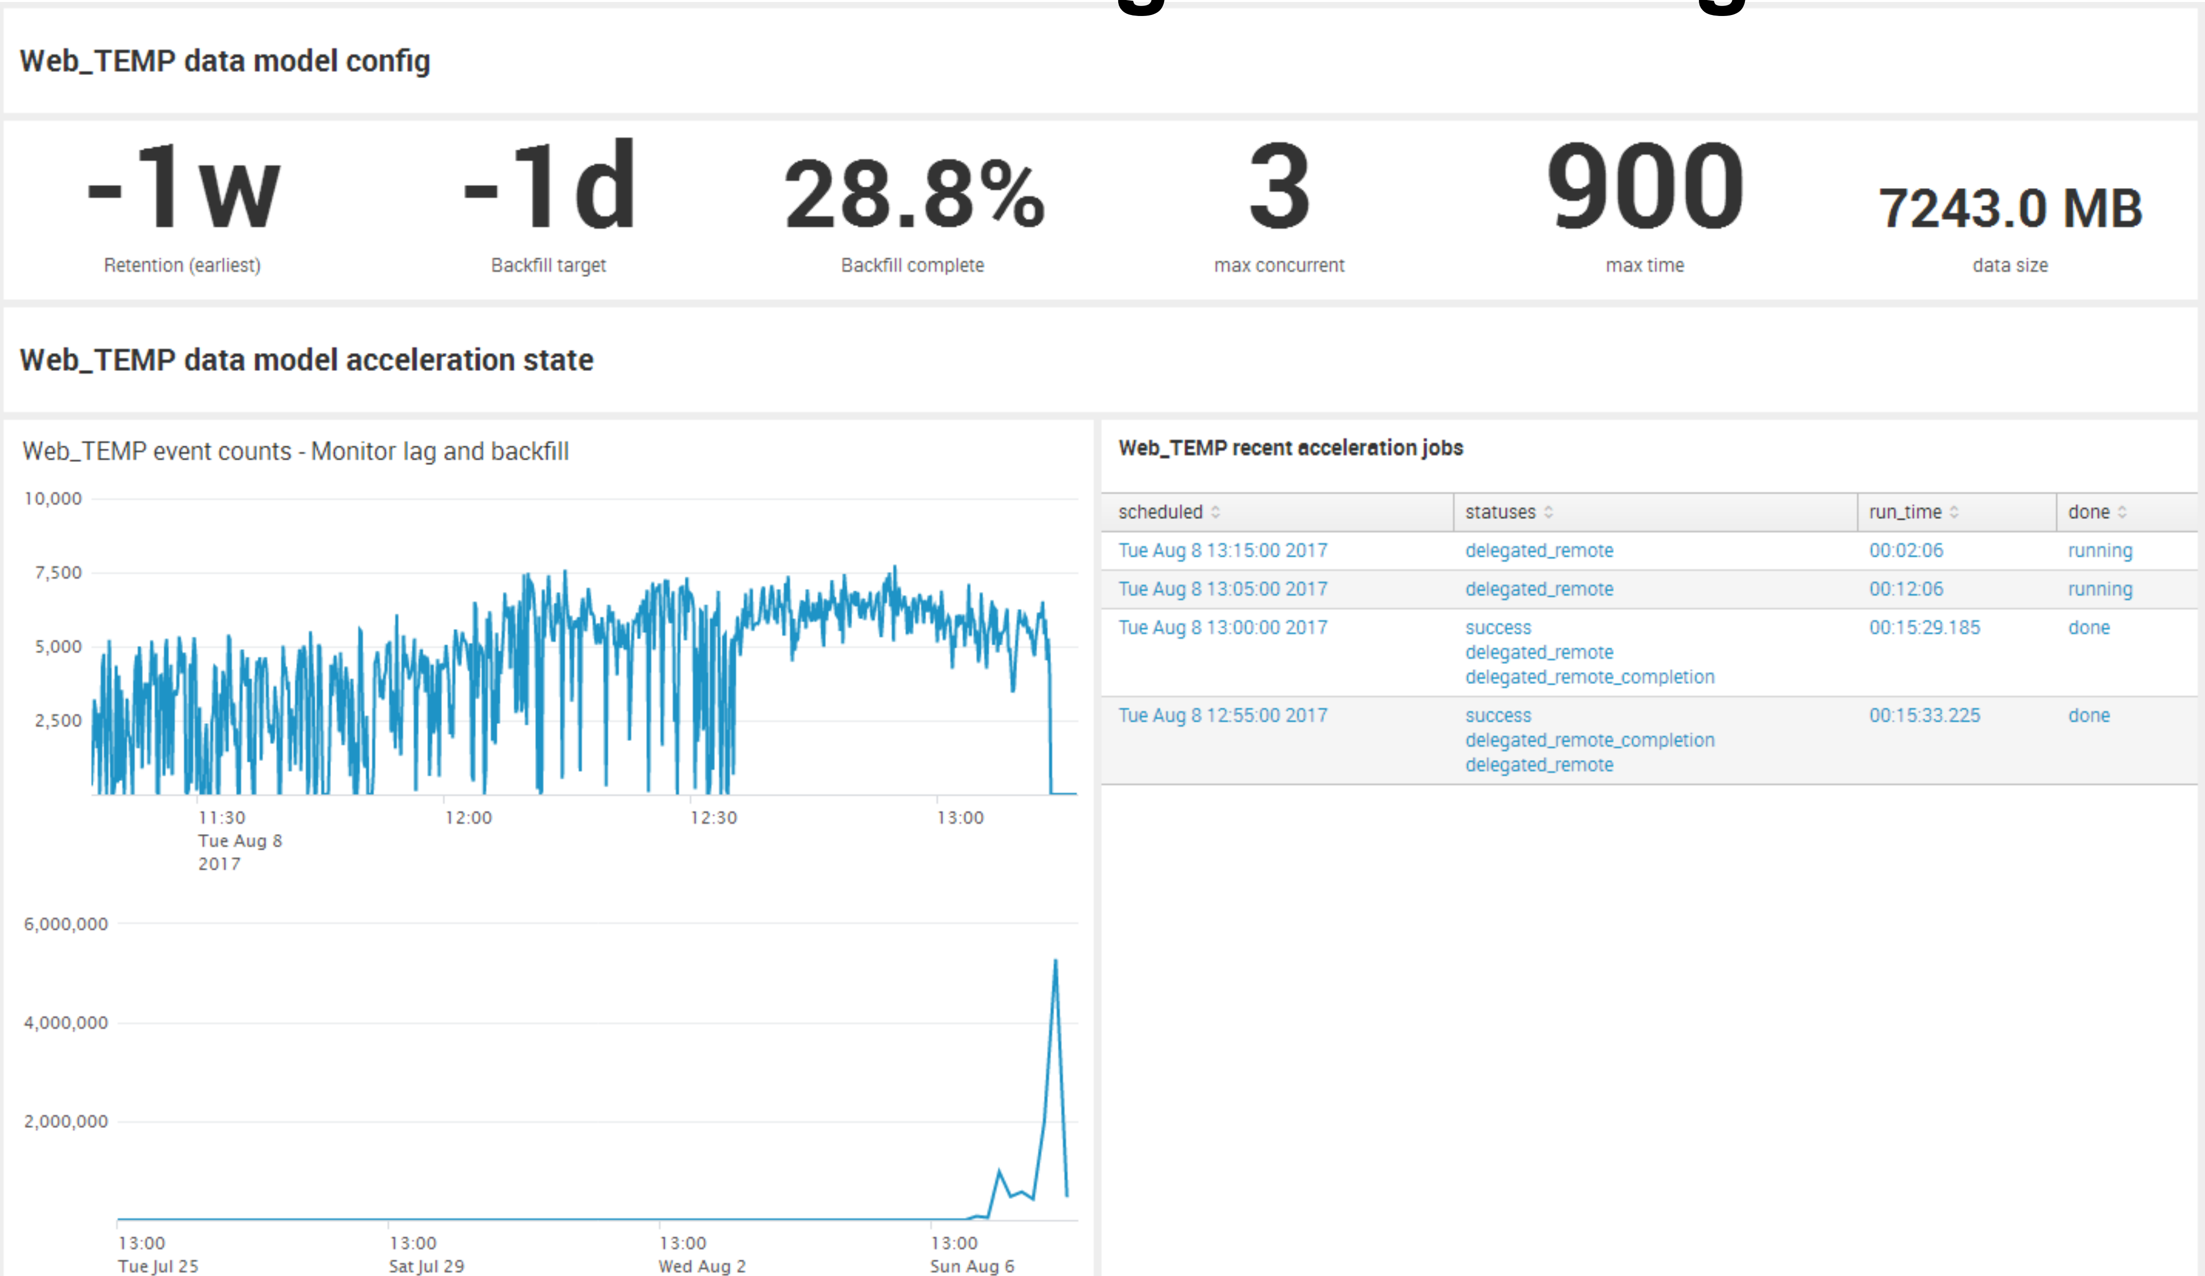

Web_TEMP data model config

-1w	-1d	28.8%	3	900	7243.0 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 13:15:00 2017	delegated_remote	00:02:06	running
Tue Aug 8 13:05:00 2017	delegated_remote	00:12:06	running
Tue Aug 8 13:00:00 2017	success delegated_remote delegated_remote_completion	00:15:29.185	done
Tue Aug 8 12:55:00 2017	success delegated_remote_completion delegated_remote	00:15:33.225	done

Backfilling Without Lag

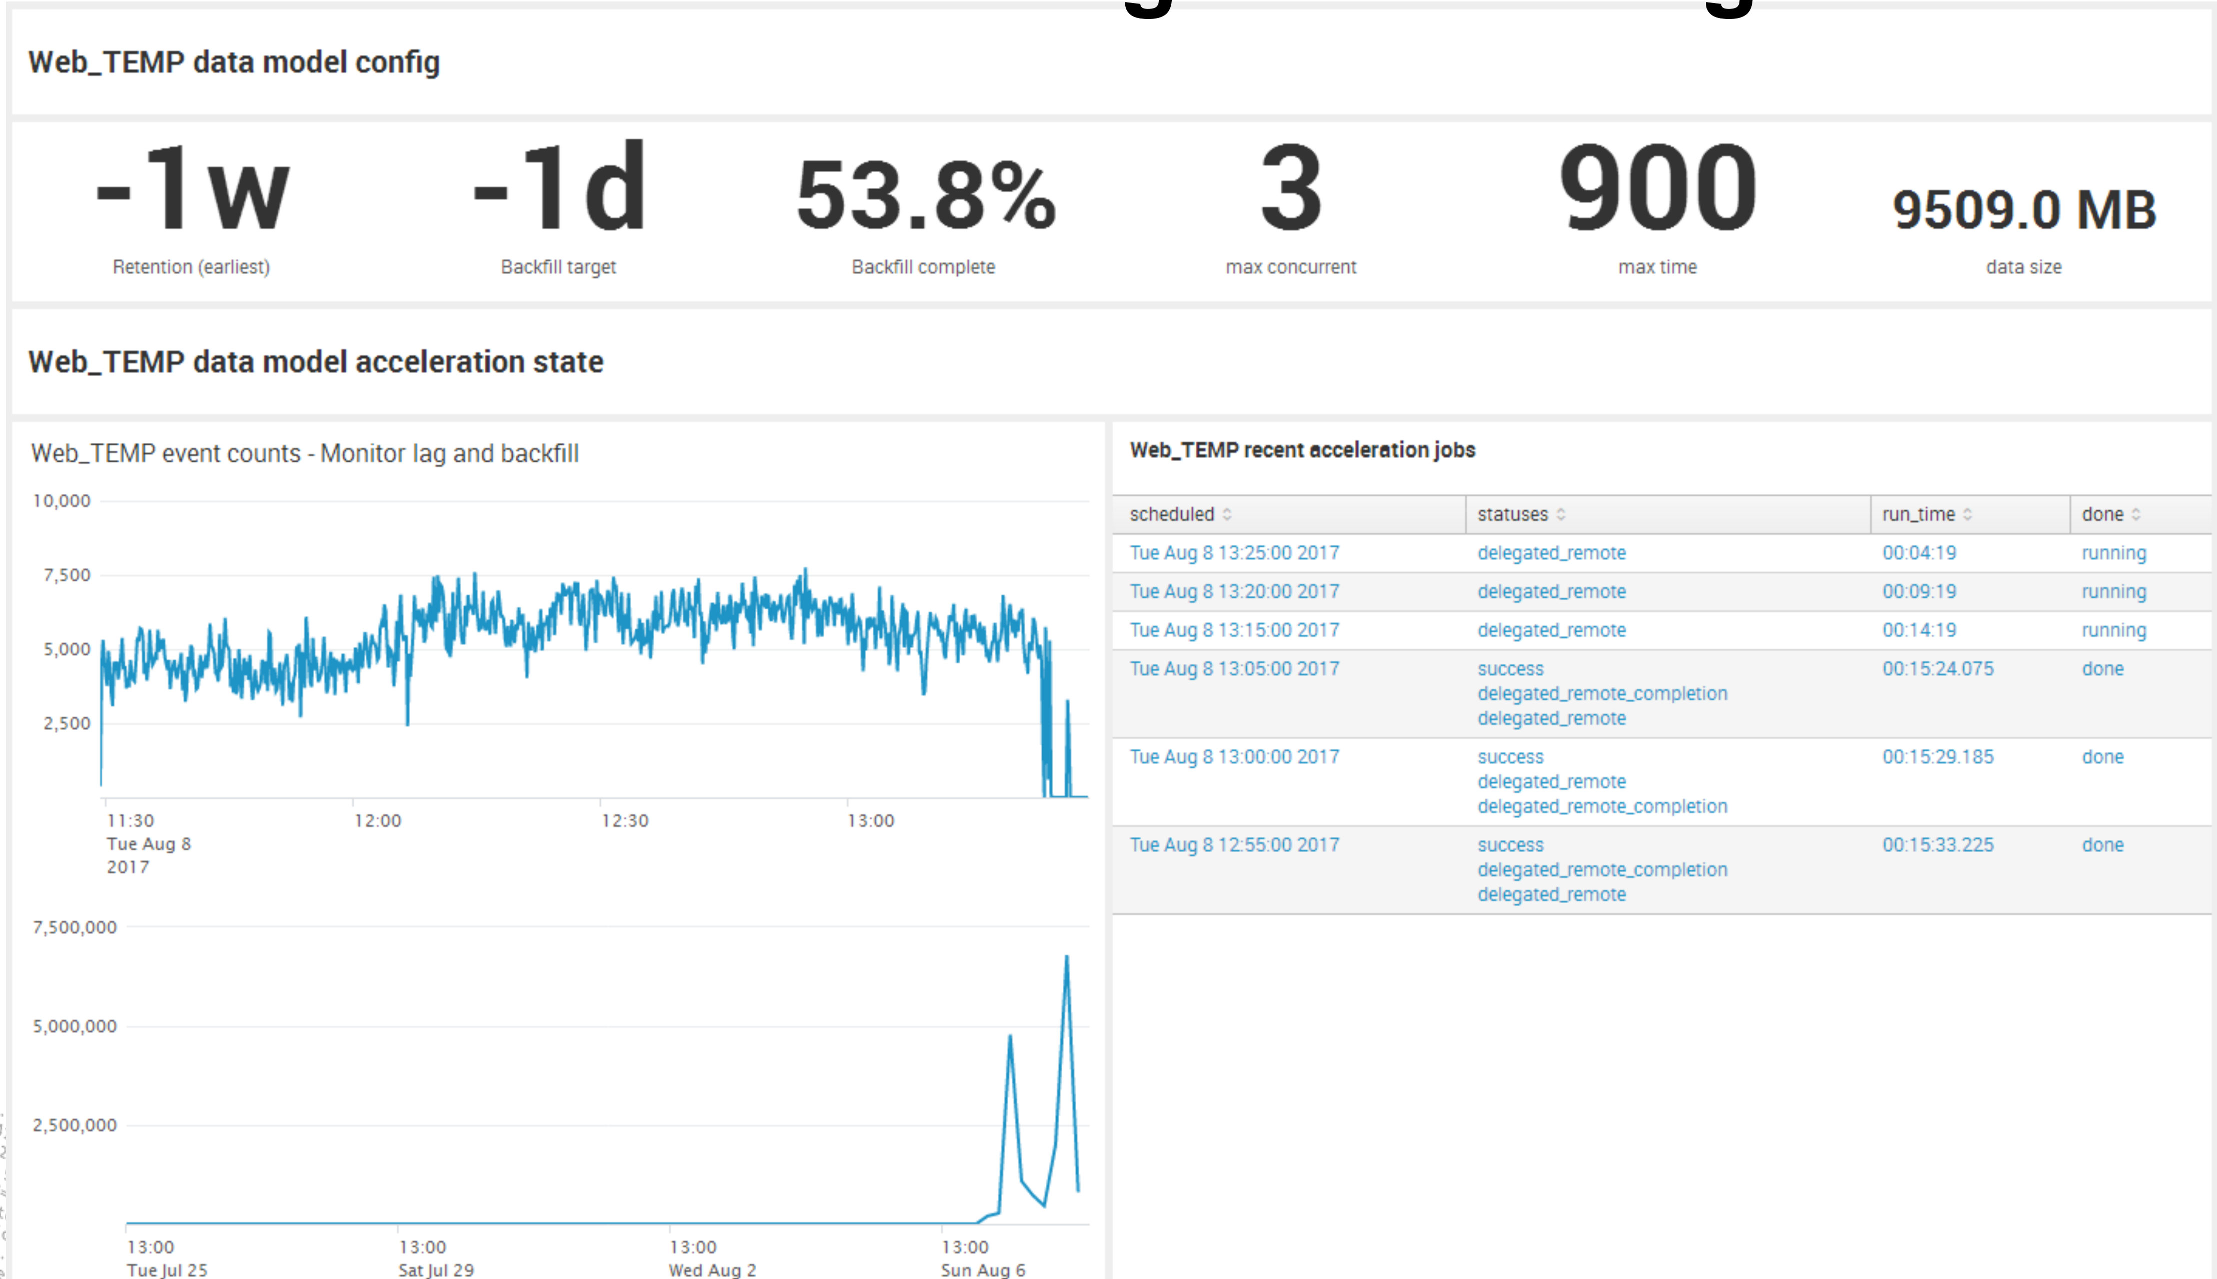

Web_TEMP data model config

-1w	-1d	53.8%	3	900	9509.0 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 13:25:00 2017	delegated_remote	00:04:19	running
Tue Aug 8 13:20:00 2017	delegated_remote	00:09:19	running
Tue Aug 8 13:15:00 2017	delegated_remote	00:14:19	running
Tue Aug 8 13:05:00 2017	success delegated_remote_completion delegated_remote	00:15:24.075	done
Tue Aug 8 13:00:00 2017	success delegated_remote delegated_remote_completion	00:15:29.185	done
Tue Aug 8 12:55:00 2017	success delegated_remote_completion delegated_remote	00:15:33.225	done

Backfilling Without Lag

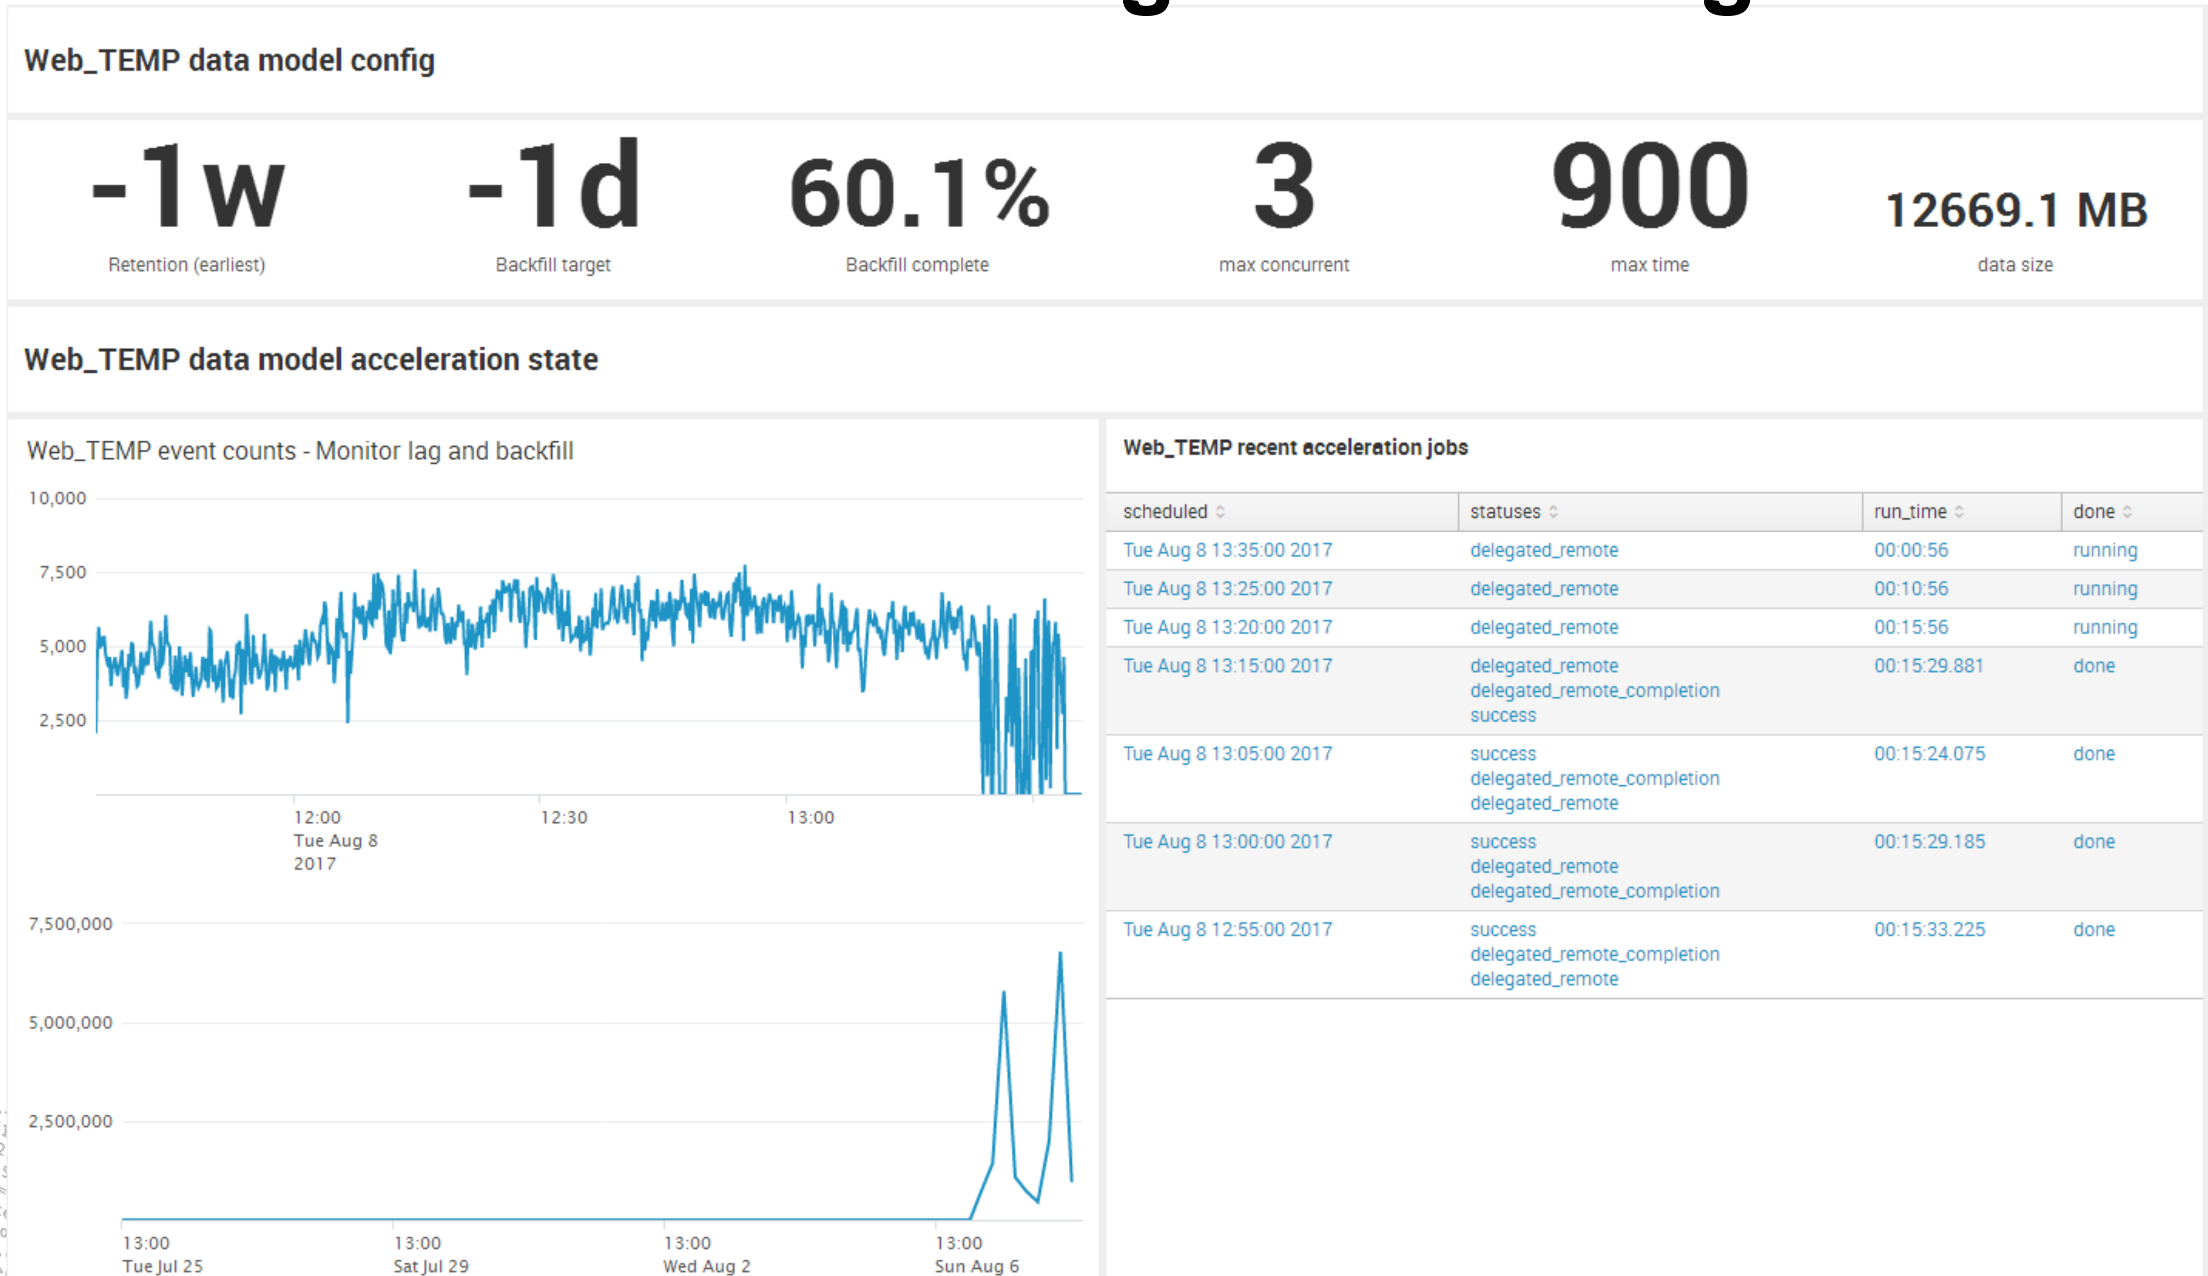

Web_TEMP data model config

-1w	-1d	60.1%	3	900	12669.1 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 13:35:00 2017	delegated_remote	00:00:56	running
Tue Aug 8 13:25:00 2017	delegated_remote	00:10:56	running
Tue Aug 8 13:20:00 2017	delegated_remote	00:15:56	running
Tue Aug 8 13:15:00 2017	delegated_remote delegated_remote_completion success	00:15:29.881	done
Tue Aug 8 13:05:00 2017	success delegated_remote_completion delegated_remote	00:15:24.075	done
Tue Aug 8 13:00:00 2017	success delegated_remote delegated_remote_completion	00:15:29.185	done
Tue Aug 8 12:55:00 2017	success delegated_remote_completion delegated_remote	00:15:33.225	done

Backfilling Without Lag

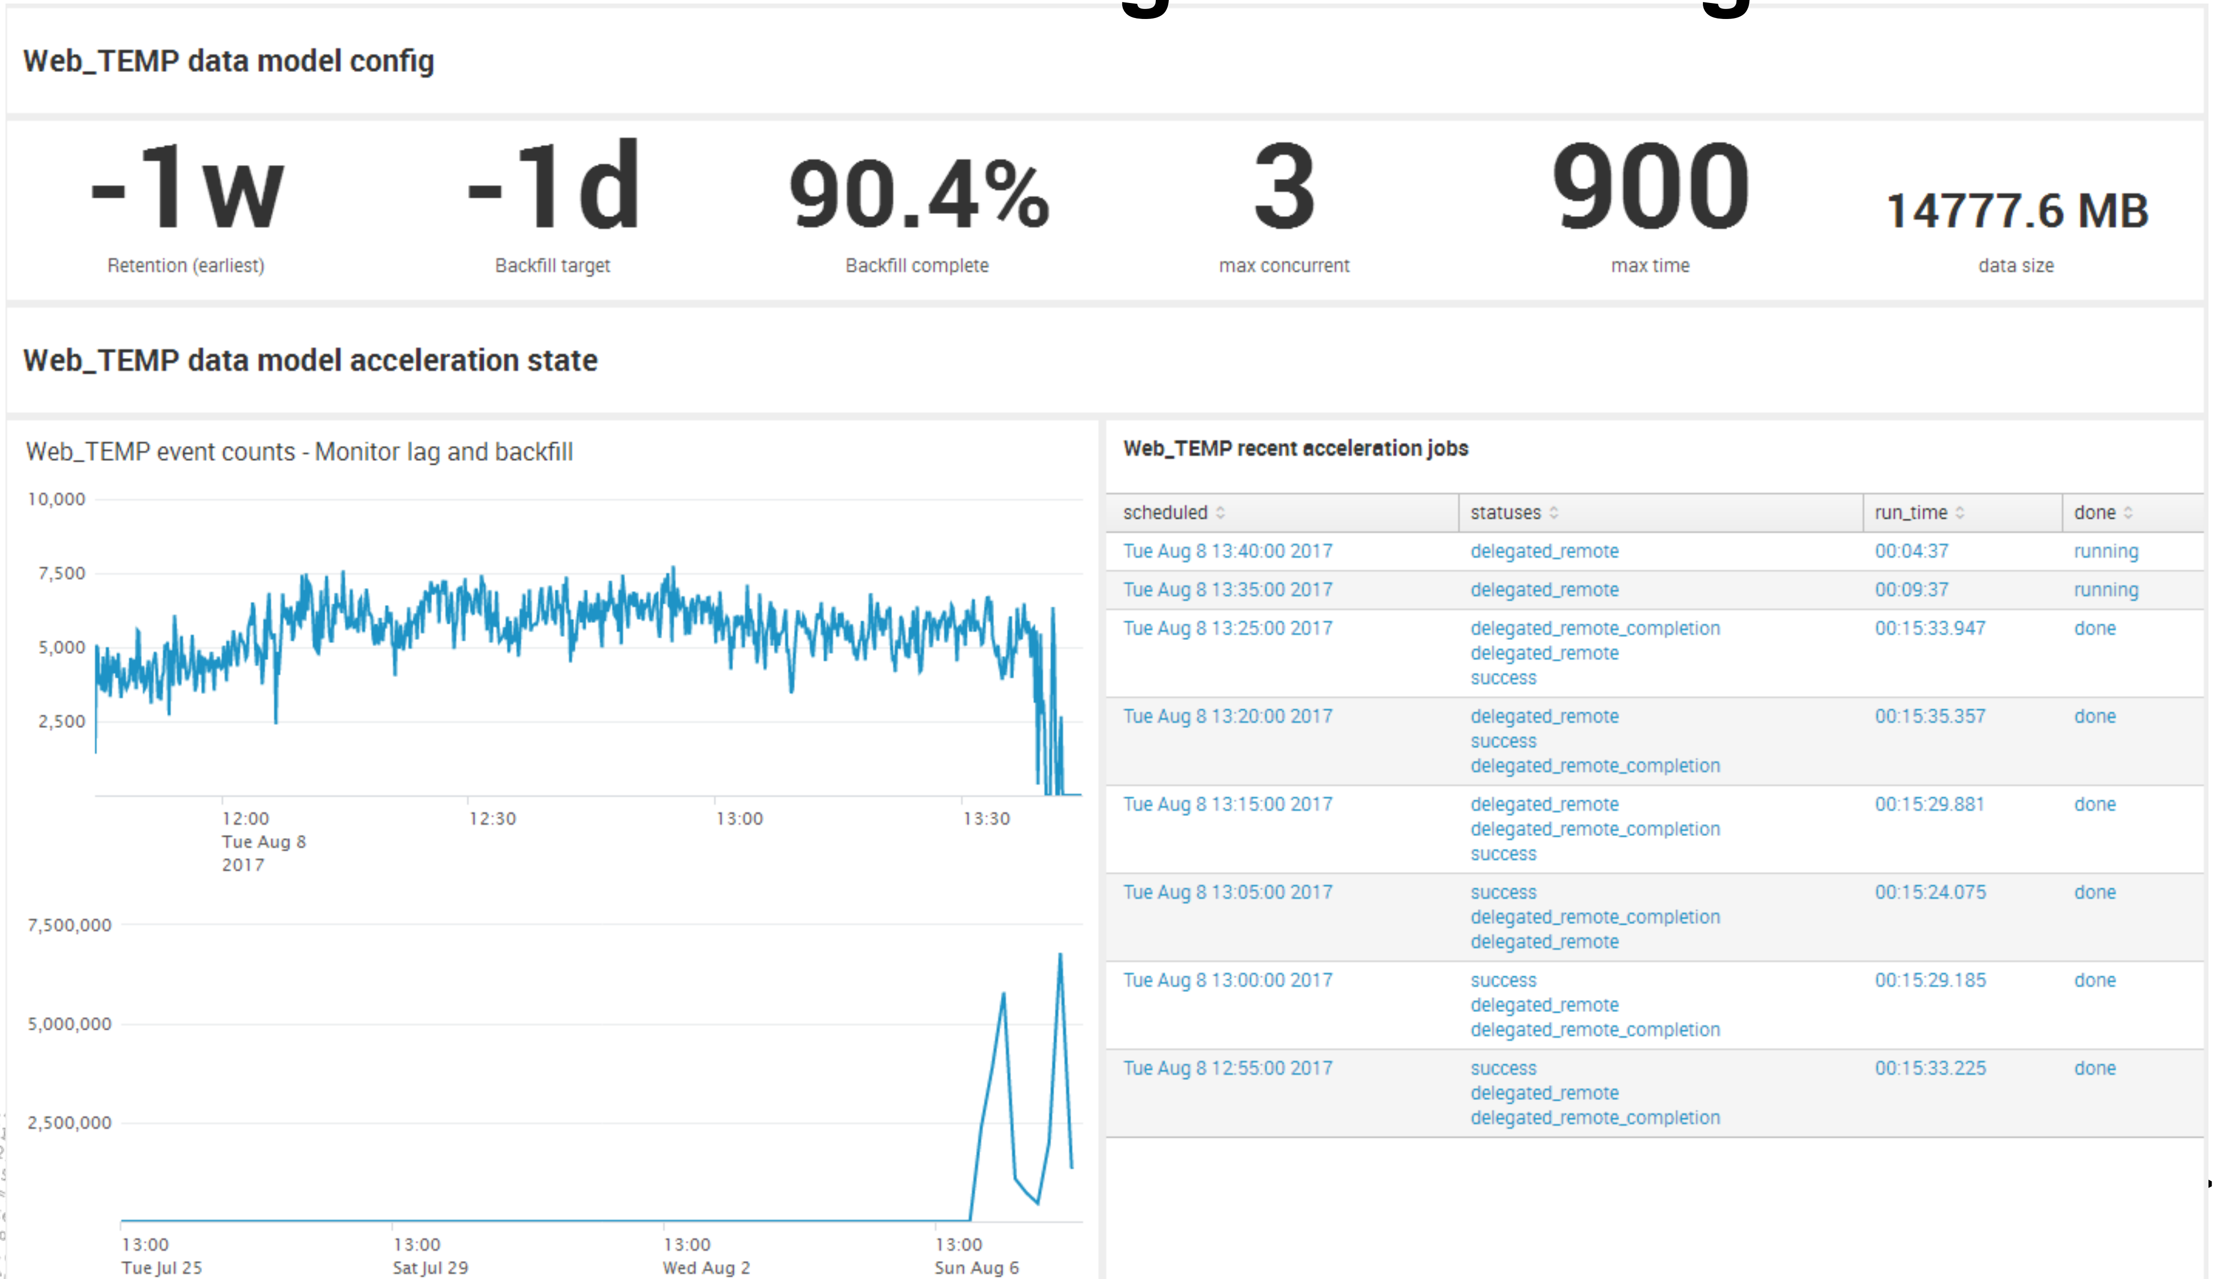

Web_TEMP data model config

-1w
Retention (earliest)

-1d
Backfill target

90.4%
Backfill complete

3
max concurrent

900
max time

14777.6 MB
data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 13:40:00 2017	delegated_remote	00:04:37	running
Tue Aug 8 13:35:00 2017	delegated_remote	00:09:37	running
Tue Aug 8 13:25:00 2017	delegated_remote_completion delegated_remote success	00:15:33.947	done
Tue Aug 8 13:20:00 2017	delegated_remote success delegated_remote_completion	00:15:35.357	done
Tue Aug 8 13:15:00 2017	delegated_remote delegated_remote_completion success	00:15:29.881	done
Tue Aug 8 13:05:00 2017	success delegated_remote_completion delegated_remote	00:15:24.075	done
Tue Aug 8 13:00:00 2017	success delegated_remote delegated_remote_completion	00:15:29.185	done
Tue Aug 8 12:55:00 2017	success delegated_remote delegated_remote_completion	00:15:33.225	done

Backfilling Without Lag

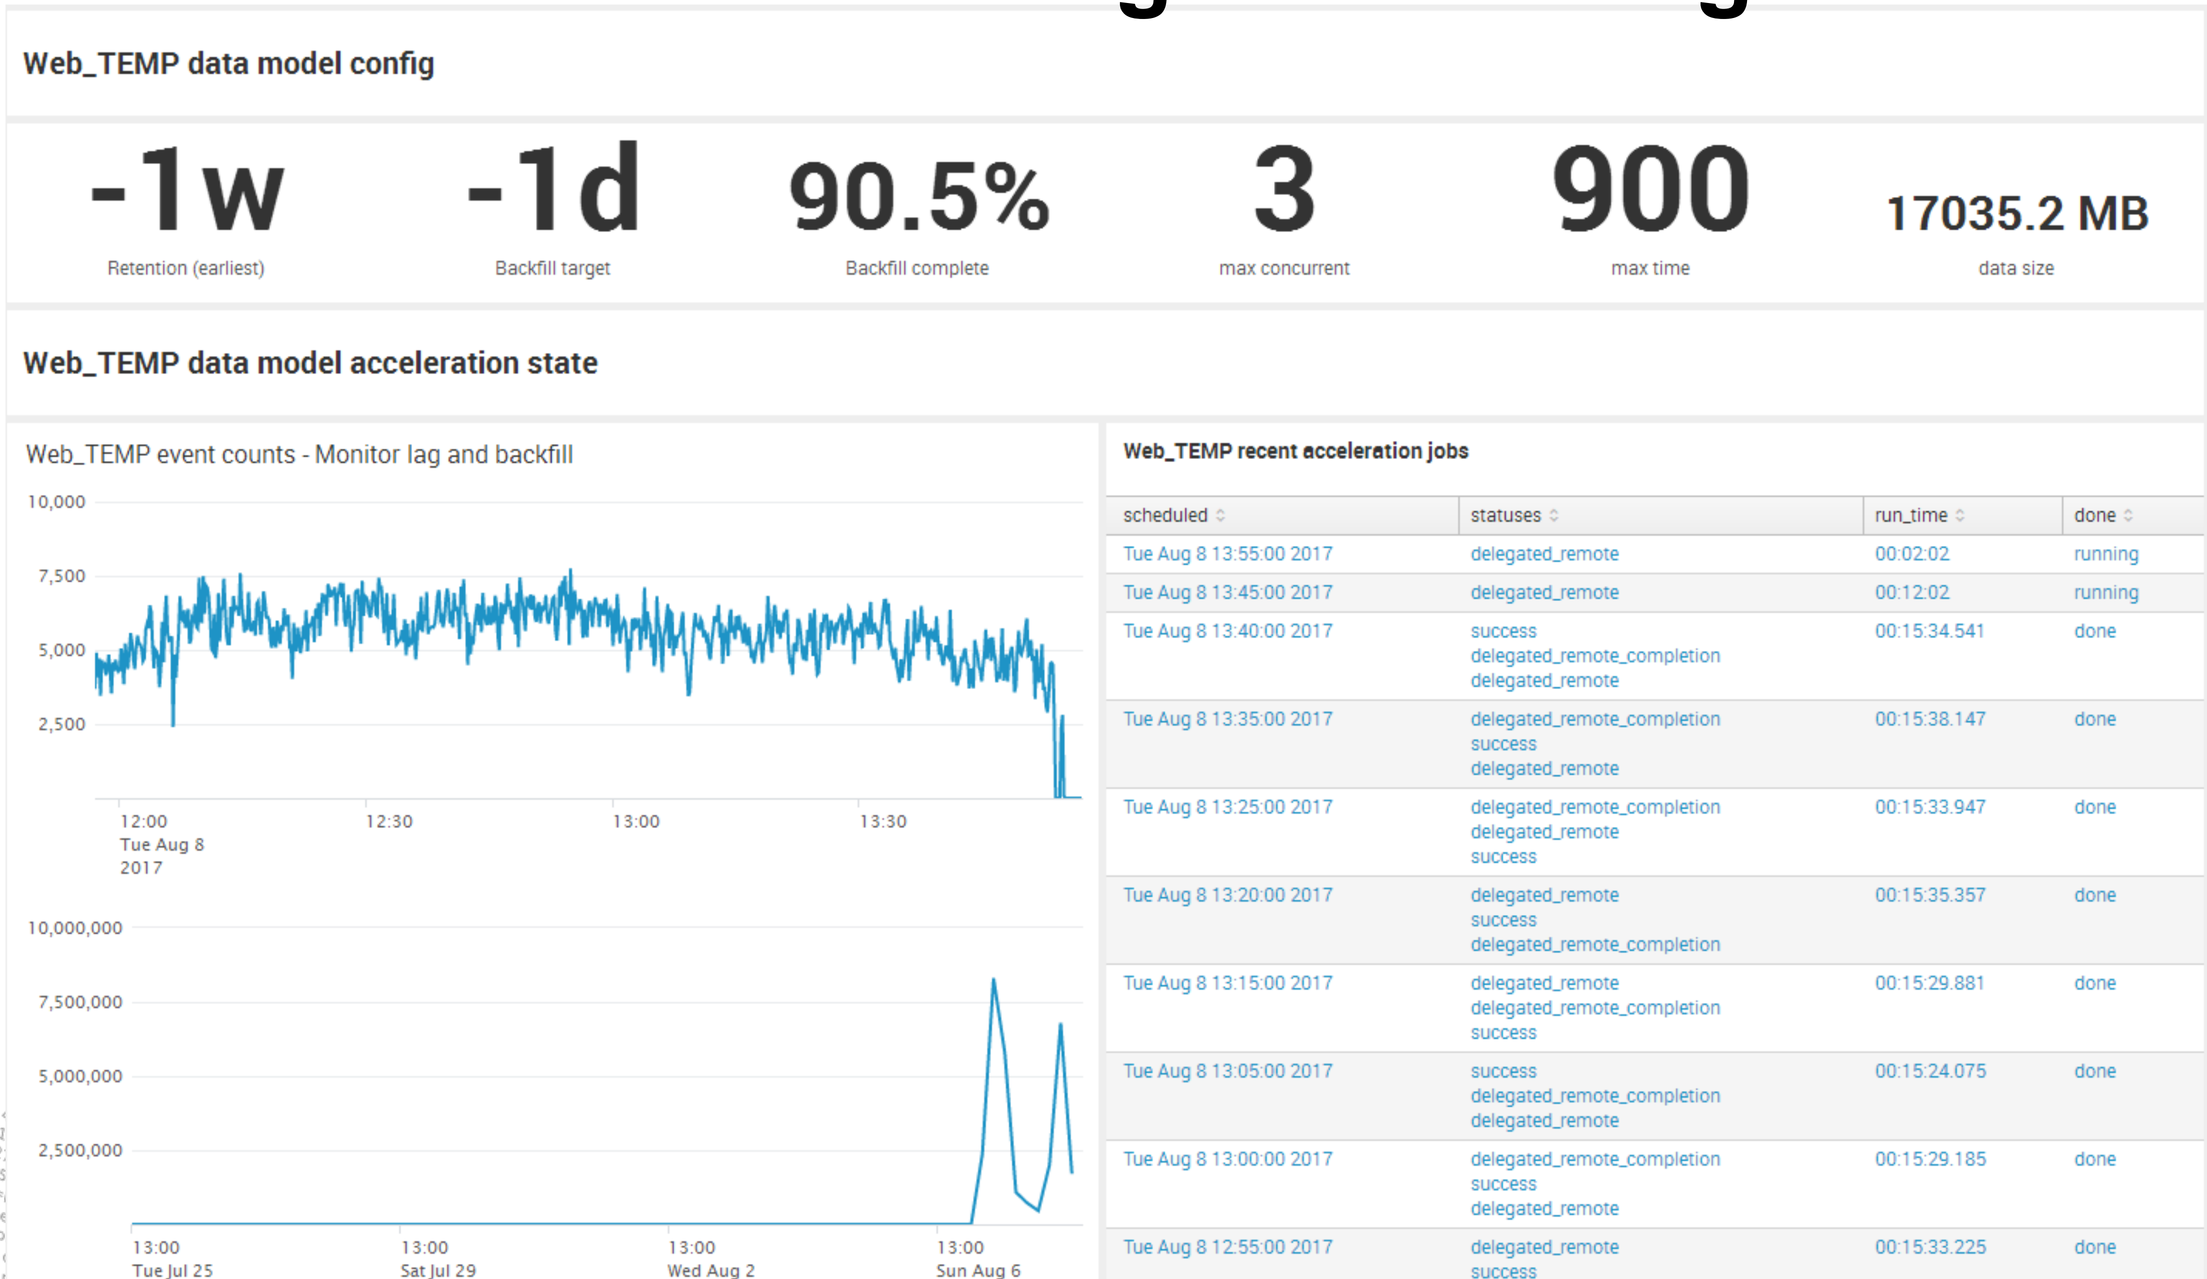

Web_TEMP data model config

-1w	-1d	90.5%	3	900	17035.2 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 13:55:00 2017	delegated_remote	00:02:02	running
Tue Aug 8 13:45:00 2017	delegated_remote	00:12:02	running
Tue Aug 8 13:40:00 2017	success delegated_remote_completion delegated_remote	00:15:34.541	done
Tue Aug 8 13:35:00 2017	delegated_remote_completion success delegated_remote	00:15:38.147	done
Tue Aug 8 13:25:00 2017	delegated_remote_completion delegated_remote success	00:15:33.947	done
Tue Aug 8 13:20:00 2017	delegated_remote success delegated_remote_completion	00:15:35.357	done
Tue Aug 8 13:15:00 2017	delegated_remote delegated_remote_completion success	00:15:29.881	done
Tue Aug 8 13:05:00 2017	success delegated_remote_completion delegated_remote	00:15:24.075	done
Tue Aug 8 13:00:00 2017	delegated_remote_completion success delegated_remote	00:15:29.185	done
Tue Aug 8 12:55:00 2017	delegated_remote success	00:15:33.225	done

Backfilling Without Lag

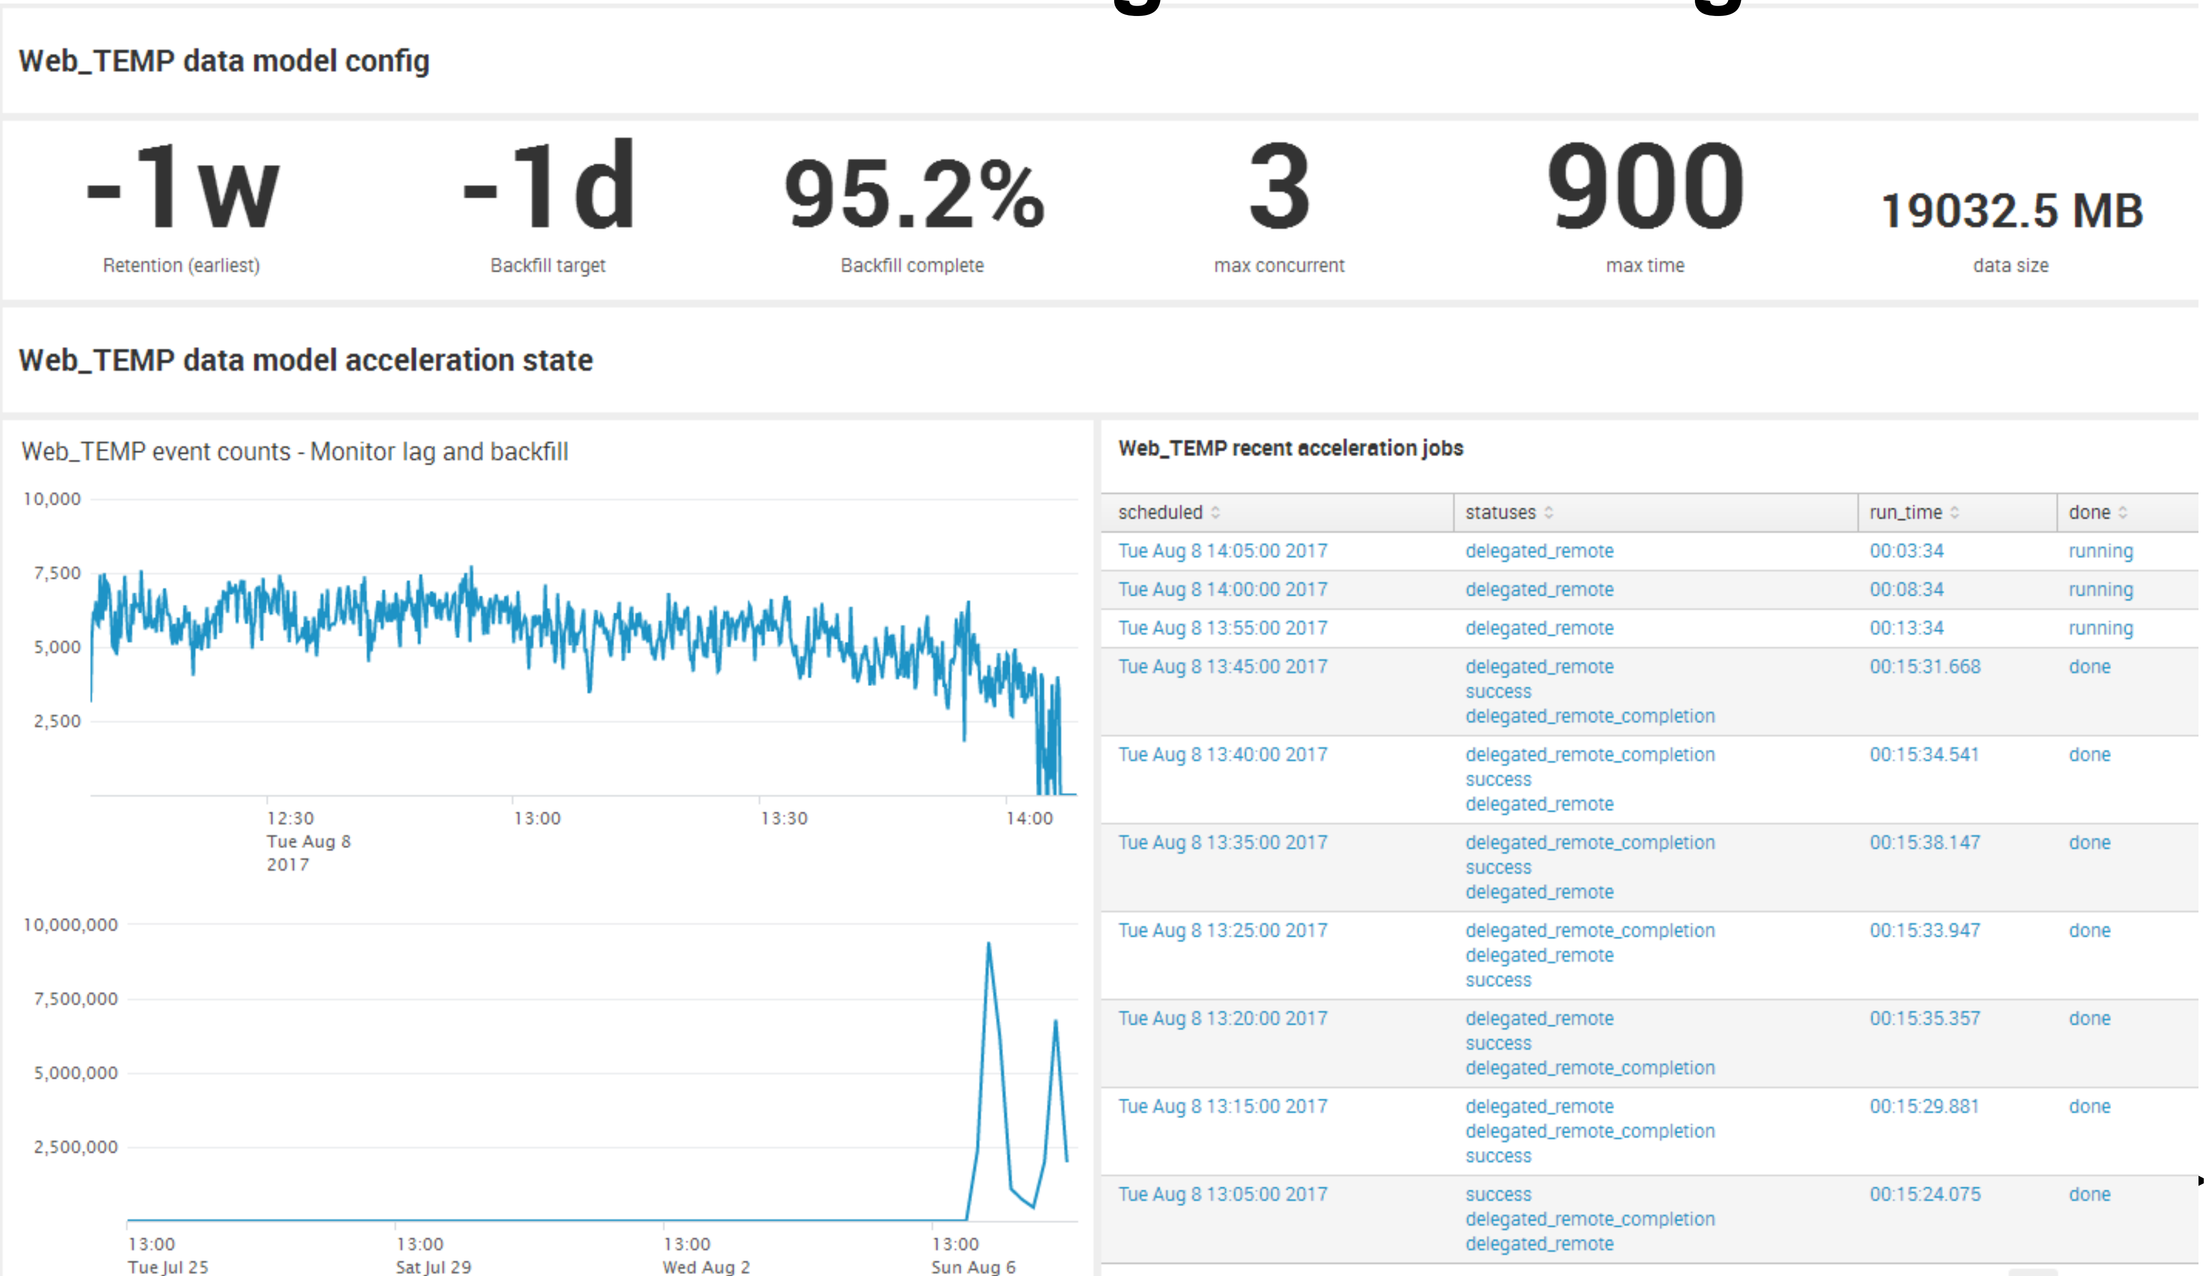

Web_TEMP data model config

-1w	-1d	95.2%	3	900	19032.5 MB
Retention (earliest)	Backfill target	Backfill complete	max concurrent	max time	data size

Web_TEMP data model acceleration state

Web_TEMP event counts - Monitor lag and backfill

Web_TEMP recent acceleration jobs

scheduled	statuses	run_time	done
Tue Aug 8 14:05:00 2017	delegated_remote	00:03:34	running
Tue Aug 8 14:00:00 2017	delegated_remote	00:08:34	running
Tue Aug 8 13:55:00 2017	delegated_remote	00:13:34	running
Tue Aug 8 13:45:00 2017	delegated_remote success delegated_remote_completion	00:15:31.668	done
Tue Aug 8 13:40:00 2017	delegated_remote_completion success delegated_remote	00:15:34.541	done
Tue Aug 8 13:35:00 2017	delegated_remote_completion success delegated_remote	00:15:38.147	done
Tue Aug 8 13:25:00 2017	delegated_remote_completion delegated_remote success	00:15:33.947	done
Tue Aug 8 13:20:00 2017	delegated_remote success delegated_remote_completion	00:15:35.357	done
Tue Aug 8 13:15:00 2017	delegated_remote delegated_remote_completion success	00:15:29.881	done
Tue Aug 8 13:05:00 2017	success delegated_remote_completion delegated_remote	00:15:24.075	done

The Big Picture

Accuracy, Load & Data Models

What Do People Really Want?

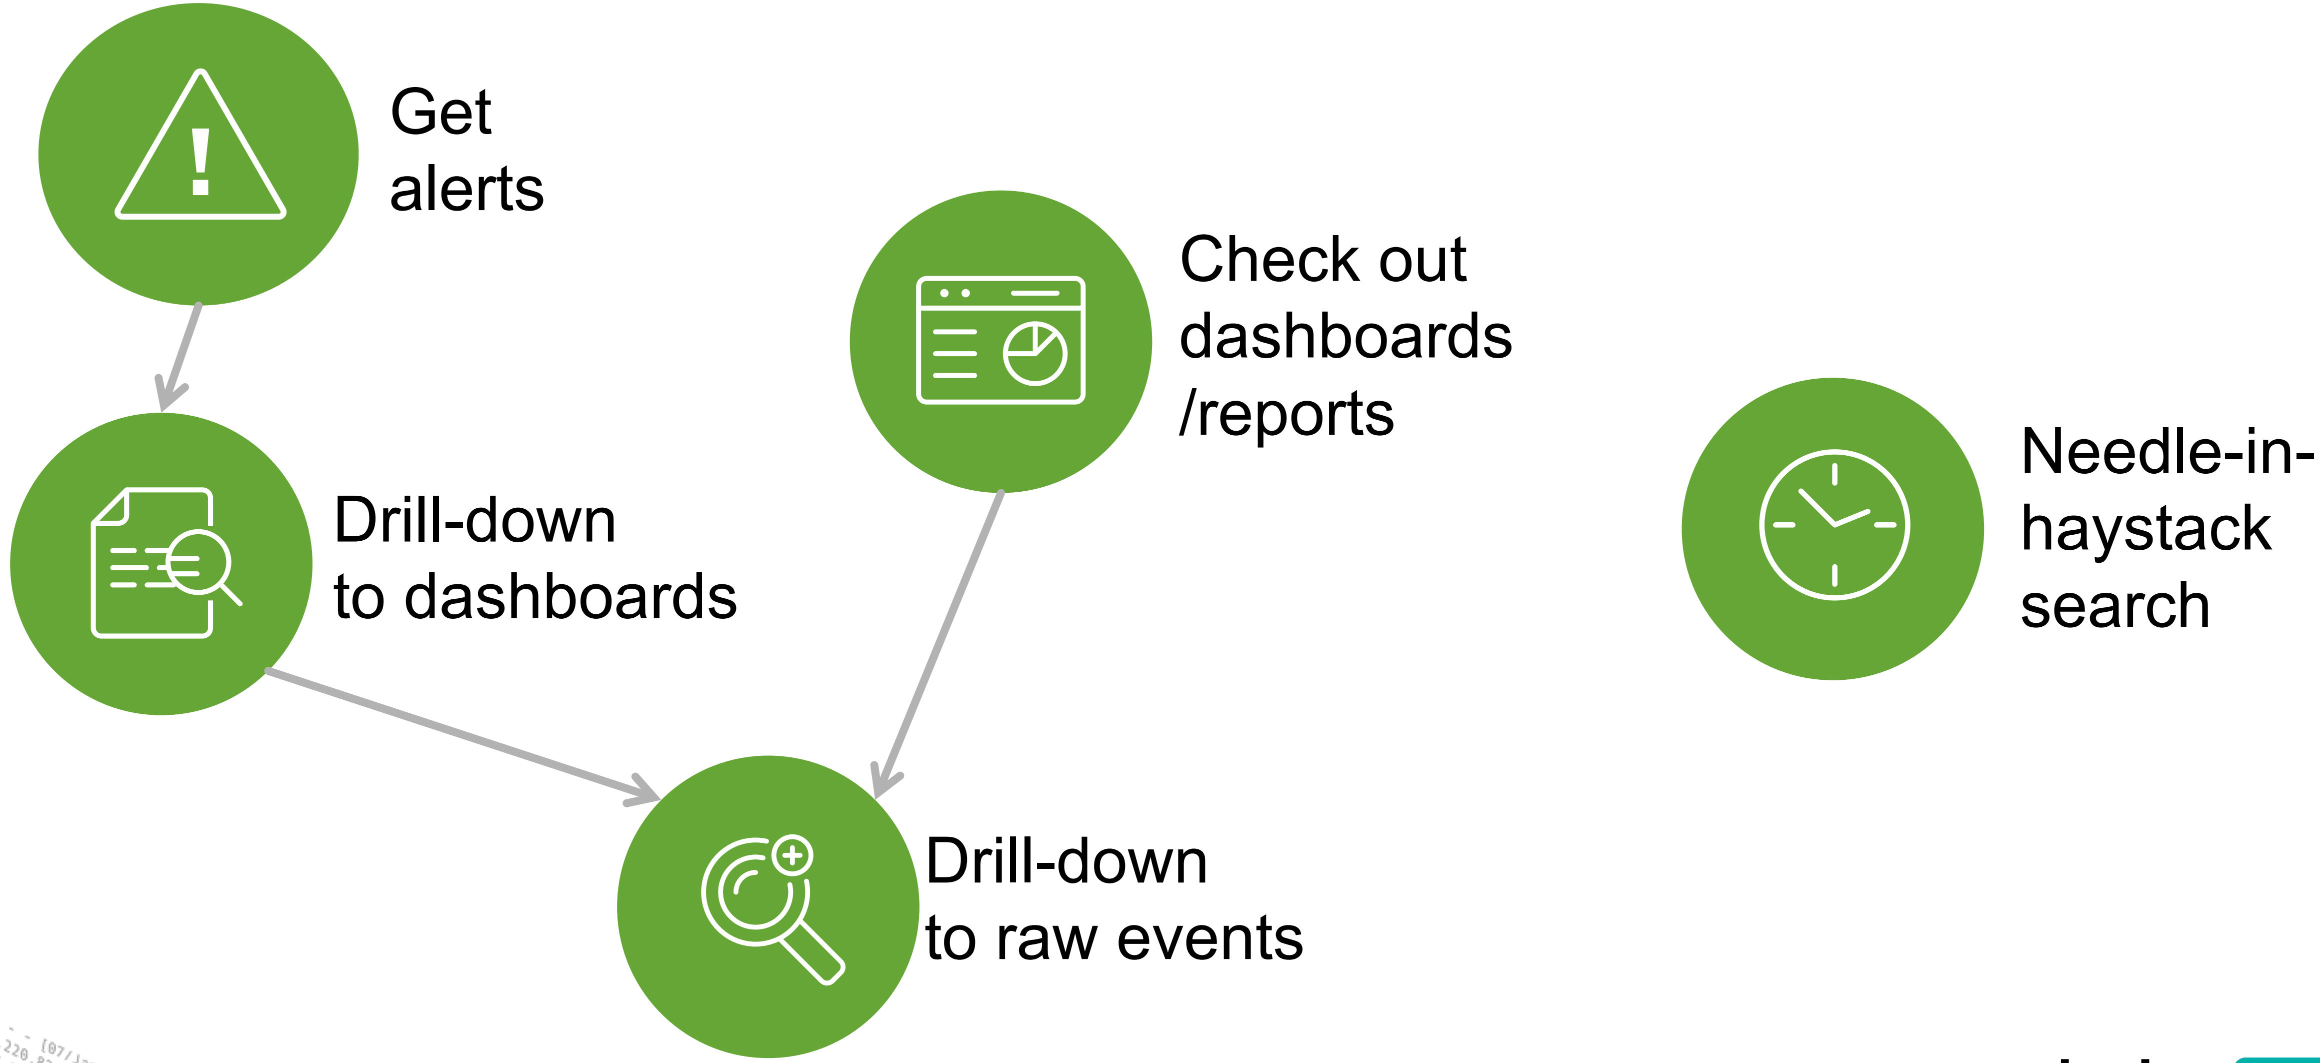

Get alerts

Drill-down to dashboards

Check out dashboards /reports

Drill-down to raw events

Needle-in-haystack search

Accuracy, Load & Data Models

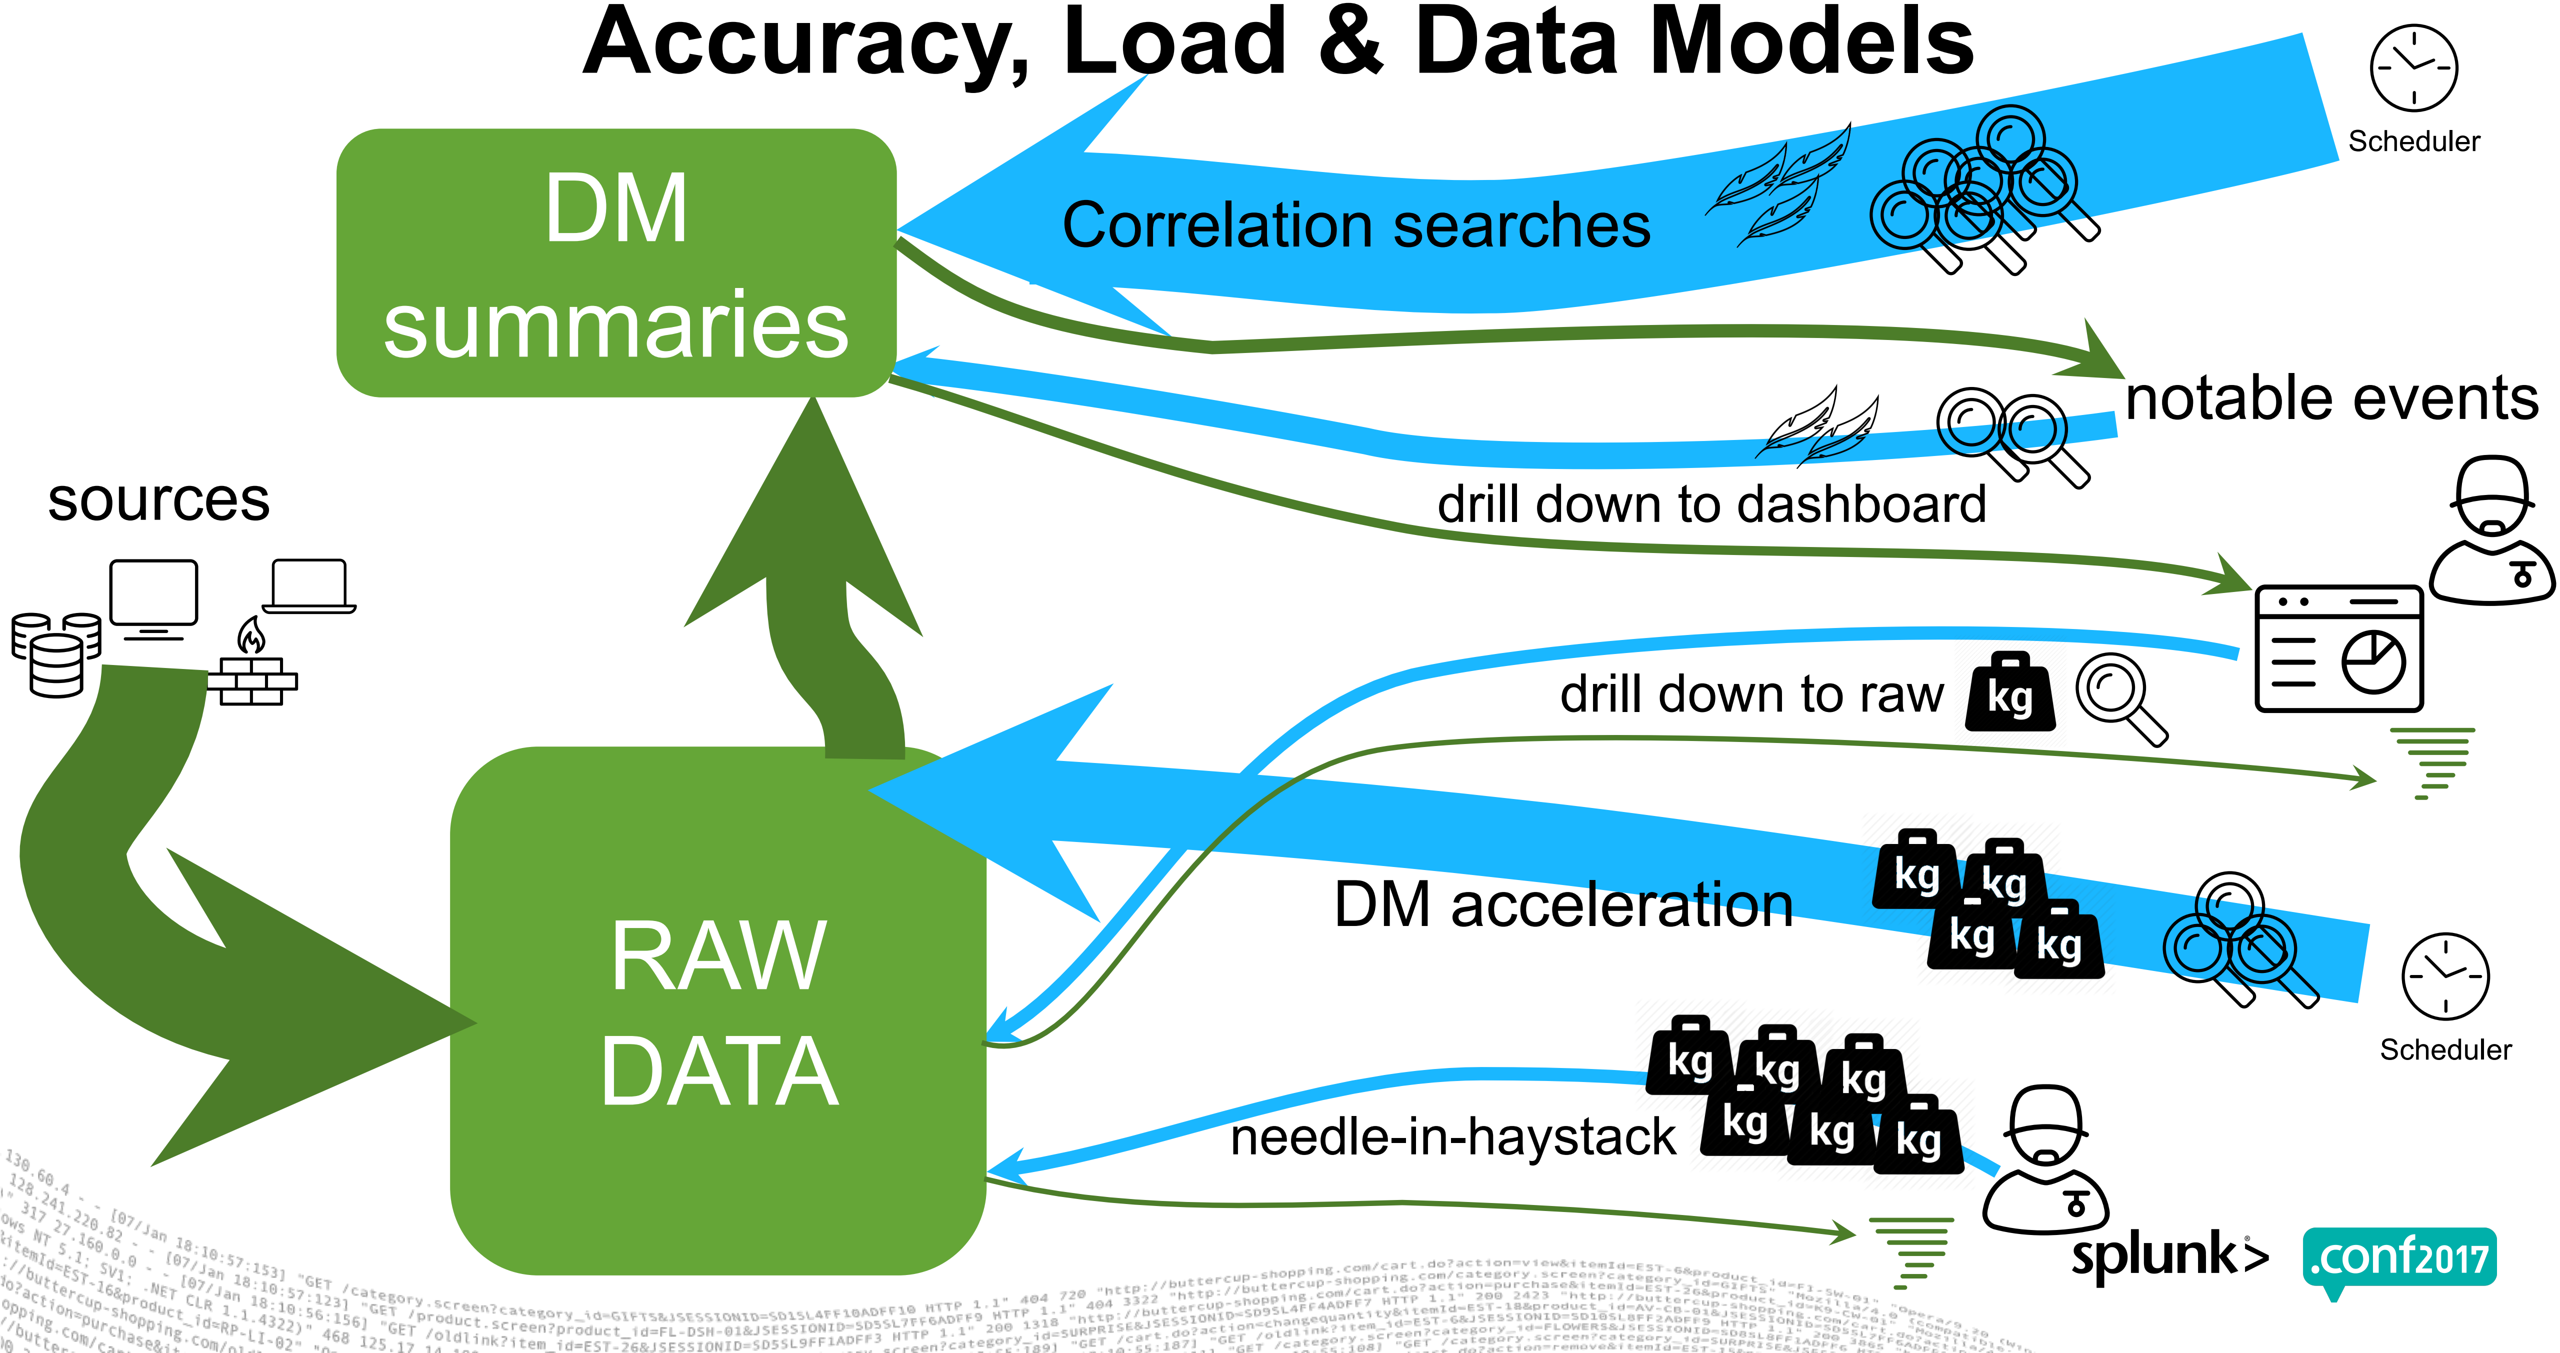

DM summaries

Correlation searches

Scheduler

notable events

drill down to dashboard

sources

drill down to raw

RAW DATA

DM acceleration

Scheduler

needle-in-haystack

Accuracy, Load & Data Models

Healthy sources

- Accurate time (NTP)
- Accurate TZ
- No lag
- No rubbish

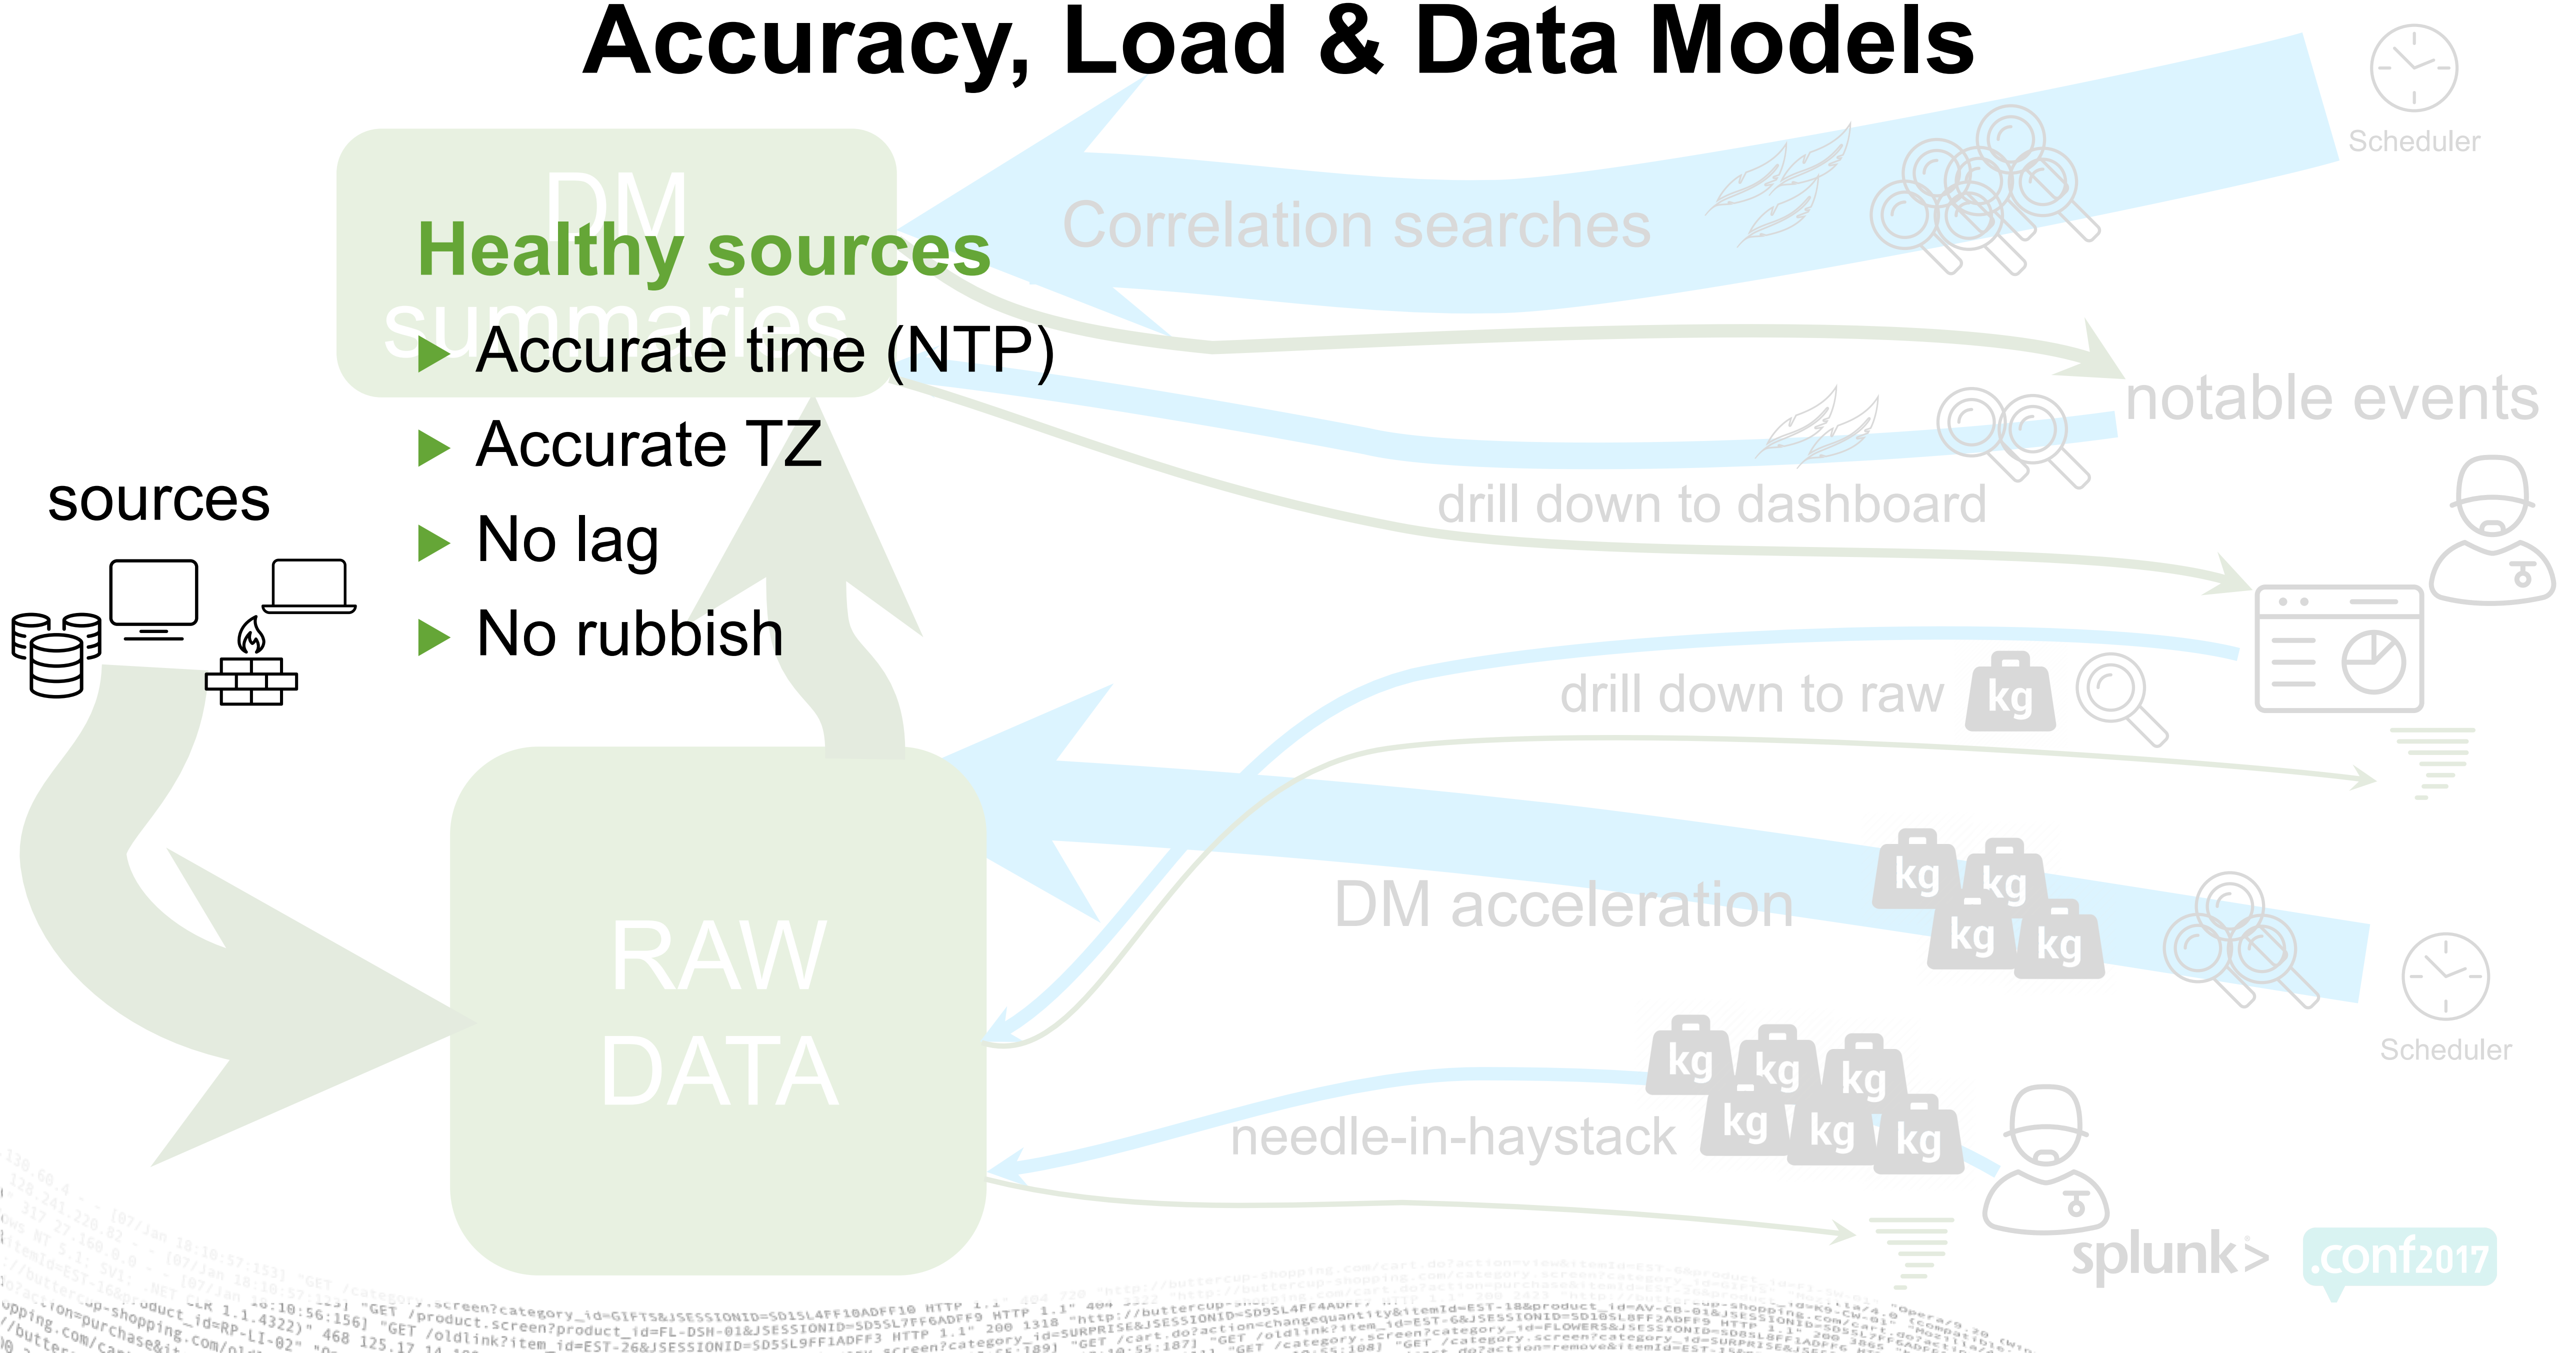

sources

DM summaries

RAW DATA

Correlation searches

Scheduler

notable events

drill down to dashboard

drill down to raw

DM acceleration

needle-in-haystack

Scheduler

Accuracy, Load & Data Models

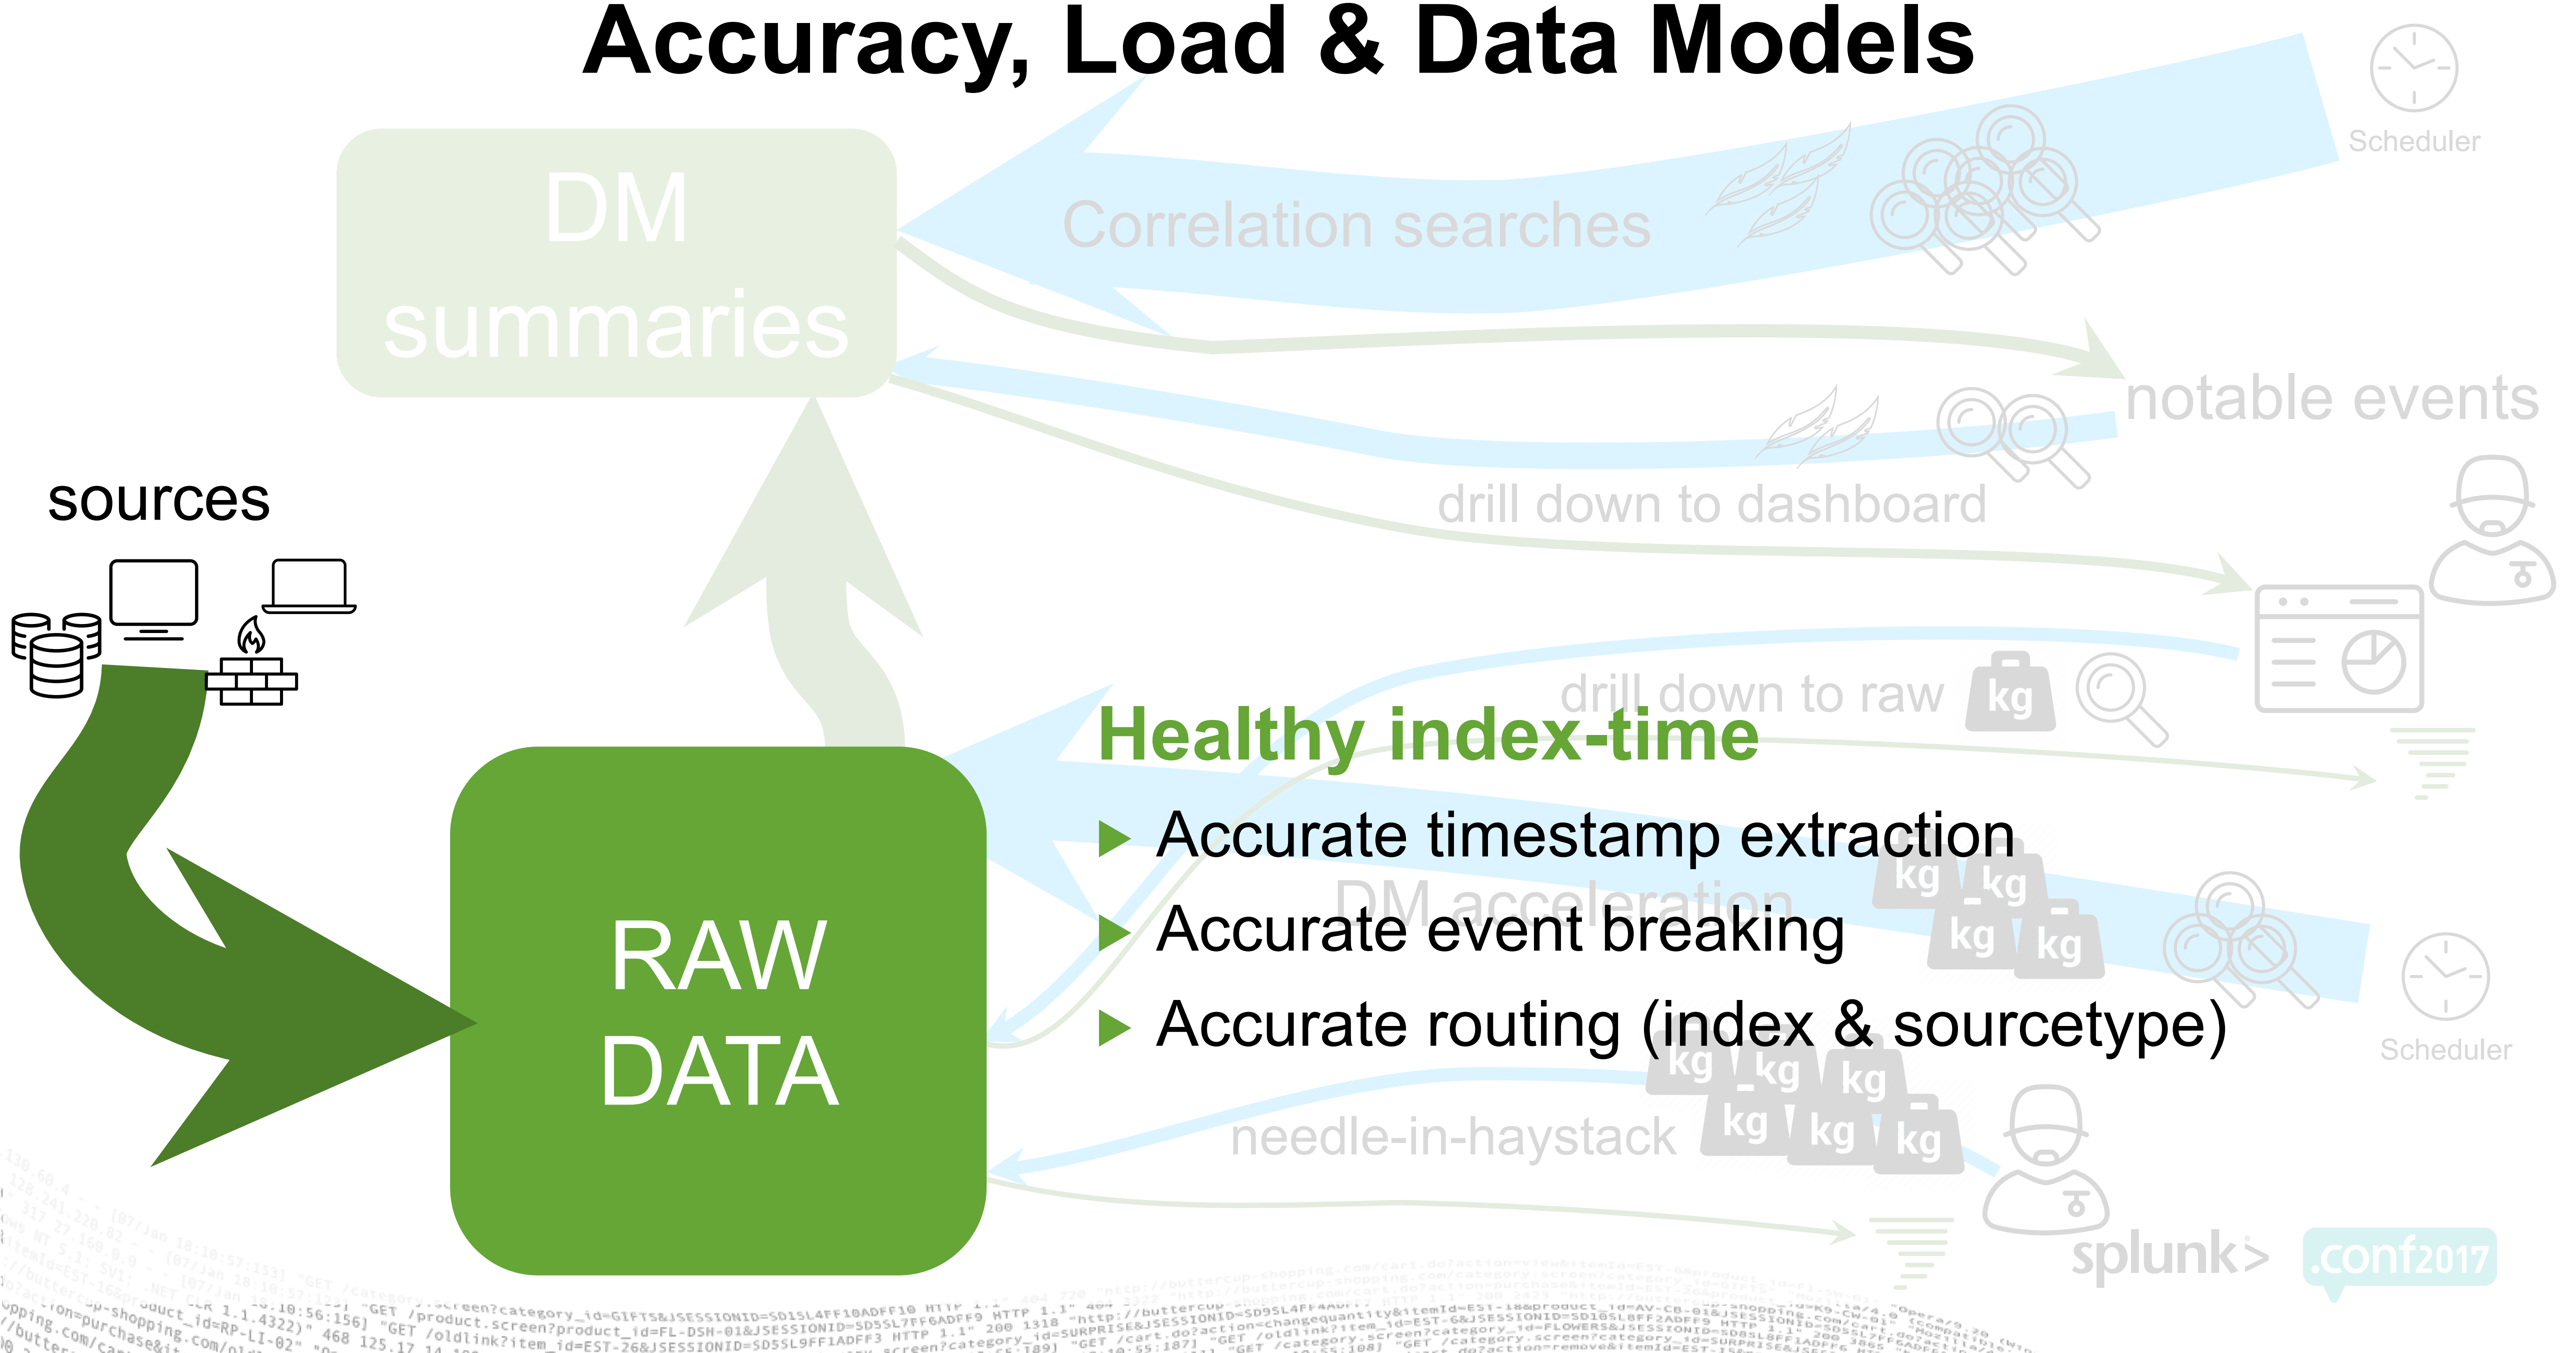

Healthy index-time

- Accurate timestamp extraction
- Accurate event breaking
- Accurate routing (index & sourcetype)

Accuracy, Load & Data Models

Healthy search-time

- CIM-compliant field extractions
- CIM-compliant eventtypes & tags

Healthy acceleration

- Timely acceleration searches (no lag)

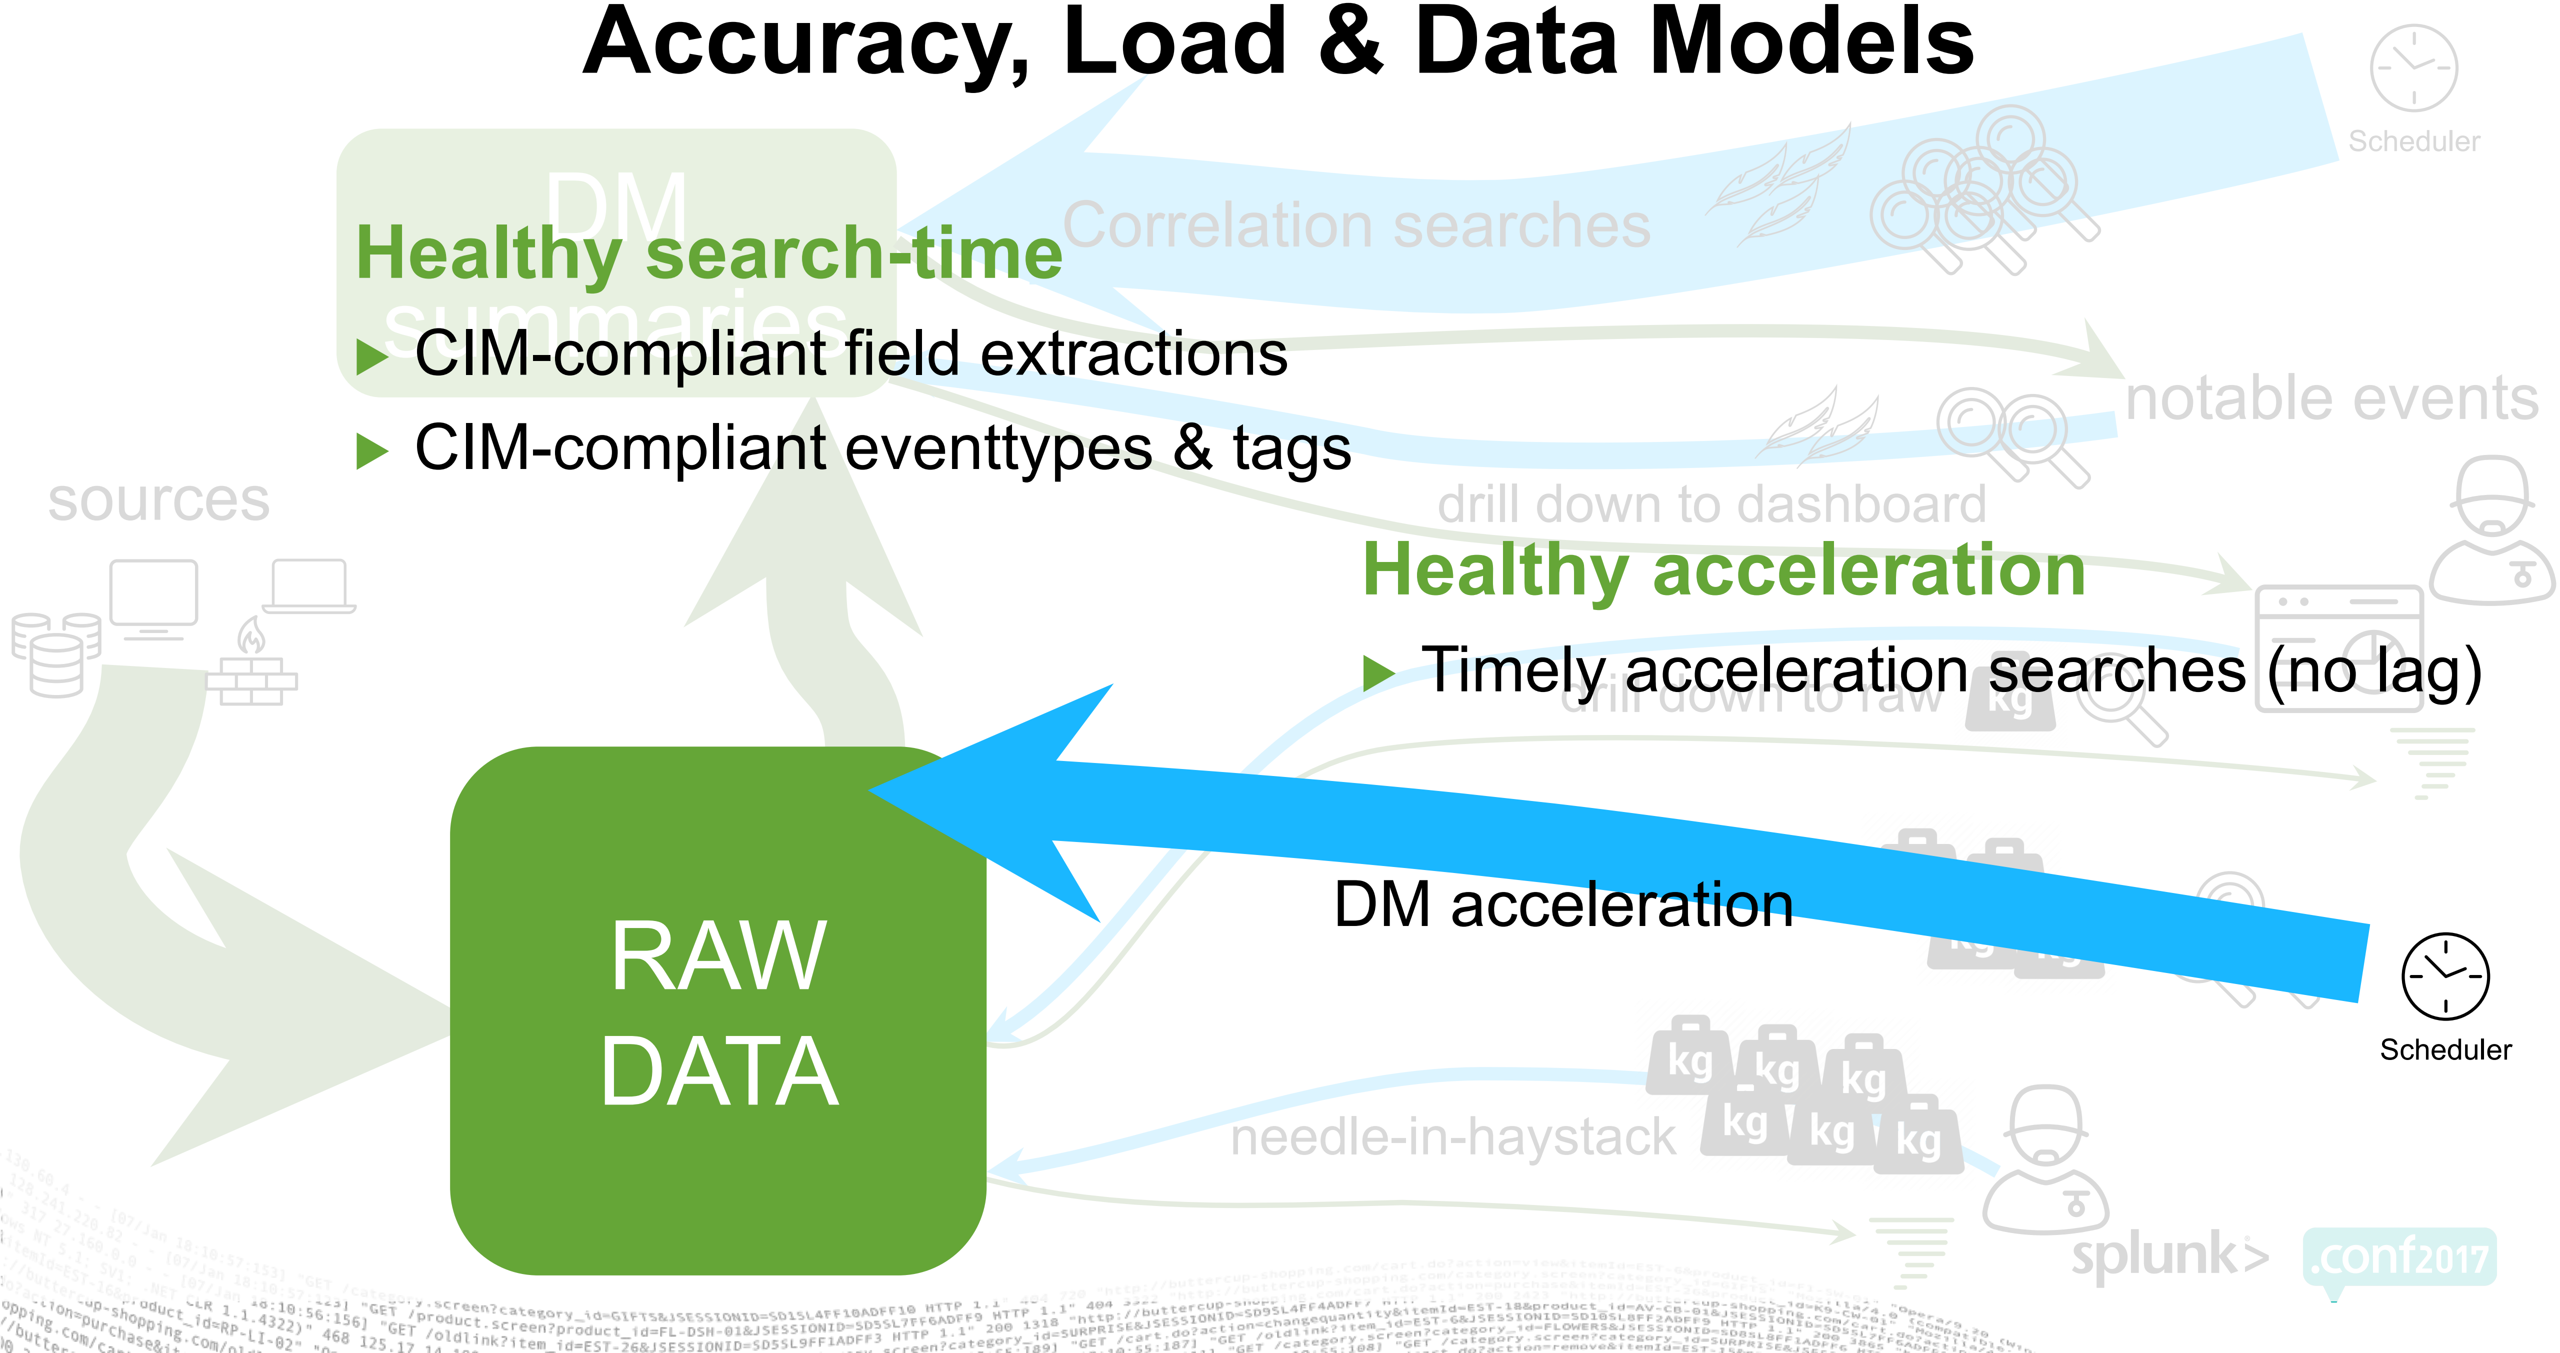

Accuracy, Load & Data Models

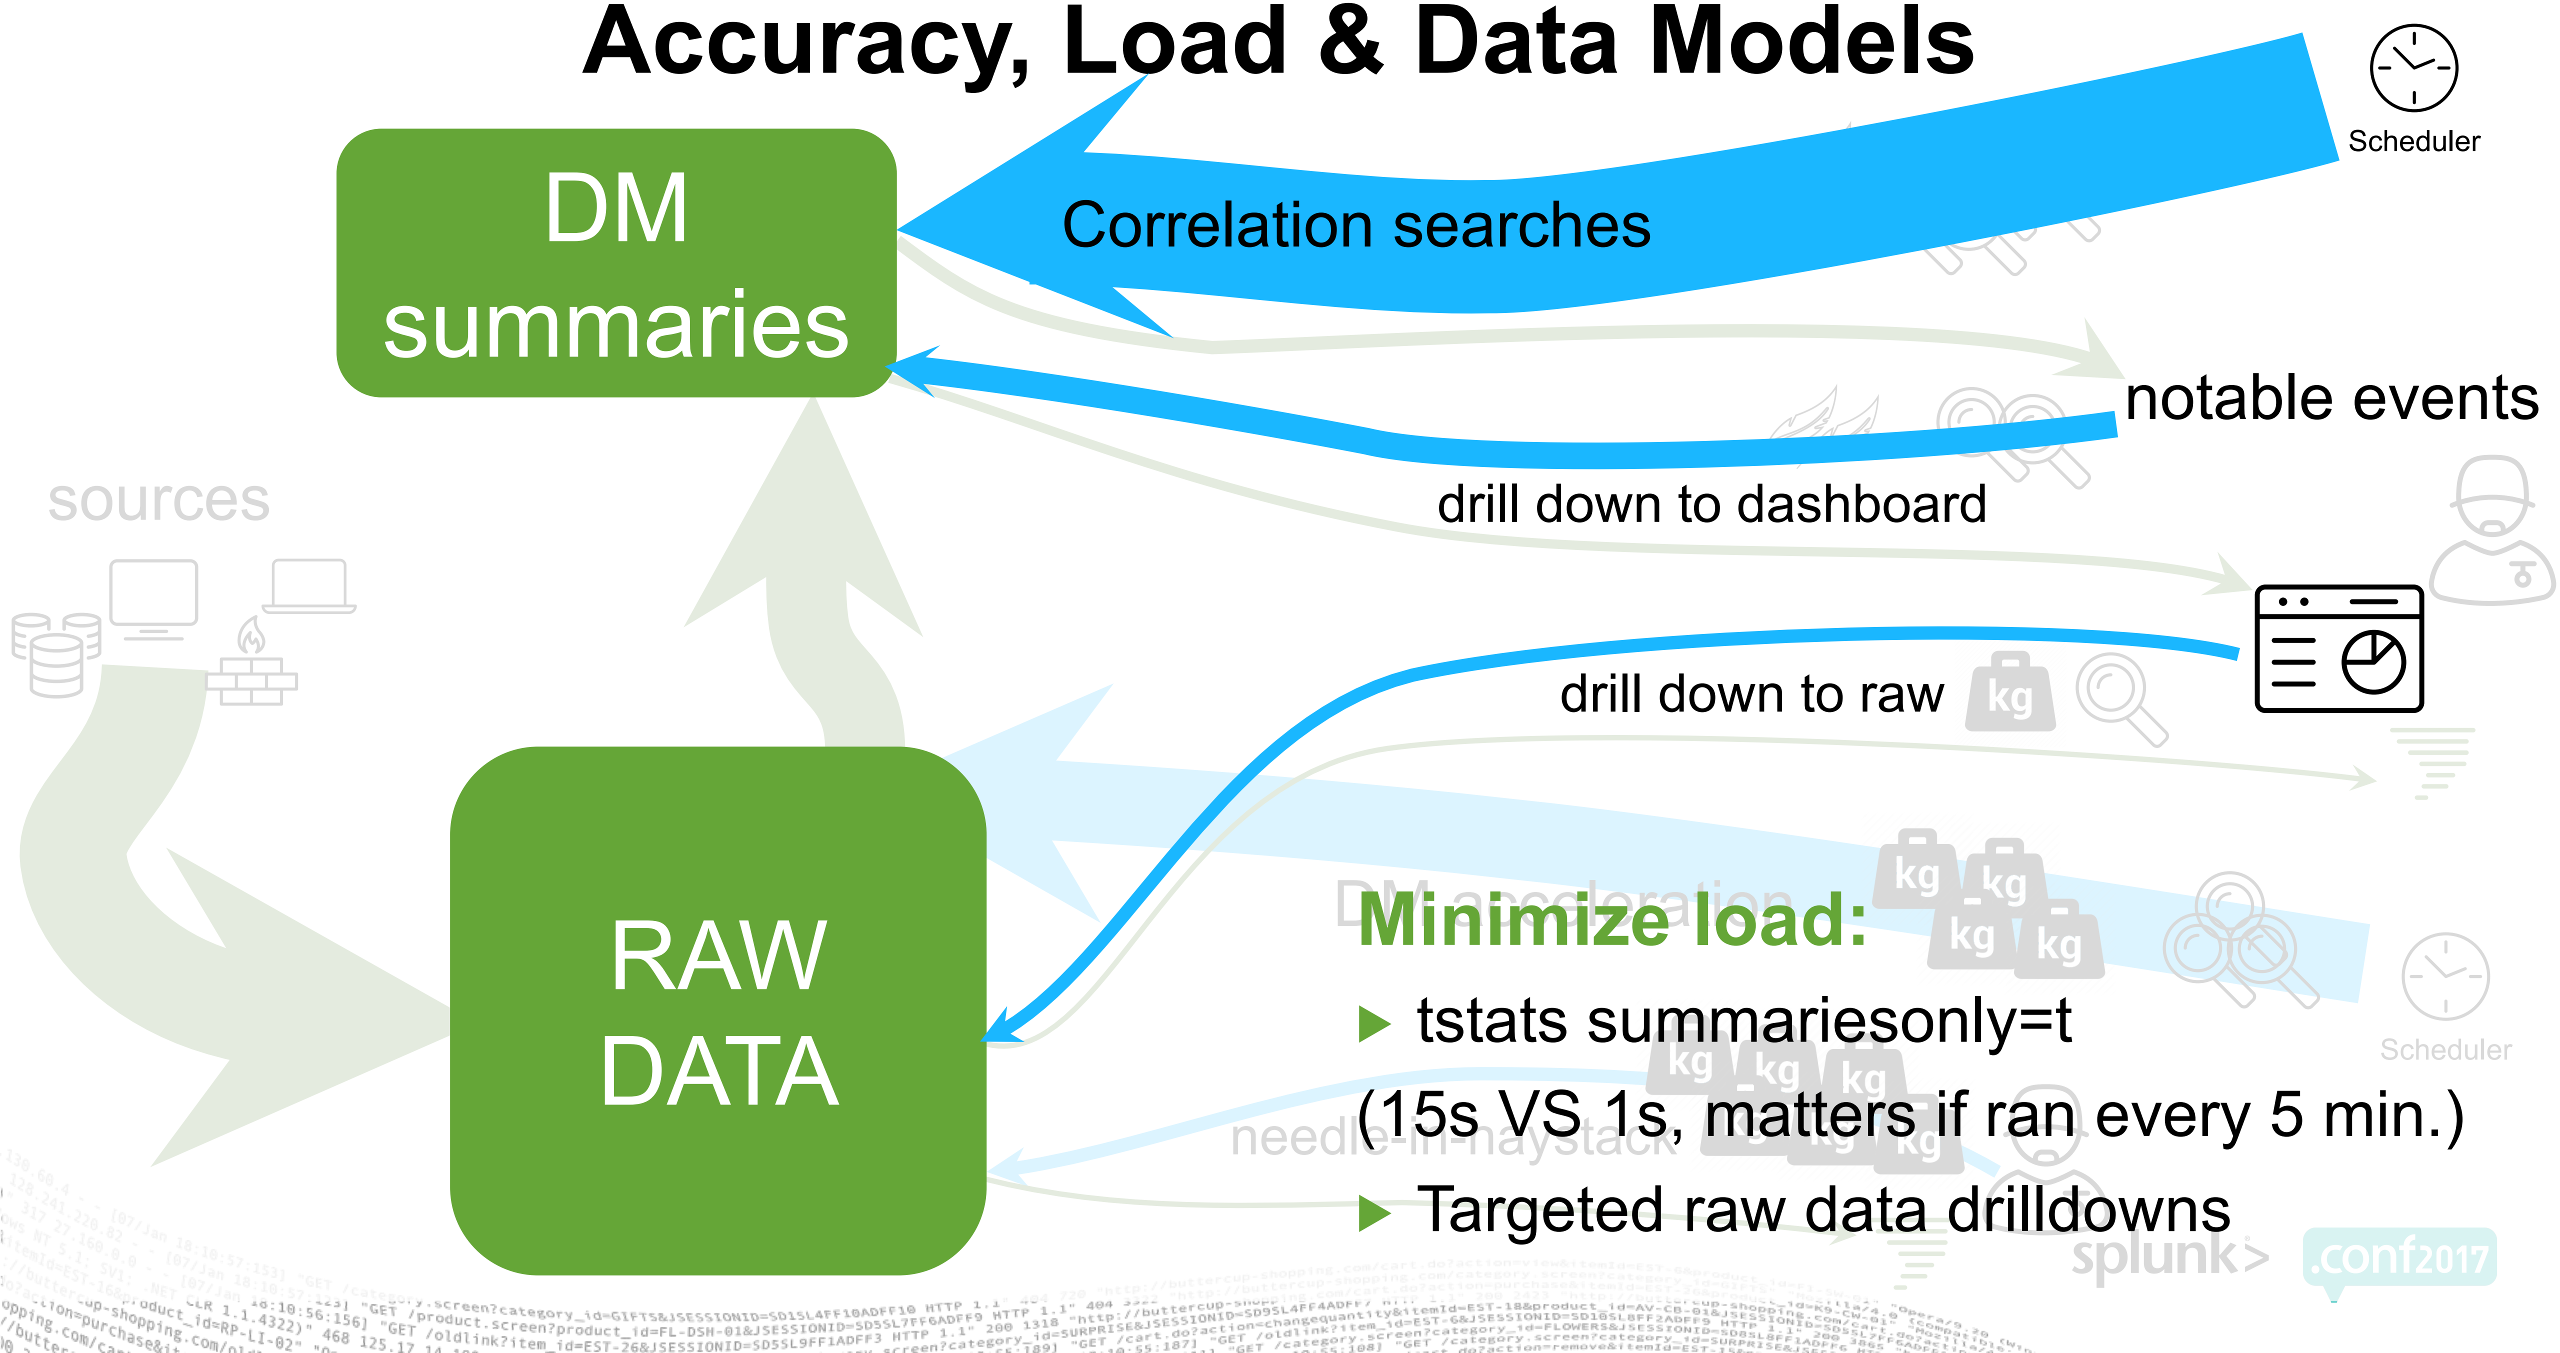

Minimize load:

- tstats summariesonly=t

(15s VS 1s, matters if ran every 5 min.)

- Targeted raw data drilldowns

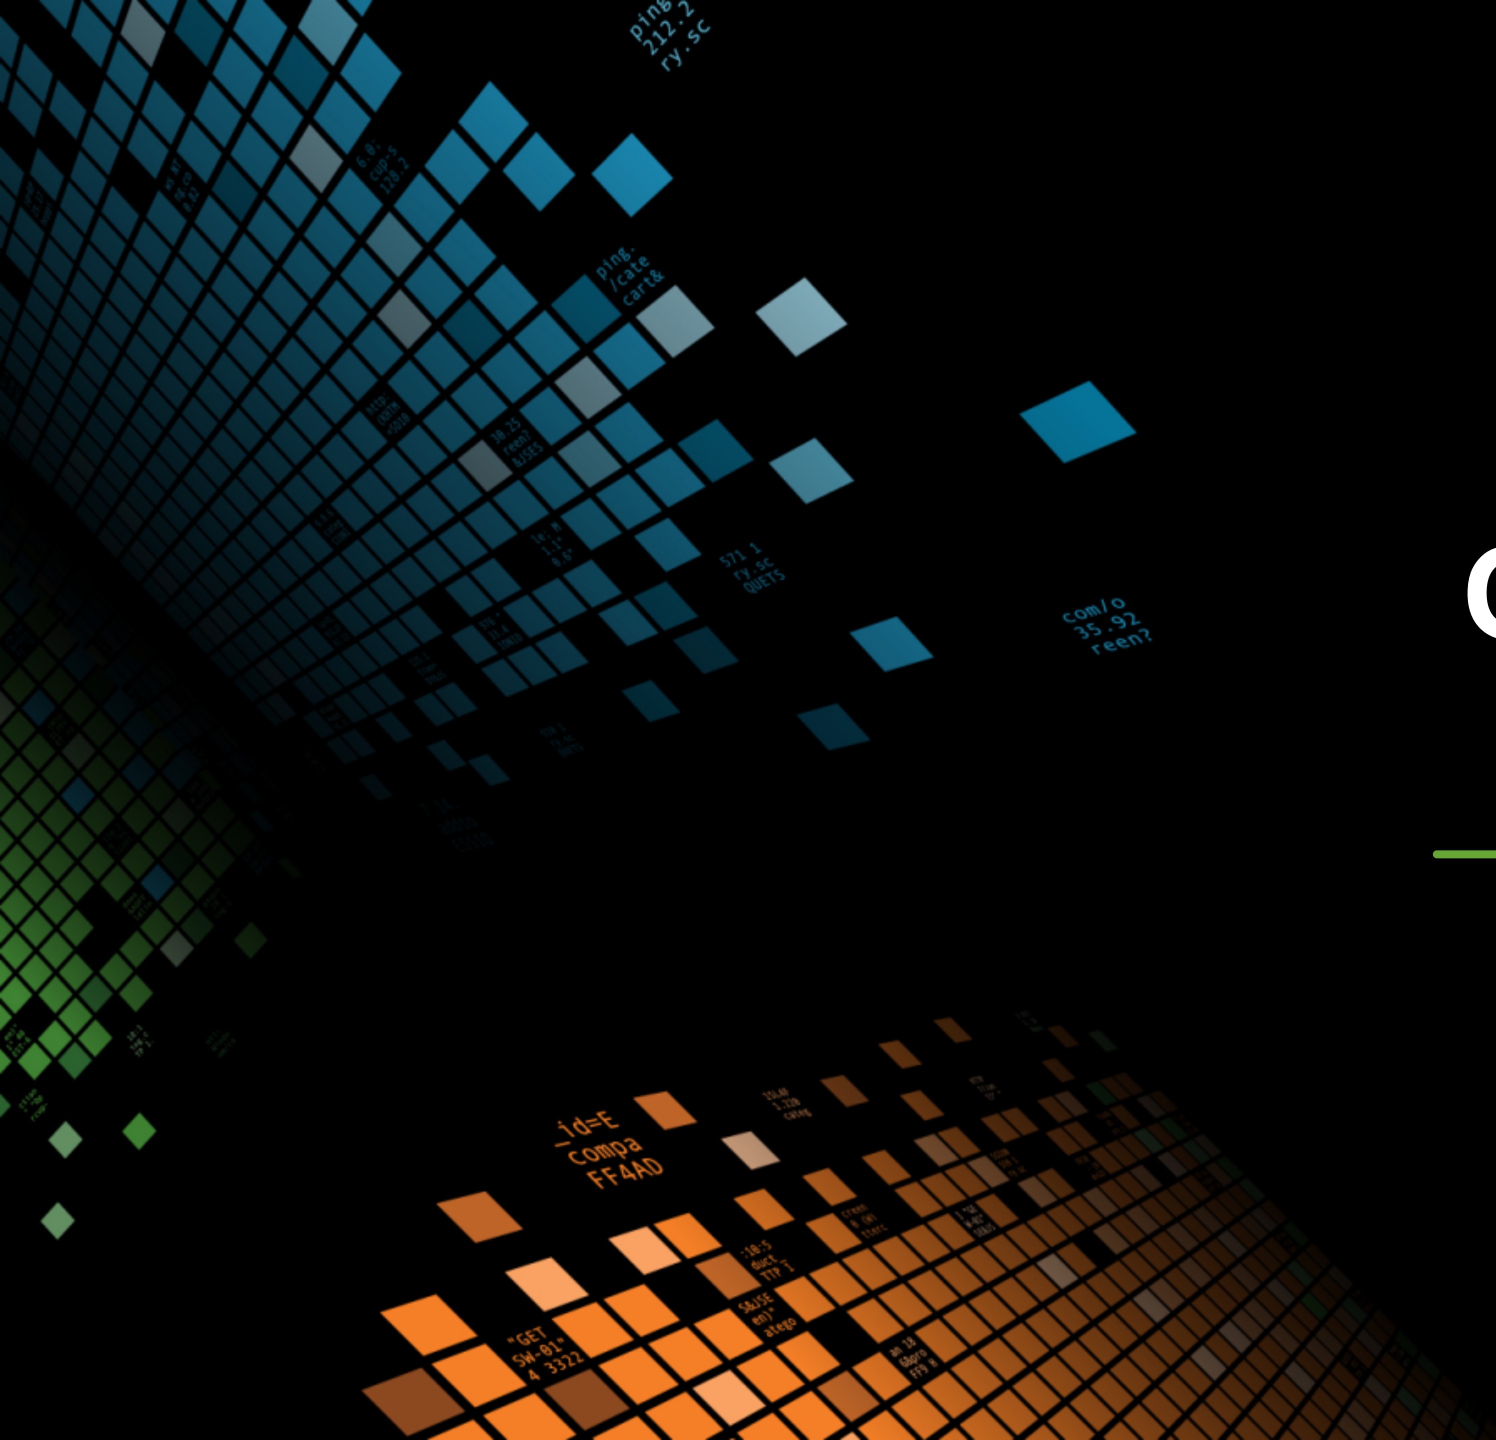

Optimizing DM Acceleration

Update a DM Without Rebuilding

- datamodels.conf

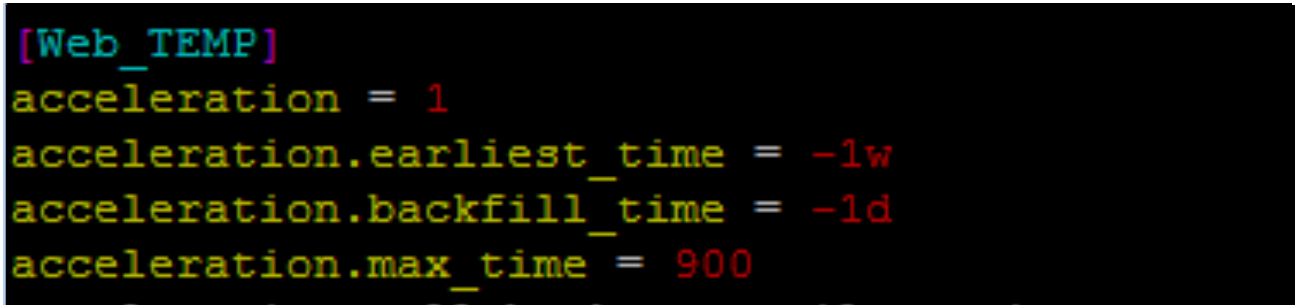

```
[Web_TEMP]
acceleration = 1
acceleration.earliest_time = -1w
acceleration.backfill_time = -1d
acceleration.max_time = 900
```

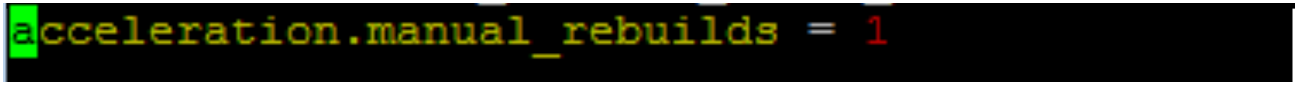

```
acceleration.manual_rebuilds = 1
```

- limits.conf

```
[tstats]
allow_old_summaries = true
```

- Splunk Web doesn't let you change an accelerated DM :-(
- Turning off the acceleration and back on again may trigger a rebuild anyway

Update a DM Without Rebuilding

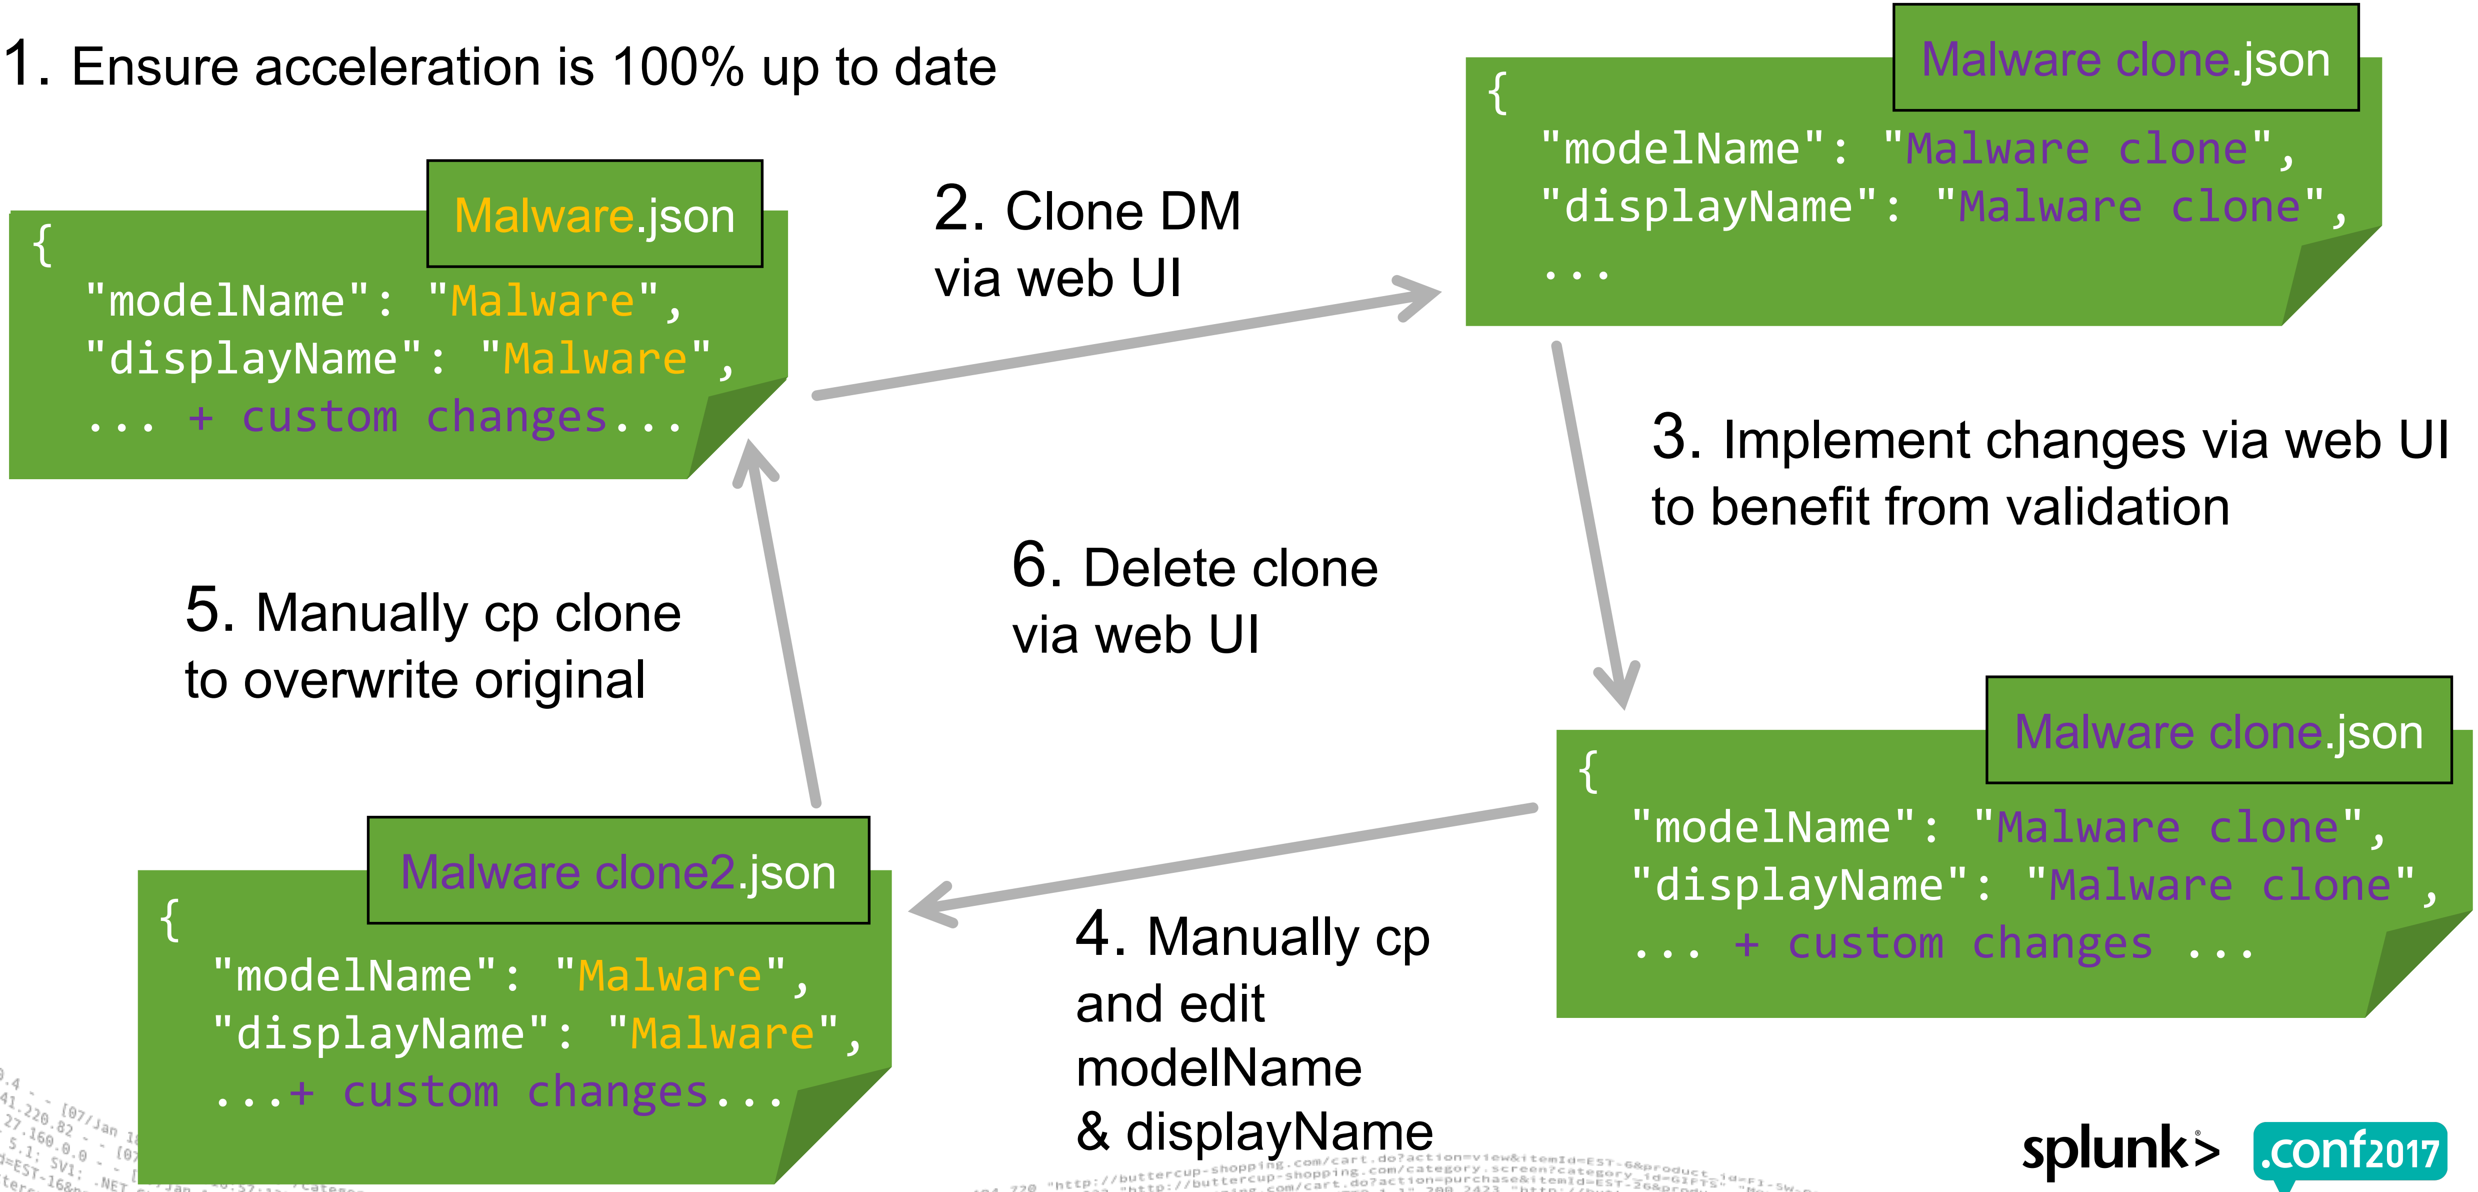

1. Ensure acceleration is 100% up to date

Malware.json

```
{
  "modelName": "Malware",
  "displayName": "Malware",
  ... + custom changes...
```

2. Clone DM via web UI

Malware clone.json

```
{
  "modelName": "Malware clone",
  "displayName": "Malware clone",
  ...
```

3. Implement changes via web UI to benefit from validation

Malware clone.json

```
{
  "modelName": "Malware clone",
  "displayName": "Malware clone",
  ... + custom changes ...
```

4. Manually cp and edit modelName & displayName

Malware clone2.json

```
{
  "modelName": "Malware",
  "displayName": "Malware",
  ...+ custom changes...
```

5. Manually cp clone to overwrite original

6. Delete clone via web UI

Optimize Constraint – Optimize Tags/Eventtypes

```
tag=malware tag=attack
```

- Look at eventtypes and tags to narrow them down
  - Remove eventtypes that you know will never happen in your data
  - Tweak eventtypes to make them faster
  - Painstaking! Only if you know what you're doing!

Optimize Constraint – Specify The Index

- Remember when I said the constraint for the Malware DM is this?

```
tag=malware tag=attack
```

- I lied! It's actually:

```
`cim_Malware_indexes` tag=malware tag=attack
```

- The CIM setup page sets that macro
- For other non splunk_SA_CIM DMs, you must manually edit the constraint in the same spirit

Optimize Constraint – Specify The Index

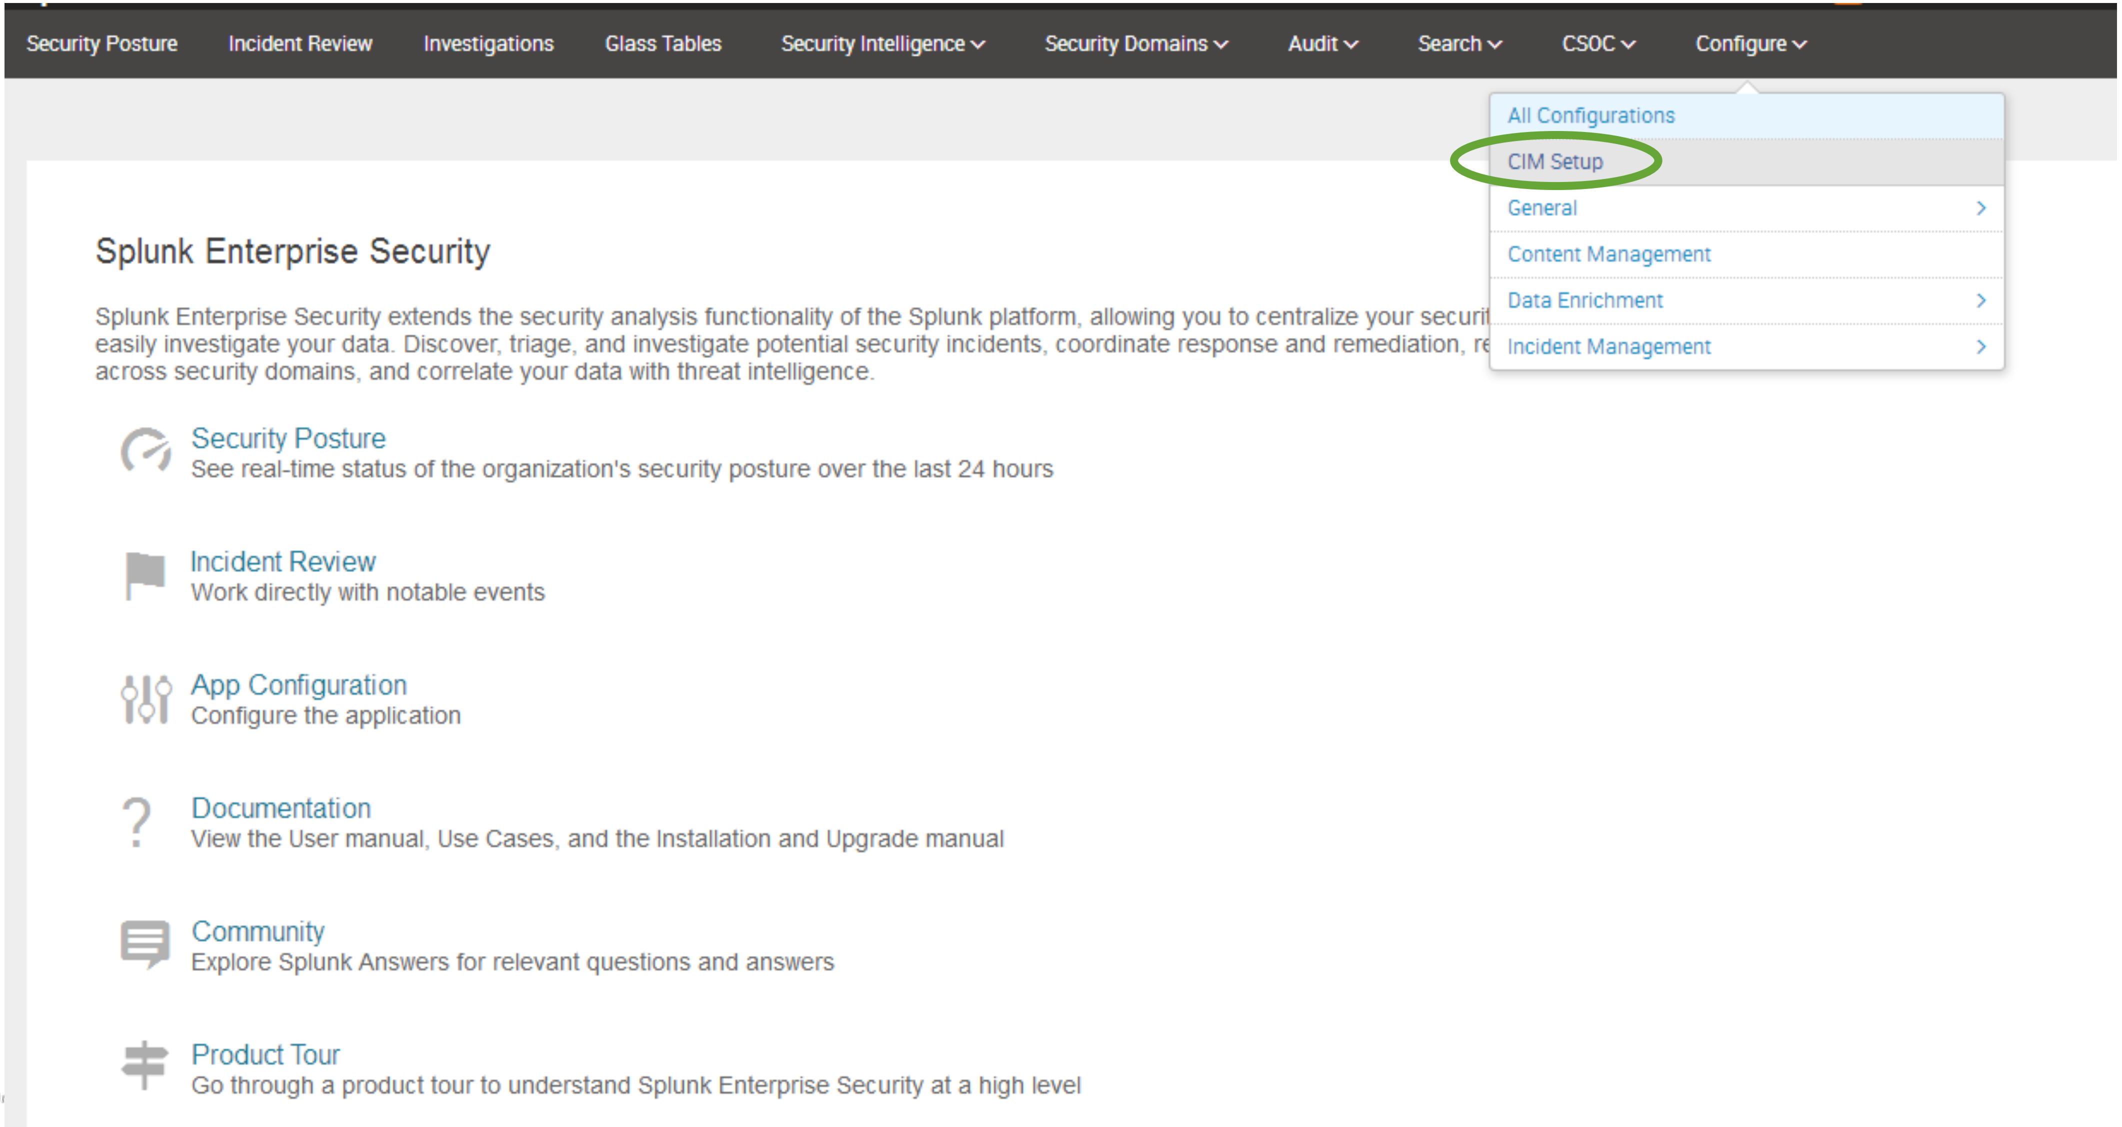

Splunk Enterprise Security

Splunk Enterprise Security extends the security analysis functionality of the Splunk platform, allowing you to centralize your securit
easily investigate your data. Discover, triage, and investigate potential security incidents, coordinate response and remediation, re
across security domains, and correlate your data with threat intelligence.

Security Posture
See real-time status of the organization's security posture over the last 24 hours

Incident Review
Work directly with notable events

App Configuration
Configure the application

Documentation
View the User manual, Use Cases, and the Installation and Upgrade manual

Community
Explore Splunk Answers for relevant questions and answers

Product Tour
Go through a product tour to understand Splunk Enterprise Security at a high level

Optimize Constraint – Specify The Index

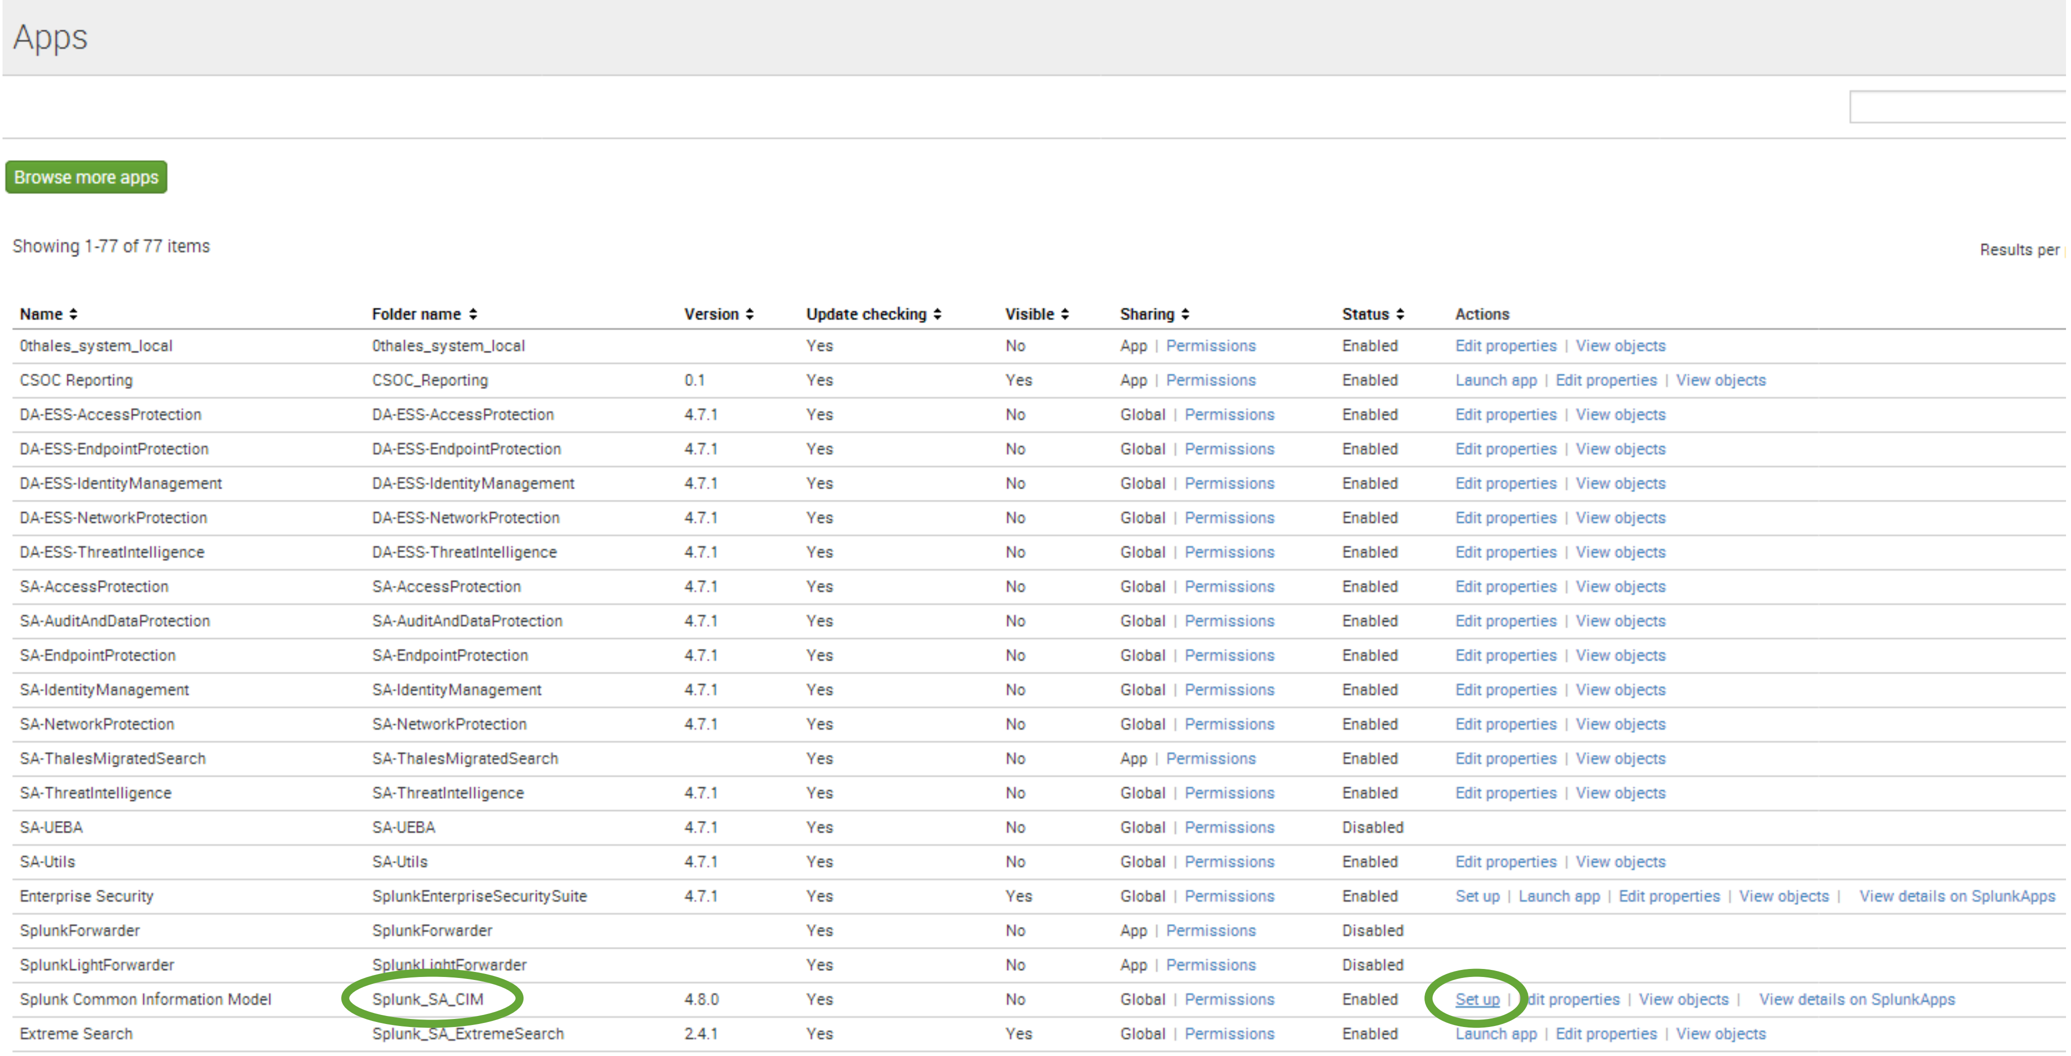

Apps

Browse more apps

Showing 1-77 of 77 items

Results per

Name ⇕	Folder name ⇕	Version ⇕	Update checking ⇕	Visible ⇕	Sharing ⇕	Status ⇕	Actions
0thales_system_local	0thales_system_local		Yes	No	App | Permissions	Enabled	Edit properties | View objects
CSOC Reporting	CSOC_Reporting	0.1	Yes	Yes	App | Permissions	Enabled	Launch app | Edit properties | View objects
DA-ESS-AccessProtection	DA-ESS-AccessProtection	4.7.1	Yes	No	Global | Permissions	Enabled	Edit properties | View objects
DA-ESS-EndpointProtection	DA-ESS-EndpointProtection	4.7.1	Yes	No	Global | Permissions	Enabled	Edit properties | View objects
DA-ESS-IdentityManagement	DA-ESS-IdentityManagement	4.7.1	Yes	No	Global | Permissions	Enabled	Edit properties | View objects
DA-ESS-NetworkProtection	DA-ESS-NetworkProtection	4.7.1	Yes	No	Global | Permissions	Enabled	Edit properties | View objects
DA-ESS-ThreatIntelligence	DA-ESS-ThreatIntelligence	4.7.1	Yes	No	Global | Permissions	Enabled	Edit properties | View objects
SA-AccessProtection	SA-AccessProtection	4.7.1	Yes	No	Global | Permissions	Enabled	Edit properties | View objects
SA-AuditAndDataProtection	SA-AuditAndDataProtection	4.7.1	Yes	No	Global | Permissions	Enabled	Edit properties | View objects
SA-EndpointProtection	SA-EndpointProtection	4.7.1	Yes	No	Global | Permissions	Enabled	Edit properties | View objects
SA-IdentityManagement	SA-IdentityManagement	4.7.1	Yes	No	Global | Permissions	Enabled	Edit properties | View objects
SA-NetworkProtection	SA-NetworkProtection	4.7.1	Yes	No	Global | Permissions	Enabled	Edit properties | View objects
SA-ThalesMigratedSearch	SA-ThalesMigratedSearch		Yes	No	App | Permissions	Enabled	Edit properties | View objects
SA-ThreatIntelligence	SA-ThreatIntelligence	4.7.1	Yes	No	Global | Permissions	Enabled	Edit properties | View objects
SA-UEBA	SA-UEBA	4.7.1	Yes	No	Global | Permissions	Disabled	
SA-Utils	SA-Utils	4.7.1	Yes	No	Global | Permissions	Enabled	Edit properties | View objects
Enterprise Security	SplunkEnterpriseSecuritySuite	4.7.1	Yes	Yes	Global | Permissions	Enabled	Set up | Launch app | Edit properties | View objects | View details on SplunkApps
SplunkForwarder	SplunkForwarder		Yes	No	App | Permissions	Disabled	
SplunkLightForwarder	SplunkLightForwarder		Yes	No	App | Permissions	Disabled	
Splunk Common Information Model	Splunk_SA_CIM	4.8.0	Yes	No	Global | Permissions	Enabled	Set up | Edit properties | View objects | View details on SplunkApps
Extreme Search	Splunk_SA_ExtremeSearch	2.4.1	Yes	Yes	Global | Permissions	Enabled	Launch app | Edit properties | View objects

Optimize Constraint – Specify The Index

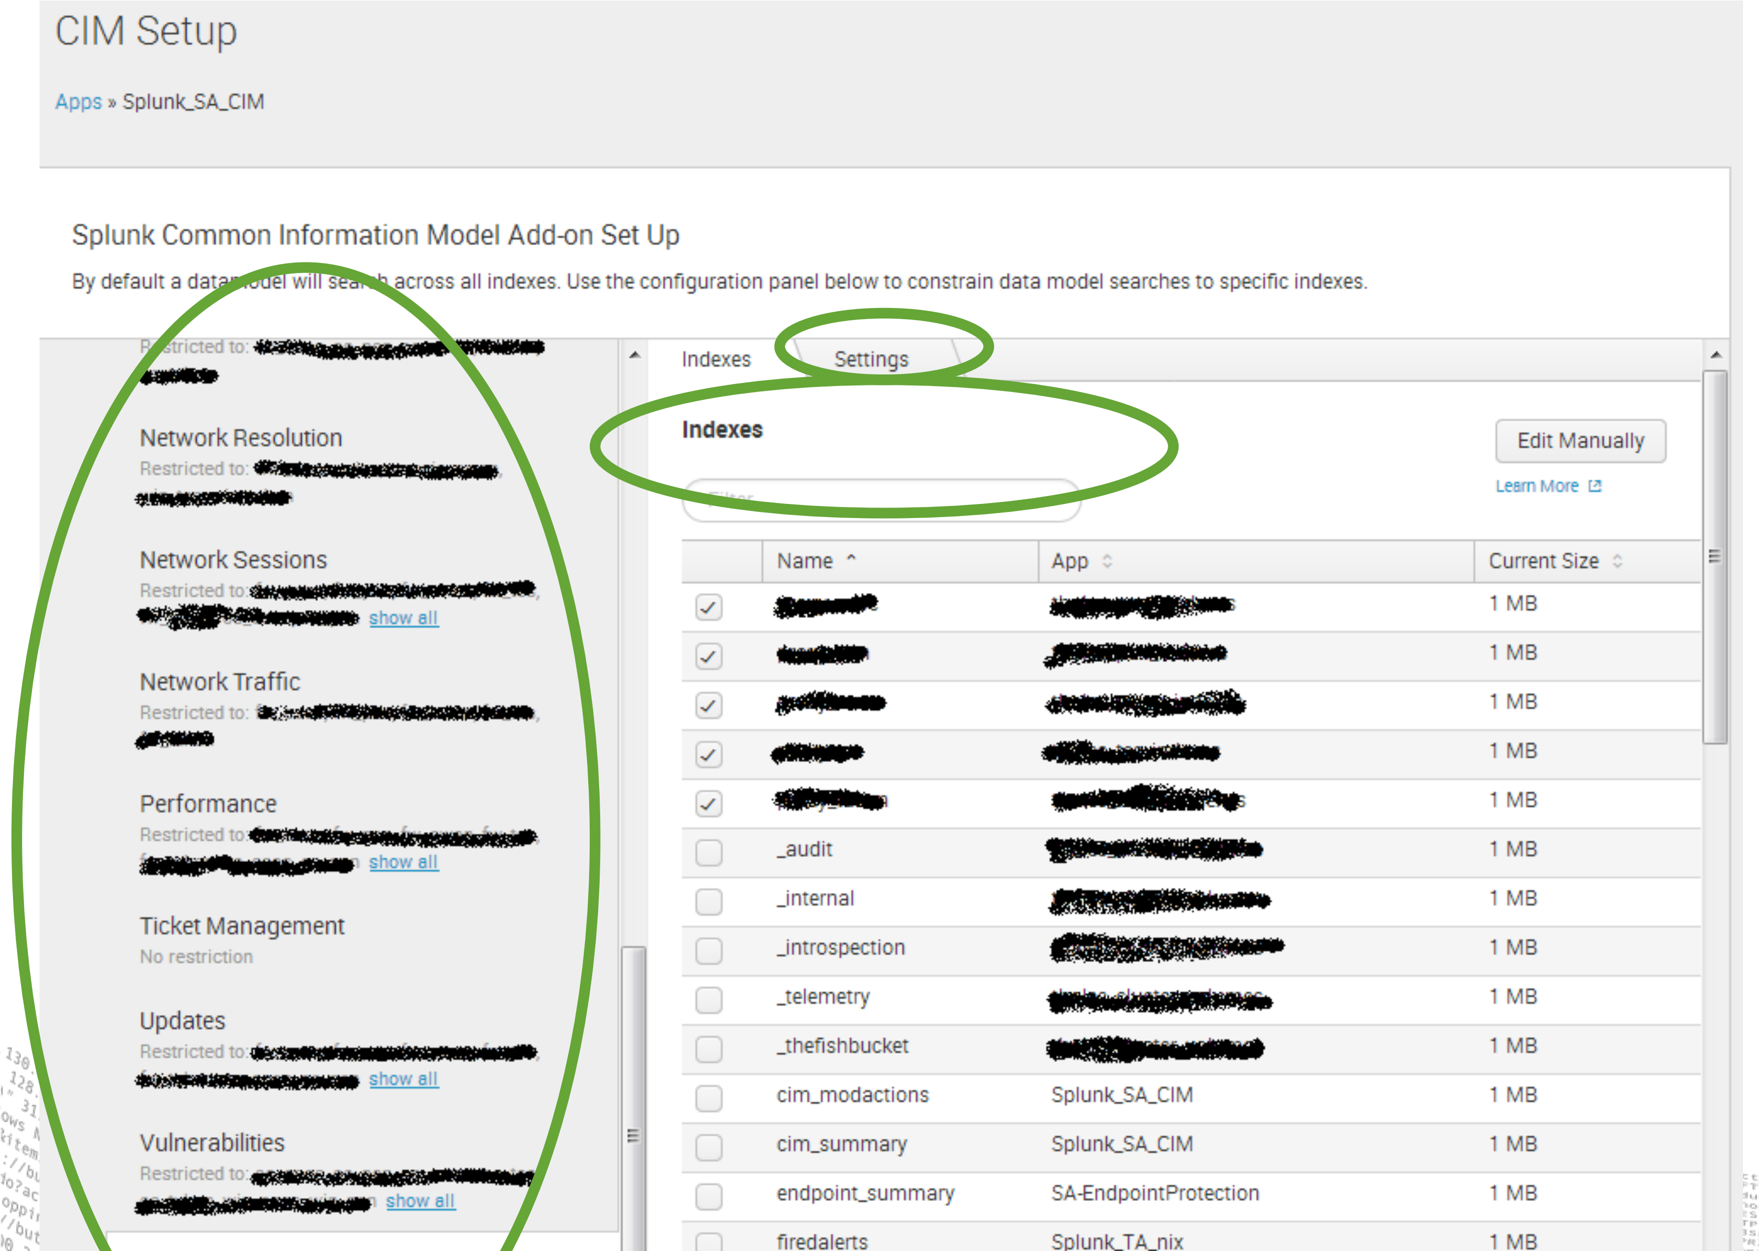

CIM Setup

Apps » Splunk_SA_CIM

Splunk Common Information Model Add-on Set Up

By default a data model will search across all indexes. Use the configuration panel below to constrain data model searches to specific indexes.

Restricted to: [redacted]

Network Resolution
Restricted to: [redacted]

Network Sessions
Restricted to: [redacted] show all

Network Traffic
Restricted to: [redacted]

Performance
Restricted to: [redacted] show all

Ticket Management
No restriction

Updates
Restricted to: [redacted] show all

Vulnerabilities
Restricted to: [redacted] show all

Indexes

	Name	App	Current Size
☑	[redacted]	[redacted]	1 MB
☑	[redacted]	[redacted]	1 MB
☑	[redacted]	[redacted]	1 MB
☑	[redacted]	[redacted]	1 MB
☑	[redacted]	[redacted]	1 MB
☐	_audit	[redacted]	1 MB
☐	_internal	[redacted]	1 MB
☐	_introspection	[redacted]	1 MB
☐	_telemetry	[redacted]	1 MB
☐	_thefishbucket	[redacted]	1 MB
☐	cim_modactions	Splunk_SA_CIM	1 MB
☐	cim_summary	Splunk_SA_CIM	1 MB
☐	endpoint_summary	SA-EndpointProtection	1 MB
☐	firedalerts	Splunk_TA_nix	1 MB

- Don't put all your data in the same index
- (nor in hundreds!)

The Default VS Local Problem

- Let's start the story with version A of the Malware data model:

/opt/splunk/etc/apps/Splunk_SA_CIM/**default**/data/models/Malware.json

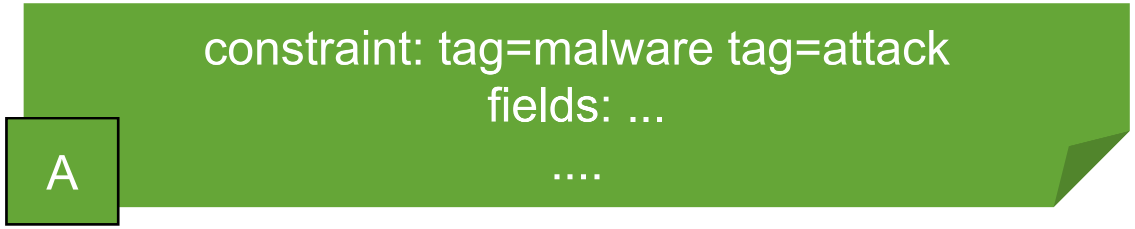

constraint: tag=malware tag=attack
fields: ...
....

The Default VS Local Problem

- One day we add a field to the data model, creating version A*:

/opt/splunk/etc/apps/Splunk_SA_CIM/**default**/data/models/Malware.json

constraint: tag=malware tag=attack
fields: ...
....

A

/opt/splunk/etc/apps/Splunk_SA_CIM/**local**/data/models/Malware.json

constraint: tag=malware tag=attack
fields: ... **+ custom field**
....

A*

The Default VS Local Problem

- Later we upgrade Splunk_SA_CIM from version A to version B:

/opt/splunk/etc/apps/Splunk_SA_CIM/**default**/data/models/Malware.json

B

constraint: **`cim_Malware_indexes`** tag=malware tag=attack
fields: ...
....

/opt/splunk/etc/apps/Splunk_SA_CIM/**local**/data/models/Malware.json

A*

constraint: tag=malware tag=attack
fields: ... **+ custom field**
....

The Default VS Local Solution

- The advice you hear: "clone before you modify"
  - I don't see the point as it doesn't solve the problem
- My advice: be paranoid with upgrades!

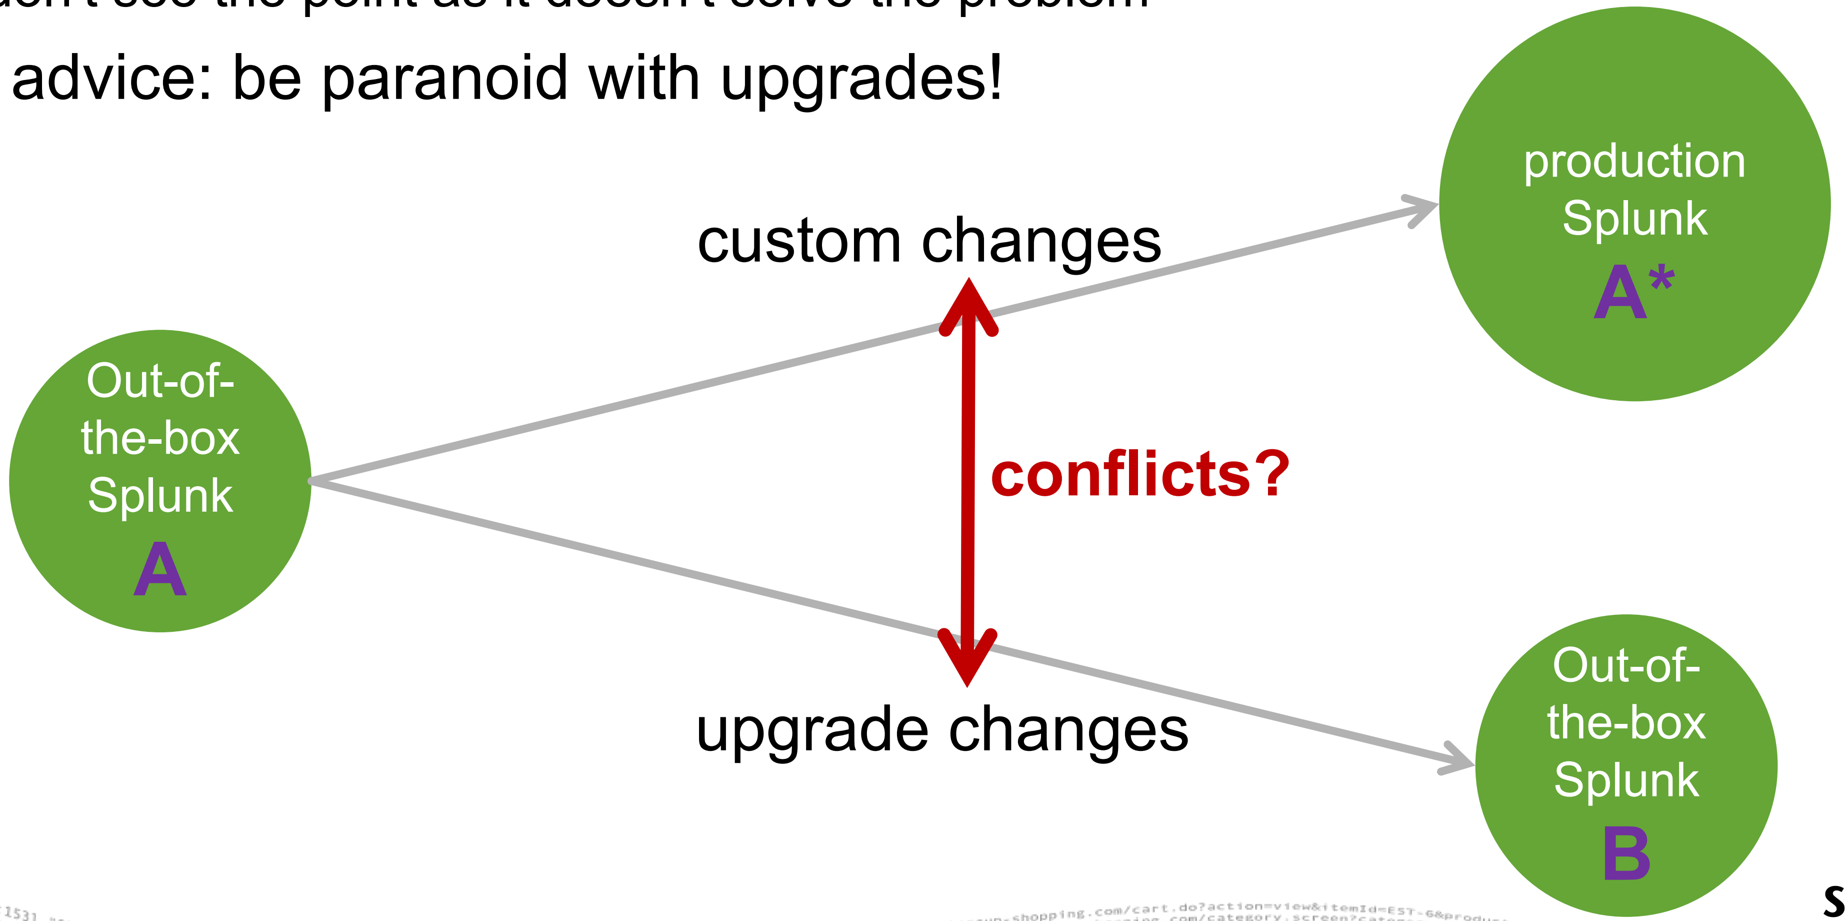

Other Optimization

- Disable DM acceleration for DM you don't use (duh!)
  - Warning: they can come back to life automatically: https://docs.splunk.com/Documentation/ES/4.7.2/Install/Datamodels#Data_model_acceleration_enforcement
- Don't enable DM acceleration on all your search heads (or move to SHC!)
- Reduce cardinality in DM:
  - E.g. session ID in network traffic
- Tweak retention:
  - Performance impact is small, space impact is potentially big

Monitoring Skipped Searches

Dashboard Source included!

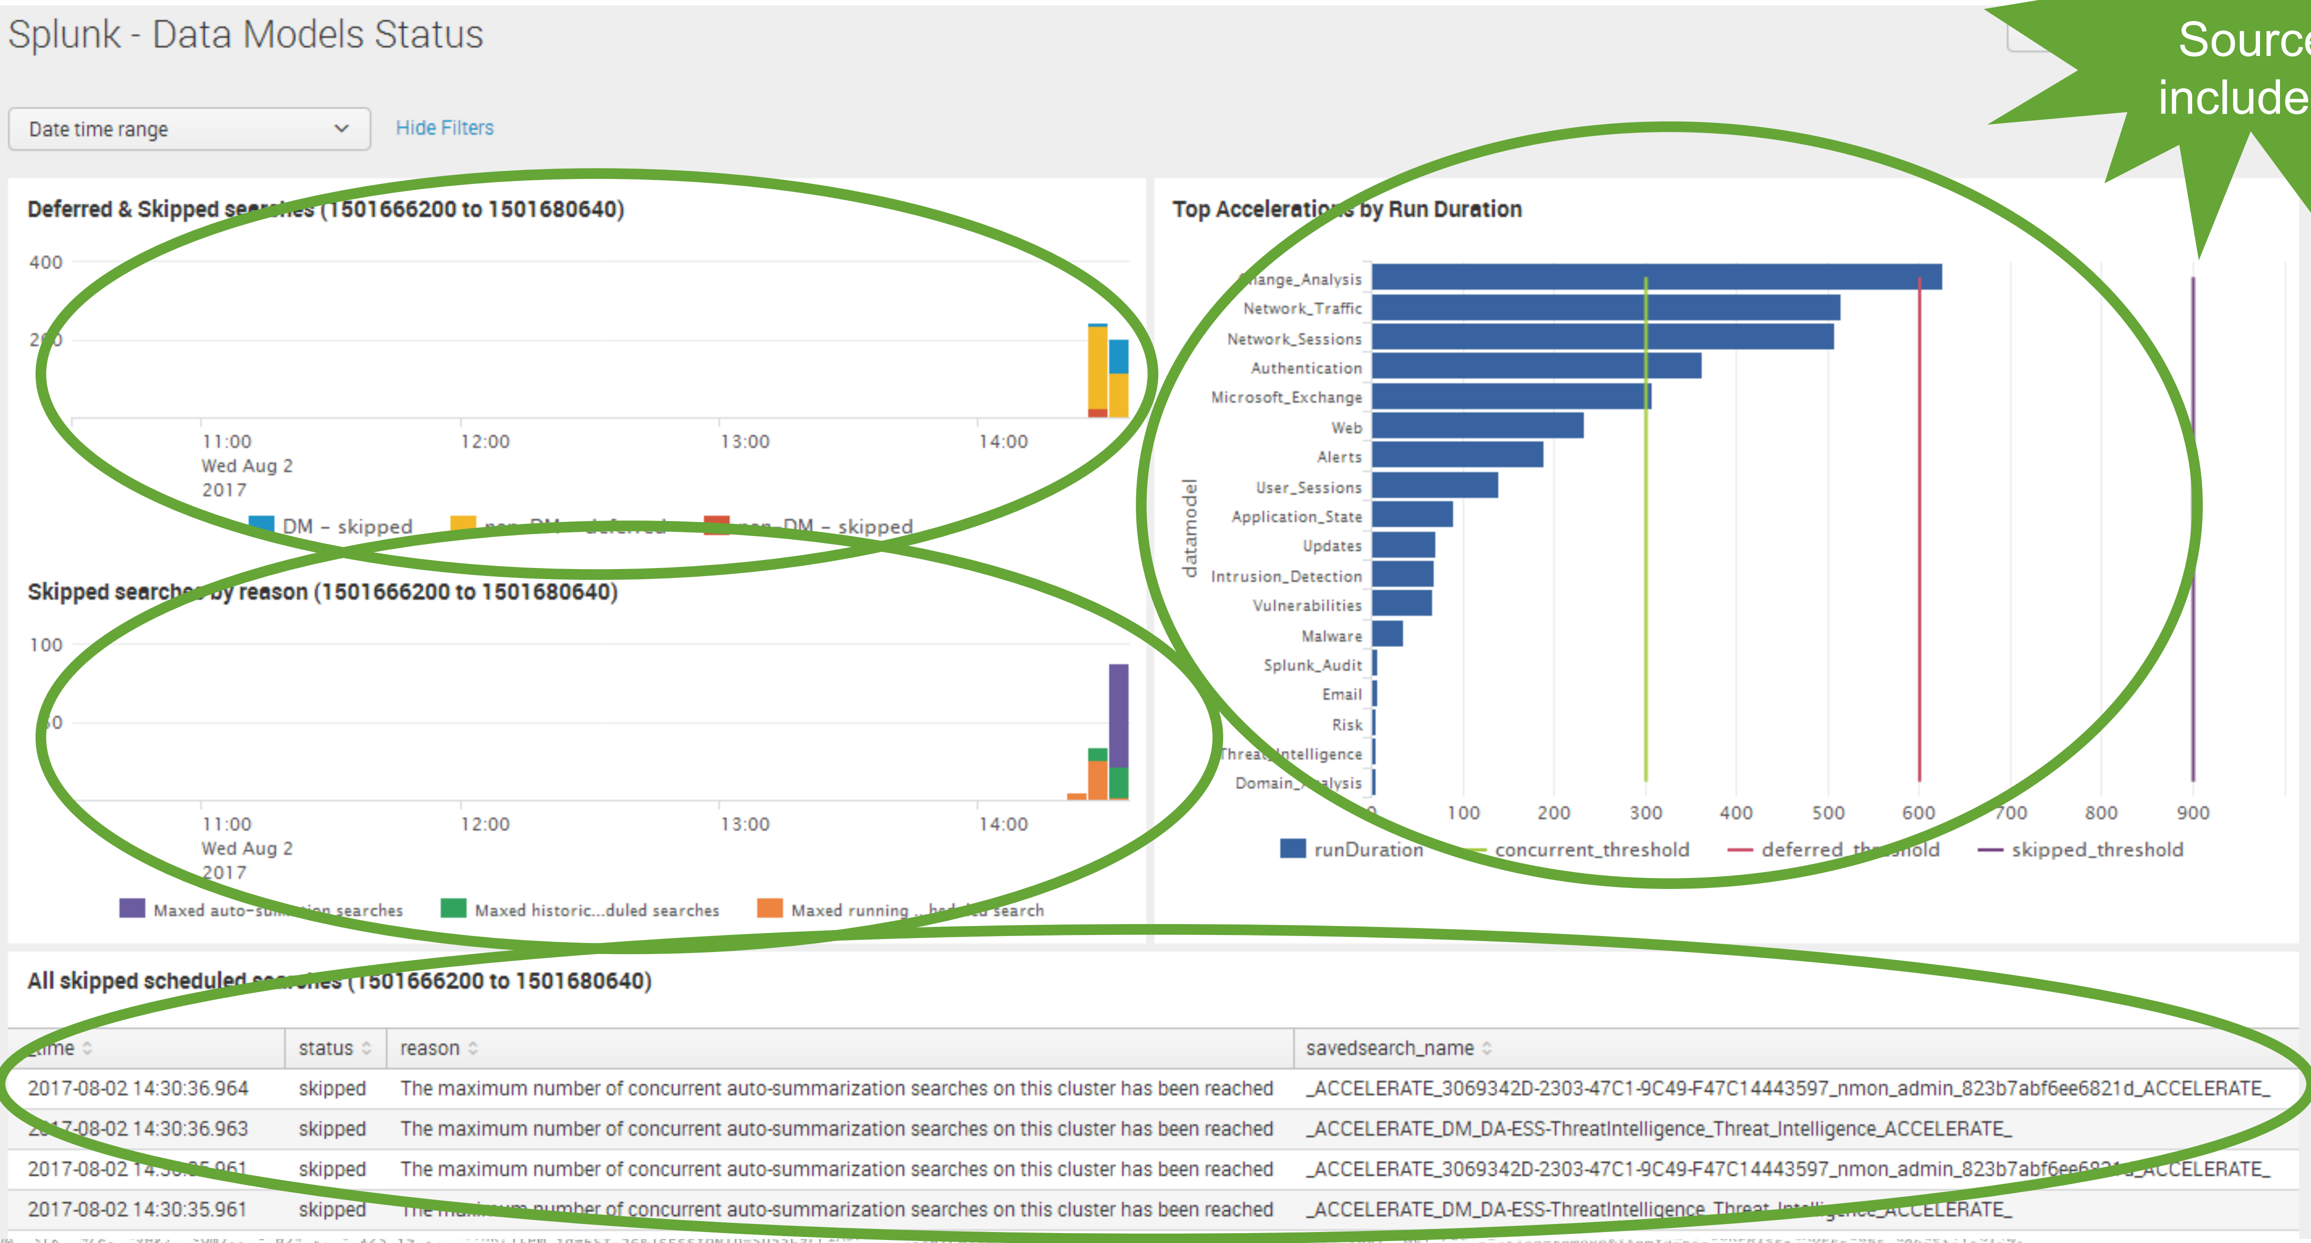

Splunk - Data Models Status

All skipped scheduled searches (1501666200 to 1501680640)

time	status	reason	savedsearch_name
2017-08-02 14:30:36.964	skipped	The maximum number of concurrent auto-summarization searches on this cluster has been reached	_ACCELERATE_3069342D-2303-47C1-9C49-F47C14443597_nmon_admin_823b7abf6ee6821d_ACCELERATE_
2017-08-02 14:30:36.963	skipped	The maximum number of concurrent auto-summarization searches on this cluster has been reached	_ACCELERATE_DM_DA-ESS-ThreatIntelligence_Threat_Intelligence_ACCELERATE_
2017-08-02 14:30:35.961	skipped	The maximum number of concurrent auto-summarization searches on this cluster has been reached	_ACCELERATE_3069342D-2303-47C1-9C49-F47C14443597_nmon_admin_823b7abf6ee6821d_ACCELERATE_
2017-08-02 14:30:35.961	skipped	The maximum number of concurrent auto-summarization searches on this cluster has been reached	_ACCELERATE_DM_DA-ESS-ThreatIntelligence_Threat_Intelligence_ACCELERATE_

Reducing Skipped Searches

- Consider increasing quota of scheduled OR auto-summary searches (last resort!)
  - limits.conf:

```
[scheduler]
max_searches_perc = 50
max_auto_summary_searches = 50
```

- **Careful:** this might affect non-summary searches negatively!

- **Careful:** don't allow a powerful Search Head tier to overwhelm a struggling Indexer tier

Last Word: Protect Your Accuracy

- Enable data summary replication on your indexer cluster
  - Not default, but recommended by splunk
  - Don't do it for performance, do it for accuracy!
  - https://conf.splunk.com/files/2016/slides/replication-of-summary-data-in-indexer-cluster.pdf
  - In server.conf on Cluster Master :

```
[clustering]
summary_replication=true
```

- Consider increasing max_concurrent for a lagging DM:
  - In datamodels.conf:

```
acceleration.max_concurrent = 3
```

  - Only do so if you have spare resources! (last-ish resort)

Want To Tinker More?

- That's all for Data Model Acceleration!
- There is a lot more you should do to optimize **everything else** in splunk (see further reading)

The End...

Thanks!

I hope you liked it!

1. **Slides** and **recording** available on http://conf.splunk.com/sessions/2017-sessions.html in a few weeks
2. Slides available **now** if you email me at gabriel.vasseur@uk.thalesgroup.com Includes **source code** & **further reading**!
3. **Rate** this session in the app :-)
4. **Poke** me on the .conf app if you want to hang out!

Thank You

Don't forget to rate this session in the .conf2017 mobile app

splunk> .conf2017

Thank You

Don't forget to rate this session in the .conf2017 mobile app

splunk> .conf2017

Further Reading

On Data Models

- **Acceleration docs**

http://docs.splunk.com/Documentation/Splunk/latest/Knowledge/Acceleratedatamodels

- **ES & Data Models docs**

https://docs.splunk.com/Documentation/ES/4.7.2/Install/Datamodels

- conf 2016 talk @ http://conf.splunk.com/sessions/2016-sessions.html

"The power of data normalization: a look at CIM under the hood"

Mark Bonsack & Vladimir Skoryk

On Searches

conf 2016 talks @ http://conf.splunk.com/sessions/2016-sessions.html

- **"Behind the magnifying glass: how search works"**

Jeff Champagne

- **"How to scale: from _raw to tstats"**

David Veuve

- **"Fields, indexed tokens and you"**

Martin Muller

On Optimizing Everything Else

conf 2016 talks @ http://conf.splunk.com/sessions/2016-sessions.html

- **Jiffy-lube quick tune-up for your splunk environment**

Jeff Champagne & Sean Delaney

- **Architectural anti-pattern: it seemed like a good idea at the time**

David Paper & Duane Waddle

- **Worst practices and how to fix them**

Jeff Champagne

Source Code

Search Source

- Shared for inspiration only – no guarantee!
- No support is included :-)
- But I might help if you ask kindly and I'm not crazy busy
- Works for me and my splunk setup, but your mileage may vary

Rebuild Monitor Source

```
TODO: Get the custom_decorations.css from the "Splunk 6.x Dashboard Examples" app and place it in etc/apps/YOURAPPNAMEHERE/appserver/static/
TODO: Replace any occurrence of "=Web" below with whichever DM you are using in your dashboard.
NOTE: $value_never_set$ is not supposed to be set :-)

<form stylesheet="YOURAPPNAMEHERE:custom_decorations.css">
...
  <row>
    <panel depends="$value_never_set$">
      <input type="dropdown" token="DM_earliest_token" searchWhenChanged="false">
        <search>
          <query>| `datamodel("Splunk_Audit", "Datamodel_Acceleration")` | `drop_dm_object_name("Datamodel_Acceleration")`
               | eval retention_days=retention/(24*60*60)| eval DM_earliest=if(retention==0,0,now()-retention)
               | table datamodel retention retention_days DM_earliest | eval DM_earliest_human=strftime(DM_earliest,"%c")
               | search datamodel=Web</query>
          <earliest>0</earliest>
        </search>
        <fieldForLabel>DM_earliest</fieldForLabel>
        <fieldForValue>DM_earliest</fieldForValue>
        <selectFirstChoice>true</selectFirstChoice>
      </input>
    </panel>
  </row>
  <row>
    <panel>
      <single>
        <title>Oldest record in DM</title>
        <search>
          <query>| tstats `summariesonly` min(_time) as oldest from datamodel=Web |  eval oldest=strftime(oldest,"%c")</query>
          <earliest>0</earliest>
          <latest></latest>
        </search>
        <option name="height">50</option>
        <option name="drilldown">none</option>
      </single>
    </panel>
    <panel>
      <search>
        <query>| stats count as DM_earliest | eval DM_earliest="$DM_earliest_token$"
               | eval timepicker="$time_token.earliest$"
               | eval timepicker_earliest=if(match(timepicker,"^\d+$"),timepicker,relative_time(now(),timepicker))
               | eval is_in_range=if(timepicker_earliest&gt;=DM_earliest,1,0)
               | eval is_in_range_human=if(is_in_range==1,"YES","NO")
               | rangemap field=is_in_range low=1-1 severe=0-0</query>
        <earliest>0</earliest>
        <progress>
          <set token="dm_range_value">$result.is_in_range_human$</set>
          <set token="dm_range">$result.range$</set>
        </progress>
      </search>
      <html>
        <div class="panel-head dashboard-element-header">
          <h3 class="dashboard-element-title">Search within DM range?</h3>
        </div>
        <div class="custom-result-value icon-only $dm_range$"> </div>
      </html>
    </panel>
  </row>
...
</form>
```

Rebuild Monitor Source

```
<form>
  <label>Splunk - Data Models Rebuild Monitor</label>
  <fieldset submitButton="false">
    <input type="dropdown" token="dm">
      <label>Data model</label>
      <search>
        <query>| `datamodel("Splunk_Audit", "Datamodel_Acceleration")` | `drop_dm_object_name("Datamodel_Acceleration")` | fields datamodel |  sort 100 + datamodel</query>
      </search>
      <fieldForLabel>datamodel</fieldForLabel>
      <fieldForValue>datamodel</fieldForValue>
    </input>
    <input type="dropdown" token="earliest_token" depends="$value_never_set$">
      <label>field1</label>
      <fieldForLabel>acceleration.earliest_time</fieldForLabel>
      <fieldForValue>acceleration.earliest_time</fieldForValue>
      <search>
        <query>| rest /services/configs/conf-datamodels | search title=$dm$ | fields acceleration.earliest_time</query>
        <earliest>0</earliest>
        <latest></latest>
      </search>
      <selectFirstChoice>true</selectFirstChoice>
    </input>
  </fieldset>
  <row>
    <panel>
      <html>
      <h1>$dm$ data model config</h1>
    </html>
    </panel>
  </row>
  <row>
    <panel>
      <single>
        <search>
          <query>| rest /services/configs/conf-datamodels | search title=$dm$ | fields acceleration.earliest_time</query>
          <earliest>@d</earliest>
          <latest>now</latest>
        </search>
        <option name="colorBy">value</option>
        <option name="colorMode">none</option>
        <option name="drilldown">all</option>
        <option name="height">30</option>
        <option name="numberPrecision">0</option>
        <option name="showSparkline">1</option>
        <option name="showTrendIndicator">1</option>
        <option name="trendColorInterpretation">standard</option>
        <option name="trendDisplayMode">absolute</option>
        <option name="underLabel">Retention (earliest)</option>
        <option name="unitPosition">after</option>
        <option name="useColors">0</option>
        <option name="useThousandSeparators">1</option>
      </single>
      <single>
        <search>
          <query>| rest /services/configs/conf-datamodels | search title=$dm$ | fields acceleration.backfill_time</query>
          <earliest>@d</earliest>
          <latest>now</latest>
        </search>
        <option name="colorBy">value</option>
        <option name="colorMode">none</option>
        <option name="drilldown">all</option>
        <option name="height">30</option>
        <option name="numberPrecision">0</option>
        <option name="showSparkline">1</option>
        <option name="showTrendIndicator">1</option>
        <option name="trendColorInterpretation">standard</option>
        <option name="trendDisplayMode">absolute</option>
        <option name="underLabel">Backfill target</option>
        <option name="unitPosition">before</option>
        <option name="useColors">0</option>
        <option name="useThousandSeparators">1</option>
      </single>
      <single>
        <search>
          <query>| `datamodel("Splunk_Audit", "Datamodel_Acceleration")` | `drop_dm_object_name("Datamodel_Acceleration")` | search datamodel=$dm$ | eval complete(%)=round(complete*100,1)."%" | fields complete(%)</query>
          <earliest>0.000</earliest>
        </search>
        <option name="colorBy">value</option>
        <option name="colorMode">none</option>
        <option name="drilldown">all</option>
        <option name="height">30</option>
        <option name="numberPrecision">0</option>
        <option name="showSparkline">1</option>
        <option name="showTrendIndicator">1</option>
        <option name="trendColorInterpretation">standard</option>
        <option name="trendDisplayMode">absolute</option>
        <option name="underLabel">Backfill complete</option>
        <option name="unitPosition">after</option>
        <option name="useColors">0</option>
        <option name="useThousandSeparators">1</option>
      </single>
      <single>
        <search>
          <query>| rest /services/configs/conf-datamodels | search title=$dm$ | fields acceleration.max_concurrent</query>
          <earliest>@d</earliest>
          <latest>now</latest>
        </search>
        <option name="colorBy">value</option>
        <option name="colorMode">none</option>
        <option name="drilldown">all</option>
        <option name="height">30</option>
        <option name="numberPrecision">0</option>
        <option name="showSparkline">1</option>
        <option name="showTrendIndicator">1</option>
        <option name="trendColorInterpretation">standard</option>
        <option name="trendDisplayMode">absolute</option>
        <option name="underLabel">max concurrent</option>
        <option name="unitPosition">after</option>
        <option name="useColors">0</option>
        <option name="useThousandSeparators">1</option>
      </single>
      <single>
        <search>
          <query>| rest /services/configs/conf-datamodels | search title=$dm$ | fields acceleration.max_time</query>
          <earliest>@d</earliest>
          <latest>now</latest>
        </search>
        <option name="colorBy">value</option>
        <option name="colorMode">none</option>
        <option name="drilldown">all</option>
        <option name="height">30</option>
        <option name="numberPrecision">0</option>
        <option name="showSparkline">1</option>
        <option name="showTrendIndicator">1</option>
        <option name="trendColorInterpretation">standard</option>
        <option name="trendDisplayMode">absolute</option>
        <option name="underLabel">max time</option>
        <option name="unitPosition">after</option>
        <option name="useColors">0</option>
        <option name="useThousandSeparators">1</option>
      </single>
      <single>
        <search>
          <query>| `datamodel("Splunk_Audit", "Datamodel_Acceleration")` | `drop_dm_object_name("Datamodel_Acceleration")` | search datamodel=$dm$ | fields size | eval size=round(size/1048576,1)." MB"</query>
          <earliest>@d</earliest>
          <latest>now</latest>
        </search>
        <option name="colorBy">value</option>
        <option name="colorMode">none</option>
        <option name="drilldown">all</option>
        <option name="height">30</option>
        <option name="numberPrecision">0</option>
        <option name="rangeColors">["0x65a637","0x6db7c6","0xf7bc38","0xf58f39","0xd93f3c"]</option>
        <option name="rangeValues">[0,30,70,100]</option>
        <option name="showSparkline">1</option>
        <option name="showTrendIndicator">1</option>
        <option name="trendColorInterpretation">standard</option>
        <option name="trendDisplayMode">absolute</option>
        <option name="underLabel">data size</option>
        <option name="unitPosition">after</option>
        <option name="useColors">0</option>
        <option name="useThousandSeparators">1</option>
      </single>
    </panel>
  </row>
  <row>
    <panel>
      <html>
      <h1>$dm$ data model acceleration state</h1>
    </html>
    </panel>
  </row>
  <row>
    <panel>
      <title>$dm$ event counts - Monitor lag and backfill</title>
      <chart>
        <search>
          <query>| tstats prestats=t summariesonly=t allow_old_summaries=t count from datamodel=$dm$ by _time span=10s| timechart count span=10s</query>
          <earliest>-2h</earliest>
          <latest>now</latest>
        </search>
        <option name="height">275</option>
        <option name="charting.axisLabelsX.majorLabelStyle.overflowMode">ellipsisNone</option>
        <option name="charting.axisLabelsX.majorLabelStyle.rotation">0</option>
        <option name="charting.axisTitleX.visibility">collapsed</option>
        <option name="charting.axisTitleY.visibility">collapsed</option>
        <option name="charting.axisTitleY2.visibility">visible</option>
        <option name="charting.axisX.scale">linear</option>
        <option name="charting.axisY.scale">linear</option>
        <option name="charting.axisY2.enabled">0</option>
        <option name="charting.axisY2.scale">inherit</option>
        <option name="charting.chart">line</option>
        <option name="charting.chart.bubbleMaximumSize">50</option>
        <option name="charting.chart.bubbleMinimumSize">10</option>
        <option name="charting.chart.bubbleSizeBy">area</option>
        <option name="charting.chart.nullValueMode">gaps</option>
        <option name="charting.chart.showDataLabels">none</option>
        <option name="charting.chart.sliceCollapsingThreshold">0.01</option>
        <option name="charting.chart.stackMode">default</option>
        <option name="charting.chart.style">shiny</option>
        <option name="charting.drilldown">all</option>
        <option name="charting.layout.splitSeries">0</option>
        <option name="charting.layout.splitSeries.allowIndependentYRanges">0</option>
        <option name="charting.legend.labelStyle.overflowMode">ellipsisMiddle</option>
        <option name="charting.legend.placement">none</option>
      </chart>
      <chart>
        <search>
          <query>| tstats prestats=t allow_old_summaries=t summariesonly=t count from datamodel=$dm$ by _time span=4h| timechart count span=4h</query>
          <earliest>$earliest_token$-1w</earliest>
          <latest>now</latest>
        </search>
        <option name="height">275</option>
        <option name="charting.axisLabelsX.majorLabelStyle.overflowMode">ellipsisNone</option>
        <option name="charting.axisLabelsX.majorLabelStyle.rotation">0</option>
        <option name="charting.axisTitleX.visibility">collapsed</option>
        <option name="charting.axisTitleY.visibility">collapsed</option>
        <option name="charting.axisTitleY2.visibility">visible</option>
        <option name="charting.axisX.scale">linear</option>
        <option name="charting.axisY.scale">linear</option>
        <option name="charting.axisY2.enabled">0</option>
        <option name="charting.axisY2.scale">inherit</option>
        <option name="charting.chart">line</option>
        <option name="charting.chart.bubbleMaximumSize">50</option>
        <option name="charting.chart.bubbleMinimumSize">10</option>
        <option name="charting.chart.bubbleSizeBy">area</option>
        <option name="charting.chart.nullValueMode">gaps</option>
        <option name="charting.chart.showDataLabels">none</option>
        <option name="charting.chart.sliceCollapsingThreshold">0.01</option>
        <option name="charting.chart.stackMode">default</option>
        <option name="charting.chart.style">shiny</option>
        <option name="charting.drilldown">all</option>
        <option name="charting.layout.splitSeries">0</option>
        <option name="charting.layout.splitSeries.allowIndependentYRanges">0</option>
        <option name="charting.legend.labelStyle.overflowMode">ellipsisMiddle</option>
        <option name="charting.legend.placement">none</option>
      </chart>
    </panel>
    <panel>
      <table>
        <title>$dm$ recent acceleration jobs</title>
        <search>
          <query>index=_internal source=*scheduler.log _ACCELERATE_DM_*$dm$_ACCELERATE_
| eval scheduled=strftime(scheduled_time,"%c")
| stats values(scheduled) as scheduled, values(scheduled_time) as scheduled_time, list(status) as statuses, values(run_time) as run_time by savedsearch_name sid
| sort - scheduled_time
| eval done=if(isnull(run_time),"running","done")
| eval run_time=tostring(if(isnull(run_time),now()-scheduled_time,run_time),"duration") | fields - scheduled_time savedsearch_name sid</query>
          <earliest>@d</earliest>
          <latest>now</latest>
        </search>
        <option name="count">10</option>
        <option name="dataOverlayMode">none</option>
        <option name="drilldown">cell</option>
        <option name="rowNumbers">false</option>
        <option name="wrap">true</option>
      </table>
    </panel>
  </row>
</form>
```

Skipped Searches Status Source

NOTE: you might want to limit host=* to your search heads only.

```
<form>
  <label>Splunk - Data Models Status</label>
  <fieldset submitButton="false">
    <input type="time" token="timepicker1">
      <label></label>
      <default>
        <earliest>-4h@m</earliest>
        <latest>now</latest>
      </default>
    </input>
  </fieldset>
  <row>
    <panel>
      <chart>
        <title>Skipped searches ($timepicker1.earliest$ to $timepicker1.latest$)</title>
        <search>
          <query>index=_internal host=* sourcetype=scheduler status="skipped"             | eval type=if(match(savedsearch_name,"^_ACCELERATE_"),"DM","non-DM") | eval reason = if(isnull(reason) OR reason == "", "none", reason)  | eval combo=type . " - " . reason           | timechart span=5m count by combo</query>
          <earliest>$timepicker1.earliest$</earliest>
          <latest>$timepicker1.latest$</latest>
        </search>
        <option name="charting.axisLabelsX.majorLabelStyle.overflowMode">ellipsisNone</option>
        <option name="charting.axisLabelsX.majorLabelStyle.rotation">0</option>
        <option name="charting.axisTitleX.visibility">collapsed</option>
        <option name="charting.axisTitleY.visibility">visible</option>
        <option name="charting.axisTitleY2.visibility">visible</option>
        <option name="charting.axisX.scale">linear</option>
        <option name="charting.axisY.scale">linear</option>
        <option name="charting.axisY2.enabled">0</option>
        <option name="charting.axisY2.scale">inherit</option>
        <option name="charting.chart">column</option>
        <option name="charting.chart.bubbleMaximumSize">50</option>
        <option name="charting.chart.bubbleMinimumSize">10</option>
        <option name="charting.chart.bubbleSizeBy">area</option>
        <option name="charting.chart.nullValueMode">gaps</option>
        <option name="charting.chart.showDataLabels">none</option>
        <option name="charting.chart.sliceCollapsingThreshold">0.01</option>
        <option name="charting.chart.stackMode">stacked</option>
        <option name="charting.chart.style">shiny</option>
        <option name="charting.drilldown">all</option>
        <option name="charting.layout.splitSeries">0</option>
        <option name="charting.layout.splitSeries.allowIndependentYRanges">0</option>
        <option name="charting.legend.labelStyle.overflowMode">ellipsisMiddle</option>
        <option name="charting.legend.placement">bottom</option>
        <option name="height">200</option>
      </chart>
      <chart>
        <title>Deferred &amp; Skipped searches ($timepicker1.earliest$ to $timepicker1.latest$)</title>
        <search>
          <query>index=_internal sourcetype=scheduler status=continued OR status=skipped  | eval type=if(match(savedsearch_name,"^_ACCELERATE_"),"DM","non-DM") | eval status=replace(status,"continued","deferred") | eval combo=type . " - " . status| timechart span=5m count by combo</query>
          <earliest>$timepicker1.earliest$</earliest>
          <latest>$timepicker1.latest$</latest>
        </search>
        <option name="charting.axisLabelsX.majorLabelStyle.overflowMode">ellipsisNone</option>
        <option name="charting.axisLabelsX.majorLabelStyle.rotation">0</option>
        <option name="charting.axisTitleX.visibility">collapsed</option>
        <option name="charting.axisTitleY.visibility">visible</option>
        <option name="charting.axisTitleY2.visibility">visible</option>
        <option name="charting.axisX.scale">linear</option>
        <option name="charting.axisY.scale">linear</option>
        <option name="charting.axisY2.enabled">0</option>
        <option name="charting.axisY2.scale">inherit</option>
        <option name="charting.chart">column</option>
        <option name="charting.chart.bubbleMaximumSize">50</option>
        <option name="charting.chart.bubbleMinimumSize">10</option>
        <option name="charting.chart.bubbleSizeBy">area</option>
        <option name="charting.chart.nullValueMode">gaps</option>
        <option name="charting.chart.showDataLabels">none</option>
        <option name="charting.chart.sliceCollapsingThreshold">0.01</option>
        <option name="charting.chart.stackMode">stacked</option>
        <option name="charting.chart.style">shiny</option>
        <option name="charting.drilldown">all</option>
        <option name="charting.layout.splitSeries">0</option>
        <option name="charting.layout.splitSeries.allowIndependentYRanges">0</option>
        <option name="charting.legend.labelStyle.overflowMode">ellipsisMiddle</option>
        <option name="charting.legend.placement">bottom</option>
        <option name="height">200</option>
      </chart>
    </panel>
    <panel>
      <chart>
        <title>Top Accelerations by Run Duration</title>
        <search>
          <query>| `datamodel("Splunk_Audit", "Datamodel_Acceleration")` | `drop_dm_object_name("Datamodel_Acceleration")` | join type=outer last_sid [| rest splunk_server=* count=0 /services/search/jobs reportSearch=summarize* | rename sid as last_sid | fields last_sid,runDuration] | eval size(MB)=round(size/1048576,1) | eval retention(days)=if(retention==0,"unlimited",retention/86400) | eval complete(%)=round(complete*100,1) | eval runDuration(s)=round(runDuration,1) |  sort 18 - runDuration | table datamodel,runDuration | eval concurrent_threshold=300 | eval deferred_threshold=600 | eval skipped_threshold=900</query>
          <earliest>0.000</earliest>
          <latest></latest>
        </search>
        <option name="height">400</option>
        <option name="charting.axisLabelsX.majorLabelStyle.overflowMode">ellipsisNone</option>
        <option name="charting.axisLabelsX.majorLabelStyle.rotation">0</option>
        <option name="charting.axisTitleX.visibility">visible</option>
        <option name="charting.axisTitleY.visibility">visible</option>
        <option name="charting.axisTitleY2.visibility">visible</option>
        <option name="charting.axisX.scale">linear</option>
        <option name="charting.axisY.scale">linear</option>
        <option name="charting.axisY2.enabled">0</option>
        <option name="charting.axisY2.scale">inherit</option>
        <option name="charting.chart">bar</option>
        <option name="charting.chart.bubbleMaximumSize">50</option>
        <option name="charting.chart.bubbleMinimumSize">10</option>
        <option name="charting.chart.bubbleSizeBy">area</option>
        <option name="charting.chart.nullValueMode">gaps</option>
        <option name="charting.chart.overlayFields">concurrent_threshold,deferred_threshold,skipped_threshold</option>
        <option name="charting.chart.showDataLabels">none</option>
        <option name="charting.chart.sliceCollapsingThreshold">0.01</option>
        <option name="charting.chart.stackMode">default</option>
        <option name="charting.chart.style">shiny</option>
        <option name="charting.drilldown">all</option>
        <option name="charting.layout.splitSeries">0</option>
        <option name="charting.layout.splitSeries.allowIndependentYRanges">0</option>
        <option name="charting.legend.labelStyle.overflowMode">ellipsisMiddle</option>
        <option name="charting.legend.placement">bottom</option>
      </chart>
    </panel>
  </row>
  <row>
    <panel>
      <table>
        <title>All skipped scheduled searches ($timepicker1.earliest$ to $timepicker1.latest$)</title>
        <search>
          <query>index=_internal host=* sourcetype=scheduler status="skipped"  | table _time status savedsearch_name | sort - _time</query>
          <earliest>$timepicker1.earliest$</earliest>
          <latest>$timepicker1.latest$</latest>
        </search>
        <option name="drilldown">none</option>
      </table>
    </panel>
  </row>
  <row>
    <panel>
      <chart>
        <title>Indexing rate (kbps)</title>
        <search>
          <query>index=_internal source=*license_usage.log type=Usage | eval kbps=(b/1024)/(5*60) | timechart span=5m sum(kbps) as kbps by idx partial=f</query>
          <earliest>$timepicker1.earliest$</earliest>
          <latest>$timepicker1.latest$</latest>
        </search>
        <option name="height">290</option>
        <option name="charting.axisLabelsX.majorLabelStyle.overflowMode">ellipsisNone</option>
        <option name="charting.axisLabelsX.majorLabelStyle.rotation">0</option>
        <option name="charting.axisTitleX.visibility">visible</option>
        <option name="charting.axisTitleY.visibility">visible</option>
        <option name="charting.axisTitleY2.visibility">visible</option>
        <option name="charting.axisX.scale">linear</option>
        <option name="charting.axisY.scale">linear</option>
        <option name="charting.axisY2.enabled">0</option>
        <option name="charting.axisY2.scale">inherit</option>
        <option name="charting.chart">column</option>
        <option name="charting.chart.bubbleMaximumSize">50</option>
        <option name="charting.chart.bubbleMinimumSize">10</option>
        <option name="charting.chart.bubbleSizeBy">area</option>
        <option name="charting.chart.nullValueMode">gaps</option>
        <option name="charting.chart.showDataLabels">none</option>
        <option name="charting.chart.sliceCollapsingThreshold">0.01</option>
        <option name="charting.chart.stackMode">stacked</option>
        <option name="charting.chart.style">shiny</option>
        <option name="charting.drilldown">all</option>
        <option name="charting.layout.splitSeries">0</option>
        <option name="charting.layout.splitSeries.allowIndependentYRanges">0</option>
        <option name="charting.legend.labelStyle.overflowMode">ellipsisMiddle</option>
        <option name="charting.legend.placement">right</option>
      </chart>
    </panel>
  </row>
  <row>
    <panel>
      <chart>
        <title>Event count in 2 data models ($timepicker1.earliest$ to $timepicker1.latest$)</title>
        <search>
          <query>| tstats summariesonly=true  prestats=t count    from datamodel=Network_Traffic by _time span=1m | eval datamodel="Network_Traffic" | tstats append=t  summariesonly=true  prestats=t count    from datamodel=Authentication by _time span=1m | eval datamodel=if(isnull(datamodel),"Authentication",datamodel) | timechart span=1m count by datamodel</query>
          <earliest>$timepicker1.earliest$</earliest>
          <latest>$timepicker1.latest$</latest>
        </search>
        <option name="charting.axisLabelsX.majorLabelStyle.overflowMode">ellipsisNone</option>
        <option name="charting.axisLabelsX.majorLabelStyle.rotation">0</option>
        <option name="charting.axisTitleX.visibility">collapsed</option>
        <option name="charting.axisTitleY.visibility">visible</option>
        <option name="charting.axisTitleY2.visibility">visible</option>
        <option name="charting.axisX.scale">linear</option>
        <option name="charting.axisY.scale">linear</option>
        <option name="charting.axisY2.enabled">0</option>
        <option name="charting.axisY2.scale">inherit</option>
        <option name="charting.chart">line</option>
        <option name="charting.chart.bubbleMaximumSize">50</option>
        <option name="charting.chart.bubbleMinimumSize">10</option>
        <option name="charting.chart.bubbleSizeBy">area</option>
        <option name="charting.chart.nullValueMode">gaps</option>
        <option name="charting.chart.showDataLabels">none</option>
        <option name="charting.chart.sliceCollapsingThreshold">0.01</option>
        <option name="charting.chart.stackMode">default</option>
        <option name="charting.chart.style">shiny</option>
        <option name="charting.drilldown">all</option>
        <option name="charting.layout.splitSeries">0</option>
        <option name="charting.layout.splitSeries.allowIndependentYRanges">0</option>
        <option name="charting.legend.labelStyle.overflowMode">ellipsisMiddle</option>
        <option name="charting.legend.placement">bottom</option>
      </chart>
    </panel>
  </row>
</form>
```

“Don't believe everything you read on the internet.”

Charles Babbage