

# FlightAware: a modern enterprise built around Tcl

The screenshot displays the FlightAware website interface. At the top, there's a navigation bar with the FlightAware logo and 'LIVE FLIGHT TRACKING' text. Below this, a search bar allows users to track flights by Airline (United) and Flight # (450), or by Origin (New York) and Destination (New York). A 'Track' button is visible. To the right, there's a Breitling advertisement and a map of the United States with green dots indicating flight locations. Below the search bar, a 'RECENT' flights table lists several flights:

Airline	Flight #	Departure	Arrival
American Eagle	EGF4328	Tue 19:50 (KORD)	Tue 22:15 (CYKF)
ExpressJet	ASQ4328	Tue 16:45 (KEWR)	Tue 19:13 (KMSY)
United	UAL972	Tue 18:00 (KORD)	Wed 09:00 (EBBR / BRU)

Below the table, there's a 'PHOTOS' section with several images of aircraft. To the right, there's a 'SQUAWKS' section with several news articles, including 'F-16 Flies without a pilot', 'What does TSA do with your lost, confiscated stuff?', 'Blackhawk Flight Crew Gets A Parking Ticket', and 'Kids with autism get flying practice at NYC's JFK'. On the far right, there's a weather section for various airports, including Newark, NJ, Shenzhen, Guangdong, Houston, TX, College Station, TX, Monterey, London, England, and Atlanta, GA.

# About FlightAware - Brief History

- Founded in 2005 by three pilots/hackers, based in Houston, Texas
- Originally focused on general aviation, expanded into airlines in 2008
- Largest aviation web site in the world
  - 3,000,000 registered users
  - 150,000,000 page views/month
  - Issue 250,000 flight alerts/day
  - We are a top 100 site based on page views.
- Largest customers include United, NetJets, ARINC Direct, Boeing
- 40 employees at offices in Houston and New York City
- Data centers in Houston, Dallas, and London

# FlightAware at a Glance

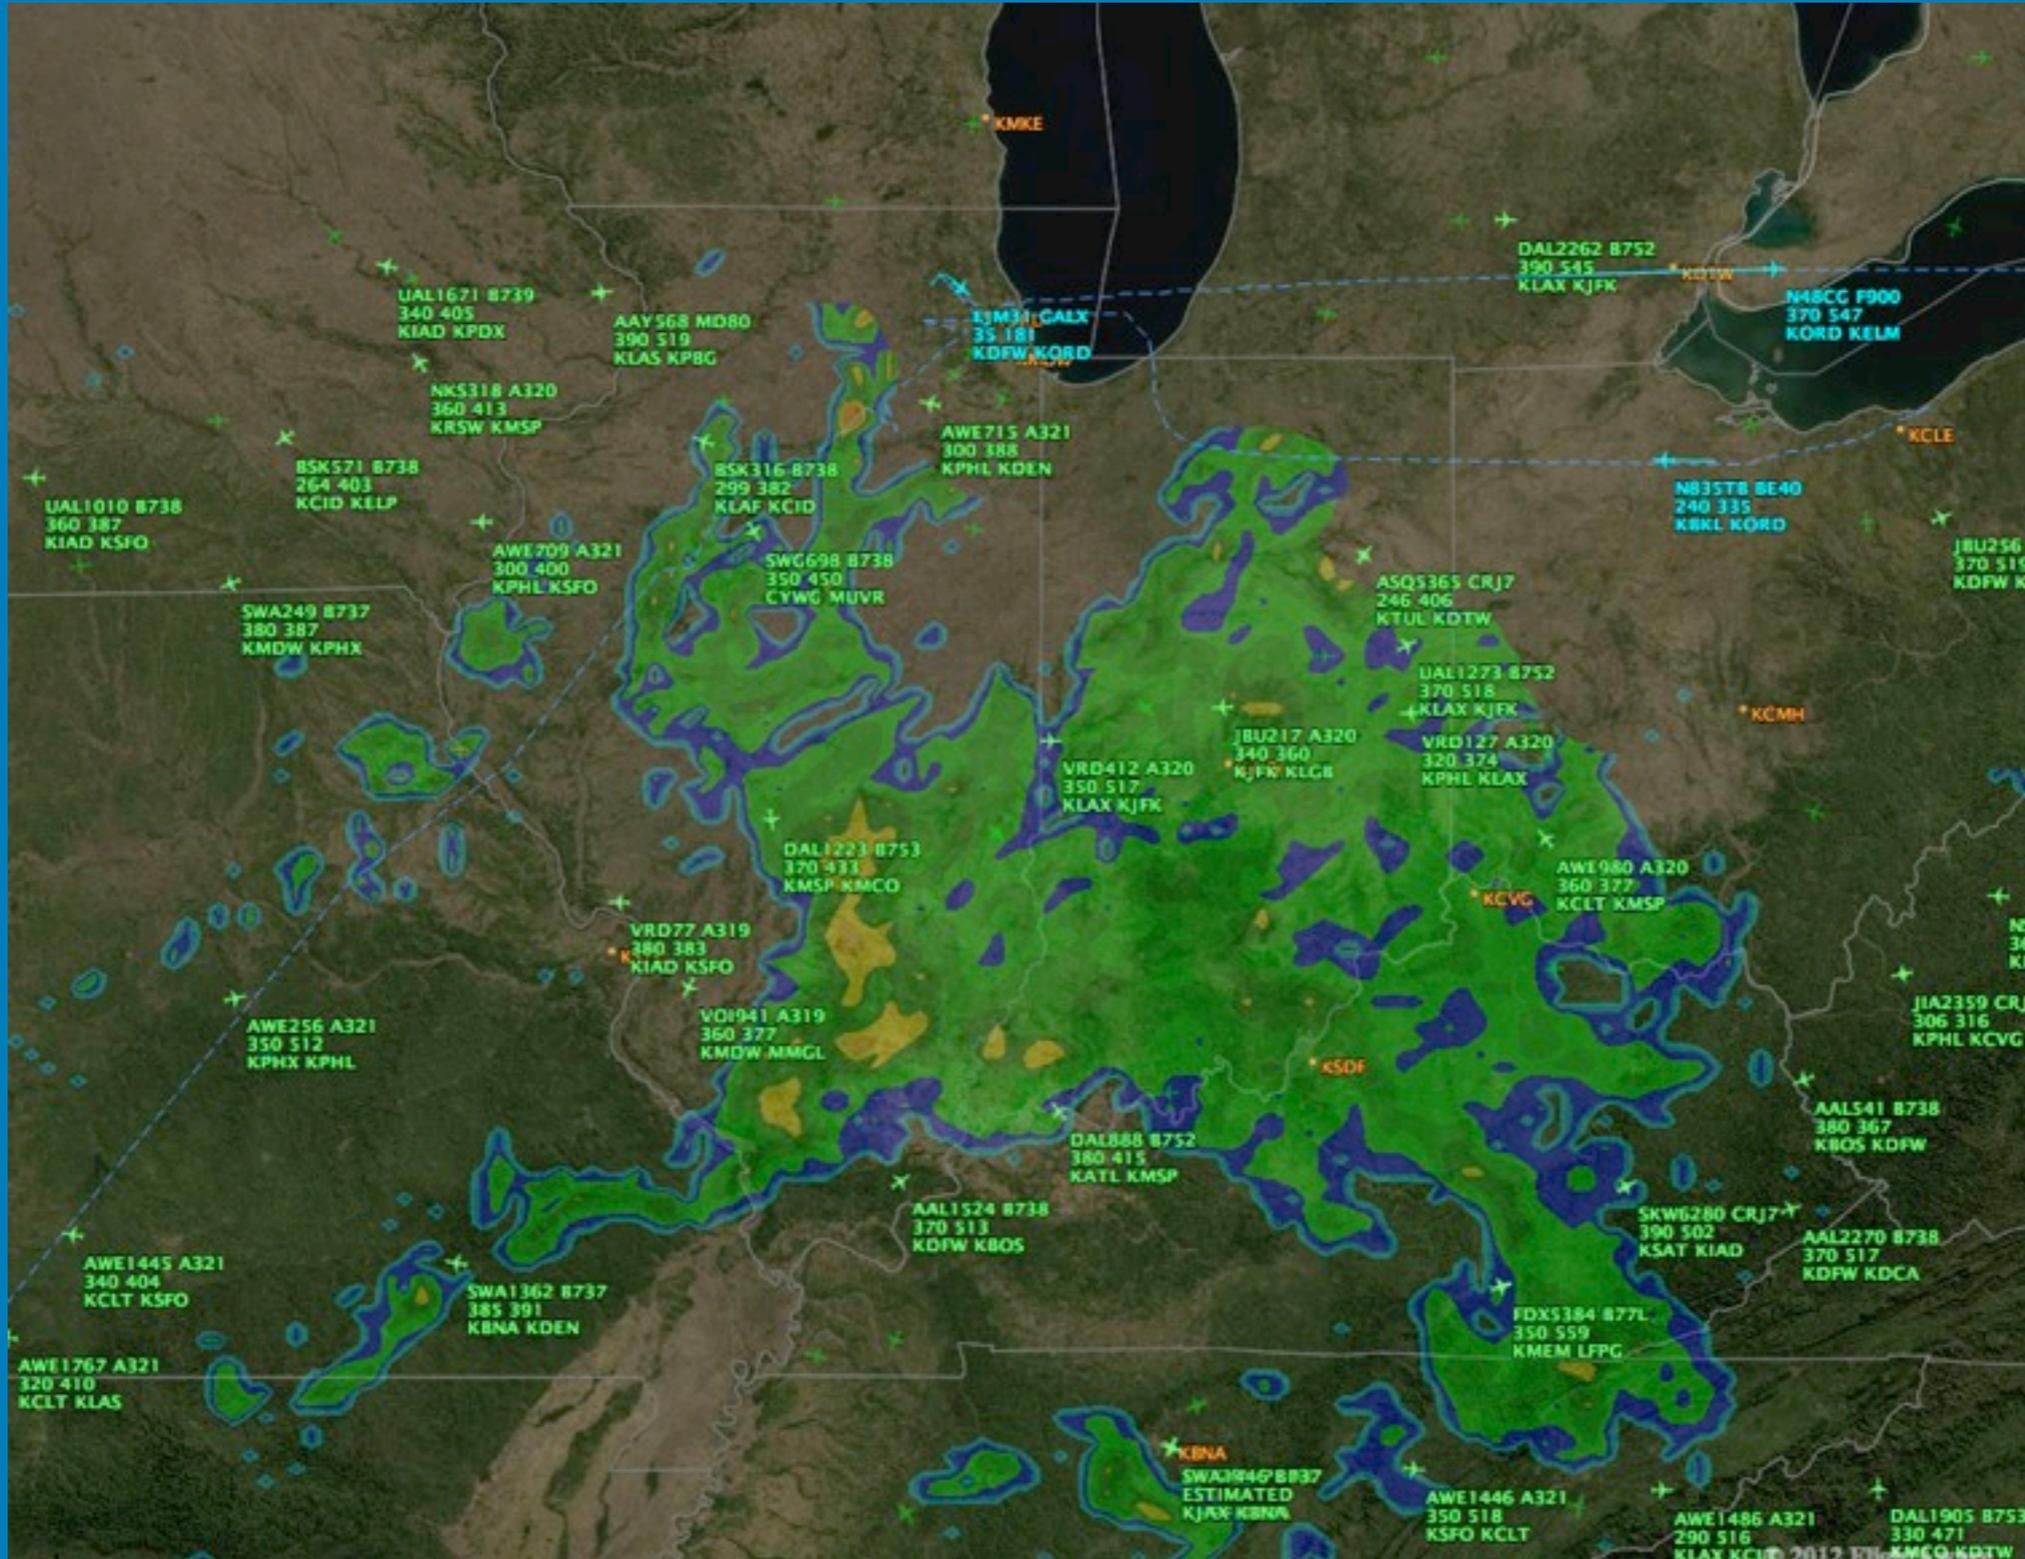
Provides real-time flight status to:

- Travelers
- Aircraft operators
- Airport operators
- Other tracking apps (TripAdvisor)
- Apps for every mobile app platform
- Localized in 15 languages

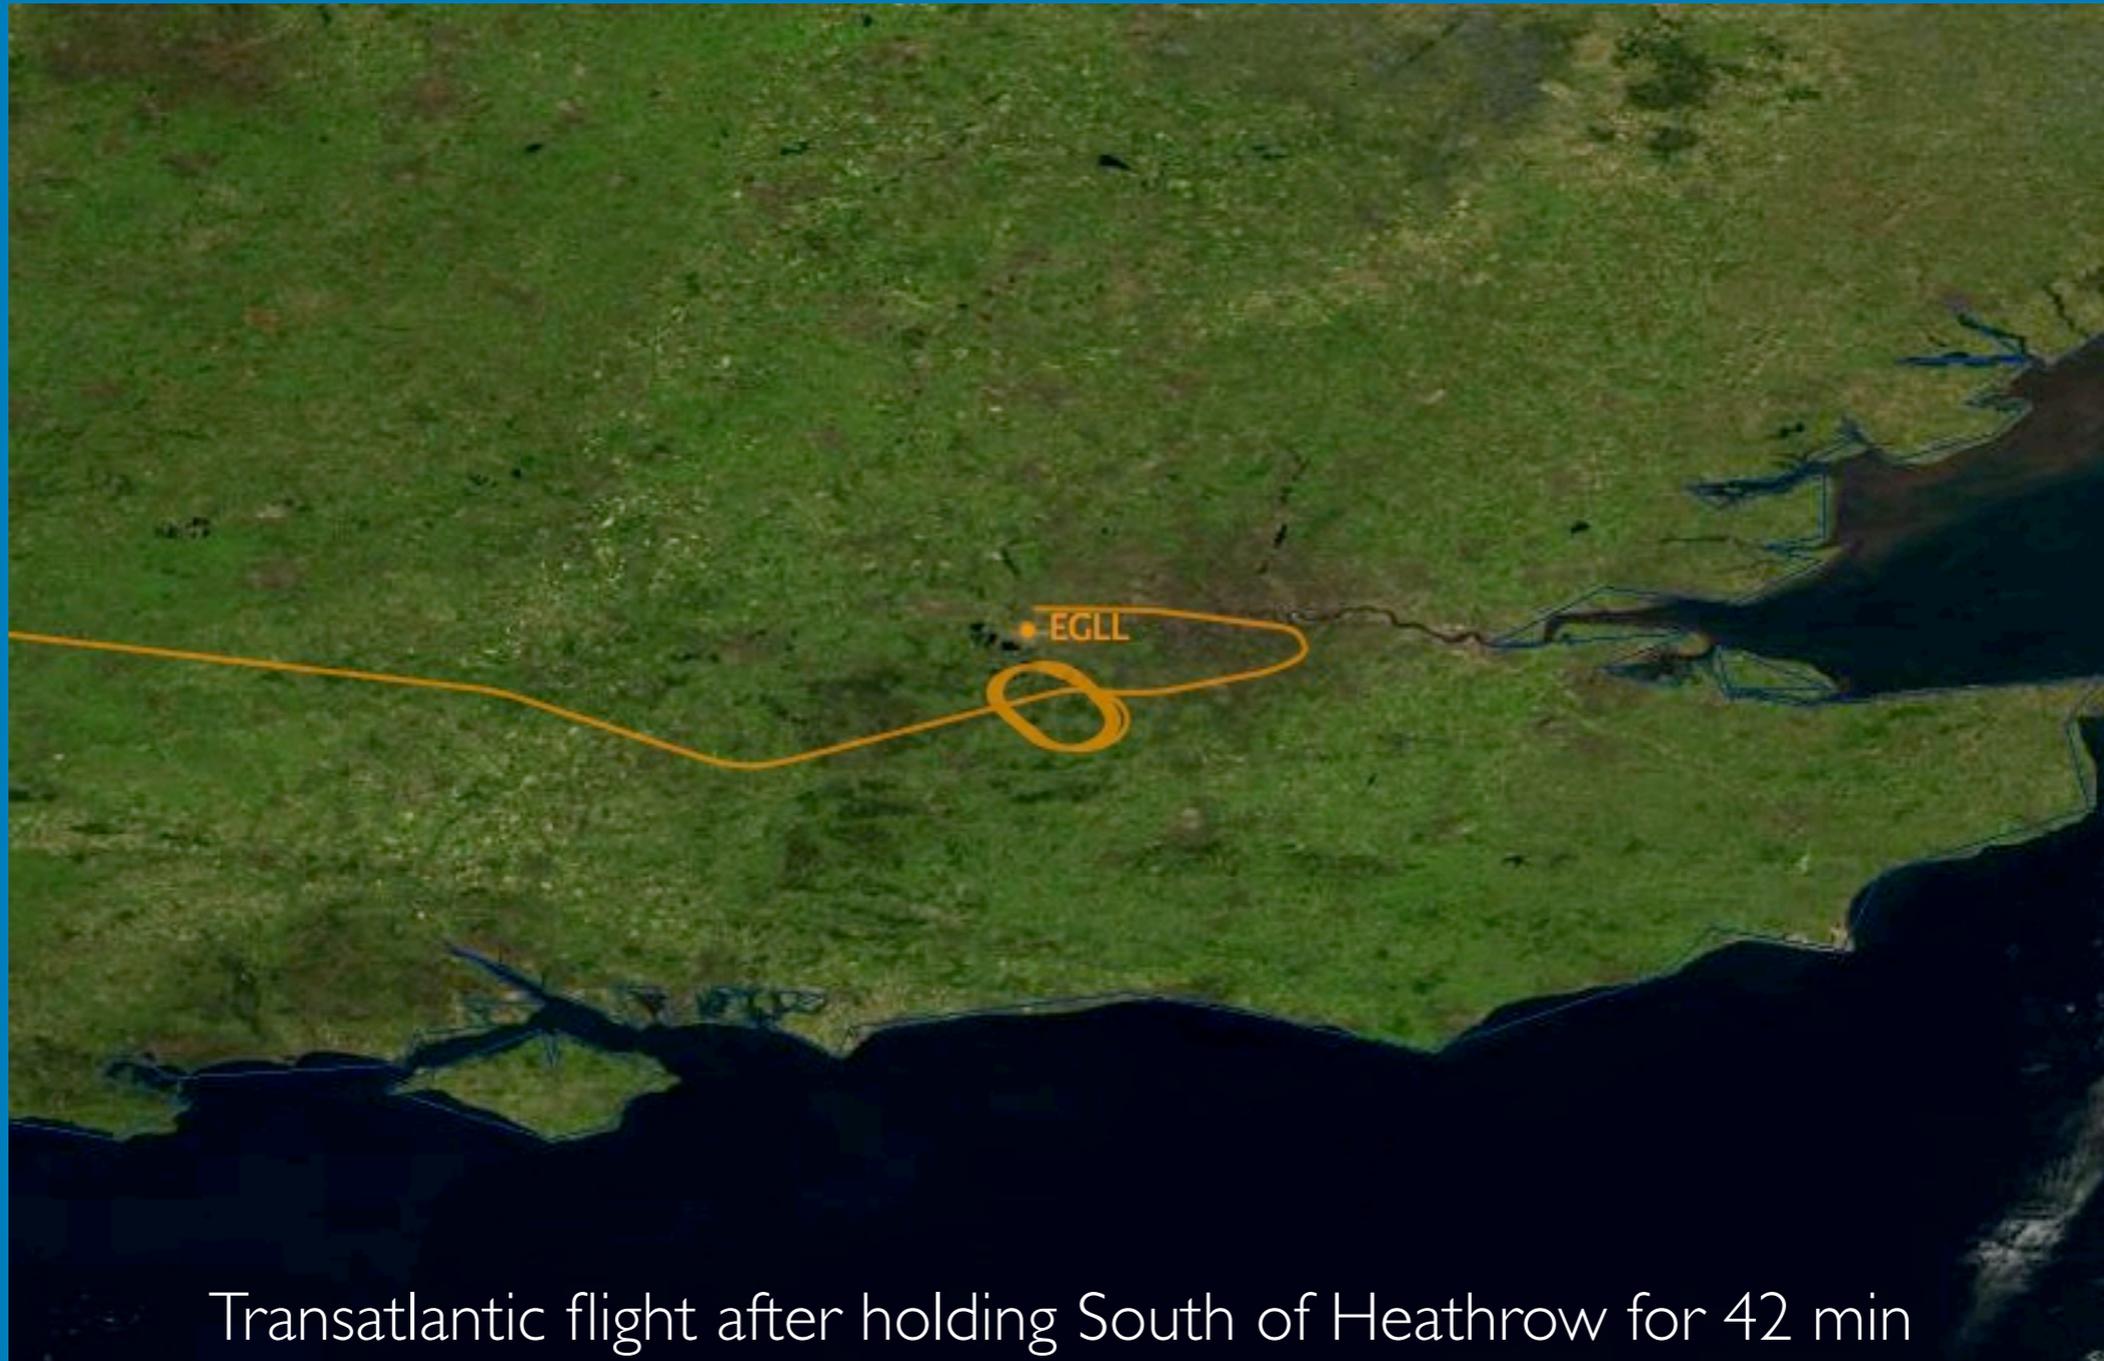
The screenshot displays the FlightAware website interface for tracking flight A6-DAB. The page is titled "Live Flight Tracking" and shows the flight's path on a map from London Heathrow (EGLL) to Dubai International (OMDB). Key flight details include the aircraft type (Boeing 747-200), speed (498 kts), and altitude (42,000 ft). An activity log table shows the flight's history, and a graph at the bottom tracks altitude and speed over time.

Date	Aircraft	Origin	Destination	Departed	Arrival	Duration
17-Oct-2012 *	B742	London Heathrow (EGLL / LHR)	Dubai Int'l (OMDB / DXB)	09:13AM	06:22PM	6:09
15-Oct-2012 *	B742	Dubai Int'l (OMDB / DXB)	London Heathrow (EGLL / LHR)	09:10AM	01:06PM	6:56
14-Oct-2012 *	B742	London Heathrow (EGLL / LHR)	Dubai Int'l (OMDB / DXB)	11:04AM	08:22PM	6:18
13-Oct-2012 *	B742	Dubai Int'l (OMDB / DXB)	London Heathrow (EGLL / LHR)	10:16AM	02:18PM	7:02

# About FlightAware - Airspace Views

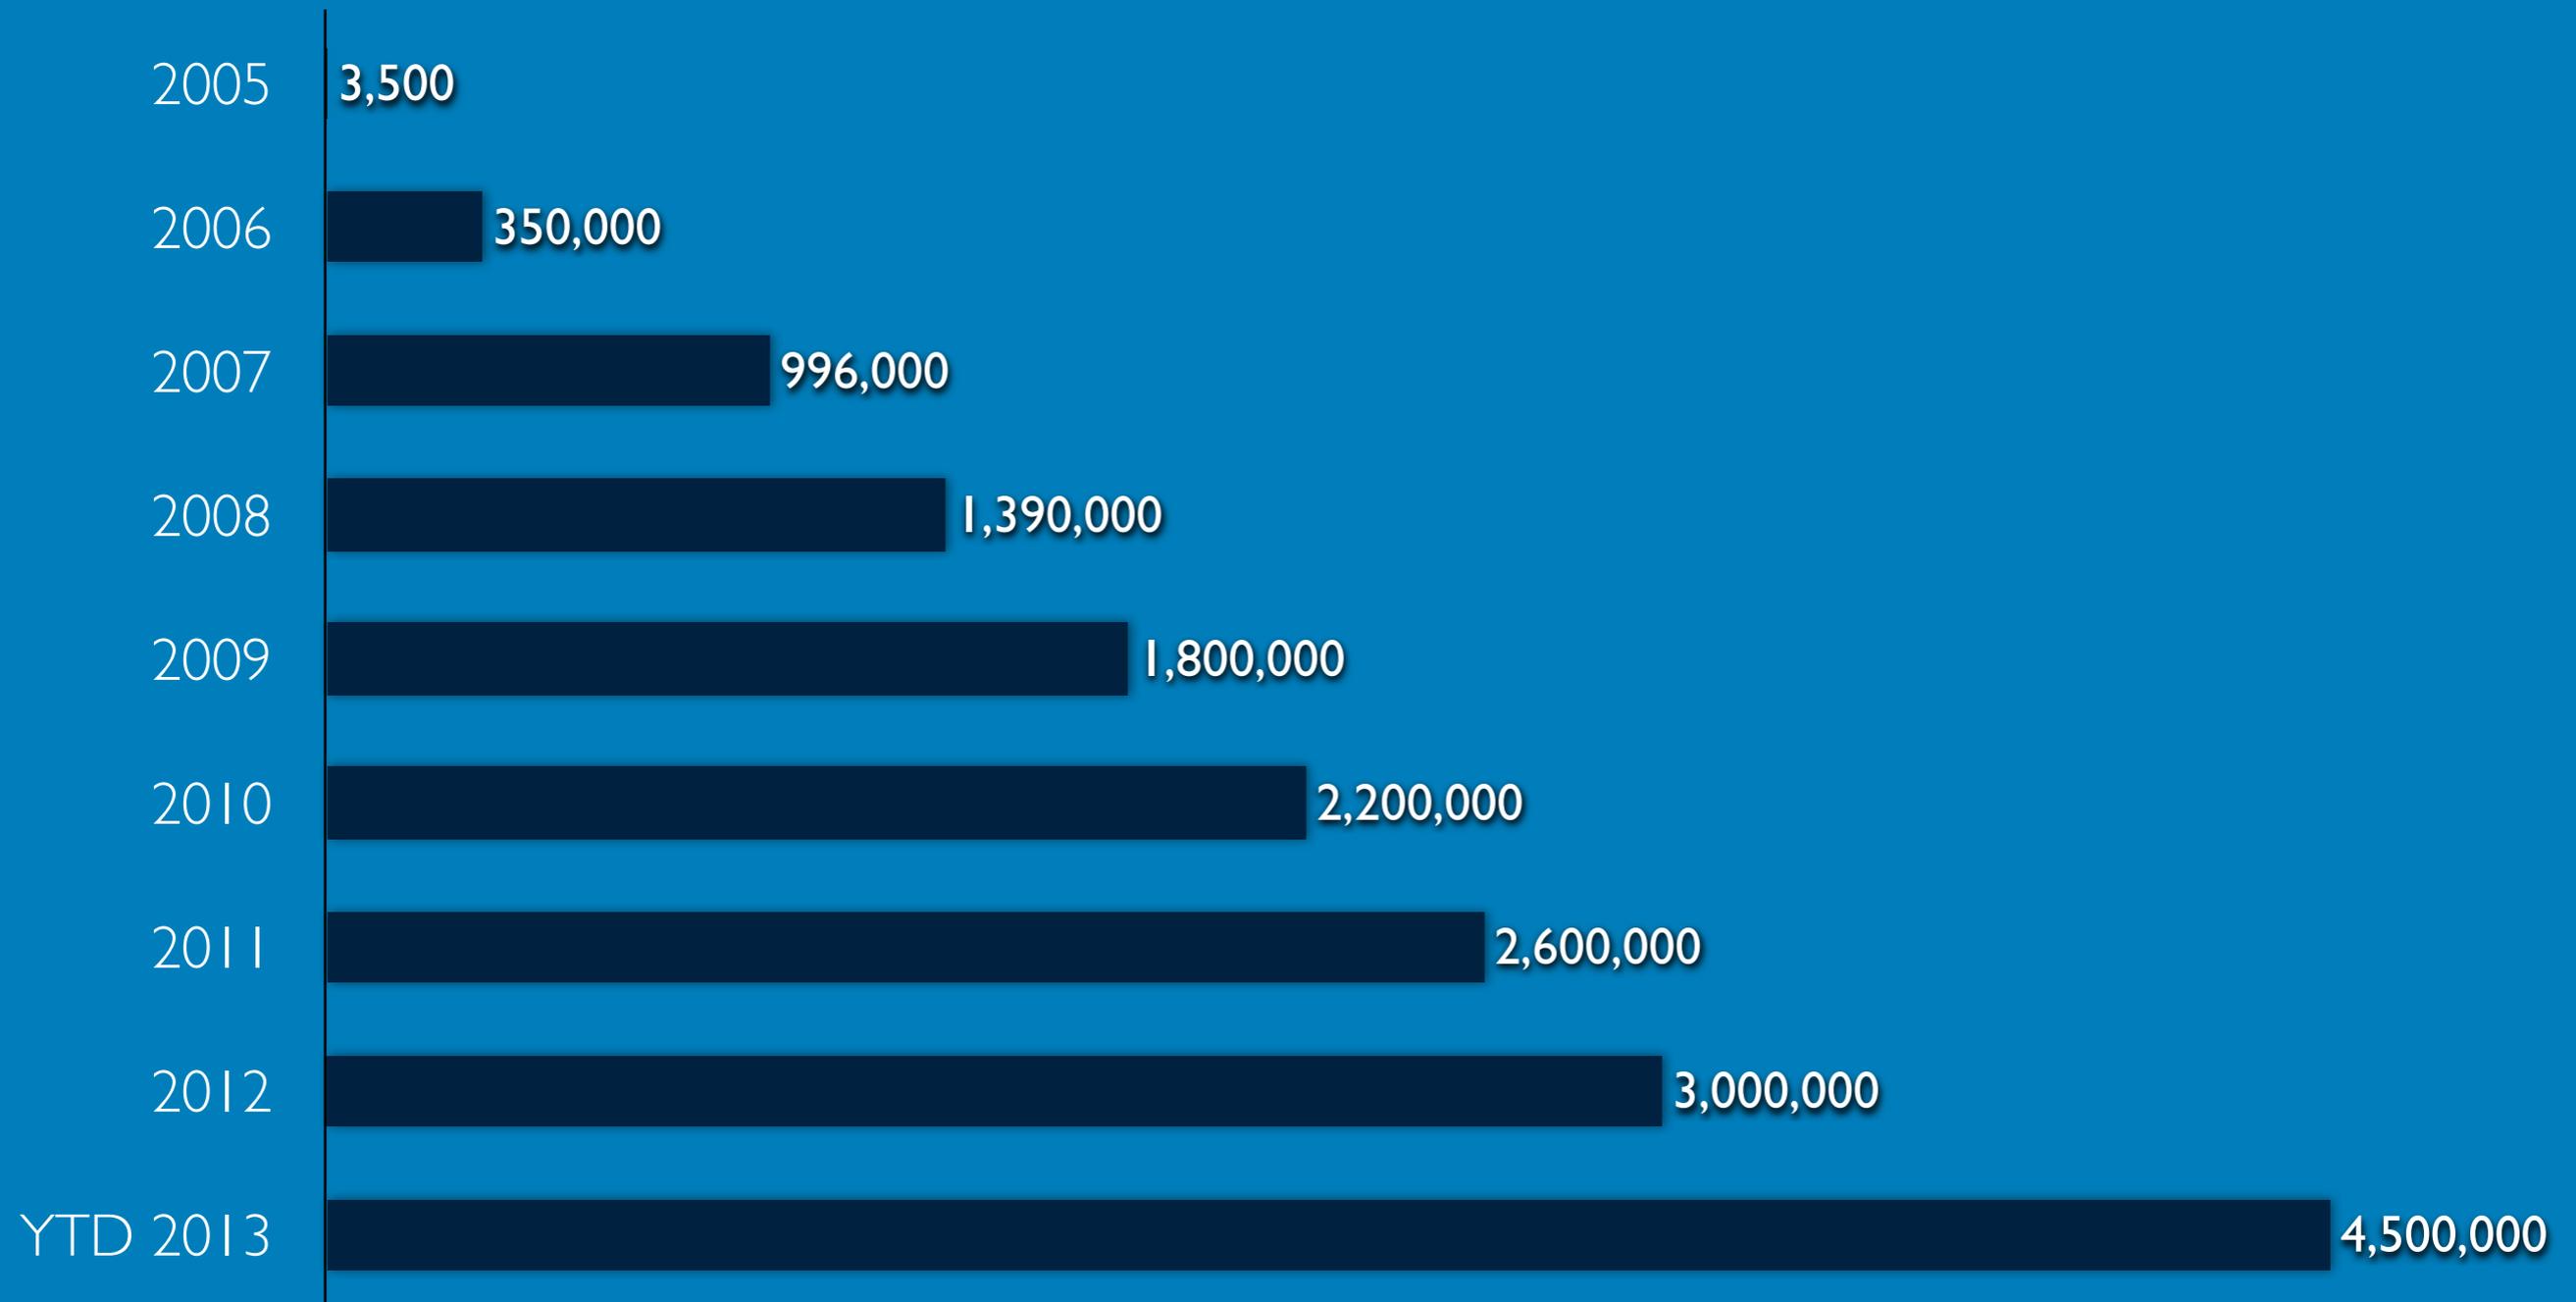


# High-Resolution Zoom Explains Delays



Transatlantic flight after holding South of Heathrow for 42 min

# Monthly Unique Web Site Visitors



2013 uniques up 50%  
YOY, beating projections.

# About FlightAware - Worldwide Airport Delays



London Heathrow (EGLL) is currently experiencing arrival delays for airborne aircraft an average of 55 minutes. ([all delays](#))

# FlightAware - Airspace Aware - Delays

Flight #

[TRACK FLIGHT](#)

[FORGOT THE FLIGHT NUMBER](#)

**AIRPORT TRACKER/INFO**

Airport Code

Airport City

[VIEW ACTIVITY](#)

[VIEW INFO](#)

<a href="#">KEFK</a>	Venango Rgnl	3	09:58:00	15:23:00	01:51:00	4	08:00:00	14:51:00	01:11:00
<a href="#">KJST</a>	Johnstown-Cambria Co	1	12:28:00	12:28:00	12:28:00	1	11:36:00	11:36:00	11:36:00
<a href="#">KLBE</a>	Arnold Palmer Rgnl	1	10:51:00	10:51:00	10:51:00	0	—	—	—
<a href="#">KLNS</a>	Lancaster	1	00:25:00	00:25:00	00:25:00	1	02:36:00	02:36:00	02:36:00
<a href="#">KMDT</a>	Harrisburg Intl	3	07:00:00	11:06:00	00:18:00	2	05:29:00	10:21:00	00:36:00
<a href="#">KPNE</a>	Northeast Philadelphia	0	—	—	—	1	00:51:00	00:51:00	00:51:00
<a href="#">KSEQ</a>	Penn Valley	0	—	—	—	1	14:36:00	14:36:00	14:36:00

**Number of Delayed Flights per Hour**

BCAT1    LIFR    IFR    MVFR    VFR    Future (unknown)

**KPHL (Philadelphia Intl)**

■ arrival delays = 0  
■ dept delays = 0

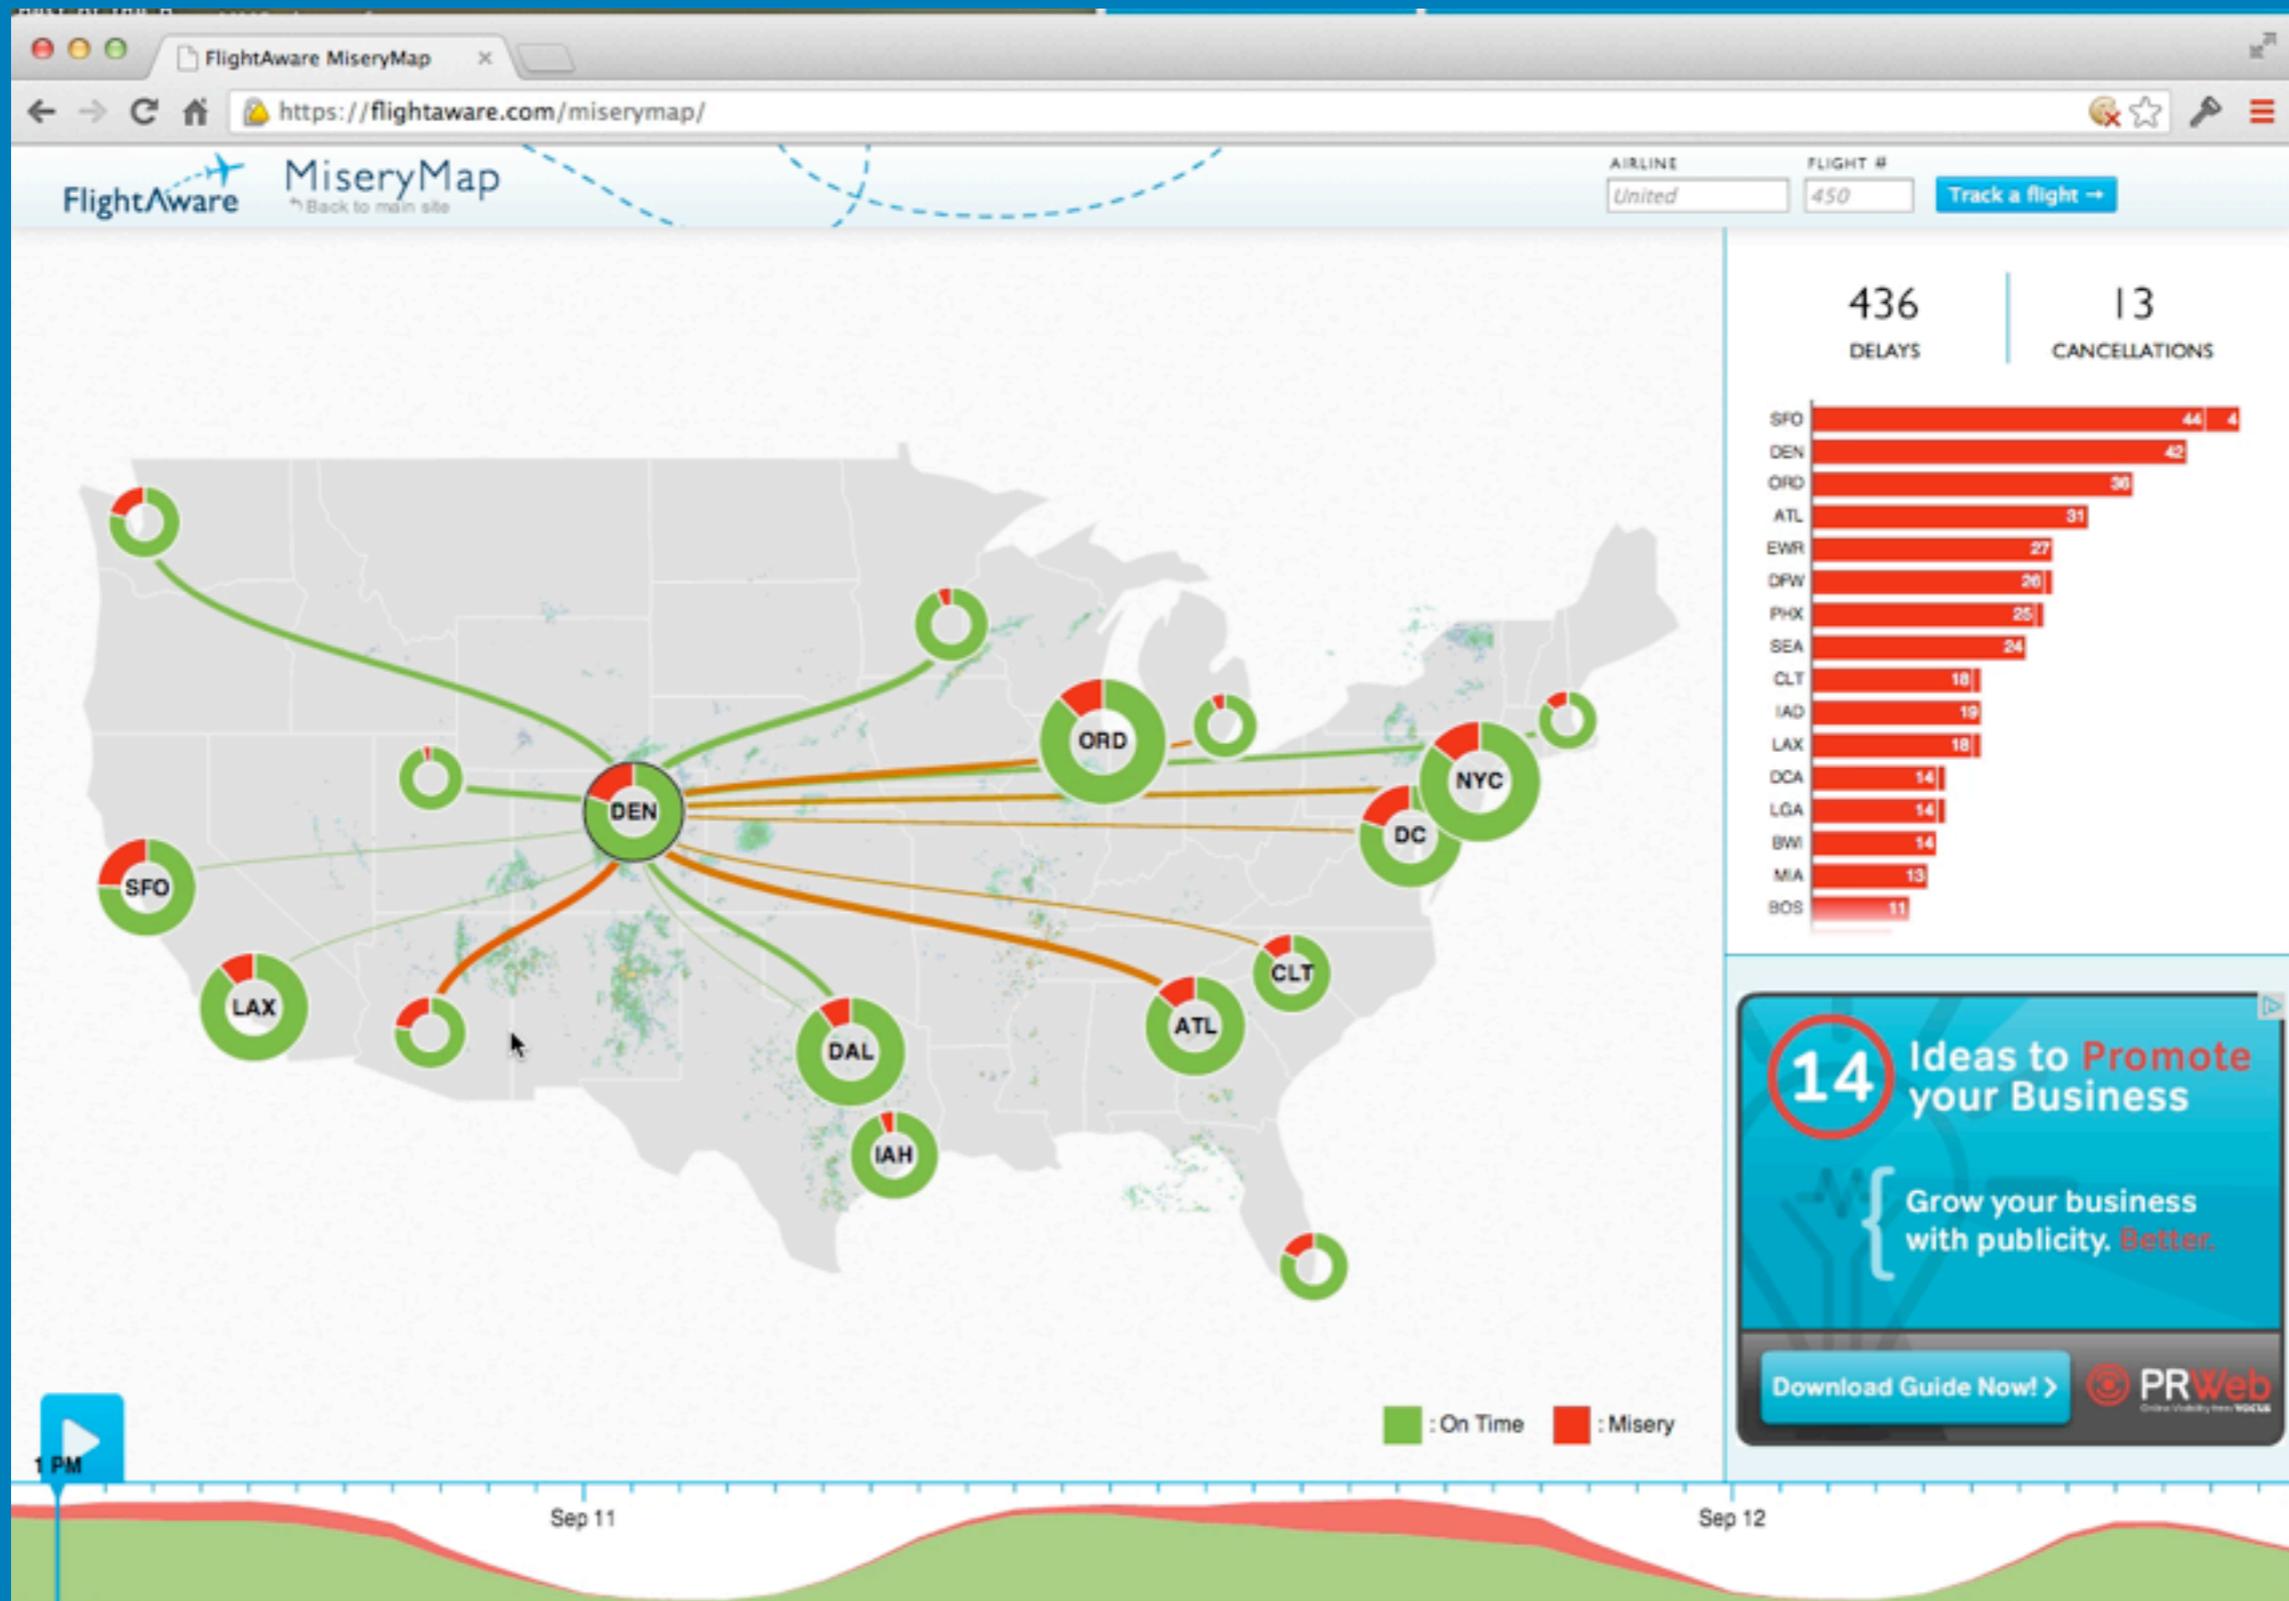
**KPIT (Pittsburgh Intl)**

■ arrival delays = 1  
■ dept delays = 0

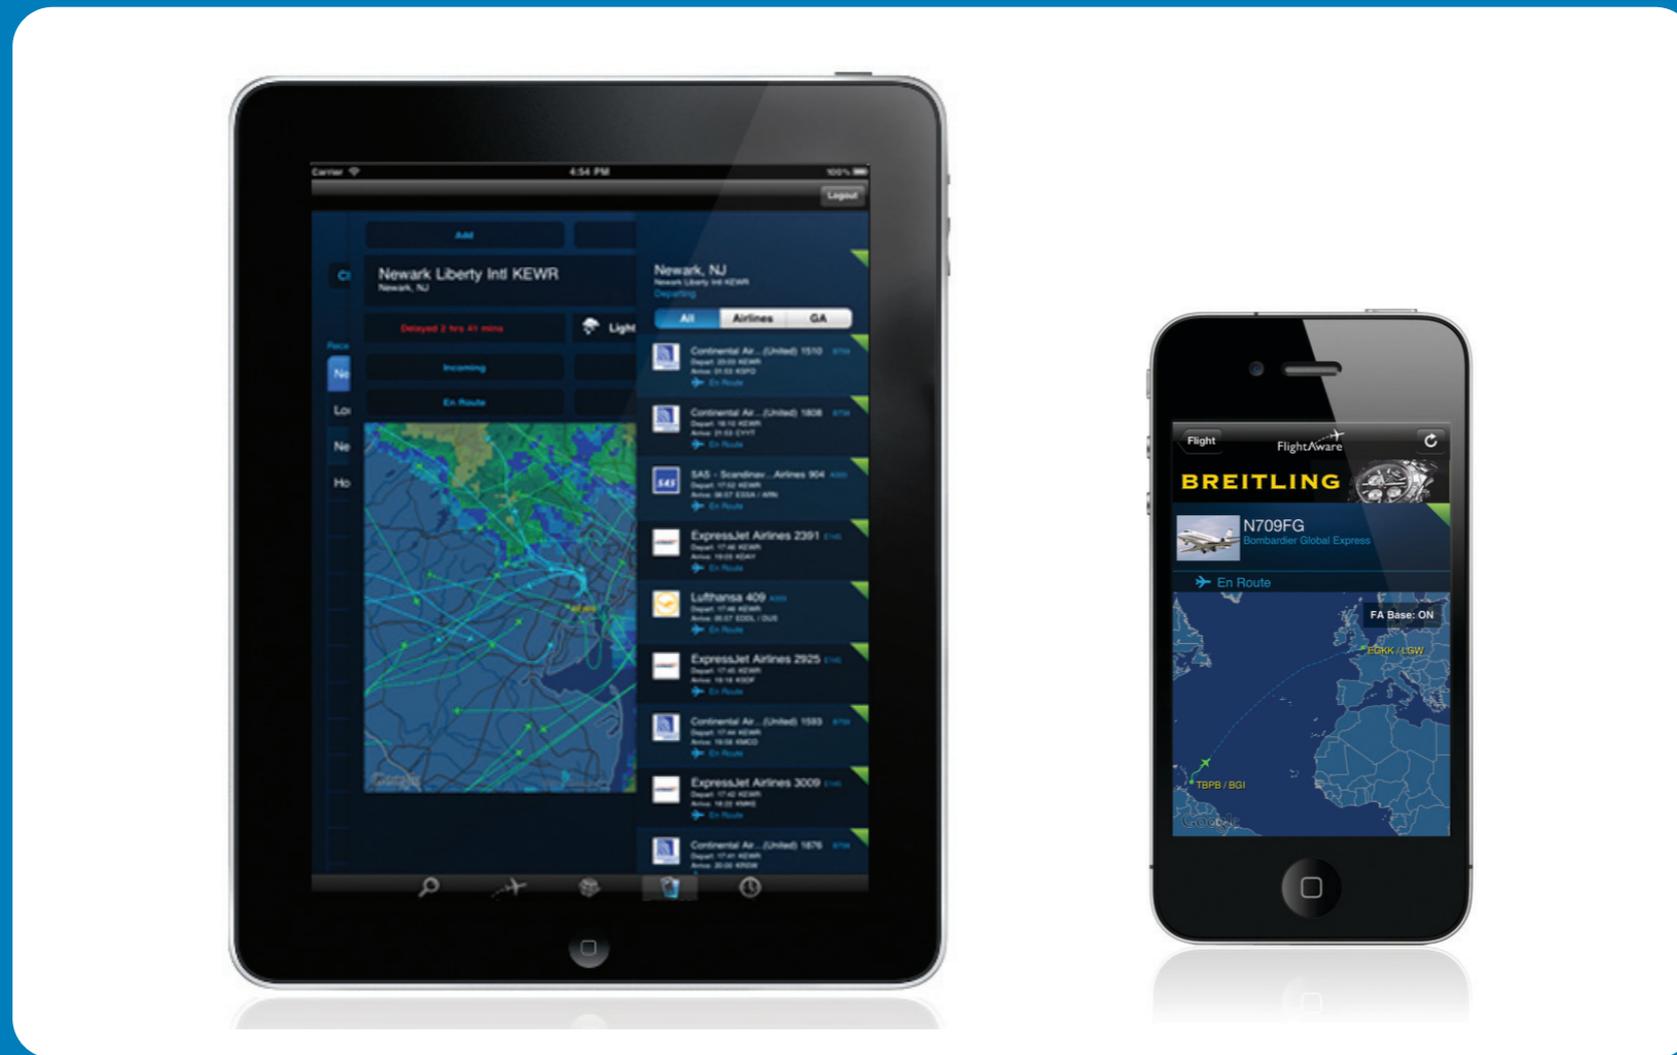
العربية | Deutsch | English (UK) | English (USA) | Español (España) | Français | العربية | Italiano | 日本語 | 한국어 | Nederlands | Português | Русский | Svenska | Türkçe | 中文(简体) | 中文(台灣)

© 2013 FlightAware | [Terms & Conditions](#) | [Privacy](#) | [FAQs](#)  
 2560 ms et | 19 ms cpu | 19 sql (11 b) (hide) (debug) (nolog+hide)

# Misery Map - Interactive Delay Explorer



# Mobile Apps



iPhone, iPad, Android, BlackBerry, PlayBook, Symbian, Windows Phone 7, Windows 8

# Next-Generation Flight Tracking - Datalink Coverage

**ARINC**

**ARINC Direct<sup>SM</sup>**

Global Data Center | **Honeywell**

spidertracks

satcom  
direct

 **UNIVERSAL**  
Weather & Aviation, Inc.

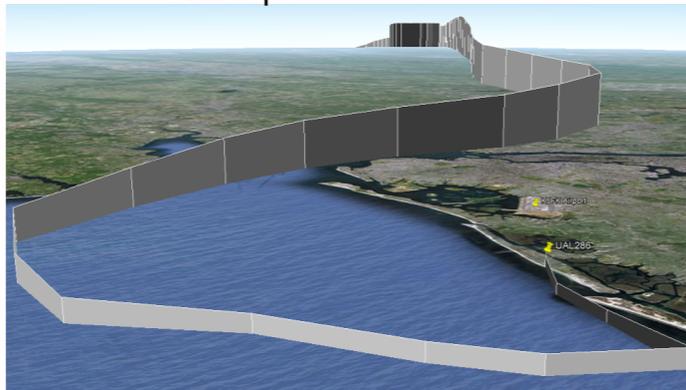
FlightAware is the datalink flight tracking provider for over 4,000 aircraft worldwide.

It took years to develop the credibility in the industry to be able to go after these relationships.

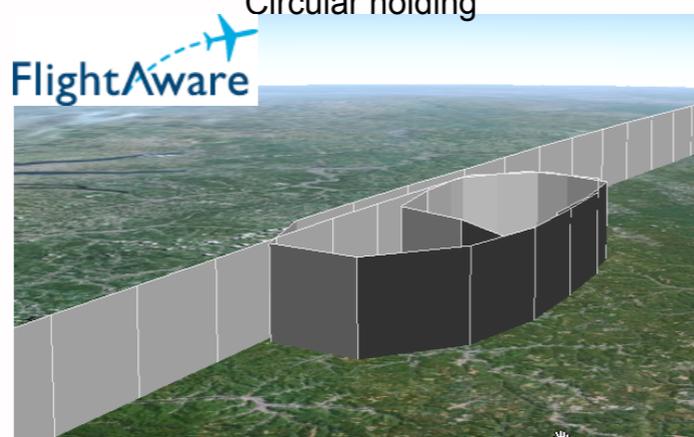
# FlightAware and United - Currently Flight Operations

## Data via external vendors

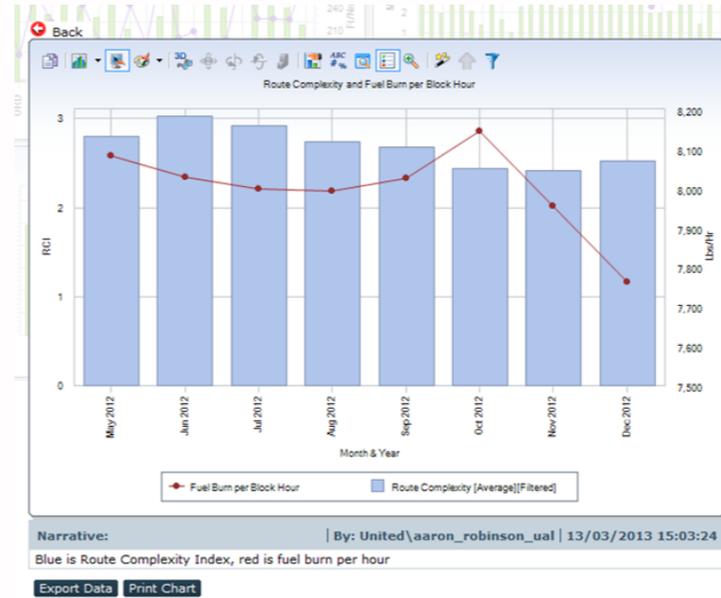
ATC procedure issues



Circular holding



Route Complexity Index



Expert **Insight** Delivered



Operational Analysis & Efficiency Automation



# Push Technology



Live, streaming flight push for all flights on [FlightAware.com](https://FlightAware.com)

# About FlightAware - Products

Operational	Analytics	Content
<p>Web-Based Flight Tracking</p> <p>Premium Accounts, FlightAware Global</p>	<p>Historical Flight Data</p> <p>Over 325M archived flights, billions of positions</p>	<p>Worldwide Airport, Navaid, and Oceanic Track Database</p>
<p>FlightXML - API</p> <p>Used by dispatch software, billing software, air carriers, airport authorities, etc.</p>	<p>Fuel Price Market Analysis</p> <p>Reporting, graphing, comparisons, and analysis</p>	<p>Worldwide Fuel Price Data Feed</p>
<p>Data Feeds</p> <p>Consolidated "FlightAware Controlstream" data feed that incorporates dozens of data feeds.</p>	<p>Traffic Flow Management Simulations</p>	<p>Weather Products</p> <p>International text and composite imagery</p>
<p>Flight Planning</p> <p>ICAO Flight Plan Filing, over 100 aircraft auto-setup.</p>	<p>Monthly Industry Aviation Report</p>	<p>Aviation News &amp; E-Mail Newsletter</p> <p>Monthly Circulation: 3,000,000</p>
<p>Two-Way Messaging</p> <p>Web Interface to ACARS messaging</p>	<p>Airspace Optimizations</p>	<p>Aircraft Photo Database</p> <p>375,000 photos 100,000 aircraft</p>
<p>Flight Alerts</p> <p>via email, text message (SMS), etc.</p>	<p>Commercial Data Services</p>	<p>FlightAware News Network</p>



# About FlightAware - Updates

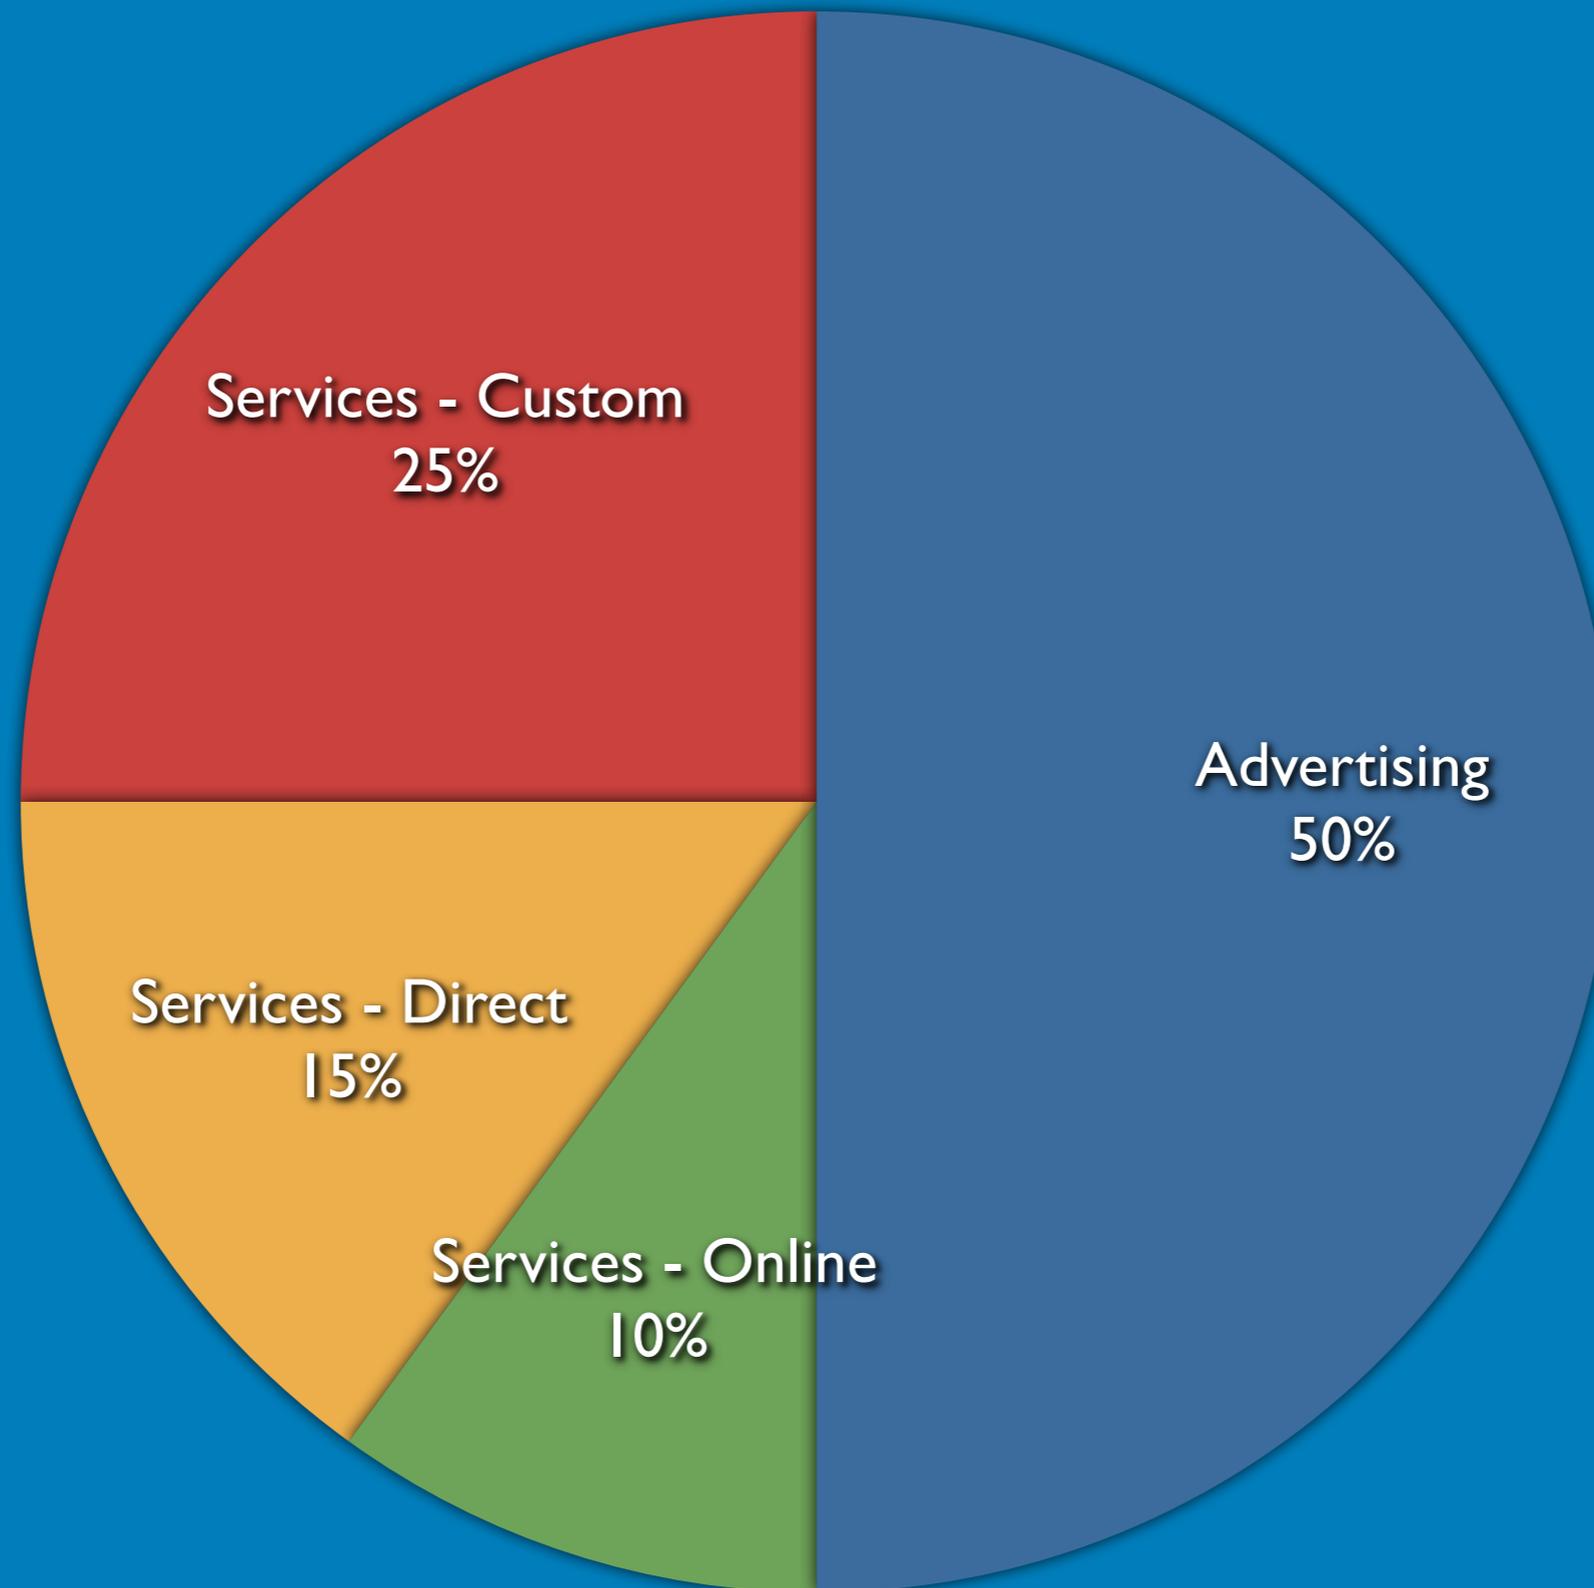
Coming Soon (Launching at NBAA in October 2013)

- Integration with ForeFlight electronic flight bag iPad App
- View any FlightAware route in ForeFlight
- Updated FlightAware routes pushed to iPad app.



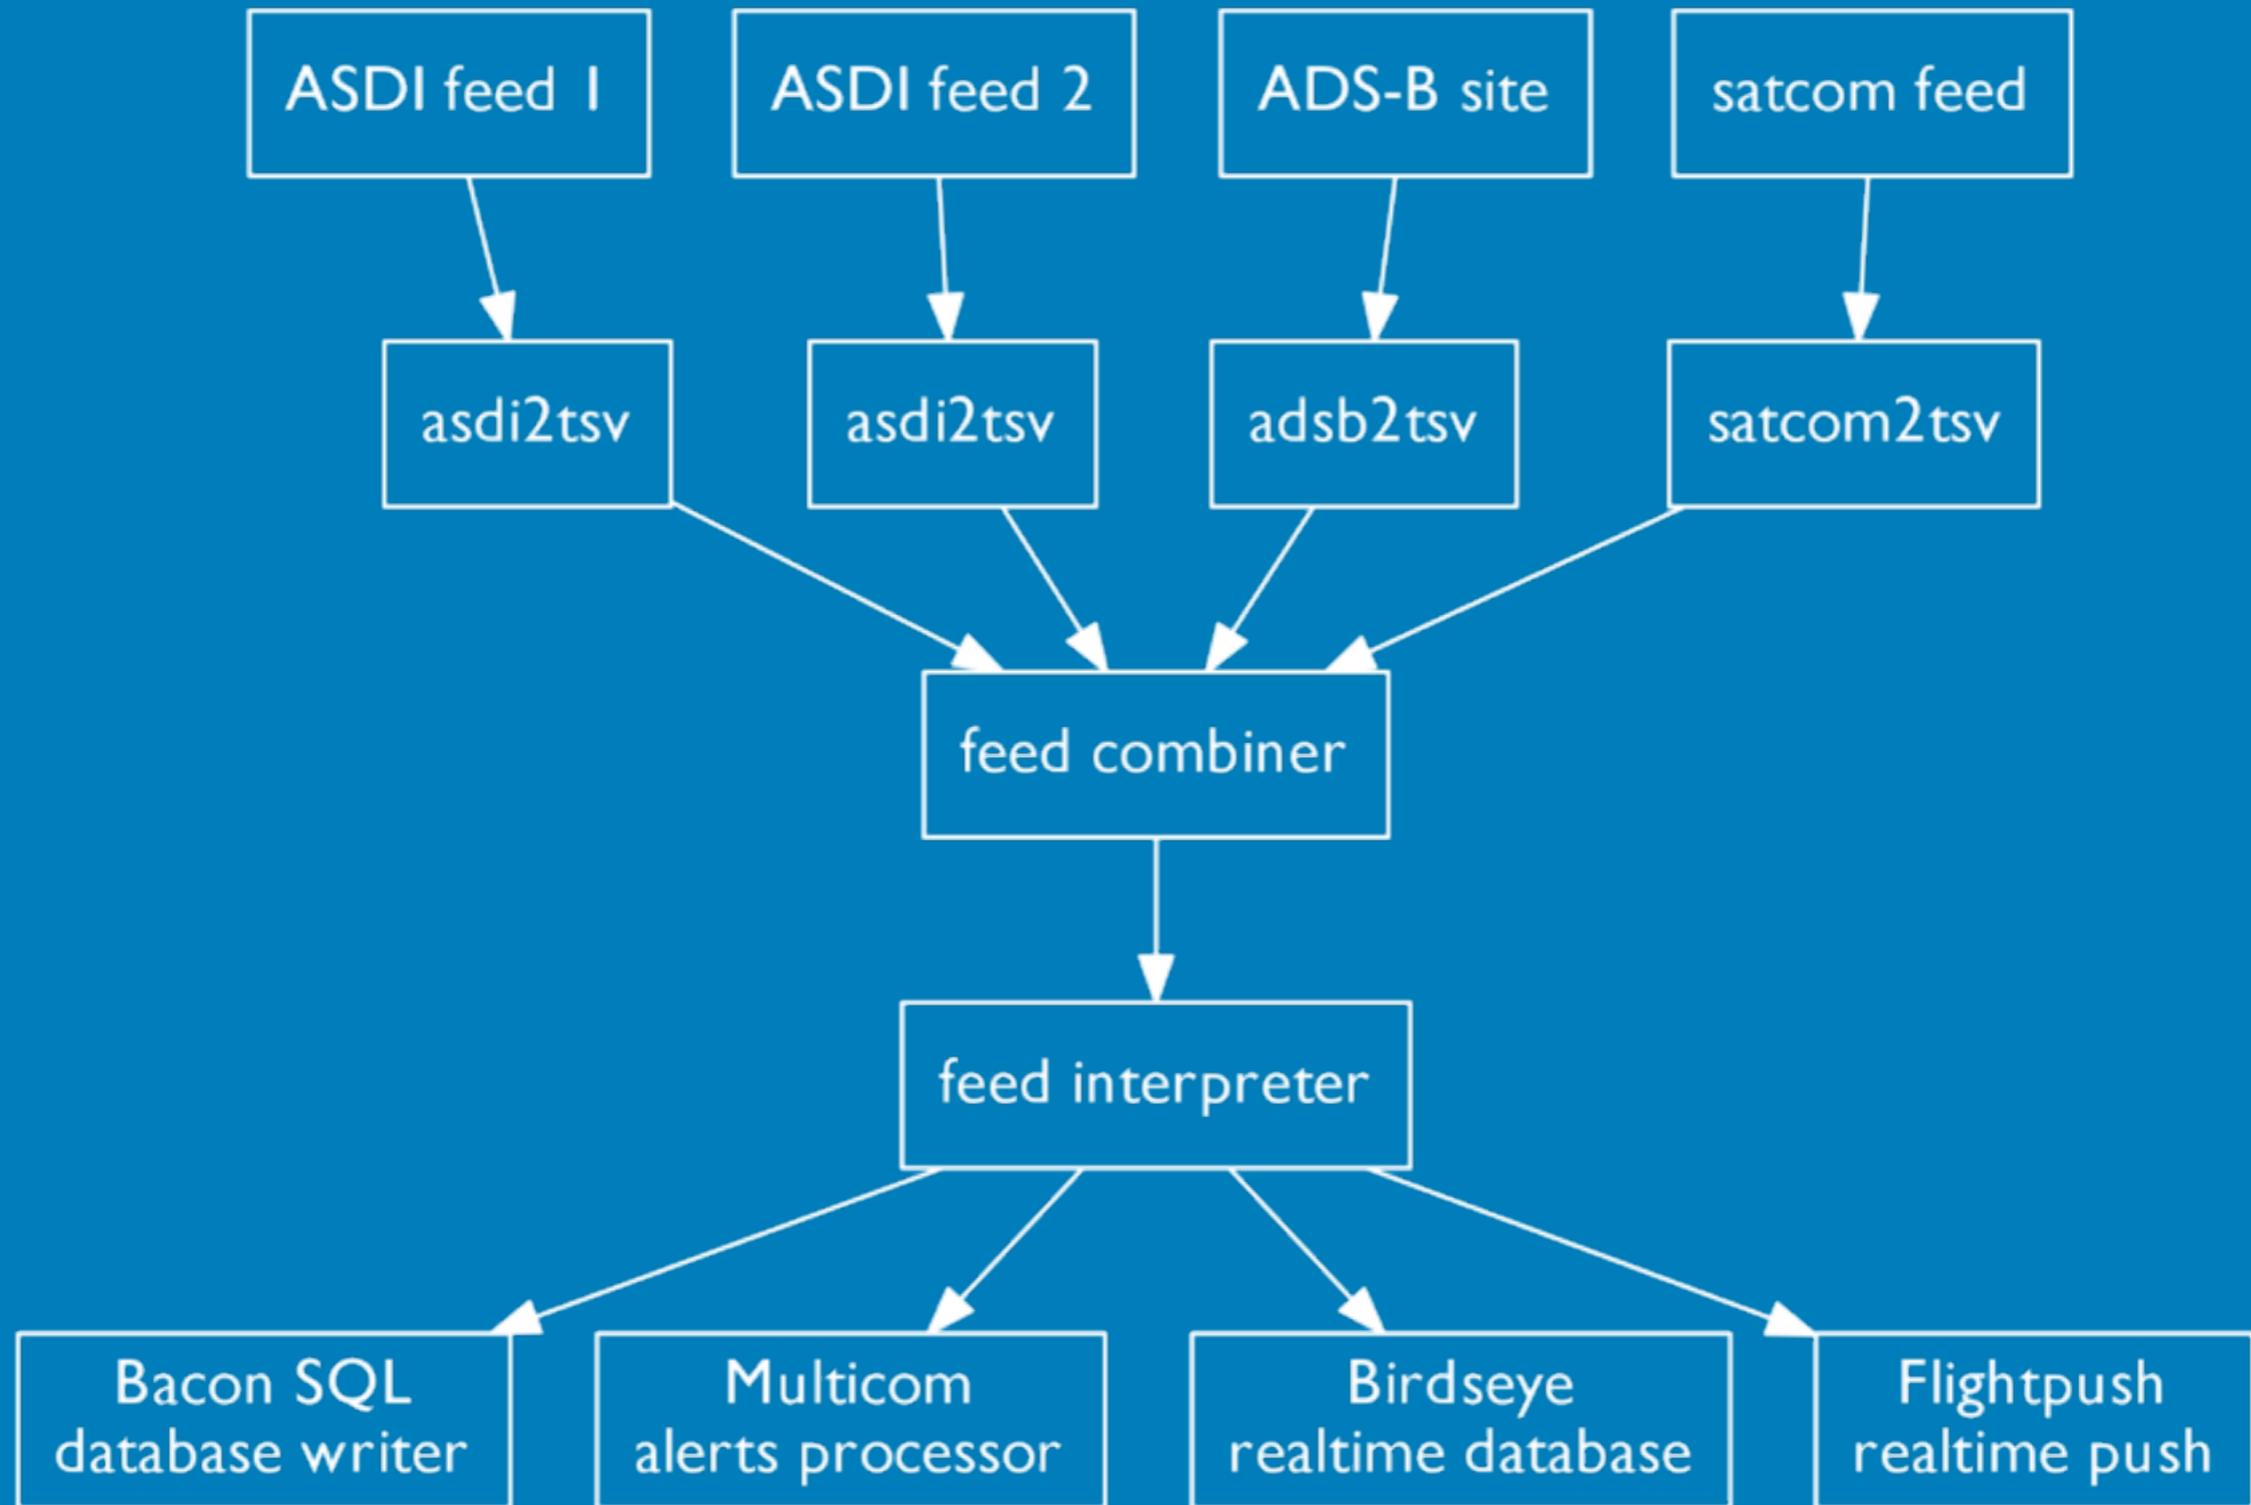
# FlightAware at a Glance - Behind the Scenes

Where the revenue comes from...



# FlightAware Technology

# FlightAware Technology - Flight Tracking



# FlightAware Technology - Behind the Scenes

- Feed Conditioners - recipients of any format data feed to be normalized for internal processing
- Daystream - temporal cloud technology for storing and retrieving terabytes of real-time or historical data
- Combiner - aggregator of over 30 real-time data feeds
- Feed Interpreter - stateful decision engine that analyzes “Combiner” and produces “Controlstream”
- Controlstream - normalized, internal data feed available to all internal programs (140 msgs/sec)
- Birdseye - custom 10GB shared memory real-time 4D position/flight plan database
- Archiver - stores positions from “Birdseye” and proxies real-time requests for over 250M historical flights
- MULTICOM - event alerting service between data events (e.g., departure) and delivery channels (e.g., email)
- Flightpush - web push server for streaming, real-time updates to user maps
- Balancer - redundant cloud manager of over two dozen map and weather servers (real time and historical)

# What the Feed Interpreter does...

- Collect flightplans from a variety of sources and assigns unique Flight IDs to them
  - ident (*UAL972*), orig (*EGLL*), dest (*KIAH*), registration (*N769UA*), EDT, ETA, route, waypoints, etc.
- Collect positions from a variety of sources (ANSP radar, ADS-B, ACARS via satellite and/or VHF)
- Collect OOOI (Out, Off, On, In) where available
- Match positions and other messages with flightplans
- Do extensive work to disambiguate the noisy signals
- Generate departures, arrivals, minutes-out messages, change messages (gate, bag claim, equipment, etc)
- Move aircraft along projected path when positions aren't coming in
- Update ETAs
- Maintain multiple "forks" of flights with different input sources to abide by data provider covenants
- Cash the checks

# FlightAware Technology - Error Correction

## Seventeen Simple Examples:

- Lost messages (e.g., missing departure, missing arrival, etc.)
- Out-of-order messages
- Inter-facility messaging / erroneous flight plan cancellations
- Non-standard ATC idents and aircraft types
- Mistyped ATC idents
- Ambiguous/duplicate airport codes
- Erroneous departures or arrival messages
- Conflicting positions reported by multiple radars
- Inconsistent data (two sources disagree)
- Duplicate flight plans
- Duplicate transoceanic positions or post-arrival positions
- False “airspace protection” positions
- Accidental or double diversions
- Wrong dates (+/- 24/48/72 hours)
- Invalid/incorrect destinations
- Airborne departures from a VOR/waypoint/fix (no origin airport)
- Flights without flightplans

# FlightAware Technology - Out of range RADAR

Positions from FAA ASDI



Positions after FlightAware processing



# FlightAware Technology - Projections



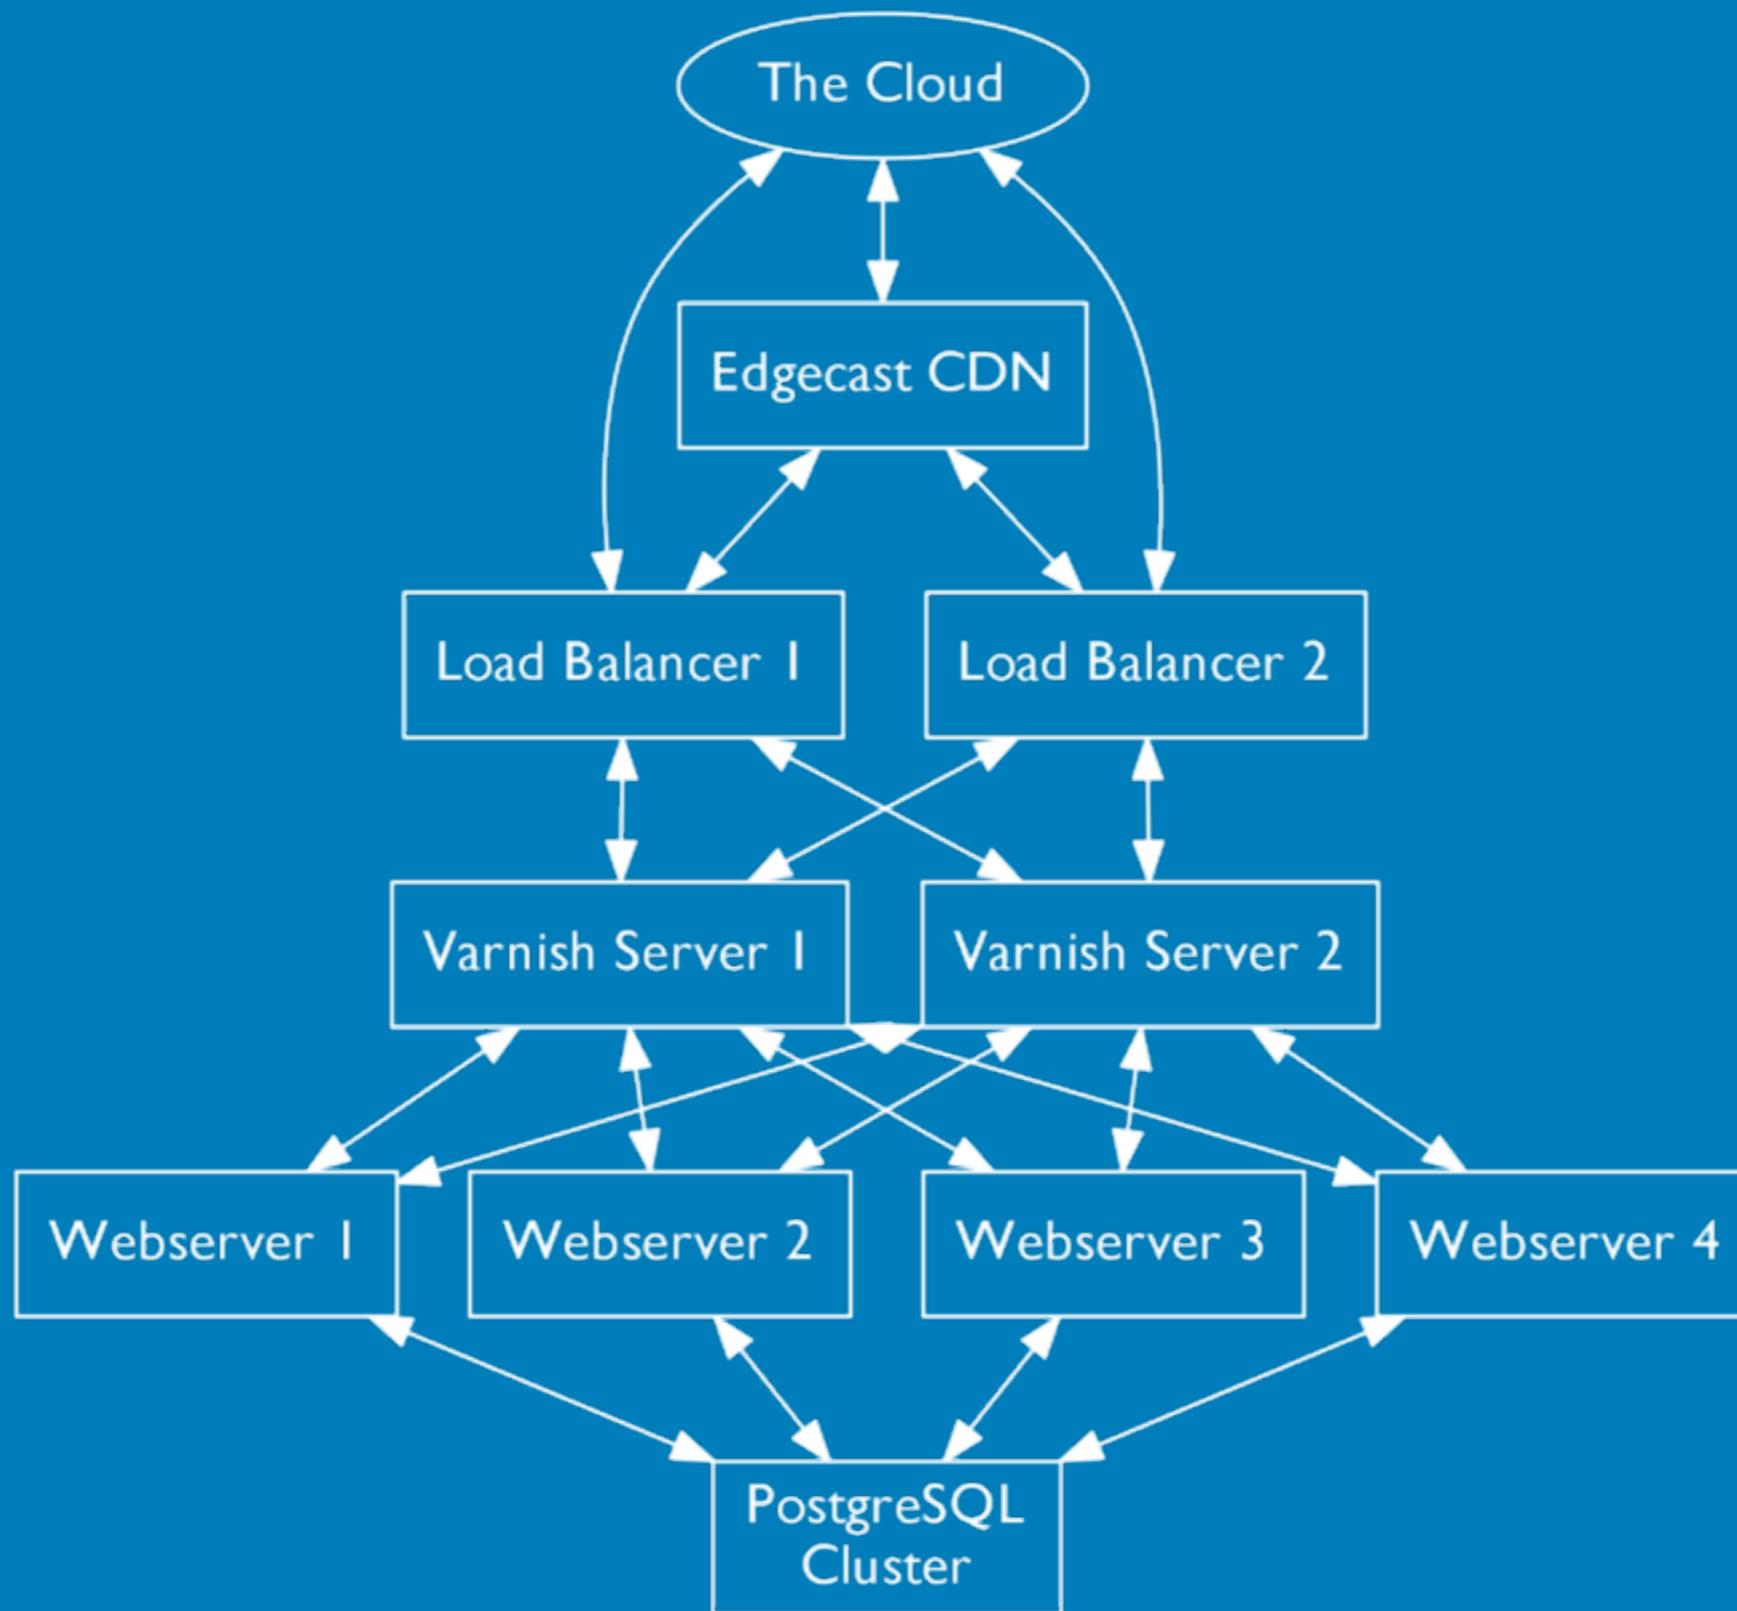
01:11AM	48.5333	-47.6833	76°	East	480	552	40,000		📍 Gander Center
01:12AM	48.5667	-47.4833	68°	East	480	552	40,000		📍 Gander Center
01:13AM	48.6167	-47.3000	71°	East	490	564	40,000		📍 Gander Center
(... time gap ...)									
01:52AM	50.4270	-39.9531	86°	East					📍 FlightAware Approximate
01:53AM	50.4804	-39.7648	88°	East					📍 FlightAware Approximate
01:54AM	50.5334	-39.5761	91°	East					📍 FlightAware Approximate
(... time gap ...)									
04:33AM	51.6070	-3.2658	101°	East	393	452	28,900		📍 FlightAware Europe
04:33AM	51.5958	-3.1738	102°	East	389	448	28,000	-1,620 ↓	📍 FlightAware Europe
04:34AM	51.5798	-3.0542	102°	East	379	436	26,900	-1,560 ↓	📍 FlightAware Europe
04:34AM	51.5688	-2.9718	102°	East	373	429	26,100	-1,620 ↓	📍 FlightAware Europe

Useful when no satellite/VDL positions are being received.

# Feed Interpreter - Some Details

- Feed interpreter comprises about 21,000 lines of Tcl code.
- It would have been basically impossible to write and have unacceptable performance without speedtables.
- If the feed interpreter isn't running, no flights depart, arrive or even move.
- Feed interpreter knows about airline flights from two days before the flight until a day after.
  - Consequently, feed interpreter is stateful.
  - All of its state is maintained in speedtables.
- FI's speedtables are checkpointed every 30 minutes to provide a restart point in the event of a crash.
- Feed interpreter forks a child to do the checkpoint so it doesn't stop processing.
- Checkpoints are rsynced to guard against a hard machine failure.
- Despite constant software upgrades, the feed interpreter's state has been maintained continuously for over five years.
- You can telnet into the feed interpreter from localhost and directly execute Tcl commands.
- All feed interpreter's procs reside in 70 source files that contain nothing other than proc definitions.
  - Consequently, updated procs without data structure changes can be sourced-in live.
    - *Yeah, we really do that, but we test first.*

# FlightAware Web-Serving Architecture



# Server Capacity

We own and operate almost all of our server capacity

- ~100 servers in six racks, mostly IUs and blades
- Three datacenters (Houston, Dallas, London)
- Originally whitebox; migrated to Dell due to ease of ordering/managing/repair/financing.
- Banks will finance server purchases.
- Electricity is more expensive than physical space or bandwidth.

“Haven’t you guys heard of AWS?”

*“Yes, frequently during their multi-hour outages.”*

# FlightAware from the inside

- We release FlightAware four times a week.
- Staff uses tomorrow's FlightAware today.
- On a typical day we're doing over 150 pages/sec most of the day.
- Every uncaught error is recorded and investigated.
- On a typical day about one page in a million suffers a partial failure.
- No amount of testing will find all the bugs that our 950,000 daily users will uncover.

# FlightAware and Tcl

- We use Apache with mod\_rivet to script in webpages
- We use Pgtcl to talk to PostgreSQL
- We write database stored procedures in PL/Tcl
- We use scotty for datagrams, ICMP, and SNMP.
- We used huddle to generate JSON until the performance became too big of an issue, then we wrote yajl-tcl, a C extension to talk to yajl
- We use Itcl, SOAP, http, ncgi, sgml, and sha1.
- We use ftp, htmlparse, ldap, and TLS, smtp, soundex and zlib.
- The Apache parent process loads 546 Tcl packages, which are inherited by the 300 Apache children spawned on each webserver.

# FlightAware from the inside

Source code to the Homepage

```
<?
flightaware_startpage 1 300

set ::pagedata(leftAd) "1 main_left"

if {[userIsMobile]} {
    flightaware_openlayers_emit delay
}

flightaware_header "" "" 0

flightaware_framework_mainpage

flightaware_footer
?>
```

# FlightAware Lessons Learned

- Do what you love and what you would do for free.
- Build something great; worry about business model later.
- Spend what you earn.
  - Started out with \$25K in 2005
  - You don't need VC to start a software/services company.
- You can build a legit / viable enterprise around Tcl.
- If you can predict the future with reasonable accuracy for something people care about, they will pay.
- Node.js is crashy and hard to code for.

# Major insights from forty years developing software

...that I don't hear people talking about very much

# Major insights from forty years developing software

## Build iteratively

- Spending two years writing a spec and five years implementing it is a recipe for failure
  - You don't really understand the application until you've implemented it
- But how can I implement it if I don't understand it?
- Get something working and start trying to use it
  - If you're building an alarms subsystem, start with a call to queue an alarm that just dumps the call parameters to the terminal
  - If you need to support alarm priorities, allow the priority to be specified but ignore it until you're ready to handle it
- Create flexible, extensible interfaces
  - key-value pairs, etc

# Major insights from forty years developing software

Text is King

Most applications, even math-intensive things like the feed interpreter, spend a lot more of their time manipulating text than performing calculations.

# Major insights from forty years developing software

## Build monitoring into your apps

- Relying on your customers to tell you when something isn't working is awful and if that's all you've got, you are asking for it.
- Monitoring based on machines being up and programs existing in the process table, etc, is inadequate.
- Build monitoring into your apps.
  - Send an "I'm OK" message to your monitoring software upon successfully processing an event, etc.
  - The absence of "I'm OK" messages tells you something is wrong, and will work, regardless of the cause.
  - Make it zero-config so programmers will use it.
  - *fa\_watchdog\_reset united\_delay\_extractor -interval 3600 -description "United delay extractor" -class feed -tags reports*
- We monitor 545 points (213 unique) in our software throw.

# Major insights from forty years developing software

## Solving super hard problems

If the only solution you have been able to come up with, you know it won't work, but after tons of thinking you have come up with no alternative, write the thing you know won't work and see what you learn from it.

# Major insights from forty years developing software

Take a long view

- If you are trying to build a software business, you had better be thinking about a twenty-year timeline for your platform.
- We have already been at FlightAware for eight years -- a ten-year timeline would be insufficient.

# Major insights from forty years developing software

Save all your data as received if you can afford to

- The data is complex; our understanding was incomplete.
- Had we not saved the raw data we would have lost fidelity.

# Major insights from forty years developing software

## Avoid Ground-Up Rewrites (like the plague)

- Almost never is it worth it to throw everything away and start over
  - Far less than the typical programmer's impulse
  - You don't actually understand the thing you're throwing away
    - You don't understand how complicated it is
    - You don't understand all the problems that it solves
- It only makes sense for rare, epic events where massive modernization is essential.
  - Getting off of proprietary hardware or a proprietary operating system (1980s)
  - Getting away from something horrible like systems programming in Fortran (1980s)
  - Moving from text files to a SQL database; moving to the Internet (1990s)
- History is littered with the corpses of solid companies that committed suicide by committing to a ground-up rewrite.

# Major insights from forty years developing software

Fashion: A popular trend, especially in styles of dress and ornament or manners of behavior

- It is the nature of fashion that it is both unpredictable and capricious.
- Fashion is a sort of epiphenomena of a community of thinkers, publicizers, followers, and fools.
- Fashion applies to programming languages, too.
- Languages go out of fashion through no fault of their own.
- Proponents of the new hotness always oversell the capabilities, strength, maturity, usefulness and applicability of their language.
- They are advocates, even if they pretend to be otherwise.

# Major insights from forty years developing software

Immature development environments are schedule-killers, or even business killers.

- Malfunctions at any layer below the layer you're working on can cause multi-day schedule hits, or worse.
  - Hardware
  - Operating System
  - Language Implementation
- It is extremely valuable to have the source code to the entire stack.

# Tcl - the Pluses

- Extremely stable
- Featureful
- Reliable
- Huge library of useful packages

# Tcl - stuff I'd like to see

- Higher performance: Javascript gets faster with each release;Tcl gets slower.
- Better capabilities to represent complex data hierarchies, like the DOM.
  - If you are using upvar to alias arrays from a bunch of arrays in a namespace, you are working around Tcl's lack of facility with complex data structures.
  - For me, dict doesn't cut it.
  - Speedtables really help, but they're not a first-class part ofTcl, so they move data in and out of arrays.
  - I am jealous of Javascript's ability to do things like

```
var projected_circle = new OpenLayers.Feature.Vector(new
OpenLayers.Geometry.Polygon.createRegularPolygon(layer.features[0].geometry,
radius, circle_segments));
```
- A general ability to invoke procs while explicitly naming the arguments as in Tk and iTcl.

# Can we expand the Tcl userbase and if so, how?

- The website is outdated, and it shows.
- The wiki is outdated and it super shows.
- There is no single place to go to find manpages, etc.
- Not hosting on github seems like we're willfully driving users away.

# Where Things Are Going

- Scripting languages will continue to be (very) important.
- I expect Javascript to be the dominant scripting language for the next ten years.
- Javascript has some pretty nasty warts but overall it is quite good.
- Node is holding Javascript back on the server side.
- There is no actual end in sight -- Tcl will remain viable as long as it has a community of active developers.
- Tcl New and Proven? Absolutely!