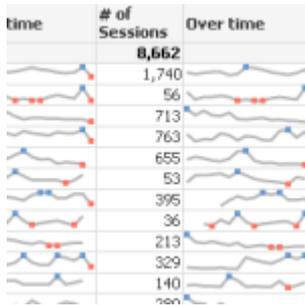


Visualizing End User Adoption

written by DaltonRuer | April 3, 2015



This past week I had my annual review. This time of year always makes me envious of those that produce widgets. I would love to be able to show my boss a list of all of the widget producers and say “See boss I’m in the top 2% of all of the widget producers in the company and the top 5% of widget producers around the world please compensate me accordingly.”

Since you are reading this post the odds are high that like me you produce Business Intelligence Applications and aren’t producing widgets either. So how do we evaluate our work? How should management evaluate us?

One way to evaluate our work might be to simply count the number of applications that we build. Of course I could barely contain a laugh just writing that. Obviously that is wrought with problems so let’s not even consider this option.

In a strictly financial sense many types of business can measure the return on investment (ROI.) But perhaps the application we spent 9 months building is intended to help resolve bottlenecks in the company that will lead to improved patient satisfaction. The resolutions that surface may cost the company more money. Does that mean we failed? Certainly not. So we can’t measure ourselves by dollars spent and dollars saved either.

If you follow industry pundits, tweets and other social media you might be familiar with the focus of many in the industry to focus on “user adoption.” Evaluating to what degree users

actually utilize our applications is probably a good way to measure ourselves. It could be argued that it isn't a perfect measure of our efforts, however, it does seem to be a pretty good measure of our effectiveness. Because whether we like it or not, our jobs involve more than just slapping an application together. End user adoption, or the lack thereof, measures our ability to brand, market and support our application. It is also a pretty good representation of how trustworthy the data in our application is. One of the most important things that end user adoption will measure is our ability to effectively visualize the data in ways that encourage usage.

[Tweet "End user adoption measures our ability to brand, market and support our BI applications"]

Taking advantage of Qlikview logfiles

One of the nice features of Qlikview is that it retains a log file in the background on the server that retains information about every single end user session that is invoked. Since the introduction for this post was so long I will spare you the pain of reading the raw data of a session log file and skip right to ways to effectively visualize end user adoption using the data that those logfiles contain. Please refer to other posts and discussions directly in the Qlik Community about where to find and how to access these log files.

The session log files contain information that would let us look at things like "how many users used the application" "how many times were sessions invoked" and "how many minutes were used." Thus the first chart I present contains all 3 of those measures.

Document	# of Users	# of Sessions	# of Sessions	# of Minutes	# of Minutes
	265	8,662		181,493	
Application 69	91	1,740		29,539	
Application 85	39	56		445	
Application 49	30	713		16,468	
Application 56	30	763		11,369	
Application 33	29	655		17,701	
Application 58	27	53		977	
Application 28	26	395		6,788	
Application 16	25	36		187	
Application 77	23	329		4,800	
Application 62	23	213		3,550	
Application 39	21	140		2,316	
Application 42	20	200		0	

The first point I want to make is that I've masked the real document names. I did this for two reasons. First you don't need to know what my real document names are. The more important reason is that I don't want know what the real document names are. At least for the duration of the time I'm trying to figure out how to effectively measure "end user adoption." That seems rather odd so let me explain.

Overcoming bias when choosing how to measure

I believe that we all have biases. I haven't developed all of my companies applications and frankly I have some favorites of those that I have developed and some that I was forced against my will to develop. If I knew what the application names were I could be inclined to choose and recommend the metrics that make "my" applications look the best.

If you refer back to the chart you will see that Application 69 has the greatest number of users by a large margin. If I knew that Application 69 was written by me I could immediately come to the conclusion that our end user adoption should be based on the number of users that use the application. If I also wrote Application 85 I would probably really push for that policy. "Show me the money."

But wait someone else on my team seems to have an objection because it appears that Application 85 has a lot of distinct

users but only has a tiny amount of Sessions and very tiny amount of minutes. Hard for me to argue with that, and I put my outstretched hand back into my pocket.

A discussion ensues for several minutes and perhaps we re-sort the chart by number of sessions. Then by Number of Minutes.

Document	# of Users	# of Users	# of Sessions	# of Sessions	# of Minutes	# of Minutes
	265		8,662		181,493	
Application 69	91		1,740		29,539	
Application 33	29		655		17,701	
Application 49	30		713		16,468	
Application 56	30		763		11,369	
Application 29	18		208		8,318	
Application 43	20		289		8,080	
Application 7	10		205		7,848	
Application 20	19		369		7,533	
Application 28	26		395		6,788	
Application 25	6		182		6,206	
Application 10	16		230		5,594	
Application 30	10		222		5,270	

The author of Applications 33, 49 and 56 now suggests that we evaluate end user adoption by the number of minutes used. I'd like to vote for that since I was the author of Application 69 but I also authored applications at the bottom of the chart for number of minutes. I'm kind of in a no win situation on this.

Can you understand my point for masking the document names so that we don't really know which application was developed by whom? If we are choosing a method of evaluation we need to hide the real document names so that nobody pushes for a choice just because it is better for them.

Perhaps of equal importance can you appreciate the beauty of having all 3 columns displayed with numbers as well as bar charts? Obvious patterns jump off the page that help you avoid jumping to quick conclusions just based on 1 value or the other. If we are going to devise the method of coming up with an evaluation method we need the visualization to be really crisp, and this method provides that.

You might be screaming "You rotten Qlik Dork ... just tell me

which of the measures is the right one to use!” To which my reply is a resounding “None of them and yet all of them.”

You see nobody said we had to use a single value to do the evaluation of end user adoption and there is so much more that we can do with Qlikview to present a more complete picture. The chart below slices and dices the data a few other ways that presents a different picture.

Document	# Minutes per Session ✓	Median Minutes	# Sessions per User
	21.0	7.0	32.7
Application 53	89.4	5.0	10.2
Application 38	45.0	21.5	18.7
Application 29	40.0	13.0	11.6
Application 7	38.3	35.0	20.5
Application 22	38.0	38.0	1.0
Application 42	37.0	10.0	2.2
Application 23	35.8	39.0	2.0
Application 25	34.1	19.5	30.3
Application 43	28.0	14.0	14.4
Application 52	27.9	10.0	6.6
Application 33	27.0	12.0	22.6
Application 6	26.9	14.5	4.7
Application 11	25.7	30.0	3.0
Application 44	25.1	9.0	15.5
Application 50	24.9	30.0	1.6
Application 19	24.4	9.5	7.0
Application 10	24.3	12.0	14.4
Application 26	24.2	13.0	16.3
Application 30	24.1	5.0	11.7
Application 49	23.1	8.0	23.8
Application 8	23.0	31.0	1.5
Application 87	23.0	33.0	1.0
Application 2	22.8	0.0	2.0
Application 60	22.3	9.5	4.0
Application 45	22.0	11.0	12.6
Application 48	21.9	30.0	4.5
Application 9	21.3	8.5	1.6
Application 21	21.1	4.0	8.0
Application 76	21.1	3.0	3.5
Application 20	20.4	8.0	19.4
Application 51	19.3	4.0	10.5
Application 58	18.4	12.0	2.0
Application 31	17.4	6.0	2.6
Application 28	17.2	6.0	15.2
Application 69	17.0	7.0	19.1

The first column interprets the average number of minutes per session. I might argue that value really represents user adoption of data analytics applications. Regardless if the application was built for a team of 5 or 50 to consume it reflects how long users stay engaged with the application. If we believe that is the goal then perhaps this is the perfect

measure. Woo hoo. I think I wrote application 53.

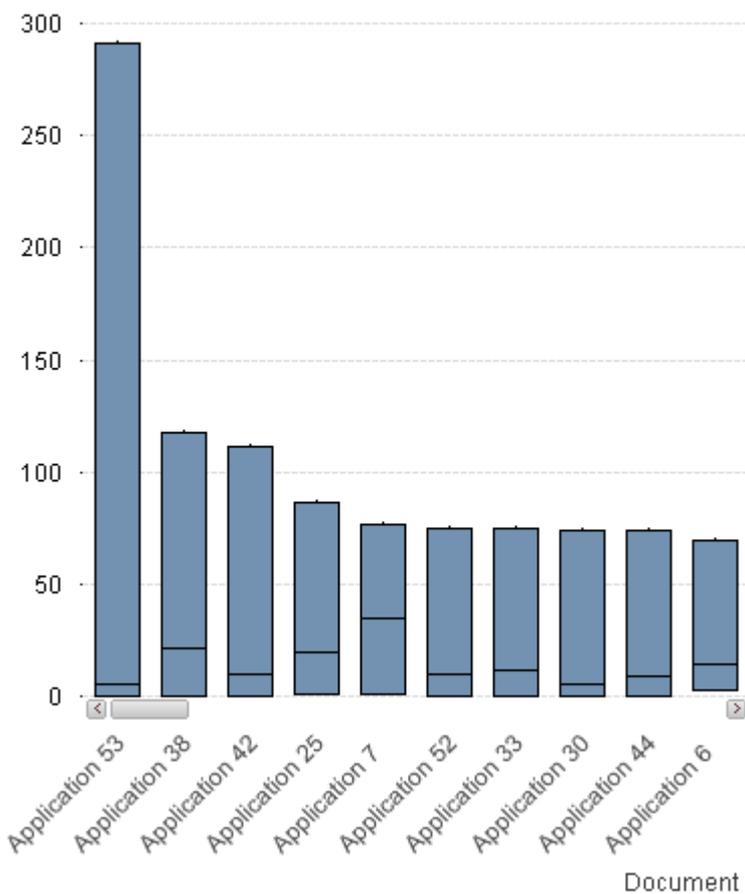
Oh wait a second the other developer raises their hand to complain yet again, and points out that average is a really poor statistical indicator and that Median is a better measurement because it isn't so swayed by outliers. In theory I agree, but as the author of application 53 it appears this statistics mumbo jumbo is costing me a big fat raise because while the average number of minutes per session is the highest, the median number of minutes per session is a measly 5. Phooey on heat maps I say, because if it weren't color coded nobody would have spotted the 5.

Whether we used average number of minutes, or median number of minutes both point out something very interesting. If you look at the very bottom and see the numbers for Application 69 it appears that any of the single measurements like # of Users/Sessions/Minutes alone didn't show a complete picture. Lots of total users and minutes, just not many minutes per session. Quantity for sure but not necessarily much analytical quality.

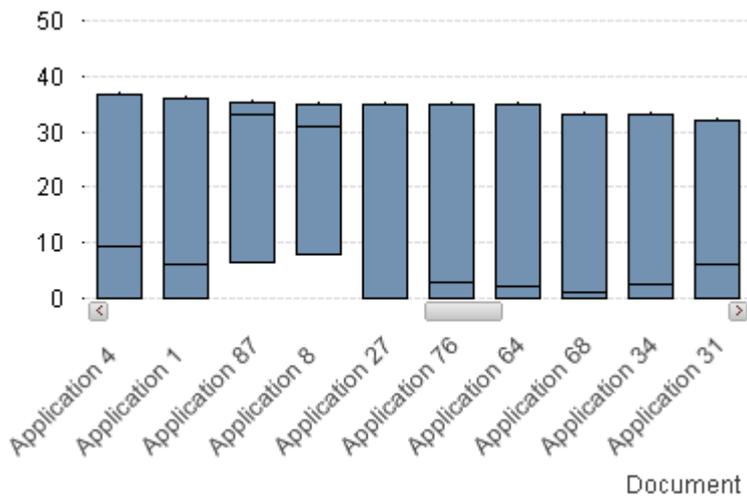
The third column illustrates a completely different measurement, the number of sessions per user. In other words how frequently are users engaging with our application. Like the raw data displayed in chart 1, displaying all 3 of these combined measurements helps paint a broader picture: Is our application engaging users for a very long time? Are they engaging once every 6 months, or are they coming back every other day and working?

Box plots to the rescue?

If we produce a box plot and make a few minor tweaks we can see that in fact Application 53 does in fact have a very high max value but the very low median of 5.



But the beauty of what a box plot visualizes for us can best be seen as I scroll to the right a bit. Notice for applications 87 and 8 both have pretty high medians, which we would see in the heat map chart, but more importantly you can see that even their lowest values are near 10 minutes per session. Meaning when these applications are used they are used for a good amount of time and the time is pretty consistent in a predictable range. Perhaps we could measure the end user adoption based on the predictability and consistency with which users engage?



Of course any kind of visualization of end user adoption would be incomplete if we didn't look at the values over time so that we could see if things were getting better, stabilizing or getting worse.

Document	# of Users	Over time	# of Sessions	Over time	# of Minutes	Over time
	265		8,662		181,493	
Application 69	91		1,740		29,539	
Application 85	39		56		445	
Application 49	30		713		16,468	
Application 56	30		763		11,369	
Application 33	29		655		17,701	
Application 58	27		53		977	
Application 28	26		395		6,788	
Application 16	25		36		187	
Application 62	23		213		3,550	
Application 77	23		329		4,800	
Application 39	21		140		2,316	
Application 42	20		280		8,000	

The wonderful thing is that while I focused on each method individually the great thing about visualizing data in Qlikview is that we can keep the entire picture together so that we get a true overview. A scorecard of sorts for each application.

Document	# of Users	# of Users	Over time	# of Sessions	# of Sessions	Over time	% of Sessions	# of Minutes	# of Minutes	Over time	% Minutes	Avg Minutes	Median Minutes	# Sessions per User
	265			8,662			100.0%	181,493			100.0%	21.0	7.0	32.7
Application 23	3			6			0.1%	215			0.1%	35.8	39.0	2.0
Application 22	1			1			0.0%	38			0.0%	38.0	38.0	1.0
Application 7	10			205			2.4%	7,848			4.3%	38.3	35.0	20.5
Application 87	3			3			0.0%	69			0.0%	23.0	33.0	1.0
Application 8	2			3			0.0%	69			0.0%	23.0	31.0	1.5
Application 11	1			3			0.0%	77			0.0%	25.7	30.0	3.0
Application 50	16			26			0.3%	647			0.4%	24.9	30.0	1.6
Application 48	6			27			0.3%	591			0.3%	21.9	30.0	4.5
Application 38	6			112			1.3%	5,035			2.8%	45.0	21.5	18.7
Application 25	6			182			2.1%	6,206			3.4%	34.1	19.5	30.3
Application 24	2			2			0.0%	30			0.0%	15.0	15.0	1.0

The truth about visualizing end user adoption

In the end the truth about measuring end user adoption is simple – every application is unique. Sorry to break this to you 7 pages in but you can't compare an application that was built to surface a small set of data in a scorecard fashion to an application that is really meant to be used as an ad hoc interface to find a cohort of patients. Perhaps 80 supervisory people per month use the scorecard for 1-2 minutes each time. While only 5 people use the other application for 2 hours per at a time multiple times per month.

Don't fret though? Just because I can't compare Application 69 to Application 53 doesn't mean I can't apply what I know about Application 69 and the intended audience and take steps to interact with the users and figure out how to improve the appropriate set of numbers. I may never get more than 10 users for Application 22, but if I can address issues that users have perhaps I can get them to engage 3 times per week instead of 3 times per year. I can add value to Application 72 and instead of end users engaging for 2 minutes per session I can increase their engagement to 10 minutes per session.