



# IPv6 Statistics on Vyncke.org

An overview

Eric Vyncke

[evyncke@cisco.com](mailto:evyncke@cisco.com)

[eric@vyncke.org](mailto:eric@vyncke.org) [eric.vyncke@ipv6council.be](mailto:eric.vyncke@ipv6council.be) [@evyncke](https://twitter.com/evyncke)

*“It is a capital mistake to theorize before one has data.”*

Sherlock Holmes, “A Study in Scarlet” (Arthur Conan Doyle)

# What to Measure? Which data sources?

# Type of Data

- How many IPv6 Internet users?
  - In app measurement => requires a specific app, app bias
  - Browser measurement => no specific requirements
- How many IPv6 ISP worldwide?
  - How many prefix allocated by Regional Internet Registries
  - How many IPv6 routes in global routing table?
  - How important are those ISP? Tier one, tier two...
- How much IPv6 contents?
  - How many web servers can be reached over IPv6 access?
- How much IPv6 traffic over the Internet ?
  - Internet Exchange Points ?

# Data Sources

- Browser data => need to have big site and make data public
  - Thank you Google, APNIC
- Allocated IPv6 prefix => easy thank you RIPE, ARIN, ...
- Global Routing Table => easy thank you routeview.org
- Content, get the list from Alexa, then query DNS and issue real HTTP requests
- Traffic: AMS-IX does not see Google, Facebook, Netflix traffic ...
  - Little information from ISP

# Even for Content, not so easy

- Free Alexa top 1-million global view not per country
  - Manual process to assign the .com, .net, .org to countries
  - With too many false positives for .cd, .to, .io, ... Or evenyoutu.be !
  - Attempt to measure the content **SERVED** by the country
- Pay Alexa a fee to get the top-500 per country
  - Nearly no false information
  - Measuring the content **SEEN** by the country
- Possible bias because Alexa is a voluntary plugin mainly on desktop

# Difference for the Netherlands

## Content in country <-> Content seen

- Booking.com
- Wetransfer.com
- Usenet.nl
- Ero-advertising.com
- Dumpert.nl
- Markplaats.nl
- Tweakers.net
- Ing.nl
- Google.nl
- Youtube.com
- Google.com
- Vk.com
- Facebook.com
- Wikipedia.org
- Reddit.com
- Ok.ru


# Let's Start Showing Data



# One Country Users by Google: Netherlands




# Multiple Countries Users by Google



# IPv6 Content Served by Netherlands (top-50)



# IPv6 Content Seen by Netherlands (top-100)



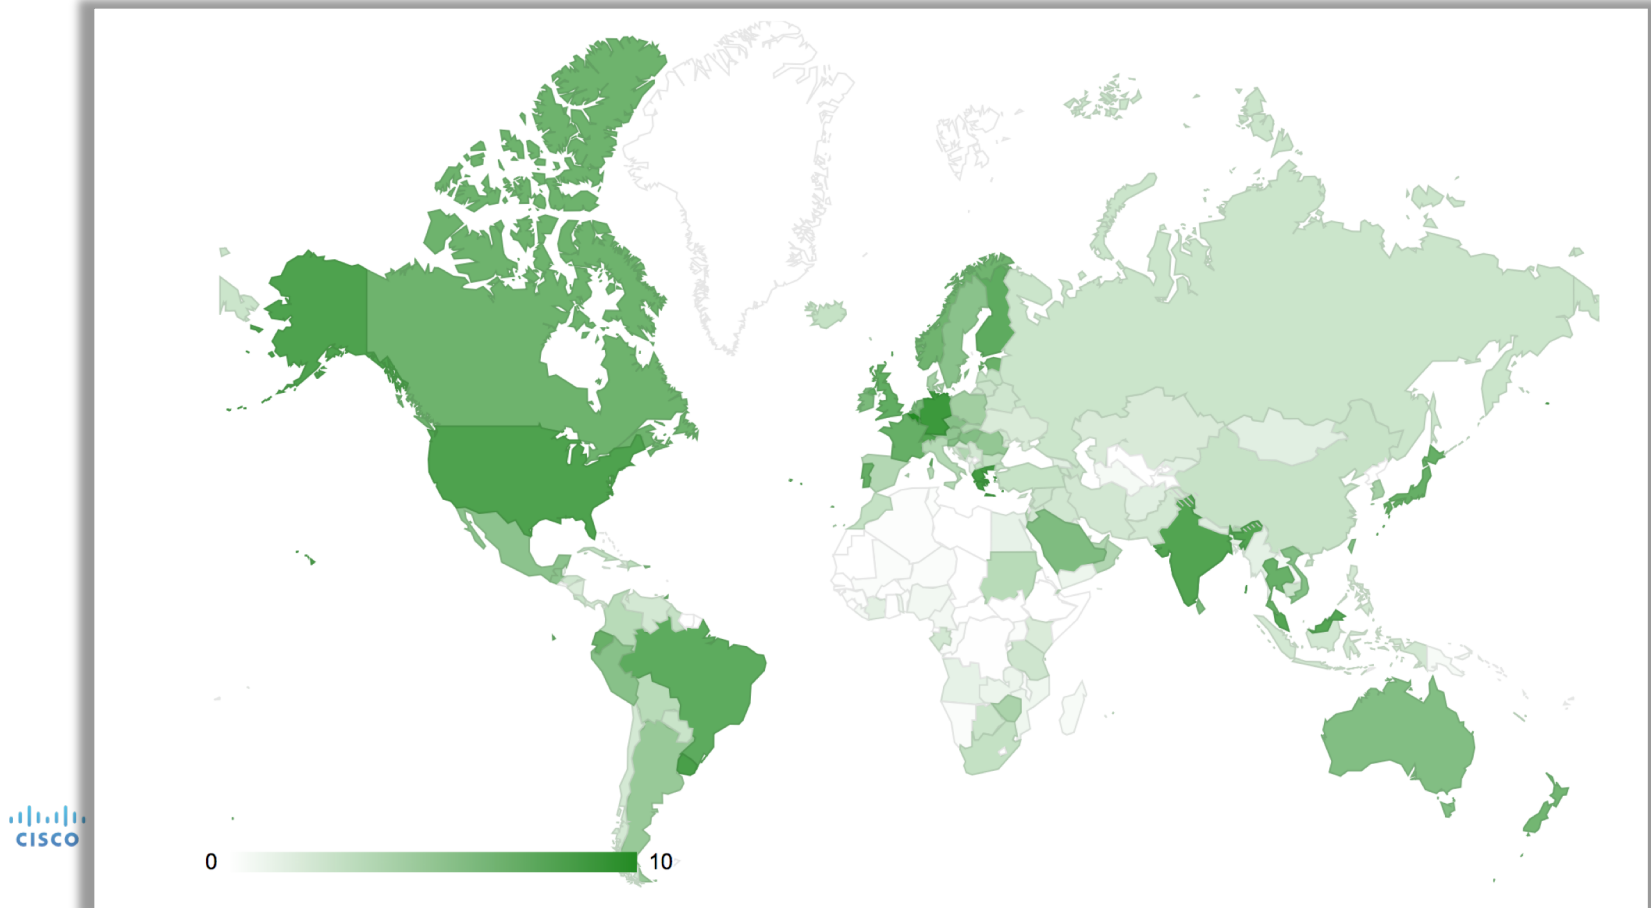
# Too Many Metrics ?

- Let's blend them !

$$RelativeMean = \frac{1}{4} \frac{\%TransitAS}{\max_{world} (\%TransitAS)} + \frac{3}{4} \frac{\sqrt{\%content \times \%user}}{\max_{world} (\sqrt{\%content \times \%user})}$$



# 6lab.cisco.com






# What about tomorrow?

Can past predict the future?

<https://www.vyncke.org/ipv6status/project.php?metric=p>

# Using Google Users Data: linear regression

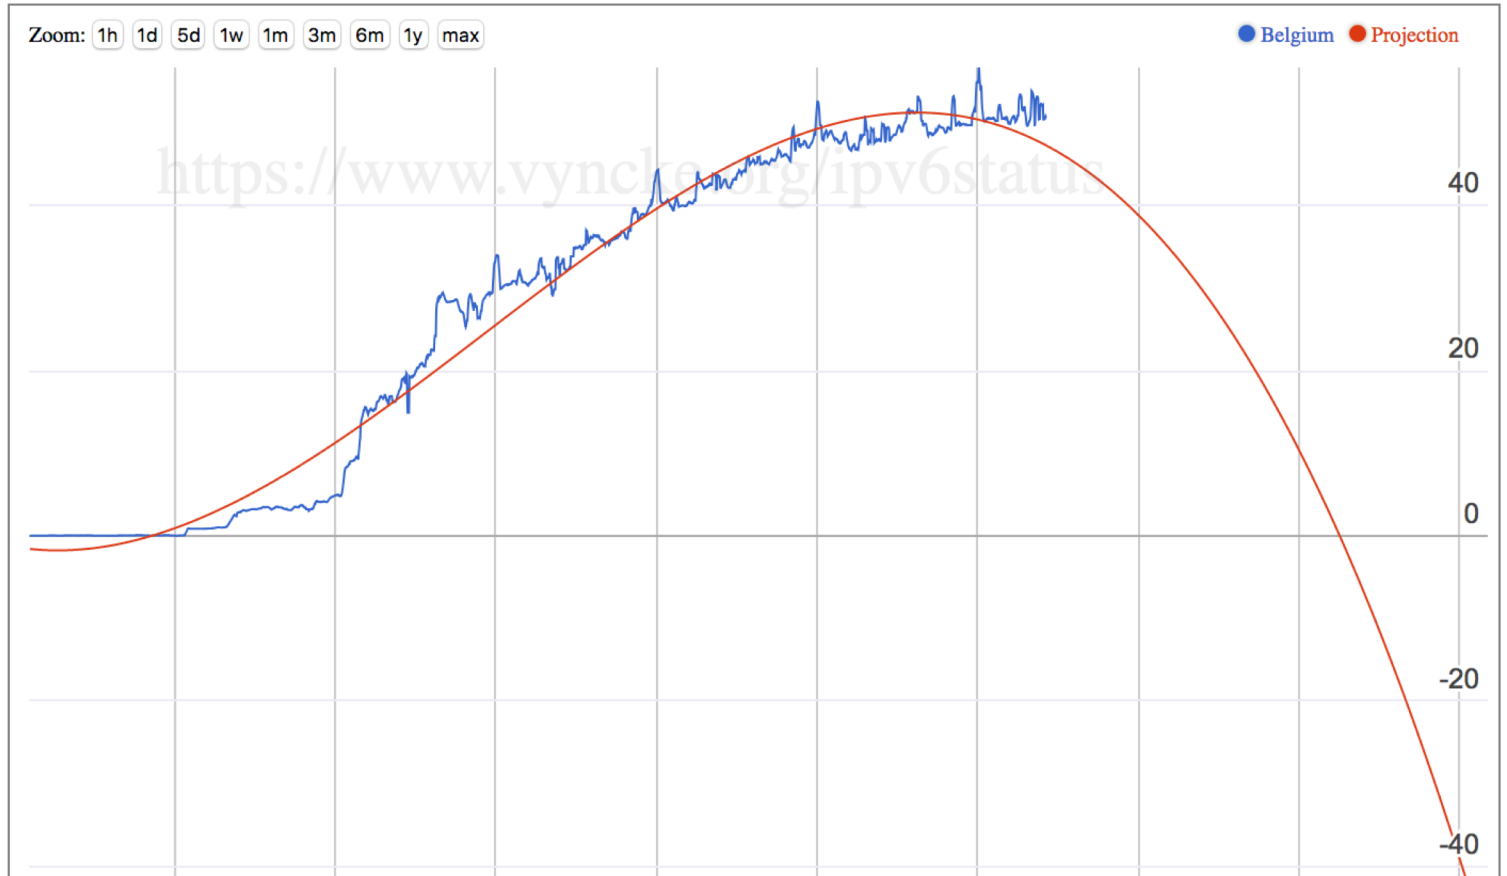




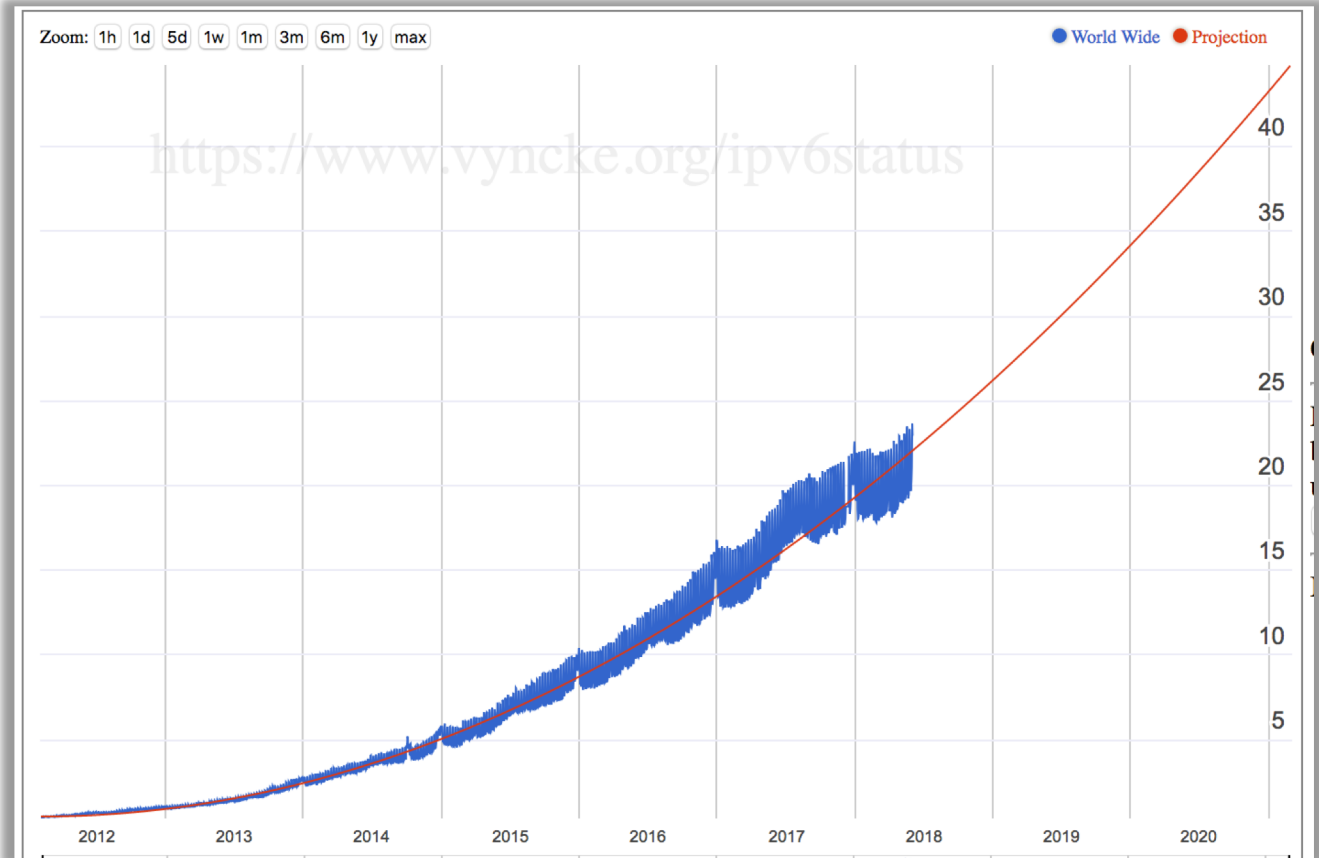
# Using Google Users Data: quadratic



# 3<sup>rd</sup> Order Regression is strange...



# Logistic Curve 'S-curve' is probably the most sensible



*“Do not forget that data can be biased, manipulated, shown on a specific light, ...”*