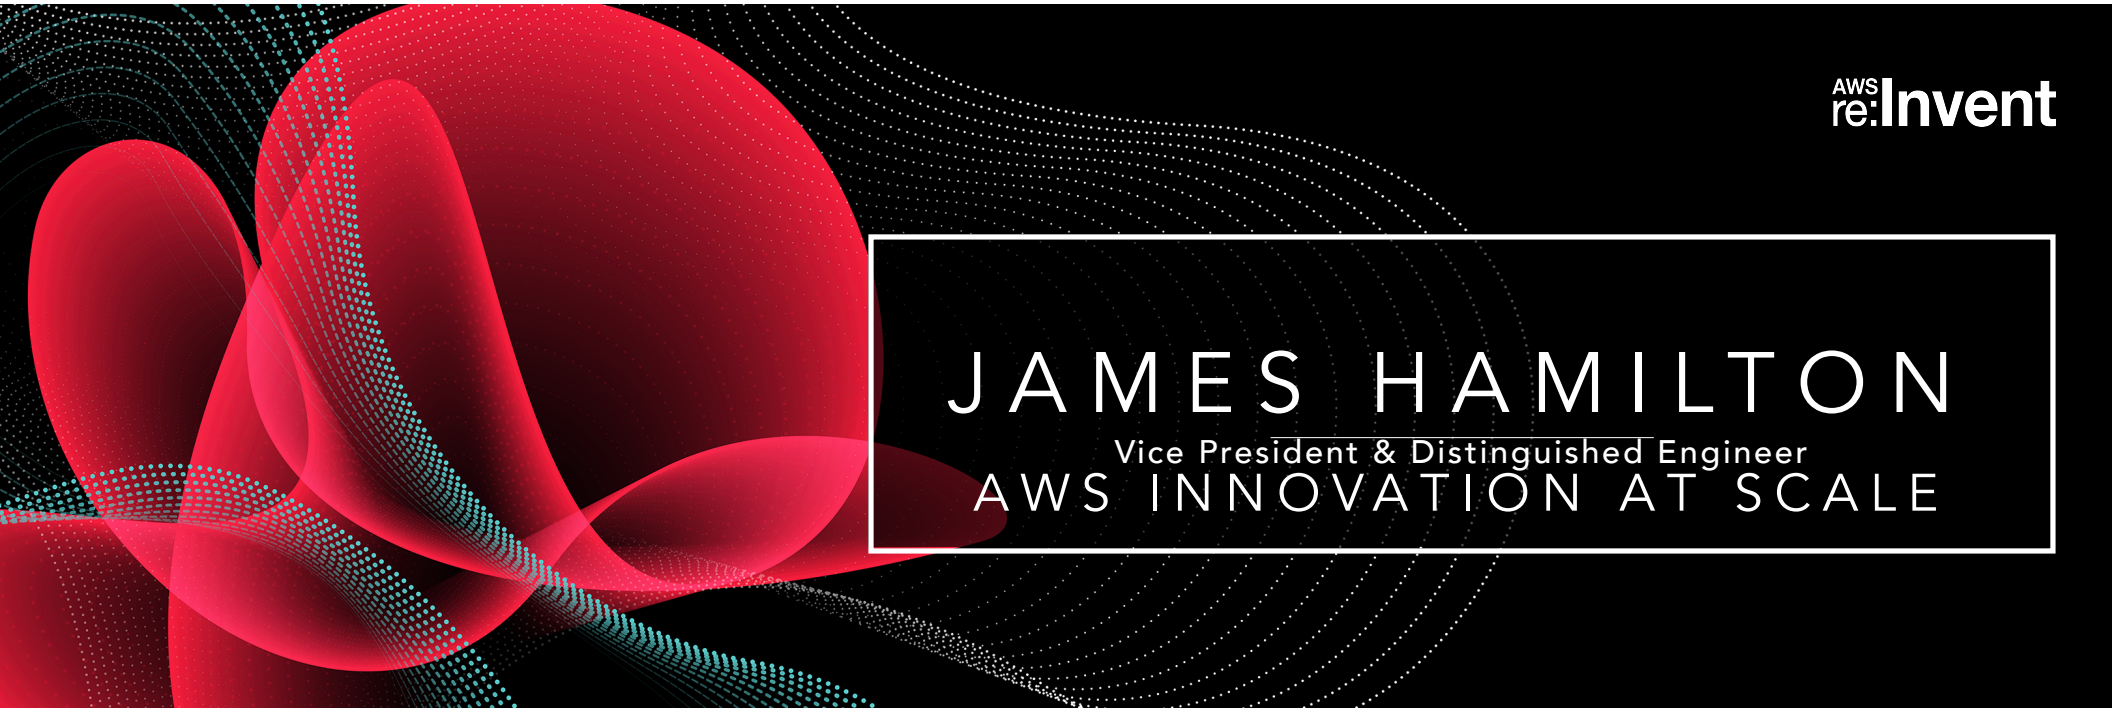
AWS re:Invent
JAMES HAMILTON
Vice President & Distinguished Engineer
AWS INNOVATION AT SCALE

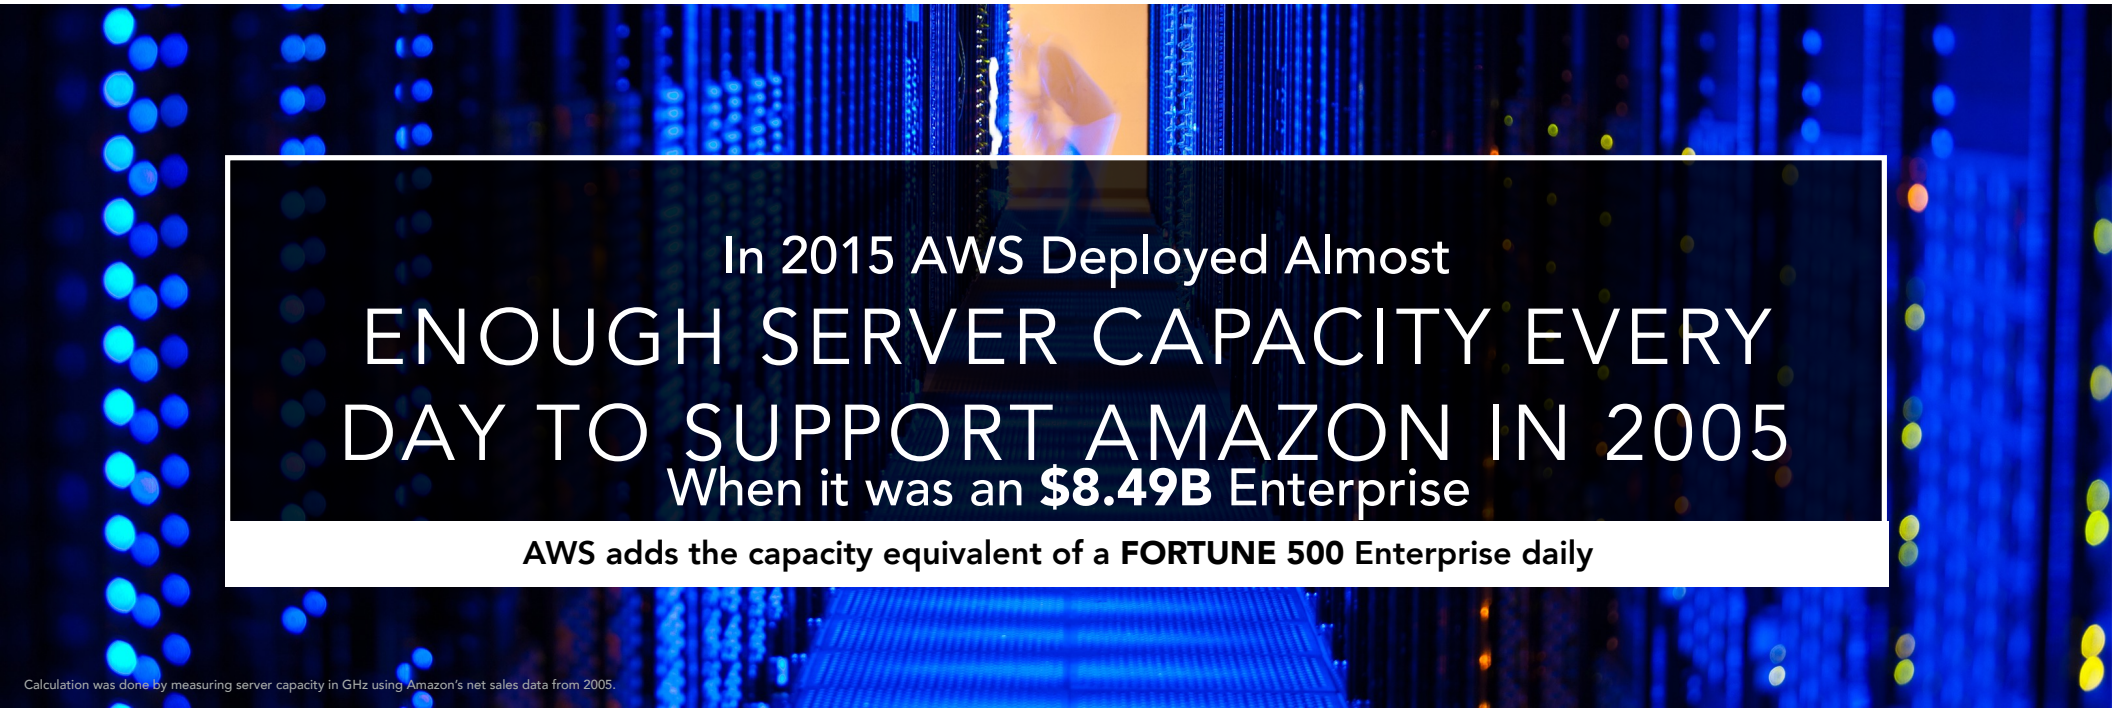
In 2015 AWS Deployed Almost
ENOUGH SERVER CAPACITY EVERY DAY TO SUPPORT AMAZON IN 2005
When it was an $8.49B Enterprise
AWS adds the capacity equivalent of a FORTUNE 500 Enterprise daily
Calculation was done by measuring server capacity in GHz using Amazon's net sales data from 2005.

ELASTIC
IS THE NEW
NORMAL

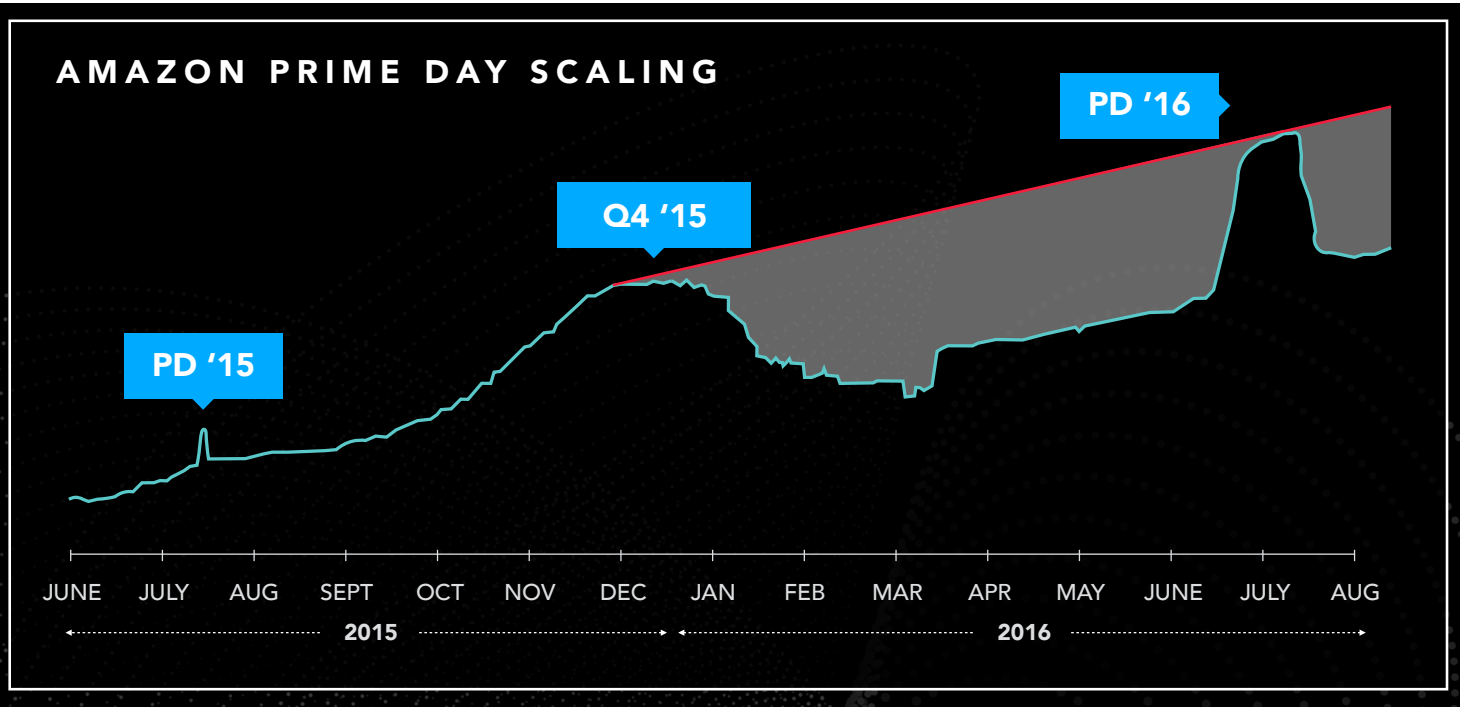


Oregon
N. California
AWS Govcloud
Ohio
Montreal
N. Virginia
Ireland
London
Frankfurt
Paris
Ningxia
Seoul
Beijing
Tokyo
Mumbai
Singapore
São Paulo
Sydney
14 Regions worldwide
growing to 18 regions
in the next year

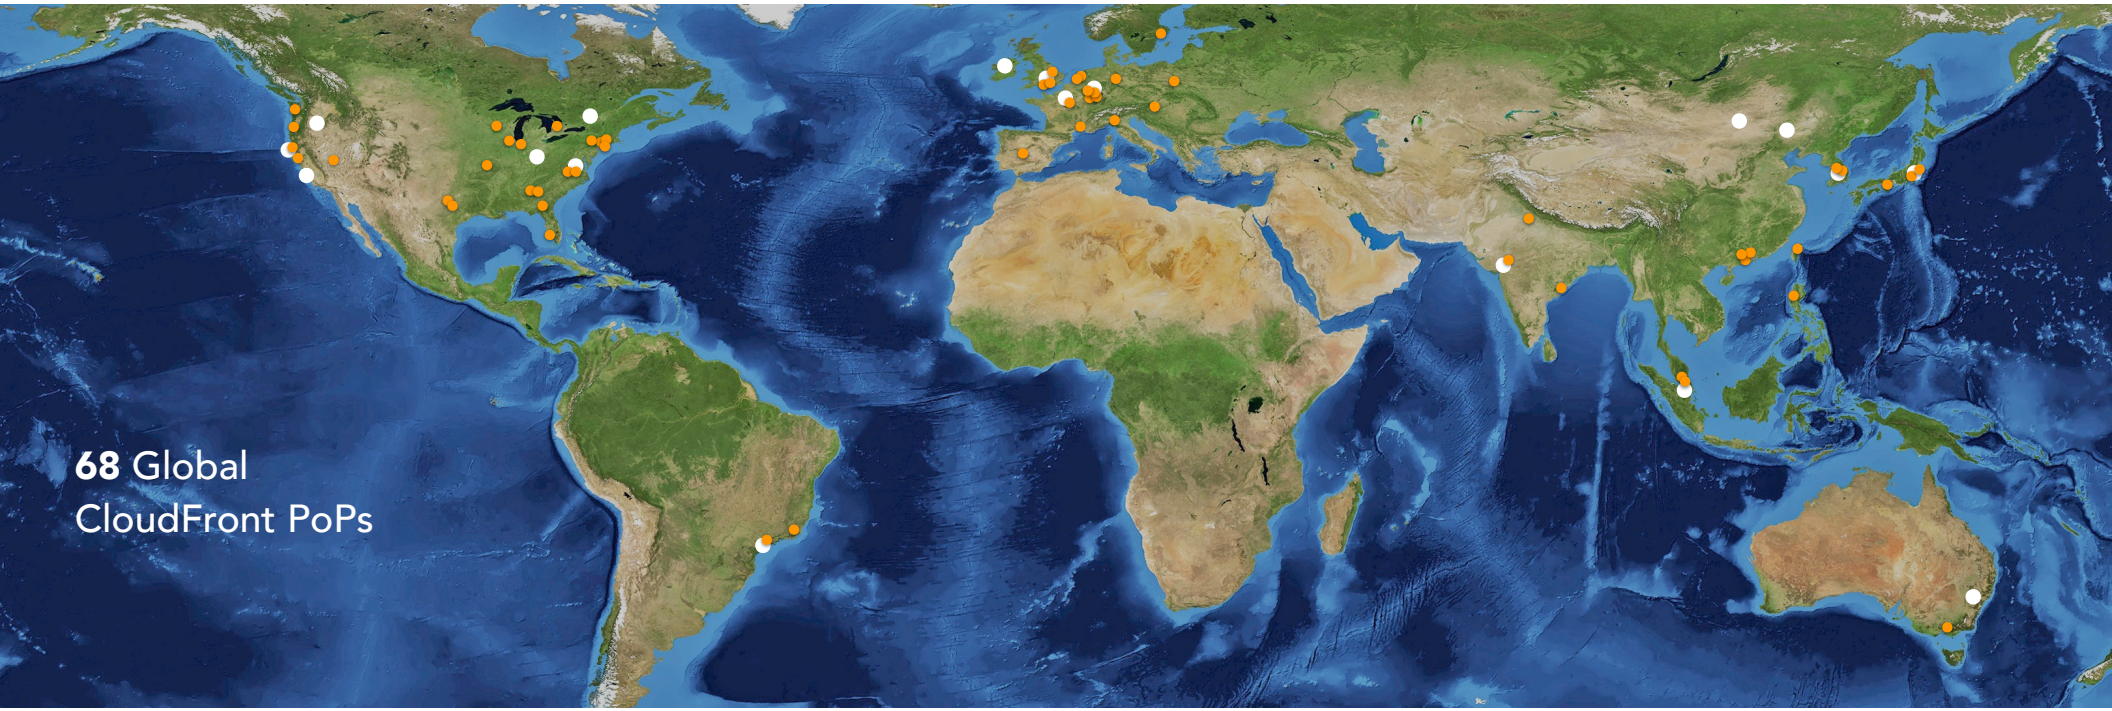
68 Global
CloudFront PoPs

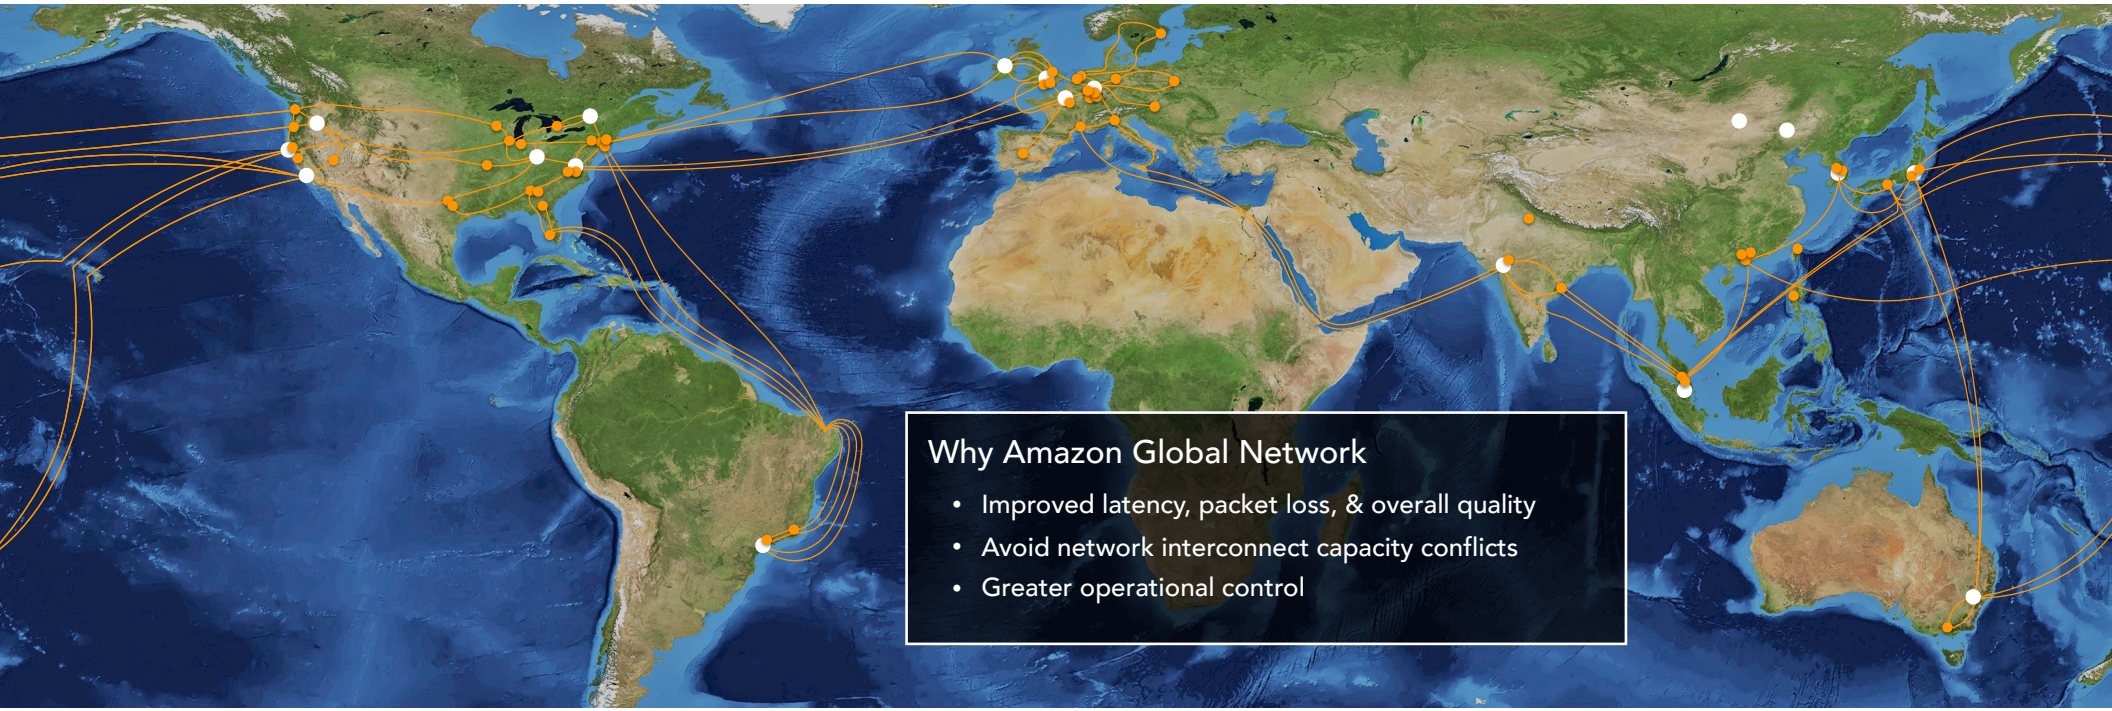
Why Amazon Global Network
• Improved latency, packet loss, & overall quality
• Avoid network interconnect capacity conflicts
• Greater operational control

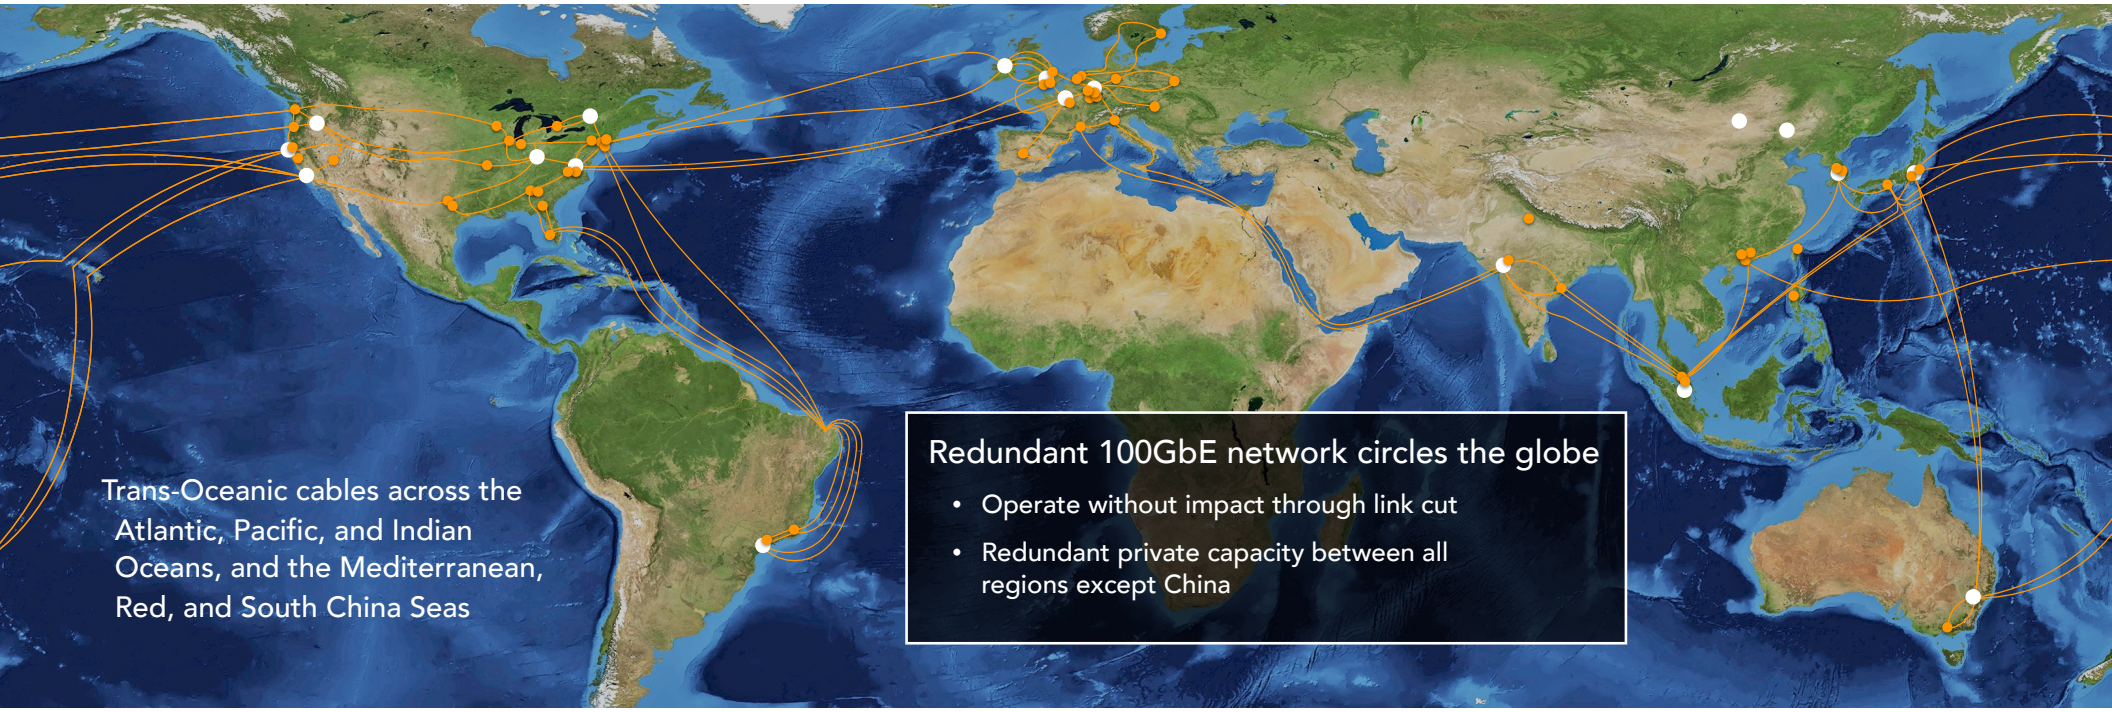
Trans-Oceanic cables across the Atlantic, Pacific, and Indian Oceans, and the Mediterranean, Red, and South China Seas
Redundant 100GbE network circles the globe
• Operate without impact through link cut
• Redundant private capacity between all regions except China

NEWEST PROJECT

- Hawaii trans-pacific cable
- 14,000km linking Australia, New Zealand, Hawaii, and Oregon
- 3 fiber pairs
- 100 waves @ 100G
- New Zealand shore side ground breaking last week

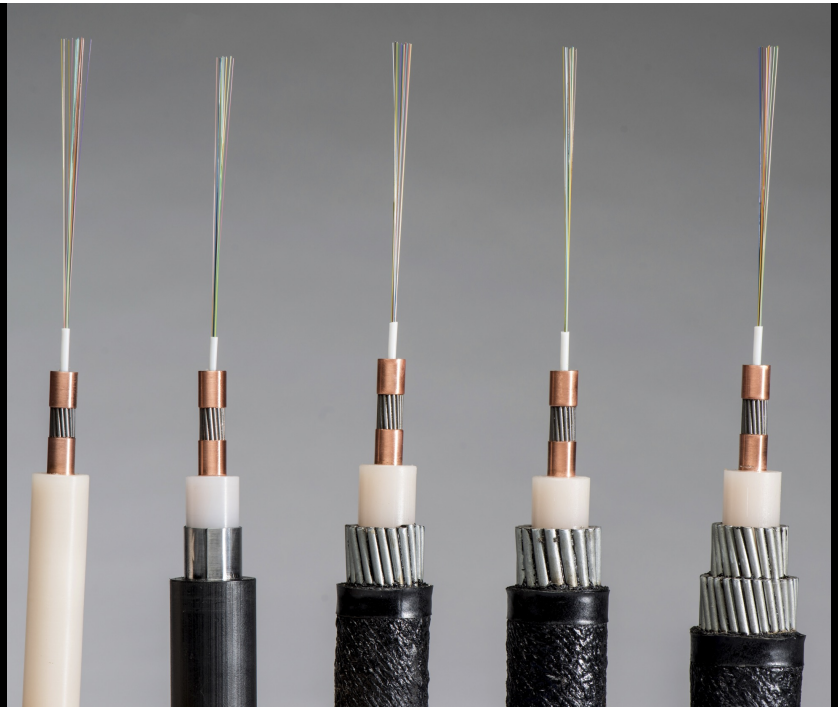

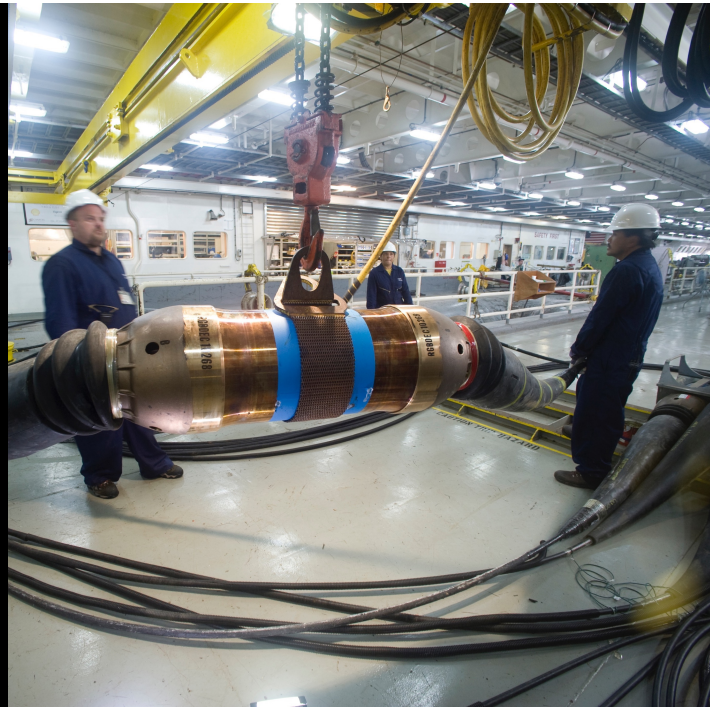

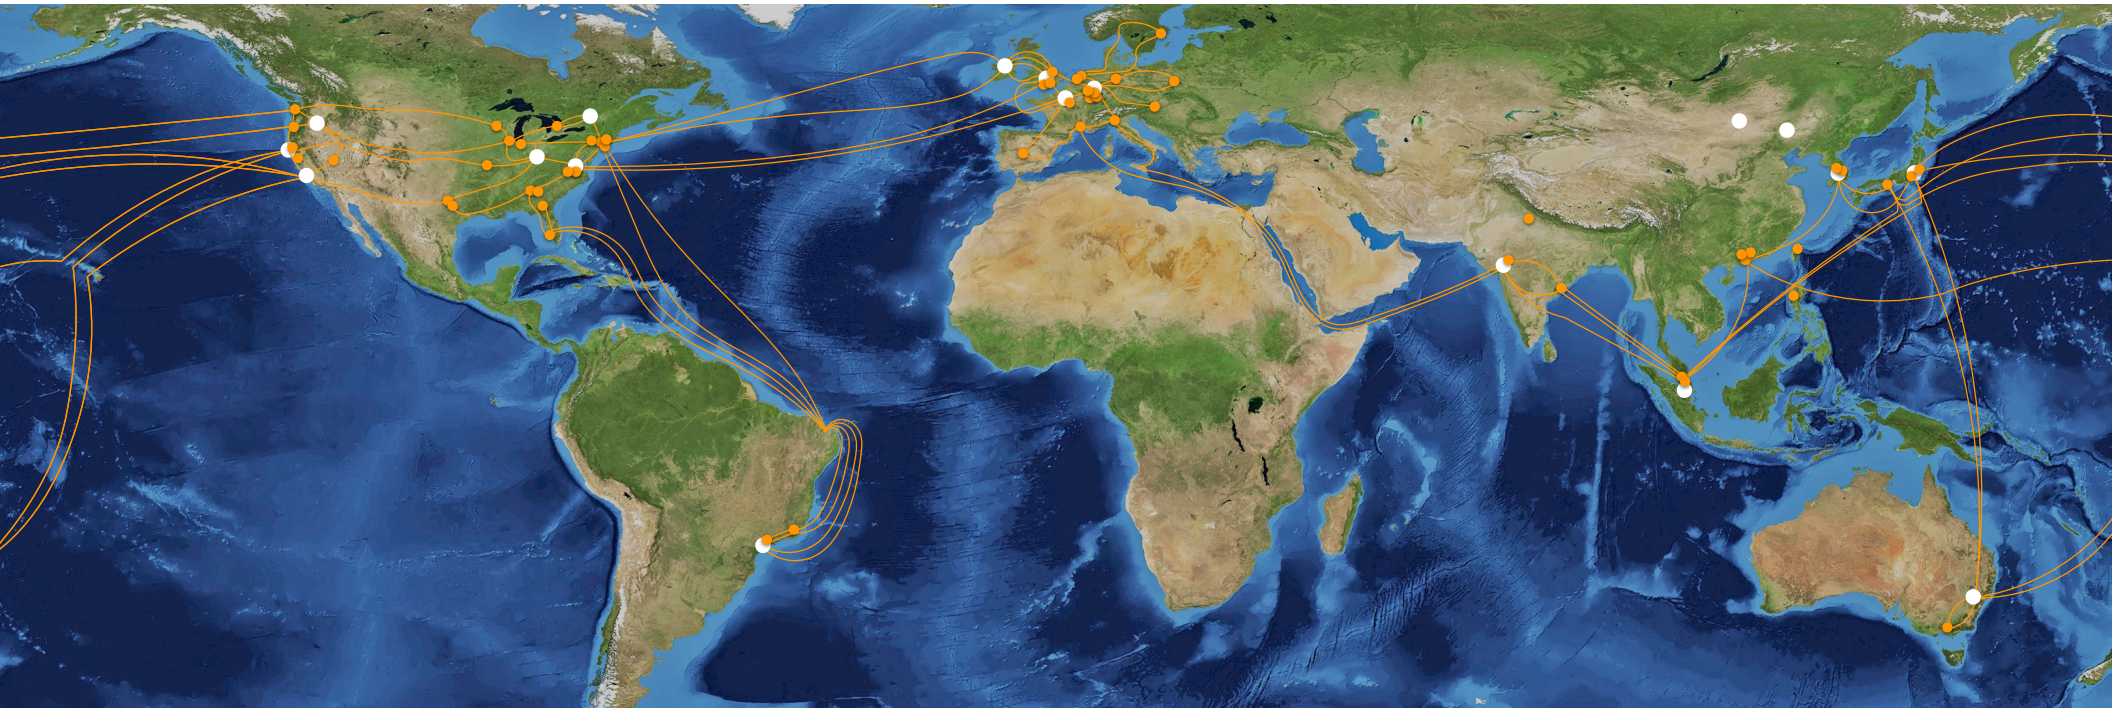

ACTUAL AWS REGION

- 1 of 14 worldwide
- All with 2+ AZs
- New builds all 3+
- As many as 5 AZs

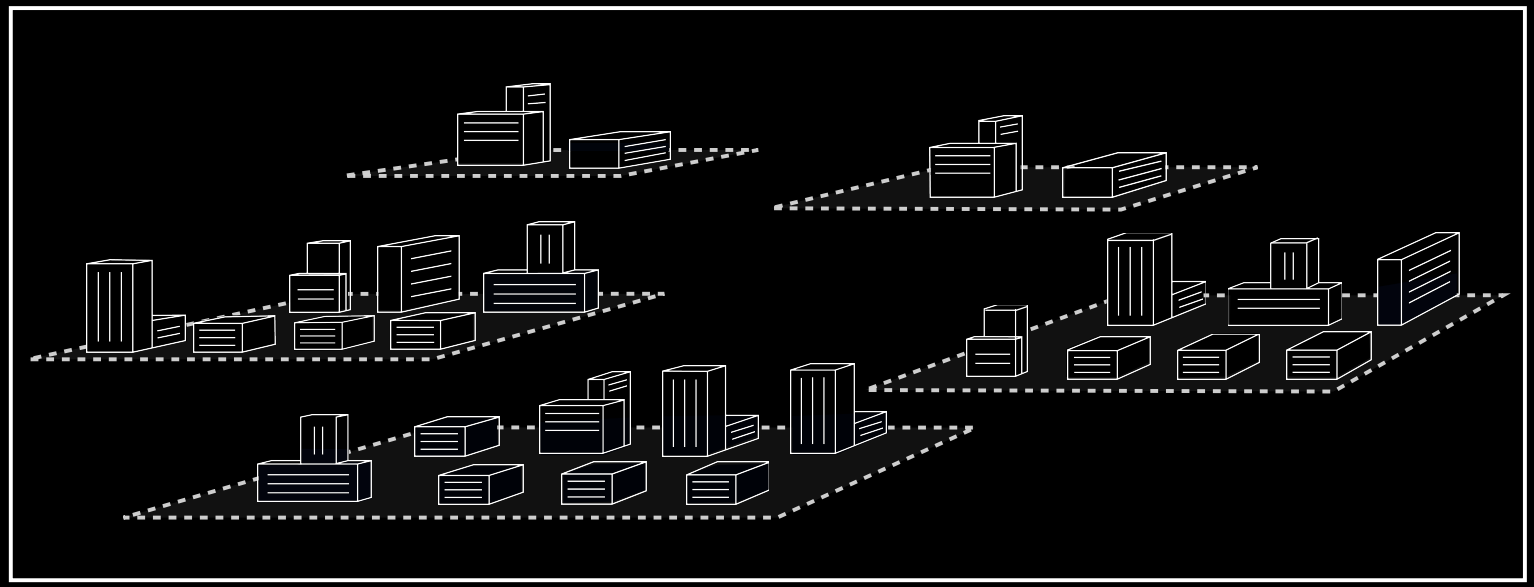

TRANSIT CENTERS

- 2 redundant Transit Centers
- Highly peered & connected facilities

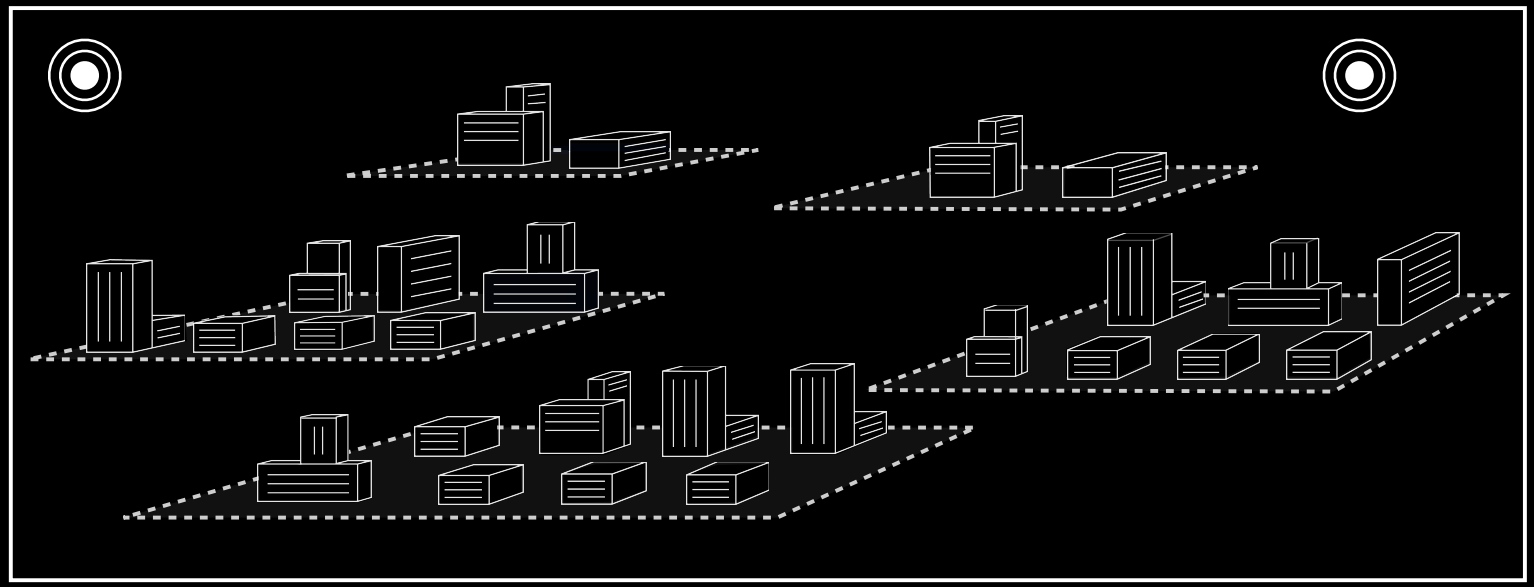

METRO FIBER

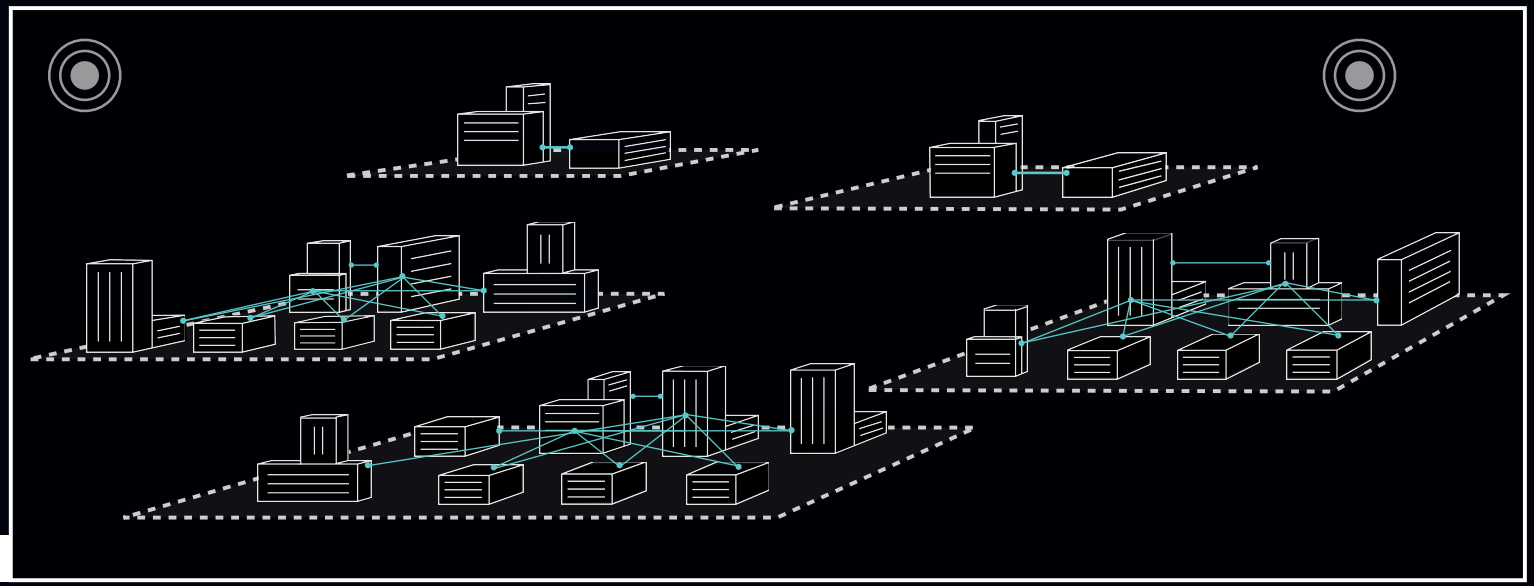

Intra-AZ connections

METRO FIBER

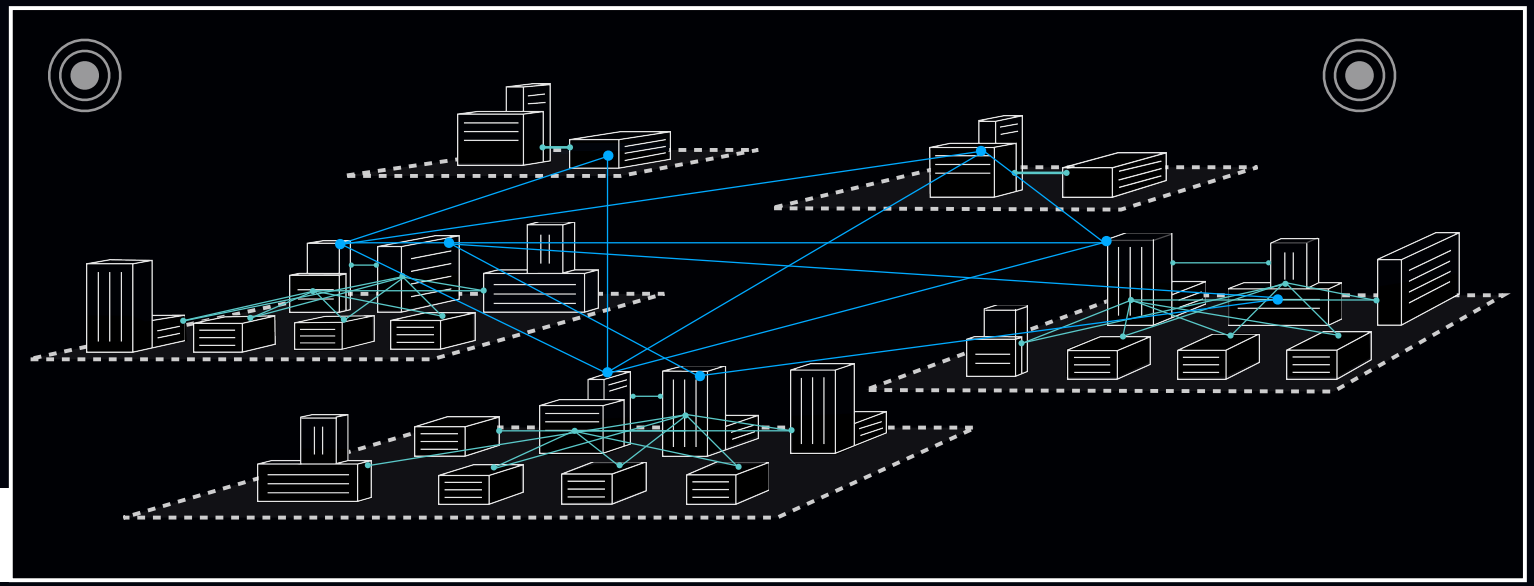

Inter-AZ connections

Intra-AZ connections

METRO FIBER

- **126** unique spans
- **242,472** total fiber strands
- AWS is the first company to deploy **3,456** fiber count cable

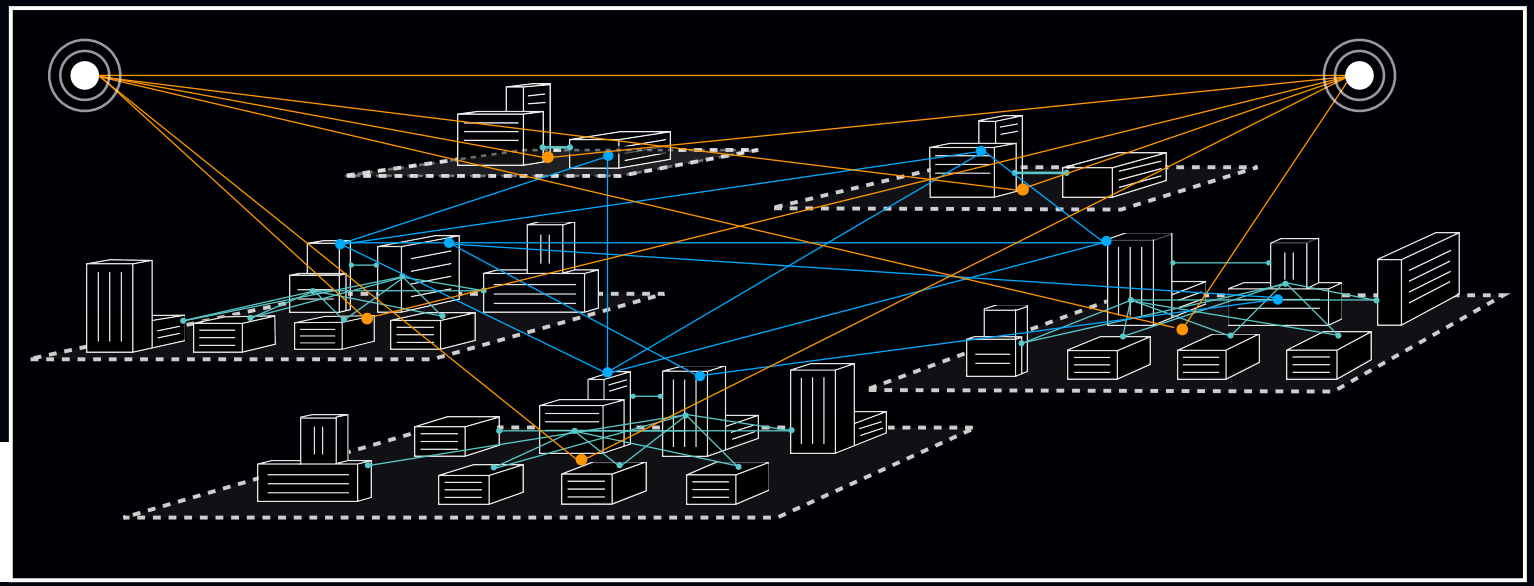

FULLY SCALED AZ

- Each AZ is 1+ data center
- Some with as many as 8
- Redundant network links
- Several AZs over 300K servers

DATA CENTER

- **60-120MW** or even larger are easy to build
- Larger scale drops cost slowly
- Larger scale increases blast radius quickly
- Redundant & concurrently maintainable

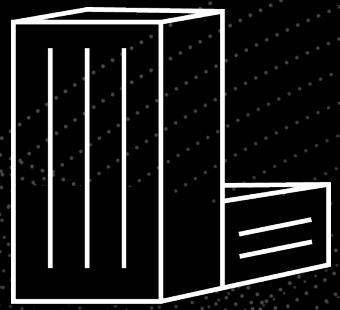

25–32MW

50k–80k servers

AWS CUSTOM ROUTERS

- Old school routers
  - Complex & unreliable
  - Expensive
  - ~6 months to correct issues
- AWS custom built routers
  - H/W built to spec
  - AWS protocol development team
- Committed to 25GbE early
  - Industry at 10GbE & 40GbE at time
  - Optics availability was tight
- 40GbE is actually 4x 10GbE
- 50GbE (2x 25GbE) is cheaper than 40GbE

AWS CUSTOM ROUTERS

- Old school routers
  - Complex & unreliable
  - Expensive
  - ~6 months to correct issues
- AWS custom built routers
  - H/W built to spec
  - AWS protocol development team
- Committed to 25GbE early
  - Industry at 10GbE & 40GbE at time
  - Optics availability was tight
- 40GbE is actually 4x 10GbE
- 50GbE (2x 25GbE) is cheaper than 40GbE

AWS CUSTOM ROUTERS

- Old school routers
  - Complex & unreliable
  - Expensive
  - ~6 months to correct issues
- AWS custom built routers
  - H/W built to spec
  - AWS protocol development team
- Committed to 25GbE early
  - Industry at 10GbE & 40GbE at time
  - Optics availability was tight
- 40GbE is actually 4x 10GbE
- 50GbE (2x 25GbE) is cheaper than 40GbE

CUSTOM ROUTERS

- AWS custom Broadcom Tomahawk ASIC
  - 7B transistors
  - 128 ports of 25GbE
  - 1RU, 22lbs, < 310W
- Core ASIC can be used in a variety of form factors
- Healthy ecosystem: Cavium, Mellanox, Broadcom, Innovium, Barefoot, Marvell

SOFTWARE DEFINED NETWORKING

- AWS SDN-based since beginning of EC2
- 2012 moved to hardware offload
  - Custom 10GbE NIC
  - Custom processor with AWS software
- Offload server network virtualization overhead
- Lower latency & less server jitter
  - SR-IOV & Enhanced Networking
  - < 70 µsec avg RTT in Placement Group

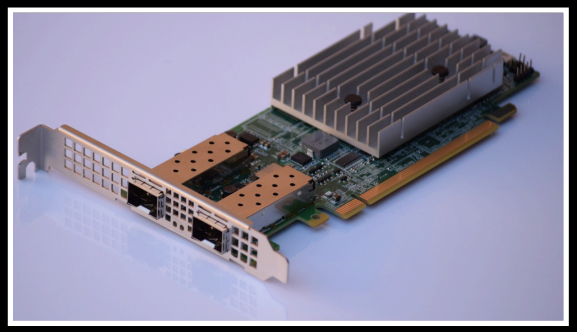

2016 CUSTOM SILICON

- Custom Si & 25GbE
  - 2x 25GbE cheaper & higher bandwidth than 40GbE
- Amazon Annapurna ASIC
  - Second generation Enhanced Networking
  - AWS controls silicon, hardware & software
  - AWS pace of innovation
- Instance peak bandwidth to 20GbE
  - Small instance peak bandwidth at 10GbE
  - Most instance types going forward

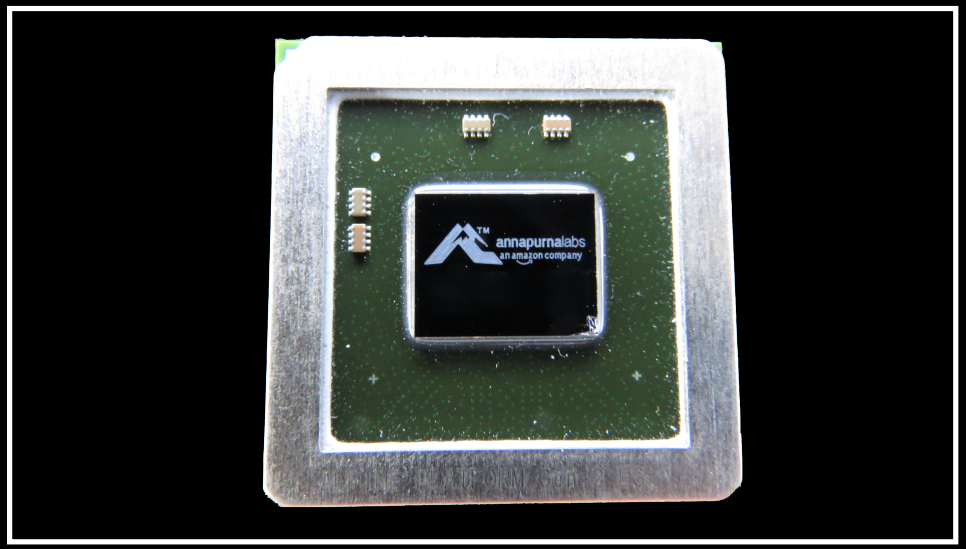

RARE POWER EVENTS

- Major US airline world wide outage
  - Some servers failed over & some lost power
  - $100M lost revenue (~2% for the month)
  - Switchgear failed & locked out reserve generators
- Customer impact
  - Monday: 1,000 flights cancelled
  - Tuesday: 775 flights cancelled
  - Wednesday: 90 flights cancelled
- Likely first time this fault seen by this operator
  - No compression algorithm for experience

CUSTOM UTILITY SWITCHGEAR

- Airline fault mode same as 2013 Super Bowl
- Switch gear locks out backup power
- The data center goes black in 5 to 10 minutes
- Amazon custom firmware protects the load
  - If faults outside, the full facility continues
  - If inside, only branch breaker opens, but no load dropped

CUSTOM STORAGE SERVER

- 2014: I showed 880 disks/rack
- Next design supported:
  - 1,110 disks/rack
  - 8.8PB at design time (would be 11PB today)
  - 2,778 lbs of storage
- More advanced designs now in production

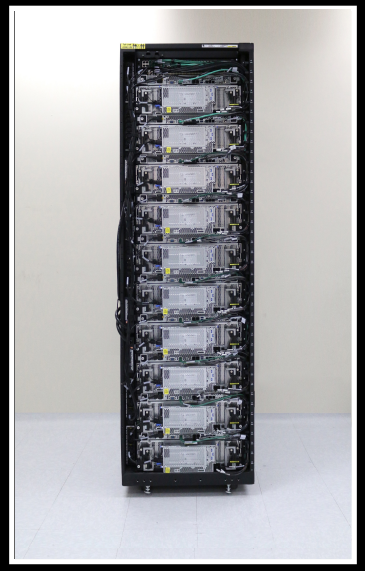

CUSTOM COMPUTE SERVER

- Simple, no-frills 1RU server
- Thermal & power efficiency favored over density
- PSU & VRD >90% efficiency
- Replaced by newer design
  - Still compares favorably to some recently blogged cloud servers

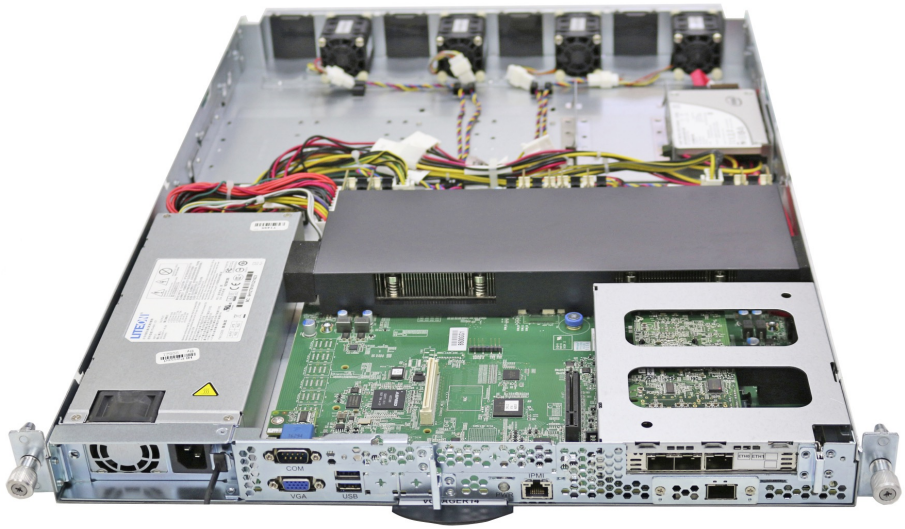


Long-term commitment to achieve
100%
RENEWABLE ENERGY


25%
April 2015
0%
100%

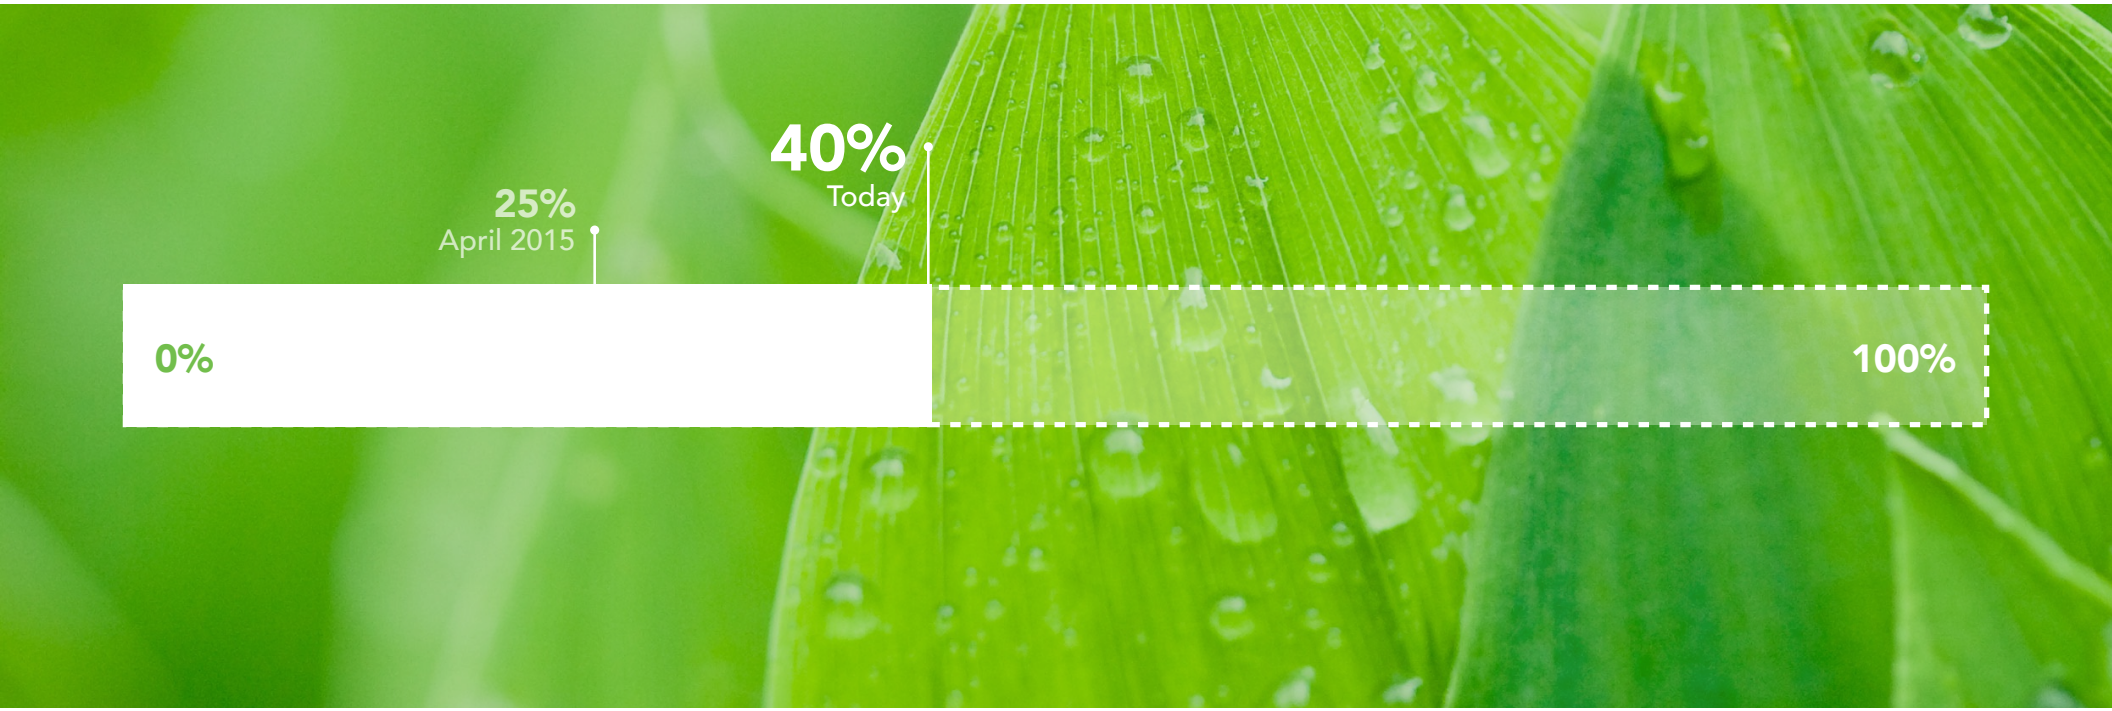
40%
Today
25%
April 2015
0%
100%

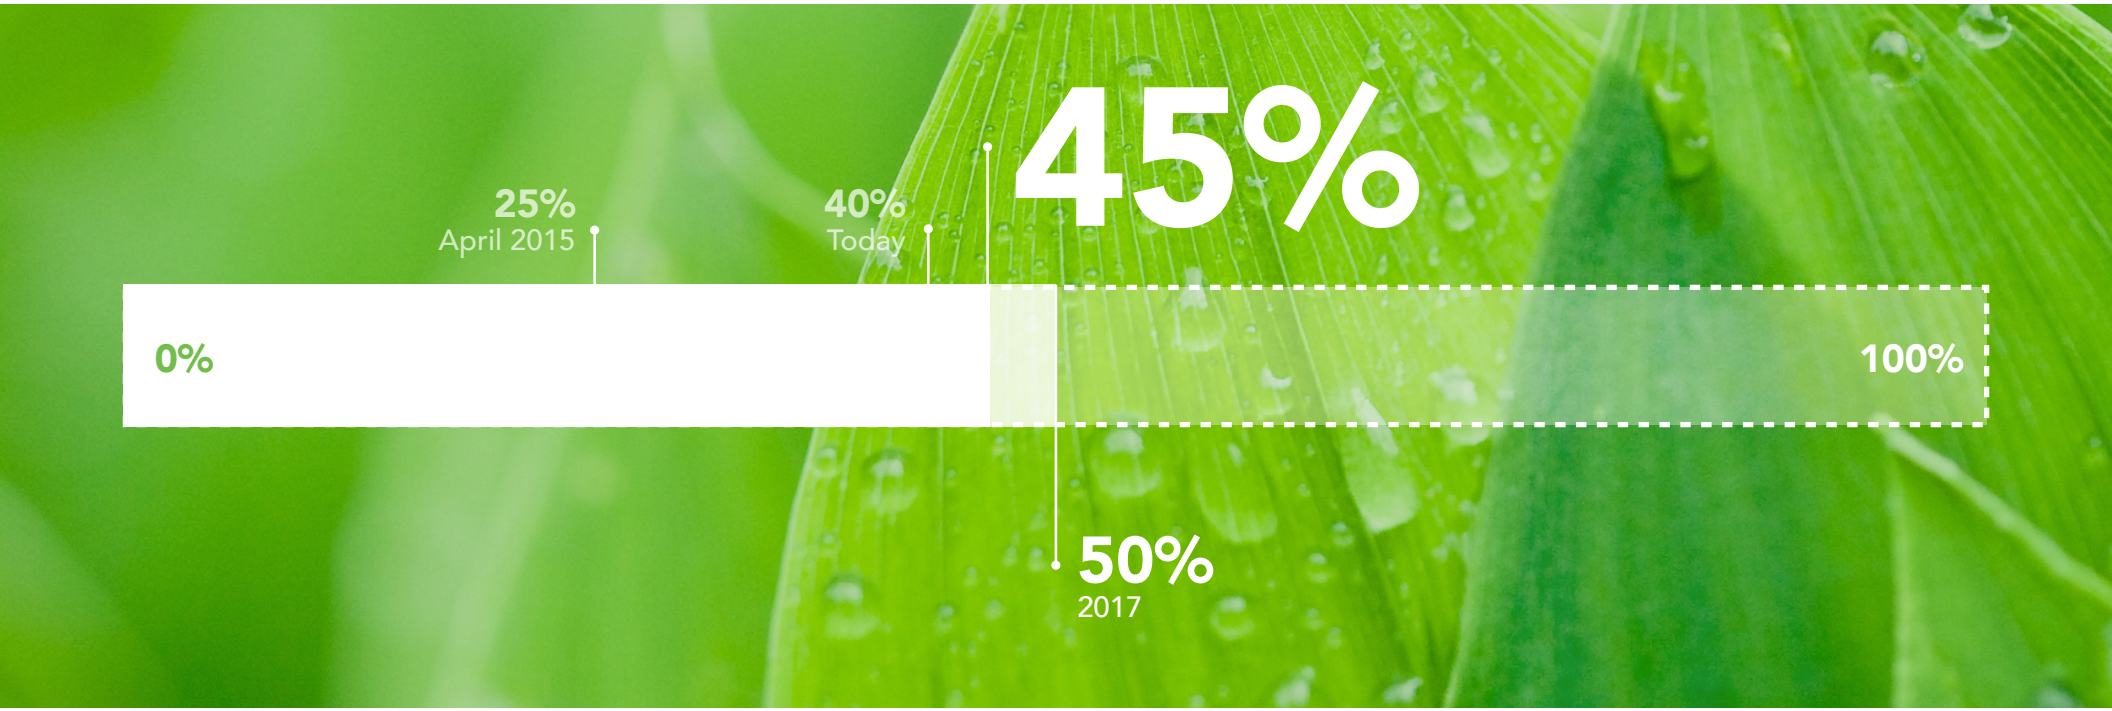
45%
25%
April 2015
40%
Today
0%
100%
50%
2017


AMAZON
WIND FARM FOWLER RIDGE
150MW
Jan '15


AMAZON
WIND FARM
FOWLER
RIDGE
150 MW
AMAZON
SOLAR FARM EAST
80MW
Jan '15
June '15

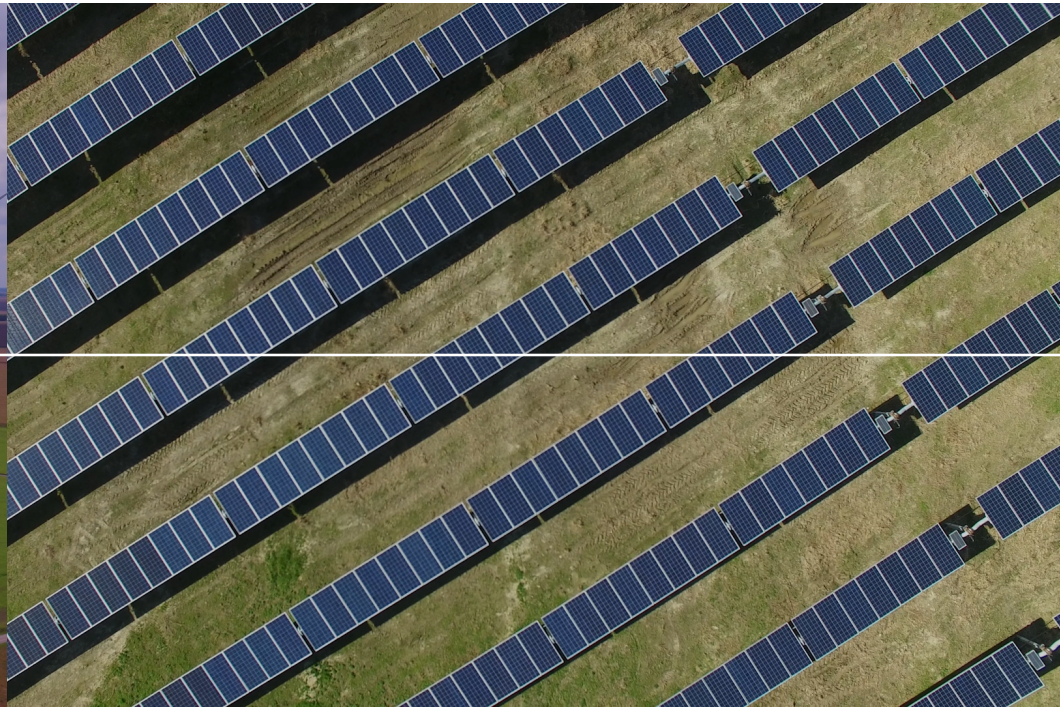


AMAZON
WIND FARM
FOWLER
RIDGE
150 MW
Jan '15
AMAZON
SOLAR
FARM
EAST
80MW
June '15
AMAZON
WIND FARM US EAST
208MW
July '15


AMAZON
WIND FARM
FOWLER
RIDGE
150 MW
Jan '15
AMAZON
SOLAR
FARM
EAST
80MW
June '15
AMAZON
WIND
FARM
US EAST
208MW
July '15
AMAZON
WIND FARM US CENTRAL
100MW
Nov '15


AMAZON
WIND FARM
FOWLER
RIDGE
150 MW
Jan '15
AMAZON
SOLAR
FARM
EAST
80MW
June '15
AMAZON
WIND
FARM
US EAST
208MW
July '15
AMAZON
WIND
FARM US
CENTRAL
100MW
Nov '15
AMAZON
WIND FARM US CENTRAL 2
189MW
Oct '16

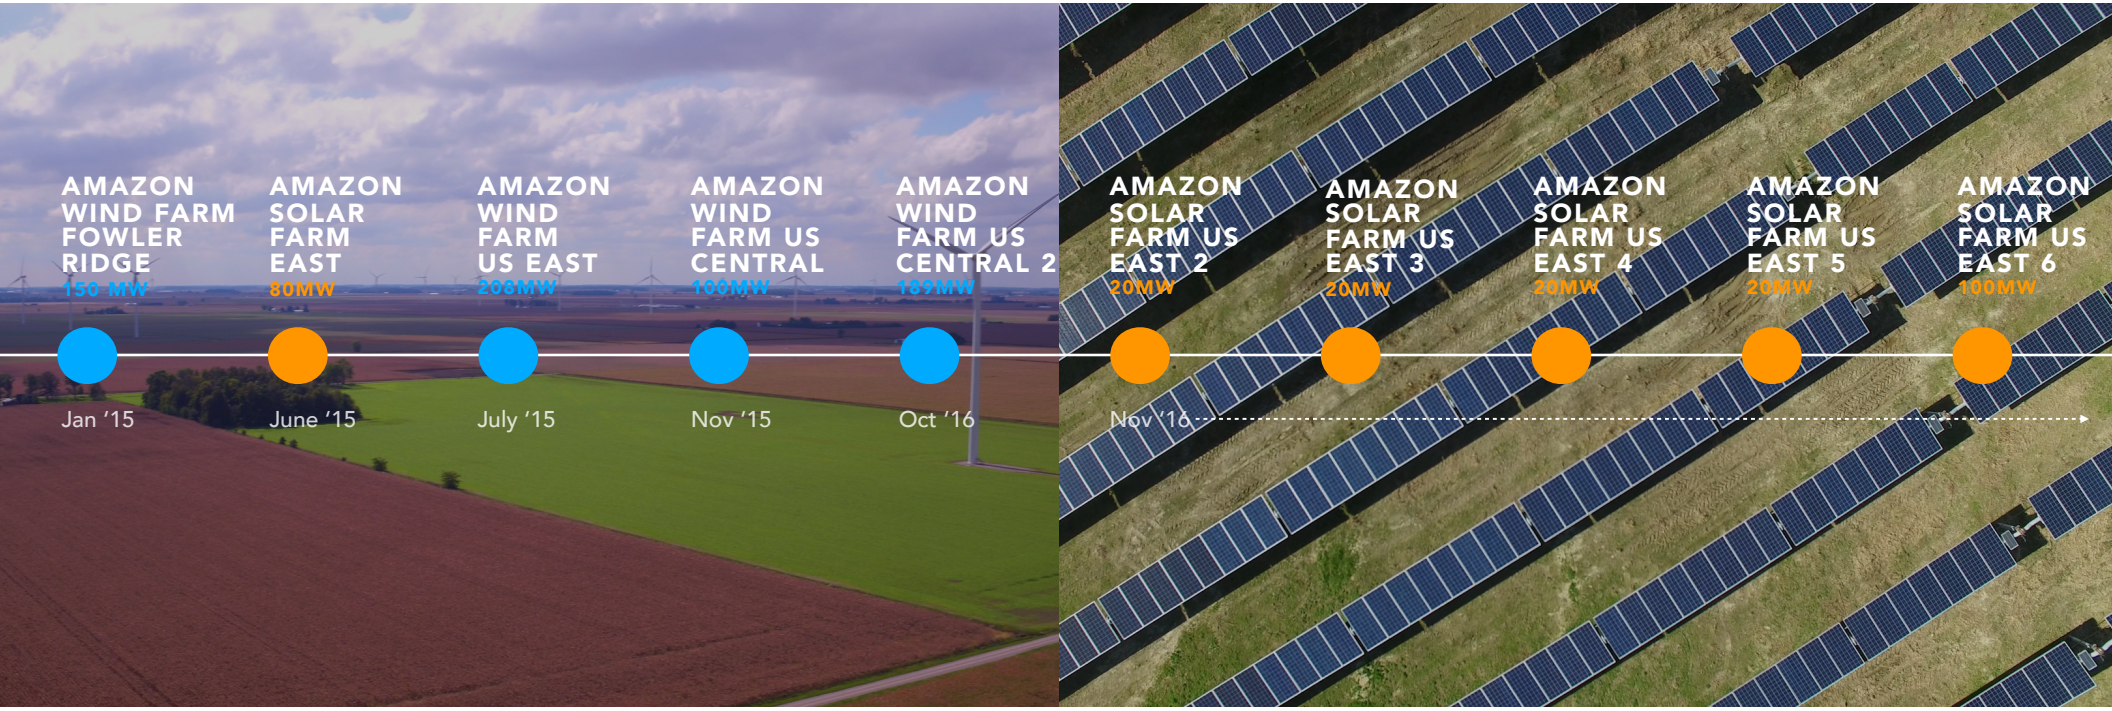
AMAZON WIND FARM FOWLER RIDGE
150 MW
Jan '15
AMAZON SOLAR FARM EAST
80MW
June '15
AMAZON WIND FARM US EAST
208MW
July '15
AMAZON WIND FARM US CENTRAL
100MW
Nov '15
AMAZON WIND FARM US CENTRAL 2
189MW
Oct '16
AMAZON SOLAR FARM US EAST 2
20MW
Nov '16
AMAZON SOLAR FARM US EAST 3
20MW
AMAZON SOLAR FARM US EAST 4
20MW
AMAZON SOLAR FARM US EAST 5
20MW
AMAZON SOLAR FARM US EAST 6
100MW


AWS Projects bring online **907MW**
of new renewable generation


AWS Projects bring online
Projects will generate
2.6 million MWhr of energy annually


AWS
re:Invent
THANK YOU