Elements of a Scalable Infrastructure for Weather Forecaster Access to Joint Polar Satellite System (JPSS) Data

John D. Evans, Ph.D.
Global Science & Technology, Inc. (GST)
NOAA Joint Polar Satellite System
Algorithm Management Project
Lanham / Greenbelt, MD

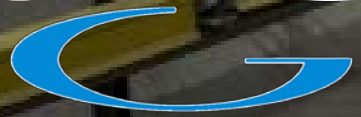

99th American Meteorological Society Annual Meeting
15th Annual Symposium on New Generation Operational Environmental Satellite Systems
5B Special Session on JPSS Series Satellite System – Part II – Jan. 9, 2019

Joint Polar Satellite System

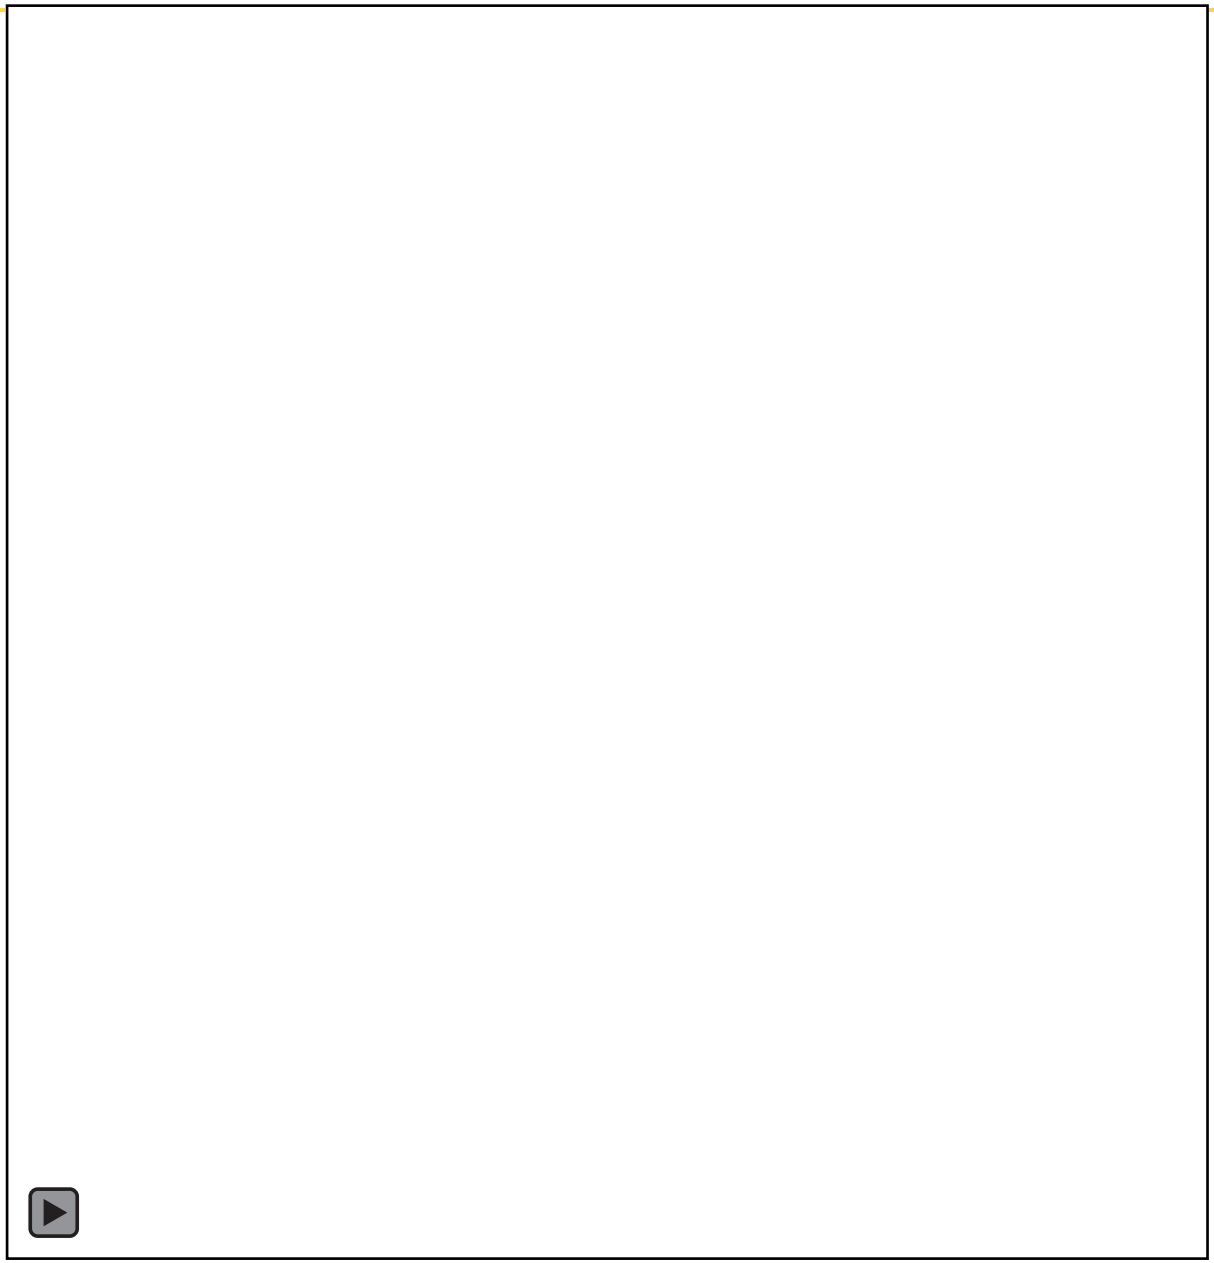

- Polar-orbiting satellites, S-NPP and NOAA-20 (a.k.a. JPSS-1)
- Instruments:
  - Visible/Infrared Imaging Radiometer Suite (VIIRS)
  - Advanced Technology Microwave Sounder (ATMS)
  - Cross-track Infrared Sounder (CrIS)
  - Ozone Mapping and Profiler Suite (OMPS)
- Each satellite orbits 14x/day
- Each images the globe 2x/day
- Each produces 2TB/day globally
- Challenge: providing forecasters the data they need in a timely fashion

Most weather forecasters don't need all of that data

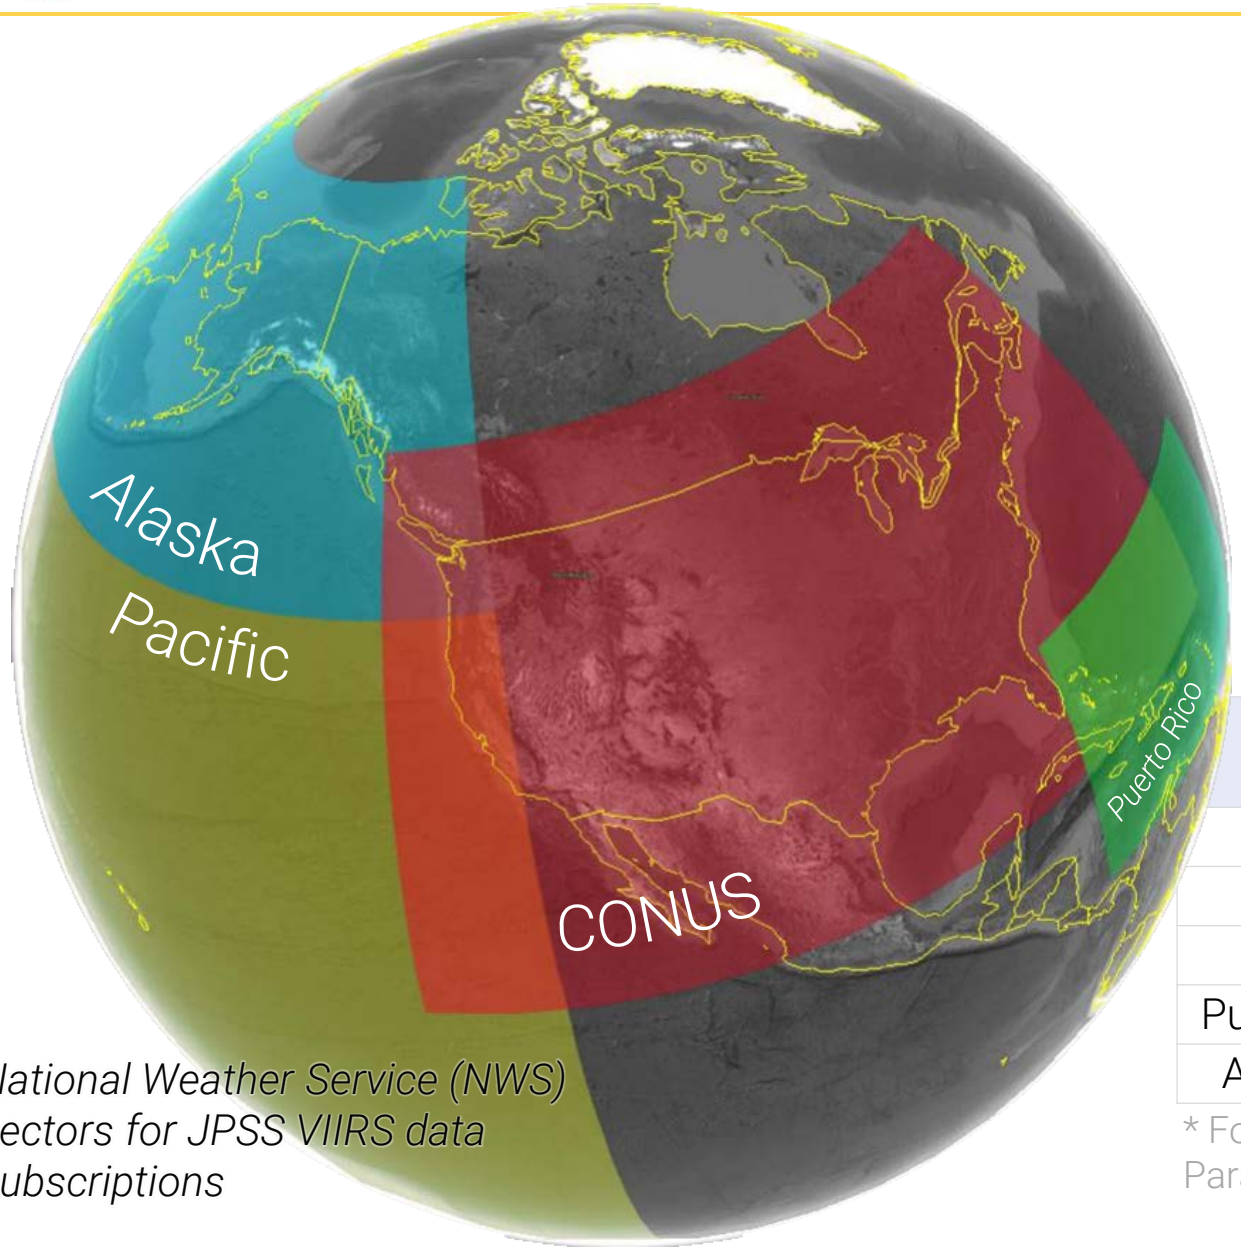

National Weather Service (NWS) sectors for JPSS VIIRS data subscriptions

- Forecasters in the field generally need regional data, rather than global.
- They also don't need every single format / aggregation / variant of data
- No one forecasting office will need every data product.

Sector	*Daily VIIRS granules, per satellite*	*GB/day, per satellite **	*NWS sites*
Alaska	128	74	5
Pacific	216	124	3
CONUS	78	45	128
Puerto Rico	23	13	1
All sectors	368	212	

* For all products tagged by NWS as Key Performance Parameter (KPP), Critical, Supplemental High, or Suppl. Low

JPSS satellites: daily CONUS overpasses

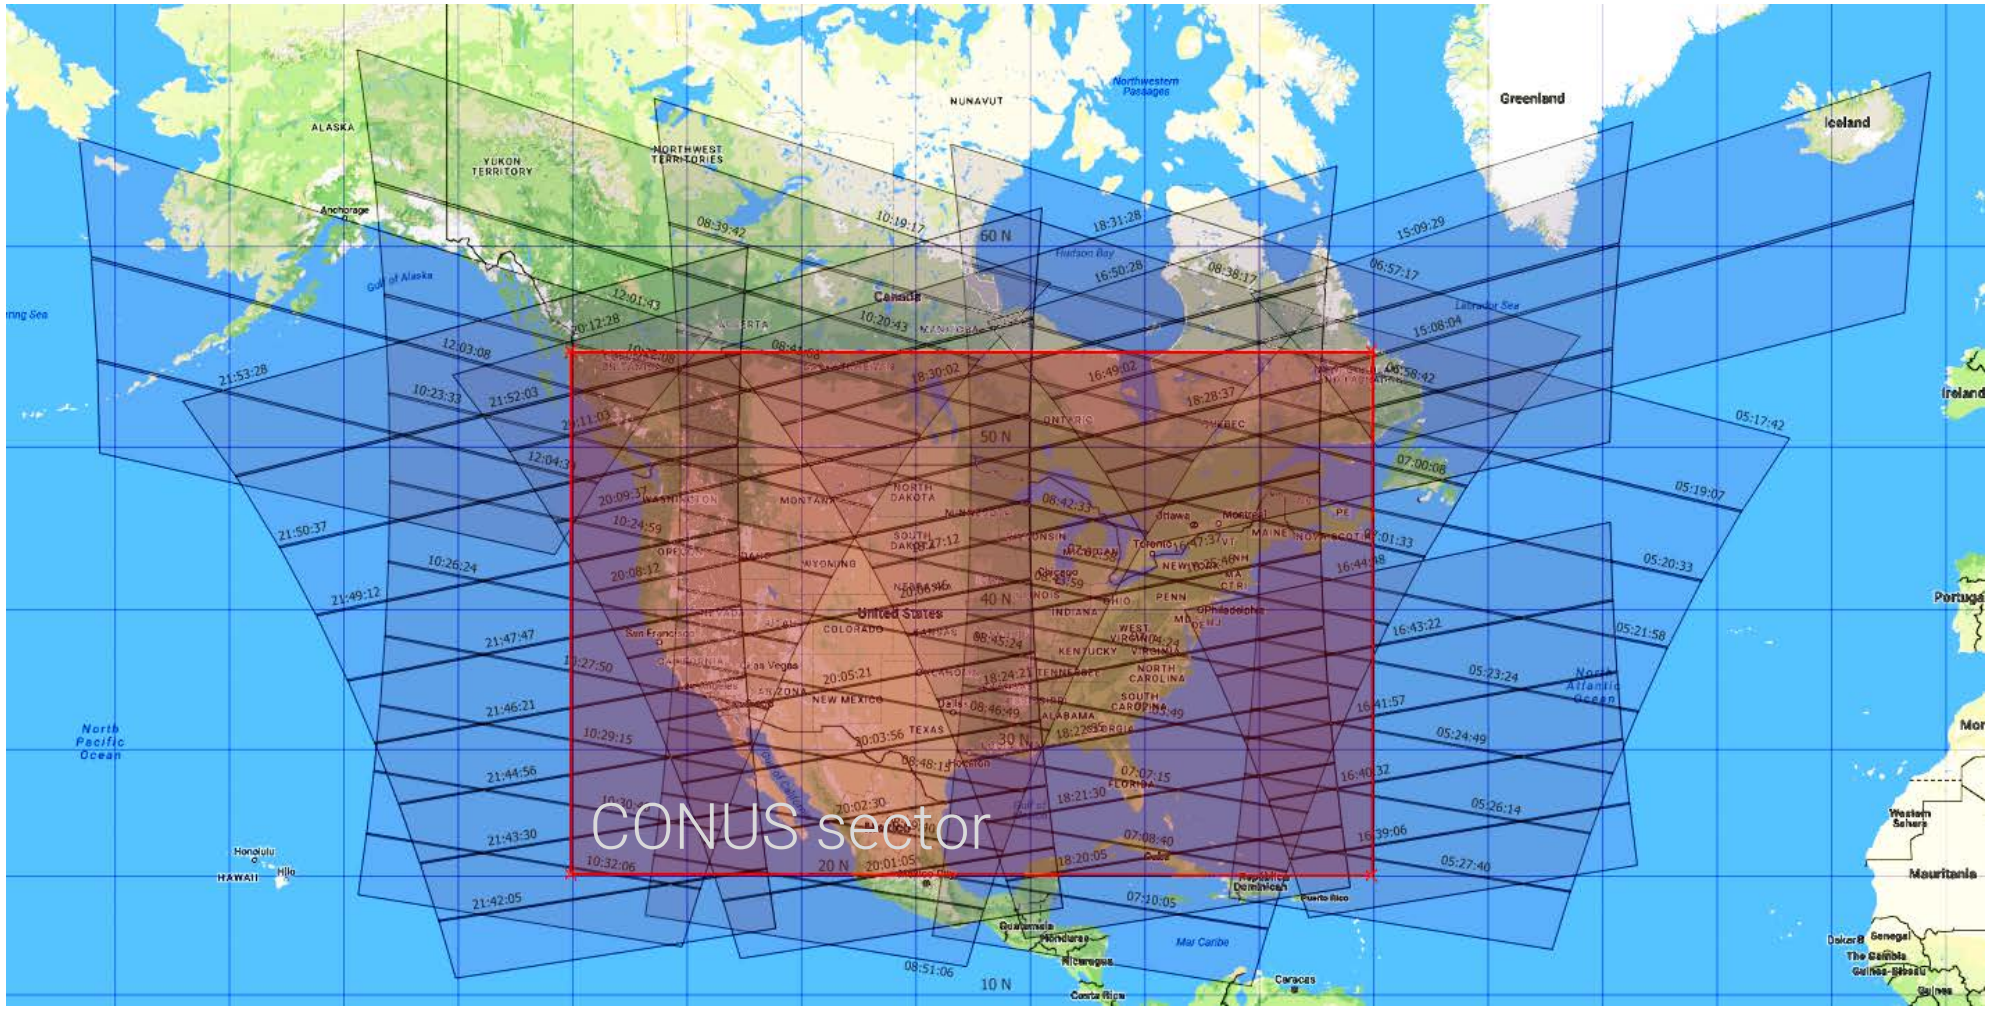

Each JPSS satellite sees (some of) CONUS about 10 times per day (78-80 VIIRS granules / day)

Near-real-time users: one overpass at a time

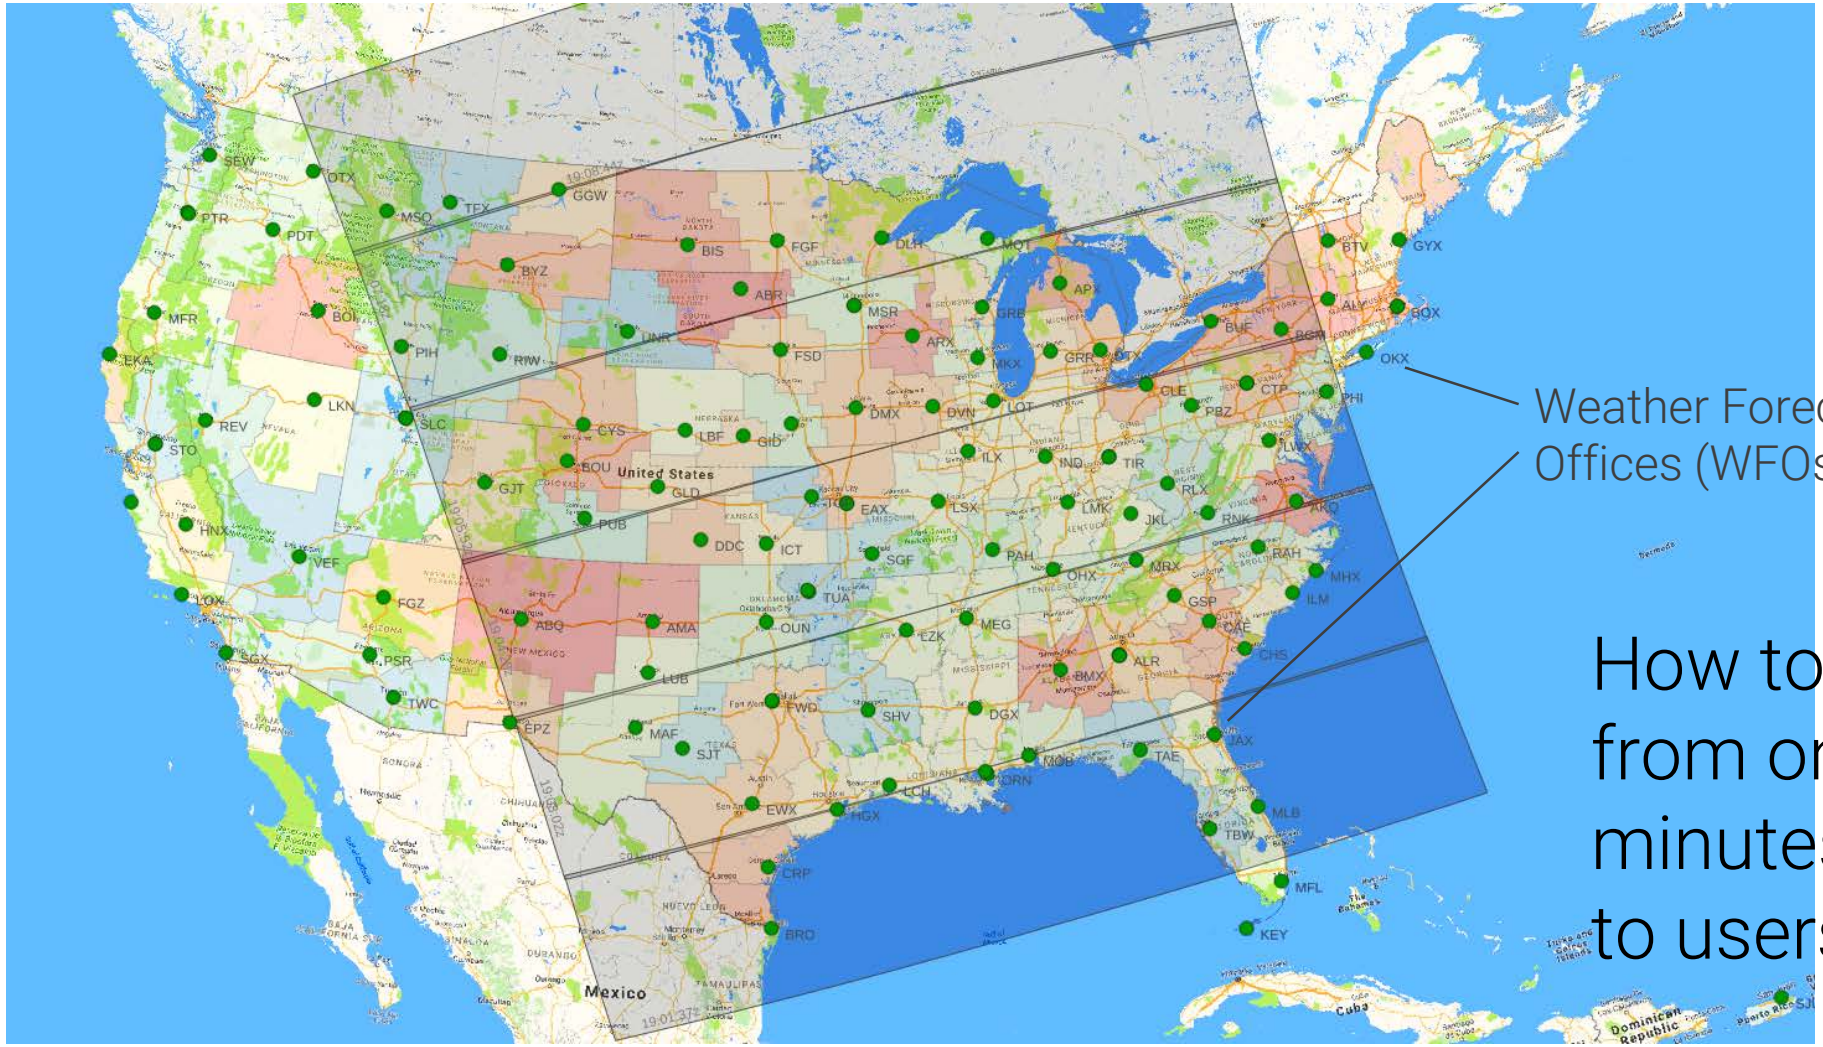

How to get data products from one overpass (~7 minutes of observations) to users quickly enough?

JPSS Enterprise Products: File size and structure

VIIRS Cloud Composition
(33 MB)

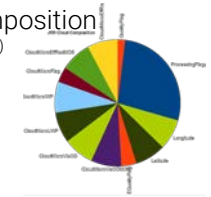

VIIRS Cloud Phase
(7-11 MB)

VIIRS Cloud Height
(48-96 MB)

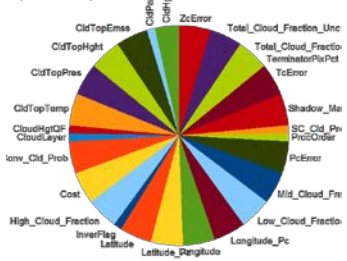

VIIRS Cloud Mask
(14-22 MB)

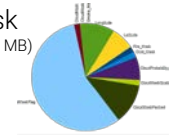

VIIRS Cloud Base
(16-33 MB)

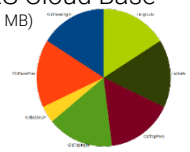

VIIRS I5 Imagery
(thinned for NWS:
14-19MB)

Latitude, Longitude

VIIRS NCC Imagery
(thinned for NWS:
1.4-1.9 MB)

VIIRS Active Fires
(0.05-0.5MB)

CrIS / ATMS NUCAPS
(3 MB)

ATMS MiRS Imagery
(7 MB)

CrIS / ATMS NUCAPS
(thinned for NWS)
(0.2 MB)

VIIRS Aerosol Detection
(7-17 MB)

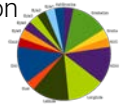

VIIRS Volcanic Ash
(40-110 MB)

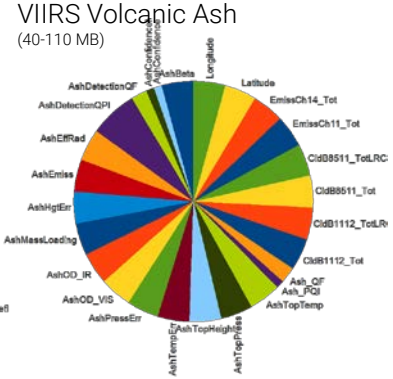

ATMS MiRS soundings
(49 MB)

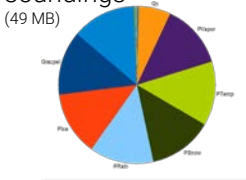

VIIRS Aerosol Optical Depth
(9-150 MB)

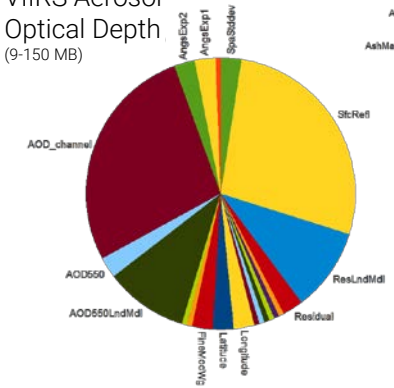

VIIRS Snow Cover
(20-36 MB)

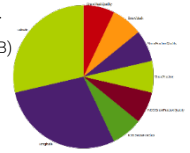

VIIRS Ice Concentration
(8-17 MB)

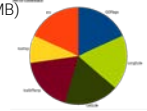

VIIRS Ice Age & Thickness
(7-15 MB)

VIIRS Polar Winds
(8 MB)

VIIRS Surface Reflectance
(208 MB)

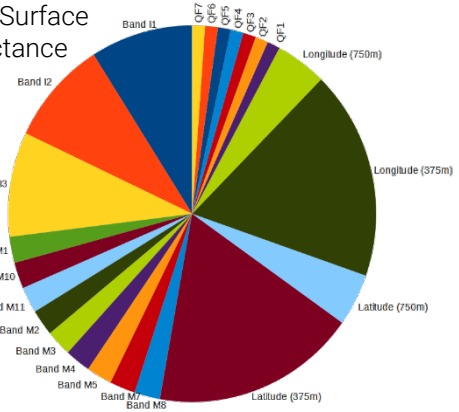

ACSPO VIIRS Sea Surface Temp. L2
(175 MB)

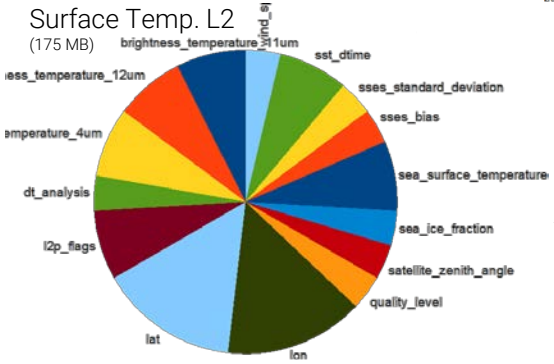

ACSPO VIIRS SST L3
(18 MB)

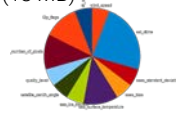

From each JPSS overpass, how many of these products can we afford to send to forecast offices?

Which products can we afford to get from each overpass?

Product volumes (MB) for a 7-minute JPSS overpass; and
Cumulative bandwidth needs for timely (10-minute) delivery (==> "bandwidth budget").

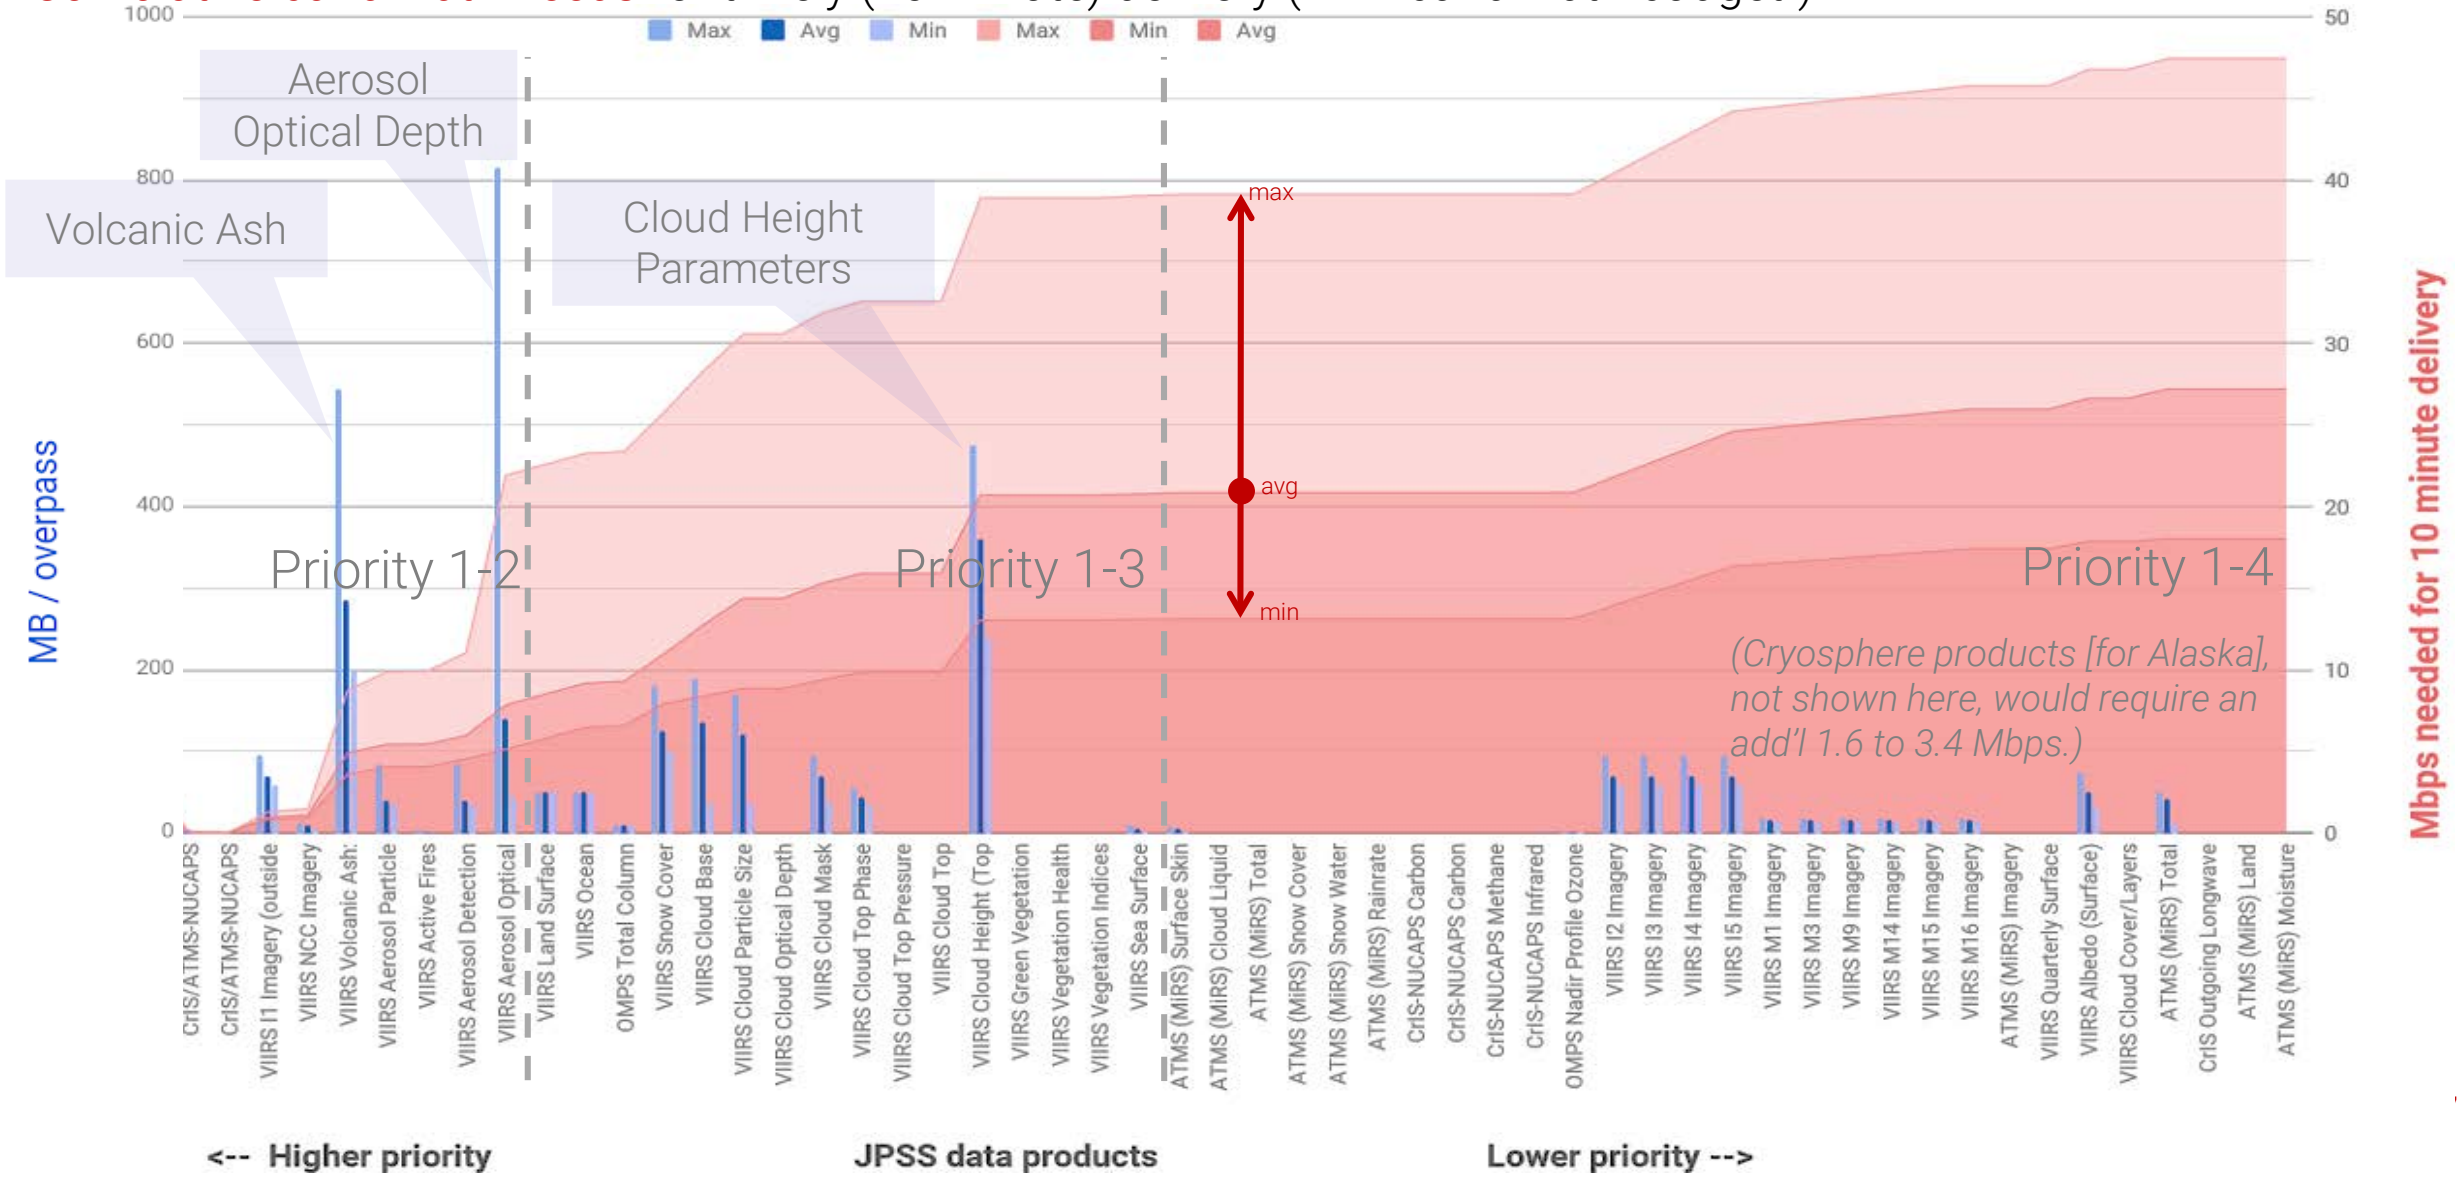

Getting JPSS data products to AWIPS

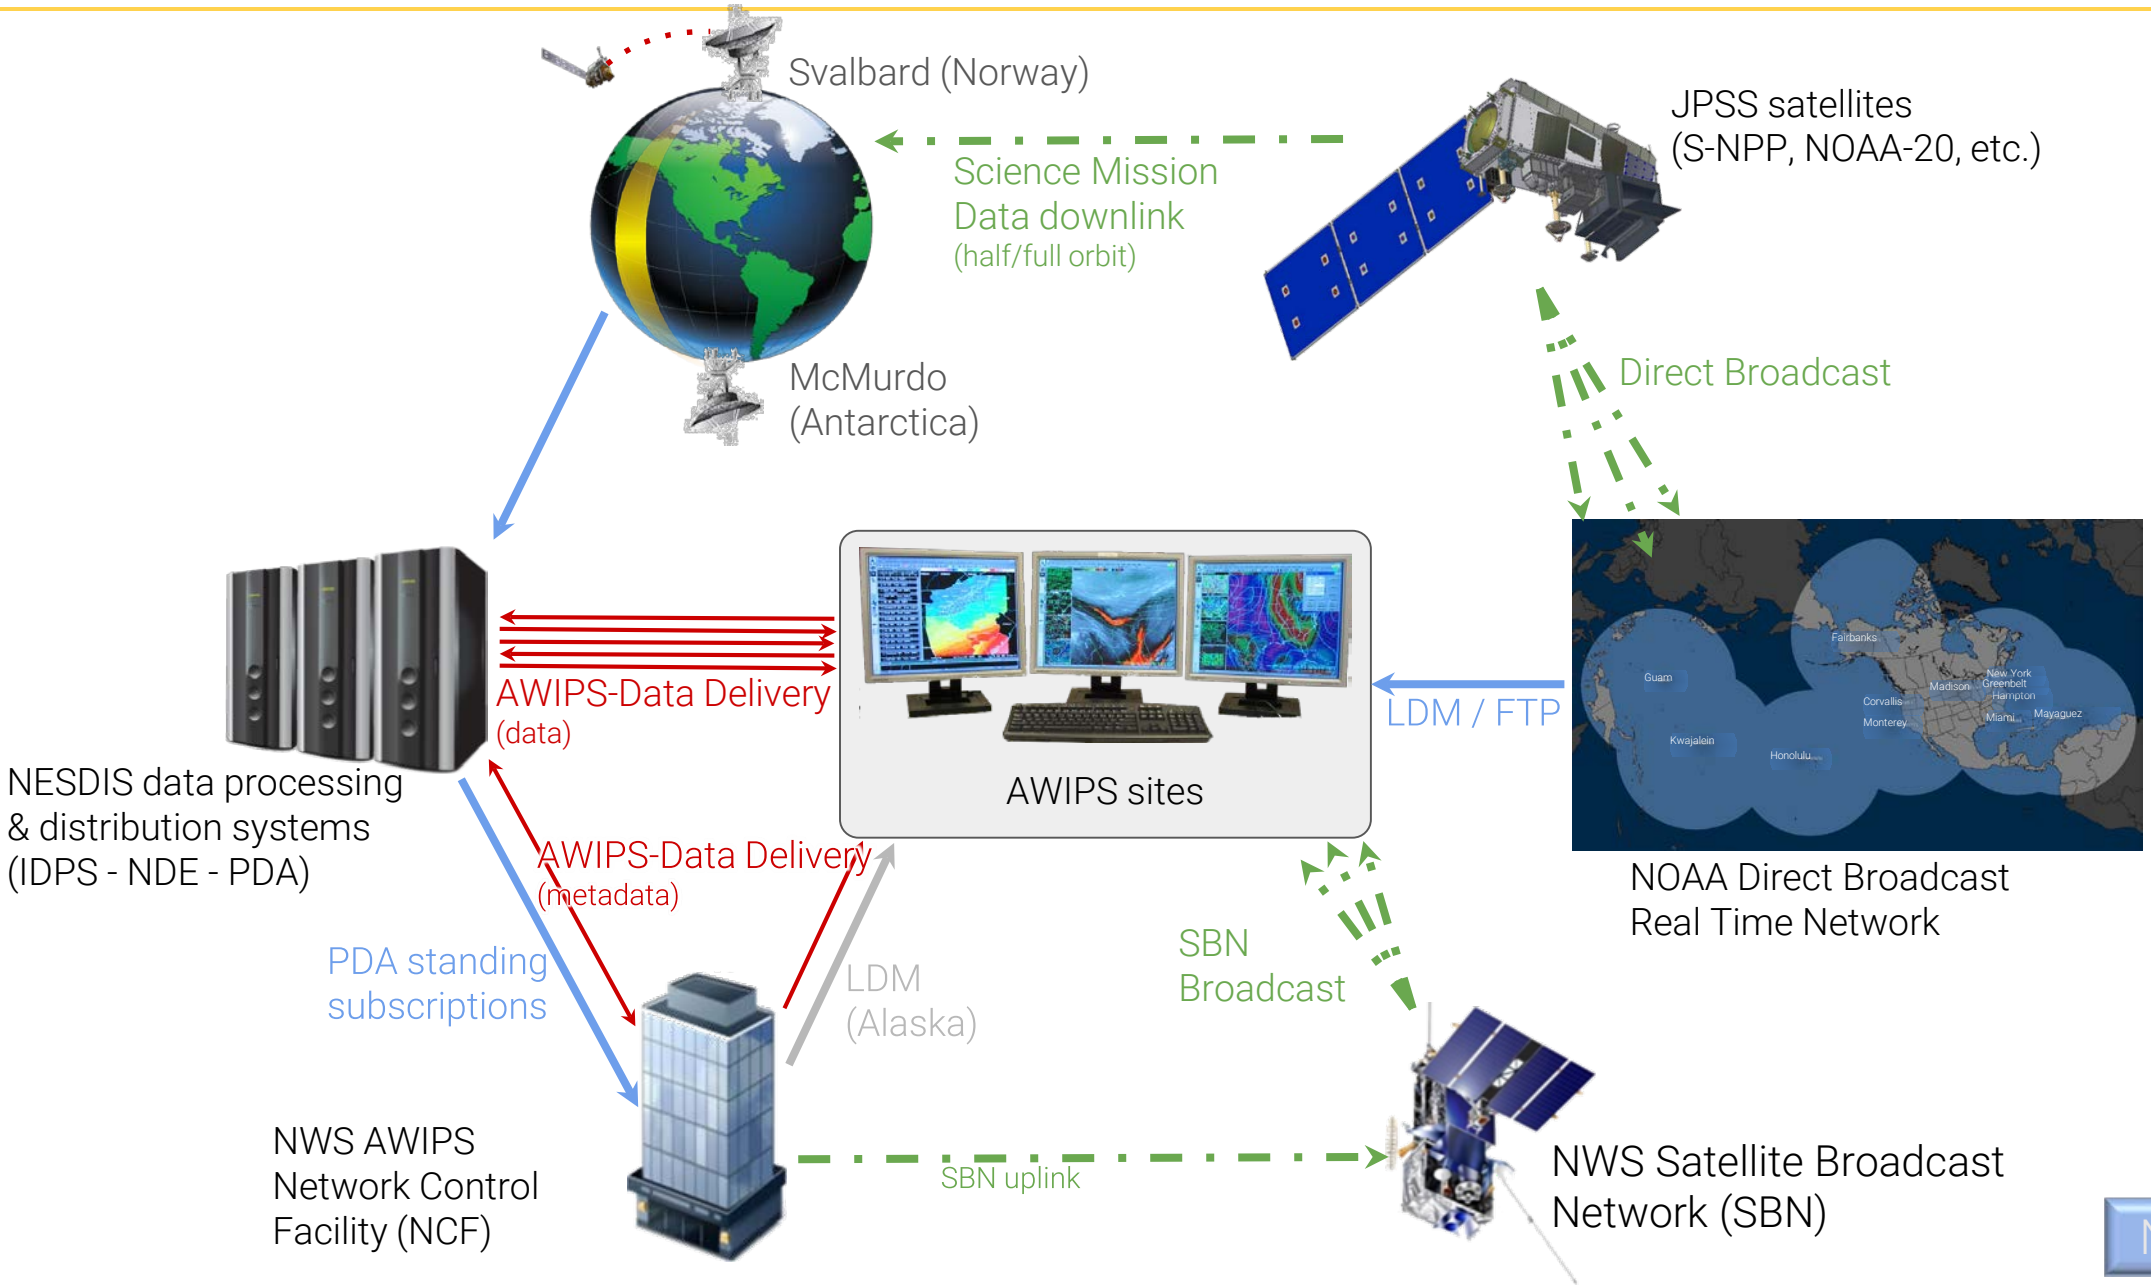

NWS Satellite Broadcast Network (SBN)

- "Mainstream" path into AWIPS forecaster workstations
- 69+ Mbps total; 6 Mbps available for polar satellite data
- Everyone receives the same thing – e.g., for JPSS:
  - VIIRS Near-Constant Contrast (NCC) Day-Night Band imagery for Alaska, Pacific, CONUS, Puerto Rico
  - NUCAPS soundings for Americas + Pacific & East Asia
  - VIIRS bands I1, I4, I5 imagery for Alaska region

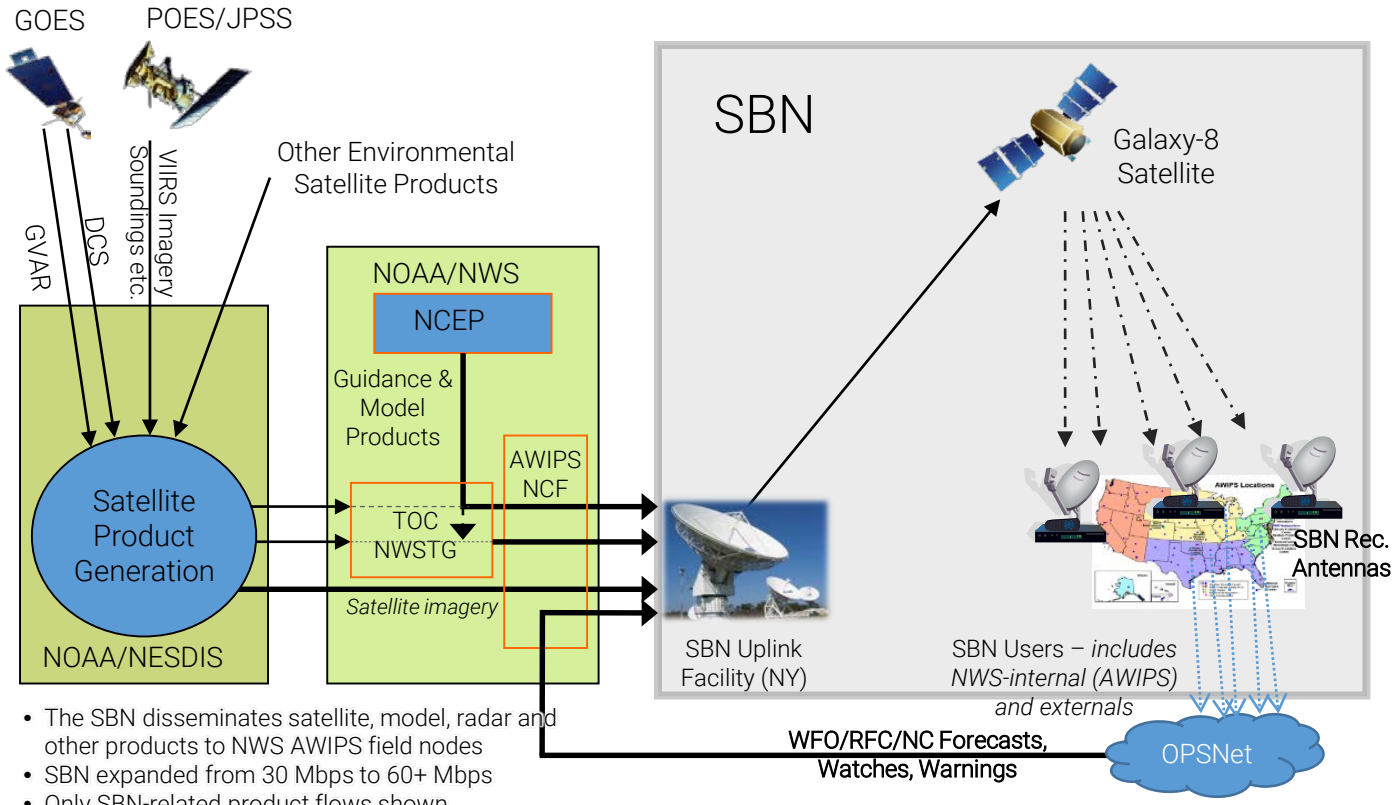

- The SBN disseminates satellite, model, radar and other products to NWS AWIPS field nodes
- SBN expanded from 30 Mbps to 60+ Mbps
- Only SBN-related product flows shown

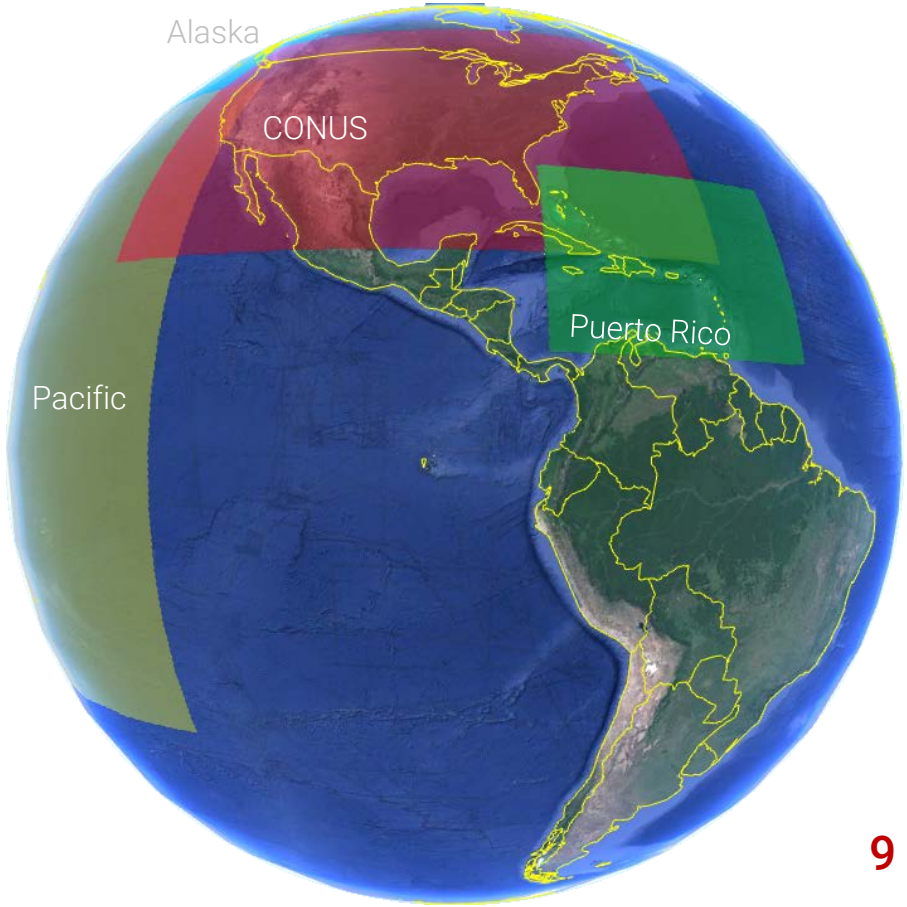

SBN is nearly full already

Global Science & Technology, Inc.

SBN data substreams (Mbps / 10-minute aggregation)

Stacked area chart

Mbps needed for 10-minute delivery of model and satellite data over a 24-hour period

(Not shown: NEXRAD radar data)

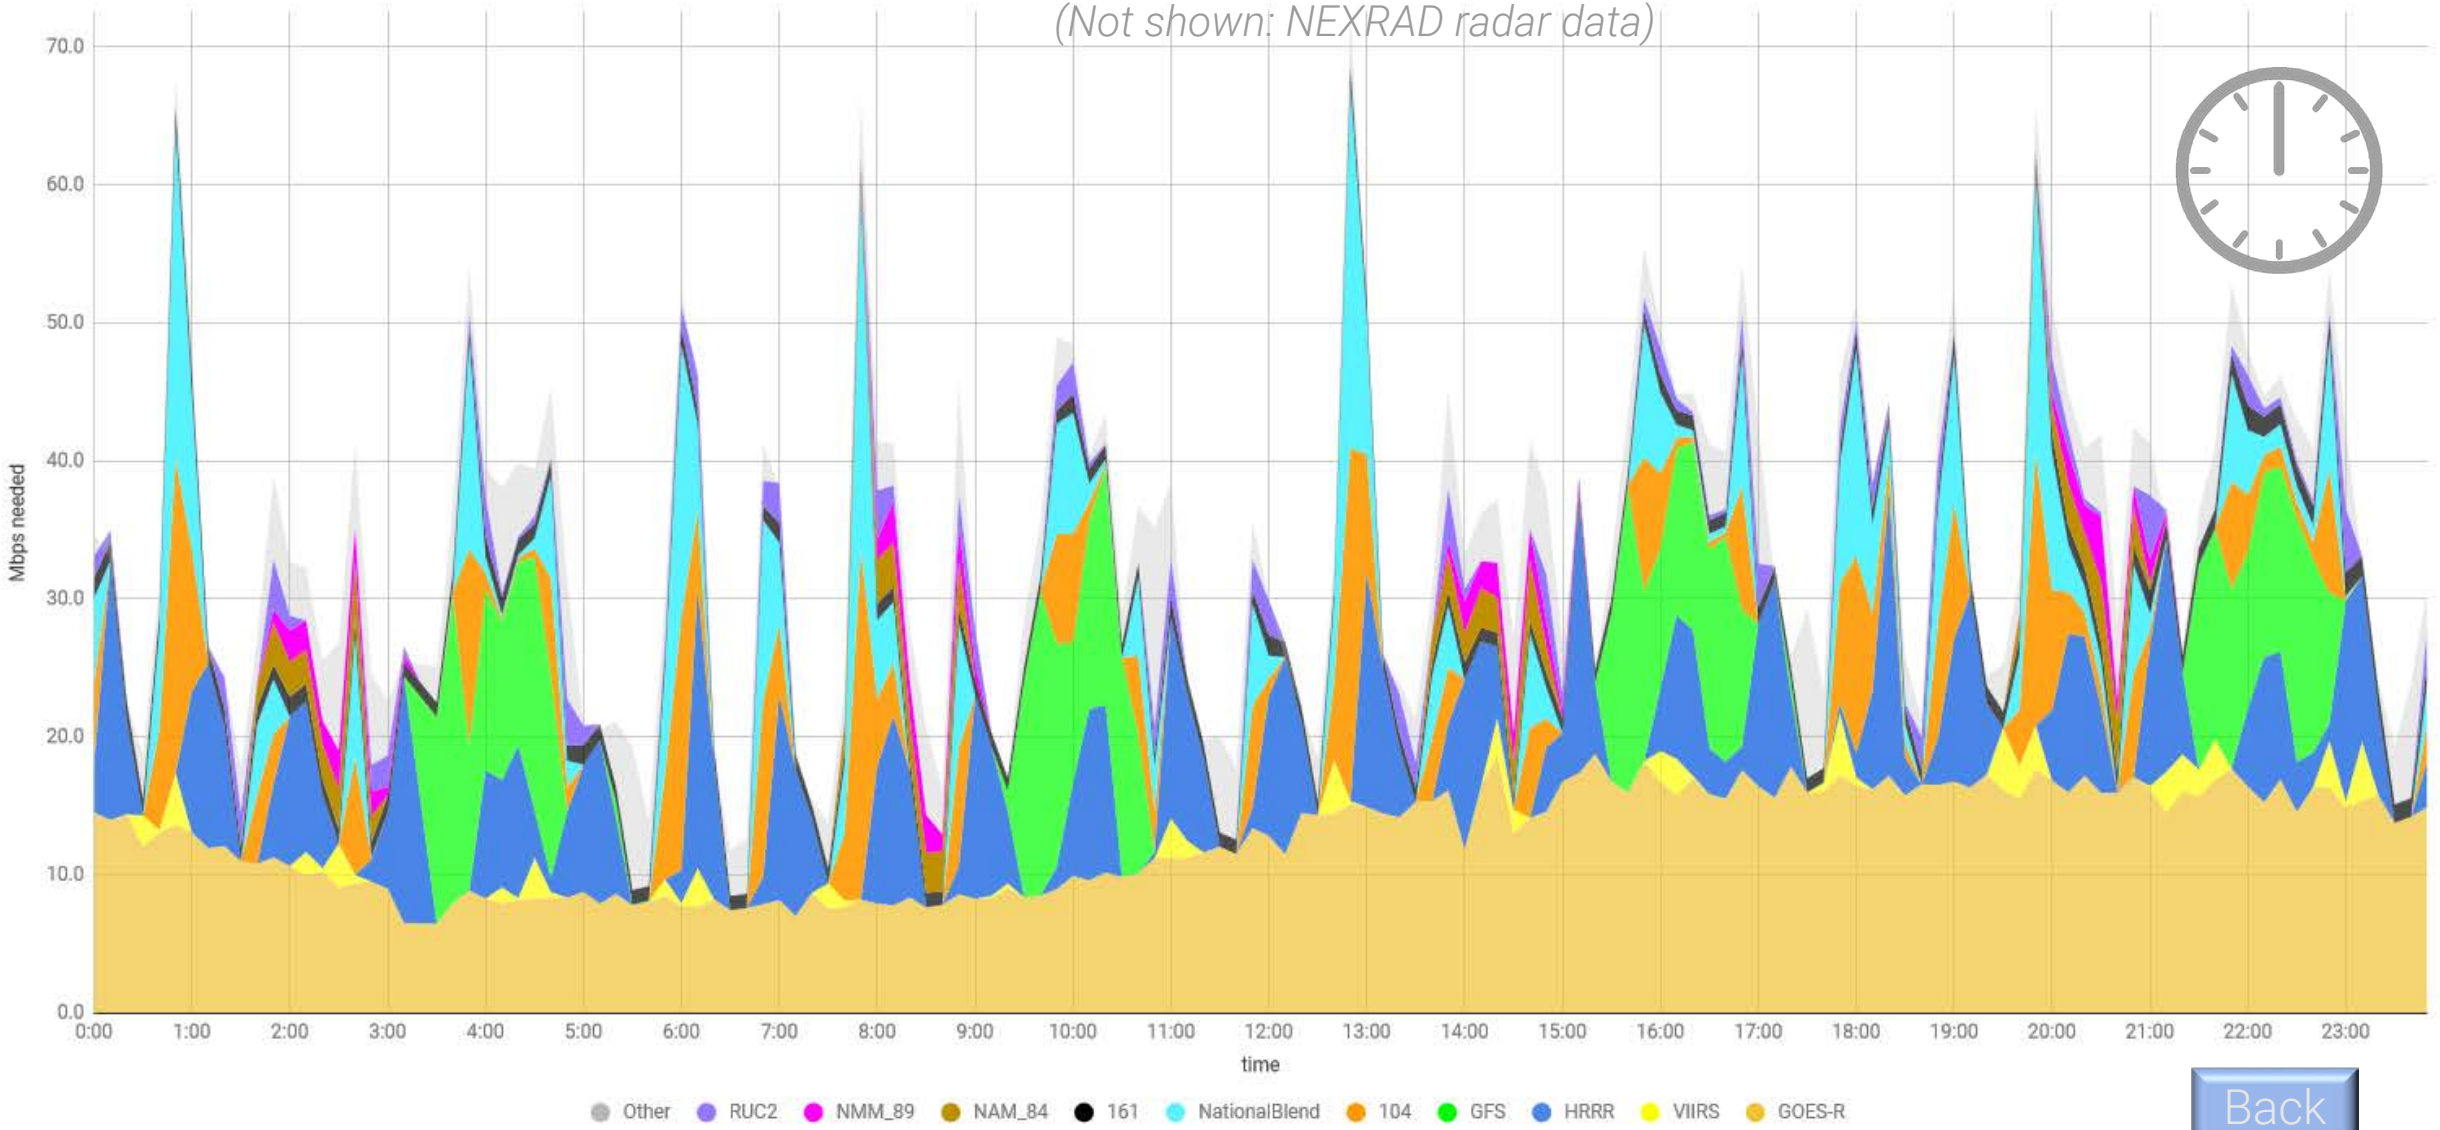

Back

AWIPS Data Delivery (AWIPS-DD)

On-demand services, connecting AWIPS to ESPDS Product Distribution and Access (PDA) and others

- Deliver only what users request
- Less need to pinpoint end-user needs
  - Less predictable usage patterns

AWIPS-DD & PDA use an asynchronous protocol

- Loosely based on the Open Geospatial Consortium Web Coverage Service (WCS) standard
- A fairly complex protocol, but has been shown to work with GOES-R data; AWIPS team is now adapting it to polar data
- Fetches discrete JPSS product files (no on-demand "tailoring")

Data, metadata, requests, and responses travel via OneNWSnet TCP/IP fiber-optic network

- Now 100Mbps => a workable "bandwidth budget"

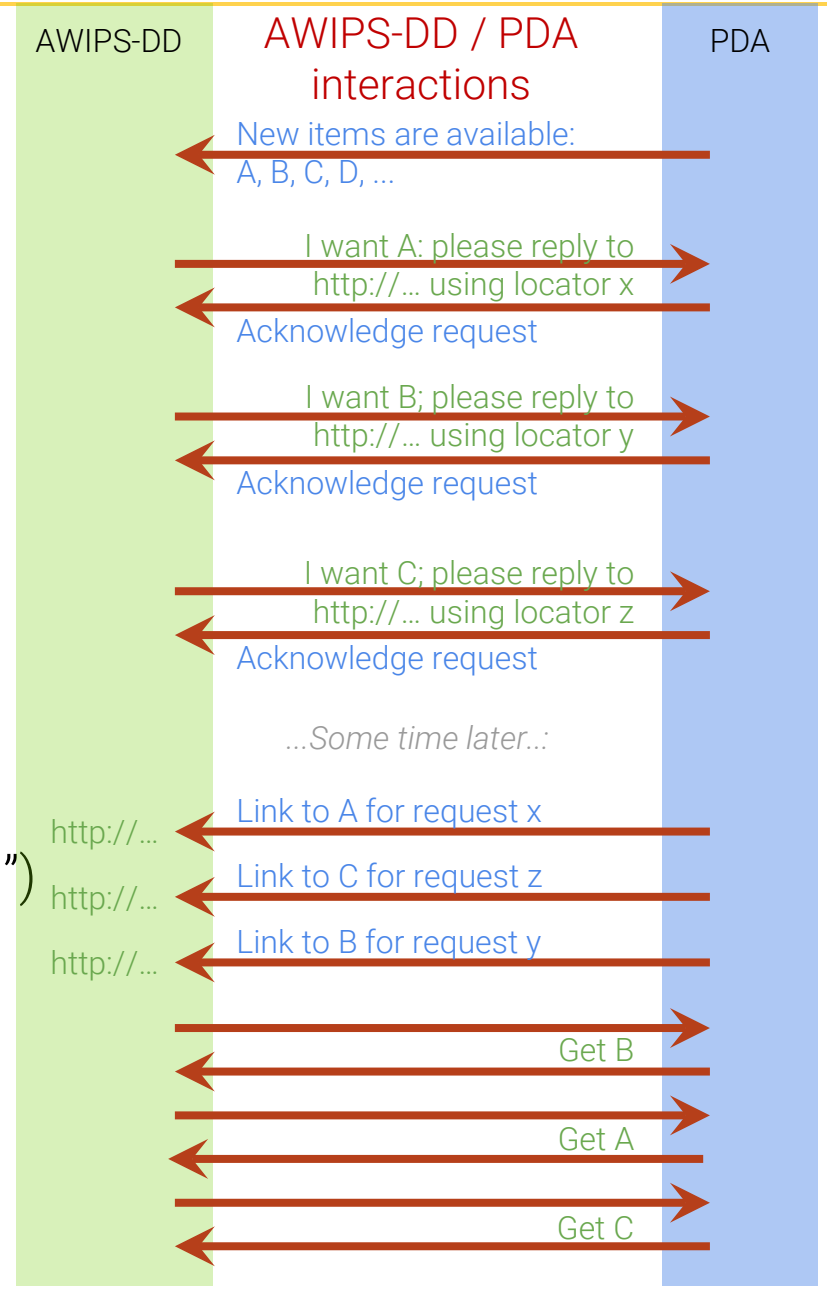

Back

NOAA Direct Broadcast Real Time Network

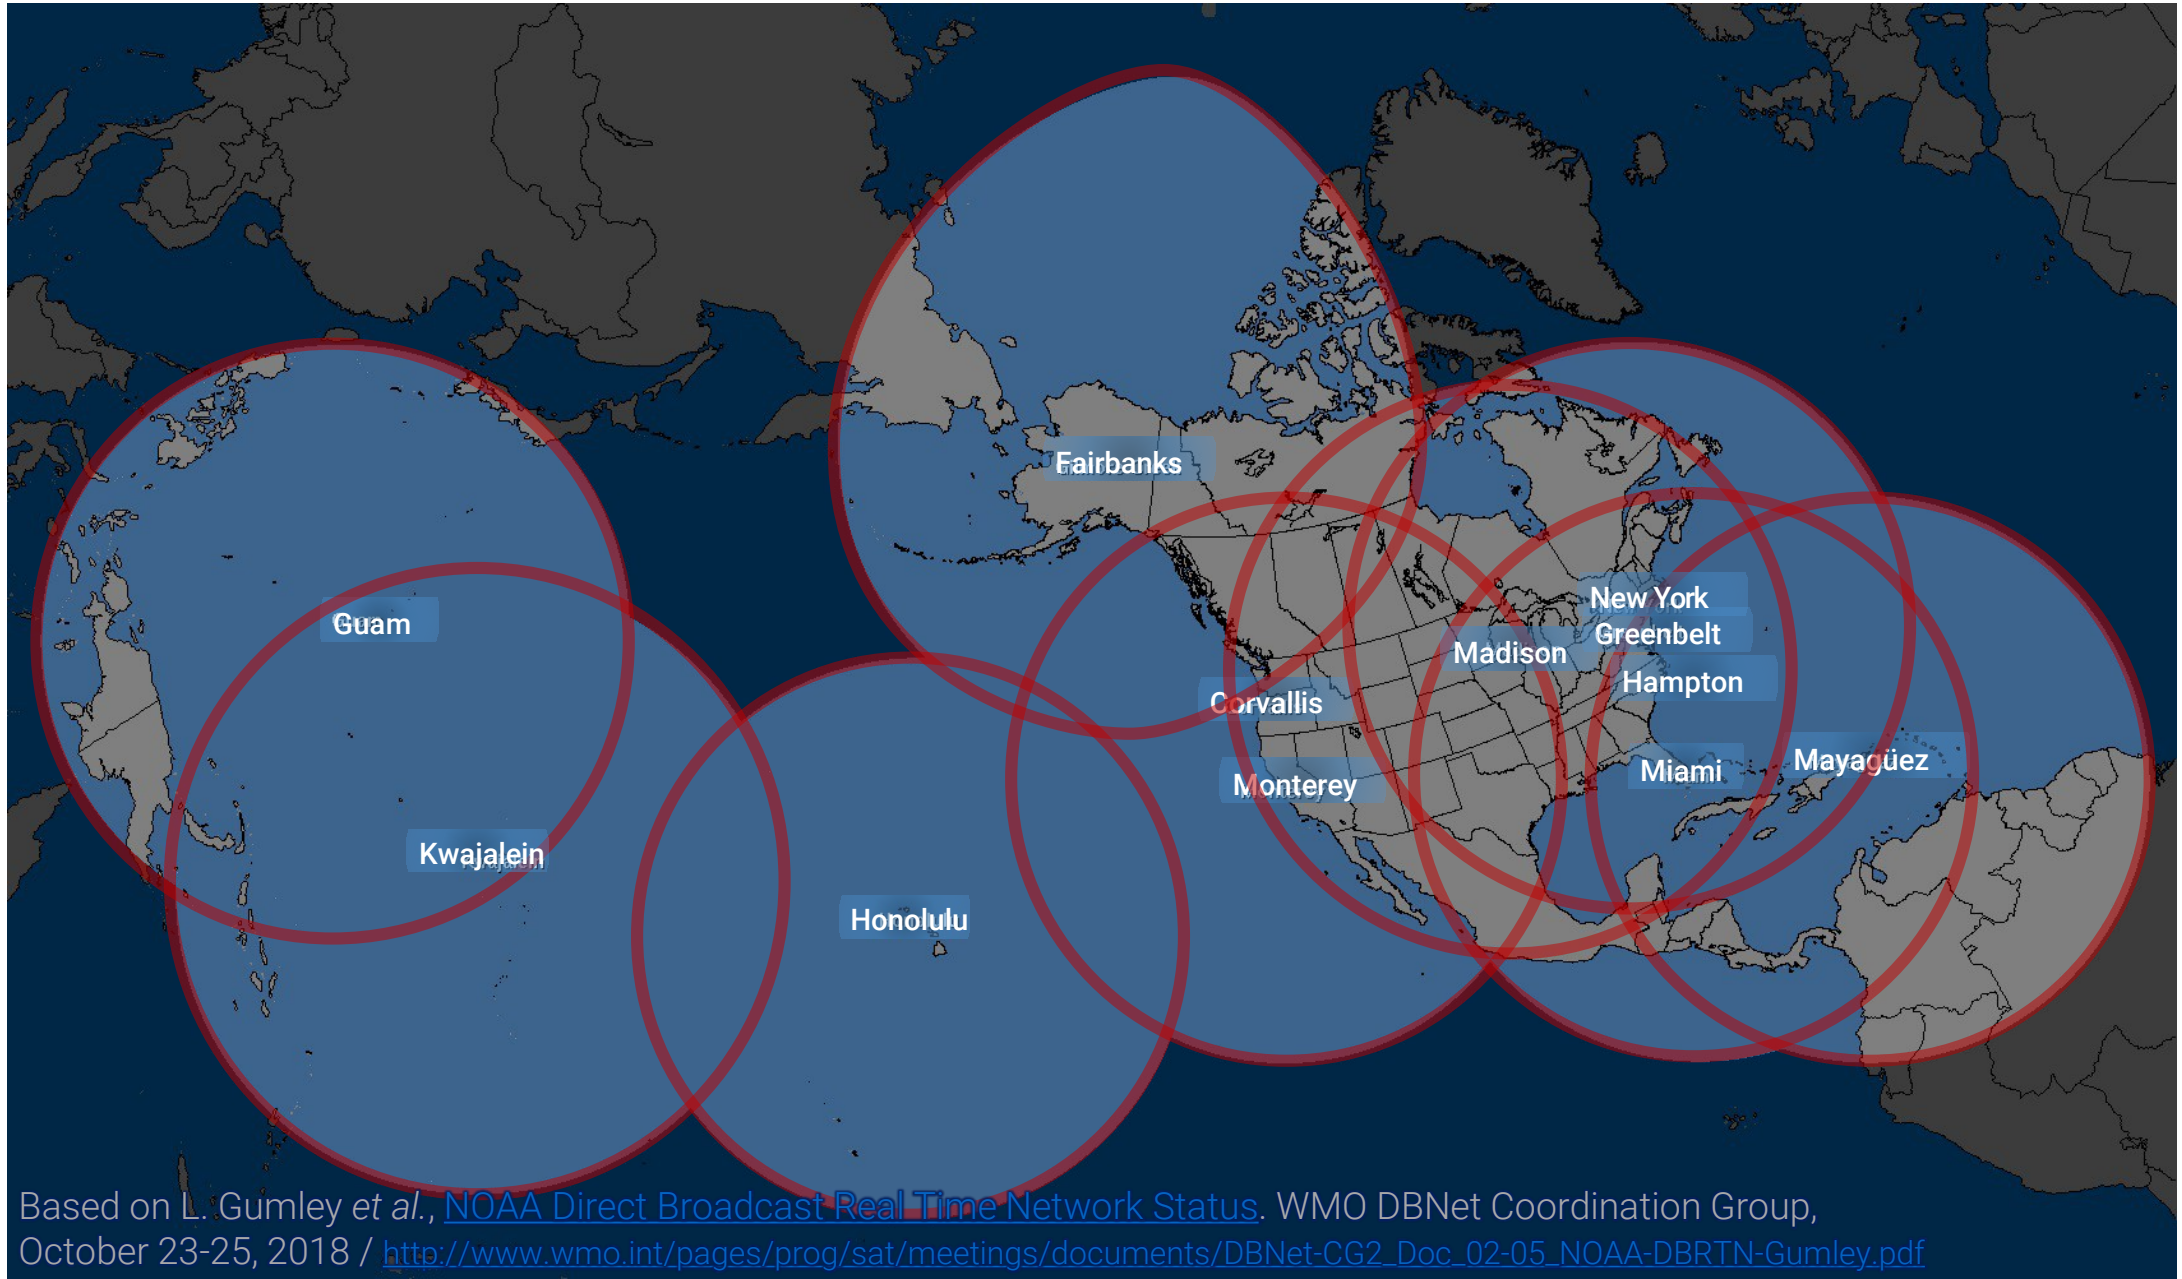

Based on L. Gumley *et al.*, NOAA Direct Broadcast Real Time Network Status. WMO DBNet Coordination Group, October 23-25, 2018 / http://www.wmo.int/pages/prog/sat/meetings/documents/DBNet-CG2_Doc_02-05_NOAA-DBRTN-Gumley.pdf

Direct Broadcast: Game-changing latency

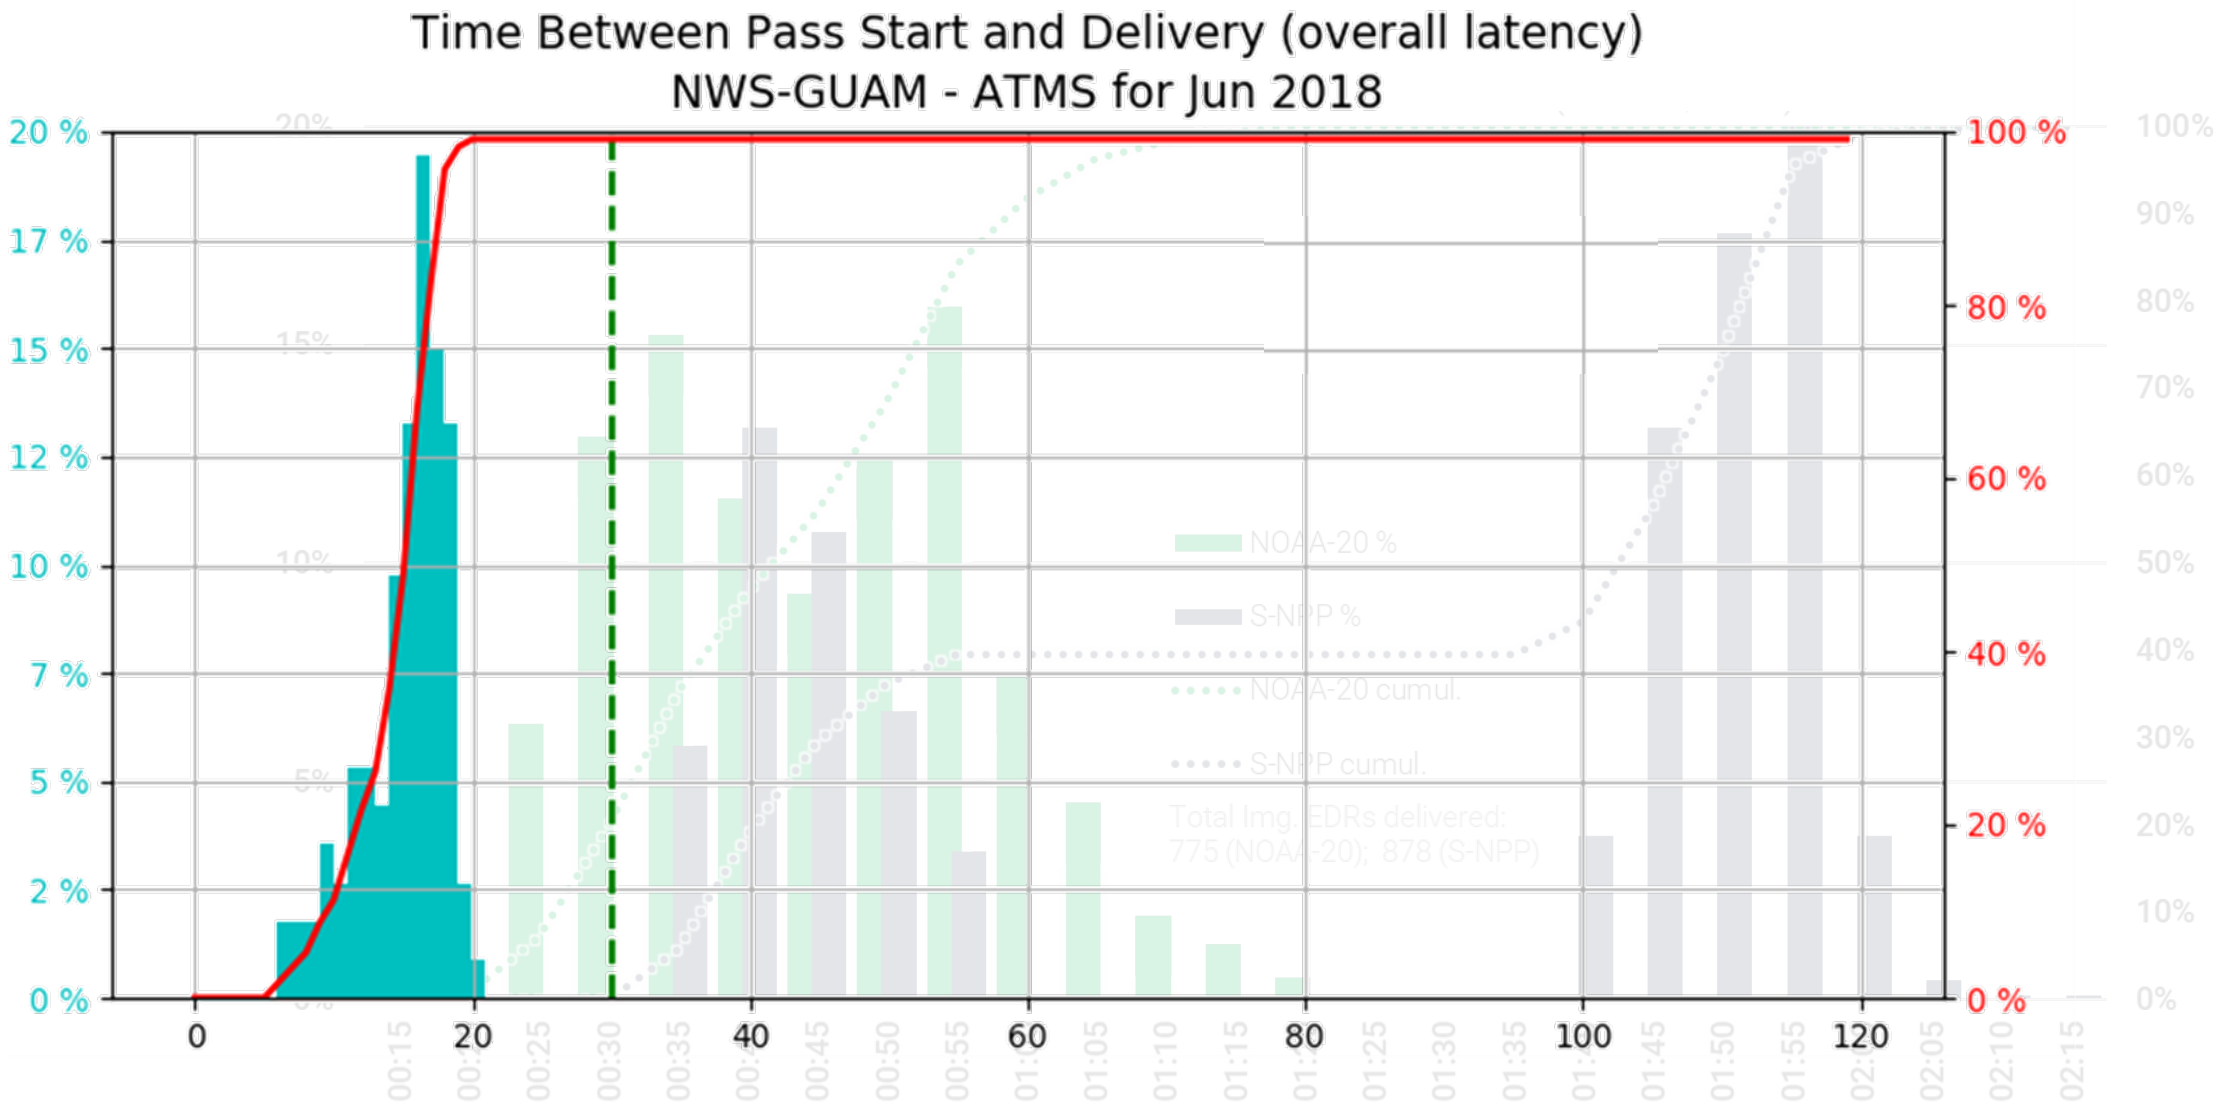

Based on L. Gumley *et al.*, NOAA Direct Broadcast Real Time Network Status. WMO DBNet Coordination Group, October 23-25, 2018 / http://www.wmo.int/pages/prog/sat/meetings/documents/DBNet-CG2_Doc_02-05_NOAA-DBRTN-Gumley.pdf

For comparison: latencies via Svalbard / McMurdo

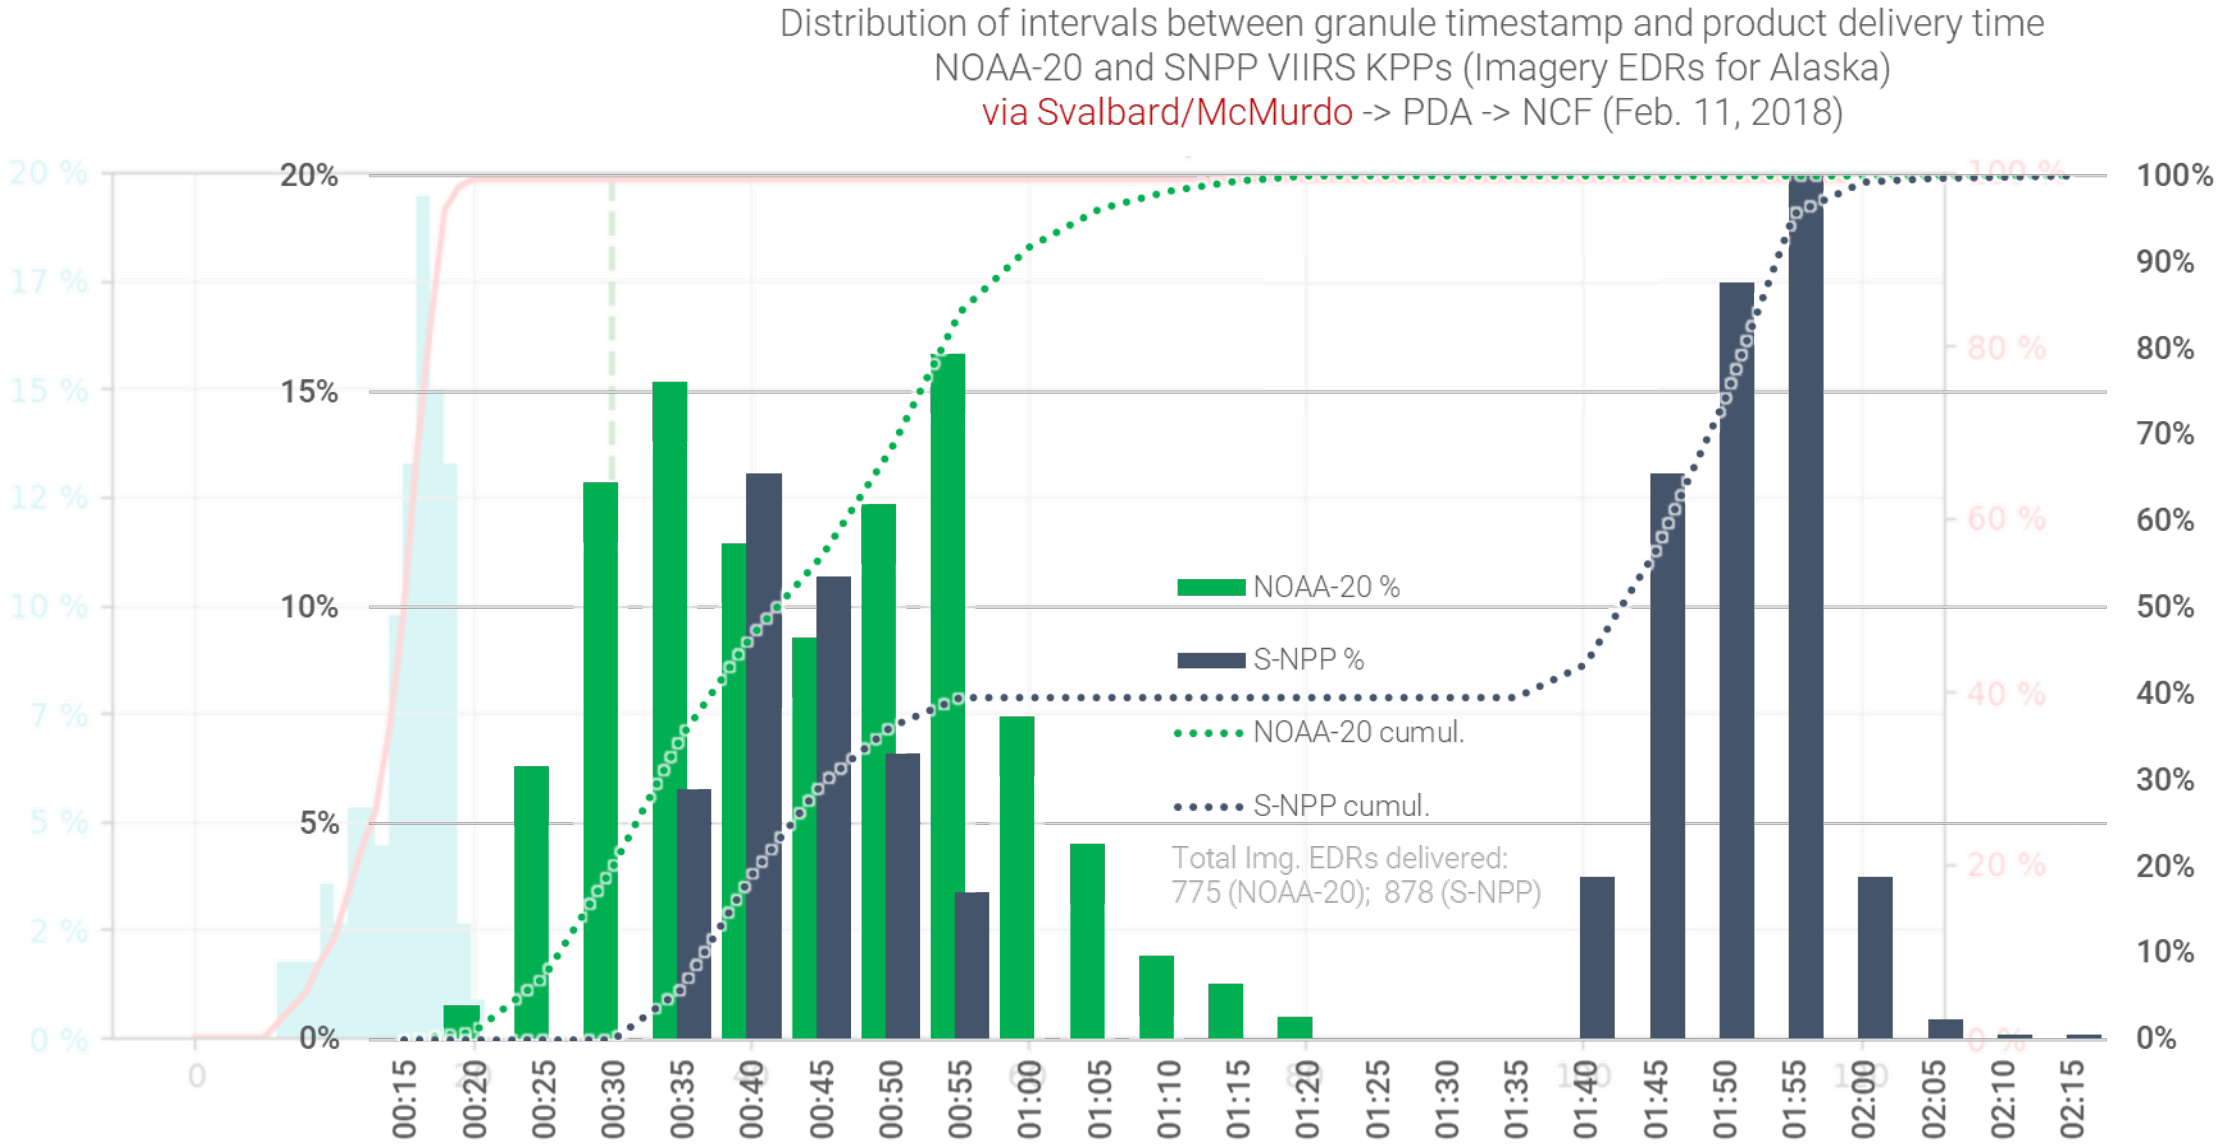

Ensuring scalability

- Current technologies (esp. AWIPS-DD / PDA via OneNWSnet's 100 Mbps) will give forecasters timely access to the data products they need
  ... For now. Handling many more users and more data will require
- Reducing unnecessary data movement – *e.g.*,
  - Produce and disseminate smaller ("thinned") versions of products
  - Subset data on demand by location, time, or parameter
- Limiting server loads – *e.g.*,
  - Tiered Content Distribution Network
  - Conditional (or on-demand) processing
- Emphasizing simplicity, fault tolerance, interoperability

Reducing data movement:

Thinned products for NWS forecasters

VIIRS Volcanic Ash
(6-120 MB)

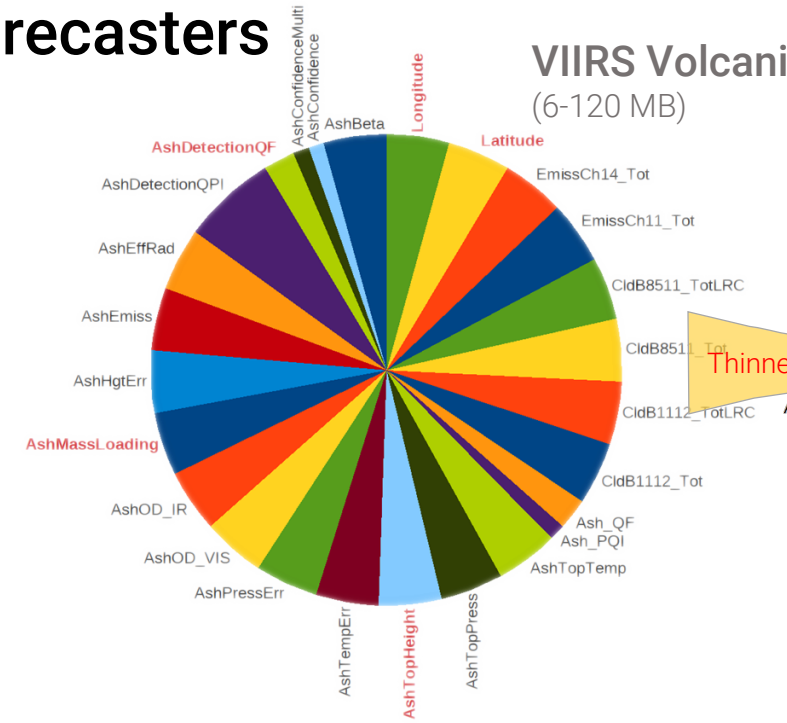

Thinned product

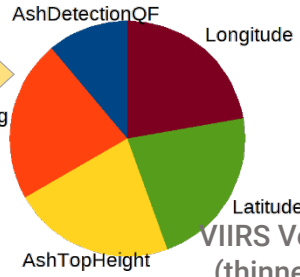

VIIRS Volcanic Ash
(thinned for NWS)
(6-24 MB)

VIIRS Aerosol Optical Depth
(9-171 MB)

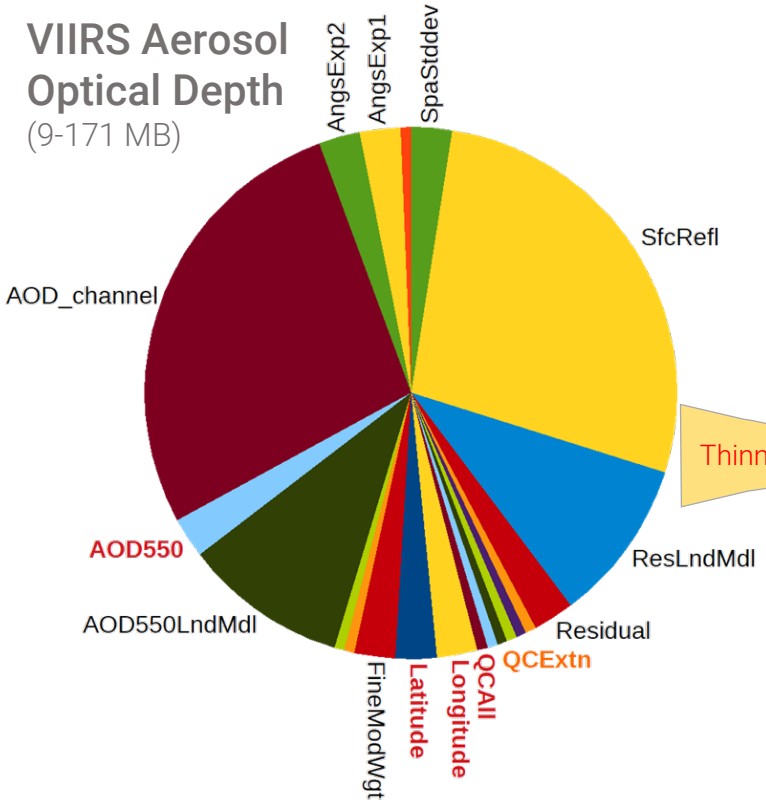

Thinned product

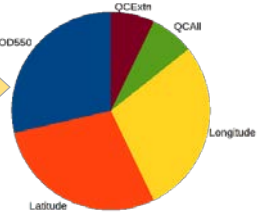

VIIRS Aerosol Optical Depth
(thinned for NWS)
(6-14 MB)

CrIS / ATMS NUCAPS
(3 MB)

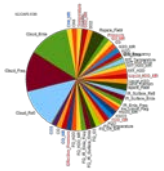

Thinned product

CrIS / ATMS NUCAPS
(thinned for NWS)
(0.2 MB)

Reducing data movement:

Streamlined, interoperable Web services

Full OGC Web Coverage Service (WCS) capability would include on-demand selection and subsetting by location, time, and field/parameter;

- ➢ Perhaps also Resampling, Aggregation, and Reprojection

This may further reduce data transfer volumes (4x ~ 100x)

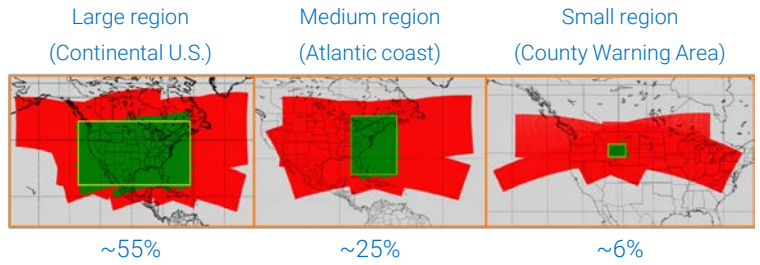

Use of actual OGC WCS protocol would also bring

- Simpler client-server interaction
- Possibility of COTS solutions
- Interoperable & reusable service

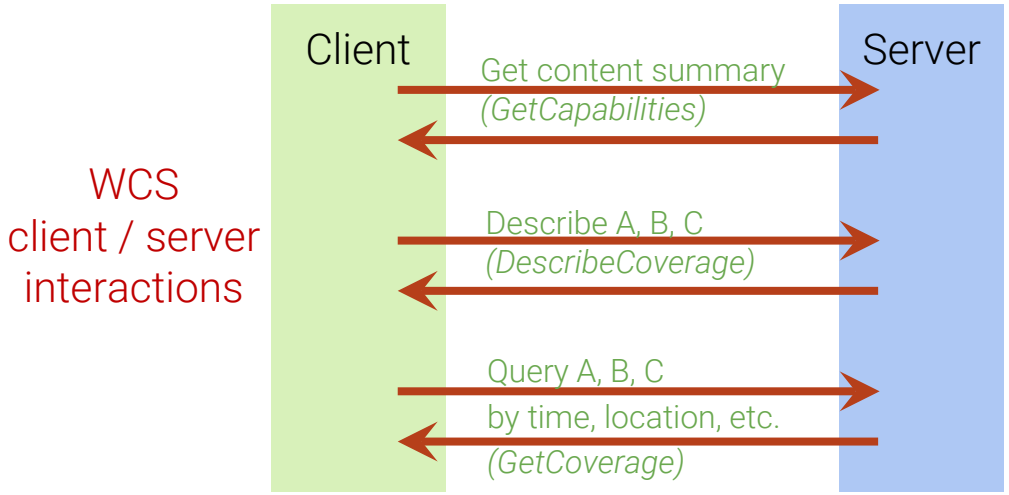

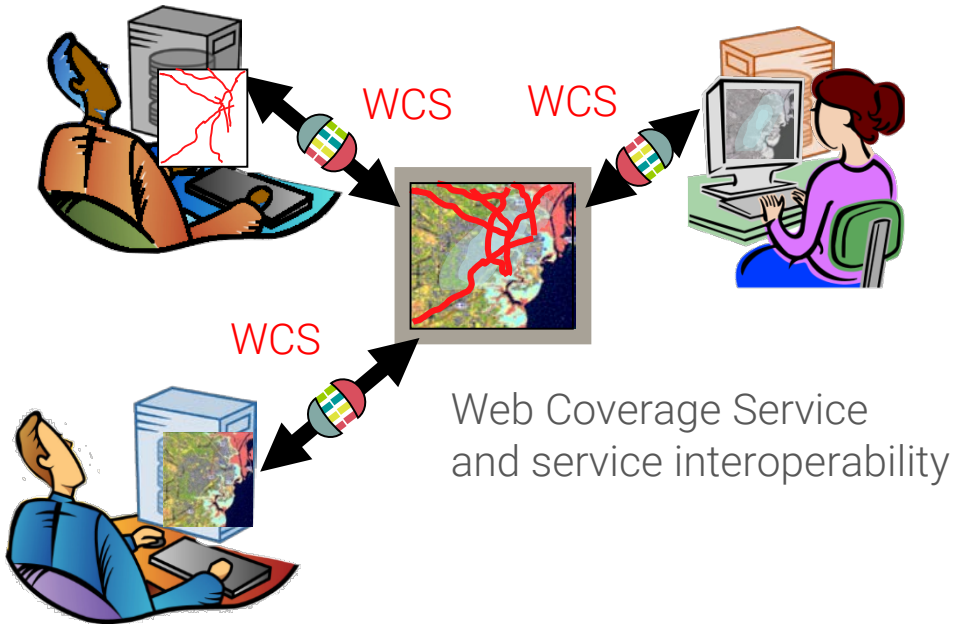

Limiting server loads:

Tiered Content Distribution Network

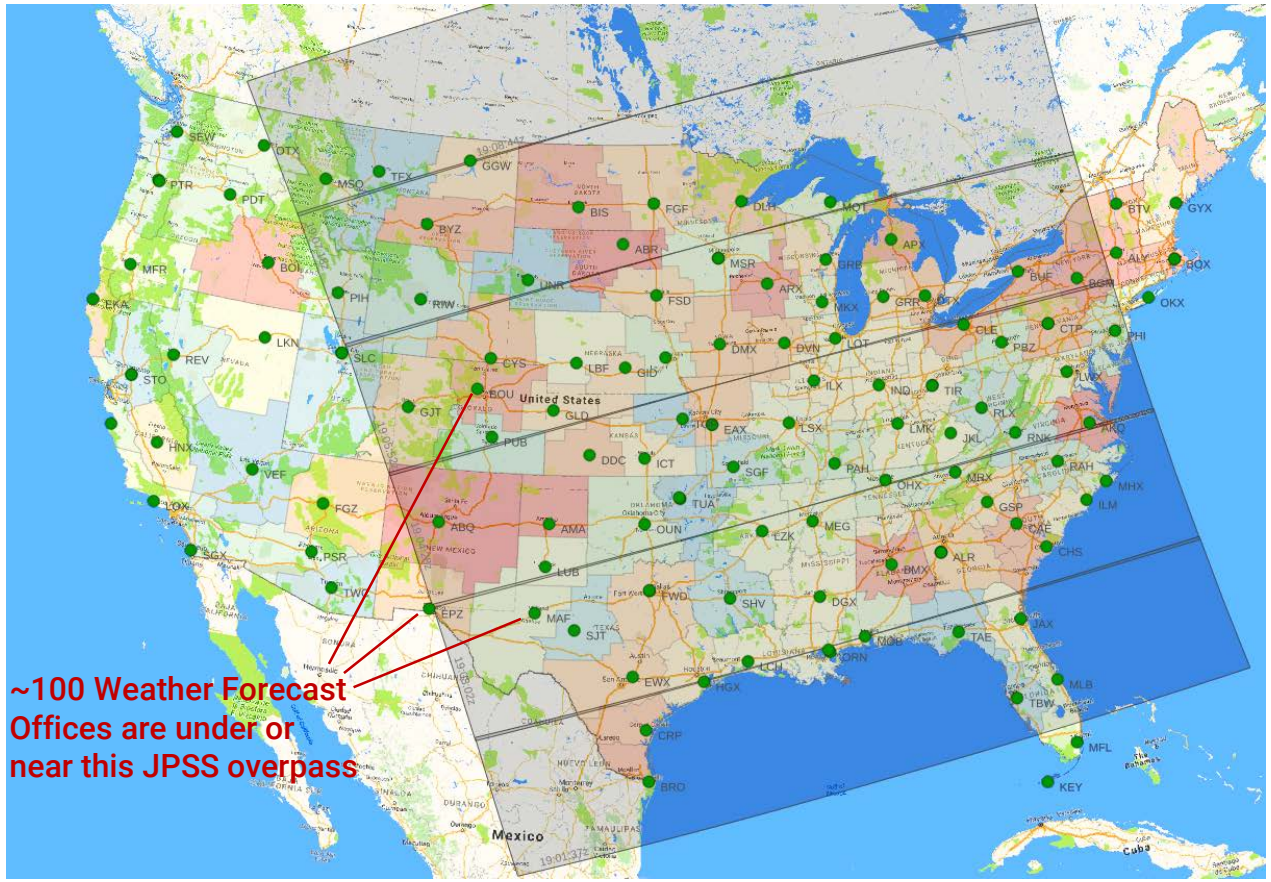

After each CONUS overpass of a JPSS satellite, up to ~100 sites may opt to fetch ~5-10 products from each of ~5 granules, within a short time.

- 5,000 concurrent file transfers!?

Instead of having all end users fetch products from a single site, maybe distribute to one or more tiers of "edge" servers – likely via Cloud Computing.

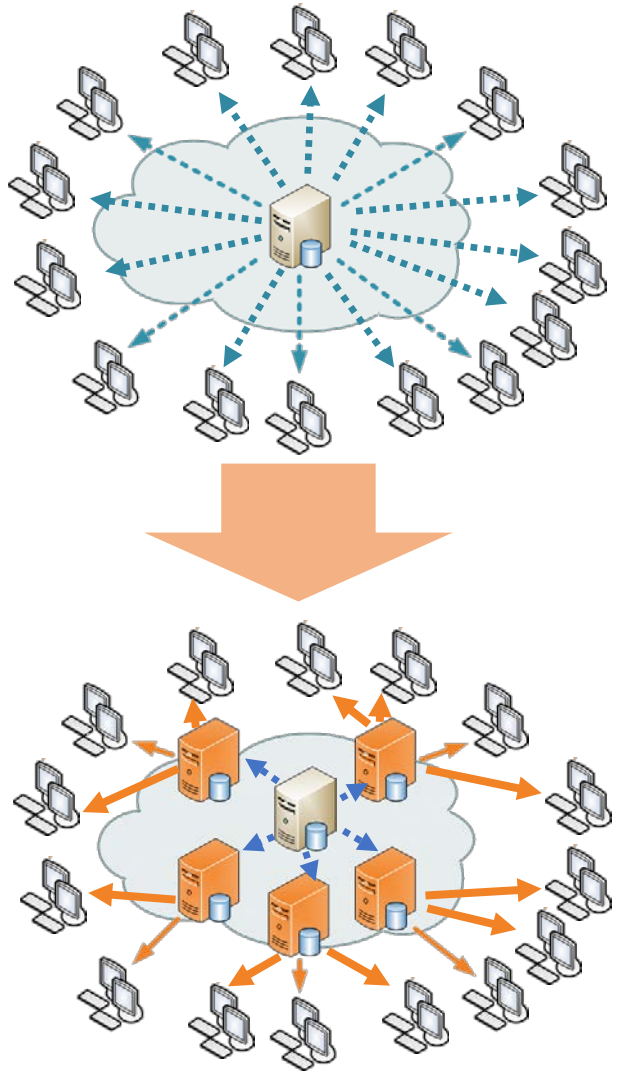

Limiting server loads:

Conditional / On-demand product generation

- Instead of running the full suite of algorithms on every data granule, maybe generate some less-frequently used products only under certain circumstances, or only when requested by users.
- Cloud Computing would allow rapid and temporary “scale-out” of processing resources when needed.

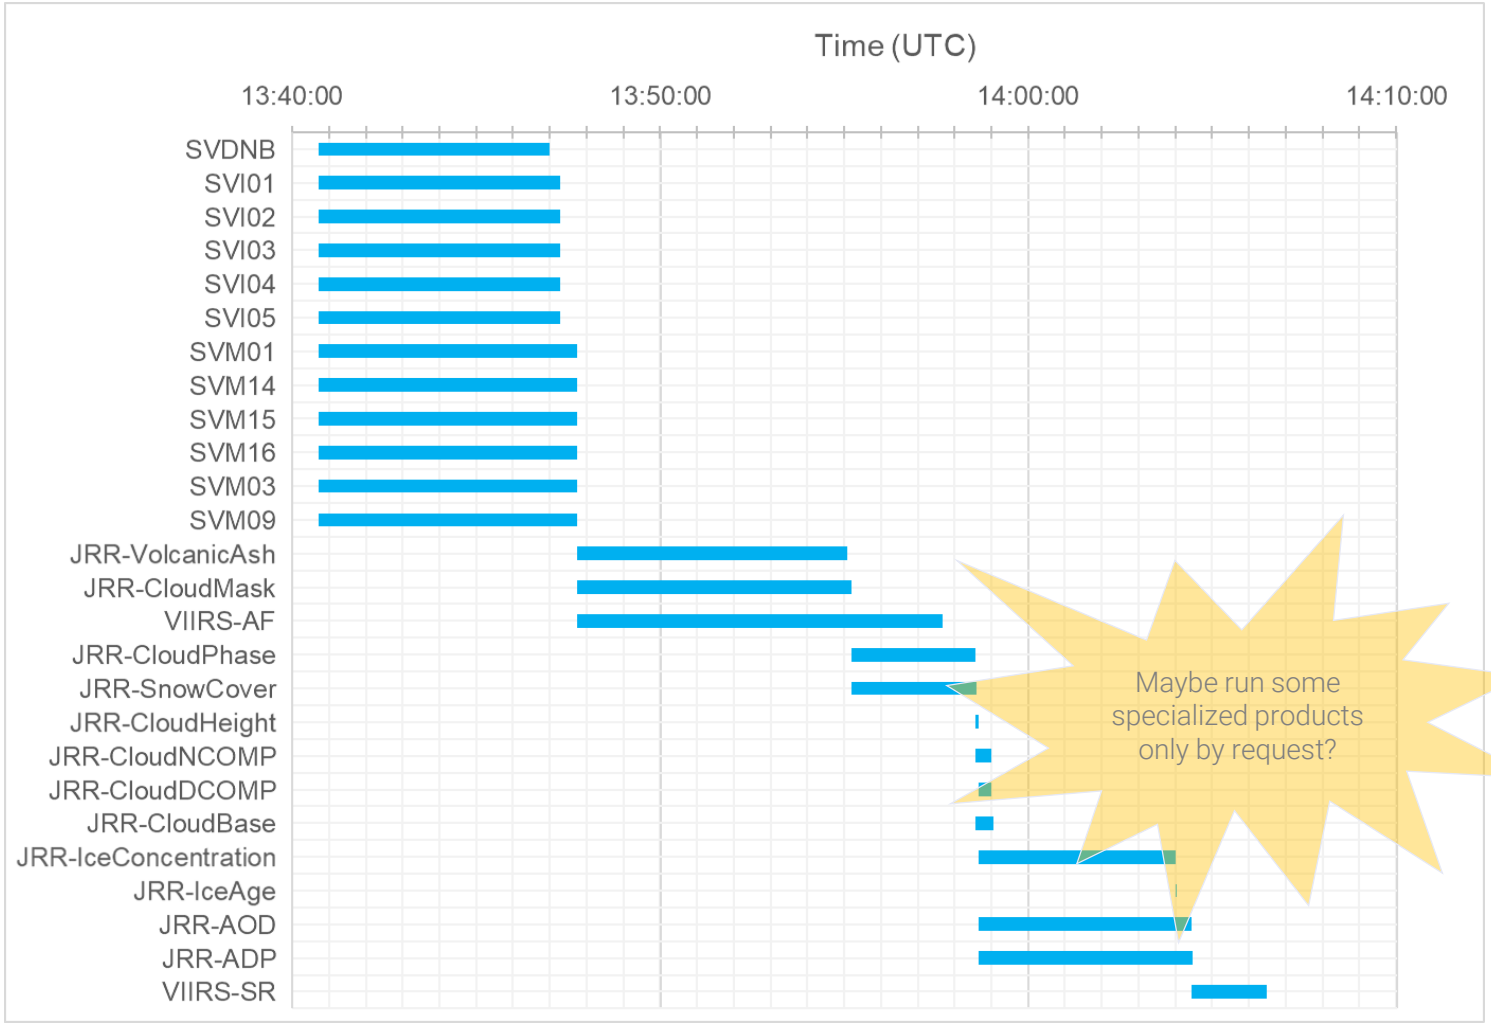

JPSS algorithm run times and interdependencies

Summary

- Providing forecaster access to Joint Polar Satellite System data products is a significant challenge;
- However, given expected patterns of data supply and demand from polar orbiting satellites, and the new OneNWSnet bandwidth, PDA and AWIPS-Data Delivery will be able *(for now)* to provide forecasters with timely access to the JPSS data products they need.
- As data volumes and usage grow, we will need more scalable approaches to product generation, distribution, and access – *for example:*
  - Reduce data movement via thinned products, improved data access services, *etc.*;
  - Use industry-standard protocols for simplicity, versatility, and resilience; and
  - Limit server loads via tiered content distribution and conditional processing.

Near-real-time users: Mb/s, not GB/day

Mbps needed for 10-minute delivery of top 20 JPSS and AMSR-2 products from each CONUS overpass over a 24-hour period

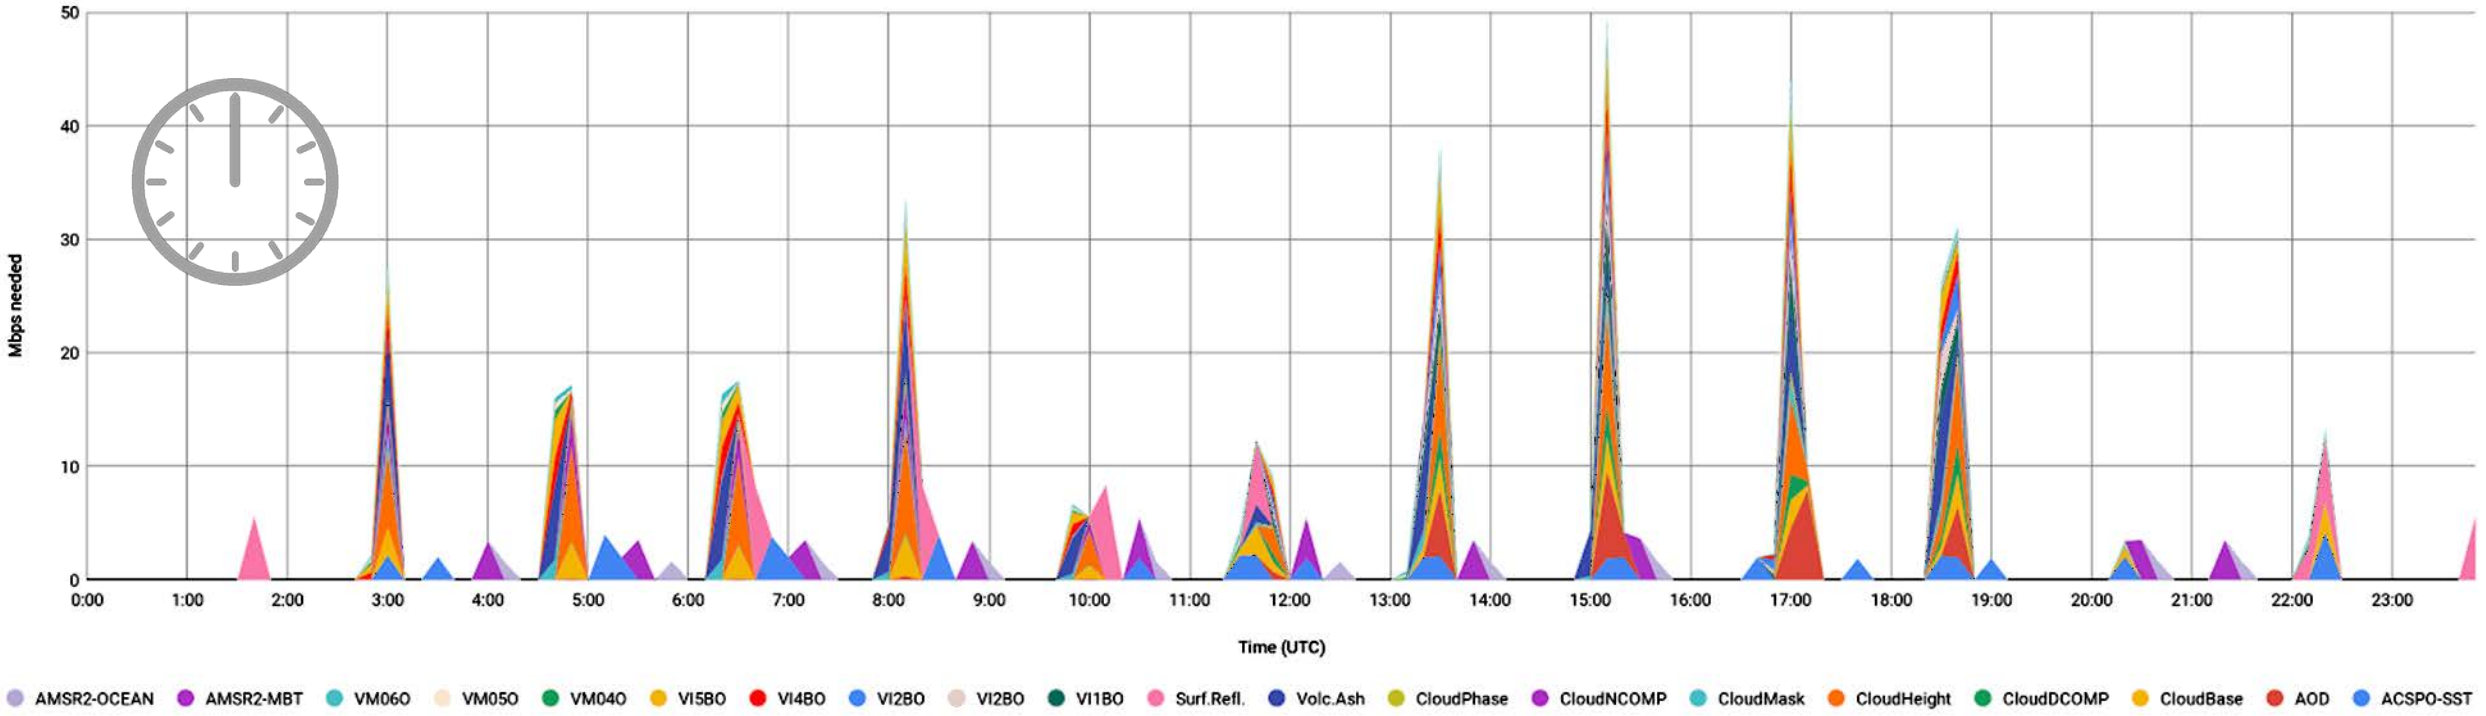

Method: Collected file sizes and file creation times for 44 SNPP and AMSR2 products of interest to AWIPS, intersecting the NWS CONUS region on July 18, 2018. Summed MB (received / sent) per 10-minute interval to infer Mbps needed to "clear the buffer" in each interval. Stacked-area chart above shows Mbps needed for the 20 largest products (~90% of daily data volume).
Findings: Total daily data volume for this region: 45GB / day; but bandwidth needs vary from 0 Mbps to over 58 Mbps (e.g., at 15:10 above).

Sidebar: WCS extensions under development

WCS extensions for swath data (EOX @ OGC Testbed 14)

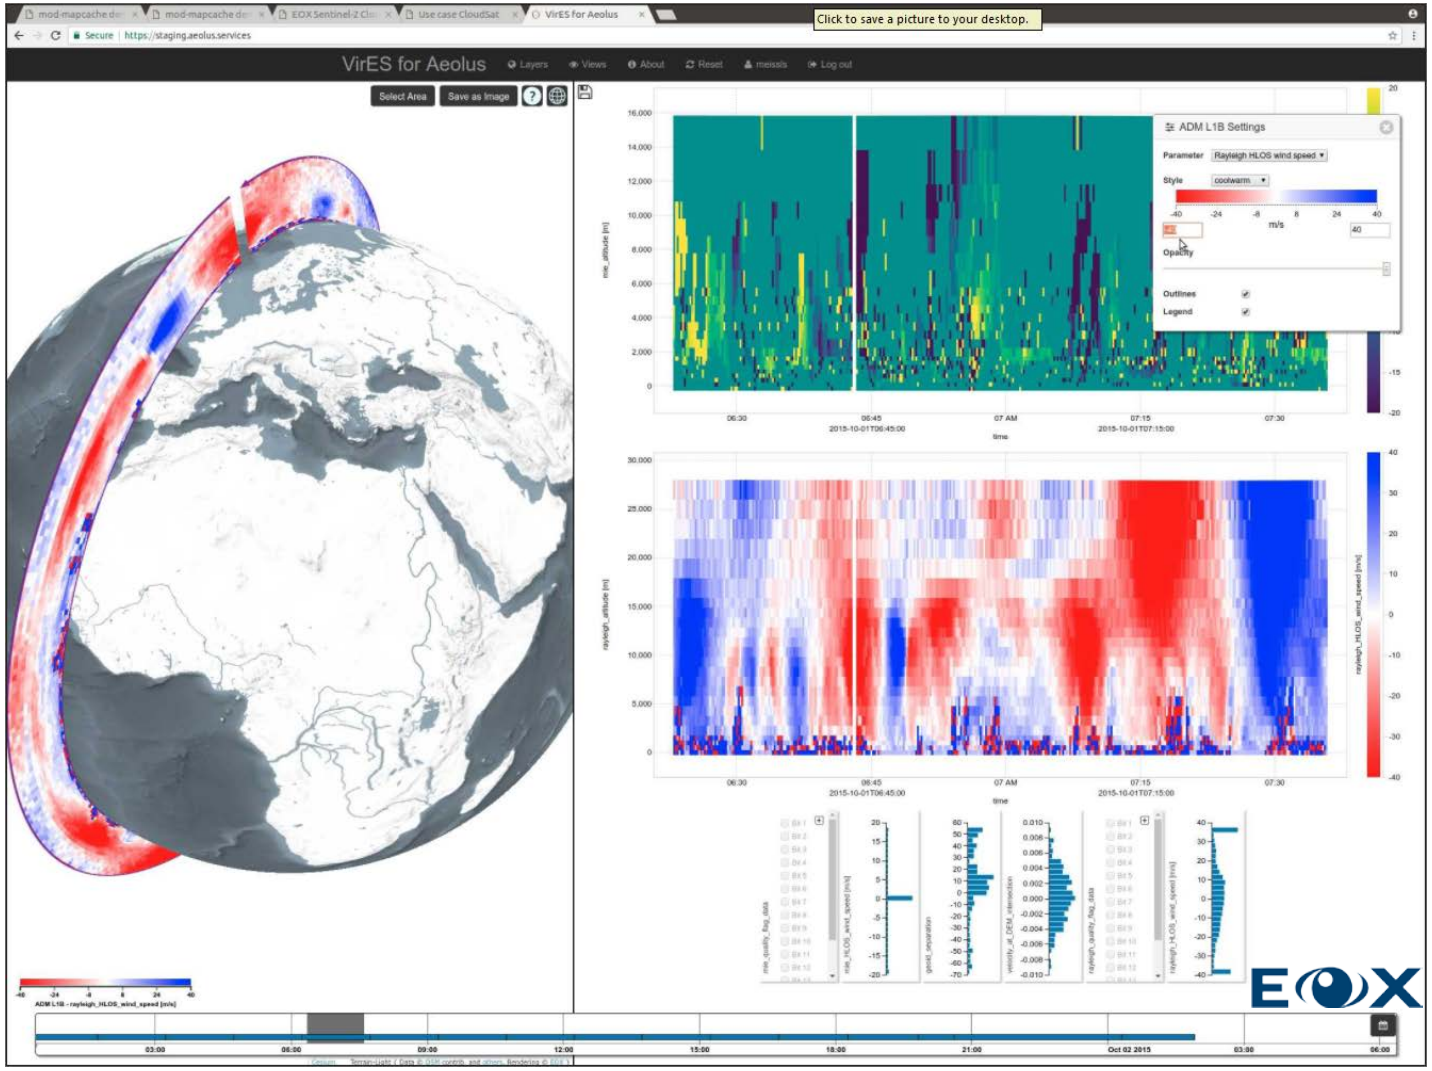

WCS GetCorridor extension (UK Met Office)