

# Xrootd S3 Gateway for WLCG Storage

Computing in High Energy & Nuclear Physics  
CHEP May 8-12, 2023

---

Andrew Hanushevsky, SLAC

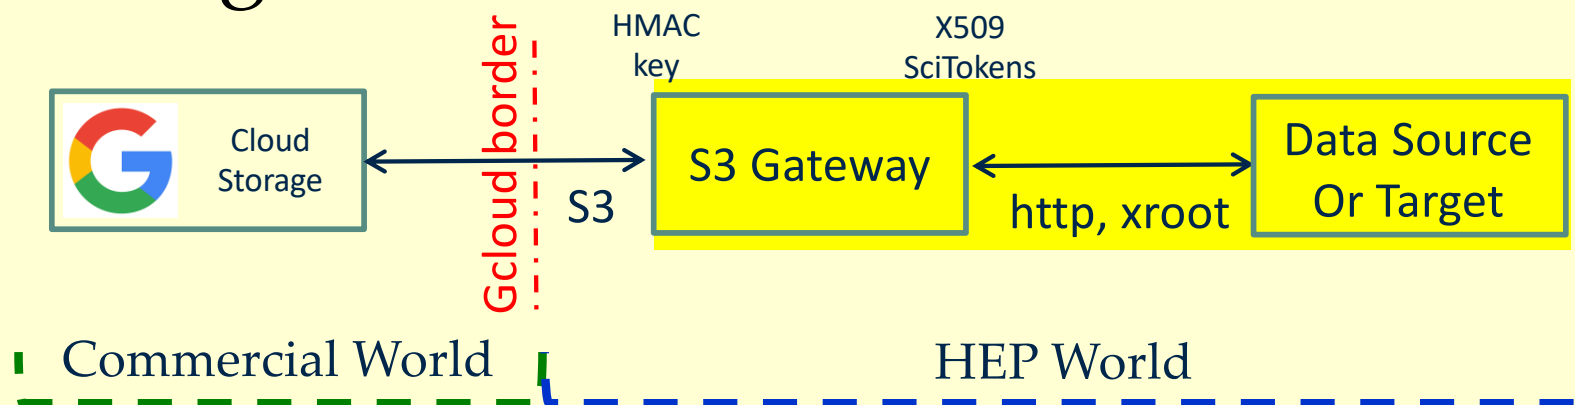


SLAC

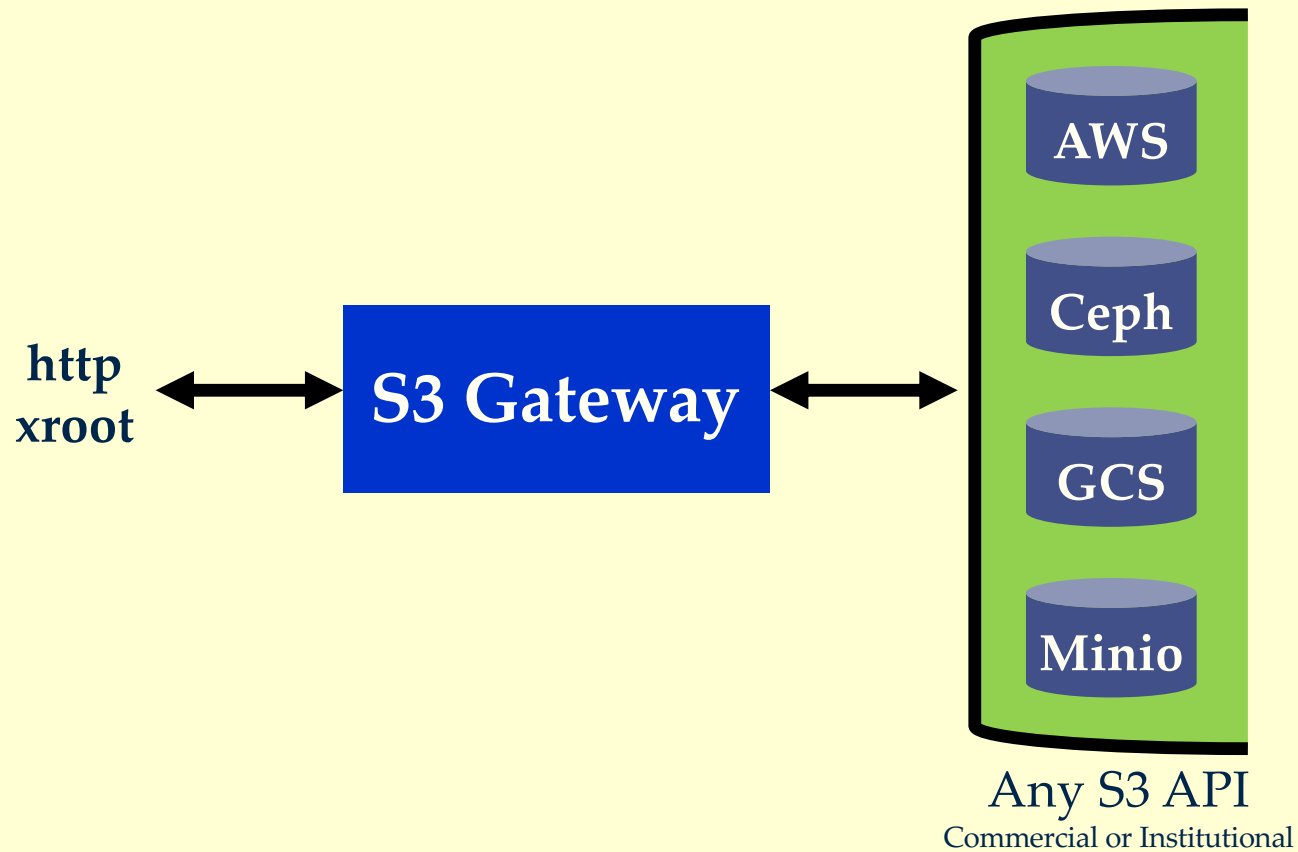


# S3 Gateway Architecture

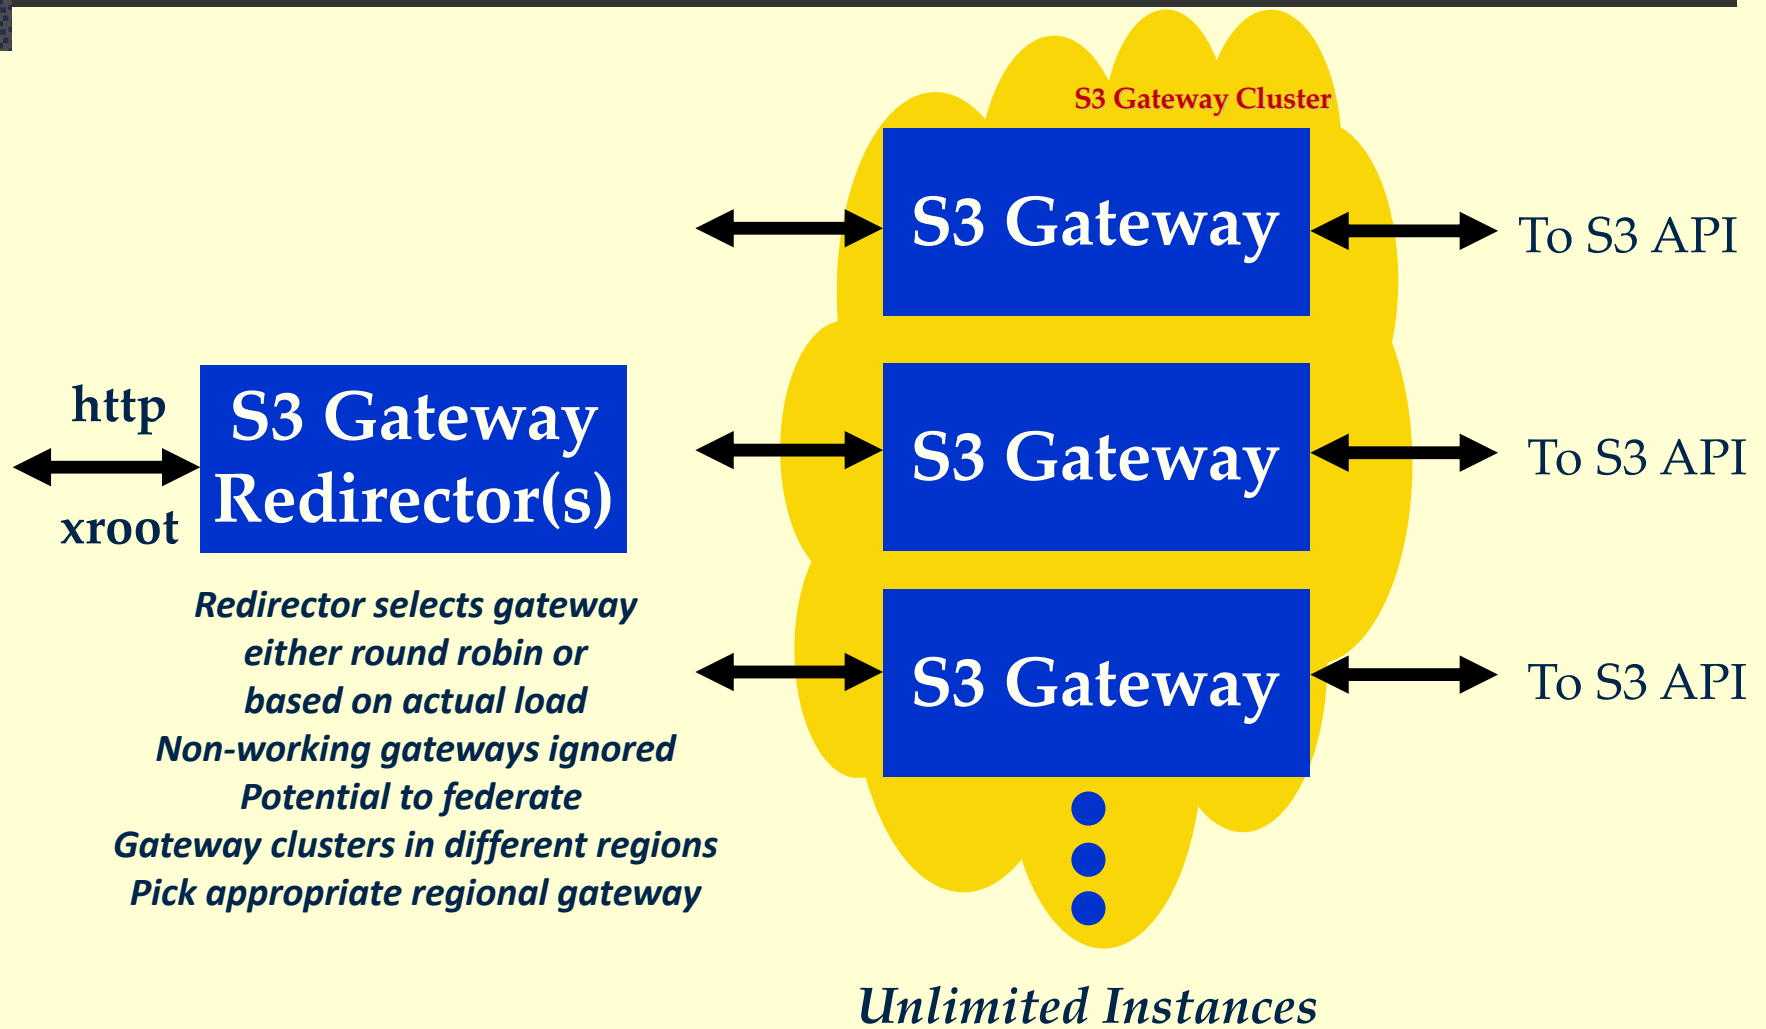
- # Based on **XrdCl** (xrootd client) http plug-in
  - Uses Davix, an HTTP SDK developed at CERN
    - Very reliable and supported
    - Performs better than most commercial SDK's
- # Bridges HEP & commercial world security



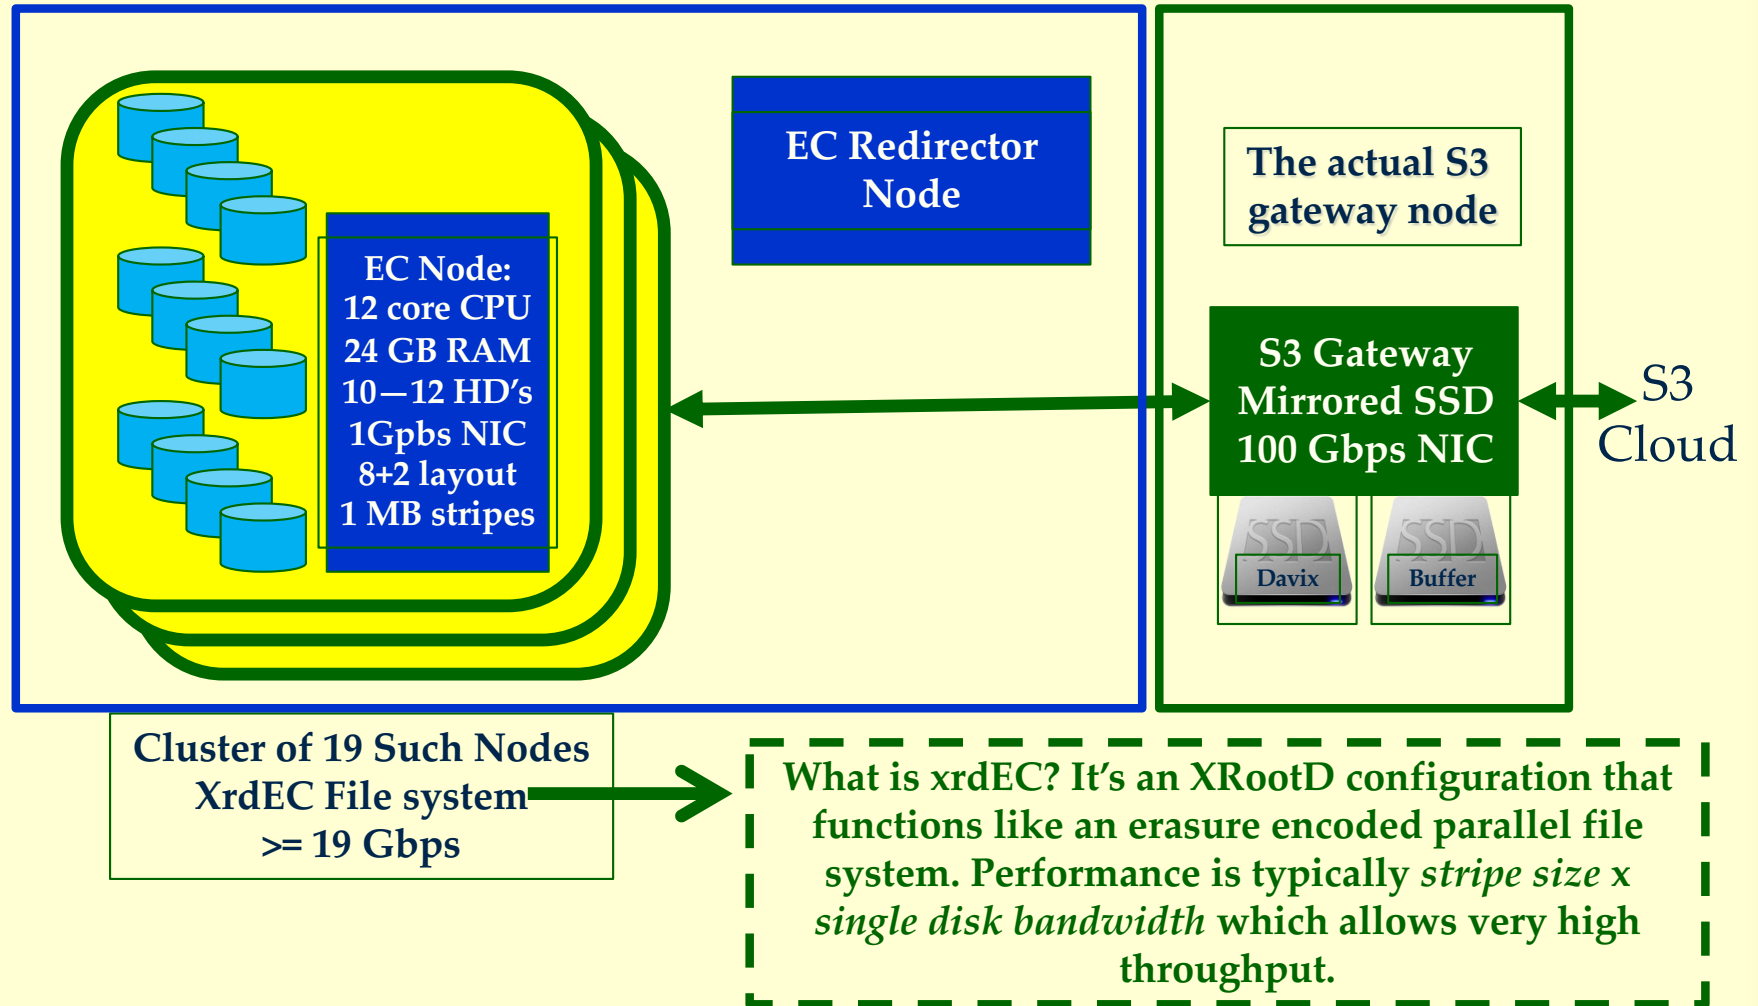
# S3 Gateway Applicability



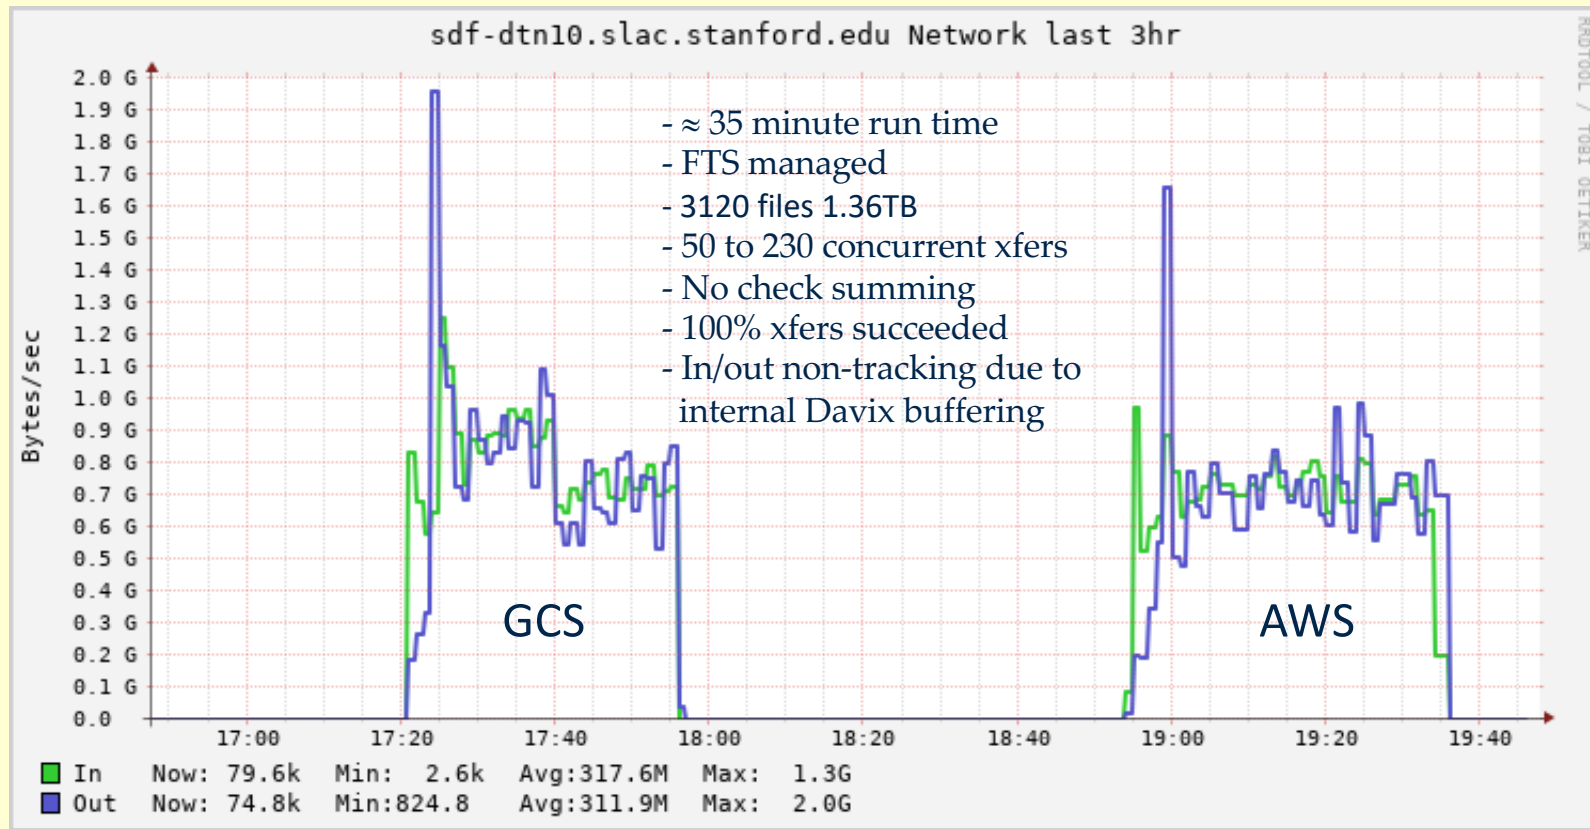
# S3 Gateway Scaling



# S3 Gateway Test Setup



# Ingress Performance

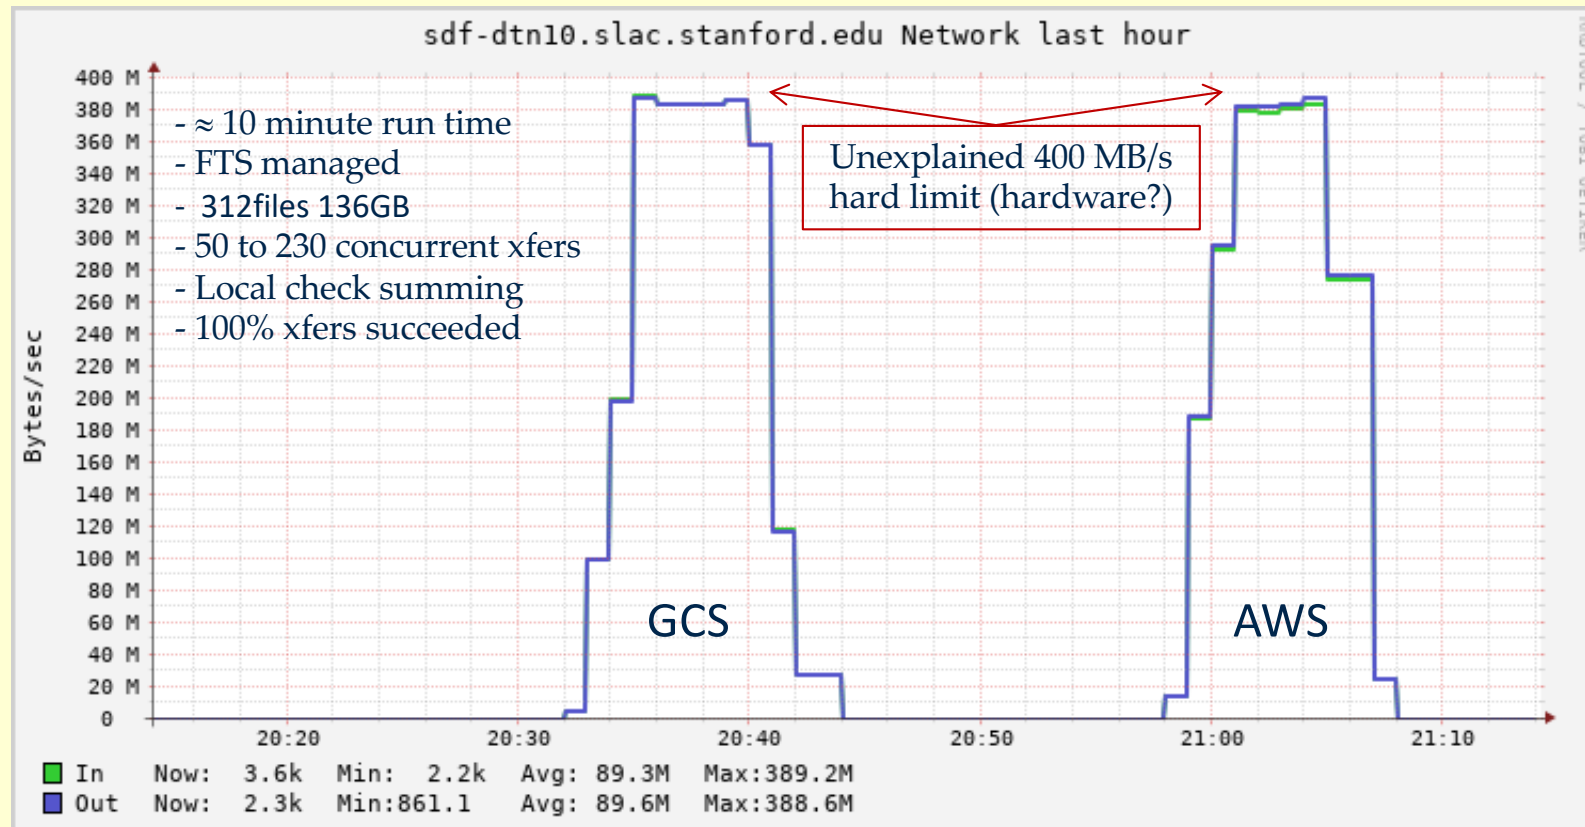


# Internal Davix Buffering

---

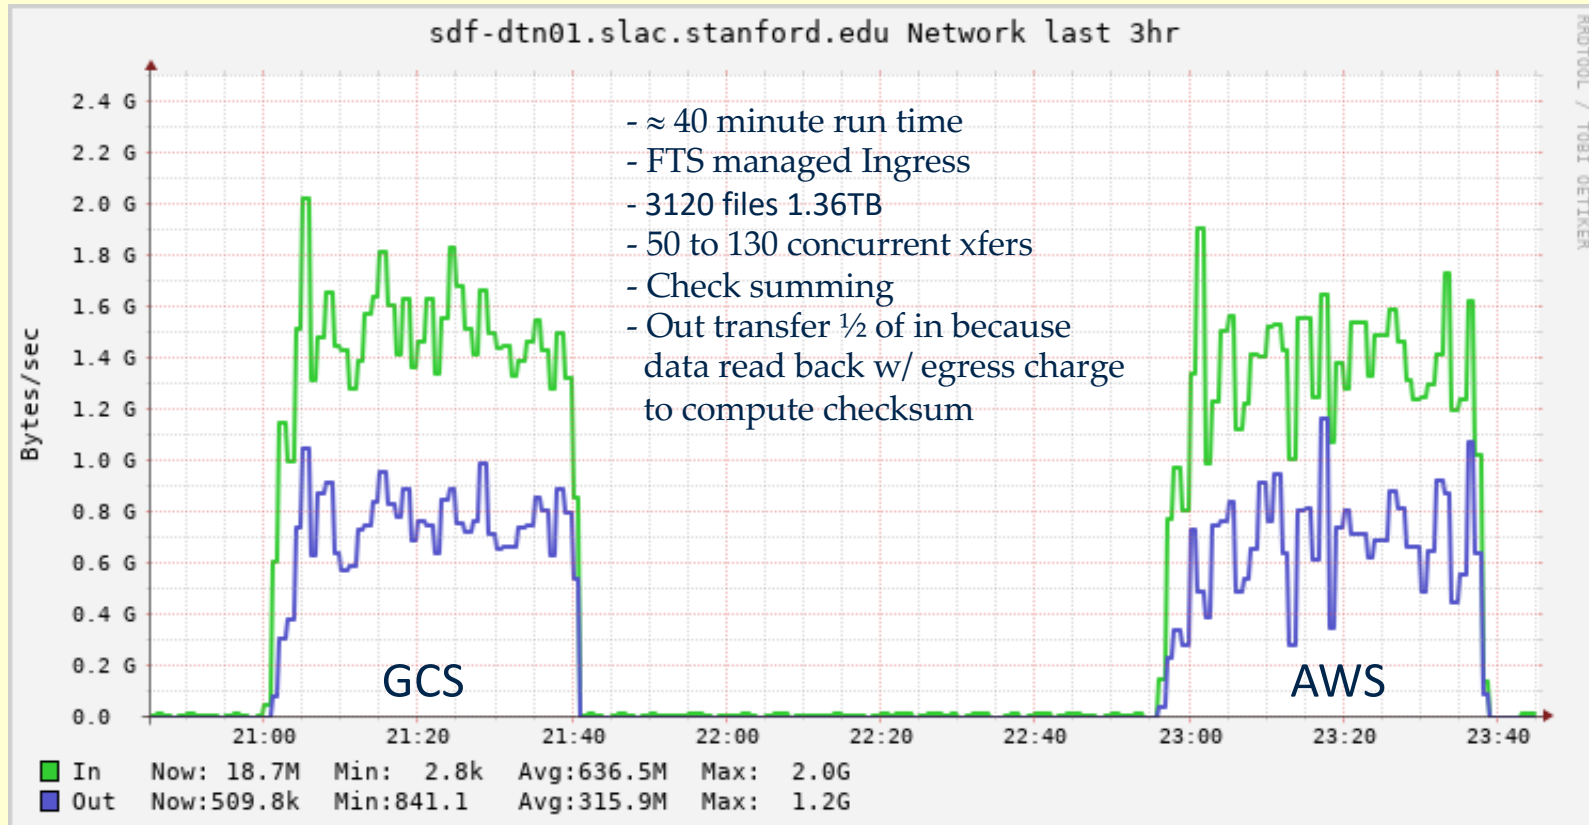
- # Current Davix fully buffers stream I/O
  - Stream I/O creates a physical file
  - File is then forwarded to endpoint
- # Future Davix eliminates this buffering
  - Working with Davix team to implement this
    - Will be part of final product

# Egress Performance





# S3 Gateway CKS Performance



# Avoiding CKS Egress Charge I

- # AWS & GCS provide server-less computing
  - AWS via lambda
    - Python, Java, Google Go and C#
  - GCS via Google Cloud Functions
    - Python, Java, Google Go, .NET, Ruby, and PHP
- # Leverage these to compute checksum
  - S3 Gateway triggers server-less cks program
    - Checksum computed in the cloud (no egress)
    - Result transmitted back to S3 gateway

# S3 Gateway for multiple API's

- # The S3 Gateway is universal
  - Work with all S3 storage flavors we tested
    - work with both s3v4 and older s3v2
  - S3 credentials: different names, but same thing
    - AWS: ACCESS\_KEY\_ID & SECRET\_ACCESS\_KEY
    - Ceph: HMAC key pair
    - GCS: HMAC key pair
    - MinIO: username & password
    - S3 Gateway/Davix uses AWS naming convention

# Conclusion

- # S3 Gateway is extremely economical
  - Avoids most cloud charges
    - Except egress when fetching data from the cloud
      - Built-in authorization can restrict who can do this
- # S3 Gateway provides uniform access
  - Regardless of S3 provider access is the same
    - Automatic conversion of HEP auth to S3 auth
- # Proven scalability and performance
- # Doc on [Xrootd-HowTo](#)