



XRootD at Qualys

Artem Harutyunyan,
Engineering Director

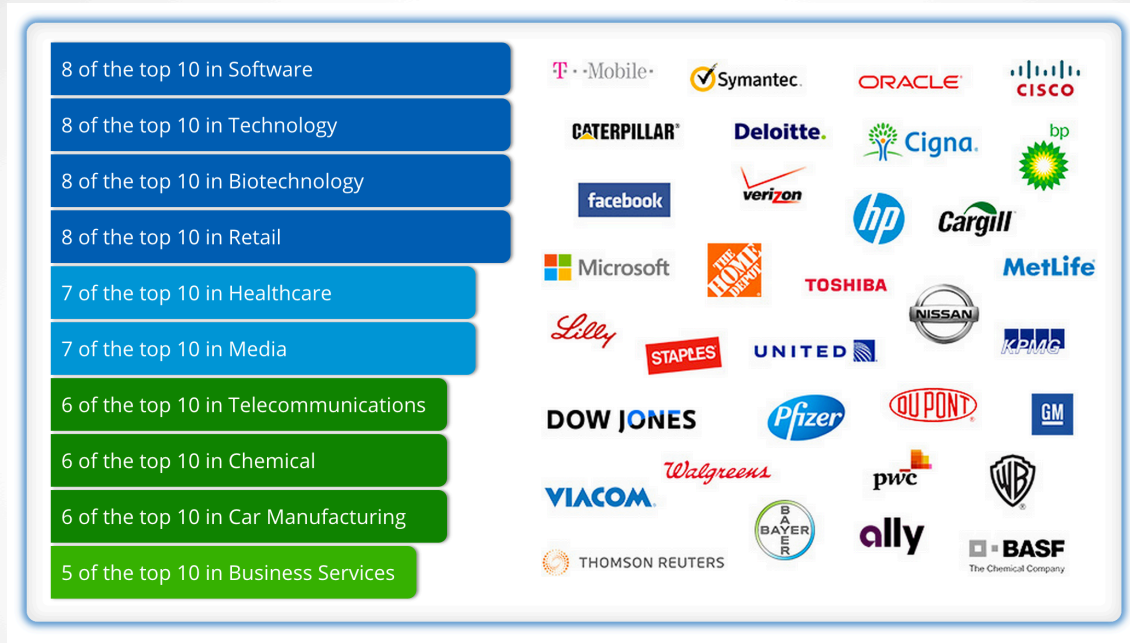
Who am I?

- ALICE Offline (2004 – 2010)
 - AliEn (porting to Windows, authentication, scheduling, etc.)
- Fellow at PH/SFT (2010 – 2012)
 - CernVM Co-Pilot – distributed computing framework powering LHC@home 2.0
 - CernVM development and infrastructure
 - CernVM File System infrastructure
- Architect/Engineering Director at Qualys (2012 – present)

About Qualys

- Provider of cloud security, compliance and related services for large corporations as well as small and medium-sized enterprises
- 6700 customers in more than 100 countries
- Engineering offices across US (Silicon Valley, Denver, Wisconsin, Seattle), in France (Paris), Switzerland (Geneva), India (Pune)


About Qualys Customers



- Some of the challenges
 - Scanning an entire class A network in under 4 hours
 - Manage O(1,000,000) assets per company
 - Manage O(10,000) web applications per company

How it works?

Network scanning
and data collection



Storage and
Processing



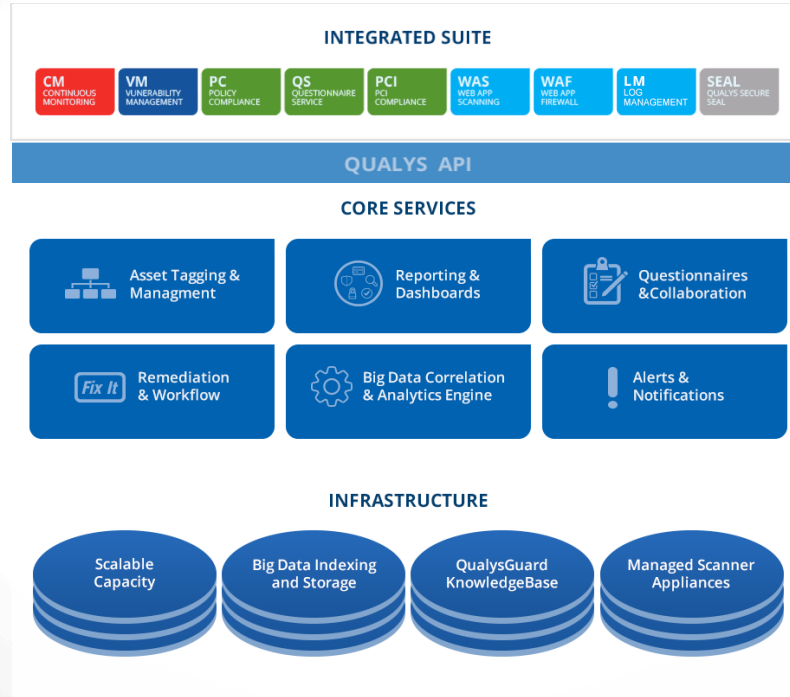
Analysis and
Correlation



Customer's network

Qualys data center

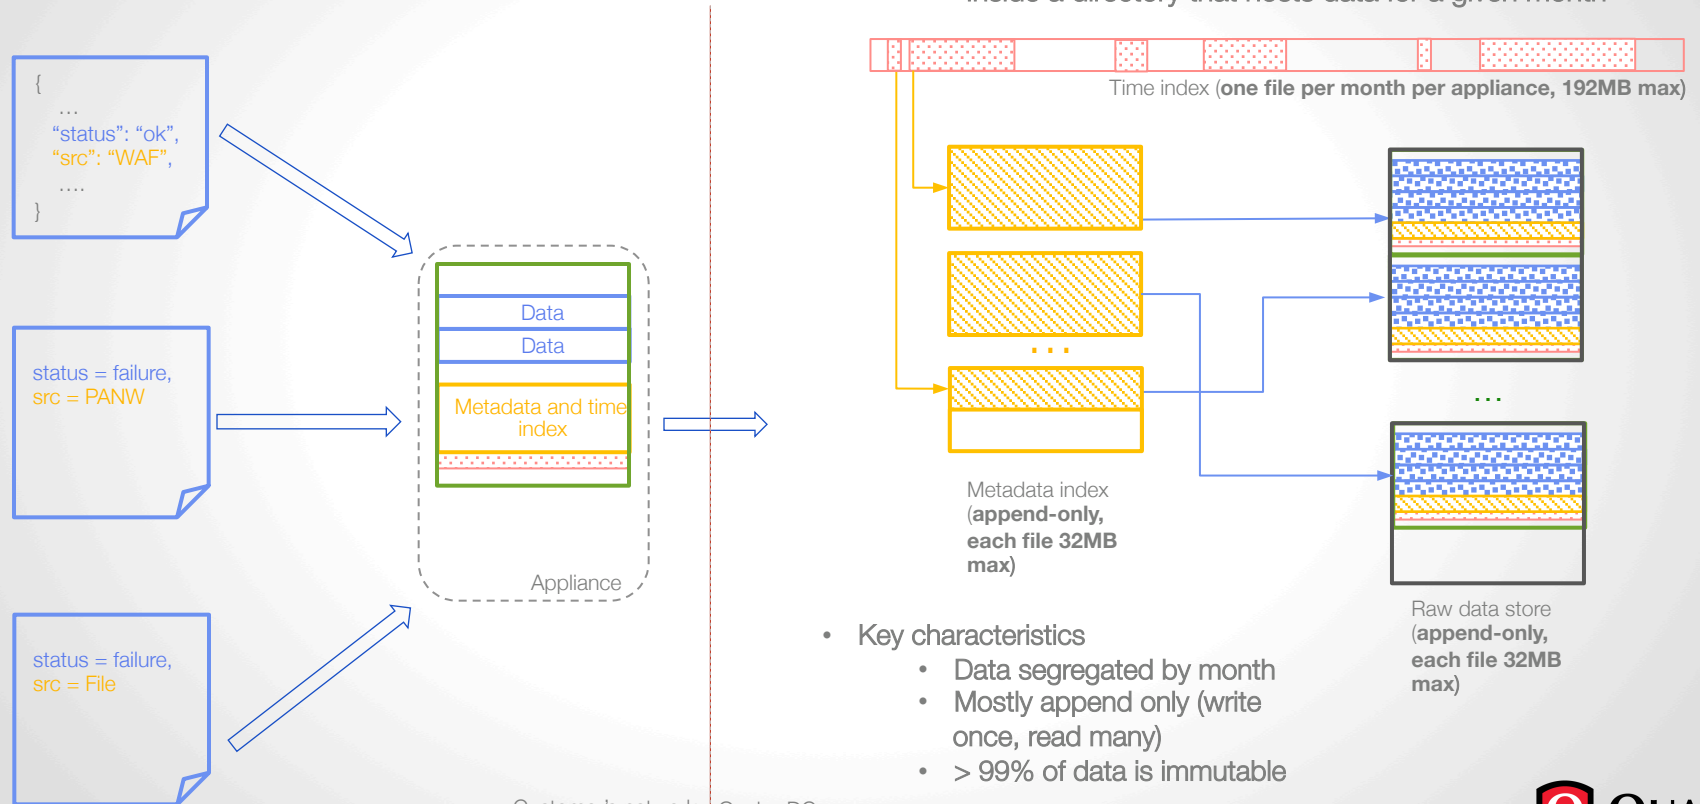
What does it do?



How do we use XRootD?

- XRootD is used to store
 - Compressed raw data
 - Time and metadata indices
- Data is accessed in parallel by a set of distributed processing nodes
- All the data resides within the same cluster/datacenter


Data collection and storage



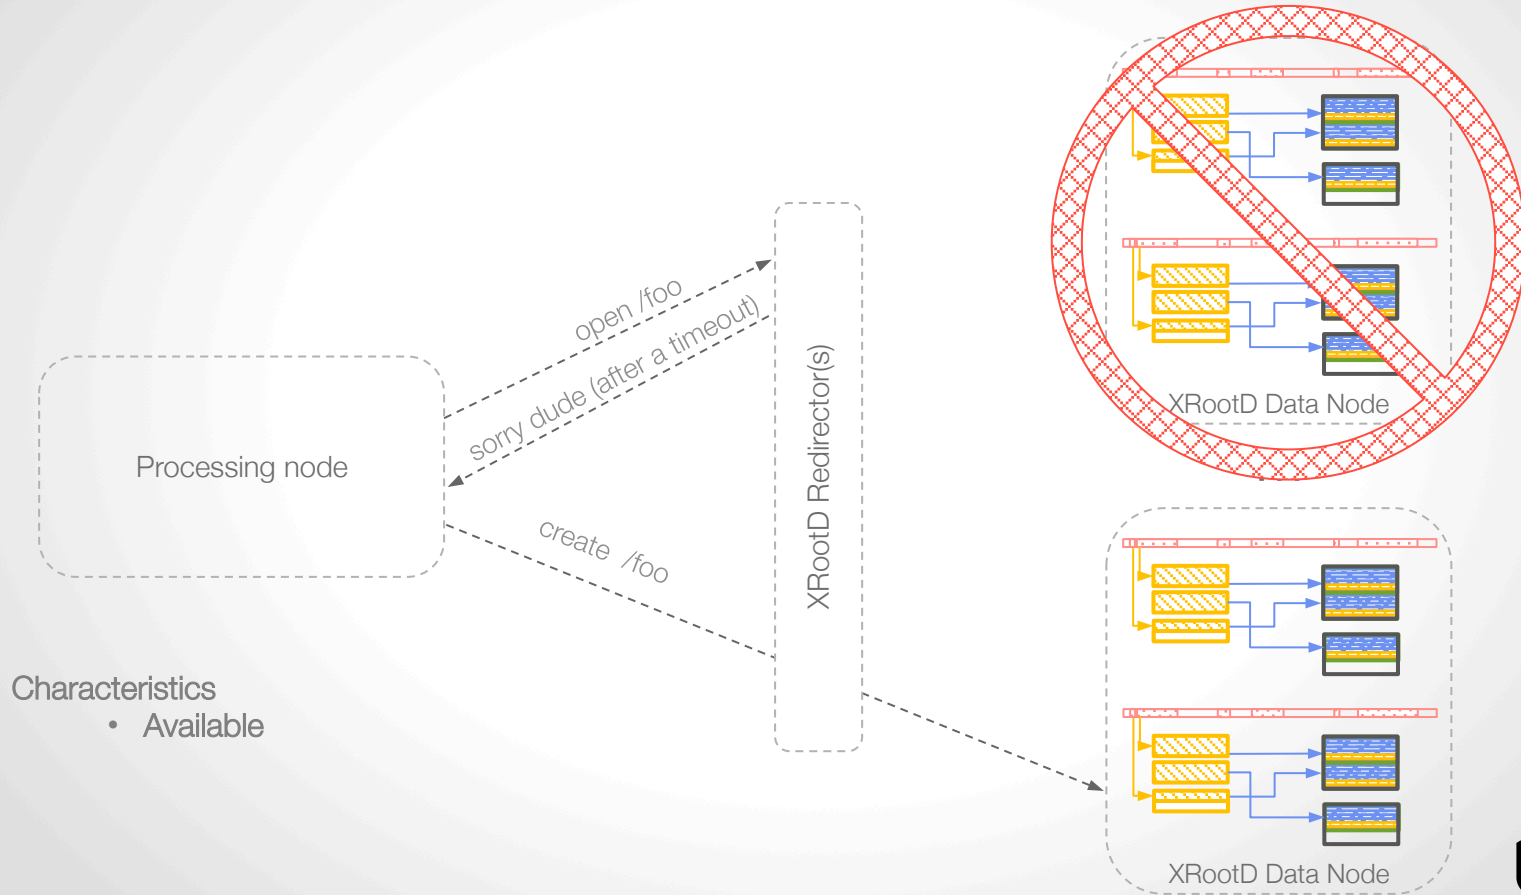
Customer's network | Qualys DC

- Key characteristics
 - Data segregated by month
 - Mostly append only (write once, read many)
 - > 99% of data is immutable

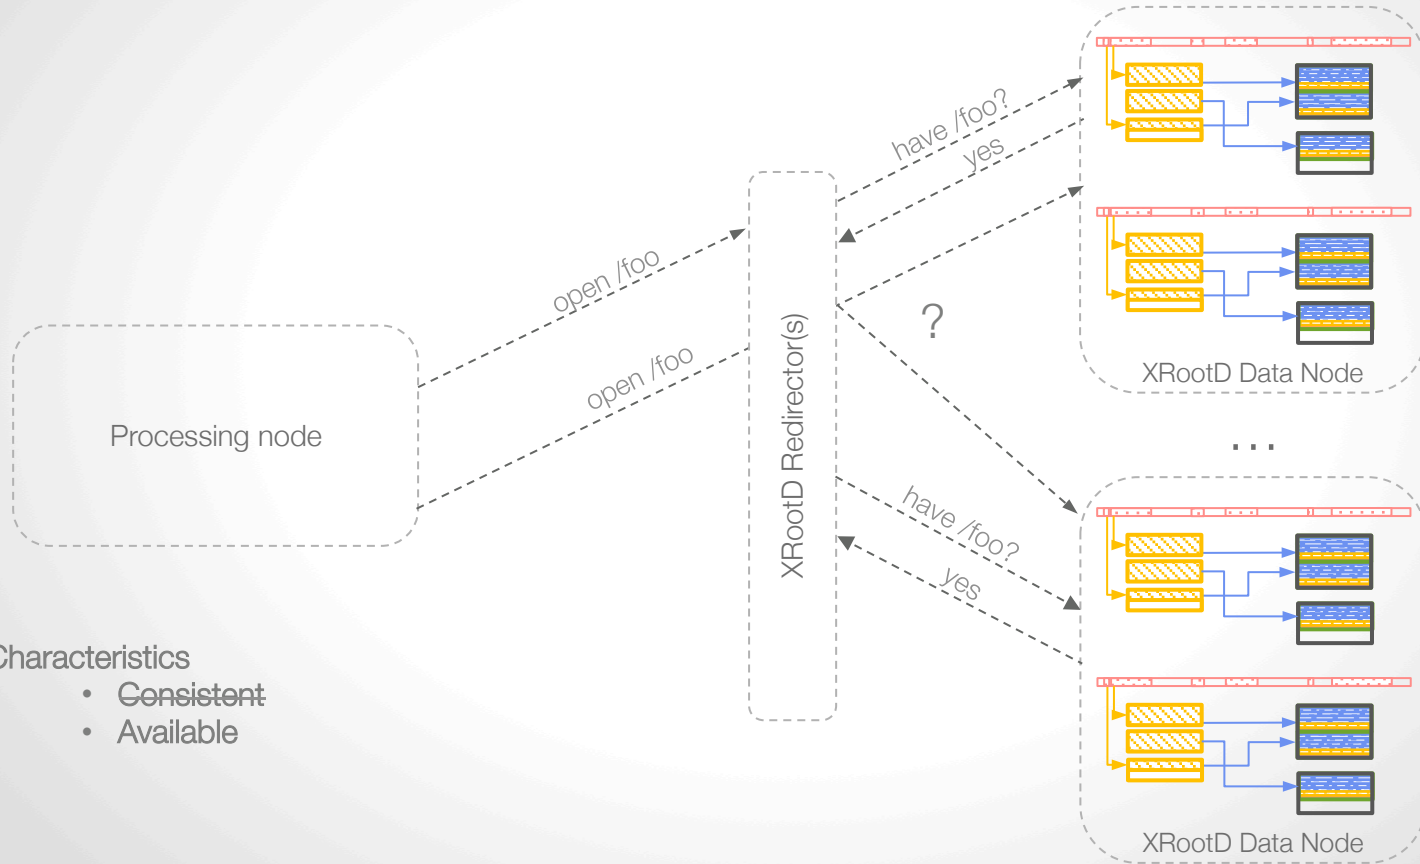
Data access with XRootD



Data access with XRootD



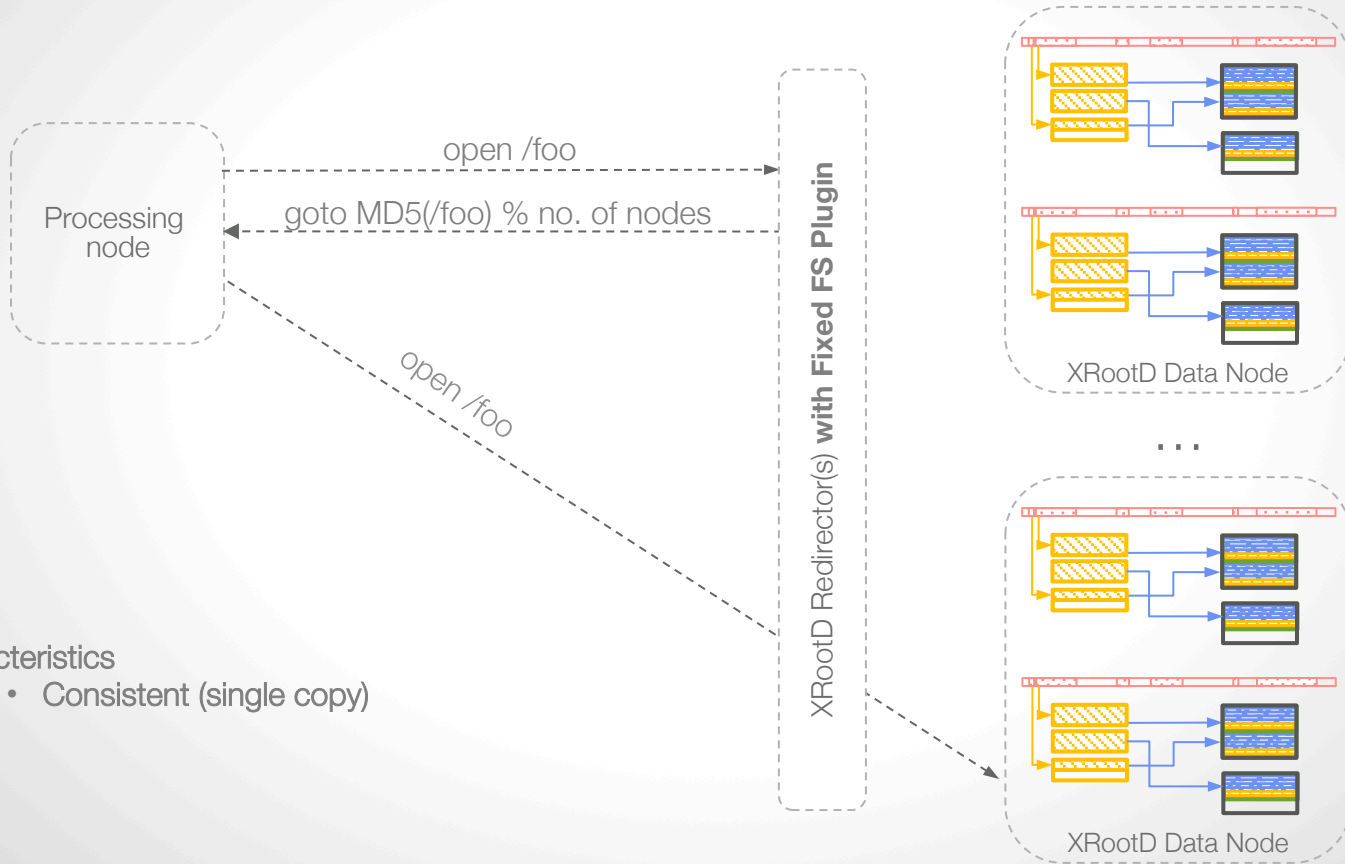
Data access with XRootD



Characteristics

- Consistent
- Available

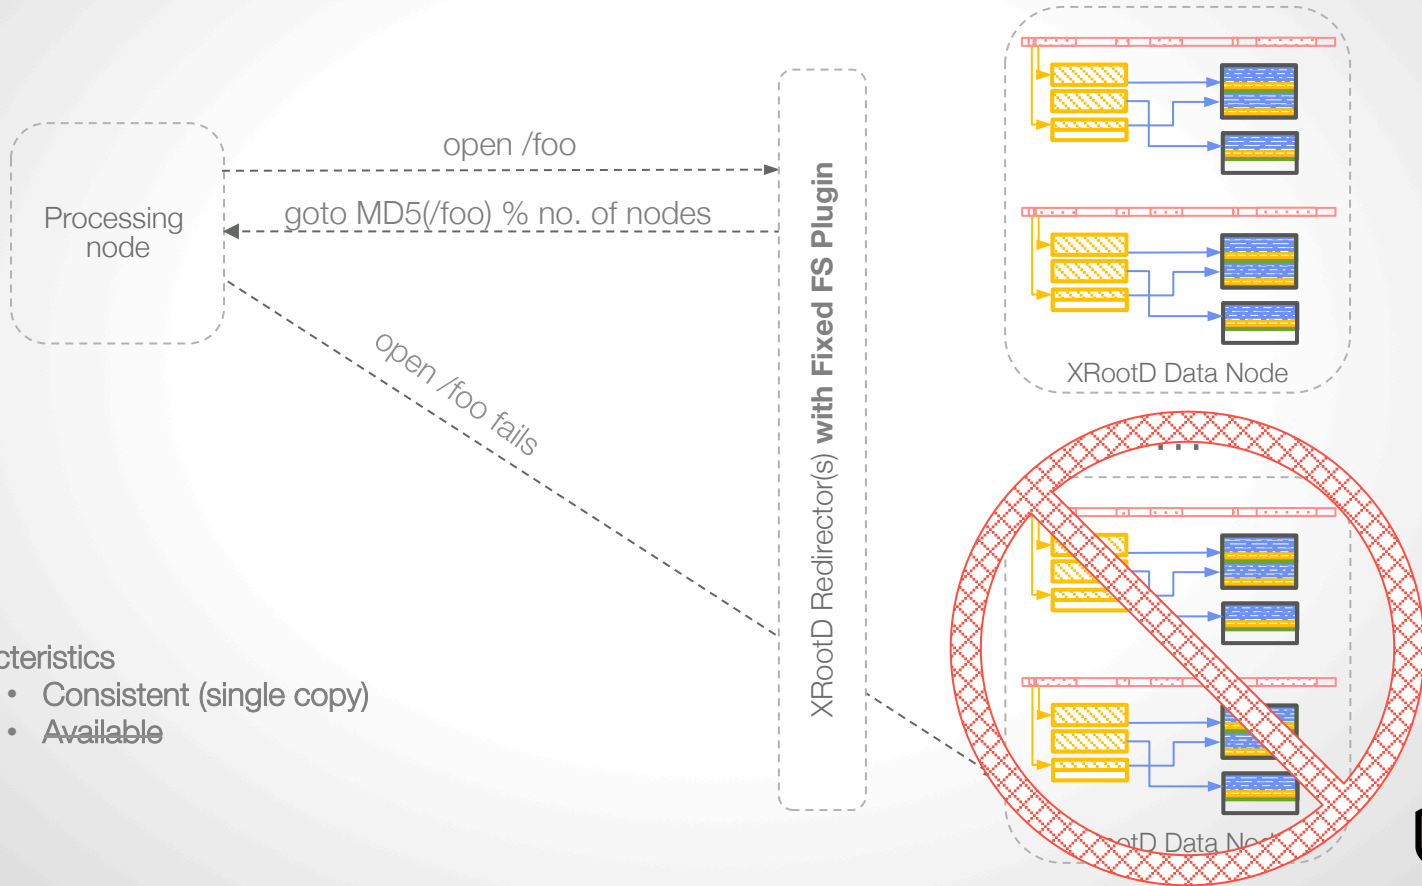
Data access with XRootD: 'Fixed' FS plugin



Characteristics

- Consistent (single copy)

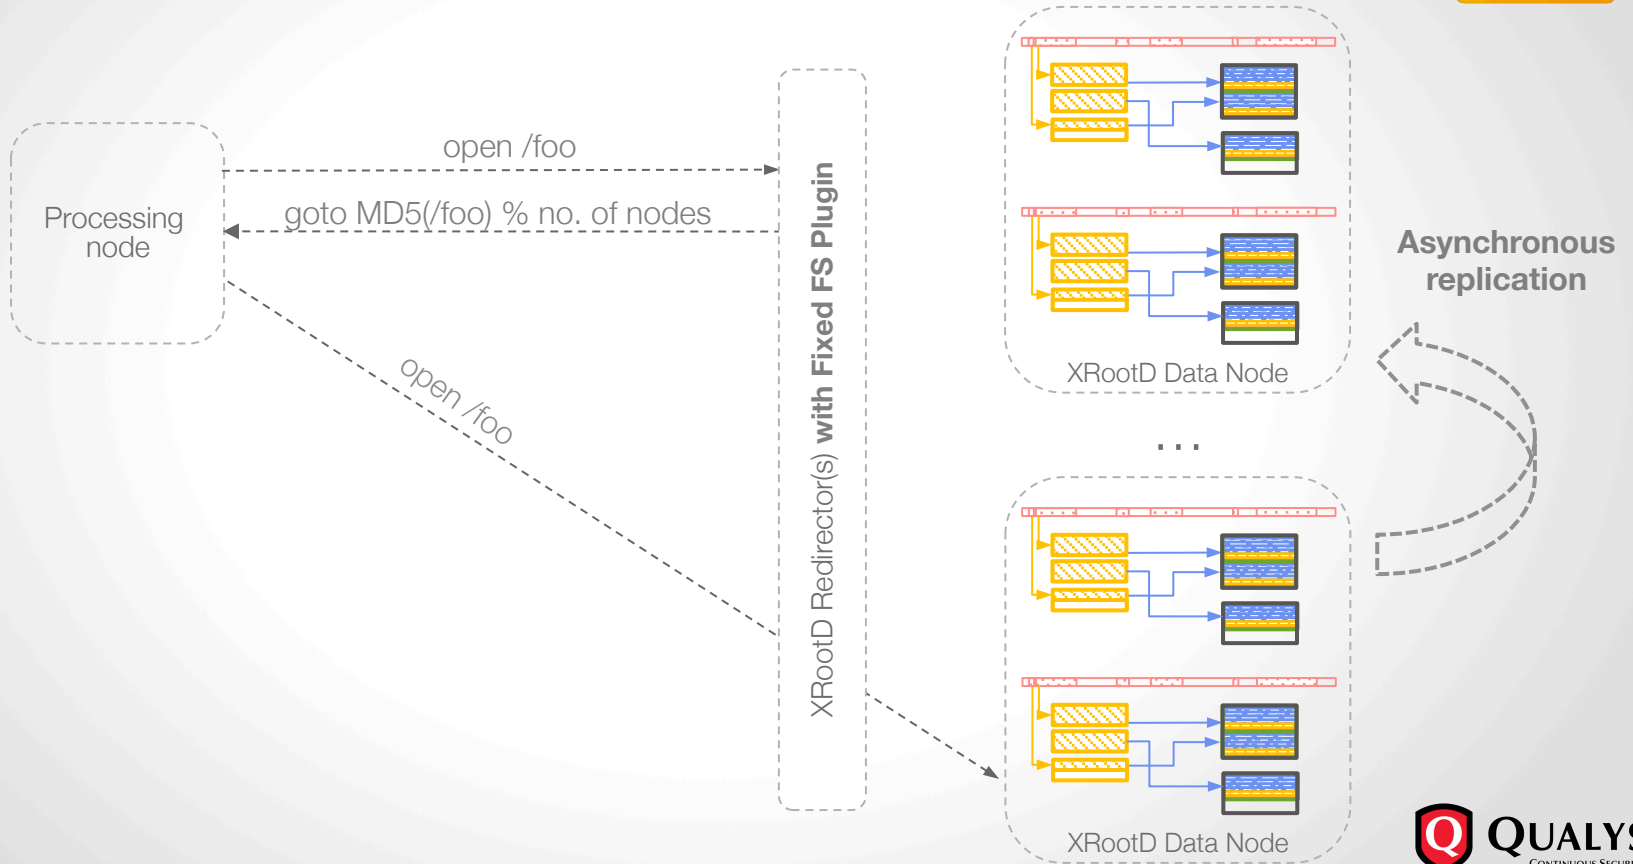
Data access with XRootD: 'Fixed' FS plugin



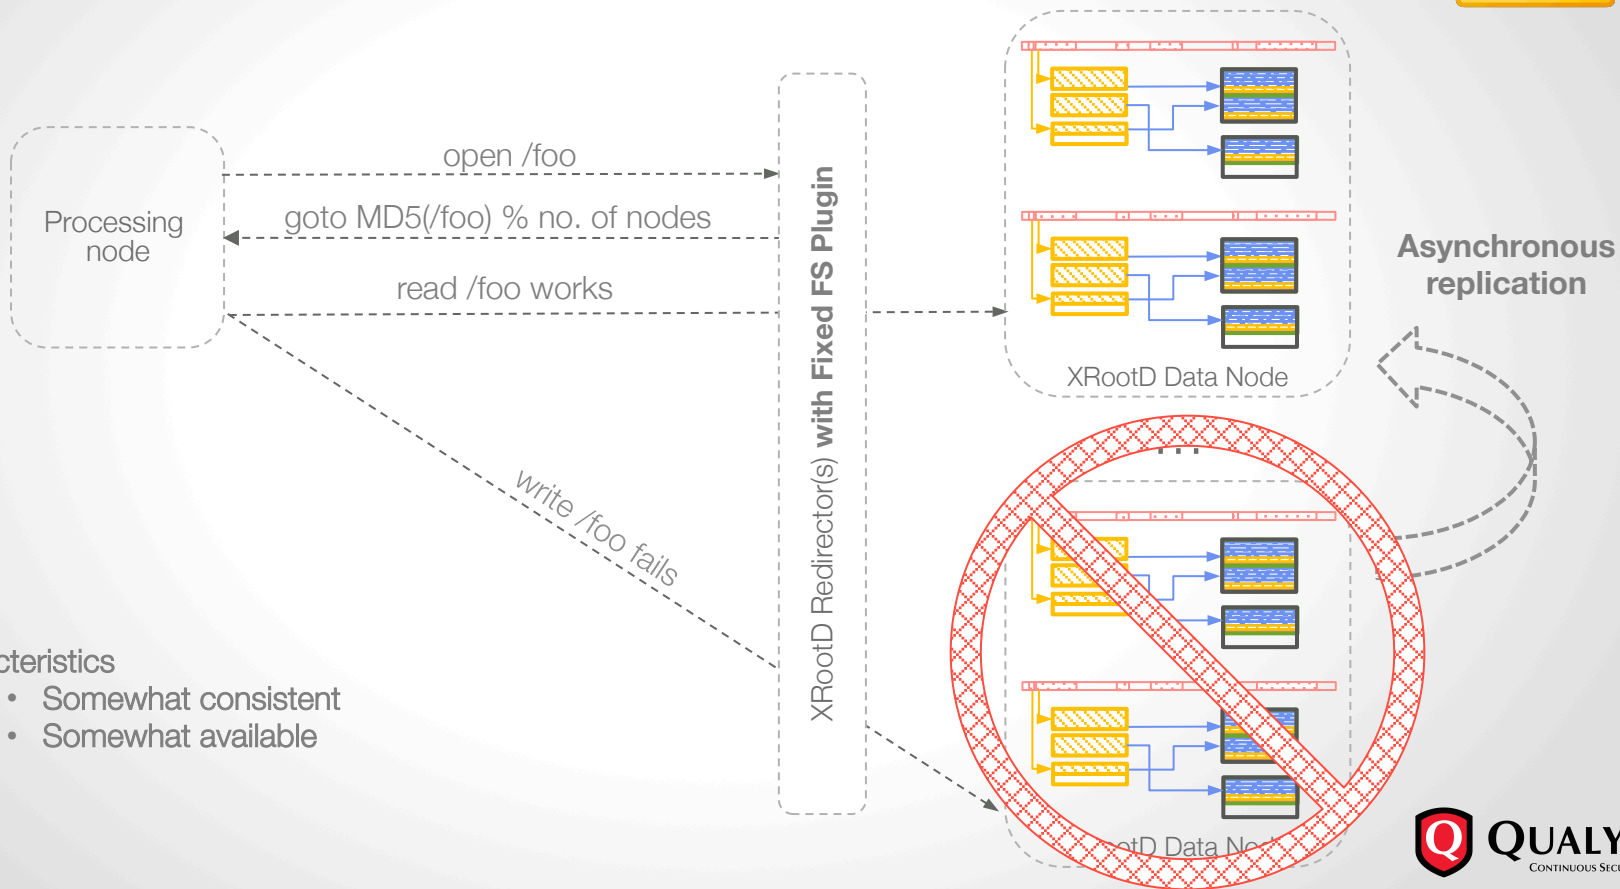
Characteristics

- Consistent (single copy)
- Available

Data access with XRootD: 'Fixed' FS plugin + replication



Data access with XRootD: 'Fixed' FS plugin + replication



What about scaling?

- Changing the number of data nodes will require relocation of the (entire) dataset. Can be done, but better be avoided
- When the current cluster gets near it's storage capacity, the new one will be created and Fixed FS redirector plugin will redirect requests to the right cluster (based on the month/year which is always present in the path)

Questions?



QUALYS®

Thank You

hartem@qualys.com

twitter.com/hartem

[linkedin.com/in/artemharutyunyan](https://www.linkedin.com/in/artemharutyunyan)