



ESO lessons learned implementing VOSpace 1.0 with NGAS backend

Paul Harrison
ESO


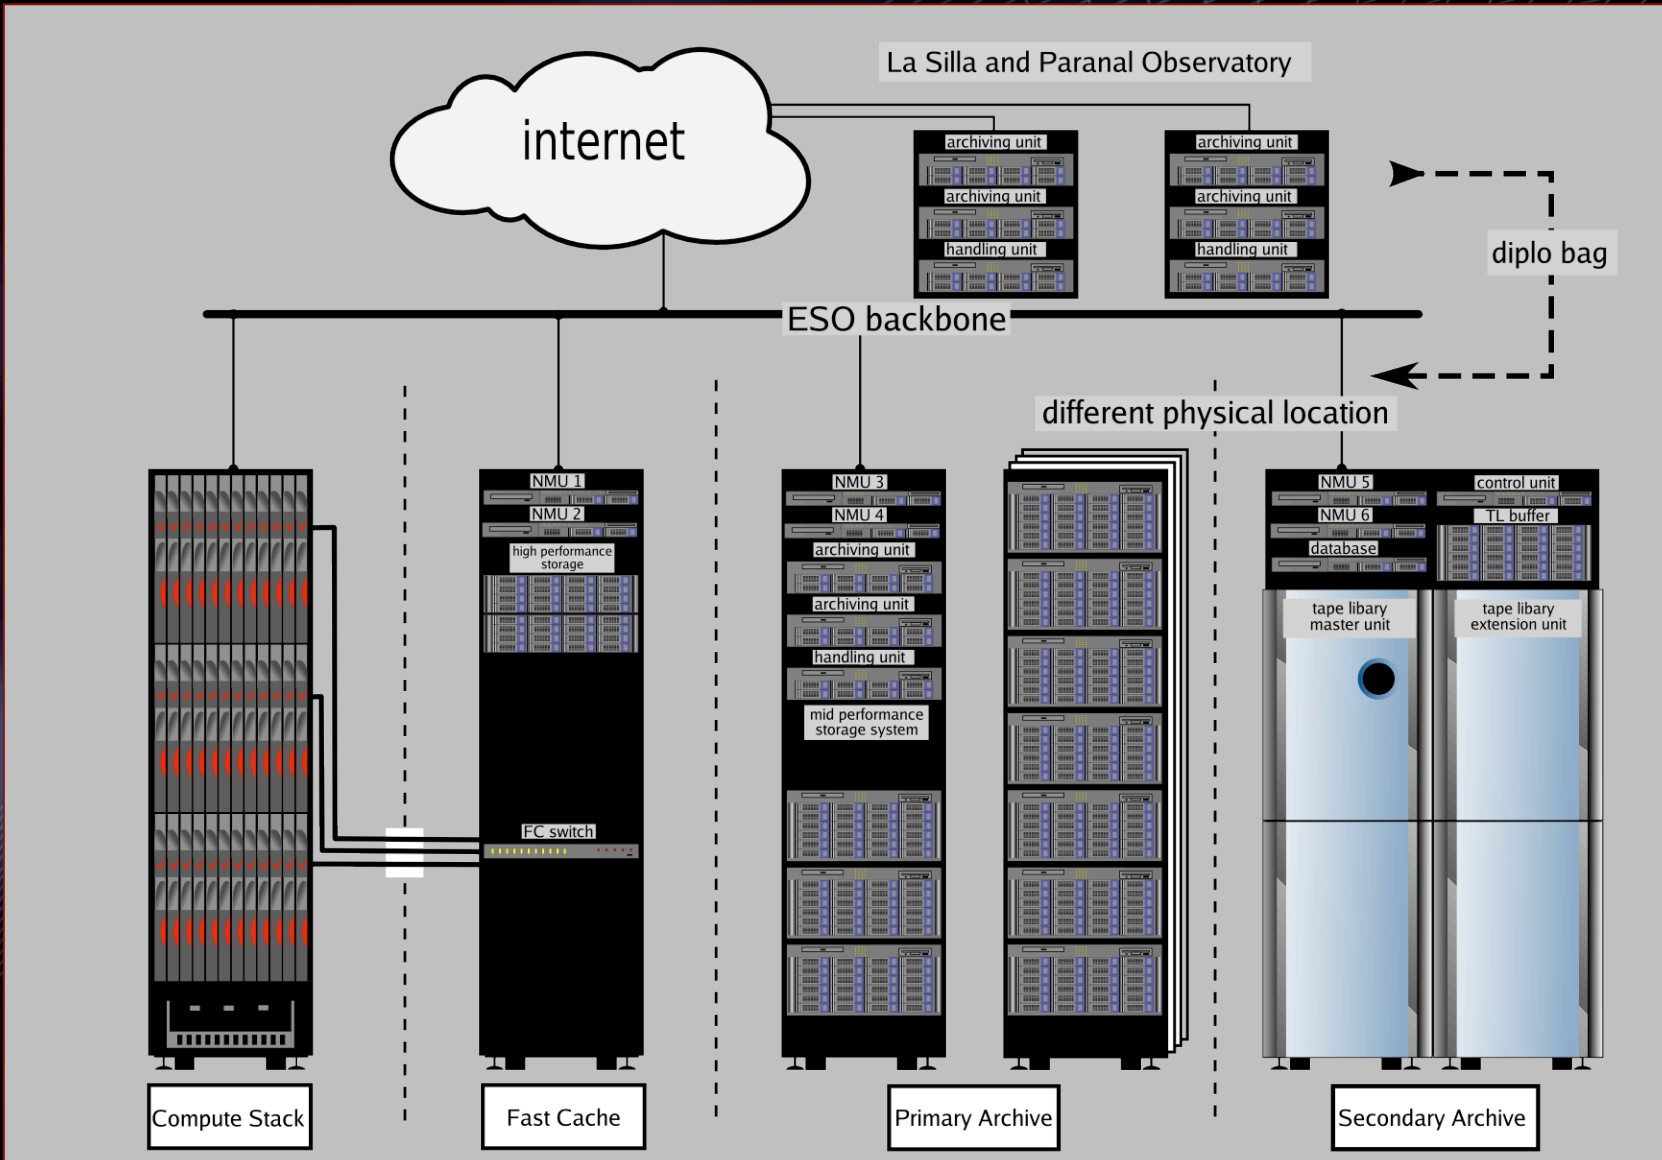
Introduction

- Motivation
 - What is NGAS?
- Implementation Details
- Integration testing
 - Registry
 - Scaling
 - Authentication & Authorization
 - Interoperability

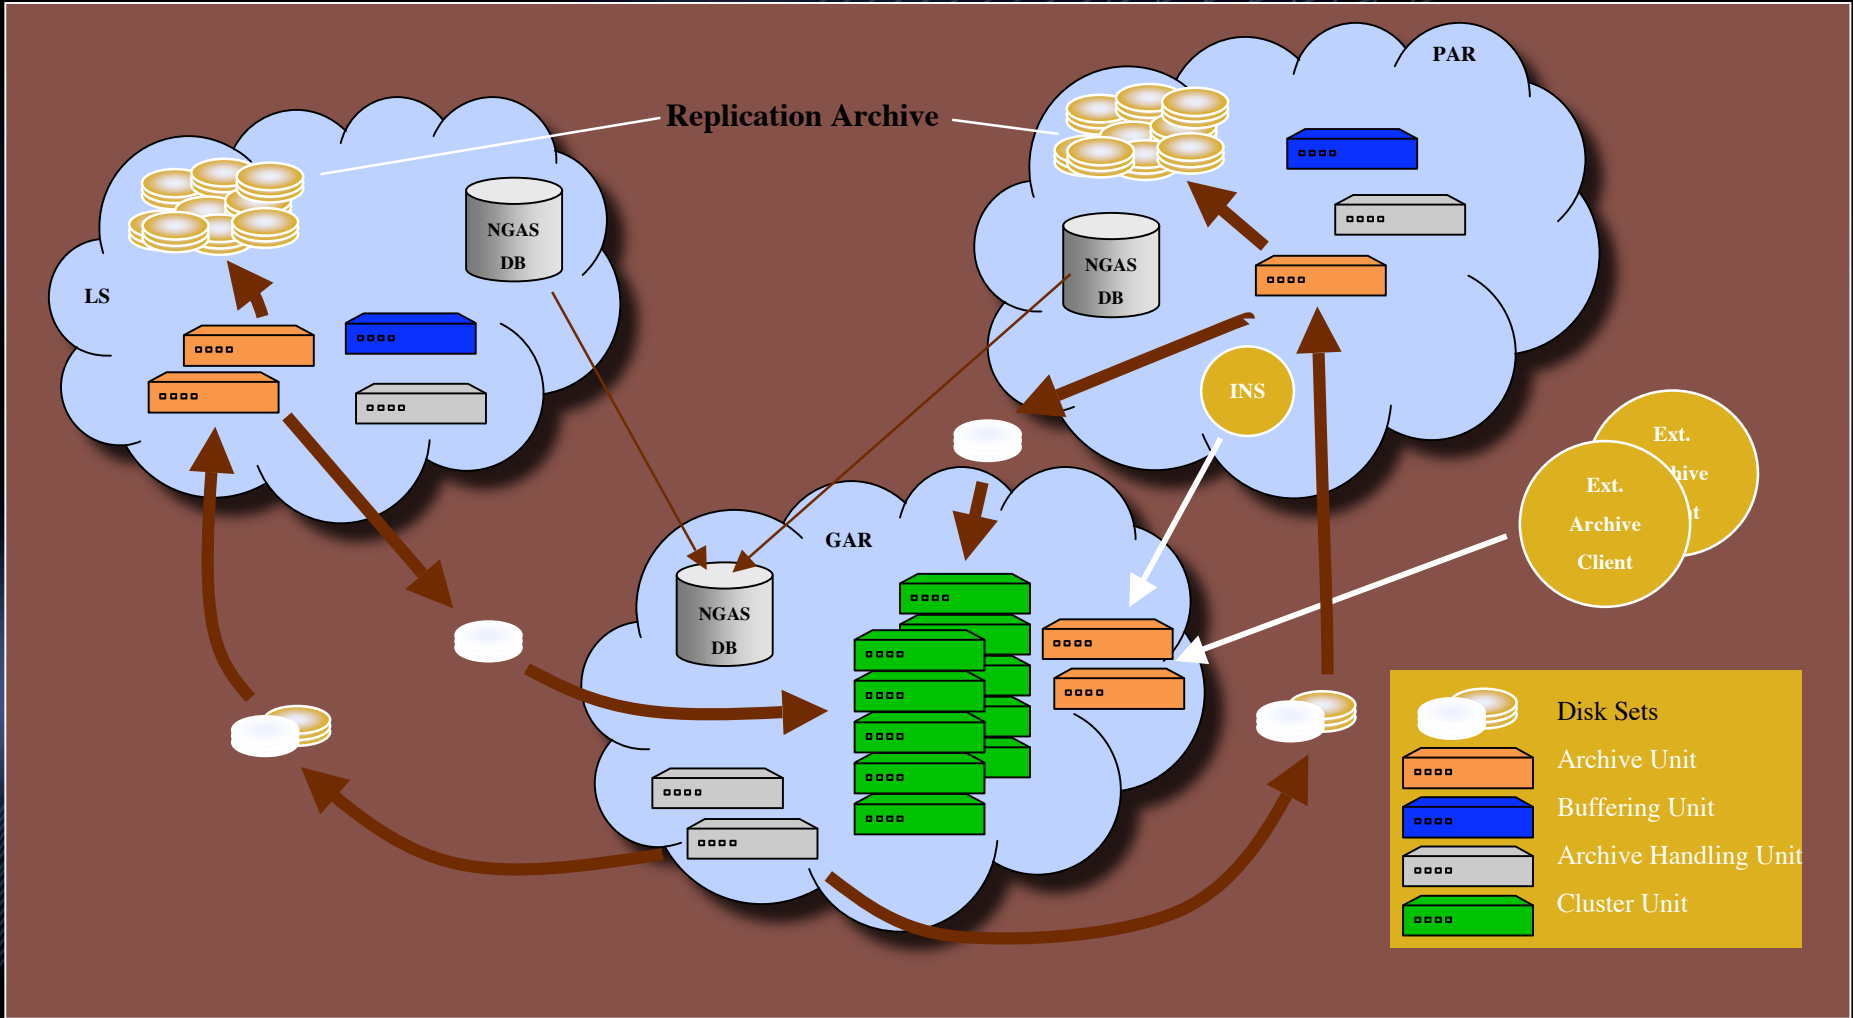
NGAS

- Handle large data streams from Telescopes
- Low Cost
 - Standard hardware
- Written in Python
 - Multi-threaded
 - Plug-in architecture
- Features
 - Communication with server over http
 - Metadata synchronization
- <http://archive.eso.org/NGAST/>

Hardware



NGAS Schematic



Hardware - some figures

- Primary Archive: currently NGAS
 - 25 NCUs x 6.4 TB/NCU RAID 5 2x XEON
 - 1 GBit network connection
- Secondary Archive: 335 SAIT tapes, each with 500 GB max.
 - currently 158 TB total capacity 4 drives, each drive 30 MB/s includes archiving software, but total is less than 4x 30 MB/s, CPU limited
 - 1 GBit network connection
- Compute Stack: 20 blades with 2x Opterons
- Fast Cache: 2x 16 x 130 GB disk RAID 0 approx. 4 TB, max 300-400 MB/s 2 controller x 2 ports x 2 GBit Fibre Channel

VOSpace Use Cases/Motivation



- In ESO there is a case for two VOSpaces
 - Read-only access on the ESO Archive.
 - General Read/Write VOSpace for users as a result of running VO services.
- The first of these is much more urgently desired than the second.

VOSpace

- Implemented the rc6 version of the VOSpace 1.0 interface (recently promoted to WD status) with NGAS as the bulk storage back end
 - <http://vo.eso.org:8080/vospace>
 - Published October 2006
 - Implements WS-security
 - Has http put/get transport with simplistic one-time password.
 - No client - except the java delegate library

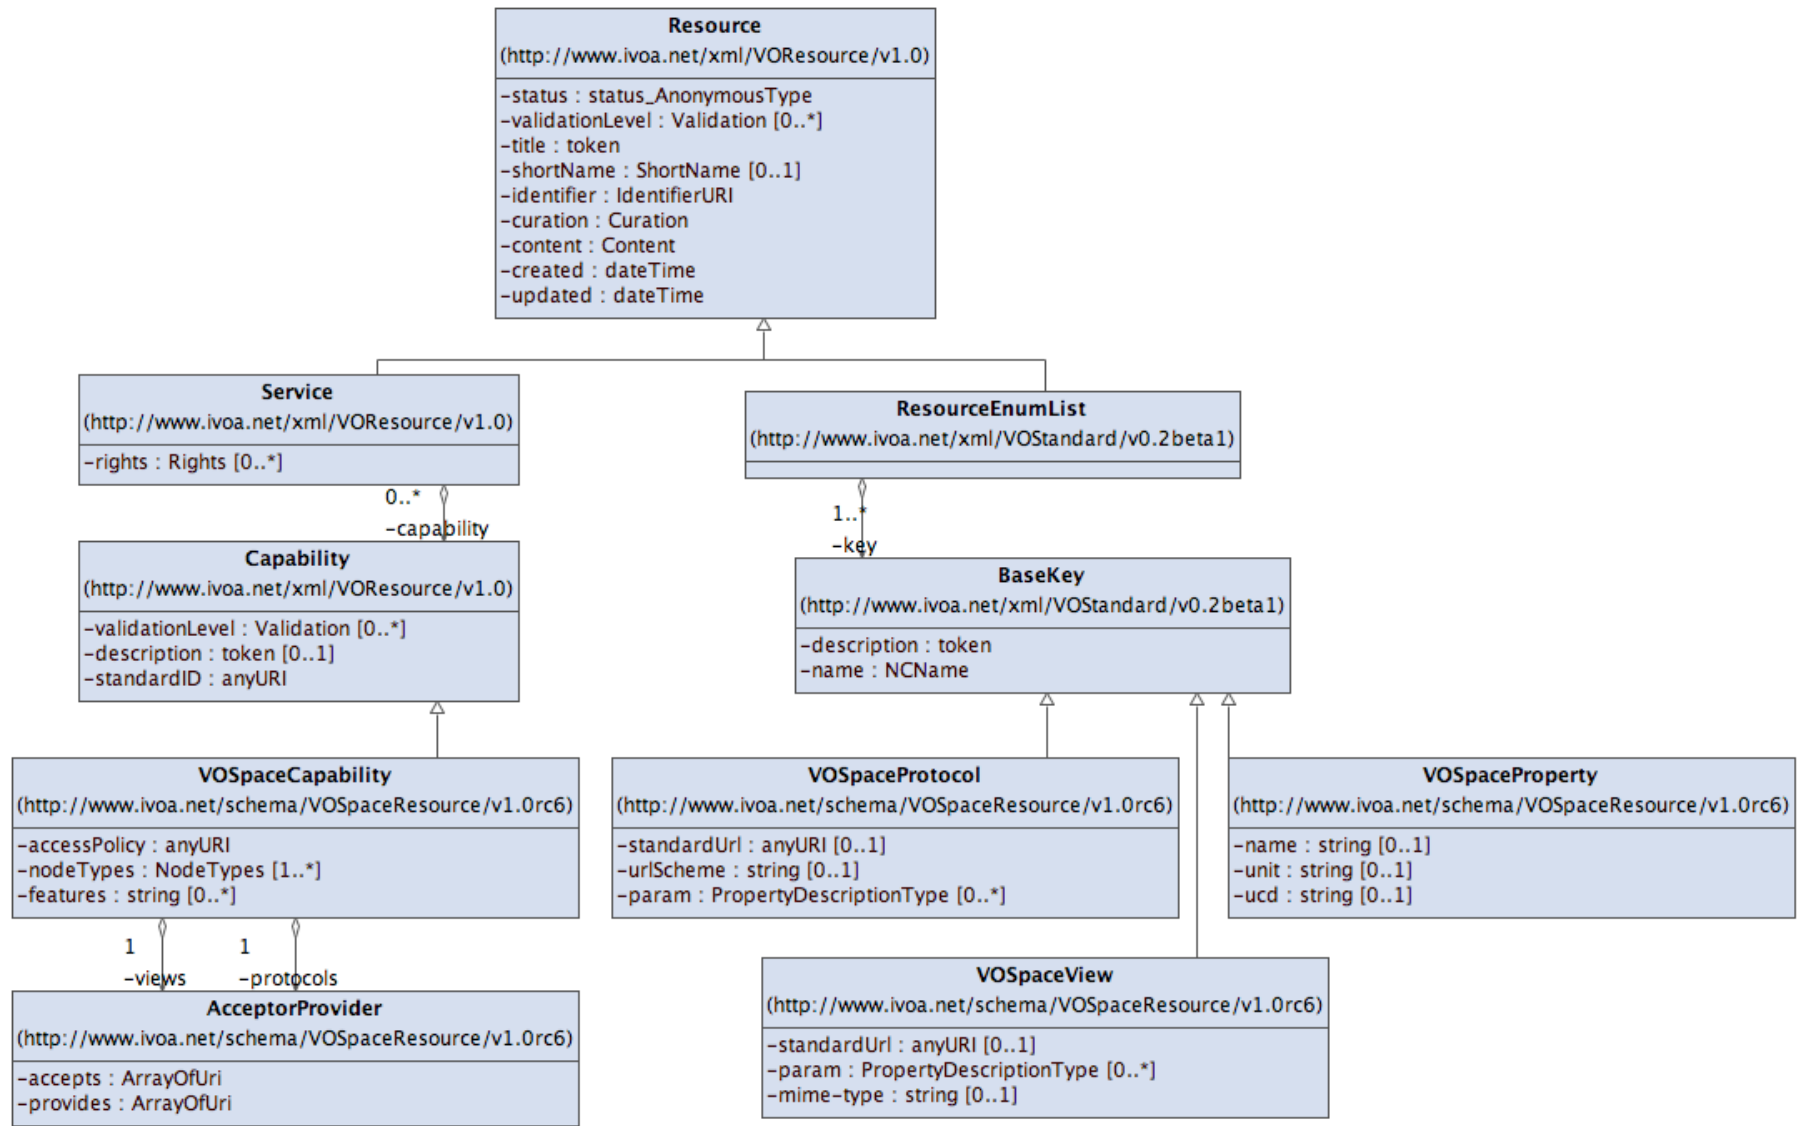
Implementation details

- VOSpace interface layered on top of NGAS
 - NGAS core software independent of VOSpace
 - Minor changes necessary to NGAS layer have been incorporated into production codebase.
- Java - J2EE - deployed in Apache Tomcat
- Axis for web services - not ideal
 - Not streaming
 - Produces “difficult to use” objects
- Interfaces directly with the Sybase NGAS metadata store
- Uses the standard NGAS http based API to access the data.

Implementation Problems

- “Flat” containerless structure makes listNodes() impractical
 - 25 million files in ESO archive.
 - A “search” interface is needed to interact directly with VOSpace
- Authentication & Authorization
 - Set up to use X509 Authentication
 - Need to be a registered ESO user - only I am configured. -> SSO interoperability.
 - Lack of individual ACLs on data holdings
 - Cannot distinguish proprietary data.
- Lack of GUI clients

VOSpace Registry Schema



Interoperability Issues

- Registration
 - Have not tested discovering other VOSpaces.
 - vos: URI scheme
 - VOSpace properties.
 - `ivo://net.ivoa.vospace/properties#size`
 - `ivo://net.ivoa.vospace/properties#owner`
 - VOSpace protocols
 - `ivo://net.ivoa.vospace/protocols#http-1.1-get`
- SSO

Conclusions

- 1.0 VOSpace “just about” sufficient for read-only archive access.
 - Authentication & Authorization - but that is SSO issue
 - Interoperability not really tested
- 1.x urgently needed for “User VOSpace”