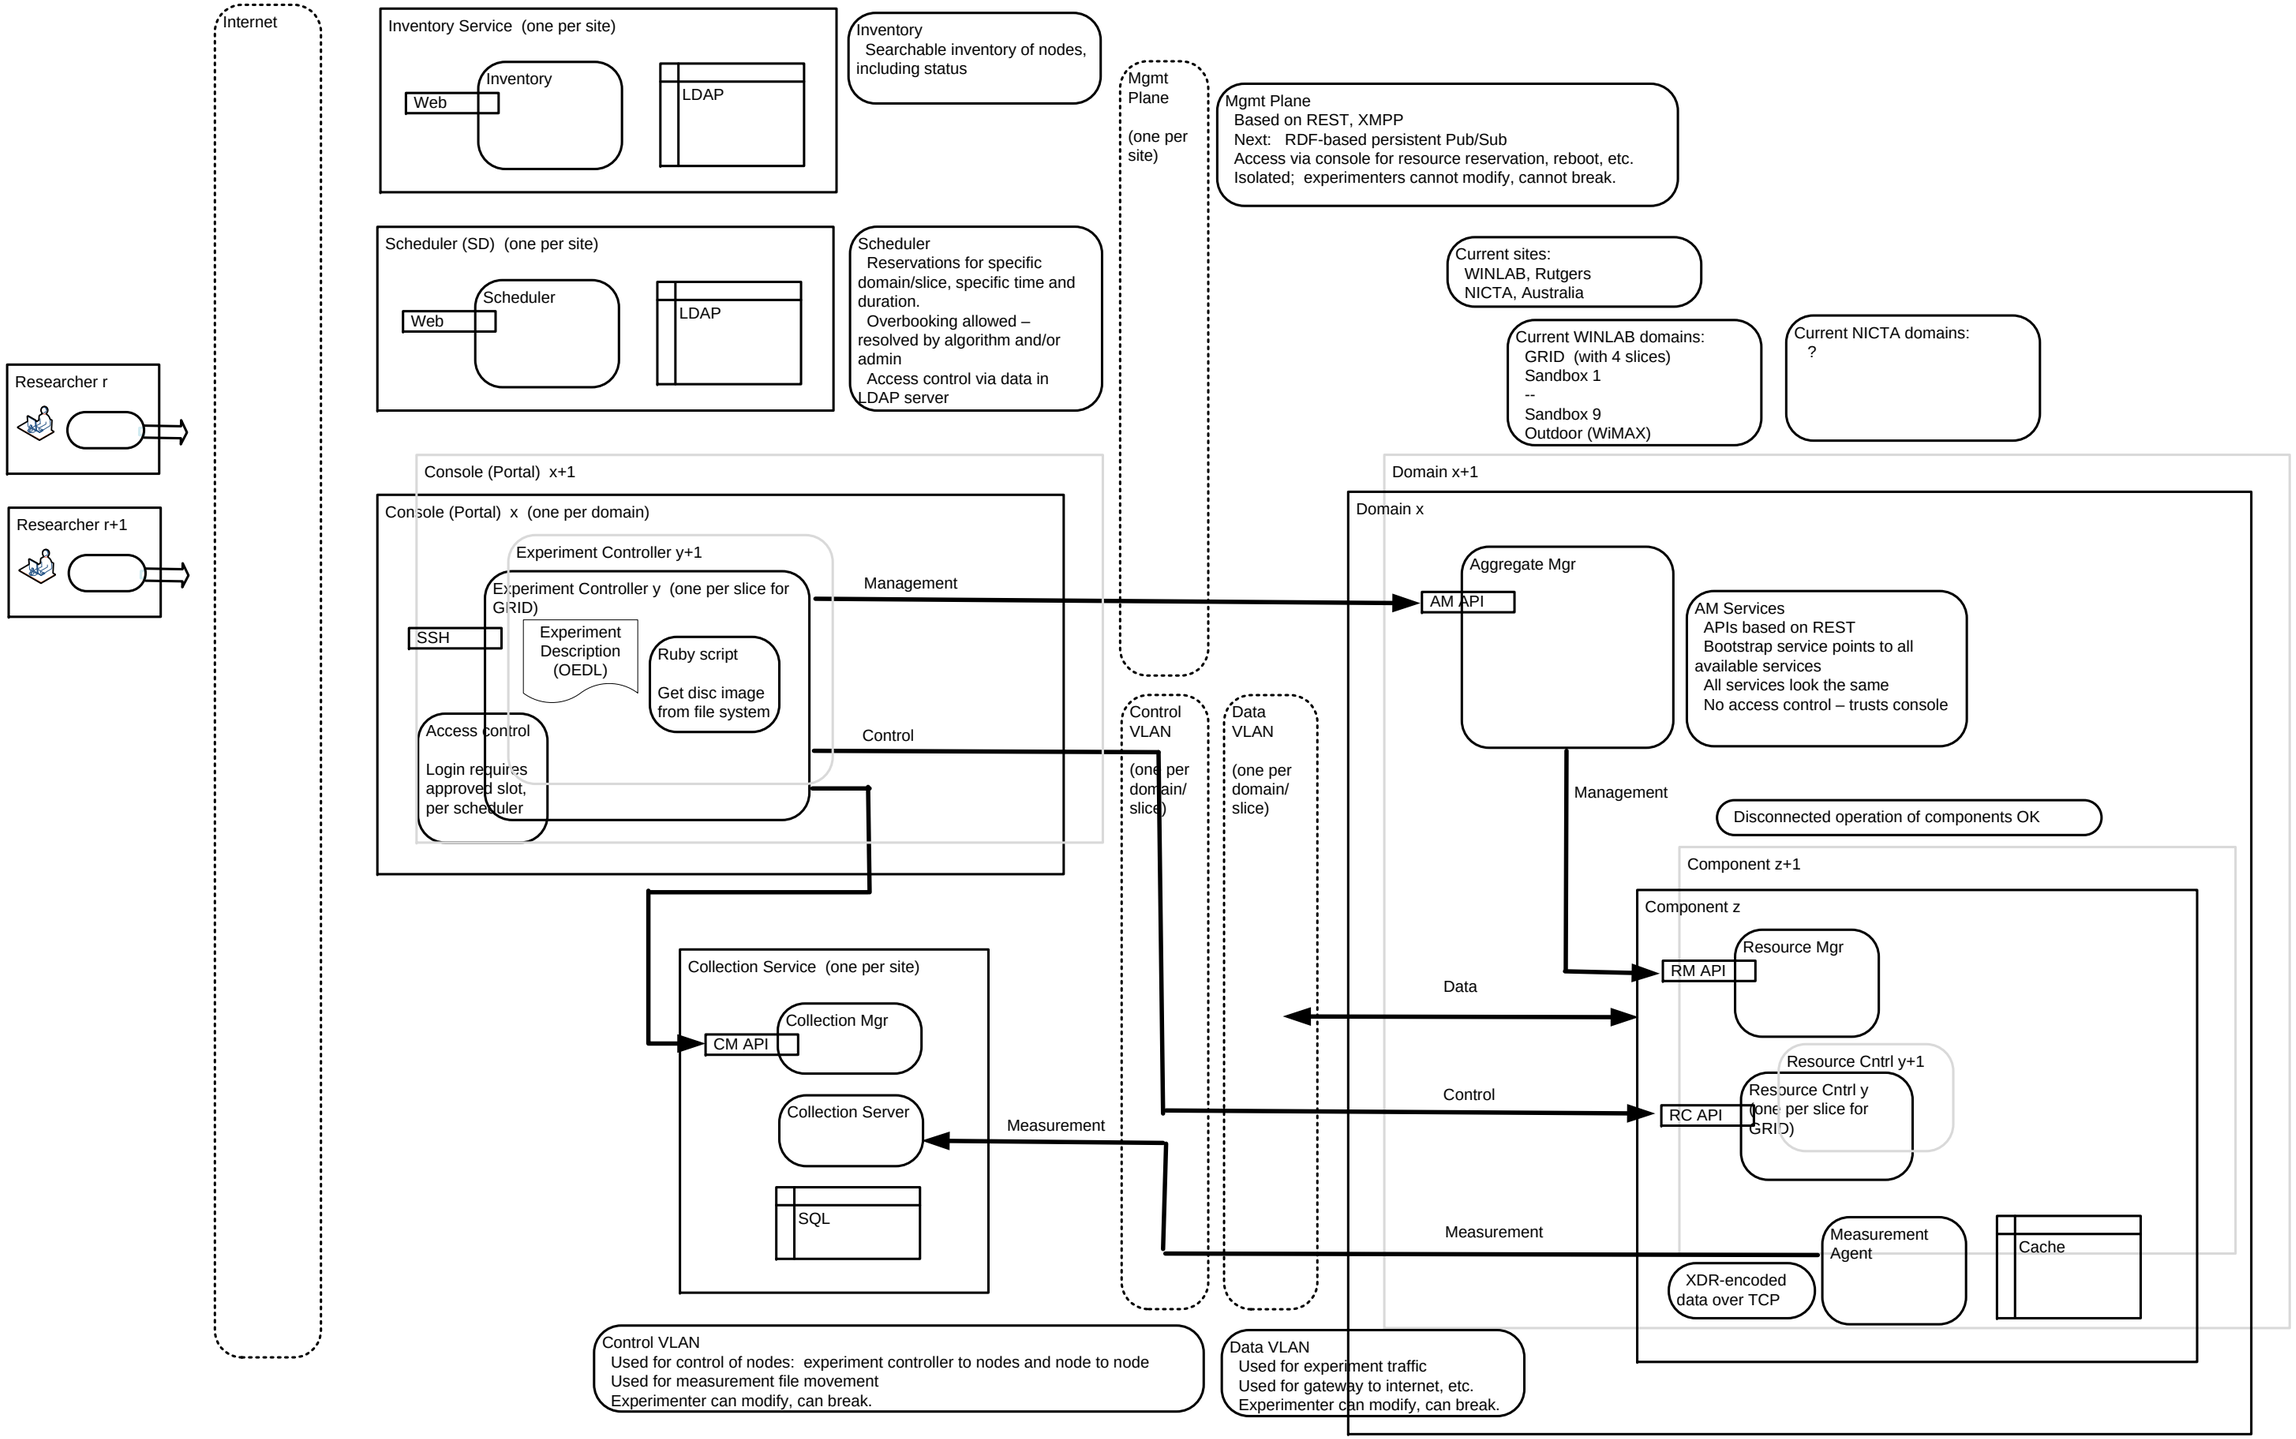
Internet
Researcher r
Researcher r+1
Inventory Service (one per site)
Web
Inventory
LDAP
Inventory
Searchable inventory of nodes, including status
Scheduler (SD) (one per site)
Web
Scheduler
LDAP
Scheduler
Reservations for specific domain/slice, specific time and duration.
Overbooking allowed – resolved by algorithm and/or admin
Access control via data in LDAP server
Mgmt Plane
(one per site)
Mgmt Plane
Based on REST, XMPP
Next: RDF-based persistent Pub/Sub
Access via console for resource reservation, reboot, etc.
Isolated; experimenters cannot modify, cannot break.
Current sites:
WINLAB, Rutgers
NICTA, Australia
Current WINLAB domains:
GRID (with 4 slices)
Sandbox 1
--
Sandbox 9
Outdoor (WiMAX)
Current NICTA domains:
?
Console (Portal) x+1
Console (Portal) x (one per domain)
Experiment Controller y+1
Experiment Controller y (one per slice for GRID)
SSH
Experiment Description (OEDL)
Ruby script
Get disc image from file system
Access control
Login requires approved slot, per scheduler
Management
Control
Control VLAN
(one per domain/ slice)
Data VLAN
(one per domain/ slice)
Domain x+1
Domain x
Aggregate Mgr
AM API
AM Services
APIs based on REST
Bootstrap service points to all available services
All services look the same
No access control – trusts console
Management
Disconnected operation of components OK
Component z+1
Component z
Resource Mgr
RM API
Data
Resource Cntrl y+1
Resource Cntrl y (one per slice for GRID)
RC API
Control
Collection Service (one per site)
Collection Mgr
CM API
Collection Server
SQL
Measurement
Measurement
Measurement Agent
Cache
XDR-encoded data over TCP
Control VLAN
Used for control of nodes: experiment controller to nodes and node to node
Used for measurement file movement
Experimenter can modify, can break.
Data VLAN
Used for experiment traffic
Used for gateway to internet, etc.
Experimenter can modify, can break.

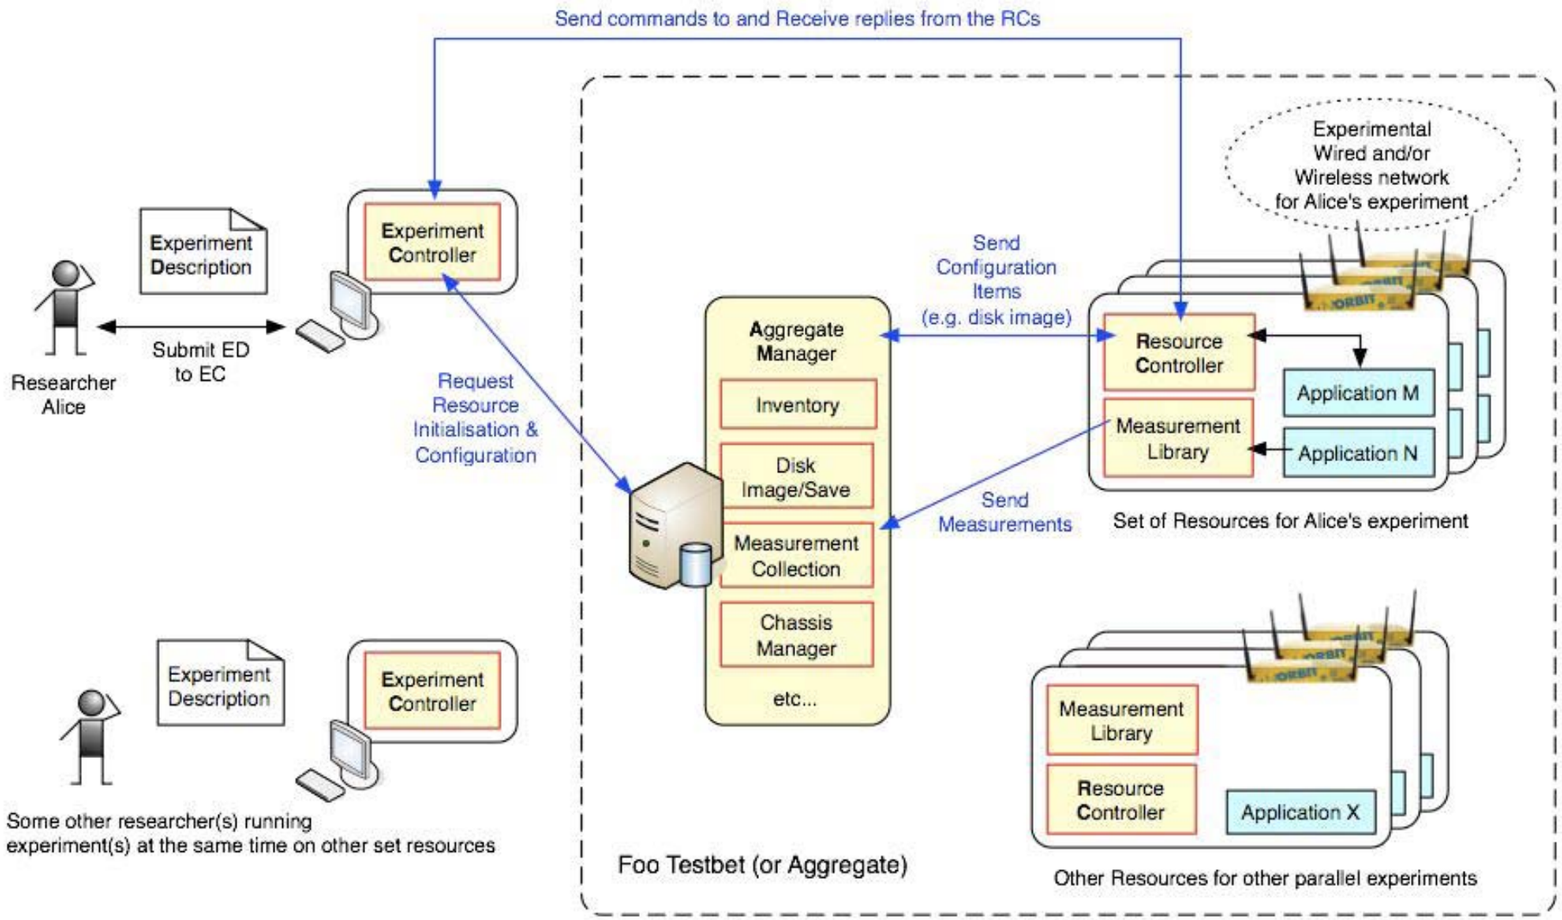
Send commands to and Receive replies from the RCs
Experiment Description
Experiment Controller
Researcher Alice
Submit ED to EC
Request Resource Initialisation & Configuration
Send Configuration Items (e.g. disk image)
Aggregate Manager
Inventory
Disk Image/Save
Measurement Collection
Chassis Manager
etc...
Experimental Wired and/or Wireless network for Alice's experiment
Resource Controller
Application M
Measurement Library
Application N
Send Measurements
Set of Resources for Alice's experiment
Measurement Library
Resource Controller
Application X
Other Resources for other parallel experiments
Foo Testbet (or Aggregate)
Experiment Description
Experiment Controller
Some other researcher(s) running experiment(s) at the same time on other set resources

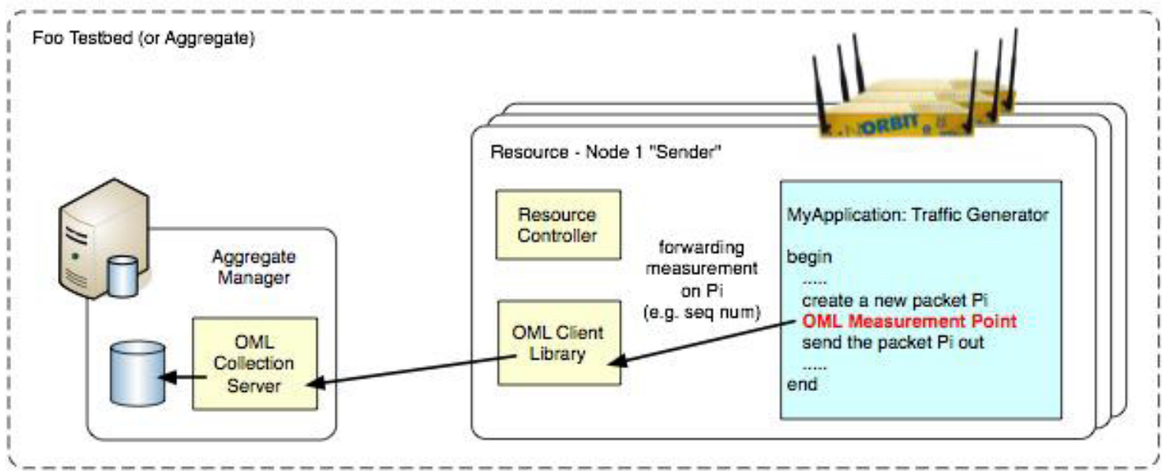
Foo Testbed (or Aggregate)
Resource - Node 1 "Sender"
Resource Controller
forwarding measurement on Pi (e.g. seq num)
MyApplication: Traffic Generator
begin
.....
create a new packet Pi
OML Measurement Point
send the packet Pi out
.....
end
OML Client Library
Aggregate Manager
OML Collection Server

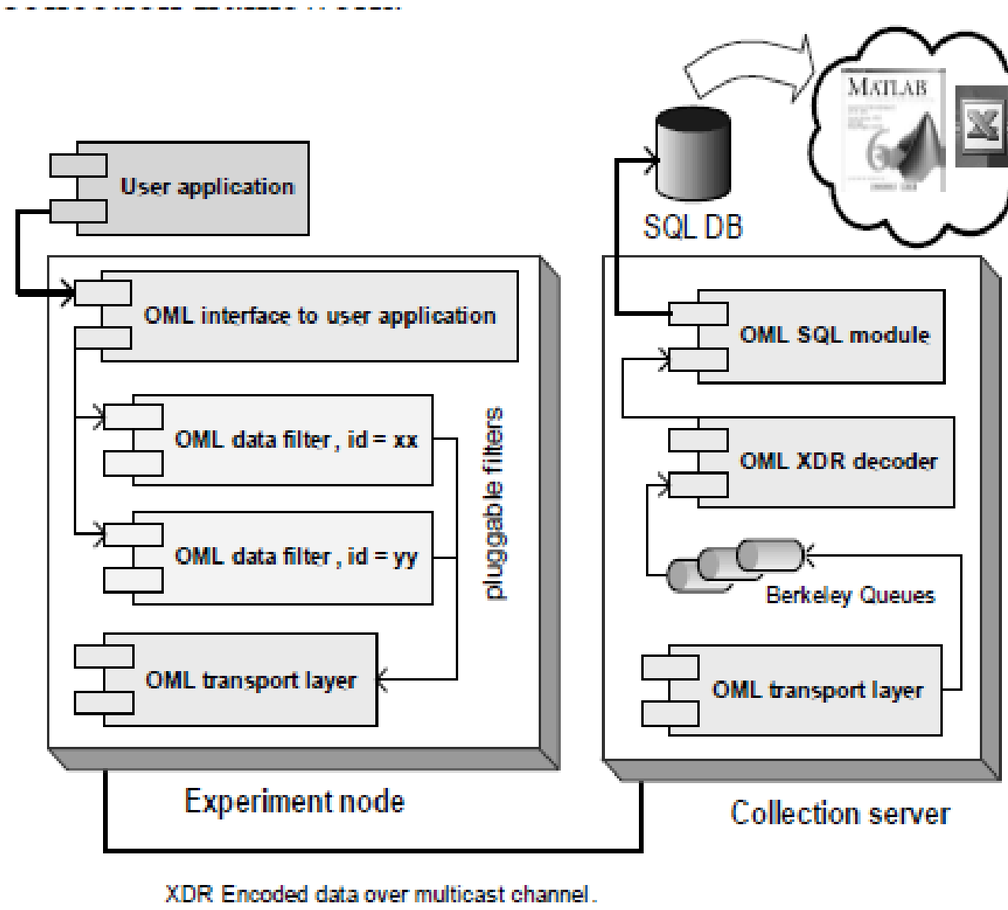

Figure 1. OML component architecture

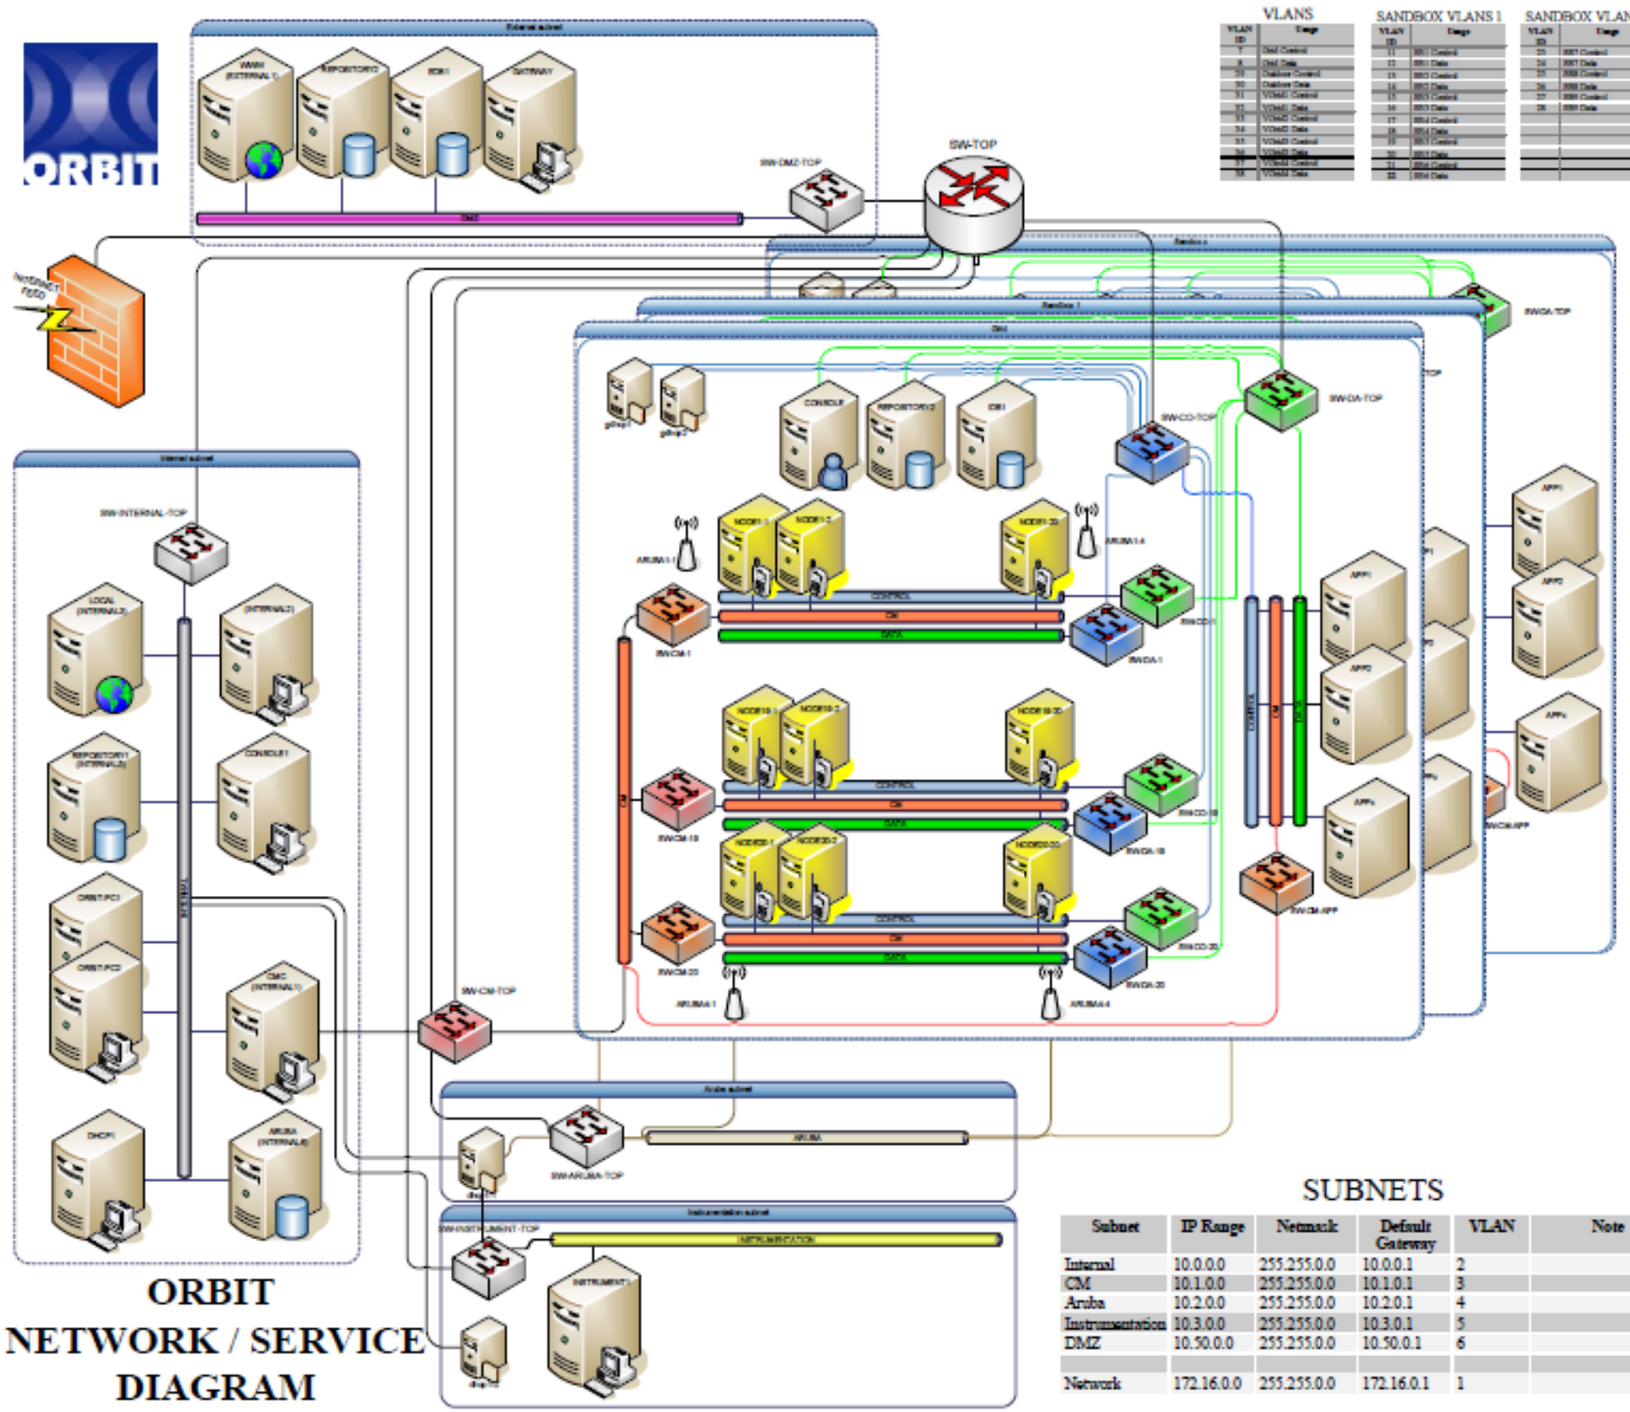

SUBNETS

Subnet	IP Range	Netmask	Default Gateway	VLAN	Note
Internal	10.0.0.0	255.255.0.0	10.0.0.1	2	
CM	10.1.0.0	255.255.0.0	10.1.0.1	3	
Aruba	10.2.0.0	255.255.0.0	10.2.0.1	4	
Instrumentation	10.3.0.0	255.255.0.0	10.3.0.1	5	
DMZ	10.50.0.0	255.255.0.0	10.50.0.1	6	
Network	172.16.0.0	255.255.0.0	172.16.0.1	1	

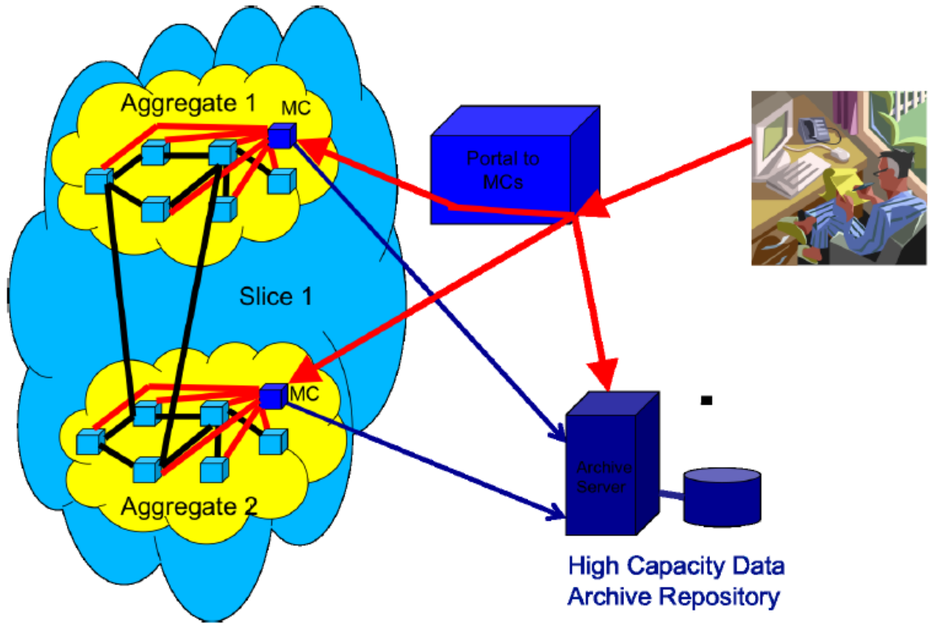

Figure 2: The MC Portal provides users with a single point of entry to their measurement data.

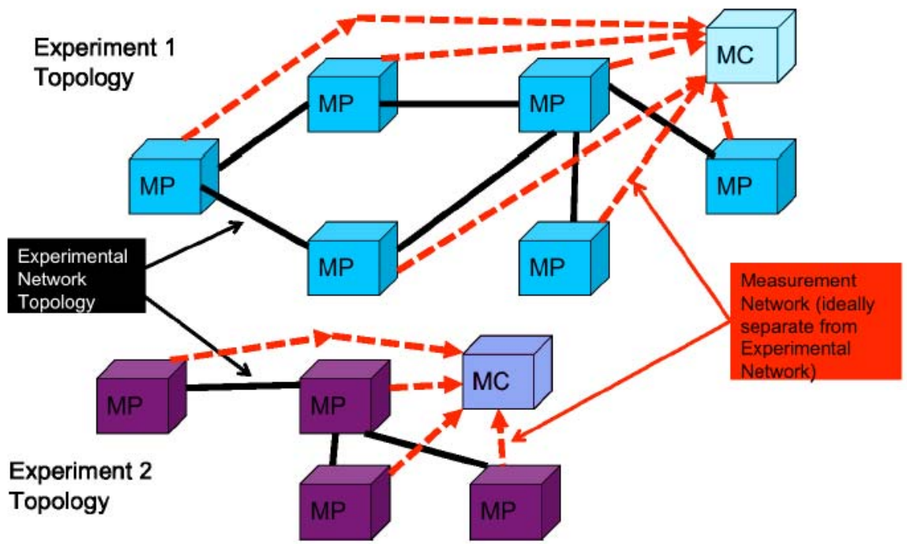

Figure 1: Each experiment/slice has its own MC and instrumentation and measurement network.

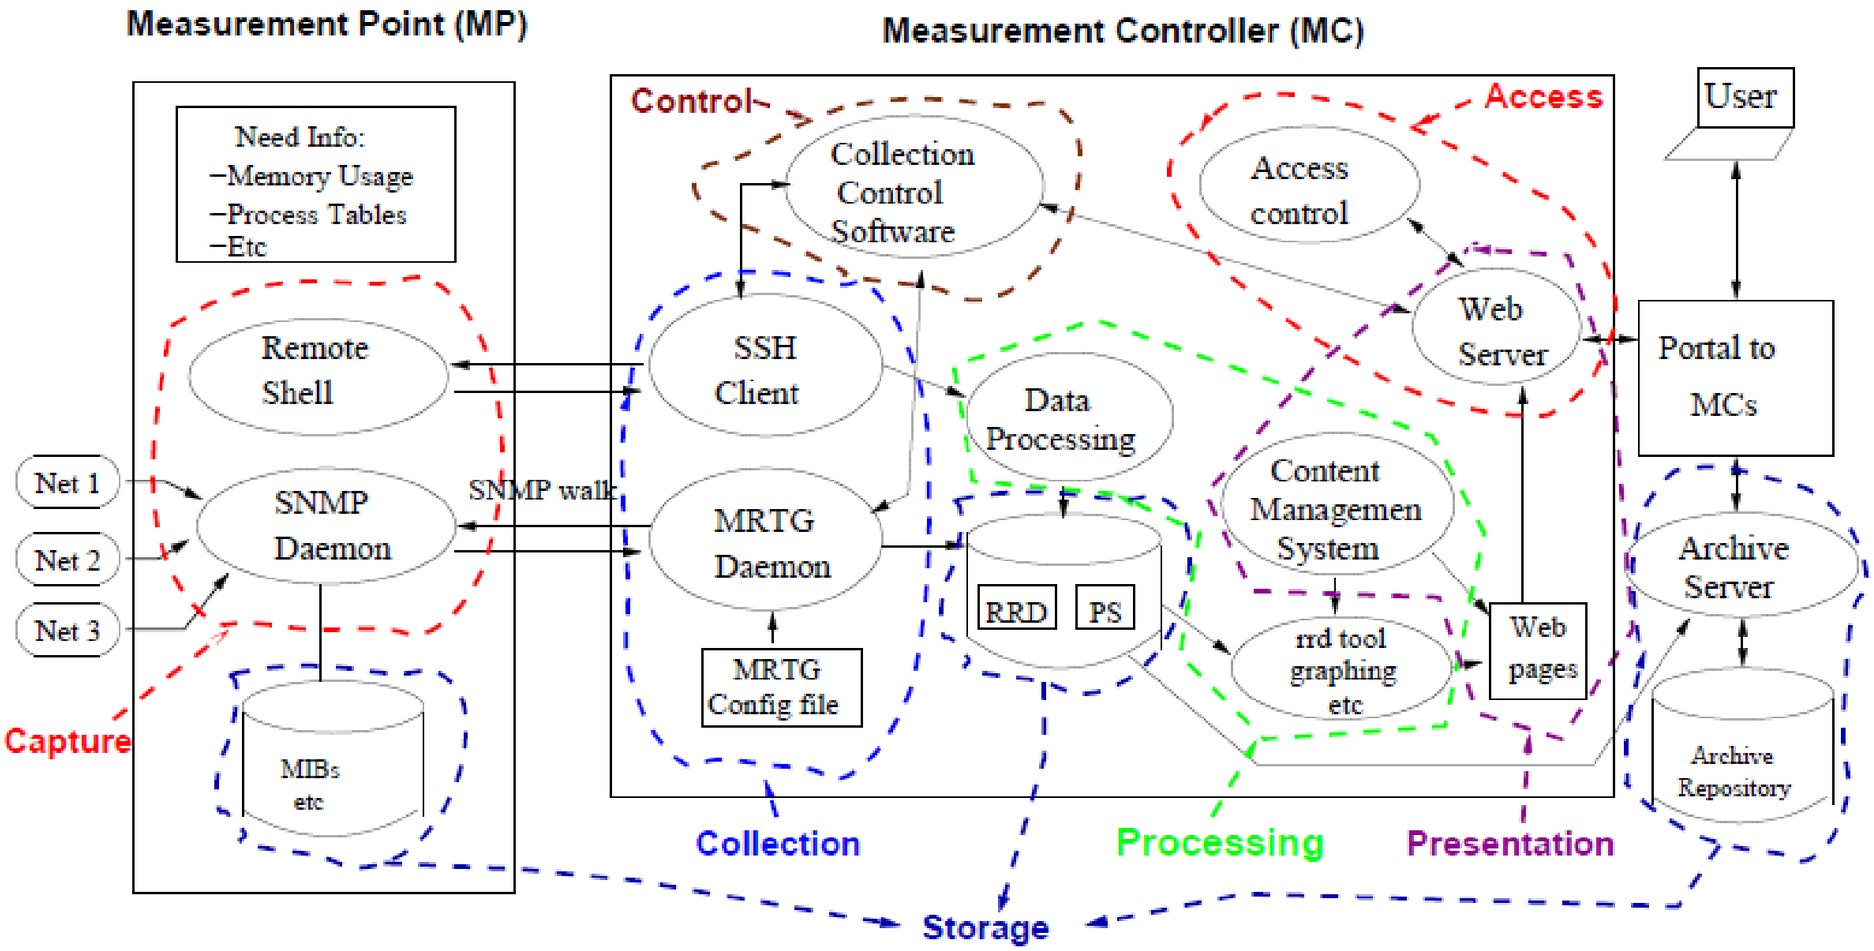

Figure 3: The Architectural Components of the INSTOOLS Toolset

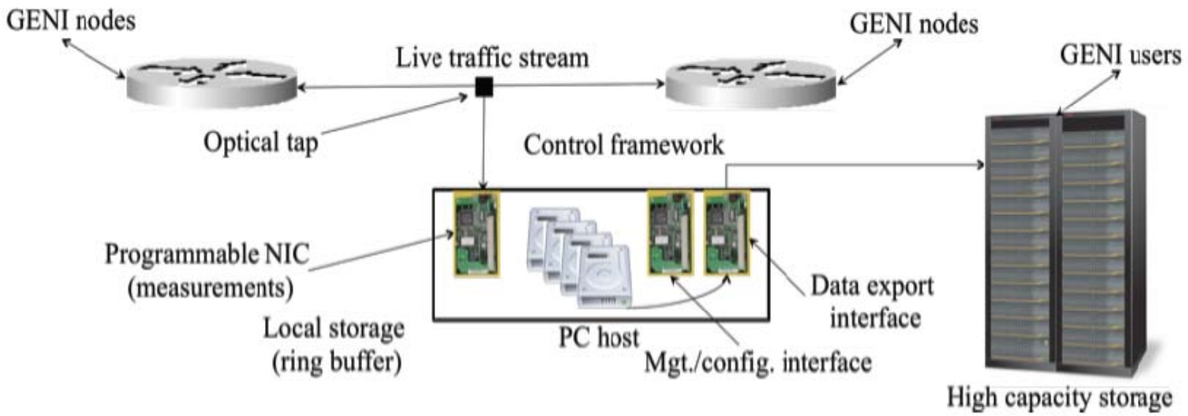

Figure 1. Basic Physical Components for Instrumentation and Measurement.

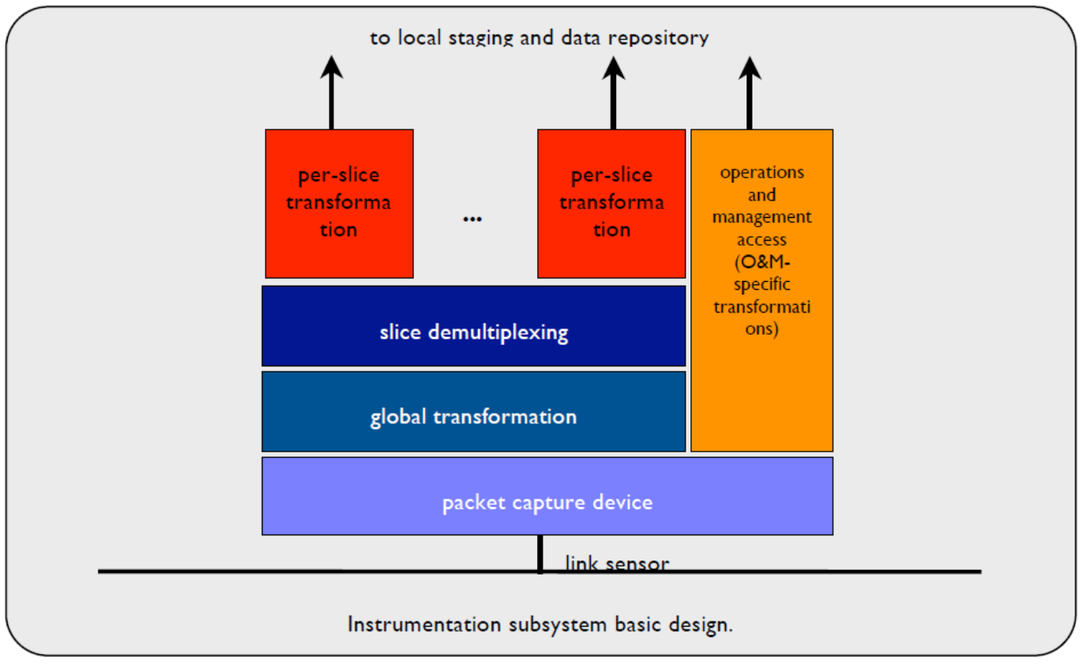

Instrumentation subsystem design block diagram.

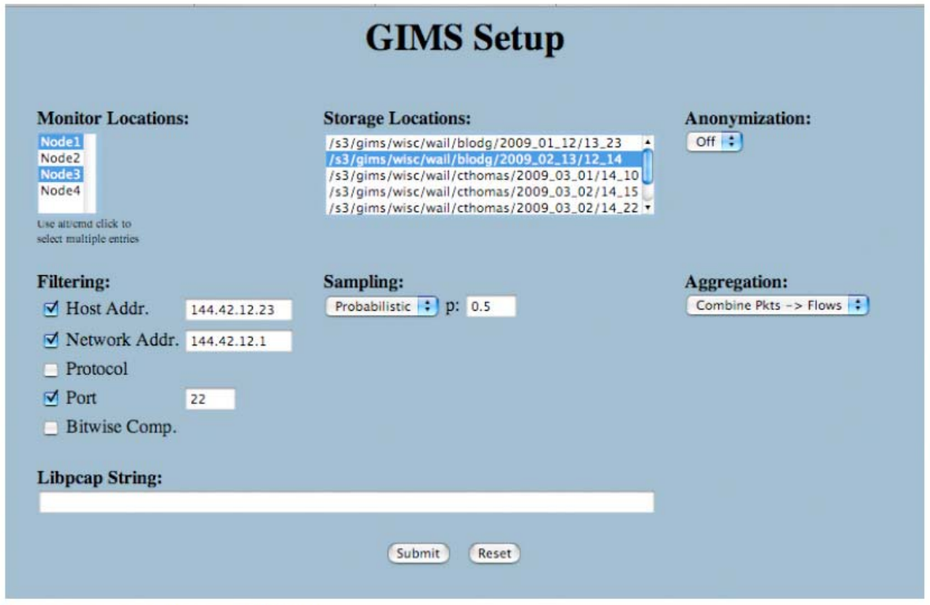

Figure 4. Mockup of GIMS UI

perfSONAR Architecture

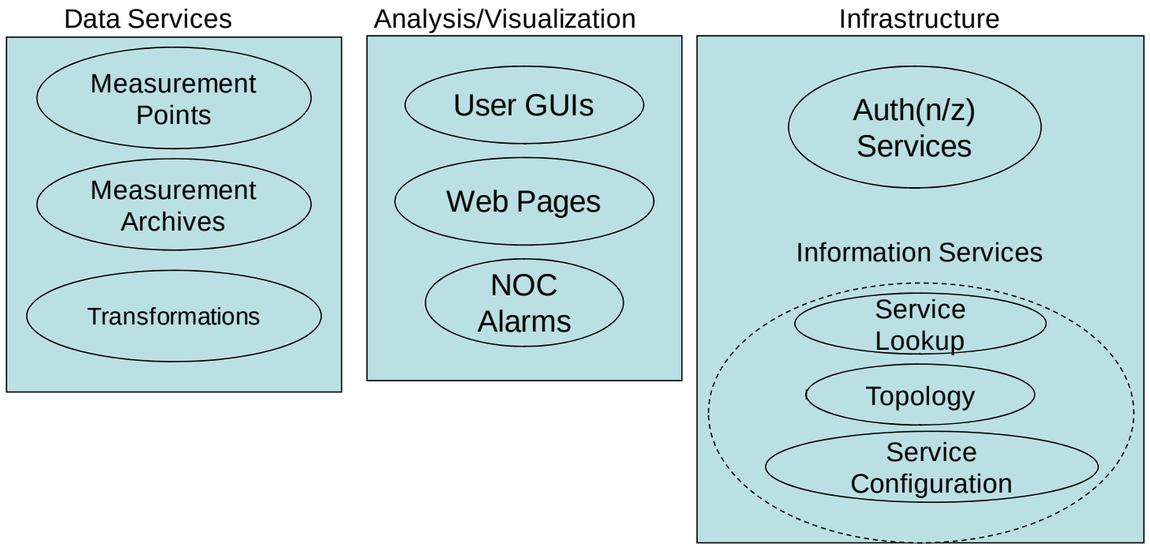

Example perfSonar client interaction

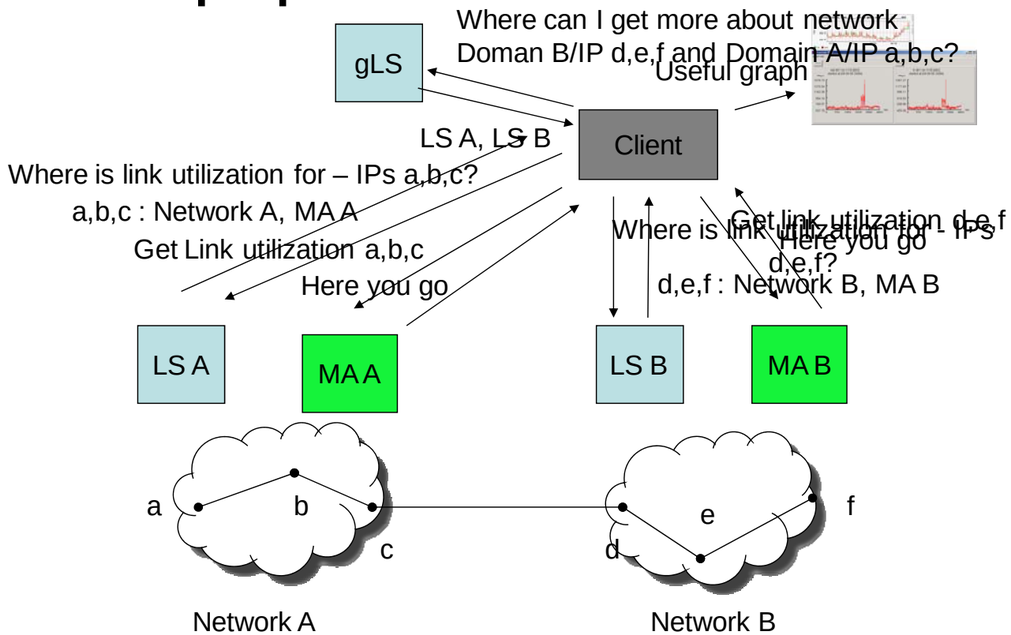

perfSONAR Client Developments

- Most tuned to specific services currently
- Different user focus (micro vs macro view)

- Client applications
  - perfSONAR-UI (acad.bg)
- Web Based
  - GMAPS (SLAC)
  - Domain Utilization Browser (ESnet)
  - pS-PS Weathermap (Internet2)
  - pingER Analysis (FNAL)
  - perfAdmin (Internet2)
  - E2EMon (DFN)

Sensor Pod

Web-Service (WS) enabled collection of sensors

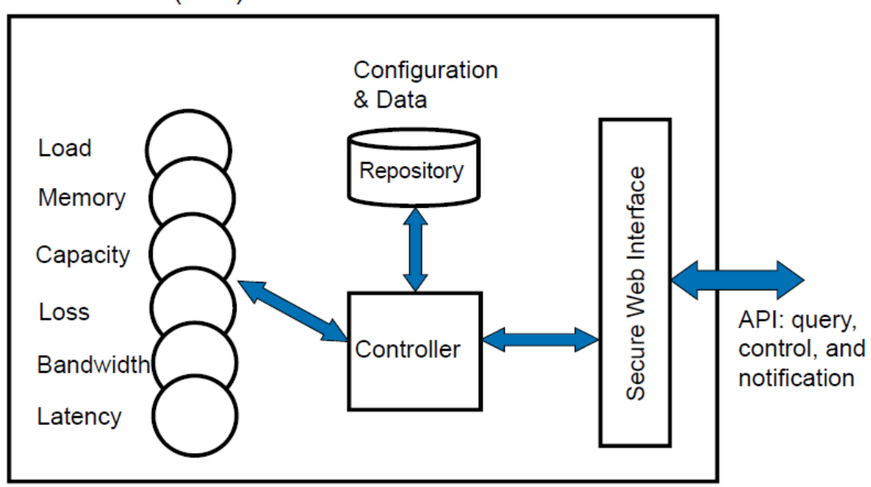

Sensing Backplane

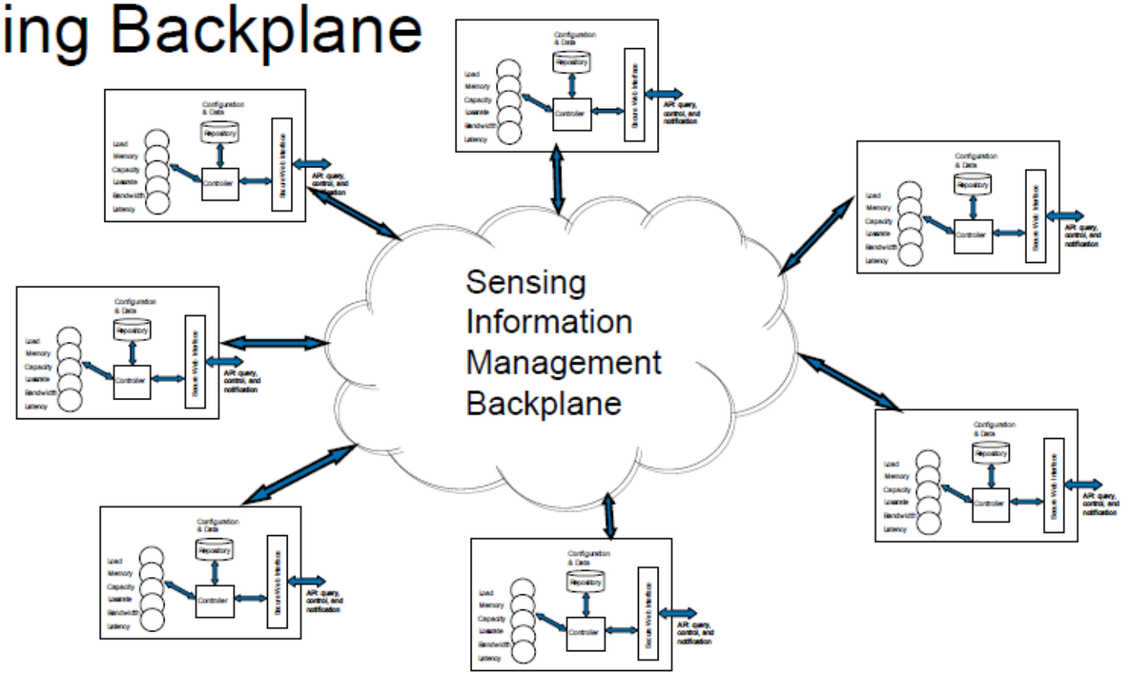

Scalable Inference Engines

- Large overhead for probing and data exchange
  - $O(N^2)$ measurements in a network of N nodes
  - Dynamically changing ➔ Need frequent probing
- Measurement/Monitoring failures
  - Failed or slow end machines
  - Measurement tool failures
- Inference based on incomplete information
  - Exploit properties such as triangular inequality
  - A coarse estimate may suffice for many applications
- Prediction based on archived information
- Tradeoff between accuracy and overhead
- When and where to use inference? [Blanton et al., ICDCS09]

- Aggregate data from end-points
- Configurable and self-managing
- E.g., SDIMS [SIGCOMM 2004]

6

PURDUE UNIVERSITY

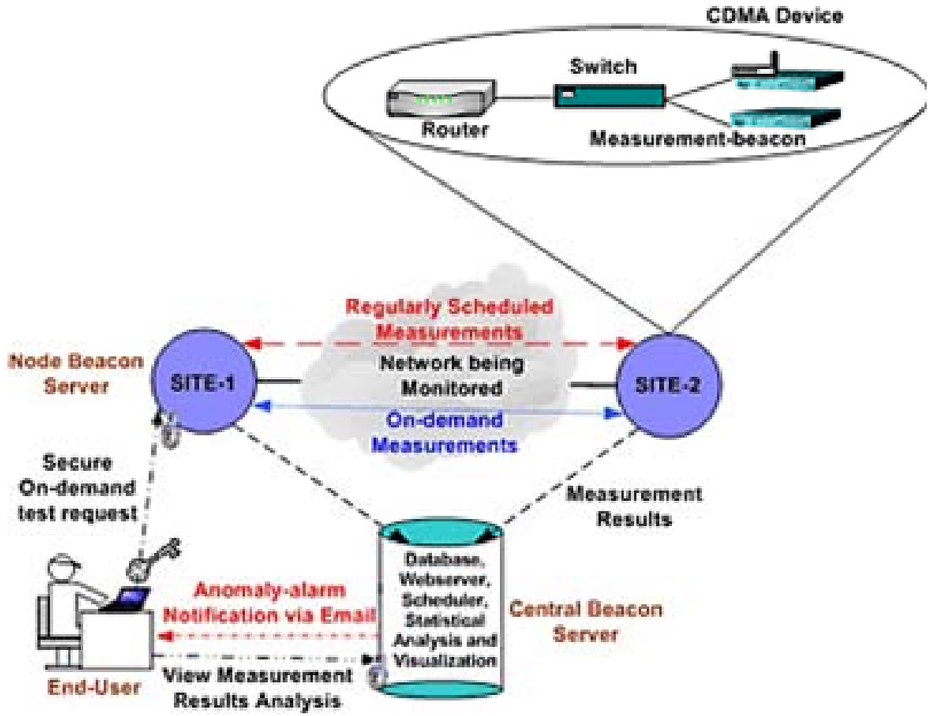
CDMA Device
Switch
Router
Measurement-beacon
Regularly Scheduled Measurements
Node Beacon Server
SITE-1
Network being Monitored
SITE-2
On-demand Measurements
Secure On-demand test request
Measurement Results
Database, Webserver, Scheduler, Statistical Analysis and Visualization
Anomaly-alarm Notification via Email
Central Beacon Server
End-User
View Measurement Results Analysis

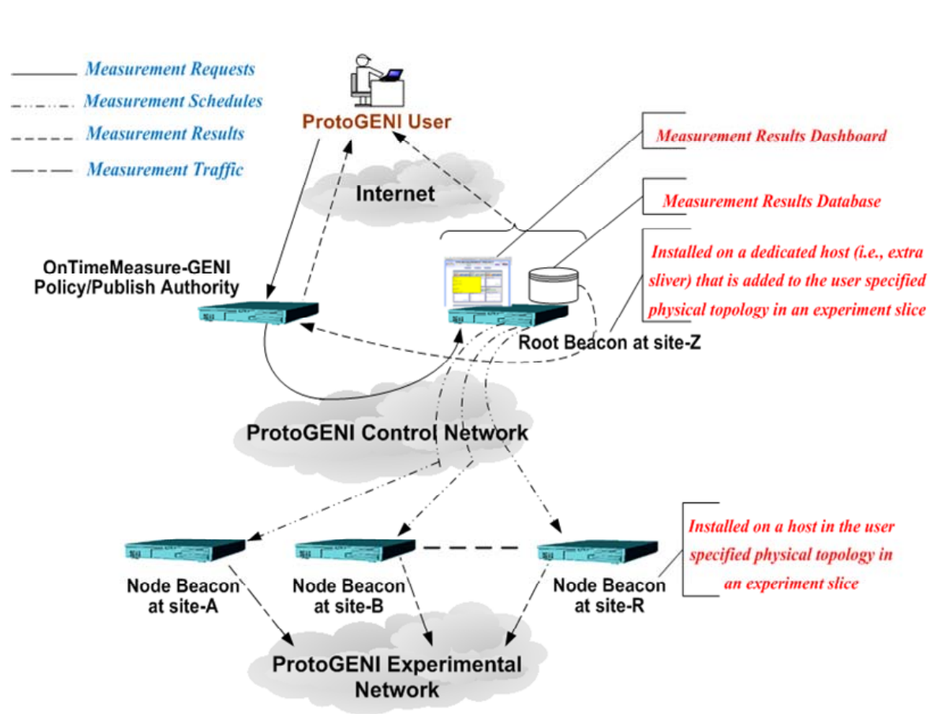
Measurement Requests
Measurement Schedules
Measurement Results
Measurement Traffic
ProtoGENI User
Internet
Measurement Results Dashboard
Measurement Results Database
Installed on a dedicated host (i.e., extra sliver) that is added to the user specified physical topology in an experiment slice
OnTimeMeasure-GENI Policy/Publish Authority
Root Beacon at site-Z
ProtoGENI Control Network
Installed on a host in the user specified physical topology in an experiment slice
Node Beacon at site-A
Node Beacon at site-B
Node Beacon at site-R
ProtoGENI Experimental Network

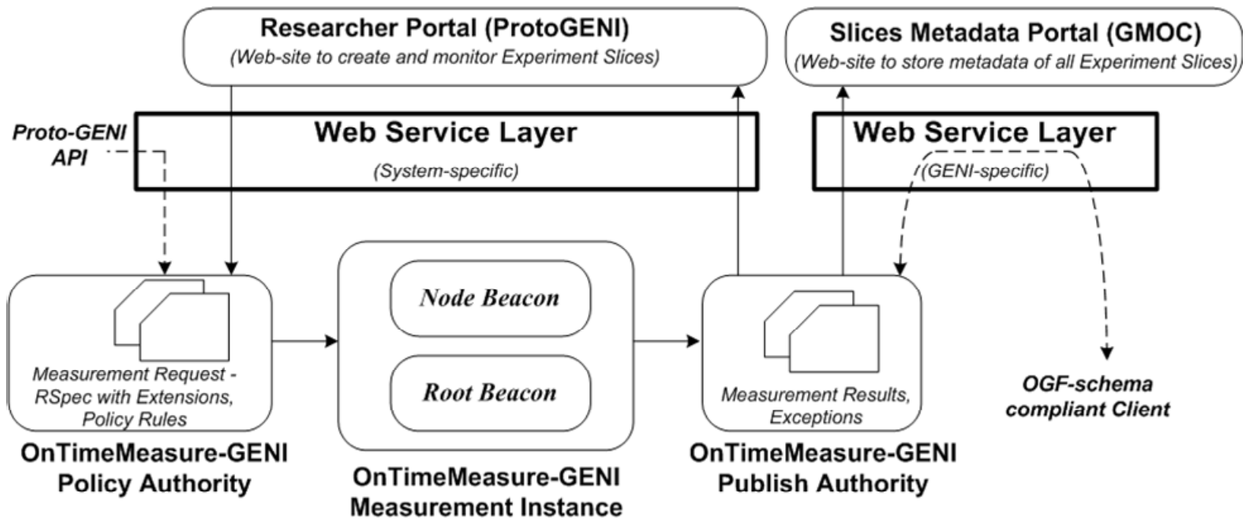
Researcher Portal (ProtoGENI)
(Web-site to create and monitor Experiment Slices)
Slices Metadata Portal (GMOC)
(Web-site to store metadata of all Experiment Slices)
Proto-GENI API
Web Service Layer
(System-specific)
Web Service Layer
(GENI-specific)
Node Beacon
Root Beacon
Measurement Request - RSpec with Extensions, Policy Rules
Measurement Results, Exceptions
OGF-schema compliant Client
OnTimeMeasure-GENI Policy Authority
OnTimeMeasure-GENI Measurement Instance
OnTimeMeasure-GENI Publish Authority

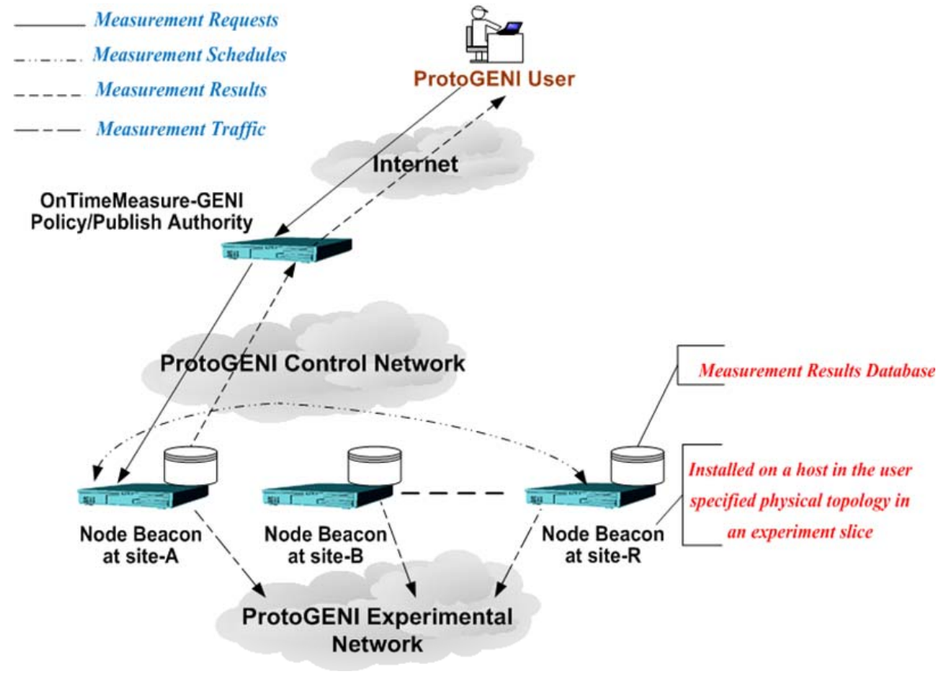
Measurement Requests
Measurement Schedules
Measurement Results
Measurement Traffic
ProtoGENI User
Internet
OnTimeMeasure-GENI Policy/Publish Authority
ProtoGENI Control Network
Measurement Results Database
Installed on a host in the user specified physical topology in an experiment slice
Node Beacon at site-A
Node Beacon at site-B
Node Beacon at site-R
ProtoGENI Experimental Network

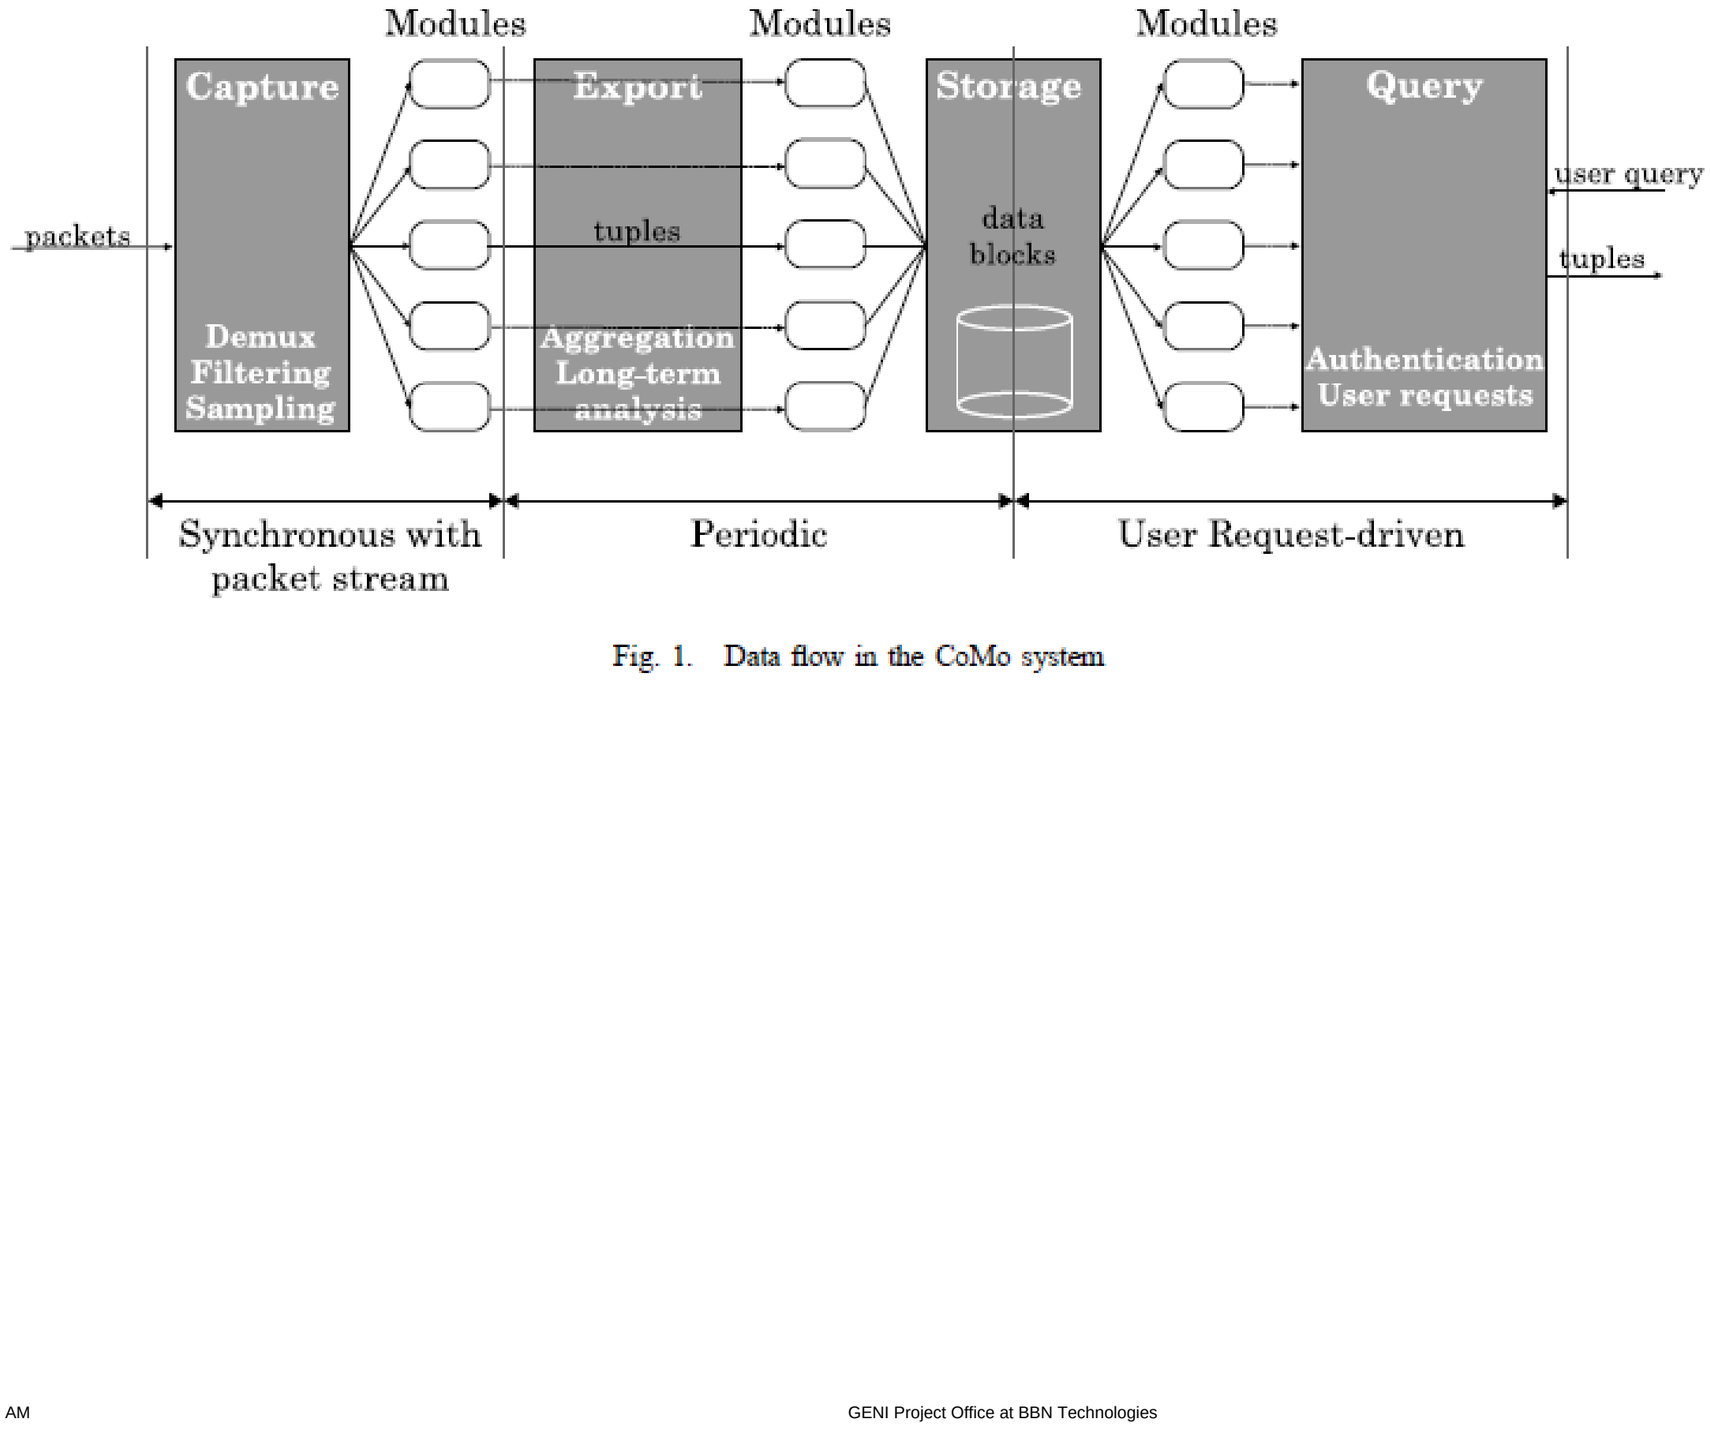

Fig. 1. Data flow in the CoMo system