


Instrumentation Tools for Researchers (Students)

- * Each instance contained in one Aggregate (site)
- * Integrated with protoGENI CF
- * Meas Orch (MO) svc:
 - * Adds MO svc to protoGENI Exp Cntrl Svc that runs in Researcher's browser
 - * When Res specifies RSpec for experiment slice, it adds host for *Meas Controller* (MC) sliver
 - * Uses info from manifest to identify slivers to be monitored, and dynamically creates config files for SNMP daemon and other capture software
 - * Then, it downloads software for *MC* sliver, and it adds *MP* svc to each *Experiment Sliver*



Manage Svcs

- * *MC Svc* has collection control software
- * *MP Svc* includes remote access daemon to execute capture software
- * How?

MP Svc is automatically loaded into each Experiment Sliver

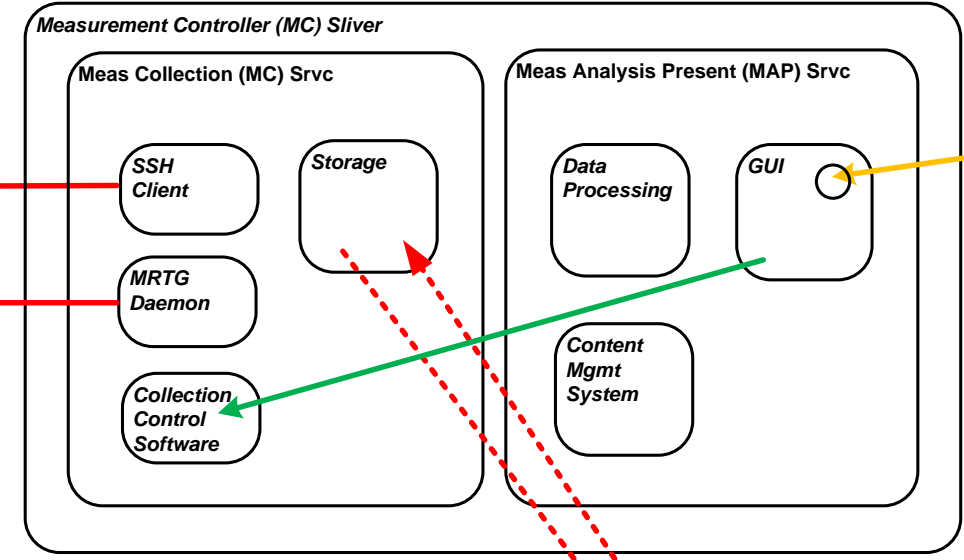
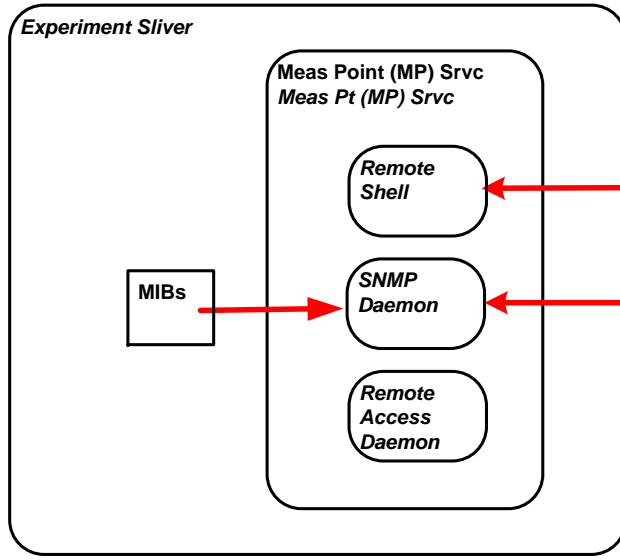
MC Sliver is automatically loaded into each Experiment Slice

Authorization

- * Emulab (ssh) key distribution mechanism used to authorize MC to get data from MPs

Security Concern

- * SNMP walk of MIBs



Software	Information	Display Type
SNMPd	Routing Table	Table
	IP Traffic	Graph
	ICMP Traffic	Graph
	TCP Traffic	Graph
	UDP Traffic	Graph
	CPU Utilization	Graph
	Memory Utilization	Graph
	Total Network Traffic	Graph
	Link-specific Traffic	Graph
	Link-specific Unicast Traffic	Graph
ssh/arp	ARP cache	Table
ssh/netstat	TCP streams	Table
ssh/netstat	UDP listeners	Table
ssh/ps	Process list	Table
ssh/lsmod	Installed Kernel Modules	Table

Portal to MCs

- * How does this work?

MP Svc includes:

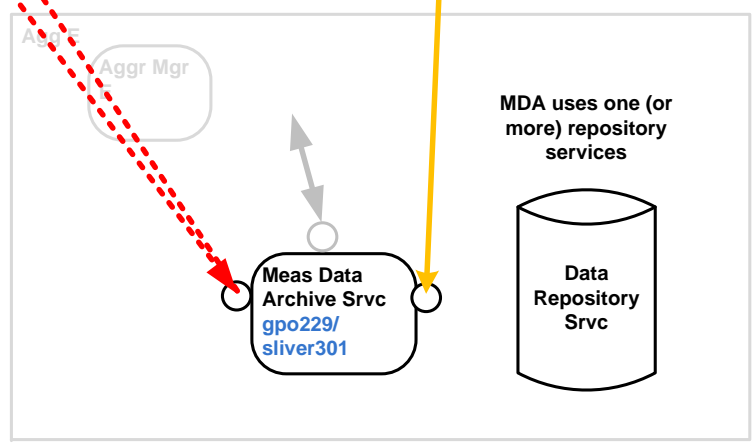
- * SNMP daemon, using existing MIBs, or added MIBs
- * tcpdump
- * netflow
- * custom monitoring code based on pcap library to collect packet stats not captured by SNMP daemon
- * ps, vmstat, and SNMP, to capture OS info, such as CPU load, memory load, routing table configs, ARP caches, loaded modules, etc.
- * remote access daemon to to execute capture software

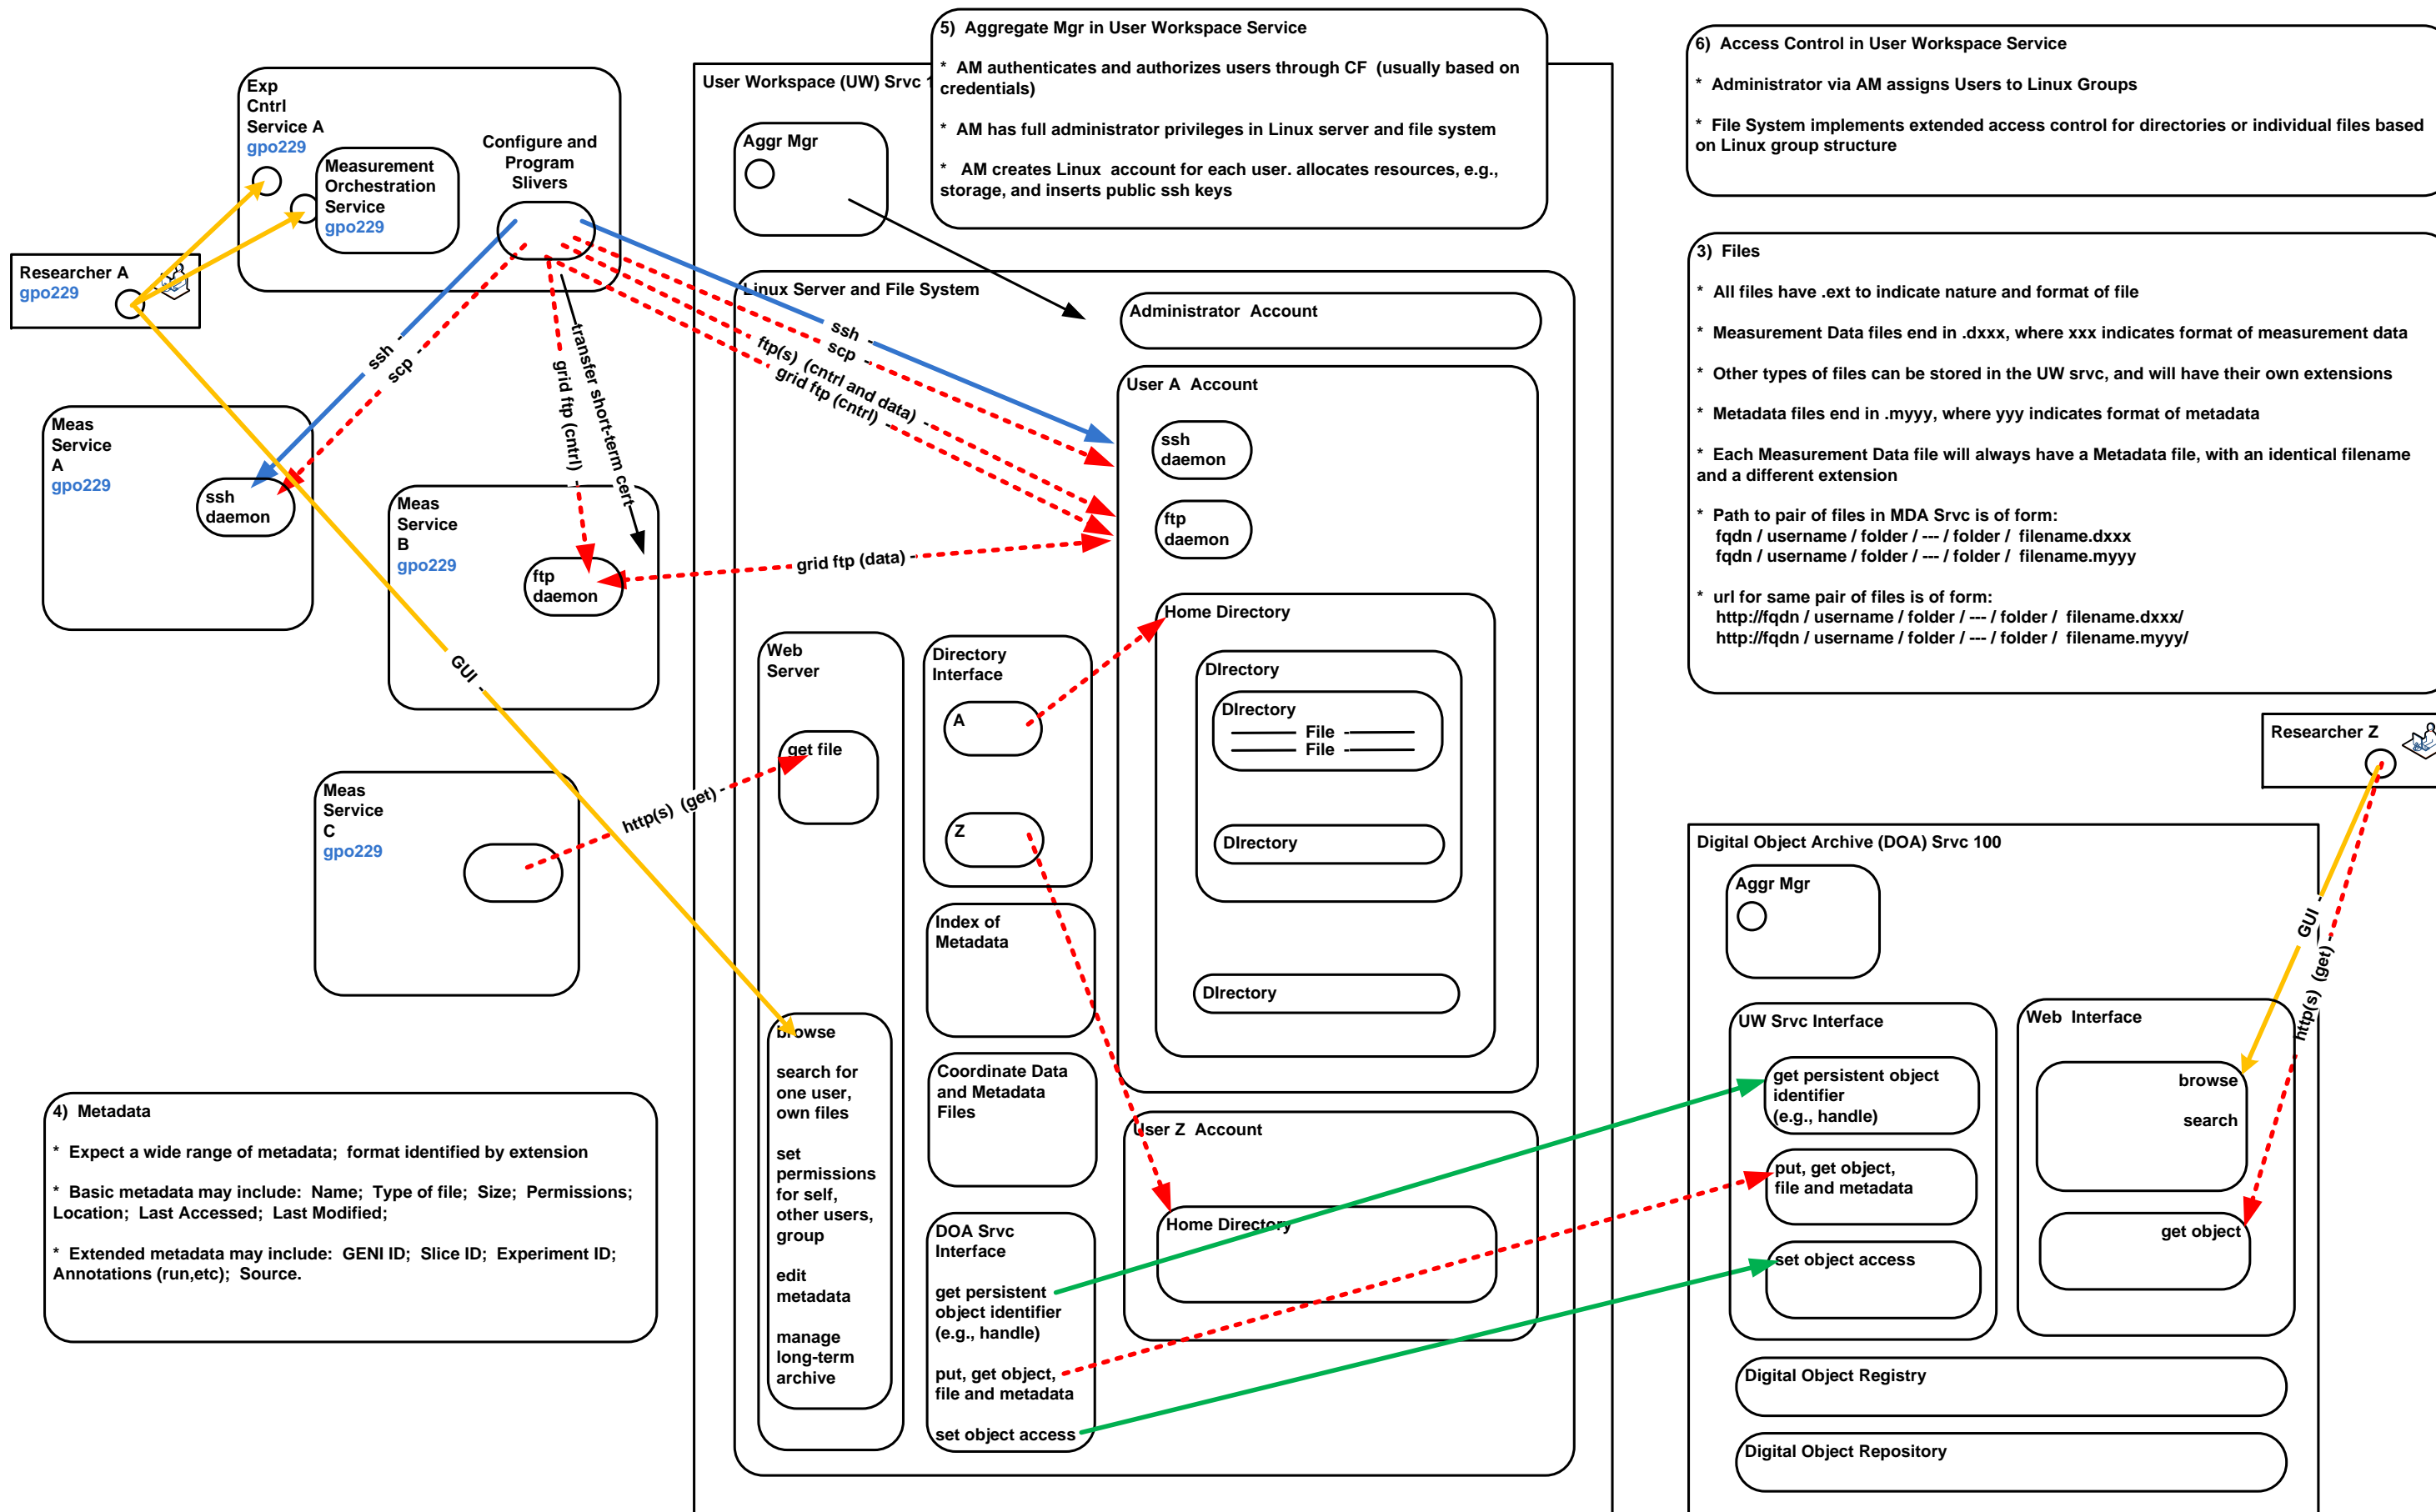
Meas Data Schema

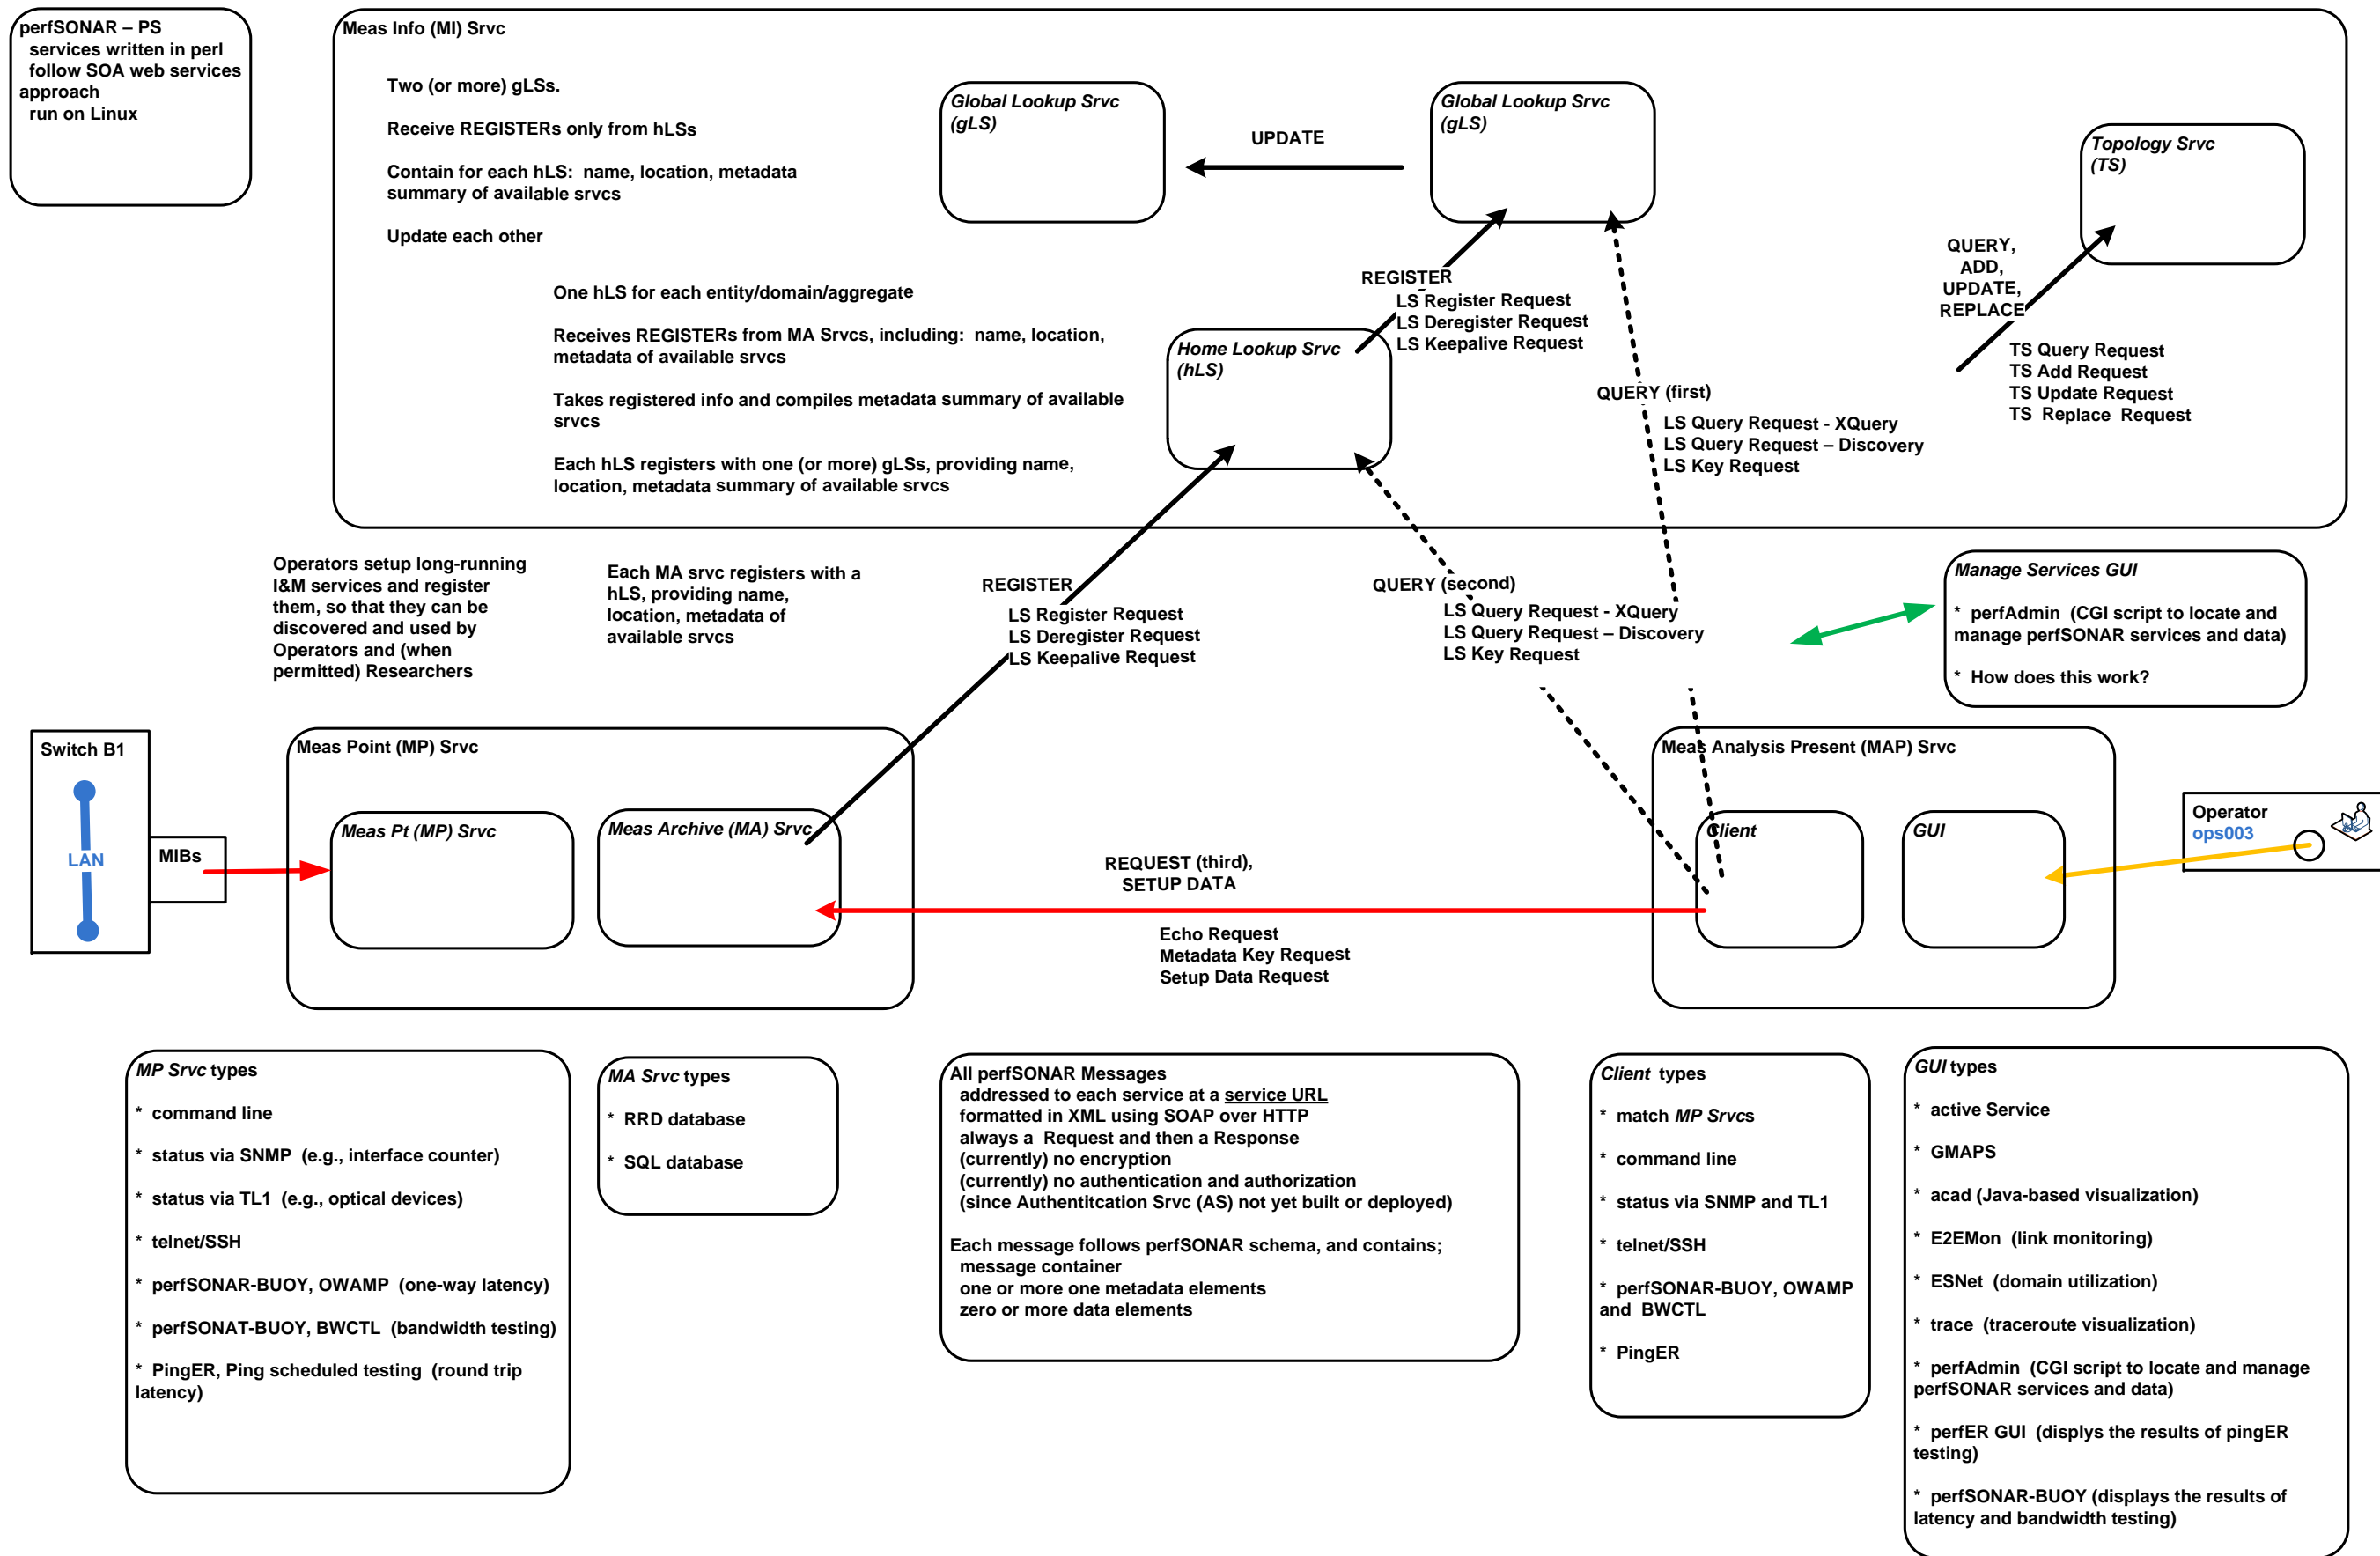
- * none defined
- * metadata not defined

Meas Data Arch Svc

- * svc not yet defined
- * messages not yet defined

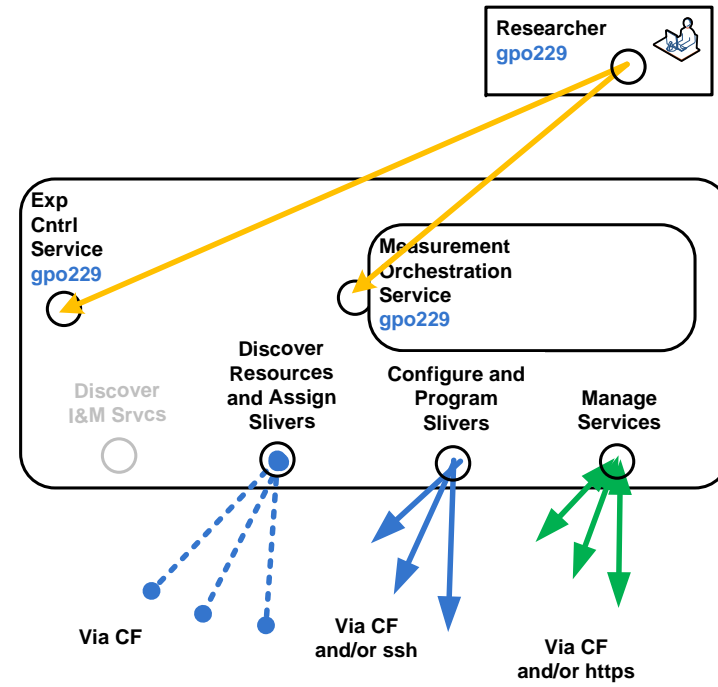






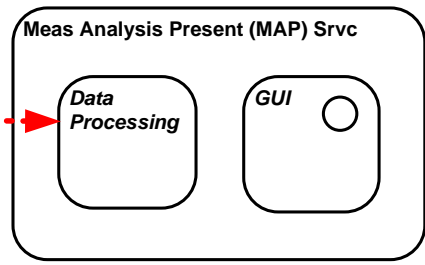
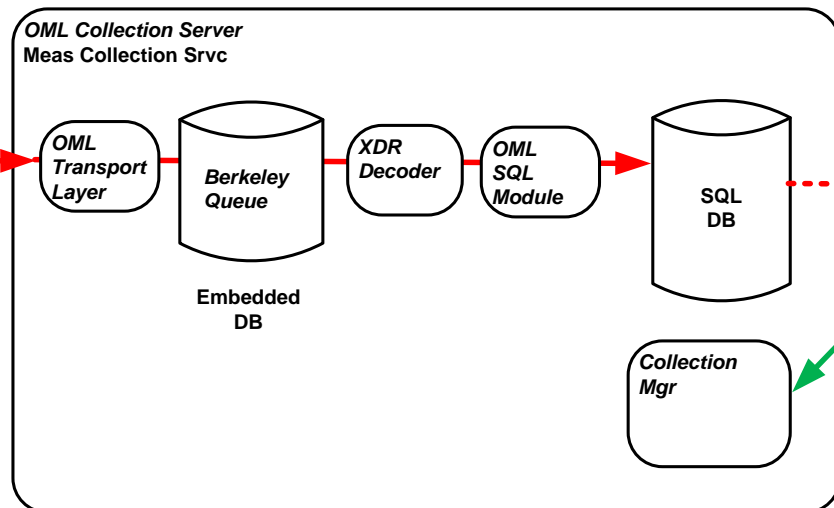
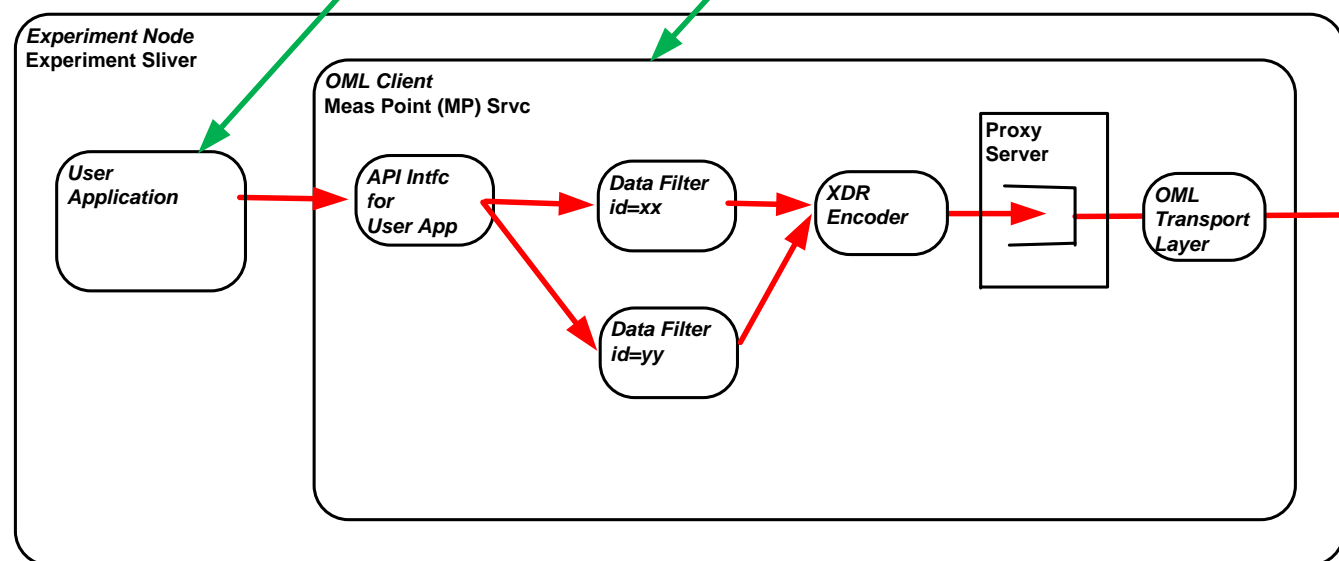
OMF/OML Structure

- * One OML Client included in each Experiment Node
- * One OML Collection Server provided for each OMF site
- * Researcher programs User Application and OML Client to gather and filter desired measurement data, and define stream to Collection Server
- * Meas Data streamed from OML Client to OML Collection Server



Manage (and Control) Srvcs

- * Via HTTP to all srvcs's, with APIs based on REST.
- * Via HTTP to OML Client srvcs, to config files specifying filtering and streaming, which are then compiled into code



Meas Analysis Present Srvc

- * Running outside of OMF/OML.
- * Can import directly from SQL DB
- * EC can arrange to convert tables into graphs

Client API

- * Researcher uses web interface to define measurement points and parameters through a web interface, saved as XML config file; causes XSLT-based code generator to generate source code for measurement client.
- * Researcher includes measurement-point id's, and metric id's, which dynamically creates a schema.

Meas Data Filters

- * Researcher defines filters, separate from User App, so that can easily change data collection behavior
- * Can be triggered by various properties, e.g., time, by the no of data values collected, etc.
- * Can be triggered by events, e.g., passing a trip line in a mobile application

Meas Data Streams

- * Researcher defines measurement streams, gathering data samples and averaging, etc.
- * Meas data is series of typed vectors, XDR coded, and then streamed from client to collection server using proprietary OML protocol, on top of TCP, over dedicated Control VLAN
- * Considering using IPFIX instead of prop OML protocol; IPFIX typically uses SCTP for transport
- * If path becomes disconnected from time-to-time. data is cached in Proxy Server FIFO, and then forwarded when path is reestablished

Meas Data Schema

- * Meas data follows schema defined by researcher, including: measurement-point id's, metric id's, etc.
- * A sensor (or application, or service) define a set of measurement points, with each measurement point defined by a name and a typed vector (sensor schema).
- * At runtime, the experimenter (or operator) provides a streams spec which defines what measurement points are going to be activated and what initial processing is going to be performed - that defines the actual schema going over the wire and/or ending up in the collection database

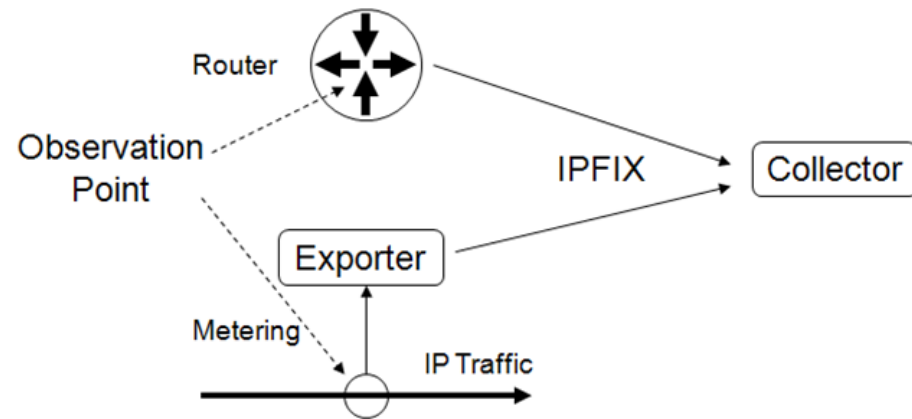
DB Schema

- * Application definition is used to create DB schema for experiment, using XSLT.
- * DB table is created for each measurement point, names based on id attribute of the group element.
- * Includes mandatory fields for name/id, timestamp, sequence number
- * Protocol is self describing
- * Server automatically creates a table for every distinct stream (distinct in terms of schema not source).
- * Streams carry their own name which is translated into a database using a simple naming convention.

Experiment Portal

- * Early prototype
- * Each experiment results in a separate page containing all the experiment related information (script, parameter, resources used, time) as well as a pointer to the measurement database.
- * Where?


IPFIX Architecture



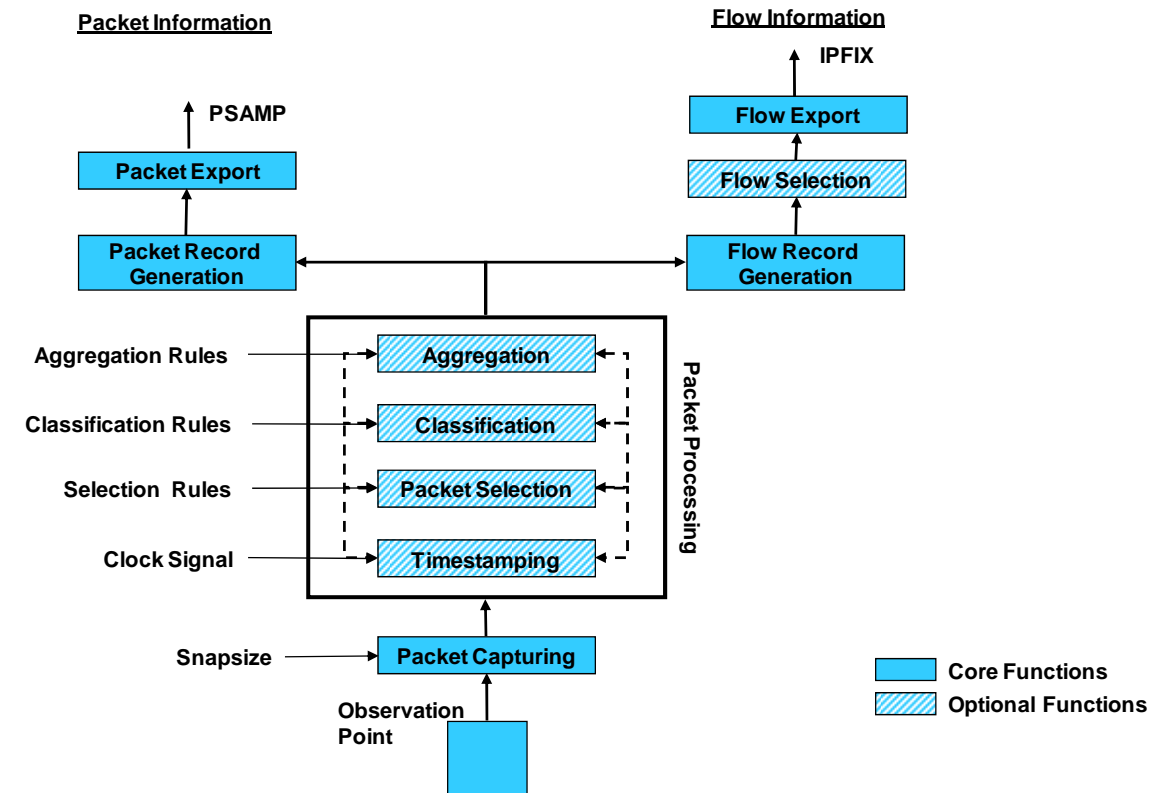
- push protocol: periodically IPFIX messages to configured receivers
- Transport protocols: SCTP (, UDP, TCP)

Data Representation

- Templates in the message stream describe the data sets
- Allows flexible and efficient (binary) representation of flows on the wire



IPFIX/PSAMP Measurement Model



Information Model

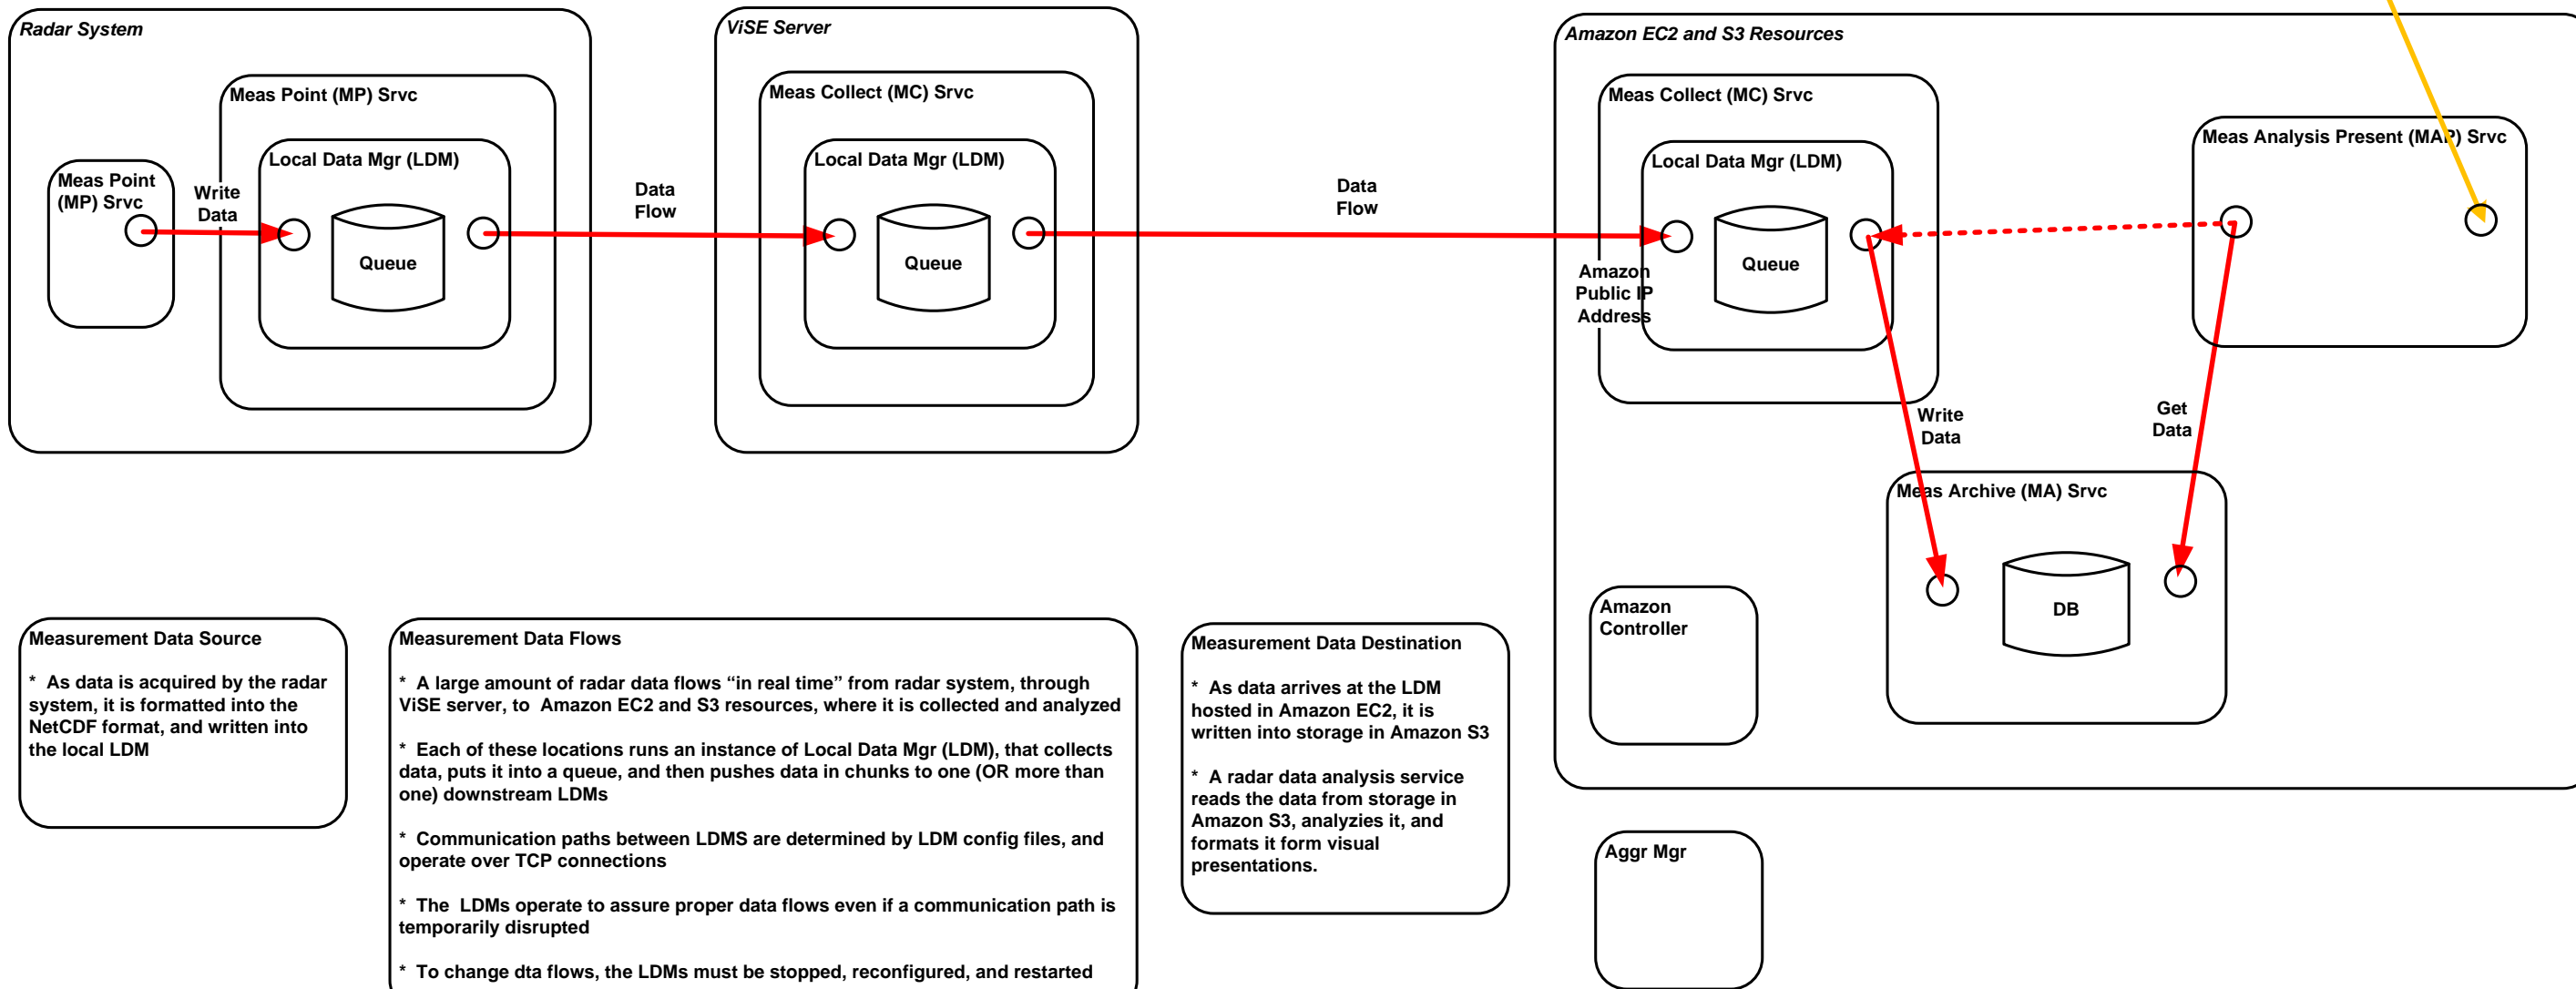
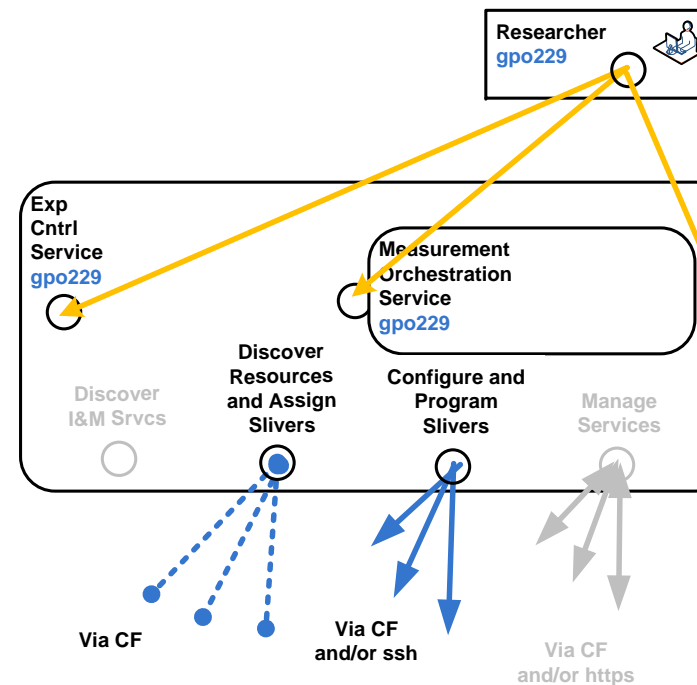
- The information model supports reporting a wide variety of information elements (IEs):
 - "Five-tuple" (IPv4, IPv6 header fields) and standard packet/byte counters
 - All ICMP, TCP, UDP header fields
 - Layer 2, VLAN, MPLS, and other sub-IP information
 - Timestamps down to nanosecond resolution
 - Packet treatment: e.g., routed next hop and AS
 - Detailed counters: e.g., sum of squares, flag counters
- New IEs registered with IANA
- Enterprise-specific IEs for private extensions
- New defined IEs
 - location / GPS information, QoS parameters, spectrum measurements, ...

Measurement Data Schema

- * Radar data follows NetCDF format.
- * NetCDF is: "self-describing, machine independent, binary, data format, that supports the creation, access and sharing of array-oriented scientific data."
- * "Self-describing" means that there is a header that describes the rest of the file, in particular the data arrays, as well as arbitrary file metadata in the form of name/value attributes.
- * Various users of NetCDF (such as Climate and Forecast (CF)) have guidelines for metadata.

References:

- * NetCDF (network Common Data Form) is a set of software libraries and machine-independent data formats that support the creation, access, and sharing of array-oriented scientific data.
<http://www.unidata.ucar.edu/software/netcdf/>
- * Unidata Local Data Manager (LDM) is a collection of cooperating programs that select, capture, manage, and distribute arbitrary data products. The system is designed for event-driven data distribution, and is currently used in the Unidata Internet Data Distribution (IDD) project. The LDM system includes network client and server programs and their shared protocols. An important characteristic of the LDM is its support for flexible, site-specific configuration.
<http://www.unidata.ucar.edu/software/ldm/>
<http://www.unidata.ucar.edu/software/ldm/ldm-6.6.5/basics/generic-LDM.html>



Measurement Data Source

- * As data is acquired by the radar system, it is formatted into the NetCDF format, and written into the local LDM

Measurement Data Flows

- * A large amount of radar data flows "in real time" from radar system, through VISE server, to Amazon EC2 and S3 resources, where it is collected and analyzed
- * Each of these locations runs an instance of Local Data Mgr (LDM), that collects data, puts it into a queue, and then pushes data in chunks to one (OR more than one) downstream LDMS
- * Communication paths between LDMS are determined by LDM config files, and operate over TCP connections
- * The LDMS operate to assure proper data flows even if a communication path is temporarily disrupted
- * To change data flows, the LDMS must be stopped, reconfigured, and restarted

Measurement Data Destination

- * As data arrives at the LDM hosted in Amazon EC2, it is written into storage in Amazon S3
- * A radar data analysis service reads the data from storage in Amazon S3, analyzes it, and formats it form visual presentations.