

Online remote monitoring facilities for the ATLAS experiment

Serguei Kolos (University of California, Irvine)

Of behalf of the ATLAS remote monitoring team


ATL-DAQ-SLIDE-2010-396
14 October 2010





Outline

- Introduction
 - ATLAS collaboration
 - Motivation for remote monitoring
- Remote monitoring requirements
- Remote monitoring services:
 - ATLAS online monitoring architecture
 - Remote monitoring services design and implementations
 - Experience of the exploitation
- Summary and conclusion

ATLAS collaboration



- ATLAS is one of the 4 major LHC experiments at CERN
- ATLAS detector and TDAQ system construction took about 15 years:
 - 140 000 000 channels
- Over 3000 people from:
 - 169 institutes
 - 38 countries



The motivation for remote monitoring

- After the first year of the ATLAS exploitation the DAQ efficiency for 7TeV data is about 93%
- Availability of expertise from the people who were participating in the ATLAS construction is the key point to the success:
 - Many experts are not permanently located at CERN
 - Coming to CERN is possible but not too often due to various reasons (other commitments, budget limitations, teaching duties, etc.)
 - Even those who are at CERN can not be present in the Control Room all day

ATLAS security restrictions

- ATLAS experiment has its private set of networks:
 - Data, Technical and Control
 - Connection from the CERN Global Public Network is provided via dedicated gateway server:
 - Closed during data taking
 - Experts can log-in only in case of serious issues after the approval of the ATLAS Shift Leader
- Even experts based at CERN don't have direct access to the monitoring data, unless sitting in one of the control rooms

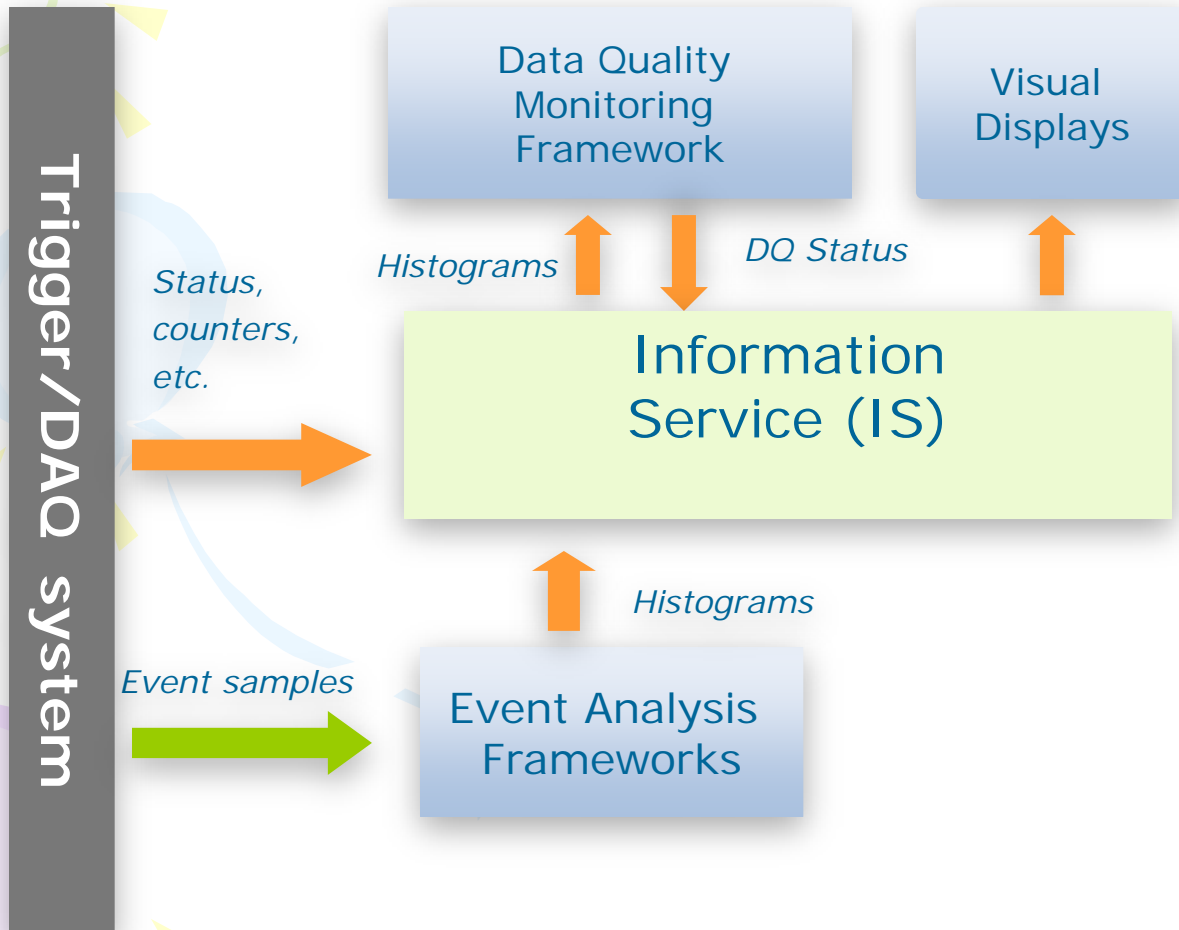
ATLAS Remote Monitoring Requirements

- Any member of the ATLAS collaboration shall be able to:
 - View overall high-level status of the currently ongoing data taking activities
 - Small amount of information is made available to everybody
 - Updated at fixed time intervals (every couple of minutes)
- ATLAS sub-system expert shall be able to:
 - Request and obtain an up-to-date value of a given monitoring information
- ATLAS sub-system shifter shall be able to:
 - Permanently monitor the status of the given sub-system in real-time while data are being taken
 - Reduces the burden on the sub-system experts, who are contacted only in case of problems

(-) Amount of data (+)

(-) Number of users (+)

ATLAS Online Monitoring Architecture



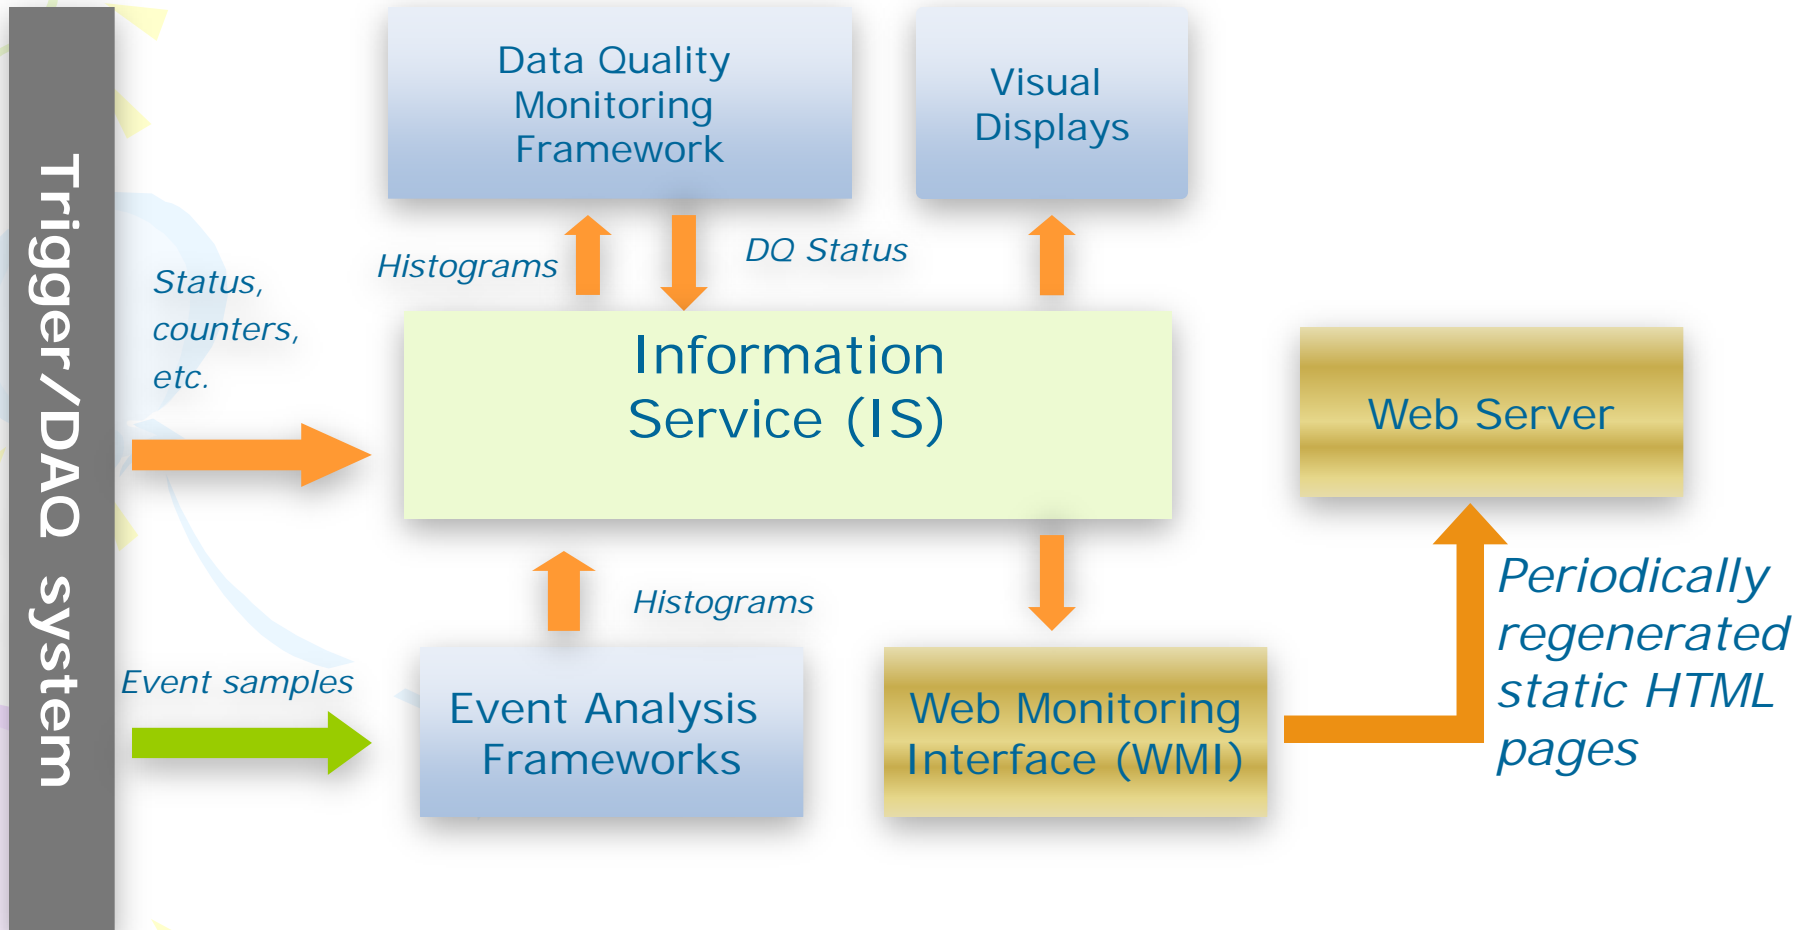
- All monitoring information is kept in the Information Service:
 - It's complete
 - It's up-to-date
- IS provides read-write and subscribe-callback APIs:
 - Java
 - C++
 - Python

General public remote monitoring



- Very limited set of data
- Practically unlimited number of users

General public remote monitoring: Design



General public remote monitoring: Implementation

- WMI is a C++ framework:
 - It is running in the ATLAS private network
 - It fires plug-ins execution at configurable time intervals
 - It provides data-out C++ API for the plug-ins
 - It converts monitoring data chosen by the plug-ins to HTML files
 - It copies those files to the ATLAS Web server
- WMI plug-ins:
 - Select information to be put to HTML files

WMI DAQ pages

File Edit View History Bookmarks Tools Help

Run Status

Partition ATLAS

TTC Partitions: Pix Barrel - Pix Disk - Pix B-Layer - SCT BA - SCT BC - SCT EA - SCT EC - TRT BA - TRT BC - TRT EA - TRT EC - LAr EMBA - LAr EMBC - LAr EMCA - LAr EMECC - LAr HECA - LAr HECC - LAr FCALA - LAr FCALC - TIB BA - TIB BC - TIB EA - TIB EC - MDT BA - MDT BC - MDT EA - MDT EC - RPC BA - RPC BC - TGC EA - TGC EC - CSC EA - CSC EC - L1 calo preprocessor - L1 calo cluster DAQ - L1 calo cluster Rol - L1 calo jet/E DAQ - L1 calo jet/E Rol - MUCTPI - CTP - L2SV - SFI - SFO - LVL2 - EF - BCM - Lucid - ZDC -

[Check today's program here!](#) [Data taking efficiency](#)
[Other active partitions can be seen here.](#)

Run Info		Run Statistics		Trigger Info		Beam Info	
Run State	RUNNING	RunTime	11:54:57	Master & Prescale Keys	920, 2284, 2281	Beam Mode	INJECTION PROBE BEAM
Run Tag	data10_7TeV	Luminosity Block	367	L1 Bunch Group	106	Beam 1 Status	Present & Safe
Run Type	Physics	LB changes every	120 seconds	Simple Deadtime	5	Beam 2 Status	Present & Safe
Run Number	166150	Average Event Size [MB]	1.408	Complex Deadtime	7/415	Stable Beams	FALSE
Run Mode	Standby	Throughput to Disk [MB/s]	62.0425	HLT Release Version	15.6.9.28	Beam Energy	450.12

Busy Status

CTPMI	CTPCORE	CTPOUT 12	CTPOUT 13	CTPOUT 14	CTPOUT 15
VME 0%	Backplane 1.192%	CTP(LUCID) 0%	BCM 0%	LHCf OUT 0%	CSC 0%
ECR 0.041%	Result 1.192%	Pixel 0%	ZDC 0%	MDT B 0%	ALFA OUT 0%
Veto 0 0%		SCT 0%	LAr H/F 0%	MDT EC 0%	TGC 0%
Veto 1 0%		TRT 1.152%	LAr EMEC 0%	Tile EB 0%	RPC 0%
Backplane 1.192%		L1Calo 0%	LAr EMB 0%	Tile LB 0%	MUCTPI 0%

Global Rates (Hz)

Legend:

- L1 out (TRP)
- L2out (TRP)
- EFout (TRP)
- L2SV: L2 In
- L2PU:L2 Int Ev Rate
- SFI:Ev build Rate
- SFO:Ev Save Rate
- L1_MBTS_1_1

WMI Data Quality pages

ATLAS Data Quality Monitoring

Left Sidebar (Monitoring Categories):

- IDGL
- TileRegion
- CSC
- ourGlobal
- Led Events
- TGC_DQR
- MBTSRegion
- LCD
- RPC
- Scalers
- IDBCM
- HLT
- ZDC
- TRT
- BCM_L2_DQRegion
- Pixels
- L1CAL
- BCM_PM_DQRegion
- CaloGlobal
- SCT
- LAr
- MDT_DQR
- DataFlow

Main Content Area (Pedestal Details):

Summary: NumDisabled = 0; NumExcluded = 4; NumGreen = 0; NumRed = 60; NumUndefined = 0; NumYellow = 0;

Channel 0: ped.adc.avg.chan000 (4/10/10 02:29:34)

- Actual Mean Value = 73.1426; Actual RMS Value = 0;
- Histograms: ZdcRodDqm.ZdcSpyHist_DQM.pedADCAvg.ppm0.chan00
- Reference
- Algorithm: ZdcPedestalStatCheck
- Parameters: MinStat = [1,]; TargetValue = [40,];
- Thresholds: 10 < DistanceFromTarget < 20 ; 5 < RMSThreshold < 10 ;

Channel 1: ped.adc.avg.chan001 (4/10/10 02:29:34)

- Actual Mean Value = 73.1426; Actual RMS Value = 0;
- Histograms: ZdcRodDqm.ZdcSpyHist_DQM.pedADCAvg.ppm0.chan01
- Reference
- Algorithm: ZdcPedestalStatCheck
- Parameters: MinStat = [1,]; TargetValue = [40,];
- Thresholds: 10 < DistanceFromTarget < 20 ; 5 < RMSThreshold < 10 ;

Channel 2: ped.adc.avg.chan002 (4/10/10 02:29:34)

- Actual Mean Value = 73.1426; Actual RMS Value = 0;
- Histograms: ZdcRodDqm.ZdcSpyHist_DQM.pedADCAvg.ppm0.chan02
- Reference

Bottom Left: Last updated: Mon Oct 4 09:40:43 2010

WMI HTML pages usage statistics (first week of October)

File Edit View History Bookmarks Tools Help

cern.ch https://atlasop.cern.ch/apache/webalizer/usage_201010 Google

Most Visited Google


ATLAS TDAQ Monitorin... x Usage Statistics for atla... x

Top 30 of 9361 Total URLs

#	Hits	KBytes	URL
1	196322 3.43%	16144 0.02%	/sysadmin/ldap_roles/scripts/ajax/ldap.py/getShiftLeaderBoxContent
2	113878 1.99%	732569 0.95%	/atlas-point1/wmi/current/Run Status_wmi/ATLAS.html
3	113467 1.98%	93211 0.12%	/atlas-point1/wmi/current/Data Quality Monitoring_wmi/ATLAS.html
4	88691 1.55%	7121 0.01%	/sysadmin/ldap_roles/scripts/ajax/rolesdb.py/showRolesInWaitingState
5	55252 0.97%	8871 0.01%	/atlas-point1/dcs/css/atldcs.css
6	53432 0.93%	101594 0.13%	/atlas-point1/script-new/iquery-1.2.1.min.js
7	53426 0.93%	65026 0.08%	/atlas-point1/script-new/iquery.js
8	53350 0.93%	5798 0.01%	/atlas-point1/script-new/collapDiv.js
9	53348 0.93%	2767 0.00%	/atlas-point1/script-new/menu.js
10	53304 0.93%	60435 0.08%	/atlas-point1/css-new/style.css
11	53224 0.93%	1859 0.00%	/atlas-point1/css-new/reset.css
12	42358 0.74%	315423 0.41%	/atlas-point1/dcs/dcs/process.php

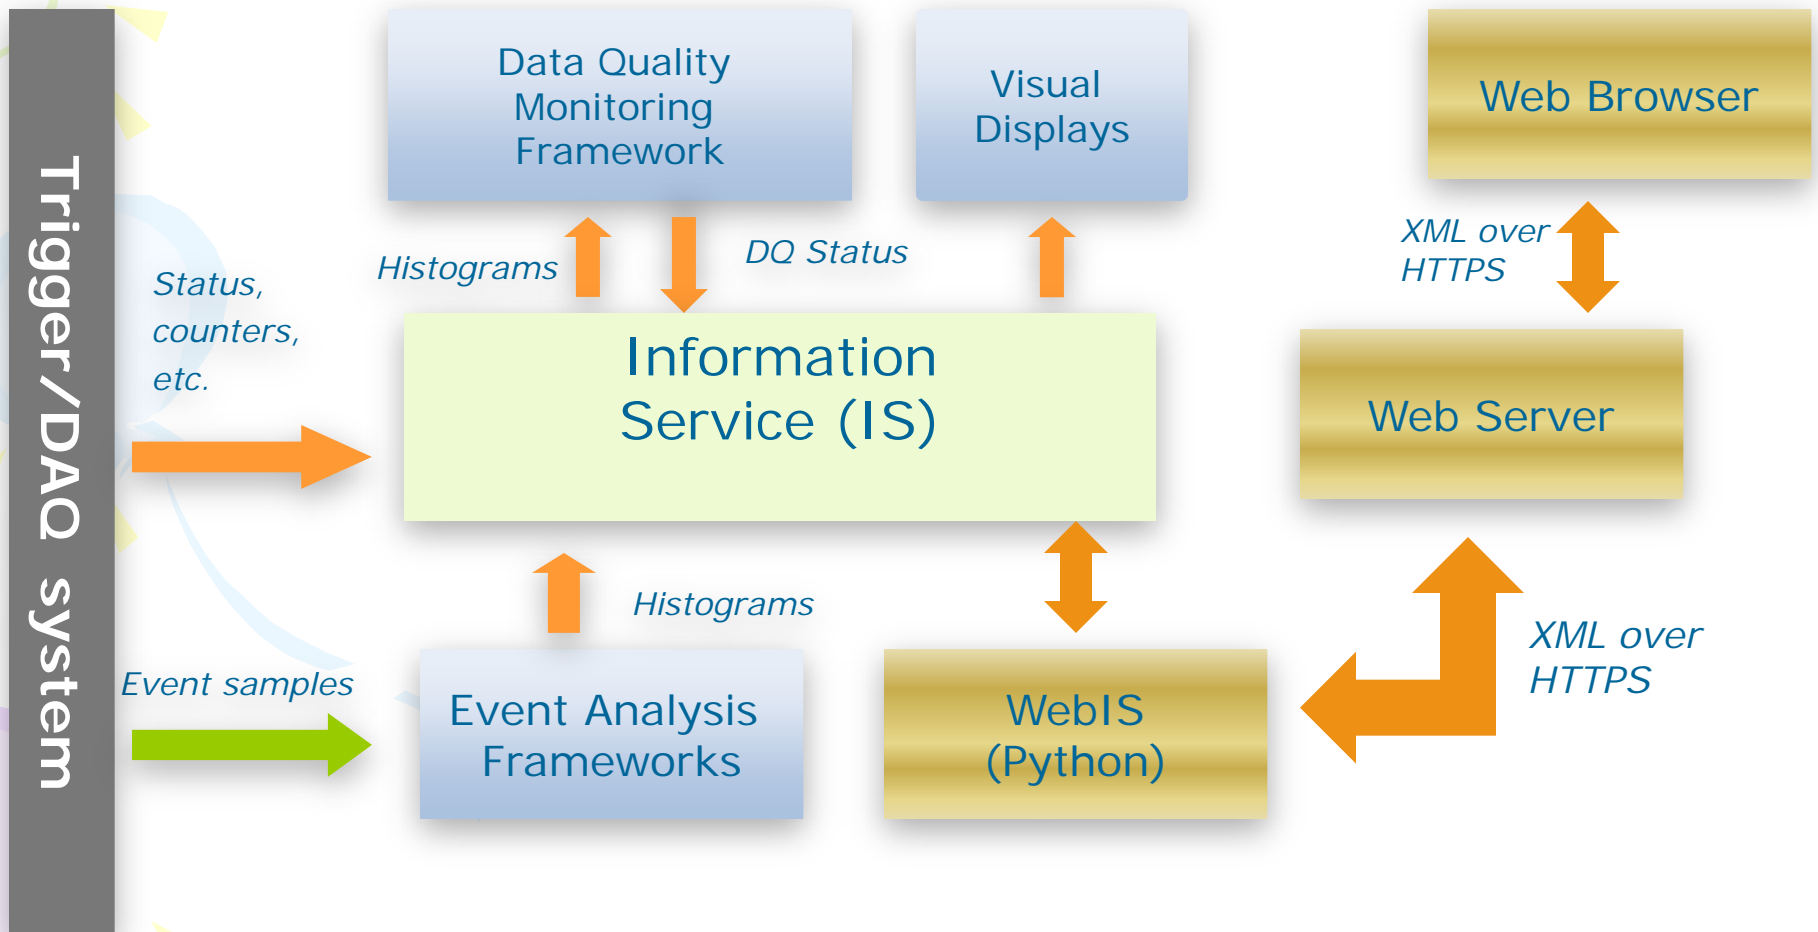
Done

Expert mode remote monitoring




- Any monitoring information can be provided on request
- Overall request rate is limited

Expert mode remote monitoring: Design



Expert mode remote-monitoring: Implementation



- Each object in IS has a unique URL:
 - <https://atlasop.cern.ch/info/current/ATLAS/is/RunParams/RunParams.RunParams>
 - One can type this URL in the Web browser to get the specific info
- WebIS – Python HTTP server:
 - Send back the content of IS objects with respect to given URLs
 - Apache cache is used to reduce the number of requests
- Generic facility which can be used to construct complex WEB based GUIs using:
 - CSS, JavaScript, Java applets, standalone Java applications

WebIS: Histogram presenter

https://atlasop.cern.ch/atlas-point1/tdaq/web_is/ohp/ATLAS.html

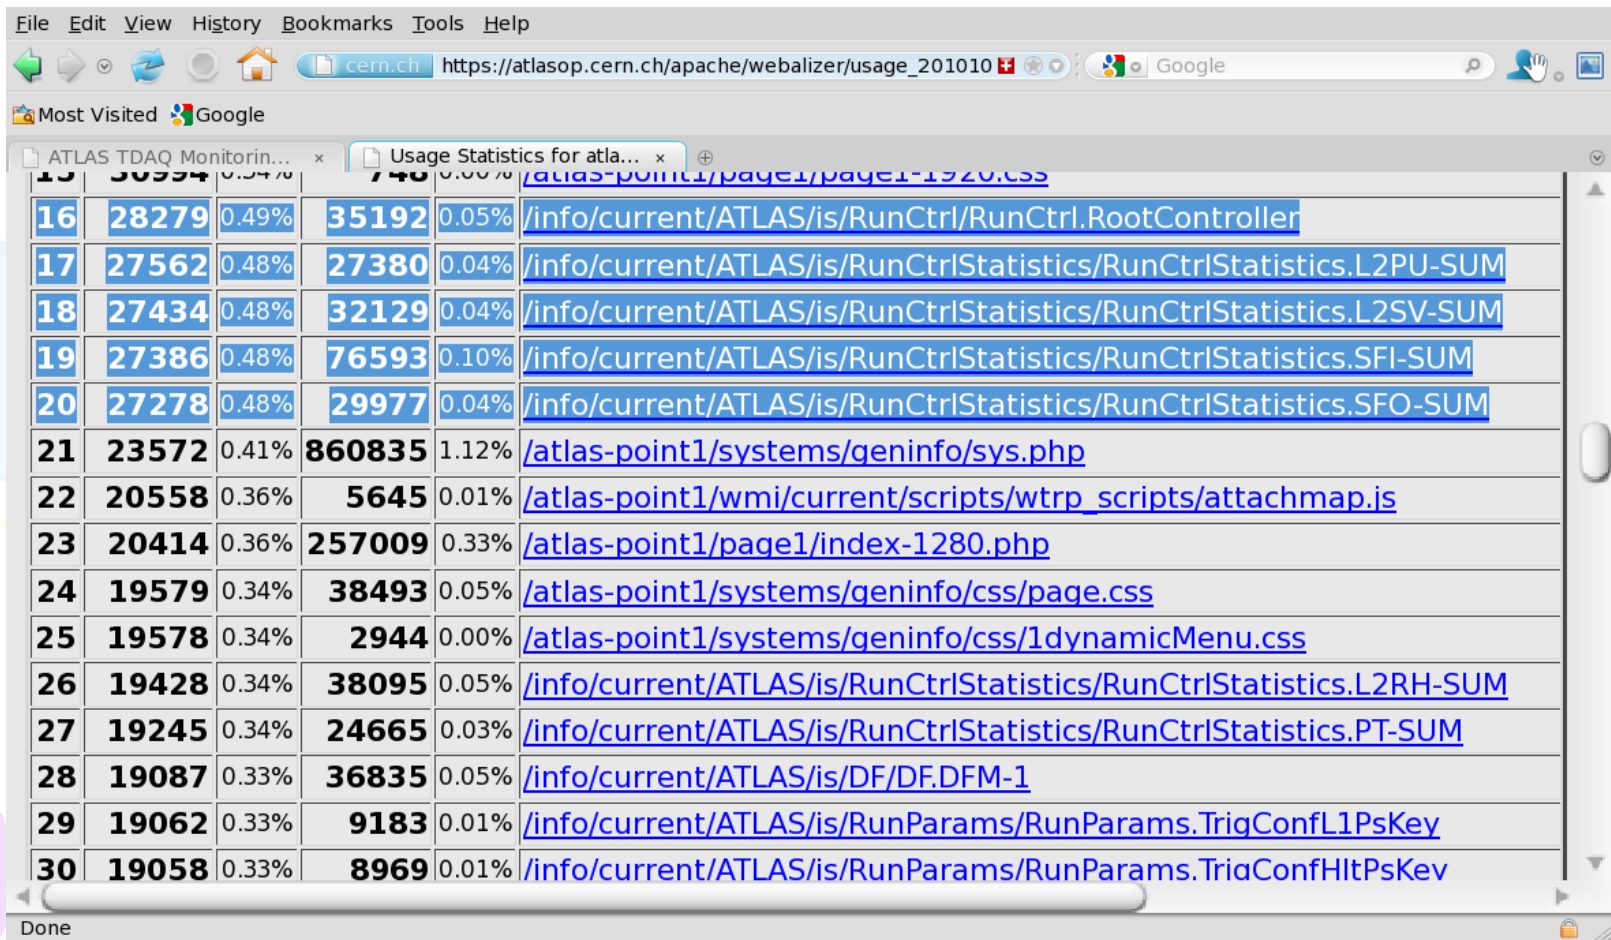
The screenshot shows a web browser window displaying the ATLAS WebIS interface. The browser's address bar shows the URL https://atlasop.cern.ch/atlas-point1/tdaq/web_is/ohp/ATLAS.html. The page title is "ATLAS: RUNNING". On the left side, there is a navigation menu with the following items: CollisionVIPlots, InDet, Tile, Muon, Global, CombPerf, Coincidence, LArEnergyFlow, egammaOther, GlobalOther, JetsOther, TRTOther, Rates, WTRP, and Browser. The main content area is titled "TILE" and contains four plots arranged in a 2x2 grid. The top-left plot is titled "Run 166296 Trigger AnyPhysTrig: Tile 2D Cell Energy Average deposition (MeV)" and shows a scatter plot of energy deposition in MeV versus pseudorapidity η (x-axis, -2 to 2) and pseudorapidity η (y-axis, -3 to 3). The top-right plot is titled "Run 166296 Trigger AnyPhysTrig: Tile Cell Position of cells over threshold 300 MeV" and shows a heatmap of cell positions over the same axes, with a color scale from 0 to 500. The bottom-left plot is titled "Run 166296 Trigger bld_RNDM: Tile 2D Cell Energy Average deposition (MeV)" and shows a scatter plot similar to the top-left plot. The bottom-right plot is titled "Run 166296 Trigger bld_RNDM: Tile Cell Position of cells over threshold 300 MeV" and shows a heatmap similar to the top-right plot, with a color scale from 0 to 350. The browser's status bar at the bottom shows "Done".

WebIS: L1 Calorimeter remote monitoring application (Java)

<https://twiki.cern.ch/twiki/bin/view/Atlas/LevelOneCaloRemoteMonitoring>



Expert mode Web access statistics (first week of October)



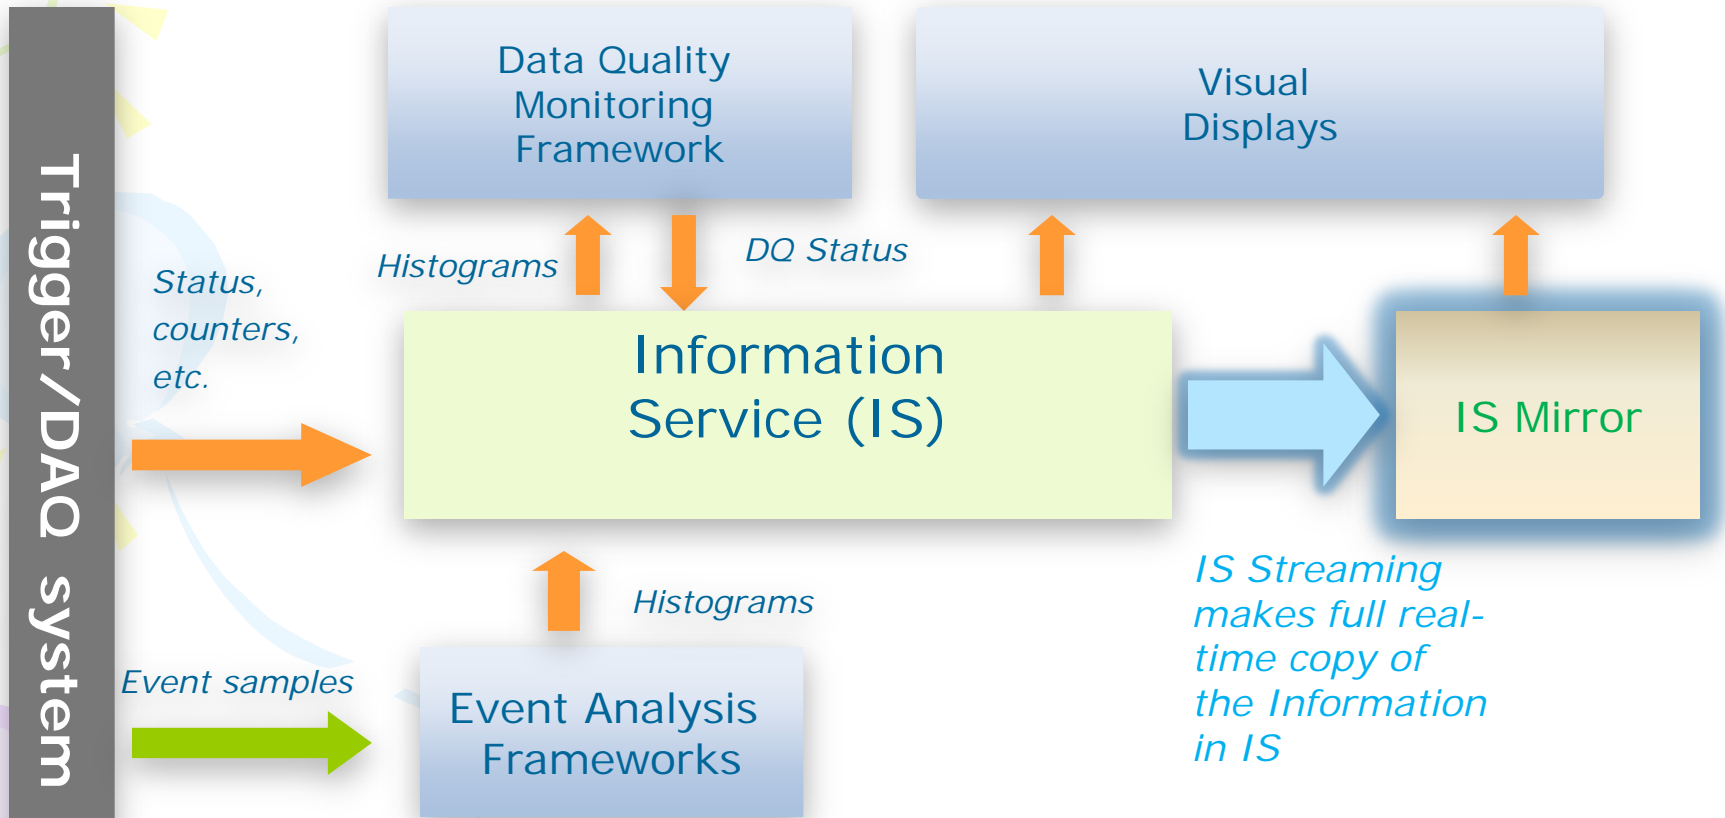
Rank	Hits	Percentage	Hits	Percentage	URL
16	28279	0.49%	35192	0.05%	/info/current/ATLAS/is/RunCtrl/RunCtrl.RootController
17	27562	0.48%	27380	0.04%	/info/current/ATLAS/is/RunCtrlStatistics/RunCtrlStatistics.L2PU-SUM
18	27434	0.48%	32129	0.04%	/info/current/ATLAS/is/RunCtrlStatistics/RunCtrlStatistics.L2SV-SUM
19	27386	0.48%	76593	0.10%	/info/current/ATLAS/is/RunCtrlStatistics/RunCtrlStatistics.SFI-SUM
20	27278	0.48%	29977	0.04%	/info/current/ATLAS/is/RunCtrlStatistics/RunCtrlStatistics.SFO-SUM
21	23572	0.41%	860835	1.12%	/atlas-point1/systems/geninfo/sys.php
22	20558	0.36%	5645	0.01%	/atlas-point1/wmi/current/scripts/wtrp_scripts/attachmap.js
23	20414	0.36%	257009	0.33%	/atlas-point1/page1/index-1280.php
24	19579	0.34%	38493	0.05%	/atlas-point1/systems/geninfo/css/page.css
25	19578	0.34%	2944	0.00%	/atlas-point1/systems/geninfo/css/1dynamicMenu.css
26	19428	0.34%	38095	0.05%	/info/current/ATLAS/is/RunCtrlStatistics/RunCtrlStatistics.L2RH-SUM
27	19245	0.34%	24665	0.03%	/info/current/ATLAS/is/RunCtrlStatistics/RunCtrlStatistics.PT-SUM
28	19087	0.33%	36835	0.05%	/info/current/ATLAS/is/DF/DF.DFM-1
29	19062	0.33%	9183	0.01%	/info/current/ATLAS/is/RunParams/RunParams.TrigConfL1PsKey
30	19058	0.33%	8969	0.01%	/info/current/ATLAS/is/RunParams/RunParams.TriConfHitPsKey

Shifter mode remote monitoring



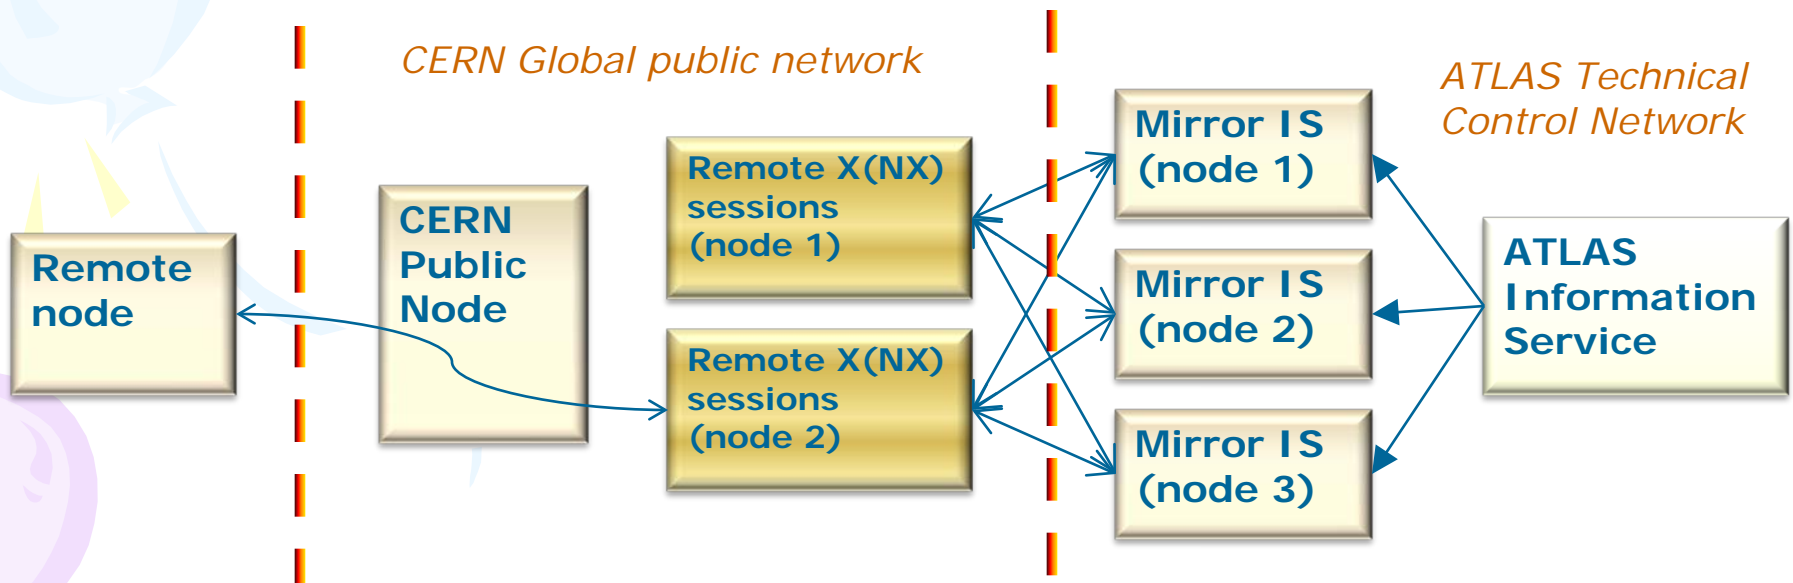
- Provides the same desktop and GUI monitors as in the ATLAS Control Room
- Supports limited number of users

Shifter mode remote monitoring: Design



Shifter mode remote monitoring: Implementation

- Mirror IS nodes have 2 NICs:
 - One can be connected only from Atlas Private Network
 - Another one is restricted to “Remote X session” nodes
- Remote users:
 - Open ssh-tunnel on a CERN public node
 - Open X (NX) session on one of the “Remote X session” nodes
 - Use the same GUI displays as in the ATLAS Control Room



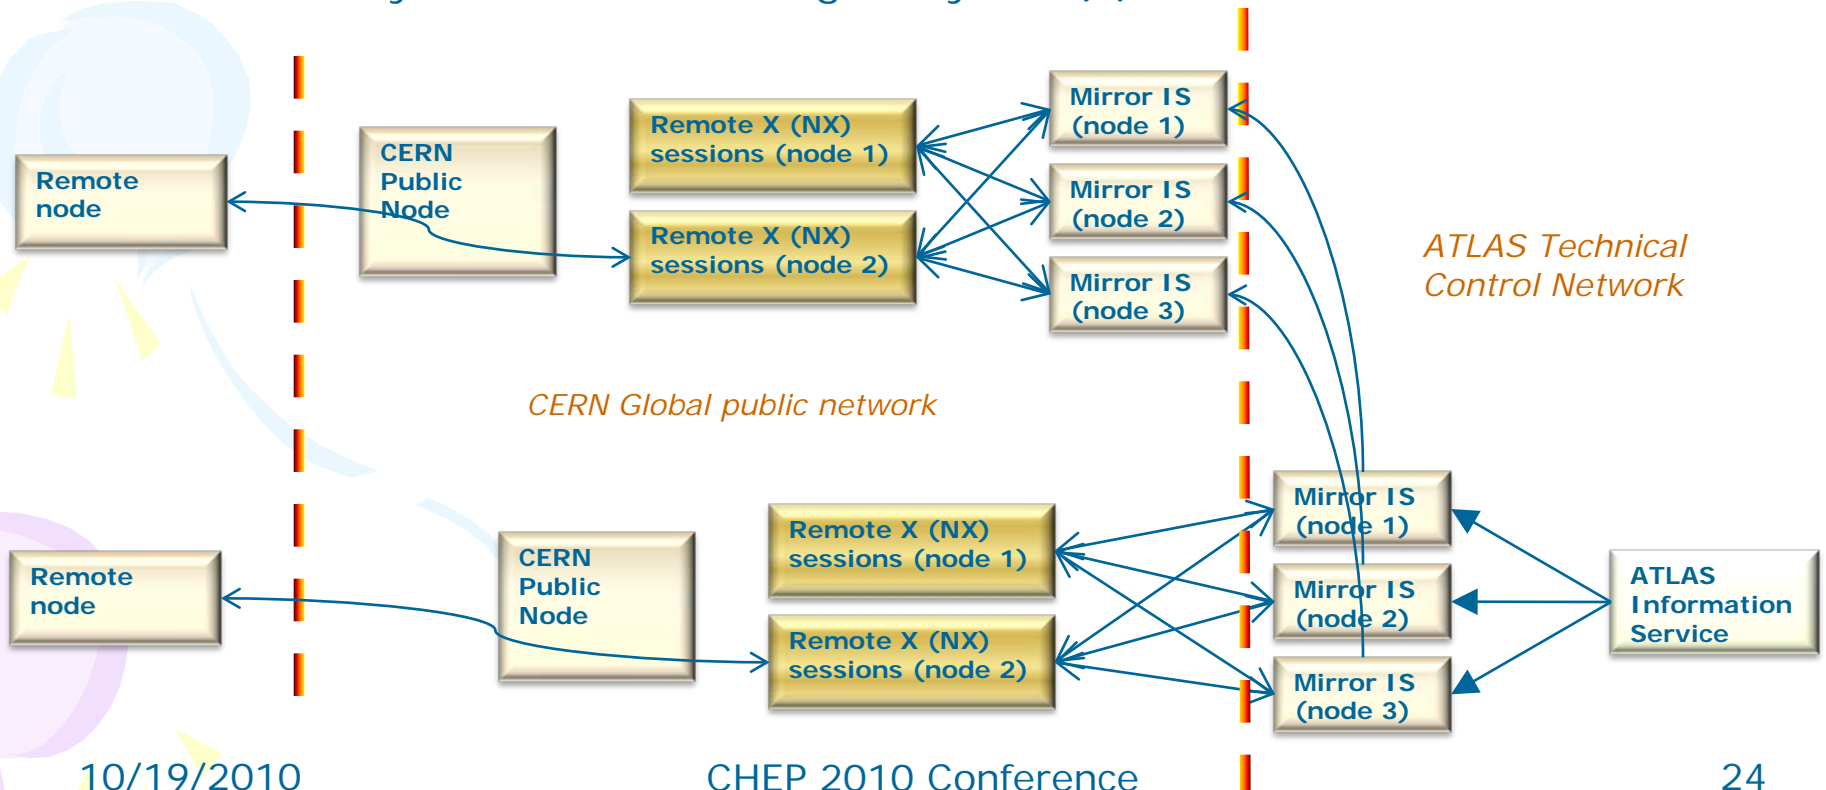


NX technology

- Plain X11 is very slow over long distance connections
- Several technologies were evaluated:
 - VNC, Sun Secure Global Desktop, NX
- Finally NX technology has been chosen:
 - The fastest, the cheapest and with small CPU/Memory overhead
- What is NX:
 - Initially proposed and implemented by NoMachine company
 - Provides compression and caching over standard X11 protocol
 - We are using FreeNX server (GPL)
 - Clients can download NoMachine NX client:
 - It is free for non-commercial use
 - It is available on Linux, Windows, MacOS

Shifter mode remote monitoring: Scalability

- This approach is almost “infinitely” scalable:
 - The information in the “IS mirror” service can be streamed to the “Secondary IS mirror”, which in turn can stream to the “Tertiary IS mirror”, etc.
 - At every level the streaming delay is $O(1)$ ms



Shifter mode remote monitoring: The highlights

- Mirrors full status of the ATLAS data taking session in real-time:
 - Information transfer delay is at the order of a few milliseconds
- Provides the same GUI applications as being used in the ATLAS Control Room:
 - No additional learning curve involved
 - No overhead for the SW development and maintenance
- Current configuration:
 - Two “Remote X session nodes” handle now up to 20 concurrent remote users sessions
- With recurrent mirroring, any practical number of remote sessions can be easily handled:
 - It’s just a matter of available HW resources



Summary and conclusion

- Three types of services have been implemented to provide remote monitoring for ATLAS
 - Each service is oriented to a specific group of ATLAS collaboration members
- The information available via WEB covers most of the experts needs for the monitoring
 - The number of log-in requests to the ATLAS private network was drastically reduced when the services were put in place
- Remote NX sessions are working nicely for people logged in from USA, Japan, Russia and other long distance sites:
 - Remote shifts become regular for some ATLAS sub-systems