FermiGrid
The Fermilab Campus Grid
28-Oct-2010

Keith Chadwick

Work supported by the U.S. Department of Energy under contract No. DE-AC02-07CH11359

Fermilab – Pre Campus Grid

- Multiple "siloed" clusters, each dedicated to a particular stakeholder:
  - CDF – 2 clusters, ~2,000 slots
  - D0 – 2 clusters, ~2,000 slots
  - CMS – 1 cluster, ~4,000 slots
  - GP – 1 cluster, ~500 slots
- Difficult to share:
  - When a stakeholder needed more resources, or did not need all of their currently allocated resources, it was extremely difficult to move jobs or resources to match the demand.
- Multiple interfaces and worker node configurations:
  - CDF – Kerberos + Condor
  - D0 – Kerberos + PBS
  - CMS – Grid + Condor
  - GP – Kerberos + FBSNG

FermiGrid – the Fermilab Campus Grid

- Site Wide Globus Gatekeeper (FNAL_FERMIGRID).
- Centrally Managed Services (VOMS, GUMS, SAZ, MySQL, MyProxy, Squid, Gratia Accounting, etc.)
- Compute Resources are "owned" by various stakeholders:

Compute Resources	# Clusters	# Gatekeepers	Batch System	# Batch Slots
CDF	3	5	Condor	5685
D0	2	2	PBS	5305
CMS	1	4	Condor	6904
GP	1	3	Condor	1901
Total	**7**	**15**	**n/a**	**~19,000**
Sleeper Pool	1	2	Condor	~14,200

FermiGrid - Architecture

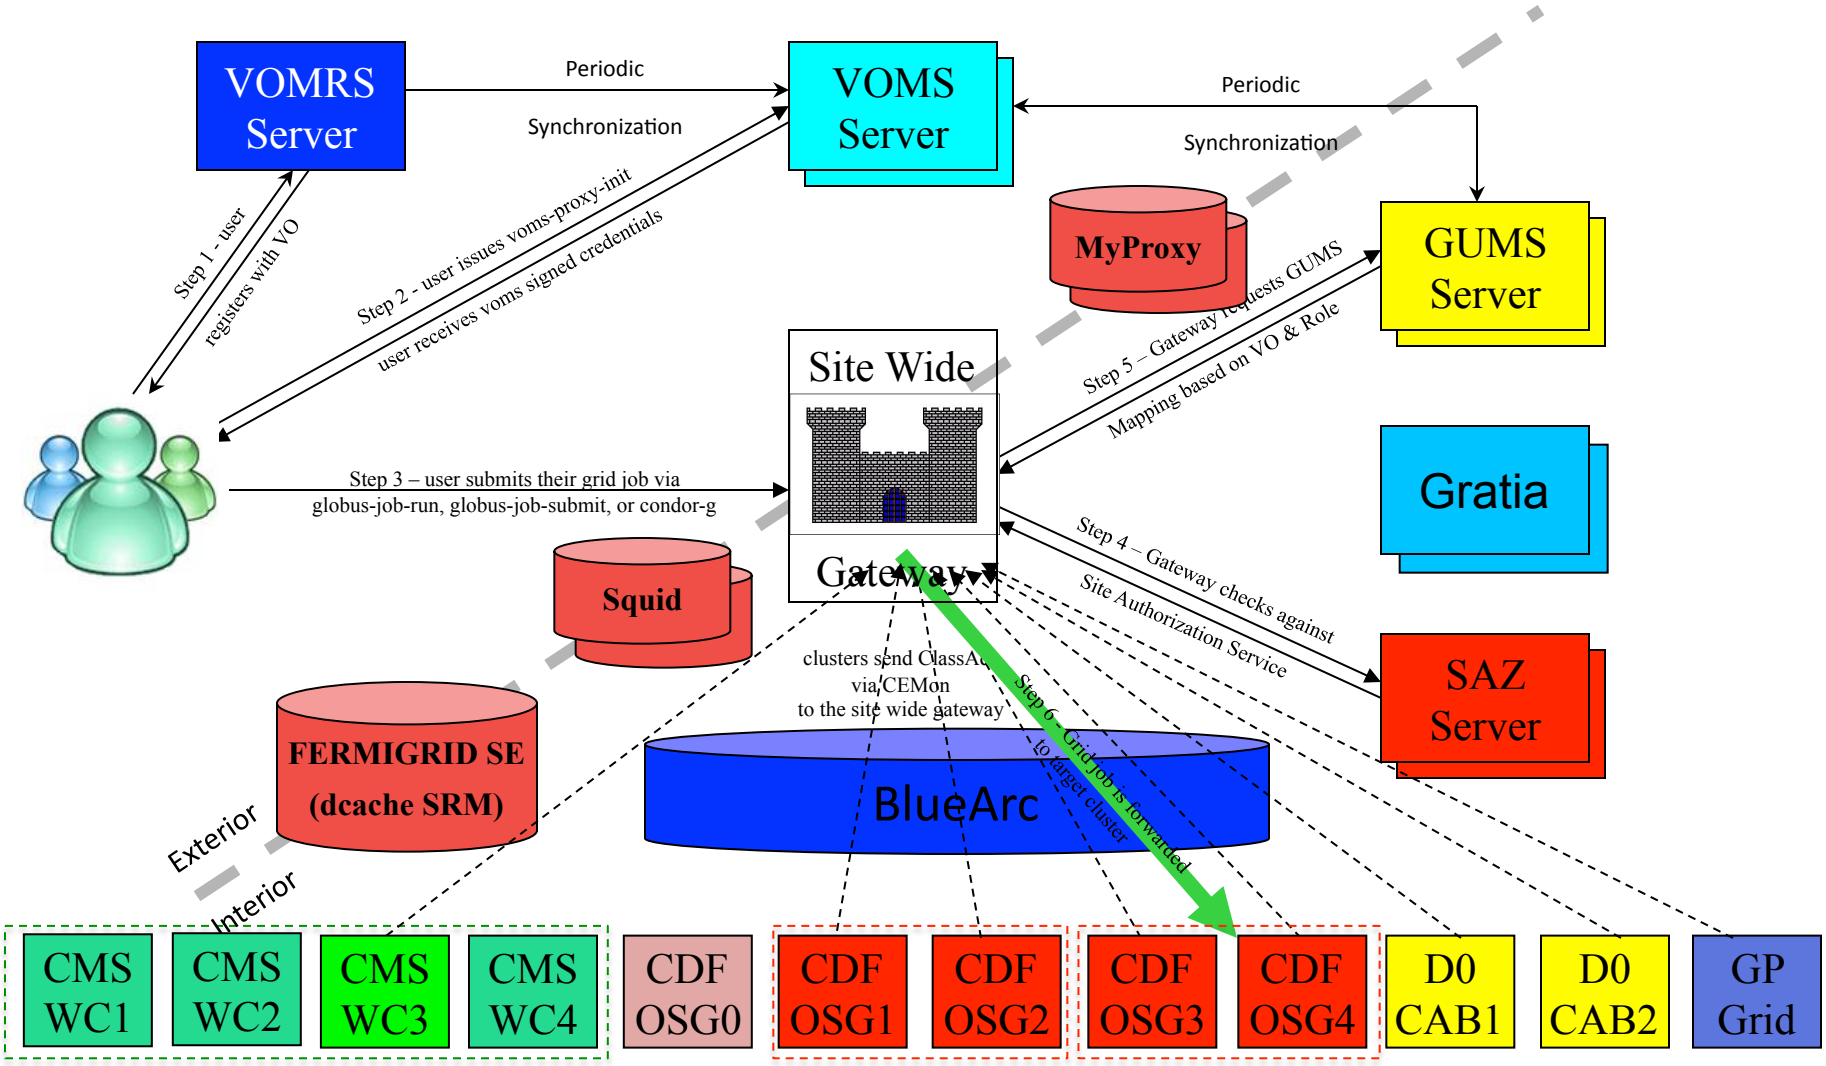

Measured Service Availability for the Past Year*

Service	Availability	Downtime
VOMS-HA	100%	0m
GUMS-HA	100%	0m
SAZ-HA (gatekeeper)	100%	0m
Squid-HA	99.988%	59.97m
MyProxy-HA	99.985%	78.80m
ReSS-HA	99.979%	107.96m
Gratia-HP	99.616%	2,014.32m
Database-HA	99.867	697.73m

Measured Average Occupancy and Utilization for the Past Year

Cluster	Average # of Batch Slots	Raw Occupancy Avg/Max	Effective Utilization Ang/Max
CDF	5,574	88.9%/100%	63.1%/97.4%
CMS	7,194	67.8%/99.9%	55.0%/99.9%
D0	5,284	70.8%/98.8%	46.5%/98.5%
GP	1,793	74.9%/99.2%	69.8%/99.9%
Total	19,849	74.9%/99.3%	56.2%/93.9%

FermiGrid - Evolution not Revolution

- We did not start with a massive project to transition to a [Campus] Grid infrastructure overnight.

- FermiGrid was commissioned over roughly a 18 month interval:
  - Ongoing discussions with stakeholders,
  - Establish initial set of central services based on these discussions [VOMS, GUMS],
  - We worked with each stakeholder to transition their cluster(s) to use the Grid infrastructure,
  - Periodically review the set of central services and deploy additional services as necessary/appropriate [SAZ, MyProxy, Squid, etc.].

FermiGrid -> FermiGrid-HA -> FermiGrid-HA2 Evolution

- If you have centrally provided services [eg. GUMS, SAZ] to multiple independent clusters, then you will eventually need to implement some sort of high availability service configuration.
  - Don't have to do this right off the bat, but it is useful to keep in mind when designing and implementing services

- FermiGrid -> FermiGrid-HA:
  - The high availability (HA) services were implemented when the growth and demand for the services showed that they were needed;

- We are now planning for FermiGrid-HA2:
  - HA2 will allow us to offer a geographically separated set of HA services that will survive an individual building power outage.

Conclusions

- Campus Grids offer significant cost savings.
  - FermiGrid has been operating as a Campus Grid for ~5 years.
  - We have made extensive use of "onsite" opportunistic running.
- Campus Grids do require a bit more infrastructure to establish and support.
  - This can be added incrementally.
  - Virtualization can be used to significantly lower the hardware cost of the additional infrastructure.
- FermiGrid took advantage of the onsite Kerberos Certificate Authority infrastructure to implement the "fermilab" umbrella VO.
- Many large higher education and research organizations have already deployed similar infrastructure (Shibboleth, InCommon, CiLogin-CA, etc.)
  - These can be leveraged to provide the necessary underpinnings for a Campus Grid.
- Campus Grids can also be integrated into larger Grid organizations (such as the Open Science Grid or TeraGrid) to give your community access to larger or specialized resources.
  - Of course it's nice if you are also willing to make your unused resources available for opportunistic access.

Fin

- Any Questions?

Other Considerations

- You will likely want to tie your (? centrally managed?) administration/staff/ faculty/student computer account data into your Campus Grid resources.
  - FermiGrid has implemented automated population of the "fermilab" virtual organization (VO) from our Central Name and Address Service (CNAS).

Why Do Campus Grids ?

- Improve utilization of (existing) resources – don't purchase resources when they are not needed.
  - Cost savings.
- Provide common administrative framework and user experience.
  - Cost savings.
- Buy resources (clusters) in "bulk" @ lower costs.
  - Cost savings.
- Lower maintenance costs.
  - Cost savings.
- Unified user interface will reduce the amount of user training required to make effective use of the resources.
  - Cost savings.

What are the drawbacks ?

- Additional centralized infrastructure to provision and support.
  - Additional costs.
  - Can be provisioned incrementally to manage buy-in costs.
  - Virtual machines can be used to lower buy-in costs.
- Can make problem diagnosis somewhat more complicated.
  - Correlation of multiple logs across administrative boundaries.
  - A central log repository is one mechanism to manage this.
- Not appropriate for all workloads.
  - Don't want financials running on the same resources as research.
- Have to learn (and teach the user community) how to route jobs to the appropriate resources.
  - Trivially parallel jobs require different resources than MPI jobs.
  - I/O intensive jobs require different resources than compute intensive jobs.
- Limited stakeholder buy-in may lead to a campus grid that's less interoperable than you might like.

Outline

- Definition of a Campus Grid
- FermiGrid
  - Pre-Grid Situation
  - Today
  - Architecture
  - Services
  - Metrics
- Evolution & Other Considerations
- Conclusions

Definition

- A Campus Grid is a distributed collection of [compute and storage] resources, provisioned by one or more stakeholders, that can be seamlessly accessed through one or more [Grid] portals.

FermiGrid HA Services - 1

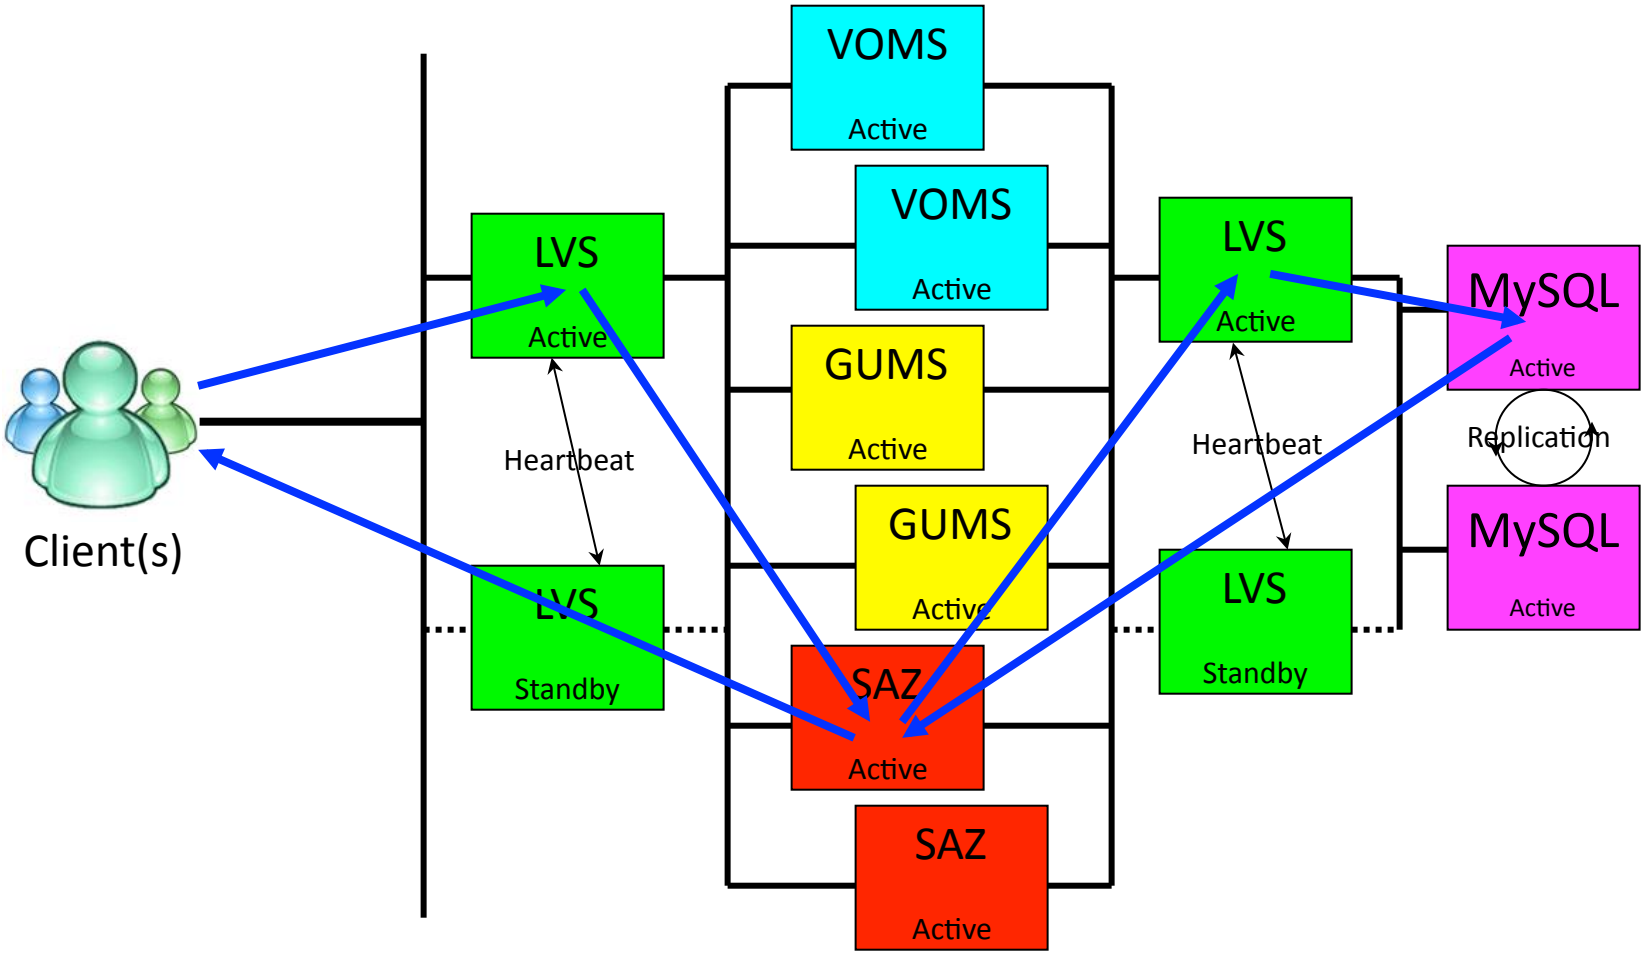

FermiGrid HA Services - 2

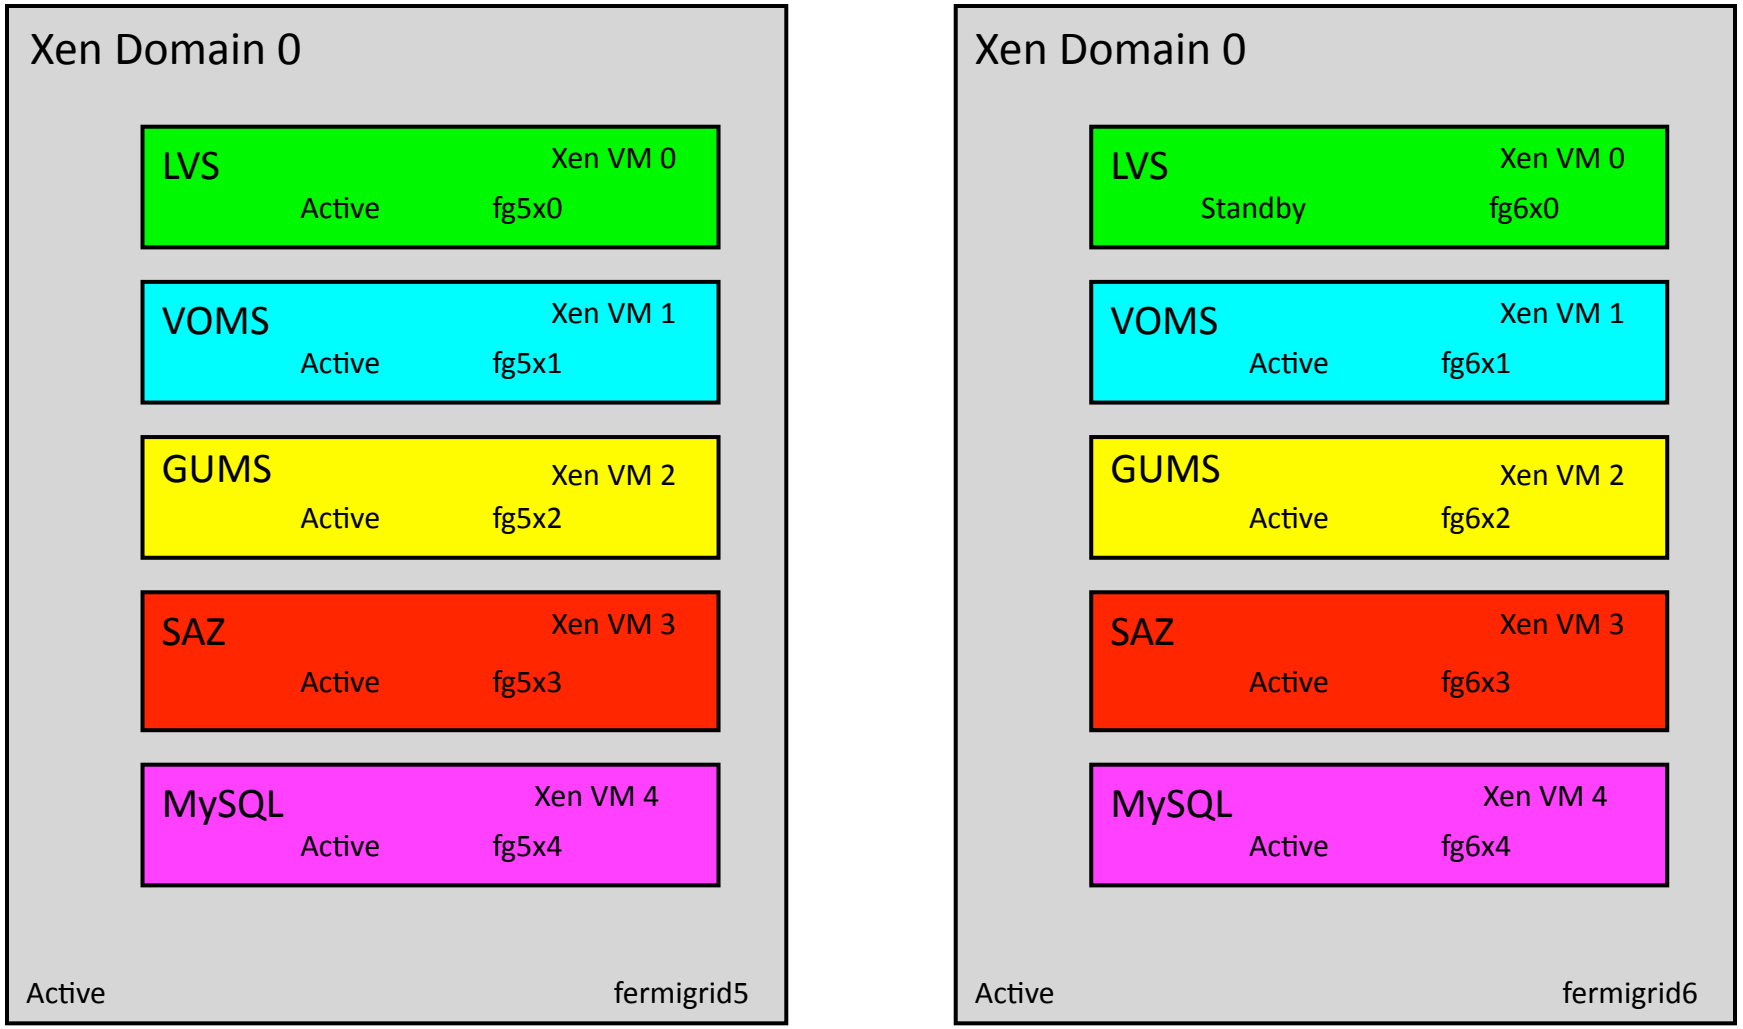

FermiGrid Utilization

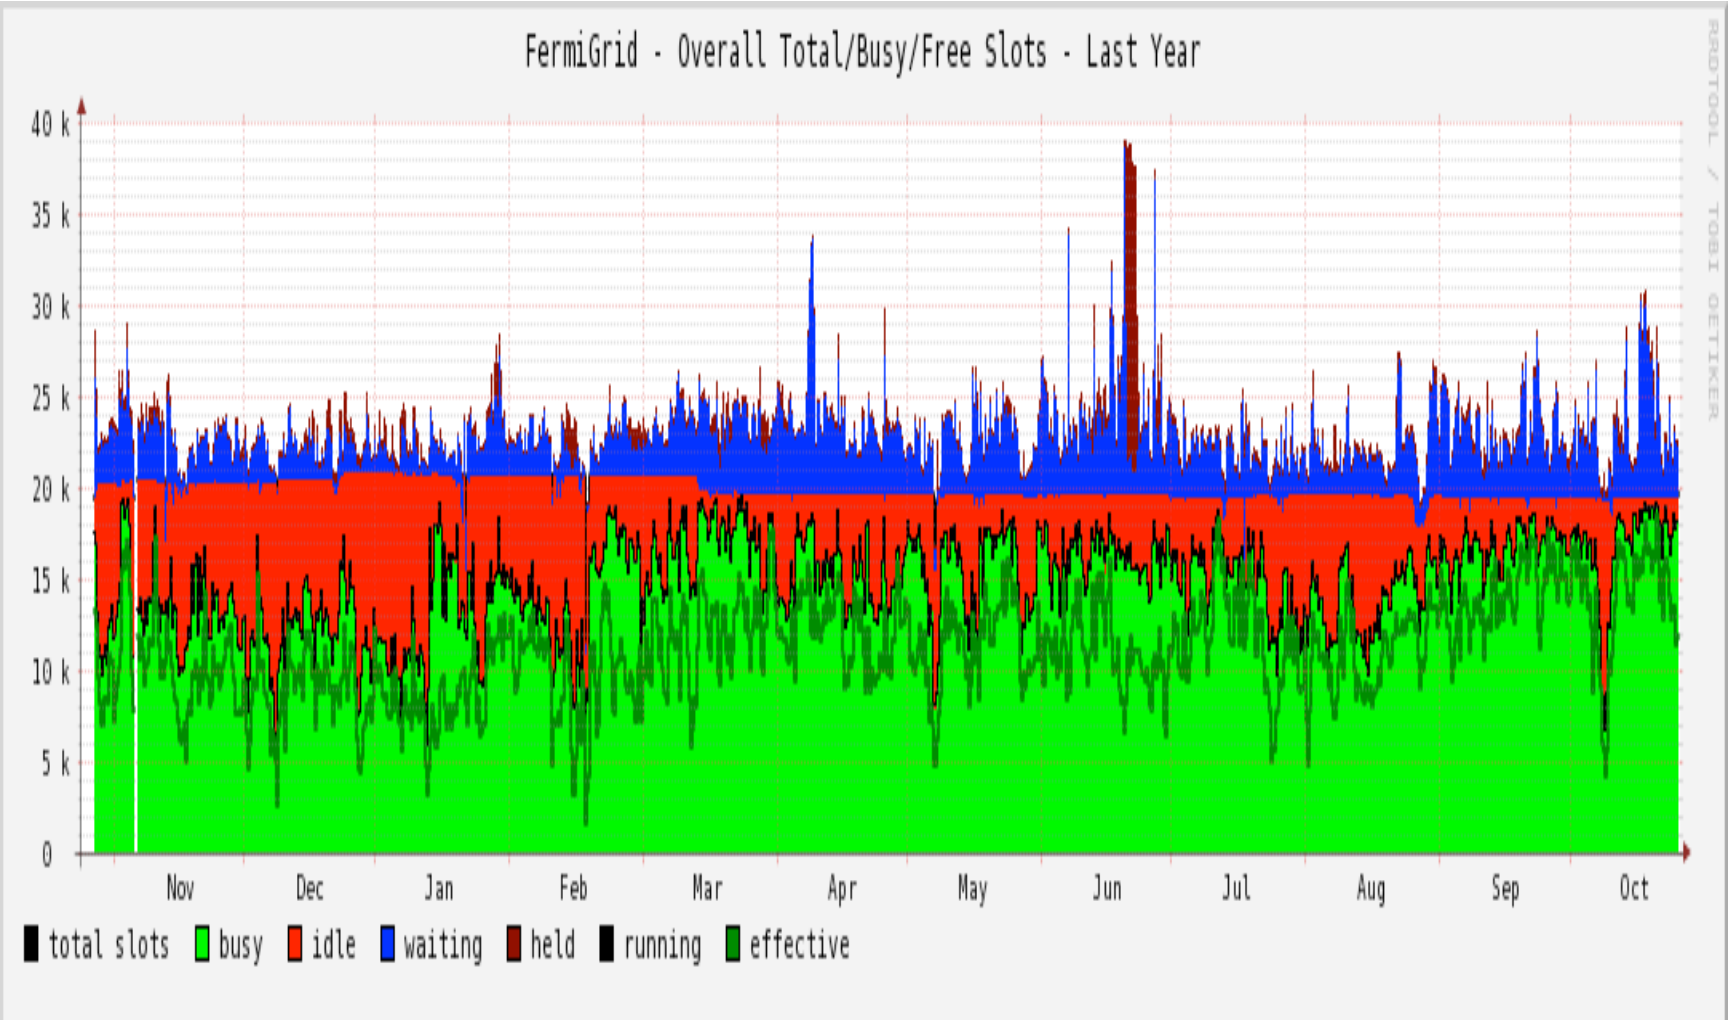

FermiGrid – Occupancy by VO

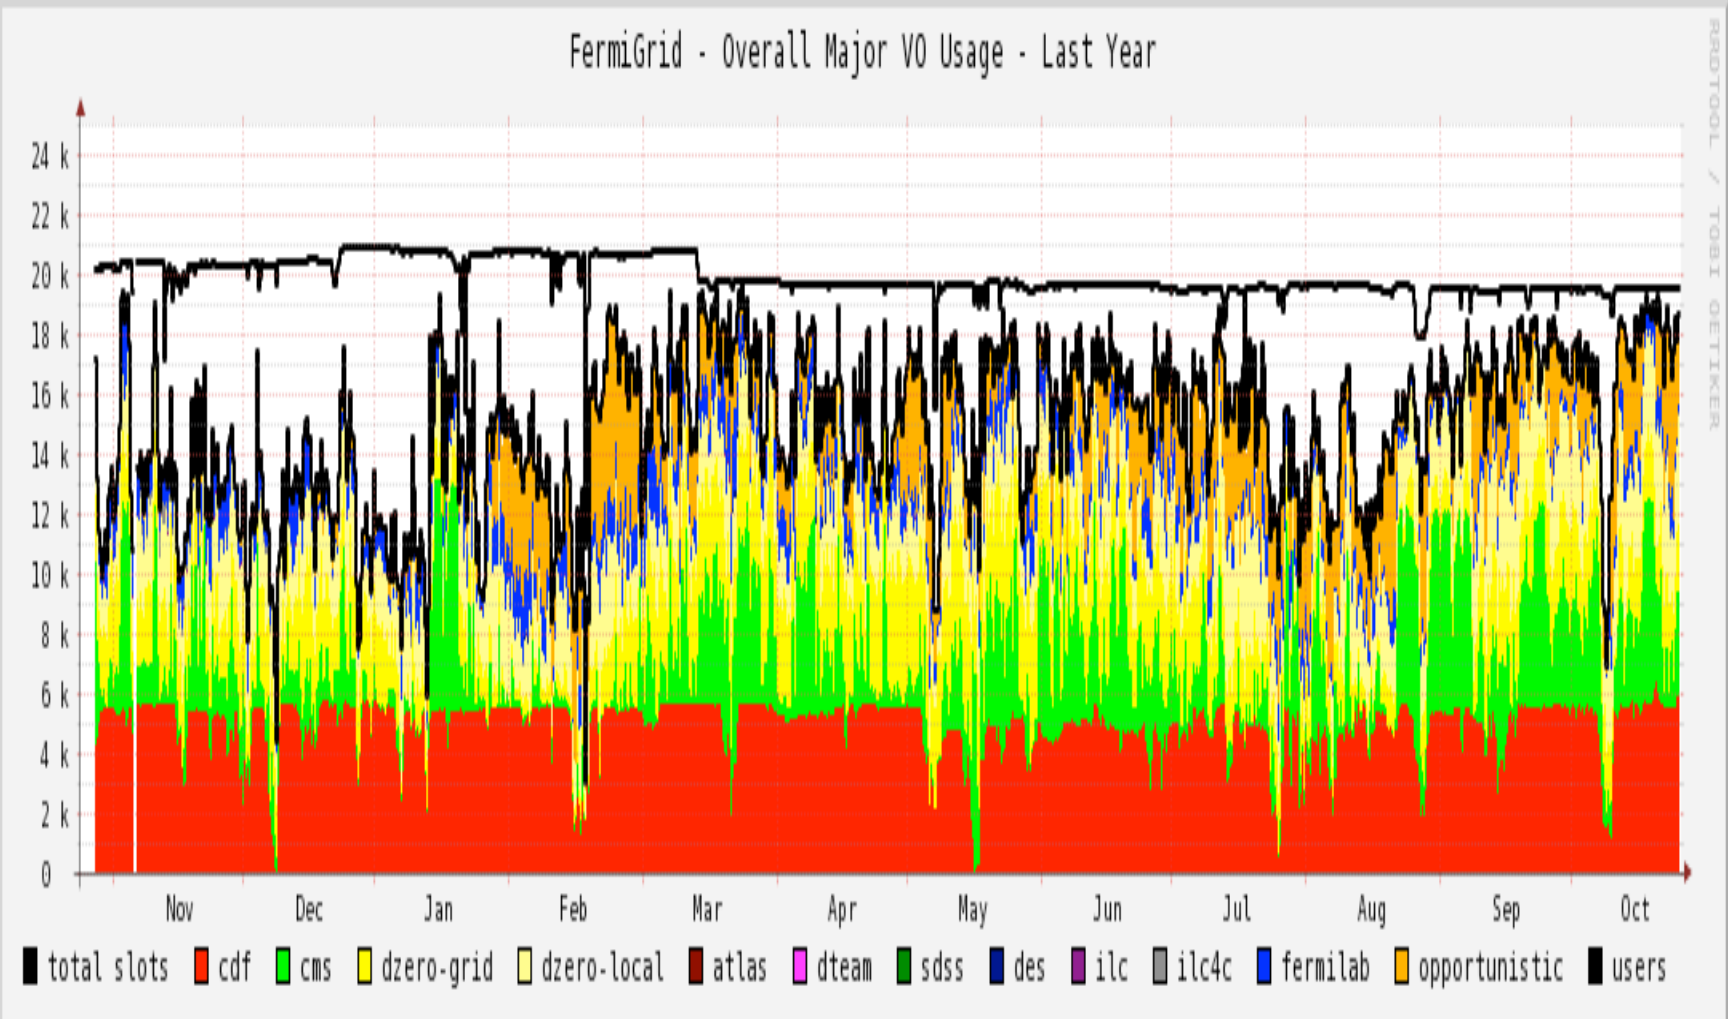

CMS – Occupancy by VO

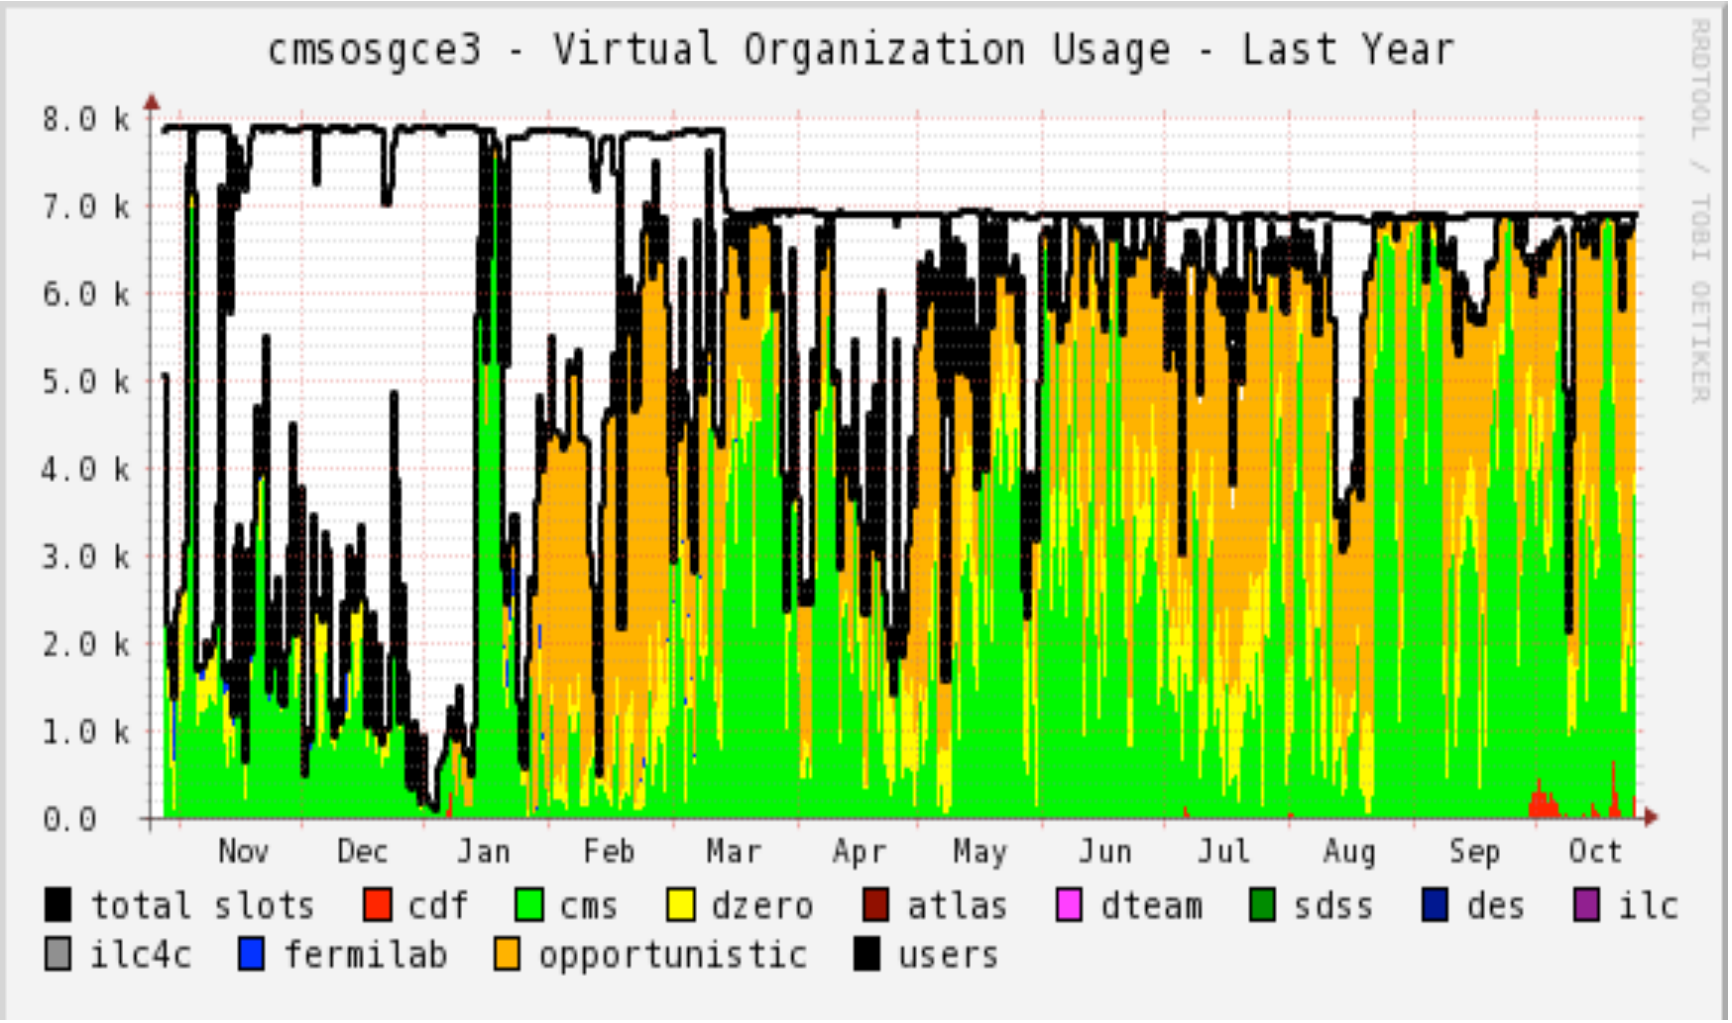

GP – Occupancy by VO

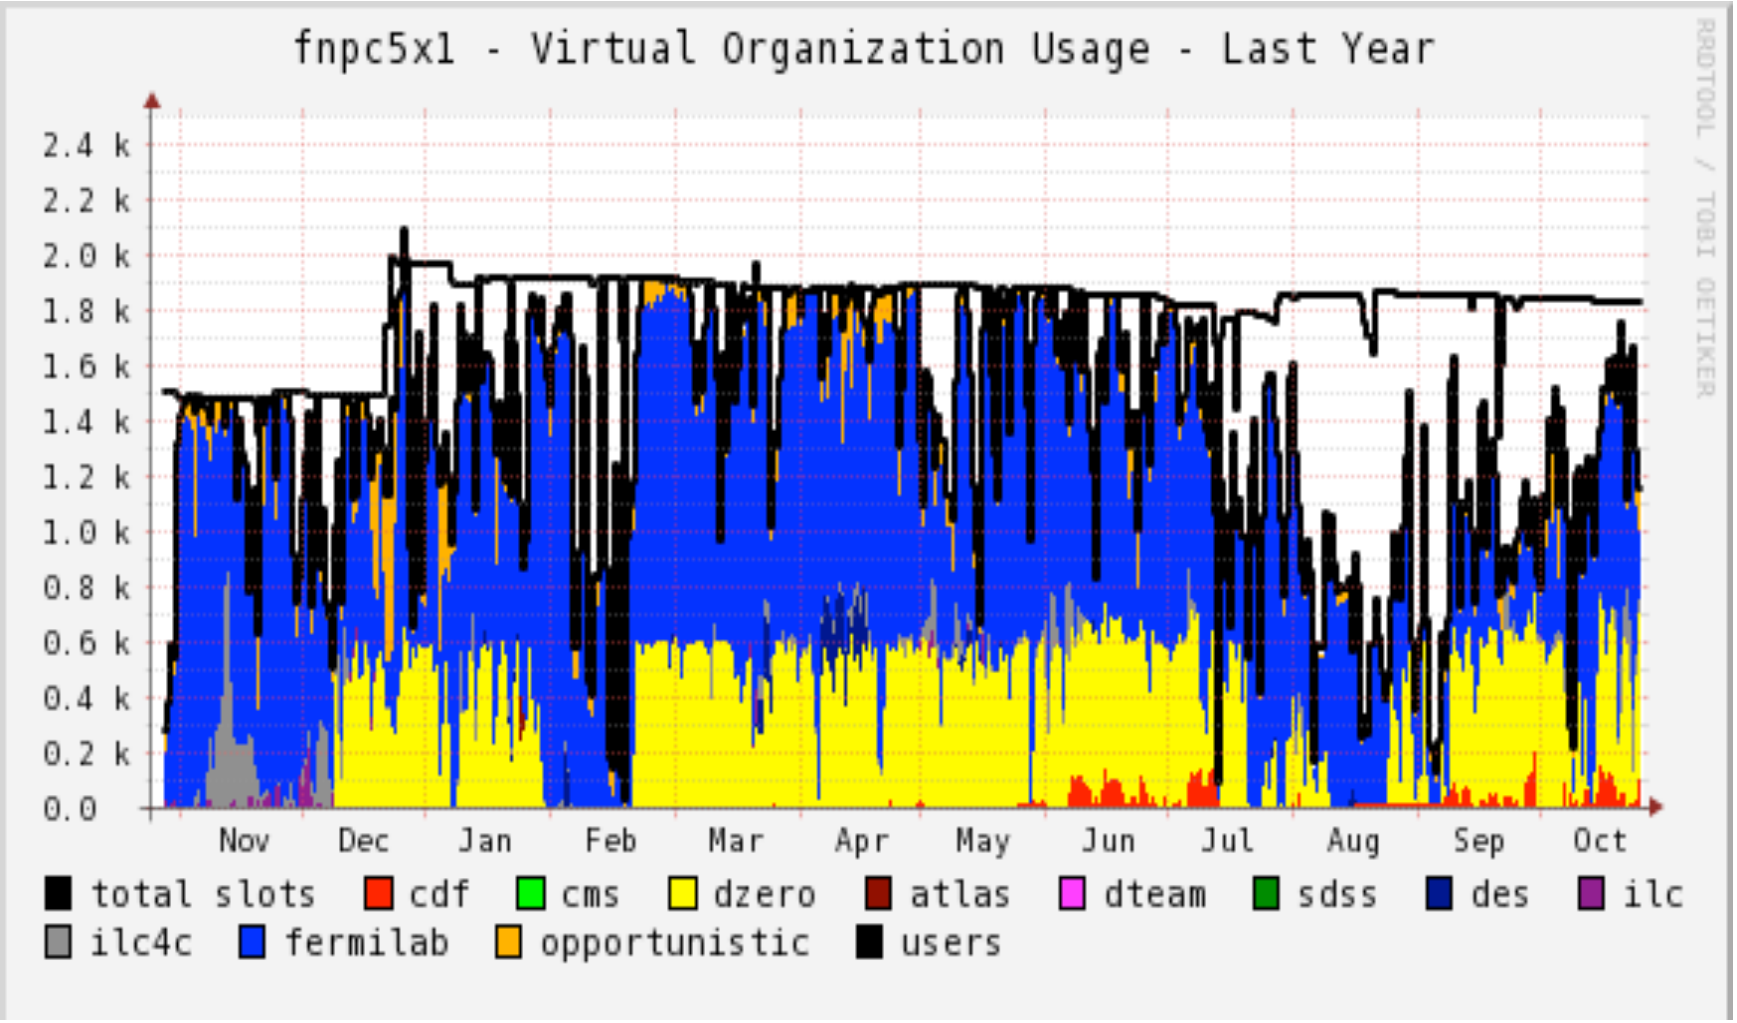

VOMS-PROXY-INIT calls

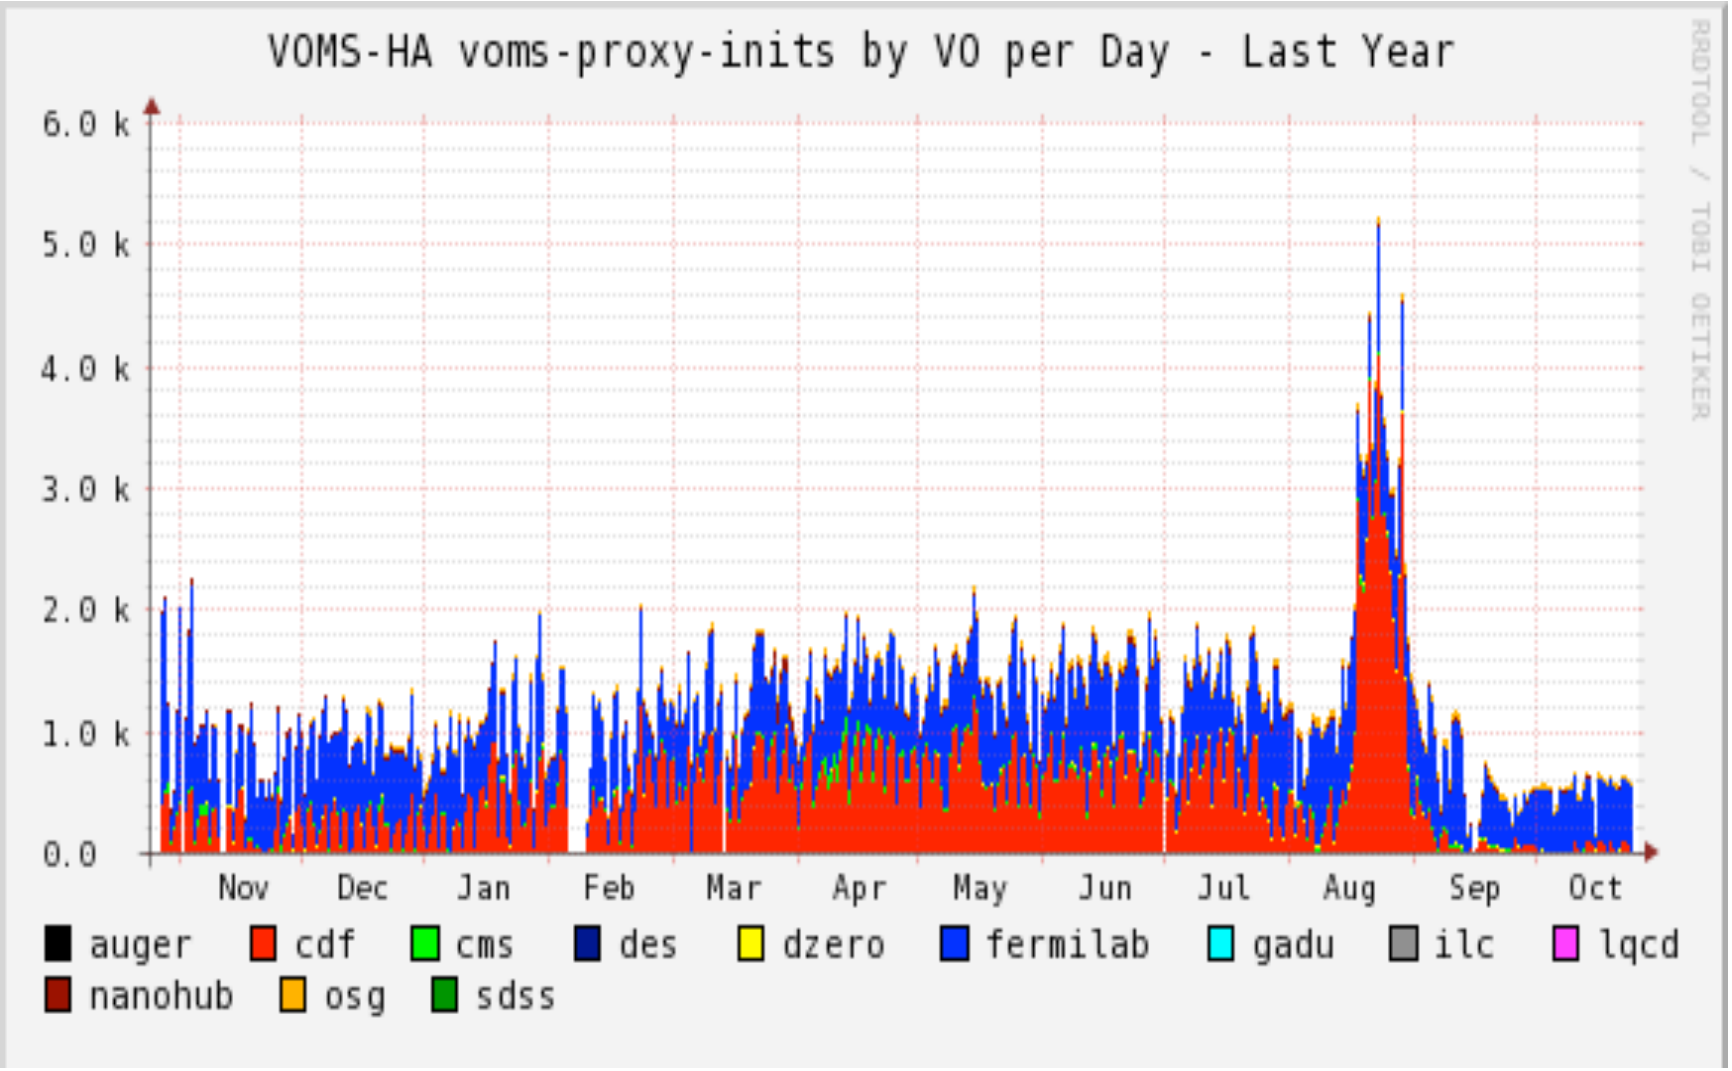

GUMS calls

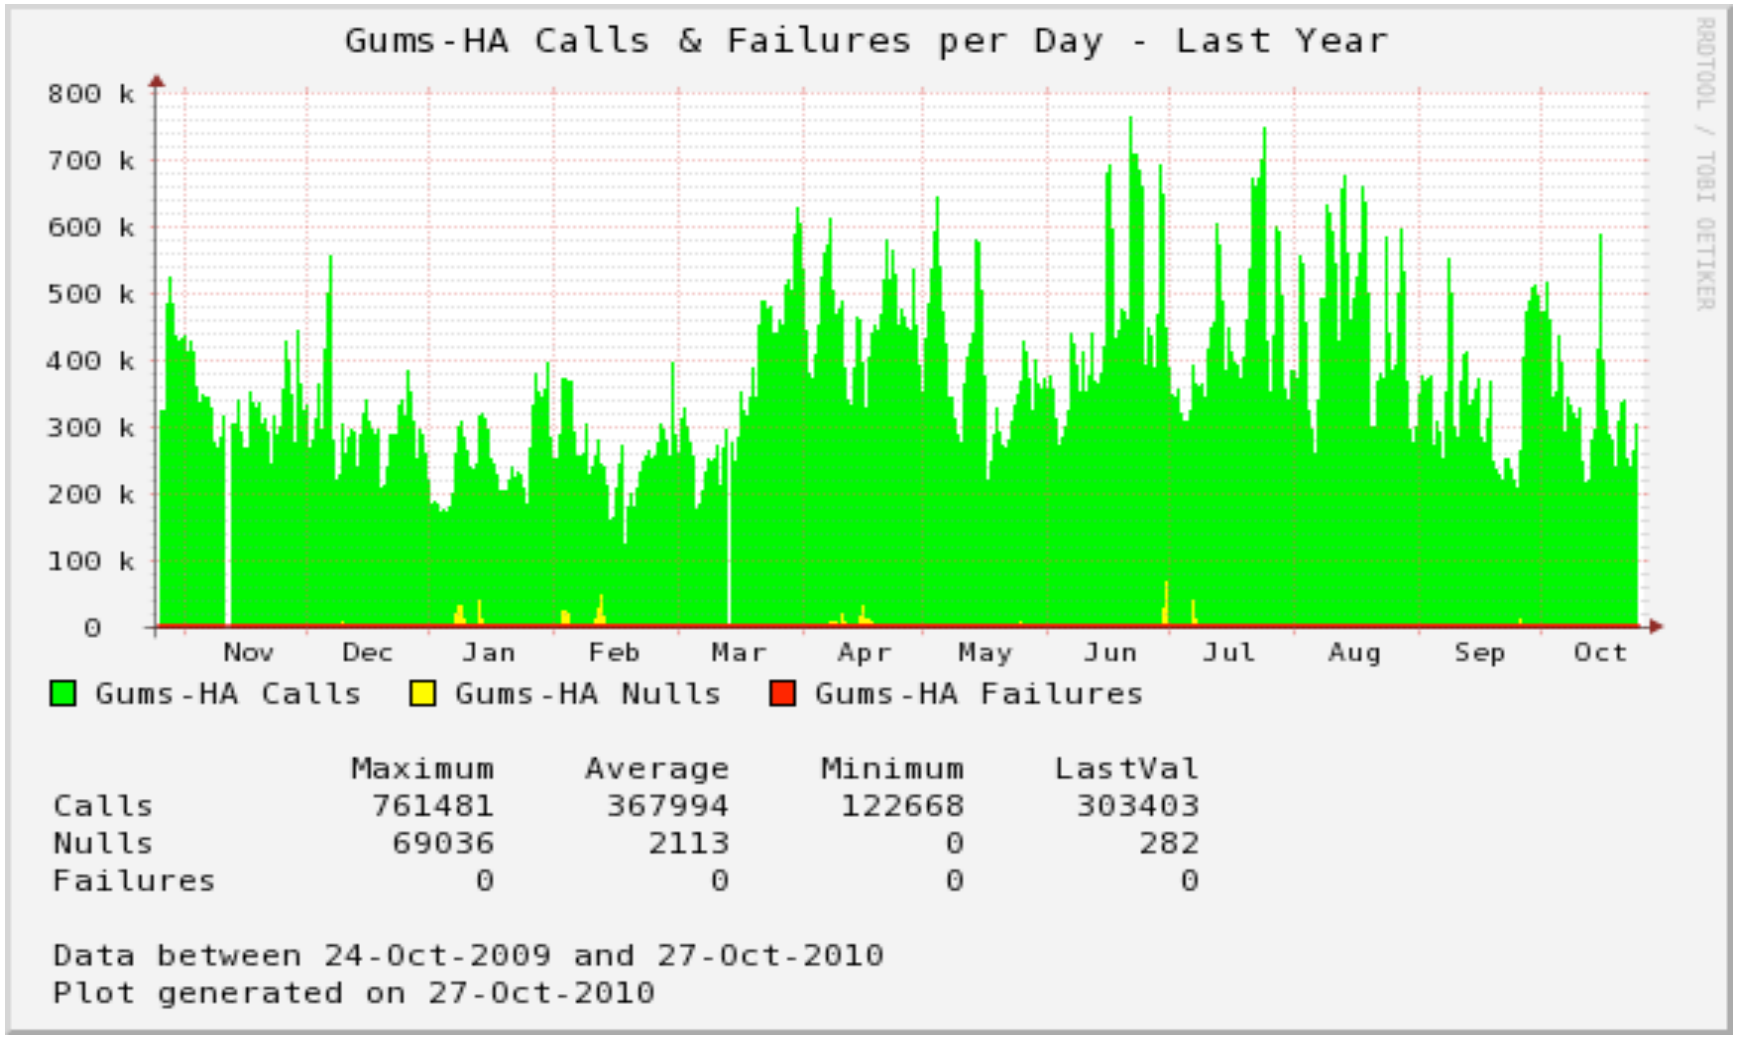

	Maximum	Average	Minimum	LastVal
Calls	761481	367994	122668	303403
Nulls	69036	2113	0	282
Failures	0	0	0	0

Data between 24-Oct-2009 and 27-Oct-2010
Plot generated on 27-Oct-2010

SAZ Calls

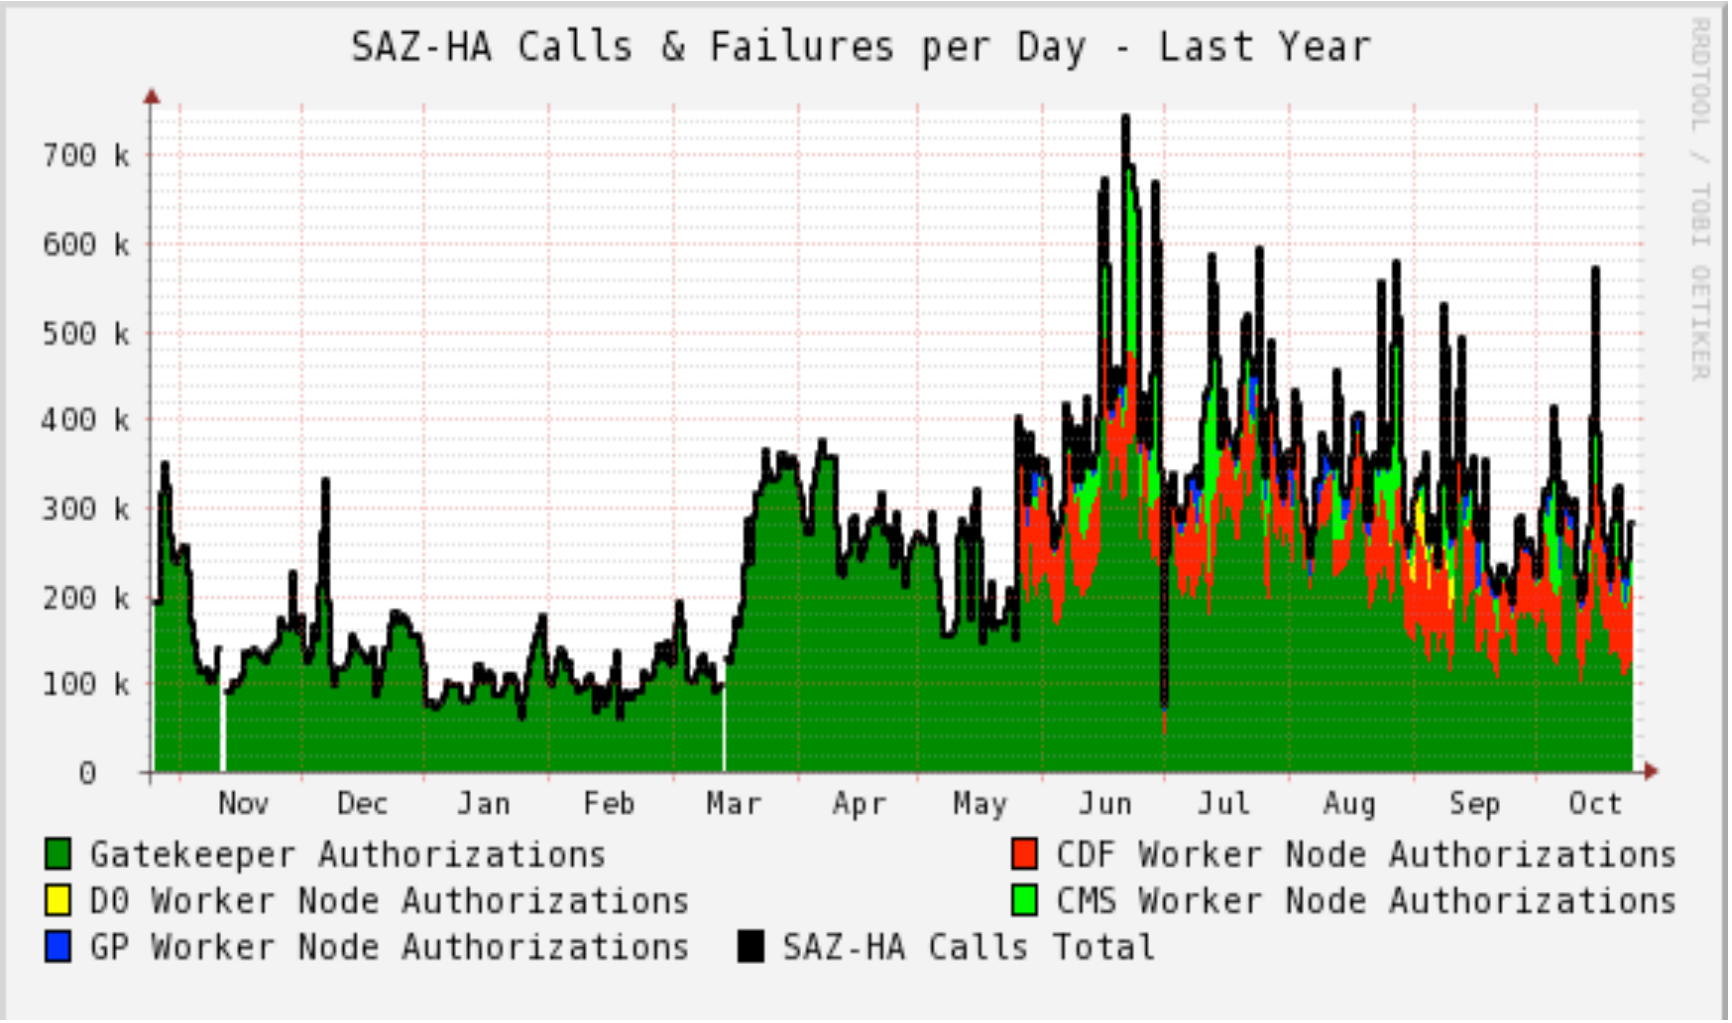
SAZ-HA Calls & Failures per Day - Last Year
700 k
600 k
500 k
400 k
300 k
200 k
100 k
0
Nov
Dec
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
Gatekeeper Authorizations
D0 Worker Node Authorizations
GP Worker Node Authorizations
CDF Worker Node Authorizations
CMS Worker Node Authorizations
SAZ-HA Calls Total
RRDTOOL / TOBI OETIKER

Squid Calls

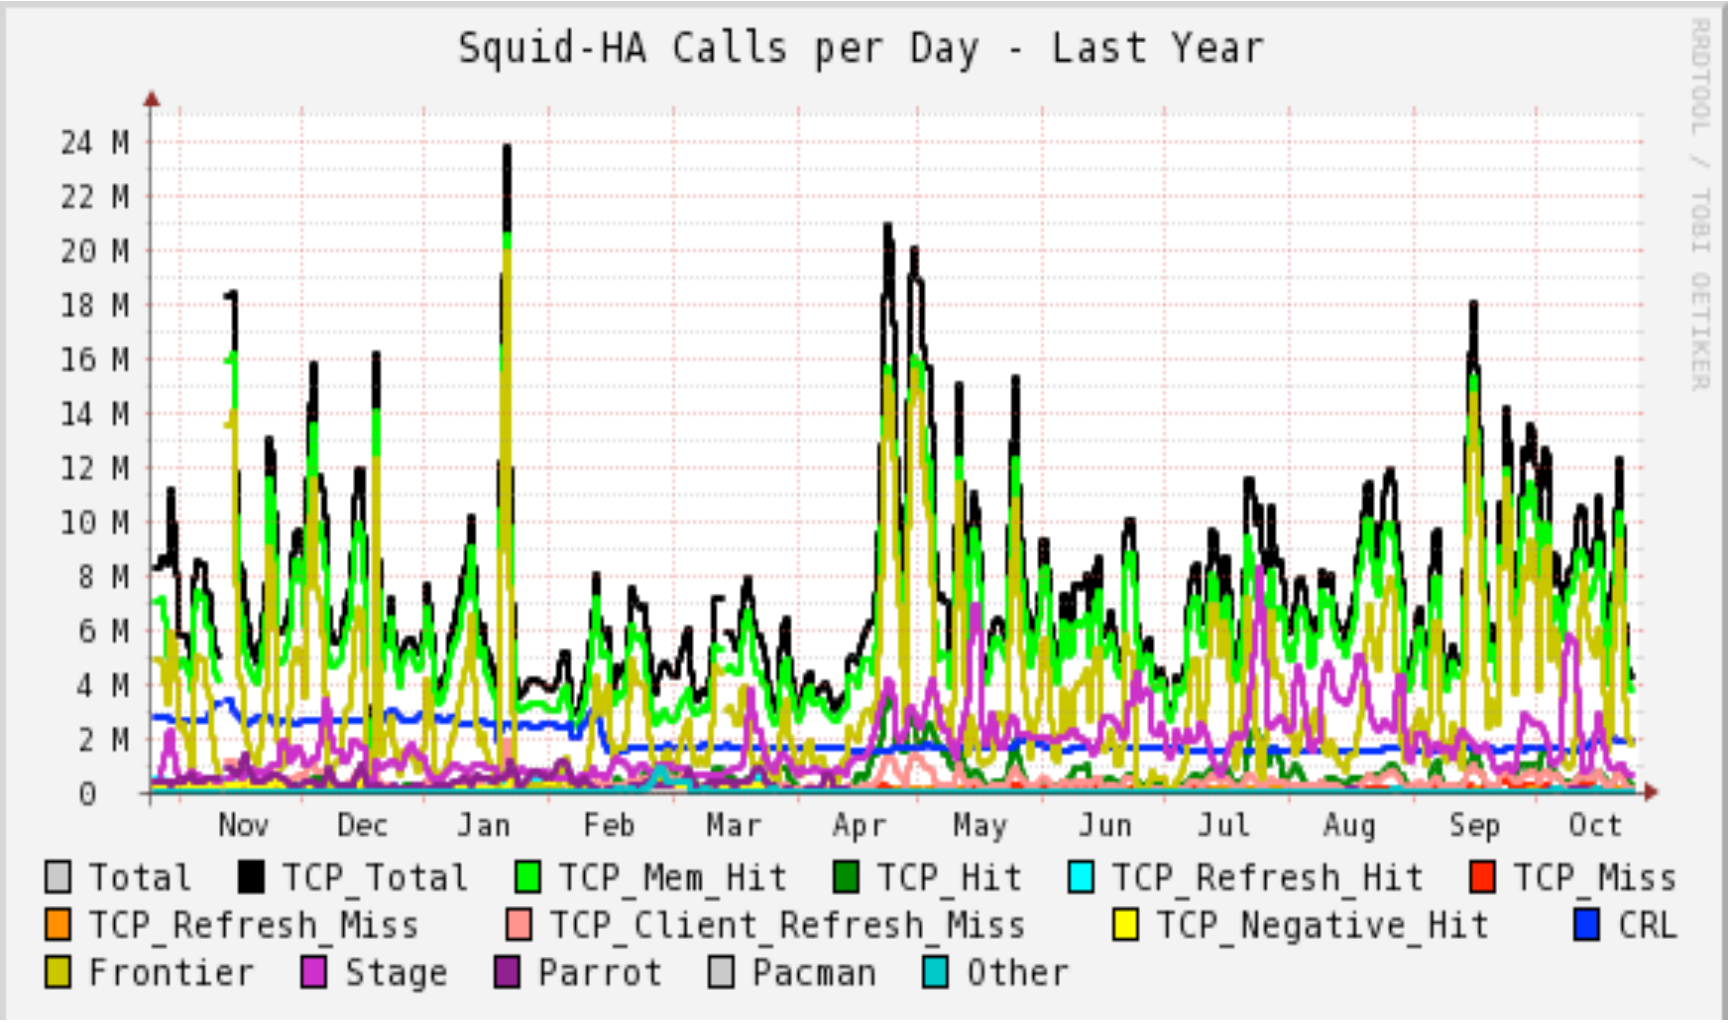
Squid-HA Calls per Day - Last Year
RRDTOOL / TOBI OETIKER
24 M
22 M
20 M
18 M
16 M
14 M
12 M
10 M
8 M
6 M
4 M
2 M
0
Nov
Dec
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
Total
TCP_Total
TCP_Mem_Hit
TCP_Hit
TCP_Refresh_Hit
TCP_Miss
TCP_Refresh_Miss
TCP_Client_Refresh_Miss
TCP_Negative_Hit
CRL
Frontier
Stage
Parrot
Pacman
Other