ATLAS EXPERIMENT

UCIRVINE

DESIGNING A NEW INFRASTRUCTURE FOR ATLAS ONLINE WEB SERVICES

□ **General description**
□ **Current and future infrastructure for Web services**
□ **Security and requirements**
□ **Conclusions**

Diana Scannicchio
University of California, Irvine
on behalf of
Ballestrero S, University of Johannesburg
Brasolin F, INFN Bologna
Sanchez Pineda A, University and INFN Udine, ICTP, CERN
Twomey MS, University of Washington

Introduction - ATLAS

- ❑ ATLAS (A Toroidal LHC Apparatus) is one of the major experiments at the Large Hadron Collider (LHC) at CERN
  - ★ the computing farm is composed of ~4000 servers processing the data read out from ~100 million detector channels through multiple trigger levels
- ❑ ATLAS servers are connected to a dedicated network, ATLAS Technical and Control Network (ATCN) separated from the CERN General Purpose Network (GPN)
  - ★ the ATLAS Gateway servers, connected to both ATCN and GPN, allow access to ATCN from GPN and the reverse
    - ➢ user access is restricted
- ❑ One of the requirements is to be able to continue data taking for a couple of days in case of a disconnection from GPN
  - ★ all core services are duplicated inside ATCN
    - ➢ Active Directory, DHCP, DNS, NTP, repositories

The capability to monitor the ongoing data taking and all the involved applications is essential to debug and intervene promptly to ensure efficient data taking

Introduction - Web Services

- ❑ The base of the current web service architecture was designed a few years ago, at the beginning of the ATLAS operation (Run 1 - 2009)
  - ★ it has shown its limits, as the trend towards Web User Interfaces continues
  - ★ its increasing complexity became an issue for maintenance, growth and security

- ❑ A review of the overall web services architecture has become necessary
  - ★ taking into account the current and future needs of the upcoming LHC Run 3 - 2021

- ❑ Investigation and road map started in order to re-design the web services system to better operate and monitor the ATLAS detector maintaining
  - ★ the security of critical services, such as Detector Control System (DCS)
  - ★ the separation of remote monitoring and on-site control according to ATLAS policies

Current web services architecture

- ❑ The current architecture was intended to serve primarily static content from a Network Attached Storage (NAS) in ATCN
  - ★ to enhance security, separate web servers are used for internal (ATCN) and external (GPN and public internet) access
- ❑ An increasing number of dynamic web-based UIs have been added to the static content
  - ★ to provide new functionalities and replace legacy desktop UIs
  - ★ typically served by applications on VMs inside ATCN and made accessible externally via chained Apache reverse HTTP proxies

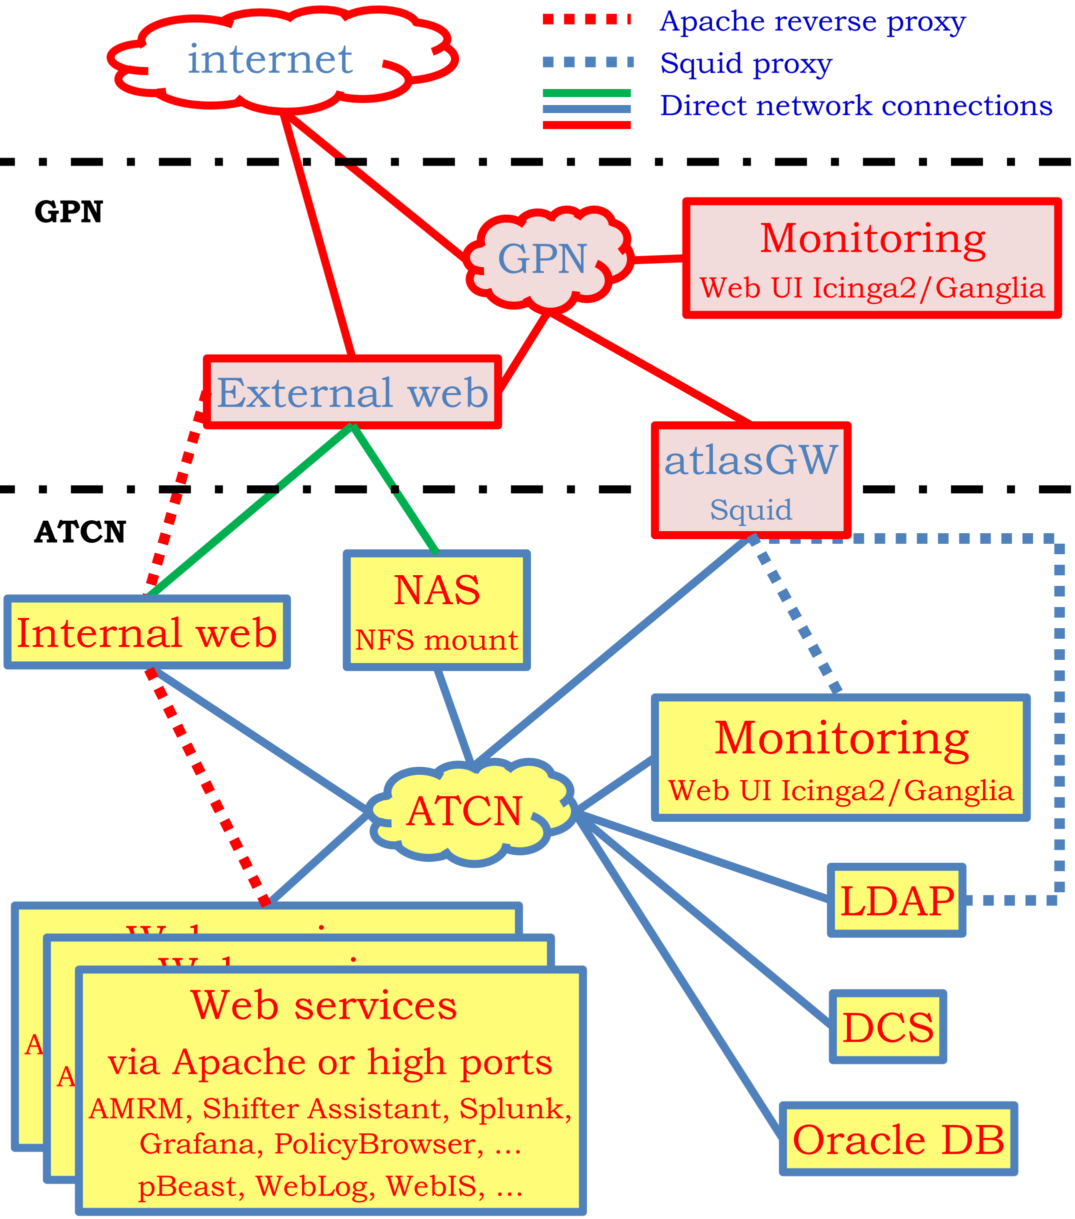

Current web services architecture

- ❑ Chained (and entangled) reverse proxies configured in Apache
  - ★ from the external web server (GPN) to the internal one (ATCN)
  - ★ from the internal web server to the VMs providing specific application
  - ★ from the VM itself to the specific port provided by the users' application
- ❑ Squid on the ATLAS Gateways is used to give access to the web content not managed correctly by the Apache reverse proxy (e.g. Icinga2)
  - ★ it is a caching and forwarding HTTP web proxy

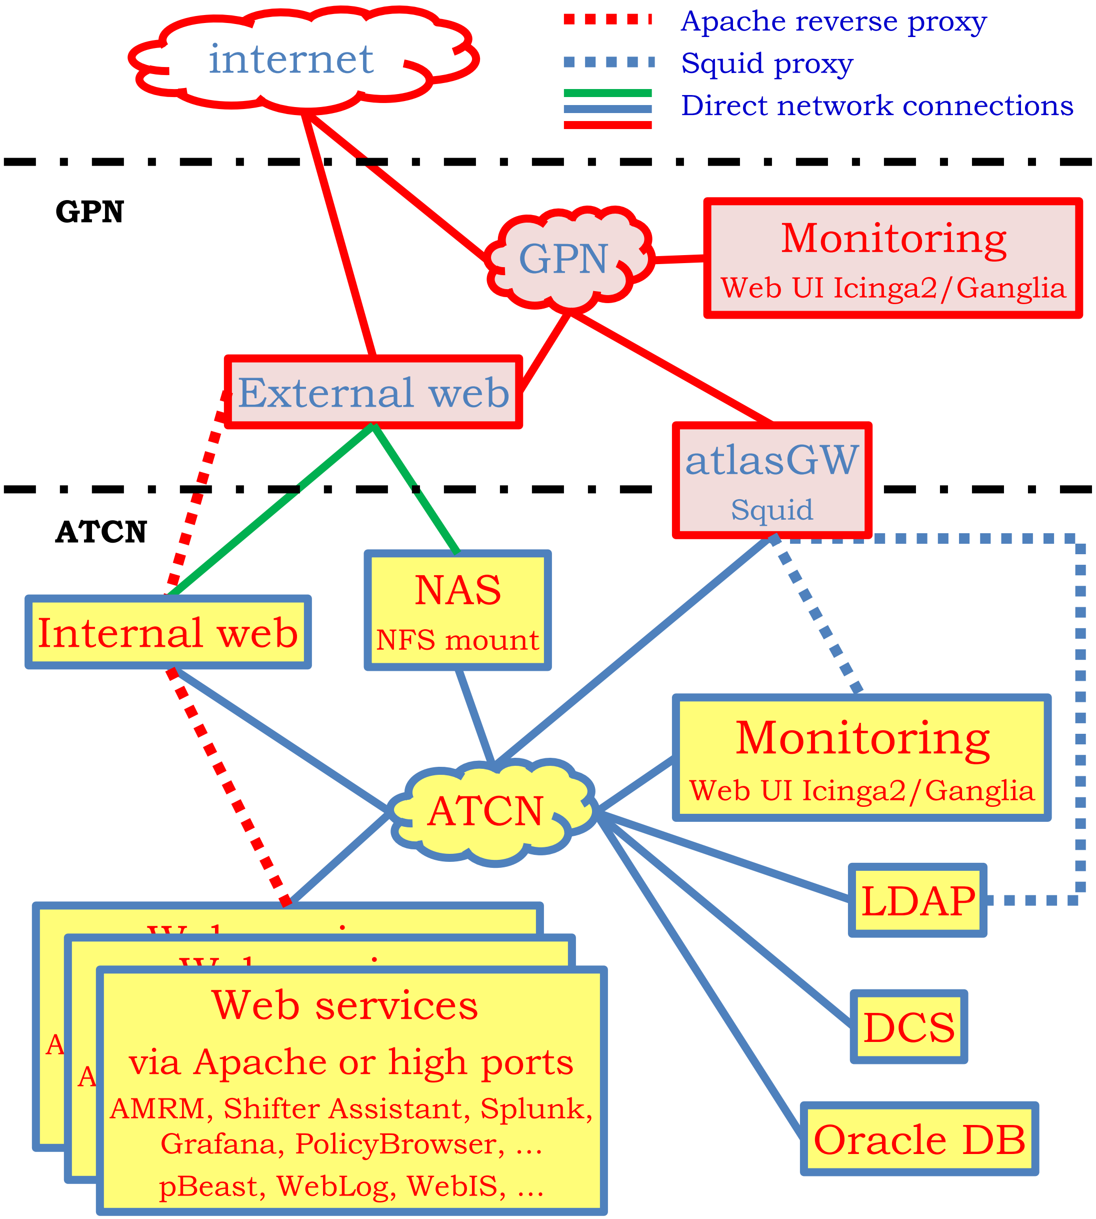

Current requirements / features

- ❑ Redundancy
  - ★ the service was not considered as critical
  - ★ currently cold spare nodes are available for both the internal and external web servers, the switch should take ~30 minutes because of manual intervention needed
    - ➢ aliases to be moved
    - ➢ Apache configuration to be adjusted
- ❑ Availability from ATCN, GPN and outside CERN domain
  - ★ control room shift crew access the needed information from inside ATCN
  - ★ on-call experts may need to access information from GPN and/or outside the CERN domain
  - ★ Single-Sign On (SSO) on the external web server for access restriction
- ❑ Survive a GPN disconnection
  - ★ the internal web server grants access from inside ATCN
- ❑ Ability to expose or not a service outside ATCN
  - ★ some services are accessible only from inside ATCN
  - ★ some applications (e.g. twiki) writable only from inside ATCN

Improvements

❑ Improve redundancy
- ★ usage of `keepalived` would grant a smooth switch over between two servers

❑ Improve security removing the direct network link
- ★ from the external web server to the internal one
- ★ from the external web server to the NAS

❑ Clean up and simplify Squid configuration
- ★ used on the ATLAS Gateway to grant access from/to specific sites

❑ Simplify the overall Apache configuration removing the complex chained reverse proxies

New design

- ❑ Our investigation to re-design the web services system led us to look at HAProxy
  - ★ it is a free, very fast and reliable solution offering
    - ➢ high availability, load balancing and proxying for TCP and HTTP-based applications
  - ★ "Its mode of operation makes its integration into existing architectures easy and riskless, while offering the possibility not to expose fragile web servers to the net"

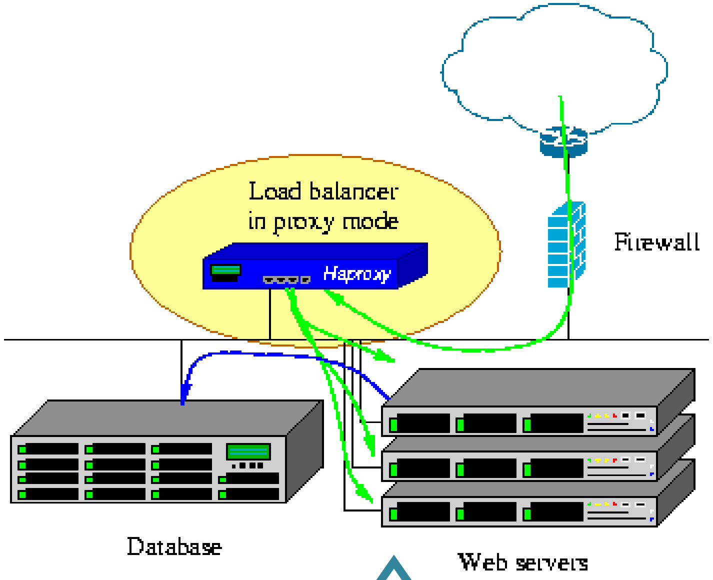

From http://www.haproxy.org/

- ❑ One server has been installed: configuration and testing are ongoing
  - ★ one network interface connected to GPN and one to ATCN
  - ★ Virtual IPs configured for both to be used by `keepalived`
- ❑ A second server will be added for redundancy

New schema and details

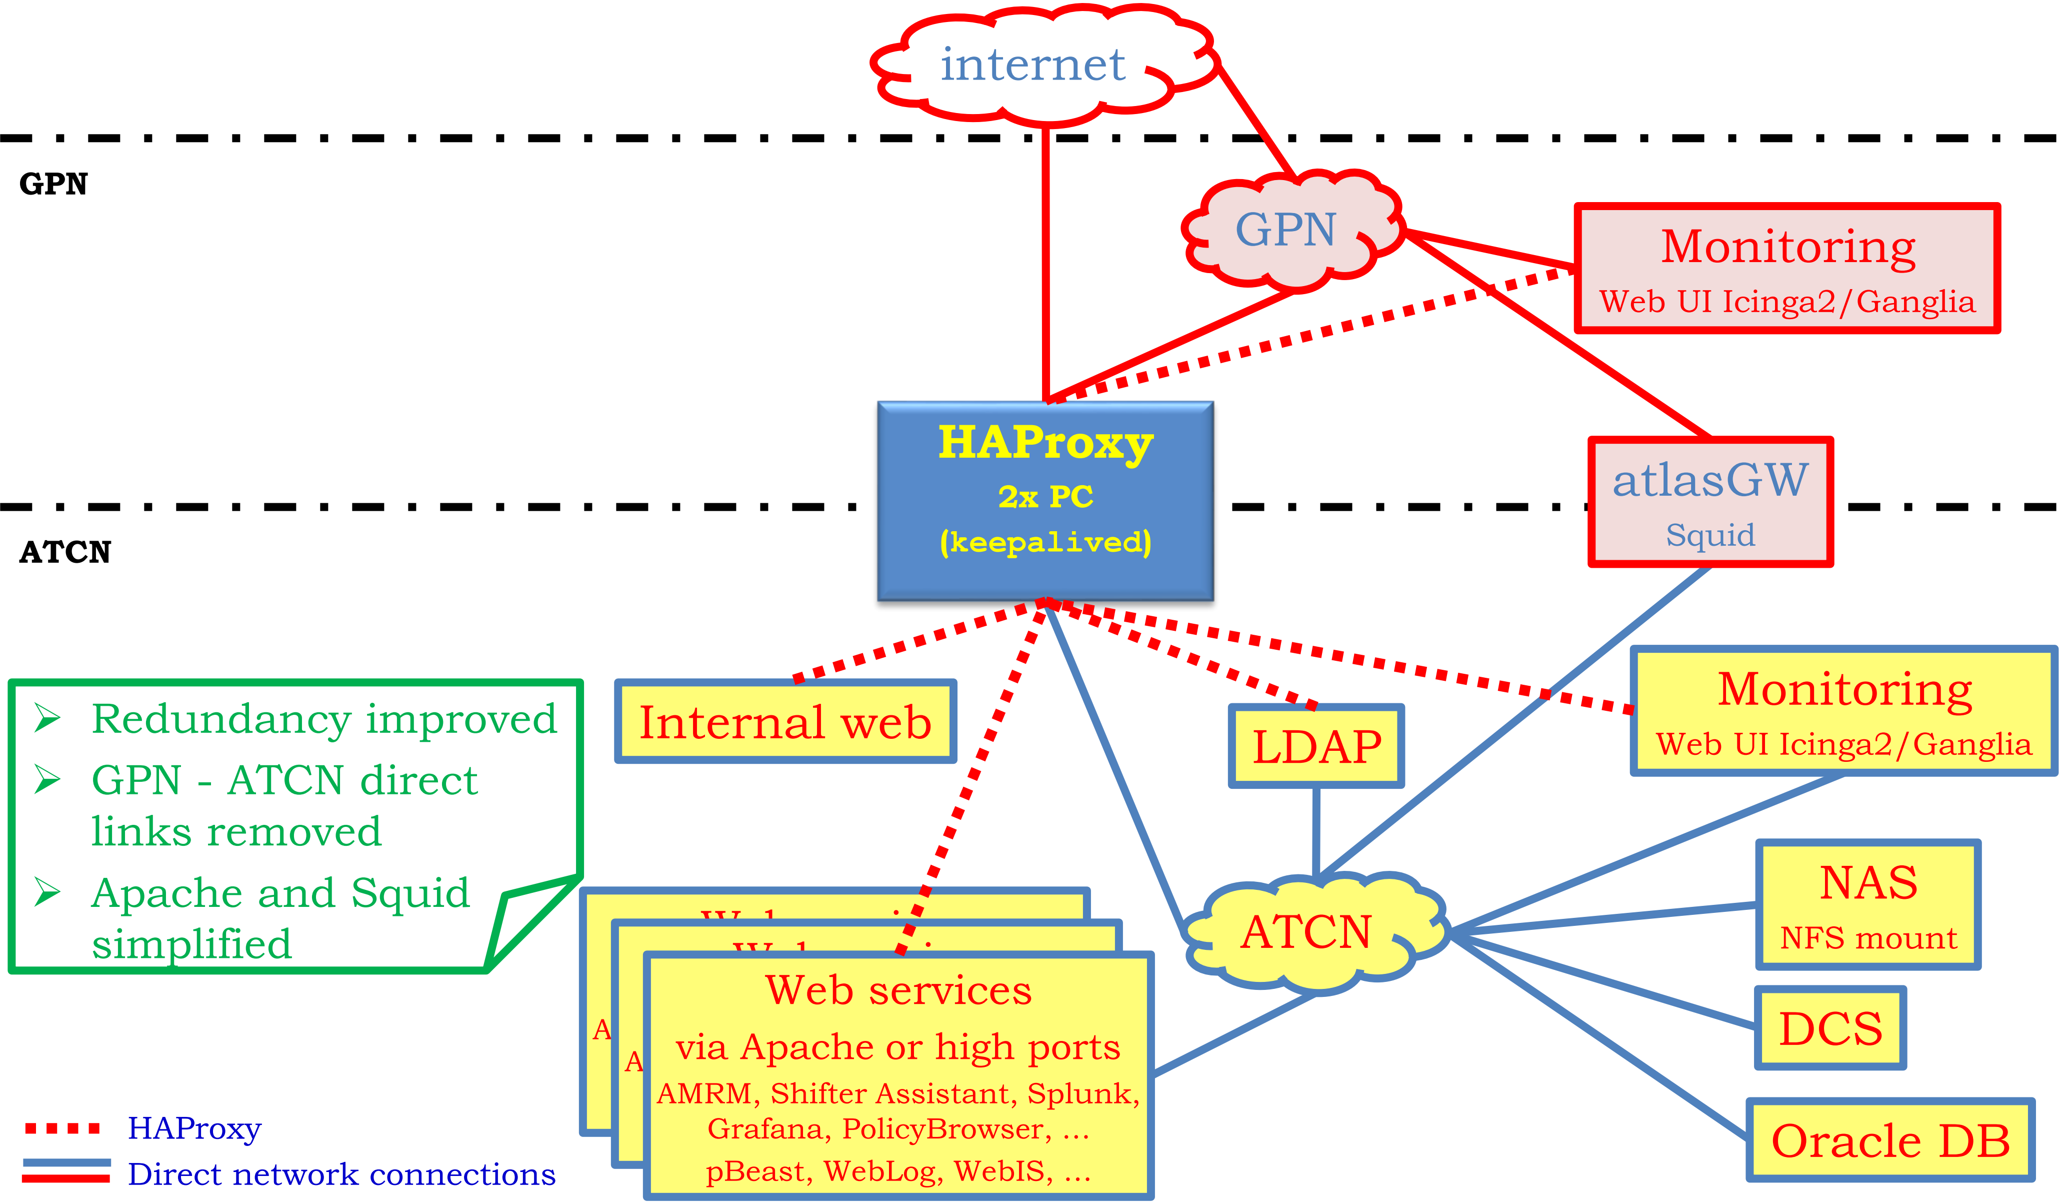

- ➢ Redundancy improved
- ➢ GPN - ATCN direct links removed
- ➢ Apache and Squid simplified

Advantages /disadvantages

Advantages

- ❑ Security improved
  - ★ direct connection removed
  - ★ Squid cleaned
- ❑ Redundancy improved
  - ★ via `keepalived`
- ❑ Apache configuration simplified on the web server
  - ★ multiple and chained reverse proxies removed
- ❑ Usage of a modern tool made for proxying and load balancing

Disadvantages

- ❑ Aliases corresponding to the Apache Virtual Hosts configuration needs to be defined and used by `HAProxy`
  - ★ new URLs are defined to access the services
    - ➢ users will need to get used to them
    - ➢ temporary re-direct could be added

Security perspective

- ❑ On all the nodes providing web services
  - ★ iptables
  - ★ updates from the CERN live repositories are applied immediately
  - ★ Rootkit Hunter
    - ➢ Unix based tool that scans for rootkits, backdoors and possible local exploits
  - ★ no user access
  - ★ read-only mounted web areas from the NAS
    - ➢ read-write mount only from a dedicated VM in ATCN with limited user access

- ❑ On the GPN node exposed to outside CERN domain
  - ★ security scans by CERN IT are performed
  - ★ log files are collected and sent to CERN IT to be analysed

- ❑ Samhain scans the NAS web areas exposed by the web servers
  - ★ integrity checker and intrusion detection system

Future and conclusions

- ❑ Proof of concept
  - ★ we managed to prove that the `HAProxy` is a viable solution
    - ➢ monitoring URL accessible without the need of Squid
    - ➢ basic static web pages available
    - ➢ some of the web applications (VMs based) are accessible
- ❑ HAProxy configuration is performed in steps
  - ★ current production environment cannot be disrupted
  - ★ the current and the new configuration coexist
- ❑ Features and requirements still to be addressed and ironed out
  - ★ Twiki pages are read-only from GPN/outside and read-write from ATCN
    - ➢ documentation for shifters and experts
  - ★ ATLAS DNS to be modified to answer with the internal (ATCN) IP to the requests sent to the external (GPN) one
    - ➢ to survive a disconnection from GPN
  - ★ access to applications provided via a specific network port only
  - ★ Single-Sign On will soon not be available, replacement to be investigated
- ❑ The transition should be completed before Run 3 - 2021

Complementary material

New and current schemas

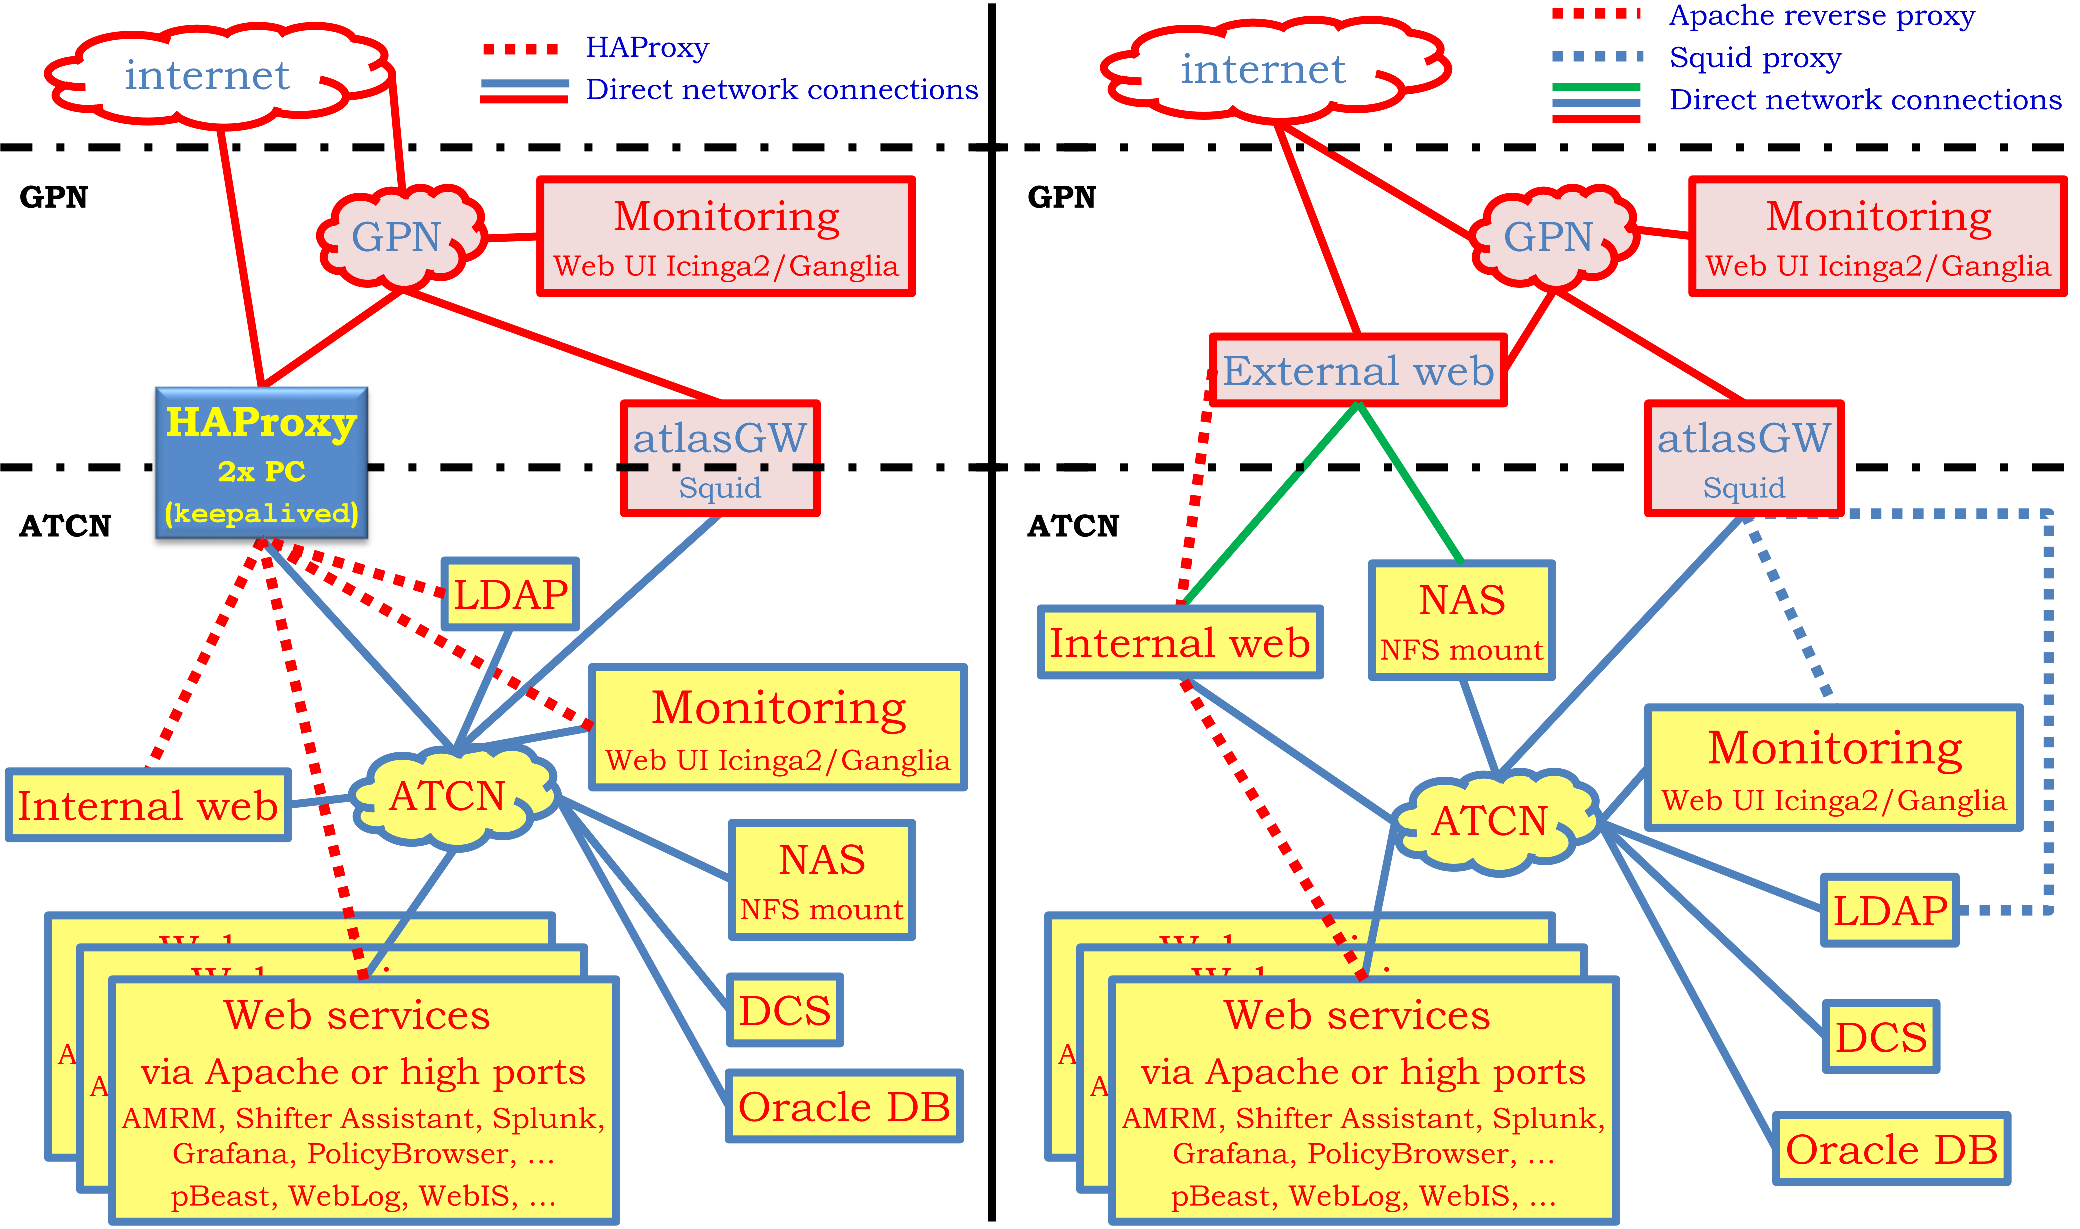

■■■■ HAProxy
〓 Direct network connections

internet

GPN

GPN

Monitoring
Web UI Icinga2/Ganglia

HAProxy
2x PC
(keepalived)

atlasGW
Squid

ATCN

LDAP

Monitoring
Web UI Icinga2/Ganglia

Internal web

ATCN

NAS
NFS mount

DCS

Oracle DB

Web services
via Apache or high ports
AMRM, Shifter Assistant, Splunk,
Grafana, PolicyBrowser, …
pBeast, WebLog, WebIS, …

■■■■ Apache reverse proxy
■■■■ Squid proxy
〓 Direct network connections

internet

GPN

GPN

Monitoring
Web UI Icinga2/Ganglia

External web

atlasGW
Squid

ATCN

Internal web

NAS
NFS mount

Monitoring
Web UI Icinga2/Ganglia

ATCN

LDAP

DCS

Oracle DB

Web services
via Apache or high ports
AMRM, Shifter Assistant, Splunk,
Grafana, PolicyBrowser, …
pBeast, WebLog, WebIS, …

Current schema and details

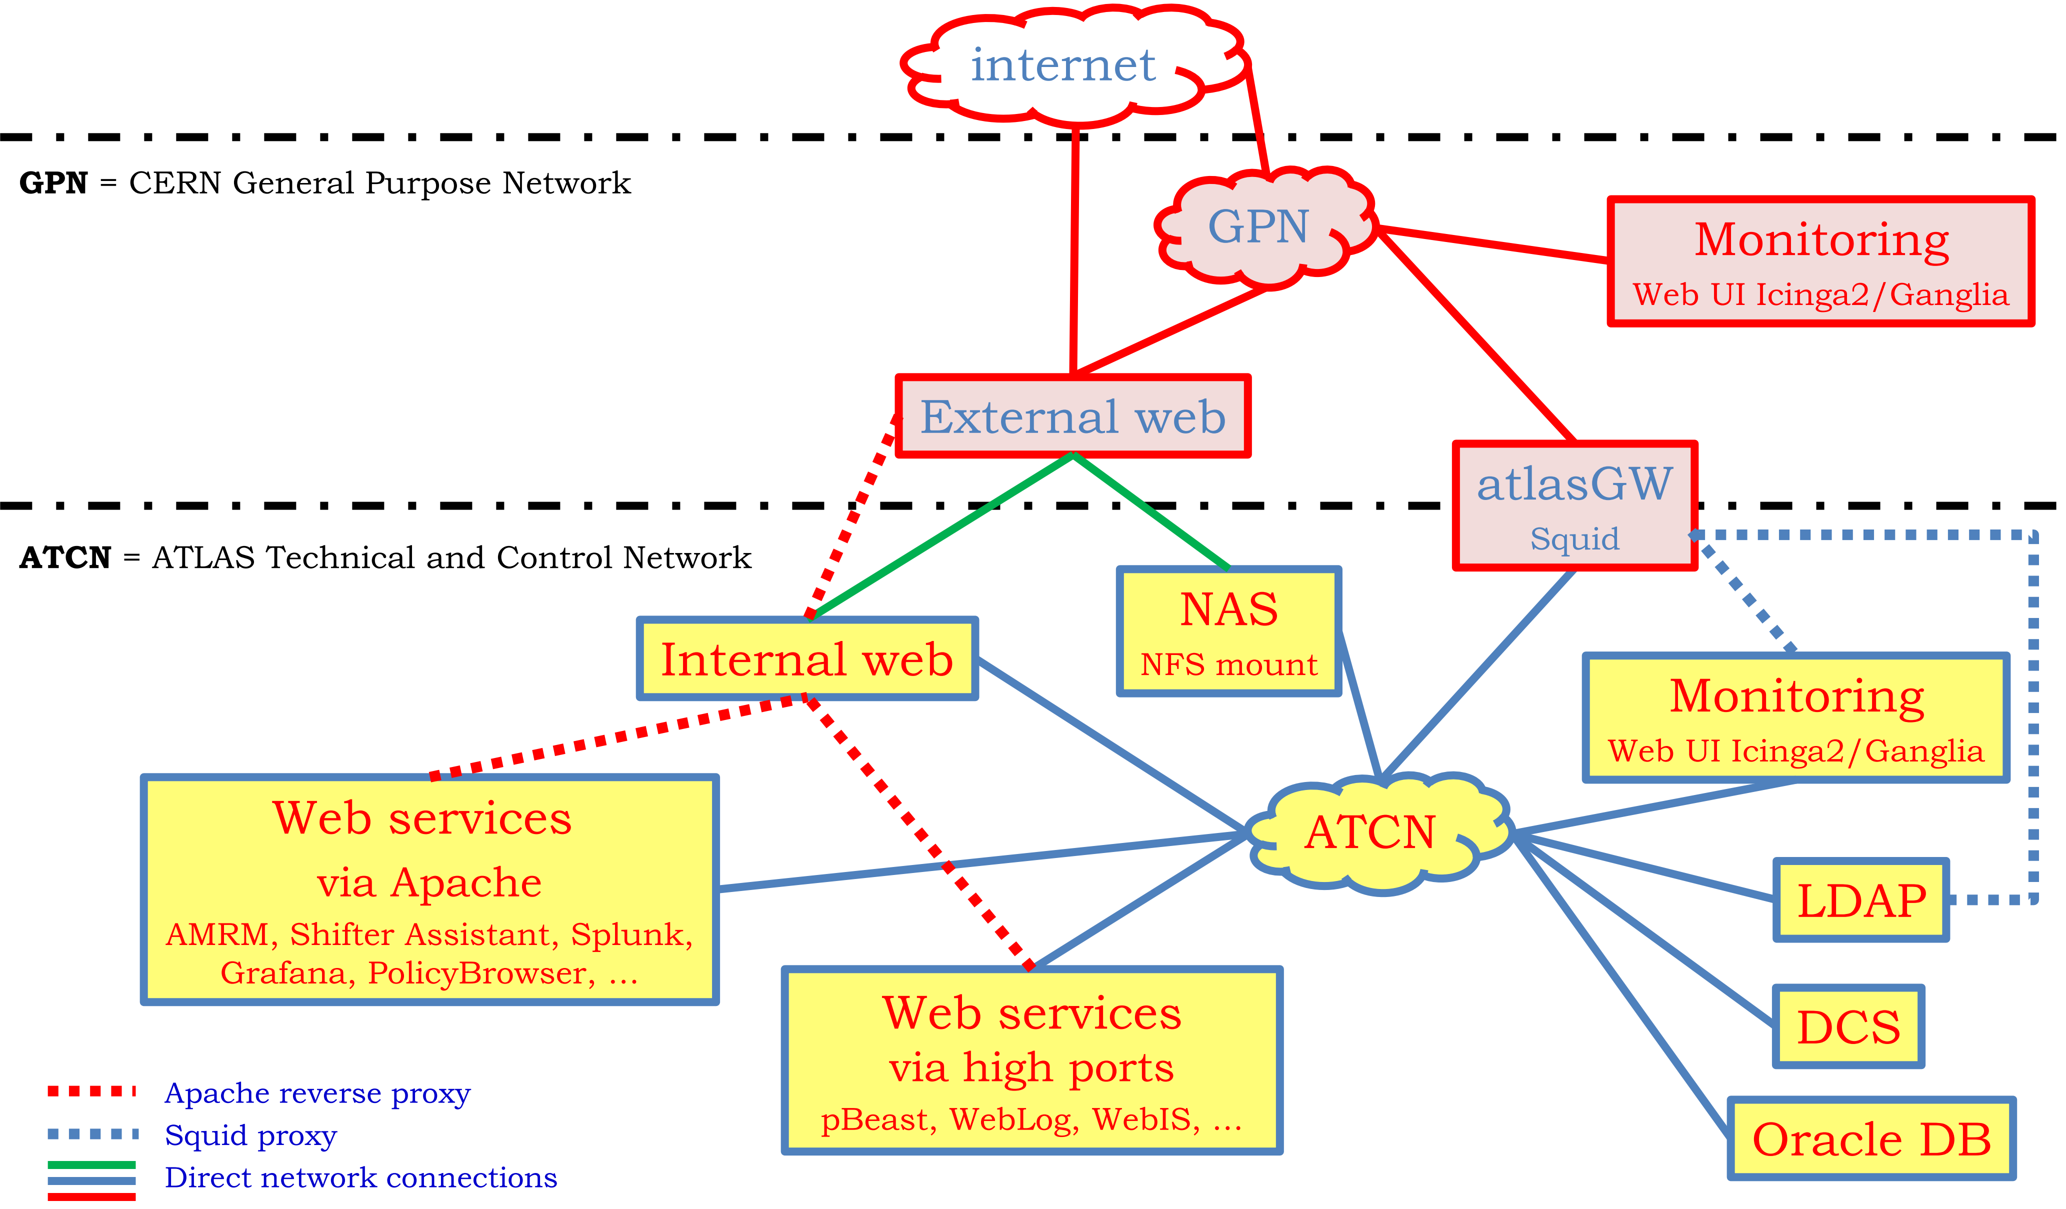

New schema and details

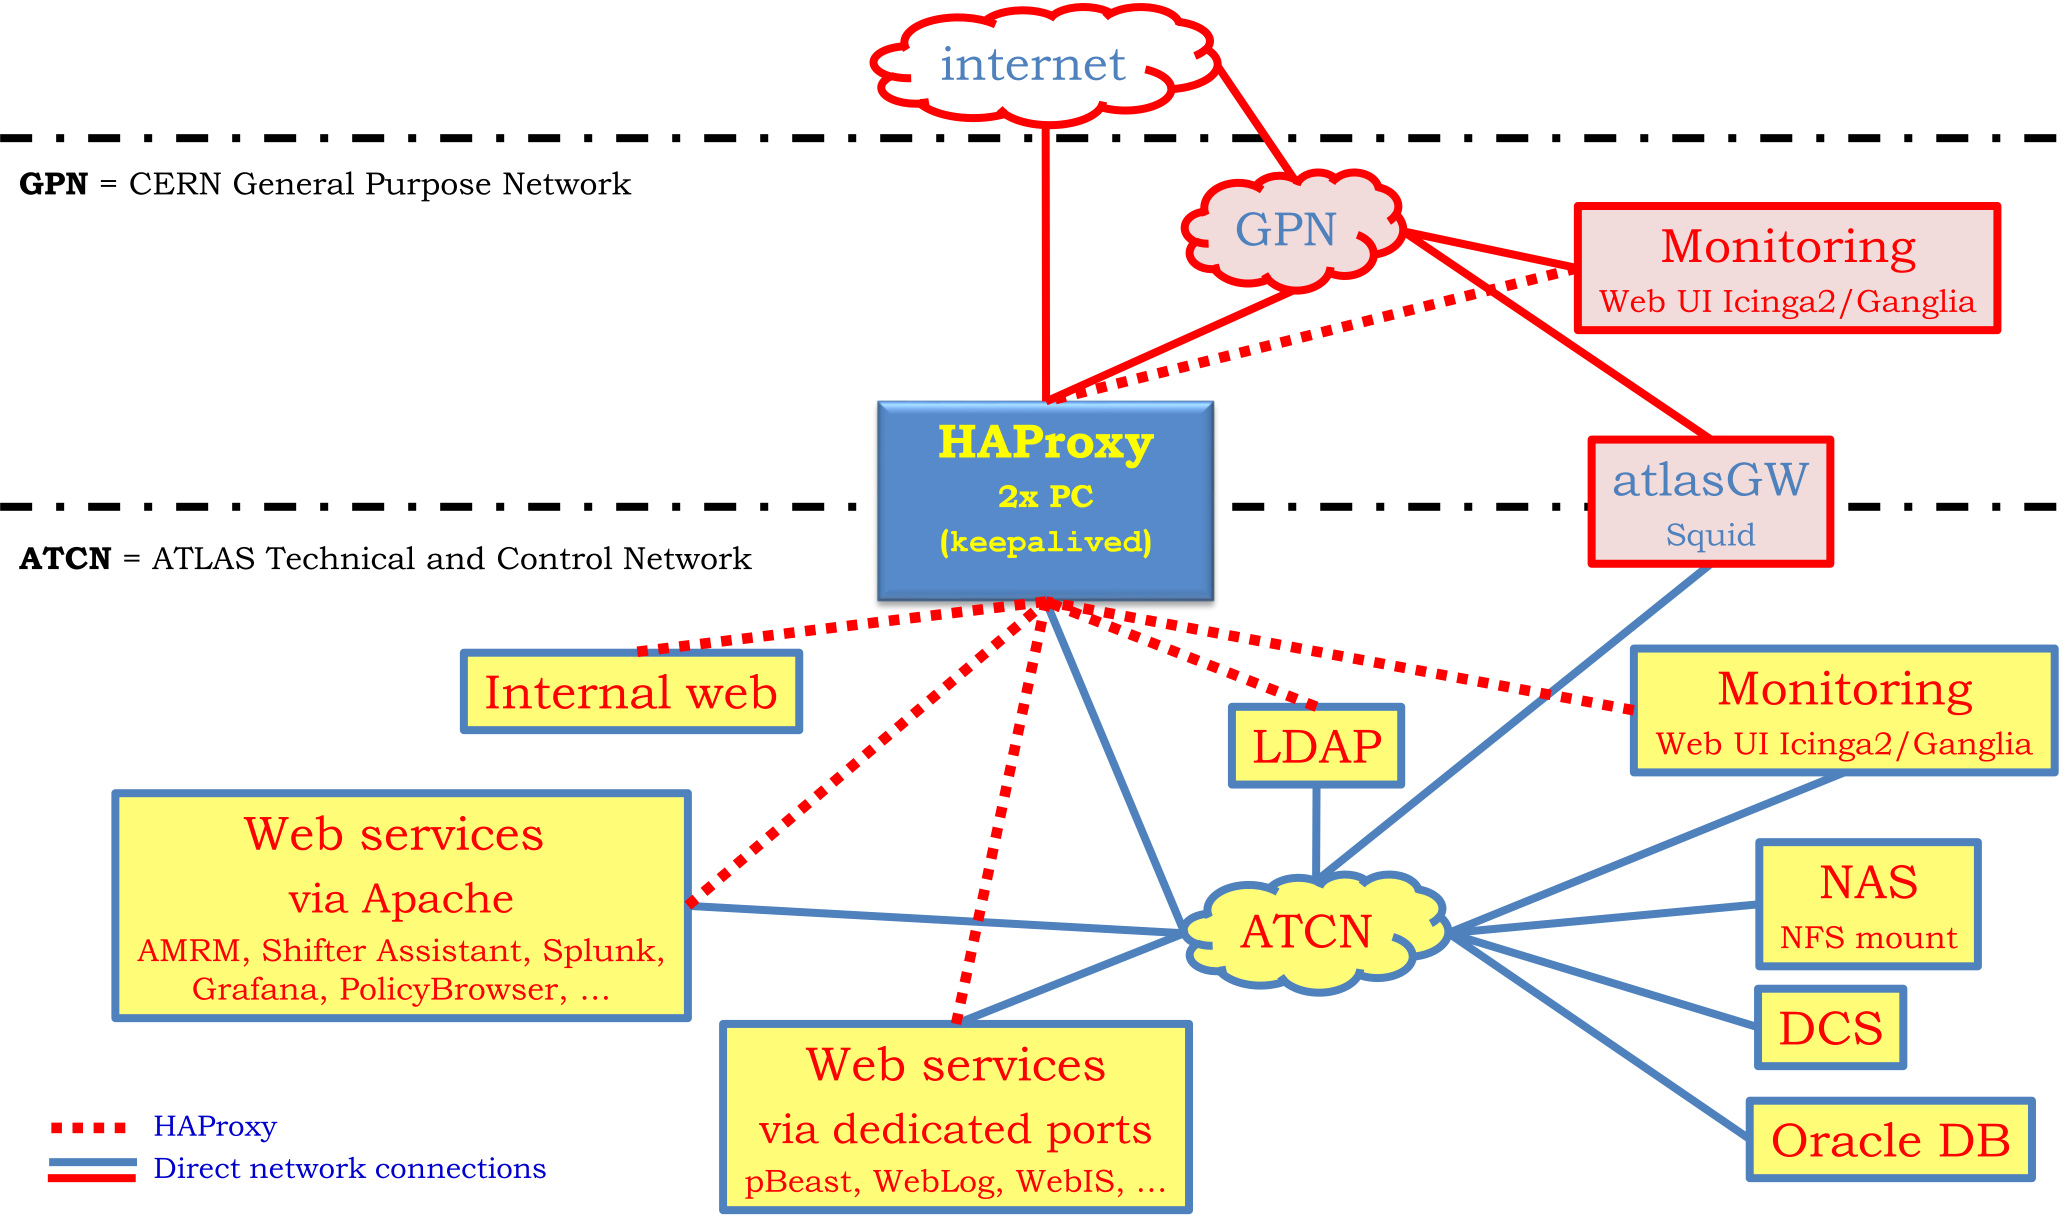

GPN = CERN General Purpose Network

ATCN = ATLAS Technical and Control Network

▪▪▪▪ HAProxy

═ Direct network connections