CHEP 2007

Addressing the Pilot Security Problem With gLExec

by

I. Sfiligoi (Fermilab)

in collaboration with
D. Yocum (Fermilab)
O. Koeroo (NIKHEF)
G. Venekamp (NIKHEF)

Security in the Grid

Gatekeepers decide who to let in

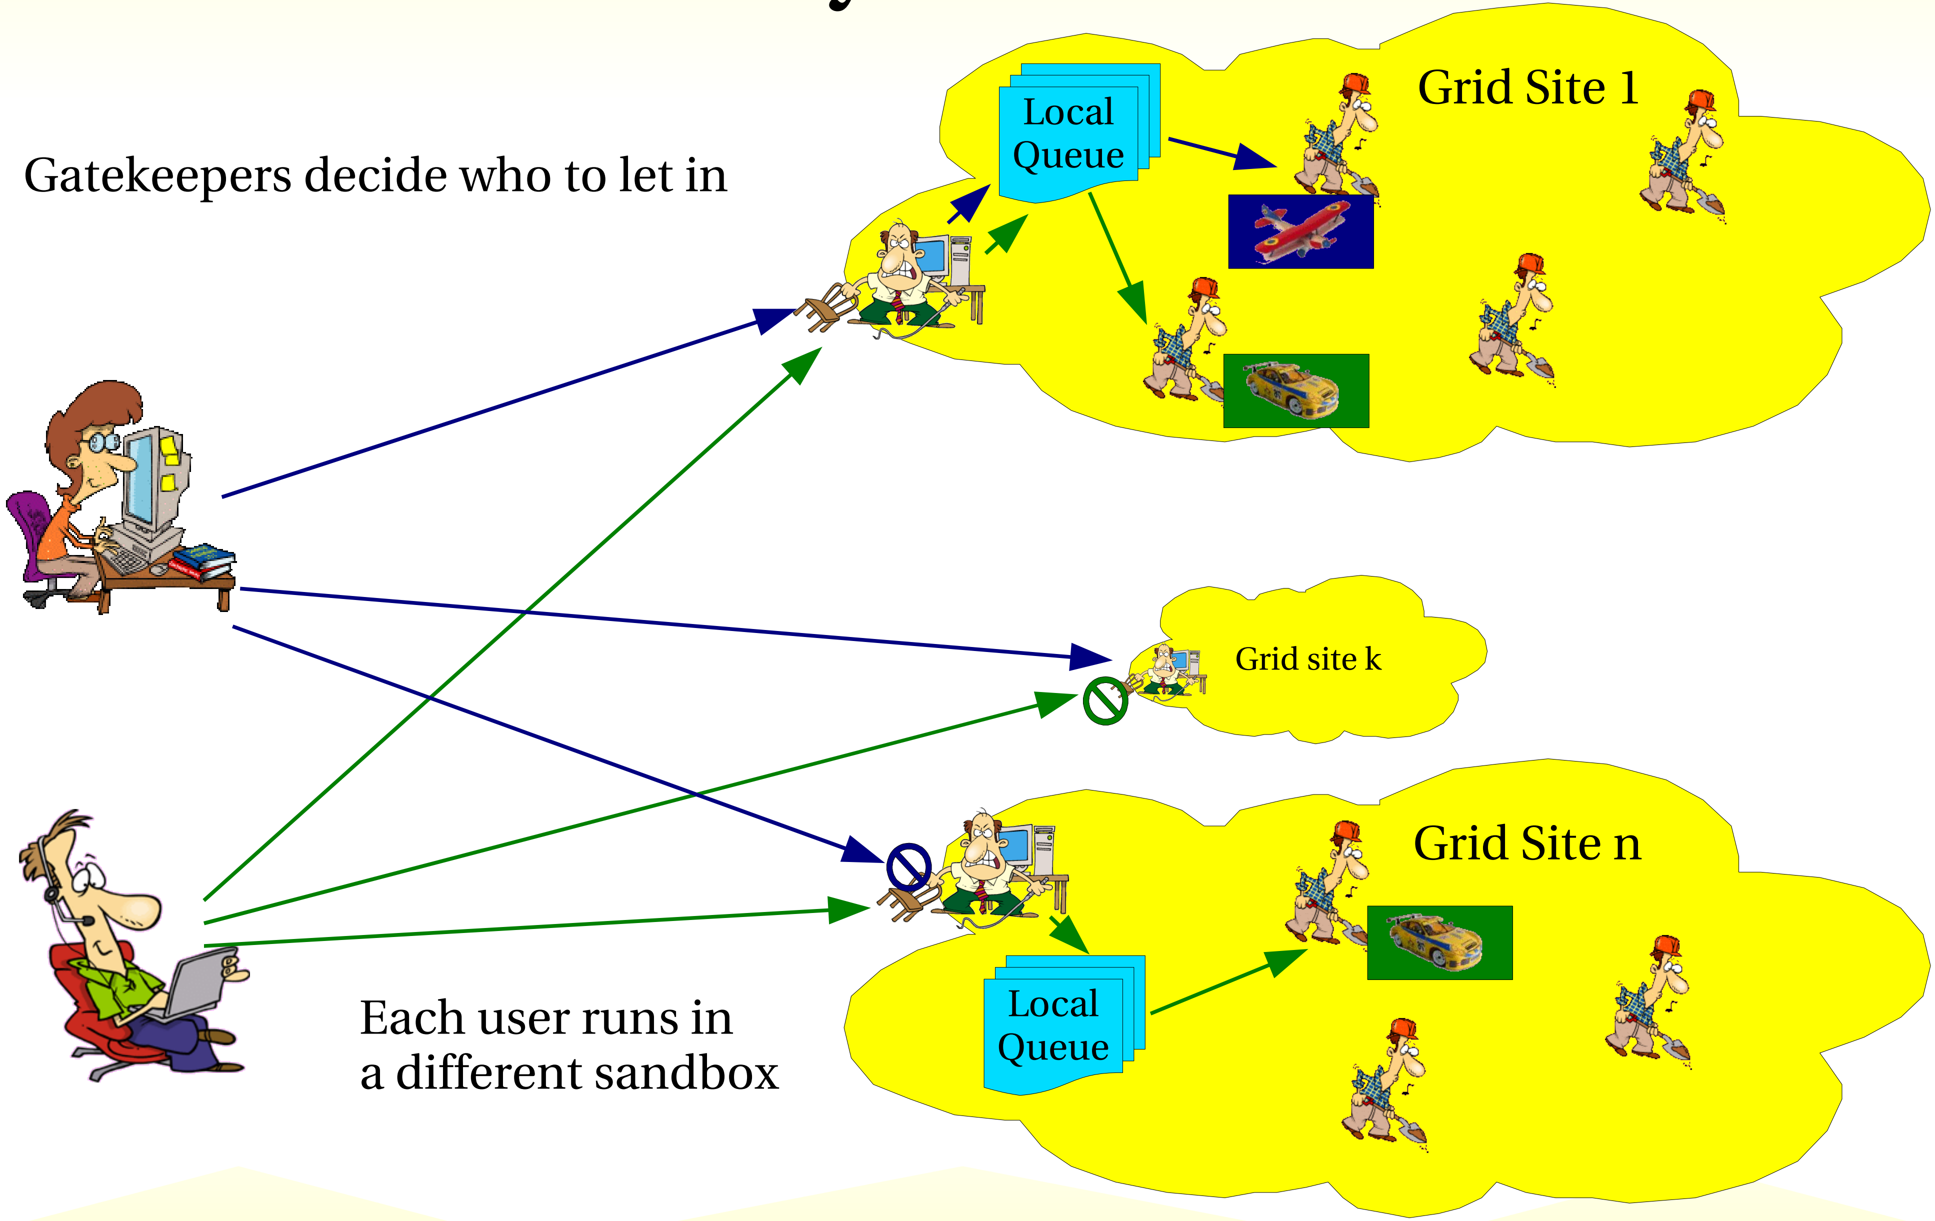

Each user runs in
a different sandbox

Pilots and the Grid (1)

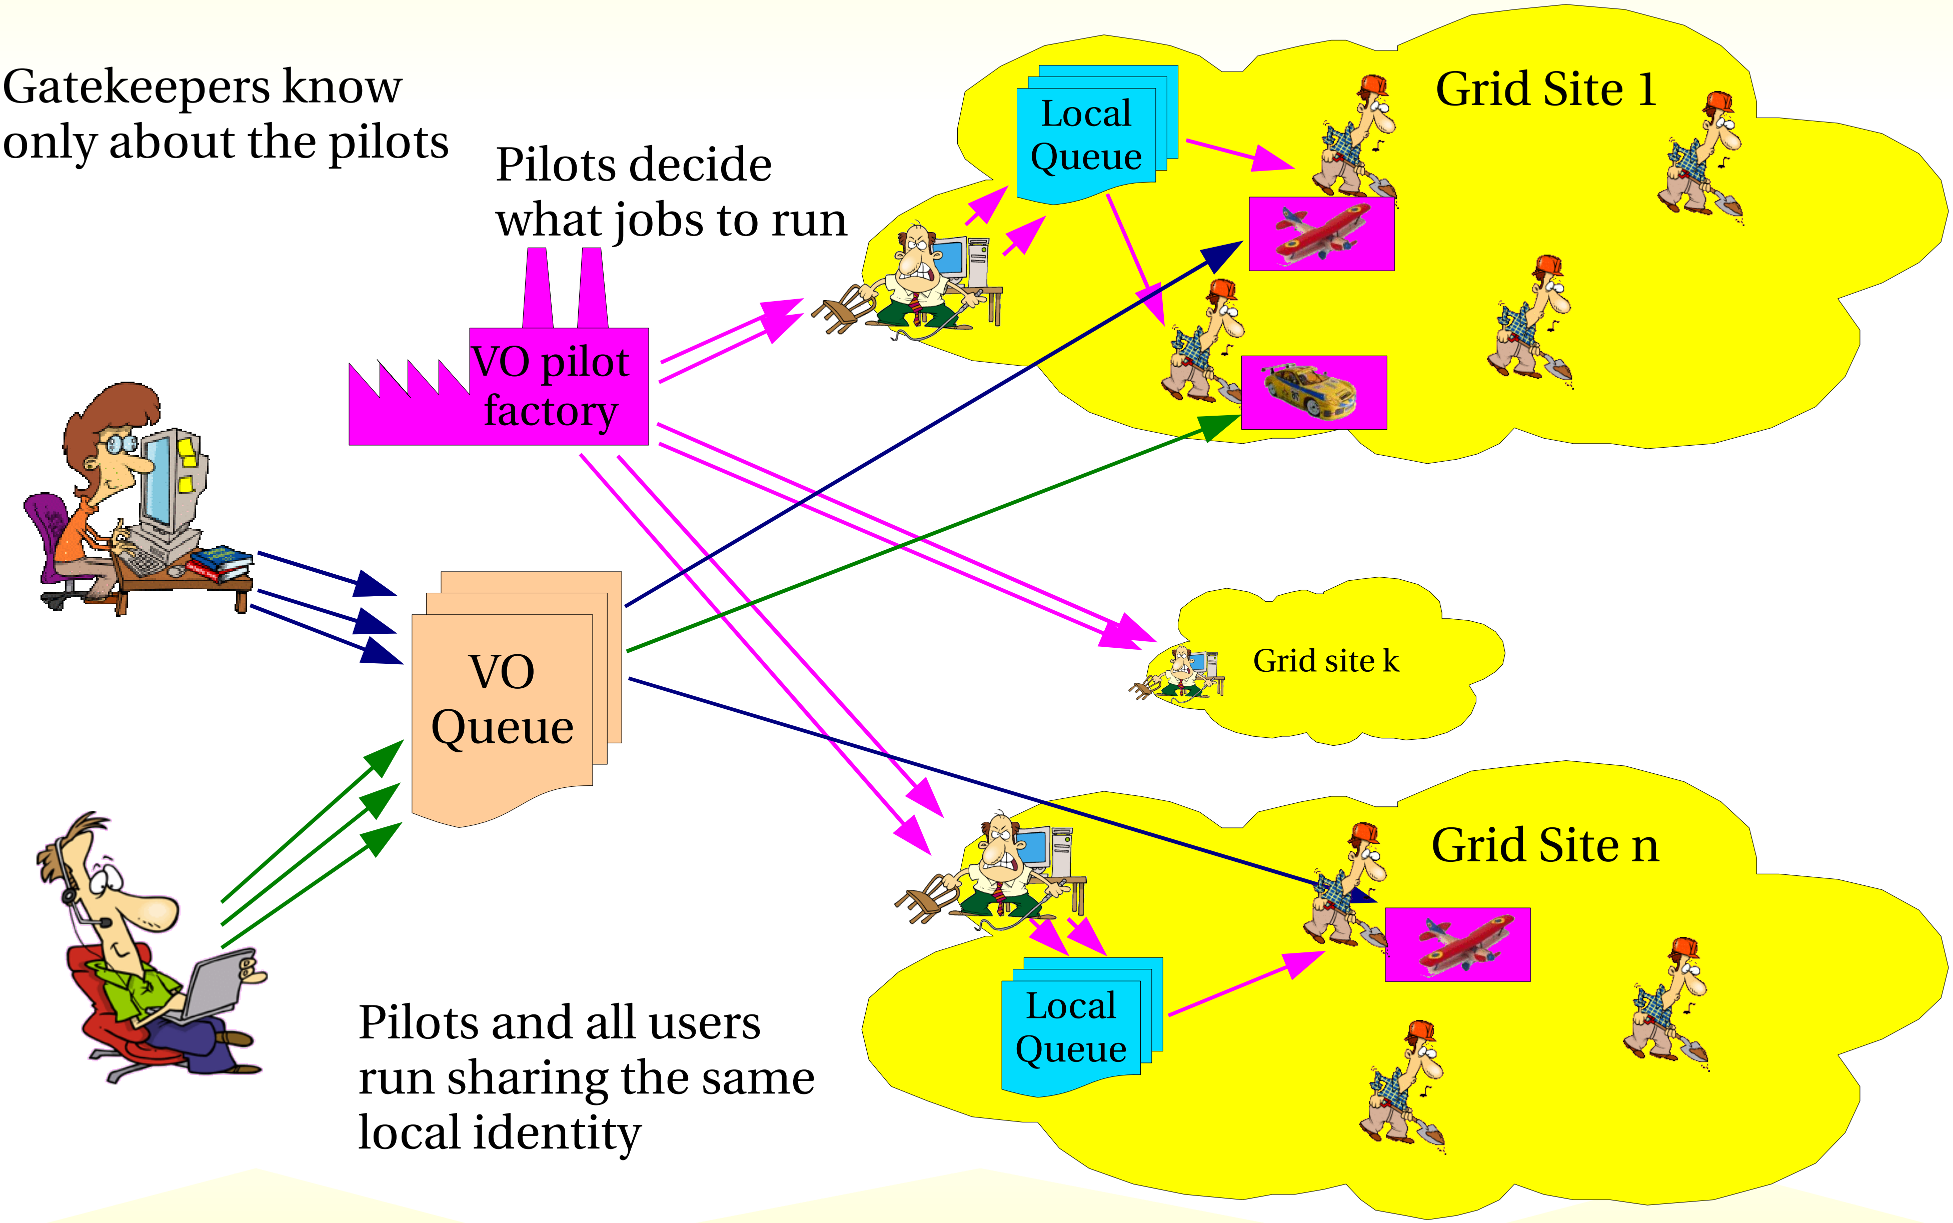

Gatekeepers know only about the pilots

Pilots decide what jobs to run

Pilots and all users run sharing the same local identity

Pilots and the Grid (2)

Gatekeepers know only about the pilots

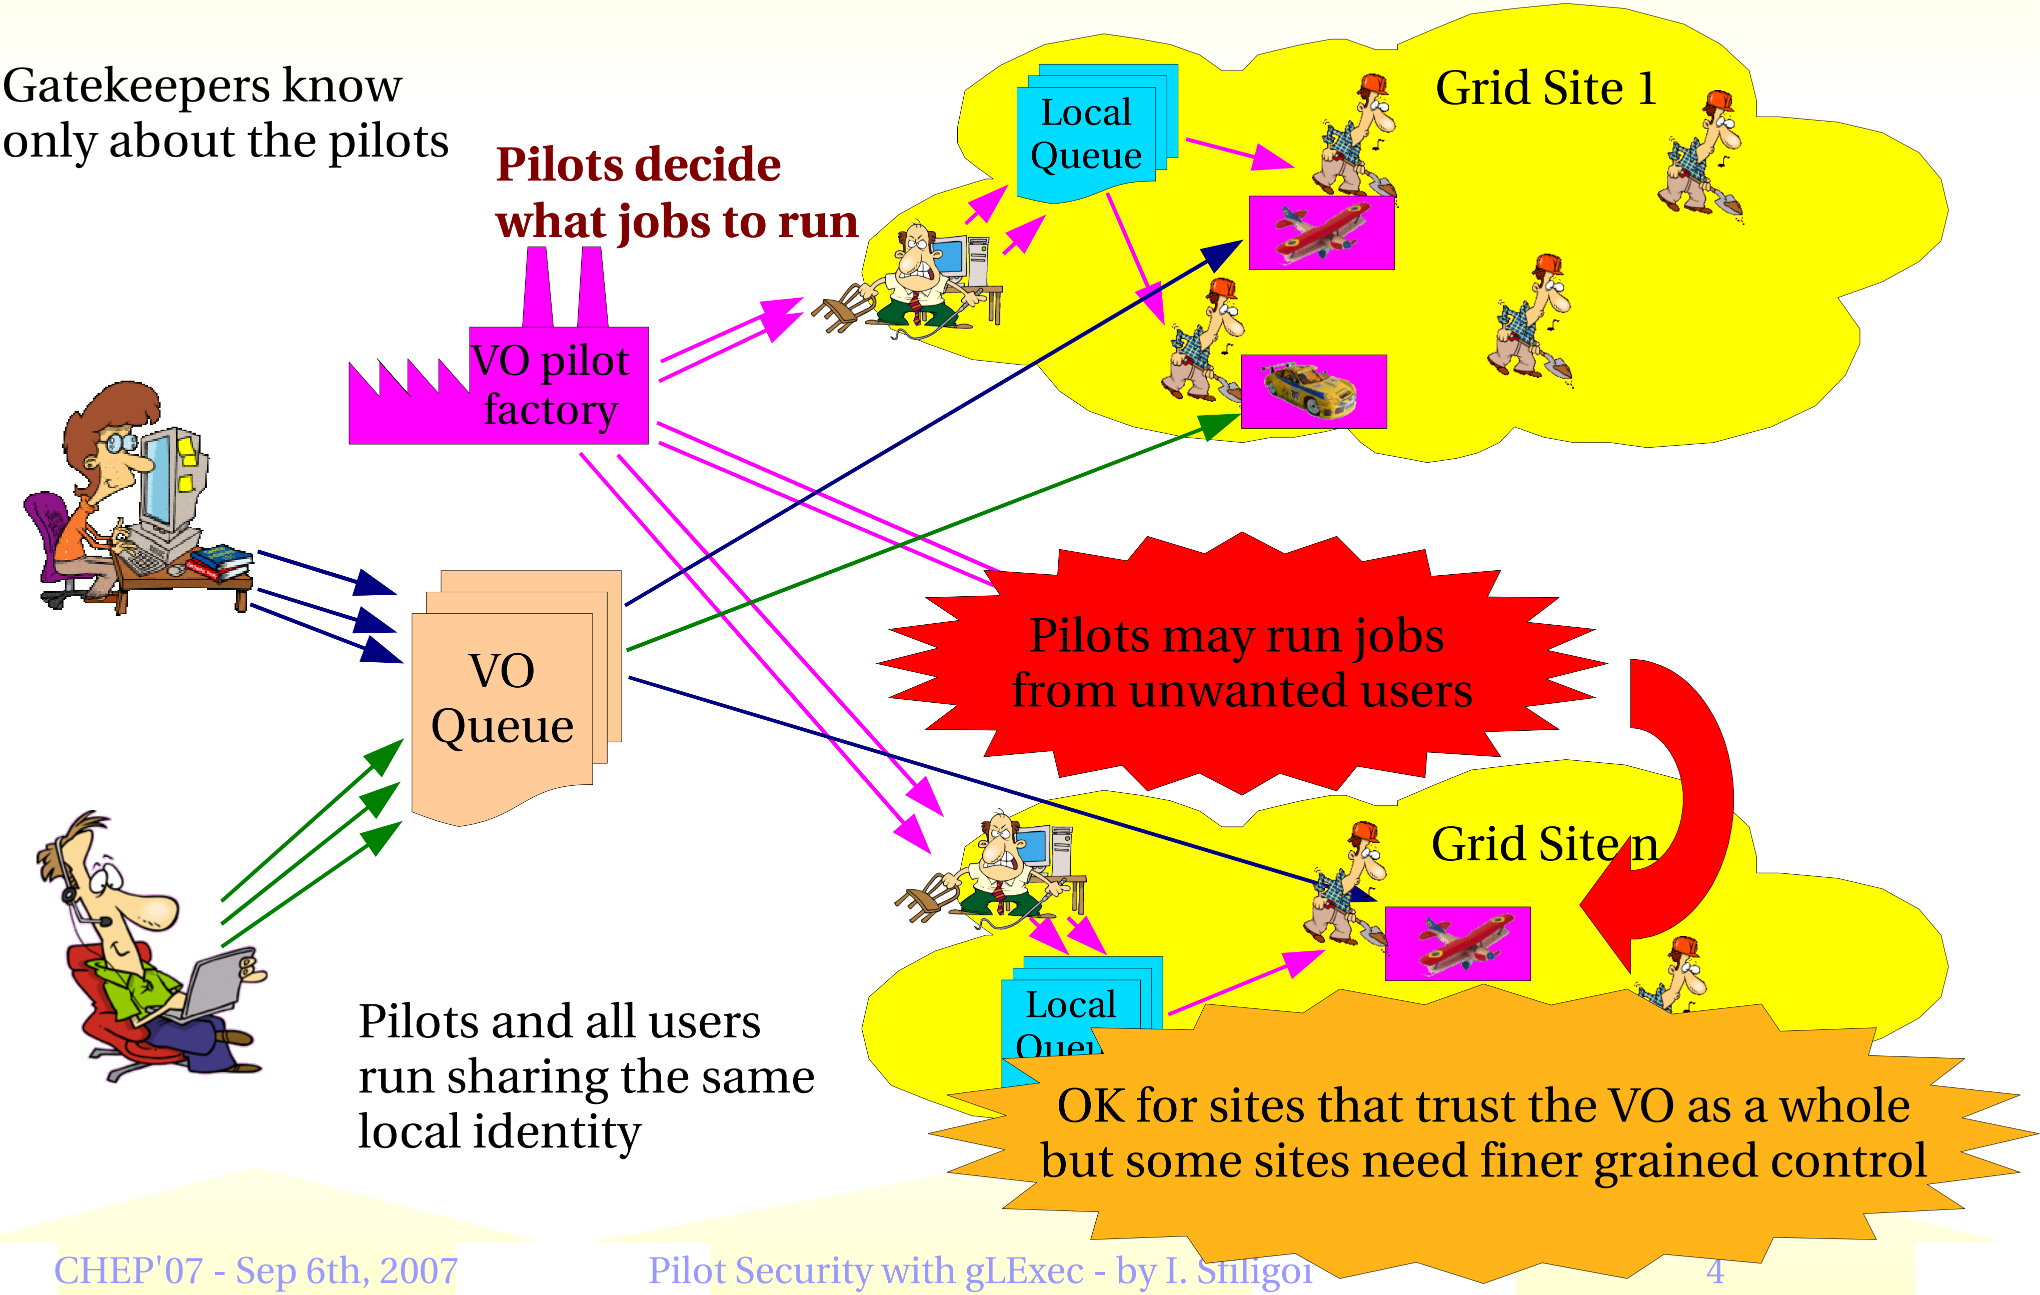

Pilots may run jobs from unwanted users

OK for sites that trust the VO as a whole but some sites need finer grained control

Pilots and all users run sharing the same local identity

Pilots and the Grid (3)

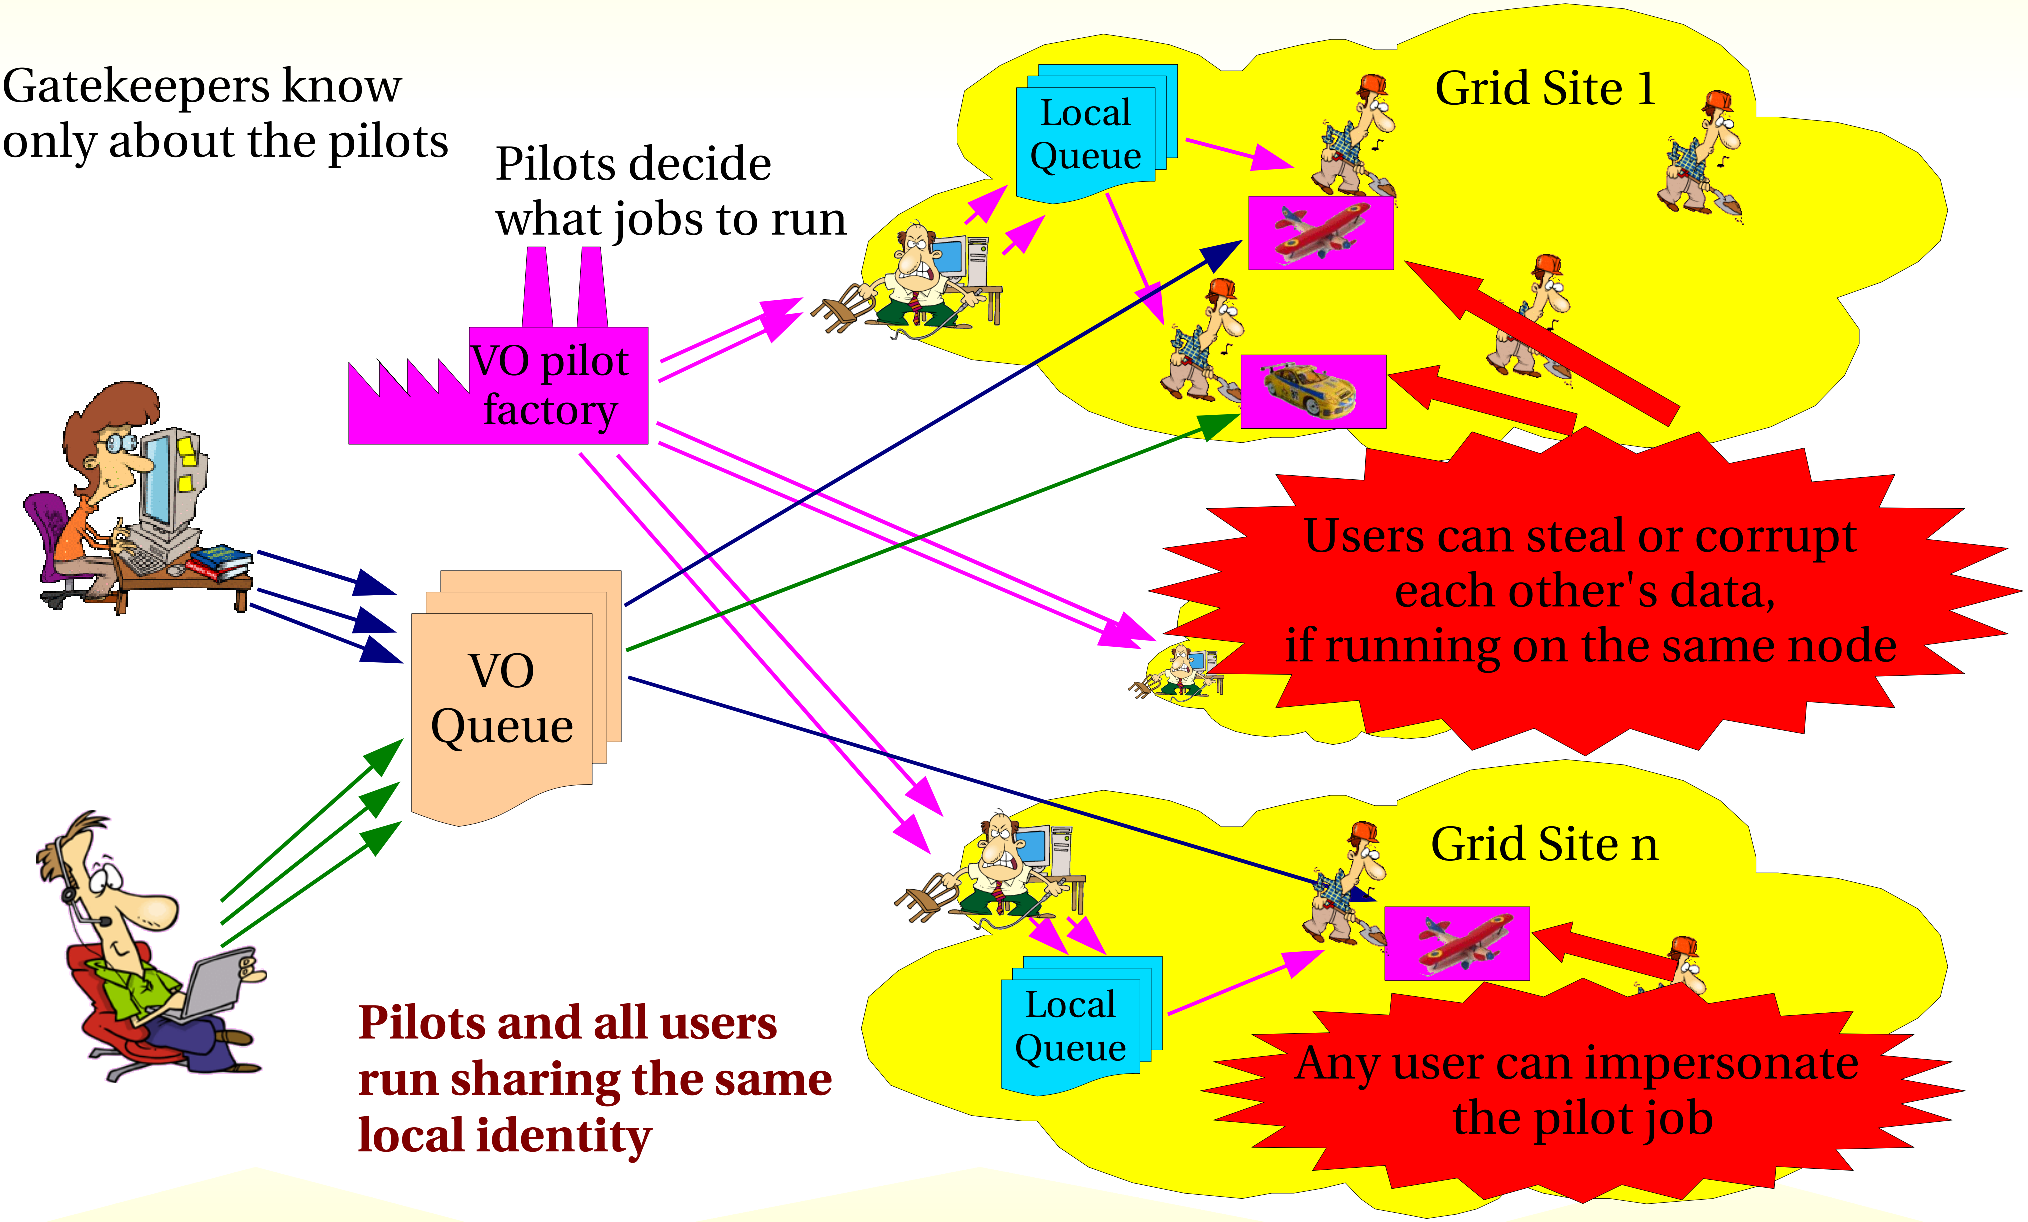
Gatekeepers know only about the pilots
Pilots decide what jobs to run
VO pilot factory
VO Queue
Local Queue
Grid Site 1
Users can steal or corrupt each other's data, if running on the same node
Grid Site n
Local Queue
Any user can impersonate the pilot job
Pilots and all users run sharing the same local identity

Pilots and the Grid (4)

- Pilots can be a problem for both the Grid sites and pilot owners
- Grid sites
  - Who is running
  - Who to hold accountable
  - Reject jobs from unwanted users
- Pilot owners
  - Protect its infrastructure from malicious users
  - Protect users from each other
  - Avoid being held accountable for user's actions

The (simplified) life of a pilot job

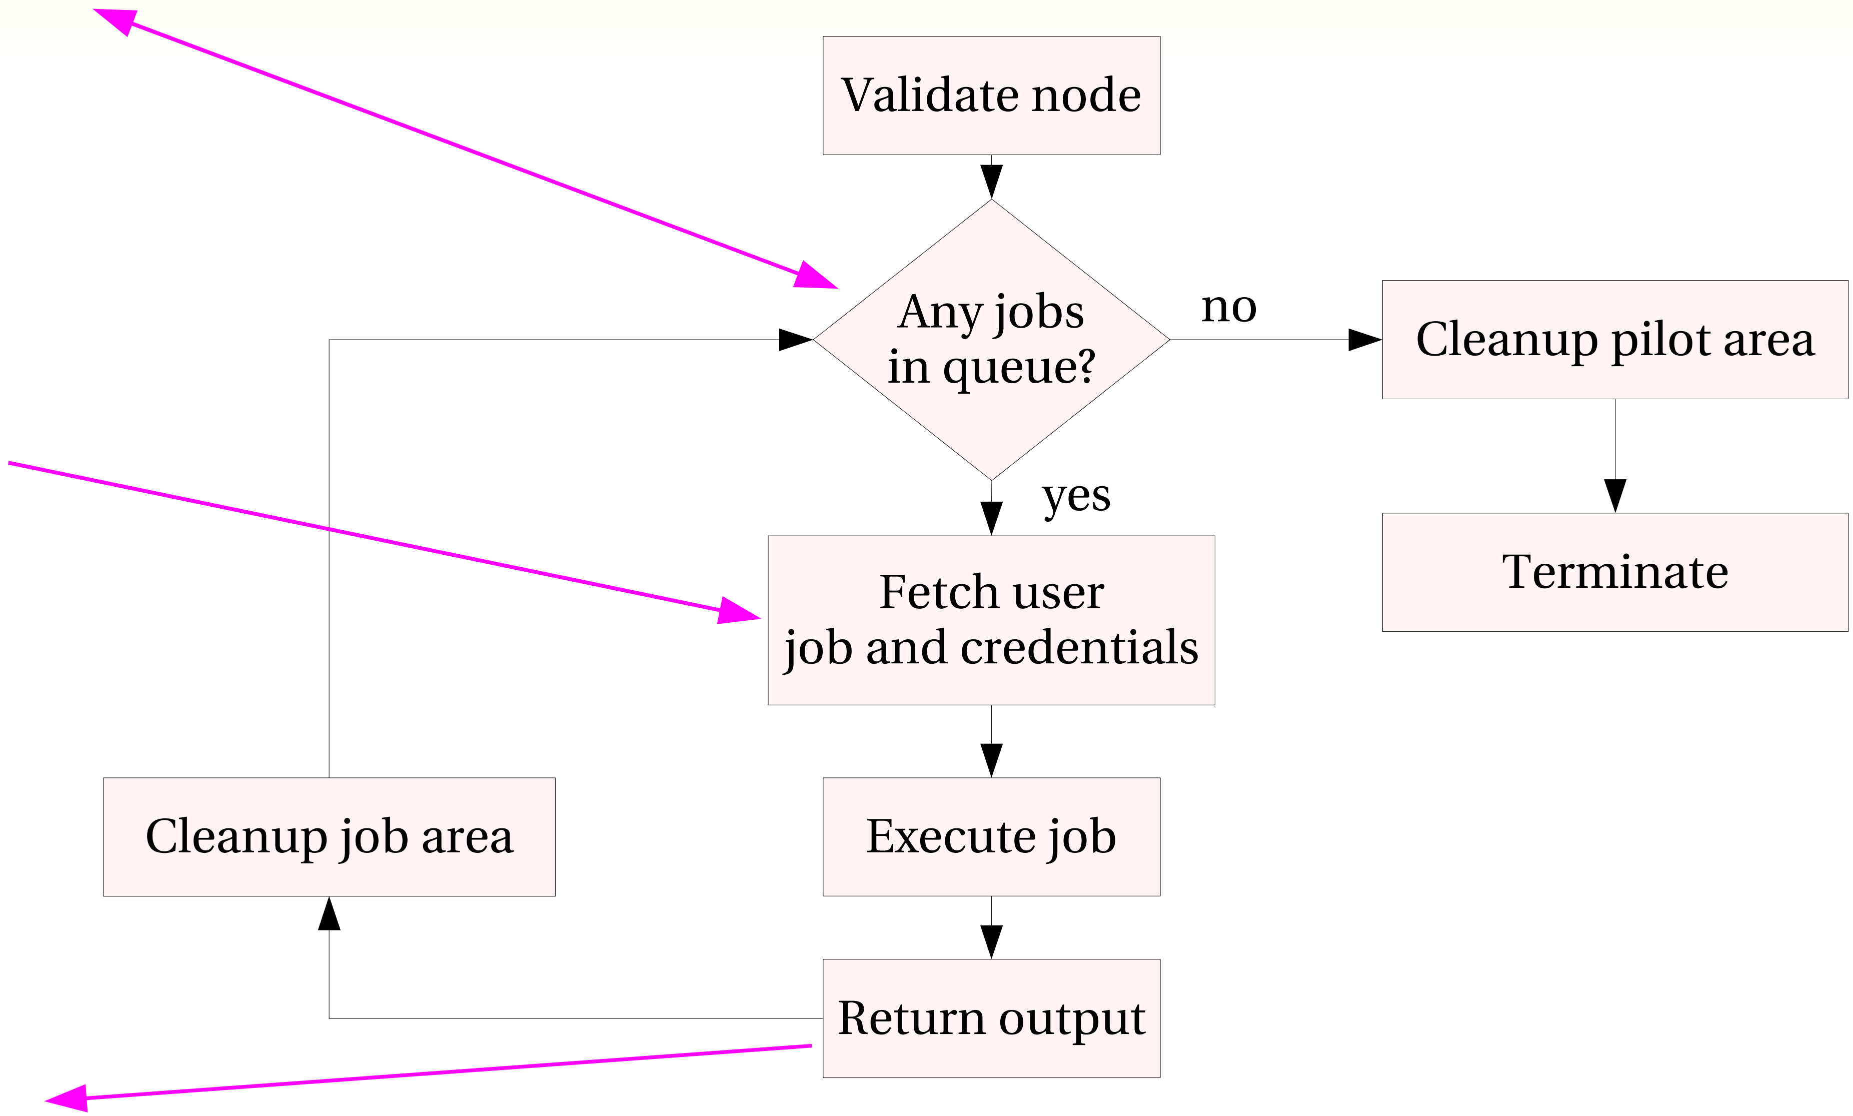

Introducing gLExec

- A x509 aware (apache) suExec equivalent
- Given a proxy and a binary
  - authenticates and authorizes
  - changes UID/GID
  - executes the binary
- Necessarily owned by root and suid-ed
  - **Not a daemon**

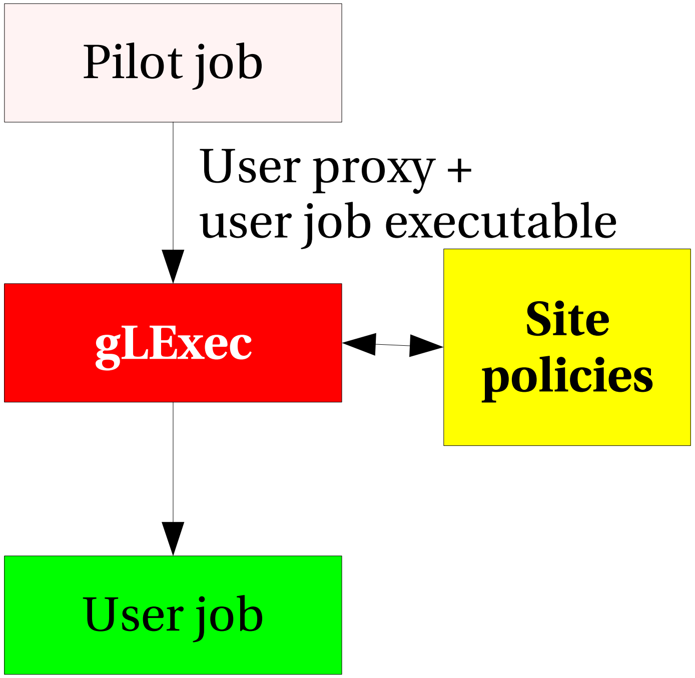

Pilots with gLExec (1)

Sites decide who to let run

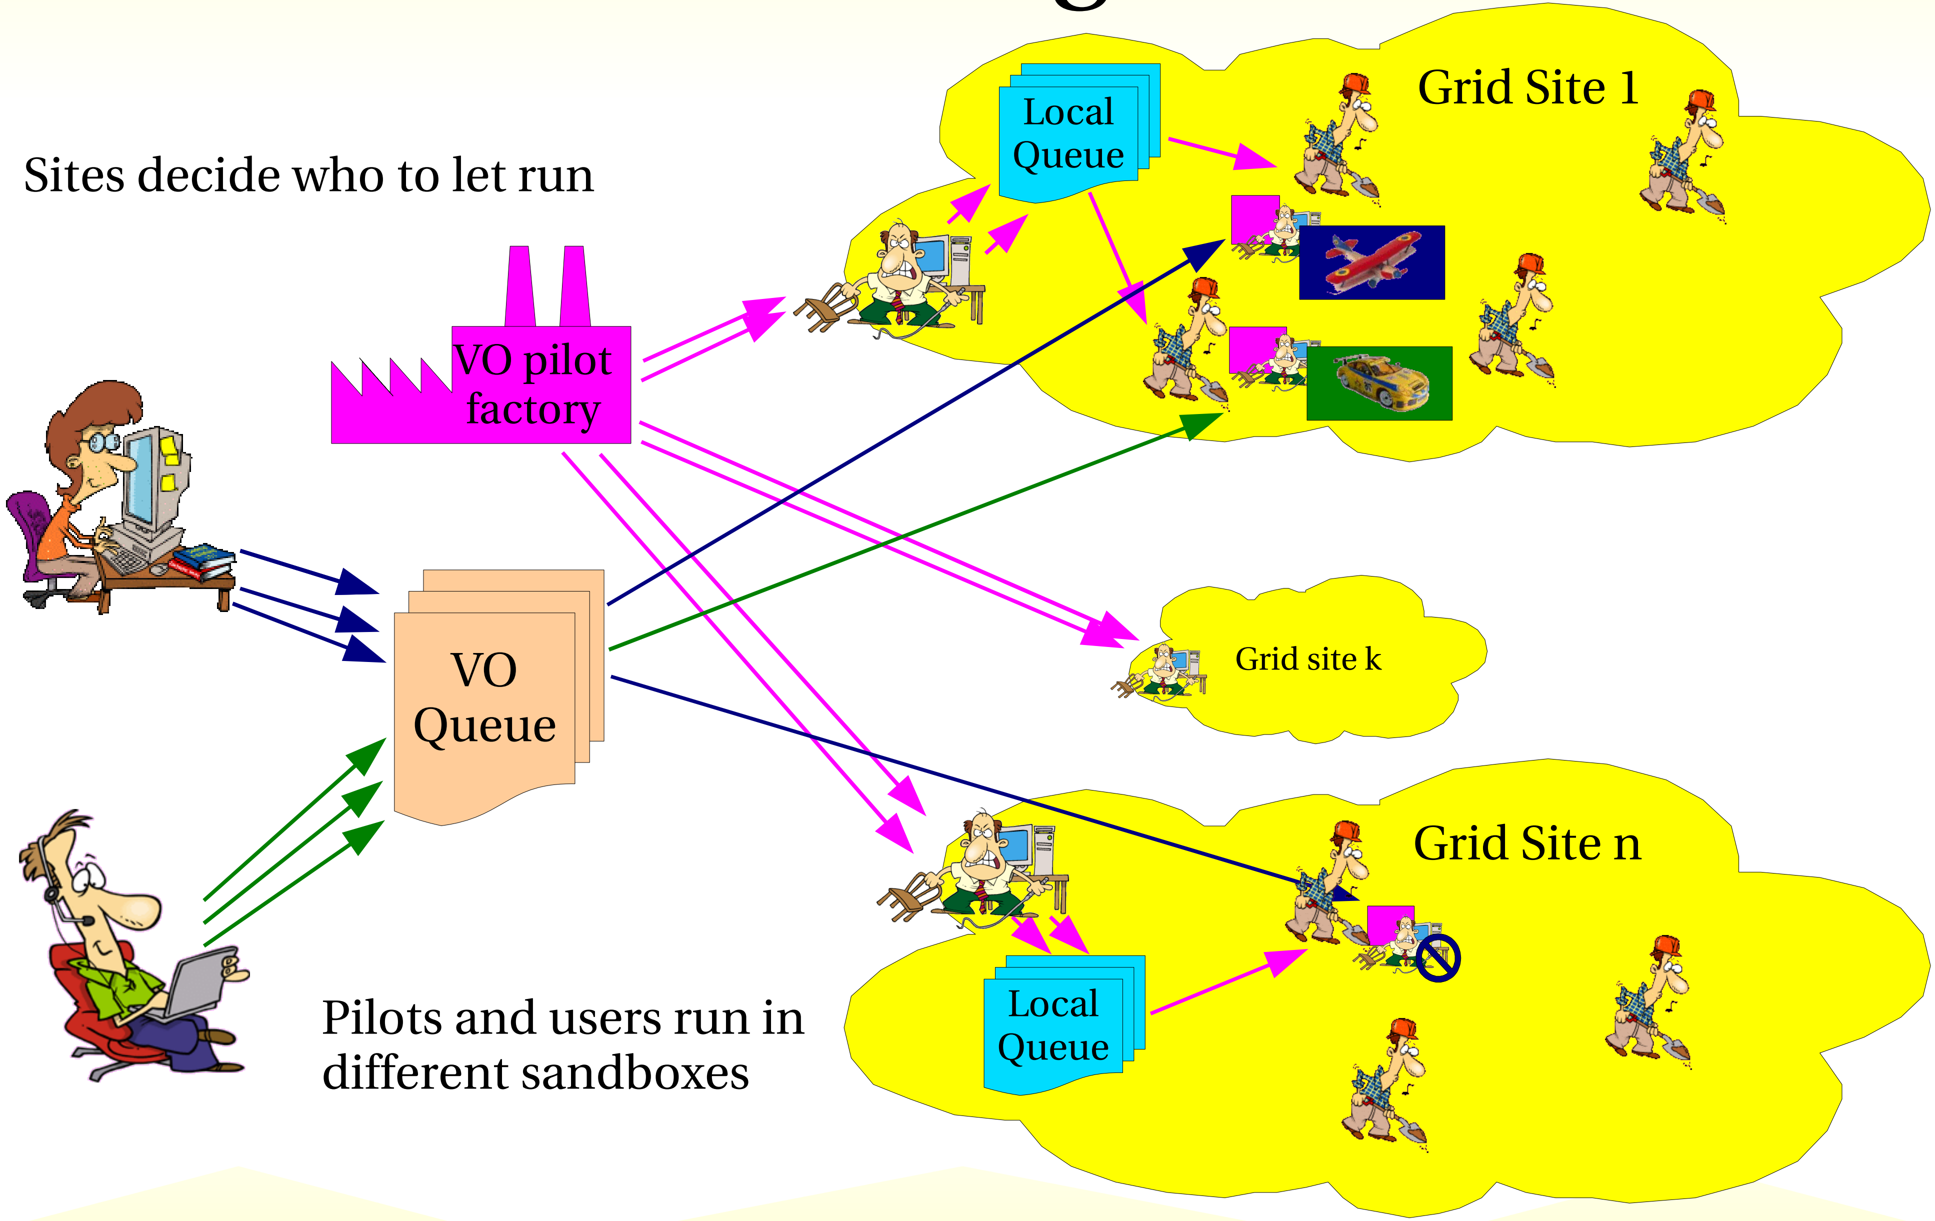

Pilots and users run in different sandboxes

Pilots with gLExec (2)

Sites decide who to let run

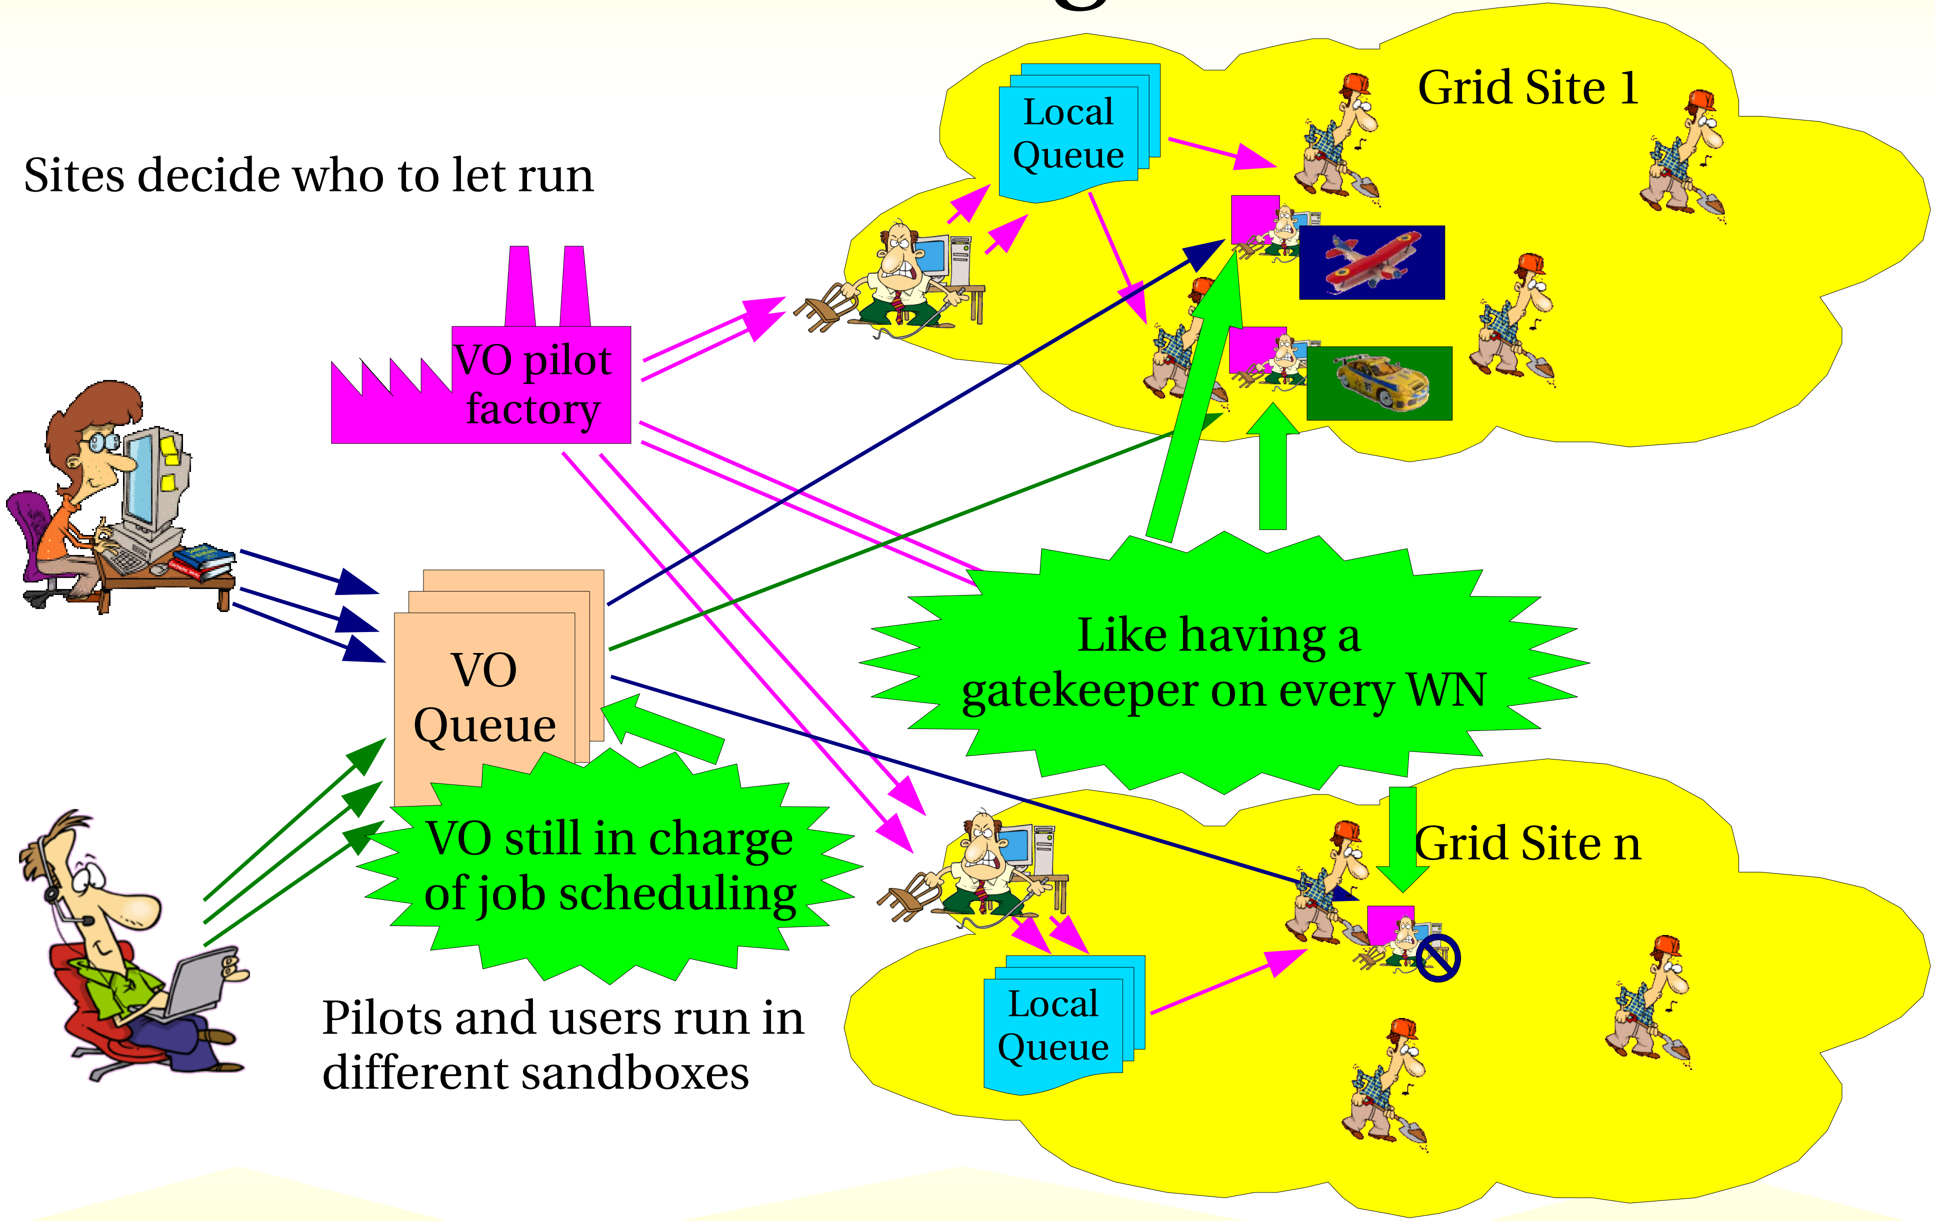

Like having a gatekeeper on every WN

VO still in charge of job scheduling

Pilots and users run in different sandboxes

The need for centralized authorization

- With gLExec used by Pilots, sites can have thousands of decision points
  - A centralized authorization infrastructure helps keep consistency
- Open Science Grid (OSG) has GUMS/PRIMA
  - A PRIMA gLExec module exists

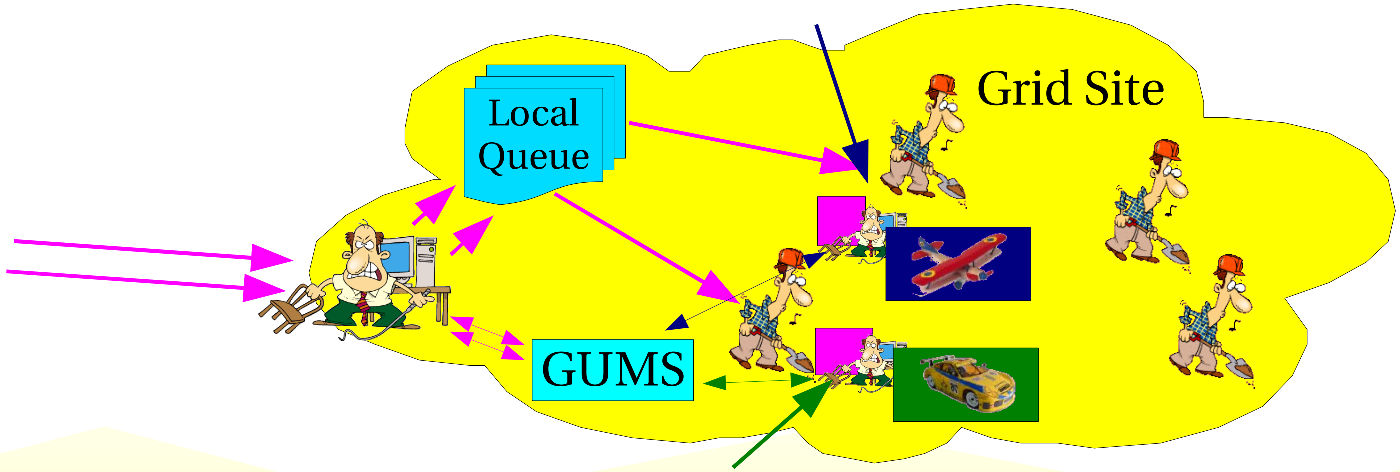

gLExec deployment @ Fermilab

- The PRIMA plugin for gLExec was jointly developed by the NIKHEF and Fermilab groups
- Fermilab was the first OSG site to install gLExec on the worker nodes
  - First release in Oct'06
  - Current release deployed in May'07

gLExec/GUMS @ Fermilab

- gLExec deployed and used on more than 500 nodes
  - Primary user is CDF
- GUMS shown to be very scalable
  - Routinely handling more than 500k requests per day
  - gLExec contributing for at most 20k

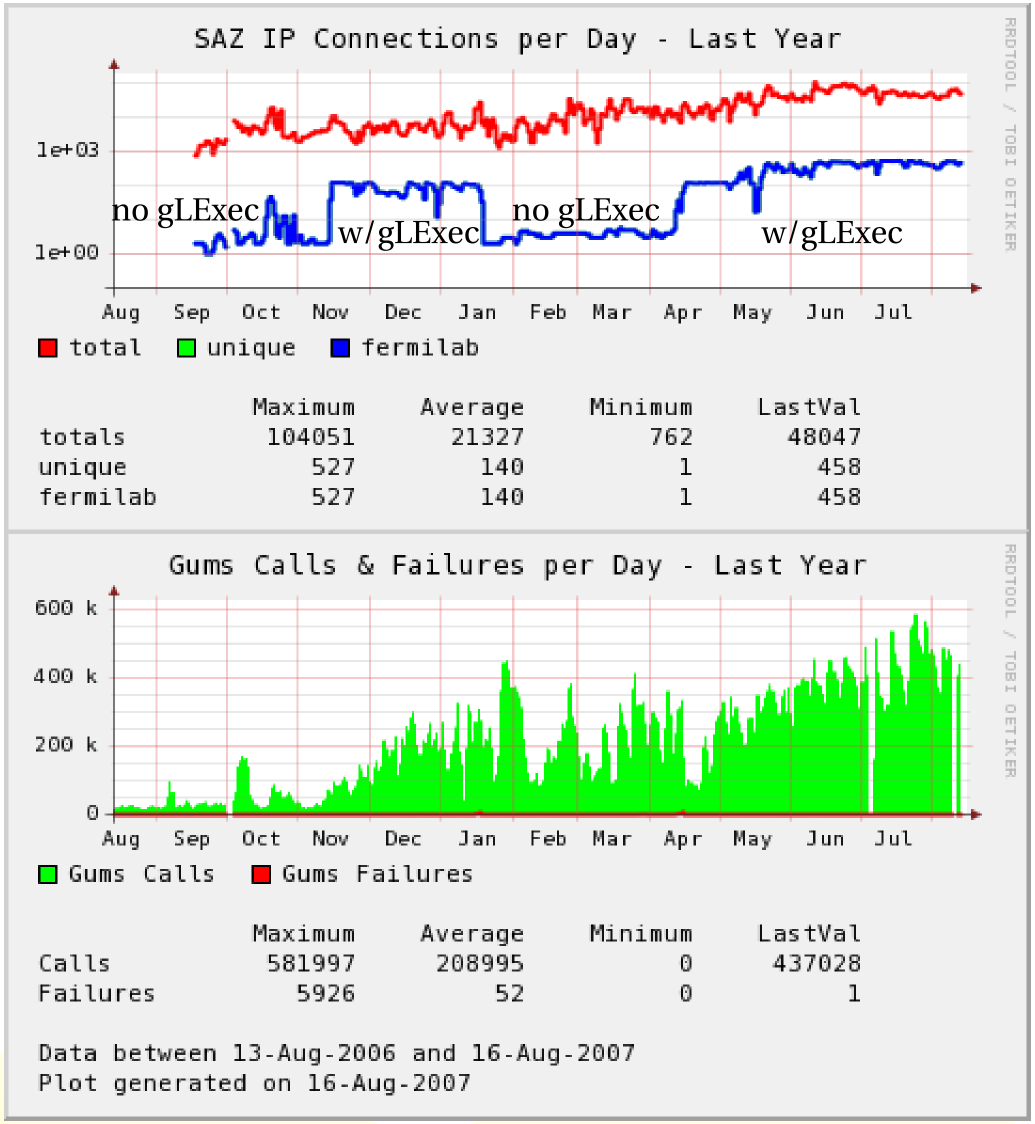

	Maximum	Average	Minimum	LastVal
totals	104051	21327	762	48047
unique	527	140	1	458
fermilab	527	140	1	458

	Maximum	Average	Minimum	LastVal
Calls	581997	208995	0	437028
Failures	5926	52	0	1

Data between 13-Aug-2006 and 16-Aug-2007
Plot generated on 16-Aug-2007

gLExec in OSG

- gLExec with PRIMA plugin has been incorporated in VDT 1.8
- OSG sites will be encouraged to deploy it with the upcoming OSG 0.8.0 release

gLExec monitoring

- gLExec keeps a log of all authorization requests
  - The DN and FQAN of the proxy
  - PID of the calling process
  - (Sub-)processes spawned
  - Start times and end times
- Gratia probe uses this data
  - Gives accounting the real user identity

Summary

- Pilots are becoming a reality
  - But are introducing new problems
- Yesterday's Grid security mechanisms are not enough anymore
  - Site admins want fine grained control
  - Pilot owners want OS level protection
- gLExec can help solve the technical problems
- Fermilab experience positive,
  waiting for wide OSG deployment

CHEP 2007
Pilot Security with gLExec

Backup slides

gLExec jobs a day @ FNAL

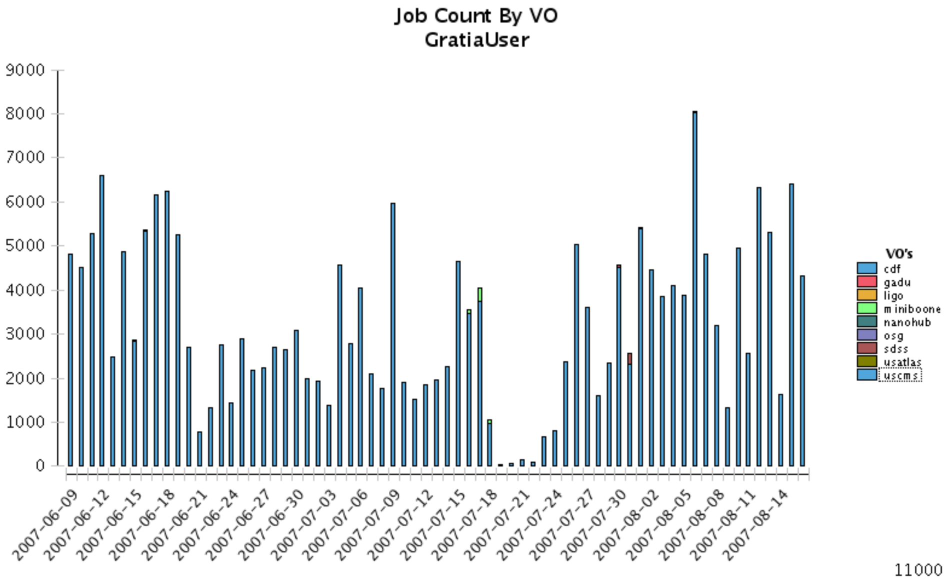

fcdfosg1

fcdfosg2

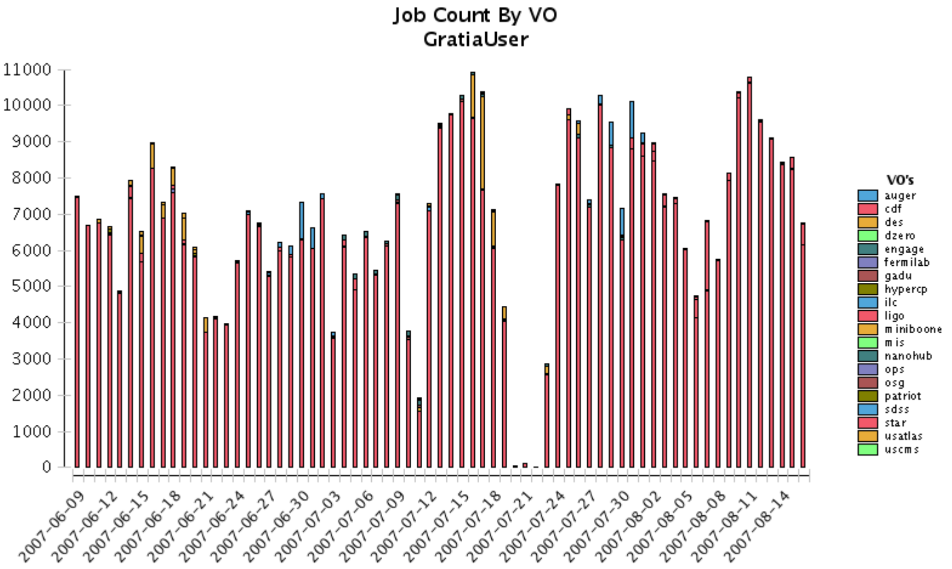
Job Count By VO
GratiaUser
VO's
auger
cdf
des
dzero
engage
fermilab
gadu
hypercp
ilc
ligo
miniboone
mis
nanohub
ops
osg
patriot
sdss
star
usatlas
uscms