



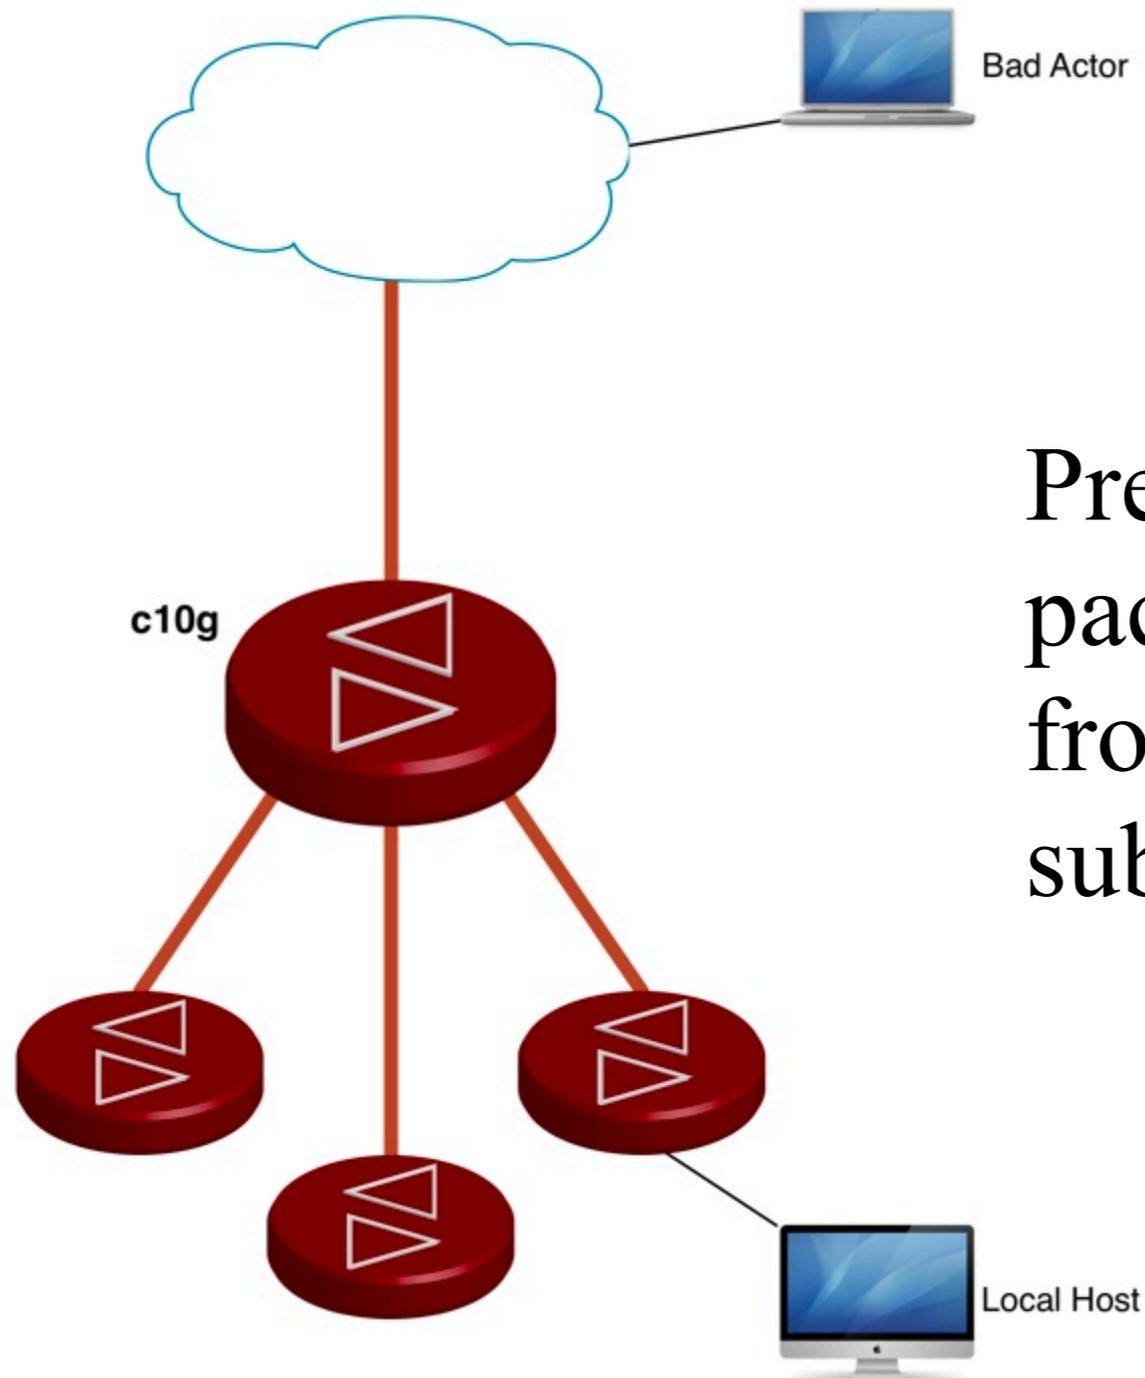
Black Hole Filtering

NERSC's New Blocking
Infrastructure

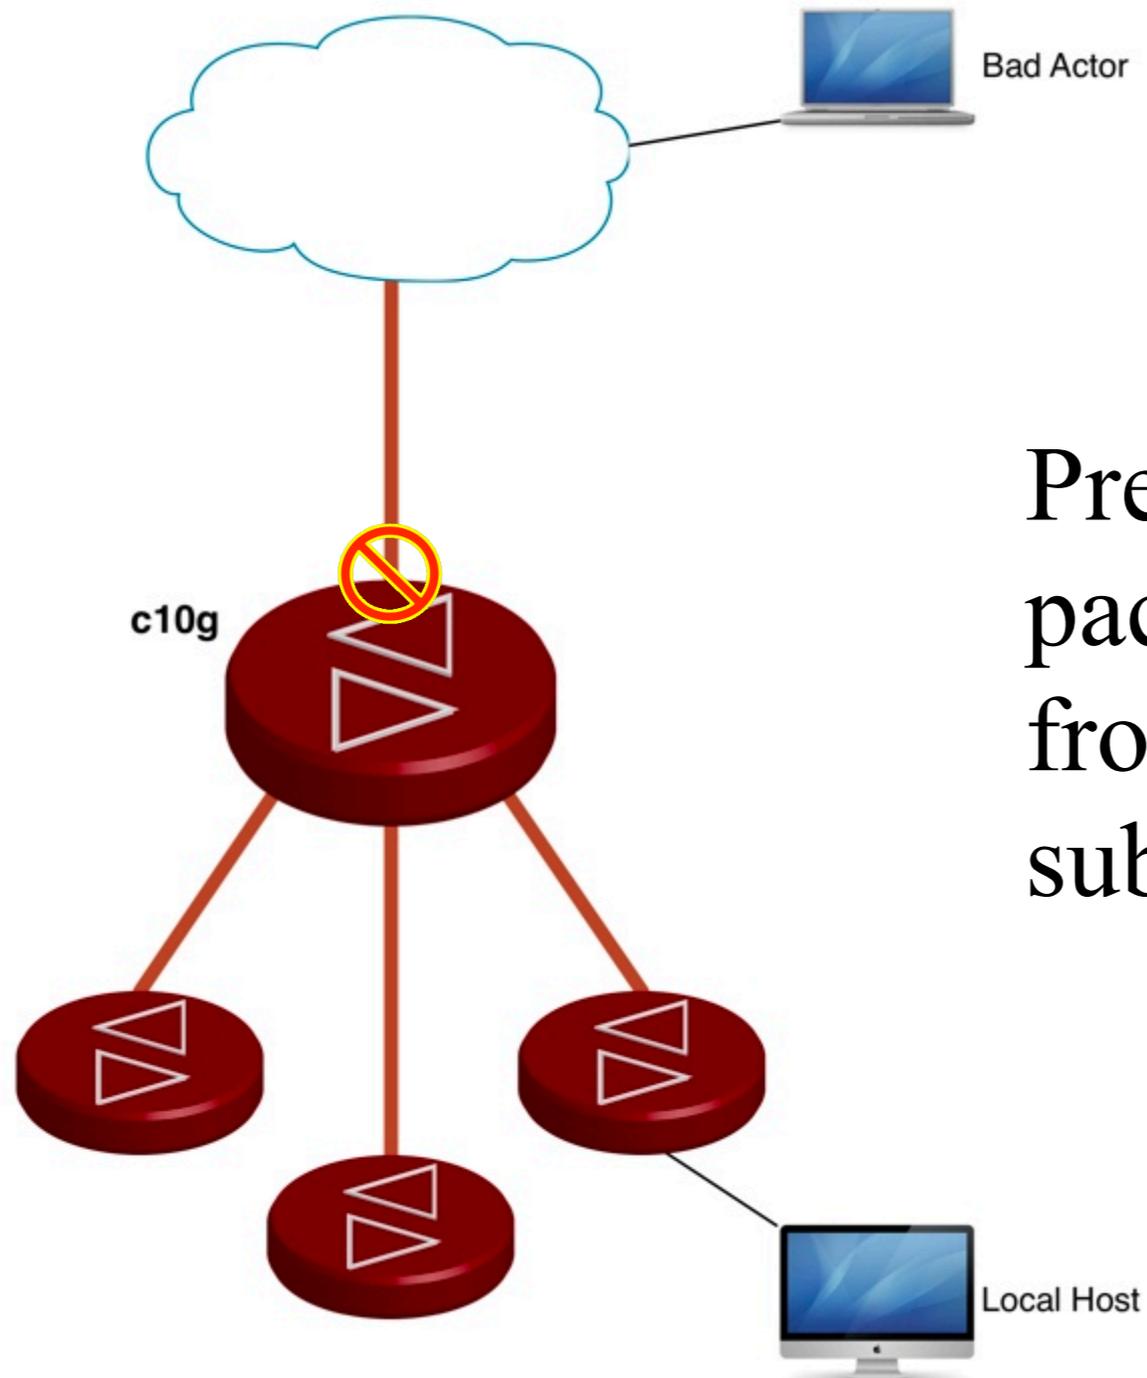


National Energy Research
Scientific Computing Center





Preventing the delivery of packets destined for or coming from specific addresses (or subnets)



Preventing the delivery of packets destined for or coming from specific addresses (or subnets)



Why Block?

Some addresses get blocked because they are actively attacking.

The most common reason for blocking an address is scanning.

If a systems isn't doing any harm, why block it?



Why Block?

Some addresses get blocked because they are actively attacking.

The most common reason for blocking an address is scanning.

If a systems isn't doing any harm, why block it?

- to limit reconnoissance information
- to reduce the likelihood of return



Blocking at NERSC

The vast majority of blocks at NERSC are thrown automatically by our intrusion detection system (Bro).

Some blocks are made manually in response to specific situations.

Automatic blocking is, of course, the most effective mitigation and the faster the better.



Blocking at NERSC

Bro can throw blocks for:

- Scanning
- Suspicious SSH activity
- Root logins
- Suspicious URLs
- Etc.



How Much does NERSC Block?

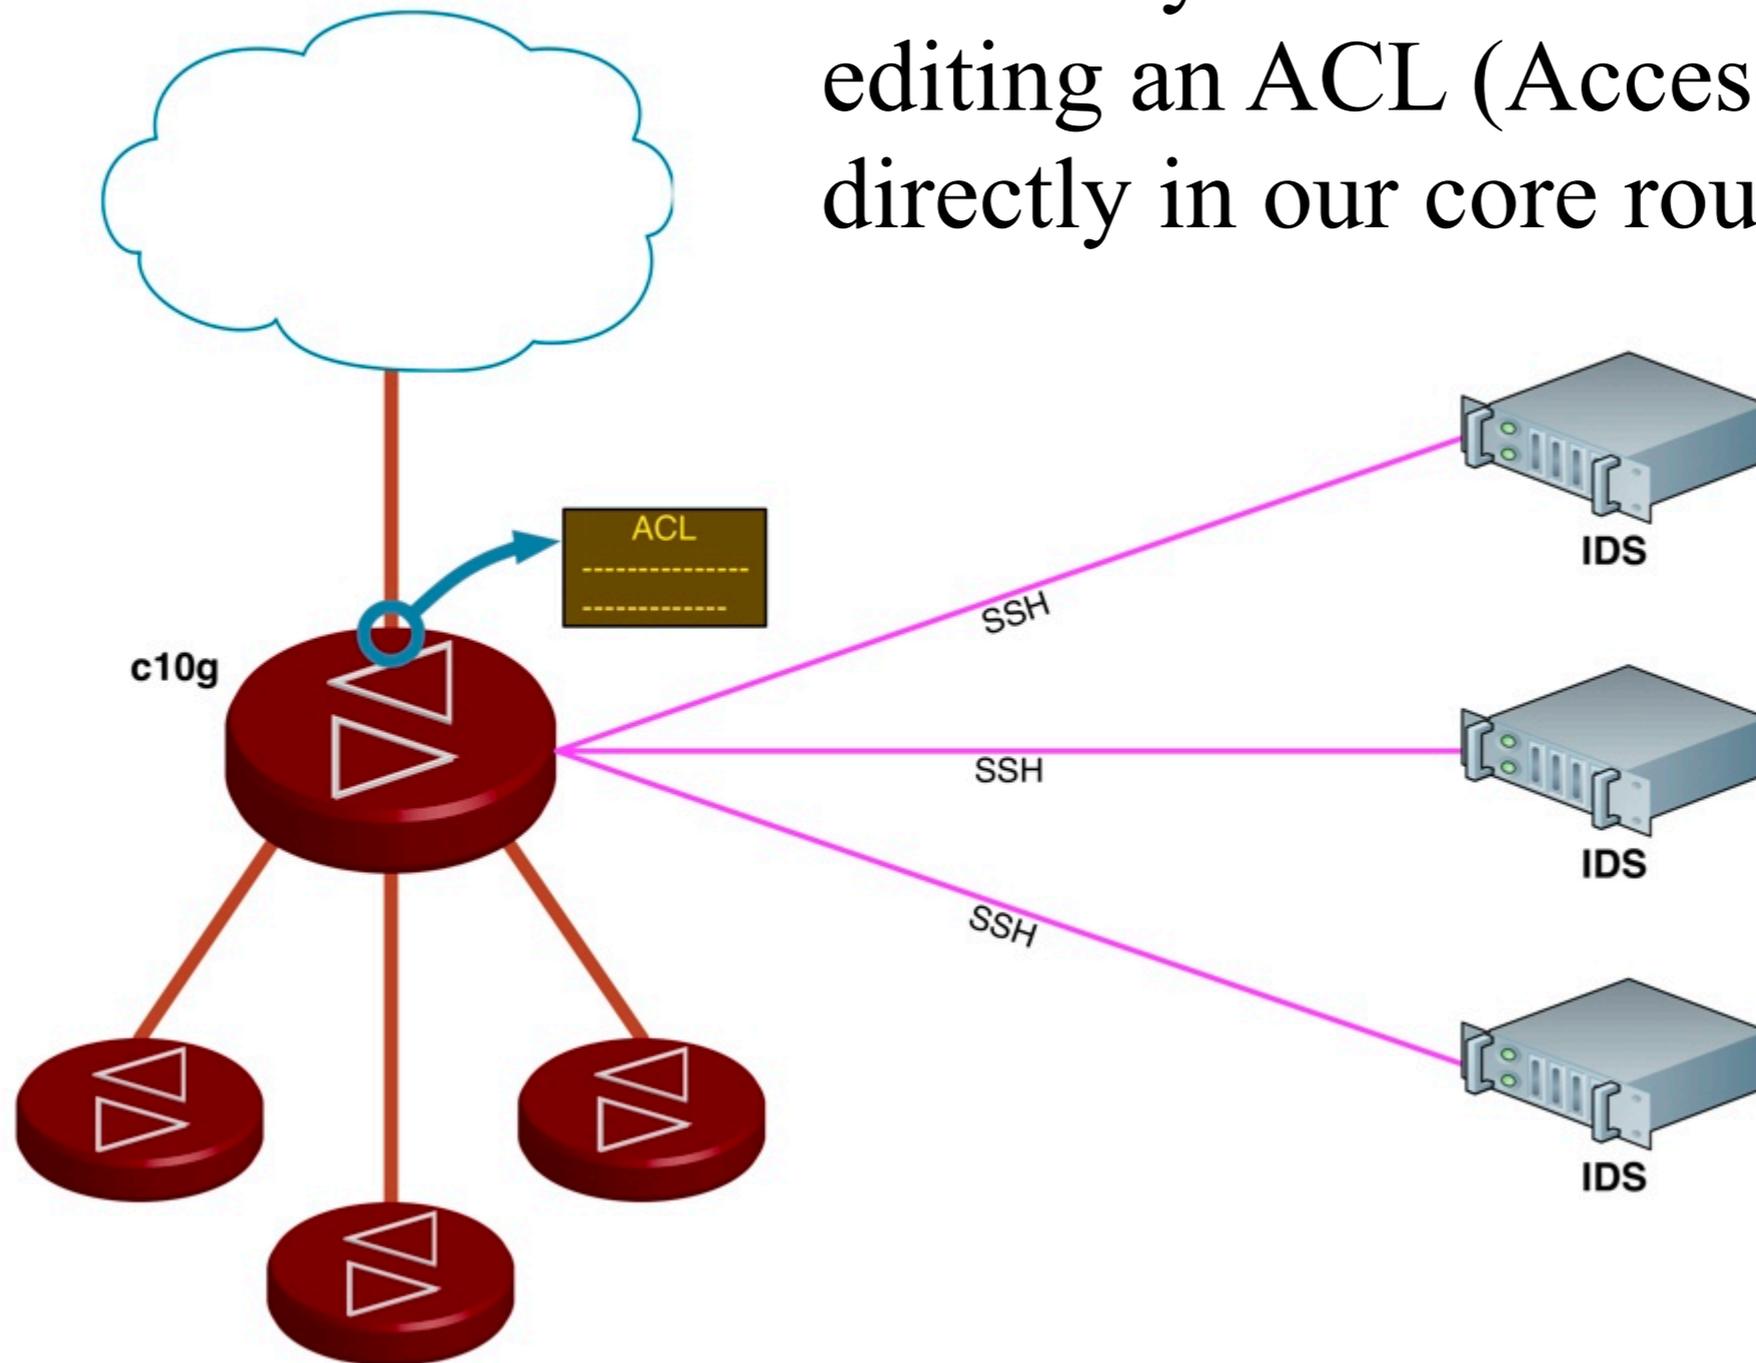
NERSC typically blocks anywhere from a hundred to several hundred addresses every day.

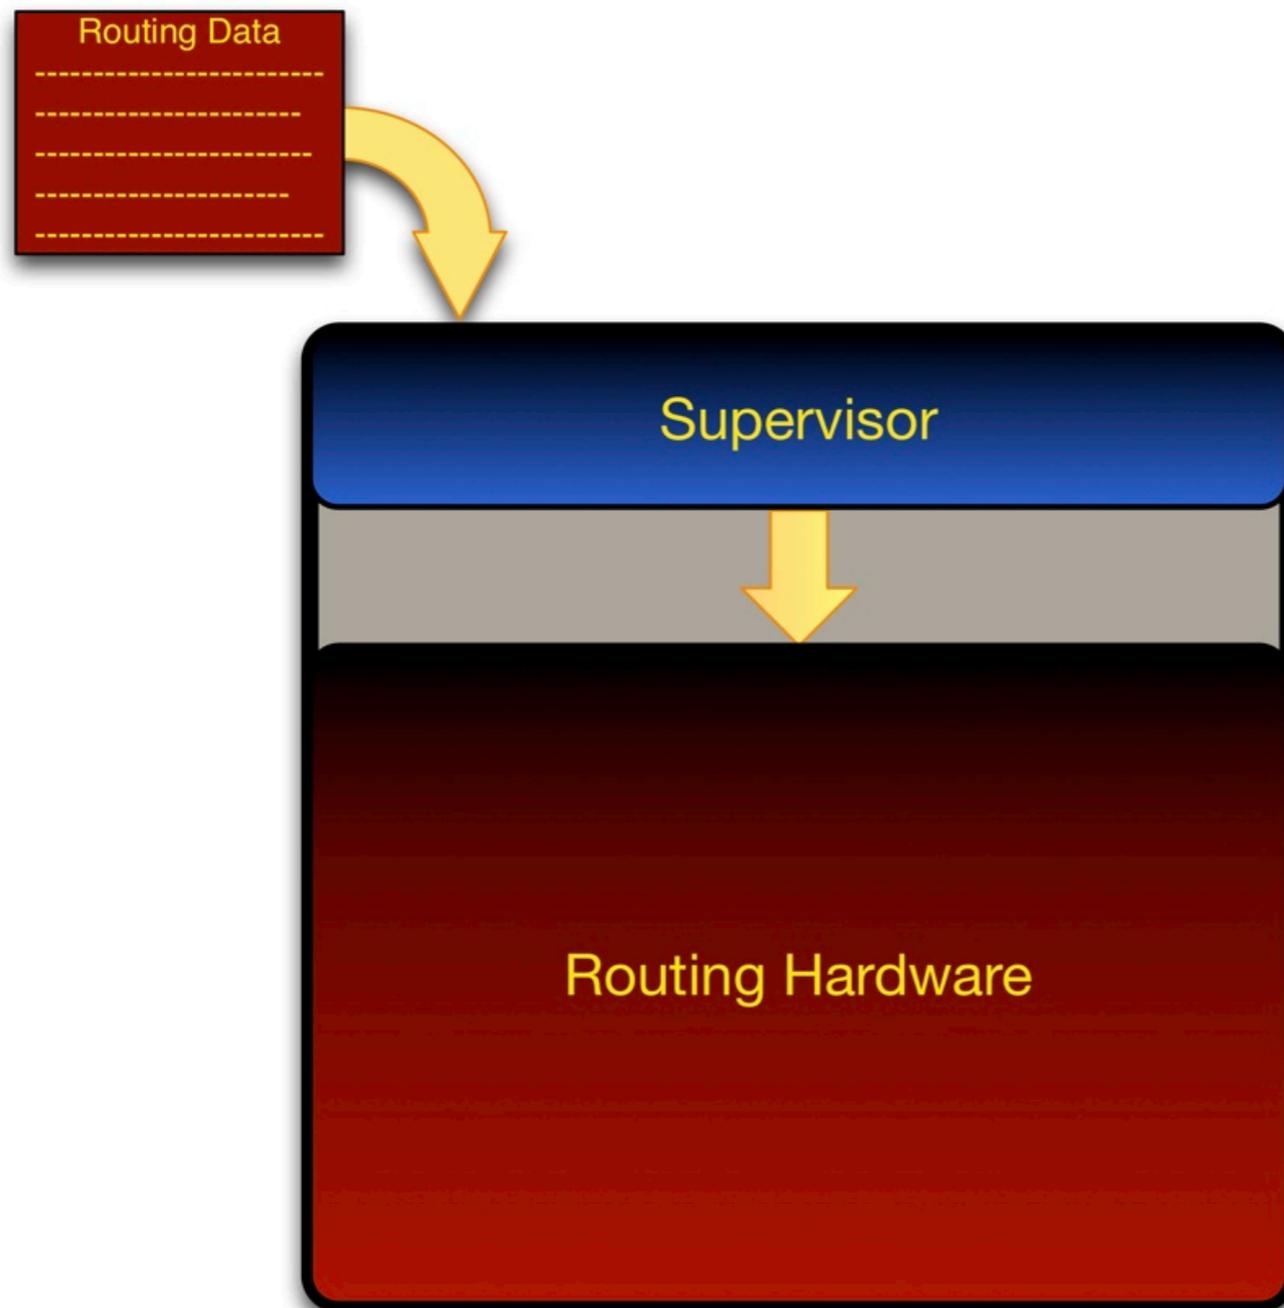
Rarely 1,000 or more address blocks may be thrown in just one hour.



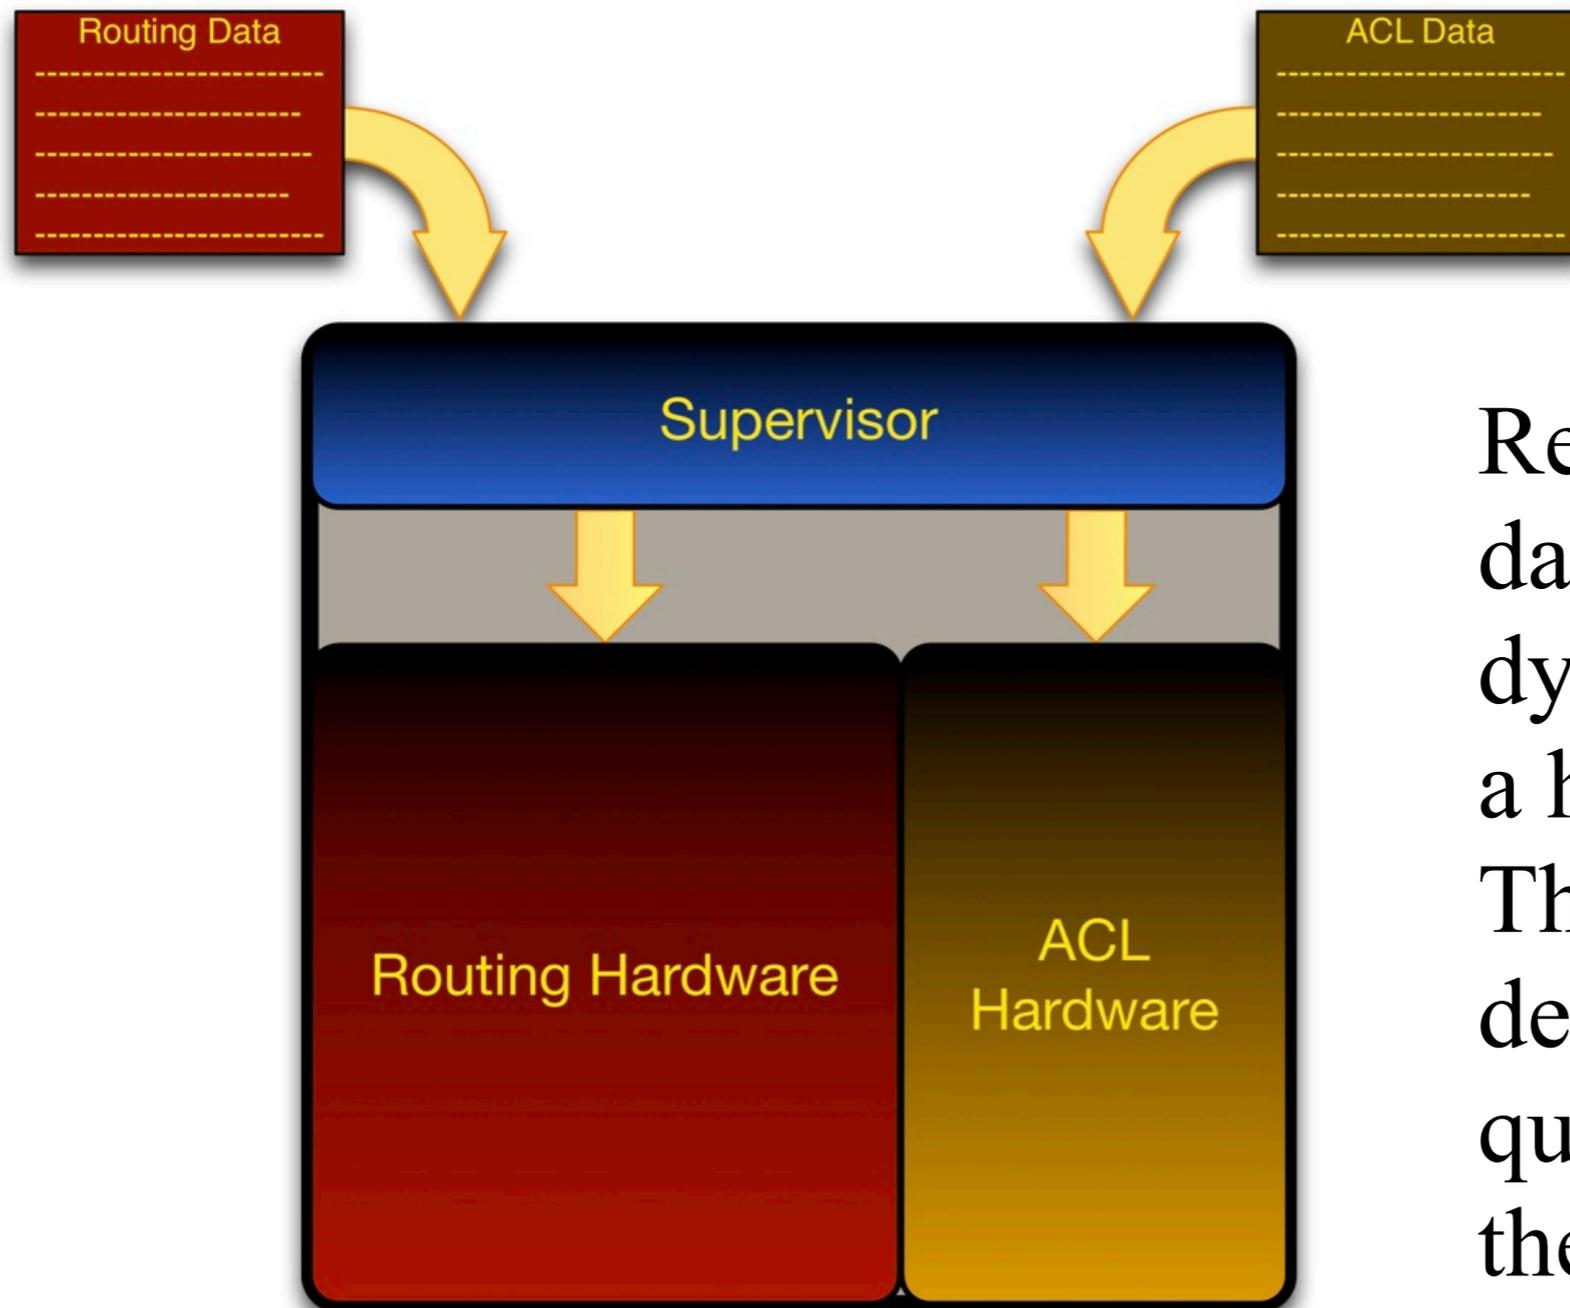
How Does NERSC Block?

Currently NERSC blocks addresses by editing an ACL (Access Control List) directly in our core router.





Relatively simple routing data can be quickly and dynamically encoded into a hardware decision tree. This allows routing decisions to happen very quickly without burdening the CPU (Supervisor).



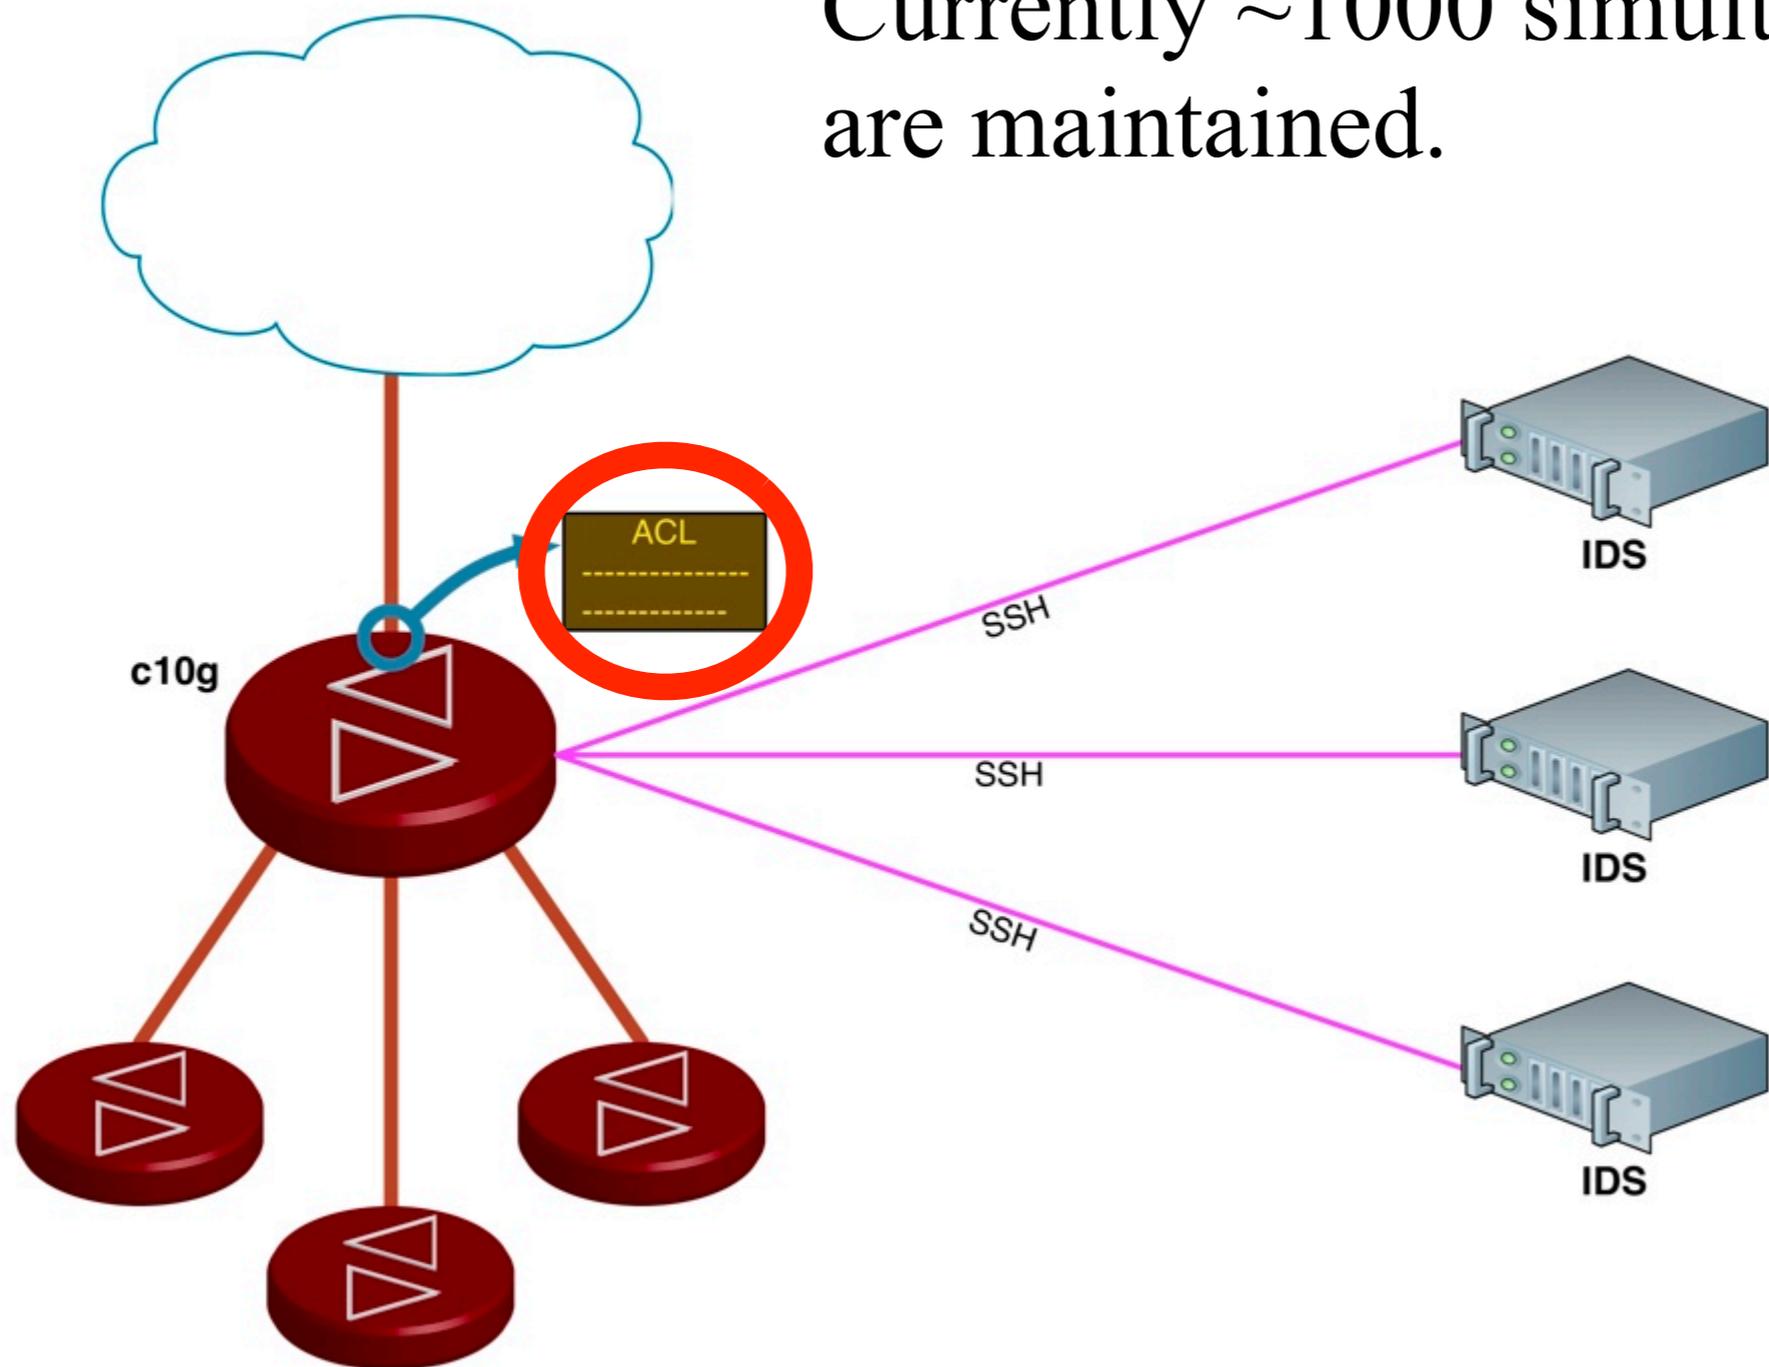
Relatively simple routing data can be quickly and dynamically encoded into a hardware decision tree. This allows routing decisions to happen very quickly without burdening the CPU (Supervisor).



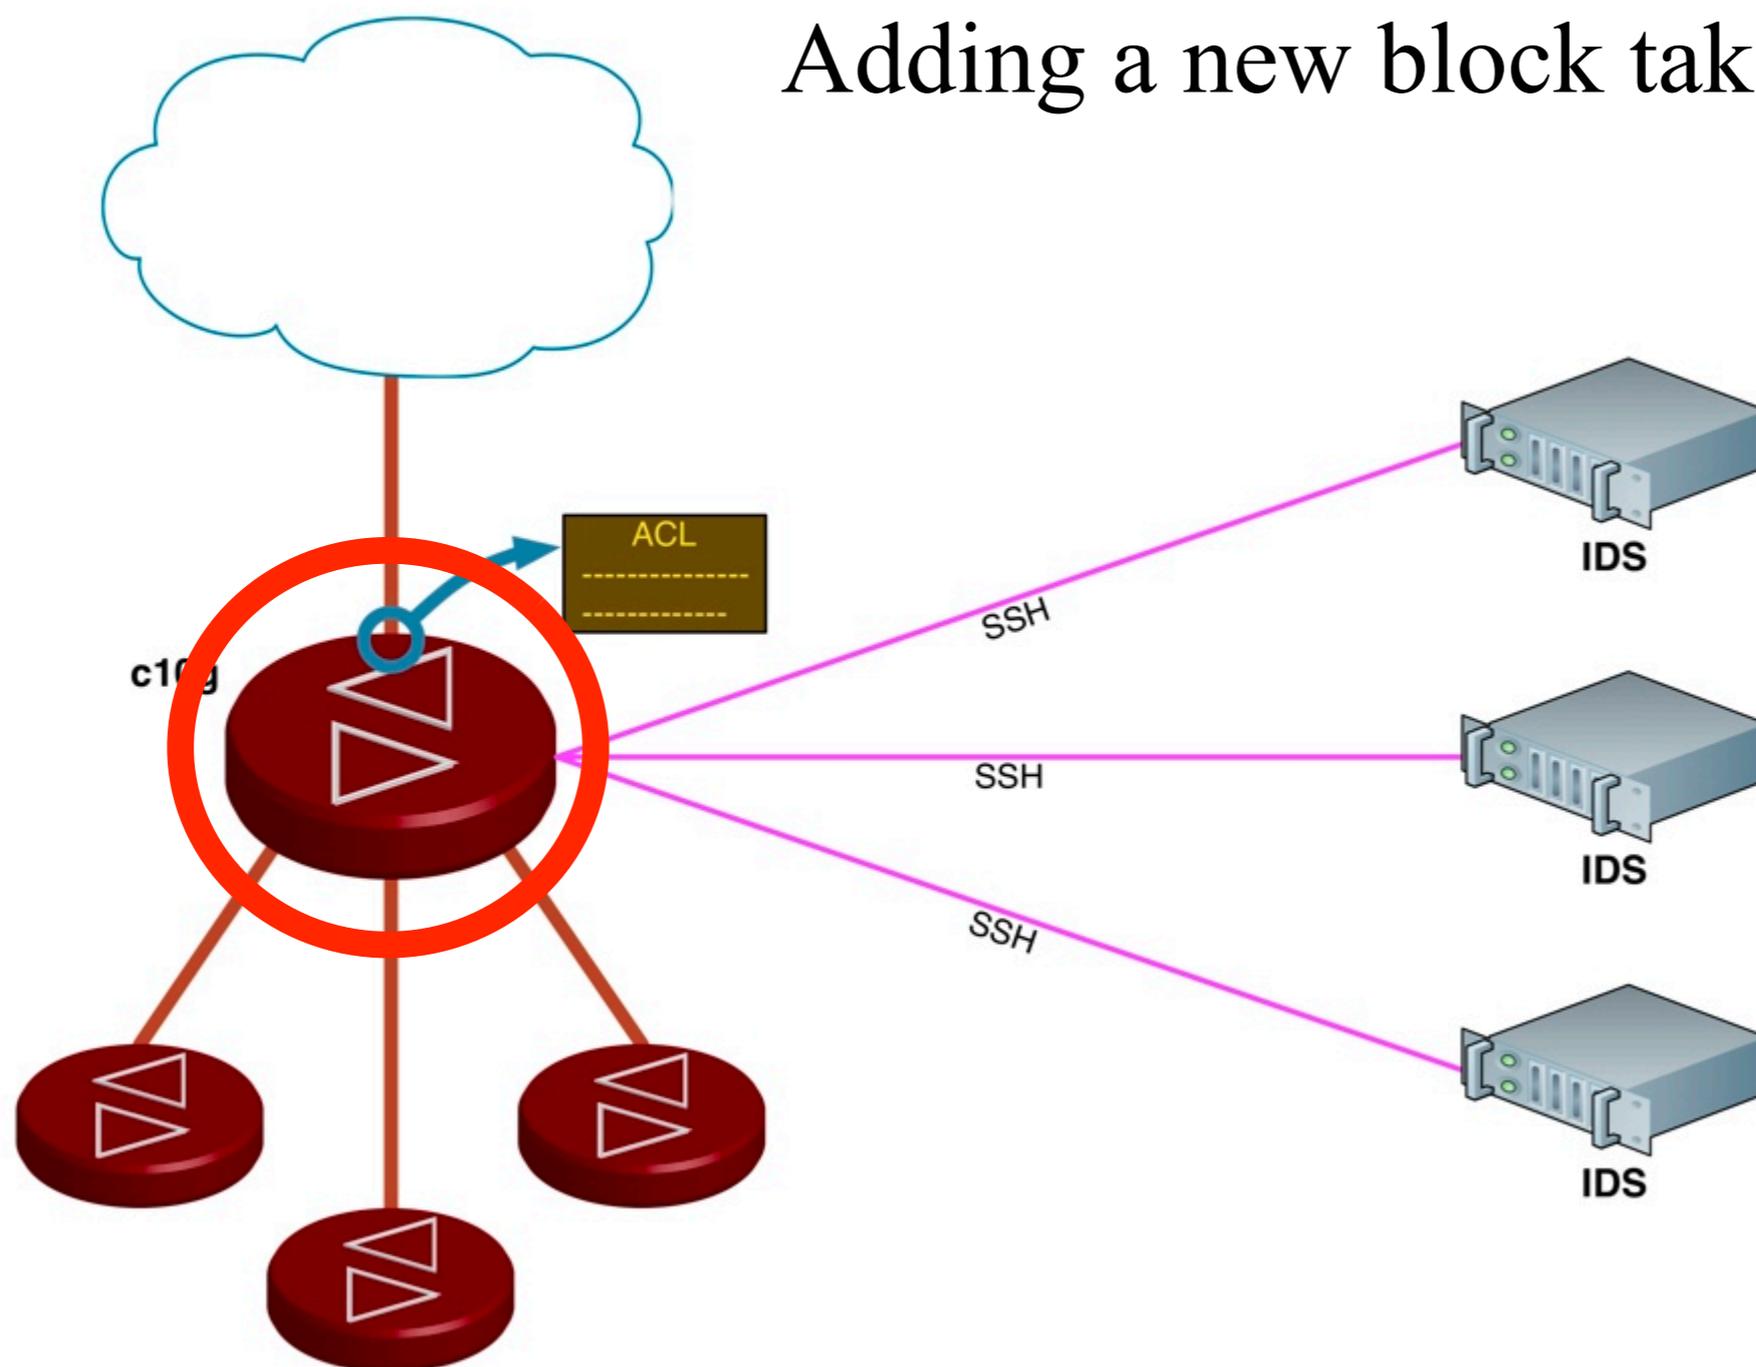
Problems with ACL Editing

- Limited number of blocks can be held
- Blocks take time to add
- Blocks cannot be added while any other configuration is being done on the core router

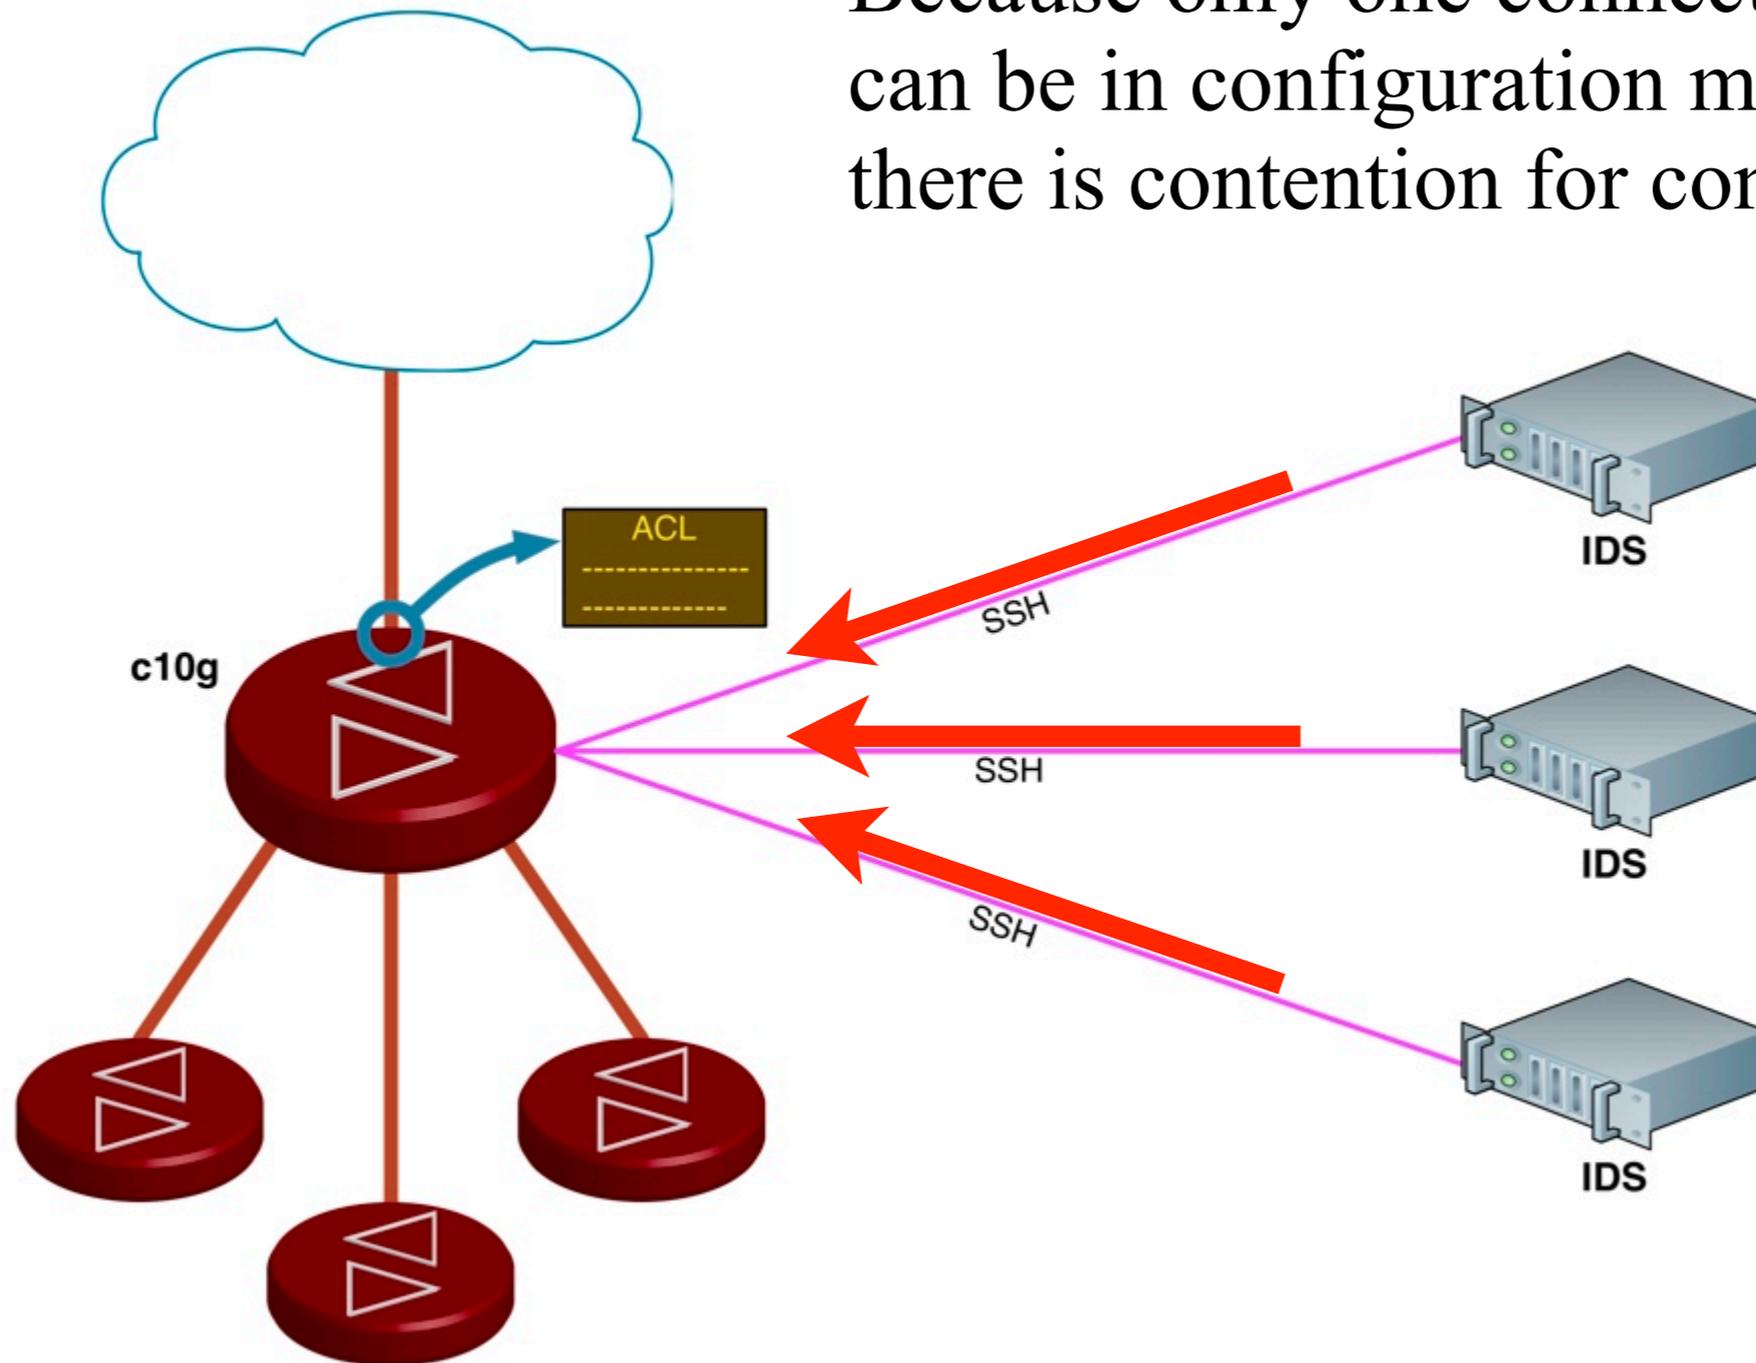
Currently ~1000 simultaneous blocks are maintained.



Adding a new block takes ~5-8 seconds.



Because only one connection to the router can be in configuration mode at a time, there is contention for configuration mode.



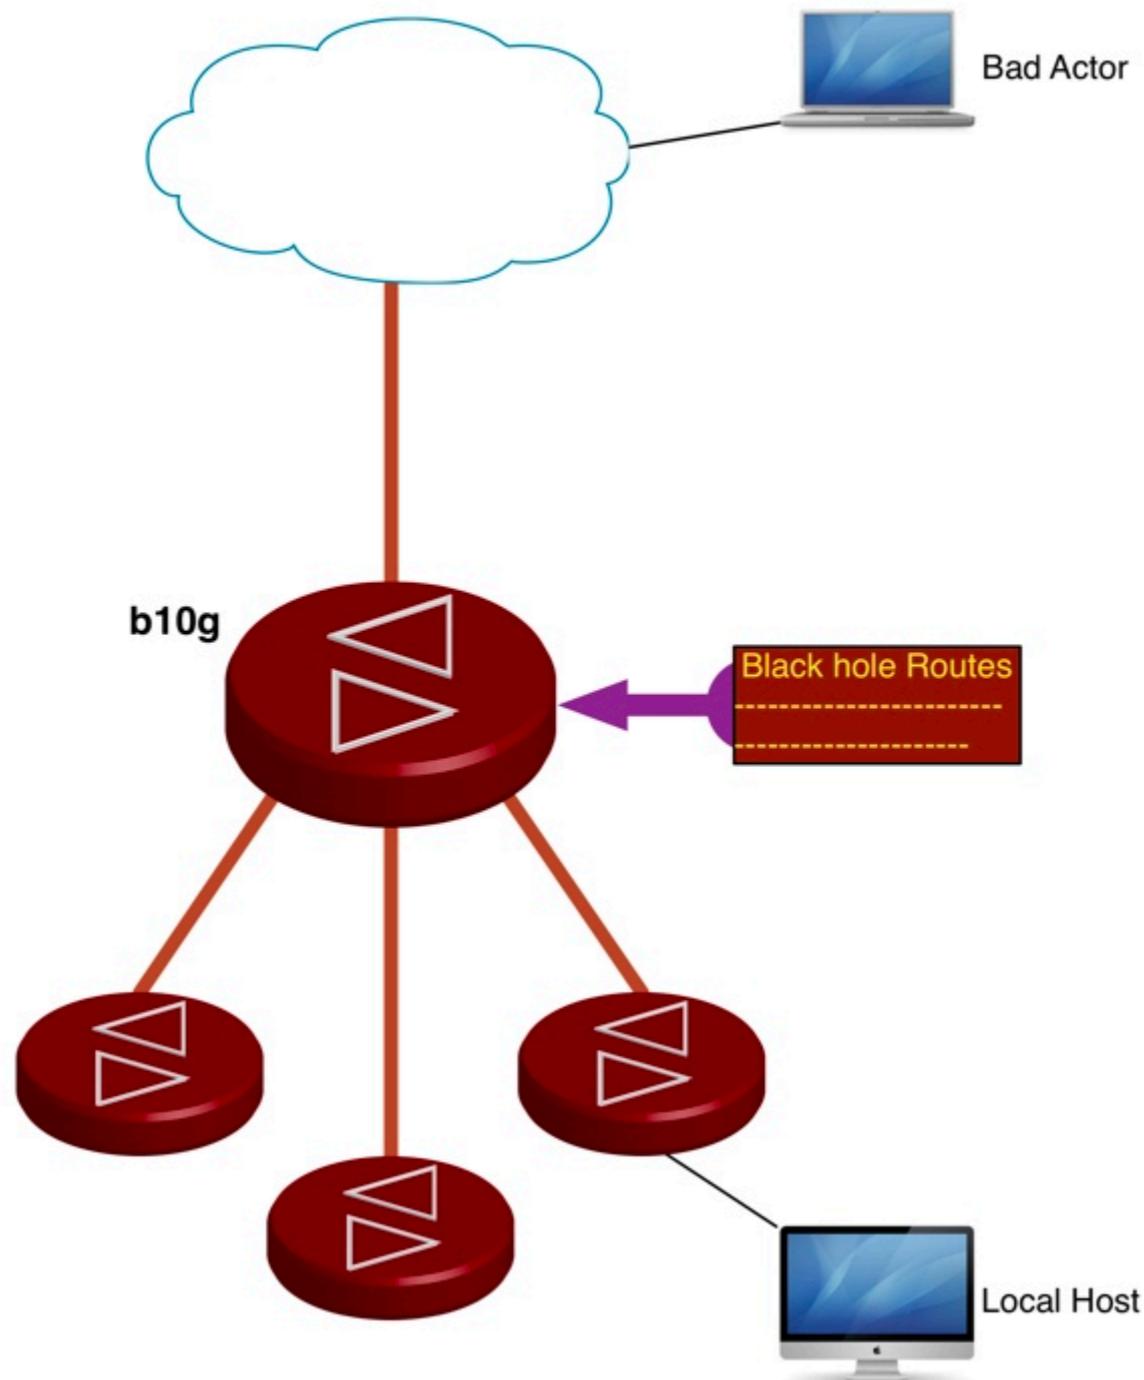


New Method: RTBH Filtering

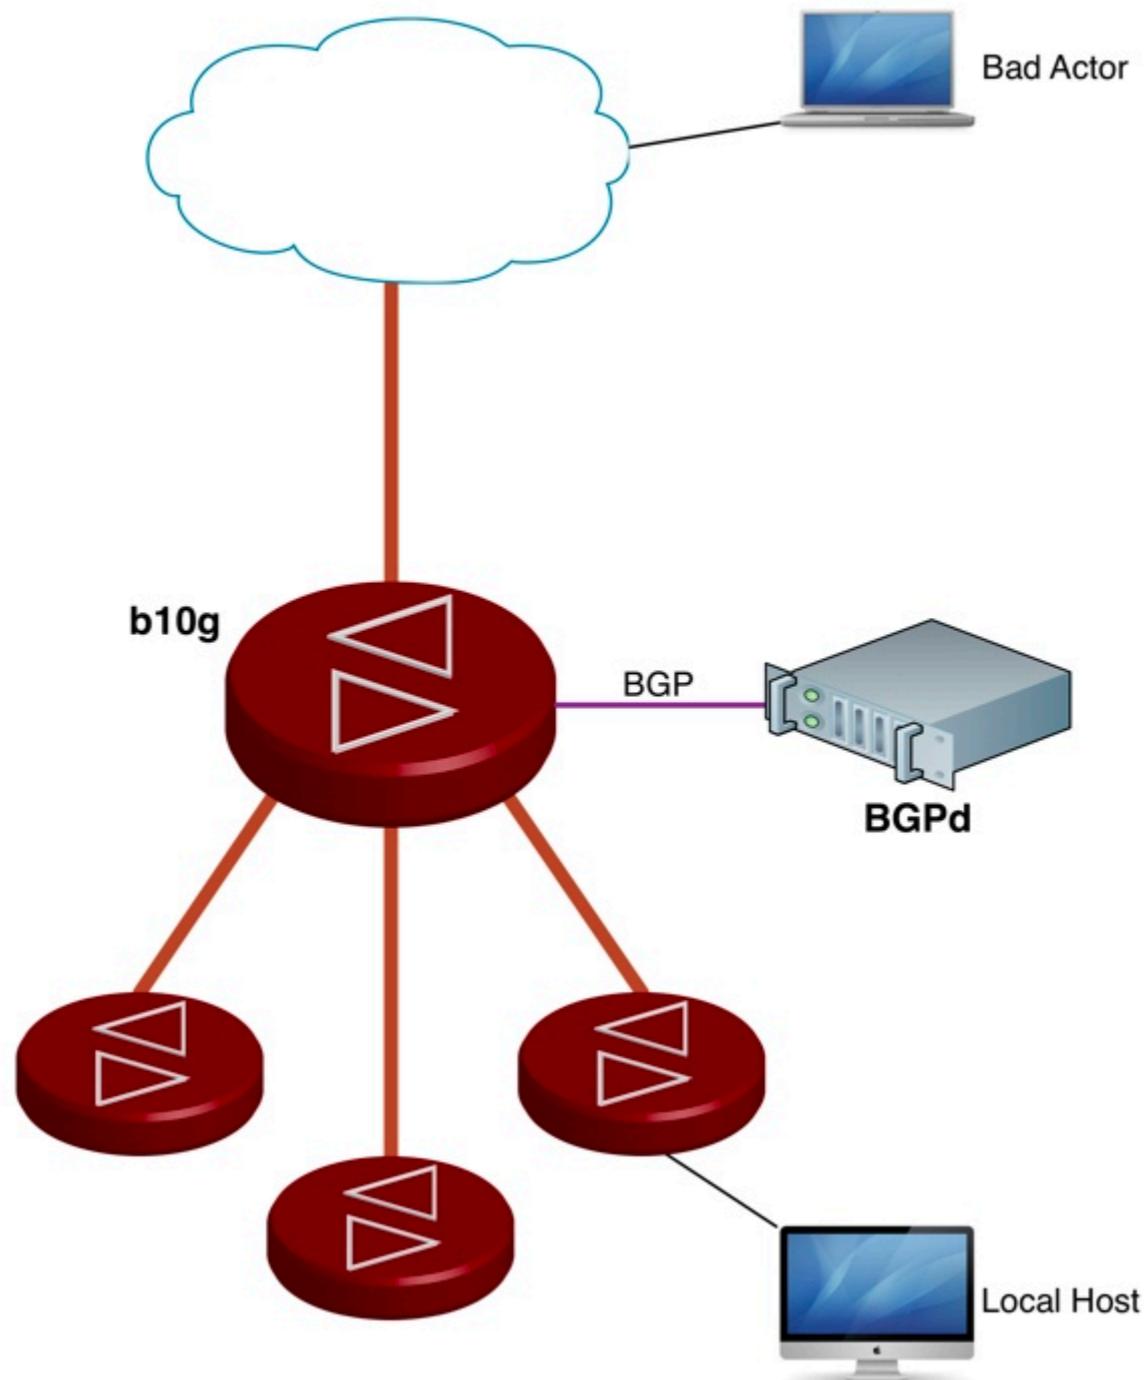
Our new method of throwing blocks is known as RTBH (Remotely Triggered Black Hole) filtering.

The idea is that rather than adding an ACL rule, we inject a special route called a “black hole route”.

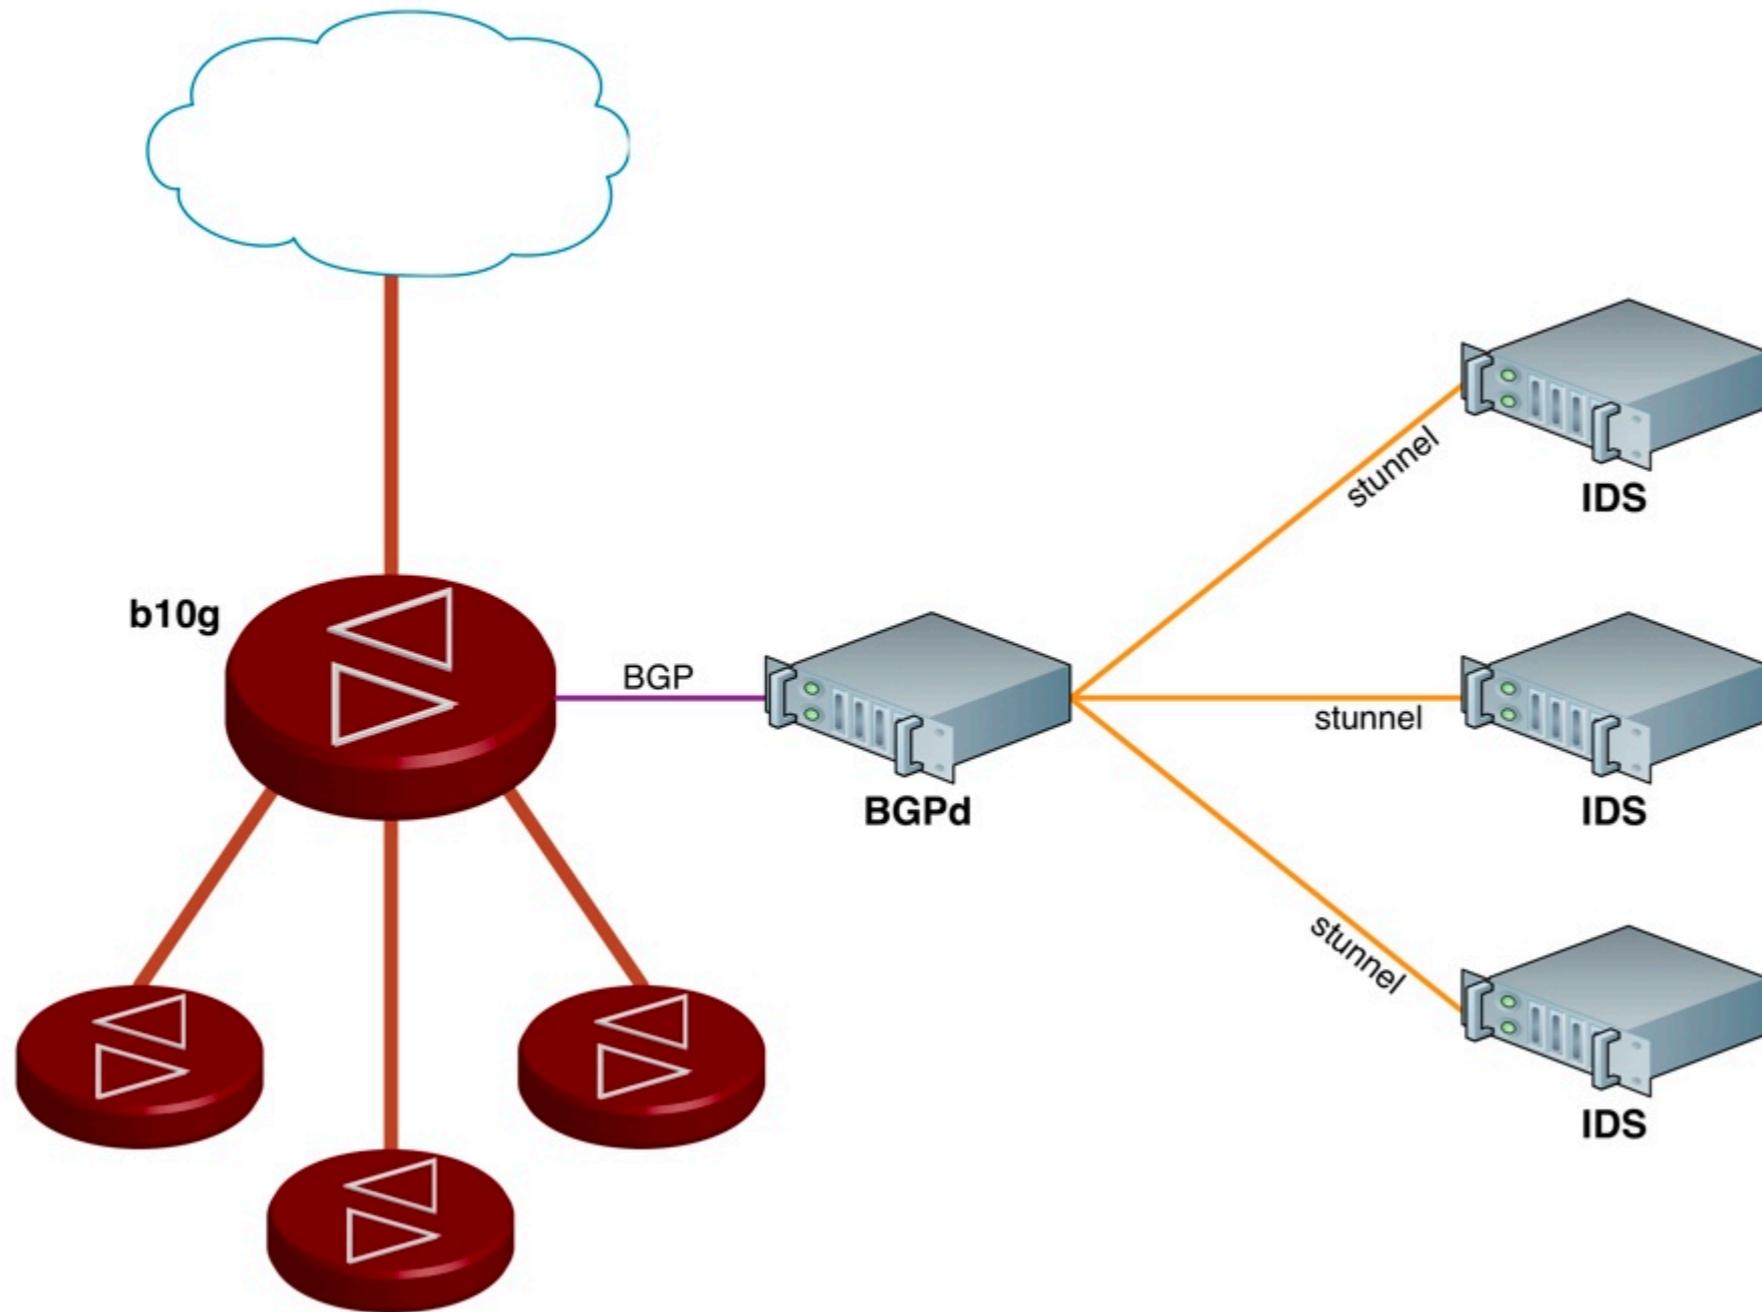
A black hole route is basically a route that declares the “next hop” to a specific destination is /dev/null.



Packets destined for the blocked IP address are silently dropped by re-routing them (into a black hole).



BGP is a routing protocol that allows us to quickly inject black hole routes without the need to reconfigure the router.





Now How Many Blocks?

Since we're now injecting *routes* rather than *ACL rules*, we can maintain **many** more concurrent blocks.

In a recent test I was able to add just under two million black hole routes before the router simply ran out of memory.

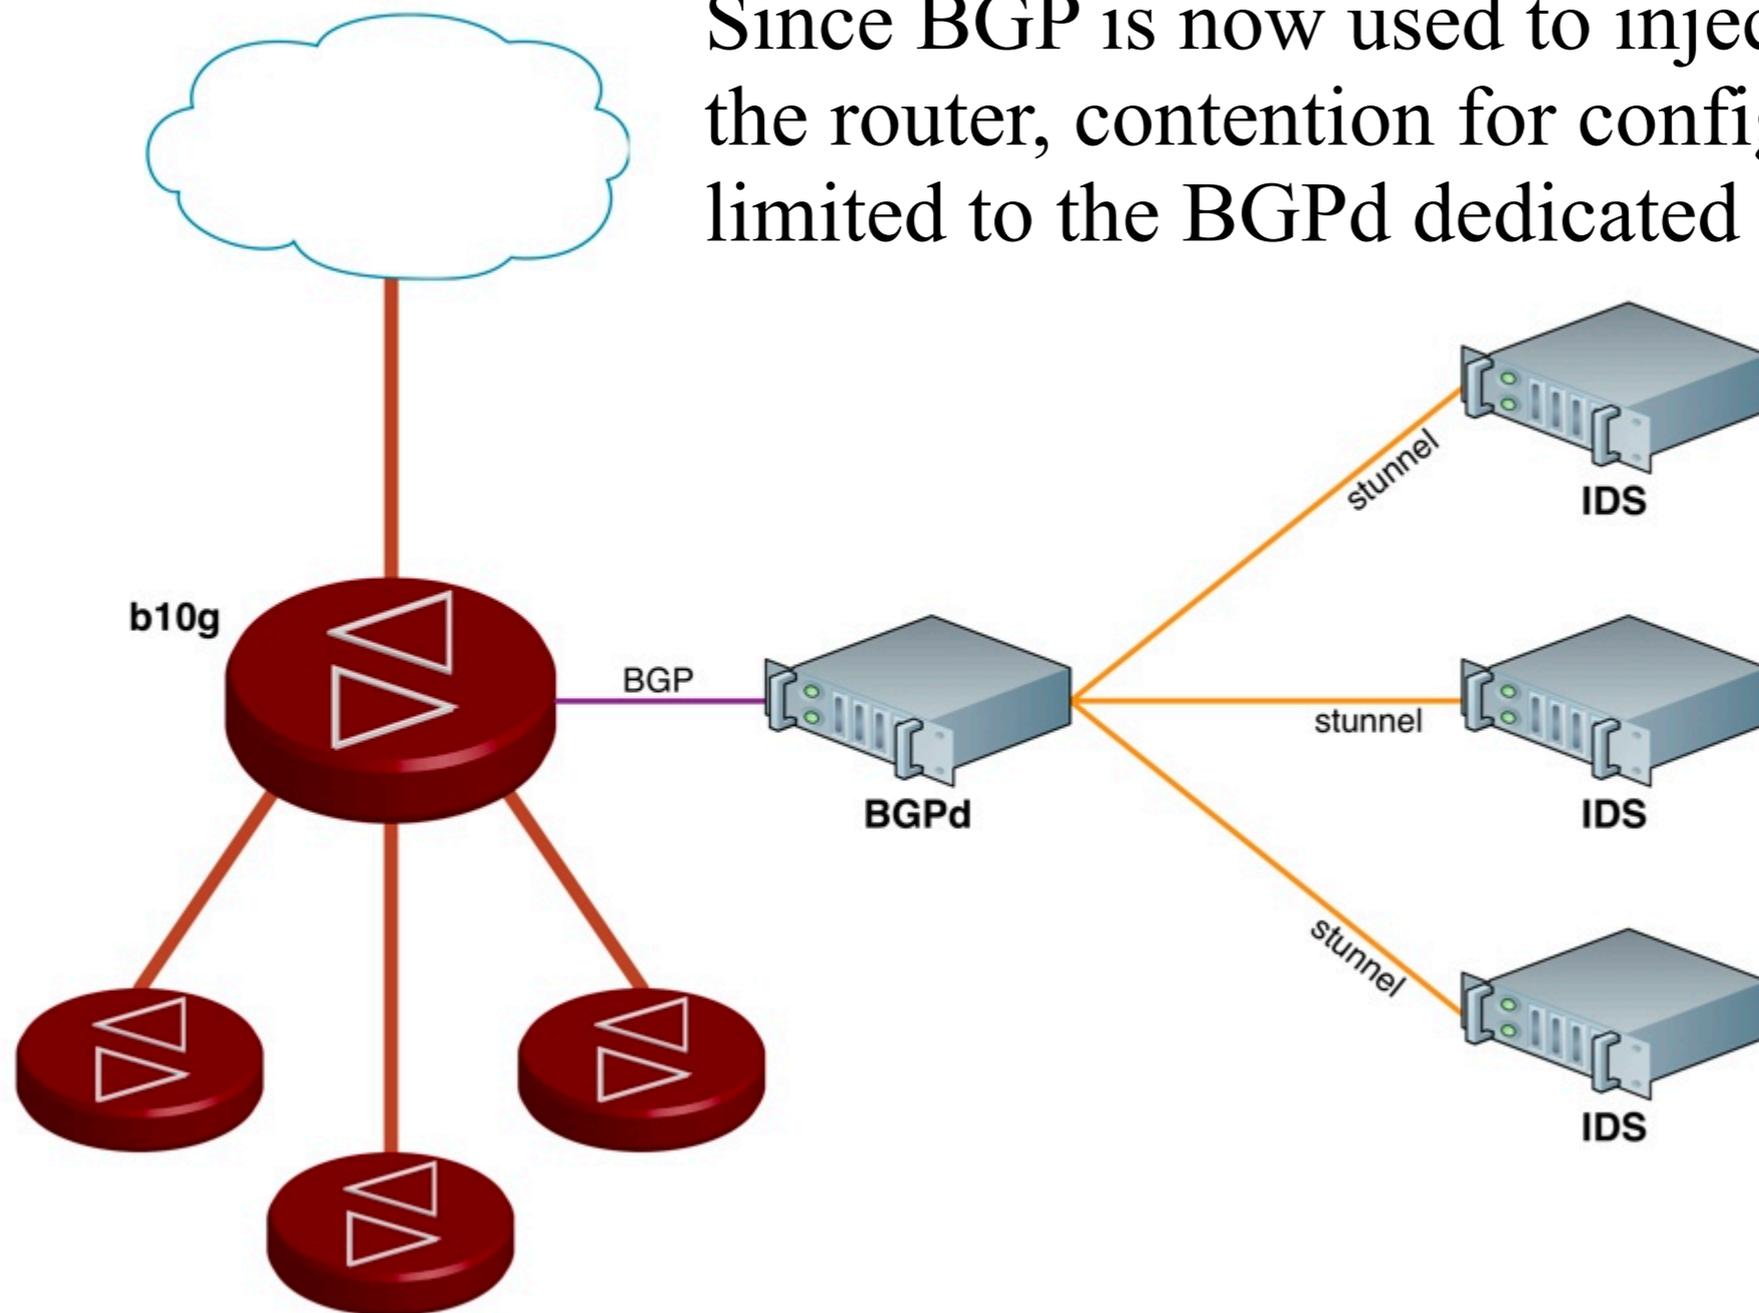


Now How Fast?

Since the router can dynamically add new routes to the routing hardware, adding a new black hole route is significantly faster.

We're seeing "sub-second" times to drop an address.

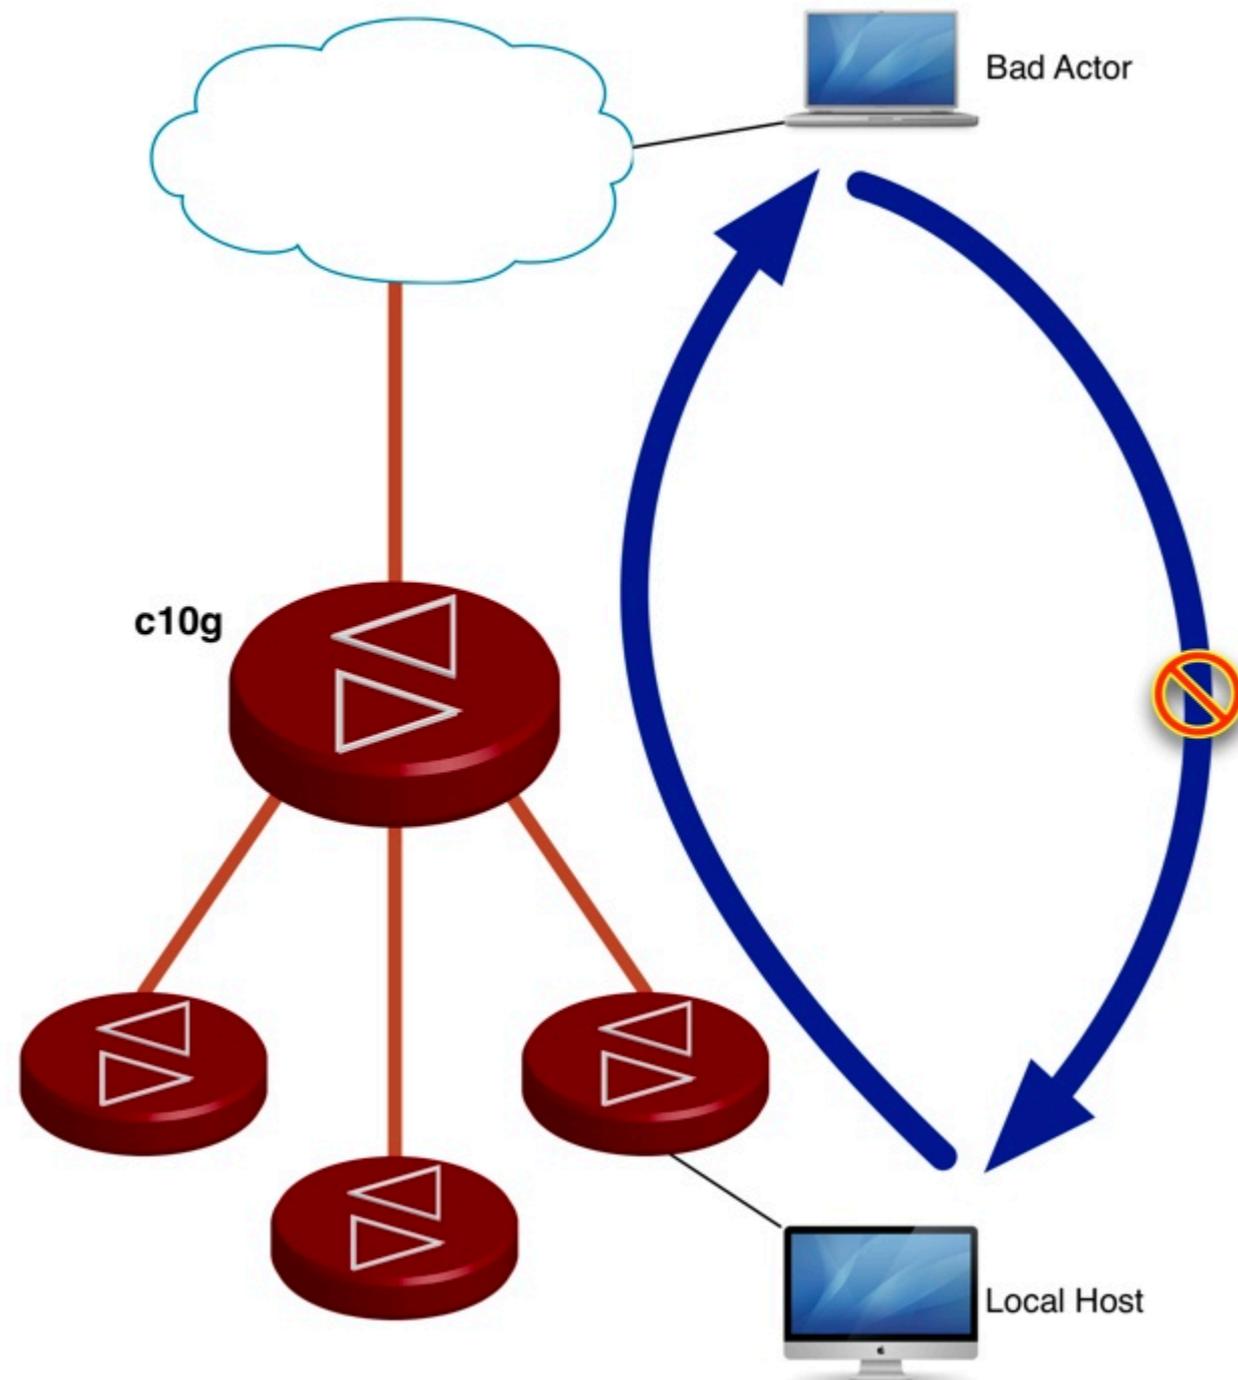
Since BGP is now used to inject new routes on the router, contention for configuration mode is limited to the BGPd dedicated to this function.





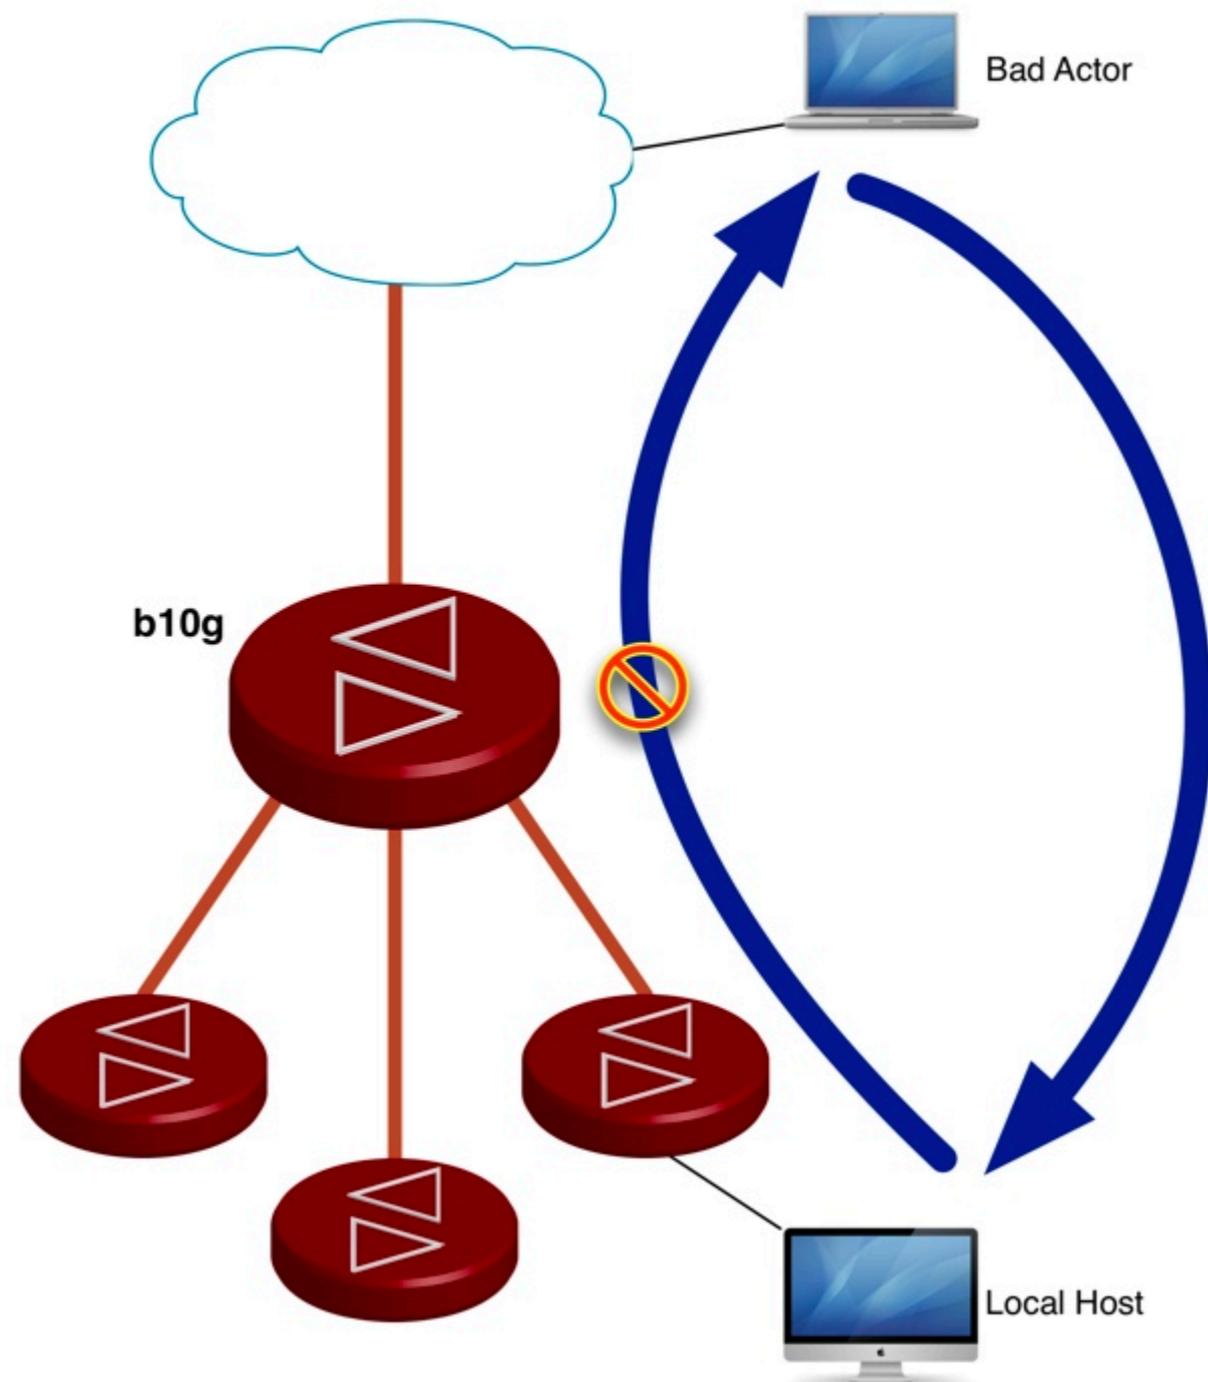
Issues with RTBH Filtering

- Outbound vs. inbound blocking
- Potential for failing open



Our ACL blocking mechanism only drops inbound packets from the target.

Outbound packets are still allowed.



The RTBH filtering mechanism only drops outbound packets to the target. Inbound packets are still allowed.



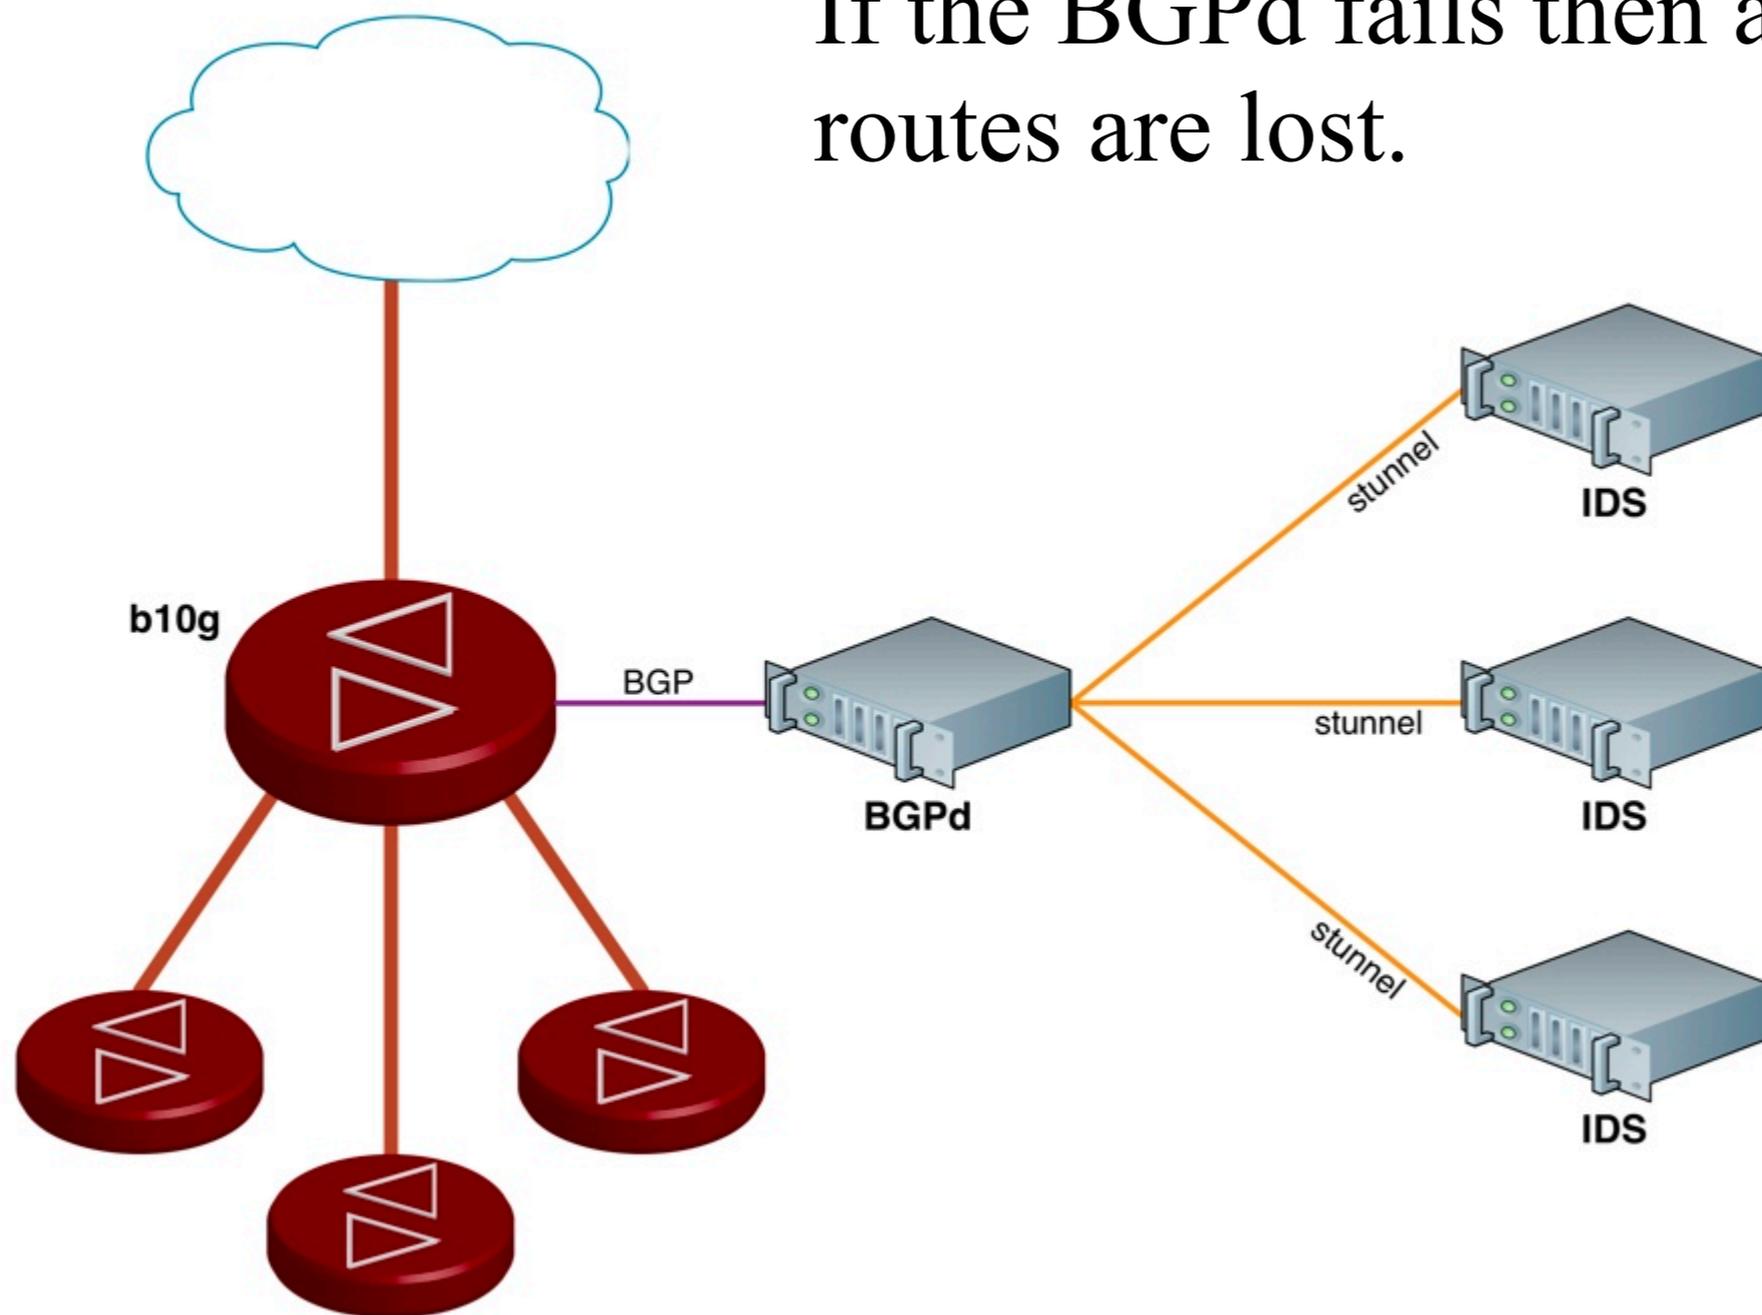
Ingress vs. Egress Filtering

For scanning egress filtering is better.

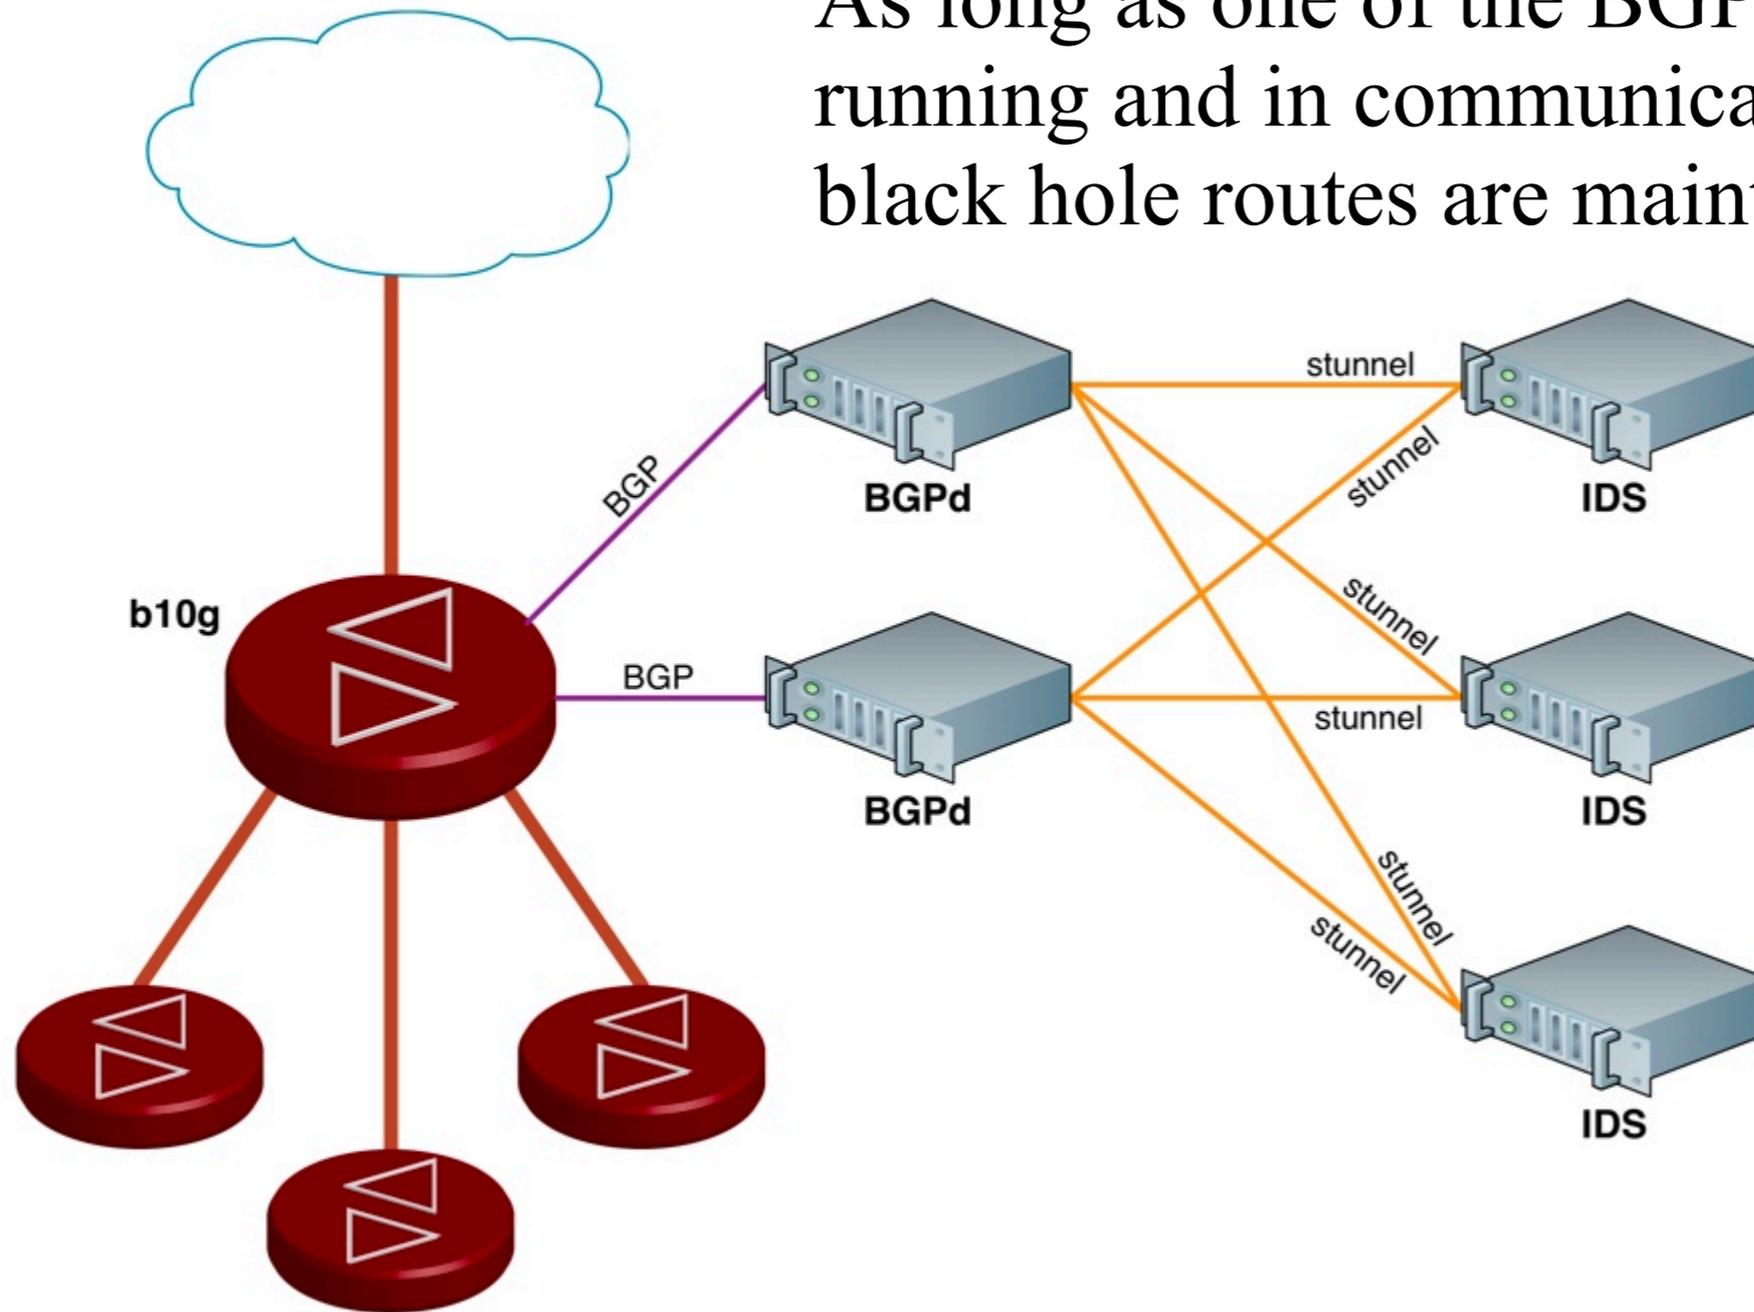
For DOS/single packet attacks ingress filtering is required.

Given NERSC's threat model, egress filtering is probably more effective than ingress filtering.

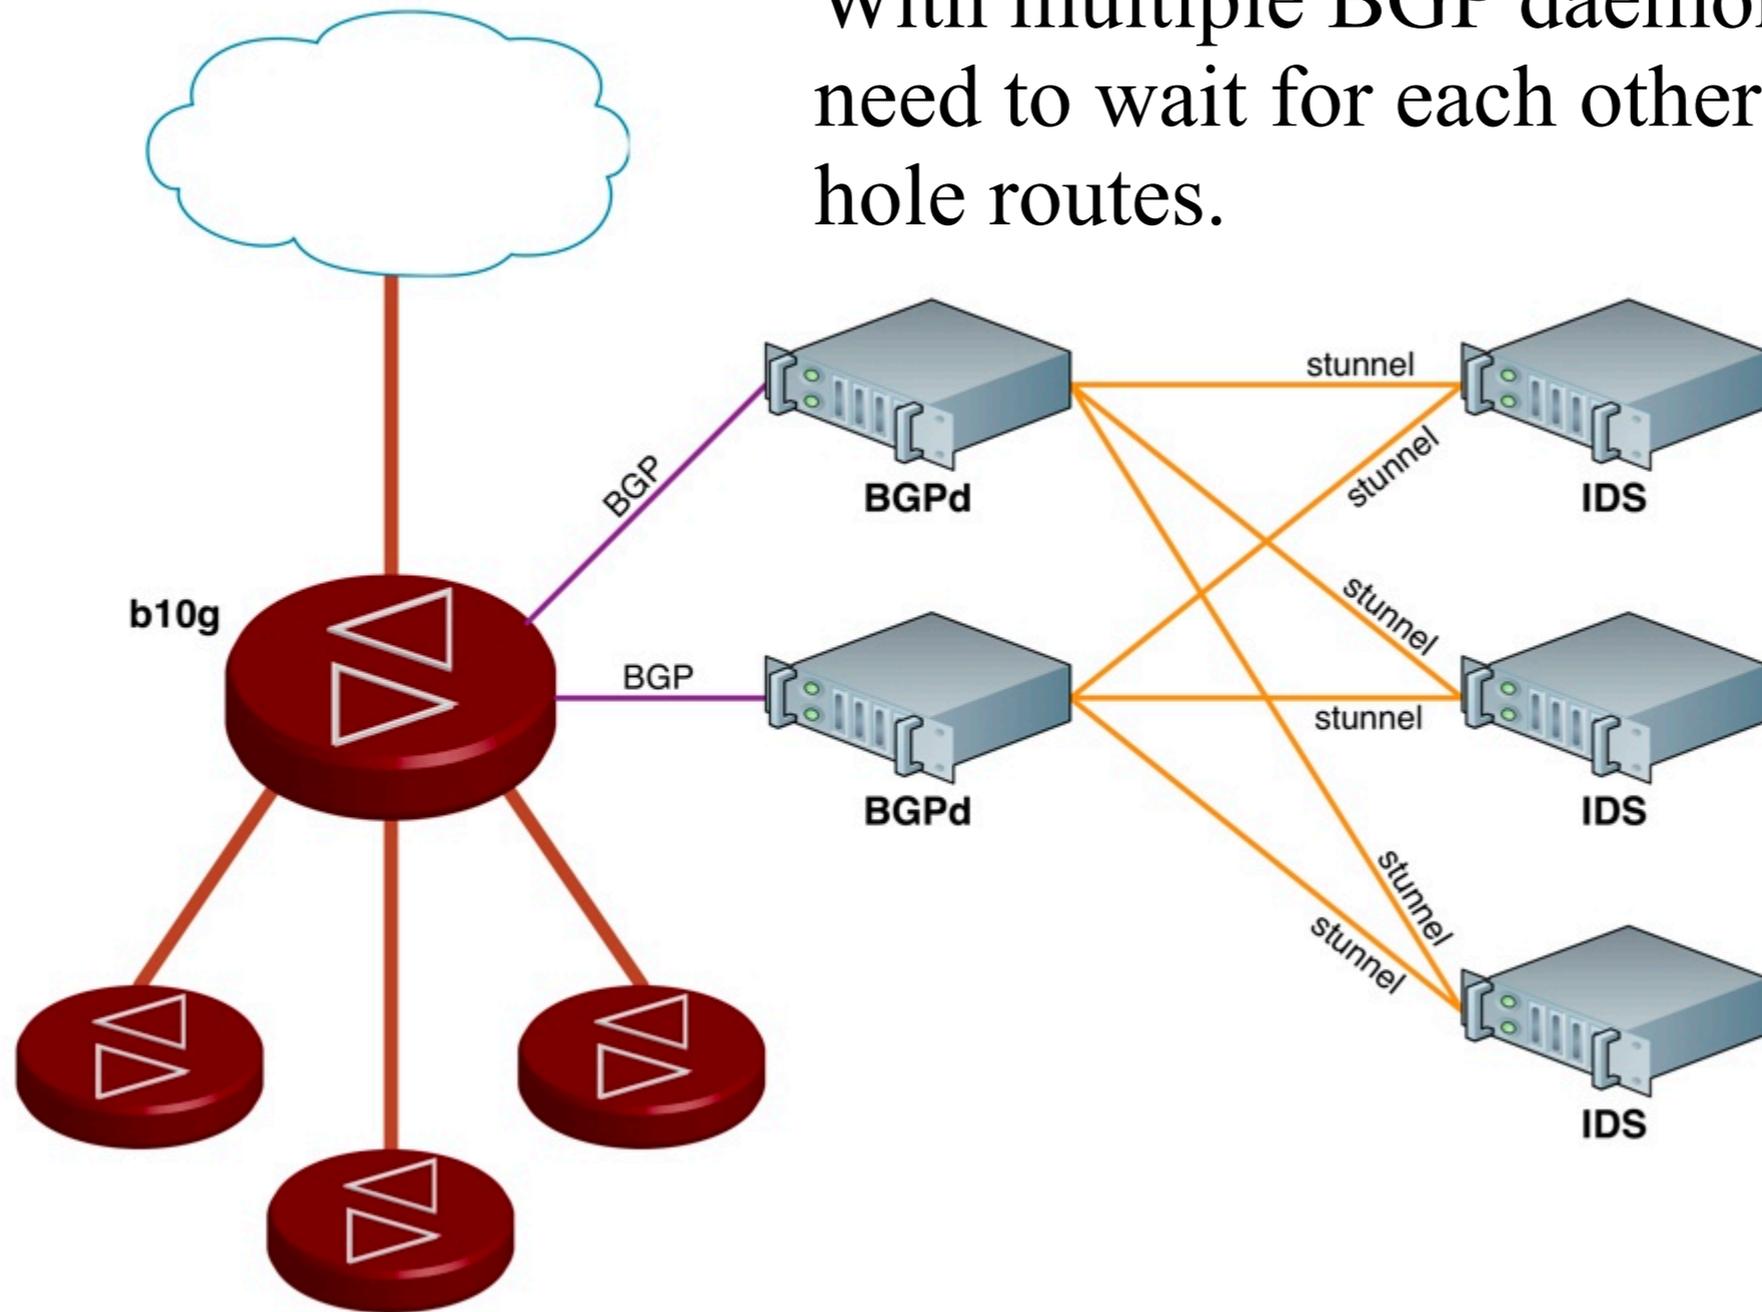
If the BGPd fails then all advertised routes are lost.



As long as one of the BGP daemons is running and in communication with b10g, all black hole routes are maintained.



With multiple BGP daemons, IDSs don't need to wait for each other to install black hole routes.





ACL->RTBH Summary

- Concurrent Blocks: 1,000 -> 2,000,000
- Blocking Speed: 5s -> 0.75s
- Configuration Mode Contention Mitigated