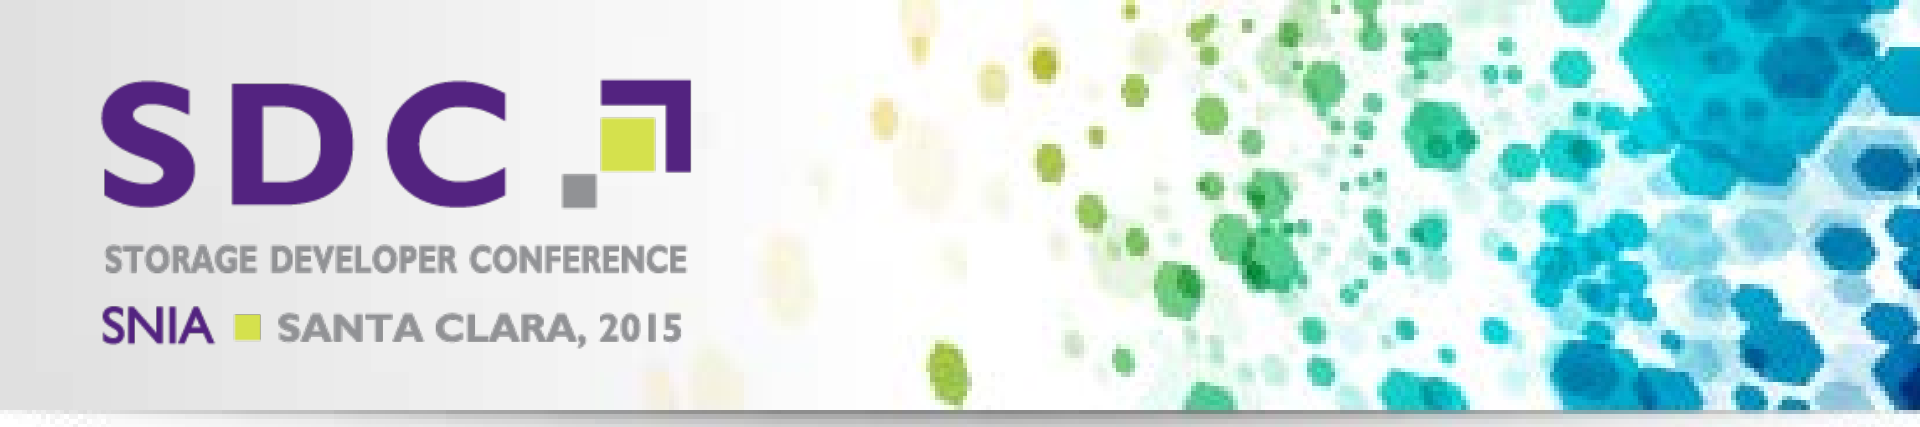

War Stories of Building a Multi-Node AD Client

Oliver Jones

EMC ISILON

What is OneFS?

- A clustered scale out file system
  - Each node has a full view of the file system

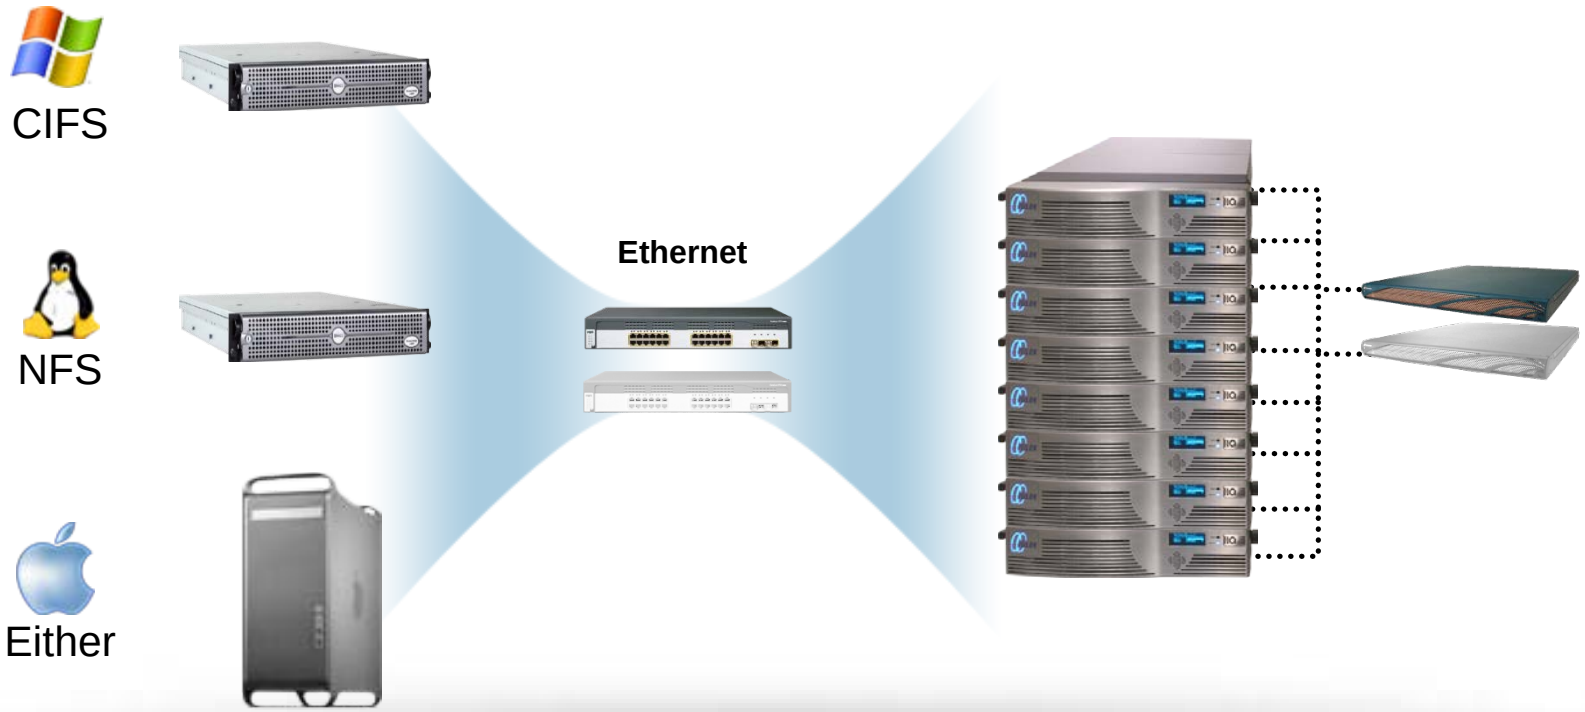

So you want to Join a Clustered Server to AD

- Solution A: Create Individual account per Node
- Pros
  - Don’t have to share passwords
- Cons
  - 144 Machine accounts to manage
  - What happens when we want to raise our maximum node size?

So you want to Join a Clustered Server to AD

- Solution B: Create a Shared Account
- Pros:
  - Single Account to Keep track of and manage
- Cons:
  - Have to share password
    - Who Updates that password?
    - All nodes must affinitize to a single AD when password is changed

Solution

- Creating a single machine account sounds hard lets just have an account per node

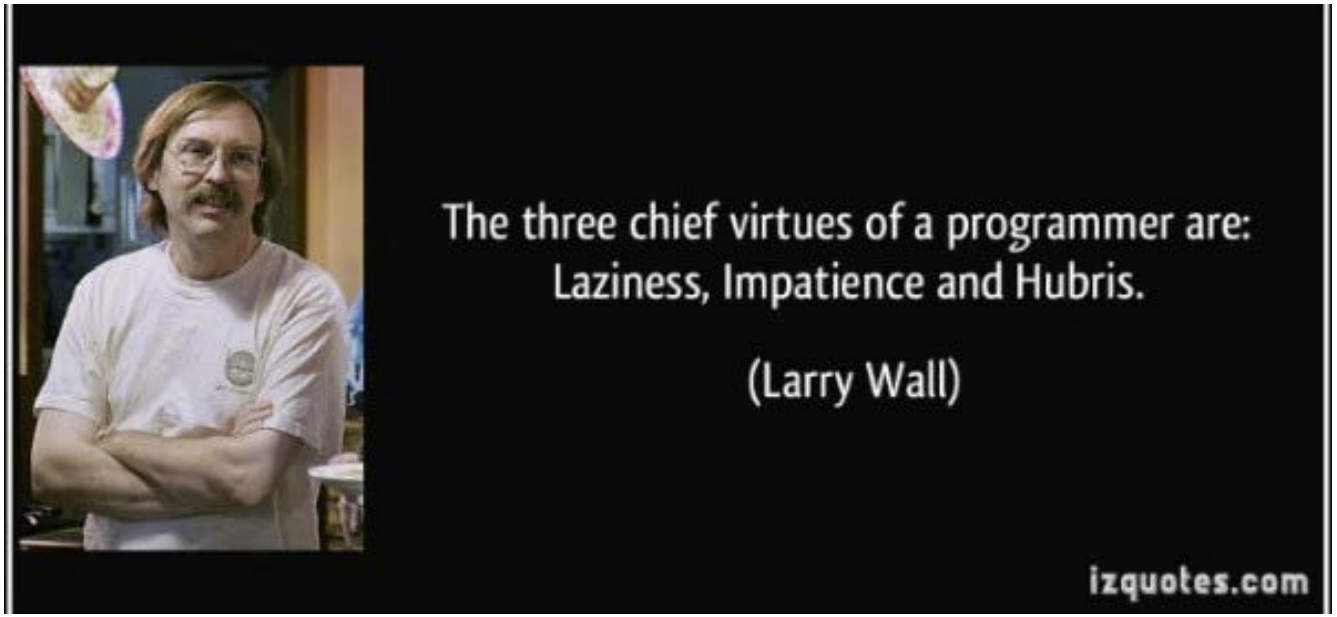

Problem:
We created on hell of a DDOS platform

1. We use CLDAP pings to find the fastest DC
2. All nodes tend to select a single DC
3. Nodes send a torrent of ldap traffic

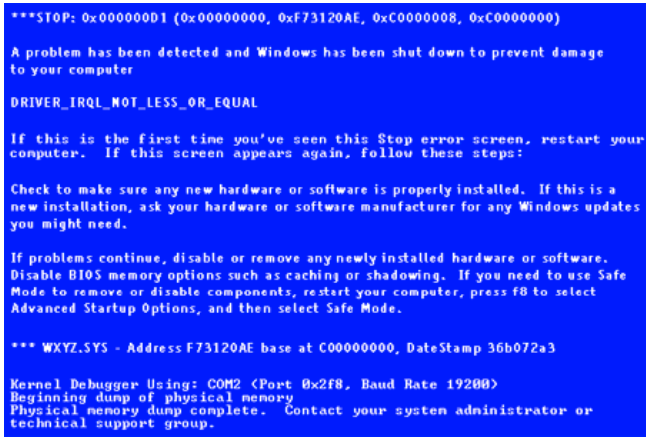

4. Rinse and repeat

Solution

- Use semi random algorithm to select DC's
- Fix Caching holes
- Take the opportunity to select DC's on something more than their ping time.

Problem:

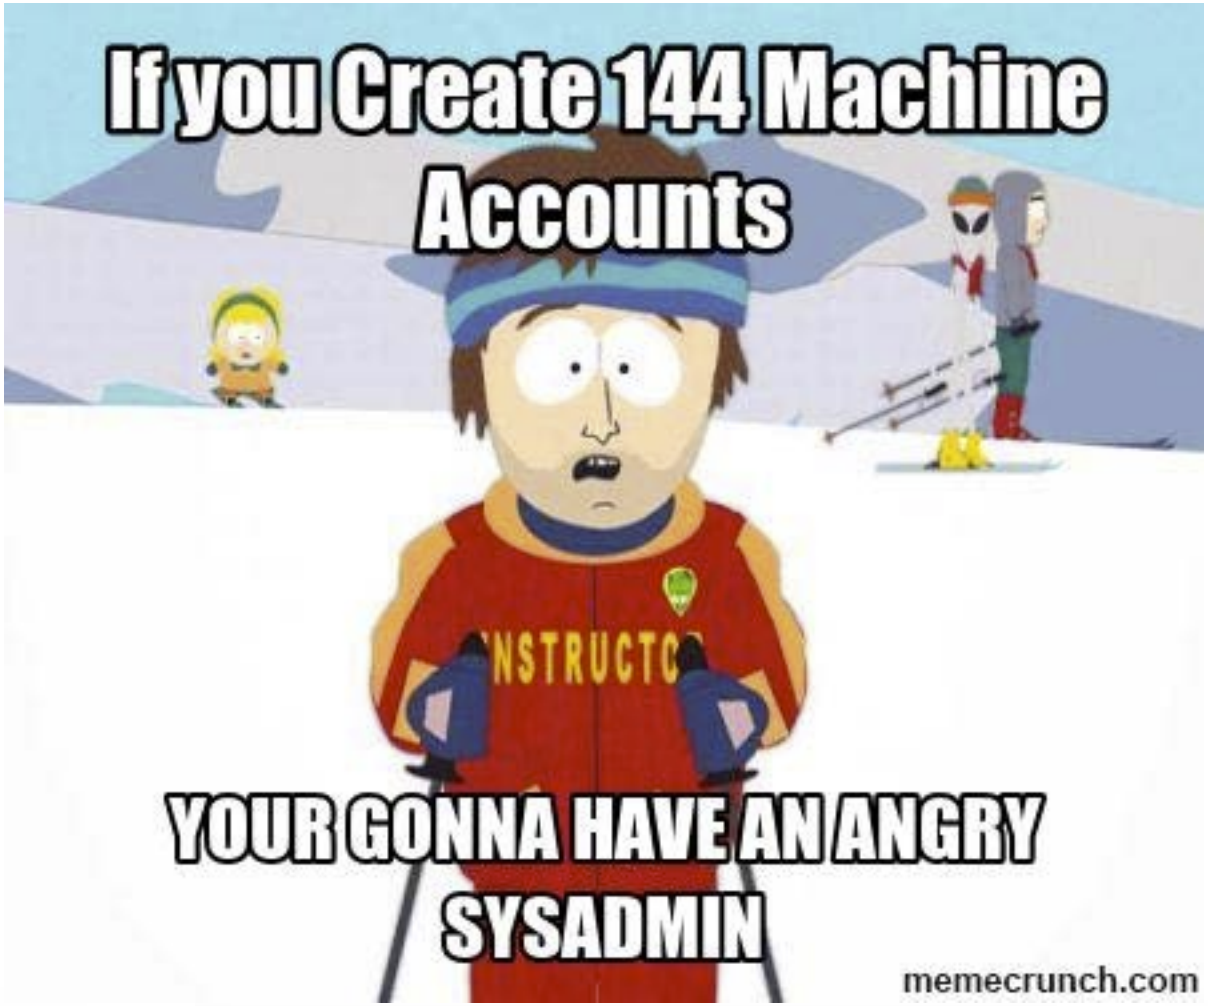

Solution

- Create a shared Machine Account for all nodes

Problem:
We have to share a password

- Remember that con list?:
  - Have to share password across the cluster
    - Who Updates that password?
    - All nodes must affinitize to a single AD when password is changed

Solution

- We are still dealing with machine password issues to this day

Problem:
SMB Signing is broken

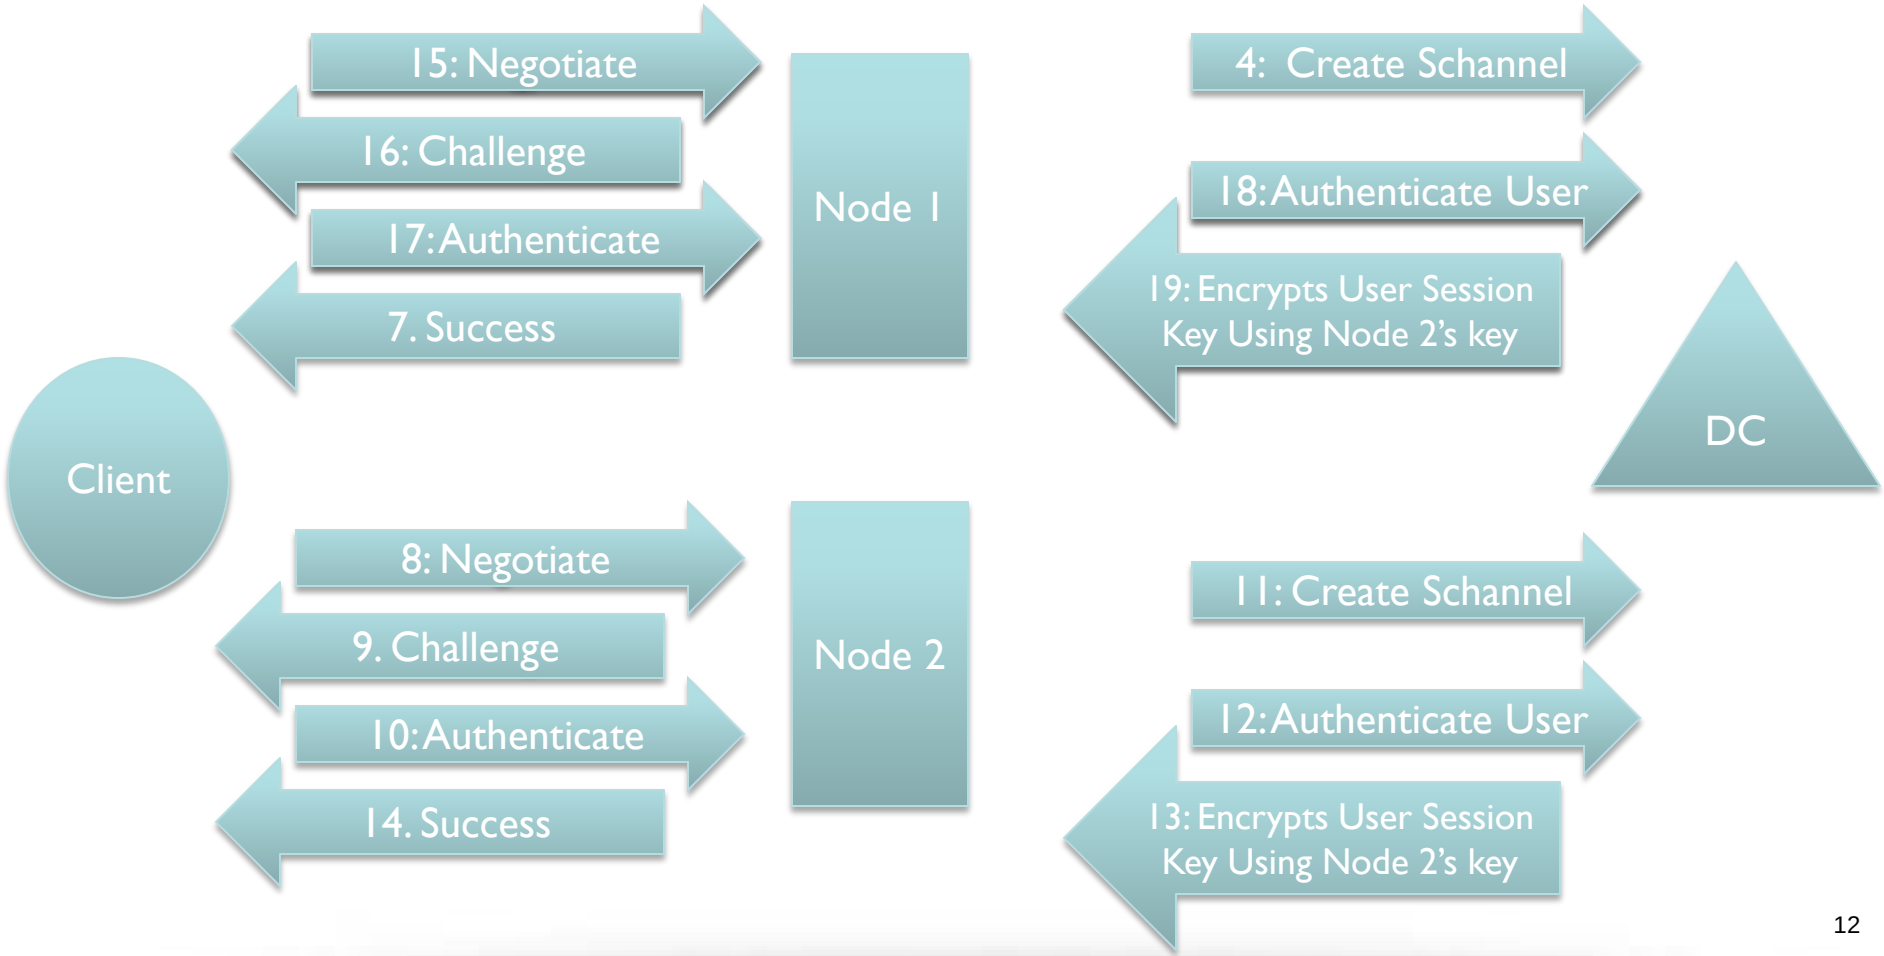

Solution

- The AD Servers were using the computer name from the Schannel creation to store the session key.

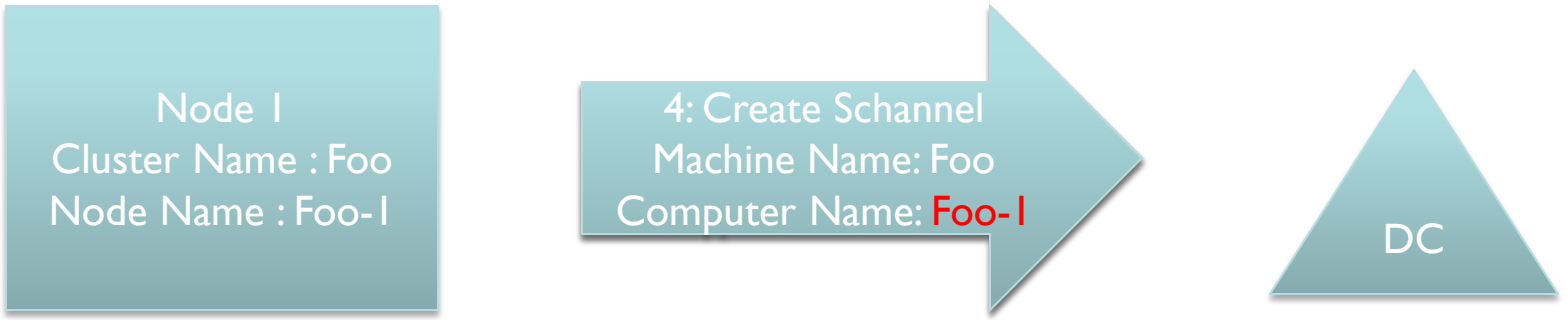

- Send the node name instead.

Problem:
MS15-027 Security Patch

- Patch requires that the “Machine account name” used to create the Schannel and “target name” in the Challenge message must be the same.

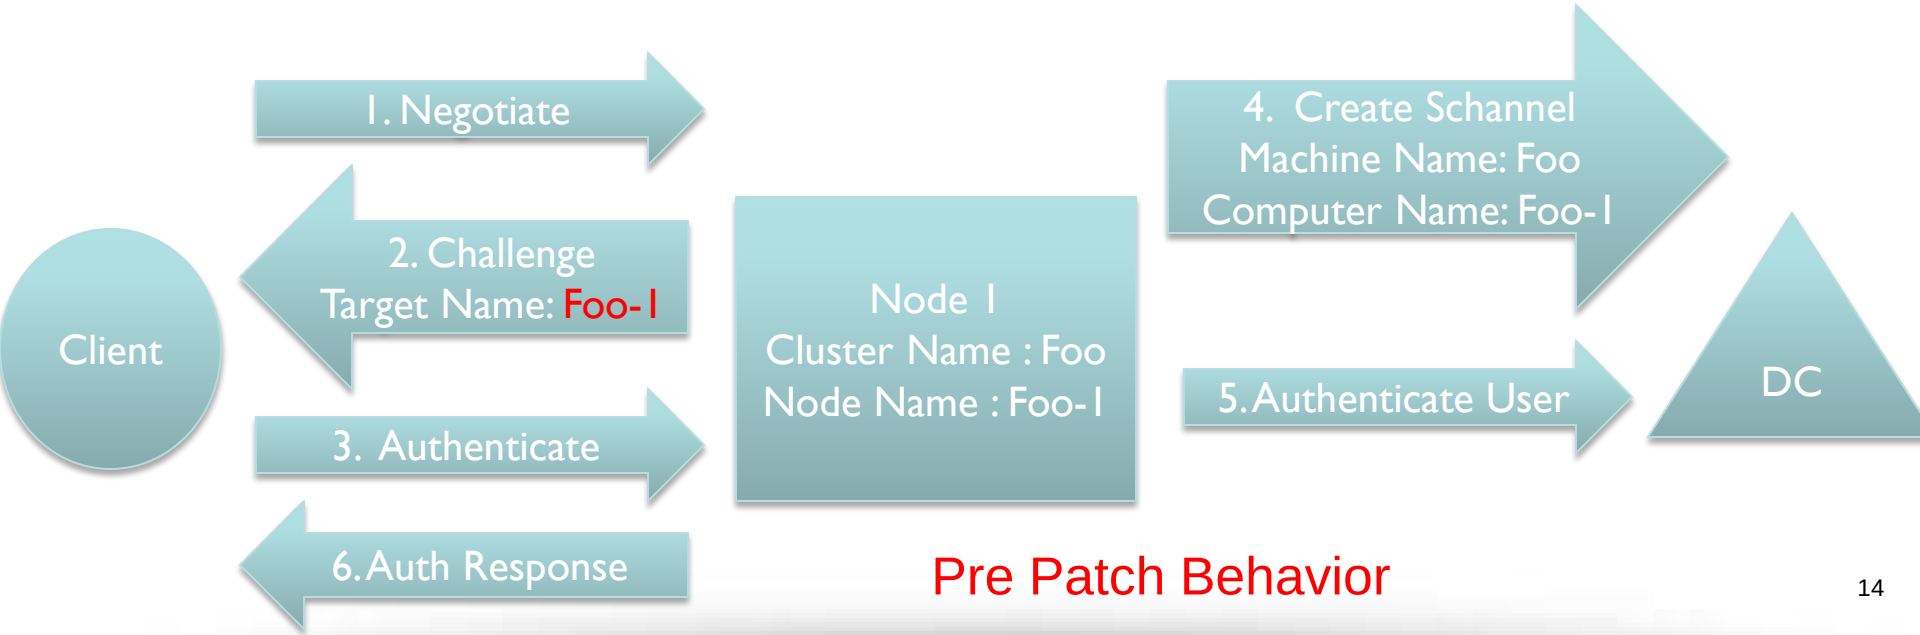

Solution

- Use cluster name as target name

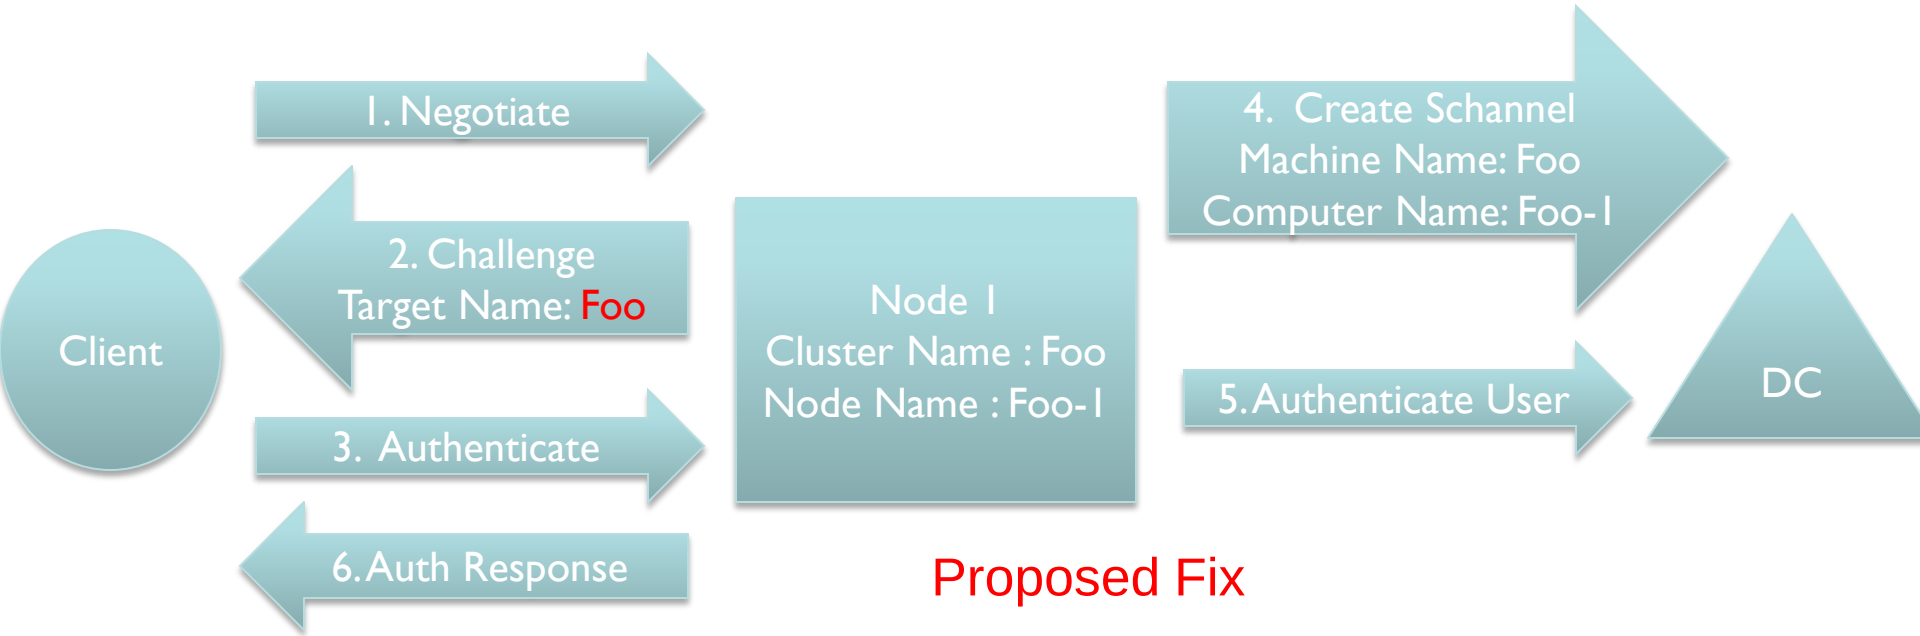

More Problems:
MS15-027 Security Patch

- Pre patch DC's require that "Computer Name" used to create Schannel must match "Target Name"
- #$@&!! We have a mutually exclusive scenario

Solution

- Every time we start using a new DC determine if it is pre or post patch

Questions?