SIPeerior
Technologies
A superior way to connect

Emerging IETF Standards Work on P2PSIP

David A. Bryan

Outline

- What is P2PSIP vs. Client Server SIP?
- Work done to date in the IETF
- What will the new working group do?
- What are the big challenges and questions?

Client/Server Session

In a Client/Server session, two Peers must use a central server to communicate

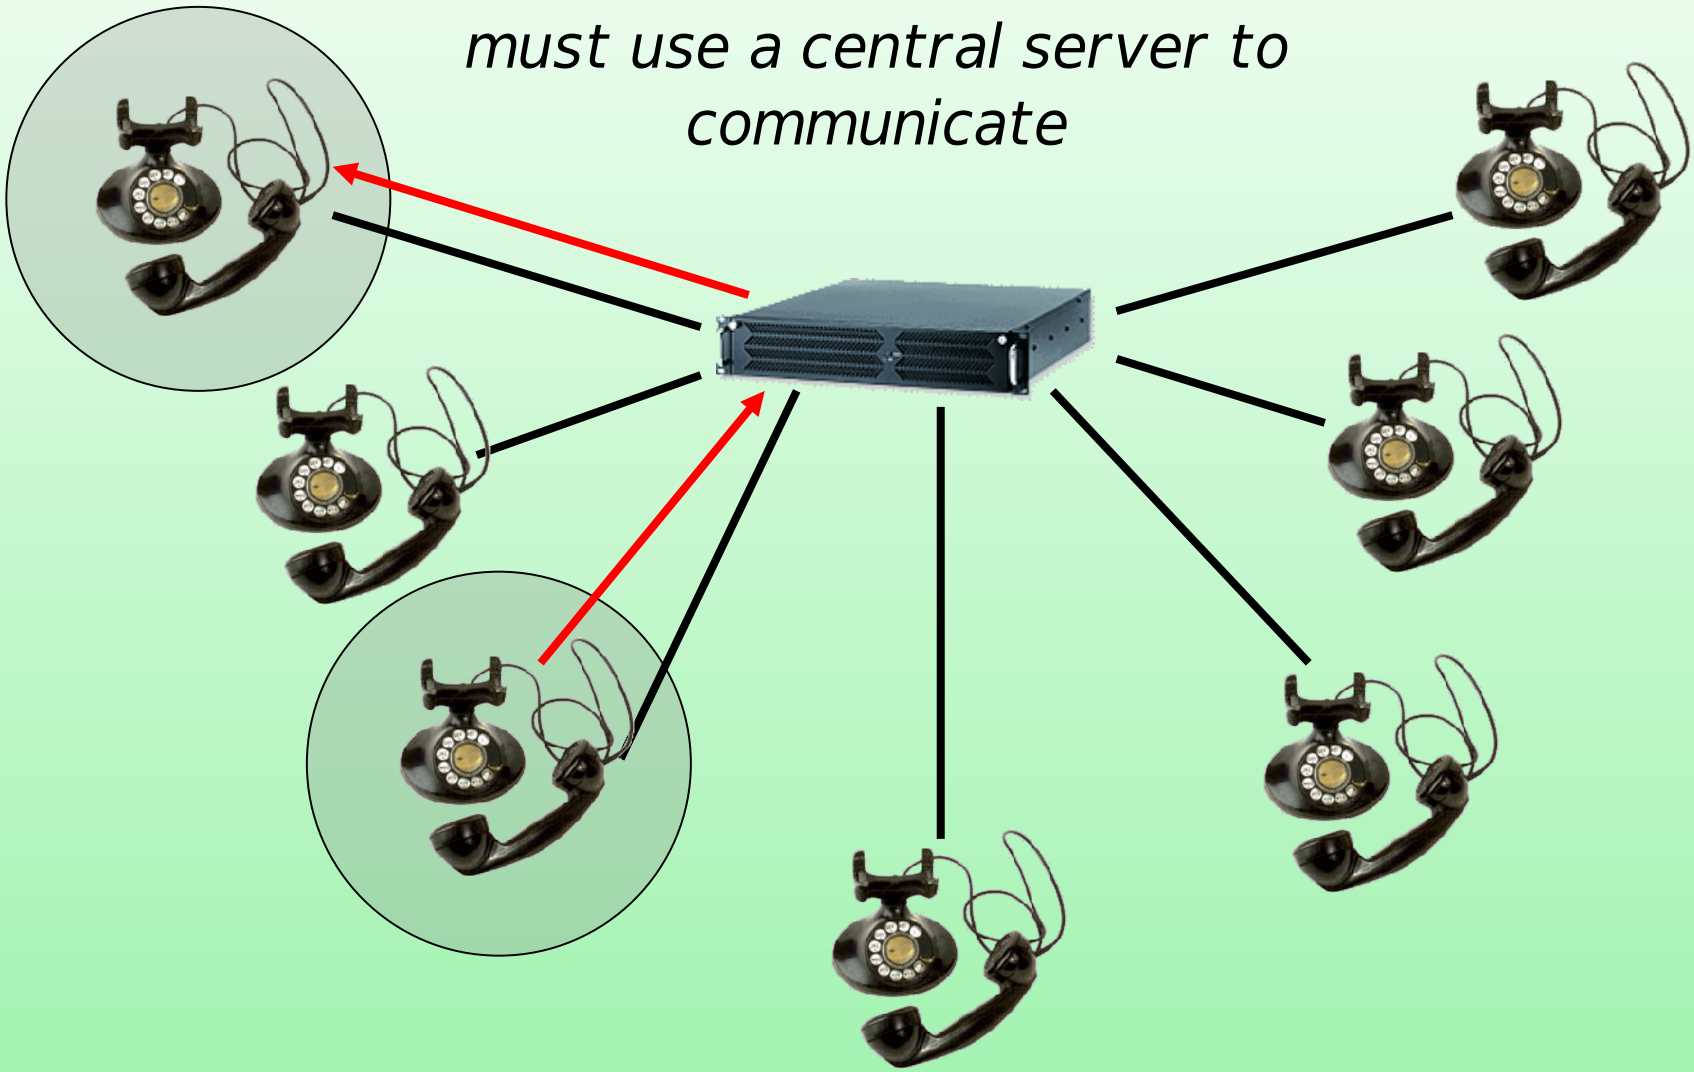

P2P Session

In a Peer-to-Peer session, when two Peers communicate, a few other Peers, rather than a central server, help complete the call

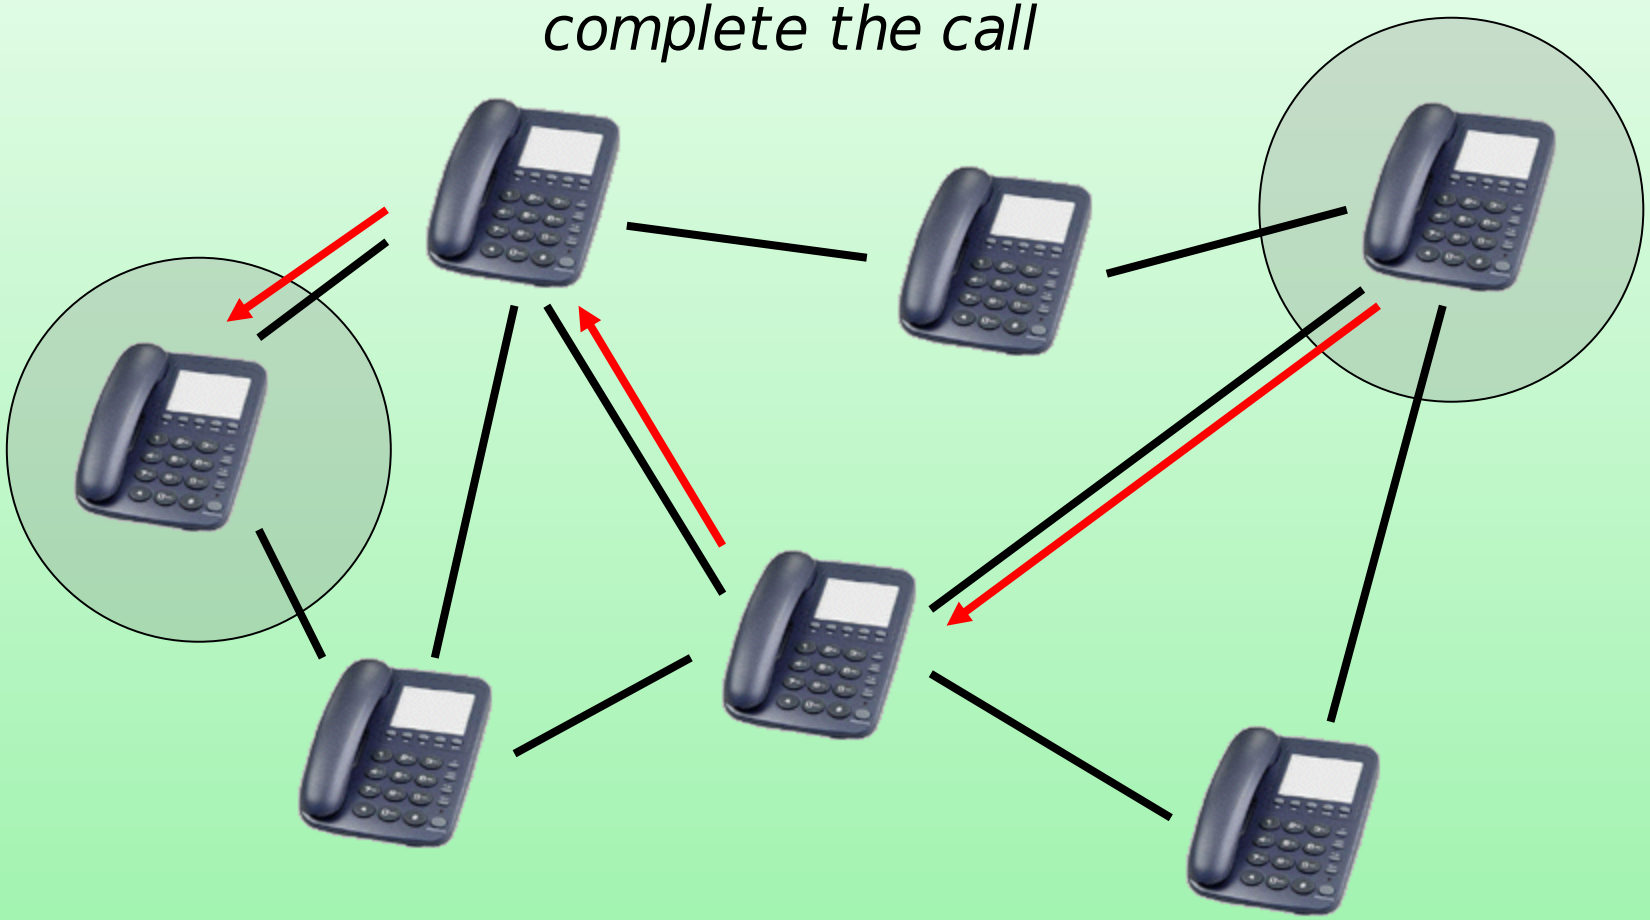

Why use P2P?

- Infrastructure independence
  - No central servers or administration needed
  - Don’t need connectivity to Internet
  - Scalable - new Peers bring more resources!
- Simple discovery and setup
  - Peers find each other
- Privacy
  - No need for information between nearby peers to flow offsite. Hashing of DHT can enhance privacy for some cases.

Why P2PSIP?

- Widely established protocol
  - Standards based, interact with others
  - Compatibility with existing equipment
  - Reuse existing software components
  - Many problems already solved by community
    - Can simply use proven solutions for SIP
    - Support for IM (SIMPLE) and VoIP, video, etc.

Motivating Scenarios

- Why is this important and where can this technology be used?
- The range of uses span the range from the smallest of offices to global user communities

Motivating Scenarios

- Small deployments
  - Security (don't want to use a central provider)
  - Lack of resource (can't run servers)
- Limited or no Internet connectivity
  - Emergency scenarios (turn on WiFi, start endpoints, and go!), remote locations
- Sharing media among portable devices
- Ad-Hoc and ephemeral groups
- Large scale decentralized communications
  - Works with the rest of the (SIP) world

IETF Work to Date

- “Group” has been very busy!
- Have had meetings at the last 6 IETFs
  - First meeting at IETF-62 in March 2005
  - Best attended meeting at IETF-67
- Almost 2 dozen submitted drafts
- ~200 messages/month on mailing list
- Problem is, haven’t been able to move forward, since not an official working group

IETF-67 BoF

- At IETF-67, BoF to discuss forming Working Group (WG)
- Group will be chartered!
  - Should be meeting for first time at IETF-68 in Prague next week!
  - Charter of the groups goal in place
  - Chairs announced (Brian Rosen and myself)
- Finally can start working on official drafts!

What will the WG do?

- Primary (top level) chartered goals:
  - Submit an overview document
  - Create protocol drafts
    - Peer Protocol
    - (Maybe) Client Protocol
  - P2PSIP applicability document

Overview Document

- What is a P2PSIP architecture going to look like?
- What are the design constraints?
- What are the requirements?
- In many ways, this can be thought of as a design and requirements document for the protocols
- Some work already begun on this
  - draft-willis-p2psip-concepts-03
- Sometime in 2007

Protocol Documents

- One or two documents defining the actual protocol
- Here is where the hard questions get answered
  - More on these in a minute
- Technical details of the protocol
- What bits go on the wire?

Applicability Document

- A bit fuzzy today
- How do we use this?
  - What scenarios does this address and how is it used to address them
  - How does this interact/use other protocols and mechanisms (SIP, ICE, various security constructs, configuration) to actually do something with it

Other Documents

- While these look to be the likely primary WG items, will likely be many other drafts in P2PSIP (many individual efforts)
  - Positions on the tough questions
  - Security
  - BCPs or “Best Current Practices” on how to do certain things using the protocols
  - Proposals for the protocols

“Hard Questions”

- There are a number of hard questions for the WG to decide
  - How do we handle NATs, and does this mean we need clients and peers?
  - How do we encode/transport the bits (SIP, XML, something else?)
  - How do we secure this for large deployments while leaving flexibility for smaller ones

How Do We Handle NATs?

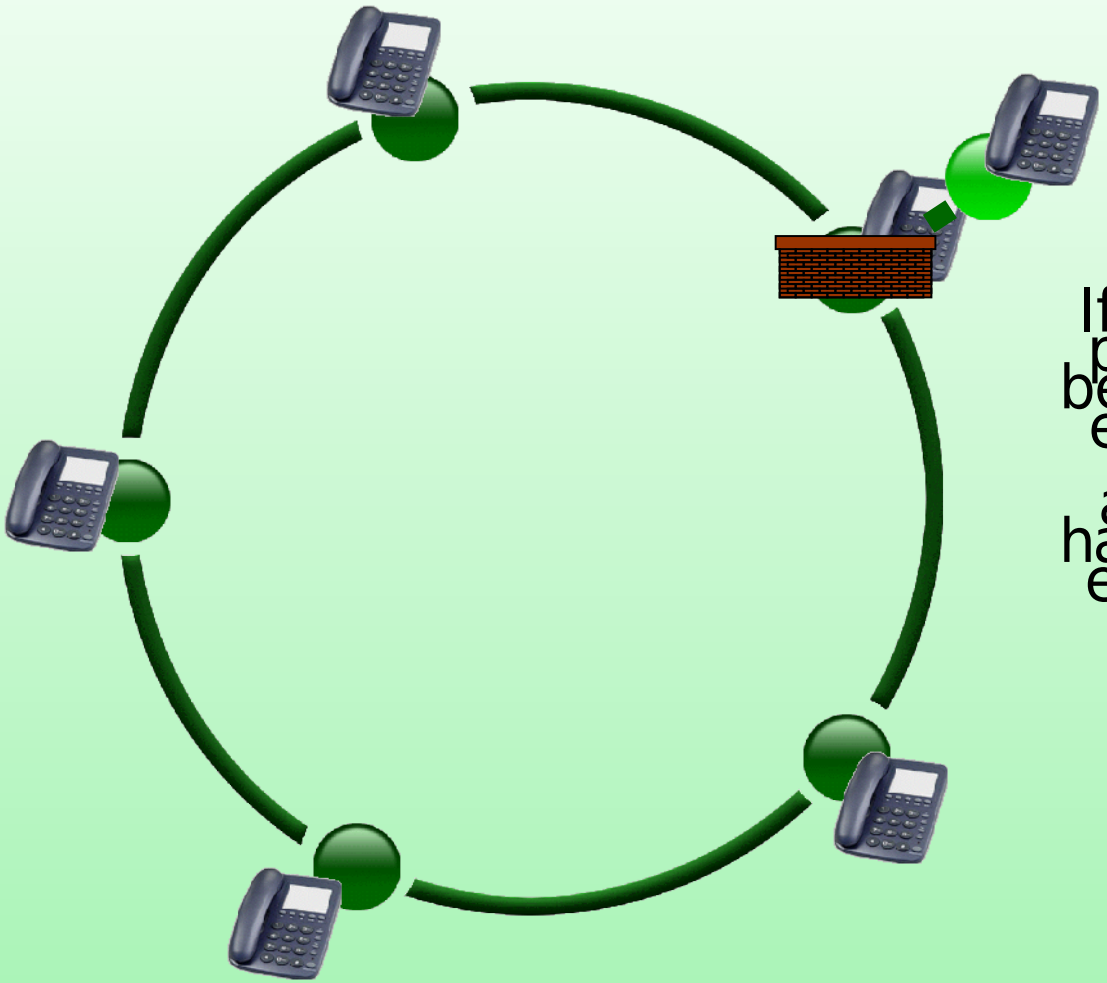

If all peers are in the public Internet, they can easily communicate with all other peers, and each share storing some of the information...

If one or more peers are behind NATs, they may be unreachable, and may have trouble storing some of the information...

Architecture for NATs

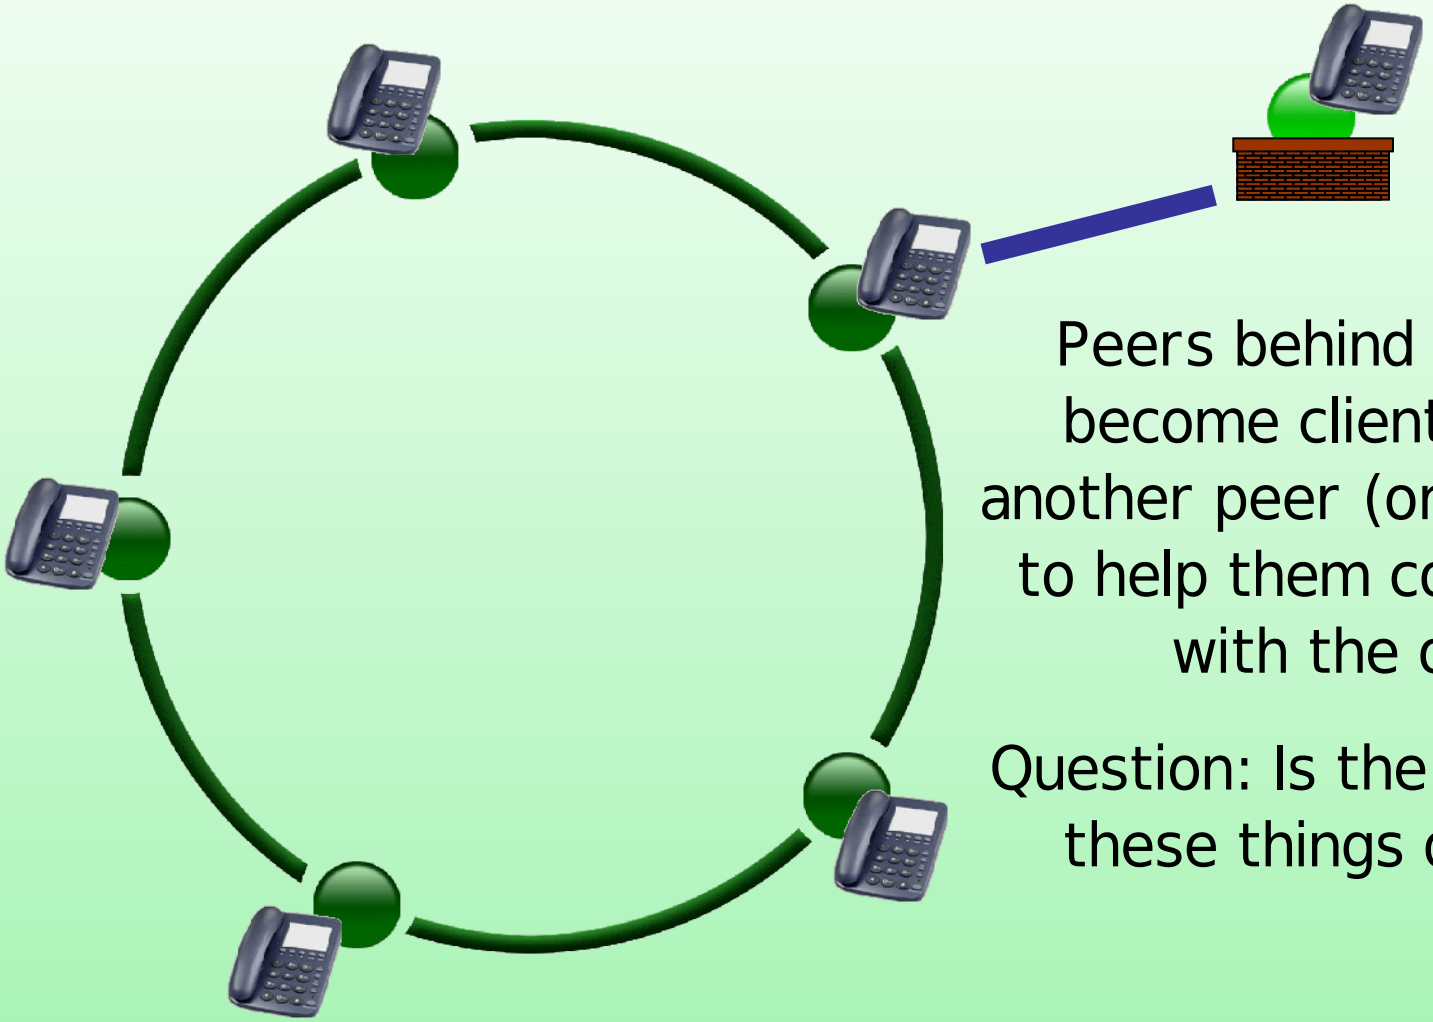

Peers behind NATs may become clients and use another peer (or super-peer) to help them communicate with the others

Question: Is the protocol for these things different?

Peers and Clients

- Some debate about idea of clients and peers
- Most agree that things that are fully participating w/o NATs are peers, clients less clear
- As a result, working group may define two protocols:
  - P2PSIP Peer Protocol: between the peers
  - P2PSIP Client Protocol: between clients and the peers

Peers and Clients

- Clients will likely only retrieve/place information into the overlay
- Clients could be pure SIP endpoints (just a phone) – then P2PSIP Peer Protocol would just be conventional SIP or not exist
- P2PSIP Client Protocol could also be different, and likely a subset of the P2PSIP Peer Protocol
- Crux of the question: 1 or 2 protocols?

What goes on the wire?

- Debate over the format of the messages on the wire
  - SIP messages with special headers
  - SIP messages with XML bodies
  - HTTP messages
  - XML based over SOAP or HTTP
  - Binary format
  - Something else?
- Some believe using SIP improves compatibility, others that separating P2P from signaling makes more sense...

Is any of this SIP?

- Which (if any) of the protocols will be some flavor of SIP? Several options:
  - Clients speak pure SIP, Peers speak modified SIP
  - Clients speak pure SIP, Peers speak something new
  - Clients and Peers speak something new, but clients speak a subset of the new protocol
  - Clients and Peers both speak something new but each is different (not being seriously considered)

Security Problems

- P2PSIP presents very unique security challenges, different than conventional SIP
- One example:
  - Each peer needs a unique ID, or PeerID, which controls where in the cluster of peers it is
  - PeerID also determines which resources are stored by each peer
  - What security challenges does this pose?

PeerID Attacks

- A fundamental assumption of a structured P2P network is that PeerIDs are randomly distributed
- If attackers are able to select PeerIDs, they can mount a variety of attacks
- Two attacks if you can select PeerIDs:
  - Partition the P2P network
  - Block Access Information

Partitioning a Network

- If an attacker can gain control of all routes between two complete, disjoint neighbor sets, they can partition the network

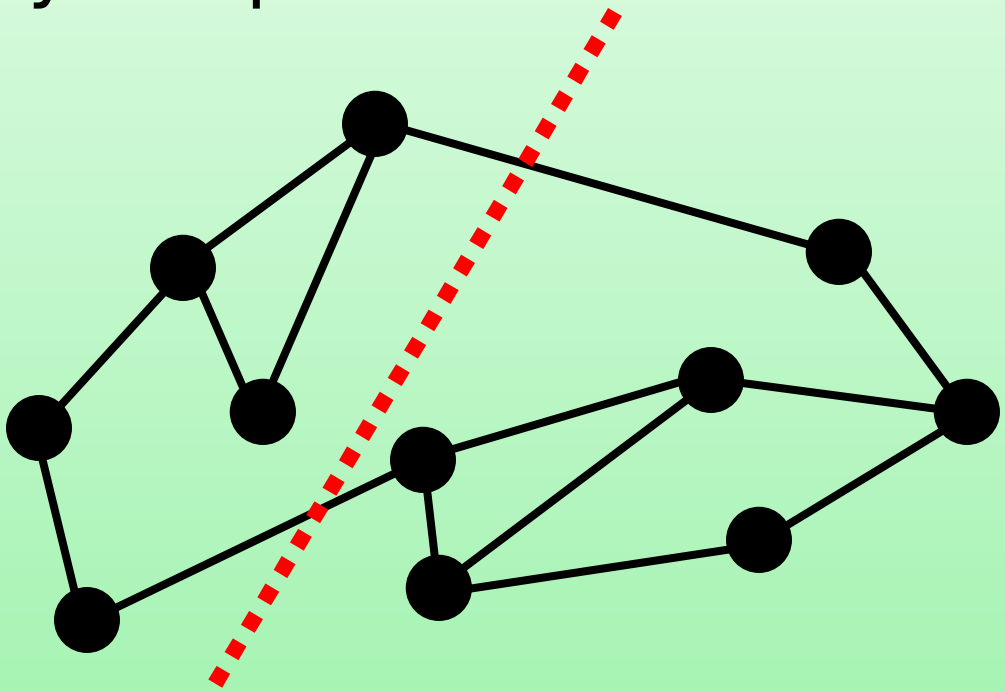

Blocking by PeerID

- An attacker who can insert nodes with particular values, can “censor” data or split a node from the overlay

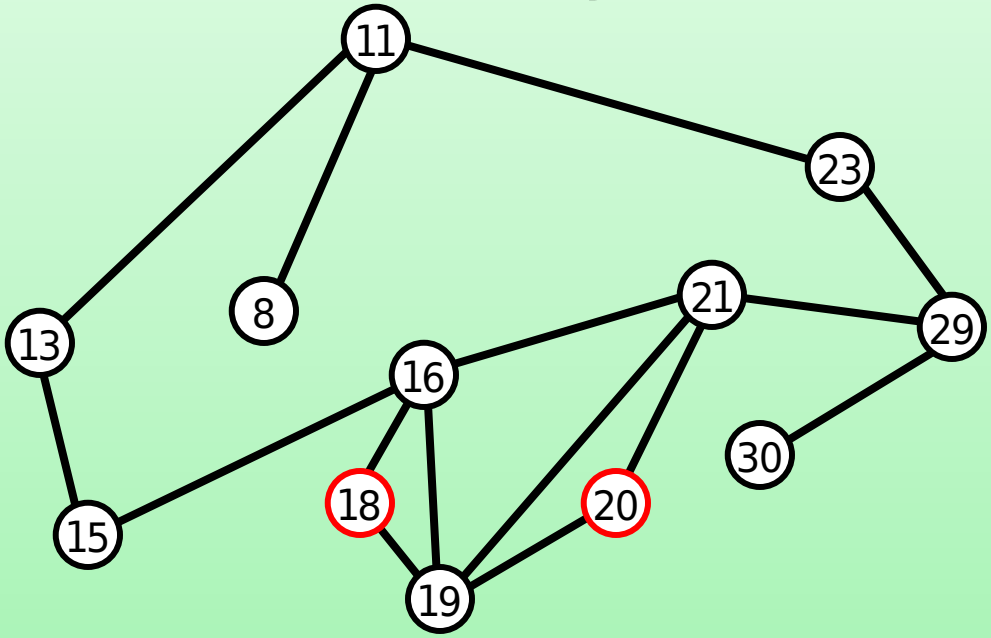

Sybil Attack

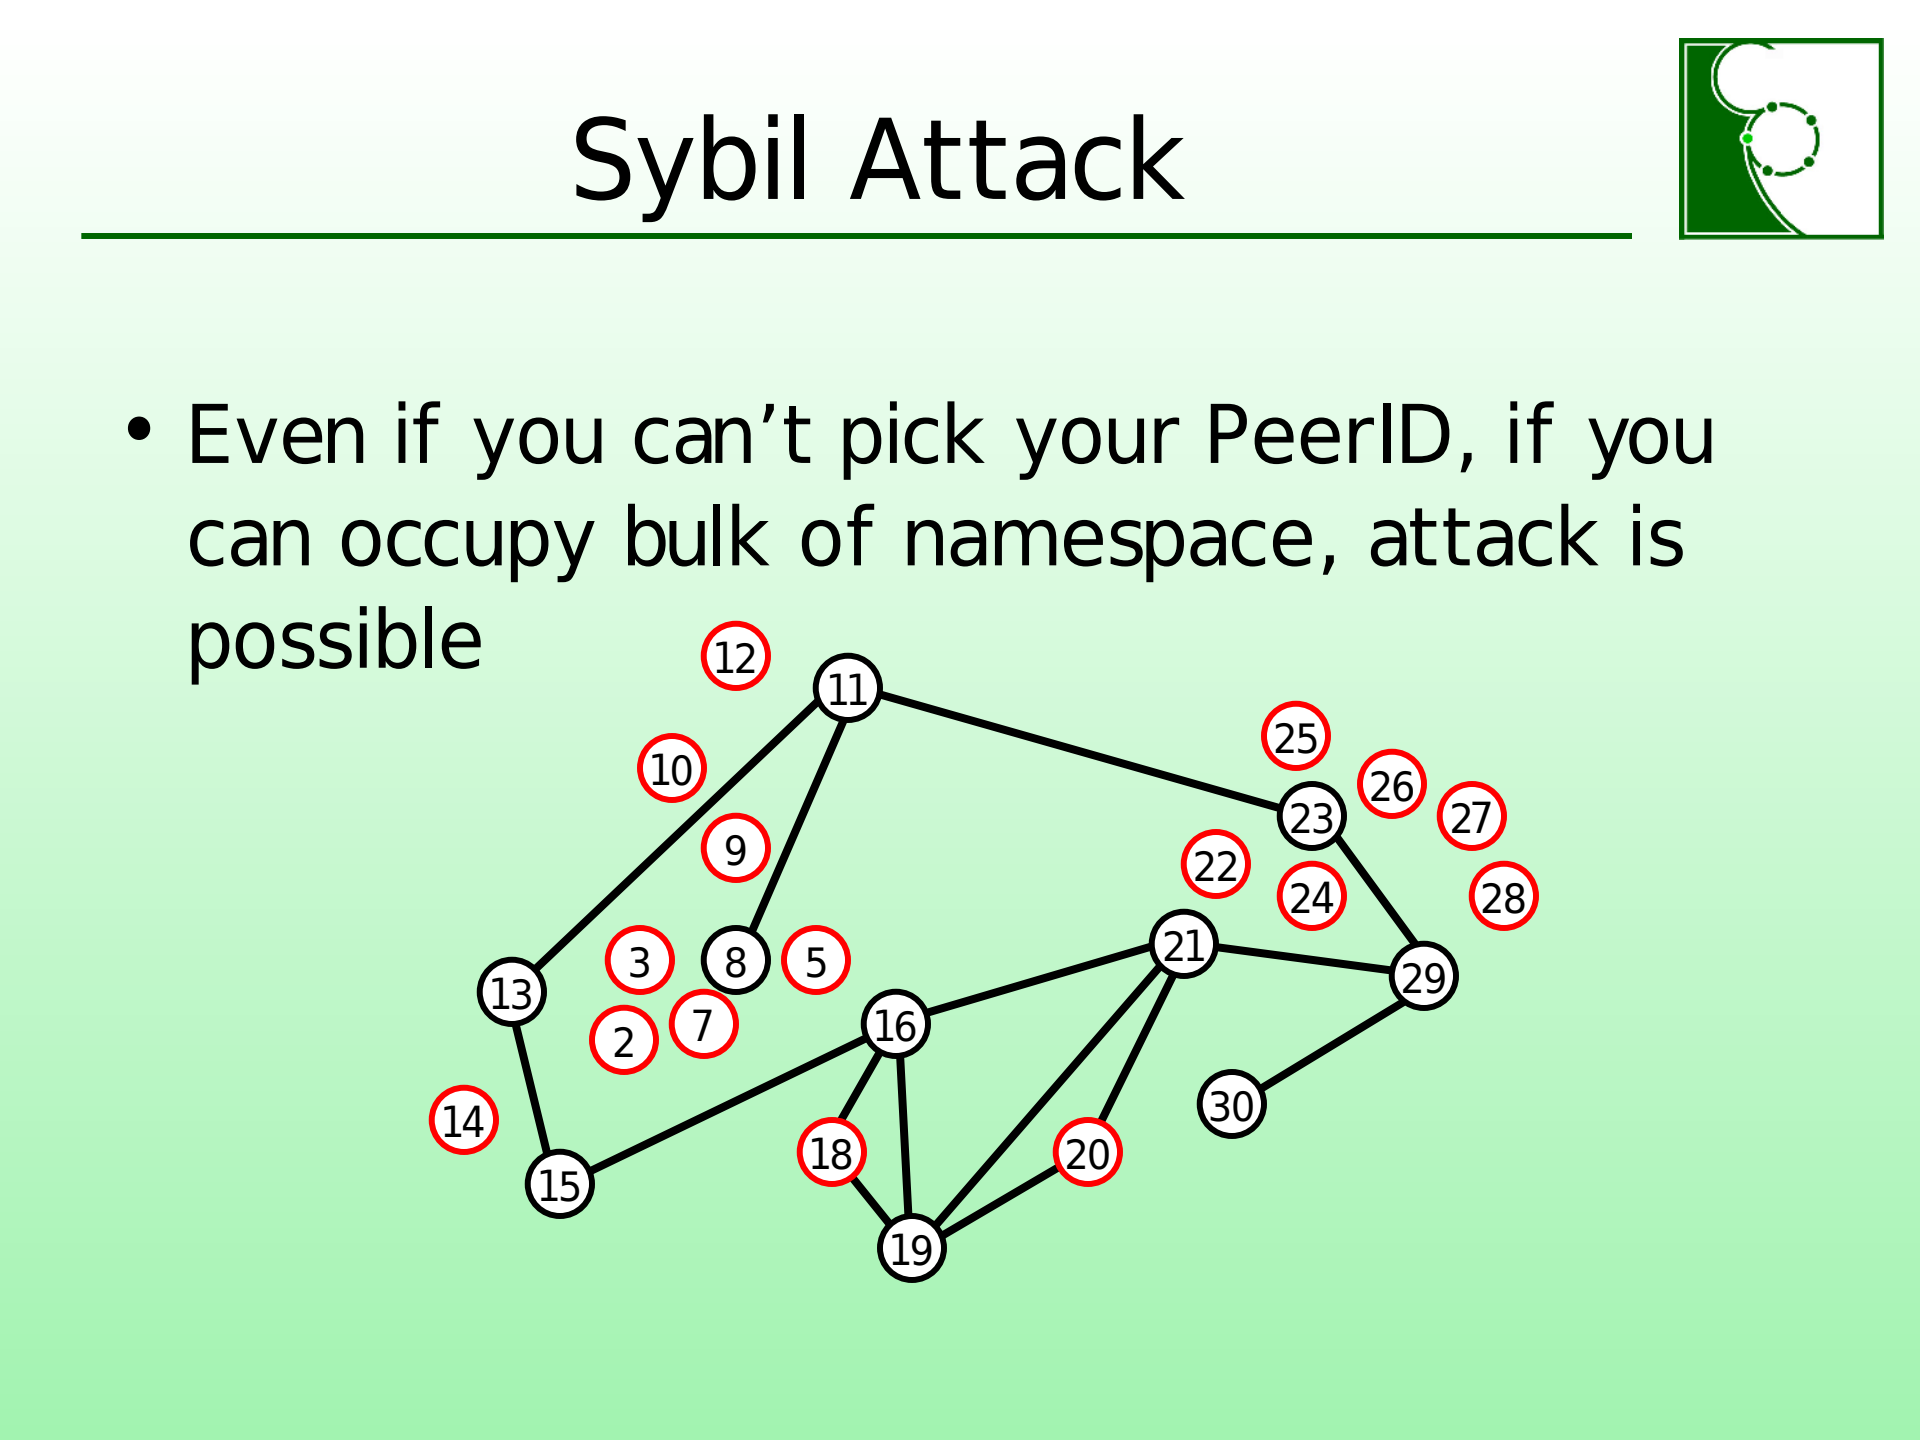

- Even if you can't pick your PeerID, if you can occupy bulk of namespace, attack is possible

How to Secure?

- Most can be solved by a central server that issues certificates, but only contacted at enrollment in the service
- May be very acceptable for a global telecom type system, may not be for an ad-hoc meeting
- Group must balance security with the flexibility P2P provides
  - Want a continuum of choices for different deployments

Conclusion

- P2PSIP is a very interesting area with a great deal of work being done on emerging standard
- MUCH work left to do
- There are many areas people are looking at deploying P2PSIP
  - What "P2PSIP" means to each is different, and all interests need to be balanced in protocol defined.
- P2PSIP is in someways revolutionary, but also evolutionary
- P2PSIP shares the same issue with its predecessors: It much harder to agree on global standards than it is to build a new (isolated) communication system.

References

Here is the main repository for references:

http://p2psip.org

You can link to most other things (IETF work, published papers, mailing lists) from there...

Questions?