Most of the time, evidence of a problem within our networks consists of small amounts of previously-unseen communication, rather than big splashy attacks. This traffic may consist of infected machines reaching back to a command and control network, sending low levels of attack traffic or spam, or finding more vulnerable machines to attack.

How do we find these things? At some point, an administrator gets a NetFlow summary report and sees something odd. At some point, they may dig in a little further; maybe it's an attack, maybe some user is just doing something odd. The next time they're out with a friend, they might ask if they've seen the same thing. Maybe they'll shoot an email to the person nominally in charge of the machine(s). Eventually, the team will figure out what it is, and what to do about it. But it'll be awhile.

Why does this take so long? Because it is low traffic, and because the task of diagnosing the problem requires a great deal of manual effort. In this talk, I'll talk about a tool that tries to assist administrators by automate much of this process. 3DCoP works at all levels of the problem: it scans local flow data for anomalous flows, checks them against known issues, shares determinations about the cause of the traffic across sites, and can even track traffic across reflection / spoofing points. Many of these tasks are explicitly designed to mimic exactly the same steps described above, automatically. The result is that instead of waking up and finding just a suspicious flow in an email, they wake up and learn about the flow, its possible causes and mitigations, and who else amongst their peers is seeing the same thing. There may even be some suggestions about what to do about it, ready and waiting.