

MIT 6.858 - Computer Systems Security “Lecture 5 - Privilege Separation”

Lecture 5 - Privilege Separation

Privilege separation is the practice of dividing your software into multiple independent services which each run in isolation. This way one service being compromised minimizes the chance of the attacker being able to compromise another service. It also reduces the attack surface.

It is easy to describe the idea of privilege separation, but hard to actually implement it in practice. It is similar to how it is always recommended to make your software modular, but writing modular code is easier said than done.

This week’s paper describes the working of a web server - OKWS - which is written with the principle of privilege separation at its core.