## DRAM cell re-mapping and cell fencing

Enterprise memory module supply chain fraud leverage both non-factory performed cell re-mapping and cell fencing to pass off defective, degraded, or scavenged memory as premium, enterprise-grade hardware. By exploiting the gap between what the hardware reports and what physically exists, threat actors maximize their profit margins on counterfeit components.

Re-mapped cells are seamlessly bypassed at the hardware level using internal spare resources to fix errors, while fenced cells (like "BadRAM" or page retirement) leave the damaged physical area offline and force the operating system to map around the damaged regions.

1. **Re-mapping (Spare Substitution )**
   **How it works:** When a cell or row begins to fail, the memory controller or DRAM logic detects it and silently swaps the damaged physical location with a factory-reserved spare cell or row.
   a) **Result:** The total available memory capacity remains unchanged. The address space appears completely healthy to the operating system, requiring no software intervention.
   b) **Usage:** Common in advanced enterprise GPUs (e.g., NVIDIA Row Remapping) and certain server-grade processors.

2. **Fenced / Retired Cells (Software Avoidance)**
   **How it works:** The operating system or memory controller receives an error report (like ECC errors or diagnostic readouts). Instead of fixing it physically, the OS blocks out those specific Page Frame Numbers (PFN) via kernel parameters or memory maps (e.g., Windows {badmemorylist} or Linux BadRAM).
   a) **Result:** The system permanently hides those exact cells from the OS and applications, reducing your total usable RAM. If software touches a fenced area, it is forcefully redirected.
   b) **Usage:** Not common with enterprise modules, primarily used in consumer applications when physical memory chips develop errors and the user chooses to keep the module functional instead of throwing it away.