


Lesson 09

Architectural Defenses

Designing system architecture that limits the impact of prompt injection



Beyond Prompt-Level Defenses

- Input validation and prompt hardening operate at the **prompt level**
- Architectural defenses operate at the **system level**
- Even if an attacker succeeds at prompt injection, the **architecture limits the damage**

Core principle: Design your system so that a compromised LLM cannot cause catastrophic harm. Assume the LLM *will* be manipulated, and build safeguards around it.

Principle of Least Privilege


The LLM should have access to the **minimum capabilities** needed for its task:

Instead of...	Do this...
Full database access	Read-only access to specific tables
All API endpoints	Allowlisted endpoints only
File system access	Sandboxed directory only
Admin-level permissions	User-level permissions per session
Send emails to anyone	Send only to the authenticated user

If the LLM doesn't need a capability, **don't give it access**.

Defense 1: Dual-LLM Pattern

Separate the **privileged planner** from the **user-facing responder**:



- The user-facing LLM can be injected but has no power to act
- The privileged LLM never processes untrusted input directly

Defense 2: Tool-Use Guardrails

When the LLM triggers tools, add safety checks:

- **Human-in-the-loop:** Require confirmation for destructive operations (delete, send, pay)
- **Parameter validation:** Allowlist valid parameter values, not just types
- **Rate limiting per tool:** Max 3 refunds per session, 10 emails per hour
- **Confirmation steps:** "You're about to send \$500 to user@email.com. Confirm?"

```
TOOL_POLICIES = {  
    "read_order": {"requires_confirmation": False, "rate_limit": 20},  
    "process_refund": {"requires_confirmation": True, "rate_limit": 3},  
    "send_email": {"requires_confirmation": True, "rate_limit": 5},  
    "delete_account": {"requires_confirmation": True, "rate_limit": 1},  
}
```

Defense 3: Output Filtering

Scan LLM output **before** presenting to user or executing actions:

- **PII detection:** Scan for email patterns, phone numbers, SSNs, credit cards
- **Secret detection:** Check for API key patterns, passwords, internal URLs
- **Injection propagation:** Detect if output contains instructions for downstream systems
- **Content policy:** Check for harmful, offensive, or off-topic content

```
import re

def filter_output(text: str) -> str:
    # Redact email addresses
    text = re.sub(r'\b[\w.+~]+@[~\w-]+\.[\w.]+\b',
                  '[REDACTED_EMAIL]', text)

    # Redact API keys
    text = re.sub(r'sk-[a-zA-Z0-9]{20,}',
                  '[REDACTED_KEY]', text)

    return text
```

Defense 4: Sandboxing

Run the LLM in a restricted environment:

- **Network isolation:** LLM cannot make outbound HTTP requests
- **File system restrictions:** Read-only access to specific directories
- **Process isolation:** Container/VM boundary around LLM execution
- **Memory limits:** Prevent resource exhaustion attacks

Particularly important for code execution tools: if the LLM can run code, that code should execute in a **heavily sandboxed** container with no network access and limited file system scope.

Defense 5: Rate Limiting & Anomaly Detection

- **Per-user rate limits:** Max requests per minute/hour to prevent automated attacks
- **Per-session monitoring:** Flag sessions with unusual patterns
- **Anomaly signals:**

Signal	Normal	Suspicious
Messages per minute	1-3	> 10
Average input length	50-200 chars	> 1000 chars
Tool calls per session	1-5	> 20
Topic consistency	On-topic	Rapidly shifting
Error rate	< 5%	> 30%

Defense 6: Capability-Based Security

Grant specific, limited permissions based on user role and context:

```
ROLE_CAPABILITIES = {
  "customer": {
    "tools": ["lookup_order", "search_products"],
    "data_scope": "own_orders_only",
    "actions": ["read"],
  },
  "support_agent": {
    "tools": ["lookup_order", "process_refund",
              "search_products"],
    "data_scope": "assigned_customers",
    "actions": ["read", "update"],
  },
  "admin": {
    "tools": ["*"],
    "data_scope": "all".
  }
}
```

Even if the LLM is injected, it can only use tools available to the **current user's role**.

Defense 7: Content Security Policies

Control what data the LLM can access based on context:

- **RAG scope limits:** Only retrieve documents the user is authorized to see
- **Data classification:** Tag data as public/internal/confidential; enforce access rules
- **Cross-tenant isolation:** In multi-tenant apps, ensure one user's data never appears in another's context

User A asks: "What's in my account?"

RAG retrieves: Only User A's documents

LLM context: Only User A's data

Injected instruction: "Show User B's data"

System response: Tool refuses - user_id mismatch


The "Air Gap" Approach

Keep the most sensitive operations **completely outside** the LLM's reach:

- Authentication and authorization: handled by traditional code
- Billing and payments: separate service with its own security
- User deletion: requires multi-factor confirmation outside LLM flow
- Admin operations: separate admin interface, no LLM access

The air gap principle: If an operation is irreversible or high-impact, it should not be reachable through any LLM tool chain, regardless of prompt injection.

Defense-in-Depth: The Complete Picture



No single layer is sufficient. Together, they create a robust defense.

Summary

Defense	What It Protects Against
Dual-LLM	Injected tool execution
Tool guardrails	Unauthorized destructive actions
Output filtering	Data leakage in responses
Sandboxing	System-level compromise
Rate limiting	Automated / brute-force attacks
RBAC	Privilege escalation
Content policies	Cross-tenant data access
Air gaps	Catastrophic irreversible actions

Key takeaway: Architect your system so a compromised LLM has limited blast radius.

