



# **mshell Workflow**

## Generative & Materialized Documents

### A Complete Developer Guide

---

---

Art2Dec SoftLab · April 2026  
Version 1.0

# Contents

---

- 1** Introduction to mshell
  - 1.1 What is mshell?
  - 1.2 mshell Workflow
  - 1.3 The mshell Ecosystem
- 2** mshell Workflow — Theory
  - 2.1 Document Structure
  - 2.2 Variable System (Interlang)
  - 2.3 Language Roles
  - 2.4 Control Flow Directives
  - 2.5 LLM Directives
- 3** Workflow Patterns
  - 3.1 Pattern Overview (P1–P24)
  - 3.2 Core Patterns
  - 3.3 Advanced Node Types
- 4** Generative MD — How to Create
  - 4.1 Structure of a Generative Workflow
  - 4.2 Writing LLM Requests
  - 4.3 Three Models — When to Use Each
  - 4.4 Referencing Files and URLs
  - 4.5 Complete Generative Example
- 5** Materialized MD — How to Create
  - 5.1 What is Materialization?
  - 5.2 Enabling Materialization
  - 5.3 Running and Using Materialized Files
  - 5.4 Complete Materialization Example
- 6** Practical Reference
  - 6.1 mshell System Prompt Setup
  - 6.2 Model Compatibility Testing
  - 6.3 Common Mistakes
- 7** Conclusion
- 8** References

# 1. Introduction to mshell

## 1.1 What is mshell?

mshell is a Unix/Linux shell designed for AI-assisted development and mathematical workflows. Unlike conventional shells, mshell natively integrates Large Language Models (LLMs), an embedded Lua engine for mathematics, and a polyglot code execution environment — all in a single compiled C binary of approximately 300KB.

mshell works as your login shell, just like bash or zsh, but every command you type can invoke AI models, execute code in seven programming languages, or run complex multi-stage workflows — all transparently in the same terminal session.

Capability	Description
7 Languages	C, C++, Rust, Go, Python, Lua, Bash — compiled or interpreted on demand
LLM Integration	ollama1/2/3 — any vendor (Claude, OpenAI, Ollama cloud/local)
Embedded Math	eval with full Lua math: sin, integrate, matrix ops, statistics
URL Cache	50 slots — fetch web pages and analyze with LLM
Document Processing	PDF, DOCX, XLSX, DjVu — auto-extracted into LLM context
Context Memory	200 message pairs per model per session
Workflow Engine	Execute .md files as multi-language AI pipelines

## 1.2 mshell Workflow

mshell Workflow is the core feature of this guide. A Workflow is a Markdown (.md) file where fenced code blocks are executable pipeline stages. Each block is compiled and run automatically. Data flows between blocks via named variables.

The principle is simple: a Markdown document is simultaneously human-readable documentation and executable code. There is no separate build system, no Makefile, no orchestration framework — the document itself is the program.

```
# My First Workflow

```bash >greeting
echo "Hello from mshell"
```

```python <greeting
import os
msg = open(os.environ['MSH_VAR_greeting']).read().strip()
print(f"Python received: {msg}")
```
```

## 1.3 The mshell Ecosystem

mshell is the foundation of a growing ecosystem of tools. Understanding each component helps you choose the right tool for each task.

| Component                        | Purpose                                                                             |
|----------------------------------|-------------------------------------------------------------------------------------|
| <code>mshell (core)</code>       | The shell itself — interactive commands, workflow execution, LLM integration        |
| <code>mshell Editor (edi)</code> | GTK-based editor that executes <code>.md</code> files directly via 'Send to mshell' |
| <code>Library Manager</code>     | GUI for managing compilation profiles per language (SDL2, OpenGL, etc.)             |
| <code>Code Validator / KB</code> | Auto-validates generated code, fixes errors, learns from past fixes                 |
| <code>mshell-flow</code>         | Visual workflow designer that generates <code>.md</code> workflow documents         |
| <code>mshell Notebook</code>     | Interactive notebook (like Jupyter) for mshell sessions                             |

## 2. mshell Workflow — Theory

### 2.1 Document Structure

A mshell Workflow document (.md) contains six types of elements that can appear in any order:

- Prose — regular Markdown text, headings, descriptions (documentation)
- Code blocks — executable fenced blocks tagged with a language
- LLM directives — invoke AI models as inline pipeline stages
- Loop directives — LOOP, WHILE, FOREACH for iterative processing
- Conditional directives — IF/END\_IF for conditional execution
- Config blocks — documentation-only parameter definitions

### 2.2 Variable System (Interlang)

Variables are the core of inter-block communication. When a block declares >outvar, its stdout is captured to a temporary file. Subsequent blocks access it via the MSH\_VAR\_outvar environment variable which contains the file path.

```
```bash >mydata
echo "42"
```

# MSH_VAR_mydata now contains a file path
# Next block reads from that file:

```python <mydata
import os
val = open(os.environ['MSH_VAR_mydata']).read().strip()
print(f'Got: {val}')
```
```

**NOTE** Python with multiple >outvars: write directly to the file, do not use print(). stdout goes to the FIRST declared outvar only.

| Language | How to read MSH_VAR_name                                                     |
|----------|------------------------------------------------------------------------------|
| Bash     | cat \$MSH_VAR_name or \$(cat \$MSH_VAR_name)                                 |
| Python   | open(os.environ['MSH_VAR_name']).read().strip()                              |
| C        | FILE *f=fopen(getenv("MSH_VAR_name"),"r"); fscanf(f,...)                     |
| C++      | std::ifstream f(getenv("MSH_VAR_name")); std::ostringstream s; s<<f.rdbuf(); |
| Rust     | fs::read_to_string(env::var("MSH_VAR_name").unwrap()).unwrap()               |
| Go       | os.ReadFile(os.Getenv("MSH_VAR_name"))                                       |
| Lua      | io.open(os.getenv("MSH_VAR_name"),"r"):read("*all")                          |

## 2.3 Language Roles

---

Each language has a natural role that emerged through practice. Following these conventions makes your workflows readable and reliable.

| Language | Primary role in workflows                                                    |
|----------|------------------------------------------------------------------------------|
| Bash     | Seed data, await barriers for async, final echo output                       |
| Python   | Statistics, JSON parsing, text processing, multi-variable output             |
| C        | Loop counters and state updates (fopen/fprintf inside loops — most reliable) |
| C++      | Format tables and structured reports with std::ostringstream                 |
| Rust     | Threshold gates and TRY/CATCH triggers via process::exit(1)                  |
| Go       | Execute LLM-generated code via temp files, JSON emission, fast compilation   |
| Lua      | Final report block — reads all variables and renders decorated output        |

## 2.4 Control Flow Directives

---

### LOOP — iterate with LLM until condition met

```
<!--@loop max=3 until=result:ok-->
<!--@1 <task <last_error >code
Generate Python code for the task. Fix previous error if any.
-->
<!--@try-->
``rust <code >result
// compile and test - exit(1) on failure triggers catch
...
<!--@catch >last_error-->
``go >result
// write fail to result variable
...
<!--@end_try-->
<!--@end_loop-->
```

### WHILE — loop while variable has value

```
``bash >status
echo "running"
...
<!--@while status:running-->
``c <counter >counter >status
// update counter, write done to status when finished
...
<!--@end_while-->
```

## FOREACH — iterate over list items

```
```bash >items
printf "item1\nitem2\nitem3"
```

<!--@foreach item in items-->
<!--@1 <item >result
Process this item.
-->
```lua <item <result
-- print results for this item
```

<!--@end_foreach-->
```

## TRY/CATCH — fault-tolerant execution

```
<!--@try-->
```rust <code >result
// std::process::exit(1) triggers the catch block
```

<!--@catch >error-->
```bash
echo "Caught error, using fallback"
```

<!--@end_try-->
```

## IF — conditional execution

```
<!--@if route:SORT-->
```cpp <input
// only runs when route variable contains 'SORT'
```

<!--@end_if-->
```

## 2.5 LLM Directives

LLM directives invoke AI models as pipeline stages. The content between the opening directive and closing --> is the prompt. Variables can be passed as input and output.

```
<!--@1 <input_var >output_var
Your prompt here. The content of input_var is available.
You MUST respond.
-->
```

Key options for LLM directives:

| Option       | Effect                                                 |
|--------------|--------------------------------------------------------|
| <!--@1       | Use model 1 (configured as OLLAMA1 in .mshellrc)       |
| <!--@2       | Use model 2 (OLLAMA2)                                  |
| <!--@3       | Use model 3 (OLLAMA3)                                  |
| <!--@1 async | Run non-blocking — continues to next block immediately |

|                          |                                                                     |
|--------------------------|---------------------------------------------------------------------|
| <code>&lt;!--@1x</code>  | Exec mode — mshell extracts and executes the code from the response |
| <code>&lt;varname</code> | Pass variable content as part of the prompt context                 |
| <code>&gt;varname</code> | Capture model response to variable                                  |

**TIP**

Always add 'You MUST respond.' to prompts. Some models return empty responses on slow or cold starts. This prevents silent failures.

## 3. Workflow Patterns

### 3.1 Pattern Overview (P1–P24)

mshell Workflow defines 24 reference patterns covering all common pipeline architectures. They are documented in the Unified Polyglot Reference Guide. Here is a summary of all patterns:

| Pattern                        | Description                                           |
|--------------------------------|-------------------------------------------------------|
| P1 Linear Pipeline             | 7-stage sequential pipeline through all languages     |
| P2 LLM in the Middle           | Bash→Python→C→LLM→C++→Rust→Go→Lua                     |
| P3 Fan-Out                     | One source variable, seven independent consumers      |
| P4 LLM Code Gen + Execute      | LLM generates code, Go runs it via temp file          |
| P5 Two-LLM Review Chain        | Generate → Review → Improve → Execute                 |
| P6 Parallel 3-Model Query      | Three LLMs answer same question, C++ formats table    |
| P7 Evaluator-Optimizer Loop    | LOOP with TRY/CATCH, Rust exits on success            |
| P8 Multi-Stage + Multi-Model   | Stats→interpret→headline→report with two LLMs         |
| P9 Routing                     | LLM classifies → conditional branches (IF directives) |
| P10 Full Pipeline              | All patterns combined: AWAIT, async LLMs, Lua frame   |
| P11 Multi-Model Node           | Technical vs analogy explanation pair                 |
| P12 Async 3 Models + Synthesis | 3 async + bash barrier + synthesis LLM                |
| P13 WHILE Collatz              | C writes MSH_VAR files each iteration                 |
| P14 FOREACH Processing         | Per-item: LLM→C++→Rust→Go→Lua card                    |
| P15 TRY/CATCH                  | Python ZeroDivisionError → Go fallback                |
| P16 SPLIT + MERGE              | Python splits → async LLMs → synthesis                |
| P17 CONFIG Pipeline            | CONFIG docs, bash sets runtime values                 |
| P18 FOREACH + Async            | Async pair per item + bash barrier                    |
| P19 WHILE Quality Gate         | Rust writes done only when score >= threshold         |
| P20 Map-Reduce                 | Python splits → 3 async MAP → REDUCE                  |
| P21 TRY/CATCH + LOOP           | Rust TRY + Go CATCH + bash final status               |
| P22 Multi-Variable Output      | Python 3-outvar, writes directly to files             |
| P23 CONFIG + WHILE + 3 Models  | Rust gate + LLM @3 polish after loop                  |
| P24 FOREACH + TRY/CATCH        | Fault-tolerant batch JSON processing                  |

## 3.2 Core Patterns

---

Start with these patterns. They cover the vast majority of real-world use cases.

### P1 — Linear Pipeline (recommended starting point)

The simplest pattern. Data flows in one direction through a sequence of blocks. Each block processes the output of the previous one.

```
```bash >raw
echo "42"
```
```python <raw >stats
import os
val = int(open(os.environ['MSH_VAR_raw']).read().strip())
print(f'square={val*val} cube={val**3}')
```
<!--@1 <stats >analysis
Interpret these statistics in one sentence. You MUST respond.
-->
```lua <raw <stats <analysis
local function read(env)
    local f=io.open(os.getenv(env),"r"); if not f then return "(n/a)" end
    local s=f:read("*all"):gsub("%s+$",""); f:close(); return s
end
print("Raw: "..read("MSH_VAR_raw"))
print("Stats: "..read("MSH_VAR_stats"))
print("Analysis: "..read("MSH_VAR_analysis"))
```
```

### P5 — Two-LLM Review Chain (for quality-critical output)

Model @1 generates, Model @2 reviews and scores, Model @1 improves based on the review. This pattern is used in production systems (e.g., email agent) for defense-in-depth.

```
<!--@1 <task >draft
Write a professional email response. You MUST respond.
-->
<!--@2 <draft >review
Review this draft. Score 1-10. Reply: SCORE: N STATUS: APPROVED/REVISE
You MUST respond.
-->
<!--@1 <draft <review >final
Improve the draft based on the review. Return only the improved text.
-->
```

### P19 — WHILE Quality Gate (iterative improvement)

The loop runs until a numeric quality score reaches a threshold. Rust reads the score and writes 'done' to the status variable when the threshold is met. This is the correct way to exit a quality loop — not through the LLM's own judgment.

```
```bash >status
echo "running"
```

```

...
<!--@while status:running-->
```c <counter >counter
// increment counter via fopen/fprintf
...

<!--@1 <task >code
Write the Python function. You MUST respond.
-->

<!--@2 <code >score
Score this code 1-10. Reply with ONLY the integer. You MUST respond.
-->
```rust <score <counter >status
let score: u32 = s.trim().parse().unwrap_or(0);
let done = score >= 8 || iter >= 5;
fs::write(env::var("MSH_VAR_status").unwrap(),
    if done {"done"} else {"running"}).unwrap();
...
<!--@end_while-->

```

### 3.3 Async Patterns

Async LLM directives allow multiple models to run in parallel. A bash await barrier synchronizes all of them before continuing.

```

<!--@1 async <topic >answer1
Explain from technical perspective. You MUST respond.
-->

<!--@2 async <topic >answer2
Explain with a real-world analogy. You MUST respond.
-->

<!--@3 async <topic >answer3
Explain for a beginner. You MUST respond.
-->
```bash await=answer1,answer2,answer3
for p in "$(printenv MSH_VAR_answer1)" "$(printenv MSH_VAR_answer2)"
"$$(printenv MSH_VAR_answer3)"; do
    i=0; while [ ! -s "$p" ] && [ $i -lt 120 ]; do sleep 1; i=$((i+1)); done
done
echo "all ready"
...

<!--@1 <answer1 <answer2 <answer3 >synthesis
Synthesize these three perspectives into one definitive answer.
-->

```

**NOTE** Await barriers must be in bash blocks only. Other languages recompile on each iteration and can miss file creation events.

## 4. Creating Generative Workflows

### 4.1 What is a Generative Workflow?

A generative workflow is a .md file that uses LLM directives to generate code, text, or data at runtime. Every time you run it, the LLM produces fresh output. The workflow orchestrates multiple models, languages, and data transformations into a coherent pipeline.

Generative workflows are ideal for: rapid prototyping, generating multiple variants of code, AI-assisted data analysis, and creating the source material that will later be materialized.

### 4.2 Structure of a Generative Workflow

A well-structured generative workflow follows this pattern:

1. Configuration — bash blocks that set parameters and seed data
2. Generation — LLM directives that create code or content
3. Validation — code blocks that test the generated output
4. Transformation — additional processing in compiled languages
5. Report — Lua block that renders the final result

### 4.3 Writing Effective LLM Requests

The quality of your prompt directly determines the quality of the output. Follow these principles:

#### Be explicit about format

Tell the model exactly what format to use. Vague prompts produce inconsistent output.

```
<!-- BAD: vague -->
<!--@1 >code
Write a sorting function.
-->

<!-- GOOD: explicit -->
<!--@1 >code
Write a Python function sort_numbers(lst) that sorts a list of integers.
Return ONLY the function - no imports, no explanation, no markdown fences.
You MUST respond.
-->
```

#### Always add 'You MUST respond.'

Some models return empty responses on cold starts or ambiguous prompts. Adding this phrase prevents silent failures.

#### Pass context through variables

Use <varname to include previous results in the prompt context. The model receives the file content automatically.

```
```python >data_summary
```

```

# compute statistics, print to stdout
print(f'mean=25.3 stddev=4.1 count=100')
...

<!--@1 <data_summary >interpretation
The input contains statistical data about a dataset.
Write one paragraph interpreting what these numbers suggest.
You MUST respond.
-->

```

## 4.4 Three Models — When to Use Each

With three models configured, you can assign each a specific role in your workflow. A common pattern from production systems:

Model	Recommended Role
@1 (primary)	Main generation — code writing, content creation, primary LLM tasks
@2 (reviewer)	Quality review — code scoring, legal/security checks, second opinion
@3 (specialist)	Specific tasks — final polish, specialized domain, alternative perspective

Example: generating and reviewing code with two models:

```

<!--@1 <task >code
Write the implementation. Return only code. You MUST respond.
-->

<!--@2 <code >review
You are a senior developer. Review this code.
Score 1-10 for correctness, edge cases, style.
Reply: SCORE: N ISSUES: description or none
You MUST respond.
-->

```python <review
import os, re
review = open(os.environ['MSH_VAR_review']).read()
score = re.search(r'SCORE: (\d+)', review)
if score and int(score.group(1)) >= 8:
    print('APPROVED')
else:
    print('NEEDS_REVISION')
...

```

## 4.5 Referencing Files and URLs

### PDF and document files

mshell automatically extracts text from PDF, DOCX, XLSX, and DjVu files when they are referenced in an LLM prompt. Simply mention the filename in the prompt text:

```

<!--@1 >summary
Summarize the key findings from report.pdf

```

```
Focus on the executive summary and recommendations.
You MUST respond.
-->
```

mshell detects the file reference, extracts text via pdftotext, and includes it in the LLM context automatically. The FILE INCLUSION SUMMARY in the output confirms what was included.

## Web URLs

Use urlcache to fetch and cache web pages, then reference the cache in LLM prompts:

```
# In interactive mode:
urlcache load https://docs.python.org/3/library/json.html

# Reference in workflow:
<!--@1 cached=1 >explanation
Based on the Python JSON documentation, explain how to handle nested
objects.
You MUST respond.
-->
```

## 4.6 Complete Generative Example

---

This example generates a system monitoring tool using two models: one generates Go code for data collection, another generates Python for analysis, and both are validated before use.

```
# System Monitor Generator

```bash >task_go
echo "Write a Go program that reads /proc/meminfo and prints MemTotal in
GB."
```

```bash >task_python
echo "Write a Python script that reads /proc/mounts and shows disk usage."
```

<!--@1 <task_go >monitor_go
Write ONLY the Go program described. Use only standard library.
Return complete working code in a ```go code block.
You MUST respond.
-->

<!--@2 <task_python >monitor_py
Write ONLY the Python script described. Use only standard library.
Return complete working code in a ```python code block.
You MUST respond.
-->

```bash
echo "=== System Monitor Generated ==="
echo "Go component: $(wc -l < $MSH_VAR_monitor_go) lines"
echo "Python component: $(wc -l < $MSH_VAR_monitor_py) lines"
```
```



## 5. Creating Materialized Workflows

### 5.1 What is Materialization?

Materialization is the process of converting a generative workflow into a deterministic, LLM-free version. When you run a generative `.md` file with `MSHELL_MATERIALIZE=1`, `mshell` creates a `_materialized.md` file alongside it.

The materialized file contains the actual code generated by the LLMs during that run — with all LLM directives replaced by the concrete code they produced. Running the materialized file executes the same logic every time, without any LLM calls.

| Generative <code>.md</code>                             | Materialized <code>.md</code>           |
|---------------------------------------------------------|-----------------------------------------|
| Contains LLM directives<br><code>&lt;!--@1--&gt;</code> | Contains actual generated code blocks   |
| Different output each run                               | Identical output every run              |
| Requires LLM API / model                                | No LLM needed — pure code execution     |
| Slow (LLM inference time)                               | Fast (direct compilation and execution) |
| Source of truth for intent                              | Production-ready artifact               |

### 5.2 Enabling Materialization

Set the environment variable before running your workflow:

```
MSHELL_MATERIALIZE=1
edi my_workflow.md

# mshell creates: my_workflow_materialized.md
# in the same directory
```

To run the materialized version (no LLM needed):

```
MSHELL_MATERIALIZE=0
edi my_workflow_materialized.md
```

**NOTE** Variables and attributes (`<var >var`) are preserved in the materialized file. Pipeline data flow works identically — only the LLM generation is removed.

### 5.3 What Gets Materialized

During materialization, `mshell` records:

- LLM responses — the actual code generated by `<!--@1-->` directives, with `>outvar` attributes
- Code blocks — preserved as-is with all `<invar >outvar` attributes intact
- Text between blocks — headings and prose from the original document
- Loop iterations — all iterations are recorded, each code block preserved

What is NOT materialized:

- LLM directive syntax itself — replaced by the generated code
- Loop/end\_loop directives — replaced by the actual iterated blocks

## 5.4 Materialization Workflow in Practice

---

6. Write your generative .md with LLM directives
7. Run with MSHELL\_MATERIALIZE=1 — review the generated code
8. If the output is good, keep the materialized file
9. Run the materialized file without LLM for fast, reproducible execution
10. Commit both files: the generative .md as source, materialized as artifact

## 5.5 Complete Materialization Example

---

Start with this generative workflow:

```
# Data Pipeline - Generative

<!--@1 >data_generator
Write a Python script that generates 20 rows of CSV data:
columns: name, age, salary. Use random but realistic values.
Write to stdout as valid CSV with header.
Return only code in a ```python code block. You MUST respond.
-->

```bash >data
echo 'Data generated'
```

<!--@1 <data_generator >analyzer
Write a Python script that reads CSV from MSH_VAR_data_generator file path,
computes count, average age, average salary, min and max salary.
Print results as key=value lines. Return only code in ```python block.
You MUST respond.
-->

```python <analyzer <data_generator
import os, subprocess
exec(open(os.environ['MSH_VAR_analyzer']).read())
```
```

After running with MSHELL\_MATERIALIZE=1, the materialized version contains the actual Python code instead of LLM directives — ready to run repeatedly without any model.

# 6. Practical Reference

## 6.1 mshell System Prompt Setup

To give LLMs knowledge of mshell Workflow syntax automatically (without explaining it in every prompt), configure a system prompt file. This is especially valuable when working with models that have no prior knowledge of mshell.

### Step 1: Create the system prompt file

```
# Create ~/.mshell_system.txt with mshell syntax reference
# Content should cover: code block syntax, variable system,
# LLM directives, language roles, critical rules
chmod 600 ~/.mshell_system.txt
```

### Step 2: Reference it in ~/.mshellrc

```
# Add this line to ~/.mshellrc:
MSHELL_SYSTEM_FILE=~/.mshell_system.txt
```

mshell reads this file at startup and injects it as a system prompt into every LLM call — for Claude as a top-level 'system' field, for OpenAI/Ollama as the first 'system' role message.

## 6.2 Model Compatibility Testing

Before building complex workflows, verify that your configured models understand mshell syntax. Run the provided test files:

```
validator disable
MSHELL_MATERIALIZE=0
edi mshell_test_model1.md # test @1
edi mshell_test_model2.md # test @2
edi mshell_test_model3.md # test @3
```

Each test checks four capabilities:

| Test                      | What it checks                                      |
|---------------------------|-----------------------------------------------------|
| Test 1: Bash block        | Model generates correct ```bash block format        |
| Test 2: Variable passing  | Model uses MSH_VAR_ correctly in generated code     |
| Test 3: Structured format | Model follows exact multi-line output format        |
| Test 4: Exec mode         | Model wraps code in fences for <!--@Nx--> execution |

Interpretation of results:

| Result               | Recommendation                                           |
|----------------------|----------------------------------------------------------|
| FULLY COMPATIBLE     | Use all 24 patterns including async and FOREACH          |
| MOSTLY COMPATIBLE    | Avoid complex multi-variable output; use simple patterns |
| PARTIALLY COMPATIBLE | Use P1-P4 linear patterns only                           |
| NOT COMPATIBLE       | Switch to a more capable model                           |

## 6.3 Common Mistakes and How to Avoid Them

| Mistake                                                       | Correct approach                                                    |
|---------------------------------------------------------------|---------------------------------------------------------------------|
| Python prints to stdout with multiple <code>&gt;outvar</code> | Write directly to <code>open(os.environ['MSH_VAR_out'], 'w')</code> |
| Await barrier in Python/Go/Rust                               | Always use bash blocks for await barriers                           |
| Loop counter in Python inside <code>WHILE</code>              | Use C with <code>fopen/fprintf</code> — more reliable in loops      |
| No 'You MUST respond.' in prompt                              | Always add this phrase to prevent empty LLM responses               |
| <code>\uXXXX</code> escapes in Lua                            | Use actual UTF-8 characters or <code>\xNN</code> byte sequences     |
| Running test with validator on                                | Disable validator: <code>validator disable</code>                   |
| Materializing with <code>MSHELL_MATERIALIZE=1</code> left on  | Always reset: <code>MSHELL_MATERIALIZE=0</code> after use           |
| Async without await barrier                                   | Always add bash await block after async directives                  |

## 6.4 Recommended Settings for Development

```
# Before writing a new generative workflow:
validator verbose on
validator status
MSHELL_MATERIALIZE=0

# To create materialized version:
MSHELL_MATERIALIZE=1
edi my_workflow.md
MSHELL_MATERIALIZE=0

# To test materialized version:
validator disable
edi my_workflow_materialized.md
validator enable
```

## 7. Conclusion

mshell Workflow represents a fundamentally different approach to AI-assisted development. Instead of separate orchestration frameworks, API clients, and build systems, everything lives in a single Markdown document that is simultaneously executable code and human-readable documentation.

The generative/materialized workflow pair is the key insight: you design with AI assistance (generative), then freeze the result for reliable production use (materialized). This gives you the creative power of LLMs during development and the determinism of pure code in production.

### Key principles to remember

- Start simple — P1 Linear Pipeline covers most use cases
- Follow language roles — Bash seeds, C counts, Rust gates, Lua reports
- Always add 'You MUST respond.' to prompts
- Test your models before building complex workflows
- Use materialization for production — no LLM dependency at runtime
- The system prompt file makes models understand mshell from the first call

The 24 reference patterns, the Interlang variable system, and the dual-LLM review architecture give you everything needed to build production-quality agentic systems — on any hardware from a Raspberry Pi to a datacenter server, using any LLM vendor or local model.

## 8. References

mshell Workflow. NxN Matrix Multiplication Performance Study:

<https://www.appservgrid.com/paw92/index.php/2026/04/07/mshell-workflow-nxn-matrix-multiplication-performance-study/>

Anatomy of High-Performance Matrix Multiplication. Kazushige Goto, R. V. D. Geijn, 2008:

[https://www.cs.utexas.edu/~flame/pubs/GotoTOMS\\_final.pdf](https://www.cs.utexas.edu/~flame/pubs/GotoTOMS_final.pdf)

HPL Linpack Benchmark: <https://top500.org/project/linpack/>

OpenBLAS Optimized BLAS library: <https://github.com/OpenMathLib/OpenBLAS>

Go compiler documentation: <https://pkg.go.dev/cmd/compile>

Resources: - Common examples Part I (P1–P12):

<https://www.appservgrid.com/paw92/index.php/2026/02/26/mshell-workflow-patterns-reference-guide-part-i-p1-p13/>

Resources: - Common examples Part II (P13–P24):

<https://www.appservgrid.com/paw92/index.php/2026/03/11/mshell-workflow-patterns-reference-guide-part-ii-p13-p24/>

- mshell v1.4.1 cheatsheet: <https://www.appservgrid.com/paw92/index.php/2026/02/04/mshell-v-1-4-1-cheatsheet-january-26th-2026/>

Unified language Patterns for mshell Workflow — Complete Reference Guide (p1–p24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/24/unified-language-patterns-for-mshell-workflow-complete-reference-guide-p1-p24/>

Pure Python language Patterns for mshell Workflow — CompleteReference Guide (p1–p24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/18/pure-python-language-patterns-for-mshell-workflow-completereference-guide-p1-p24/>

Pure Bash language Patterns for mshell Workflow — CompleteReference Guide (P1–P24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/17/pure-bash-language-patterns-for-mshell-workflow-completereference-guide-p1-p24/>

Pure Lua language Patterns for mshell Workflow — Complete Reference Guide (p1–p24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/24/pure-lua-language-patterns-for-mshell-workflow-complete-reference-guide-p1-p24/>

Pure Go language Patterns for mshell Workflow — Complete Reference Guide (p1–p24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/23/pure-go-language-patterns-for-mshell-workflow-complete-reference-guide-p1-p24/>

Pure Rust language Patterns for mshell Workflow — Complete Reference Guide (p1–p24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/22/pure-rust-language-patterns-for-mshell-workflow-complete-reference-guide-p1-p24/>

Pure C++ language Patterns for mshell Workflow — CompleteReference Guide (p1–p24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/21/pure-c-language-patterns-for-mshell-workflow-completereference-guide-p1-p24-2/>

Pure C language Patterns for mshell Workflow — CompleteReference Guide (p1–p24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/19/pure-c-language-patterns-for-mshell-workflow-completereference-guide-p1-p24/>

Pure mshell language Patterns for mshell Workflow — CompleteReference Guide (P1–P24)

permanent link: <https://www.appservgrid.com/paw92/index.php/2026/03/17/pure-mshell-language-patterns-for-mshell-workflow-completereference-guide-p1-p24/>