

/skill-creator

Source: `~/ .claude/skills/skill-creator/SKILL.md`

name: skill-creator version: "2.0" level: 3 trigger: "create skill, new skill, update skill, skill creator, SKILL.md" author: john updated: 2026-03-16 description: Guide for creating Level 3+ skills. IF new skill request THEN scaffold SKILL.md with metadata + if/then workflow + verification + MAX TURNS. Update skill-registry.db on completion. license: Complete terms in LICENSE.txt

Skill Creator

This skill provides guidance for creating effective skills.

About Skills

Skills are modular, self-contained packages that extend Claude's capabilities by providing specialized knowledge, workflows, and tools. Think of them as "onboarding guides" for specific domains or tasks—they transform Claude from a general-purpose agent into a specialized agent equipped with procedural knowledge that no model can fully possess.

What Skills Provide

1. Specialized workflows - Multi-step procedures for specific domains
2. Tool integrations - Instructions for working with specific file formats or APIs
3. Domain expertise - Company-specific knowledge, schemas, business logic
4. Bundled resources - Scripts, references, and assets for complex and repetitive tasks

Core Principles

Concise is Key

The context window is a public good. Skills share the context window with everything else Claude needs: system prompt, conversation history, other Skills' metadata, and the actual user request.

Default assumption: Claude is already very smart. Only add context Claude doesn't already have. Challenge each piece of information: "Does Claude really need this explanation?" and "Does this paragraph justify its token cost?"

Prefer concise examples over verbose explanations.

Set Appropriate Degrees of Freedom

Match the level of specificity to the task's fragility and variability:

High freedom (text-based instructions): Use when multiple approaches are valid, decisions depend on context, or heuristics guide the approach.

Medium freedom (pseudocode or scripts with parameters): Use when a preferred pattern exists, some variation is acceptable, or configuration affects behavior.

Low freedom (specific scripts, few parameters): Use when operations are fragile and error-prone, consistency is critical, or a specific sequence must be followed.

Think of Claude as exploring a path: a narrow bridge with cliffs needs specific guardrails (low freedom), while an open field allows many routes (high freedom).

Anatomy of a Skill

Every skill consists of a required SKILL.md file and optional bundled resources:

```
skill-name/
├─ SKILL.md (required)
│  └─ YAML frontmatter metadata (required)
│     └─ name: (required)
│     └─ description: (required)
│     └─ compatibility: (optional, rarely needed)
├─ Markdown instructions (required)
└─ Bundled Resources (optional)
   ├─ scripts/          - Executable code (Python/Bash/etc.)
   ├─ references/       - Documentation intended to be loaded into context as needed
   └─ assets/           - Files used in output (templates, icons, fonts, etc.)
```

SKILL.md (required)

Every SKILL.md consists of:

- **Frontmatter** (YAML): Contains `name` and `description` fields (required), plus optional fields like `license`, `metadata`, and `compatibility`. Only `name` and `description` are read by Claude to determine when the skill triggers, so be clear and comprehensive about what the skill is and when it should be used. The `compatibility` field is for noting environment requirements (target product, system packages, etc.) but most skills don't need it.
- **Body** (Markdown): Instructions and guidance for using the skill. Only loaded AFTER the skill triggers (if at all).

Bundled Resources (optional)

Scripts (`scripts/`)

Executable code (Python/Bash/etc.) for tasks that require deterministic reliability or are repeatedly rewritten.

- **When to include:** When the same code is being rewritten repeatedly or deterministic reliability is needed
- **Example:** `scripts/rotate_pdf.py` for PDF rotation tasks
- **Benefits:** Token efficient, deterministic, may be executed without loading into context
- **Note:** Scripts may still need to be read by Claude for patching or environment-specific adjustments

References (`references/`)

Documentation and reference material intended to be loaded as needed into context to inform Claude's process and thinking.

- **When to include:** For documentation that Claude should reference while working
- **Examples:** `references/finance.md` for financial schemas, `references/mnda.md` for company NDA template, `references/policies.md` for company policies, `references/api_docs.md` for API specifications
- **Use cases:** Database schemas, API documentation, domain knowledge, company policies, detailed workflow guides
- **Benefits:** Keeps SKILL.md lean, loaded only when Claude determines it's needed
- **Best practice:** If files are large (>10k words), include grep search patterns in SKILL.md
- **Avoid duplication:** Information should live in either SKILL.md or references files, not both. Prefer references files for detailed information unless it's truly core to the skill—this keeps SKILL.md lean while making information discoverable without hogging the context window. Keep only essential procedural instructions and workflow guidance in SKILL.md; move detailed reference material, schemas, and examples to references files.

Assets (`assets/`)

Files not intended to be loaded into context, but rather used within the output Claude produces.

- **When to include:** When the skill needs files that will be used in the final output
- **Examples:** `assets/logo.png` for brand assets, `assets/slides.pptx` for PowerPoint templates, `assets/frontend-template/` for HTML/React boilerplate, `assets/font.ttf` for typography
- **Use cases:** Templates, images, icons, boilerplate code, fonts, sample documents that get copied or modified
- **Benefits:** Separates output resources from documentation, enables Claude to use files without loading them into context

What to Not Include in a Skill

A skill should only contain essential files that directly support its functionality. Do NOT create extraneous documentation or auxiliary files, including:

- README.md
- INSTALLATION_GUIDE.md
- QUICK_REFERENCE.md
- CHANGELOG.md
- etc.

The skill should only contain the information needed for an AI agent to do the job at hand. It should not contain auxiliary context about the process that went into creating it, setup and testing procedures, user-facing documentation, etc. Creating additional documentation files just adds clutter and confusion.

Progressive Disclosure Design Principle

Skills use a three-level loading system to manage context efficiently:

1. **Metadata (name + description)** - Always in context (~100 words)
2. **SKILL.md body** - When skill triggers (<5k words)
3. **Bundled resources** - As needed by Claude (Unlimited because scripts can be executed without reading into context window)

Progressive Disclosure Patterns

Keep SKILL.md body to the essentials and under 500 lines to minimize context bloat. Split content into separate files when approaching this limit. When splitting out content into other files, it is very important to reference them from SKILL.md and describe clearly when to read them, to ensure the reader of the skill knows they exist and when to use them.

Key principle: When a skill supports multiple variations, frameworks, or options, keep only the core workflow and selection guidance in SKILL.md. Move variant-specific details (patterns, examples, configuration) into separate reference files.

Pattern 1: High-level guide with references

```
# PDF Processing

## Quick start

Extract text with pdfplumber:
[code example]

## Advanced features

- Form filling: See [FORMS.md](FORMS.md) for complete guide
- API reference: See [REFERENCE.md](REFERENCE.md) for all methods
- Examples: See [EXAMPLES.md](EXAMPLES.md) for common patterns
```

Claude loads FORMS.md, REFERENCE.md, or EXAMPLES.md only when needed.

Pattern 2: Domain-specific organization

For Skills with multiple domains, organize content by domain to avoid loading irrelevant context:

```
bigquery-skill/
├─ SKILL.md (overview and navigation)
├─ reference/
│   ├─ finance.md (revenue, billing metrics)
│   ├─ sales.md (opportunities, pipeline)
│   └─ product.md (API usage, features)
```

```
└─ marketing.md (campaigns, attribution)
```

When a user asks about sales metrics, Claude only reads sales.md.

Similarly, for skills supporting multiple frameworks or variants, organize by variant:

```
cloud-deploy/  
└─ SKILL.md (workflow + provider selection)  
└─ references/  
    └─ aws.md (AWS deployment patterns)  
    └─ gcp.md (GCP deployment patterns)  
    └─ azure.md (Azure deployment patterns)
```

When the user chooses AWS, Claude only reads aws.md.

Pattern 3: Conditional details

Show basic content, link to advanced content:

```
# DOCX Processing  
  
## Creating documents  
  
Use docx-js for new documents. See [DOCX-JS.md](DOCX-JS.md).  
  
## Editing documents  
  
For simple edits, modify the XML directly.  
  
**For tracked changes**: See [REDLINING.md](REDLINING.md)  
**For OOXML details**: See [OOXML.md](OOXML.md)
```

Claude reads REDLINING.md or OOXML.md only when the user needs those features.

Important guidelines:

- **Avoid deeply nested references** - Keep references one level deep from SKILL.md. All reference files should link directly from SKILL.md.
- **Structure longer reference files** - For files longer than 100 lines, include a table of contents at the top so Claude can see the full scope when previewing.

Skill Creation Process

Skill creation involves these steps:

1. Understand the skill with concrete examples
2. Plan reusable skill contents (scripts, references, assets)
3. Initialize the skill (run `init_skill.py`)
4. Edit the skill (implement resources and write `SKILL.md`)
5. Package the skill (run `package_skill.py`)
6. Iterate based on real usage

Follow these steps in order, skipping only if there is a clear reason why they are not applicable.

Step 1: Understanding the Skill with Concrete Examples

Skip this step only when the skill's usage patterns are already clearly understood. It remains valuable even when working with an existing skill.

To create an effective skill, clearly understand concrete examples of how the skill will be used. This understanding can come from either direct user examples or generated examples that are validated with user feedback.

For example, when building an image-editor skill, relevant questions include:

- "What functionality should the image-editor skill support? Editing, rotating, anything else?"
- "Can you give some examples of how this skill would be used?"
- "I can imagine users asking for things like 'Remove the red-eye from this image' or 'Rotate this image'. Are there other ways you imagine this skill being used?"
- "What would a user say that should trigger this skill?"

To avoid overwhelming users, avoid asking too many questions in a single message. Start with the most important questions and follow up as needed for better effectiveness.

Conclude this step when there is a clear sense of the functionality the skill should support.

Step 2: Planning the Reusable Skill Contents

To turn concrete examples into an effective skill, analyze each example by:

1. Considering how to execute on the example from scratch
2. Identifying what scripts, references, and assets would be helpful when executing these workflows repeatedly

Example: When building a `pdf-editor` skill to handle queries like "Help me rotate this PDF," the analysis shows:

1. Rotating a PDF requires re-writing the same code each time
2. A `scripts/rotate_pdf.py` script would be helpful to store in the skill

Example: When designing a `frontend-webapp-builder` skill for queries like "Build me a todo app" or "Build me a dashboard to track my steps," the analysis shows:

1. Writing a frontend webapp requires the same boilerplate HTML/React each time
2. An `assets/hello-world/` template containing the boilerplate HTML/React project files would be helpful to store in the skill

Example: When building a `big-query` skill to handle queries like "How many users have logged in today?" the analysis shows:

1. Querying BigQuery requires re-discovering the table schemas and relationships each time
2. A `references/schema.md` file documenting the table schemas would be helpful to store in the skill

To establish the skill's contents, analyze each concrete example to create a list of the reusable resources to include: scripts, references, and assets.

Step 3: Initializing the Skill

At this point, it is time to actually create the skill.

Skip this step only if the skill being developed already exists, and iteration or packaging is needed. In this case, continue to the next step.

When creating a new skill from scratch, always run the `init_skill.py` script. The script conveniently generates a new template skill directory that automatically includes everything a skill requires, making the skill creation process much more efficient and reliable.

Usage:

```
scripts/init_skill.py <skill-name> --path <output-directory>
```

The script:

- Creates the skill directory at the specified path
- Generates a SKILL.md template with proper frontmatter and TODO placeholders
- Creates example resource directories: `scripts/`, `references/`, and `assets/`
- Adds example files in each directory that can be customized or deleted

After initialization, customize or remove the generated SKILL.md and example files as needed.

Step 4: Edit the Skill

When editing the (newly-generated or existing) skill, remember that the skill is being created for another instance of Claude to use. Include information that would be beneficial and non-obvious to Claude. Consider what procedural knowledge, domain-specific details, or reusable assets would help another Claude instance execute these tasks more effectively.

Learn Proven Design Patterns

Consult these helpful guides based on your skill's needs:

- **Multi-step processes:** See `references/workflows.md` for sequential workflows and conditional logic
- **Specific output formats or quality standards:** See `references/output-patterns.md` for template and example patterns

These files contain established best practices for effective skill design.

Start with Reusable Skill Contents

To begin implementation, start with the reusable resources identified above: `scripts/`, `references/`, and `assets/` files. Note that this step may require user input. For example, when implementing a `brand-guidelines` skill, the user may need to provide brand assets or templates to store in `assets/`, or documentation to store in `references/`.

Added scripts must be tested by actually running them to ensure there are no bugs and that the output matches what is expected. If there are many similar scripts, only a representative sample needs to be tested to ensure confidence that they all work while balancing time to completion.

Any example files and directories not needed for the skill should be deleted. The initialization script creates example files in `scripts/`, `references/`, and `assets/` to demonstrate structure, but most skills won't need all of them.

Update SKILL.md

Writing Guidelines: Always use imperative/infinite form.

Frontmatter

Write the YAML frontmatter with `name` and `description`:

- `name`: The skill name
- `description`: This is the primary triggering mechanism for your skill, and helps Claude understand when to use the skill.
 - Include both what the Skill does and specific triggers/contexts for when to use it.
 - Include all "when to use" information here - Not in the body. The body is only loaded after triggering, so "When to Use This Skill" sections in the body are not helpful to

Claude.

- Example description for a `docx` skill: "Comprehensive document creation, editing, and analysis with support for tracked changes, comments, formatting preservation, and text extraction. Use when Claude needs to work with professional documents (.docx files) for: (1) Creating new documents, (2) Modifying or editing content, (3) Working with tracked changes, (4) Adding comments, or any other document tasks"

Do not include any other fields in YAML frontmatter.

Body

Write instructions for using the skill and its bundled resources.

Step 5: Packaging a Skill

Once development of the skill is complete, it must be packaged into a distributable `.skill` file that gets shared with the user. The packaging process automatically validates the skill first to ensure it meets all requirements:

```
scripts/package_skill.py <path/to/skill-folder>
```

Optional output directory specification:

```
scripts/package_skill.py <path/to/skill-folder> ./dist
```

The packaging script will:

1. **Validate** the skill automatically, checking:
 - YAML frontmatter format and required fields
 - Skill naming conventions and directory structure
 - Description completeness and quality
 - File organization and resource references
2. **Package** the skill if validation passes, creating a `.skill` file named after the skill (e.g., `my-skill.skill`) that includes all files and maintains the proper directory structure for distribution. The `.skill` file is a zip file with a `.skill` extension.

If validation fails, the script will report the errors and exit without creating a package. Fix any validation errors and run the packaging command again.

Step 6: Iterate

After testing the skill, users may request improvements. Often this happens right after using the skill, with fresh context of how the skill performed.

Iteration workflow:

1. Use the skill on real tasks
 2. Notice struggles or inefficiencies
 3. Identify how SKILL.md or bundled resources should be updated
 4. Implement changes and test again
-

Revision #5

Created 2026-02-18 08:42:04 UTC by John

Updated 2026-06-21 20:01:34 UTC by John